An Evidential Knowledge Representation for Drug-mechanisms and its Application to Drug Safety

Richard David Boyce

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2008

Program Authorized to Offer Degree: Medical Education and Biomedical Informatics

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

ProQuest LLC
789 E. Eisenhower Parkway
PO Box 1346
Ann Arbor, MI 48106-1346

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Richard David Boyce

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

Ira J. Kalet

Ira Kalet

Reading Committee:

Ira J. Kalet

Ira Kalet

Carol J. Collins

Carol J. Collins

Thomas K. Hazlet

Thomas K. Hazlet

Date: 14-Jul-2008

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Proquest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature 

Date 7/14/08

University of Washington

Abstract

An Evidential Knowledge Representation for Drug-mechanisms and its Application to Drug Safety

Richard David Boyce

Chair of the Supervisory Committee:
Professor Ira Kalet
Biomedical and Health Informatics

A major challenge to designers of informatics tools that help alert clinicians to potential drug-drug interactions (DDIs) is how to best assist clinicians when they must infer the potential risk of an adverse event between medication combinations that have not been studied together in a clinical trial. The central thesis of this dissertation is that DDI prediction using drug mechanism knowledge can help drug-interaction knowledge bases expand their coverage beyond what has been tested in clinical trials while avoiding prediction errors that occur when individual drug differences are not recognized. This dissertation describes a knowledge representation system, called the Drug Interaction Knowledge Base (DIKB), that uses a novel approach to linking and assessing evidence support for drug-mechanism assertions. The DIKB is the first knowledge-representation system we are aware of to use a computable model of evidence and a Truth Maintenance System to manage assertions in its knowledge-base. The novel approach to evidence management implemented in the DIKB enables its prediction accuracy and coverage to be optimized to a particular body of evidence; a feature that is very desirable for clinical decision support. The DIKB is also novel for its computable representation of the conjectures behind a specific application of evidence. These *evidence-use assumptions* enable the system to flag when a conjecture has become invalid and alert knowledge-base maintainers to the need to reassess their original interpretation of what assertions a piece of evidence supports. They are also used as evi-

dence is input into the system to help identify a pattern, called a *circular line of evidence support*, that is indicative of fallacious reasoning by evidence-base curators. The DIKB has been shown capable of accurately predicting clinically-relevant DDIs using only pharmacokinetic drug-mechanism knowledge and development of the system has helped to identify and evaluate potential informatics solutions to the challenges of representing, synthesizing, and maintaining drug mechanism knowledge.

TABLE OF CONTENTS

Page

List of Figures vi

List of Tables ix

Chapter 1: Introduction 1
1.1 Problem Statement 1
1.2 Outline of This Dissertation 4
1.2.1 A Clarification on the Use of the Word "we" Throughout This Text . 6

Chapter 2: An Evidential Knowledge-representation for Drug Mechanisms 7
2.1 Qualitative Pharmacokinetic Modeling of Drugs 8
2.2 Major Challenges and Related Work 12
2.2.1 Work Addressing Uncertain Mechanism Knowledge 13
2.2.2 Work Addressing Missing Mechanism Knowledge 15
2.2.3 Work Addressing the Dynamic Nature of Mechanism Knowledge . . . 16
2.3 Modeling with Evidence and Truth Maintenance 17
2.3.1 The `evidence-model` 19
2.3.2 The `ddi-theory` 26
2.4 Implementation and Examples 34
2.5 Discussion 40
2.5.1 The DIKB as an System for Research 40
2.5.2 Expanding the DIKB 41
2.6 Conclusion 41

Chapter 3: A Knowledge Representation for Predicting Clinically Meaningful Drug-drug Interactions by Pharmacokinetic Inhibition 43
3.1 Introduction 43
3.1.1 A Significant Problem and a Potential Solution 43
3.1.2 The "Class-based" Reasoning Strategy 44
3.1.3 The "Mechanism-based" Reasoning Strategy 45

3.1.4 Mechanism-based Reasoning Presents Informatics Challenges 47
3.2 A Computable Representation of how DDIs Occur by Metabolic Inhibition . 48
3.2.1 How DDIs Occur by Metabolic Inhibition 49
3.2.2 The Set of Inferences Made by the Knowledge Representation 50
3.2.3 Assumptions Made by the Knowledge Representation 52
3.3 The KR's Rules and Semantics 58
3.3.1 The Machinery for Reasoning - Declarative Rules 58
3.3.2 The KR Supports Default Reasoning 59
3.3.3 The Current Set of Rules and Assertion Types 60
3.3.4 Precise Definitions Provide KR Semantics 61
3.4 Validation and Evaluation 65
3.4.1 A Non-trivial Validation Test 65
3.4.2 The Hypothetical Drug "C-cure" 65
3.5 Discussion 69
3.5.1 The KR is a Very Simplistic Model 69
3.5.2 The KR's Ontological Commitments Have Strengths and Limitations . 69
3.5.3 The KR's Reasoning System Does Not Track Uncertainty as it Performs Inference 72
3.5.4 Related Work 73
3.5.5 Conclusion 75

Chapter 4: The Collection and Classification of Drug-mechanism Evidence 76
4.1 A Novel Method for Representing Evidence 76
4.2 Considerations for an Evidence Taxonomy Oriented Toward Confidence Assignment 79
4.2.1 PharmGKB's "Categories of Pharmacogenetics Evidence" 80
4.2.2 Medical Subject Headings Publication Types 80
4.2.3 Gene Ontology Evidence Codes 81
4.2.4 The Pathway Tool's Evidence Ontology 84
4.2.5 Curator Inferences and Default Assumptions 86
4.3 An Appropriate Evidence Collection and Maintenance Process 91
4.3.1 Step One: Seek Evidence for and Against Each Relevant Assertion . . 91
4.3.2 Step Two: Decide Each Evidence Item's Type Based on Definitions in the Evidence Taxonomy and Evaluate if an Evidence Item Meets the Inclusion Criteria for its Type 95
4.3.3 Step Three: Decide if There are any "Evidence-use Assumptions" . . 97

4.3.4 Step Four: Enter a Representation of the Evidence Item into the DIKB 98
4.3.5 Step Five: Computer-supported Evidence Maintenance Processes . . . 101
4.4 Our Experience Using the Method to Represent a Body of Drug-mechanism Evidence . . . 102
4.4.1 The DIKB Evidence Taxonomy . . . 103
4.4.2 The DIKB Inclusion Criteria . . . 103
4.4.3 The Evidence Collection Process . . . 105
4.4.4 The Current Evidence-base . . . 106
4.4.5 Discussion of the Current Evidence-base . . . 118
4.4.6 Limitations . . . 118
4.5 Conclusion . . . 118

Chapter 5: An Experiment with Levels-of-evidence and Belief Criteria . . . 120
5.1 Methods . . . 121
5.1.1 Criteria for Confirmed Interactions and Non-interactions . . . 122
5.1.2 The Collection of Pharmacokinetic Data . . . 124
5.1.3 Expert-defined Belief Criteria . . . 132
5.1.4 Automatically-generated Belief Criteria . . . 138
5.1.5 Searching the Adverse Event Reporting System . . . 141
5.2 Results . . . 144
5.2.1 Evaluation of the DIKB's Novel DDI Predictions Made Using the Best-performing Belief Criteria Strategies . . . 152
5.3 Discussion . . . 155
5.3.1 Other Interesting Features of the Prediction Sets . . . 157
5.3.2 Why was the DIKB's Coverage of the Validation-set Interactions Always Incomplete? . . . 159
5.3.3 Comparing the DIKB Predictions to Labeling Statements . . . 160
5.3.4 The JTMS Could be Leveraged to Optimize the Search for High-performing Strategies . . . 161
5.3.5 Limitations . . . 162
5.4 Conclusion . . . 163

Chapter 6: Contributions, Future Work, and Concluding Comments . . . 165
6.1 Research Contributions . . . 165
6.2 Future Work . . . 167
6.3 Secondary Results . . . 168

6.4 Concluding Remarks 170

Bibliography 172

Appendix A: How Big is the Gap in Scientific Knowledge About Drug-drug Interactions? 188

Appendix B: The DIKB's Rule-based Model of DDIs Occurring by Metabolite Inhibition 190
B.1 Rules that Model Metabolic Inhibition 191
B.2 Rules for Linking Metabolites to Active Ingredients and Ancestor Compounds 202
B.3 Modeling the Effect of Inhibition Through a Graph of Catalytic Reactions . . 203
B.4 Rules for Disjunctive Cases 213

Appendix C: Definitions for Each Assertion Type Used in the DIKB's Rule-base . . 216
C.1 The `primary-total-clearance-mechanism` Assertion 217
C.2 The `bioavailability` Assertion 217
C.3 The `first-pass-effect` Assertion 218
C.4 The `fraction-absorbed` Assertion 219
C.5 The `maximum-concentration` Assertion 220
C.6 The `inhibits` Assertion 220
C.7 The `does-not-inhibit` Assertion 220
C.8 The `in-vitro-selective-inhibitor-of-enzyme` Assertion 221
C.9 The `in-viVo-selective-inhibitor-of-enzyme` Assertion 221
C.10 The `substrate-of` Assertion 221
C.11 The `in-vitro-probe-substrate-of-enzyme` Assertion 221
C.12 The `is-not-substrate-of` Assertion 221
C.13 The `primary-total-clearance-enzyme` Assertion 222
C.14 The `primary-metabolic-clearance-enzyme` Assertion 223
C.15 The `inhibition-constant` Assertion 223
C.16 The `has-metabolite` Assertion 224
C.17 The `controls-formation-of` Assertion 224
C.18 The `polymorphic-enzyme` Assertion 224
C.19 The `pceut-entity-of-concern` Assertion 225
C.20 The `sole-PK-effect-alter-metabolic-clearance` Assertion 225
C.21 The `permanently_deactivates_catalytic_function` Assertion 225
C.22 The `does_not_permanently_deactivate_catalytic_function` Assertion . . 226

Appendix D: The DIKB Evidence Taxonomy 227

Appendix E: Inclusion Criteria and Required Actions for Evidence Types in the DIKB 233
E.1 Inclusion Criteria for Reviews (EV_Review) and Sub-classes 234
E.2 Inclusion Criteria for Published Observation Reports (EV_Obs_DI_CR) and Sub-classes 235
E.3 Inclusion Criteria for Pharmacokinetic Studies (EV_CT_Pharmacokinetic) and Sub-classes 235
E.4 Inclusion Criteria for Pharmacokinetic DDI Studies (EV_CT_DDI) and Sub-classes 236
E.5 Inclusion Criteria for Non-traceable Statements in Drug Product Labeling (Non_traceable_Drug_Label_Statement) and Sub-classes 237
E.6 Inclusion Criteria for Drug Enzyme Inhibition Experiments (EV_EX_Met_Enz_-Inhibit) and Sub-classes 238
E.7 Inclusion Criteria for Metabolic Enzyme Identification Experiments (EV_EX_-Met_Enz_ID) and Sub-classes 239

Appendix F: Entering and Viewing Evidence Using the DIKB's Web Interface . . . 241

Appendix G: The Final Validation Set of Drug-drug Interactions and Non-interactions 250

Appendix H: A Belief Criteria Questionnaire 270

Appendix I: The AERS Implementation and Our Use of It 280
I.1 How to Query 281
I.2 Queries Used to Search for Adverse Event Reports 282
I.2.1 A Template SQL for Efficient Queries of AERS for DDIs 282
I.2.2 Generic and Trade-names for Drug Products Containing Active Pharmaceutical Ingredients in the DIKB 284
I.2.3 Adverse-event Terms Used to Query AERS 286

Appendix J: A Sample AERS Report Returned from our Queries 293

Appendix K: DIPS Evaluations of Case Reports 305

LIST OF FIGURES

Figure Number — Page

2.1 Rules written in mock Prolog from the pilot knowledge base used to predict metabolic drug-drug interactions in the pilot knowledge base. 8

2.2 Queries for any drugs that inhibit or induce the primary clearance enzyme for another NTI drug whose clearance is primarily metabolism 10

2.3 An *evidential* knowledge-based system that links assertions about object properties to the evidence for and against those properties 17

2.4 A set of levels of evidence (LOE) used while developing the DIKB 24

2.5 An example inference rule for when a precipitant inhibits the metabolic clearance of an object drug 27

2.6 A small JTMS dependency network 28

2.7 A change in the belief state in one of the justifications propagates to dependant consequents 28

2.8 A JTMS dependency network showing "`bc-satisfied` assumptions" 30

2.9 A *test* DDI theory 36

3.1 The KR's model of metabolic clearance pathways is isomorphic to an acyclic graph of catalytic reactions 55

3.2 A very simple model that we designed for inferring the fraction of total clearance contributed by an enzyme 56

3.3 A rule declaring some, possibly negligible, inhibition of the metabolic clearance of a drug or drug metabolite 60

3.4 A rule declaring that some drug or drug metabolite, `?x`, inhibits the enzyme responsible for at least 50% of the *metabolic* clearance of another drug or drug metabolite, `?z`, whose *total* clearance by metabolism is at least 50% . . 61

3.5 A rule declaring that some drug or drug metabolite, `?x`, inhibits the enzyme responsible for at least 50% of the *total* clearance of another drug or drug metabolite, `?z` 62

3.6 Elements from the DIKB's structured vocabulary 63

3.7 A model of the metabolic clearance pathways for the hypothetical drug *C-cure* 66

3.8 Two rules that the KR uses infer that two drugs will not interaction by inhibition of a specific enzyme 72

4.1 DIKB maintainers can specify that assertions be justified without any evidence support 88

4.2 DIKB users can see when an assertion is a *default assumption* when they review its evidence support. 89

4.3 A circular line of evidence support that indicates circular reasoning within the evidence-base 100

4.4 The 16 drugs and 19 drug metabolites chosen for DIKB experiments 102

4.5 The number of evidence sources that provide a given number of evidence items in the current DIKB evidence-base 107

F.1 The welcome page of the DIKB's Web interface 242

F.2 DIKB curators begin the process of entering in evidence that supports or rebuts an assertion from this page 243

F.3 Non-quantitative assertion types have a pre-specified range of values that can be chosen from the drop-down box on this page 244

F.4 Curators are taken to this page if they select "Add assumptions" from the page shown in Figure F.3 244

F.5 Curators can confirm the assertions they want to link as *evidence-use assumptions* from this page or use the browser's "Back" button to make a change 245

F.6 Curators use the forms like the one provided on this page to enter more information on an evidence item 245

F.7 This figure shows a portion of the list that a curator has to choose from when specifying an evidence item's type 246

F.8 The last step of the evidence entry process requires curators to confirm all of the information they have entered for an evidence item 247

F.9 Curators and expert users can review the evidence linked to an assertion by clicking on a link to its information from a hyper-linked index of assertions. . 248

F.10 Clicking on the "DIKB data" link from the DIKB welcome page takes curators to a page that shows the evidence linked to all assertions in the DIKB and allows them to change classification status of any assertion 249

K.1 DIPS evaluation of a paper involving multiple case reports published by Auclair *et al* providing evidence of an interaction between itraconazole and clarithromycin [11] 306

K.2 DIPS evaluation of a paper involving multiple case reports published by Huynh *et al* providing evidence of an interaction between diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [89] 307

K.3 DIPS evaluation of a case report published by Akram *et al* providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [5] 308
K.4 DIPS evaluation of a case report published by Itakura *et al* providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [91] 309
K.5 DIPS evaluation of a paper involving multiple case reports published by Gladding *et al* providing evidence of seperate interactions between 1) diltiazem and atorvastatin, and 2) diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [67] 310
K.6 DIPS evaluation of a case report published by Shaukat *et al* providing evidence of a DDI between fluconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [157] 311
K.7 DIPS evaluation of a case report published by Peces and Pobes providing evidence of a DDI between diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [140] 312
K.8 DIPS evaluation of a case report published by Lewin *et al* providing evidence of a DDI between diltiazem and atorvastatin [113] 313
K.9 DIPS evaluation of two case reports published by Gilad and Lampl providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [66] 314
K.10 DIPS evaluation of a case report published by Spach *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [161] 315
K.11 DIPS evaluation of a case report published by Wong *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [179] 316
K.12 DIPS evaluation of a case report published by Ayanian *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [28] 317
K.13 DIPS evaluation of two case reports published by Stein *et al* providing evidence of a DDI between ketoconazole and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [162] 318
K.14 DIPS evaluation of a case report published by Kahri *et al* providing evidence of a DDI between fluconazole and atorvastatin [101] 319
K.15 DIPS evaluation of a case report published by Grunden and Fisher providing evidence of a DDI between clarithromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [77] 320

LIST OF TABLES

Table Number — Page

2.1 The mapping between strength of inhibition in Reference A and Reference B that we used when collecting drug facts for the pilot knowledge base (KB). . 9

2.2 Predicted inhibition interactions from our pilot drug-drug interaction knowledge-base not documented in any of four online references 11

2.3 The `Pceut_Entity` class description 20

2.4 The `Drug` class description 21

2.5 The `Metabolite` class description 21

2.6 The `Pharmaceutical_Preparation` 21

2.7 The `Enzyme` class description 22

2.8 Slots in class `Assertion` 22

2.9 Slots in class `Evidence` 23

2.10 Every evidence item entered into the `DIKB` receives a label from a taxonomy of evidence types 25

2.11 Relevant slots from a hypothetical instance of class `Evidence` that possesses an *evidence-use assumption* 32

2.12 The belief state of a subset of drug property assertions and inferences in the development version of the DIKB 38

2.13 A hypothetical example showing how the addition of new evidence to the DIKB can effect a change in the predictions that are considered believable . . 39

4.1 A partial listing of the many sources for drug-mechanism knowledge 90

4.2 The evidence board used only one-third of the 36 types in the evidence taxonomy to classify all the 257 non-redundant evidence items 111

4.3 The distribution of evidence types among evidence items used to ***support*** drug mechanism assertions. 112

4.4 The distribution of evidence types among evidence items used to ***refute*** drug mechanism assertions. 113

4.5 A partial listing of the 74 evidence items that the evidence board rejected for use as support or rebuttal for at least one assertion in the DIKB. 115

4.6 A sample of five of the 58 evidence items in the DIKB's evidence-base that were entered with *evidence-use assumptions* 117

5.1 Nine drug/drug or drug/drug-metabolite pairs that were excluded from our analysis of the DIKB's accuracy 127
5.2 A small sample of the 65 statements the evidence board located in drug-product labeling that mentioned a pharmacokinetic interaction or non-interaction between one of the drug/drug and drug/drug-metabolite combinations shown in Appendix G 128
5.3 The evidence-board accepted none of the 35 case reports it found that were relevant for use in the validation set 130
5.4 Neither drug or drug metabolite in each of the pairs shown in this table are expected to be the victim of a metabolic inhibition interaction effected by the other drug or drug metabolite in the pair 132
5.5 The *interactions* in the validation set used to characterize the DIKB's prediction performance 133
5.6 The evidence-board used the *ranking categories* shown in this table 134
5.7 The levels of evidence (LOEs) defined by the evidence-board as described in Chapter 5, Section 5.1.4 135
5.8 This table shows the evidence board's *belief criteria* strategy 139
5.9 The set of assertion types not labeled as a *default assumption* shown with the number of LOEs that were defined for each of them 140
5.10 Seventeen interaction and non-interaction predictions made by the DIKB using the evidence-board's *belief criteria strategy* 145
5.11 Summary statistics for each performance parameter we analyzed over 35,599 prediction sets 146
5.12 1,152 (3%) strategies caused the DIKB to perform optimally in terms of sensitivity, positive predictive value, and agreement with the validation set as measured by Cohen's kappa 147
5.13 Each of the 1,152 strategies that caused the DIKB to perform optimally in terms of sensitivity, positive predictive value, and agreement with the validation set caused the system to make the same 65 predictions 148
5.14 The 31 interaction and nine non-interaction predictions that are unknown, but not refuted, in the validation set 149
5.15 The range of LOEs used by 1,152 *belief criteria strategies* that, like the evidence-board's strategy, had perfect sensitivity and positive predictive value but also maximum coverage of and agreement with the validation set 154
5.16 27,648 strategies led the DIKB to predict an interaction or non-interaction countered by the validation set 159
5.17 Pairs in the validation set (Tables 5.5 and 5.4)for which the DIKB made no prediction using any the 35,599 *belief criteria strategies.* 160

ACKNOWLEDGMENTS

This project was funded in part by a National Library of Medicine Biomedical and Health Informatics Training Program grant (T15 LM07442) and an award from the Elmer M. Plein Endowment Research Fund (University of Washington School of Pharmacy). The DIKB ontology and evidence taxonomy were developed using the Protégé resource, which is supported by grant LM007885 from the United States National Library of Medicine.

The author expresses sincere appreciation to each person on my dissertation committee. The project that this dissertation recounts started with a discussion between Drs. Ira Kalet and Tom Hazlet on how to improve drug-drug interaction alerting systems. This discussion led to a very interesting class project for Dr. Kalet's "Life and Death Computing" course where I programmed a system that predicted drug-drug interactions based on drug indications and contra-indications. From there I began collaborating with Drs. John Horn and Carol Collins whom narrowed my focus to a narrow range of interactions occurring by metabolic inhibition. These two drug experts have spent many hours over a four year time span selecting and discussing the knowledge-base's content. Shortly after I began working with Drs Horn and Collins I observed that, to them, general features of a particular evidence item seemed to have significant influence on the confidence they would have in any hypothesis the evidence supported or refuted. I also noticed that they often disagreed considerably on the body of evidence required to establish the validity of a drug mechanism assertion. These observations led me to design the evidential knowledge representation this thesis describes. Dr. Kalet made a major contribution when he suggested that we could iterate through different belief criteria to empirically explore the inferential force of each evidence class. I implemented this idea and its results are, in my opinion, some of the most interesting of this entire dissertation. Drs. John Gennari and Efthimis Efthimiadis have contributed many methodological suggestions that helped to shape my analysis of the

knowledge base's prediction accuracy. I also acknowledge Dr. Brent Louie for helping with the pilot version of the knowledge base (when he was a student) by entering some drug facts into the system and comparing the system's predictions with information present in four online drug information databases. Last, but not least, I thank Andrea Civan for suggesting a large number of grammar and spelling corrections.

DEDICATION

This document is dedicated to my wife Kriss Boyce. I love, respect, and admire you Kriss; thank you for you love and support as we have journeyed toward this goal. Now that we begin a new chapter in our lives together I am so glad you are here to share this moment with me.

May the words of my mouth and the meditation of my heart
be pleasing in your sight,
O LORD, my Rock and my Redeemer.

Psalm 19:14 (NIV)

Chapter 1

INTRODUCTION

1.1 Problem Statement

A 2001 report from the Agency for Healthcare Research and Quality estimates that adverse drug events result in more than 770,000 injuries and deaths each year in the United States and cost up to $5.6 million per hospital, depending on the size of the hospital [1]. A more recent report from the US Institute of Medicine estimates that over 1.5 million preventable adverse drug events (ADEs) occur each year in America [137]. Preventable ADEs include situations where a patient is harmed because a clinician fails to avoid, or properly manage, an interacting drug combination. Multiple studies indicate that these *drug-drug interactions* are a significant source of preventable ADEs. For example, Gurwitz *et al*, in their cohort study of ADEs among older Americans receiving ambulatory care, found that 13.3% of preventable errors leading to an ADE involved the co-prescription of drugs for which their was a "...well established, clinically important interaction" [79]. Also, three separate case-control studies conducted by Juurlink *et al* using population-level health records from Ontario Canada[a] found that patients were considerably more likely to be hospitalized while taking a drug combination known to be capable of causing a specific interaction than when when not [100].

Factors contributing to the occurrence of preventable drug-drug interactions (DDIs) include a lack of knowledge of the patient's concurrent medications and inaccurate or inadequate knowledge of drug interactions by health care providers [41, 142]. Information technology, especially electronic prescribing systems with clinical decision support features,

[a] Juurlink *et al* conducted three case-control studies, each focusing on a different drug-drug interaction: 1) Digoxin toxicity while taking clarithromycin, 2) ACE inhibitor hyperkalemia while taking a potassium sparing diuretic, and 3) glyburide hypoglycemia while taking co-trimoxazole.

can help address each of these factors to varying degrees and there is currently a great deal of interest from both government and private organizations in expanding the use of information technology during medication prescribing and dispensing [137, 122]. Unfortunately, studies have found the DDI components of a wide variety of clinical decision-support tools to be sub-optimal in both the accuracy of their predictions and the timeliness of their knowledge. For example, a 2001 study of retail pharmacy alerting software found that, on average, the systems missed clinically relevant DDIs one-third of the time [85]. A 2005 study of hand-held prescribing guides found that all 11 systems in the study failed to inform users of at least one life-threatening DDI [123]. Another study published in 2005 found that one widely used drug interaction database could identify less than 15% of the clinically relevant DDIs involving the metabolism of 5 immuno-suppressive drugs [160].

What all of the systems in these studies have in common is that they rely upon some representation of drug knowledge to infer DDIs. Throughout this dissertation I will refer to any store of drug knowledge used by human or computer system for inferring DDIs as a "drug-interaction knowledge base." Instances of drug-interaction knowledge bases (KBs) range from small databases of drug-interaction studies to large systems that combine trained experts and informatics tools to address the difficult task of acquiring, maintaining and distributing drug-interaction information. Currently, a handful of large drug information databases are used as drug-interaction KBs in a large range of drug interaction alerting products and electronic prescribing tools [122]. Examples include First DataBank's National Drug Data File®, Thompson Micromedex's DRUGDEX® system, and Cerner's Multum® system.

I have learned from discussions with clinicians, commentaries by DDI experts [142], reports by researchers [88, 160], and my own testing of systems that the basic service most drug-interaction KBs provide is to catalog drug pairs found to interact in clinical trials or reported as such in clinician-submitted case reports. One major limitation of this approach is that it contrains drug-interaction KBs, and the tools that utilize them, to covering only interacting drug pairs that KB maintainers find in the literature and think important to include. Clinicians often must infer the potential risk of an adverse event between medication combinations that have not been studied together in a clinical trial [142]. Systems that only

catalog DDI studies involving drug pairs can offer little or no support in these situations.

Some contemporary drug interaction KBs supplement their DDI knowledge by generalizing interactions involving some drug to all other drugs within its therapeutic class. For example, Tyken Hsieh describes the KB supporting the hospital prescription order entry system at Brigham and Women's Hospital as

> ...class-based hierarchical ingredient knowledge (e.g. ampicillin is penicillin) as well as cross-reactivity mapping (e.g. penicillins cross-react with cephalosporins) [88]

While clinically relevant class-based interactions exist (for example, the SSRIs and NSAIDs), this approach has been criticized for leading to some DDI predictions that are either false or are likely to have little clinical relevance [82]. The main reason class-based prediction can lead to false alerts is because drugs within a therapeutic class do not necessarily possess the same mechanistic properties. For example, drugs within a class can be metabolized by different enzymes and thus, have distinct metabolic interactions. False predictions can have a negative effect on electronic prescribing systems by triggering false or irrelevant DDI alerts that can markedly impede the workflow of care providers [168]. A high rate of irrelevant alerting is a potential barrier to widespread adoption of CPOE systems [147] and stands as a major obstacle to improving patient safety.

The central thesis of this dissertation is that DDI prediction using drug-mechanism knowledge can help drug-interaction KBs expand their coverage beyond what has been tested in clinical trials while avoiding prediction errors that occur when individual drug differences are not recognized. Mechanism-based DDI prediction itself is not novel; the mechanistic principles of drug-drug interactions can be found in several sources including pharmaceutics text books [68, 112]. There are a few basic pharmacologic principles by which drug-mechanism knowledge can be synthesized to make *mechanism-based* drug-drug interaction predictions. Pharmacodynamic interactions can occur when the pharmacodynamic effects of two drugs combine in additive, subtractive or synergistic ways. Pharmacokinetic interactions can occur when the binding, metabolic or physical and chemical properties of one drug affect the absorption, distribution, metabolism and/or excretion (ADME) of an-

other drug. Modulations in the ADME or the pharmacodynamic effects of a drug can lead to the possible negative outcomes of drug toxicity or loss of efficacy in patients.

Part of pre-clinical drug development is the use of mechanism-based DDI prediction to predict interactions between a new drug candidate and drugs currently on the market [26]. Most systems that model drug mechanisms are being developed and applied in the pre-clinical and pre-market phases of drug-development to identify optimal drug candidates, predict drug properties, assess the efficacy and safety of new drugs, and estimate dose to concentration relationships [56]. These early-phase modeling efforts are geared towards identifying interactions between a new drug and drugs with which it might be co-administered early on, before much time and money is invested [167, 180]. The predictions made using drug-mechanisms are generally qualitative; they indicate that two drugs might interact via a mechanism but offer no estimate of the magnitude of the interaction. Scientists can use qualitative predictions to select the set of clinical trials necessary to establish a new drug's safety profile [131].

The same knowledge that is useful for predicting DDIs in the premarket setting can help clinicians in the post-market setting assess the possibility of a DDI occurring between two drugs that have never been studied together in clinical trials [82]. In spite of this fact, and the position of the FDA that all relevant information on mechanisms from pre-market investigations be included in drug product labeling [26], little research has been done on how to best represent and synthesize drug-mechanism knowledge to support clinical decision making. This dissertation will fill in part of this knowledge gap by exploring novel informatics methods for representing drug mechanism knowledge for the purpose of making clinically relevant DDI predictions.

1.2 Outline of This Dissertation

- Chapter 2, *An evidential knowledge-representation for drug mechanisms*, proposes that correctly linking and assessing the evidence support for drug-mechanism assertions can enable knowledge-based systems to make clinically relevant drug-drug interactions in spite of the uncertain, incomplete, and dynamic nature of drug-mechanism knowledge. This chapter focuses on a set of novel informatics methods designed to test this

proposition and a prototype system, the Drug Interaction Knowledge Base (DIKB), that implements the new methods.

- Chapter 3, *A knowledge representation for predicting clinically meaningful drug-drug interactions by pharmacokinetic inhibition*, presents the new rule-based theory used in the DIKB. The theory predicts metabolic inhibition interactions and non-interactions between drug active ingredients and/or drug metabolites and categorizes its predictions into three discrete levels so that clinicians can assess the clinical relevance of each prediction. Experiments demonstrate that 1) the rule-based theory makes accurate predictions for an important class of DDIs using only knowledge of drug-mechanisms and 2) the system's prediction accuracy and coverage varies depending on the belief criteria strategy being used.

- Chapter 4, *The collection and classification of drug-mechanism evidence*, explores the DIKB's evidence representation method from a knowledge-base maintenance perspective. It begins with a brief summary of the method's goals and key assumptions contrasting it with other biomedical informatics systems that link evidence to their assertions. It then relates how the method was used to represent drug-mechanism evidence for 16 active ingredients and 19 active metabolites.

- Chapter 5, *An experiment with levels-of-evidence and belief criteria*, recounts an experiment conducted to characterize the effect of different *belief criteria strategies* on the system's accuracy and coverage of DDIs present in a reference set of interactions and non-interactions. This chapter also examines 31 novel predictions made by the DIKB using the best performing strategies and discusses an attempt to find evidence for these interactions in published case reports and data in the FDAs Adverse Event Reporting System.

- Chapter 6, *Contributions, future work, and concluding comments*, concludes with a review of this project's research contributions and a discussion of possible future directions.

1.2.1 A Clarification on the Use of the Word "we" Throughout This Text

I use the word *we* frequently throughout the text to acknowledge the fact that this work would not have been possible without collaboration and guidance of the persons mentioned in the Acknowledgement section. My specific contributions to this work include the DIKB and its Web interface, the DIKB's rule-based theory of metabolic inhibition interactions and non-interactions, the evidence collection process used to build the DIKB's evidence-base, and the design, implementation, and analysis of an experiment characterizing the effect of different *belief criteria strategies* on the DIKB's prediction accuracy. There are several locations in later chapters where I refer to members of an evidence board consisting of two drug experts and an informaticist. The two drug experts are Drs Carol Collins and John Horn and the informaticist is myself. I use this language to convey to the reader the interdisciplinary approach that we used to construct the evidence base.

Chapter 2

AN EVIDENTIAL KNOWLEDGE-REPRESENTATION FOR DRUG MECHANISMS

The same knowledge about drug mechanisms that is useful for predicting drug-drug interactions (DDIs) in the pre-market setting can help clinicians in the post-market setting assess the possibility of a DDI occurring between two drugs that have never been studied together in clinical trials [82]. In spite of this fact, and the position of the FDA that all relevant information on mechanisms from pre-market investigations be included in drug product labeling [26], little research has been done on how to best represent, utilize, and maintain drug-mechanism knowledge for the purpose of making DDI predictions in the post-market setting.

This chapter describes the development of a knowledge representation system, called the Drug Interaction Knowledge Base (DIKB), that has been shown to be capable of accurately predicting clinically-relevant DDIs using pharmacokinetic drug-mechanism knowledge. The process of developing the DIKB has helped to identify and evaluate potential informatics solutions to the challenges of representing and synthesizing drug-mechanism knowledge for post-market use. The system's design is based on the idea that, for a knowledge resource with drug-mechanism knowledge to be of clinical use, it is essential that it explicitly link each of its drug-mechanism facts to their evidence support. As a result, the DIKB implements a rich representation of evidence for and against propositions and uses that representation to support or refute assertions in its knowledge base. The system's prediction performance has been characterized and its development has led to other informatics contributions including a novel rule-based DDI prediction theory and an ontology of research evidence types. Later chapters will discuss these contributions in detail; this chapter focuses on the motivation for developing the new methods, what they are, and the design and implementation aspects of the DIKB.

```
(<- (metabolic-inhibit-interact ?precip ?object ?enz)
    (and (inhibits-primary-clearance-enzyme ?precip ?object ?enz)
         (narrow-ther-index ?object yes)))

(<- (inhibits-primary-clearance-enzyme ?precip ?object ?enz)
    (and (inhibits-partial-clearance ?precip ?object ?enz)
         (major-pathway ?object ?enz)))

(<- (inhibits-partial-clearance ?precip ?object ?enz)
    (and (inhibits-effectively ?precip ?enz)
         (substrate-of ?object ?enz)
         (primary-clearance-mechanism ?object metabolic)))

(<- (inhibits-effectively ?drug ?enz)
    (and (inhibits ?drug ?enz)
         (or (inhibit-strength ?drug ?enz strong)
             (inhibit-strength ?drug ?enz moderate))))
```

Figure 2.1: Rules written in mock Prolog from the pilot knowledge base used to predict metabolic drug-drug interactions in the pilot knowledge base.

2.1 Qualitative Pharmacokinetic Modeling of Drugs

Previously, in order to better understand the issues of formally representing DDI knowledge we constructed a First Order Logic model of the mechanisms underlying DDIs from the lectures and class notes of a graduate class on drug-interactions.[a] Several categories of DDIs were covered in the class including DDIs involving changes to liver or kidney function, gastro-intestinal motility and absorption, transport protein function, and metabolism. We selected for further experimentation rules from this representation that model the jointly sufficient conditions for DDIs that occur via metabolic inhibition or induction. These rules were interesting because a large number of DDIs can be explained by metabolic mechanisms, especially for drugs metabolized by the Cytochrome-P450 (CYP450) enzymes, and considerable research data exists on the metabolic properties of many drugs.

We then constructed a database containing the necessary drug facts for inference with

[a]We have previously published the body of work that this section describes in a conference paper titled *Qualitative Pharmacokinetic Modeling of Drugs* [36].

Table 2.1: The mapping between strength of inhibition in Reference A and Reference B that we used when collecting drug facts for the pilot knowledge base (KB).

Mechanism	*Drug KB*	*Reference A [81]*	*Reference B [44]*
Inhibition	weak	weak	weak, very weak
	moderate	n/a	moderate
	strong	strong	n/a
Induction	weak	n/a	weak
	moderate	n/a	moderate
	strong	strong	n/a

the selected rules. Facts on the important metabolic enzymes for 249 currently prescribed drugs were input into the knowledge base (KB) from a widely used pocket reference on clinically significant drug interactions [81]. This reference (Reference A) also included facts on each drug's potential for inhibition or induction of CYP450 enzymes. We then augmented the KB with information from a Continuing Education Module containing pharmacokinetic information on drugs commonly prescribed to elderly epileptic patients [44]. In addition to facts on potential CYP450 modulation, this reference (Reference B) listed the relative importance of each drug's clearance enzymes. Several drugs not found in Reference A were also added. Since terms regarding the strength of enzyme inhibition and induction varied between the resources, we constructed the mapping shown in Table 2.1. When completed, the KB contained facts useful for mechanism based inference for 267 drugs.

We implemented both the rules and the database in Lisp. The implementation uses a simple pattern matching and backward chaining program taken verbatim from chapter 15 of Paul Graham's popular Common Lisp book [72]. Graham's code uses a Prolog-like syntax, where the macro "`<-`" is analogous to the Prolog connector "`:-`", but as usual in Lisp, prefix notation is used. So, the list expressions have the macro "`<-`" followed by a head expression and optionally a tail expression. Rules that have multiple terms in the tail use combinations of the operators `and`, `or`, and `not` to combine them. Figure 2.1 shows the rules that we developed pertaining to inhibition of clearance; we developed a similar set of rules for metabolic induction.

Since our system could provide no quantitative estimates of its DDI predictions we mod-

```
(with-answer
    (metabolic-induce-interact ?precip ?object ?enz)
    (format t "Drug ~A induces ~A, a primary clearance enzyme, of drug ~A~%"
            (generic-name ?precip) ?enz (generic-name ?object)))
(with-answer
    (metabolic-inhibit-interact ?precip ?object ?enz)
    (format t "Drug ~A inhibits ~A, a primary clearance enzyme, of drug ~A~%"
            (generic-name ?precip) ?enz (generic-name ?object)))
```

Figure 2.2: Queries for any drugs that inhibit or induce the primary clearance enzyme for another NTI drug whose clearance is primarily metabolism

ified the rules to apply only to drugs with a *narrow-therapeutic index*, meaning that there is a small gap between the toxic dose of a drug and the dose at which the drug is ineffective. Any interaction involving a narrow-therapeutic index (NTI) drug could potentially result in harm to a patient, so this change made it more likely that predictions would be clinically relevant. We then applied the selected rules by performing two queries against the database (Figure 2.2) for any drugs that inhibit or induce the primary clearance enzyme of another NTI drug whose clearance is primarily metabolism. The queries returned a total of 90 predicted DDIs out of 71,022 possible pairwise combinations. We then checked the 90 predictions against four online drug reference databases.[b] A predicted drug-drug interaction was considered clinically viable if it was reported in any of the four sources.

Seventy-four of our ninety predicted DDIs were found in at least one drug reference while sixteen could not be found in any online reference (see Table 2.2). We recognized that the sixteen predicted interactions not found in any drug reference were not necessarily false predictions. It is not possible to test every possible drug combination in a clinical trial and the effects of drug interactions can be very hard to recognize so that some drug interactions escape notice in the scientific literature until years after a drug comes to market. The pilot system's predictions were based on pharmacokinetic principles that are considered valid indicators of potential interactions in FDA guidelines [26]. The clinical relevance of these

[b]The four online drug reference databases were 1) First Data Bank's Micromedex, 2) WebMD's Medscape, 3) Discovery health.discovery.com, and 4) Cerner Multum's Drugs.com.

Table 2.2: Predicted inhibition interactions from our pilot drug-drug interaction knowledge-base not documented in any of four online references. In each row, the precipitant drug inhibits the primary clearance enzyme (PCE) of the object drug based on information in our pilot system.

Precipitant	*PCE*	*Object*
amiodarone	CYP2C9	phenobarbital
disulfiram	CYP2C9	phenobarbital
fluorouracil	CYP2C9	phenobarbital
fluconazole	CYP2C9	phenobarbital
gemfibrozil	CYP3A4	carbamazepine
gemfibrozil	CYP2C9	phenobarbital
gemfibrozil	CYP2C9	phenytoin
leflunomide	CYP2C9	phenobarbital
miconazole	CYP3A4	carbamazepine
sulfamethizole	CYP2C9	phenobarbital
sulfamethoxazole	CYP2C9	phenobarbital
sulfinpyrazone	CYP2C9	phenobarbital
sulphaphenazole	CYP2C9	phenobarbital
zafirlukast	CYP3A4	carbamazepine
zafirlukast	CYP2C9	phenytoin
zafirlukast	CYP2C9	phenobarbital

predictions was based on the assumption that any change in the exposure of a patient to an NTI drug is of clinical interest. Thus, it is possible that some of these predictions are valid interactions that have not been studied.

We looked carefully at the evidence behind each fact in our database that supported any of the sixteen novel interaction predictions and found that several facts in the drug KB had varying degrees of support from the scientific literature. For example, the drug product label for zafirlukast notes that *in vitro* experiments have found zafirlukast [9] to be an inhibitor of the CYP3A4 enzyme. Unfortunately, this evidence leaves unanswered the question of whether zafirlukast will effect an clinically relevant, *in vivo*, interaction with drugs that are primarily metabolized by CYP3A4 because, even with very solid *in vitro* evidence, a pharmacokinetic drug property might not have much clinical relevance at the doses in which drugs are prescribed [92, 114]. In contrast, there is stronger evidence in the form of a clinical trial [33] that fluconazole will effect a measurable *in vivo* drug-drug interaction with drugs metabolized by CYP2C9.

Examination of the evidence for other drug properties in this pilot system revealed that

important drug-mechanism knowledge is sometimes missing. For example, all 11 of the interactions involving phenobarbital in Table 2.2 are predicted to occur by inhibition of phenobarbital's primary metabolic clearance pathway which we listed as CYP2C9 based on one source [44]. We could only indirectly support the hypothesis that phenobarbital is a CYP2C9 substrate with two studies from the early 1980s that identified an apparent metabolic interaction between sodium valproate and phenobarbital [104, 105] and an *in vitro* study conducted years later showing that sodium valproate is a CYP2C9 inhibitor [176]. This pattern of inference is weak since it assumes that the interaction could only occur by means of CYP2C9 inhibition while pharmacology research has exposed other means by which apparent metabolic interactions can occur including inhibition of transport proteins. A the time we conducted our analysis we could find no studies, such as an *in vitro* assay, designed to examine directly whether phenobarbital is metabolized by CYP2C9. Without this missing information, the validity of the pilot system's DDI predictions involving phenobarbitol remained considerably uncertain.[c]

2.2 Major Challenges and Related Work

The initial experiment described in Section 2.1 helped identify three major challenges to representing drug-mechanism knowledge. First, there is often considerable uncertainty behind claims about a drug's properties and this uncertainty affects the confidence that someone knowledgeable about drugs places on mechanism-based DDI predictions. Another challenge is that mechanism knowledge is sometimes missing; a fact that can make it difficult to assess the validity of some claims about a drug's mechanisms. Finally, mechanism knowledge is dynamic and any repository for drug-mechanism knowledge is faced with the non-trivial task of staying up to date with science's rapid advances. This section reviews related work that is relevant to overcoming each of these challenges.

---

[c]The work we are describing in this section was conducted in late 2004 and early 2005. A study involving Japanese epileptics [71] published in 2007 provides further evidence that CYP2C9 is a significant pathway for phenobarbitol though, as the investigators note, their results need to be confirmed in other ethnic populations.

2.2.1 Work Addressing Uncertain Mechanism Knowledge

Perhaps the most significant challenge identified in the initial study is that knowledge about drug mechanisms is often *uncertain.* The pilot database had no way to represent uncertainty or determine how much confidence one should have in predictions made using uncertain drug facts. There are many methods to support computational reasoning with uncertain knowledge including symbolic methods such as incidence calculus [39], purely numerical approaches such as Bayesian networks [139] and hybrid approaches such as the Certainty Factors that were attached to rules in MYCIN [52] and similar expert system shells. An interesting informatics research project would be to build a drug-mechanism knowledge system using one of these methods. In fact various models of drug pharmacokinetics and pharmacodynamics employing some of these methods are being used during early drug development for reasons that include assessing the efficacy and safety of new drugs, estimating dose to concentration relationships, and identifying optimal drug candidates [56]. However, the results of the pilot study suggest that, for a drug-mechanism KB to be of clinical use, it is essential that it explicitly link each assertion in the KB to its evidence support. So, a more immediate, and perhaps more important, research question is – what are the strengths and weaknesses of explicitly linking drug-mechanism knowledge to its evidence support and how is that evidence support best modelled?

Other knowledge-based systems link evidence to their drug facts, including DRUGDEX®[d], Q-DIPS [34], and PharmGKB[e] [109], however, there are potential limitations to the methods used by these system. One potential limitation of these systems is that they tend to collect evidence only in support of assertions. This bias towards collecting only supporting evidence could undermine attempts to evaluate how believable an assertion is. Psychological studies have shown that people tend to search for evidence that confirms their hypotheses and that this can sometimes lead them to overestimate the likelihood that a hypothesis is true. When subjects were asked to think of situations where their hypotheses would not be true, their confidence estimations were more accurate [76]. The evidence component

---

[d]http://www.micromedex.com/products/drugdex/

[e]http://www.pharmgkb.org/

of the second drug interaction KB that Section 2.3 discusses accumulates evidence both for and against object property assertions allowing exploration of the possible benefits and drawbacks of this approach.

Another potential limitation of existing systems that link evidence to their drug facts is that they rarely or never provide their criteria for selecting or excluding evidence (inclusion criteria). Research methodology can influence a study's ability of overcome biases and can weaken the validity of its results. Inclusion criteria help ensure that all evidence within a collection meet some basic quality standards. More effort is required to evaluate the evidence support for drug-mechanism assertions when expert users cannot trust that each item of evidence in the system meets some clearly stated standard for research quality. For example, the evidence selection guidelines for content in the Thomson Micromedex product DRUGDEX® are documented internally [21] but not accessible to DRUGDEX® users. The authors of Q-DIPS, a system designed to help clinicians identify and manage DDIs that occur by metabolic mechanisms, list a set of factors affecting the quality of *in vitro* studies [34] but make no mention of using these factors when selecting evidence for their system. In the former system, the criteria used for selecting evidence is not explicit while, in the latter system, it is unknown if criteria have been rigorously applied to all evidence.

The new drug interaction KB this chapter describes uses an evidence type taxonomy that defines distinct kinds of evidence based on their source and methodology. The system requires a set of explicit inclusion criteria for each evidence type that defines a lower-bounds on the quality of the methods used by instances of that type. Knowledge-base maintainers follow a set of processes for acquiring and evaluating evidence designed to ensure that inclusion criteria are met by every piece of evidence entered into the system. This treatment of evidence should make it possible for expert users to quickly assess the strength of evidence for or against assertions in the knowledge-base.

Even more intriguing is the possibility that experts can prospectively map their confidence in each assertion type to some arrangement of one or more abstract evidence types. Rather than requiring that the expert review specific evidence items, it might then be possible for a knowledge-based system to automatically determine the users confidence in its assertions using evidence meta-data. Section 2.3 details how the new drug interaction KB

supports this novel treatment of evidence while Chapter 5 describes an experiment that explores the new method's strengths and limitations using real-world data.

2.2.2 Work Addressing Missing Mechanism Knowledge

Missing drug-mechanism knowledge includes facts that are unavailable or require tests that are impossible or impractical to perform. In the pilot study, the absence of research examining directly whether phenobarbital is metabolized by CYP2C9 left the clinical relevance of 11 of the pilot system's DDI predictions unvalidated. One way to handle missing knowledge when it is important for reasoning is to assume some truth state for the knowledge until proven otherwise. This is a form of *default reasoning* whose various forms include inheritance in semantic networks, circumscription, default logic, and several methods discussed by Goldszmidt and Pearl that utilize qualitative probabilities [69]. Implementing default reasoning in a system that performs logical inference requires that the system be *non-monotonic*. Conceptually, this means that the system can retract or reinstate inferences as the belief state of assertions change. One type of non-monotonic logic system is a Justification-based Truth Maintenance System (JTMS) [61]. Typically, a JTMS system works in conjunction with a rule engine to manage assumptions and their effects on inference. Section 2.3 present a novel use of a JTMS in the second drug interaction KB; the remainder of this section describes how a JTMS works.

Many rule engines, including the pilot knowledge-base described in the beginning of this chapter, model theories as `IF-THEN` rules consisting of one or more clauses forming an *antecedent* and zero or one clauses forming a *consequent*. The *antecedent*, or `IF` portion of an `IF-THEN` rule, must be true for the *consequent*, or the `THEN` portion of the rule, to be true. This is not the case in a system using a JTMS; rather, a consequent can depend on other clauses in addition to the ones in its antecedent. The set of clauses a consequent depends on is called its *justifications*. In order for a consequent to hold true, all of its justifications must hold true.

The JTMS represents every clause in the rule engine as a node possessing a *label* reflecting its current belief state. Every rule in the rule engine specifies a set of justifications

its consequent depends on for belief. The JTMS labels a consequent `IN` when all of its justifications are `IN`. If any of a consequent's justifications are `OUT`, then the JTMS labels the consequent is `OUT`. Justifications can include clauses or *assumptions*. *Assumptions* are clauses that can be `IN` or `OUT` by assignment; they do not require any supporting justifications. The JTMS labels an assumption node `IN`, or *enabled*, when the rule engine *assumes* belief in it, and `OUT` when the rule engine *retracts* that belief. In this way, the JTMS maintains a dependency network of clauses and justifications. A change in belief in any clause or assumption node recursively propagates through the dependency network, changing the belief state of any other node that contains the changed node in its set of justifications.

2.2.3 Work Addressing the Dynamic Nature of Mechanism Knowledge

Any repository for drug-mechanism knowledge is faced with the non-trivial task of staying up to date with science's rapid advance. Unfortunately, one of the most widely used sources of drug-mechanism knowledge, the drug product label, often fails to stay up to date with emerging drug mechanism knowledge. For example, since the late 1990s regulatory agencies have recommended both *in vitro* and *in vivo* investigations into the pharmacokinetic, and especially metabolic, mechanisms a new drug during its early stages of development. However, labeling for older drugs is often missing this emerging knowledge. For example, a study in 1999 found that 10% of the drugs approved between 1992 and 1997 did not include findings from existing *in vitro* metabolic studies [183]. Others have noted that very few labels for drugs approved in the early 1980s provide pharmacokinetic information such as mechanisms of hepatic elimination and the percentage of drug eliminated by renal excretion [116].

A more effective approach might be to track and evaluate both drug label and primary research evidence using editorial boards consisting of domain experts. DRUGDEX®, Facts and Comparisons®, and other comparable systems, use some form of an editorial board approach to stay current with knowledge from clinical trials or case reports [18, 21] and they have scaled the approach to the thousands of drug products listed in these sources. Q-DIPS [34], though possibly no longer an active project, demonstrated that the editorial

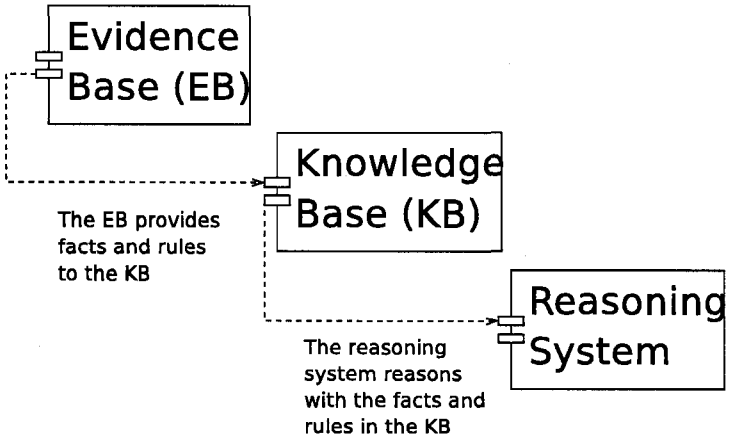

Figure 2.3: An *evidential* knowledge-based system that links assertions about object properties to the evidence for and against those properties. The system's reasoning system makes inferences with assertions in the knowledge-base whose evidence meets belief criteria defined using evidence meta-data.

board approach is feasible for drug-mechanism knowledge. The maintainers of the Q-DIPS system curated a database of *in vitro* and *in vivo* studies supporting assertions about the enzymes a drug is a substrate of or modulates. Users of Q-DIPS could identify DDIs by viewing tables showing the metabolic properties of the set of drugs they are interested in. The tables in Q-DIPS were *dynamic* meaning that their content changed as knowledge about each drug's metabolic profile is updated in the study database.

2.3 Modeling with Evidence and Truth Maintenance

Based on the results of the pilot study and the review of related work we designed a new *evidential* DDI knowledge-base called the *Drug Interaction Knowledge Base*, or DIKB. Figure 2.3 shows an architectural model of the system. The DIKB enables knowledge curators to link each assertion about a drug property to both supporting and refuting evidence. DIKB maintainers place evidence for, or against, each assertion about a drug's mechanistic properties in an *evidence-base* that is kept current through an editorial board approach. Maintainers attach meta-data describing the source and study type of each piece of evidence in the *evidence-base* and users of the system can define specific belief criteria for each

assertion in the *evidence-base* using combinations of the evidence meta-data. The system has a separate *knowledge-base* that contains only those assertions in the *evidence-base* that meet belief criteria. The DIKB's reasoning system uses assertions in this *knowledge-base* and so only makes DDI predictions using those facts considered current by the system's maintainers and believable by users.

Section 2.2.2 introduces a type of non-monotonic logic system called a Justification-based Truth Maintenance System (JTMS) [61] that works in conjunction with a rule engine to manage assumptions and their effects on inference. The DIKB uses a JTMS to handle both default reasoning and the effects on inference of changes in the *knowledge-base* as new evidence causes assertions in the *evidence-base* to meet, or fail to meet, belief criteria. This latter application appears to be a novel use of a JTMS within the field of biomedical informatics.

The DIKB implements the three modules shown in the Figure 2.3 using two software components; one called the `ddi-theory` and the other called the `evidence-model`. The `evidence-model` implements the *evidence-base* component of the model in Figure 2.3. It models the evidence for and against each assertion in the *knowledge-base* and communicates to the `ddi-theory` which assertions it can use for inference. The `ddi-theory` implements both the *knowledge-base* and the *reasoning system* components of the model in Figure 2.3. It consists of a JTMS, an inference engine, and a novel rule-based DDI prediction theory.

An explicit function in the *reasoning system* executes a forward chaining inference algorithm that applies the rule-based DDI theory to assertions in the *knowledge-base*. Any new assertions that result from inference are added to the *knowledge-base*. Users can pose queries against the *knowledge-base* and the system will return any assertions about drugs, including drug-drug interaction predictions, that match a user's query. It will also return links to the evidence for and against each assertion used to satisfy the query. The next two sections set forth in greater detail each component of this system.

2.3.1 The evidence-model

The purpose of the `evidence-model` component of the DIKB is to manage evidence for and against assertions about the attributes of objects that are relevant for predicting DDIs (e.g. active ingredients, metabolites, and enzymes). The `evidence-model` stores instances of frame-based representations of these objects and communicates the current state of belief for their attributes to the `ddi-theory`. The set of assertions in `ddi-theory` and their belief state changes as the `evidence-model` accumulates evidence for and against object property assertions. To satisfy its purpose, the `evidence-model`:

1. stores evidence and evidence meta-data for and against each object attribute (Section 2.3.1.1)

2. tests the evidence for and against each object attribute against user-defined belief criteria (Section 2.3.1.2)

3. exports statements that tell the `ddi-theory` to add assertion nodes to the *knowledge-base* and change their belief state (Section 2.3.2.2)

2.3.1.1 Storing Evidence and Evidence Meta-data

The `evidence-model` represents objects of interest to the `ddi-theory` and assertions about their attributes as instances of classes derived from an abstract class called `Frame`. A simple class, called `KB`, performs storage and retrieval functions for these class instances. This class has two sub-classes, or children; `DrugKB` for objects whose properties are important for inference and `EvidenceBase` for assertions about the properties of these objects.

The singleton `DrugKB` can contain instances of type `Drug` (Table 2.4), `Metabolite` (Table 2.5), and `Enzyme` (Table 2.7), and `Pharmaceutical_Preparation` (Table 2.6). These classes contains two types of slots, *categorical* slots that store plainly factual knowledge such as a drug's generic and trade names and *evidence* slots that model knowledge that rests on conclusions from research.

Table 2.3: The `Pceut_Entity` class description – a frame-based model for an abstract pharmaceutical entity. The classes that model active ingredients (Table 2.4) and metabolites (Table 2.5) inherit all of the slots of this class. All slots are of type `Evidence` (see Table 2.3.1)

Slot	*Description*
`substrate-of`	enzymes that metabolize the entity
`is-not-substrate-of`	an incomplete list of enzymes that *do not* catalyze the entity
`in-vitro-probe-substrate-of-enz`	enzymes that meet the FDA's definition of a preferred or acceptable chemical substrate for *in vitro* studies with the entity [26]
`primary-total-clearance-enzyme`	an enzyme responsible for 50% of the entity's total clearance from the body (if one exists)
`primary-total-clearance-mechanism`	the entity's primary route clearance: metabolic, renal, biliary, or exhalation
`primary-metabolic-clearance-enz`	an enzyme responsible for 50% of the entity's total *metabolic* clearance from the body (if one exists).
`inhibits`	an incomplete list of enzymes this entity inhibits
`does-not-inhibit`	an incomplete list of enzymes this entity *does not* inhibit
`in-viVo-selective-inhibitor-of-enz`	enzymes for which the entity meets the FDA's definition of a preferred or acceptable chemical substrate for *in vivo* studies [26]
`in-vitro-selective-inhibitor-of-enz`	enzymes for which the entity meets the FDA's definition of a preferred or acceptable chemical substrate for *in vitro* studies [26]
`permanently-deactivates-catalytic-function`	enzymes the entity inhibits in such a way that they are no longer available for catalysis
`does-not-permanently-deactivate-catalytic-function`	enzymes for which the entity is a competitive inhibitor
`has-metabolite`	an incomplete list of biochemical entities that the entity is transformed to via catalysis
`pceut-entity-of-concern`	true or false depending on whether a small change in the systemic concentration of the entity would be of clincial interest
`sole-PK-effect-alter-metabolic-clearance`	asserts that the entity's sole pharmacokinetic effect on another entity is alteration of its metabolic clearance
`maximum_concentration`	the observed *in vivo* maximum concentration, C_{max}, in grams/L of the entity at various doses
`inhibition-constant`	rate constant(s) in grams/L for enzymes the entity has been shown to inhibit in `in vitro` studies
`increases-auc`	the set of active ingredients or metabolites for which this entity, when co-administered, causes an increase in AUC.

Table 2.4: The `Drug` class description – a frame-based model for an abstract active ingredient. Because the `Drug` class is a child-class of the `Pceut_entity` class it inherits all of the slots in Table 2.3

Slot	*Type*	*Description*
`active-ingredient-name`	categorical	an active ingredient name from the VA-NDF-RT vocabulary [37]
`prodrug`	categorical	is this drug a prodrug? true or false
`bioavailability`	evidence	the percentage of the drug available for systemic distribution by formulation and dose. Each evidence item refers to the dose and formulation of the drug that is associated with the bioavailability value.
`first-pass-effect`	evidence	the proportion of drug that is cleared by first-pass metabolism
`fraction-absorbed`	evidence	the percentage of drug that is absorbed in the gastro-intestinal tract
`fraction-cleared-by`	evidence	the fraction of the active ingredient's dose that is cleared by various enzymes

Table 2.5: The `Metabolite` class description – a frame-based model for metabolite objects. Because the `Metabolite` class is a child-class of the `Pceut_entity` class it inherits all of the slots in Table 2.3

Slot	*Type*	*Description*
`metabolite-name`	categorical	an name for this metabolite that links it to other data in the NCBI's PubChem database
`metabolite`	categorical	(always `True`) maintains that an instance of this class models a metabolite

Table 2.6: The `Pharmaceutical_Preparation` class description – a frame-based model for abstract pharmaceutical preparations. This class currently possesses only *categorical* slots; slots that store knowledge that is plainly factual about drugs such as its generic and trade names.

Slot	*Type*	*Description*
`prep-name`	categorical	the name of the preparation from the VA-NDF-RT terminology [37]
`form`	categorical	route of administration - oral, transdermal, or IV
`dose`	categorical	the dose the drug is given
`preparation`	categorical	states if the entity is a normal or extended release formulation
`ingredients`	categorical	a list of active pharmaceutical ingredients

Table 2.7: The **Enzyme** class description – a frame-based model for objects an abstract enzyme entity

Slot	*Type*	*Description*
`enzyme-name`	categorical	the symbol for the enzyme in the HUGO Gene Nomenclature Committee (HGNC) database
`polymorphic-enzyme`	evidence	True if this enzyme has multiple drug-catalysis phenotypes due to genetic polymorphisms
`controls-formation-of`	evidence	a list of biochemical entities that this enzyme catalyzes

The singleton `EvidenceBase` stores instances of class `Assertion` (Table 2.8); a class that models the evidence both for and against an attribute of some instance in `DrugKB`. When users find evidence for or against the property represented by an *evidence* slot they create a new instance of the `Evidence` class shown in Table 2.9 and enter values for its slots. These instances are then placed in either the `evidence-for` or `evidence-against` slot of the `Assertion` instance associated with the property's *evidence* slot.

Table 2.8: Slots in class `Assertion`

Slot	*Description*
`object`	the object's name in `DrugKB`
`slot`	name of the slot
`value`	an allowable value for this slot
`evidence-for`	a list of `Evidence` types
`evidence-against`	a list of `Evidence` types
`ready-for-classification`	is this assertion ready to classify
`assert-by-default`	true if this assertion should be considered valid by default
`evidence-rating`	the result of testing the evidence for this assertion against user belief criteria; one of `assume!`, `retract!`, `none-assigned`, or `can't decide`
`cont-val`	the discretized value of a continuous value assertion (e.g. `low`, `medium`, `high`); depends on a method for discretizing `numeric-val`
`numeric-val`	the simple numerical value of a continuous value assertion; depends on a method for combining the values of each continuous-valued evidence item
`id`	a unique identifier for the assertion instance

Table 2.9: Slots in class `Evidence`

Slot	*Description*
`evidence-type`	a meta-data label from the DIKB evidence taxonomy
`doc-pointer`	a pointer to the evidence document
`quote`	a short summary of the evidence
`reviewer`	person entering this evidence
`assumptions`	a list of *evidence-use assumptions* – assertions upon which the current use of evidence depends; the evidence instance is not used in establishing the validity of assertions unless all assertions in the assumptions list meet belief criteria
`timestamp`	a timestamp for when evidence item was entered into the system

2.3.1.2 User-defined Criteria for Belief and Disbelief:

The description of pilot work in Section 2.1 relates how the evidence support for the facts in the pilot database were useful for assessing the validity of the system's predictions. For example, it was explained that, based on evidence, the proposition that zafirlukast inhibits CYP3A4 was less justified than the proposition that fluconazole inhibits CYP2C9. This case suggests that the confidence someone knowledgeable about drugs has in the clinical validity of a DDI prediction can vary depending on the *type* of evidence that supports or refutes each of the facts leading to the prediction. To explore this idea further, the DIKB supports using evidence types to track the level of certainty users have in the system's drug-mechanism assertions.

The types of evidence that can support drug-mechanism facts include, among others, labeling statements, results from *in vitro* studies, expert interpretation of case reports, and various pharmacokinetic trials involving volunteers. A novel feature of the DIKB is that expert users can define combinations of evidence that they believe lend different degrees of certainty to the assertion types that the DIKB uses to predict DDIs. Different combinations of evidence types might confer different levels of certainty in an assertion and these can be rank ordered to produce "levels of evidence" (LOEs).

The DIKB distinguishes between assertion *instances* and assertion *types*. An assertion

instance is a specific fact about a particular object such as a drug or protein. For example, the assertion `(substrate-of 'CARB 'CYP3A4)` is an instance of the generic `(substrate-of X Y)` assertion type. DIKB users define one or more LOEs for each generic assertion type by creating logical statements listing the level's required evidence types and their multiplicity. The LOEs for an assertion type apply to any instance of that type. Users can also place evidence types that they feel have similar levels of validity into a group called a *ranking category*. They can then use the ranking category just like other evidence types to define LOEs.

For every assertion type users select one LOE as the *belief criteria*. The `evidence-model` will tell the `ddi-theory` to use a particular assertion instance in inference if, and only if, there exists a body of evidence *for* the assertion that satisfies the *belief criteria* for the assertion's type and the evidence *against* the object property does not satisfy *belief criteria*. The DIKB allows the *belief criteria* for evidence supporting an assertion type to be different from the *belief criteria* for evidence refuting an assertion type. Table G show the evidence-types and ranking categories used while developing the DIKB while Figure 2.4 shows test LOEs. Chapter 4 presents a more rigorous evidence ontology that was used for labeling evidence in the final DIKB evidence-base.

LOE-1 ::=	**RCT**+ |
	FDA Guidances+
LOE-2 ::=	LOE-1 |
	Drug Labeling+
LOE-3 ::=	LOE-2 |
	Drug Labeling+ |
	(**in vitro**+ and **Non-random**+)
LOE-4 ::=	LOE-2 |
	in vitro+ |
	Non-random+

Figure 2.4: A set of levels of evidence (LOE) used while developing the DIKB. The symbol '::=' means the term to the left "is defined as" the term to the right, | means "or", and '+' means that "one or more occurrences of" of the symbol to the left are allowed. So item one reads "LOE-1 is defined as one or more **RCT** OR one or more **FDA Guidance** evidence types."

Table 2.10: Every evidence item entered into the DIKB receives a label from a taxonomy of evidence types. This table shows a *test* evidence type taxonomy used for designing the DIKB. Chapter 4 presents a mature evidence ontology used for labeling evidence in the final DIKB evidence-base.

Ranking Category	*Evidence Type*
• **RCT**	• a randomized, controlled, clinical trial
• **Non-random**	• a cohort study • a case-control study • a non-randomized trial with concurrent or historical controls • a retrospective study looking a clinical records over time • a fixed order study with non-randomized healthy volunteers
• **Case Reports**	• a single case report • a case series
• **FDA Guidances**	• a statement in an FDA guidance to industry
• ***in vitro***	• *in vitro* evidence from human tissue, microsomal • *in vitro* evidence from human tissue, recombinant
• **Drug Labeling**	• *in vitro* or *in vivo* information found in drug product labeling that provides no citation of its source

2.3.2 *The* `ddi-theory`

It was mentioned earlier in this chapter that the `ddi-theory` implements both the *knowledge-base* and the *reasoning system* components of the model in Figure 2.3. The *reasoning system* component of the `ddi-theory` consists of two parts – a rule engine and a JTMS that maintains the belief state of clauses in the rule engine. Section 2.2.2 relates how a JTMS works; let's now examine how the `ddi-theory` uses the JTMS to handle both default reasoning and the effects on inference of changes in the *knowledge-base* as new evidence causes assertions in the *evidence-base* to meet, or fail to meet, belief criteria.

Figure 2.5 shows an example inference rule applicable when a precipitant drug inhibits the metabolic clearance of an object drug.[f] The first line declares that this is a `rule`, the next line specifies a pattern for when one object inhibits another. The `:IN` before the pattern declares that this antecedent must be believed in order to evaluate as true. The consequent in Figure 2.5 says to assert that `?x` inhibits the metabolic clearance of `?z` by `?y` when the antecedents evaluate true. Then follows a list containing a series of justifications for the consequent. The justifications represent clauses or assumptions that must be `IN` in order for the consequent to be `IN`. When the rule engine makes an assertion, the JTMS creates a node for it in the *knowledge-base* and then looks to see if the consequent's justifications are `IN`; if so, the JTMS labels the node `IN`.

2.3.2.1 *Default Knowledge in the DIKB*

The `ddi-theory` models default knowledge as JTMS assumptions. Belief in the truth of default information causes the assumption representing that information to be *enabled.* Assumption nodes receive an `:IN` label when they are enabled. The system can later *retract* that belief causing the nodes to receive an `:OUT` label. A change in belief in any assumption node recursively propagates through the other nodes that contain the changed node in its set of justifications changing their belief state.

For example, the test rule in Figure 2.5 requires the default assumption that all precip-

[f]Please note that the rules shown in this chapter were used during the development of the DIKB. A more sophisticated rule-based theory of metabolic inhibition was used for experiments that tested the accuracy and coverage of mechanisms-based prediction (See Chapter 3).

```
(rule
  ((:IN (inhibits ?x ?y))
   (:IN (substrate-of ?z ?y)))
  (rassert!
   (inhibit-metabolic-clearance ?x ?z ?y)
   (
    (inhibits ?x ?y)
    (substrate-of ?z ?y)
    (inhibitory-concentration ?x ?y)
  )))
```

Figure 2.5: An example inference rule for when a precipitant inhibits the metabolic clearance of an object drug

itants reach a concentration sufficient to cause a clinically significant effect on the enzymes they are known to inhibit. The following listing tells the `ddi-theory` to create an assertion that carbamazepine (CARB) is a substrate-of the drug metabolizing enzyme Cytochrome P-450 3A4 (CYP3A4).

```
(assert!
  '(substrate-of 'CARB 'CYP3A4))
```

If the system also asserts that clarithromycin (CMYN) inhibits CYP3A4, it would need to create an enabled assumption declaring CMYN to be at a concentration sufficient to inhibit CYP3A4.

```
(assert!
  '(inhibits 'CMYN 'CYP3A4))

(assume!
  '(inhibitory-concentration 'CMYN 'CYP3A4))
```

Figure 2.6 shows how the JTMS dependency network would look at this point in the example. If further data causes the belief state of our default assumption to change to false, then the program can retract belief:

```
(retract!
 '(inhibitory-concentration 'CMYN 'CYP3A4))
```

The effect of changing this belief is shown in Figure 2.7. The JTMS changes the node labels for both the assumption `(inhibitory-concentration 'CMYN 'CYP3A4)` and the assertion `(inhibit-metabolic-clearance 'CMYN 'CARB 'CYP3A4)` to `:OUT` meaning that this assertion is no longer believed true. It is important to note that *any other assertions or inferences* that depend directly, or indirectly, on either of these assertions will now also be labeled `OUT`.

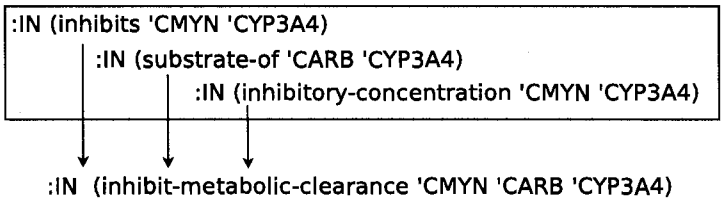

Figure 2.6: A small JTMS dependency network; justifications are shown in the box

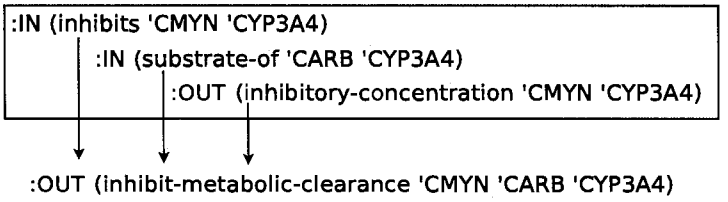

Figure 2.7: A change in the belief state in one of the justifications propagates to dependant consequents

2.3.2.2 The `evidence-model` Uses Assumptions to Affect the `ddi-theory`

The previous section describes how the DIKB uses JTMS assumptions to handle default reasoning. The system also uses assumptions to effect changes in the *knowledge-base* as new evidence causes assertions in the *evidence-base* to meet, or fail to meet, belief criteria. The `evidence-model` accomplishes this by adding a special justification, called a "`bc-satisfied` assumption", to the list of justifications belonging to an assumption. The `evidence-model` tells the `ddi-theory` to enable a `bc-satisfied` assumption when the evidence support for an assertion meets belief criteria and to retract the same enabled `bc-satisfied` assumption when the evidence support for the assertion no longer meets belief criteria.

For example, the following listing represents an assertion The assertion declaring that clarithromycin inhibits CYP3A4 is similar in form to assertions that the `evidence-model` would send to the `ddi-theory`.

```
(assert!
  '(inhibits 'CMYN 'CYP3A4)
       '((bc-satisfied 'assertion_40)))
```

The JTMS component of the `ddi-theory` would create a node for this assertion in the *knowledge-base*. Notice the `bc-satisfied` assumption in the previous listing. The `(bc-satisfied 'assertion_40)` assumption must be `:IN` in order for the belief state of the `inhibits` assertion to be `:IN`. When the evidence support for this particular `inhibits` assertion meets belief criteria the `evidence-model` could tell the `ddi-theory` to enable the `bc-satisfied` assumption that is in the assertion's list of justifications as follows:

```
(assume!
  '((bc-satisfied 'assertion_40)))
```

This change causes the `(inhibits 'CMYN 'CYP3A4)` assertion to also receive an `:IN` label because the assumption `(bc-satisfied 'assertion_40)` is its only justification.

All assertions in the DIKB that depend on evidence for justification require one or more `bc-satisfied` statements in their list of justifications. Extending the example from the

previous section, the following listing shows the information the `evidence-model` would export to the `ddi-theory` when the evidence support for the assertion `(substrate-of 'CARB 'CYP3A4)` meets belief criteria.

```
(assert!
  '(substrate-of 'CARB 'CYP3A4)
      '((bc-satisfied 'assertion_10)))

(assume!
  '((bc-satisfied 'assertion_10)))
```

Figure 2.8 shows how the JTMS dependency network would look at this point in our example. Recall that, in this example, the system assumes by default that enzyme inhibitors are at sufficient concentration to affect metabolism and so automatically enables the `inhibitory-concentration` assumption.

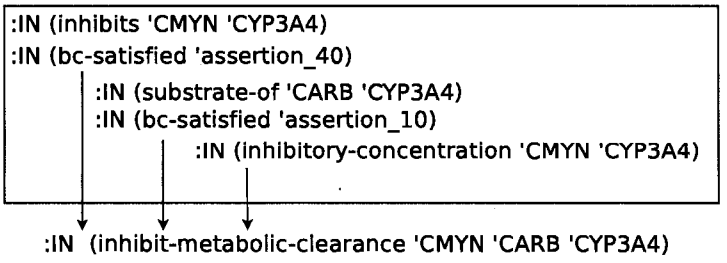

Figure 2.8: A JTMS dependency network showing "`bc-satisfied` assumptions" in the set of justifications shown inside of the box. The `evidence-model` uses `bc-satisfied` assumptions to affect the belief state of assertions as their evidence support meets, or fails to meet, belief criteria.

There are two situations where the `evidence-model` will re-assess the evidence for an assertion about one of its objects – when an assertion's evidence support changes or `DIKB` users change the LOE that they have selected as the belief criteria for the assertion's type. In both cases, the `evidence-model` compares the evidence for and against the assertion.

If the evidence *for* the assertion satisfies the belief criteria currently assigned to the assertion's type, and the evidence *against* the assertion does not satisfy belief criteria, then the `evidence-model` will cause the assertion's `bc-satisfied` justification to be enabled (labeled `:IN`) in the `ddi-theory`. The `evidence-model` will retract (label `:OUT`) the assertion's `bc-satisfied` justification when either 1) evidence against the assertion meets its belief criteria, 2) the belief criteria changes and the evidence for an assertion is no longer sufficient, or 3) the system calls into question the use of an evidence item as support for the assertion (see Section 2.3.2.3). The `evidence-model` keeps track of state so that, if the `evidence-model` has already triggered an assertion or placed an assumption in a desired state, then the system will make no change.

2.3.2.3 Representing Conjectures can Help Maintain the System's Knowledge

Interpreting the results of a scientific investigation as support for a particular assertion can sometimes require making conjectures that scientific advance might later prove to be invalid. If such conjectures are later shown to be false, it is important to re-consider how much support the scientific investigation lends to any assertion it was once thought to support. One unique feature of the DIKB is that it can represent the conjectures behind a specific application of evidence. These representations are called *evidence-use assumptions* and they facilitate keeping knowledge in the system current.

For example, let's say that a pharmacokinetic study involving healthy patients finds a significant increase in the systemic concentration of simvastatin in the presence of diltiazem. If the study meets inclusion criteria, and it is thought that that diltiazem is a selective inhibitor of the CYP3A4 enzyme in humans, then an evidence-base curator might apply this evidence as support for the assertion `(substrate-of 'simvastatin 'CYP3A4)`. This use of the diltiazem-simvastatin study as supporting evidence for the assertion `(substrate-of 'simvastatin 'CYP3A4)` depends on the conjecture that diltiazem is an *in vivo* selective inhibitor of the CYP3A4. The curator should reconsider this use of evidence if future work reveals that diltiazem increases patient exposure to simvastatin by some other mechanism than reducing CYP3A4's catalytic function (e.g. transport protein modulation).

Unlike systems that just cite evidence, the DIKB's formal model of evidence enables it to flag when a conjecture has become invalid and alert knowledge-base curators to the need to reassess their original interpretation of what assertions a piece of evidence supports. Currently, DIKB curators make an evidence-use assumption known to the DIKB by first identifying the label of an assertion in the `evidence-model` that represents the evidence-use assumption. They then add the label to the `assumptions` list in the instance of the `Evidence` class used to represent the evidence item.

Continuing the current example, Table 2.11 shows what an `Evidence` instance would look like if a curator identified the assertion `(in-vivo-selective-inhibitor 'diltiazem 'CYP3A4)` as an evidence-use assumption for the use of the evidence item as support for the assertion `(substrate-of 'simvastatin 'CYP3A4)`.

Table 2.11: This table shows relevant slots from a hypothetical instance of class `Evidence` that possesses an *evidence-use assumption* representing the conjecture `(in-vivo-selective-inhibitor 'diltiazem 'CYP3A4)`. If the system places this instance in the `evidence-for` list for some assertion the system would cause the belief state of the assertion to depend on the belief state of the evidence-use assumption.

Slot	*Description*
`doc-pointer`	PubMed ID: 10741630
`evidence-type`	`Non-random`
`assumptions`	`'(in-vivo-selective-inhibitor 'diltiazem 'CYP3A4)`
`quote`	"Diltiazem significantly increased the mean peak serum concentration of simvastatin by 3.6-fold ($P < .05$) and simvastatin acid by 3.7-fold ($P < .05$)" [125]

JTMS assumption nodes provide an elegant method for notifying the `ddi-theory` of evidence-use assumptions and changes in their belief state. When the `evidence-model` exports an assertion to the `ddi-theory` it checks the `assumptions` list of each `Evidence` instance used to meet belief criteria. Each assertion in the `assumptions` list is added to the set of justifications for the assertion being exported. Extending the present example, assume that the evidence item in Table 2.11 is the only item in the `evidence-for` list for the assertion `(substrate-of 'simvastatin 'CYP3A4)`. Then, `evidence-model` would sent an assertion statement similar to the one below to the `ddi-theory`:

```
(assert!
  '(substrate-of 'simvastatin 'CYP3A4)
     '((bc-satisfied 'assertion_id_X)
       (in-vivo-selective-inhibitor 'diltiazem 'CYP3A4)))
```

The assertion `(substrate-of 'simvastatin 'CYP3A4)` cannot be used by the *reasoning system* for inference unless its evidence meets belief criteria and the belief state of the evidence-use assumption held by its supporting evidence also meets belief criteria.

A more complicated scenario occurs when an `evidence-model` assertion has more than one item in its `evidence-for` list that meets belief criteria and each item requires evidence-use assumptions. In this case, the `evidence-model` will send one assertion instance for each combination of evidence that satisfies belief criteria. The set of justifications belonging to each assertion instance exported to the `ddi-model` will contain the set of evidence-use assumptions belonging to each evidence item used to meet belief criteria.

For example, assume that `(substrate-of 'simvastatin 'CYP3A4)` has the two items of evidence in its `evidence-for` list. The `evidence-for` lists for these two evidence items each have one evidence-use assumption:

```
Item 1:
  assumptions:'(in-vivo-selective-inhibitor 'diltiazem 'cyp3a4)

Item 2:
  assumptions:'(in-vivo-selective-inhibitor 'ketoconazole 'cyp3a4)
```

If both evidence items individually satisfy the belief criteria for this assertion type then, since there are two different ways that `(substrate-of 'simvastatin 'CYP3A4)` can be established by evidence, the `evidence-model` must send two different assertion instance to the `ddi-theory`.

```
(assert!
  '(substrate-of 'simvastatin 'cyp3a4)
```

```
    (bc-satisfied 'substrate-of-simvastatin-cyp3a4)
    (in-vitro-selective-inhibitor 'diltiazem 'cyp3a4))

(assert!
  '(substrate-of 'simvastatin 'cyp3a4)
    (bc-satisfied 'substrate-of-simvastatin-cyp3a4)
    (in-vitro-selective-inhibitor 'ketoconazole 'cyp3a4))
```

If at some point after the `ddi-theory` creates these assertion nodes the evidence for the assertion `(substrate-of 'simvastatin 'CYP3A4)` meets belief criteria and both `in-vitro-selective-inhibitor` assertions are labeled `:IN` then, the system will give both nodes for the assertion `(substrate-of 'simvastatin 'CYP3A4)` an `:IN` label. In this state, the evidence for the assertion, as well as all relevant evidence-use assumptions, meet belief criteria. Each `(substrate-of 'simvastatin 'CYP3A4)` assertion node will maintain its `:IN` label until the evidence for the assertion no longer meets belief criteria or the system retracts one of the `in-vitro-selective-inhibitor` assertions. Put another way, so long as belief criteria and evidence-use assumptions for one `(substrate-of 'simvastatin 'cyp3a4)` assertion are met, then the assertion will be available for use in inference.

2.4 Implementation and Examples

Users can use a Web interface to both enter evidence for drug-mechanism assertions into the `evidence-model` and view previously entered evidence. Both the Web interface and the `evidence-model` are implemented in Python.[g] The latter is implemented as a set of Python classes and shell scripts while the former uses the HTMLGen library[h] for creating Web pages and the Twisted networking frame work[i] for serving them.

The `ddi-theory` uses Forbus and de Kleer's ANSI Common Lisp rule engine (JTRE) and

---

[g]http://www.python.org

[h]http://starship.python.net/crew/friedrich/HTMLgen/html/main.html

[i]http://twistedmatrix.com/

JTMS from [61] with no modification. This implementation was chosen because it is both open source and well-documented. The rules shown Figures 2.5 and 2.9 are enclosed in a Lisp function that initializes globally accessible JTRE and JTMS objects. The `evidence-model` writes asserted and retracted assumptions to a file stored on disk. This file is manually loaded by the user from an interactive Lisp session each time the `evidence-model` re-assesses its evidence.

Figures 2.5 and 2.9 show the rule-based DDI theory used to design the DIKB. The theory represented in these figures is not validated, its only purpose was to test the DIKB during development. Chapter 3 presents a more sophisticated DDI theory that we validated during the experiment described in Chapter 5 of this dissertation.

The rule in Figure 2.5 and the first rule in Figure 2.9 capture inhibition of a clearance enzyme of a drug that is primarily cleared by metabolism. The third and fourth rule serve to capture a disjunctive state when a drug has a narrow-therapeutic index and/or is considered a *sensitive substrate.*[j] These rules are necessary because our JTMS implementation can only accept single literal positives and cannot directly assert disjunctive clauses. The final rule in Figure 2.9 specifies conditions that, if present, greatly increase the likelihood of a clinically significant inhibition interaction.

Output from an example run of the system is shown in Table 2.12. The example run was conducted using the *test* rule-based theory, a subset of the drug properties and evidence in an experimental version of the `evidence-model`, and three of the four levels of evidence shown in Figure 2.4. The example illustrates one advantage of using evidence meta-data to specify belief criteria for assertions in the knowledge-base – the system can provide different views of its knowledge and inferences to users who might not agree about what combination of evidence makes an assertion believable.

Three different levels of evidence were chosen as belief criteria; Table 2.12 shows output of the system at each level. LOE-1 accepts only one or more evidence items from either the **RCT** or **FDA Guidance** categories as evidence. LOE-2 adds to this a very significant

[j]The FDA defines a sensitive substrate as a substrate that exhibits a 5-fold or greater increase in exposure with the addition of an inhibitor. There are currently several drugs on the FDAs published list including buspirone, eletriptan, felodipine, lovastatin, midazolam, sildenafil, simvastatin and triazolam [26]

```
(rule
 ((:IN
   (inhibit-metabolic-clearance ?x ?z ?y)
                :TEST (not (equal ?x ?z))
   ))
 (rassert!
  (increase-drug-exposure ?x ?z ?y)
  (nil
   ;;justifications
   (inhibit-metabolic-clearance ?x ?z ?y)
   (primary-clearance-mechanism
                          ?z 'METABOLISM)
   )))

(rule
 ((:IN (narrow-therapeutic-range ?z)))
 (rassert!
  (nti-or-sensitive-substrate ?z)
  ;;justifications
  (nil
   (narrow-therapeutic-range ?z))
  ))

(rule
 ((:IN (sensitive-substrate ?z)))
 (rassert!
  (nti-or-sensitive-substrate ?z)
  ;;justifications
  (nil
   (sensitive-substrate ?z))
  ))

(rule
 ((:IN (increase-drug-exposure ?x ?z ?y))
  (:IN (primary-clearance-enzyme ?z ?y))
  (:IN (nti-or-sensitive-substrate ?z)))
 (rassert!
  (metabolic-inhibition-interaction ?x ?z ?y)
  ;;justifications
  (nil
   (increase-drug-exposure ?x ?z ?y)
   (primary-clearance-enzyme ?z ?y)
   (nti-or-sensitive-substrate ?z)
   )))
```

Figure 2.9: A *test* DDI theory consisting of these rules plus the one shown in Figure 2.5 was used for developing the DIKB.

source of evidence; **labeling**. As is shown, changing the level of evidence from LOE-1 to LOE-2 has a dramatic effect on the belief state of predicted DDIs. Only one change in predicted DDIs occurs when moving from LOE-2 to LOE-4; the prediction that fluvastatin will inhibit the metabolic clearance of rosuvastatin by some, possibly negligible, amount. This is because the experimental version of the `evidence-model` contained only one item of evidence supporting the claim that fluvastatin inhibits CYP2C9 [60]; an ***in vitro*** type acceptable only at LOE-4.

Another example shows how a JTMS can efficiently handle the effects on inference of changes the *knowledge-base* as new evidence causes assertions in the *evidence-base* to meet, or fail to meet, a user's belief criteria. Table 2.13 shows that the system made the prediction that fluvastatin will inhibit the metabolic clearance of rosuvastatin via CYP2C9 at the LOE-2 level instead of the LOE-4 level when an evidence item that mapped to the **RCT** ranking category was added to the evidence-base. It is important to note that *any other assertions or inferences* that depended directly, or indirectly, on this inference would now also be labeled `IN` provided that all of their other justification are `IN`.

Table 2.12: The belief state of a subset of drug property assertions and inferences in the development version of the DIKB when three different levels of evidence from Figure 2.4 were chosen as belief criteria. Inferences were made using the rules in Figures 2.5 and 2.9. The evidence *for* each drug property is shown; no drug property shown had evidence *against* it.

Drug property assertions

Drug Property	**Evidence for**	***LOE-1***	***LOE-2***	***LOE-4***
(INHIBITS 'DILTIAZEM 'CYP3A4)	RCT,Non-random,*in vitro*	IN	IN	IN
(INHIBITS 'FLUVASTATIN 'CYP2C9)	*in vitro*	OUT	OUT	IN
(SUBSTRATE-OF 'SIMVASTATIN 'CYP3A4)	RCT, Non-random, *in vitro*	IN	IN	IN
(SUBSTRATE-OF 'ROSUVASTATIN 'CYP2C9)	label	OUT	IN	IN
(PRIMARY-CLEARANCE-ENZYME 'SIMVASTATIN 'CYP3A4)	label	OUT	IN	IN
(PRIMARY-CLEARANCE-MECHANISM 'SIMVASTATIN 'METABOLISM)	label	OUT	IN	IN
(SENSITIVE-SUBSTRATE 'SIMVASTATIN 'CYP3A4)	FDA Guidance	IN	IN	IN

Predicted Drug-drug interactions

Drug Property	***LOE-1***	***LOE-2***	***LOE-4***
(INHIBIT-METABOLIC-CLEARANCE 'DILTIAZEM 'SIMVASTATIN 'CYP3A4)	OUT	IN	IN
(INCREASE-DRUG-EXPOSURE 'DILTIAZEM 'SIMVASTATIN 'CYP3A4)	OUT	IN	IN
(METABOLIC-INHIBITION-INTERACTION 'DILTIAZEM 'SIMVASTATIN 'CYP3A4)	OUT	IN	IN
(INHIBIT-METABOLIC-CLEARANCE 'FLUVASTATIN 'ROSUVASTATIN 'CYP2C9)	OUT	OUT	IN
(INCREASE-DRUG-EXPOSURE 'FLUVASTATIN 'ROSUVASTATIN 'CYP2C9)	OUT	OUT	OUT

Table 2.13: A hypothetical example showing how the addition of new evidence to the DIKB can effect a change in the predictions that are considered believable. Here, we show that, if an evidence item of the RCT type supporting (INHIBITS 'FLUVASTATIN 'CYP2C9) were added to the system, one of the system's inferences would be believed at *LOE-2* instead of *LOE-4*. †– items different from Table 2.12.

Drug property assertions

Drug Property	**Evidence for**	***LOE-1***	***LOE-2***	***LOE-4***
(INHIBITS 'FLUVASTATIN 'CYP2C9)	*in vitro* , RCT†	IN	IN†	IN
(SUBSTRATE-OF 'ROSUVASTATIN 'CYP2C9)	label	OUT	IN	IN

Predicted Drug-drug interactions

Drug Property	***LOE-1***	***LOE-2***	***LOE-4***
(INHIBIT-METABOLIC-CLEARANCE 'FLUVASTATIN 'ROSUVASTATIN 'CYP2C9)	OUT	IN†	IN

2.5 *Discussion*

The examples in Section 2.4 demonstrate that important features of the DIKB are functional. Chapter 5 discusses in detail how the DIKB was used to successfully predict drug-drug interactions and non-interactions between 595 drug/drug and drug/drug-metabolite pairs. The results section of that chapter discusses the strengths and limitations of the DIKB's methods based on real-world application. The remainder of this section highlights some observations about the system's design.

2.5.1 The DIKB as an System for Research

It is important to note that the DIKB's reasoning system, like that of the pilot system (Section 2.1), is unable to track uncertainty through inference. Rather, the DIKB automatically selects assertions that meet user defined belief criteria assuming that these assertions are certain from the user's perspective. If the user selects belief criteria that represent full confidence in each assertion type, and each assertion the DIKB uses meets the user's belief criteria, then the system's DDI predictions will also meet belief criteria. This arrangement is useful for researching how evidence can be used to establish the certainty of drug-mechanism knowledge but it does not address how to handle assertions that do not meet belief criteria.

The DIKB does not prevent users from assigning as belief criteria LOEs that do not inspire their full confidence in an assertion. However, the system has no way of establishing the certainty a user should have in a DDI prediction that depends on such assertions. One can imagine scenarios where having knowledge of even uncertain DDI predictions could be valuable. For example, if the perceived risk of death to a patient is high, a clinician might want to be extra cautious while determining a drug therapy and avoid, if possible, every predicted DDI, regardless of the certainty of its occurrence. In such cases, selecting as belief criteria an LOE that does confer complete confidence in an assertion might be justified if it had the effect of producing more, though possibly less certain, DDI predictions.

It might be possible to assign a numerical value to each LOE that represents the user's certainty in any assertion possessing the combination of evidence the LOE models. Then, the system could arrive at a final certainty value for any inference by combining the confidence

value for all the assertions it depends on using some theory of reasoning under uncertainty. Section 5.1.4 of Chapter 5 presents an experiment testing the system's prediction accuracy using over 36,000 different belief criteria strategies. This experiment might form the basis for future work exploring the feasibility of assigning a numerical representation of user certainty to each LOE.

2.5.2 Expanding the DIKB

An important concern is how feasible it would be to expand the system's DDI prediction ability to more pharmaceutical entities and mechanisms. Some insight was gained into this question when the rule-based DDI theory used to design the DIKB (Figures 2.5 and 2.9) was replaced by a more sophisticated and validated DDI theory. Chapter 3 presents the new theory in detail but it is appropriate to mention here that its development was an iterative process that took two drug experts (Drs Carol Collins and John Horn) and myself several months to complete. New pharmaceutical entities could be easily added to the system but each addition increased the number of assertions for which evidence had to be collected. Adding more mechanisms was considerably more difficult since it required the development and validation new DDI prediction rules. Once the new rules were developed, it was very simple to add them to the `ddi-theory`. It was also simple to add any new objects and attributes required by the new rules to the `evidence-model`. Occasionally the drug-experts and myself had to develop new evidence types, levels-of-evidence, and belief criteria (or revise existing ones) before we could begin to collect evidence for the new attributes. To summarize, expanding the system to more drugs and mechanisms is feasible, but non-trivial. It is also important to recognize that any future expansions will need to take place while keeping existing knowledge in the system up to date.

2.6 Conclusion

This chapter begins by proposing that, in spite of the uncertain, incomplete, and dynamic nature of drug-mechanism knowledge, a system that correctly links and assesses the evidence support for drug-mechanism assertions can make clinically relevant drug-drug interactions.

This chapter then describes in detail a functional system that links drug-mechanism assertions to their supporting evidence and allows users to define, and vary, the criteria for belief in an assertion. In comparison with other knowledge-based systems that link evidence to their drug facts the DIKB is unique in that 1) it collects evidence both for and against assertions, 2) it enables users to define belief criteria for assertions using evidence meta-data, and 3) it can provide different views of its knowledge and inferences to users who might not agree about what combination of evidence makes an assertion believable. The design of the DIKB is intended to address several of the issues with modeling drug-mechanism knowledge. Later chapters in this dissertation explore the strengths and limitations of the system's design by attempting to predict real-world DDIs using only mechanistic assertions.

Chapter 3

A KNOWLEDGE REPRESENTATION FOR PREDICTING CLINICALLY MEANINGFUL DRUG-DRUG INTERACTIONS BY PHARMACOKINETIC INHIBITION

3.1 Introduction

This chapter presents a computable representation of a theory on how drugs interact with each other by metabolic inhibition. The new knowledge representation enables a computer to predict metabolic inhibition interactions and non-interactions between drugs[a] and/or drug metabolites using only pharmacokinetic drug-mechanism knowledge. The knowledge representation and its inference machinery compose the *reasoning system* component the Drug Interaction Knowledge Base (DIKB) shown in Figure 2.3 of Chapter 2. Experiments (see Chapter 5) with the DIKB demonstrate that it is capable of accurately predicting clinically-relevant drug-drug interactions (DDIs) for an important class of therapeutic agents and avoids making the kinds of false predictions that occur when individual drug differences are not recognized.

3.1.1 A Significant Problem and a Potential Solution

There are many drug combinations whose combined effects have never been investigated in clinical trials.[b] Information systems that only catalog DDI studies involving drug pairs can provide little or no guidance on the safety of unstudied drug combinations. This fact presents a difficult obstacle to clinicians who often must assess the potential risk of an adverse event between medication combinations that have not been studied together in a clinical

---

[a]Throughout this chapter we use the term *drug* to mean an active pharmaceutical ingredient – a molecular substance that is a component of a drug product or formulation and has pharmacologic properties. We use the term *drug metabolite* to mean a molecule that is the product of enzymatic processes involving some active pharmaceutical ingredient.

[b]A very rough estimate of the minimum number of missing clinical trials investigating DDIs for drug pairs can be found in Appendix A

trial [142]. There are at least two possible strategies that can compensate for the significant knowledge gaps that exist within the domain of drug-drug interactions. One strategy is to generalize interactions involving some drug to all other drugs within its therapeutic class. The other strategy combines knowledge about biochemical and physiological mechanisms of drug absorption, distribution, metabolism, and excretion with an understanding of the how drugs interact with each other to make mechanism-based predictions.

3.1.2 The "Class-based" Reasoning Strategy

Prescribers tend to think about drugs in terms of therapeutic class and disease [142] and generalizing interactions involving some drug to all other drugs within its therapeutic class fits that perspective well. While clinically relevant class-based interactions exist (for example, the SSRIs and NSAIDs [115, 124]), this approach has been criticized for leading some drug information systems to catalog DDI predictions that are either false or are likely to have little clinical relevance [82]. Class-based prediction can lead to false DDI predictions because many interactions occur by metabolic mechanisms and drugs within the same therapeutic class can vary widely in their metabolic characteristics.

For example, the following statement from the current drug product labeling for erythromycin [2] extrapolates an interaction observed between the macrolide antibiotic and one or more HMG-CoA reductase inhibitors to all drugs in that class:

> Erythromycin has been reported to increase concentrations of HMG-CoA reductase inhibitors (e.g., lovastatin and simvastatin). Rare reports of rhabdomyolysis have been reported in patients taking these drugs concomitantly.

Since *rosuvastatin* is member of the HMG-CoA reductase inhibitor drug class, it is reasonable to infer from this labeling statement that there is the potential for a pharmacokinetic interaction between erythromycin and rosuvastatin. However, a randomized clinical trial could find no increase in rosuvastatin concentrations in the presence of erythromycin [47]. The results of this clinical trial appear in the current product labeling for rosuvastatin [10]. A clinician reading both the erythromycin and rosuvastatin product

labels would see contradictory statements and might be left wandering if they should feel safe prescribing erythromycin to a patient taking rosuvastatin. The next section shows how a mechanism-based reasoning strategy might help resolve this dilemma.

3.1.3 The "Mechanism-based" Reasoning Strategy

A complementary approach to class-based DDI prediction is to apply a theory of how drugs can interact at the level of biochemical and physiological mechanisms to knowledge of the mechanistic properties of each drug in a drug combination. This approach is complementary to class-based DDI prediction because scientific knowledge is often incomplete making it likely that will always be some interactions that are known to occur but cannot be explained by a given state of mechanistic knowledge. However, the mechanisms for many DDIs are understood and mechanism-based reasoning is currently an important part of the pre-clinical investigation of new drug candidates. In this setting, the metabolic mechanisms of new drug candidates are compared with the known mechanisms of existing drugs to predict combinations that might result in pharmacokinetic DDIs [26]. Mechanism-based predictions made during pre-clinical drug development can be followed by clinical trials to determine the clinical significance of the predicted DDIs [131]. Knowledge derived from pre-clinical drug-mechanism investigations and pre-market clinical trials can be used to construct mathematical and computational models that map pharmacokinetic interactions to pharmacodynamic effect [56]. These kinds of models help pharmaceutical and regulatory organizations assess the efficacy and safety of new drugs before they are released on the market.

We can apply mechanism-based reasoning to consider the likelihood that erythromycin will cause an increase in the concentration of rosuvastatin. First note that the process by which the body removes externally-introduced molecular compounds (xenobiotics) is called *clearance*. Inhibition of an enzyme that is important for the clearance of a xenobiotic can result in an increase in its systemic concentration [112]. The literature indicates that erythromycin inhibits the drug-metabolizing enzyme Cytochrome P-450 3A4 (CYP3A4) [73, 184]. CYP3A4 is important for the transformation of *some* HMG-CoA

reductase inhibitors into molecules that the body can easily excrete. Evidence from both clinical trials and *in vitro* studies suggests this mechanism to be the one responsible for the observed concentration increase of the two drugs named in the previously shown labeling statement (lovastatin and simvastatin) in the presence of erythromycin [128, 129, 143].

The aforementioned statement from erythromycin's product label does not clarify the mechanism underlying the observed pharmacokinetic interaction between erythromycin and some HMG-CoA reductase inhibitors. The statement does, however, imply what the effect of such an interaction could be in rare circumstances – rhabdomyolysis, a potentially fatal condition involving the destruction of muscle fiber. Patients taking a drug from the HMG-CoA reductase inhibitor class are at a higher risk for rhabdomyolysis and another muscle disorder, myopathy, if they are also taking a drug that is itself myotoxic or that reduces the clearance of the HMG-CoA reductase inhibitor [86].

Since HMG-CoA reductase inhibitors can be metabolized by different enzymes, some of them might not be subject to the same mechanism of reduced metabolic clearance as lovastatin and simvastatin. In fact, a randomized clinical trial could find no metabolic interaction between rosuvastatin and ketoconazole [48], a selective inhibitor of the CYP3A4 enzyme [26]. This study is strong evidence that, unlike lovastatin and simvastatin, CYP3A4 plays no clinically significant role in the clearance of rosuvastatin. A mechanisms-based explanation for why no metabolic interaction was observed between erythromycin and rosuvastatin during the previously mentioned clinical trial [47] is that rosuvastatin has little or no clearance by the enzyme that erythromycin inhibits.

Combining knowledge of the metabolic properties of erythromycin and rosuvastatin with an understanding of how metabolic inhibition effects an increase in drug concentration leads to the conclusion that erythromycin will likely not increase the systemic concentration of rosuvastatin by CYP3A4 inhibition. The validity of this non-interaction prediction is supported by a randomized clinical trial. The same theory of metabolic inhibition applied to knowledge of the metabolic properties of simvastatin and lovastatin would lead to the conclusion, also supported by clinical trials, that erythromycin should cause an increase in these drugs. The clinical implication of these two inferences is that a patient taking the erythromycin - rosuvastatin combination should be at a lower risk for developing a

muscle disorder than if they took erythromycin and an HMG-CoA reductase inhibitor whose primary clearance pathway is CYP3A4.

3.1.4 Mechanism-based Reasoning Presents Informatics Challenges

As the previous example illustrates, mechanism-based reasoning can be used to infer both clinically relevant interactions and non-interactions and is an improvement over therapeutic class-based reasoning alone because it avoids the kind of prediction errors that occur when individual drug differences are not recognized. We think that, because of these qualities, mechanism-based reasoning should be able to help expand the coverage of drug information systems to include accurate interaction predictions for unstudied drug combinations. Unfortunately, there are significant obstacles to this goal.

One potential obstacle is that mechanism-based reasoning requires knowledge of the mechanistic properties of each drug and drug metabolite but this knowledge is often missing or uncertain (see Chapter 2, Section 2.2 for further discussion). Another obstacle is that collecting and maintaining even a basic set of drug-mechanism knowledge for all drugs and drug metabolites of interest would require a significant amount effort. To illustrate this challenge, consider that the previous analysis of the metabolic interactions occurring by CYP3A4 inhibition between erythromycin and HMG-CoA reductase inhibitors (Section 3.1.3) only focused on two mechanistic properties: *metabolic clearance pathway* and *enzyme inhibition*. A recent query of the Federal Drug Administrations (FDA) `drugs@fda` database [58] of all currently approved prescription and over-the-counter drugs identified about 1300 unique drugs used in more than 7000 drug products.[c] Therefore, one would need to seek and maintain knowledge on 2600 drug mechanism properties if they would like to repeat a similar analysis with all 1300 drugs.

Another potential obstacle is that expanding the coverage of drug information systems using mechanism-based reasoning requires translating how drugs interact with each other at the level of biochemical and physiologic mechanisms into information that is useful for

[c]We made this estimate by first searching the drugs@fda database on 06/24/2006 for all the unique active pharmaceutical ingredients used in drug products currently on the US market then reducing this list manually by collapsing multiple versions of individual active pharmaceutical ingredients to a single entry.

clinical decision making. Just indicating that a metabolic inhibition interaction is possible does not provide much assistance to clinicians who must decide how to reduce the risk of an adverse outcome [82]. The clinical relevance of an interaction can require consideration of a number of factors including the evidence supporting the interaction, the potential effect of the interaction on a patient, the existence of special risk factors in particular patients, and the frequency of specific adverse events in patients taking the suspected interacting combination [169].

Chapter 2 introduced the DIKB, a system designed to test informatics methods for overcoming these challenges. The remainder of the current chapter describes one component of this system – the computable model of mechanism-based reasoning that the DIKB currently uses to predict interactions and non-interactions occurring by metabolic inhibition.

3.2 A Computable Representation of how DDIs Occur by Metabolic Inhibition

There are a few basic pharmacologic principles by which one can make mechanism-based DDI predictions. Pharmacodynamic interactions can occur when the pharmacodynamic effects of two drugs combine in additive, subtractive, or synergistic ways. Pharmacokinetic interactions can occur when the binding, metabolic or physical and chemical properties of one drug affect the absorption, distribution, metabolism and/or excretion (ADME) of another drug. Modulations in the ADME or the pharmacodynamic effects of a drug can lead to the possible negative outcomes of drug toxicity or loss of efficacy in patients [68, 112].

The computable representation we have built focuses on a narrow subset of mechanism-based interactions – metabolic *inhibition* interactions. A number of known DDIs occur by metabolic inhibition and we believe that the theory of metabolic inhibition has much in common with other mechanism-based DDI theories. For example, a recent FDA guidance includes discussions of how DDIs can occur by metabolic inhibition, induction, and transport protein modulation [26]. All three theories involve interactions between drugs, metabolites, enzymes, and routes of elimination that can effect changes in systemic concentration and metabolite formation. These commonalities make it reasonable that a clinically useful computational representation of metabolic inhibition will be extendable to other mechanisms.

This section begins by briefly summarizing the theory of how drug-drug interactions

occur by metabolic inhibition. It then explains the set of inferences our new knowledge representation (KR) is designed to support and the assumptions the KR makes about the process and effects of metabolic inhibition.

3.2.1 How DDIs Occur by Metabolic Inhibition

The biochemical process of enzyme inhibition can be classified into three major types: *rapidly reversible*, *slowly reversible*, and *irreversible* [112]. Rapidly reversible inhibition occurs when two substrates of an enzyme compete for the enzyme's active site (*competitive inhibition*), when an inhibitor binds to a substrate-enzyme complex (*uncompetitive inhibition*), or when a substrate causes an enzyme catalyst to lose its catalytic function (*non-competitive inhibition*). Slowly reversible inhibition occurs when an enzyme inhibitor forms a complex with the enzyme and the product of a catalytic reaction involving the enzyme. Irreversible inhibition occurs when an inhibitor covalently bonds to the enzyme forming a stable complex that permanently eliminates the enzyme's original catalytic function.

Any of the three major types of metabolic inhibition can cause a DDI by reducing the clearance of another drug whose metabolic clearance depends, at least in part, on the inhibited enzyme [112]. A decrease in the clearance of a drug by metabolic inhibition can lead to an increase of its systemic concentration potentially leading to drug toxicity and harmful side-effects [44]. The magnitude of a metabolic inhibition DDI is affected by several factors including the importance the victim drug or drug metabolite's non-enzymatic clearance routes and the number and importance of its metabolic clearance pathways [112].

For example, if the percentage of a drug's clearance by metabolism is less than 10%, then even complete metabolic inhibition should effect a relatively small increase in systemic concentration assuming all other non-enzymatic routes of clearance remain functional. Conversely, if an inhibitor reduces the function of an enzyme that accounts for more than 50% of a drug's total clearance, then the effect could be quite significant. Active metabolites can be formed by several metabolic pathways for some drugs or drug metabolites. An inhibitor might have a negligible effect on the total concentration of some active metabolite if it affects a minor pathway that leads to its formation. The exact opposite might be true

if an inhibitor reduces catalysis along the sole metabolic pathway responsible for forming a particular active metabolite.

3.2.2 The Set of Inferences Made by the Knowledge Representation

Knowledge of drug interaction mechanisms can assist a clinician in predicting the time course of an interaction or in deriving ways to minimize the risk of patient harm as a result of the interaction [82]. The DIKB's knowledge representation (KR) supports this kind of reasoning by providing clear details about the mechanisms by which two drugs or drug metabolites could interact via metabolic inhibition. The KR can make the following set of inferences designed to help a clinician assess the potential effect of a predicted interaction on a given patient.

3.2.2.1 Inference One

> **Inference One:** For some drug or drug metabolite, *D1*, is there another drug or drug metabolite, *D2*, that will reduce the clearance of *D1* by inhibition of some enzyme *E*? If so, what is the anticipated increase in concentration of *D1*?

Just indicating that a metabolic inhibition interaction is possible does not provide much assistance to clinicians who must decide how to reduce the risk of an adverse outcome [82]. We believe that, where possible, mechanism-based predictions should indicate the anticipated increase in concentration of the victim drug or drug metabolite. This information should increase the clinical value of a prediction because it is often true that systemic concentrations of a drug that are either too high or too low precede harmful effects. Undesirable side-effects ranging from loss-of-efficacy to death are concentration-related for many drugs including some anti-depressants, antiarrhythmics, and blood thinners.

Inference One predicts when a drug or drug metabolite pair will interact by metabolic inhibition and then classifies the interaction into one of three discrete categories based on the anticipated magnitude increase in concentration of the victim drug. Section 3.2.3.5. defines these levels and presents the logic behind their justification.

3.2.2.2 Inference Two

> **Inference Two:** For some drug or drug metabolite, *D1*, is there another drug or drug metabolite, *D2*, that *will not* reduce the clearance of *D1* by inhibition of some enzyme *E*?

The discussion in Section 3.1.1 of the potential for erythromycin to reduce the clearance of rosuvastatin demonstrates that non-interaction predictions can also be of clinical value. *Inference Two* predicts when a drug or drug metabolite pair will *not* interact by a *specific* metabolic clearance pathway. The inference does not exclude the possibility that such pairs might interact by alternative mechanisms but can help eliminate some mechanisms from consideration.

3.2.2.3 Inference Three

> **Inference Three:** Which drugs or drug metabolites will cause a *decrease* or *increase* in the formation of a drug metabolite by enzyme inhibition? If so, will a decrease or increase in the formation of an drug metabolite have a non-ambiguous effect on a descendent metabolite?

DDIs occurring by metabolic inhibition can affect the concentration of active or toxic drug metabolites in clinically relevant ways. For example, both lovastatin and simvastatin are administered in lactone forms that have little or no HMG-CoA reductase inhibition activity but that are readily converted by the body to pharmacodynamically active metabolites [119, 120]. Clinical trial data indicates that metabolism by CYP3A4 is a clinically relevant clearance pathway for these metabolites [128, 129]. Similarly, *in vitro* evidence indicates that CYP3A4 is the primary catalyst for the conversion of the HMG-CoA reductase inhibitor atorvastatin into its two active metabolites [94]. *Inference Three* allows the prediction of drug - metabolite and metabolite - metabolite interactions that could affect the concentration of some active metabolite.

3.2.3 Assumptions Made by the Knowledge Representation

The KR computes over a mechanism-based theory of DDIs and produces simple, qualitative, DDI predictions that might be useful in drug therapy planning and management. It makes several simplifying assumptions that are important to consider when assessing the clinical relevance of its predictions because the assumptions might not hold for many drug combinations.

3.2.3.1 The KR Defines Inhibition Qualitatively

The KR defines enzyme inhibition qualitatively as a measurable *in vivo* occurrence:

> *inhibits:* A drug or drug metabolite, X, is said to `inhibit` some enzyme, E, if X effects a measurable reduction in the catalytic function of E in humans.

Evidence of enzyme inhibition can come from multiple sources; this definition specifically excludes the *direct* use of evidence from *in vitro* experiments.[d] Evidence from *in vitro* experiments is especially common for many drugs and considerable interest from both industrial and academic researchers has been focused on how to make quantitative estimates of *in vivo* effects from *in vitro* evidence. Unfortunately, there is currently no general method for making accurate quantitative estimates of the magnitude of a metabolic inhibition DDI using *in vitro* data [131].

The KR implements a method for using *in vitro* evidence to *indirectly* support a measurable *in vivo* effect in humans. In the KR, a drug or drug metabolite is labeled an *in vivo* inhibitor for some drug metabolizing enzyme at the concentrations it is expected to reach during drug therapy if the following relationship holds:

$$\frac{C_{max}}{K_i} > 0.1 \tag{3.1}$$

[d] WordNet [121] lists the definition of *in vitro* as "in an artificial environment outside the living organism" and *in vivo* as "in the living organism." The DIKB excludes all data from non-human animal models so, throughout this dissertation, the terms *in vitro* and *in vivo* refer to experiments with human tissue or clinical trials respectively.

Where C_{max} is the maximum observed concentration the inhibitor has reached in patients at normal therapeutic doses and K_i is an inhibition constant for *reversible* inhibition derived from a well-designed *in vitro* enzyme inhibition experiment involving the inhibitor. This relationship applies to inhibition of members of the Cytochrome P-450 enzyme family and is not applicable if the inhibitor is thought to permanently remove the affected enzyme from further participation in catalysis by any means. The basis for this relationship can be found in a recent FDA guidance to industry that includes the recommendation that a clinically relevant effect from competitive enzyme inhibition be considered possible if the following relationship holds (see [26], p.33):

$$\frac{[I]}{K_i} > 0.1 \tag{3.2}$$

Where $[I]$ is the estimated concentration of the inhibitor at the enzyme binding site.

The KR also allows *in vitro* evidence to indirectly *refute* that a drug or drug metabolite effects a measurable *in vivo* effect in humans. In the KR, a drug or drug metabolite is labeled an *in vivo* non-inhibitor for some drug metabolizing enzyme at the concentrations it is expected to reach during drug therapy if the following relationship holds:

$$\frac{C_{max}}{K_i} <= 0.1 \tag{3.3}$$

Where C_{max} is the maximum observed concentration the inhibitor has reached in patients at normal therapeutic doses and K_i is an inhibition constant for *reversible* inhibition derived from a well-designed *in vitro* enzyme inhibition experiment involving the inhibitor. The KR will not allow *in vitro* evidence to support or refute that a drug or drug metabolite is an *in vivo* inhibitor of some enzyme if that entity is known to permanently deactivate the enzyme's catalytic function. Modeling the effects of such inhibitors *in vivo* requires sophisticated reasoning that is outside of the scope of the current DIKB.

3.2.3.2 The KR Has No Concept of Time and Does Not Distinguish Between Types of Inhibition

The KR makes no estimates of the time-course of its interaction predictions. It assumes that any drug or drug metabolite that inhibits an enzyme will cause some, possibly negligible, reduction in the clearance of any drug the enzyme catalyzes at some, non-specified, time after the inhibitor is administered.

Since the KR has no concept of time, it does not bother to distinguish between the various types of inhibition (Section 3.2.1). It can be important to distinguish reversible from irreversible inhibition when predicting the time-course of metabolic inhibition because the effects of irreversible inhibition are both time and dose-dependent while reversible inhibition is generally only dose-dependent [114]. These distinctions are less important when the time-course of an interaction is ignored because metabolic inhibition by any type, when supported by *in vivo* evidence, should always lead to some increase in the plasma concentration of the effected drug [114].

3.2.3.3 The Percentage of Drug Clearance by All Major Routes is Fixed

The KR models four possible major routes of clearance for a drug or drug metabolite – metabolism, renal excretion, biliary excretion, and exhalation. The KR assumes that the percentage of a drug or drug metabolite cleared by each of these major clearance routes remains fixed. For example, assume that 20% of a drug is cleared by a single metabolic pathway and 80% by renal excretion. In this situation the KR would assume that the percentage of drug cleared by renal excretion will remain at 80% even if the metabolic clearance pathway is completely inhibited. The KR uses this assumption to reason that the 20% of drug once cleared by metabolism will contribute to an increase in systemic concentration.

3.2.3.4 Assumptions About Metabolic Inhibition and Metabolic Clearance Pathways

The KR can model drug and drug metabolites that have either single or multiple metabolic clearance pathways. The KR's model of metabolic clearance pathways is isomorphic to an

acyclic graph of catalytic reactions where nodes are drugs or drug metabolites and branches are the specific catalytic enzymes (Figure 3.1). A connection between two nodes represents the conversion of one substrate (a drug or drug metabolite) to one metabolite by some enzyme. The root node of a metabolic clearance pathway is the starting drug or drug metabolite and the leaf nodes are the pathway's final metabolic products.

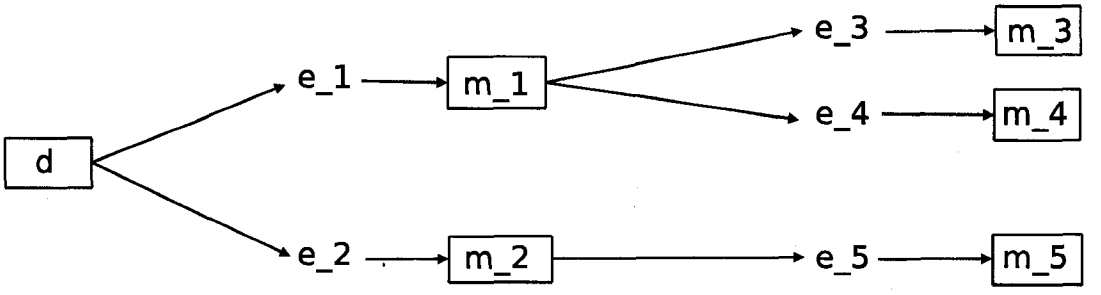

Figure 3.1: The KR's model of metabolic clearance pathways is isomorphic to an acyclic graph of catalytic reactions where nodes are drugs or drug metabolites and branches are the specific catalytic enzymes. In this image, text enclosed by a box represents drug and metabolite nodes. There are five catalytic reactions represented in the image; two sibling reactions converting drug "d" to "m_1" and "m_2", two sibling reactions converting "m_1" to "m_3" and "m_4", and one reaction converting "m_2" to "m_5"

The KR assumes that the effect of inhibiting any catalytic reaction within a metabolic clearance pathway is to increase the concentration of the drug or drug metabolite that is the substrate of the reaction and to decrease the concentration of all metabolites produced by any downstream catalytic reactions. In other words, the KR assumes that metabolic inhibition is transitive along a metabolic clearance pathway so that the concentration of metabolites downstream from an inhibited catalytic reaction will also experience some, possibly negligible, decrease in concentration.

When a drug or drug metabolite has multiple metabolic clearance pathways, the set of alternate catalytic reactions involving it are called *sibling catalytic reactions* (for an example see Figure 3.1). The KR assumes that inhibition of one sibling catalytic reaction will influence the formation of metabolites produced by all other sibling reactions provided that 1) the enzymes in sibling catalytic reactions are known, 2) they do not catalyze the

same reaction as the inhibited enzyme and 3) they are not also inhibited.

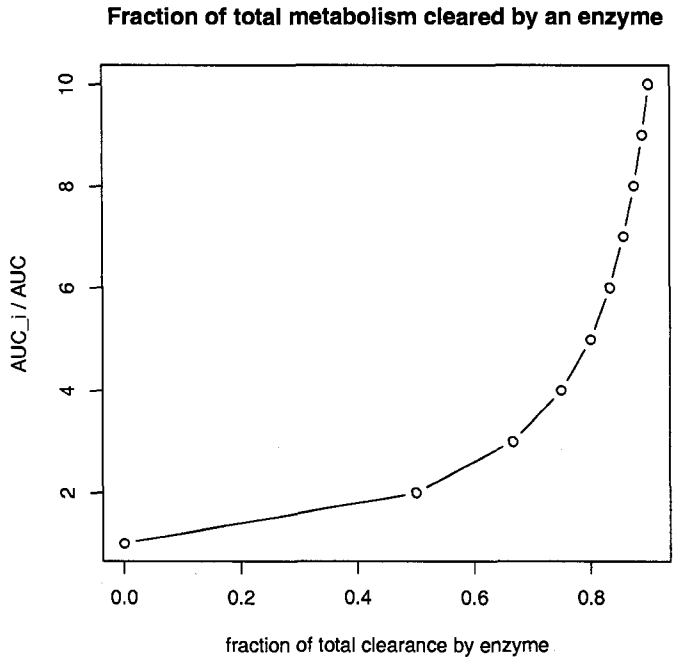

Figure 3.2: A very simple model that we designed for inferring the fraction of total clearance contributed by an enzyme from the AUC data provided in a pharmacokinetic clinical trial. The model is based on a number of assumptions that ignore many of the factors that can contribute to an increase in the AUC of an object drug in pharmacokinetic DDI study. For example, the model assumes that the data is from a pharmacokinetic study with a inhibitor that is selective for the enzyme *in vivo* and that linear inhibition kinetics hold. Please see Section 3.2.3.5 for a more detailed discussion.

3.2.3.5 Estimating Concentration Increases

Inference One (Section 3.2.2.1) predicts when a drug or drug metabolite pair will interact by metabolic inhibition and then classifies the interaction into one of three discrete categories that we defined based on the the anticipated magnitude increase in concentration of the object drug:

1. *PKI-1* indicates that the concentration of the affected drug or drug metabolite should

increase by some, possibly negligible, amount.

2. *PKI-2* indicates that the concentration of the affected drug or drug metabolite should increase by at least 33% (1.3 fold).

3. *PKI-3* indicates that the concentration of the affected drug or drug metabolite should increase by at least 100% (2 fold).

One typical measure taken during pharmacokinetic clinical trials is the Area Under the Concentration time curve (AUC) of some drug or drug metabolite before and after the administration of another drug. Some researchers, such as Ohno *et al* [133], have developed sophisticated mathematical models for inferring the fraction of a drug that is cleared by a particular enzyme (f_{enz}) from AUC data. The three levels defined above were chosen based on a very simple model that we designed for doing the same thing. The model is based on a number of assumptions that ignore many of the factors that can contribute to an increase in the AUC of an object drug in pharmacokinetic DDI study. We will defer discussion of the accuracy of this very simplistic model to Chapter 5 when we discuss the results of an experiment we conducted to characterize the prediction accuracy of the DIKB.

Assume that a well-designed clinical trial investigates the pharmacokinetics of drug `X` in the presence of drug `Y`. Assume also that drug `Y` has no measurable effect on `X`'s clearance by renal clearance, biliary clearance, or exhalation and is a *selective* inhibitor of some enzyme, `ENZ`.[e] Finally, assume that the amount to which drug `Y` inhibits `ENZ` is the same regardless of its unbound systemic concentration and that linear inhibition kinetics hold. We propose that the following equation will provide a rough estimate of the fraction of total clearance contributed by `ENZ` under these assumptions:

$$f_{enz} = 1 - \frac{1}{\frac{AUC_i}{AUC}} \tag{3.4}$$

Where AUC is the baseline Area Under the Concentration time curve for `X`, AUC_i is the Area Under the Concentration time curve for `X` when `Y` is co-administered and f_{enz} is

[e]In other words, `Y` inhibits no other enzyme besides `ENZ` at the doses it is given to participants of the study

the fraction of total clearance ENZ contributes to X.

Figure 3.2 shows a plot of Equation 3.4 at AUC ratios ranging from zero to ten. The plot shows that if a selective inhibitor causes a drug's AUC to increase more that 2-fold then one can infer that the inhibited enzyme is responsible for at least 50% of the affected drug's clearance. This logic is reversed in the KR so that it predicts a 2-fold or greater increase in the concentration of any drug or drug metabolite if an enzyme responsible for at least 50% of the entity's total clearance is inhibited. The KR applies this reasoning to establish that a predicted interaction is at the *PKI-3* level.

Many drugs or drug metabolites have no single metabolic pathway responsible for more than 50% of their total clearance. Figure 3.2 shows that selective inhibition of an enzyme responsible for 25% or more of a drug or drug metabolites clearance should result in an AUC increase of at least 1.3 fold. The KR applies this reasoning to establish that a predicted interaction is at the *PKI-2* level.

Finally, if an inhibited enzyme is not known to contribute more than 25% or more of a drug or drug metabolites clearance, then the KR will predict the interaction to be at the *PKI-1* level. While the percent increase in concentration of a victim drug or drug metabolite at the *PKI-1* level is small it might be of clinical interest if the entity is a "pharmaceutical entity of concern" – an active ingredient or metabolite for which even a small change in the system concentration would be of concern to a clinician. Such entities might include drug or drug metabolites for which therapeutic drug monitoring is required or for which the ratio between the toxic systemic concentration of the entity and the concentration at which the entity is therapeutic is less than or equal to 2.0.

3.3 The KR's Rules and Semantics

This section provides several technical details of how the KR infers drug-drug interactions and non-interactions occuring by metabolic inhibition.

3.3.1 The Machinery for Reasoning - Declarative Rules

The KR's mechanism-based theory of drug-drug interactions and non-interactions is represented as a set of declarative rules – structured logic sentences that explicitly convey the

implication of a certain body of knowledge about drug mechanisms. The logical form of the KR's rules is restricted to being a disjunction of logically static constants (called *literals*) of which *exactly one* is positive. The following example is typical of the logical form:[f]

$$\neg A \vee \neg B \vee C \tag{3.5}$$

This logical form is also known as a *definite clause* and it can be shown using truth tables that it is logically equivalent to the following implication which is both easy to read and write:

$$A \wedge B \Rightarrow C \tag{3.6}$$

The definite clause form is a restricted version of another kind of clause called a *Horn clause* defined as a disjunction of literals of which *at most* one is positive. There exists inference algorithms that are proven to perform *sound* and *complete* inference with Horn clauses very efficiently [154]. *Sound* inference algorithms derive only the set of inferences entailed by a knowledge base. *Complete* inference algorithms derive all inferences that are entailed by a knowledge base. The DIKB's reasoning system applies a forward-chaining inference algorithm to the KR's rules and assertions that is sound and complete for Horn clauses and has a computational complexity that grows linearly with the size of the knowledge-base.

3.3.2 The KR Supports Default Reasoning

Mechanism-based reasoning requires knowledge of the mechanistic properties of drugs and drug metabolites but this knowledge is often missing (see Chapter 2, Section 2.2 for further discussion). The KR represents missing knowledge that is important for mechanism-based reasoning as assumptions whose truth state can change. Each rule in the KR is written as a

[f]Here we use '$\wedge$' to represent conjunction (e.g. X *and* Y), '$\vee$' to represent disjunction (e.g. X *or* Y), '$\Rightarrow$' to represent implication (e.g. X *implies* Y), and $\neg$ to imply negation.

definite clause whose predicates can contain *default assumptions* – knowledge whose truth state is assigned by default. The KR can retract or reinstate inferences that depend on such assumptions as appropriate depending on their truth state. This is a form of *default reasoning* whose various forms include inheritance in semantic networks, circumscription, default logic, and several methods discussed by Goldszmidt and Pearl that utilize qualitative probabilities [69]. This feature expands the kinds of drug knowledge that the KR can represent without significantly affecting how such knowledge appears in the knowledge-base. Section 2.2.2 of Chapter 2 explains the details of how the KR supports default reasoning.

3.3.3 The Current Set of Rules and Assertion Types

The KR's rule-base consists of 38 rules and is written to execute on Forbus and de Kleer's ANSI Common Lisp rule engine (JTRE) and Justification-based Truth Maintenance System (JTMS) [61]. Forbus and de Kleer's JTRE/JTMS was chosen because its an open-source implementation of the simplest family of Truth Maintenance System that is also and well documented. Appendix B contains a complete listing of the rules that comprise the KR at the time of this writing. Figures 3.3, 3.4, and 3.5 show the three rules that the KR uses to make *Inference One* (Section 3.2.2.1). The rules are written in a slightly different, but more readable, syntax than that used by the KR.

```
IF ?x INHIBITS ?y AND
   ?z is-SUBSTRATE-OF ?y
   :TEST (NOT (EQUAL ?x ?z))
THEN
   ?x INHIBITS-METABOLIC-CLEARANCE-of ?z via ?y
```

Figure 3.3: The rule shown in this figure declares that some, possibly negligible, inhibition of the metabolic clearance of a drug or drug metabolite, `?z`, will occur if another drug or drug metabolite, `?x`, inhibits the catalytic function of some enzyme `?y`. The KR maps interaction predictions made using this rule to the *PKI-1* level (Section 3.2.3.5).

The rule in Figure 3.3 declares that some, possibly negligible, inhibition of metabolic clearance of a drug or drug metabolite, `?z`, will occur if another drug or drug metabolite `?x`

inhibits the ability of some enzyme ?y to catalyze ?z. The KR maps interaction predictions made using this rule to the *PKI-1* level (Section 3.2.3.5).

```
IF ?x INHIBITS-METABOLIC-CLEARANCE-of ?z via ?y AND
  PRIMARY-TOTAL-CLEARANCE-MECHANISM-of ?z
                              'METABOLIC-CLEARANCE AND
  PRIMARY-METABOLIC-CLEARANCE-ENZYME-of ?z is ?y
THEN
 ?x INHIBITS ?y the-PRIMARY-METABOLIC-ENZYME-of ?z
```

Figure 3.4: The rule shown in this figure declares that some drug or drug metabolite, ?x, inhibits the enzyme responsible for at least 50% of the *metabolic* clearance of another drug or drug metabolite, ?z, whose *total* clearance by metabolism is at least 50%. The KR maps interaction predictions made using this rule to the *PKI-2* level (Section 3.2.3.5).

The rule that predicts DDI interactions at the *PKI-2* level is shown in Figure 3.4. This rule declares that some drug or drug metabolite, ?x, inhibits the enzyme responsible for at least 50% of the *metabolic* clearance of another drug or drug metabolite, ?z, whose *total* clearance by metabolism is at least 50%. This rule can be useful if no pharmacokinetic clinical trial has been conducted investigating the importance of a particular enzymatic pathway to a drug or drug metabolite. The predicates in this rule were chosen because the percentage of a drug or drug metabolite's clearance by metabolism is generally easy to find and *in vitro* data, which tends to be more readily available, can sometimes indicate if one enzyme dominates metabolic clearance.

Figure 3.5 shows the rule that the KR uses to predict DDIs at the *PKI-3* level. The rule declares that some drug or drug metabolite, ?x, inhibits the enzyme responsible for at least 50% of the *total* clearance of another drug or drug metabolite, ?z. Section 3.2.3.5 explains the logic behind the qualitative estimate that this rule and the rule in Figure 3.4 makes.

3.3.4 Precise Definitions Provide KR Semantics

The theory of how DDIs occur by metabolic inhibition (Section 3.2) involves drugs, metabolites, enzymes, routes of elimination, and changes in systemic concentration and metabolite formation. The KR represents each of these entities so that a computer can infer DDIs.

```
IF ?x INHIBITS-METABOLIC-CLEARANCE-of ?z via ?y AND
   PRIMARY-TOTAL-CLEARANCE-ENZYME-of ?z is ?y
THEN
 ?x INHIBITS ?y the-PRIMARY-TOTAL-CLEARANCE-ENZYME-of ?z
```

Figure 3.5: The rule shown in this figure declares that some drug or drug metabolite, `?x`, inhibits the enzyme responsible for at least 50% of the *total* clearance of another drug or drug metabolite, `?z`. The KR maps interaction predictions made using this rule to the *PKI-3* level (Section 3.2.3.5).

The semantics of each entity represented by the KR are found in two sources – a structured vocabulary that we designed specifically for supporting mechanism-based DDI reasoning and a simple dictionary. Both sources help add precision to the assertions in the DIKB's knowledge-base which, in turn, clarifies the meaning of the KR's inferences.

3.3.4.1 The DIKB's Structured Vocabulary

The DIKB's structured vocabulary helps to add precision to the KR by providing clear definitions to many of the objects and processes that the KR models. Figure 3.6 shows the taxonomic relationships between terms in two portions of the vocabulary. The KR uses many of the terms in the structured vocabulary as values for the predicates used by the KR's DDI prediction rules. For example, the vocabulary defines three specific terms for the pharmacokinetic process of drug excretion – `Biliary_Excretion`, `Exhalation_Excretion`, and `Renal_Excretion`. The KR uses these three symbols as potential values for the `primary--total-clearance-mechanism` assertion shown in the antecedent portion of the rule in Figure 3.4.

An important example of how the vocabulary adds precision to the KR is the definition of the seemingly simple term "drug." When people typically speak about a "drug" they are often referring to pharmaceutical preparations such as a drug formulation ("250mg clarithromycin tablets") or product ("Biaxin Filmtab"). Pharmaceutical preparations are entities that can have several components such as active and inactive ingredients, dyes, buffer, and sweeteners. In contrast, the DIKB's definition of a "drug" is:

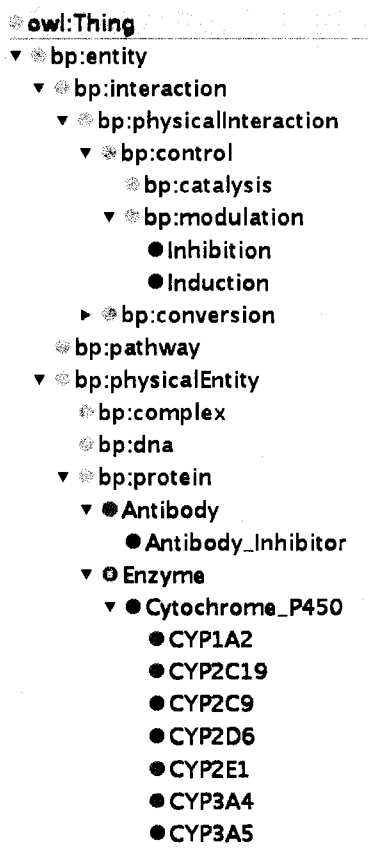

(a) A portion of the DIKB's vocabulary showing the relationships between the biochemical entities and the pharmacokinetic interactions of `inhibition` and `induction`. This portion of the ontology extends the `interaction` and `entity` subhierarchies of the BioPAX ontology [45]

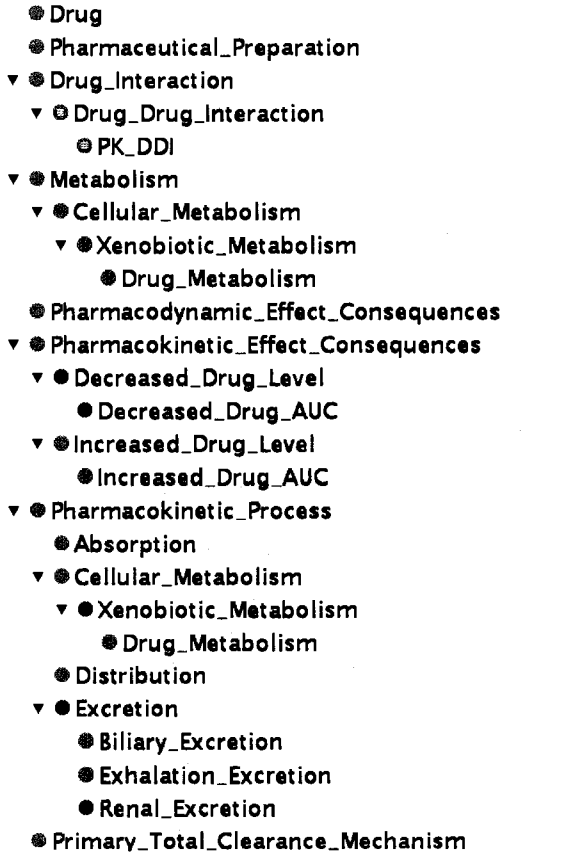

(b) A portion of the DIKB's vocabulary defining the relationships between many of the pharmacologic components used in the rule-based theory of metabolic inhibition DDIs

Figure 3.6: Elements from the DIKB's structured vocabulary

> *Drug:* a specific molecular substance that has pharmacologic properties and is the component of a pharmaceutical preparation such as a drug product or formulation

The definition of "drug" and "pharmaceutical preparation" in the DIKB's vocabulary maps to types defined in the Veteran's Administration's National Drug File Reference Terminology (NDF-RT) [37]. The type `Drug` is a direct sub-type of NDF-RT concept code C178 *"active ingredients"* while `Pharmaceutical_Preparation` is unified with concept code NDF-RT C176 "pharmaceutical preparation." These types reside in distinct taxonomies within the NDF-RT and so are disjoint within the DIKB's vocabulary and within the KR. This arrangement makes sense when making mechanism-based DDI predictions because the pharmaceutical preparation that an active pharmaceutic ingredient belongs to can affect the set of interactions that it can be involved in. For example, an intravenous preparation of an acid-labile drug would not be susceptible to interactions related to changes in pH in the gut while an oral preparation of the same drug might be. The KR can represent these differences using rules that revise the predictions it makes for an active pharmaceutic ingredient based on the specific pharmaceutical preparation it belongs to.

The DIKB's vocabulary is implemented in the OWL-DL language [46]; a *description logic* that provides a formal semantics for representing taxonomic relationships in a manner that can be automatically checked to ensure consistent classification. We used the Protégé ontology editor[g] to create the vocabulary and the RACER inference engine [80] to test it for consistent type definitions. The vocabulary is incorporates several existing type and concept definitions from other biomedical terminologies such as the NCI Thesaurus (NCI) [90], Gene Ontology (GO) [64], the National Library of Medicine's Medical Subject Headings (MeSH) [132], and BioPAX [45]. It also defines many new types and concepts that were not found in other terminologies. The current version is of the DIKB's vocabulary is available on the Web [35].

---

[g] http://protege.stanford.edu/

3.3.4.2 Precise Definitions for Rule Predicates

The DIKB's structured vocabulary does not yet include all terms that the KR uses. A simple dictionary (Appendix C) provides an additional set of definitions; many of which define predicates found in the KR's rules. We plan to incorporate these definitions into the DIKB's structured vocabulary at sometime in the future.

3.4 Validation and Evaluation

3.4.1 A Non-trivial Validation Test

We designed a non-trivial validation test involving a hypothetical drug having multiple metabolic clearance pathways and descendant metabolites as a test case for verifying the KR's rule-base. The example is shown here to demonstrate the kinds of inference the KR is capable of and to clarify the effect on inference of the KR's assumptions.

3.4.2 The Hypothetical Drug "C-cure"

Figure 3.7 shows a state of scientific knowledge about a *hypothetical* cancer drug *C-cure*. The figure shows that *C-cure* has many metabolic clearance pathways including conversion to the *major-met-1*, *major-met-2*, and *entity-of-concern-B* metabolites. The KR considers these three catalytic reactions to be siblings since each catalyst has *C-cure* as a substrate. Notice that several enzymes catalyze *C-cure*'s conversion to *major-met-1*; the KR considers the reactions that these enzymes catalyze to also be siblings. In this hypothetical example, the conversion of *C-cure* to the *major-met-1* metabolite by the CYP2B6 enzyme is assumed to contribute at least 50% to *C-cure*'s total clearance.

The three metabolites of *C-cure* (*minor-met-1*, *entity-of-concern-A*, and *minor-met-2*) are each substrates of other catalytic reactions, three of which have unknown enzyme catalysts. The hypothetical metabolite of *entity-of-concern-B* (*minor-met-3*) is a substrate of yet another catalytic reaction involving the same enzyme as the parent catalytic reaction (CYP2A6). In this hypothetical example *entity-of-concern-A* and *entity-of-concern-B* are pharmaceutical entities of concern because they both have concentration-dependent toxic effects.

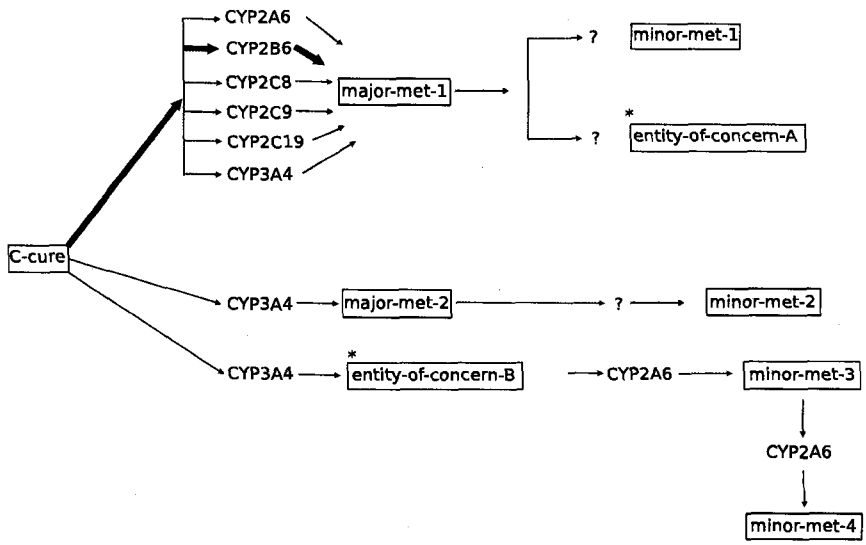

Figure 3.7: A model of the metabolic clearance pathways for the hypothetical drug *C-cure*. Boxes are drawn around *C-cure* and its metabolites. Arrows show the enzymes that control transformation from parent to child compound. The thickest arrows indicate that the pathway responsible for at least 50% of *C-cure*'s total clearance is conversion to the *major-met-1* metabolite by the CYP2B6 enzyme. The '?' symbol is used to represent a case where the transformation is known but no controlling enzyme has been identified. A star or red box by a metabolite means that is of clinical interest.

We entered this knowledge about *C-cure* into the KR and then added to the system the following default assumptions:

- itraconazole is a selective inhibitor of CYP3A4 *in vivo*
- sulfinpyrazone is a selective inhibitor of CYP2C9 *in vivo*
- clopidogrel is a selective inhibitor of CYP2B6 *in vivo*

While the example is hypothetical there is some evidence for these three assumptions. For example, Itraconazole and sulfinpyrazone are listed as potent *in vivo* inhibitors of the respective enzymes in a recent FDA guidance document [26] and there is both *in vitro* and *in vivo* evidence suggesting that clopidogrel inhibits CYP2B6 [166, 149]. The system made

the following set of inferences, all of which are in accord with the KR's assumptions about metabolic inhibition:

3.4.2.1 Interaction Predictions for Clopidogrel:

- Clopidogrel would inhibit CYP2B6 causing a reduction in the the conversion of *C-cure* to its *major-met-1* metabolite. This would cause *C-cure*'s concentration to increase at the *PKI-3* level because conversion of *C-cure* to the *major-met-1* metabolite by the CYP2B6 enzyme contributes at least 50% to *C-cure*'s total clearance.

- The inhibited conversion of *C-cure* to its *major-met-1* metabolite would reduce the conversion of *C-cure* to *minor-met-1* and *entity-of-concern-A*. The system noted that *entity-of-concern-A* is a pharmaceutical entity of concern.

- The increased concentration of *C-cure* would cause an increase in the concentration of *C-cure*'s *major-met-2* and *entity-of-concern-B* metabolites since these are sibling catalytic reactions that do not catalyze the same reaction as the inhibited enzyme and are not themselves inhibited. The system notes that *entity-of-concern-B* is a pharmaceutical entity of concern.

- Since, in this example, clopidogrel is assumed to have no effect on CYP2A6, the increased concentration of *entity-of-concern-B* would lead to an increase in the concentration of *minor-met-3* and *minor-met-4*.

- clopidogrel would cause an increase in the concentration of *major-met-2* but the KR would label the effect of this increase on *minor-met-2* as unknown because the enzyme that catalyzes its formation is not identified.

3.4.2.2 Interaction Predictions for Sulfinpyrazone:

- Sulfinpyrazone would reduce the transformation of *C-cure* to its *major-met-1* metabolite by inhibiting CYP2C9. This would cause an increase in *C-cure*'s concentration

at the *PKI-1* level. The increase in *C-cure* will cause the concentration of the *major-met-2* and *entity-of-concern-B* to increase by some, possibly negligible magnitude.

- The inhibited conversion of *C-cure* to its *major-met-1* metabolite would reduce the conversion of *C-cure* to *minor-met-1* and *entity-of-concern-A*. The system noted that *entity-of-concern-A* is a pharmaceutical entity of concern.

- Since, in this example, sulfinpyrazone is assumed to have no effect on CYP2A6, the KR infers that the increased formation of *entity-of-concern-B* would lead to an increased formation of *minor-met-3* and *minor-met-4*.

- The KR would infer that sulfinpyrazone would cause an increase in the formation of *major-met-2* but would label the effect of this increased formation on *minor-met-2* as unknown because the enzyme that catalyzes its formation is unknown.

3.4.2.3 Interaction Predictions for Itraconazole:

- Itraconazole would inhibit CYP3A4 and reduce the catalysis of *C-cure* to its *major-met-1*, *major-met-2*, and *entity-of-concern-B* metabolites. This would cause an increase in *C-cure*'s concentration at the *PKI-1* level.

- Even though CYP3A4 inhibition causes an increase in *C-cure* concentration, its effect on the concentration of *major-met-1*, *major-met-2*, and *entity-of-concern-B* is uncertain. This is because the conversion of *C-cure* to each metabolite is also inhibited. Since the KR cannot predict which will be greater, the increase in *C-cure* concentration or the reduction in formation of the other metabolites, it labels these effects as unknown.

- The KR considers the effect of an increase in *C-cure* concentration on *major-met-1* due to inhibition of CYP3A4 to be unknown. Since *minor-met-1* and *entity-of-concern-A* are formed by downstream catalytic reactions the KR considers the effect of CYP3A4 inhibition on these metabolites to be unknown. Similarly, the KR considers the effect

of CYP3A4 inhibition on the formation of *minor-met-2*, *minor-met-3*, and *minor-met-4* to be unknown because the effect of CYP3A4 inhibition on upstream metabolites (*major-met-2* and *entity-of-concern-B*) is unknown.

3.5 Discussion

3.5.1 The KR is a Very Simplistic Model

The KR was designed to make accurate, qualitative, predictions for a wide range of drugs or drug metabolites using strictly a mechanism-based DDI inference strategy. As a result, the KR's model of metabolic inhibition includes assertions about mechanisms that we thought would be relatively easy to find in the literature for most drugs. We deliberately excluded concepts that can be found in the DDI literature such as the "strength" of an inhibitor (c.f. [26], p.22), the sensitivity of an enzyme substrate (c.f. [26], p.22), and multiple enzyme binding sites (c.f. [114], p.311). Also, it is important to note that the KR is not designed to support pharmacokinetic simulations. It has no representation of time or stoichiometry and it presumes no knowledge about drug dose, order of administration, the drug metabolizing phenotype of individual patients, or what specific adverse events could occur for any of its predictions.

3.5.2 The KR's Ontological Commitments Have Strengths and Limitations

Davis, Shrobe, and Szolovits offer the view that one of the most important contributions a knowledge representation can make is its *ontological commitments* [53] – "a set of decisions about how and what to see in the world." These commitments help reduce the overwhelming complexity of reality to a finite set of objects and relationships thought to be relevant by the representation's designers. They also limit the methods possible for expressing knowledge and the strategies available for inferring new knowledge. These limitations can be useful because they clarify what kinds of knowledge the representation can model and the soundness, completeness, and efficiency of its inferences. One of the KR's major ontological commitments is the use of definite clauses to represent drug-mechanism knowledge and the theory of how DDIs occur by metabolic inhibition. Here we note two significant implications

of this choice.

3.5.2.1 Definite Clauses Limit Expressivity

The KR and its inference machinery compose the *reasoning system* component of the DIKB shown in Figure 2.3 of Chapter 2. The *reasoning system* employs a novel use of a Truth Maintenance System to handle both default reasoning and the effects on inference of changes in the *knowledge-base* as new evidence causes assertions in the *evidence-base* to meet, or fail to meet, belief criteria (see Chapter 2). There are several families of Truth Maintenance Systems that have been developed over the years and we decided to use the simplest type, a Justification-based Truth Maintenance System (JTMS), as a test platform for exploring how evidence could be linked to a rule-based DDI theory.

Earlier in this chapter it was noted that the logical form of the KR's rules is restricted to being *definite clauses*; Horn clauses with exactly one positive literal (see Section 3.3.1). This is because, by design, the JTMS formalism is only capable of representing definite clauses [61]. While this constraint retains the benefits of Horn clauses mentioned previously (they are easy to write and read and there exists inference algorithms that are proven to perform *sound* and *complete* inference over them very efficiently [154]) they limit the *expressivity* of rules that can be entered into the KR. Some knowledge states are difficult, or impossible, to represent as definite clause's. For example, one might like to represent rule statements like the following in the KR:

```
IF NOT PRIMARY-TOTAL-CLEARANCE-MECHANISM-of ?x Biliary-Excretion AND
   NOT PRIMARY-TOTAL-CLEARANCE-MECHANISM-of ?x Renal-Excretion AND
   NOT PRIMARY-TOTAL-CLEARANCE-MECHANISM-of ?x Exhalation-Excretion
THEN
   PRIMARY-TOTAL-CLEARANCE-MECHANISM-of ?x Metabolic-Clearance
```

This rule has the logical form:

$$\neg A \wedge \neg B \wedge \neg C \Rightarrow D \tag{3.7}$$

It can be shown using truth tables that this logical form is equivalent to the following disjunction:

$$A \vee B \vee C \vee D \tag{3.8}$$

Since the disjunction contains four positive literals, it is not a definite clause (or even a Horn clause). Therefore, the rules like the one above cannot be represented in the KR.

A real example should illustrate how the JTMS requirement that all knowledge be represented as definite clauses constrains the kinds of rules that can exist in the KR. The JTMS does not allow the negation of logical statements such as `(?z is-SUBSTRATE-OF ?y)` so, we could not represent the following rule in the KR:

```
IF ?x INHIBITS ?y AND
   NOT (?z is-SUBSTRATE-OF ?y)
THEN
   NOT (?x INHIBITS-METABOLIC-CLEARANCE-of ?z via ?y)
```

The KR gets around this limitation by using predicates that represent the inverse state of other predicates in the rule-base along with additional rules that identify contradictions.[h] For example, the KR uses the first rule in Figure 3.8 to declare that a drug or drug metabolite, `?x`, that inhibits some enzyme, `?y`, will not reduce the clearance of another drug, `?z`, if `?z` is not a substrate of `?y`. The second rule in Figure 3.8 ensures that contradictory predicates do not enter the knowledge base by triggering a function called `CONTRADICTION` if some drug `?x` both is, and is not, a substrate of some enzyme `?y`.

The necessity of the work-around just mentioned is a limitation of representing knowledge using the JTMS and will likely make the KR more difficult to scale to other kinds of DDI mechanisms. Therefore, future expansion of the DIKB might require that the JTMS be replaced with a rule engine that allows more expressive logical statements. Fortunately,

[h]The reader might have noticed that the rule Figure 3.3 includes the statement `:TEST (NOT (EQUAL ?x ?z))`. This statement is actually a Lisp function that operates on the variables such as `?x` and `?y` but not logical statements such as `(?z is-SUBSTRATE-OF ?y)`.

```
IF ?x INHIBITS ?y AND
   ?z is-not-SUBSTRATE-OF ?y
THEN
   ?x does-not-INHIBIT-THE-METABOLIC-CLEARANCE-of ?z via ?y

IF ?x is-SUBSTRATE-OF ?y AND
   ?x is-not-SUBSTRATE-OF ?y
THEN
   (CONTRADICTION '(?x is-not-SUBSTRATE-OF ?y))
```

Figure 3.8: Two rules that the KR uses infer that two drugs will not interact by inhibition of a specific enzyme. The first rule says that a drug or drug metabolite, `?x`, that inhibits some enzyme `?y` will not reduce the clearance of another drug, `?z`, if `?z` is not a substrate of `?y`. The second rule is necessary to ensure that contradictory predicates do not enter the knowledge base.

more expressive families of Truth Maintenance Systems exist including Logic-based Truth Maintenance Systems that allow rules to be constructed using any propositional clause including non-definite clauses [61].

3.5.3 The KR's Reasoning System Does Not Track Uncertainty as it Performs Inference

There is often considerable uncertainty behind claims about a drug's mechanistic properties and this uncertainty affects the confidence that someone knowledgeable about drugs places on mechanism-based DDI predictions (see Section 2.1). It is important to note that the KR has no method for modeling or tracking uncertainty as it performs inference. Rather, a separate component of the DIKB called the `evidence-model` (Section 2.3.1) automatically selects assertions that meet user defined belief criteria and assumes that these assertions are certain from the user's perspective. If the user has selected belief criteria that represent full confidence in each assertion type, and each assertion the DIKB uses meets the user's belief criteria, then so will the KR's predictions.

This arrangement is useful for researching how evidence can be used to establish the certainty of drug-mechanism knowledge but it does not address how to handle assertions that do not meet belief criteria. In spite of this limitation, the experiment with the DIKB that

we will describe in Chapter 5 shows that the KR can make both accurate and well-supported novel predictions for an important class of DDIs. However, the experiment also shows that the KR's coverage of known interactions is incomplete. We believe that integrating an appropriate method for modeling and tracking uncertainty will help increase the system's coverage of known interactions but are postponing this investigation for future work.

3.5.4 Related Work

3.5.4.1 The KR and Drug-mechanism Tables

The KR is a rule-based representation of a strategy for reasoning about the potential occurrence of a metabolic DDI between unstudied drug combinations. The strategy combines knowledge about biochemical and physiological mechanisms of drug absorption, distribution, metabolism, and excretion with an understanding of the how drugs interact with each other to make mechanism-based predictions. Another tool that supports mechanism-based reasoning are tables published in paper or computer drug-interaction references that list the known metabolic properties of drugs. For example, a set of tables providing facts that clinicians can use to infer both metabolic inhibition and induction interactions can be found in a pocket reference for clinicians called *The Top 100 Drug Interactions* [81]. Also, the computer program Q-DIPS [34] provided a similar set of tables to help assist pharmaceutics researchers in selecting the optimal set of clinical trials needed to establish a new drug's safety profile.

While the current KR does not reason about metabolic induction, its representation of metabolic inhibition has several advantages over drug mechanism facts represented in tabular form.

- The KR directly infers DDIs from an explicit mechanism-based DDI theory while clinicians have to apply their own knowledge to information spread over two or more drug-mechanism tables.

- The KR can provide an estimate of the magnitude of a metabolic DDI – something not supported by any drug-mechanism tables that we are aware of.

- To the best of our knowledge, drug-mechanism tables tend to focus only on supporting interaction predictions and do not attempt to support non-interaction predictions. The KR model includes rules for predicting when a drug or drug metabolite pairs should not interact by inhibition of a particular enzyme.

- The KR uses its knowledge of the relationship between a drug and its active metabolites to infer how a change in concentration of a parent compound might effect a downstream product of metabolism. This kind of reasoning is not practical with drug-mechanism tables because it would require that the clinician synthesize information in several tables such as those listing the metabolites of each drug, the metabolites of each drug metabolite, and drug mechanisms.

3.5.4.2 Other Rule-based Representations

There is a long history of rule-based systems of various kinds being used to predict or detect drug interactions. One early system that, like the KR, was designed to support clinical reasoning, is that reported by Roach *et al* in 1985 [150]. Their system used rules and frames to organize pharmacologic information, including mechanisms, for retrieval by clinicians. Rule-based drug interaction systems have since become very common. For example, a report by Resetar *et al* describes their work with a commercial rule-base containing nearly 77,000 drug-drug interaction rules [148] and the Drug Ordering Decision Support System developed by Del Fiol *et al* imported drug-interaction rules from two different hospital systems [55].

The KR's focus is much smaller than rule-based systems in many contemporary clinical decision support systems or even some very early systems like Roach's. Though small, the KR performs a novel range of metabolic inhibition DDI predictions that includes inferring how a change in concentration of a parent compound might effect a downstream product of metabolism. The KR is a component of a larger system, the DIKB, that implements a rich representation of evidence for and against the drug-mechanism "facts" that the KR uses during inference (Chapter 2, Section 2.3). To the best of our knowledge, the DIKB is unique among rule-based systems that represent drug-mechanism knowledge because it not

only predicts mechanism-based metabolic drug-drug interactions and non-interactions but also can represent missing knowledge using *default reasoning* (Chapter 2, Section 2.2.2) and provide clear evidence support for each of the assertions it uses to arrive at its predictions.

3.5.5 Conclusion

The DIKB is a research system designed to evaluate novel informatics solutions to the challenges of representing and synthesizing drug-mechanism knowledge for post-market use. This chapter has described the computational representation of metabolic inhibition DDIs that the DIKB currently uses. The KR offers several advantages over simple tables of drug-mechanism facts (the tool that is probably most available to clinicians for inferring mechanism-based interactions) because it is able to infer concentration changes and the effect on downstream products of metabolism. The KR's method for representing DDI knowledge is easy to use and extend but has some limitations on what knowledge it can easily represent. Its model of metabolic inhibition makes interesting and accurate predictions in spite of a number of simplifying assumptions about the process of metabolic inhibition. We believe that the theory of metabolic inhibition has much in common with other mechanism-based DDI theories such as how DDIs occur by metabolic induction or transport protein modulation. These commonalities make it reasonable that the approaches used in the KR will be extendable to modeling DDIs that occur by other mechanisms.

Chapter 4

THE COLLECTION AND CLASSIFICATION OF DRUG-MECHANISM EVIDENCE

Our first effort to predict clinically relevant DDIs from drug-mechanism knowledge (Chapter 2, Section 2.1) convinced us that, for a knowledge resource with drug-mechanism knowledge to be of clinical use, it is essential that it explicitly link each of its drug-mechanism facts to their evidence support. In other words, we believe that a drug-mechanism knowledge-base should be able to clarify what clinical trials, *in vitro* experiments, or other forms of scientific evidence confirm or refute each of its assertions. One major benefit we expect to come from this arrangement is that expert users will be able to assess their confidence in a mechanism-based drug-drug interaction (DDI) prediction by viewing the evidence support for each drug property used to make the prediction. We also anticipate that this arrangement will make it possible to implement a set of computer-supported evidence maintenance processes that help keep a body of drug-mechanism knowledge up to date with current research.

Chapter 2 presented the design of the novel method for representing and computing with evidence that we implemented in the Drug Interaction Knowledge Base (DIKB). This chapter explores the DIKB's evidence representation method from a knowledge-base maintenance perspective. We begin with a brief summary of the method's goals and key assumptions along the way contrasting it with other biomedical informatics systems that link evidence to their assertions. We then relate our experience using the method to represent drug-mechanism evidence for 16 active ingredients and 19 active metabolites.

4.1 A Novel Method for Representing Evidence

The DIKB approach to evidence-modeling is the result of discussions among our research group while reaching consensus on the validity of some of the novel DDI predictions made by our pilot drug interaction system (Chapter 2, Section 2.1). The two drug experts in our

group were very knowledgeable about the drugs we had entered into the pilot system and had significant expertise on how to assess the clinical relevance of results from pharmacokinetic clinical trials. However, they had different opinions on the clinical relevance of results from *in vitro* drug-mechanism experiments. This led them to have different criteria for the kinds of scientific evidence that would convince them that a drug possessed certain mechanistic properties that would be measurable in humans (*in vivo*) at the drug's therapeutic doses. One expert felt that *in vitro* experiments of any kind had little utility for determining *in vivo* mechanistic properties while the other felt that, for some kinds of assertions, well-designed *in vitro* experiments were of some utility for making *in vivo* inferences.

Both drug experts could provide sensible justifications for their opinions about *in vitro* studies. For one expert, the clinical relevance of a drug-mechanism property derived from *in vitro* studies was always suspect until proven in a clinical trial because *in vitro* conditions do not accurately reflect the complex interplay of physiology, genetics, disease, and environment in humans. To this expert, the role of such studies was appropriate only in pre-clinical drug development where the results of such studies could be followed up by clinical trials. The other expert could provide examples where some drug-mechanism properties derived from *in vitro* studies seem to map to robust, clinically relevant, findings. This expert could define some situations where a well-designed *in vitro* experiment might be sufficient to support some drug-mechanism properties.

The various groups of users of any large-scale drug-mechanism knowledge-base will likely have similar disagreements about the kinds of scientific support that justify belief in drug-mechanism properties. This would be consistent with the fact that the science underlying drug-mechanism knowledge is dynamic and it can take years before a new experimental tool or method is understood well enough to define the range of inferences that can be made from its results. For example, many researchers have tried to develop a robust method for making quantitative *in vivo* DDI predictions solely from the results of *in vitro* experiments [92, 175]. Obach *et al* note how this approach to DDI prediction is feasible in principal but has only been partially successful so that no general method exists for making accurate, quantitative, estimates of the magnitude of a metabolic inhibition DDI using *in vitro* data [131]. Since the theory is still being developed, the *in vivo* relevance of data from *in vitro* experiments

is a matter for expert interpretation.

The DIKB's knowledge-representation method recognizes that experts can have sensible reasons for disagreeing on what evidence makes a drug-mechanism assertion believable. The new method assumes that it is possible to map a user's confidence in an drug-mechanism assertion to some arrangement of one or more abstract *evidence types*. These abstract evidence types are simply labels from a taxonomy of the kinds of evidence that might support or rebut a drug-mechanism assertion. The DIKB distinguishes between assertion *instances* and assertion *types*. An assertion instance is a specific fact about a particular object such as a drug or protein. For example, the generic `(X substrate-of Y)` is an assertion type whose instances might include `(carbamazepine substrate-of CYP3A4)` and `(s-warfarin substrate-of CYP2C9)`. Expert users map their confidence in drug-mechanism assertions by first defining combinations of evidence types from an evidence taxonomy that represent the kinds of evidence that might support or refute instances of each assertion type. They then rank the evidence-type combinations by the relative amount of confidence that they would have in an assertion instance of the given assertion type if it were supported by the types of evidence present in the definition.

We call rank-ordered combinations of evidence types *levels-of-evidence* (LOEs) and use them in the DIKB to provide customized views of a comprehensive body of drug-mechanism knowledge to different users. Defining LOEs is as simple as listing the evidence types that, based on expert opinion and/or scientific considerations, confer similar levels of justification to a given assertion type. One important principal is that any single LOE should not consist of the conjunction of two or more non-independent pieces of evidence. For example, a non-traceable statement in drug product labeling might repeat the same data that is present in a randomized pharmacokinetic study so it would be incorrect to say that an assertion the study might support is more justified when one combines these evidence items than when one considers them separately. To guard against this, expert users should never define an LOE that requires both a non-traceable statement *and* any evidence type that represents an actual study of some kind such as an experiment or clinical trial. Defining the LOE as

the disjunction of the respective types should not lead to the same error.[a]

There are two lists for every assertion instance in the DIKB's knowledge-base; one for evidence that supports the validity of the assertion and another for evidence that detracts. For each assertion type in the system, expert users define two, possibly identical, sets of LOEs. One for the types of evidence that can support an assertion type, the other for the types of evidence that refute it. They then select one LOE for each set of LOEs as *belief criteria*. A query of the DIKB's knowledge-base for valid drug-mechanism assertions will return only those assertions whose body of evidence *for* satisfies the *belief criteria* and whose body of evidence *against* does not satisfy *belief criteria*.

The DIKB's method for modeling and computing with evidence depends on an evidence taxonomy oriented toward confidence assignment. The evidence taxonomy must have sufficient coverage of all the kinds of evidence that might be relevant including various kinds of experiments, clinical trials, observation-based reports, and statements in product labeling or other resources. Another important requirement for the taxonomy is that users must be able to assess their confidence in each type either by itself or in combination with other types. The next section of this chapter examines these requirements in detail while considering the relevance of other biomedical evidence taxonomies to the task of representing drug-mechanism evidence in the way that the DIKB proposes.

4.2 Considerations for an Evidence Taxonomy Oriented Toward Confidence Assignment

Only a handful of biomedical informatics systems exist that attempt to label or categorize evidence; these include the PharmGKB's *categories of pharmacogenetics evidence* [152], Medical Subject Headings' *Publication Types* [27], Gene Ontology's *evidence codes* [65], and Pathway Tools' *evidence ontology* [106].

---

[a] A similar situation can occur when the same piece of evidence has been entered into the evidence-base more than once but under different identifiers. If an LOE requires two evidence items of a particular type and a different LOE requires only one, then the repeated evidence item could falsely increase in the amount of justification given to assertion that the items support. This situation would likely be rare and could be avoided by applying an algorithm that can identify repeated evidence items by methods other than comparing their unique identifiers.

4.2.1 PharmGKB's "Categories of Pharmacogenetics Evidence"

The PharmGKB is a Web-based knowledge repository for pharmacogenetics and pharmacogenomics research. Scientists upload into the system data supporting phenotype relationships among drugs, diseases, and genes. All data in the PharmGKB is tagged with labels from one or more of five non-hierarchical categories called *categories of pharmacogenetics evidence* [152]. The *categories of pharmacogenetics evidence* are different from the DIKB's evidence types because the latter represent specific *sources* of scientific inference such as experiments and clinical trials while the former are designed to differentiate the various kinds pharmacogenetic gene-drug *findings* by the specific phenotypes they cover (e.g. clinical, pharmacokinetic, pharmacodynamic, genetic, etc). In other words, the *categories* are oriented toward data integration rather than confidence assignment. The designers of the PharmGKB used this approach because they hypothesized that it would be capable of coalescing the results of a range of methods and study types within the field of pharmacogenetics into a single data repository that would be useful to all researchers in the field [7].

4.2.2 Medical Subject Headings Publication Types

One of the most used biomedical evidence taxonomies is the publication-type taxonomy that is a component of the Medical Subject Headings (MeSH) controlled vocabulary [43]. The MeSH controlled vocabulary is a set of over 20,000 terms used to index a very broad spectrum of medical literature for the National Library of Medicine's PubMed database (formerly MEDLINE). Each article in PubMed is manually indexed with several MeSH terms and additional descriptors including the article's publication type. The MeSH publication type taxonomy is designed to provide a general classification for the very wide range of articles indexed in PubMed. Hence, the taxonomy is very broad but relatively shallow. For example, publication types in the 2008 MeSH taxonomy [27] include types as varied as `Controlled Clinical Trial` and `Sermons` but only one type, `In Vitro`, that represents the wide range of experiments that are done with excised tissue.

In knowledge representation terms, the coverage by MeSH publication types of the ev-

idence types relevant for validating drug-mechanism knowledge is too *coarse-grained.* This is because the design of some *in vitro* experiments makes them better support for certain drug-mechanism assertions than others. For example, a recent FDA guidance to industry on drug interaction studies distinguishes three different *in vitro* experimental methods for identifying which, if any, specific Cytochrome P-450 enzymes metabolize a drug [26]. The three experiment types are different from the *in vitro* experiment type that the FDA suggests is appropriate for identifying if a drug inhibits a drug metabolizing enzyme. The next two sections will discuss two systems whose coverage of *in vitro* evidence is less coarse than MeSH publication types – the Gene Ontology evidence codes [65] and the Pathway Tool's evidence ontology [106].

4.2.3 Gene Ontology Evidence Codes

The Gene Ontology (GO) is a system of three separate ontologies defining relationships between biological objects in micro- and cellular biology [63]. GO is a consortium-based effort that has gained wide acceptance in the bioinformatics community because it supports consistent descriptions of the cellular location of a gene product, the biological process it participates in, and its molecular function. Authors of GO annotations are expected to specify an *evidence code* that indicates how a particular annotation is supported. GO evidence codes [65] are labels representing the kinds of support that a biologist might use to annotate the molecular function, cellular component, or biological process (s)he is assigning to a gene or gene product. GO has over a dozen evidence codes including codes that indicate that a biological inference is supported by experimental evidence, computational analysis, traceable and non-traceable author statements, or the curators' judgement based on other GO annotations.

In the DIKB, the user's confidence in an assertion rests on some arrangement of one or more evidence types. This means that the user must trust the validity of each instance of evidence that the system uses to meet the belief criteria without necessarily reviewing the evidence for herself. In contrast with these requirements, the authors of the GO evidence codes are very clear that they cannot be used as a measure of the validity a GO annotation

as indicated by their following statement:

> Evidence codes are **not** statements of the quality of the annotation. Within each evidence code classification, some methods produce annotations of higher confidence or greater specificity than other methods, in addition the way in which a technique has been applied or interpreted in a paper will also affect the quality of the resulting annotation [65].

This quote from GO evidence code documentation mentions two possible characteristics of GO evidence codes that preclude them from serving as a measure of the justification for biological annotations. First, GO evidence codes seem to represent evidence types that vary in terms of their appropriateness for justifying hypotheses. In our view, the evidence codes represent evidence *families* rather than evidence *types*. The distinction is that the kinds of evidence that an evidence type represents should be fairly homogeneous in terms of their appropriateness for justifying hypotheses. Like MeSH publication types, GO evidence codes are too coarse-grained for use as a tool for confidence assignment. Second, GO evidence codes do not address the fact that there are many possible problems with studies, experiments, author statements, and other types of evidence that can effect their validity. In other words, even if GO evidence codes were granular enough for decision support, the user would have to assess the quality of each evidence item directly or else place blind faith in the annotator's judgment.

4.2.3.1 The Need for Inclusion Criteria

This analysis of GO evidence codes indicates that there is at least one other dimension to biomedical evidence assessment besides the confidence that a particular group of methods or sources inspire in some hypothesis. A discussion of the *quality* of scientific evidence should help identify the necessary factors to consider when assessing scientific evidence.

The Agency for Healthcare Research and Quality, in their report assessing a substantial collection of systems for rating scientific evidence [177], defines the *quality* of a research study to be "the extent to which a study's design, conduct, and analysis have minimized

selection, measurement, and confounding biases" ([177], p.1). According to this definition there are three components of a study that contribute or detract from its quality - its design, how it is conducted, and how its results are analyzed. While it is possible to create meta-data labels that accurately reflect a study's design, it is intractable to abstract the full range of issues that affect a study's conduct and analysis.

Take for example a research study by Ford *et al* on the effect of fluoxetine on the clotting effect of warfarin [62]. The study's purpose was to see if fluoxetine would cause a pharmacodynamic interaction with warfarin. In this small correlation study, patients given fluoxetine for three weeks while on a low-dose of warfarin experienced no significant change in the amount of time it took for their blood to clot. Previous studies have shown that metabolism via the CYP450 enzymes is the primary clearance mechanism of warfarin. If our focus were on metabolic mechanisms, one possible interpretation of these results is that fluoxetine and its metabolite norfluoxetine do not effect the metabolic clearance of warfarin. However, as Ford and colleagues acknowledge in their discussion, fluoxetine and its metabolite both have long half-lives making it possible that a three week study was not adequate time to see the effect of a metabolic interaction between either of them and warfarin.

In terms of evidence, we might classify this study as an *uncontrolled drug-drug interaction study* but this will leave open a number of questions for the expert user who sees this evidence label such as:

- Was the dosing of both drugs sufficient to allow accurate measurements of a pharmacodynamic or pharmacokinetic effect?

- Were there certain attributes of the study's participants that could bias results? For example, were all participants very ill? Were they all elderly?

- What was the route of administration for both drugs?

Our approach to ensuring that the user can use evidence types to establish confidence assignments is to develop and consistently apply *inclusion criteria* for each type of evidence

in the DIKB. Inclusion criteria help ensure that all evidence within a collection meet some minimum standard in terms of quality. They are complimentary to evidence type definitions which should represent evidence classes that are fairly homogeneous in terms of their appropriateness for justifying hypotheses. The criteria are designed to help answer the kinds of methodology questions that expert users have when told that an evidence item is of a certain type.

4.2.4 The Pathway Tool's Evidence Ontology

One other currently used biomedical evidence taxonomy is found in the Pathway Tools system of pathway/genome databases (PGDBs) [106]. The Pathway Tools `evidence ontology` is both a computable evidence taxonomy and a set of data-structures designed so that PGDB maintainers can attach 1) the types of evidence that support an assertion in the PGDB, 2) the source of each evidence item, and 3) a numerical representation of the degree of confidence a scientist has in an assertion. The taxonomy component of the `evidence ontology` shares several of the types defined in GO evidence codes (Section 4.2.3) but adds a number of sub-types that define more specific kinds of experiments and assays than GO. The data-structure component of the "evidence ontology" enables PGDB maintainers to record the source of an evidence item, the accuracy of a given method for predicting specific hypotheses (e.g., the accuracy of an operon prediction algorithm, if it is known), and the scientist's confidence in a PGDB assertion given the full complement of evidence supporting an assertion.

PGDB users are presented with a visual summary of the kinds of evidence support for a given assertion in the form of icons representing top-level evidence-types from the Pathway Tools evidence taxonomy (e.g. "computational" or "experimental"). Users can click on the icons to view more detailed information of the specific evidence items represented by the top-level icons including the sources of each item and its specific evidence type. This approach enables Pathway Tools to provide an overview of the kinds of evidence support for an assertion so that users might make their own judgements on the amount of confidence they should have in a PGDB assertion.

4.2.4.1 Addressing Confirmation Bias

Pathway Tools evidence types serve a similar function as DIKB evidence types by helping users assess their confidence in knowledge-base assertions. However, an important distinction must be made between the evidence modeling approach of Pathway Tools and that of the DIKB. PGDB maintainers use the Pathway Tools `evidence ontology` to represent only *supporting* evidence while DIKB, maintainers use evidence types to represent both supporting and *refuting* evidence.

Griffin in his review of research in the domain *probability judgement calibration* [76] lists several robust findings from a considerable body of research exploring biases people have when estimating the likelihood of uncertain hypotheses. Among them is the finding that people tend to exhibit various forms of over-confidence when estimating the probability that some hypothesis is true. Among the possible explanations for this tendency put forth by some calibration researchers is that over-confidence is a result of *confirmation bias* – "...people tend to search for evidence that supports their chosen hypothesis" [76]. Under this model, confidence estimations should be more accurate when people consider situations where their hypotheses might not be true. Griffin reports that the results of some research studies are consistent with this model but that confirmation bias does not seem to be the sole cause of over-confidence during probability judgement.

We think that these results are relevant to representing drug-mechanism knowledge because one of the fundamental goals of a drug-mechanism knowledge-base should be to facilitate the maintenance of a coherent body of knowledge that has minimal bias. For every assertion in the DIKB knowledge-base there are two lists; one for evidence that contributes support to the validity of the assertion, another for evidence that detracts. Maintainers use an editorial board process to seek evidence both for, and against drug-mechanism assertions. The intent of this arrangement is to help knowledge-base maintainers avoid any tendency to collect evidence that only supports knowledge-base assertions and to help expert users create unbiased criteria for judging their confidence in the system's assertions.

4.2.5 Curator Inferences and Default Assumptions

One final remark on GO evidence codes and the Pathways Tools `evidence ontology`: in both systems, there is an evidence type called `Inferred by Curator` which curators use for knowledge they infer from other assertions or annotations in the respective systems [65, 106]. The `Inferred by Curator` evidence type provides a convenience to the maintainers of these systems by enabling them to quickly support some knowledge element based on information already in the system. The following example of how this occurs is paraphrased from an example given in the section titled "IC: Inferred by Curator" in the GO evidence code guide [65].

> The experiment described in (Noel et al. 1998) provides evidence that the protein encoded by the *S. cerevisiae* gene UGA3 has the function `specific RNA polymerase II transcription factor activity` (GO:0003704). The curator deduces from the functional annotation that UGA3 is located in the nucleus because 1) *S. cerevisiae* is a eukaryote, 2) RNA polymerase II is a nuclear polymerase, and 3) UGA3 is a gene product associated with RNA polymerase II. The curator annotates UGA3 with the cellular-component term `nucleus` (GO:0005634) and applies the evidence code `Inferred by Curator` to record the evidence support for the new annotation.

This example makes it apparent that evidence codes like `Inferred by Curator` are not evidence types at all, but rather a record of why a particular assertion exists within a knowledge-base. The DIKB requires a set of evidence types that users can use to judge their confidence in the system's assertions. An evidence code such as `Inferred by Curator` indicates that some curator, quite likely unknown to the user, inferred the knowledge that the code is linked to. In this situation users might apply the level of trust that they have for the knowledge source based on previous experiences. If they have found the knowledge source trustworthy, then they might consider the unknown curator's inference trustworthy as well. In such a case, the expert would be assessing their confidence in a knowledge-curation system rather than a scientific proposition. Alternatively, the expert might attempt to

explicitly trace the curators' judgement so as to decide for themselves if the inference was reasonable. This process might be straightforward as in the above example, or confusing depending on the complexity of the logic the used by the curator when making the inference in question.

In constructing the DIKB we have also found situations where it was desirable to assert some knowledge element based on our knowledge of other assertions in the system. As a trivial example, when evidence in the DIKB supports the assertion that some enzyme, E, is responsible for 50% or more of some drug or drug metabolite's total clearance from the body, then the system should also contain an assertion that more than 50% of a drug's clearance is by metabolism. A more complex example can be seen in the rules that the DIKB uses to infer a drug or drug metabolite's metabolic clearance pathway shown in Appendix B, Section B.2. In both of these cases the DIKB is able to use declarative rules and Truth Maintenance System (TMS) justifications (Chapter 2, Section 2.3.2) to automatically add the needed assertions to knowledge-base. The system's TMS links each automatically-inferred assertion to the assertions and rules from which it was inferred. Procedural code leverages the DIKB's TMS and *evidence-base* components to create a report showing the logic and evidence support for any automatically inferred assertions.

The advantage of the DIKB's approach becomes apparent when one considers that the construction and maintenance of a large knowledge-base is a collaborative effort. GO and the PGDBs in the Pathway Tools system require curation by many domain experts and we think it reasonable to expect that, in spite of the best of intentions, curators will sometimes make mistakes or not be entirely consistent in how they enter knowledge or assign evidence. Furthermore, as a knowledge-based system grows it becomes less tractable for curators to know all of the inferences supported directly by other knowledge in the system. In contrast, once a rule is added to the DIKB that makes an assertion based on other assertions present in the system, it will always be applied *consistently* and across *all* possible instances where it is applicable.

It turns out that there are other occasions where an evidence type like `Inferred by Curator` might seem applicable within the DIKB. The system's curators sometimes face situations where they are justified in entering an assertion without linking it to evidence.

Such an event can occur when the curator is unable to find evidence for an assertion or when (s)he decides that an assertion does not need to be justified by evidence. In both cases the curator can decide to enter it as a *default assumption. Default assumptions* are a special kind of assertion introduced in Section 2.3.2.1 of Chapter 2 that is considered justified by default, but that can be retracted either manually by curators or automatically by the system as it proceeds with inference.

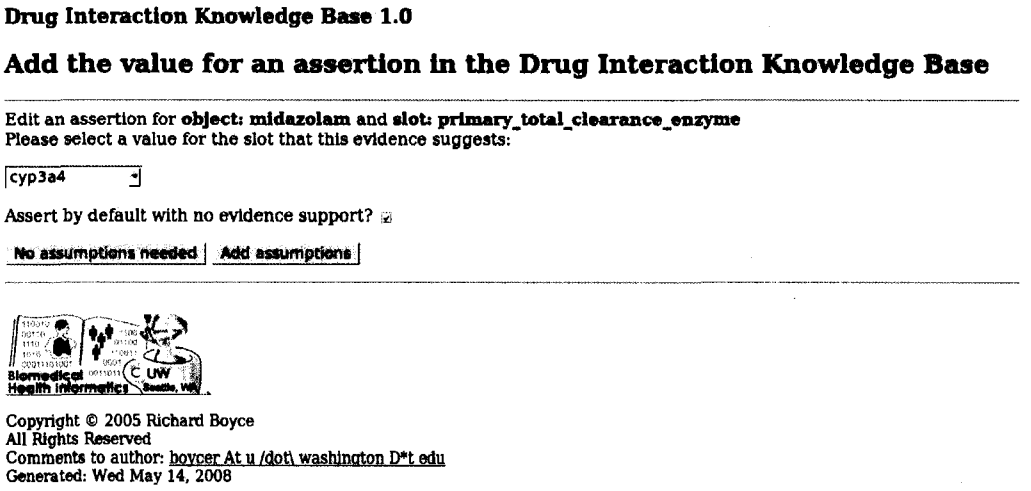

Figure 4.1: DIKB maintainers can specify that assertions be justified without any evidence support. These assertions are called *default assumptions.* This figure shows a user specifying that the assertion (`midazolam primary-total-clearance-enzyme CYP3A4`) is a default assumption. Curators can still link evidence items to assertions labeled *default assumptions* though the system will not asses if the evidence items meet belief criteria until the assertion is no longer justified by default.

For example, the current DIKB policy is that any enzyme that the FDA considers a drug or drug metabolite to be an *in vivo* probe substrate for should be labeled its *primary total clearance enzyme* (see Appendix C, Section C.13). The FDA suggests several drugs and drug metabolites that can serve as probe substrates for *in vivo* pharmacokinetic metabolism identification studies [26]. Since the guidance lists midazolam as an *in vivo* probe substrate of CYP3A4, a curator who sees this evidence would designate the assertion (`midazolam primary-total-clearance-enzyme CYP3A4`) a default assumption

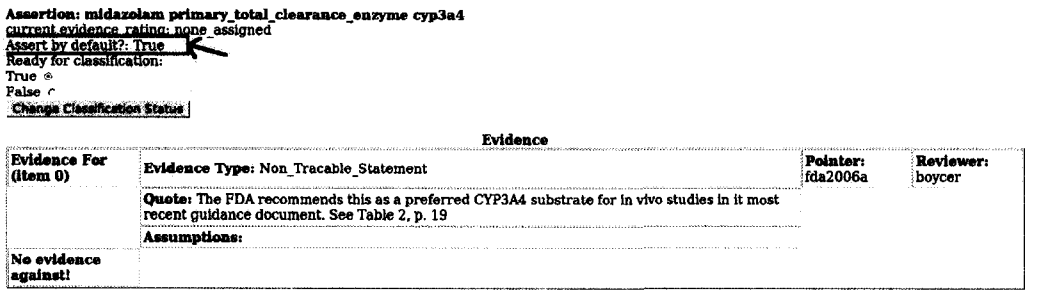

Figure 4.2: DIKB users can see when an assertion is a *default assumption* when they review its evidence support.

(Figure 4.1). DIKB users will see that it is a default assumption when they attempt to review its evidence support (Figure 4.2). If the user does not agree with this default assumption it is possible for the system to retract both the assumption and all assertions and inferences that depend on it for justification.

Table 4.1: A partial listing of the many sources for drug-mechanism knowledge. The DIKB curators searched these evidence sources for mechanism knowledge on 16 drugs and 35 active metabolites.

Knowledge source	*Example*	*Comments*
pharmaceutics and pharmacology text books	*Metabolically-Based Drug-Drug Interactions: Principles and Mechanisms* [112]	detailed drug-mechanism knowledge often with references
drug product labels	The NLM's DailyMed database[a]	a wide variety of information written by the drug's manufacturer; most statements are non-traceable; updated infrequently
drug information references	*Goodman & Gilman's the Pharmacological Basis of Therapeutics* [38]	quick source of basic pharmacokinetic data often with references
commercially licensed drug information databases	*The Metabolism and Transport Drug Interaction Database*[b]	searchable drug-mechanism knowledge selected by experts; all data refers to their original source
primary research article databases	PubMed[c]	comprehensive sources of indexed scientific evidence
regulatory guidelines	FDA Center for Drug Evaluation and Research Guidelines[d]	authoritative but often non-traceable consensus statements
continuing modules of education focusing on drugs and drug interactions	*Drug-Drug Interaction in the Elderly with Epilepsy: Focus on Antiepileptic, Psychiatric, and Cardiovascular Drugs* [44]	succinct, traceable, summaries of various drugs
unpublished pre-market and drug approval data	drugs@fda[e] or the drug manufacturer	the source of many non-traceable statements found in drug product labeling; often difficult to access
personal bibliographies belonging to drug experts	We built a search engine for the personal bibliography of one drug expert that spanned almost 30 years of work	a collection of high quality evidence specific to drug interactions

[a]http://www.dailymed.nlm.nih.gov

[b]http://www.druginteractioninfo.org/

[c]http://www.ncbi.nlm.nih.gov/PubMed/

[d]http://www.fda.gov/cder/guidance/

[e]http://www.accessdata.fda.gov/scripts/cder/drugsatfda/

4.3 An Appropriate Evidence Collection and Maintenance Process

The quality and coverage of the DIKB's drug-mechanism knowledge will depend a great deal on what process is used to collect and maintain evidence. The system requires an evidence collection and maintenance process that is geared toward building a coherent body of knowledge that has minimal bias and is up-to-date. This section examines the essential steps of a process that we believe meets these requirements.

4.3.1 Step One: Seek Evidence for and Against Each Relevant Assertion

There are number of sources of drug-mechanism evidence where curators might search including pharmaceutics and pharmacology text books, drug product labels, drug information references, primary research articles, regulatory guidelines, continuing modules of education, and unpublished pre-market studies (see Table 4.1). We will discuss later the specific search process that we used to collect evidence for 16 active ingredients and 19 active metabolites. We stress here the intent of the search for evidence – to acquire a minimally biased body of relevant evidence that can be evaluated *using LOEs and belief criteria.*

We propose that knowledge-base maintainers can reduce bias within the DIKB's evidence-base by collecting sufficient evidence to support *two* propositions for every drug-mechanism assertion; the first proposition being that the assertion is true, the second being that it is false. Using this approach, the curator would seek all sources of evidence that are relevant for supporting or refuting a drug-mechanism assertion and enter both kinds of evidence *even if* items contradict each-other. It is an open research question how effective this approach will be in reducing bias in a body of evidence. We will relate our experience applying this method toward the end of this chapter.

4.3.1.1 Quantitative Assertions

There are some attributes of pharmaceutical entities that are quantitative in nature such as the maximum systemic concentration that a drug will reach when administered at normal therapeutic doses. Quantitative assertion types are statistical in nature and the DIKB treats them differently than declarative assertion types such as `(X substrate-of Y)` .

- Curators collect only supporting evidence for *some* quantitative assertion types. For example, it makes little sense to collect evidence against an assertion about a drug's maximum concentration (C_{max}) at therapeutic doses because all drugs are sure to possess *some* C_{max}. A similar issue occurs with the measure of a drug or drug metabolite's bioavailability.

- Curators can collect evidence against a quantitative assertion type when the value does not exist for some entities within a class. For example, since only a sub-set of all possible drugs or drug metabolites will be found to inhibit a particular enzyme *in vitro*, it would make sense to include an experiment showing a zero or non-significant inhibition constant as evidence against an `inhibition-constant` assertion (Appendix C, Section C.15) for some pharmaceutical entity.

- If the body of evidence *for* a quantitative assertion satisfies the *belief criteria* and the body of evidence *against* the same assertion does not satisfy *belief criteria* then the DIKB exports to its knowledge-base a single value derived from the body of supporting evidence. The exported value can be numerical (e.g. "8.8", "0.99", etc) or qualitative representations of the value's magnitude (e.g. `LOW`, `MEDIUM`, `HIGH`).

The DIKB can derive values from a body of supporting evidence using a method that is customized to a particular assertion type. By default, the system exports the maximum value present in the list of supporting evidence. There are numerous potential alternative approaches including taking the minimum value, taking the simple average of the values provided by each item of supporting evidence, or weighting their values and combining them to derive a weighted-average. Our approach has been to derive values using methods that are pragmatic and that we think will increase the system's predictive sensitivity.

To clarify, consider that the DIKB labels a drug or drug metabolite an *in vivo* inhibitor for some drug metabolizing enzyme if the following relationship holds:

$$\frac{C_{max}}{K_i} > 0.1 \tag{4.1}$$

Where C_{max} is the maximum observed concentration of the inhibitor in patients at normal, therapeutic, doses and K_i is an inhibition constant derived from a well-designed *in vitro* enzyme-inhibition experiment involving the inhibitor.[b]

The system is programmed to take the *maximum* value found in the body of evidence supporting a drug or drug metabolite's C_{max} and the *minimum* value found in the body of evidence for a drug's K_i. This method should lead to more sensitive DDI predictions because the standard for qualifying as an *in vivo* inhibitor is lower than if the system chose the minimum value for C_{max} and the maximum value for K_i. The method is also pragmatic because C_{max} values are often based on pre-market studies that cannot be found in the literature but whose results are published in drug product labeling. In these cases, the simple average of C_{max} values should not be taken because this would assume that each study had a roughly equal number of participants but this information will be unknown to DIKB curators.

The DIKB can also map a numerical value to a qualitative representation of the value's magnitude using a function that is customized to a particular assertion type. For example, the system is programmed to map bioavailability values to the following discrete categories:

- `LOW`: $[0.0, 0.20]$
- `MEDIUM`: $(0.201, 0.50]$
- `HIGH`: $(0.501, 1]$

The motivation for choosing these categories is based on simple conjectures about what the maximum increase in AUC can be at various bioavailability levels. The AUC of a drug with a bioavailability of 50% should increase no more than 2-fold AUC if whatever is blocking the drug from entering systemic circulation is completely removed. The maximum possible magnitude increase at the 20% level is approximately 5-fold while there is no limit for drugs with bioavailability values near zero.

---

[b]Please see Chapter 3, Section 3.2.3.1, for further explanation of why the DIKB uses this technique.

A different mapping is used for a qualitative statement of the degree to which an active pharmaceutic ingredient is cleared from the body before entering systemic circulation. This "first-pass effect" can be important because some orally-administered drugs are heavily metabolized before, or while, passing from the intestine to the liver. This value is mapped to the following discrete categories:

- `LOW`: $[0.0, 0.50]$
- `MEDIUM`: $(0.501, 0.80]$
- `HIGH`: $(0.801, 1]$

The motivation for choosing these categories is based on simple conjectures about what the maximum increase in AUC can be at various first-pass-effect levels. For example, the AUC of a drug with a first-pass effect of 50% should increase no more than 2-fold increase if the first-pass effect is completely removed. The maximum possible magnitude increase at the 80% level is approximately 5-fold while there is no limit for drugs with first-pass effect values near 100%. Appendix C presents the method that the DIKB uses for each quantitative assertion type to derive and/or map its value.

4.3.1.2 Enough is Enough

Some assertions have numerous pieces of evidence to support them of many different types. For example, the assertion `(itraconazole inhibits CYP3A4)` is supported by at least three randomized clinical trials [136, 170, 181], drug product labeling [96], and an FDA guidance [26] (Appendix A, Table 2, p. 19). An interesting question in this case is – when should the curator stop collecting evidence for an assertion?

DIKB curators are charged with collecting a minimally biased body of relevant evidence that can be evaluated *using LOEs and belief criteria.* This goal is different than the task of acquiring sufficient evidence to prove that an assertion is true. To see how, let us consider once more how the DIKB uses LOEs. At the present time, they are used as rank-ordered grading scales for the kinds of evidence that are relevant to a particular assertion type.

Users must choose two, possibly identical, *belief criteria*; one from the LOEs for supporting evidence and one for the LOEs for refuting evidence. The *belief criteria* act as filters specifying the minimum evidence criteria that must be met for an assertion to be supported or refuted. Given this view of LOEs, DIKB curators should only collect evidence until it meets the user's belief criteria for each assertion type. Once evidence meets the belief criteria, any additional evidence will have no influence an whether the system asserts or retracts an assertion.

Such an approach makes sense when the LOEs and belief criteria are known to people collecting evidence and putting it into the DIKB. Unfortunately, it will not scale if the DIKB has multiple groups of expert users because each group will likely define a different set of LOEs and belief criteria. The DIKB will be more scalable if curators attempt to collect all available items of each evidence type that is relevant for supporting or refuting each assertion. To be practical, this approach will certainly require the use of advanced informatics tools to ease the curators task. Research in machine learning and artificial intelligence provides several examples of machine classifiers that accurately identify relevant articles from indexed research abstracts [153] and automatically extract biomedical relationships [145]. We think that DIKB curators should always make the final decision as to how to apply a given item of evidence but automated tools have the potential to greatly ease their task.

4.3.2 Step Two: Decide Each Evidence Item's Type Based on Definitions in the Evidence Taxonomy and Evaluate if an Evidence Item Meets the Inclusion Criteria for its Type

Once evidence has been collected, DIKB curators must tag *all* evidence items with a label specifying its type from the DIKB's evidence taxonomy. It is important to note that we define an *evidence item* to be a single research result within some *evidence source* (e.g. a specific journal article) rather than the evidence source itself. This distinction is necessary because a single evidence source might have multiple evidence items each of a different type. For example, a single journal article published by Jacobsen *et al* [94] reports the results of a variety of *in vitro* assays characterizing the metabolism of atorvastatin and its

metabolites by enzymes in the Cytochrome P450 family. The paper also reports the results of experiments identifying *in vitro* DDIs. There are several evidence items in this single source including some items that identify specific atorvastatin metabolites and others that indicate metabolic DDIs between atorvastatin and CYP3A4 inhibitors. Since the purpose and methods of the assays were different than those of the DDI experiments, the curator would classify the *evidence items* into different *evidence types* even though they are from the same *evidence source.*

At some point after an evidence item's type is classified, DIKB curators must decide if it meets inclusion criteria. Usually, this is as simple as reviewing the full-text source of the evidence item and ensuring that the item meets all the requirements for its evidence type. However, sometimes information is not available in the evidence source and the curator must rely on his or her judgement to decide if the evidence item meets criteria. For example, our inclusion criteria for certain *in vitro* enzyme inhibition experiments requires that an NADPH regenerating system be added to the enzyme system. The curator may decide that this requirement was met by a relatively recent report, even if it makes no mention of the addition of NADPH, since this procedure has become standard protocol for such experiments in recent years.

The previous example brings up an important point – in no sense are inclusion criteria tools for automating the evidence collection and curation process. The human curator is in the loop at all times and has full power to accept or reject an evidence item *even if it meets inclusion criteria.* Inclusion criteria help ensure that all evidence within a collection meet some minimum standard in terms of quality. As Section 4.2.3.1 states, it is intractable to abstract the full range of issues that affect a study's conduct and analysis using evidence types. Neither do we think it feasible that inclusion criteria will address all potential quality issues. So, an evidence item can proceed to the next step of the curation process if, in the curator's judgement, it meets inclusion criteria and there are no other quality issues that curator is aware of.

4.3.3 Step Three: Decide if There are any "Evidence-use Assumptions"

Interpreting the results of a scientific investigation as support for a particular assertion can sometimes require making conjectures that scientific advance might later prove to be invalid. If such conjectures are later shown to be false, it is important to re-consider how much support the scientific investigation lends to any assertion it was once thought to support. One unique feature of the DIKB is that it can represent the conjectures behind a specific application of evidence. These representations are called *evidence-use assumptions* in the DIKB. Chapter 2, Section 2.3.2.3, explains how these assumptions facilitate keeping knowledge in the system up-to-date and provides the technical details of how the DIKB models them. Here we briefly discuss how they are defined and what steps a curator takes to use them.

In our experience, *evidence-use assumptions* are an attribute of a particular class of evidence. For example, pharmacokinetic drug-drug interaction studies often involve administering a drug or drug metabolite (the precipitant) that is considered a selective inhibitor *in vivo* for some drug-metabolizing enzyme to study participants taking another drug (the object drug) that has reached a steady-state concentration. If the systemic concentration of the object drug increases, then it is strong evidence that the object drug's metabolic clearance depends significantly on the inhibited enzyme. However, this inference depends on the assumption that the precipitant has no measurable effect on any other clearance route of the object drug. This is an *evidence-use assumption* that applies to all pharmacokinetic drug-drug interaction studies using selective inhibitors.

DIKB maintainers attempt to define *evidence-use assumptions* for each new type that is added to the DIKB's evidence taxonomy. Like the previous example, these assumptions are written as general statements that apply to one or more evidence types. The maintainers add such statements to inclusion criteria documentation so that curators will know what specific assumption(s) should be declared when adding an item of evidence to the system. After curators have approved an evidence item, they identify assertions within the DIKB that match each specific *evidence-use assumption.* In many cases, a suitable assertion will not be present in the DIKB. If so, curators must add the new assertion to the DIKB then

link it as an *evidence-use assumption* for the evidence item.

4.3.4 Step Four: Enter a Representation of the Evidence Item into the DIKB

There are two ways that an evidence item can be entered into the DIKB; from within an interactive Python[c] session or using a simple Web interface. The Python interface is a powerful tool for querying the DIKB's evidence base but is not suitable for adding evidence because it requires that a curator be very familiar with the DIKB application program interface. The Web interface simplifies the task of evidence entry a great deal and is also useful for viewing evidence items and assertions within the DIKB. Appendix F shows how the user can enter evidence items and view then using the DIKB's Web interface.

Another advantage of the Web interface is that the system will perform several validation tests on a new evidence entry before it is stored in the DIKB's evidence-base. System tests include identifying if the evidence entry is redundant or has been rejected by DIKB curators as support or rebuttal for certain assertions. The system also checks if entering the item will create an evidence pattern that is indicative of circular reasoning by evidence-base curators. The next few sections provide more details on these validation checks.

4.3.4.1 Redundant Evidence Entry

A *redundant evidence entry* is defined as the exact same application of an evidence item as is currently existing in the system. The system tests for this occurrence by scanning the evidence for and against the assertion that the curator is attempting to link an evidence item to. If the new evidence item shares the same external document pointer (e.g PubMed identifier) *and* the same evidence position (support or refute) as an evidence item in the two bodies of evidence, then the system will warn the curator that they seem to be applying an evidence item redundantly. The system will not prevent the user from entering the evidence item because the system only knows that the same evidence *source* is being linked to an assertion more than once. The curator may have found multiple, independent, evidence items within the same source that should all be connected to the same assertion.

---

[c]http://www.python.org

4.3.4.2 Rejected Evidence Usage

An interesting fact is that it is possible for an evidence *source* to contain an evidence *item* that is not suitable as support or rebuttal for one assertion but perfectly acceptable for another. This can happen when an evidence source describes multiple studies or experiments and only a subset of them meet inclusion criteria. It can also happen when a curator considers only some of the results from a single study or experiment to be valid. In either case, we believe that it is important for evidence-base curators to keep track of every assertion that an evidence-source cannot support or refute. Doing so helps curators avoid redundant effort by alerting them when a particular evidence source contains items that should not be linked to a particular assertion.

The DIKB helps curators manage rejected evidence items by informing all participants in the evidence collection process when an evidence item has been rejected for some use. Each time curators reject an evidence item, they add an entry into a simple database indicating the item's source, a short description explaining why they rejected the item, and which assertion or assertions the item should never be used to support or rebut. The system scans the contents of this database each time a curator attempts to add a new evidence item to the evidence-base. Curators can still link an evidence item to any assertion that it has not yet been rejected from supporting or refuting. However, the system will not allow any rejected use of an evidence item.

4.3.4.3 Circular Support

Evidence-use assumptions were designed so that the DIKB could alert curators when one or more conjectures that a particular application of evidence depends on fail to meet belief criteria. They can also help identify a pattern, called a *circular line of evidence support*, that is indicative of fallacious reasoning by evidence-base curators. A hypothetical example should help clarify the kind of situation we are describing and its implications.

Let's say some evidence item, E, exists in the evidence-base as support for the assertion `(diltiazem inhibits CYP3A4)` and that `(simvastatin primary-clearance-enzyme CYP3A4)` is an *evidence-use assumption* for this application of E. In addition, assume that E

also acts as support for `(simvastatin primary-clearance-enzyme CYP3A4)` and that this *other* use of E depends on the validity of the assertion `(diltiazem inhibits CYP3A4)`. If there is no evidence against either assertion and E meets both assertions' supporting belief criteria, then the system will consider both assertions to be valid.

Figure 4.3 makes apparent the problem here – the conjecture, `(simvastatin primary--clearance-enzyme CYP3A4)`, is necessary for evidence item E to act as support for the assertion `(diltiazem inhibits CYP3A4)` but is being justified by the same evidence item, E, that *assumes the same proposition E is supposed to justify.* Intriguingly, the same unsound reasoning would be present even if evidence item E is being used to refute the assertion `(diltiazem inhibits CYP3A4)`. Neither kind of circular reasoning should be allowed in the DIKB's evidence-base.

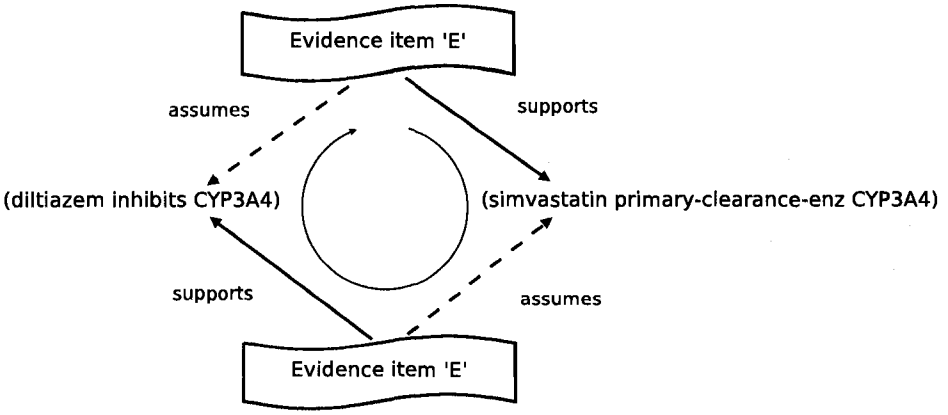

Figure 4.3: A circular line of evidence support that indicates circular reasoning within the evidence-base

The DIKB implements an algorithm that we have designed for detecting when an new evidence item would cause a *circular line of evidence support.*

> Let E be an evidence item that is being considered as evidence for or against some assertion, A. Assume that the use of E as evidence for or against A is contingent on the validity of one or more other assertions in the set $A_L = as_1, as_2, ..., as_n$. The set of assertions in A_L are the *evidence-use assumptions*

> for E. If E is currently being used as evidence for or against some assertion, as_i, in A_L and the use of E to support or refute as_i depends on the assumption A, then the use of E to support or refute A would create a *circular line of evidence support.*

The DIKB will not allow a curator to enter an evidence item that passes this test into its evidence base.

Circular reasoning might be present in the evidence-base *anytime* an *evidence-use assumption* is supported by the same evidence item that the assumption is linked to. We can create an algorithm to detect this form of circular evidence support by simplifying the previous algorithm.

> Let E be an evidence item and let the set $A_L = as_1, as_2, ..., as_n$ be the set of *evidence-use assumptions* for E. If E is currently being used as evidence for or against some assertion, as_i, in A_L, then circular reasoning might be present in the evidence-base.

The DIKB does not currently implement this algorithm in its validation tests but will in future versions.

4.3.5 Step Five: Computer-supported Evidence Maintenance Processes

Many drug-mechanism facts that we consider well-supported today will need revision to account for scientific progress. Hence, collecting and maintaining a drug-mechanism evidence-base should be an ongoing process by design. The DIKB, as a research system, was built for a specific set of experiments and so maintenance of its evidence-base is currently suspended. However, there are many tools and methods that could be useful for maintaining the evidence-base if work on it becomes active again. Section 6.2 of Chapter 6 discusses these possibilities in greater detail.

4.4 *Our Experience Using the Method to Represent a Body of Drug-mechanism Evidence*

We applied the evidence collection process and novel evidence representation method that the previous sections describe to the task of representing drug-mechanism evidence for six members of a family of drugs called HMG-CoA reductase inhibitors (*statins*) and ten drugs with which they are sometimes co-prescribed. Members of the statin drug family are very commonly used to help patients manage their LDL-cholesterol levels. While statins have a relatively wide therapeutic range, patients taking a drug from this class are at a higher risk for a damaging muscle disorder called myopathy if they take another drug that reduces the statin's clearance [86].

The sixteen drugs we chose are all currently sold on the US market, popularly prescribed by physicians, and have been the subject of numerous *in vivo* and *in vitro* pharmacokinetic studies. Many of them are known to be cleared, at least partly, by drug metabolizing enzymes that are susceptible to inhibition. DDIs that occur by metabolic inhibition can affect the concentration of active or toxic drug metabolites in clinically relevant ways (see Chapter 3, Section 3.2.2.3). For this reason, we also collected and entered drug-mechanism evidence for 19 active metabolites of the drugs we had chosen. Figure 4.4 lists the 16 drugs and 19 drug metabolites we chose to represent in the DIKB.

active ingredients:
atorvastatin, clarithromycin, diltiazem alprazolam, erythromycin, fluconazole, fluvastatin, itraconazole, ketoconazole, lovastatin, midazolam, nefazodone, pravastatin, rosuvastatin, simvastatin, triazolam

metabolites:
1'-hydroxymidazolam, 14-hydroxyclarithromycin, 4-hydroxyalprazolam, 4-hydroxymidazolam, 4-hydroxytriazolam, 6'-exomethylene-lovastatin, 6'-exomethylene-simvastatin, 6'-hydroxy-simvastatin, 6'-hydroxymethyl--simvastatin, 6'beta-hydroxy-lovastatin, N-demethyldesacetyl-diltiazem, N-demethyldiltiazem, N-desmethyl-rosuvastatin, alpha-hydroxyalprazolam, beta-hydroxy-lovastatin, beta-hydroxy-simvastatin, desacetyldiltiazem, ortho-hydroxy-atorvastatin, para-hydroxy-atorvastatin

Figure 4.4: The 16 drugs and 19 drug metabolites chosen for DIKB experiments

4.4.1 The DIKB Evidence Taxonomy

Appendix D shows the current DIKB evidence taxonomy. It contains 36 evidence types arranged under seven groupings representing evidence from retrospective studies, clinical trials, metabolic inhibition identification, metabolic catalysis identification, statements, reviews, and observational reports. We developed the taxonomy iteratively by collecting evidence for the drugs and drug metabolites shown in Figure 4.4, identifying the attributes of each evidence item, and deciding on evidence-type definitions.

We were able to incorporate some definitions from WordNet [121], MeSH [43], and NCI Thesaurus [54] but the majority of the taxonomy consists of new definitions. The structure of the taxonomy and granularity of its definitions is similar to the Pathway Tools' *evidence ontology* [106] however, the only definitions that the two resources share are for traceable and non-traceable author statements. Also, we deliberately excluded the "Inferred by Curator" evidence type present in the Pathway Tools' *evidence ontology* [106] and Gene Ontology's *evidence codes* [65] because we consider it to be a record of why a particular assertion exists within a knowledge-base rather than an evidence type (see Section 4.2.5).

We implemented the taxonomy in the OWL-DL language [46]; a *description logic* that provides a formal semantics for representing taxonomic relationships in a manner that can be automatically checked to ensure consistent classification. We used the Protégé ontology editor[d] to create the taxonomy and the RACER inference engine [80] to test it for consistent type definitions. We integrated the evidence taxonomy into the DIKB's structured vocabulary (Chapter 3, Section 3.3.4.1) the current version of which is available on the Web [35].

4.4.2 The DIKB Inclusion Criteria

We designed the set of seven inclusion criteria shown in Appendix E to compliment a sub-set of evidence type definitions from the DIKB's evidence taxonomy. Like the evidence taxonomy, we developed the inclusion criteria iteratively during the early stages of collecting evidence for the drugs and drug metabolites shown in Figure 4.4. This meant that changes

[d] http://protege.stanford.edu/

to inclusion criteria would sometimes require that evidence previously thought acceptable be discarded. The criteria became stable after making progress collecting evidence on several drugs. In their current form, shown in Appendix E, the seven criteria define the minimum quality standards for 21 evidence types in the taxonomy.

Six of the seven inclusion criteria apply to two or more evidence types within a sub-hierarchy of the evidence taxonomy. For example, the evidence taxonomy uses four evidence types to represent different kinds of clinical trials that test for pharmacokinetic DDIs. A single set of inclusion criteria apply to all four of the evidence types. Similarly, only one set of inclusion criteria apply to all eight of the evidence types that represent different *in vitro* experiments capable of identifying the specific enzymes responsible for a drug's metabolism. The one remaining criteria is specific to the evidence type representing non-traceable drug-label statements. This evidence type is a leaf node in a hierarchy of five types representing various traceable and non-traceable statements (see Appendix D).

There was a total of 12 evidence types for which we did not define inclusion criteria. Seven of these are general evidence types: *Statement*, *Non-traceable Statement*, *An observation-based report*, *An observation-based ADE report*, *A clinical trial*, *A DDI clinical trial*, and *A retrospective study*. We preferred to use more specific evidence types within the taxonomic sub-hierarchies that these five types reside in over the use of these general types and so defined inclusion criteria accordingly.

The other five evidence types with no inclusion criteria represent classes of evidence that we decided not to include for this study. We excluded the two types of author statements in the taxonomy (*A traceable author statement* and *A traceable drug-label statement*) because our evidence collection policy requires that curators retrieve and evaluate the evidence source that an author's statement references rather than rely strictly on the author's interpretation of that evidence source. We excluded the type *A retrospective population pharmacokinetic study* because we thought evidence of this class would be difficult to acquire and interpret. We also neglected to define inclusion criteria for the type *A retrospective DDI study* because we did not come across evidence of this type while defining inclusion criteria. Finally, the evidence collection process described in Section 4.4.3 did not include public adverse-event reporting databases so we did not define inclusion criteria for the type *An observation-based*

adverse-drug event report in a public reporting database.

4.4.3 The Evidence Collection Process

One informaticist and two drug-experts formed an evidence board that was responsible for collecting and entering all evidence into the DIKB. The informaticist, who was also the person who designed and wrote the DIKB software, led the evidence collection process which started in November of 2006 and ended in January of 2008. The process was iterative for the first few months while evidence types and inclusion criteria were being developed. The evidence board would choose a particular drug to model then collect a set of journal articles, drug labels, and authoritative statements that seemed relevant to each of the various drug-mechanism assertions defined in Appendix C. The evidence board would then meet together and discuss each evidence item and the issues that affected its use in the DIKB. By the time all members of the evidence board committed to using the evidence types and inclusion criteria shown in Appendices D and E the following evidence collection process had become routine:

1. The evidence board chose a particular drug to model.

2. The informaticist then received from each drug expert references to specific evidence sources that they thought would support or rebut one or more drug-mechanism assertions.

3. The informaticist did his own search of the literature that included seeking information in the various sources for drug-mechanism knowledge shown in Table 4.1. One of the drug experts was affiliated with the proprietary *The Metabolism and Transport Drug Interaction Database*[e] and performed searches of that resource then forwarded the results to the informaticist.

4. The informaticist would then summarize all evidence items from each source, classify their evidence types, and check if they met inclusion criteria. The evidence board

[e]http://www.druginteractioninfo.org/

would then meet and decide as a group whether each evidence item should enter the DIKB's evidence-base or be rejected as support or rebuttal for a specific assertion.

5. The informaticist would enter accepted evidence items into the DIKB using the DIKB's simple Web interface (Section 4.3.4). He entered rejected evidence items into a simple database used by the DIKB during evidence validation tests (Section 4.3.4.2).

4.4.4 The Current Evidence-base

Work on the evidence-base stopped in January 2008. In its present state it consists of evidence from 102 unique sources applied as evidence for or against 222 drug-mechanism assertions. In this section we will characterize some features of the evidence-base and the evidence items it includes while considering the goals of the DIKB's evidence representation method.

4.4.4.1 The Number of Evidence Items in the Evidence-base

In Section 4.3.2 we defined an *evidence item* to be a research result presented or referred to within a single *evidence source* (e.g. a specific journal article). Using this definition, each link from an evidence source to a single assertion in the DIKB represents a single evidence item. There are 272 links from individual evidence sources from the 222 assertions in the DIKB. Therefore, the current evidence-base consists of 272 evidence items taken from 102 evidence sources. However, some of the 272 evidence items are actually identical evidence items applied twice; once as support for some assertion and then again as refutation for the assertion's inverse. For example, the evidence-base uses an article presenting the results of several metabolism identification experiments [94] to support the assertion `(atorvastatin is-not-substrate-of CYP2C19)` and refute the assertion's inverse `(atorvastatin substrate-of CYP2C19)`.[f]

While it was right for the evidence board to apply supporting evidence for some assertions as refuting evidence for inverse assertions, only one instance of these evidence items should

[f]Chapter 3, Section 3.5.2.1, explains that inverse assertions are sometimes necessary in the DIKB because its knowledge representation formalism does not support the negation of predicates.

be included in the present analysis to prevent double counting identical evidence items. There are 15 evidence items that will be excluded from the remaining analysis for this reason; 11 evidence items linked as *supporting* evidence for 11 `is-not-substrate-of` assertions and four items linked as *supporting* evidence for four `does-not-inhibit` assertions. These 15 evidence items are also linked as *refuting* evidence for their respective inverse assertion (`substrate-of` and `inhibits`). Excluding these evidence items brings the current number of evidence items in the evidence-base to 257 items linked to 207 assertions.

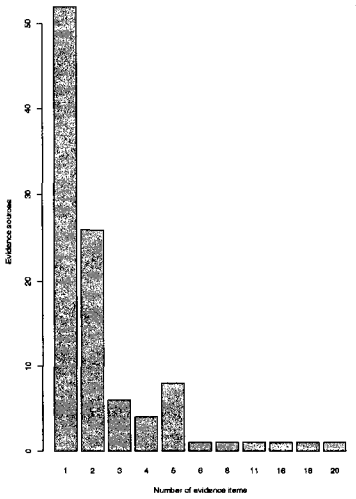

Figure 4.5: The number of evidence sources that provide a given number of evidence items in the current DIKB evidence-base. The figure shows that the evidence board found only one or two evidence items in the majority of the evidence sources. The number of evidence items in any one the 102 evidence sources ranges from one to 20 with an mean of 2.5

4.4.4.2 The Number of Evidence Items Found per Evidence Source

Figure 4.5 shows the counts of evidence items found in each evidence source. The figure shows that the evidence board found only one or two evidence items in the majority of the evidence sources that it reviewed. In fact, there is more than a ten-fold difference between the percentage of evidence sources that provide one evidence item (51%) and that of sources

providing 10 or more evidence items in (4%). Evidence sources that provide an unusually large number of evidence items include one article publishing the results of a number of metabolism identification experiments by Williams *et al* [178] (20 evidence items) and an single FDA guidance to industry on pre-market drug-interaction studies [26] (18 evidence items).

4.4.4.3 The Classification of Evidence within the Evidence-base

Another interesting fact is that the evidence board used only one-third of the 36 types in the evidence taxonomy to classify all the 257 non-redundant evidence items. Section 4.4.2 explains that five evidence types were not used because of specific evidence collection policies. This leaves 19 evidence types that were never used to classify any evidence item. Several of the types were not used because no acceptable evidence in their class could be found. For example, even though the evidence board collected numerous case reports describing adverse drug events in patients taking two or more of the drugs in our study, none of the five observation-based evidence types were entered into the system. This was because none of the reports measured the systemic concentrations of the purported object drug in a way that would support or refute an assertion about its metabolic properties.[g]

The 12 evidence types that were used to classify evidence items are shown in Table 4.2 along with the number of supporting or refuting evidence items each type was assigned to. One can calculate from Table 4.1 that evidence types assignments in the current DIKB are biased toward clinical trial types (42%) followed by a relatively similar distribution of *in vitro* studies (27%) and non-traceable statements in drug labeling and FDA guidance documents (30%).

The distribution of evidence types shown in Table 4.1 is partially a result of the evidence board's bias towards collecting certain kinds of evidence. For example, the evidence board's informaticist usually looked for evidence in drug-product labeling first because it was the most easily accessible. Not surprisingly, the most commonly assigned evidence type was

[g]We will discuss in Chapter refchap:novel-ddis, Section 5, that a significant proportion of these case reports qualified as evidence of a drug-drug interaction and were later used to explore the feasibility of novel DDI predictions made by the DIKB.

A non-traceable drug-label statement. Similarly, the informaticist generally searched for randomized DDI studies before searching for non-randomized ones. If one or two randomized trials could be found involving a drug or drug metabolite, he would generally not make an effort to seek non-randomized trials. As a result, the most frequently assigned clinical trial evidence type in the current DIKB is *A randomized DDI clinical trial.*

Generally-defined evidence types were often used when an evidence item did not fit one of the more specific evidence-types within a particular sub-hierarchy. Eleven of the twelve types shown in Table 4.2 are sub-types of some other, more general, evidence types within the greater evidence taxonomy (Appendix D). This indicates that the evidence taxonomy was broad enough to classify most of the drug mechanism evidence that the evidence board found into fairly specific categories. One exception to this trend was the most general *in vitro* evidence type *A drug metabolism identification experiment* that is assigned four times in the current DIKB evidence-base. All four uses of the evidence type were to classify metabolite identification experiments that could not be classified using the more specific types within the hierarchy.

It is clear from Table 4.2 that some evidence types are present in the evidence-base much more often than other types even though the experiments they represent have relatively similar purposes. For example, the evidence-base has almost eight-fold more evidence items of the type *A CYP450, human microsome, metabolic enzyme inhibition experiment* then the type *A CYP450, recombinant, metabolic enzyme inhibition experiment* even though the purpose of both kinds of experiments is to test a drug or drug metabolite's ability to inhibit some enzyme *in vitro.* Similarly, the system has three-fold more items of the type *A CYP450, recombinant, drug metabolism identification experiment with possibly NO probe enzyme inhibitor(s)* than the type *A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors* even though both experiments attempt to identify the CYP450 enzymes capable of metabolizing a drug or drug metabolite *in vitro.* It is possible that their relative occurrence in the DIKB is a reflection of bias in the research literature that causes one experiment type to be performed or published more frequently than another since the evidence board had no preference in searching for these types. Our methods are not capable of answering this question definitively because we do not claim to

have collected the exhaustive set evidence within any of our evidence classifications.

4.4.4.4 Observed Biases

Tables 4.3 and 4.4 show the relative distribution of evidence types stratified by each assertion type in the DIKB. These tables make it apparent that the distribution of evidence types among individual assertion types is much more diverse than that of the evidence-base as a whole. For example, all 15 evidence items linked to `inhibition-constant` assertions are from *in vitro* evidence types while no *in vitro* evidence is currently linked to a `maximum-concentration` assertion. Likewise, two-thirds of the evidence items linked to `maximum-concentration` assertions are instances of clinical trial types while the one-third are instances of non-traceable statement types. Approximately the opposite distribution of evidence types is present in items linked to `bioavailability` assertions (38% clinical trial types and 62% non-traceable statements). In comparison, evidence items across all assertion types are biased towards clinical trial types (42%) followed by a relatively similar distribution of *in vitro* studies (27%) and non-traceable statements in drug labeling and FDA guidance documents (30%).

Even though the evidence board attempted to collect both supporting and refuting evidence for each assertion, the current evidence-base is strongly biased toward supporting evidence. 82% of the 102 evidence sources provide evidence items that are used strictly as support for one or more assertions. In comparison, only 3% of sources provide strictly refuting evidence items and only 15% of sources provide both supporting and refuting evidence items. Of the 257 non-redundant evidence items, 229 (89%) support, and 28 (11%) refute, some drug mechanism assertion. In terms of the 20 assertions types that the DIKB currently represents, only four (20%) have any assertions with refuting evidence; `substrate-of`, `inhibits`, `increases-auc`, and `primary-metabolic-enzyme` (see Table 4.4).

Table 4.2: The evidence board used only one-third of the 36 types in the evidence taxonomy to classify all the 257 non-redundant evidence items. The 12 evidence types are shown in this table along with the number of supporting or refuting evidence items each type was assigned to. Indented evidence types are sub-types of the type in the previous row.

Clinical trial types

Evidence type	*Evidence for*	*Evidence against*
A pharmacokinetic clinical trial	31	0
A genotyped pharmacokinetic clinical trial	5	1
A randomized DDI clinical trial	49	11
A non-randomized DDI clinical trial	8	0
A parallel groups DDI clinical trial	4	0
Total	97	12

***in vitro* experiment types**

Evidence type	*Evidence for*	*Evidence against*
A CYP450, recombinant, metabolic enzyme inhibition experiment	2	0
A CYP450, human microsome, metabolic enzyme inhibition experiment	13	2
A drug metabolism identification experiment	4	0
A CYP450, recombinant, drug metabolism identification experiment with possibly NO probe enzyme inhibitor(s)	31	6
A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors	8	4
Total	58	12

Non-traceable statement types

Evidence type	*Evidence for*	*Evidence against*
A non-traceable, but possibly authoritative, statement	22	0
A non-traceable drug-label statement	52	4
Total	74	4

Table 4.3: The distribution of evidence types among evidence items used to ***support*** drug mechanism assertions.

Qualitative assertion types

Assertion type	*Assertions*	*Defaults*	*Evidence Items*	*Percent of evidence items by type*		
				clinical trial	*in vitro*	*non-traceable*
increases-auc	41	0	43	93	0	7
substrate-of	28	0	29	31	48	21
has-metabolite	20	0	27	19	37	44
controls-formation-of	16	0	17	0	94	6
primary-total-clearance-enzyme	12	4	12	58	0	42
inhibits	11	0	11	91	0	9
primary-total-clearance-mechanism	10	0	14	21	0	79
does-not-permanently-deactivate-catalytic-function[g]	8	8	0	0	0	0
in-vitro-selective-inhibitor-of-enzyme	8	8	8	0	0	100
in-vitro-probe-substrate-of-enzyme	5	5	5	0	20	80
sole-PK-effect-alter-metabolic-clearance	5	5	0	0	0	0
polymorphic-enzyme	4	4	0	0	0	0
in-viVo-selective-inhibitor-of-enzyme	4	4	4	0	0	100
primary-metabolic-clearance-enzyme	3	0	3	0	33	67
permanently-deactivates-catalytic-function[g]	1	1	1	0	100	0
Total	176	39	174			

Quantitative assertion types

Assertion type	*Assertions*	*Defaults*	*Evidence Items*	*Percent of evidence items by type*		
				clinical trial	*in vitro*	*non-traceable*
maximum-concentration	11	0	24	67	0	33
inhibition-constant	9	0	15	0	100	0
bioavailability	8	0	13	38	0	62
fraction-absorbed	2	0	2	100	0	0
first-pass-effect	1	0	1	0	0	100
Total	31	0	55			

[g] The inverse assertions `permanently-deactivates-catalytic-function` and `does-not-permanently-deactivate-catalytic-function` do not share any of the same evidence items so are included in the analysis (see Section 4.4.4.1)

Table 4.4: The distribution of evidence types among evidence items used to ***refute*** drug mechanism assertions.

Qualitative assertion types

Assertion type	*Assertions*	*Defaults*	*Evidence Items*	*Percent of evidence items by type*		
				clinical trial	*in vitro*	*non-traceable*
increases-auc	41	0	12	75	0	25
substrate-of	28	0	11	0	91	9
has-metabolite	20	0	0	0	0	0
controls-formation-of	16	0	0	0	0	0
primary-total-clearance-enzyme	12	4	0	0	0	0
inhibits	11	0	4	50	50	0
primary-total-clearance-mechanism	10	0	0	0	0	0
does-not-permanently-deactivate-catalytic-function[g]	8	8	0	0	0	0
in-vitro-selective-inhibitor-of-enzyme	8	8	0	0	0	0
in-vitro-probe-substrate-of-enzyme	5	5	0	0	0	0
sole-PK-effect-alter-metabolic-clearance	5	5	0	0	0	0
polymorphic-enzyme	4	4	0	0	0	0
in-viVo-selective-inhibitor-of-enzyme	4	4	0	0	0	0
primary-metabolic-clearance-enzyme	3	0	1	100	0	0
permanently-deactivates-catalytic-function[g]	1	1	0	0	0	0
Total	176	39	28			

Quantitative assertion types

Assertion type	*Assertions*	*Defaults*	*Evidence Items*	*Percent of evidence items by type*		
				clinical trial	*in vitro*	*non-traceable*
maximum-concentration	11	0	0	0	0	0
inhibition-constant	9	0	0	0	0	0
bioavailability	8	0	0	0	0	0
fraction-absorbed	2	0	0	0	0	0
first-pass-effect	1	0	0	0	0	0
Total	31	0	0			

[g] The inverse assertions `permanently-deactivates-catalytic-function` and `does-not-permanently-deactivate-catalytic-function` do not share any of the same evidence items so are included in the analysis (see Section 4.4.4.1)

4.4.4.5 Default Assumptions

The evidence-board labeled approximately one-fifth (39) of the assertions in the DIKB *default assumptions* (see Table 4.3). Nearly half (17) of the *default assumptions* were entered because of DIKB policies regarding information in FDA guidances (see Appendix C, Sections C.8, C.9, C.11, and C.13). The 17 assertions are linked to evidence items that refer to the FDA guidance that prompted the evidence-board's decision to make them *default assumptions.*

Another 17 assertions are labeled *default assumptions* but have no evidence items linked to them at all. Five of these were entered by the evidence board because of actions specified in the inclusion criteria for pharmacokinetic DDI studies (Appendix E, Section E.4). The remaining 12 were entered without evidence based on the knowledge of one or more members of the evidence-board. These were entered as *default assumptions* out of convenience with the intent that a DIKB curator would seek evidence for and against the assertions at a later time.

4.4.4.6 Rejected Evidence

The evidence board rejected 74 evidence items for use as support or rebuttal for at least one assertion in the DIKB. A partial listing of the evidence items with rejected use-cases are shown in Table 4.5 along with an explanation for why the rejection occurred. Nearly half (33) of the rejections involved observational case reports that the evidence-board felt could not support or refute a specific pharmacokinetic DDI. Often this was because the reports contained missing data or there were viable alternate explanations for the reported adverse event besides the occurrence of a pharmacokinetic interaction. The other 41 rejections involved a variety of *in vitro* and clinical trial study types that did not meet explicit inclusion criteria or had other flaws detected by the evidence-board.

The rejecting evidence included clinical trials with too few participants or that used drug dosing schemes that the board considered inappropriate for inferring DDIs. Several *in vitro* experiments were rejected because they used animal models and the DIKB requires data derived from humans. Other experiments were rejected because their methods were

not considered sufficiently accurate (e.g. immuno-chemical quantization for determining the fraction of a drug cleared by a particular enzyme) or used novel microsomal systems (e.g. intestinal microsomes). A couple of studies were rejected because the evidence board found the publications that reported them too unclear about important details.

There were three evidence sources for which the evidence board accepted some evidence items and rejected others. One article in particular [94] had 16 accepted evidence items and only one rejection. In this case, the evidence-board rejected the item because its results were specific to forms of atorvastatin and its metabolites that the DIKB does not represent[h]

Table 4.5: A partial listing of the 74 evidence items that the evidence board rejected for use as support or rebuttal for at least one assertion in the DIKB.

Source	*Rejected for/against assertion(s)*	*Comments*
[97]	all assertions involving diltiazem	The study did not use in vitro selective inhibitors
[93]	all assertions involving lovastatin	The study used an intestinal microsome system; these microsomal systems are not currently accepted in the DIKB
[173]	simvastatin's primary total-clearance mechanism is metabolic clearance	The study relied on an animal model; the DIKB only accepts human-based evidence
[161]	erythromycin increases the AUC of lovastatin	The lovastatin level was drawn after the patient developed renal failure plus, the patient was on numerous concomitant medications
[77]	clarithromycin increases the AUC of lovastatin	The data provided in this case report not sufficient for inferring a PK interaction (drug levels were not taken before and after challenge)

4.4.4.7 Evidence-use Assumptions

Four different DIKB inclusion criteria contain statements notifying curators of specific *evidence-use assumptions* (Section 4.3.3) that they should declare when adding evidence

[h] The DIKB does not distinguish between the lactone and acid forms of atorvastatin and its metabolites; the results of one of the experiments conducted by Jacobson *et al* [94] were lactone-specific.

of a certain type to the system (Appendix E, Sections E.3, E.4, E.6, and E.7). The evidence board followed these statements when entering them into the DIKB. As a result, 58 (23%) of the evidence items in the current evidence-base have at least one *evidence-use assumption.* Table 4.6 provides a sample of five of these evidence items. Fifty-three evidence items are linked to one *evidence-use assumption* and five evidence items are linked to two bringing the total number of *evidence-use assumptions* in the current DIKB to 63. Only twenty-three (11%) of the 207 assertions in the DIKB comprise all 63 *evidence-use assumptions.* The number of times the evidence board used any specific assertion as an *enabled assumption* ranged from once to nine times (mean: 2.7, median: 1).

Table 4.6: A sample of five of the 58 evidence items in the DIKB's evidence-base that were entered with *evidence-use assumptions*

Source	*Assertion evidence is supporting*	*Evidence type*	*Evidence-use assumption(s)*
[110]	diltiazem inhibits CYP3A4	*A randomized DDI clinical trial*	triazolam's primary-total-clearance enzyme is CYP3A4
[129]	simvastatin is a substrate-of CYP3A4	*A randomized DDI clinical trial*	itraconazole is a selective inhibitor of CYP3A4 *in vivo*
[138]	alprazolam is a substrate-of CYP3A5	*A genotyped pharmacokinetic clinical trial*	CYP3A5 has multiple drug-metabolizing phenotypes
[151]	clarithromycin is a substrate-of CYP3A4	*A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors*	ketoconazole is a selective inhibitor of CYP3A4 *in vitro*
[128]	lovastatin's primary-total-clearance enzyme is CYP3A4	*A randomized DDI clinical trial*	1) itraconazole's sole PK effect (in this study) is to alter the metabolic clearance lovastatin 2) itraconazole is a selective inhibitor of CYP3A4 *in vivo*

4.4.5 Discussion of the Current Evidence-base

The DIKB's current evidence-base demonstrates that the DIKB's evidence collection and representation methods can be applied to a significant body of drug mechanism evidence. Construction of the evidence-base was not trivial; the software for the evidence-base and its simple Web interface required several months of part-time programming effort by the evidence-board's informaticist to build. Also, it took the three-person evidence-board 13 months of part-time effort to collect, evaluate, and enter the contents of the evidence-base. As we will discuss in Chapter 5, these efforts were fruitful because the evidence-base presented in this chapter was used to accurately predict known and novel DDIs for an important set of drugs.

4.4.6 Limitations

One limitation of the evidence-base is that the evidence-board only searched for drug mechanism knowledge among the sources listed in Table 4.1. The evidence-board did not search for evidence in the EMBASE[i] or Web of Science® publication database. It is possible that these resources might have contained important evidence that is now missing in the DIKB.

Another limitation is that we did not independently evaluate how accurately and consistently the evidence-board classified evidence. The evidence board employed some internal consistency checks such as reviewing each evidence item multiple times before it was entered into the DIKB and using double-entry methods to track an evidence item's progress through the evidence collection process. However, it would be desirable to acquire independent verification that the evidence-board's classifications were accurate and consistent across all entries.

4.5 Conclusion

This chapter has explored the DIKB's evidence representation method from a knowledge-base maintenance perspective and presented the results of applying the method to representing drug-mechanism evidence for 16 active ingredients and 19 active metabolites. The

[i] http://www.embase.com/

DIKB's current evidence-base integrates drug mechanism evidence from a variety of sources including *in vitro* experiments, clinical trials, and statements from drug product labels. The evidence in the DIKB is of high quality because each evidence item has been screened to ensure that it meets an explicit set of quality criteria. A novel set of validation tests were used to ensure that the current evidence-base has no redundant entries, rejected evidence items, or applications of evidence that are the result of circular reasoning by evidence-base curators.

Every evidence item in the DIKB is labeled by its type from a novel evidence taxonomy. We designed the taxonomy so that each type represents scientific knowledge sources that are fairly homogeneous in terms of their appropriateness for justifying hypotheses. Expert users should be able to assess their confidence in the system's assertions relatively quickly once they are familiar with evidence type definitions and their associated inclusion criteria. We hypothesize that this process should involve less effort and be more consistent than requiring the expert to review the original sources for each evidence item. We will show in in Chapter 5 how we were able to use evidence types to define specific belief criteria for each assertion type in the system and have the DIKB automatically determine our confidence in its DDI predictions. Chapter 5 will also report on how the evidence-base presented in this chapter was used to accurately predict a subset of DDIs for an important class of drugs using only knowledge of drug-mechanisms.

Chapter 5

AN EXPERIMENT WITH LEVELS-OF-EVIDENCE AND BELIEF CRITERIA

When drug experts define and rank *levels-of-evidence* (LOEs) or choose *belief criteria*[a] they are making subjective judgements about the inferential force of an abstract body of evidence. An important question is whether the experts' choices have any relationship to the empirical prediction accuracy of the system. Our hypothesis is that the system will make fewer, but more accurate, predictions when using *belief criteria* that inspire complete confidence in a drug expert than when using criteria that the expert believes to be less trustworthy. As the expert user relaxes their criteria for including assertions, the DIKB should predict a larger number of true interactions; possibly at the expense of also making more false predictions. In terms of decision support, the DIKB's *sensitivity* should go up but its *specificity* should go down. Similarly, if the user tightens their criteria, the system should make fewer predictions and have a lower sensitivity but its specificity should increase. These features would be very desirable for supporting clinical decision making because the system's prediction performance could be customized to perform at the most optimal level possible given the contents of its evidence-base.

This chapter recounts an experiment that we conducted to characterize the effect of varying *belief criteria* on the system's accuracy and coverage of DDIs present in a reference set of interactions and non-interactions. The experiment was designed to answer the following research questions:

1. What is the DIKB's accuracy and coverage of reference set interactions and non-interactions when using a particular set of evidence items and expert-defined LOEs and *belief criteria*?

[a]Chapter 4, Section 4.1 provides a definition of *levels-of-evidence* and *belief criteria* in its explanation of the DIKB's novel approach to evidence-modeling.

2. Does changing the LOEs that are selected as *belief criteria* alter the systems prediction accuracy or coverage?

3. Do computational experiments imply a particular *belief criteria strategy* that optimizes the system's prediction performance using a particular set of evidence items and expert-defined LOEs? If so, what is the optimal strategy and how is the system's accuracy and coverage of reference set interactions and non-interactions when using the optimal strategy different from when it uses expert-defined *belief criteria*?

4. When the system is using the optimal *belief criteria strategy*, how does its accuracy and coverage of reference set interactions and non-interactions compare with statements in drug product labeling?

5.1 Methods

An evidence-board consisting of two clinician drug-experts and one informaticist from our research group collected sufficient evidence on the pharmacokinetic drug properties of 16 drugs and 19 drug metabolites to perform this experiment. Chapter 4 describes in detail the methods used to collect, classify, and enter evidence it into the DIKB. Figure 4.4 of that chapter lists the specific drugs and drug metabolites we chose to represent in the DIKB. Once the evidence-base was complete except for minor revisions, the evidence board attempted to identify all known pairwise *metabolic inhibition* interactions and non-interactions between 35 pharmaceutical entities in the DIKB's evidence-base. The evidence board's intent was to use the interactions and non-interactions that they found as a *validation set* for determining the accuracy and coverage of the DIKB's DDI predictions. The DIKB predicts interactions using knowledge of drug mechanisms and a theory of how interactions occur by metabolic inhibition (Chapter 3). Therefore, the method that the evidence board used to confirm interactions and non-interactions had to be independent from the the one used by the DIKB to predict DDIs.

5.1.1 Criteria for Confirmed Interactions and Non-interactions

The evidence board considered a metabolic inhibition *interaction* between any pair of pharmaceutical entities in the DIKB to be independently confirmed if *any* of following criteria were satisfied:

1. A pharmacokinetic DDI study provided data showing a *statistically significant* (see below) increase in the *Area Under the concentration-time Curve* (AUC) of the study's object drug or drug metabolite in the presence of the precipitant's drug or drug metabolite.

2. An observation-based case report provided data showing a measurable increase in the systemic concentration of a drug or drug metabolite in the presence of another drug or drug metabolite and the evidence board could find no viable alternate explanations for the observed increase.

The evidence board considered a metabolic inhibition *non-interaction* between any pair of pharmaceutical entities in the DIKB to be independently confirmed if *all* the following criteria were satisfied:

1. A pharmacokinetic DDI study provided data showing *no statistically significant* (see below) increase in the AUC of the study's object drug or drug metabolite in the presence of the study's precipitant drug or drug metabolite.

2. None of the criteria listed above as independently confirming a metabolic inhibition interaction were met.

5.1.1.1 "Unknowns" - Pairs with no Known Interaction or Non-interaction

If neither a metabolic inhibition interaction or non-interaction could be confirmed for any pair of pharmaceutical entities in the DIKB, then the pair was labeled as having no known interaction or non-interaction.

5.1.1.2 *AUC Ratios and Statistical Significance*

The evidence board defined a statistically significant increase in AUC to be:

$$\frac{AUC_i}{AUC} > 0 \quad (p \leq .05) \tag{5.1}$$

Where AUC is the baseline AUC for a DDI study's object drug or drug metabolite and AUC_i is the AUC for the object drug in the presence of the study's precipitant drug or drug metabolite. AUC values could be derived from concentration measurements taken from the time of the patient's initial exposure (t_0) to the time that the drug reached its maximum systemic concentration or from an estimate of the AUC as the concentration-time curve approaches its asymptote ($t_0 \rightarrow \infty$). Often studies do not provide p-values, in such cases an AUC increase was considered statistically significant if the study provided 95% confidence intervals for the relationship in Equation 5.1 that did not include 1.0. If the study's results did not satisfy Equation 5.1 or, the 95% confidence intervals for the AUC ratio provided by the study ($\frac{AUC_i}{AUC}$) included 1.0, then the evidence board defined the metabolic inhibition interaction to *not* be statistically significant.

5.1.1.3 *Inclusion Criteria for Validation Set Data*

There were three sources of pharmacokinetic data where the evidence board sought evidence for confirming interactions and non-interactions – published research articles, drug product labeling, and published observation-based case reports. If the data came from a research article, then the study must have satisfied the definition of the evidence type *A DDI clinical trial* or any of its sub-types in the DIKB evidence taxonomy (Appendix D). The research article must also have met the inclusion criteria shown for the evidence type *A DDI clinical trial* and its sub-types (Appendix E, Section E.4). If the evidence board found the data in drug product labeling then it must have met the inclusion criteria for the evidence type *A non-traceable drug-label statement* (Appendix E, Section E.5). Finally, case reports needed to meet the inclusion criteria for evidence of their type listed in Appendix E, Section E.2 *and* provide quantitative measurements of the systemic concentration of the purported victim

drug before and after administration of the purported precipitant drug.

5.1.2 The Collection of Pharmacokinetic Data

The criteria for confirming or refuting a metabolic-inhibition DDI were designed before the evidence-board began collecting evidence. The board used the process explained in Chapter 4, Section 4.4.3 to build the DIKB's evidence-base. Once the DIKB's evidence-base was complete, the evidence-board began building the validation set. They started by enumerating all pairwise combinations of the 35 drugs and drug metabolites in the DIKB's final evidence-base; a total 1190 (Equation 5.2) excluding same-compound combinations (e.g. simvastatin-simvastatin):

$$35 * 34 = 1190 \tag{5.2}$$

In the validation set, the evidence board used a single label to represent both possible ways that an interaction could occur between a drug or drug metabolite pair. For example, the evidence board considered two possible interactions involving diltiazem and atorvastatin (that diltiazem effects a change in the systemic concentration of atorvastatin and *vice versa*) but represented both possible interactions by the single label *diltiazem-atorvastatin* in the validation set's table of interactions. Appendix G lists 595 drug/drug and drug/drug-metabolite pairs representing all 1190 pairwise interaction and non-interactions between the drugs and drug metabolites in the DIKB.

5.1.2.1 Avoiding Biased Measures of the DIKB's Accuracy

Throughout the evidence collection process the evidence board often found clinical trials providing data that were relevant to both establishing the validity of an assertion about some metabolic property and confirming a metabolic interaction or non-interaction. The board's initial policy was to not allow evidence items to be applied in both the DIKB's evidence-base and the validation set because doing so might introduce bias into calculations of the DIKB's prediction accuracy. To avoid bias, evidence items that could be applied to both places would be placed *only* in the DIKB's evidence-base. Shortly after implementing

this policy the board found that it weakened the validation set by excluding known interactions and non-interactions. For example, the simvastatin product label [119] reports data from clinical trial where no increase was observed in systemic concentration of midazolam, a CYP3A4 probe substrate, in the presence of simvastatin. This statement meets the validation set's inclusion criteria for confirming a non-interaction by metabolic inhibition between the two drugs. It also can be used in the DIKB to support the drug mechanism assertion `(simvastatin does-not-inhibit CYP3A4)`. Not including the assertion in the validation set neglects an important finding but including the assertion in both the DIKB and the validation set might bias the DIKB predictions towards the validation set and could inflate calculations of the system's accuracy.

To clarify, assume that only one evidence item in the validation set supports the claim that there is *no* pharmacokinetic interaction between simvastatin and midazolam. Assume also that this same evidence item is in the DIKB as support for the assertion `(simvastatin does-not-inhibit CYP3A4)` and that this assertion is used along with others by the DIKB to predict that simvastatin *will not* interact with midazolam by inhibiting its primary total clearance enzyme CYP3A4. The problem is that the same evidence item is used in both the validation set and the DIKB to come to the same conclusion - that midazolam and simvastatin will not interact via metabolic inhibition. This effectively biases the DIKB toward the validation set even though, by itself, the study does not cause the system to predict the non-interaction (the DIKB would need to combine the assertion that the study supports with other assertions to arrive at its prediction).

The evidence board's solution was to allow evidence items that could be applied to both the DIKB's evidence-base and the validation set to be placed in both places. Then, before assessing the systems accuracy, the board would identify any evidence item that supported a claim made by the validation set and was also used to support a DIKB assertion that leads to the same conclusion. The interaction or non-interaction that such evidence was linked to was dropped from further analysis. The board applied this approach to the DIKB's evidence-base before running the experiment described in this chapter. In total, seven pairs were dropped from further analysis for this reason. These seven drug/drug or drug/drug-metabolite pairs are shown in Table 5.1 along with two other pairs that were

accidentally excluded from the experiment described in this chapter due to a transcription error. Excluding these nine pairs brought the total number of drug/drug and drug/drug-metabolite pairs used for characterizing the DIKB's accuracy down to 586.

Table 5.1: The tables below show nine drug/drug or drug/drug-metabolite pairs that were excluded from our analysis of the DIKB's accuracy. Seven pairs were excluded because a validation set interaction or non-interaction involving the pair was supported by a single clinical trial that was also present in DIKB assertions that the system could use to infer the interaction or non-interaction. These clinical trials are referred to as "dual-use" evidence items. Two other pairs were accidentally excluded due to a transcription error.

Drug pairs excluded because of "dual-use" evidence items

drug/drug or drug/drug-metabolite pair	*source*	*confirms*	*also supports*
clarithromycin - beta-hydroxy-simvastatin	[95]	interaction	beta-hydroxy-simvastatin is a substrate of CYP3A4
clarithromycin - simvastatin	[95]	interaction	simvastatin is a substrate of CYP3A4
itraconazole - beta-hydroxy-simvastatin	[129]	interaction	beta-hydroxy-simvastatin's primary total clearance enzyme is CYP3A4, beta-hydroxy-simvastatin is a substrate-of CYP3A4
itraconazole - simvastatin	[129]	interaction	simvastatin's primary total clearance enzyme is CYP3A4, simvastatin is a substrate of CYP3A4
midazolam - atorvastatin	[118]	interaction	atorvastatin inhibits CYP3A4
midazolam - beta-hydroxy-simvastatin	[144]	non-interaction	beta-hydroxy-simvastatin does not inhibit CYP3A4
midazolam - simvastatin	[144]	non-interaction	simvastatin does not inhibit CYP3A4

Drug pairs excluded due to a transcription error

itraconazole - para-hydroxy-atorvastatin
midazolam - desacetyldiltiazem

Table 5.2: A small sample of the 65 statements the evidence board located in drug-product labeling that mentioned a pharmacokinetic interaction or non-interaction between one of the drug/drug and drug/drug-metabolite combinations shown in Appendix G. Three of these statements were not used in the validation set because they did not provide quantitative data. The arrows point to the drug or drug metabolite that, based on the statement, would be the victim of a pharmacokinetic DDI involving the pair.

Statement	*interaction/non-interaction*	*validation set*
"Human pharmacokinetic data suggest that SPORANOX [itraconazole] inhibits the metabolism of atorvastatin, cerivastatin, lovastatin, and simvastatin, which may increase the risk of skeletal muscle toxicity, including rhabdomyolysis." [96]	itraconazole - atorvastatin $\rightarrow$ itraconazole - lovastatin $\rightarrow$ itraconazole - simvastatin $\rightarrow$	no
"In a small pharmacokinetic study involving HIV infected patients, clarithromycin was shown to increase plasma concentrations of itraconazole." [96]	clarithromycin - itraconazole $\rightarrow$	no
"Similarly, following administration of 1 gram of erythromycin ethyl succinate and 200 mg itraconazole as single doses, the mean C_{max} and AUC $0-\infty$ of itraconazole increased by 44% (90% CI: 119-175%) and 36% (90% CI: 108-171%), respectively." [96]	itraconazole - erythromycin $\leftarrow$	yes
""When single 40 mg doses of simvastatin or atorvastatin, both substrates of CYP3A4, were given to healthy adult volunteers who had received nefazodone hydrochloride, 200 mg BID for 6 days, approximately 20 fold increases in plasma concentrations of simvastatin and simvastatin acid and 3 to 4 fold increases in plasma concentrations of atorvastatin and atorvastatin lactone were seen. These effects appear to be due to the inhibition of CYP3A4 by nefazodone because, in the same study, nefazodone had no significant effect on the plasma concentrations of pravastatin, which is not metabolized by CYP3A4 to a clinically significant extent." [164]	atorvastatin - nefazodone $\leftarrow$ nefazodone - simvastatin $\rightarrow$ nefazodone - beta-hydroxy-simvastatin $\rightarrow$ nefazodone - pravastatin $\nrightarrow$	yes
"As with other macrolides, clarithromycin has been reported to increase concentrations of HMG-CoA reductase inhibitors (e.g., lovastatin and simvastatin). Rare reports of rhabdomyolysis have been reported in patients taking these drugs concomitantly." [3]	clarithromycin - HMG-CoA reductase inhibitors[a] $\rightarrow$	no

[a]The evidence board considered the following drugs and drug metabolites listed in the DIKB to be HMG-CoA reductase inhibitors:

- atorvastatin, ortho-hydroxy-atorvastatin, para-hydroxy-atorvastatin
- simvastatin, beta-OH-simvastatin (simvastatin acid), 6'-exomethylene-simvastatin, 6'-hydroxy-simvastatin, 6'-hydroxymethyl-simvastatin
- lovastatin, beta-OH-simvastatin (lovastatin acid), 6'beta-hydroxy-lovastatin
- fluvastatin, pravastatin, rosuvastatin

5.1.2.2 Searching for Validation-set DDIs in Drug-product Labeling

After completing an intensive search for all relevant clinical trials, the evidence board conducted a search in drug product labeling for evidence that could be used define interactions and non-interactions in the validation set. The board conducted its search so that drug-product labeling, clinical trial literature, and case reports could be compared for their agreement on validation set interactions and non-interactions. All statements that mentioned a pharmacokinetic interaction or non-interaction between one of the drug/drug and drug/drug-metabolite combinations shown in Appendix G were noted and then filtered so that only statements providing quantitative data were used to support interactions in the validation set. All searches were done using product labeling in the NLM's DailyMed database.[b]

DailyMed provides labeling information for drug products containing various combinations of active and inactive ingredients in several possible formulations including capsule, liquid, intravenous, instant, or extended release. The evidence board searched all labels written for each drug product whose *only* active pharmaceutical ingredient was a drug in the DIKB. The number of qualifying product labels for each drug ranged from one (atorvastatin, fluvastatin, and rosuvastatin) to 18 (diltiazem) but a large proportion of the statements in one product label were repeated in all of the other available labels. The evidence board found it efficient to identify all relevant statements in multiple product labels by closely reading four main sections (contraindications, warnings, precautions, adverse reactions) and two sub-sections (drug-drug interactions, and drug-interactions) from a single label then looking for differences in the other available labels. To make the work even more efficient, the evidence-board's informaticist used a computer program he wrote that highlights deletions, insertions, and replacements between labels.

In total, the evidence-board found 65 statements in drug-product labeling that mentioned a pharmacokinetic interaction or non-interaction between one of the drug/drug or drug/drug-metabolite pairs. Only 21 (31%) statements reported the quantitative results of a pharmacokinetic clinical trial. The evidence-board approved these 21 for use in the valida-

[b] http://www.dailymed.nlm.nih.gov

Table 5.3: The evidence-board accepted none of the 35 case reports it found that were relevant for use in the validation set. Most case reports did not provide adequate measurements of the purported victim drug's systemic concentration. The three reports cited here failed to meet inclusion criteria for other reasons.

Case report	*Reported interaction*	*Reason for rejection*
[11]	itraconazole - clarithromycin →	The patients in the report were taking concomitant medications that could have played a role in high clarithromycin levels
[161]	erythromycin - lovastatin →	The lovastatin level was drawn after the patient developed renal failure plus the patient was taking concomitant medications
[87]	midazolam - erythromycin ←	The indicted effect of IV erythromycin on the first pass metabolism of oral midazolam is unusual; use of unknown fruit juice in preop leaves open the possibility of a CYP3A4 inhibition by grapefruit juice

tion set; the remaining 44 statements were retained so that drug-product labeling, clinical trial literature, and case reports could be compared for their agreement on validation set interactions and non-interactions at a later time. Table 5.3 shows a small sample of some of the labeling statements that were accepted or rejected.

5.1.2.3 Searching for Validation-set DDIs in Case Report Literature

Having completed searches within clinical trial literature and drug-product labeling, the evidence-board then did an intensive search search of *The Metabolism and Transport Drug Interaction Database*[c] and PubMed[d] for published case reports claiming the occurrence of a DDI between any pair of the active ingredients or metabolites in our study. One of the drug experts on the evidence-board was affiliated with the proprietary *The Metabolism and Transport Drug Interaction Database* and performed searches of that resource then forwarded the results to the informaticist. The informaticist did an exhaustive search of PubMed for abstracts involving the pairs of interest using a computer program he wrote that executed multiple queries for each drug or drug metabolite in the pair. The program

[c]http://www.druginteractioninfo.org/

[d]http://www.ncbi.nlm.nih.gov/PubMed/

issued queries to PubMed of the following form:

```
Case Reports [PT] AND
(Drug Interactions [MeSH Terms] OR interaction [Text Word]) AND
<d1> AND ("<d1>" [MeSH Terms] OR <d1> [Text Word]) AND
"<d2>" [MeSH Terms] OR <d2> [Text Word])
```

The program iteratively substituted the variables <d1> and <d2> with the generic and trade names of the drug or drug metabolites in the pair (listed in Appendix I, Section I.2.2) and their pharmacologic actions as specified in the Medical Subject Headings (MeSH) controlled vocabulary [43].

The evidence-board found abstracts for 35 relevant case reports in its search of *The Metabolism and Transport Drug Interaction Database* and PubMed. The board retrieved full-text articles for all 35 reports and evaluated each report to see if they met the criteria stated in Section 5.1.1.3. Unfortunately, none of the 35 reports were accepted for use in the validation set. The board rejected most case reports because they did not provide quantitative measurements of the systemic concentration of the purported victim drug before and after administration of the purported precipitant drug. Three case reports provided adequate measurements of systemic concentration but failed to meet inclusion criteria for other reasons. Table 5.3 cites the reports and provides an explanation for why they were rejected.

5.1.2.4 The Final Validation Set

The interactions and non-interactions in the final validation are shown in Tables 5.5 and 5.4. The validation set claims that some DDI will occur by metabolic inhibition for 41 drug/drug and drug/drug-metabolite pairs and that no DDI will occur by metabolic inhibition for seven pairs. No interaction or non-interaction could be identified for 537 pairs in the validation set[e] using its criteria (Section 5.1.1). It is important to stress that many of these pairs might

[e]These 537 pairs are listed in Appendix G along with the nine pairs that were excluded from this experiment.

Table 5.4: Neither drug or drug metabolite in each of the pairs shown in this table are expected to be the victim of a metabolic inhibition interaction effected by the other drug or drug metabolite in the pair. These are the validation set *non-interactions* that we used to characterize the DIKB's prediction performance. Arrows with a line through them indicate which drug or drug metabolite should not be affected by a metabolic inhibition interaction involving the other drug in the pair.

Pair	*Source*
diltiazem - pravastatin $\not\rightarrow$	[32]
erythromycin - rosuvastatin $\not\rightarrow$	[47]
fluconazole - 14-hydroxyclarithromycin $\not\rightarrow$	[3]
fluconazole - pravastatin $\not\rightarrow$	[102]
fluconazole - rosuvastatin $\not\rightarrow$	[49]
itraconazole - fluvastatin $\not\rightarrow$	[108]
nefazodone - pravastatin $\not\rightarrow$	[164]

have clinically-relevant DDIs with each-other that were missed by our evidence collection process or that have not been reported in the sources we searched.

5.1.3 Expert-defined Belief Criteria

Once work on the evidence-base and the validation set was complete, the evidence-board then defined which combinations of evidence that they believed lent different degrees of certainty to assertion types in the DIKB. The DIKB distinguishes between assertion *instances* and assertion *types*. An assertion instance is a specific fact about a particular object such as a drug or protein. For example, the assertion `(midazolam substrate-of CYP3A4)` is an instance of the generic `(X substrate-of Y)` assertion type. The evidence-board defined one or more LOEs for each generic assertion type by creating logical statements listing the level's required evidence types and their multiplicity. This was a two step process; the evidence-board's informaticist first identified all evidence types from the DIKB's evidence taxonomy (Appendix D) that were applicable to each assertion type, then he helped the board's two drug experts define which combinations of evidence they believed lent different degrees of certainty to each assertion type. Appendix H shows a "belief criteria questionnaire" that helped all members of the evidence-board reach consensus on LOEs for each assertion type.

Table 5.5: The *interactions* in the validation set used to characterize the DIKB's prediction performance. The arrows indicate the drug or drug metabolite that the validation set considers the victim of a metabolic inhibition interaction that occurs between the pair.
†- The noted interaction occurs by inhibition of the metabolic clearance of a parent compound.

Pair	*Source*
alprazolam - erythromycin ←	[182]
alprazolam - itraconazole ←	[181]
alprazolam - ketoconazole ←	[74], [156]
alprazolam - nefazodone ←	[57], [75]
atorvastatin - erythromycin ←	[159]
atorvastatin - nefazodone ←	[164]
clarithromycin - atorvastatin →	[8], [95]
clarithromycin - fluconazole ←	[3]
clarithromycin - pravastatin →	[95]
diltiazem - beta-hydroxy-lovastatin →	[29]
diltiazem - lovastatin →	[32]
diltiazem - midazolam →	[30]
diltiazem - simvastatin →	[125]
diltiazem - triazolam →	[171]
erythromycin - beta-hydroxy-simvastatin →	[103]
erythromycin - simvastatin →	[103]
fluconazole - 1'-hydroxymidazolam → †	[4]
fluconazole - fluvastatin →	[102]
itraconazole - atorvastatin →	[117]
itraconazole - beta-hydroxy-lovastatin →	[108]
itraconazole - erythromycin ←	[96]
itraconazole - lovastatin →	[108]
itraconazole - ortho-hydroxy-atorvastatin → †	[117]
itraconazole - pravastatin →	[117]
itraconazole - rosuvastatin →	[50]
ketoconazole - simvastatin →	[42]
midazolam - clarithromycin ←	[70], [78]
midazolam - erythromycin ←	[135]
midazolam - fluconazole ←	[4], [134]
midazolam - itraconazole ←	[136]
midazolam - ketoconazole ←	[136]
midazolam - nefazodone ←	[111]
nefazodone - 4-hydroxyalprazolam → †	[75]
nefazodone - beta-hydroxy-simvastatin →	[164]
nefazodone - simvastatin →	[164]
triazolam - clarithromycin ←	[73]
triazolam - erythromycin ←	[141]
triazolam - fluconazole ←	[172]
triazolam - itraconazole ←	[130], [170]
triazolam - ketoconazole ←	[170], [174]
triazolam - nefazodone ←	[31]

Table 5.6: The evidence-board used the *ranking categories* shown in this table so that multiple evidence types that they felt conferred the same degree of justification for certain assertion types could be represented by a single symbol. †- A *ranking criterium* that was created to represent evidence types that the drug experts on the evidence-board felt would not be applicable to supporting or refuting particular assertions.

Ranking category	*Evidence types*
iv-met-enz-id-Cyp-450-with-inh	*A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors* *A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors* *A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors* *A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors*
iv-met-enz-id-Cyp-450-recombinant	*A CYP450, recombinant, drug metabolism identification experiment with possibly NO probe enzyme inhibitor(s)*
iv-met-enz-id-Cyp-450-microsomal	*A CYP450, human microsome, drug metabolism identification experiment*
iv-met-inh-recombinant	*A CYP450, recombinant, metabolic enzyme inhibition experiment*
iv-met-inh-microsomal	*A CYP450, human microsome, metabolic enzyme inhibition experiment*
pk-ct-pk	*A randomized DDI clinical trial* *A genotyped pharmacokinetic clinical trial* *A phenotyped pharmacokinetic clinical trial*
pk-ct-pk-genotype	*A genotyped pharmacokinetic clinical trial*
pk-ct-pk-phenotype	*A phenotyped pharmacokinetic clinical trial*
pk-ddi-non-rndm	*A non-randomized DDI clinical trial:* *A parallel groups DDI clinical trial*
pk-ddi-rndm	*A randomized DDI clinical trial*
label-statement	*A non-traceable drug-label statement*
nt-statement	*A non-traceable, but possibly authoritative, statement*
obs-eval	*A published and evaluated observation-based ADE report*
na-primary-total-clearance-enz †	*A non-traceable, but possibly authoritative, statement*
na-primary-metabolic-clearance-enzyme †	*A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors*
na-substrate-of †	*A CYP450, recombinant, drug metabolism identification experiment with possibly NO probe enzyme inhibitor(s)*

Table 5.7: The levels of evidence (LOEs) defined by the evidence-board as described in Chapter 5, Section 5.1.4. The LOEs use the *ranking categories* defined in Table 5.6 so that multiple evidence types can be represented by a single symbol. In this table, the symbol '::=' means the term to the left "is defined as" the term to the right, | means "or", and '+' means that "one or more occurrences of" of the symbol to the left are allowed. For example, LOE-1 of set A reads "LOE-1 is defined as one or more evidence types from the **pk-ct-pk** *ranking category* OR one or more evidence types from the **label-statement** *ranking category*."

ID	*LOEs*	*Applies to*
A	LOE-1 ::= **pk-ct-pk+** | **label-statement+**	`bioavailability` `maximum-concentration`
B	LOE-1 ::= **pk-ct-pk+** LOE-2 ::= **label-statement+**	`first-pass-effect` `fraction-absorbed`
C	LOE-1 ::= **pk-ct-pk+** LOE-2 ::= **label-statement+** LOE-3 ::= **na-primary-total-clearance-enz+**	`primary-total-clearance-mechanism`
D	LOE-1 ::= **pk-ct-pk-genotype+** | **pk-ct-pk-phenotype+** LOE-2 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** LOE-3 ::= **iv-met-enz-id-Cyp450-with-inh+** LOE-4 ::= **label-statement+** LOE-5 ::= **na-substrate-of+**	`controls-formation-of` `substrate-of` `is-not-substrate-of`
E	LOE-1 ::= **pk-ct-pk-phenotype+** | **pk-ct-pk-genotype+** LOE-2 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** LOE-3 ::= **label-statement+**	`primary-total-clearance-enzyme`

This table is continued on next page...

Table 5.7: continued from the previous page...

ID	*LOEs*	*Applies to*
F	LOE-1 ::= **pk-ct-pk-phenotype+** | **pk-ct-pk-genotype+** LOE-2 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** LOE-3 ::= **label-statement+** LOE-4 ::= **na-primary-metabolic-clearance-enzyme+**	`primary-metabolic-clearance-enzyme`
G	LOE-1 ::= **pk-ct-pk+** LOE-2 ::= **iv-met-enz-id-Cyp450-with-inh+** LOE-3 ::= **label-statement+** LOE-4 ::= **na-substrate-of+**	`has-metabolite`
H	LOE-1 ::= **iv-met-inh-microsomal+** | **iv-met-inh-recombinant+** LOE-2 ::= **label-statement+**	`inhibition-constant`
I	LOE-1 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** LOE-2 ::= **iv-met-inh-microsomal+** | **iv-met-inh-recombinant+** | **label-statement+**	`inhibits`
J	LOE-1 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** | **iv-met-inh-microsomal+** | **iv-met-inh-recombinant+** LOE-2 ::= **label-statement+**	`does-not-inhibit`

K	LOE-1 ::= **iv-met-enz-id-Cyp450-microsomal+** | **iv-met-enz-id-Cyp450-recombinant+**
	LOE-2 ::= **label-statement+**
	LOE-3 ::= **nt-statement+**
		`does-not-permanently-deactivate-catalytic-function`
		`permanently-deactivates-catalytic-function`
		`in-vitro-probe-substrate-of-enzyme`

This table is continued on next page...

Table 5.7: continued from the previous page...

ID	*LOEs*	*Applies to*
L	LOE-1 ::= **iv-met-inh-recombinant+** LOE-2 ::= **iv-met-inh-microsomal+** LOE-3 ::= **label-statement+** LOE-4 ::= **nt-statement+**	`in-vitro-selective-inhibitor-of-enzyme`
M	LOE-1 ::= **pk-ddi-rndm+** LOE-2 ::= **pk-ddi-non-rndm+** LOE-3 ::= **label-statement+** LOE-4 ::= **nt-statement+**	`in-viVo-selective-inhibitor-of-enzyme` `sole-PK-effect-alter-metabolic-clearance`
N	LOE-1 ::= **pk-ddi-rndm+** | **pk-ddi-non-rndm+** | **obs-eval+** | **label-statement+** LOE-2 ::= **nt-statement+**	`pceut-entity-of-concern`
O	LOE-1 ::= **pk-ct-pk-phenotype+** | **pk-ct-pk-genotype+** LOE-2 ::= **label-statement+** LOE-3 ::= **nt-statement+**	`polymorphic-enzyme`

5.1.3.1 The Evidence-board's Levels-of-evidence and Ranking Criteria

Table 5.7 shows the 15 LOE groups defined by the evidence-board. The evidence-board used the *ranking categories* shown in Table 5.6 when defining LOEs. This enabled them to use a single symbol to represent multiple evidence types that they felt conferred the same degree of justification for certain assertion types. There were some evidence types that the drug experts on the evidence-board felt would not be applicable to supporting or refuting particular assertions. The informaticist defined specific *ranking criteria* for these LOEs (see Table 5.6) and added them, where appropriate, under the lowest-ranking LOE defined by the drug experts.

5.1.3.2 The Evidence-board's Belief Criteria Strategy'

In the DIKB, expert users select one LOE for every assertion type as the *belief criteria* and the system will use a particular assertion instance in inference if, and only if, there exists a body of evidence *for* the assertion that satisfies the *belief criteria* for the assertion's type and the evidence *against* the object property does not satisfy *belief criteria.* The DIKB allows the *belief criteria* for evidence supporting an assertion type to be different from the *belief criteria* for evidence refuting an assertion type. Table 5.8 shows the belief criteria that the evidence-board chose for each assertion type.

5.1.4 Automatically-generated Belief Criteria

We expanded the DIKB so that the system could select any LOE belonging to an assertion type as its *belief criteria.* With this added functionality, the system could create every combination of *belief criteria* possible with the LOEs shown in Table 5.7 and write each of them to separate files on disk. The system could then apply the rule-based theory for predicting DDIs that occur by metabolic inhibition described in Chapter 3 to each distinct set of *belief criteria* and output prediction results to a separate file. We also wrote other small computer programs to assist with calculating the accuracy and coverage of the system using each *belief criteria strategy* and save the results in a single table. We were now able to characterize the effect of varying *belief criteria* on the system's accuracy and coverage.

Table 5.8: This table shows the evidence board's *belief criteria* strategy. The columns next to each assertion type indicate the LOE group and specific LOE from Table 5.7 that the evidence-board chose as the type's *belief criteria*. There are two columns because the DIKB allows *belief criteria* for supporting evidence to be different than *belief criteria* for refuting evidence. Assertion types indicated as *default assumptions* are noted separately because the system uses them for inference even if their evidence does not meet *belief criteria*

Assertions not used as *default assumptions*		
Assertion type	*Evidence for*	*Evidence against*
bioavailability	A, LOE-1	A, LOE-1
maximum-concentration	A, LOE-1	A, LOE-1
first-pass-effect	B, LOE-1	B, LOE-1
fraction-absorbed	B, LOE-1	B, LOE-1
primary-total-clearance-mechanism	C, LOE-1	C, LOE-1
controls-formation-of	D, LOE-1	D, LOE-1
substrate-of	D, LOE-1	D, LOE-1
is-not-substrate-of	D, LOE-1	D, LOE-1
primary-total-clearance-enzyme	E, LOE-1	E, LOE-1
primary-metabolic-clearance-enzyme	F, LOE-1	F, LOE-1
has-metabolite	G, LOE-1	G, LOE-1
inhibition-constant	H, LOE-1	H, LOE-1
inhibits	I, LOE-1	J, LOE-1
does-not-inhibit	J, LOE-1	I, LOE-1
Assertions used as *default assumptions*		
permanently-deactivates-catalytic-function	K, LOE-1	K, LOE-1
does-not-permanently-deactivate-catalytic-function	K, LOE-1	K, LOE-1
in-vitro-probe-substrate-of-enzyme	K, LOE-1	K, LOE-1
in-vitro-selective-inhibitor-of-enzyme	L, LOE-1	L, LOE-1
in-viVo-selective-inhibitor-of-enzyme	M, LOE-1	M, LOE-1
sole-PK-effect-alter-metabolic-clearance	M, LOE-1	M, LOE-1
pceut-entity-of-concern	N, LOE-1	N, LOE-1
polymorphic-enzyme	O, LOE-1	O, LOE-1

Initially, we were going to generate all *belief criteria strategies* by varying the LOEs chosen as *belief criteria* for every assertion type that was not labeled as a *default assumption.* We excluded *default assumptions* because the system does not evaluate the evidence items that are linked to them to see if they meet *belief criteria* and varying their *belief criteria* would have no effect on the DIKB's predictions. Table 5.9 lists the set of assertion types *not* labeled as *default assumptions* along with the number of LOEs that the evidence-board defined for each of them. As Equation 5.3 shows, the total number of *belief criteria strategies* that the DIKB would generate for these assertion types is 576,000.

$$1 * 5 * 2 * 2 * 2 * 4 * 2 * 2 * 5 * 1 * 4 * 3 * 3 * 5 = 576,000 \tag{5.3}$$

Table 5.9: The set of assertion types not labeled as a *default assumption* shown with the number of LOEs that were defined for each of them. The specific LOEs for each assertion type are shown in Table 5.7 †- An assertion type for which varying the LOE chosen as *belief criteria* would have no effect on the DIKB's prediction performance

Assertion type	*LOE count*
`controls-formation-of`	5
`substrate-of`	5
`is-not-substrate-of`	5
`has-metabolite`	4
`primary-metabolic-clearance-enzyme`	4
`primary-total-clearance-enzyme`	3
`primary-total-clearance-mechanism`	3
`does-not-inhibit` †	2
`first-pass-effect` †	2
`fraction-absorbed` †	2
`inhibition-constant` †	2
`inhibits`	2
`maximum-concentration` †	1
`bioavailability` †	1

However, inspection of the DIKB's evidence-base revealed that there were six assertion types for which *all* of the evidence items, for or against, belonged to the highest-ranked LOE for the type. This meant that varying the LOE chosen as *belief criteria* for any of these types (see Table 5.9) would have no effect on the DIKB's prediction performance. Excluding

these six assertion types and the eight assertion types labeled as *default assumptions* meant that there were eight assertion types for which varying the LOE chosen as *belief criteria would* have an effect on prediction performance. The DIKB generated 36,000 different *belief criteria strategies* by altering the LOEs chosen as *belief criteria* for these eight assertion types while keeping the highest-ranking LOE as *belief criteria* for the remaining 14 assertion types.

5.1.5 Searching the Adverse Event Reporting System

Based on our previous experience predicting DDIs using drug mechanisms (see Chapter 2, Section 2.1), we anticipated that the DIKB would make several predictions for which we could find no evidence in clinical trials, case reports, or drug product labeling. Another potential source of observational reports that might indicate the occurrence of a DDI is the Federal Drug Administration's Adverse Event Reporting System (AERS); a public database of population-level post-market safety data [14]. AERS contains more than two million reports of adverse events involving a couple of thousand drugs and biologics [6]. Spontaneous adverse-event reports are not conclusive evidence that the drugs named in the report are causing harm to individuals. However, they can alert drug-safety researchers at the FDA, and elsewhere, of potential safety issues. For example, Jones and Davidson combined AERS reporting data with descriptive statistics on fibrate and statin dispensing and estimates of the rate of fibrate/statin combination therapies to suggest that fenofibrate/statin combination therapy results in fewer adverse event reports per-million prescriptions than does gemfibrozil/statin combination therapy [98]. Also, several researchers have had varying success identifying drug-interactions and DDIs by applying dis-proportionality analysis (data-mining) to AERS [6, 83, 84, 163].

AERS is a volunteer reporting system and suffers from the fact that relevant reporting data is often missing or contains errors. A recent FDA guidance to industry on pharmacovigilance practices and pharmaco-epidemiologic assessment describes several biases that negatively affect AERS data:

> voluntary adverse event reporting systems such as AERS or VAERS are subject

> to a variety of reporting biases (e.g., some observations could reflect concomitant treatment, not the product itself, and other factors, including the disease being treated, other co-morbidities or unrecorded confounders, may cause the events to be reported). In addition, AERS or VAERS data may be affected by the submission of incomplete or duplicate reports, under-reporting, or reporting stimulated by publicity or litigation [25].

With these issues in mind, we planned to search AERS for reports involving the drugs and drug metabolite pairs named in the DIKB's novel interaction predictions for which we could not find case reports. We would extract reports from AERS that included each drug and at least one adverse event term indicative of a toxic effect caused by the purported object drug of the metabolic inhibition interaction. For novel predictions involving drug metabolites we would search for reports containing the metabolite's parent drug. We would then attempt to perform a simple clinical evaluation of each relevant report to see if it provided evidence for the occurrence of the novel interaction prediction.

5.1.5.1 Setting up a Local Copy of AERS

In order to prepare for this analysis, we first acquired AERS data from 1998 to 2007 and set it up in a local SQL database. While AERS data includes reports going back to 1968 we chose to only focus on the last nine years of data since all of the drugs in the DIKB have been sold in the US over this time period. AERS data comes in an structured format called SGML so, we designed and implemented an efficient database schema to store and retrieve AERS data and then wrote a computer program to automatically translate the AERS data from SGML into our schema. The program made it trivial to add more data as new releases are made from the FDA. We made our local copy of the FDA's Adverse Event Reporting System (AERS) accessible to interested researchers at the University of Washington. Currently, the database contains all available database records from November of 1997 through June of 2007. Our implementation is setup on a secure server accessible from the Internet [20].

5.1.5.2 Selecting Drug Product Names for AERS Queries

Our initial tests of AERS indicated that the system uses both generic and trade names for the drug products present in its reports. This meant that we needed to compile a comprehensive list of generic and trade names for drug products containing each of the drugs present in the DIKB's novel interactions in order to query AERS. We collected drug product names for the 16 drugs in the DIKB from the `drugs@fda` database [58] , the FDA's "Orange Book" [15], and RxNorm [19]. Each name represented a drug product that 1) is oral or injectable, 2) contains only one active ingredient and so is *not* a combined therapy, and 3) was present, as of September 2007, in DRUGDEX Tradenames®. Section I.2.2 of Appendix I shows our final list of generic and trade names for the 16 drugs in the DIKB.

5.1.5.3 Choosing Adverse Event Terms for AERS Queries

Persons who submit a report to AERS are required to note the adverse events that prompted them to send the report. FDA personnel code each adverse event using the MedDRA [22] terminology before entering the report into AERS. We compiled a list of MedDRA terms representing the kinds of adverse events that might be observed in patients experiencing toxic side-effects from a victim drug in one of the DIKB's novel DDI predictions which we will present in Chapter 5, Section 5.2.1. We first attempted to utilize the so-called "Standardised MedDRA queries" [126] to build our term sets. These queries are provided by the MedDRA vendor to aid in retrieving cases of interest from databases using the vocabulary. However, we found these to be of little help for the drugs we were interested in with the exception of members of the HMG-CoA reductase inhibitor family. So, we employed the following process to derive a list of terms we thought more appropriate for querying AERS for DDIs:

1. The two drug experts in our group sent the informaticist a list of words describing the effect of a pharmacokinetic interaction for each relevant drug. The informaticist also scanned through drug labels to identify other words that might be useful.

2. The informaticist searched the UMLS Meta-Thesaurus [13] for each of the words

found in Step One to identify concepts in the meta-thesaurus and their mapping to the MedDRA vocabulary.

3. The informaticist created a list of MedDRA "preferred terms" (PTs) from the terms identified in Step Two then used the program shown in Appendix I, Section I.2.3.1 to expand the PT lists to include all MedDRA "LLTs"

4. The two drug experts reviewed the resulting list and removed all LLTs that they did not think relevant to our search task

Appendix I, Section I.2.3 shows the seven sets of adverse-event terms we used to query AERS.

5.1.5.4 Statistical Analysis and Programming

We used the R statistical language [146] to calculate all descriptive and performance statistics. Bruno Falissard's psy package [59] was used to calculate three-valued Cohen's kappa scores as a measure of the degree over random chance to which the DIKB and validation set agreed on interactions, non-interactions, and unknowns. Both R and the Python programming language[f] were used extensively to write various small programs that aided our analysis.

5.2 Results

We began the experiment by testing the accuracy and coverage of the DIKB using the evidence board's *belief criteria strategy* (Section 5.1.3.2). Using this strategy, the DIKB predicted that 15 drug/drug or drug/drug-metabolite pairs would interact by metabolic inhibition and that two would not (see Table 5.10). Fourteen of interaction predictions were identified by the validation set to be true positives, the remaining interaction prediction and the two non-interaction predictions were classified as "unknown" in the validation set. Taken together, these three pairs represent interactions and non-interactions that our review of the literature indicates have never been studied together.

[f] www.python.org

The predicted pharmacokinetic magnitude of all 14 confirmed predictions corresponded with levels observed in clinical trial data. While the system's predictions and magnitude estimates using the evidence-board's strategy had perfect accuracy, its coverage of known interactions was poor. Only 14 (34%) of the 41 pairs known in the validation set to interact by metabolic inhibition were predicted to interact by the DIKB. Also, the system failed to predict any of the seven pairs known in the validation set to not interact by the same mechanism. The system's poor coverage was because only few drug-mechanism assertions were linked to sufficient evidence to meet the evidence board's *belief criteria.*

Table 5.10: Seventeen interaction and non-interaction predictions made by the DIKB using the evidence-board's *belief criteria strategy.* The arrows point to the drug or drug metabolite that the system predicts will be the victim of the interaction. The DIKB makes interaction predictions at three levels: PKI-1, PKI-2, PKI-3 (see Chapter 3, Section 3.2.3.5). All 15 interaction predictions are at the PKI-3 level indicating that the concentration of the affected drug or drug metabolite should increase by at least 100% (2 fold). †- a pair classified as "unknown" in the validation set. ‡- AUC rations are from values present in at least one evidence item in the validation set that supports the interaction.

Pair	*DIKB prediction*	AUC_i/AUC ‡	*correct level?*
diltiazem - midazolam	PKI-3 →	≥ 3	Y
diltiazem - triazolam	PKI-3 →	2.83	Y
midazolam - clarithromycin	PKI-3 ←	≥ 7	Y
midazolam - erythromycin	PKI-3 ←	4.4	Y
midazolam - fluconazole	PKI-3 ←	≥ 2.6	Y
midazolam - itraconazole	PKI-3 ←	10.8	Y
midazolam - ketoconazole	PKI-3 ←	15.9	Y
midazolam - nefazodone	PKI-3 ←	4.6	Y
triazolam - clarithromycin	PKI-3 ←	5.3	Y
triazolam - erythromycin	PKI-3 ←	2.1	Y
triazolam - fluconazole	PKI-3 ←	2.5	Y
triazolam - itraconazole	PKI-3 ←	≥ 3.1	Y
triazolam - ketoconazole	PKI-3 ←	≥ 9.2	Y
triazolam - nefazodone	PKI-3 ←	≥ 3.9	Y
triazolam - atorvastatin †	PKI-3 ←	n/a	n/a
triazolam - simvastatin †	NO-PKI	n/a	n/a
triazolam - beta-hydroxy-simvastatin †	NO-PKI	n/a	n/a

We then tested if using alternate *belief criteria strategies* had any influence on the accuracy and coverage DIKB's predictions. The DIKB failed to make any predictions using one strategy due to some unknown error that occurred during the experiment. This results

from this strategy[g] were not used in further analysis. We analyzed the remaining 35,999 different strategies for accuracy, coverage, and agreement with the validation set. Table 5.11 shows summary statistics for each performance parameter we analyzed over all prediction sets. The DIKB's sensitivity ranged from 0.88 to 1.0 with 19,583 (54%) of the *belief criteria strategies* causing the system to operate at maximum sensitivity. The systems specificity ranged from 0.0 to 1.0 with 6,912 (19%) of the *belief criteria strategies* causing the system to operate at maximum specificity. The system had excellent positive predictive value (range: 0.94 to 1) across all *belief criteria strategies.* However, we could not characterize the system's negative predictive value in a meaningful way because the DIKB never predicted more than two of the seven validation set non-interactions.

Table 5.11: Summary statistics for each performance parameter we analyzed over 35,599 prediction sets. The columns labeled "*n*" show the number of *belief criteria strategies* whose predictions shared each minimum and maximum value.

statistic	*Min.*	*n*	*Median*	*Mean*	*Max.*	*n*
true positives	14.0	1,440	33.0	30.8	34.0	17,280
false positives	0.0	17,279	1.0	0.6	2.0	2,880
true negatives	0.0	21,599	0.0	0.7	2.0	11,520
false negatives	0.0	19,583	0.0	0.6	2.0	4,032
sensitivity	0.88	576	1.00	0.98	1.00	19,583
specificity	0.00	11,232	0.50	0.44	1.00	6,912
positive predictive value	0.94	240	0.97	0.98	1.00	17,279
DIKB-only unknown	10.0	5,760	13.0	15.3	34.0	576
validation-set-only unknown	3.0	864	40.0	42.4	62.0	2,880
kappa	0.41	576	0.47	0.48	0.52	768

We calculated three-valued kappa scores for every prediction set using Cohen's kappa to see how the agreement between the DIKB and the validation set compared with random chance. The DIKB's predictions across all prediction sets had moderate agreement (0.4 to 0.5) with the validation set and never reached levels typically considered indicative of significant agreement (≥ 0.7) or disagreement (≤ 0.3).

[g]To be specific, we threw out the results of the *belief criteria strategy* that used the following LOEs for *belief criteria*: `controls-formation-of`:LOE-1, `has-metabolite`:LOE-4, `inhibits`:LOE-2, `is-not-substrate-of`:LOE-1, `primary-metabolic-clearance-enzyme`:LOE-1, `primary-total-clearance-enzyme`:LOE-1, `primary-total-clearance-mechanism`:LOE-3, `substrate-of`:LOE-3

Table 5.12: 1,152 (3%) strategies caused the DIKB to perform optimally in terms of sensitivity, positive predictive value, and agreement with the validation set as measured by Cohen's kappa. This table shows all measured performance characteristics for these "best-performing" strategies.

statistic	*value*
true positives	34
false positives	0
true negatives	0
false negatives	0
sensitivity	1
specificity	n/a
positive predictive value	1
DIKB-only unknown	14
validation-set-only unknown	40
kappa	0.52

A fascinating result of this experiment is that 8,351 (23%) of the 35,599 tested strategies caused the DIKB to have equal or better performance in terms of sensitivity, positive predictive value, and agreement with the validation set than the evidence board's strategy. Table 5.12 shows the performance characteristics for 1,152 (3%) strategies that performed at the top level in these three performance categories. All of these strategies caused the DIKB to make the same set of 65 interaction and non-interaction predictions.

These strategies predicted a metabolic inhibition interaction for 34 (83%), of the 41 interacting pairs in the validation while making no false positive and no false negative predictions. As Table 5.13 shows, the pharmacokinetic magnitude of 30 of the 34 confirmed (88%) predictions made using the best performing belief criteria strategies matched levels observed in clinical trial data. While the coverage of the DIKB using these strategies was greater than with the evidence-board's strategies and there was no loss of accuracy, the system's magnitude estimates were not as accurate and its coverage of validation set data remained incomplete – it missed seven interactions and made no predictions for the seven non-interactions listed in the validation set.

Table 5.13: Each of the 1,152 strategies that caused the DIKB to perform optimally in terms of sensitivity, positive predictive value, and agreemen with the validation set caused the system to make the same 65 predictions shown here. This table shows the 34 interaction prediction that are confirmed by the validation set. The arrows point to the drug or drug metabolite that the system predicts will be the victim of the interaction. ‡- The DIKB also predicted a metabolic inhibition interaction at the PKI-1 level with clarithromycin as the victim.

pair	*DIKB level*	AUC_i/AUC	*correct level?*
alprazolam - erythromycin	PKI-1 ←	1.61	Y
alprazolam - itraconazole	PKI-1 ←	2.7	N
alprazolam - ketoconazole	PKI-1 ←	3.98	N
alprazolam - nefazodone	PKI-1 ←	1.98	Y
atorvastatin - erythromycin	PKI-3 ←	1.4	N
atorvastatin - nefazodone	PKI-3 ←	3-4	Y
clarithromycin - atorvastatin ‡	PKI-3 →	$\geq$ 1.8, max 5.4	Y
clarithromycin - fluconazole	PKI-1 ←	$\geq$ 1.18, max 1.33	Y
diltiazem - lovastatin	PKI-3 →	$\geq$ 3	Y
diltiazem - beta-hydroxy-lovastatin	PKI-3 →	3.57	Y
diltiazem - midazolam	PKI-3 →	3.75	Y
diltiazem - simvastatin	PKI-3 →	4.8	Y
diltiazem - triazolam	PKI-3 →	2.83	Y
erythromycin - simvastatin	PKI-3 →	6.3	Y
erythromycin - beta-hydroxy-simvastatin	PKI-3 →	3.9	Y
fluconazole - fluvastatin	PKI-3 →	1.83	N
itraconazole - atorvastatin	PKI-3 →	2.5	Y
itraconazole - lovastatin	PKI-3 →	14.8	Y
itraconazole - beta-hydroxy-lovastatin	PKI-3 →	8.56	Y
ketoconazole - simvastatin	PKI-3 →	12.55	Y
midazolam - clarithromycin	PKI-3 ←	$\geq$ 7	Y
midazolam - erythromycin	PKI-3 ←	4.4	Y
midazolam - fluconazole	PKI-3 ←	$\geq$ 2.6	Y
midazolam - itraconazole	PKI-3 ←	10.8	Y
midazolam - ketoconazole	PKI-3 ←	15.9	Y
midazolam - nefazodone	PKI-3 ←	4.6	Y
nefazodone - simvastatin	PKI-3 →	20	Y
nefazodone - beta-hydroxy-simvastatin	PKI-3 →	20	Y
triazolam - clarithromycin	PKI-3 ←	5.3	Y
triazolam - erythromycin	PKI-3 ←	2.1	Y
triazolam - fluconazole	PKI-3 ←	2.5	Y
triazolam - itraconazole	PKI-3 ←	$\geq$ 3.1	Y
triazolam - ketoconazole	PKI-3 ←	$\geq$ 9.2	Y
triazolam - nefazodone	PKI-3 ←	$\geq$ 3.9	Y

Table 5.14: This large table shows the 31 interaction and nine non-interaction predictions made by the DIKB that are unknown, but not refuted, in the validation set. We list next to each interaction its predicted magnitude and any relevant case report that we found along with our estimate (using the DIPS scale [99]) of the likelihood that the adverse event in the case report was caused by the interaction . We also show the results of searching the FDA's spontaneous adverse-event reporting database (AERS) for any interaction prediction for which we could find no case reports. The DIPS forms used to evaluate these 15 case reports are shown in Appendix K.

pair	*DIKB level*	*Case report/DIPS analysis*	*AERS search*
alprazolam - atorvastatin	PKI-1 ←	none found	no data set sought; the interaction was deemed possibly too hard to identify due due to the nature of alprazolam usage
alprazolam - clarithromycin	PKI-1 ←	none found	no search conducted
alprazolam - fluconazole	PKI-1 ←	none found	no data set sought; the interaction was deemed possibly too hard to identify due to the nature of alprazolam usage
atorvastatin - lovastatin	PKI-3 →	none found	no data set sought; the interaction was deemed possibly too hard to identify because both drugs have similar therapeutic and toxic effects
atorvastatin - beta-hydroxy-lovastatin	PKI-3 →	none found	no data set sought; the interaction was deemed possibly too hard to identify because both drugs have similar therapeutic and toxic effects
atorvastatin - simvastatin	PKI-3 →	none found	no data set sought; the interaction was deemed possibly too hard to identify because both drugs have similar therapeutic and toxic effects
atorvastatin - beta-hydroxy-simvastatin	PKI-3 →	none found	no data set sought; the interaction was deemed possibly too hard to identify because both drugs have similar therapeutic and toxic effects
atorvastatin - fluconazole	PKI-3 ←	[101]: *possible*; also on proton pump inhibitor	n/a
clarithromycin - erythromycin	PKI-1 ←	none found	no data set sought; the interaction was deemed possibly too hard to identify because both drugs have similar therapeutic and toxic effects
clarithromycin - lovastatin	PKI-3 →	[77]: *possible* for two cases; both patients also on another medication with interaction potential	n/a

Table continued on the next page...

Table 5.14: continued from the previous page...

pair	*DIKB level*	*Case report/DIPS analysis*	*AERS search*
clarithromycin - beta-hydroxy-lovastatin	PKI-3 →	see clarithromycin - lovastatin	n/a
clarithromycin - nefazodone	PKI-1 ←	none found	six clinician-submitted reports with insufficient data to evaluate
diltiazem - alprazolam	PKI-1 →	none found	no data set sought; the interaction was deemed possibly too hard to identify due to the nature of alprazolam usage
diltiazem - atorvastatin	PKI-3 →	[67]: *possible*, all patients on proton pump inhibitors that could interact and have been reported to cause rhabdomyolysis [113] : *possible*; Pt had CHF, valve disease and decr CO may cause decr renal function. CK elevation was mild and no myoglobin in urine.	n/a
diltiazem - beta-hydroxy-simvastatin	PKI-3 →	[140]: *probable* [67]: *possible*; all patients on proton pump inhibitors that could interact and have been reported to cause rhabdomyolysis [89]: *possible*; patients on other medications with interaction potential	n/a
diltiazem - clarithromycin	PKI-1 →	none found	one report related to *torsades de pointes*, seven related to arrhythmia, and 21 related to general macrolide side-effects. All reports had insufficient data to evaluate.
erythromycin - lovastatin	PKI-3 →	[161]: *probable* [179]: *probable* [28]: *probable*; note erythromycin + diltiazem combination	n/a

Table continued on the next page...

Table 5.14: continued from the previous page...

pair	*DIKB level*	*Case report/DIPS analysis*	*AERS search*
erythromycin - beta-hydroxy-lovastatin	PKI-3 →	see erythromycin - lovastatin	n/a
fluconazole - lovastatin	PKI-3 →	none found	one clinician-submitted reports with insufficient data to evaluate
fluconazole - beta-hydroxy-lovastatin	PKI-3 →	none found	see fluconazole -lovastatin
fluconazole - simvastatin	PKI-3 →	[157]: *possible*	n/a
fluconazole - beta-hydroxy-simvastatin	PKI-3 →	see fluconazole - simvastatin	n/a
itraconazole - clarithromycin	PKI-1 →	[11]: *possible* for two cases; used historic controls for blood-level data when Pt not taking interacting drugs, clinical observational study	n/a
ketoconazole - atorvastatin	PKI-3 →	none found	11 clinician-submitted reports with insufficient data to evaluate
ketoconazole - clarithromycin	PKI-1 →	none found	16 reports related to *torsades de pointes*, two related to arrhythmia, and one related to general macrolide side-effects. All reports had insufficient data to evaluate.
ketoconazole - lovastatin	PKI-3 →	[162]: *probable*	n/a
ketoconazole - beta-hydroxy-lovastatin	PKI-3 →	(see ketoconazole - lovastatin)	n/a
ketoconazole - beta-hydroxy-simvastatin	PKI-3 →	[66]: *probable* [5]: *probable* [91]: *probable*	n/a
nefazodone - lovastatin	PKI-3 →	none found	one clinician-submitted reports with insufficient data to evaluate
nefazodone - beta-hydroxy-lovastatin	PKI-3 →	none found	see nefazodone - lovastatin
triazolam - atorvastatin	PKI-3 ←	none found	no search conducted
alprazolam - beta-hydroxy-simvastatin	NO-PKI ↚	n/a	n/a
alprazolam - simvastatin	NO-PKI ↚	n/a	n/a
lovastatin - beta-hydroxy-simvastatin	NO-PKI ↚	n/a	n/a
simvastatin - beta-hydroxy-lovastatin	NO-PKI ↛	n/a	n/a
simvastatin - beta-hydroxy-simvastatin	NO-PKI ↮	n/a	n/a
simvastatin - lovastatin	NO-PKI ↛	n/a	n/a
triazolam - beta-hydroxy-simvastatin	NO-PKI ↚	n/a	n/a
triazolam - simvastatin	NO-PKI ↚	n/a	n/a
beta-hydroxy-simvastatin - beta-hydroxy-lovastatin	NO-PKI ↛	n/a	n/a

5.2.1 *Evaluation of the DIKB's Novel DDI Predictions Made Using the Best-performing Belief Criteria Strategies*

The system also predicted 31 metabolic inhibition interactions and nine non-interactions using the "best-performing" *belief criteria strategies* whose validity was unknown by the validation set. These *novel* interaction predictions, shown in Table 5.14, represent potentially interacting drug combinations that our review of the literature indicate have not been studied. After running the experiment, we used a similar method to the one we used to search for case reports for the validation set (see Section 5.1.2.3) to search PubMed for for clinical trials involving these pairs. We could only find one clinical trial that was not already included in the validation set [158] unfortunately, we judged this study's methodology too poor to use it as evidence for or against any interactions.

Fifteen of the published case reports we had collected while constructing the validation set claimed the occurrence of a DDI that matched one of the 31 novel predictions. Each report was reviewed using the *Drug Interaction Probability Scale* (DIPS) [99] by a clinician co-investigator. The DIPS defines four qualitative levels (*Doubtful*, *Possible*, *Probable*, and *Highly Probable*) representing the degree to which the information provided by the report supports the proposition that a specific drug combination effected an adverse event or events. Six novel predictions were matched with case reports that met the DIPS *Probable* level; meaning that the predicted interactions were the likely cause of an adverse event occurring in a patient. Seven novel predictions were matched with reports that met the DIPS *Possible* level; meaning that the predicted interactions could not be excluded from consideration as the cause of an adverse event in a patient. The DIPS forms used to evaluate these 15 case reports are shown in Appendix K.

5.2.1.1 *Querying AERS*

This left 18 novel interaction predictions for which we could find no supporting or refuting published data. We anticipated the evaluation of AERS reports to be time consuming so we prioritized the remaining 18 remaining novel interactions and chose the six interactions we thought would be the most important to investigate. For each of the six predictions we

extracted all AERS reports that were submitted by a clinician, involve both pharmaceutical entities in the novel DDI prediction, and have at least one adverse event term that is indicative of the metabolic inhibition interaction involving either drug in the prediction. Appendix I, Section I.2 provides details on how we executed these queries.

Our queries returned one or more clinician-submitted AERS report for each of the six novel DDI predictions. Unfortunately, *none* of the reports provided sufficient data on the administration dates of the drugs listed in any report for us to be able to evaluate how (or even *if*) the drugs listed in the report were co-administered. Without this information it was impossible for the drug experts to assess if a DDI was the cause of an adverse event listed in the report.

Drugs labeled in the report as "suspect" were far more likely to have administration dates than drugs labeled as "concomitant." This could be due to the fact that the forms used to submit reports [16, 17] do not provide separate boxes for dates of administration for concomitant medications. A few reports provided medication dates but they seemed to indicate the usage history of a particular medication over different regimens rather than the administration order of the drugs that the patient was taking at the time of the adverse event. For example, one report (shown in Appendix J) provided multiple non-overlapping dates for oxyconton administration but no dates for 33 medications listed as concomitant making it impossible for use to determine the medications that the patient was taken before experiencing the adverse event.

Table 5.15: This table shows the range of LOEs used by 1,152 *belief criteria strategies* that, like the evidence-board's strategy, had perfect sensitivity and positive predictive value but also maximum coverage of and agreement with the validation set. The LOE chosen as *belief criteria* for supporting and refuting evidence was always the same for the assertions shown in the table. The columns *Evidence for* and *Evidence against* show the distribution of evidence items across the LOEs. Not shown are the default assumptions note in Table 5.8 and the six assertions for which varying the LOE chosen as *belief criteria* would have no effect on the DIKB's prediction performance (Section 5.1.4).

Assertion type	*Belief criteria*	*Evidence for*	*Evidence against*
`inhibits`	I: LOE-1, LOE-2	n=11 ∼ LOE-1(91%), LOE-2(9%)	n=4 ∼ LOE-1(50%), LOE-2(50%)
`substrate-of`	D: LOE-3	n=29 ∼ LOE-1(10%), LOE-2(21%) LOE-3(17%), LOE-4(21%) LOE-5(31)	n=11 ∼ LOE-1(0%), LOE-2(0%) LOE-3(36%), LOE-4(9%) LOE-5(55%)
`is-not-substrate-of`	D: LOE-1 → LOE-3	n=11 ∼ LOE-1(0%), LOE-2(0%) LOE-3(36%), LOE-4(9%) LOE-5(55%)	
`primary-total-clearance-mechanism`	C: LOE-1 → LOE-3	n=14 ∼ LOE-1(21%), LOE-2(64%), LOE-3(14%)	
`primary-total-clearance-enzyme`	E: LOE-3	n=8 ∼ LOE-1(25%), LOE-2(62%) LOE-3(13%)	
`primary-metabolic-clearance-enzyme`	F: LOE-1 → LOE-4	n=3 ∼ LOE-1(0%), LOE-2(0%) LOE-3(67%), LOE-4(33%)	n=1 ∼ LOE-1(100%), LOE-2(0%) LOE-3(0%), LOE-4(0%)
`controls-formation-of`	D: LOE-1, → LOE-3, LOE-4[g], LOE-5[g]	n=17 ∼ LOE-1(0%), LOE-2(0%), LOE-3(12%), LOE-4(6%), LOE-5(82%)	
`has-metabolite`	G: LOE-1 → LOE-2 LOE-3[g], LOE-4[g]	n=27 ∼ LOE-1(19%), LOE-2(0%) LOE-3(44%), LOE-4(37%)	

[g] No *belief criteria strategy* in the set of 1,152 used LOE-4 or any lower-ranking LOE for the `controls-formation-of` assertion type *and* LOE-3 or any lower-ranking LOE for the `has-metabolite` assertion type. There were 288 Strategies that used these LOEs as belief criteria along with the other possible LOES for each assertion type shown in the table above. They all made the prediction that fluconazole would reduce the metabolic clearance of rosuvastatin by inhibiting the CYP2C9 enzyme - a prediction that the validation set considered to be false.

5.3 Discussion

It's important to note that our experiment only looked at binary performance criteria – predictions were classified as true or false according to the validation set and the goal was to maximize true predictions and minimize false predictions. An entirely different set of *belief criteria strategies* than the best-performing strategies of this experiment would be relevant if our goal was to optimize the accuracy of the system's magnitude predictions. This would be a very worthwhile experiment because, as Section 5.2 mentioned, the system is capable of accurate magnitude predictions – the DIKB's magnitude estimates for all 14 interactions known in the validation set were at the correct level. A set *belief criteria strategies* that focused on optimizing magnitude would seek to expand the DIKB's coverage of known interactions past these 14 while still making correct magnitude predictions. This kind of analysis might also indicate the limitations of the very simplistic model that we used to infer the fraction of a drug that is cleared by a particular enzyme from AUC data (Chapter 3, Section 3.2.3.5).

Although our experiment's clinical-relevance is likely to be less than if we had used magnitude-based performance criteria, the binary performance criteria were sufficient for us to conclude some important findings. First, changing the LOEs selected as *belief criteria* does alter the systems prediction accuracy and coverage in the way that we had anticipated. We found for this data set that, as the the criteria for including assertions was relaxed, the DIKB predicted a larger number of true interactions; sometimes at the expense of also making more false predictions. By having the computer iterate through a large set of possible *belief criteria strategies* we found that a significant proportion (23%) of the *belief criteria strategies* we looked at predicted a larger number of true interactions than the most rigorous strategy while still making *no* known false predictions.

Our experiment also found a particular family *belief criteria strategies* that optimized the system's prediction accuracy and coverage to the body of evidence present in the DIKB's evidence-base. Table 5.15 shows the range of LOEs used by 1,152 *belief criteria strategies* that, like the evidence-board's strategy, had perfect sensitivity and positive predictive value but also maximum coverage of and agreement with the validation set. Analyzing the table

is difficult because there is a complex interplay between the kinds of evidence present in the knowledge-base, how it is linked to each assertion instance, and the relationship between each assertion type and the variables chosen for scoring the system's prediction performance.

To better understand the information in Table 5.15 it is useful to note how different combinations of the assertion types that we varied *belief criteria* for could directly or indirectly cause the system to make an interaction or non-interaction prediction. The types `substrate-of` and `inhibits` were used by the DIKB to establish an interaction at the `PKI-1` level (see the rule shown in Figure 3.3), the type `is-not-substrate-of` was used with the type `inhibits` to predict a non-interaction (see the rule shown in Figure 3.8), and the types `controls-formation-of` and `has-metabolite` were used together to establish that a drug or drug metabolite is a substrate of a particular enzyme which would then form an antecedent to the rule just mentioned that predicted interactions at the `PKI-1` level. Similarly, the type `primary-total-clearance-enzyme` was used both in rules that made predictions at the `PKI-3` level (see Figure 3.5) and by the system to infer `substrate-of` assertion instances used by the rule that predicted interactions at the `PKI-1` level.

Only one assertion type present in Table 5.15, `primary-metabolic-clearance-enzyme`, was exclusively used by the system for magnitude estimation[h] so, varying its *belief criteria* should not affect the accuracy of the system. Indeed, Table 5.15 shows that all this assertion type's LOEs were used by the 1,152 "best-performing" *belief criteria strategies.*

Another way that multiple LOEs can be chosen as *belief criteria* for a some assertion types without changing the accuracy and coverage of the DIKB (assuming *belief criteria* for other assertion types are static) is for there to be no evidence items that map to a particular LOE. This is the case for some LOEs belonging to the `controls-formation-of`, `is-not-substrate-of`, and `primary-metabolic-clearance-enzyme` assertion types. For example, since none of the 17 evidence items that were linked to `controls-formation` instances map to the assertions type's LOE-1 and LOE-2, the system's predictions were not affected when LOE-2 was chosen as *belief criteria* instead of LOE-1.

[h]That this was true is actually the result of a mistake. We made an oversight by not having the system infer `substrate-of` assertion instances from justified `primary-metabolic-clearance-enzyme` instances. As with the `primary-total-clearance-enzyme` assertion instances, a drug or drug metabolite is clearly a substrate of any enzyme that is believed to have a significant role in its metabolic clearance.

A less obvious way for different *belief criteria* to lose their influence on DIKB predictions is for all of the assertion instances of a particular assertion type to have already been justified by *belief criteria* using higher-ranking LOEs. For example, during our experiment, the assertion `(ketoconazole-inhibits-cyp3a4)` had three evidence items in the evidence-base – two that met LOE-1 for the `inhibits` assertion type (see Table 5.1.3) and one that met LOE-2. Since there were evidence items that mapped to both LOEs, the assertion `(ketoconazole-inhibits-cyp3a4)` was justified no matter which LOE the system chose as belief criteria. Also, the evidence item that mapped to LOE-2 was the only one in the entire evidence-base *supporting* an `inhibits` assertion instance that did not meet LOE-1 for its type. As a result, no *new* assertions could be introduced when the system relaxed the type's *belief criteria* partially explaining why there was no difference in the DIKB's predictions between strategies that used LOE-1 as *belief criteria* for the `inhibits` assertion and those that used and LOE-2. The situation is made more complicated by the fact that two evidence items *refuting* `inhibits` assertions did map to LOE-2. Had either of these evidence items been linked to the assertion `(ketoconazole-inhibits-cyp3a4)`, then both evidence for and against the assertion would have met `belief criteria` and the DIKB would have retracted the assertion along with any predictions made using it. However, this did not occur because these evidence items were linked to the assertions `(pravastatin does-not-inhibit CYP2C8)` and `(rosuvastatin does-not-inhibit CYP2C8)`.

5.3.1 Other Interesting Features of the Prediction Sets

So far in our discussion we have focused only on the performance of the DIKB using the "best-performing" strategies. However, 27,648 (77%) of the strategies we tested caused the DIKB to predict at least one interaction or non-interaction considered invalid by the validation set. Both kinds of invalid predictions were made using 7488 (21%) of the strategies and the maximum number of interaction or non-interaction predictions countered by the validation set for any single strategy was three (either two invalid interactions and one invalid non-interaction or vice versa).

Table 5.16 shows the four invalid predictions that appeared in various combinations

among the predictions made using a wide range of strategies. For the two interaction predictions countered by the validation set, the table indicates which drug the DIKB considers the victim and the specific enzyme whose inhibition should lead to the interaction. The itraconazole-fluvastatin interaction prediction occurred when the system used strategies that accept drug labeling statements as *belief criteria* because 1) the assertion `(itraconazole inhibits CYP3A4)` was a *default assumption* and 2) the evidence-base recorded one labeling statement (based on a non-cited *in vitro study*) proposing that fluvastatin is a minor substrate (<20% of total clearance) of CYP3A4 [51]. The DIKB predicted the fluconazole-rosuvastatin interaction using strategies that allow statements in product labeling to justify the `controls-formation` and `has-metabolite` assertion types and non-randomized clinical trial data to justify the `inhibits` assertion type. In this case, the system inferred that rosuvastatin is a substrate of CYP2C9 because the two assertions `(rosuvastatin has-metabolite N-desmethylrosuvastatin)` and `(CYP2C9 controls-formation-of N--desmethylrosuvastatin)` were each supported by evidence items based on labeling information [10] and the assertion `(fluconazole inhibits cyp2c9)` was supported by a non-randomized clinical trial [33].

Considering now the two countered *non*-interaction predictions; Table 5.16 indicates which drug the DIKB predicted would *not* be affected by inhibition of a specific metabolic enzyme. The DIKB predicted a non-interaction between itraconazole and rosuvastatin via CYP3A4 inhibition using strategies that allow statements in product labeling to justify the `is-not-substrate-of` assertion type. This was because the evidence-base contained one evidence item, based on a labeling statement, declaring CYP3A4 to *not* have a role in the metabolic clearance of rosuvastatin [10]. The system predicted a non-interaction between fluconazole and clarithromycin using strategies that considered `inhibits` type justified by non-randomized clinical trial data and the `is-not-substrate-of` assertion type justified by *in vitro* metabolism identification studies using human microsomes and chemical inhibitors. In these strategies the system could apply one evidence item [33] to justify the assertion `(fluconazole inhibits CYP2C9)` and another item [151] to justify the assertion `(clarithromycin is-not-substrate-of CYP2C9)`. Interestingly, this non-interaction prediction was overruled when the system used strategies that, in addition to the previously

mentioned *belief criteria*, also considered the `substrate-of` assertion type justified by evidence from *in vitro* metabolism identification experiments. Using these strategies, the DIKB predicted an interaction between fluconazole and clarithromycin CYP3A4 because an evidence item based on a *in vitro* metabolism identification experiment justified the assertion `(clarithromycin substrate-of CYP3A4)` [151].

Table 5.16: 27,648 strategies led the DIKB to predict an interaction or non-interaction countered by the validation set. The countered predictions produced by the DIKB using each strategy consisted of one or more of the four pairs shown in this table. For countered interactions, the arrows indicate the drug that the DIKB considers the victim of a metabolic inhibition interaction via inhibition of the enzyme shown in parentheses. For countered non-interactions, they point to the drug that should *not* be affected by inhibition of the enzyme shown in parentheses.

Countered interaction	*Countered non-interaction*
itraconazole - fluvastatin (CYP3A4) $\rightarrow$	itraconazole - rosuvastatin (CYP3A4) $\rightarrow$
fluconazole - rosuvastatin (CYP2C9) $\rightarrow$	clarithromycin - fluconazole (CYP2C9) $\leftarrow$

5.3.2 Why was the DIKB's Coverage of the Validation-set Interactions Always Incomplete?

Table 5.17 shows six interactions and four non-interactions present in the validation set that were never predicted by the DIKB using any the 35,599 *belief criteria strategies*. The system did not make these predictions for the following reasons:

- Two missing interactions and three missing non-interactions are accounted for by the fact that there were no assertions or evidence items in the system indicating that pravastatin is cleared by, or inhibits, a metabolic enzyme.

- Similarly, three missing interactions and one non-interaction are accounted for by the fact that were no assertions in the system indicating which enzymes do or do not metabolize 1'-hydroxymidazolam, ortho-hydroxy-atorvastatin, 4-hydroxyalprazolam, 14-hydroxyclarithromycin. Neither are there assertions indicating that these metabolites inhibit a drug-metabolizing enzyme present in the system. The DIKB's data model is capable of predicting when inhibition of the parent compound will or will

not affect the formation of these metabolites made but we did not include these kinds of predictions in the study

- There were two evidence items in the system that supported the assertion `(erythromycin inhibits CYP3A4)` [73, 184] but no assertion or evidence in the system claiming that itraconazole is a substrate of that enzyme. Conversely, the system had three *default assumptions* that separately established itraconazole to be both an *in vivo* and *in vitro* selective inhibitor of CYP3A4 and erythromycin to be an *in vitro* probe substrate. However, the system had no rule that could infer *substrate-of* assertions from *in-vitro-probe-substrate-of* assertions. If it had, the system would have predicted the itraconazole-erythromycin interaction to occur at the `PKI-1` level.

Table 5.17: Pairs in the validation set (Tables 5.5 and 5.4)for which the DIKB made no prediction using any the 35,599 *belief criteria strategies.*

missing interactions	*missing non-interactions*
clarithromycin - pravastatin	diltiazem - pravastatin
fluconazole - 1'-hydroxymidazolam	fluconazole - pravastatin
itraconazole - pravastatin	nefazodone - pravastatin
itraconazole - erythromycin	fluconazole - 14-hydroxyclarithromycin
itraconazole - ortho-hydroxy-atorvastatin	
nefazodone - 4-hydroxyalprazolam	

5.3.3 Comparing the DIKB Predictions to Labeling Statements

We would have liked to have done a quantitative comparison of the system's predictions with drug-drug interaction statements from product labeling but could not because a significant proportion of the validation set was constructed from labeling statements. We did look over the statements that were not used in the validation set and found one statement that specified an interaction countered by clinical trial data present in the validation set. This statement extrapolated an interaction observed between erythromycin and one or more HMG-CoA reductase inhibitors to all drugs in that class:

> Erythromycin has been reported to increase concentrations of HMG-CoA reductase inhibitors (e.g., lovastatin and simvastatin). Rare reports of rhabdomyolysis have been reported in patients taking these drugs concomitantly [2].

The active ingredient rosuvastatin is among the HMG-CoA reductase inhibitors included in our study. The labeling statement indirectly declares a potential pharmacokinetic interaction between erythromycin and rosuvastatin that is countered by a randomized clinical trial in the validation set [47]. None of the interaction predictions made by the DIKB using the evidence board or the best performing belief criteria strategies were countered by the validation set (i.e. false positives or false negatives). While the system made no prediction involving erythromycin and rosuvastatin with these strategies, it correctly predicted a non-interaction between erythromycin and rosuvastatin using other, lower specificity, belief criteria strategies. These results indicate that, depending on belief criteria strategies, DDI prediction using drug-mechanism knowledge can be very accurate and avoid making the kinds of false predictions that occur when individual drug differences are not recognized.

5.3.4 The JTMS Could be Leveraged to Optimize the Search for High-performing Strategies

It took more than three days of non-stop computation on two modern desktop computers to generate all 36,000 prediction sets for the experiment this chapter describes.[i] This lengthy amount of time is more reflective of the process we used to generate prediction sets than the computational complexity of the DDI prediction task because we chose to have the DIKB reset its knowledge-base every time we generated a prediction set using a new *belief criteria strategy*. This forced the system to rebuild the JTMS dependency network for each new strategy, a computationally expensive task, but also enabled us to divide the work onto different machines and easily recover from any computer crashes with no loss of data.

The DIKB's `evidence-model` component (Chapter 2, Section 2.3.1) keeps track of the justification state of each assertion in the system. For example, if the `evidence-model` has already justified an assertion and the assertion's evidence continues to meet *belief criteria*,

[i]For those interested, about 70% of the prediction sets were generated on a computer with a single AMD Athlon 64-bit processor with 2 Gig of RAM. It took about the same amount time for that another computer with a 1.3 GHz Pentium processor and 1 Gig of RAM to generate the other 30%.

then the system will make no change. One could leverage this feature to generate all prediction sets without resetting the knowledge-base or re-running rules but only tracking the assumptions and prediction results that change between each strategy. The method would be quicker than the brute force technique that we used in our experiment for arriving at optimal *belief criteria strategies* and could form the basis for a special search algorithm that minimized the time required to search for the best-performing strategies. Such an algorithm could also take into account knowledge about the existing evidence-base to identify which assertion types will not be affected by changing *belief criteria.* For example, it could exclude LOEs from analysis that will make no difference because no evidence maps to them or because all evidence items map to higher-ranking LOEs.

5.3.5 Limitations

This problem of a biased-use of evidence items introduced in Section 5.1.2.1 can be more complex when considering drug-labeling statements. Drug labeling statements almost always provide no citation to published studies. Since both labeling information and published studies are included in the DIKB it is possible for conclusions from the same study to appear in the system as different evidence items. For example, support for the assertion that active ingredient X inhibits enzyme E could include a pharmacokinetic study and a non-traceable drug-label statement that echoes the results from the *same* study. In our experience, one can usually only conjecture whether the non-traceable statement echoes a specific study or if it refers to a different study. This ambiguity can affect the development of LOEs because an expert's confidence in an assertion whose evidence support includes one or more non-traceable statements along with actual study evidence should be no different than if the assertion rested on only study evidence *unless* it can be shown that the non-traceable statements refer to distinct studies than the ones already included.

In terms of calculating the DIKB's prediction accuracy, it seems reasonable that the same bias mentioned in Section 5.1.2.1 will occur if a clinical trial is applied as support for an interaction or non-interaction in the validation set while a labeling statement echoing, but not citing, the study supports an assertion that the DIKB uses to predict the same

interaction. In such cases, the validation set interaction or non-interaction should be excluded from calculations of the systems accuracy. Conversely, if a validation set interaction or non-interaction rests on only a single non-traceable statement and there are assertions in the DIKB used to predict the interaction or non-interaction that depend on the study that inspired the statement, then the interaction or non-interaction should also be excluded from calculations of the system's accuracy. Unfortunately, we did not implement any strategy to avoid this kind of bias so it is possible that some labeling data was used to support mechanistic assertions that led to predictions validated by the same data but appearing in a different source. Future work will examine if this bias was present and, if so, what effect removing the affected interactions or non-interactions has on the calculations of DIKB accuracy.

5.4 Conclusion

This chapter has described a novel experiment characterizing the effect of varying *belief criteria* on the system's accuracy and coverage of DDIs present in a reference set of interactions and non-interactions. The experiment's results demonstrate that the DIKB can make accurate predictions for an important class of DDIs using only knowledge of drug-mechanisms and that the system's prediction accuracy and coverage varies depending on the *belief criteria strategy* being used. We were able to use LOEs and *belief criteria* to optimize the system's prediction performance to the contents of its evidence-base. Though we only looked at binary performance criteria, we know from the success of evidence board's strategy that the same optimization approach will work for maximizing the accuracy and coverage of the system's magnitude estimates. We conclude from these results that the evidential knowledge representation approach used by the DIKB has features that are very desirable for supporting clinical decision making.

The central thesis of this this dissertation is that DDI prediction using drug-mechanism knowledge can help drug-interaction KBs expand their coverage beyond what has been tested in clinical trials while avoiding prediction errors that occur when individual drug differences are not recognized. The fact that nearly half (42%) of the novel DDI predictions has some degree of support from published case reports is strong evidence that drug-mechanism

knowledge can help drug-interaction KBs expand their coverage of DDIs beyond what has been tested in clinical trials. The system also correctly avoided predicting a pharmacokinetic interaction between erythromycin and rosuvastatin even though class-based reasoning present in drug product labeling suggested the interaction could occur. This shows that the system's mechanism-based knowledge can help avoid errors that occur when making class-based inferences that do not respect individual drug differences.

Our exploration of AERS was prompted by a desire to seek evidence for drug and drug/metabolite combinations that the DIKB predicted were likely putting people at risk. To the best of our knowledge no other investigators have tested if it is possible to gather evidence from a public reporting database for DDIs predicted to occur based on well-supported pharmacologic mechanisms. While our attempt did not yield evidence for or against the DIKB's novel predictions, it does suggest changes the national spontaneous reporting system that would make its data more useful for the new drug safety methods that we are proposing. The data elements present in the system conform to the ICH E2b/M2 standard for transmitting post-market safety report information [24]; a model that we consider sufficient for representing the date needed to assess reports using a tool like DIPS [99]. The real issue is that the needed data is not being entered into the AERS reports. We think that one reason for this is that the data entry forms used by spontaneous reporters do not specifically request the dates that patients were prescribed or administered concomitant medications.

Finally, this experiment has helped to identify some current technical limitations of the evidential approach to knowledge representation. There is an opportunity for research into new computational methods to help support analysis of *belief criteria strategies*; a task that is currently very difficult because of the complex interplay between the kinds of evidence present in the knowledge-base, how it is linked to each assertion instance, and the relationship between each assertion type and the variables chosen for scoring the system's prediction performance. Also, research on new search algorithms that leverage the evidence-base's contents and internal state machine promise to significantly speed up the time it takes to locate an optimal set of *belief criteria strategies*.

Chapter 6

CONTRIBUTIONS, FUTURE WORK, AND CONCLUDING COMMENTS

6.1 Research Contributions

Little research has been done on how to best represent and maintain knowledge about drug mechanisms so that it can be of use in clinical decision making. We have shown that the new knowledge-representation methods employed in the DIKB enable the system to make accurate predictions for an important class of DDIs using only knowledge of drug-mechanisms. We also showed that the prediction accuracy and coverage of the DIKB can be optimized to a particular body of evidence; a feature that is very desirable for clinical decision support. The success of our new methods is not only a contribution to biomedical informatics research but also to drug safety. Using the best-performing *belief criteria strategies*, the system accurately predicted 34 (83%) of 41 interacting pairs present in a validation set while making no false positive and no false negative predictions. Thirteen (42%) of the 31 novel interaction predictions the system made at its optimal performance level had some degree of support from published case reports. The remaining 18 novel predictions could represent combinations with the potential to harm patients that have not previously been recognized or studied. These predictions are important because they are based on mechanistic assertions supported by strong evidence from studies in humans.

The DIKB is the first knowledge-representation system we are aware of to use a computable model of evidence and a Truth Maintenance System to manage assertions in its knowledge-base. We expect this approach to be generalizable to knowledge representation in other biomedical domains where an ontology of evidence types can be created and used to define rank-ordered levels of justification for assertions in some rule-based theory. For example, the method might be useful for constructing a pathway/genome database that can provide different views of its knowledge to users who might not agree about what combina-

tion of evidence confirm the existence of an biochemical entity or its relationship to other entities within a biochemical pathway.

We developed a new evidence taxonomy to support representing drug-mechanism evidence in the DIKB and contrasted it with three other evidence taxonomies in the bioinformatics domain (Medical Subject Headings' *Publication Types* [27], Gene Ontology's *evidence codes* [65], and Pathway Tools' *evidence ontology* [106]). An important finding was that none of the four taxonomies by themselves could be used to construct levels-of-evidence because their type definitions fail to ensure that all evidence within a collection meet some minimum standard in terms of quality. Our solution was to develop and consistently apply *inclusion criteria* for each type of evidence in the taxonomy. Inclusion criteria help ensure that all evidence within a collection meet some minimum standard in terms of quality and are the key to enabling expert users of a knowledge-base prospectively map their confidence in each assertion type to some arrangement of one or more abstract evidence types. We expect inclusion criteria to enable the use of evidence types from taxonomies like the Pathway Tools' *evidence ontology* within evidential knowledge-representation systems like the DIKB. This fact will be important if future work requires expanding the DIKB to include the more general biochemical pathway knowledge present in pathway/genome databases such as MetaCyc [107].

The DIKB is also novel for its computable representation of conjectures behind a specific application of evidence. The DIKB's *evidence-use assumptions* were designed so that the system could alert curators when one or more conjectures that a particular application of evidence depends on fail to meet belief criteria. They enable the system to flag when a conjecture has become invalid and alert knowledge-base maintainers to the need to reassess their original interpretation of what assertions a piece of evidence supports. We used them during the evidence collection process to help identify a pattern, called a *circular line of evidence support*, that is indicative of fallacious reasoning by evidence-base curators. The two algorithms that we proposed to identify *circular lines of evidence support* (Chapter 4, Section 4.3.4.3) are general and should be applicable in other knowledge-representation domains where fallacious reasoning can occur.

Finally, our exploration of AERS enables us to suggest changes to the national spon-

taneous reporting system that would make its data more useful for drug safety methods that identify potentially interacting drug combinations based on drug mechanisms. The data elements used by AERS conform to the ICH E2b/M2 standard for transmitting post-market safety report information [24]; a model that we consider sufficient for representing the date needed to assess reports using a tool like DIPS [99]. The real issue is that data necessary for determining if an adverse event is the result of a DDI is not being entered into the AERS report. Most noticeable is the lack of non-ambiguous administration dates for concomitant medications. We think that one reason for this might the data entry forms used by spontaneous reporters do not specifically request the dates that patients were prescribed or administered concomitant medications.

6.2 Future Work

Our experiment with the DIKB has helped to identify some current technical limitations of the evidential approach to knowledge representation. A high priority for future work will be research into new computational methods to help support analysis of *belief criteria strategies*; a task that is currently very difficult because of the complex interplay between the kinds of evidence present in the knowledge-base, how it is linked to each assertion instance, and the relationship between each assertion type and the variables chosen for scoring the system's prediction performance. Also, it will be important to research new search algorithms that leverage the evidence-base's contents and internal state machine to significantly speed up the time it takes to locate an optimal set of *belief criteria strategies.*

We acknowledged that many drug-mechanism facts that we consider well-supported today will need revision to account for scientific progress thus, collecting and maintaining a drug-mechanism evidence-base should be on ongoing process by design. An important area of future work will be on the development of new computer-supported evidence maintenance processes. For example, it is quite feasible to develop software agents that leverage the remote query and *RSS syndication* facilities of journal Web-sites, publication databases such

as PubMed and PubMed Central,[a] and drug-product labeling resources such as DailyMed.[b] Such agents could automatically identify when evidence sources in the DIKB have been updated or cited by others and alert DIKB curators. Software agents could also filter query results based on evidence existing in the DIKB so that curators could quickly retrieve similar evidence sources.

We think that DIKB curators should always make the final decision as to how to apply a given item of evidence but automated tools have the potential to greatly ease their task. Especially promising are the new methods emerging from research in machine learning and artificial intelligence. For example, Rubin *et al* developed and successfully applied a statistical classifier that accurately identified pharmacogenomics research articles from indexed research abstracts [153]. Their methods involved training the classifier on a set of abstracts collected and classified by humans using the PharmGKB's *categories of pharmacogenetics evidence* [152]. Pustejovsky and colleagues successfully applied corpus-based linguistics to extract statements describing protein inhibition from biomedical research abstracts [145]. Also notable is the work of Rzhetsky *et al* on integrating automated methods for several tasks necessary to build large-scale biological pathway knowledge-bases including the selection of relevant evidence items and extraction of concepts [155]. The DIKB's current evidence-base consists of consists of evidence from 102 unique sources applied as evidence for or against 222 drug-mechanism assertions. We think that this body of evidence could form a solid training set for testing the methods used by these and other researchers.

6.3 Secondary Results

We made several discoveries regarding the quality and accuracy of some current information resources whose purpose is to support pharmacy practice while collecting evidence for the DIKB. For example, in the process of identifying the set of generic and trade names for each active ingredient in our system we found what we believe were errors in a tool name `RxNorm` developed by the National Library of Medicine.[c] `RxNorm` claims to provide "...standard

[a] http://www.pubmedcentral.nih.gov/

[b] http://dailymed.nlm.nih.gov

[c] http://www.nlm.nih.gov/research/umls/rxnorm

names for clinical drugs (active ingredient + strength + dose form) and for dose forms as administered to a patient" [19]. All of the following are potential errors that we identified in the `BrandName` field of the September 2007 version of `RxNorm` :

1. As near as we can tell there is no such product "CAKNEMYCIN" yet it is listed a product containing erythromycin; this possibly a misspelling of "AKNEMYCIN."

2. "CARIZEM" is listed as a product of diltiazem but it is likely a misspelling of cardizem

3. "ROYMICIN" appears to be a misspelling of the erythromycin product robimycin

4. "ALTOCOR" is listed as a product containing lovastatin but there is no product by that name in the two primary sources for drug information managed by the FDA, `drugs@fda` [58] and "The Orange Book" [15]. This is possibly a misspelling caused by conflating two FDA-approved lovastatin products "ALTOPREV" and " ADVICOR."

5. The tool lists "Dermamycin" as a product containing erythromycin but searches of both the Micromedex® and UpToDate® medical knowledge databases suggest that it is a product containing diphenhydramine. Neither `drugs@fda` [58] or "The Orange Book" [15] have any record of any drug product by this name and `RxNorm` provides no references for its name assignments.

It is important to note that a natural use of `RxNorm` by software developers is as a dictionary for drug names and synonyms. Errors in the resource could results in serious consequences. For example, it is unknown to us if "Dermamycin" contains erythromycin or diphenhydramine or if it is even a real product. What is important is that that, if the assignment is in error (i.e. "Dermamycin" exists and contains diphenhydramine) and if any online drug information source repeats the `RxNorm` assignment then, some patients could be *misled* into thinking that erythromycin interactions apply to a diphenhydramine containing product.[d] We have already sent an email to the developers of `RxNorm` listing potential

[d] As of December 2007 at least one online drug information source (`http://www.herb-drug.com/drugs/dermamycin.html`) used the same assignment that was in `RxNorm` though we do not know if developers of the information resource used `RxNorm` as their tool for drug synonyms.

errors we have identified but, as of the time of this writing, do not know if they have been corrected.

We also identified several issues with the DailyMed system; a resource that the National Library of Medicine claims...

> ...provides health information providers and the public with a standard, comprehensive, up-to-date, look-up and download resource of medication content and labeling as found in medication package inserts [12].

Unfortunately, we found that at least some the labels at DailyMed are neither standard nor up-to-date. For example the label for instant release alprazolam [127] does not have a separate section titled "*drug-drug interactions*" while the label for extended release alprazolam [165] does. Also, Product labels for the same drug sometimes provided out-of-date information. In such cases the evidence-board collected the statement that seemed most up-to-date. For example, one alprazolam product label [165] states that the AUC levels of alprazolam increased 3.98-fold in patients who were exposed to ketoconazole while another, out-of-date, statement in another alprazolam label declares:

> Although in vivo interaction data with alprazolam are not available, ketoconazole and itraconazole are potent CYP 3A inhibitors and the co-administration of alprazolam with them is not recommended. [127]

We intend to send a complete list of the errors we identified in `DailyMed` to its developers as soon as possible. We also intend to write at least one conference paper or journal article to make the errors we have identified in `RxNorm` and `DailyMed` public. We think it vital to alert clinicians to the fact that they should never blindly trust these resources. It is also critical to inform the increasing number of biomedical informaticists developing pharmacy-focused tools of the serious need for building accurate and up-to-date resources.

6.4 Concluding Remarks

A recent shift in our nation's focus to patient safety has inspired a broad effort by government and industry to expand the use of electronic prescribing aids. As a results, there has been

an increased effort in researching ways to overcome many technical and socio-technical challenges to bringing sound DDI knowledge from the knowledge-base to the bedside. We envision that, over the next decade, a new generation of highly accurate tools will become available that use pharmacologic theory, drug mechanism knowledge, and patient-specific data to help clinicians assess the combined effect of multiple drugs, the effect of removing a drug from a patients drug regimen, and individual response to therapy due to enzyme polymorphisms. These tools will be a significant advance in medicine and a radical change from the functionality that current prescribing software offers. Our research on how to best represent drug mechanism knowledge for the purpose of making clinically relevant DDI predictions is a small, through important step, toward understanding how to build and deploy the highly accurate tools that we envision.

BIBLIOGRAPHY

[1] Reducing and Preventing Adverse Drug Events To Decrease Hospital Costs. Technical Report "Research in Action, Issue 1", AHRQ, March 2001. Internet, `http://www.ahrq.gov/qual/aderia/aderia.htm`. Last accessed 06/20/2008.

[2] Abbot. pce (erythromycin) tablet. FDA-approved drug product labeling, 02 2007. Last accessed on DailyMed 03/03/2008.

[3] Abbott. Biaxin filmtab (clarithromycin) tablet, film coated. FDA-approved drug product labeling, 03 2007. Last accessed on DailyMed 05/29/2008.

[4] J. Ahonen, K. T. Olkkola, and P. J. Neuvonen. Effect of route of administration of fluconazole on the interaction between fluconazole and midazolam. *Eur J Clin Pharmacol*, 51(5):415–419, 1997.

[5] K. Akram, S. Rao, and M. Parker. A lesson for everyone in drug-drug interactions. *Int J Cardiol*, 118(1):e19–e20, 2007.

[6] June Almenoff, William DuMouchel, L. Allen Kindman, and Xionghu Yang. Disproportionality analysis using empirical Bayes data mining: a tool for the evaluation of drug interactions in the post-marketing setting. *Pharmacoepidemiology and Drug Safety*, 12:517–521, 2003.

[7] R. B. Altman, D. A. Flockhart, S. T. Sherry, D. E. Oliver, D. L. Rubin, and T. E. Klein. Indexing pharmacogenetic knowledge on the World Wide Web. *Pharmacogenetics*, 13(1):3–5, 2003.

[8] G. W. Amsden, O. Kuye, and G. C. Wei. A study of the interaction potential of azithromycin and clarithromycin with atorvastatin in healthy volunteers. *J Clin Pharmacol*, 42(4):444–449, 2002.

[9] AstraZeneca. Accolate® zafirlukast tablets. FDA-approved drug product labeling, July 2004. Last accessed on DailyMed 08/06/2005.

[10] Astrazeneca. crestor (Rosuvastatin calcium) tablet, film coated for oral use. FDA-approved drug product labeling, 11 2007. Last accessed on DailyMed 04/17/2008.

[11] B. Auclair, S. E. Berning, G. A. Huitt, and C. A. Peloquin. Potential interaction between itraconazole and clarithromycin. *Pharmacotherapy*, 19(12):1439–1444, 1999.

[12] Internal Authorship. About DailyMed. Internet. http://dailymed.nlm.nih.gov/dailymed/about.cfm. Last accessed 06/01/2008.

[13] Internal Authorship. About the UMLS® Resources. Internet. http://www.nlm.nih.gov/research/umls/about_umls.html. Last accessed 06/01/2008.

[14] Internal Authorship. Adverse Event Reporting System (AERS). Internet. http://www.fda.gov/cder/aers/default.htm. Last accessed 06/01/2008.

[15] Internal Authorship. Electronic Orange Book. Internet. http://www.fda.gov/cder/ob/. Last accessed 06/01/2008.

[16] Internal Authorship. FORM FDA 3500. Internet. http://www.fda.gov/medwatch/SAFETY/3500.pdf. Voluntary adverse-event reporting form for MedWatch, last accessed 06/01/2008.

[17] Internal Authorship. FORM FDA 3500A. Internet. http://www.fda.gov/medwatch/SAFETY/3500A.pdf. Mandatory adverse-event reporting form for MedWatch v.10/05, last accessed 06/01/2008.

[18] Internal Authorship. Frequently Asked Licensing Questions (Section "What sets your data apart from the rest?"). Internet. http://www.factsandcomparisons.com/DataLicensing/FAQ.aspx. Last accessed 06/11/2008.

[19] Internal Authorship. RxNorm. Internet. http://www.nlm.nih.gov/research/umls/rxnorm/index.html. Last accessed 06/01/2008.

[20] Internal Authorship. The University of Washington BHI implementation of the FDA's AERS database. Internet. http://marigold.informatics.washington.edu:7000/phpmyadmin. Last accessed 06/04/2008.

[21] Internal Authorship. THOMSON MICROMEDEX Editorial Workflow. Internet. http://www.micromedex.com/about_us/editorial/ed_process.pdf. Last accessed 02/19/2008.

[22] Internal Authorship. Welcome to MedDRA and the MSSO. Internet. http://www.meddramsso.com/MSSOWeb/index.htm. Last accessed 06/01/2008.

[23] Internal Authorship. FDA guidance for industry - population pharmacokinetics. Technical report, Federal Drug Administration, 1999.

[24] Internal Authorship. FDA guideline: Providing regulatory submissions in electronic format – postmarketing expedited safety reports. Internet, May 2001. http://www.fda.gov/cder/guidance/4153dft.pdf. Last accessed 06/04/2008.

[25] Internal Authorship. FDA guideline: Guidance for industry good pharmacovigilance practices and pharmacoepidemiologic assessment. Internet, March 2005. http://www.fda.gov/cder/guidance/63590CC.htm. Last accessed 06/01/2008.

[26] Internal Authorship. FDA guideline: Drug interaction studies – study design, data analysis, and implications for dosing and labeling. Internet, September 2006. http://www.fda.gov/Cber/gdlns/interactstud.htm. Last accessed 09/25/2006.

[27] Internal Authorship. Medical Subject Headings Publication Types. Internet, 2008. http://www.nlm.nih.gov/mesh/pubtypes2008.html. Last accessed 05/14/2008.

[28] J. Z. Ayanian, C. S. Fuchs, and R. M. Stone. Lovastatin and rhabdomyolysis. *Ann Intern Med*, 109(8):682–683, 1988.

[29] N. E. Azie, D. C. Brater, P. A. Becker, D. R. Jones, and S. D. Hall. The interaction of diltiazem with lovastatin and pravastatin. *Clin Pharmacol Ther*, 64(4):369–377, 1998.

[30] J. T. Backman, K. T. Olkkola, K. Aranko, J. J. Himberg, and P. J. Neuvonen. Dose of midazolam should be reduced during diltiazem and verapamil treatments. *Br J Clin Pharmacol*, 37(3):221–225, 1994.

[31] R. H. Barbhaiya, U. A. Shukla, P. D. Kroboth, and D. S. Greene. Coadministration of nefazodone and benzodiazepines: Ii. a pharmacokinetic interaction study with triazolam. *J Clin Psychopharmacol*, 15(5):320–326, 1995.

[32] Biovail. cardizem (diltiazem hydrochloride) tablet, coated. FDA-approved drug product labeling, 04 2006. Last accessed on DailyMed 05/29/2008.

[33] D. J. Black, K. L. Kunze, L. C. Wienkers, B. E. Gidal, T. L. Seaton, N. D. McDonnell, J. S. Evans, J. E. Bauwens, and W. F. Trager. Warfarin-fluconazole. II. A metabolically based drug interaction: in vivo studies. *Drug Metab Dispos*, 24(4):422–428, 1996.

[34] Pascal Bonnabry, Johann Sievering, Thierry Leemann, and Pierre Dayer. Quantitative Drug Interactions Prediction System (Q-DIPS): A Dynamic Computer-Based Method to Assist in the Choice of Clinically Relevant In Vivo Studies. *Clinical Pharmacokinetics*, 40(9):631–640, 2001.

[35] Richard Boyce, Carol Collins, John Gennari, John Horn, and Ira Kalet. DIKB Ontology. Internet, 2007. http://marigold.informatics.washington.edu:25000/~boycer/DIKB_evidence_ontology_v1.0.owl. Last accessed 06/21/2008.

[36] Richard Boyce, Carol Collins, John Horn, and Ira Kalet. Qualitative pharmacokinetic modeling of drugs. In *Proceedings of the AMIA*, pages 71–75, 2005.

[37] S. H. Brown, P. L. Elkin, S. T. Rosenbloom, C. Husser, B. A. Bauer, M. J. Lincoln, J. Carter, M. Erlbaum, and M. S. Tuttle. Va national drug file reference terminology: a cross-institutional content coverage study. *Medinfo*, 11(Pt 1):477–481, 2004.

[38] Laurence Brunton. *Goodman & Gilman's the Pharmacological Basis of Therapeutics*, chapter Design and Optimization of Dosage Regimens: Pharmacokinetic Data. McGraw-Hill, New York, 2006.

[39] Alan Bundy. Incidence Calculus. DAI research paper, University of Edinburgh, 80 South Bridge, Edinburgh Eh1 1HN, Scotland, April 2004. Commisioned by the Encyclopedia of Artificial Intelligence.

[40] CDER Web page on drug interactions. Internet. `http://www.fda.gov/Cder/drug/drugInteractions/`. Last accessed 04/29/2008.

[41] Y-F Chen, AJ Avery, KE Neil, C Johnson, ME Dewey, and IH Stockly. Incidence and possible causes of prescribing potential hazardous/contraindicated drug combinations in general practice. *Drug Safety*, 28:67–80, 2005.

[42] E. Chung, A. N. Nafziger, D. J. Kazierad, and J. S. Jr Bertino. Comparison of midazolam and simvastatin as cytochrome p450 3a probes. *Clin Pharmacol Ther*, 79(4):350–361, 2006.

[43] M. H. Coletti and H. L. Bleich. Medical Subject Headings Used to Search the Biomedical Literature. *J Am Med Inform Assoc*, 8(4):317–323, 2001.

[44] Carol Collins and René Levy. Drug-drug interaction in the elderly with epilepsy: Focus on antiepileptic, psychiatric, and cardiovascular drugs. *Profiles in Seizure Management*, 3(6), 2004.

[45] BioPAX Consortium. BioPAX. Internet, 2007. `http://www.biopax.org/`. Last accessed 02/13/2007.

[46] World Wide Web Consortium. Web Ontology Languave (OWL). Internet, 2007. `http://www.w3.org/2004/OWL/`. Last accessed 02/14/2007.

[47] K. J. Cooper, P. D. Martin, A. L. Dane, M. J. Warwick, A. Raza, and D. W. Schneck. The effect of erythromycin on the pharmacokinetics of rosuvastatin. *Eur J Clin Pharmacol*, 59(1):51–56, 2003.

[48] K. J. Cooper, P. D. Martin, A. L. Dane, M. J. Warwick, A. Raza, and D. W. Schneck. Lack of effect of ketoconazole on the pharmacokinetics of rosuvastatin in healthy subjects. *Br J Clin Pharmacol*, 55(1):94–99, 2003.

[49] K. J. Cooper, P. D. Martin, A. L. Dane, M. J. Warwick, D. W. Schneck, and M. V. Cantarini. The effect of fluconazole on the pharmacokinetics of rosuvastatin. *Eur J Clin Pharmacol*, 58(8):527–531, 2002.

[50] K. J. Cooper, P. D. Martin, A. L. Dane, M. J. Warwick, D. W. Schneck, and M. V. Cantarini. Effect of itraconazole on the pharmacokinetics of rosuvastatin. *Clin Pharmacol Ther*, 73(4):322–329, 2003.

[51] Novartis Pharmaceuticals Corporation. lescol (fluvastatin sodium) capsule; lescol xl (fluvastatin sodium) tablet, extended release. FDA-approved drug product labeling, 01 2007. Last accessed on DailyMed 07/02/2008.

[52] Randall Davis, Bruce Buchanan, and Edward Shortliffe. Production rules as a basis for a knowledge-based consultation program. *Artificial Intelligence*, 8(1):15–45, 1977.

[53] Randall Davis, Howard E. Shrobe, and Peter Szolovits. What Is a Knowledge Representation? *AI Magazine*, 14(1):17–33, 1993.

[54] S. de Coronado, M. W. Haber, N. Sioutos, M. S. Tuttle, and L. W. Wright. NCI Thesaurus: using science-based terminology to integrate cancer research results. *Medinfo*, 11(Pt 1):33–37, 2004.

[55] G. Del Fiol, B. H. Rocha, G. J. Kuperman, D. W. Bates, and P. Nohama. Comparison of two knowledge bases on the detection of drug-drug interactions. In *Proc AMIA Symp*, pages 171–175, 2000.

[56] Hartmut Derendorf, Lawrence Lesko, Philip Chaikin, Wayne Colburn, Peter Lee, Raymond Miller, Robert Powell, Gerald Rhodes, Donald Stanski, and Jürgen Venitz. Pharmacokinetic/pharmacodynamic modeling in drug research and development. *J Clin Pharmacol*, 40:1399–1418, 2000.

[57] C. L. DeVane, J. L. Donovan, H. L. Liston, J. S. Markowitz, K. T. Cheng, S. C. Risch, and L. Willard. Comparative CYP3A4 inhibitory effects of venlafaxine, fluoxetine, sertraline, and nefazodone in healthy volunteers. *J Clin Psychopharmacol*, 24(1):4–10, 2004.

[58] drugs@fda website. Internet. `http://www.fda.gov/cder/drugsatfda/datafiles/default.htm`. Last accessed 03/01/2006.

[59] Bruno Falissard. *psy: Various procedures used in psychometry*, 2007. R package version 0.7.

[60] V. Fischer, L. Johanson, F. Heitz, R. Tullman, E. Graham, J. P. Baldeck, and W. T. Robinson. The 3-hydroxy-3-methylglutaryl coenzyme A reductase inhibitor fluvastatin: effect on human cytochrome P-450 and implications for metabolic drug interactions. *Drug Metab Dispos*, 27(3):410–416, 1999.

[61] Kenneth D. Forbus and Johan De Kleer. *Building Problem Solvers.* MIT Press, 1993.

[62] MA Ford, ML Anderson, JP Rindone, and DW Jaskar. Lack of effect of fluoxetine on the hypoprothrombinemic response of warfarin. *J Clin Psychopharmacol*, 17(2):110–20, 1997.

[63] Gene Ontology Consortium. The Gene Ontology (GO) project in 2006. *Nucleic Acids Research*, 34, 2006. Database Issue.

[64] Gene Ontology Consortium. The Gene Ontology. Internet, 2007. `http://www.geneontology.org`. Last accessed 2/13/2007.

[65] Gene Ontology Consortium. The GO Evidence Code Guide. Internet, 2008. `http://www.geneontology.org/GO.evidence.shtml`. Last accessed 05/12/2008.

[66] R. Gilad and Y. Lampl. Rhabdomyolysis induced by simvastatin and ketoconazole treatment. *Clin Neuropharmacol*, 22(5):295–297, 1999.

[67] P. Gladding, H. Pilmore, and C. Edwards. Potentially fatal interaction between diltiazem and statins. *Ann Intern Med*, 140(8):W31, 2004.

[68] David E. Golan, Armen Tashjihan, and Ehrin Armstrong et al. *Principles of Pharmacology - The Pathophysiologic Basis of Drug Therapy*, chapter Drug Receptor Interactions. Lippincott, Williams, & Wilkins, 2005.

[69] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 1(84):57–112, 1996.

[70] J. C. Gorski, D. R. Jones, B. D. Haehner-Daniels, M. A. Hamman, E. M. Jr O'Mara, and S. D. Hall. The contribution of intestinal and hepatic CYP3A to the interaction between midazolam and clarithromycin. *Clin Pharmacol Ther*, 64(2):133–143, 1998.

[71] S. Goto, T. Seo, T. Murata, N. Nakada, N. Ueda, T. Ishitsu, and K. Nakagawa. Population estimation of the effects of cytochrome p450 2c9 and 2c19 polymorphisms on phenobarbital clearance in japanese. *Ther Drug Monit*, 29(1):118–121, 2007.

[72] Paul Graham. *ANSI Common Lisp.* Prentice Hall, 1996.

[73] D. J. Greenblatt, L. L. von Moltke, J. S. Harmatz, M. Counihan, J. A. Graf, A. L. Durol, P. Mertzanis, S. X. Duan, C. E. Wright, and R. I. Shader. Inhibition of triazolam clearance by macrolide antimicrobial agents: in vitro correlates and dynamic consequences. *Clin Pharmacol Ther*, 64(3):278–285, 1998.

[74] D. J. Greenblatt, C. E. Wright, L. L. von Moltke, J. S. Harmatz, B. L. Ehrenberg, L. M. Harrel, K. Corbett, M. Counihan, S. Tobias, and R. I. Shader. Ketoconazole inhibition of triazolam and alprazolam clearance: differential kinetic and dynamic consequences. *Clin Pharmacol Ther*, 64(3):237–247, 1998.

[75] D. S. Greene, D. E. Salazar, R. C. Dockens, P. Kroboth, and R. H. Barbhaiya. Coadministration of nefazodone and benzodiazepines: Iii. a pharmacokinetic interaction study with alprazolam. *J Clin Psychopharmacol*, 15(6):399–408, 1995.

[76] Dale Griffin and Lyle Brenner. *The Blackwell Handbook of Judgement and Decision Making*, chapter Perspectives On Probability Judgment. Blackwell, 2004.

[77] J. W. Grunden and K. A. Fisher. Lovastatin-induced rhabdomyolysis possibly associated with clarithromycin and azithromycin. *Ann Pharmacother*, 31(7-8):859–863, 1997.

[78] B. Gurley, M. A. Hubbard, D. K. Williams, J. Thaden, Y. Tong, W. B. Gentry, P. Breen, D. J. Carrier, and S. Cheboyina. Assessing the clinical significance of botanical supplementation on human cytochrome P450 3A activity: comparison of a milk thistle and black cohosh product to rifampin and clarithromycin. *J Clin Pharmacol*, 46(2):201–213, 2006.

[79] JH Gurwitz, TS Field, J Judge, P Rochon, LR Harrold, C Cadoret, M Lee, K White, J LaPrino, J Erramuspe-Mainard, M DeFlorio, L Gavendo, J Auger, and DW Bates. The incidence of adverse drug events in two large academic long-term facilities. *Am J Med*, 118:251–258, 2005.

[80] V. Haarslev and R. Möller. Racer: A core inference engine for the semantic web. pages 27–36, 2003.

[81] Philip D. Hansten and John R. Horn. *The Top 100 Drug Interactions, A Guide to Patient Management*, chapter Cytochrome P450 Enzymes and Drug Interactions. H&H Publications, 2004.

[82] Phillip Hansten. Drug interaction management. *Pharmacy World and Science*, 25(3), Jun 2003.

[83] M. Hauben, V. Patadia, C. Gerrits, L. Walsh, and L. Reich. Data mining in pharmacovigilance: the need for a balanced perspective. *Drug Saf*, 28(10):835–842, 2005.

[84] M. Hauben, L. Reich, E. P. Van Puijenbroek, C. M. Gerrits, and V. K. Patadia. Data mining in pharmacovigilance: lessons from phantom ships. *Eur J Clin Pharmacol*, 62(11):967–970, 2006.

[85] Thomas Hazlet, Todd A. Lee, Phillip Hansten, and John R. Horn. Performance of community pharmacy drug interaction software. *J Am Pharm Assoc*, 41(2):200–204, 2001.

[86] R. J. Herman. Drug interactions and the statins. *CMAJ*, 161(10):1281–1286, 1999.

[87] A. Hiller, K. T. Olkkola, P. Isohanni, and L. Saarnivaara. Unconsciousness associated with midazolam and erythromycin. *Br J Anaesth*, 65(6):826–828, 1990.

[88] Tyken C. Hsieh, Gilad Kuperman, Tonushree Jaggi, Patricia Hojnowski-Diaz, Julie Fiskio, Deborah Williams, David Bates, and Tejal Gandhi. Characteristics and consequences of drug allergy alert overrides in a computerized physician order entry system. *JAMIA*, 11(6):482–91, Nov/Dec 2004.

[89] T. Huynh, D. Cordato, F. Yang, T. Choy, K. Johnstone, F. Bagnall, N. Hitchens, and R. Dunn. HMG CoA reductase-inhibitor-related myopathy and the influence of drug interactions. *Intern Med J*, 32(9-10):486–490, 2002.

[90] The National Cancer Institute. The National Cancer Institute (NCI) Thesaurus. Internet, 2007. `http://nciterms.nci.nih.gov/NCIBrowser/Dictionary.do`. Last accessed 02/13/2007.

[91] H. Itakura, D. Vaughn, D. G. Haller, and P. J. O'Dwyer. Rhabdomyolysis from Cytochrome P-450 interaction of ketoconazole and simvastatin in prostate cancer. *J Urol*, 169(2):613, 2003.

[92] Kiyomi Ito, Hayley S. Brown, and J. Brian Houston. Database analysis for the prediction of in vivo drug drug interactions from in vitro data. *British Journal of Clinical Pharmacology*, 57(4):473–486, 2003.

[93] W. Jacobsen, G. Kirchner, K. Hallensleben, L. Mancinelli, M. Deters, I. Hackbarth, K. Baner, L. Z. Benet, K. F. Sewing, and U. Christians. Small intestinal metabolism of the 3-hydroxy-3-methylglutaryl-coenzyme A reductase inhibitor lovastatin and comparison with pravastatin. *J Pharmacol Exp Ther*, 291(1):131–139, 1999.

[94] W. Jacobsen, B. Kuhn, A. Soldner, G. Kirchner, K. F. Sewing, P. A. Kollman, L. Z. Benet, and U. Christians. Lactonization is the critical first step in the disposition of the 3-hydroxy-3-methylglutaryl-coa reductase inhibitor atorvastatin. *Drug Metab Dispos*, 28(11):1369–1378, 2000.

[95] T. A. Jacobson. Comparative pharmacokinetic interaction profiles of pravastatin, simvastatin, and atorvastatin when coadministered with cytochrome P450 inhibitors. *Am J Cardiol*, 94(9):1140–1146, 2004.

[96] Janssen. sporanox (itraconazole) capsule. FDA-approved drug product labeling, 01 2008. Last accessed on DailyMed 05/16/2008.

[97] D. R. Jones, J. C. Gorski, M. A. Hamman, B. S. Mayhew, S. Rider, and S. D. Hall. Diltiazem inhibition of cytochrome P-450 3A activity is due to metabolite intermediate complex formation. *J Pharmacol Exp Ther*, 290(3):1116–1125, 1999.

[98] P. H. Jones and M. H. Davidson. Reporting rate of rhabdomyolysis with fenofibrate + statin versus gemfibrozil + any statin. *Am J Cardiol*, 95(1):120–122, 2005.

[99] JR Horn JR, PD Hansten, and LN Chan. Proposal for a new tool to evaluate drug interaction cases. *Ann Pharmacother*, 41(4):674–80, 2007.

[100] D. N. Juurlink, M. Mamdani, A. Kopp, A. Laupacis, and D. A. Redelmeier. Drug-drug interactions among elderly patients hospitalized for drug toxicity. *JAMA*, 289(13):1652–1658, 2003.

[101] J. Kahri, M. Valkonen, T. Backlund, M. Vuoristo, and K. T. Kivisto. Rhabdomyolysis in a patient receiving atorvastatin and fluconazole. *Eur J Clin Pharmacol*, 60(12):905–907, 2005.

[102] T. Kantola, J. T. Backman, M. Niemi, K. T. Kivisto, and P. J. Neuvonen. Effect of fluconazole on plasma fluvastatin and pravastatin concentrations. *Eur J Clin Pharmacol*, 56(3):225–229, 2000.

[103] T. Kantola, K. T. Kivisto, and P. J. Neuvonen. Erythromycin and verapamil considerably increase serum simvastatin and simvastatin acid concentrations. *Clin Pharmacol Ther*, 64(2):177–182, 1998.

[104] IM Kapetanovic and HJ Kupferberg. Stable isotope methodology and gas chromatography mass spectrometry in a pharmacokinetic study of phenobarbital. *Biomed Mass Spectrom*, 7(2):47–52, 1980.

[105] IM Kapetanovic, HJ Kupferberg, RJ Porter, W Theodore, E Schulman, and JK Penry. Mechanism of valproate-phenobarbital interaction in epileptic patients. *Clin Pharmacol Ther*, 29(4):480–6, Apr 1981.

[106] P Karp, S Paley, C Krieger, and P Zhang. An evidence ontology for use in pathway/genome databases. In *Pacific Symposium on Biocomputing 2004*, pages 190–201, 2004.

[107] P. D. Karp, M. Riley, M. Saier, I. T. Paulsen, S. M. Paley, and A. Pellegrini-Toole. The EcoCyc and MetaCyc databases. *Nucleic Acids Res*, 28(1):56–59, 2000.

[108] K. T. Kivisto, T. Kantola, and P. J. Neuvonen. Different effects of itraconazole on the pharmacokinetics of fluvastatin and lovastatin. *Br J Clin Pharmacol*, 46(1):49–53, 1998.

[109] T.E. Klein, J.T. Chang, M.K. Cho, K.L. Easton, R. Fergerson, M. Hewett, Z. Lin, Y. Liu, S. Liu, D.E. Oliver, D.L. Rubin, F. Shafa, J.M. Stuart, and R.B. Altman. Integrating Genotype and Phenotype Information: An Overview of the PharmGKB Project. *The Pharmacogenomics Journal*, 1:167–170, 2001.

[110] K. Kosuge, M. Nishimoto, M. Kimura, K. Umemura, M. Nakashima, and K. Ohashi. Enhanced effect of triazolam with diltiazem. *Br J Clin Pharmacol*, 43(4):367–372, 1997.

[111] Y. W. Lam, C. L. Alfaro, L. Ereshefsky, and M. Miller. Pharmacokinetic and pharmacodynamic interactions of oral midazolam with ketoconazole, fluoxetine, fluvoxamine, and nefazodone. *J Clin Pharmacol*, 43(11):1274–1282, 2003.

[112] René H. Levy, Kenneth E. Thummel, William F. Trager, Philip D. Hansten, and Michel Eichelbaum, editors. *Metabolic Drug Interactions - Kenneth E. Thummel and Kent L. Kunze and Danny D. Shen*, chapter "Metabolically-Based Drug-Drug Interactions: Principles and Mechanisms". Lippincott, Williams, and Wilkens, 2000.

[113] J. J. Lewin 3rd, J. M. Nappi, and M. H. Taylor. Rhabdomyolysis with concurrent atorvastatin and diltiazem. *Ann Pharmacother*, 36(10):1546–1549, 2002.

[114] Jiunn H. Lin. Sense and nonsense in the prediction of drug-drug interactions. *Current Drug Metabolism*, 1(4):305–332, 2000.

[115] Y. K. Loke, A. N. Trivedi, and S. Singh. Meta-analysis: gastrointestinal bleeding due to interaction between selective serotonin uptake inhibitors and non-steroidal anti-inflammatory drugs. *Aliment Pharmacol Ther*, 27(1):31–40, 2008.

[116] PJ Marroum and J Gobburu. The product label: how pharmacokinetics and pharmacodynamics reach the prescriber. *Clin Pharmacokinet.*, 41(3):161–9, 2002.

[117] A. L. Mazzu, K. C. Lasseter, E. C. Shamblen, V. Agarwal, J. Lettieri, and P. Sundaresen. Itraconazole alters the pharmacokinetics of atorvastatin to a greater extent than either cerivastatin or pravastatin. *Clin Pharmacol Ther*, 68(4):391–400, 2000.

[118] C. G. Mc Donnell, S. Harte, J. O'Driscoll, C. O'Loughlin, F. N. Van Pelt, and G. D. Shorten. The effects of concurrent atorvastatin therapy on the pharmacokinetics of intravenous midazolam. *Anaesthesia*, 58(9):899–904, 2003.

[119] Merck. zocor (simvastatin) tablet, film coated. FDA-approved drug product labeling, 07 2007. Last accessed on DailyMed 04/19/2008.

[120] Merck. mevacor (lovastatin) tablet. FDA-approved drug product labeling, 01 2008. Last accessed on DailyMed 04/19/2008.

[121] George A. Miller. WordNet: a lexical database for English. *Commun. ACM*, 38(11):39–41, 1995.

[122] RA. Miller, RM. Gardner, KB. Johnson, and G. Hripscak. Clinical decision support and electronic prescribing systems: a time for responsible thought and action. *JAMIA*, 12(4):365–76, 2005.

[123] F.D. Min, B. Smyth, N. Berry, H. Lee, and B.C. Knollmann. Critical evaluation of hand-held electronic prescribing guides for physicians. In *American Society for Clinical Pharmacology and Therapeutics*, volume 75. American Society for Clinical Pharmacology and Therapeutics, 2004.

[124] J. R. Mort, R. R. Aparasu, and R. K. Baer. Interaction between selective serotonin reuptake inhibitors and nonsteroidal antiinflammatory drugs: review of the literature. *Pharmacotherapy*, 26(9):1307–1313, 2006.

[125] O. Mousa, D. C. Brater, K. J. Sunblad, and S. D. Hall. The interaction of diltiazem with simvastatin. *Clin Pharmacol Ther*, 67(3):267–274, 2000.

[126] P. Mozzicato. Standardised MedDRA queries: their role in signal detection. *Drug Saf*, 30(7):617–619, 2007.

[127] Mylan. Alprazolam (alprazolam) Tablet. FDA-approved drug product labeling, 12 2006. Last accessed on DailyMed 05/29/2008.

[128] P. J. Neuvonen and K. M. Jalava. Itraconazole drastically increases plasma concentrations of lovastatin and lovastatin acid. *Clin Pharmacol Ther*, 60(1):54–61, 1996.

[129] P. J. Neuvonen, T. Kantola, and K. T. Kivisto. Simvastatin but not pravastatin is very susceptible to interaction with the CYP3A4 inhibitor itraconazole. *Clin Pharmacol Ther*, 63(3):332–341, 1998.

[130] P. J. Neuvonen, A. Varhe, and K. T. Olkkola. The effect of ingestion time interval on the interaction between itraconazole and triazolam. *Clin Pharmacol Ther*, 60(3):326–331, 1996.

[131] R. Scott Obach, Robert L. Walsky, Karthik Venkatakrishnan, J. Brian Houston, and Larry Tremaine. In vitro cytochrome P450 inhibition data and the prediction of drug-drug interactions: qualitative relationships, quantitative predictions, and the rank-order approach. *Clin Pharmacol Ther*, 78:582–92, 2005.

[132] National Library of Medicine. Medical subject headings. Internet, 2007. `http://www.nlm.nih.gov/mesh/`. Last accessed 02/13/2007.

[133] Y. Ohno, A. Hisaka, and H. Suzuki. General framework for the quantitative prediction of cyp3a4-mediated oral drug interactions based on the auc increase by coadministration of standard drugs. *Clin Pharmacokinet*, 46(8):681–696, 2007.

[134] K. T. Olkkola, J. Ahonen, and P. J. Neuvonen. The effects of the systemic antimycotics, itraconazole and fluconazole, on the pharmacokinetics and pharmacodynamics of intravenous and oral midazolam. *Anesth Analg*, 82(3):511–516, 1996.

[135] K. T. Olkkola, K. Aranko, H. Luurila, A. Hiller, L. Saarnivaara, J. J. Himberg, and P. J. Neuvonen. A potentially hazardous interaction between erythromycin and midazolam. *Clin Pharmacol Ther*, 53(3):298–305, 1993.

[136] K. T. Olkkola, J. T. Backman, and P. J. Neuvonen. Midazolam should be avoided in patients receiving the systemic antimycotics ketoconazole or itraconazole. *Clin Pharmacol Ther*, 55(5):481–485, 1994.

[137] Committee on Identifying and Preventing Medication Errors. Preventing medication errors. Technical report, Institute of Medicine, 2006. 0309102685.

[138] J. Y. Park, K. A. Kim, P. W. Park, O. J. Lee, D. K. Kang, J. H. Shon, K. H. Liu, and J. G. Shin. Effect of CYP3A5*3 genotype on the pharmacokinetics and pharmacodynamics of alprazolam in healthy subjects. *Clin Pharmacol Ther*, 79(6):590–599, 2006.

[139] Judea Pearl and Stuar Russell. Bayesian networks. Technical report, UCLA Cognitive Systems Laboratory, November 2000.

[140] R. Peces and A. Pobes. Rhabdomyolysis associated with concurrent use of simvastatin and diltiazem. *Nephron*, 89(1):117–118, 2001.

[141] J. P. Phillips, E. J. Antal, and R. B. Smith. A pharmacokinetic drug interaction between erythromycin and triazolam. *J Clin Psychopharmacol*, 6(5):297–299, 1986.

[142] Sheldon H. Preskorn. How drug-drug interactions can impact managed care. *The American Journal of Managed Care*, 10(6 Suppl):S186–S198, July 2004.

[143] T. Prueksaritanont, L. M. Gorham, B. Ma, L. Liu, X. Yu, J. J. Zhao, D. E. Slaughter, B. H. Arison, and K. P. Vyas. In vitro metabolism of simvastatin in humans [SBT]identification of metabolizing enzymes and effect of the drug on hepatic P450s. *Drug Metab Dispos*, 25(10):1191–1199, 1997.

[144] T. Prueksaritanont, J. M. Vega, J. D. Rogers, K. Gagliano, H. E. Greenberg, L. Gillen, M. J. Brucker, D. McLoughlin, P. H. Wong, and S. A. Waldman. Simvastatin does not affect CYP3A activity, quantified by the erythromycin breath test and oral midazolam pharmacokinetics, in healthy male subjects. *J Clin Pharmacol*, 40(11):1274–1279, 2000.

[145] J. Pustejovsky, J. Castano, J. Zhang, M. Kotecki, and B. Cochran. Robust relational parsing over biomedical literature: extracting inhibit relations. In *Pacific Symposium on Biocomputing*, pages 362–373, 2002.

[146] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2006. ISBN 3-900051-07-0.

[147] Richard Reichley, Terry Seaton, Ervina Resetar, Scott Micek, Karen Scott, Victoria Fraser, Clairborne Dunagan, and Thomas Bailey. Implementing a commercial rule base as a medication order safety net. *JAMIA*, 12(4):383–389, 2005.

[148] E. Resetar, R. M. Reichley, L. A. Noirot, W. C. Dunagan, and T. C. Bailey. Customizing a commercial rule base for detecting drug-drug interactions. In *AMIA Annu Symp Proc*, pages 1094–1094, 2005.

[149] T. Richter, T. E. Murdter, G. Heinkele, J. Pleiss, S. Tatzel, M. Schwab, M. Eichelbaum, and U. M. Zanger. Potent mechanism-based inhibition of human CYP2B6 by clopidogrel and ticlopidine. *J Pharmacol Exp Ther*, 308(1):189–197, 2004.

[150] J. Roach, S. Lee, J. Wilcke, and M. Ehrich. An expert system for information on pharmacology and drug interactions. *Comput Biol Med*, 15(1):11–23, 1985.

[151] A. D. Rodrigues, E. M. Roberts, D. J. Mulford, Y. Yao, and D. Ouellet. Oxidative metabolism of clarithromycin in the presence of human liver microsomes. Major role for the cytochrome P4503A (CYP3A) subfamily. *Drug Metab Dispos*, 25(5):623–630, 1997.

[152] D. L. Rubin, M. Carrillo, M. Woon, J. Conroy, T. E. Klein, and R. B. Altman. A resource to acquire and summarize pharmacogenetics knowledge in the literature. *Medinfo*, 11(Pt 2):793–797, 2004.

[153] D. L. Rubin, C. F. Thorn, T. E. Klein, and R. B. Altman. A statistical approach to scanning the biomedical literature for pharmacogenetics knowledge. *J Am Med Inform Assoc*, 12(2):121–129, 2005.

[154] Stuart Russell and Peter Norvig. *Artificial Intelligence: a Modern Approach*, chapter Logical Agents. Prentice Hall, 2003.

[155] A. Rzhetsky, I. Iossifov, T. Koike, M. Krauthammer, P. Kra, M. Morris, H. Yu, P. A. Duboue, W. Weng, W. J. Wilbur, V. Hatzivassiloglou, and C. Friedman. Geneways: a system for extracting, analyzing, visualizing, and integrating molecular pathway data. *J Biomed Inform*, 37(1):43–53, 2004.

[156] J. Schmider, J. Brockmoller, G. Arold, S. Bauer, and I. Roots. Simultaneous assessment of CYP3A4 and CYP1A2 activity in vivo with alprazolam and caffeine. *Pharmacogenetics*, 9(6):725–734, 1999.

[157] A. Shaukat, M. Benekli, G. D. Vladutiu, J. L. Slack, M. Wetzler, and M. R. Baer. Simvastatin-fluconazole causing rhabdomyolysis. *Ann Pharmacother*, 37(7-8):1032–1035, 2003.

[158] J. Shi, S. Chapel, G. Montay, P. Hardy, J. S. Barrett, D. Sica, S. K. Swan, R. Noveck, B. Leroy, and V. O. Bhargava. Effect of ketoconazole on the pharmacokinetics and safety of telithromycin and clarithromycin in older subjects with renal impairment. *Int J Clin Pharmacol Ther*, 43(3):123–133, 2005.

[159] P. H. Siedlik, S. C. Olson, B. B. Yang, and R. H. Stern. Erythromycin coadministration increases plasma atorvastatin concentrations. *J Clin Pharmacol*, 39(5):501–504, 1999.

[160] W. D. Smith, R. C. Hatton, A. L. Fann, M. A. Baz, and B. Kaplan. Evaluation of drug interaction software to identify alerts for transplant medications. *Ann Pharmacother*, 39(1):45–50, 2005.

[161] D. H. Spach, J. E. Bauwens, C. D. Clark, and W. G. Burke. Rhabdomyolysis associated with lovastatin and erythromycin use. *West J Med*, 154(2):213–215, 1991.

[162] C. A. Stein, S. Goel, and R. Ghavamian. Hepatitis and rhabdomyolysis in a patient with hormone refractory prostate cancer on ketoconazole and concurrent lovastatin therapy. *Invest New Drugs*, 25(3):277–278, 2007.

[163] Ana Szarfman, Stella Machado, and Robert O' Neill. Use of Screening Algorithms and Computer Systems to Efficiently Signal Higher-Than-Expected Combinations of Drugs and Events in the US FDAs Spontaneous Reports Database. *Drug Safety*, 25(6):381–392, 2002.

[164] Teva. nefazodone hydrochloride (Nefazodone Hydrochloride) tablet. FDA-approved drug product labeling, 11 2006. Last accessed on DailyMed 05/29/2008.

[165] Teva. Alprazolam (alprazolam) Tablet, Extended Release. FDA-approved drug product labeling, 05 2007. Last accessed on DailyMed 05/29/2008.

[166] M. Turpeinen, A. Tolonen, J. Uusitalo, J. Jalonen, O. Pelkonen, and K. Laine. Effect of clopidogrel and ticlopidine on cytochrome p450 2b6 activity as measured by bupropion hydroxylation. *Clin Pharmacol Ther*, 77(6):553–559, 2005.

[167] Han van de Waterbeemd and Eric Gifford. ADMET in silico modeling: towards prediction paradise? *Nat Rev Drug Discov*, 2(3):192–204, March 2003.

[168] H. van der Sijs, J. Aarts, A. Vulto, and M. Berg. Overriding of drug safety alerts in computerized physician order entry. *J Am Med Inform Assoc*, 13(2):138–147, 2006.

[169] E. N. van Roon, S. Flikweert, M. le Comte, P. N. Langendijk, W. J. Kwee-Zuiderwijk, P. Smits, and J. R. Brouwers. Clinical relevance of drug-drug interactions : a structured assessment procedure. *Drug Saf*, 28(12):1131–1139, 2005.

[170] A. Varhe, K. T. Olkkola, and P. J. Neuvonen. Oral triazolam is potentially hazardous to patients receiving systemic antimycotics ketoconazole or itraconazole. *Clin Pharmacol Ther*, 56(6 Pt 1):601–607, 1994.

[171] A. Varhe, K. T. Olkkola, and P. J. Neuvonen. Diltiazem enhances the effects of triazolam by inhibiting its metabolism. *Clin Pharmacol Ther*, 59(4):369–375, 1996.

[172] A. Varhe, K. T. Olkkola, and P. J. Neuvonen. Fluconazole, but not terbinafine, enhances the effects of triazolam by inhibiting its metabolism. *Br J Clin Pharmacol*, 41(4):319–323, 1996.

[173] S. Vickers, C. A. Duncan, K. P. Vyas, P. H. Kari, B. Arison, S. R. Prakash, H. G. Ramjit, S. M. Pitzenberger, G. Stokker, and D. E. Duggan. In vitro and in vivo biotransformation of simvastatin, an inhibitor of HMG CoA reductase. *Drug Metab Dispos*, 18(4):476–483, 1990.

[174] L. L. von Moltke, D. J. Greenblatt, J. S. Harmatz, S. X. Duan, L. M. Harrel, M. M. Cotreau-Bibbo, G. A. Pritchard, C. E. Wright, and R. I. Shader. Triazolam biotransformation by human liver microsomes in vitro: effects of metabolic inhibitors and clinical confirmation of a predicted interaction with ketoconazole. *J Pharmacol Exp Ther*, 276(2):370–379, 1996.

[175] L. L. von Moltke, D. J. Greenblatt, J. Schmider, C. E. Wright, J. S. Harmatz, and R. I. Shader. In vitro approaches to predicting drug interactions in vivo. *Biochem Pharmacol*, 55(2):113–122, 1998.

[176] X Wen, JS Wang, KT Kivisto, PJ Neuvonen, and JT Backman. In vitro evaluation of valproic acid as an inhibitor of human cytochrome P450 isoforms: preferential inhibition of cytochrome P450 2C9 (CYP2C9). *Br J Clin Pharmacol.*, 52(5):547–53, 2001.

[177] Suzanne West, Valerie King, Timothy Carey, Kathleen Lohr, Nikki McKoy, Sonya Sutton, and Linda Lux. Systems to rate the strength of scientific evidence. Technical Report 02-E016, Agency for Healthcare Research and Quality, 2002.

[178] J. A. Williams, B. J. Ring, V. E. Cantrell, D. R. Jones, J. Eckstein, K. Ruterbories, M. A. Hamman, S. D. Hall, and S. A. Wrighton. Comparative metabolic capabilities of CYP3A4, CYP3A5, and CYP3A7. *Drug Metab Dispos*, 30(8):883–891, 2002.

[179] P. W. Wong, T. A. Dillard, and K. Kroenke. Multiple organ toxicity from addition of erythromycin to long-term lovastatin therapy. *South Med J*, 91(2):202–205, 1998.

[180] Fumiyoshi Yamashita and Mitsuru Hashida. In silico approaches for predicting ADME properties of drugs. *Drug Metab Phamacokin*, 19(5):327–338, 2004.

[181] N. Yasui, T. Kondo, K. Otani, H. Furukori, S. Kaneko, T. Ohkubo, T. Nagasaki, and K. Sugawara. Effect of itraconazole on the single oral dose pharmacokinetics and pharmacodynamics of alprazolam. *Psychopharmacology (Berl)*, 139(3):269–273, 1998.

[182] N. Yasui, K. Otani, S. Kaneko, T. Ohkubo, T. Osanai, K. Sugawara, K. Chiba, and T. Ishizaki. A kinetic and dynamic study of oral alprazolam with and without erythromycin in humans: in vivo evidence for the involvement of CYP3A4 in alprazolam metabolism. *Clin Pharmacol Ther*, 59(5):514–519, 1996.

[183] R. Yuan, T. Parmelee, J. D. Balian, R. S. Uppoor, F. Ajayi, A. Burnett, L. J. Lesko, and P. Marroum. In vitro metabolic interaction studies: experience of the Food and Drug Administration. *Clin Pharmacol Ther*, 66(1):9–15, 1999.

[184] T. Zimmermann, R. A. Yeates, H. Laufen, F. Scharpf, M. Leitold, and A. Wildfeuer. Influence of the antibiotics erythromycin and azithromycin on the pharmacokinetics and pharmacodynamics of midazolam. *Arzneimittelforschung*, 46(2):213–217, 1996.

Appendix A

HOW BIG IS THE GAP IN SCIENTIFIC KNOWLEDGE ABOUT DRUG-DRUG INTERACTIONS?

As a thought experiment, consider that a query of the Federal Drug Administrations (FDA) `drugs@fda` database [58] of all currently approved prescription and over-the-counter drugs identified about 1300 unique drugs used in more than 7000 drug products[a]. A simple calculation reveals that there are nearly 1.7 million pairwise DDIs possible if each drug is considered as a possible cause of an interaction involving one other drug (Equation A.1).

$$2 * \binom{1300}{2} = 1,688,700 \tag{A.1}$$

As of the time of this writing, a simple query of the PubMed database of biomedical research abstracts[b] for any study investigating drug interactions [c] returns approximately 450,000 abstracts. Let's make the unrealistic assumption that each abstract represents a study exploring the possibility that each drug in a pair drugs drawn without replacement from the 1300 APIs could be the victim of a drug-drug interaction involving the other drug. There would still be nearly 800,000 $(1,688,700 - 2 * 450,000)$ unstudied potential interactions. This would be a dramatic underestimate of the total number of unstudied interactions because it fails to factor in the active metabolites of each drug, each of which might have a different interaction profile.

---

[a]We made this estimate by searching the drugs@fda database on 06/24/2006 for all the unique active pharmaceutical ingredients used drug products currently on the US market then reducing this list manually by collapsing multiple versions of individual active pharmaceutical ingredients to a single entry.

[b]http://www.ncbi.nlm.nih.gov/PubMed/

[c]Query: `(Drug Interactions [MeSH Terms] OR interaction [Text Word])`

Appendix B

THE DIKB'S RULE-BASED MODEL OF DDIS OCCURRING BY METABOLITE INHIBITION

The following listing is the complete set of rules that comprise the DIKB's current model of DDIs occurring by metabolite inhibition. Uppercase words within rule predicates represent assertion types with a defined semantics. For example, `1-is-an-IN-VITRO-SELECTIVE--INHIBITOR-of-2` contains the uppercase words `IN-VITRO-SELECTIVE-INHIBITOR` an assertion type defined in Appendix C.

B.1 Rules that Model Metabolic Inhibition

```
;; a necessary condition for being an 'in vitro
;; selective inhibitor' is that the agent is also
;; an inhibitor
(rule
 ((:IN (1-is-an-IN-VIVO-SELECTIVE-INHIBITOR-of-2 ?x ?y)))
 (rassert!
  (1-INHIBITS-2 ?x ?y)
  (nil
   (1-is-an-IN-VIVO-SELECTIVE-INHIBITOR-of-2 ?x ?y)
   )))

   ;; a necessary condition of some active ingredient
   ;; or compound having a primary total clearance
   ;; enzyme is that it is a substrate of that enzyme
   (rule
    ((:IN (primary-total-clearance-enzyme-of-1-is-2 ?x ?y)))
    (rassert!
     (1-is-substrate-of-2 ?x ?y)
     (nil
      (primary-total-clearance-enzyme-of-1-is-2 ?x ?y)
      )))
```

```
;; a necessary condition of some active ingredient
;; or compound having a primary total clearance
;; enzyme is that it is primarily cleared by metabolism
(rule
((:IN (primary-total-clearance-enzyme-of-1-is-2 ?x ?y)))
 (rassert!
  (primary-total-clearance-mechanism-of-1-is-2 ?x 'METABOLIC-CLEARANCE)
   (nil
    (primary-total-clearance-enzyme-of-1-is-2 ?x ?y)
   )))
```

```
;; a rule that makes it a contradiction for an active ingredient
;; or compound to both permanently and not permanently deactivate the catalytic
;; function of an enzyme
(rule
 ((:IN
   (1-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2 ?drug1 ?enzyme))
  (:IN
   (1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2
        ?drug1 ?enzyme)))
 (contradiction
   (eval (quotize (list
    '1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2
        ?drug1 ?enzyme)))))

;; a rule for establishing that an active ingredient or metabolite
;; *does* inhibit an enzyme based on in vitro evidence
(rule
 ((:IN (INHIBITION-CONSTANT-of-1-for-2-is-3 ?x ?y ?k_i))
  (:IN
    (1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2 ?x ?y))
  (:IN (MAXIMUM-IN-VIVO-CONCENTRATION-of-1-is-2 ?x ?c_max)
       :TEST (> (float (/ ?c_max ?k_i )) .1)))
 (rassert! (1-INHIBITS-2 ?x ?y)
           (nil
            ;;justifications
            (INHIBITION-CONSTANT-of-1-for-2-is-3 ?x ?y ?k_i)
            (1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2
                                    ?x ?y)
            (MAXIMUM-IN-VIVO-CONCENTRATION-of-1-is-2 ?x ?c_max)
            (accept-in-vitro-based-enzyme-modulation-assertions)
            )))
```

```
;; a rule for when a metabolic transformation is
;; inhibited by inhibition of a *known*
;; pathway. NOTE: This rule could explicitly ignore
;; inhibition a metabolite's own production itself
;; if a test were added to one of the antecedents:
;; :TEST (not (equal ?q ?y))
(rule
 ((:IN (1-has-metabolite-2-via-3 ?x ?y ?z))
  (:IN (1-INHIBITS-2 ?q ?z)))
 (rassert!
  (1-inhibits-transformation-of-2-to-3-via-4 ?q ?x ?y ?z)
  (nil
   (1-has-metabolite-2-via-3 ?x ?y ?z)
   (1-INHIBITS-2 ?q ?z)
   )))
```

```
;; a rule for when an active ingredient or metabolite, ?x, will
;; not inhibit the metabolic clearance  of another drug, ?z,
;; because ?x does not inhibit enzyme ?y's ability to catalyze
;; drug ?z
(rule
 ((:IN (1-DOES-NOT-INHIBIT-2 ?x ?y))
  (:IN (1-is-SUBSTRATE-OF-2 ?z ?y)))
 (rassert!
  (1-does-not-inhibit-the-metabolic-clearance-of-2-via-3 ?x ?z ?y)
  (nil
   (1-DOES-NOT-INHIBIT-2 ?x ?y)
   (1-is-SUBSTRATE-OF-2 ?z ?y)
   )))
```

```
;; a rule for when an active ingredient or metabolite, ?x, will
;; not inhibit the metabolic clearance  of another drug, ?z,
;; because ?z is not a substrate of enzyme ?y
(rule
 ((:IN (1-inhibits-2 ?x ?y))
  (:IN (1-is-not-a-substrate-of-2 ?z ?y)))
 (rassert!
  (1-does-not-inhibit-the-metabolic-clearance-of-2-via-3 ?x ?z ?y)
  (nil
   (1-inhibits-2 ?x ?y)
   (1-is-not-a-substrate-of-2 ?z ?y)
   )))
```

```
;; a rule for establishing that an active ingredient or metabolite
;; *does not* inhibit an enzyme based on in vitro evidence
(rule
 ((:IN (INHIBITION-CONSTANT-of-1-for-2-is-3 ?x ?y ?k_i))
  (:IN (1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2 ?x ?y))
  (:IN (MAXIMUM-IN-VIVO-CONCENTRATION-of-1-is-2 ?x ?c_max)
       :TEST (<= (float (/ ?c_max ?k_i )) .1)))
 (rassert! (1-DOES-NOT-INHIBIT-2 ?x ?y)
           (nil
            ;;justifications
            (INHIBITION-CONSTANT-of-1-for-2-is-3 ?x ?y ?k_i)
            (1-DOES-NOT-PERMANENTLY-DEACTIVATES-CATALYTIC-FUNCTION-of-2 ?x ?y)
            (MAXIMUM-IN-VIVO-CONCENTRATION-of-1-is-2 ?x ?c_max)
            (accept-in-vitro-based-enzyme-modulation-assertions)
            )))
```

```
;; a rule for that makes it a contradiction for an active ingredient
;; or metabolite to both inhibit and not inhibit the catalytic
;; function of an enzyme
(rule
 ((:IN (1-INHIBITS-2 ?x ?y))
  (:IN (1-DOES-NOT-INHIBIT-2 ?x ?y)))
 (contradiction
  (eval (quotize (list
      '1-DOES-NOT-INHIBIT-2 ?drug1 ?enzyme)))))
```

```
;; a rule for that makes it a contradiction for an active ingredient
;; or compound to be  and *not* be a substrate of an enzyme
(rule
 ((:IN (1-is-substrate-of-2 ?drug ?enzyme))
  (:IN (1-is-not-substrate-of-2 ?drug ?enzyme)))
 (contradiction
  (eval (quotize (list
                  '1-is-not-substrate-of-2 ?drug ?enzyme)))))
```

```
;; Some, possibly negligible, inhibition of
;; metabolic clearance of active ingredient or
;; metabolite ?z by active ingredient or metabolite
;; ?x due to ?x's inhibition of enzyme ?y's ability
;; to catalyze ?z. NOTE: this test ignores cases
;; where a drug INHIBITS itself
(rule
 ((:IN (1-INHIBITS-2 ?x ?y))
  (:IN (1-is-SUBSTRATE-OF-2 ?z ?y)
       :TEST (not (equal ?x ?z))))
 (rassert!
  (1-INHIBITS-METABOLIC-CLEARANCE-of-2-via-3 ?x ?z ?y)
  (nil
   (1-INHIBITS-2 ?x ?y)
   (1-is-SUBSTRATE-OF-2 ?z ?y)
   )))
```

```
;; A more significant inhibition of metabolic clearance
;; that should lead to a greater *minimum* increase in AUC
;; than the 1-INHIBITS-METABOLIC-CLEARANCE-of-2-via-3 assertion captures.
;; This models the effect of inhibiting an enzyme that is responsible
;; for .25 of a drug's total clearance by requiring inhibition of an enzyme
;; responsible for at least .50 of a drug's *metabolic* clearance when that
;; form of clearance is responsible for at least .50 of the drug's
;; *total* clearance
(rule
 ((:IN
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y)
   :TEST (not (equal ?x ?z)))
  (:IN
   (PRIMARY-TOTAL-CLEARANCE-MECHANISM-of-1-is-2 ?z
                                                'METABOLIC-CLEARANCE))
  (:IN
   (PRIMARY-METABOLIC-CLEARANCE-ENZYME-of-1-is-2 ?z ?y)))
 (rassert!
  (1-inhibits-3-the-primary-metabolic-enzyme-of-2 ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y)
   (PRIMARY-TOTAL-CLEARANCE-MECHANISM-of-1-is-2 ?z
                                                'METABOLIC-CLEARANCE)
   (PRIMARY-METABOLIC-CLEARANCE-ENZYME-of-1-is-2 ?z ?y)
   )))
```

```
;; This rule models inhibition of metabolic clearance that should lead to
;; a greater *minimum* increase in AUC than the
;; 1-INHIBITS-3-the-primary-metabolic-enzyme-of-2 assertion captures.
;; If one enzyme is responsible for at least .50  of the
;; metabolic clearance of a drug and another drug fully INHIBITS that enzyme
;; then, one would expect at least at least a .50 decrease in clearance and,
;; subsequently, at least a 2-fold increase in AUC.
(rule
 ((:IN
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y)
   :TEST (not (equal ?x ?z)))
  (:IN
   (PRIMARY-TOTAL-CLEARANCE-ENZYME-of-1-is-2 ?z ?y)))
 (rassert!
  (1-inhibits-3-the-primary-total-clearance-enz-of-2 ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y)
   (PRIMARY-TOTAL-CLEARANCE-ENZYME-of-1-is-2 ?z ?y)
   )))
```

```
;; This rule models inhibition of metabolic clearance that should lead to
;; a greater *maximum* increase in AUC than the
;; inhibit-primary-tot-clearance-enz assertion captures.
;; It predicts a drastic increase in AUC for active
;; ingredients that undergo a high degree first-pass metabolism
(rule
 ((:IN
   (1-inhibits-3-the-primary-total-clearance-enz-of-2 ?x ?z ?y))
  (:IN
   (FIRST-PASS-EFFECT-on-1-is-2 ?z 'HIGH)))
 (rassert!
  (met-inhibit-drug-w-high-first-pass ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-3-the-primary-total-clearance-enz-of-2 ?x ?z ?y)
   (FIRST-PASS-EFFECT-on-1-is-2 ?z 'HIGH)
   )))


;; a rule defining some, possibly negligible, inhibition
;; of clearance for a pceut-entity-of-concern
(rule
 ((:IN
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y))
  (:IN
   (1-is-PCEUT-ENTITY-OF-CONCERN ?z)))
 (rassert!
  (first-level-metabolic-inhibition-of-pceut-entity-of-concern ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-metabolic-clearance-of-2-via-3 ?x ?z ?y)
   (1-is-PCEUT-ENTITY-OF-CONCERN ?z)
   )))
```

```
;; rules defining when the inhibition of a pceut-entity-of-concern
;; clearance should lead to a more significant increase in AUC
;; than that captured by
;; first-level-metabolic-inhibition-of-pceut-entity-of-concern
;; assertions
(rule
 ((:IN
   (1-inhibits-3-the-primary-metabolic-enzyme-of-2 ?x ?z ?y))
  (:IN
   (1-is-PCEUT-ENTITY-OF-CONCERN ?z)))
 (rassert!
  (second-level-metabolic-inhibition-of-pceut-entity-of-concern ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-3-the-primary-metabolic-enzyme-of-2 ?x ?z ?y)
   (1-is-PCEUT-ENTITY-OF-CONCERN ?z)
   )))



(rule
 ((:IN
   (1-inhibits-3-the-primary-total-clearance-enz-of-2 ?x ?z ?y))
  (:IN (1-is-PCEUT-ENTITY-OF-CONCERN ?z)))
 (rassert!
  (second-level-metabolic-inhibition-of-pceut-entity-of-concern ?x ?z ?y)
  (nil
   ;;justifications
   (1-inhibits-3-the-primary-total-clearance-enz-of-2 ?x ?z ?y)
   (1-is-PCEUT-ENTITY-OF-CONCERN ?z)
   )))
```

B.2 *Rules for Linking Metabolites to Active Ingredients and Ancestor Compounds*

```
;; a rule linking an parent compound to an metabolite
(rule
 ((:IN (1-has-metabolite-2-via-3 ?x ?y ?z)))
 (rassert!
  (1-is-ANCESTOR-OF-2 ?x ?y)
  (nil
   (1-has-metabolite-2-via-3 ?x ?y ?z)
   )))


   (rule
    ((:IN (1-has-metabolite-2-via-3 ?x ?y ?z)))
    (rassert!
     (1-is-SUBSTRATE-OF-2 ?x ?z)
     (nil
      (1-has-metabolite-2-via-3 ?x ?y ?z)
      )))


   ;; a rule linking the catalysis of the formation of a
   ;; metabolite to parent compounds
   (rule
    ((:IN (1-CONTROLS-FORMATION-of-2 ?enz ?x))
     (:IN (1-HAS-METABOLITE-2 ?y ?x)))
    (rassert!
     (1-has-metabolite-2-via-3 ?y ?x ?enz)
     (nil
      (1-CONTROLS-FORMATION-of-2 ?enz ?x)
      (1-HAS-METABOLITE-2 ?y ?x)
      )))
```

```
;; a rule linking an ancestor compound to an metabolite
(rule
 ((:IN (1-has-metabolite-2-via-3 ?x ?y ?e))
  (:IN (1-is-ANCESTOR-OF-2 ?z ?x)))
 (rassert!
  (1-is-ANCESTOR-OF-2 ?z ?y)
  (nil
   (1-has-metabolite-2-via-3 ?x ?y ?e)
   (1-is-ANCESTOR-OF-2 ?z ?x)
   )))
```

B.3 Modeling the Effect of Inhibition Through a Graph of Catalytic Reactions

All of these rules assume that alternate clearance pathways are not saturated.

```
;; inhibition of the formation of a metabolite
;; upstream affects the formation of all metabolites
;; downstream
(rule
 ((:IN
   (1-inhibits-transformation-of-2-to-3-via-4 ?q ?x ?m1 ?enz))
  (:IN
   (1-is-ANCESTOR-OF-2 ?m1 ?m2)))
 (rassert!
  (1-INHIBITS-transformation-of-2-to-3-via-4-upstream ?q ?x ?m2 ?enz)
  (nil
   (1-INHIBITS-transformation-of-2-to-3-via-4 ?q ?x ?m1 ?enz)
   (1-is-ANCESTOR-OF-2 ?m1 ?m2)
   )))
```

```
;; if the formation of two different metabolites, M1
 ;; and M2, from the same agent, X, is catalyzed by
 ;; *different* enzymes then, the effect on M2 of
 ;; modulating the clearance of X by inhibiting or
 ;; inducing the catalytic function of one of the
 ;; enzymes will be an non-ambiguous increase or
 ;; decrease
 (rule
  ((:IN (1-has-metabolite-2-via-3 ?x ?m1 ?enz1))
   (:IN (1-has-metabolite-2-via-3 ?x ?m2 ?enz2)
    :TEST (and (not (equal ?m1 ?m2))
               (not (equal ?enz1 ?enz2))
               (not (equal ?enz1 'UNKNOWN)))))
  (assume!
    (eval
     (quotize
       (list
       'effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
             ?m2 ?x ?enz1)))
               'default-inference-assumption))
```

```
;; If the effect on some metabolite, M1, of
;; modulating the clearance of its parent compound,
;; X, by inhibiting or inducing the catalytic
;; function of some enzyme, E, is an unambiguous
;; increase or decrease and if M1 has a metabolite,
;; M2, and the transformation of M1 to M2 is
;; controlled by a different enzyme than E then,
;; then an increase or decrease in X will effect an
;; non-ambiguous increase M2
(rule
 ((:IN
   (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
      ?m1 ?x ?enz1))
  (:IN (1-has-metabolite-2-via-3 ?m1 ?m2 ?enz2)
   :TEST (and (not (equal ?enz1 ?enz2))
              (not (equal ?enz1 'UNKNOWN)))))
 (assume!
  (eval
   (quotize
     (list
     'effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
             ?m2 ?x ?enz1)))
     'default-inference-assumption))
```

```
(rule
 ((:IN (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
                                ?m ?x ?enz))
  (:IN (1-inhibits-2 ?q ?enz)))
 (rassert!
  (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
                                ?q ?m ?x ?enz)
  (nil
   (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
                                ?m ?x ?enz)
   (1-inhibits-2 ?q ?enz)
   )))
```

```
;; The effect of an increased formation of a parent
;; compound, X, on some metabolite, M1, due to
;; reduced clearance by an alternate pathway is to
;; increase formation of M2 when the enzymes
;; involved in the formation of M1 and M2 are both
;; different then the enzyme whose inhibition caused
;; an increase in X
(rule
 (
  (:IN
    (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
      ?q ?m1 ?x ?enz1))
  (:IN (1-has-metabolite-2-via-3 ?m1 ?m2 ?enz2)
   :TEST (and (not (equal ?enz1 ?enz2))
              (not (equal ?enz1 'UNKNOWN))))
  (:IN
   (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
     ?m2 ?x ?enz1)))
 (rassert!
  (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
     ?q ?m2 ?x ?enz1)
  (nil
   (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
      ?q ?m1 ?x ?enz1)
   (1-has-metabolite-2-via-3 ?m1 ?m2 ?enz2)
   (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
     ?m2 ?x ?enz1)
   )))
```

```
;; Ambiguous and non-ambiguous effects are mutually
;; exclusive. Since an non-ambiguous effect is the
;; default assumption, it is retracted
(rule
 ((:IN
    (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
      ?m ?q ?x ?z))
  (:IN
    (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
     ?m ?x ?z)))
 (rretract!
    (effect-on-1-of-modulating-the-clearance-of-2-via-3-is-non-ambiguous
      ?m ?x ?z)
       default-inference-assumption))


;; If the effect of reducing the clearance of
;; metabolite is uncertain for a given metabolite,
;; it will be so for all metabolites downstream in
;; the metabolic pathway
(rule
 (
  (:IN
   (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m1 ?q ?x ?enz))
  (:IN (1-is-ancestor-of-2 ?m1 ?m2)))
 (rassert!
  (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m2 ?q ?x ?enz)
  (nil
   (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m1 ?q ?x ?enz)
   (1-is-ancestor-of-2 ?m1 ?m2)
   )))
```

```
;; It is a contradiction to have an ambiguous effect and a clearly
;; identified effect
(rule
 (
  (:IN
    (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
      ?m ?q ?x ?z))
  (:IN
    (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
      ?q ?m ?x ?z)))
 (contradiction
    (eval
      (quotize
         (list '1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
            ?q ?m ?x ?z)))))
```

```
;; If the formation of two different metabolites
;; from the same agent are catalyzed by *the same enzyme*
;; then the effect of inhibiting the enzyme
;; on both metabolites is ambiguous. This is because
;; there is both an increase in parent compound due
;; to removal of one clearance pathway and a
;; decrease in the ability of the enzyme formation
;; of child compound
(rule
 ((:IN (1-has-metabolite-2-via-3 ?x ?m1 ?z))
  (:IN (1-has-metabolite-2-via-3 ?x ?m2 ?z) :TEST (not (equal ?m1 ?m2)))
  (:IN (1-inhibits-2 ?q ?z)))
 (rassert!
  (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m2 ?q ?x ?z)
  (nil
   (1-has-metabolite-2-via-3 ?x ?m1 ?z)
   (1-has-metabolite-2-via-3 ?x ?m2 ?z)
   (1-inhibits-2 ?q ?z)
   )))
```

```
;; If it is not known if the formation of two
;; different metabolites from the same agent are
;; catalyzed by *the same enzyme* then the effect of
;; inhibiting the enzyme on both metabolites is
;; ambiguous. This is because there might be both an
;; increase in parent compound due to removal of one
;; clearance pathway and a decrease in the ability
;; of the enzyme formation of child compound
(rule
 ((:IN (1-has-metabolite-2-via-3 ?x ?m1 ?z))
  (:IN (1-has-metabolite-2-via-3 ?x ?m2 'UNKNOWN)
     :TEST (not (equal ?m1 ?m2)))
  (:IN (1-inhibits-2 ?q ?z)))
 (rassert!
  (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m2 ?q ?x ?z)
  (nil
   (1-has-metabolite-2-via-3 ?x ?m1 ?z)
   (1-has-metabolite-2-via-3 ?x ?m2 'UNKNOWN)
   (1-inhibits-2 ?q ?z)
     )))
```

```
;; The effect of an increased formation of a parent
;; compound on a metabolite due to reduced clearance
;; of an alternate pathway is unclear if the same
;; enzyme is inhibited in both the alternate pathway
;; and the formation of the metabolite
(rule
 ((:IN
   (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
     ?q ?m1 ?x ?enz))
  (:IN (1-has-metabolite-2-via-3 ?m1 ?m2 ?enz)))
 (rassert!
  (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m2 ?q ?x ?enz)
  (nil
   (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
     ?q ?m1 ?x ?enz)
   (1-has-metabolite-2-via-3 ?m1 ?m2 ?enz)
   )))
```

```
;; The effect of an increased formation of a parent
;; compound on a metabolite due to reduced clearance
;; of an alternate pathway is unclear if is not
;; known whether or not the same enzyme is inhibited
;; in both the alternate pathway and the formation
;; of the metabolite
(rule
 ((:IN
    (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
     ?q ?m1 ?x ?enz))
  (:IN (1-has-metabolite-2-via-3 ?m1 ?m2 'UNKNOWN)))
 (rassert!
  (effect-on-1-of-2-reducing-the-clearance-of-3-via-4-is-ambiguous
     ?m2 ?q ?x ?enz)
  (nil
   (1-effects-an-increase-in-2-by-reducing-clearance-of-3-via-4
    ?q ?m1 ?x ?enz)
   (1-has-metabolite-2-via-3 ?m1 ?m2 'UNKNOWN)
    )))
```

B.4 Rules for Disjunctive Cases

A set of rules to for the disjunctive case where an active ingredient is ancestor to a compound that interacts with another active ingredient or metabolite.

```
(rule
 ((:IN (1-is-an-ACTIVE-INGREDIENT ?x))
  (:IN (1-is-ANCESTOR-OF-2  ?x ?y))
  (:IN (1-INHIBITS-3-the-primary-metabolic-enzyme-of-2
                                                ?y ?z ?enz)))
 (rassert!
  (active-ingredient-1-is-ancestor-to-2-and-2-interacts-with-3
                                                    ?x ?y ?z)
  (nil
   (1-is-an-ACTIVE-INGREDIENT ?x)
   (1-is-ANCESTOR-OF-2  ?x ?y)
   (1-INHIBITS-3-the-primary-metabolic-enzyme-of-2
                                           ?y ?z ?enz)
   )))


(rule
 ((:IN (1-is-an-ACTIVE-INGREDIENT ?x))
  (:IN (1-is-ANCESTOR-OF-2  ?x ?y))
  (:IN (1-inhibits-metabolic-clearance-of-2-via-3
                                         ?y ?z ?enz)))
 (rassert!
  (active-ingredient-1-is-ancestor-to-2-and-2-effects-an-interaction-with-3
                                           ?x ?y ?z)
  (nil
   (1-is-an-ACTIVE-INGREDIENT ?x)
   (1-is-ANCESTOR-OF-2  ?x ?y)
   (1-inhibits-metabolic-clearance-of-2-via-3
                                        ?y ?z ?enz)
   )))
```

A set of rules to for the disjunctive case where an active ingredient is ancestor to a compound that is the victim of an interaction with another active ingredient or metabolite.

```
(rule
 ((:IN (1-is-an-ACTIVE-INGREDIENT ?x))
  (:IN (1-is-ANCESTOR-OF-2  ?x ?z))
  (:IN (1-inhibits-3-the-primary-metabolic-enzyme-of-2
                                                ?y ?z ?enz)))
 (rassert!
  (active-ingredient-1-is-ancestor-to-2-and-2-is-affected-by-3
                                                    ?x ?z ?y)
  (nil
   (1-is-an-ACTIVE-INGREDIENT ?x)
   (1-is-ANCESTOR-OF-2  ?x ?z)
   (1-inhibits-3-the-primary-metabolic-enzyme-of-2 ?y ?z ?enz)
   )))
```

```
(rule
 ((:IN (1-is-an-ACTIVE-INGREDIENT ?x))
  (:IN (1-is-ANCESTOR-OF-2  ?x ?z))
  (:IN (1-inhibits-metabolic-clearance-of-2-via-3 ?y ?z ?enz)))
 (rassert!
  (ACTIVE-INGREDIENT-1-is-ancestor-to-2-and-2-is-affected-by-3
                                                    ?x ?z ?y)
  (nil
   (1-is-an-ACTIVE-INGREDIENT ?x)
   (1-is-ANCESTOR-OF-2  ?x ?z)
   (1-inhibits-metabolic-clearance-of-2-via-3 ?y ?z ?enz)
   )))
```

Appendix C

DEFINITIONS FOR EACH ASSERTION TYPE USED IN THE DIKB'S RULE-BASE

C.1 The `primary-total-clearance-mechanism` *Assertion*

The "primary total clearance mechanism" of some active pharmaceutic ingredient or metabolite, `X`, is the pharmacokinetic process that accounts for more than 50% of `X`'s clearance from the body. The DIKB's structured vocabulary lists four possible clearance processes:

1. `Biliary_Excretion` - Excretion of unchanged active pharmaceutic ingredient or metabolite, be it a `complex`, `protein`, or `small molecule`, via the bile and feces

2. `Exhalation_Excretion` - Excretion of unchanged active ingredient or metabolite, be it a `complex`, `protein`, or `small molecule`, via the lungs

3. `Renal_Excretion` - Excretion of unchanged active pharmaceutic ingredient or metabolite, be it a `complex`, `protein`, or `small molecule`, via the kidneys

4. `Metabolic_Clearance` - Elimination from the body of an active ingredient or metabolite, be it a `complex`, `protein`, or `small molecule`, by transformation through the biochemical reactions and pathways to substances that are inactive and/or excreted by the body

C.2 The `bioavailability` *Assertion*

This assertion specifies the proportion of an active pharmaceutical ingredient's dose that reaches systemic circulation. This assertion does not apply to drug metabolites. When the DIKB's `evidence-model` (Chapter 2, Section 2.3.1.1) exports this assertion it takes the maximum bioavailability entry found in all of the evidence items in the `evidence-for` list belonging to a given `Assertion` instance.

This value is mapped to the following discrete categories:

- `LOW`: [0.0, .20]
- `MEDIUM`: (.201, .50]
- `HIGH`: (.501, 1]

The motivation for choosing these categories is based on simple conjectures about what the maximum increase in AUC can be at various bioavailability levels. For example, a drug with a bioavailability of 50% should only be able to experience an approximate 2-fold increase in AUC if whatever is blocking the drug from entering systemic is completely removed. The maximum possible magnitude increase at the 20% level is approximately 5-fold while there is no limit for drugs with bioavailability values near zero.

Like the `maximum_concentration` assertion, `bioavailability` depends on statistical inference rather than logical induction and all drugs have some bioavailability value. Therefore, no `evidence-against` items need be collected. When different formulations of a drug have different bioavailability values (e.g. extended vs normal release) each assertion instance must refer to the dose and formulation of the pharmaceutical preparation that is associated with the bio-availability value.

C.3 The `first-pass-effect` *Assertion*

The `first-pass-effect` assertion is a qualitative statement of the degree to which an active pharmaceutic ingredient is cleared from the body before entering systemic circulation. At the time of this writing the focus is on the degree of first-pass metabolism an active pharmaceutic ingredient undergoes in the liver and gut wall before a drug reaches systemic circulation. As more becomes known about transporter proteins (e.g. P-glycoprotein) separate rules might be created to model effects on modulation of their activity. This assertion does not apply to metabolites.

This value is mapped to the following discrete categories:

- `LOW`: $[0.0, .50]$
- `MEDIUM`: $(.501, .80]$
- `HIGH`: $(.801, 1]$

The motivation for choosing these categories is based on simple conjectures about what the maximum increase in AUC can be at various first-pass-effect levels. For example, an

active pharmaceutic ingredient with a first-pass effect of 50% should only be able to experience an approximate 2-fold increase in AUC if the first-pass effect is completely removed. The maximum possible magnitude increase at the 80% level is approximately 5-fold while there is no limit for drugs with first-pass effect values near 100%.

Establishment: There are two ways that to derive a value for this assertion:

1. The value might be found in the results of a mass-balance study

2. If quantitative values are known for both the bioavailability, F, of and percent of active pharmaceutic ingredient absorbed, f_{abs}, then first pass effect can be calculated as:

$$1 - \frac{F}{f_{abs}} \tag{C.1}$$

This is a quantitative assertion that requires statistical inference. Some drugs or drug metabolits might have no first-pass effect (e.g. pharmaceutical entities with no clearance by metabolism) so it is logical to seek evidence against this assertion as well as supporting evidence.

C.4 The `fraction-absorbed` *Assertion*

This assertion is a quantitative statement of the fraction of an active ingredient's dose that gets absorbed in the gastro-intestinal tract. Such an estimate might be obtained from a study focusing on gut wall absorption. The quantitative values are maintained by the system but, they are also mapped to the following qualitative levels:

- `LOW`: $[0.0, .50]$
- `HIGH`: $(.501, 1]$

In many cases there will be no quantitative data for either the fraction of active pharmaceutic ingredient that is absorbed, its bioavailability, or both. However, one can often find, or derive, a reasonable qualitative estimate that falls within the range of either of these levels. This assertion does not apply to metabolites.

This is a quantitative assertion that requires statistical inference. Some drugs might not be absorbed in the GI tract (e.g. drugs for which there are only IV formulations) so it is logical to seek evidence against this assertion as well as supporting evidence.

C.5 The `maximum-concentration` *Assertion*

This assertion specifies the maximum concentration (C_{max}), in grams/liter, that the an active pharmaceutic ingredient or metabolite is known to reach. For active pharmaceutical ingredients, it is linked to the particular dose, in grams, of active pharmaceutic ingredients that was given in the study. For drug metabolites, it is linked to the particular dose in grams of the metabolite's ancestor active pharmaceutic ingredient that was given in the study. When the DIKB's `evidence-model` (Chapter 2, Section 2.3.1.1) exports this assertion it takes the maximum C_{max} value entry found in all of the evidence items in the `evidence-for` list belonging to a given `Assertion` instance.

When different formulations of a drug have different maximum concentration values (e.g. extended vs normal release) each assertion instance must refer to the dose and formulation of the drug that is associated with the value being entered.

This assertion, depends on statistical inference rather than logical induction and all pharmaceutical entities will have some C_{max} value. Therefore, no `evidence-against` items need be collected.

C.6 The `inhibits` *Assertion*

An active pharmaceutic ingredient or metabolite, X, is said to be a `inhibit` some enzyme, E, if X effects a measurable reduction in the catalytic function of E *in vivo*.

C.7 The `does-not-inhibit` *Assertion*

If an active pharmaceutic ingredient or metabolite, X *does not* effect a measurable reduction in the catalytic function of some enzyme E *in vivo* then the `X does-not-inhibit E` assertion applies.

C.8 The `in-vitro-selective-inhibitor-of-enzyme` *Assertion*

The FDA has provided a list of preferred and acceptable inhibitors for *in vitro* studies in [26], Appendix C-1, Table 2, and the CDER Web page on drug interactions [40]. In the DIKB, these chemicals are assumed to be *in vitro* selective inhibitors of they enzymes that they are listed with in these sources.

C.9 The `in-viVo-selective-inhibitor-of-enzyme` *Assertion*

The FDA has provided a list of preferred and acceptable inhibitors for *in vivo* studies in [26], Appendix A, Table 2, and the CDER Web page on drug interactions [40]. In the DIKB, these chemicals are assumed to be *in vivo* selective inhibitors of they enzymes that they are listed with in these sources.

Applying this assertion to some metabolite or active pharmaceutic ingredient, X, and enzyme, ENZ, implies that X `inhibits` ENZ.

C.10 The `substrate-of` *Assertion*

An active pharmaceutic ingredient or metabolite, X, is said to be a `substrate-of` some enzyme, E, if E catalyzes the transformation of the X to a metabolite, M. This assertion does not imply any quantitative information such as contribution E makes relative to other enzymes that catalyze the same reaction.

C.11 The `in-vitro-probe-substrate-of-enzyme` *Assertion*

The FDA has provided a list of preferred and acceptable chemical substrates for *in vitro* studies in [26], Table 3, and the CDER Web page on drug interactions [40]. In the DIKB, the principle chemicals involved in these reactions are *in vitro* probe substrates of the enzymes they are listed.

C.12 The `is-not-substrate-of` *Assertion*

Let X be an active pharmaceutic ingredient or metabolite and E some enzyme. If E does not catalyze the transformation of X to any known metabolite of X then the assertion `X`

`is-not-substrate-of` E applies.

C.13 The `primary-total-clearance-enzyme` *Assertion*

The "primary total clearance enzyme" of some active pharmaceutic ingredient or metabolite, `X`, is the enzyme, `ENZ`, responsible for 50% or more of the active pharmaceutic ingredient or metabolite's total clearance from the body. In other words, if at least 50% of `X` is cleared from the body by metabolic reactions catalyzed by `ENZ` then `ENZ` is the "primary total clearance enzyme" of `X`. This assertion can be established by any of the following methods (see Chapter 3, Section 3.2.3.5 for further explanation):

1. `ENZ` is polymorphic and a well-designed *in vivo* polymorphic pharmacokinetic study shows that `ENZ` is responsible for 50% or more of `X`'s clearance

2. a well-designed clinical trial investigating the pharmacokinetics of drug `X` in the presence of drug `Y` shows an increase in the AUC of `X` of at least 2-fold. NOTE: 1) drug `Y` must have no measurable effect on `X`'s clearance by renal clearance, biliary clearance, or exhalation, and 2) drug `Y` must be a *selective* inhibitor of `ENZ`

Applying this assertion to some metabolite or active pharmaceutic ingredient, X, and enzyme, ENZ, implies that:

- X is a `substrate-of` ENZ

- the `primary-total-clearance-mechanism` of X is metabolism

The current DIKB policy is that any enzyme that the FDA considers a drug or drug metabolite to be a probesubstrate of *in vivo* should be labeled its *primary total clearance enzyme*. The FDA has provided a list of preferred and acceptable probe substrates for *in vivo* studies in [26], Appendix A, Table 2, and the CDER Web page on drug interactions [40]. In the DIKB, these chemicals are assumed to be *in vivo* probe substrates of the enzymes that they are listed with in these sources.

C.14 The `primary-metabolic-clearance-enzyme` *Assertion*

The "primary metabolic clearance enzyme" of some active pharmaceutic ingredient or metabolite, `X`, is the enzyme, `ENZ`, responsible for 50% of the active pharmaceutic ingredient or metabolite's total *metabolic* clearance from the body.

C.15 The `inhibition-constant` *Assertion*

Some *in vitro* inhibition studies provide an inhibition constant, K_i, or a value that can converted to one. This assertion is the continuous value derived from such studies. When the DIKB's `evidence-model` (Chapter 2, Section 2.3.1.1) exports this assertion it takes the minimum all K_i values in the `evidence-for` list belonging to a given `Assertion` instance. When the DIKB's prediction rules are ran, this assertion is combined with the `maximum_concentration` assertion for the (see Section C.5), C_{max}, and the `permanently_deactivates_cat` assertion (Section C.22) to derive an estimate of the clinical relevance of the observed *in vitro* inhibition.

The DIKB labels a drug or drug metabolite an *in vivo* inhibitor for some drug metabolizing enzyme at the concentrations it is expected to reach during drug therapy if the following relationship holds:

$$\frac{C_{max}}{K_i} > 0.1 \tag{C.2}$$

Where C_{max} is the maximum observed concentration the inhibitor has reached in patients at normal, therapeutic, doses and K_i is an inhibition constant derived from a well-designed *in vitro* enzyme inhibition experiment involving the inhibitor. This relationship applies to inhibition of members of the Cytochrome P-450 enzyme family and is not applicable if the inhibitor is thought to permanently remove the affected enzyme from further participation in catalysis by any means. The basis for this relationship can be found in a recent FDA guidance that recommends that a clinically relevant effect from competitive enzyme inhibition be considered possible if the following relationship holds (see [26], p.33):

$$\frac{[I]}{K_i} > 0.1 \tag{C.3}$$

Where $[I]$ is the estimated concentration of the inhibitor at the enzyme binding site.

It is important to note that K_i values can vary depending on the system of enzymes used in each study. In fact, there can be a greater than 10-fold difference between the K_i found in recombinant enzyme systems compared to the K_i derived from human liver microsomes. Thus, the DIKB requires that the enzyme system used in the study from which a K_i is taken be noted in case there will be a need to distinguish K_i value by the enzyme system from which they were derived.

Like the `maximum_concentration` assertion, `inhibition_constant` depends on statistical inference rather than logical induction. Unlike the `maximum_concentration`, the value does not exist for some pharmaceutical entities. Therefore, it is logical to collect `evidence-against` items.

C.16 The `has-metabolite` *Assertion*

If an active pharmaceutic ingredient or metabolite, X, can be chemically altered to produce another compound, M, via a single chemical reaction possibly involving some enzyme, E, then, metabolite M is considered a metabolite of X and the assertion `(X has-metabolite M)` is applicable.

C.17 The `controls-formation-of` *Assertion*

If an active pharmaceutic ingredient or metabolite, X, can be chemically altered to produce another compound, M, via a single chemical reaction that requires catalysis by some enzyme, E then, E controls the formation of M and the assertion `(E controls-formation-of M)` is applicable.

C.18 The `polymorphic-enzyme` *Assertion*

A `polymorphic-enzyme` enzyme is an enzyme that has multiple drug-catalysis phenotypes due to genetic polymorphisms. By default, the DIKB assumes all enzymes to be *non-*

polymorphic.

C.19 The `pceut-entity-of-concern` *Assertion*

A "pceut-entity-of-concern" is an active pharmaceutic ingredient or metabolite for which even a small change in the system concentration would be of concern to a clinician. We assume that the criteria for a drug to meet this definition will vary for valid reasons between different groups of experts but use the following criteria in the current DIKB:

- active pharmaceutic ingredient or metabolites for which therapeutic drug monitoring is required

- active pharmaceutic ingredient or metabolites for which the ratio between the toxic systemic concentration of the agent and the concentration at which the agent is therapeutic is less than or equal to 2.0.

C.20 The `sole-PK-effect-alter-metabolic-clearance` *Assertion*

This assertion is a required assumption of evidence from a clinical pharmacokinetic DDI study involving a non-polymorphic enzyme when a curator applies the study as support for the `primary-total-clearance-enzyme` assertion (see Section C.13). It asserts that the sole pharmacokinetic effect of an active pharmaceutic ingredient or metabolite, `Y`, on an active pharmaceutic ingredient or metabolite, `X`, is alteration of `X`'s metabolic clearance. In other words, it asserts that `Y` has no measurable effect on `X`'s clearance by renal, biliary, exhalation, or efflux transport processes.

C.21 The `permanently_deactivates_catalytic_function` *Assertion*

This assertion specifies that an active pharmaceutical ingredient or metabolite is known to affect an enzyme in such a way that the enzyme is permanently removed from further participation in catalysis. For example, this assertion is applicable for the *slowly reversible* and *irreversible* inhibition mechanisms mentioned in Chapter 3, Section 3 and discussed in detail in Levy *et al* [112]. This assertion is also applicable if there is any other mechanism

by which the active pharmaceutic ingredient or metabolite could permanently remove the enzyme from further participation in catalysis.

When the DIKB's *evidence-model* (Chapter 2, Section 2.3.1.1) asserts that some active pharmaceutical ingredient or metabolite, `X`, is an inhibitor of some enzyme, `Y`, and the `permanently_deactivate_catalytic_function` assertion contains no value, the system will assert the `does_not_permanently_deactivate_catalytic_function` assertion by default reasoning.

C.22 The `does_not_permanently_deactivate_catalytic_function` *Assertion*

This is the inverse of the `permanently_deactivates_catalytic_function` assertion (Section C.21). When the DIKB's *evidence-model* (Chapter 2, Section 2.3.1.1) asserts that some active pharmaceutical ingredient or metabolite, `X`, is an inhibitor of some enzyme, `Y`, and the `permanently_deactivates_catalytic_function` assertion contains no value, the system will assert the `does_not_permanently_deactivate_catalytic_function` assertion by default reasoning.

Appendix D

THE DIKB EVIDENCE TAXONOMY

DIKB curators categorize each evidence item into one of the evidence-types from the evidence-type taxonomy shown here. The evidence types in the taxonomy are arranged into parent and child classes of evidence. A child class inherits all of the properties of the parent class and adds some specific properties of its own. The taxonomy is shown here with child evidence-types at a deeper indent-level than its parent class.

Evidence Types

`[Statement]` *A statement:* A published artifact that is "...the basis for belief or disbelief; knowledge on which to base belief" see the term "evidence" in Wordnet version 3.0 [121]

`[Non_Traceable_Statement]` *A non-traceable, but possibly authoritative, statement:* A statement that does not explicitly refer to evidence items in justification of its assertion(s) or that refers to an evidence item that is not accessible to the curator (e.g. pre-market drug studies only accessible to drug-company or FDA researchers)

`[Non_traceable_Drug_Label_Statement]` *A non-traceable drug-label statement:* An assertion found in a drug label that does not provide any traceable citations for its evidence support

`[Traceable_Statement]` *A traceable statement:* A statement that provides citation to evidence support for justification of its assertion(s)

`[Traceable_Drug_Label_Statement]` *A traceable drug-label statement:* An assertion stated in a drug label that provides citations for its evidence support

continued on next page

continued from previous page

Evidence Types

[EV_EX_Met_Enz_ID] *A drug metabolism identification experiment:* An experiment conducted with biological tissues and/or chemical compounds in a laboratory designed to identify the specific enzymes responsible for the metabolism of a drug ([26], p. 25)

[EV_EX_Met_Enz_ID_Cyp450] *A CYP450 drug metabolism identification experiment:* A metabolic **enzyme** identification experiment specifically designed to identify the Cytochrome P-450 enzymes involved in the metabolism of a drug

[EV_EX_Met_Enz_ID_Cyp450_Hum_Recom] *A CYP450, recombinant, drug metabolism identification experiment with possibly NO probe enzyme inhibitor(s)*

[EV_EX_Met_Enz_ID_Cyp450_Hum_Recom_Chem] *A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors*

[EV_EX_Met_Enz_ID_Cyp450_Hum_Recom_Antibody] *A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors*

[EV_EX_Met_Enz_ID_Cyp450_Hum_Microsome] *A CYP450, human microsome, drug metabolism identification experiment:* A Cytochrome P-450 metabolic enzyme identification experiment using human liver microsomes that have been characterized for Cytochrome P-450 activity and possibly NO probe enzyme inhibitor(s)

[EV_EX_Met_Enz_ID_Cyp450_Hum_Microsome_Chem] *A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors*

[EV_EX_Met_Enz_ID_Cyp450_Hum_Microsome_Antibody] *A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors:*

continued on next page

continued from previous page

Evidence Types

[EV_EX_Met_Enz_Inhibit] *A metabolic enzyme inhibition experiment:* An experiment conducted with biological tissues and/or chemical compounds in a laboratory designed to determine whether or not a drug inhibits a specific drug-metabolizing enzyme

[EV_EX_Met_Enz_Inhibit_Cyp450] *A CYP450 metabolic enzyme inhibition experiment:* A metabolic **inhibition** experiment specifically designed to determine whether or not a drug inhibits a specific CYP450 enzyme

[EV_EX_Met_Enz_Inhibit_Cyp450_Hum_Recom] *A CYP450, recombinant, metabolic enzyme inhibition experiment:* A Cytochrome P-450 inhibition experiment using recombinant human enzymes

[EV_EX_Met_Enz_Inhibit_Cyp450_Hum_Microsome] *A CYP450, human microsome, metabolic enzyme inhibition experiment:* A Cytochrome P-450 metabolic enzyme inhibition experiment using human liver microsomes that have been characterized for Cytochrome P-450 activity

[EV_Observation] *An observation-based report:* An observation-based report of some occurrence

[EV_Obs_ADE] *An observation-based ADE report:* An observation-based report of an adverse drug event

[EV_Obs_ADE_Public_Reported] *An observation-based ADE report in a public reporting database:* An adverse event report on file in a public adverse event reporting database such as the FDA's Adverse Event Reporting System

[EV_Obs_DI_CR] *A published observation-based ADE report:* An published observation-based case-report of a drug interaction

[EV_Obs_DI_CR_Evaluated] *A published and evaluated observation-based ADE report:* An observation-based report of a drug interaction that has been evaluated by some assessment tool

continued on next page

continued from previous page

Evidence Types

[EV_Clinical_Trial] *A clinical trial:* "a pre-planned clinical study of the safety, efficacy, or optimum dosage schedule of one or more diagnostic, therapeutic, or prophylactic drugs, devices, or techniques in humans selected according to predetermined criteria of eligibility and observed for predefined evidence of favorable and unfavorable effects." - (Medical Subject Headings (MeSH) [43] version 2008, concept code D016430, `Clinical Trial`)

[EV_CT_DDI] *A DDI clinical trial:* A study designed to quantify the pharmacokinetic and/or pharmacodynamic effects within study participants of a single drug in the presence of a purported precipitant.

[EV_PK_DDI_NR] *A non-randomized DDI clinical trial:* A pharmacokinetic DDI study where participants receive a drug in the presence of a purported precipitant (experimental group) or not (control group) but participants are not randomly assigned to experiment and control groups. This can include fixed-order studies where all participants are tested with placebo and precipitant after some period of washout

[EV_PK_DDI_Par_Grps] *A parallel groups DDI clinical trial:* A pharmacokinetic DDI study involving two groups of non-randomized participants where both groups receive the purported object drug while only one group receives the purported precipitant

[EV_PK_DDI_RCT] *A randomized DDI clinical trial:* A randomized, controlled, pharmacokinetic DDI study where participants receive a drug either in the presence of a purported precipitant (experimental group) or not (control group)

[EV_CT_Pharmacokinetic] *A pharmacokinetic clinical trial:* "A study of the process by which a drug is absorbed, distributed, metabolized, and eliminated by the body." (NCI Thesaurus [54] version 8, concept code C49663, `Pharmacokinetic Study`)

[EV_CT_PK_Genotype] *A genotyped pharmacokinetic clinical trial:* A drug pharmacokinetics study whose population consists of at least two groups known to posses distinct forms of some drug-metabolizing enzyme

[EV_CT_PK_Phenotype] *A phenotyped pharmacokinetic clinical trial:* A drug pharmacokinetics study whose population consists of at least two groups known to posses distinct drug metabolizing phenotypes

continued on next page

continued from previous page

Evidence Types

`[EV_Retrospective]` *A retrospective study:* "Studies used to test etiologic hypotheses in which inferences about an exposure to putative causal factors are derived from data relating to characteristics of persons under study or to events or experiences in their past. The essential feature is that some of the persons under study have the disease or outcome of interest and their characteristics are compared with those of unaffected persons." (Medical Subject Headings (MeSH) [43] version 2008, concept code D012189, `Retrospective Studies`)

`[EV_PK_DDI_Retro]` *A retrospective DDI study:* A retrospective study looking at the change in patient exposure of a single drug in the presence of a purported precipitant using a retrospective set of clinical records

`[EV_Population_PK]` *A retrospective population PK study:* a "...study of the sources and correlates of variability in drug concentrations among individuals who are the target patient population receiving clinically relevant doses of a drug in question." ([23], p.1)

`[EV_Review]` *A review article:* A published analysis of the evidence supporting and/or refuting some topic

`[EV_Drug_Review]` *A drug review article:* A published analysis of research on the efficacy or safety of a drug, family of drugs, or drug therapy.

`[EV_DrugClinicalReview]` *An FDA clinical review:* An FDA-sponsored review of a drug's pre-market studies and adverse event reports.

Appendix E

INCLUSION CRITERIA AND REQUIRED ACTIONS FOR EVIDENCE TYPES IN THE DIKB

Inclusion criteria specify the necessary attributes that an instance of an evidence type in the DIKB evidence taxonomy (Appendix D) must meet for it to be used to support or refute a specific instance of an assertion type in the DIKB (Appendix C). This appendix lists all of the inclusion criteria along with specific actions that DIKB curators must take when linking evidence of a particular evidence type to an assertion instance. Unless stated otherwise, inclusion criteria and required actions apply to all sub-types of the evidence type that the criteria mentions. For example, the criteria and action that apply to the `EV_CT_DDI` evidence type (Section E.4) also apply to its sub-types `EV_PK_DDI_NR`, `EV_PK_DDI_Par_Grps`, and `EV_PK_DDI_RCT`.

E.1 Inclusion Criteria for Reviews (`EV_Review`) and Sub-classes

Though not encouraged, a statement in a published review (`EV_Review` and sub-classes) can be used as evidence for or against `inhibits`, `substrate-of`, `primary-clearance-enzyme`, `fraction-cleared-by`, `primary-clearance-mechanism`. The following inclusion criteria apply:

- the statement is non-ambiguous

- the review provides clearly cited references or is from an authoritative organization such as the Federal Drug Administration

- each cited reference meets the inclusion criteria for the evidence type it belongs to

E.2 Inclusion Criteria for Published Observation Reports (EV_Obs_DI_CR) and Sub-classes

Published observation reports that been evaluated by some assessment tool (EV_Obs_DI_CR_Evaluated) can be used as support that an interaction occurred between at least two of the active ingredients or metabolites mentioned in the report. The following inclusion criteria apply:

- the report contains sufficient pharmacokinetic data to establish that the reported interaction occured by pharmacokinetic mechanisms
- the report is not about an abnormal susceptibility to some active ingredient or metabolites peculiar to an individual, otherwise known as an idiosyncratic interaction
- the report contains enough information to apply the Drug Interaction P Scale (DIPS) [99] to evaluate the interaction claimed by the case report.
- the report receives a causation rating of at least "probable" according to the DIPS scale. This means that the interaction report establishes a probable level of causation for an interaction between the two drugs of interest in the report.

E.3 Inclusion Criteria for Pharmacokinetic Studies (EV_CT_Pharmacokinetic) and Sub-classes

Instances of the pharmacokinetic study evidence types (EV_CT_Pharmacokinetic and sub-classes) can be used as evidence for or against instances of the following assertion types:

> maximum-concentration / has-metabolite / primary-total-clearance-mechanism / bioavailability / first-pass-effect / fraction-absorbed / has-metabolite / substrate-of / primary-total-clearance-enzyme

Instance of the EV_CT_PK_Genotype or EV_CT_PK_Phenotype evidence types can support or refute polymorphic-enzyme assertions. The following inclusion criteria apply to the EV_CT_Pharmacokinetic evidence type and sub-classes:

- The route of administration must stated.
- Study participants must not be exclusively under the age of 21 or over the age of 65.
- The study's design (dosing, duration, population size, and procedure for drug administration) should be sufficient to allow accurate measurements pharmacokinetic parameters.

Required Action(s):

- If evidence item's evidence-type is one of `EV_CT_PK_Genotype` or `EV_CT_PK_Phenotype`, then the curator must link a `polymorphic-enzyme` assertion (Section C.18) as an assumption for the intended use of the evidence item.
- If an instance of the `EV_CT_PK_Genotype` or `EV_CT_PK_Phenotype` evidence types is being used to support or refute that an enzyme is polymorphic then the specific genotype of the enzyme must be noted in the description of evidence.

E.4 Inclusion Criteria for Pharmacokinetic DDI Studies (`EV_CT_DDI`) and Sub-classes

Pharmacokinetic drug-drug interaction (DDI) studies (`EV_CT_DDI` and sub-classes) can be used as evidence for or against `increases-auc`, `inhibits`, and `substrate-of` assertion instances. The following inclusion criteria apply:

- The route of administration must stated.
- If the study is to be used as evidence that the precipitant active ingredient or metabolite is, or is not, an **inhibitor** of an enzyme, `ENZ`, then `ENZ` must be the "primary total clearance enzyme" of the object active ingredient or metabolite used in the study. Section C.13 defines this concept.
- If the study is to be used as evidence that the object active ingredient or metabolite is, or is not, a **substrate** an enzyme, `ENZ`, then the precipitant must be an *in vivo selective* inhibitor of that `ENZ`. Section C.8 defines this concept.

- Study participants must not exclusively under the age of 21 or over the age of 65.
- The study's duration should be long enough for precipitant, and any of its known active metabolites, to effect enzyme pool.
- The study's design (dosing, duration, population size, and procedure for drug administration) should be sufficient to allow accurate measurements of AUC change.

Required Action(s):

- If the study is to be used as evidence that the an active ingredient or metabolite is, or is not, an **inhibitor** of an enzyme, `ENZ`, then the curator must link (as an assumption for the evidence item's usage) the assertion that `ENZ` is the `primary--total-clearance-enzyme` (Section C.13) of the study's object active ingredient or metabolite.
- If the study is to be used as evidence that the object active ingredient or metabolite is, or is not, a substrate an enzyme, `ENZ`, then the curator must link the following assertions as assumptions for the evidence item's usage:
  - the study's precipitant is an `in-viVo-selective-inhibitor-of-enzyme` of `'ENZ`. (Section C.9) and,
  - the `sole-PK-effect-alter-metabolic-clearance` assertion indicating that the sole pharmacokinetic effect of the precipitant on the object drug is alteration of its metabolic clearance

E.5 Inclusion Criteria for Non-traceable Statements in Drug Product Labeling (`Non_traceable_Drug_Label_Statement`) and Sub-classes

An assertional statement found in a drug label that *does not* provide any traceable citations for its evidence support (`Non_traceable_Drug_Label_Statement` and sub-classes) can be used as evidence for or against `inhibits`, `substrate-of`, `primary-clearance-enzyme`, `fraction-cleared-by`, `primary-clearance-mechanism`. The following inclusion criteria apply:

- The labeling statement must be the most currently available for the drug
- The date of the label must be noted
- the statement cannot be accepted as evidence if its supporting evidence is based solely on non-human studies

Required Action(s):

- non-traceable and *ambiguous* author statements (such as "drug x did not increase the AUC of drug y" with no dosing or duration information) should be labeled as such.
- non-traceable, but *non-ambiguous*, author statements (such as "drug x, given at dose A, did not increase AUC of drug y, given at dose B for duration T") should be labeled as such

E.6 Inclusion Criteria for Drug Enzyme Inhibition Experiments (`EV_EX_Met_Enz_Inhibit`*) and Sub-classes*

A metabolic enzyme inhibition experiment (`EV_EX_Met_Enz_Inhibit` and sub-classes) can be used to support or refute an `inhibition-constant` assertion for an active ingredient or metabolite and some enzyme. An `inhibition-constant` assertion must be relevant to the concentration of the inhibitor as found in clinical practice. The system will ensure that this criteria is met while applying its inference algorithm to assertions in its knowledge-base. It will compare values for the `maximum_concentration` of the active ingredient or metabolite (Section C.5) with its `inhibition-constant` values (Section C.15). Instances of the `EV_EX_Met_Enz_Inhibit` evidence type and its subtypes can also support or refute that an active ingredient or metabolite is known to affect an enzyme in such a way that the enzyme is permanantly removed from further participation in catalysis (see Section C.21). The following inclusion criteria apply for all acceptible applications of these evidence types:

- The source of the enzymes must be either from human hepatocytes or human recombinant enzymes.

- NADPH must be added to the enzyme systems as part of the experiment when appropriate. In cases where no explicit statement in the evidence item mentions the use of NADPH, the curator is free to exercise judgement as to whether NADPH was added since it is considered standard protocol for studies during or after the year 2000.

- To support an `inhibition-constant` assertion for some active ingredient or metabolite and an enzyme, the substrate used in the experiment must be a *in vitro probe substrate* of the enzyme. See Section C.11 for the definition of this concept.

- Only K_i values, not $IC50$ or "percent of enzyme inhibited" values, can support an `inhibition-constant` assertion for some active ingredient or metabolite and an enzyme. The source describing the experiment must provide an appropriately derived K_i value.

Required Actions(s):

- If the study is being used to support or refute that an active ingredient or metabolite inhibits an enzyme, then the curator must link (as an assumption for the evidence item's usage) the assertion that the the substrate is an `in-vitro-probe-substrate-of-enzyme` of the target enzyme of the study (see Section C.11).

E.7 Inclusion Criteria for Metabolic Enzyme Identification Experiments (`EV_EX_Met_Enz_ID`) and Sub-classes

A drug metabolism identification experiment (`EV_EX_Met_Enz_ID` and sub-classes) can be used to support or refute that an active ingredient or metabolite is a substrate of one or more enzymes (see Section C.12). The following inclusion criteria apply:

- The source of the enzymes must be either from human hepatocytes or human recombinant enzymes.

- NADPH must be added to the enzyme system(s) as part of the experiment. In cases where no explicit statement in the evidence item mentions the use of NADPH, the

curator is free to exercise judgement as to whether NADPH was added since it is considered standard protocol for studies conducted during or after the year 2000.

- Experiments that use antibody inhibitors cannot be applied as evidence for or against the affinity of the substrate of interest for the enzyme

- the inhibitor used to determine whether an active ingredient's or metabolite's metabolism is catalyzed by a specific enzyme must be an *in vitro selective* inhibitor of that `ENZ`. Section C.8 defines this concept.

Required Action(s):

- The curator must link an assertion that the inhibitor used in the experiment is an `in-vitro-selective-inhibitor-of-enzyme` (see Section C.8) as an assumption for the particular application of evidence.

Appendix F

ENTERING AND VIEWING EVIDENCE USING THE DIKB'S WEB INTERFACE

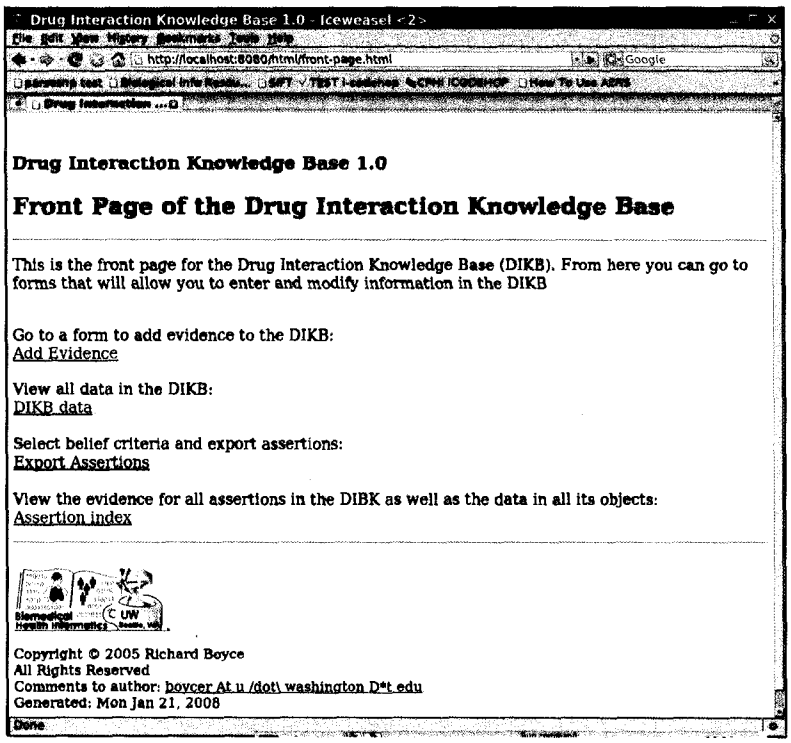

Figure F.1: This figure shows the welcome page of the DIKB's Web interface. DIKB curators and expert users have four options; clicking on the "Add Evidence" takes them to the page shown in Figure F.2 where they can begin the process of entering in evidence that supports or rebuts an assertion. Clicking on the "DIKB data" link takes them to a page (Figure F.10) that shows all assertions and evidence in the DIKB and allows them to change classification status of any assertion. Users can click on the "Export Assertions" link to start the process of defining levels-of-evidence and belief criteria. This option worked in an earlier version of the DIKB but is currently non-functional. The "Assertion Index" link loads a hyper-linked index of Web pages that summarizes the evidence for each assertion in the DIKB. Figure F.9(a) shows a sample of the index from the current DIKB while Figure F.9(b) shows of one of the assertion summary pages that the index links to.

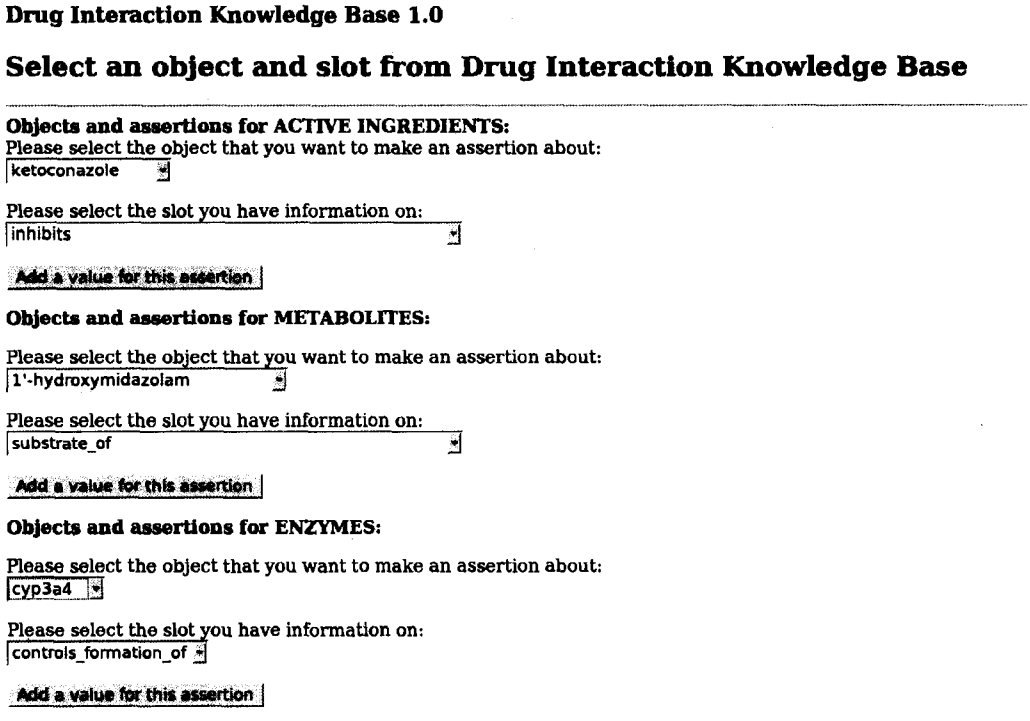

Figure F.2: DIKB curators begin the process of entering in evidence that supports or rebuts an assertion from this page by selecting the object (active ingredient, metabolite, or enzyme) and the assertion type (`inhibits`, `substrate-of`, `controls-formation-of` etc.) that the evidence item will be linked to. They make their selections using drop-down boxes then click the button labeled "Add a value for this assertion" which will take then the a page where they can select a value for the assertion (Figure F.3)

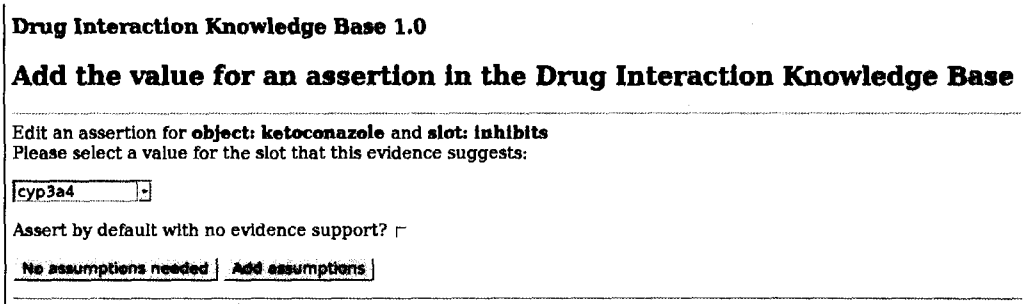

Figure F.3: Non-quantitative assertion types have a pre-specified range of values that can be chosen from the drop-down box on this page. Chapter 4, Section 4.3.1.1, discusses how the DIKB determines values for quantitative assertions. If the assertion is quantitative the only option in the drop-down box will be "continuous_value." Curators can declare that any assertion (quantitative or non-quantitative) should be considered a *default assumption* (Chapter 4, Section 4.2.5) by checking the box labeled "Assert by default with no evidence support." From this page, they can also begin the process of entering in any other assertions that should be linked as *evidence-use assumptions* (Chapter 4, Section 4.3.3) by clicking on the button labeled "Add assumptions." This action would take them to the page shown in Figure F.4. If there are no *evidence use assumptions* to add the curator clicks on the the button labeled "No assumptions needed" and proceeds to a page where they can enter more information on an evidence item (Figure F.6).

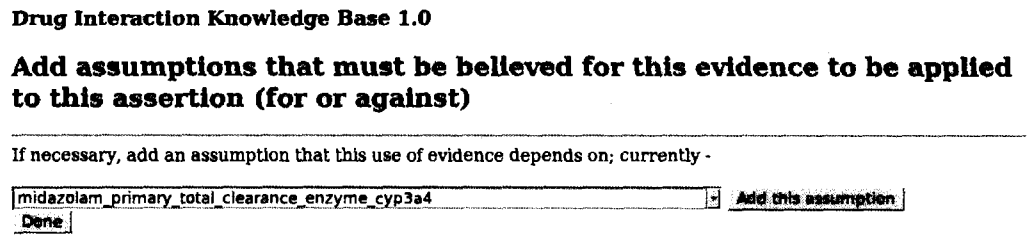

Figure F.4: Curators are taken to this page if they select "Add assumptions" from the page shown in Figure F.3. Here they can use a drop-down box to choose any assertion currently in the DIKB as a *evidence-use assumption* for the current application of an evidence item. Once they select an assertion to use as an *evidence-use assumption*, they can click the button labeled "Done" and they will proceed to a page where they can confirm their selection (Figure F.5). If they need to add more *evidence-use assumptions*, they click on the button labeled "Add this assumption" which will store their selections and re-load this page.

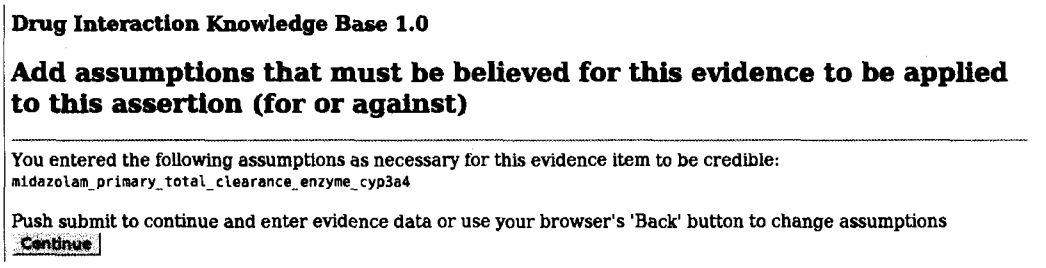

Figure F.5: Curators can confirm the assertions they want to link as *evidence-use assumptions* from this page or use the browser's "Back" button to make a change. Pressing the button labeled "Continue" will take them to a page where they can enter more information on an evidence item (Figure F.6).

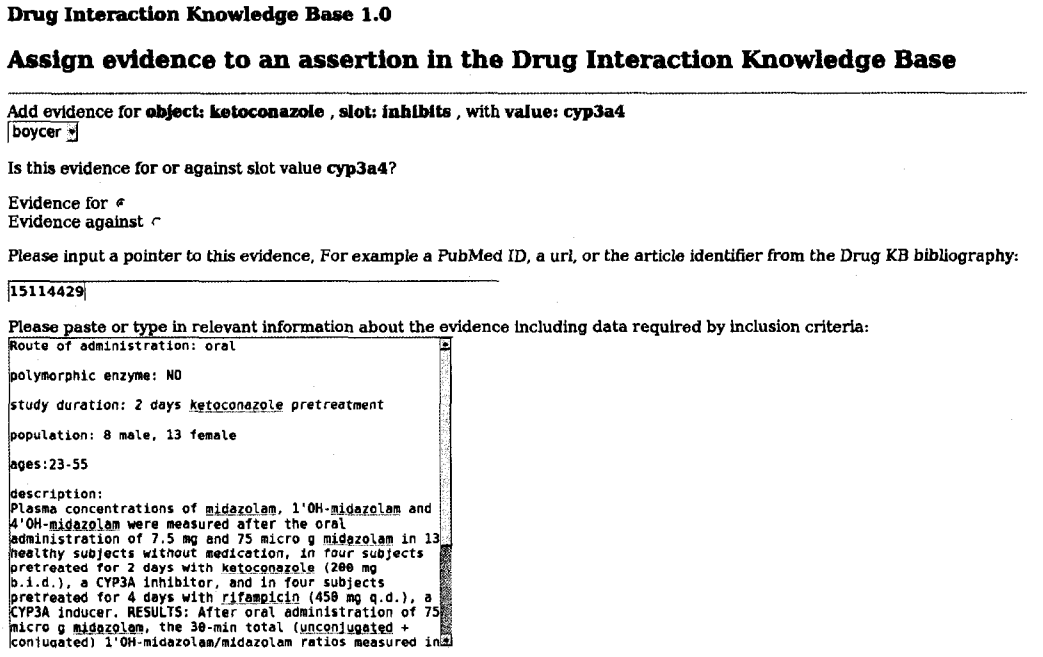

Figure F.6: Curators use the forms like the one provided on this page to enter more information on an evidence item. The Web interface provides a custom form for each assertion type. So, the form to link and evidence item to a quantitative assertion, such as a drug or drug metabolite's bioavailability, is different from the form used for an assertion about what enzyme a drug or drug metabolite is a substrate of. Curators fill out the form and then scroll the page down to select the evidence item's type (Figure F.7).

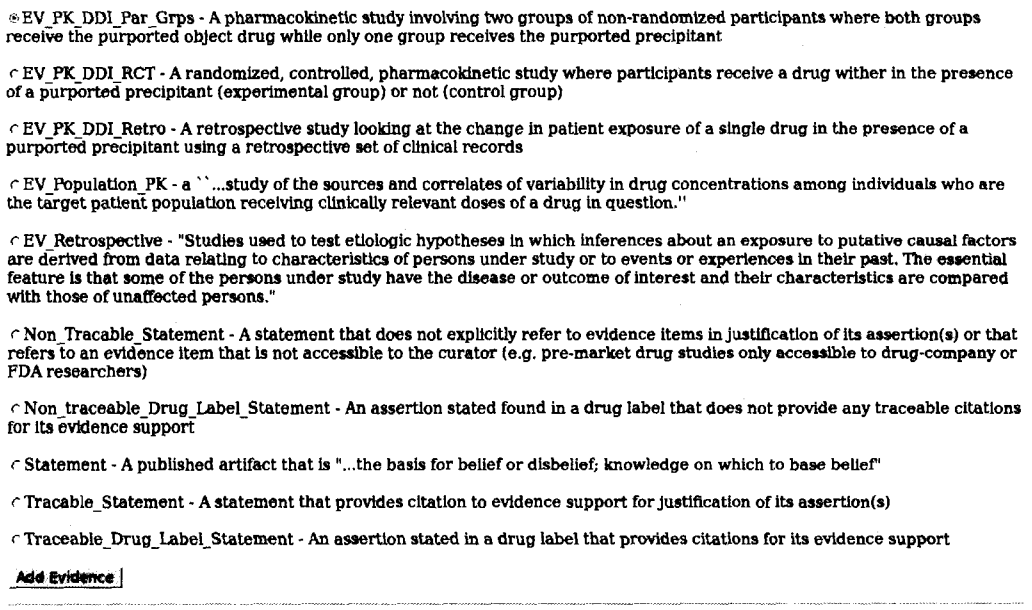

Figure F.7: This figure shows a portion of the list that a curator has to choose from when specifying an evidence item's type. A description of each evidence type is shown besides a radio button and its label in the DIKB evidence taxonomy (Appendix D). The curator clicks on a radio button to make a selection. They then click on the button labeled "Add Evidence" to proceed to a page where they can confirm all of the information they have entered for an evidence item (Figure F.8).

Drug Interaction Knowledge Base 1.0

Confirm and save new evidence to the Drug Interaction Knowledge Base

Please confirm that you want to save the following evidence:

object: ketoconazole
slot: inhibits
value: cyp3a4
assumption_picks: midazolam_primary_total_clearance_enzyme_cyp3a4
reviewer: boycer
position: for
pointer: 15114429
quote: Route of administration: oral polymorphic enzyme: NO study duration: 2 days ketoconazole pretreatment population: 8 male, 13 female ages:23-55 description: Plasma concentrations of midazolam, 1'OH-midazolam and 4'OH-midazolam were measured after the oral administration of 7.5 mg and 75 micro g midazolam in 13 healthy subjects without medication, in four subjects pretreated for 2 days with ketoconazole (200 mg b.i.d.), a CYP3A inhibitor, and in four subjects pretreated for 4 days with rifampicin (450 mg q.d.), a CYP3A inducer. RESULTS: After oral administration of 75 micro g midazolam, the 30-min total (unconjugated + conjugated) 1'OH-midazolam/midazolam ratios measured in the groups without co-medication, with ketoconazole and with rifampicin were (mean+/-SD): 6.23+/-2.61, 0.79+/-0.39 and 56.1+/-12.4, respectively. No side effects were reported by the subjects taking this low dose of midazolam. Good correlations were observed between the 30-min total 1'OH-midazolam/midazolam ratio and midazolam clearance in the group without co-medication (r(2)=0.64, P<0.001) and in the three groups taken together (r(2)=0.91, P<0.0001).
type: EV_PK_DDI_Par_Grps
has_evidence: True

Please confirm by reading through the following lists that 1) this will not be a duplicate use of this evidence and 2) that the entry of this evidence will not cause it to be linked to both an inhibits/substrate_of assertion AND an increase_auc assertion

evidence item '15114429' is linked to the following assertions as 'evidence_for':
ketoconazole_inhibits_cyp3a4

Cancel | Save

Figure F.8: The last step of the evidence entry process requires curators to confirm all of the information they have entered for an evidence item. They can use the browser's "Back" button to go back and enter new data at previous steps of the process. If the curator approves of their evidence entry, they can click on the "Save" button and the system will attempt to add the evidence item to the DIKB evidence-base. The system performs several validation tests on the data before the evidence item is entered into the DIKB's evidence-base and alerts the curator to warnings or errors. For example, in this figure, the system is informing the curator that this evidence item already exists in the system and is linked to the same assertion that the curator is trying to link it to now. Other alerts or warnings are produced if the evidence item has been rejected as evidence for or against other assertions in the DIKB (Chapter 4, Section 4.3.4.2) or, if the evidence item will form a circular evidence support pattern (Chapter 4, Section 4.3.4.3).

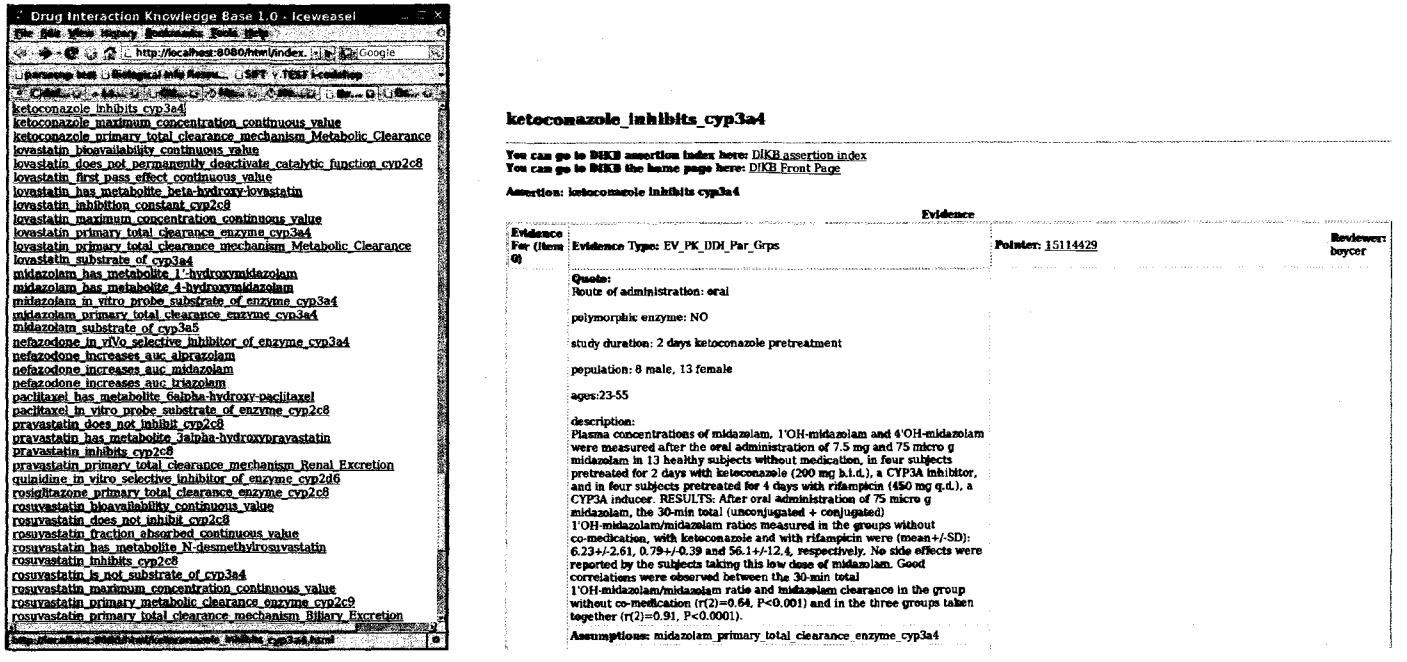

(a) A portion of the hyper-linked index of assertion Web pages from the current DIKB.

(b) A page showing a summary of all evidence linked to an assertion in the current DIKB.

Figure F.9: Curators and expert users can review the evidence linked to an assertion by clicking on a link to its information from a hyper-linked index of assertions.

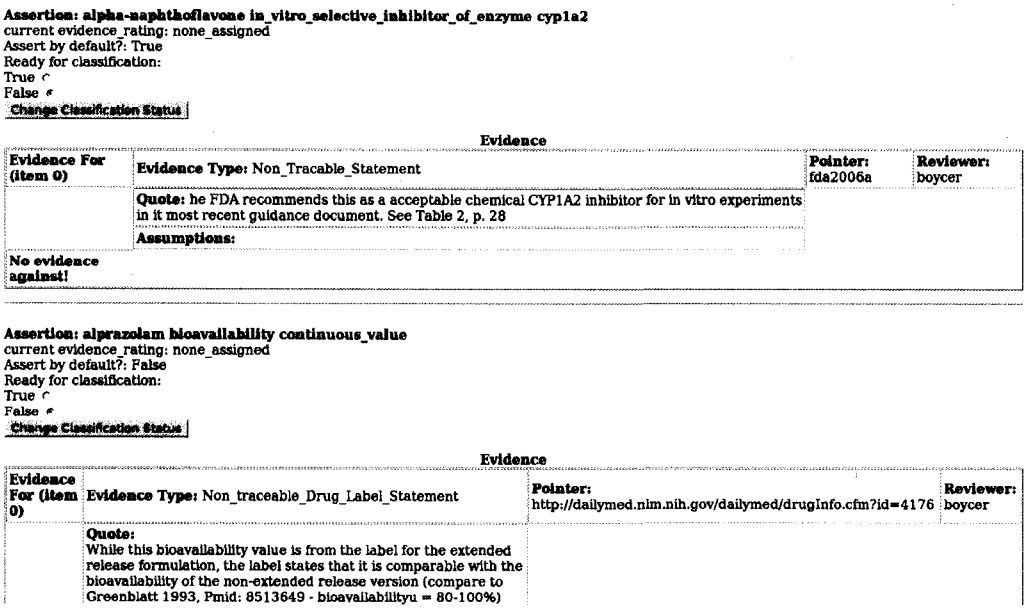

Assertion: alpha-naphthoflavone in_vitro_selective_inhibitor_of_enzyme cyp1a2
current evidence_rating: none_assigned
Assert by default?: True
Ready for classification:
True ○
False ●
Change Classification Status

Evidence

Evidence For (item 0)	**Evidence Type:** Non_Tracable_Statement	**Pointer:** fda2006a	**Reviewer:** boycer
	Quote: he FDA recommends this as a acceptable chemical CYP1A2 inhibitor for in vitro experiments in it most recent guidance document. See Table 2, p. 28		
	Assumptions:		
No evidence against!			

Assertion: alprazolam bioavailability continuous_value
current evidence_rating: none_assigned
Assert by default?: False
Ready for classification:
True ○
False ●
Change Classification Status

Evidence

Evidence For (item 0)	**Evidence Type:** Non_traceable_Drug_Label_Statement	**Pointer:** http://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?id=4176	**Reviewer:** boycer
	Quote: While this bioavailability value is from the label for the extended release formulation, the label states that it is comparable with the bioavailability of the non-extended release version (compare to Greenblatt 1993, Pmid: 8513649 - bioavailabilityu = 80-100%)		

Figure F.10: Clicking on the "DIKB data" link from the DIKB welcome page takes curators to a page that shows the evidence linked to all assertions in the DIKB and allows them to change classification status of any assertion. This figure shows a small portion of the page that is generated for the current DIKB.

Appendix G

THE FINAL VALIDATION SET OF DRUG-DRUG INTERACTIONS AND NON-INTERACTIONS

The reference set of drug-drug interactions and non-interactions used to characterize the prediction accuracy of the DIKB using a wide range of *belief criteria* including criteria chosen by the DIKB's evidence-board. An "X" in the column labeled *DDI* indicates that one of the pharmaceutical entities in the first column is the victim of a metabolic-inhibition interaction. An "X" in the *Non-DDI* column indicates that no metabolic-inhibition interaction is known to occur between the the pharmaceutical entities in the first column. Chapter 5, Section 5.1.1 explains how this validation set was created. The arrows indicate the drug or drug metabolite that the validation set considers the victim of a metabolic inhibition interaction that occurs between the pair. Arrows with a line through them indicate which drug or drug metabolite should not be affected by a metabolic inhibition interaction involving the other drug in the pair.

† The noted interaction occurs by inhibition of the metabolic clearance of a parent compound.

†† The DIKB's evidence-base uses this study to supports an drug mechanism assertion that is not related to the drug/drug or drug/drug-metabolite pair.

§ The pair was accidentally excluded from the experiment in Chapter 5 due to a trascription error.

‡ The pair was excluded because a validation set interaction or non-interaction between the two pharmaceutical entities was supported by a single clinical trial that was also present in DIKB assertions that the system could use to infer the interaction or non-interaction.

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
alprazolam - 1'-hydroxymidazolam			
alprazolam - 14-hydroxyclarithromycin			
alprazolam - 4-hydroxyalprazolam			
alprazolam - 4-hydroxymidazolam			
alprazolam - 4-hydroxytriazolam			
alprazolam - 6'-exomethylene-lovastatin			
alprazolam - 6'-exomethylene-simvastatin			
alprazolam - 6'-hydroxy-simvastatin			
alprazolam - 6'-hydroxymethyl-simvastatin			
alprazolam - 6'beta-hydroxy-lovastatin			
alprazolam - alpha-hydroxyalprazolam			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
alprazolam - atorvastatin			
alprazolam - beta-hydroxy-lovastatin			
alprazolam - beta-hydroxy-simvastatin			
alprazolam - clarithromycin			
alprazolam - desacetyldiltiazem			
alprazolam - erythromycin ←	X		[182]
alprazolam - fluconazole			
alprazolam - fluvastatin			
alprazolam - itraconazole ←	X		[181]
alprazolam - ketoconazole ←	X		[156], [74]
alprazolam - lovastatin			
alprazolam - midazolam			
alprazolam - N-demethyldesacetyl-diltiazem			
alprazolam - N-demethyldiltiazem			
alprazolam - N-desmethylrosuvastatin			
alprazolam - nefazodone ←	X		[75], [57]
alprazolam - ortho-hydroxy-atorvastatin			
alprazolam - para-hydroxy-atorvastatin			
alprazolam - pravastatin			
alprazolam - rosuvastatin			
alprazolam - simvastatin			
alprazolam - triazolam			
atorvastatin - 1'-hydroxymidazolam			
atorvastatin - 14-hydroxyclarithromycin			
atorvastatin - 4-hydroxyalprazolam			
atorvastatin - 4-hydroxymidazolam			
atorvastatin - 4-hydroxytriazolam			
atorvastatin - 6'-exomethylene-lovastatin			
atorvastatin - 6'-exomethylene-simvastatin			
atorvastatin - 6'-hydroxy-simvastatin			
atorvastatin - 6'-hydroxymethyl-simvastatin			
atorvastatin - 6'beta-hydroxy-lovastatin			
atorvastatin - alpha-hydroxyalprazolam			
atorvastatin - beta-hydroxy-lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
atorvastatin - beta-hydroxy-simvastatin			
atorvastatin - desacetyldiltiazem			
atorvastatin - erythromycin ←	X		[159]
atorvastatin - fluconazole			
atorvastatin - fluvastatin			
atorvastatin - lovastatin			
atorvastatin - N-demethyldesacetyl-diltiazem			
atorvastatin - N-demethyldiltiazem			
atorvastatin - N-desmethylrosuvastatin			
atorvastatin - nefazodone ←	X		[164]
atorvastatin - ortho-hydroxy-atorvastatin			
atorvastatin - para-hydroxy-atorvastatin			
atorvastatin - pravastatin			
atorvastatin - rosuvastatin			
atorvastatin - simvastatin			
clarithromycin - 1'-hydroxymidazolam			
clarithromycin - 14-hydroxyclarithromycin			
clarithromycin - 4-hydroxyalprazolam			
clarithromycin - 4-hydroxymidazolam			
clarithromycin - 4-hydroxytriazolam			
clarithromycin - 6'-exomethylene-lovastatin			
clarithromycin - 6'-exomethylene-simvastatin			
clarithromycin - 6'-hydroxy-simvastatin			
clarithromycin - 6'-hydroxymethyl-simvastatin			
clarithromycin - 6'beta-hydroxy-lovastatin			
clarithromycin - alpha-hydroxyalprazolam			
clarithromycin - atorvastatin →	X		[8], [95]
clarithromycin - beta-hydroxy-lovastatin			
clarithromycin - beta-hydroxy-simvastatin ‡			
clarithromycin - desacetyldiltiazem			
clarithromycin - erythromycin			
clarithromycin - fluconazole ←	X		[3]
clarithromycin - fluvastatin			
clarithromycin - lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
clarithromycin - N-demethyldesacetyl-diltiazem			
clarithromycin - N-demethyldiltiazem			
clarithromycin - N-desmethylrosuvastatin			
clarithromycin - nefazodone			
clarithromycin - ortho-hydroxy-atorvastatin			
clarithromycin - para-hydroxy-atorvastatin			
clarithromycin - pravastatin →	X		[95]
clarithromycin - rosuvastatin			
clarithromycin - simvastatin ‡			
diltiazem - 1'-hydroxymidazolam			
diltiazem - 14-hydroxyclarithromycin			
diltiazem - 4-hydroxyalprazolam			
diltiazem - 4-hydroxymidazolam			
diltiazem - 4-hydroxytriazolam			
diltiazem - 6'-exomethylene-lovastatin			
diltiazem - 6'-exomethylene-simvastatin			
diltiazem - 6'-hydroxy-simvastatin			
diltiazem - 6'-hydroxymethyl-simvastatin			
diltiazem - 6'beta-hydroxy-lovastatin			
diltiazem - alpha-hydroxyalprazolam			
diltiazem - alprazolam			
diltiazem - atorvastatin			
diltiazem - beta-hydroxy-lovastatin →	X		[29]
diltiazem - beta-hydroxy-simvastatin			
diltiazem - clarithromycin			
diltiazem - desacetyldiltiazem			
diltiazem - erythromycin			
diltiazem - fluconazole			
diltiazem - fluvastatin			
diltiazem - itraconazole			
diltiazem - ketoconazole			
diltiazem - lovastatin →	X		[32]
diltiazem - midazolam →	X		[30]
diltiazem - N-demethyldesacetyl-diltiazem			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
diltiazem - N-demethyldiltiazem			
diltiazem - N-desmethylrosuvastatin			
diltiazem - nefazodone			
diltiazem - ortho-hydroxy-atorvastatin			
diltiazem - para-hydroxy-atorvastatin			
diltiazem - pravastatin ↛		X	[32]
diltiazem - rosuvastatin			
diltiazem - simvastatin →	X		[125]
diltiazem - triazolam →	X		[171]
erythromycin - 1'-hydroxymidazolam			
erythromycin - 14-hydroxyclarithromycin			
erythromycin - 4-hydroxyalprazolam			
erythromycin - 4-hydroxymidazolam			
erythromycin - 4-hydroxytriazolam			
erythromycin - 6'-exomethylene-lovastatin			
erythromycin - 6'-exomethylene-simvastatin			
erythromycin - 6'-hydroxy-simvastatin			
erythromycin - 6'-hydroxymethyl-simvastatin			
erythromycin - 6'beta-hydroxy-lovastatin			
erythromycin - alpha-hydroxyalprazolam			
erythromycin - beta-hydroxy-lovastatin			
erythromycin - beta-hydroxy-simvastatin →	X		[103]
erythromycin - desacetyldiltiazem			
erythromycin - fluconazole			
erythromycin - fluvastatin			
erythromycin - lovastatin			
erythromycin - N-demethyldesacetyl-diltiazem			
erythromycin - N-demethyldiltiazem			
erythromycin - N-desmethylrosuvastatin			
erythromycin - nefazodone			
erythromycin - ortho-hydroxy-atorvastatin			
erythromycin - para-hydroxy-atorvastatin			
erythromycin - pravastatin			
erythromycin - rosuvastatin ↛		X	[47]

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
erythromycin - simvastatin →	X		[103]
fluconazole - 1'-hydroxymidazolam → †	X		[4]
fluconazole - 14-hydroxyclarithromycin ↛		X	[3]
fluconazole - 4-hydroxyalprazolam			
fluconazole - 4-hydroxymidazolam			
fluconazole - 4-hydroxytriazolam			
fluconazole - 6'-exomethylene-lovastatin			
fluconazole - 6'-exomethylene-simvastatin			
fluconazole - 6'-hydroxy-simvastatin			
fluconazole - 6'-hydroxymethyl-simvastatin			
fluconazole - 6'beta-hydroxy-lovastatin			
fluconazole - alpha-hydroxyalprazolam			
fluconazole - beta-hydroxy-lovastatin			
fluconazole - beta-hydroxy-simvastatin			
fluconazole - desacetyldiltiazem			
fluconazole - fluvastatin →	X		[102]
fluconazole - lovastatin			
fluconazole - N-demethyldesacetyl-diltiazem			
fluconazole - N-demethyldiltiazem			
fluconazole - N-desmethylrosuvastatin			
fluconazole - nefazodone			
fluconazole - ortho-hydroxy-atorvastatin			
fluconazole - para-hydroxy-atorvastatin			
fluconazole - pravastatin ↛		X	[102]
fluconazole - rosuvastatin ↛		X	[49]
fluconazole - simvastatin			
fluvastatin - 1'-hydroxymidazolam			
fluvastatin - 14-hydroxyclarithromycin			
fluvastatin - 4-hydroxyalprazolam			
fluvastatin - 4-hydroxymidazolam			
fluvastatin - 4-hydroxytriazolam			
fluvastatin - 6'-exomethylene-lovastatin			
fluvastatin - 6'-exomethylene-simvastatin			
fluvastatin - 6'-hydroxy-simvastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
fluvastatin - 6'-hydroxymethyl-simvastatin			
fluvastatin - 6'beta-hydroxy-lovastatin			
fluvastatin - alpha-hydroxyalprazolam			
fluvastatin - beta-hydroxy-lovastatin			
fluvastatin - beta-hydroxy-simvastatin			
fluvastatin - desacetyldiltiazem			
fluvastatin - N-demethyldesacetyl-diltiazem			
fluvastatin - N-demethyldiltiazem			
fluvastatin - N-desmethylrosuvastatin			
fluvastatin - ortho-hydroxy-atorvastatin			
fluvastatin - para-hydroxy-atorvastatin			
fluvastatin - rosuvastatin			
itraconazole - 1'-hydroxymidazolam			
itraconazole - 14-hydroxyclarithromycin			
itraconazole - 4-hydroxyalprazolam			
itraconazole - 4-hydroxymidazolam			
itraconazole - 4-hydroxytriazolam			
itraconazole - 6'-exomethylene-lovastatin			
itraconazole - 6'-exomethylene-simvastatin			
itraconazole - 6'-hydroxy-simvastatin			
itraconazole - 6'-hydroxymethyl-simvastatin			
itraconazole - 6'beta-hydroxy-lovastatin			
itraconazole - alpha-hydroxyalprazolam			
itraconazole - atorvastatin →	X		[117]
itraconazole - beta-hydroxy-lovastatin →	X		[108]
itraconazole - beta-hydroxy-simvastatin ‡			
itraconazole - clarithromycin			
itraconazole - desacetyldiltiazem			
itraconazole - erythromycin ←	X		[96]
itraconazole - fluvastatin ↛		X	[108]
itraconazole - ketoconazole			
itraconazole - lovastatin →	X		[108]
itraconazole - N-demethyldesacetyl-diltiazem			
itraconazole - N-demethyldiltiazem			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
itraconazole - N-desmethylrosuvastatin			
itraconazole - nefazodone			
itraconazole - para-hydroxy-atorvastatin §			
itraconazole - ortho-hydroxy-atorvastatin →	X		[117]
itraconazole - pravastatin →	X		[117]
itraconazole - rosuvastatin →	X		[50]
itraconazole - simvastatin ‡			
ketoconazole - 1'-hydroxymidazolam			
ketoconazole - 14-hydroxyclarithromycin			
ketoconazole - 4-hydroxyalprazolam			
ketoconazole - 4-hydroxymidazolam			
ketoconazole - 4-hydroxytriazolam			
ketoconazole - 6'-exomethylene-lovastatin			
ketoconazole - 6'-exomethylene-simvastatin			
ketoconazole - 6'-hydroxy-simvastatin			
ketoconazole - 6'-hydroxymethyl-simvastatin			
ketoconazole - 6'beta-hydroxy-lovastatin			
ketoconazole - alpha-hydroxyalprazolam			
ketoconazole - atorvastatin			
ketoconazole - beta-hydroxy-lovastatin			
ketoconazole - beta-hydroxy-simvastatin			
ketoconazole - clarithromycin			
ketoconazole - desacetyldiltiazem			
ketoconazole - erythromycin			
ketoconazole - fluconazole			
ketoconazole - fluvastatin			
ketoconazole - lovastatin			
ketoconazole - N-demethyldesacetyl-diltiazem			
ketoconazole - N-demethyldiltiazem			
ketoconazole - N-desmethylrosuvastatin			
ketoconazole - nefazodone			
ketoconazole - ortho-hydroxy-atorvastatin			
ketoconazole - para-hydroxy-atorvastatin			
ketoconazole - pravastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
ketoconazole - rosuvastatin			
ketoconazole - simvastatin →	X		[42]
lovastatin - 1'-hydroxymidazolam			
lovastatin - 14-hydroxyclarithromycin			
lovastatin - 4-hydroxyalprazolam			
lovastatin - 4-hydroxymidazolam			
lovastatin - 4-hydroxytriazolam			
lovastatin - 6'-exomethylene-lovastatin			
lovastatin - 6'-exomethylene-simvastatin			
lovastatin - 6'-hydroxy-simvastatin			
lovastatin - 6'-hydroxymethyl-simvastatin			
lovastatin - 6'beta-hydroxy-lovastatin			
lovastatin - alpha-hydroxyalprazolam			
lovastatin - beta-hydroxy-lovastatin			
lovastatin - beta-hydroxy-simvastatin			
lovastatin - desacetyldiltiazem			
lovastatin - fluvastatin			
lovastatin - N-demethyldesacetyl-diltiazem			
lovastatin - N-demethyldiltiazem			
lovastatin - N-desmethylrosuvastatin			
lovastatin - ortho-hydroxy-atorvastatin			
lovastatin - para-hydroxy-atorvastatin			
lovastatin - pravastatin			
lovastatin - rosuvastatin			
midazolam - 1'-hydroxymidazolam			
midazolam - 14-hydroxyclarithromycin			
midazolam - 4-hydroxyalprazolam			
midazolam - 4-hydroxymidazolam			
midazolam - 4-hydroxytriazolam			
midazolam - 6'-exomethylene-lovastatin			
midazolam - 6'-exomethylene-simvastatin			
midazolam - 6'-hydroxy-simvastatin			
midazolam - 6'-hydroxymethyl-simvastatin			
midazolam - 6'beta-hydroxy-lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
midazolam - alpha-hydroxyalprazolam			
midazolam - atorvastatin ‡			
midazolam - beta-hydroxy-lovastatin			
midazolam - beta-hydroxy-simvastatin ‡			
midazolam - clarithromycin ←	X		[78], [70]
midazolam - desacetyldiltiazem §			
midazolam - erythromycin ←	X		[135]
midazolam - fluconazole ←	X		[134], [4]
midazolam - fluvastatin			
midazolam - itraconazole ←	X		[136]
midazolam - ketoconazole ←	X		[136]
midazolam - lovastatin			
midazolam - N-demethyldesacetyl-diltiazem			
midazolam - N-demethyldiltiazem			
midazolam - N-desmethylrosuvastatin			
midazolam - nefazodone ← ††	X		[111]
midazolam - ortho-hydroxy-atorvastatin			
midazolam - para-hydroxy-atorvastatin			
midazolam - pravastatin			
midazolam - rosuvastatin			
midazolam - simvastatin ‡			
midazolam - triazolam			
nefazodone - 1'-hydroxymidazolam			
nefazodone - 14-hydroxyclarithromycin			
nefazodone - 4-hydroxyalprazolam → †	X		[75]
nefazodone - 4-hydroxymidazolam			
nefazodone - 4-hydroxytriazolam			
nefazodone - 6'-exomethylene-lovastatin			
nefazodone - 6'-exomethylene-simvastatin			
nefazodone - 6'-hydroxy-simvastatin			
nefazodone - 6'-hydroxymethyl-simvastatin			
nefazodone - 6'beta-hydroxy-lovastatin			
nefazodone - alpha-hydroxyalprazolam			
nefazodone - beta-hydroxy-lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
nefazodone - beta-hydroxy-simvastatin →	X		[164]
nefazodone - desacetyldiltiazem			
nefazodone - fluvastatin			
nefazodone - lovastatin			
nefazodone - N-demethyldesacetyl-diltiazem			
nefazodone - N-demethyldiltiazem			
nefazodone - N-desmethylrosuvastatin			
nefazodone - ortho-hydroxy-atorvastatin			
nefazodone - para-hydroxy-atorvastatin			
nefazodone - pravastatin ↛		X	[164]
nefazodone - rosuvastatin			
nefazodone - simvastatin →	X		[164]
pravastatin - 1'-hydroxymidazolam			
pravastatin - 14-hydroxyclarithromycin			
pravastatin - 4-hydroxyalprazolam			
pravastatin - 4-hydroxymidazolam			
pravastatin - 4-hydroxytriazolam			
pravastatin - 6'-exomethylene-lovastatin			
pravastatin - 6'-exomethylene-simvastatin			
pravastatin - 6'-hydroxy-simvastatin			
pravastatin - 6'-hydroxymethyl-simvastatin			
pravastatin - 6'beta-hydroxy-lovastatin			
pravastatin - alpha-hydroxyalprazolam			
pravastatin - beta-hydroxy-lovastatin			
pravastatin - beta-hydroxy-simvastatin			
pravastatin - desacetyldiltiazem			
pravastatin - fluvastatin			
pravastatin - N-demethyldesacetyl-diltiazem			
pravastatin - N-demethyldiltiazem			
pravastatin - N-desmethylrosuvastatin			
pravastatin - ortho-hydroxy-atorvastatin			
pravastatin - para-hydroxy-atorvastatin			
pravastatin - rosuvastatin			
rosuvastatin - 1'-hydroxymidazolam			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
rosuvastatin - 14-hydroxyclarithromycin			
rosuvastatin - 4-hydroxyalprazolam			
rosuvastatin - 4-hydroxymidazolam			
rosuvastatin - 4-hydroxytriazolam			
rosuvastatin - 6'-exomethylene-lovastatin			
rosuvastatin - 6'-exomethylene-simvastatin			
rosuvastatin - 6'-hydroxy-simvastatin			
rosuvastatin - 6'-hydroxymethyl-simvastatin			
rosuvastatin - 6'beta-hydroxy-lovastatin			
rosuvastatin - alpha-hydroxyalprazolam			
rosuvastatin - beta-hydroxy-lovastatin			
rosuvastatin - beta-hydroxy-simvastatin			
rosuvastatin - desacetyldiltiazem			
rosuvastatin - N-demethyldesacetyl-diltiazem			
rosuvastatin - N-demethyldiltiazem			
rosuvastatin - N-desmethylrosuvastatin			
rosuvastatin - ortho-hydroxy-atorvastatin			
rosuvastatin - para-hydroxy-atorvastatin			
simvastatin - 1'-hydroxymidazolam			
simvastatin - 14-hydroxyclarithromycin			
simvastatin - 4-hydroxyalprazolam			
simvastatin - 4-hydroxymidazolam			
simvastatin - 4-hydroxytriazolam			
simvastatin - 6'-exomethylene-lovastatin			
simvastatin - 6'-exomethylene-simvastatin			
simvastatin - 6'-hydroxy-simvastatin			
simvastatin - 6'-hydroxymethyl-simvastatin			
simvastatin - 6'beta-hydroxy-lovastatin			
simvastatin - alpha-hydroxyalprazolam			
simvastatin - beta-hydroxy-lovastatin			
simvastatin - beta-hydroxy-simvastatin			
simvastatin - desacetyldiltiazem			
simvastatin - fluvastatin			
simvastatin - lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
simvastatin - N-demethyldesacetyl-diltiazem			
simvastatin - N-demethyldiltiazem			
simvastatin - N-desmethylrosuvastatin			
simvastatin - ortho-hydroxy-atorvastatin			
simvastatin - para-hydroxy-atorvastatin			
simvastatin - pravastatin			
simvastatin - rosuvastatin			
triazolam - 1'-hydroxymidazolam			
triazolam - 14-hydroxyclarithromycin			
triazolam - 4-hydroxyalprazolam			
triazolam - 4-hydroxymidazolam			
triazolam - 4-hydroxytriazolam			
triazolam - 6'-exomethylene-lovastatin			
triazolam - 6'-exomethylene-simvastatin			
triazolam - 6'-hydroxy-simvastatin			
triazolam - 6'-hydroxymethyl-simvastatin			
triazolam - 6'beta-hydroxy-lovastatin			
triazolam - alpha-hydroxyalprazolam			
triazolam - atorvastatin			
triazolam - beta-hydroxy-lovastatin			
triazolam - beta-hydroxy-simvastatin			
triazolam - clarithromycin ←	X		[73]
triazolam - desacetyldiltiazem			
triazolam - erythromycin ←	X		[141]
triazolam - fluconazole ←	X		[172]
triazolam - fluvastatin			
triazolam - itraconazole ←	X		[170], [130]
triazolam - ketoconazole ←	X		[170], [174]
triazolam - lovastatin			
triazolam - N-demethyldesacetyl-diltiazem			
triazolam - N-demethyldiltiazem			
triazolam - N-desmethylrosuvastatin			
triazolam - nefazodone ←	X		[31]
triazolam - ortho-hydroxy-atorvastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
triazolam - para-hydroxy-atorvastatin			
triazolam - pravastatin			
triazolam - rosuvastatin			
triazolam - simvastatin			
beta-hydroxy-simvastatin - 6'-hydroxymethyl-simvastatin			
beta-hydroxy-simvastatin - N-demethyldiltiazem			
beta-hydroxy-simvastatin - 14-hydroxyclarithromycin			
beta-hydroxy-simvastatin - 6'beta-hydroxy-lovastatin			
beta-hydroxy-simvastatin - para-hydroxy-atorvastatin			
beta-hydroxy-simvastatin - desacetyldiltiazem			
beta-hydroxy-simvastatin - 6'-hydroxy-simvastatin			
beta-hydroxy-simvastatin - 4-hydroxytriazolam			
beta-hydroxy-simvastatin - ortho-hydroxy-atorvastatin			
beta-hydroxy-simvastatin - 4-hydroxymidazolam			
beta-hydroxy-simvastatin - 6'-exomethylene-simvastatin			
beta-hydroxy-simvastatin - 6'-exomethylene-lovastatin			
beta-hydroxy-simvastatin - N-demethyldesacetyl-diltiazem			
beta-hydroxy-simvastatin - alpha-hydroxyalprazolam			
beta-hydroxy-simvastatin - 4-hydroxyalprazolam			
beta-hydroxy-simvastatin - beta-hydroxy-lovastatin			
beta-hydroxy-simvastatin - N-desmethylrosuvastatin			
beta-hydroxy-simvastatin - 1'-hydroxymidazolam			
beta-hydroxy-lovastatin - 6'-hydroxymethyl-simvastatin			
beta-hydroxy-lovastatin - N-demethyldiltiazem			
beta-hydroxy-lovastatin - 14-hydroxyclarithromycin			
beta-hydroxy-lovastatin - 6'beta-hydroxy-lovastatin			
beta-hydroxy-lovastatin - para-hydroxy-atorvastatin			
beta-hydroxy-lovastatin - desacetyldiltiazem			
beta-hydroxy-lovastatin - 6'-hydroxy-simvastatin			
beta-hydroxy-lovastatin - 4-hydroxytriazolam			
beta-hydroxy-lovastatin - ortho-hydroxy-atorvastatin			
beta-hydroxy-lovastatin - 4-hydroxymidazolam			
beta-hydroxy-lovastatin - 6'-exomethylene-simvastatin			
beta-hydroxy-lovastatin - 6'-exomethylene-lovastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
beta-hydroxy-lovastatin - N-demethyldesacetyl-diltiazem			
beta-hydroxy-lovastatin - alpha-hydroxyalprazolam			
beta-hydroxy-lovastatin - 4-hydroxyalprazolam			
beta-hydroxy-lovastatin - N-desmethylrosuvastatin			
beta-hydroxy-lovastatin - 1'-hydroxymidazolam			
6'-hydroxy-simvastatin - 6'-hydroxymethyl-simvastatin			
6'-hydroxy-simvastatin - N-demethyldiltiazem			
6'-hydroxy-simvastatin - 14-hydroxyclarithromycin			
6'-hydroxy-simvastatin - 6'beta-hydroxy-lovastatin			
6'-hydroxy-simvastatin - para-hydroxy-atorvastatin			
6'-hydroxy-simvastatin - desacetyldiltiazem			
6'-hydroxy-simvastatin - 4-hydroxytriazolam			
6'-hydroxy-simvastatin - ortho-hydroxy-atorvastatin			
6'-hydroxy-simvastatin - 4-hydroxymidazolam			
6'-hydroxy-simvastatin - 6'-exomethylene-simvastatin			
6'-hydroxy-simvastatin - 6'-exomethylene-lovastatin			
6'-hydroxy-simvastatin - N-demethyldesacetyl-diltiazem			
6'-hydroxy-simvastatin - alpha-hydroxyalprazolam			
6'-hydroxy-simvastatin - 4-hydroxyalprazolam			
6'-hydroxy-simvastatin - N-desmethylrosuvastatin			
6'-hydroxy-simvastatin - 1'-hydroxymidazolam			
6'-hydroxymethyl-simvastatin - N-demethyldiltiazem			
6'-hydroxymethyl-simvastatin - 14-hydroxyclarithromycin			
6'-hydroxymethyl-simvastatin - 6'beta-hydroxy-lovastatin			
6'-hydroxymethyl-simvastatin - para-hydroxy-atorvastatin			
6'-hydroxymethyl-simvastatin - desacetyldiltiazem			
6'-hydroxymethyl-simvastatin - 4-hydroxytriazolam			
6'-hydroxymethyl-simvastatin - ortho-hydroxy-atorvastatin			
6'-hydroxymethyl-simvastatin - 4-hydroxymidazolam			
6'-hydroxymethyl-simvastatin - 6'-exomethylene-simvastatin			
6'-hydroxymethyl-simvastatin - 6'-exomethylene-lovastatin			
6'-hydroxymethyl-simvastatin - N-demethyldesacetyl-diltiazem			
6'-hydroxymethyl-simvastatin - alpha-hydroxyalprazolam			
6'-hydroxymethyl-simvastatin - 4-hydroxyalprazolam			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
6'-hydroxymethyl-simvastatin - N-desmethylrosuvastatin			
6'-hydroxymethyl-simvastatin - 1'-hydroxymidazolam			
6'-exomethylene-simvastatin - N-demethyldiltiazem			
6'-exomethylene-simvastatin - 14-hydroxyclarithromycin			
6'-exomethylene-simvastatin - 6'beta-hydroxy-lovastatin			
6'-exomethylene-simvastatin - para-hydroxy-atorvastatin			
6'-exomethylene-simvastatin - desacetyldiltiazem			
6'-exomethylene-simvastatin - 4-hydroxytriazolam			
6'-exomethylene-simvastatin - ortho-hydroxy-atorvastatin			
6'-exomethylene-simvastatin - 4-hydroxymidazolam			
6'-exomethylene-simvastatin - 6'-exomethylene-lovastatin			
6'-exomethylene-simvastatin - N-demethyldesacetyl-diltiazem			
6'-exomethylene-simvastatin - alpha-hydroxyalprazolam			
6'-exomethylene-simvastatin - 4-hydroxyalprazolam			
6'-exomethylene-simvastatin - N-desmethylrosuvastatin			
6'-exomethylene-simvastatin - 1'-hydroxymidazolam			
1'-hydroxymidazolam - N-demethyldiltiazem			
1'-hydroxymidazolam - 14-hydroxyclarithromycin			
1'-hydroxymidazolam - 6'beta-hydroxy-lovastatin			
1'-hydroxymidazolam - para-hydroxy-atorvastatin			
1'-hydroxymidazolam - desacetyldiltiazem			
1'-hydroxymidazolam - 4-hydroxytriazolam			
1'-hydroxymidazolam - ortho-hydroxy-atorvastatin			
1'-hydroxymidazolam - 4-hydroxymidazolam			
1'-hydroxymidazolam - 6'-exomethylene-lovastatin			
1'-hydroxymidazolam - N-demethyldesacetyl-diltiazem			
1'-hydroxymidazolam - alpha-hydroxyalprazolam			
1'-hydroxymidazolam - 4-hydroxyalprazolam			
1'-hydroxymidazolam - N-desmethylrosuvastatin			
4-hydroxymidazolam - N-demethyldiltiazem			
4-hydroxymidazolam - 14-hydroxyclarithromycin			
4-hydroxymidazolam - 6'beta-hydroxy-lovastatin			
4-hydroxymidazolam - para-hydroxy-atorvastatin			
4-hydroxymidazolam - desacetyldiltiazem			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
4-hydroxymidazolam - 4-hydroxytriazolam			
4-hydroxymidazolam - ortho-hydroxy-atorvastatin			
4-hydroxymidazolam - 6'-exomethylene-lovastatin			
4-hydroxymidazolam - N-demethyldesacetyl-diltiazem			
4-hydroxymidazolam - alpha-hydroxyalprazolam			
4-hydroxymidazolam - 4-hydroxyalprazolam			
4-hydroxymidazolam - N-desmethylrosuvastatin			
4-hydroxytriazolam - N-demethyldiltiazem			
4-hydroxytriazolam - 14-hydroxyclarithromycin			
4-hydroxytriazolam - 6'beta-hydroxy-lovastatin			
4-hydroxytriazolam - para-hydroxy-atorvastatin			
4-hydroxytriazolam - desacetyldiltiazem			
4-hydroxytriazolam - ortho-hydroxy-atorvastatin			
4-hydroxytriazolam - 6'-exomethylene-lovastatin			
4-hydroxytriazolam - N-demethyldesacetyl-diltiazem			
4-hydroxytriazolam - alpha-hydroxyalprazolam			
4-hydroxytriazolam - 4-hydroxyalprazolam			
4-hydroxytriazolam - N-desmethylrosuvastatin			
6'beta-hydroxy-lovastatin - N-demethyldiltiazem			
6'beta-hydroxy-lovastatin - 14-hydroxyclarithromycin			
6'beta-hydroxy-lovastatin - para-hydroxy-atorvastatin			
6'beta-hydroxy-lovastatin - desacetyldiltiazem			
6'beta-hydroxy-lovastatin - ortho-hydroxy-atorvastatin			
6'beta-hydroxy-lovastatin - 6'-exomethylene-lovastatin			
6'beta-hydroxy-lovastatin - N-demethyldesacetyl-diltiazem			
6'beta-hydroxy-lovastatin - alpha-hydroxyalprazolam			
6'beta-hydroxy-lovastatin - 4-hydroxyalprazolam			
6'beta-hydroxy-lovastatin - N-desmethylrosuvastatin			
6'-exomethylene-lovastatin - N-demethyldiltiazem			
6'-exomethylene-lovastatin - 14-hydroxyclarithromycin			
6'-exomethylene-lovastatin - para-hydroxy-atorvastatin			
6'-exomethylene-lovastatin - desacetyldiltiazem			
6'-exomethylene-lovastatin - ortho-hydroxy-atorvastatin			
6'-exomethylene-lovastatin - N-demethyldesacetyl-diltiazem			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
6'-exomethylene-lovastatin - alpha-hydroxyalprazolam			
6'-exomethylene-lovastatin - 4-hydroxyalprazolam			
6'-exomethylene-lovastatin - N-desmethylrosuvastatin			
4-hydroxyalprazolam - N-demethyldiltiazem			
4-hydroxyalprazolam - 14-hydroxyclarithromycin			
4-hydroxyalprazolam - para-hydroxy-atorvastatin			
4-hydroxyalprazolam - desacetyldiltiazem			
4-hydroxyalprazolam - ortho-hydroxy-atorvastatin			
4-hydroxyalprazolam - N-demethyldesacetyl-diltiazem			
4-hydroxyalprazolam - alpha-hydroxyalprazolam			
4-hydroxyalprazolam - N-desmethylrosuvastatin			
alpha-hydroxyalprazolam - N-demethyldiltiazem			
alpha-hydroxyalprazolam - 14-hydroxyclarithromycin			
alpha-hydroxyalprazolam - para-hydroxy-atorvastatin			
alpha-hydroxyalprazolam - desacetyldiltiazem			
alpha-hydroxyalprazolam - ortho-hydroxy-atorvastatin			
alpha-hydroxyalprazolam - N-demethyldesacetyl-diltiazem			
alpha-hydroxyalprazolam - N-desmethylrosuvastatin			
14-hydroxyclarithromycin - N-demethyldiltiazem			
14-hydroxyclarithromycin - para-hydroxy-atorvastatin			
14-hydroxyclarithromycin - desacetyldiltiazem			
14-hydroxyclarithromycin - ortho-hydroxy-atorvastatin			
14-hydroxyclarithromycin - N-demethyldesacetyl-diltiazem			
14-hydroxyclarithromycin - N-desmethylrosuvastatin			
desacetyldiltiazem - N-demethyldiltiazem			
desacetyldiltiazem - para-hydroxy-atorvastatin			
desacetyldiltiazem - ortho-hydroxy-atorvastatin			
desacetyldiltiazem - N-demethyldesacetyl-diltiazem			
desacetyldiltiazem - N-desmethylrosuvastatin			
N-demethyldesacetyl-diltiazem - N-demethyldiltiazem			
N-demethyldesacetyl-diltiazem - para-hydroxy-atorvastatin			
N-demethyldesacetyl-diltiazem - ortho-hydroxy-atorvastatin			
N-demethyldesacetyl-diltiazem - N-desmethylrosuvastatin			
N-demethyldiltiazem - para-hydroxy-atorvastatin			

continued on next page

continued from previous page

drug/drug or drug/metabolite pair	DDI	Non-DDI	Source
N-demethyldiltiazem - ortho-hydroxy-atorvastatin			
N-demethyldiltiazem - N-desmethylrosuvastatin			
ortho-hydroxy-atorvastatin - para-hydroxy-atorvastatin			
ortho-hydroxy-atorvastatin - N-desmethylrosuvastatin			
para-hydroxy-atorvastatin - N-desmethylrosuvastatin			

Appendix H

A BELIEF CRITERIA QUESTIONNAIRE

Questionnaire to establish belief criteria

Each assertion type in the DIKB is listed below in its own section along with DIKB evidence types that can support or refute the assertion. You can assume that all evidence, no matter what type, meets the minimum criteria for quality that we have defined in the DIKB inclusion criteria. Your task is to reflect on your experience and decide which evidence types, or combinations of evidence types, provide information you consider trustworthy for making decisions about the safe use of a drug.

For each assertion type, please list the evidence type(s) whose information. or data, that you would consider believable. For example, there are three evidence type that can support the general assertion regarding the bioavailability of some drug 'X'. Read each evidence type and ask yourself if you trust the validity of a claim about a drug's bioavailability when the information comes from such a study. Then, note which, if any, study types you would find trustworthy. If more than one evidence type meets your belief criteria then list them all separating each evidence type with a comma or and 'OR'. If some combination of the available evidence types would eliminate your doubt in an assertion then, list that combination separating each evidence type with an 'AND'.

If there is no combination of the available evidence types that would relieve your doubt as to the validity of a particular assertion then you can leave your response blank.

the bioavailability of active ingredient 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A pharmacokinetic clinical trial

Your belief criteria:

the primary_total_clearance_mechanism of active ingredient or metabolite 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A pharmacokinetic clinical trial

Your belief criteria:

the maximum_concentration of active ingredient or metabolite 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A pharmacokinetic clinical trial

Your belief criteria:

active ingredient or metabolite 'X' is a substrate_of enzyme 'E'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors

4.A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors

5.A CYP450, recombinant, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

6.A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors

7.A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors

8.A CYP450, human microsome, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

9.A randomized DDI clinical trial

10.A genotyped pharmacokinetic clinical trial

11. A phenotyped pharmacokinetic clinical trial

12. A non-randomized DDI clinical trial

Your belief criteria:

active ingredient or metabolite 'X' is is_not_a_substrate_of enzyme 'E'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors

4.A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors

5.A CYP450, recombinant, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

6.A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors

7.A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors

8.A CYP450, human microsome, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

9.A randomized DDI clinical trial

10.A genotyped pharmacokinetic clinical trial

11. A phenotyped pharmacokinetic clinical trial

12. A non-randomized DDI clinical trial

Your belief criteria:

active ingredient or metabolite 'X' has_metabolite 'M'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A pharmacokinetic clinical trial

4.A drug metabolism identification experiment

5.A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors

6.A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors

7.A CYP450, recombinant, drug metabolism identification experiment (possibly

NO probe enzyme inhibitor(s))

8.A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors

9.A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors

10. A CYP450, human microsome, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

Your belief criteria:

enzyme 'E' controls_formation_of metabolite 'M'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors

4.A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors

5.A CYP450, recombinant, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

6.A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors

7.A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors

8.A CYP450, human microsome, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

9.A randomized DDI clinical trial

10.A genotyped pharmacokinetic clinical trial

11.A phenotyped pharmacokinetic clinical trial

12. A non-randomized DDI clinical trial

Your belief criteria:

the first_pass_effect of active ingredient 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A study of the process by which a drug is absorbed, distributed, metabolized,and eliminated by the body.

Your belief criteria:

the fraction_absorbed of active ingredient 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A study of the process by which a drug is absorbed, distributed, metabolized, and eliminated by the body.

Your belief criteria:

active ingredient or metabolite 'X' increases_auc of active ingredient or metabolite 'Y'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A randomized DDI clinical trial

4.A non-randomized DDI clinical trial

Your belief criteria:

an inhibition_constant for an active ingredient or metabolite 'X' and some enzyme 'E'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, recombinant, metabolic enzyme inhibition experiment

4.A CYP450, human microsome, metabolic enzyme inhibition experiment

Your belief criteria:

active ingredient or metabolite 'X' inhibits enzyme 'E'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, human microsome, metabolic enzyme inhibition experiment

4.A CYP450, recombinant, metabolic enzyme inhibition experiment

5.A randomized DDI clinical trial

6.A non-randomized DDI clinical trial

Your belief criteria:

active ingredient or metabolite 'X' does_not_inhibit enzyme 'E'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, human microsome, metabolic enzyme inhibition experiment

4.A CYP450, recombinant, metabolic enzyme inhibition experiment

5.A randomized DDI clinical trial

6.A non-randomized DDI clinical trial

Your belief criteria:

enzyme 'E' is the primary_total_clearance_enzyme of active ingredient or metabolite 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A randomized DDI clinical trial

4.A genotyped pharmacokinetic clinical trial

5.A phenotyped pharmacokinetic clinical trial

6.A non-randomized DDI clinical trial

Your belief criteria:

enzyme 'E' is the primary_metabolic_clearance_enzyme of active ingredient or metabolite 'X'

1.A non-traceable drug-label statement

2.A non-traceable (but possibly authoritative) statement

3.A CYP450, human microsome, drug metabolism identification experiment using chemical inhibitors

4.A CYP450, human microsome, drug metabolism identification experiment using antibody inhibitors

5.A CYP450, recombinant, drug metabolism identification experiment using chemical inhibitors

6.A CYP450, recombinant, drug metabolism identification experiment using antibody inhibitors

7.A CYP450, recombinant, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

8. A CYP450, recombinant, drug metabolism identification experiment (possibly NO probe enzyme inhibitor(s))

9.A randomized DDI clinical trial

10. A genotyped pharmacokinetic clinical trial

11. A phenotyped pharmacokinetic clinical trial

12. A non-randomized DDI clinical trial

Your belief criteria:

Appendix I

THE AERS IMPLEMENTATION AND OUR USE OF IT

As of June 2008 our implementation of the AERS database is accessible as "AERS-Complete" at:

http://marigold.informatics.washington.edu:7000/phpmyadmin

Interested researcher can log in as "AERS-User" but must first get the password from Richard Boyce. This is single user access meaning that all other users are locked out while a user is using the database. It is possible to give users their own account; they will have read-only access to the database which limits them to querying, but not modifying, the database.

I.1 How to Query

Here is an example of how to query the database:

1. Log in as 'AERS-User' with the password you receive

2. At the left-hand side of the page there is a drop-down box with the word '(Databases)'; select that box and pick the 'AERS-Complete' database.

3. Select the 'SQL' tab from the set of tabs that are shown toward the top of the page (above the table showing the database structure).

4. Enter you query in SQL. For example, you can enter:

```
-- Show all reports involving patients taking FOSAMAX and LIPITOR
SELECT report.*
FROM report
WHERE EXISTS(
SELECT d1.safetyreportid
FROM `drug` AS d1, drug AS d2
WHERE (
d1.medicinalproduct
IN (
```

```
'FOSAMAX'
)
AND d2.medicinalproduct
IN (
'LIPITOR'
)
AND d2.safetyreportid = d1.safetyreportid
AND report.safetyreportid = d2.safetyreportid
)
)
```

The SQL query will be shown above the table of results. If users click the '[edit]' link (to the bottom right of the results) they can edit the query or enter a new one.

I.2 Queries Used to Search for Adverse Event Reports

I.2.1 A Template SQL for Efficient Queries of AERS for DDIs

We replaced the fields enclosed by angle brackets with the appropriate values such as the generic and trade names of drug products containing active pharmaceutical ingredients in the DIKB I.2.2, adverse-event terms representative of a toxic effect from the expected DDI I.2.3, and output-file names.

```
CREATE TEMPORARY TABLE l_reportids (rprt_id varchar(9));

INSERT INTO l_reportids
SELECT d2.safetyreportid
FROM 'drug' AS d1, drug AS d2
WHERE (
d1.medicinalproduct
IN (
<GENERIC AND TRADE NAMES FOR DRUG ONE>
)
```

```
AND d2.medicinalproduct
IN (
<GENERIC AND TRADE NAMES FOR DRUG TWO>
)
AND d2.safetyreportid = d1.safetyreportid
);

CREATE TEMPORARY TABLE rx_reportids (rprt_id varchar(9));

INSERT INTO rx_reportids
SELECT DISTINCT reaction.safetyreportid
FROM reaction INNER JOIN l_reportids ON (reaction.safetyreportid = l_reportids.rprt_id)
WHERE (
     reaction.reactionmeddrapt IN (
     <MedDRA ADVERSE EVENT TERMS REPRESENTITIVE OF THE EXPECTED TOXIC EFFECT>
     )
);


SELECT report.p_key, report.safetyreportid, report.receivedate,report.receiptdate,\
 report.serious, report.seriousness_val, report.qualification, report.patientonsetage,\
 report.patientonsetageunit, report.patientsex, report.patientweight,\
 report.patientdeathdate
FROM report INNER JOIN rx_reportids ON (report.safetyreportid = rx_reportids.rprt_id)
ORDER BY report.safetyreportid
INTO OUTFILE '/opt/Downloads/AERS_Q_Results/<NAME OF FILE TO OUTPUT REPORT META-DATA>'
FIELDS TERMINATED BY '|' OPTIONALLY ENCLOSED BY '"'
LINES TERMINATED BY '\n';


SELECT reaction.p_key, reaction.safetyreportid, reaction.reactionmeddrapt
FROM reaction INNER JOIN rx_reportids ON (reaction.safetyreportid = rx_reportids.rprt_id)
ORDER BY reaction.safetyreportid
INTO OUTFILE '/opt/Downloads/AERS_Q_Results/<NAME OF FILE TO OUTPUT REACTION DATA>'
```

```
FIELDS TERMINATED BY '|' OPTIONALLY ENCLOSED BY '"'
LINES TERMINATED BY '\n';


SELECT  drug.p_key, drug.safetyreportid, drug.drugcharacterization, drug.medicinalproduct,\
 drug.drugdosagetext, drug.drugadministrationroute, drug.drugindication, drug.drugstartdate,\
 drug.drugenddate, drug.drugtreatmentduration, drug.drugtreatmentdurationunit
FROM drug INNER JOIN rx_reportids ON (drug.safetyreportid = rx_reportids.rprt_id)
ORDER BY drug.safetyreportid
INTO OUTFILE '/opt/Downloads/AERS_Q_Results/<NAME OF FILE TO OUTPUT DRUG DATA>'
FIELDS TERMINATED BY '|' OPTIONALLY ENCLOSED BY '"'
LINES TERMINATED BY '\n';
```

I.2.2 Generic and Trade-names for Drug Products Containing Active Pharmaceutical Ingredients in the DIKB

Below is a list of names for drug products containing active pharmaceutical ingredients in the DIKB. We compiled this list from **drugs@fda** [58], the FDA's "Orange Book" [15], and/or RxNorm [19]. Each drug product is 1) oral or injectable, 2) *not* a combined therapy (contained one active ingredient), and 3) present, as of Septemter 2007, in DRUGDEX Tradenames®.

```
'alprazolam':['ALPRAZOLAM', 'ALPRAZOLAM INTENSOL', 'NIRAVAM', 'XANAX', 'XANAX XR'],

'atorvastatin':['ATORVASTATIN','CADUET','LIPITOR'],

'clarithromycin':['CLARITHROMYCIN', 'CLARITHROMYCIN EXTENDED RELEASE',
   'BIAXIN', 'BIAXIN XL'],

'diltiazem':['DILTIAZEM', 'DILTIAZEM HYDROCHLORIDE', 'CARDIZEM','CARDIZEM CD',
  'CARDIZEM LA', 'CARDIZEM LYO-JECT', 'CARDIZEM MONOVIAL', 'CARDIZEM SR',
  'CARTIA', 'CARTIA XT', 'DILACOR XR', 'DILT', 'DILT-CD', 'DILT-XR', 'DILTIA XT',
```

```
  'DILTZAC',  'TAZTIA', 'TAZTIA XT', 'TECZEM', 'TIAMATE', 'TIAZAC'],

'erythromycin':['ERYTHROMYCIN',  'AKNEMYCIN', 'BRISTAMYCIN', 'E-SOLVE-2',
  'E-BASE', 'E-MYCIN', 'E-MYCIN E', 'E-SOLVE 2', 'E.E.S', 'E.E.S. 400 FILMTAB',
  'E.E.S. GRANULES', 'E.E.S.-200', 'E.E.S.-400', 'EMGEL', 'ERY-SOL', 'ERY-TAB',
  'ERYPED', 'ERYC', 'ERYC 125', 'ERYC SPRINKLES', 'ERYMAX', 'ERYPAR', 'ERYPED',
  'ERYTHROCIN', 'ERYTHROCIN STEARATE', 'ERYTHROMYCIN ESTOLATE ', 'ERYTHROMYCIN
   ETHYLSUCCINATE', 'ERYTHROMYCIN LACTOBIONATE', 'ERYTHROMYCIN STEARATE', 'ETHRIL
   250', 'ETHRIL 500', 'ERYZOLE', 'ILOSONE', 'ILOTYCIN', 'PCE', 'PCE BRAND OF
   ERYTHROMYCIN', 'PEDIAMYCIN', 'PEDIAMYCIN 400', 'ROBIMYCIN', 'ROMYCIN',
   'WYAMYCIN E', 'WYAMYCIN S'],

fluconazole':['FLUCONAZOLE', 'DIFLUCAN', 'DIFLUCAN IN DEXTROSE 5% IN PLASTIC
   CONTAINER', 'DIFLUCAN IN SODIUM CHLORIDE 0.9%', 'DIFLUCAN IN SODIUM CHLORIDE
   0.9% IN PLASTIC CONTAINER', 'FLUCONAZALE', 'FLUCONAZOLE IN DEXTROSE 5% IN
   PLASTIC CONTAINER', 'FLUCONAZOLE IN SODIUM CHLORIDE 0.9%', 'FLUCONAZOLE IN
   SODIUM CHLORIDE 0.9% IN PLASTIC CONTAINER'],

'itraconazole':['ITRACONAZOLE','SPORANOX','SPORANOX-PULSE'],

'ketoconazole':['KETOCONAZOLE'],

'lovastatin':['LOVASTATIN', 'ADVICOR', 'ALTOPREV', 'MEVACOR'],

'midazolam':['MIDAZOLAM', 'MIDAZOLAM HYDROCHLORIDE', 'MIDAZOLAM HYDROCHLORIDE
   PRESERVATIVE FREE', 'VERSED'],

'pravastatin':['PRAVASTATIN', 'PRAVASTATIN SODIUM', 'PRAVACHOL'],

'rosuvastatin':['ROSUVASTATIN', 'CRESTOR'],

'simvastatin':['SIMVASTATIN', 'ZOCOR'],

'triazolam':['TRIAZOLAM', 'HALCION'],
```

```
'nefazodone':['NEFAZODONE', 'NEFAZODONE HYDROCHLORIDE', 'SERZONE']
```

I.2.3 Adverse-event Terms Used to Query AERS

Persons who submit a report to AERS are required to note the adverse events that prompted them to send the report. FDA personel code each adverse event using the MedDRA [22] terminology before entering the report into AERS. We compiled a list of MedDRA terms representing the kinds of adverse events that might be observed in patients experiencing toxic side-effects from a victim drugs in one of the DIKB's novel DDI predictions (Chapter 5, Section 5.2.1). We first attempted to utilize the so-called "Standardised MedDRA queries" to build our term sets. These queries are provided by the MedDRA vendor to aid in retrieving cases of interest from databases using the vocabulary [126]. However, we found these to be of little help for the drugs in our system with the exception of the HMG-CoA reductase inhibitors. So, we employed the following process to derive a list of terms we thought more appropriate for querying AERS for DDIs:

1. The two drug experts in our group sent the informaticist a list of words describing the effect of a pharmacokinetic interaction for each relevant drug class. The informaticist also scanned through drug labels to identify other words that might be useful.

2. The informaticist searched the UMLS Meta-Thesaurus [13] for each of the words found in Step One to identify concepts in the meta-thesaurus and their mapping to the MedDRA vocabulary.

3. The informaticist created a list of MedDRA "preffered terms" (PTs) from the terms identified in Step Two then used the program shown in Section I.2.3.1 to expand the PT lists to include all MedDRA "LLTs"

4. The two drug experts reviewed the resulting list and removed all LLTs that they did not think relevant to our search task

Here are the results of this process:

- nefazodone:'ACUTE LIVER DAMAGE', 'BILIOUS VOMITING', 'DAMAGE LIVER', 'DAYTIME SLEEPINESS', 'DISEASE HEPATOCELLULAR', 'DROWSINESS', 'DROWSY ON AWAKENING', 'EMESIS', 'EXCESSIVE DAYTIME SLEEPINESS', 'FEELING OF RESIDUAL SLEEPINESS', 'FEELING QUEASY', 'GROGGY', 'GROGGY AND SLUGGISH', 'GROGGY ON AWAKENING', 'HARD TO AWAKEN', 'HEPATIC DAMAGE', 'HEPATIC DAMAGE (NOS)', 'HEPATIC INTRACELLULAR DEPOSIT OF BILIRUBIN', 'HEPATIC INTRACELLULAR PIGMENTATION', 'HEPATOCELLULAR ABNORMALITY', 'HEPATOCELLULAR DAMAGE', 'HEPATOCELLULAR DAMAGE AGGRAVATED', 'HEPATOCELLULAR DAMAGE NOS', 'HEPATOCELLULAR DISTURBANCES', 'HEPATOCELLULAR INJURY', 'HYPEREMESIS', 'LESS ALERT ON ARISING', 'LIVER CELL DAMAGE', 'LIVER DAMAGE', 'LIVER DAMAGE AGGRAVATED', 'NAUSEA', 'NAUSEA AGGRAVATED', 'NAUSEA ALONE', 'NAUSEA AND VOMITING', 'NAUSEA POST CHEMOTHERAPY', 'NAUSEA VOMITING AND DIARRHEA', 'NAUSEA VOMITING AND DIARRHOEA', 'NAUSEA WITH VOMITING', 'NAUSEATED', 'NAUSEOUS', 'PERSISTENT VOMITING', 'POSTPRANDIAL EMESIS', 'POSTPRANDIAL NAUSEA', 'QUEASY', 'SEROTONIN SYNDROME', 'SICKNESS/NAUSEA', 'SLEEPINESS', 'SLEEPY', 'SOMNOLENCE', 'VOMITED', 'VOMITING', 'VOMITING AGGRAVATED', 'VOMITING ALONE', 'VOMITING NOS'

- clarithromycin, erythromycin (arrhythmia):'CARDIAC ARREST', 'BRADYCARDIA', 'CARDIAC ARREST', 'CARDIAC DEATH', 'CARDIAC TELEMETRY ABNORMAL', 'CARDIO-RESPIRATORY ARREST', 'ELECTROCARDIOGRAM ABNORMAL', 'ELECTROCARDIOGRAM AMBULATORY ABNORMAL', 'ELECTROCARDIOGRAM CHANGE', 'ELECTROCARDIOGRAM REPOLARISATION ABNORMALITY', 'HEART RATE ABNORMAL', 'HEART RATE DECREASED', 'LOSS OF CONSCIOUSNESS', 'PALPITATIONS', 'SUDDEN CARDIAC DEATH', 'SUDDEN DEATH', 'SYNCOPE'

- clarithromycin, erythromycin (Torsade de Pointes):'LONG QT SYNDROME', 'TORSADE DE POINTES', 'VENTRICULAR TACHYCARDIA', 'ELECTROCARDIOGRAM QT PROLONGED'

- clarithromycin, erythromycin (hepato-toxicity): 'HEPATOCELLULAR DAMAGE'

- clarithromycin, erythromycin (general side effects):'ABDO. DISCOMFORT', 'ABDOMEN BURNING SENSATION OF', 'ABDOMINAL DISCOMFORT', 'ABDOMINAL DISTRESS', 'ABDOMINAL PAIN LOWER', 'ABDOMINAL PAIN PEPTIC ULCER TYPE', 'ABDOMINAL PAIN UPPER', 'ACHE STOMACH', 'ACUTE DIARRHEA', 'ACUTE DIARRHOEA', 'ACUTE GASTRIC PAIN', 'BOWEL DISCOMFORT', 'BURNING IN ABDOMEN', 'BURNING SENSATION IN ABDOMEN', 'CHRONIC DIARRHEA', 'CHRONIC DIARRHOEA', 'CHRONIC EPIGASTRIC PAIN', 'CHURNING OF STOMACH', 'CRAMP IN LOWER ABDOMEN', 'DIARRHEA', 'DIARRHEA AGGRAVATED', 'DIARRHEA NOS', 'DIARRHEA RECURRENT',

'DIARRHOEA', 'DIARRHOEA AGGRAVATED', 'DIARRHOEA NOS', 'DIARRHOEA RECURRENT', 'DISCOMFORT ABDOMINAL', 'DISTRESS ABDOMINAL', 'DISTRESS GASTROINTESTINAL', 'EMESIS', 'EPIGASTRALGIA', 'EPIGASTRIC ACHE', 'EPIGASTRIC PAIN', 'EPIGASTRIC PAIN EPIGASTRALGIA', 'EPIGASTRIC PAIN NOT FOOD-RELATED', 'EXPLOSIVE DIARRHEA', 'EXPLOSIVE DIARRHOEA', 'GASTRALGIA', 'GASTRIC PAIN', 'GASTRIC SPASM', 'GASTROINTESTINAL DISCOMFORT', 'GASTROINTESTINAL IRRITATION', 'GASTROINTESTINAL UPSET', 'GI IRRITATION', 'GI UPSET', 'HYPEREMESIS', 'HYPOCHONDRIAL PAIN', 'HYPOCHONDRIUM PAIN LEFT', 'HYPOCHONDRIUM PAIN RIGHT', 'HYPOGASTRIC PAIN', 'IATROGENIC DIARRHEA', 'IATROGENIC DIARRHOEA', 'ILIAC FOSSA PAIN', 'IRRITATION GASTROINTESTINAL', 'LOOSE BOWEL', 'LOOSE BOWELS', 'LOOSE MOTIONS', 'LOOSE STOOLS', 'LOWER ABDOMINAL PAIN', 'MUCOUS DIARRHEA', 'MUCOUS DIARRHOEA', 'MUSHY DIARRHEA', 'MUSHY DIARRHOEA', 'NAUSEA', 'NAUSEA AGGRAVATED', 'NAUSEA ALONE', 'NAUSEA AND VOMITING', 'NAUSEA VOMITING AND DIARRHEA', 'NAUSEA VOMITING AND DIARRHOEA', 'NAUSEA WITH VOMITING', 'NAUSEATED', 'NAUSEOUS', 'NOCTURNAL DIARRHEA', 'NOCTURNAL DIARRHOEA', 'PAIN EPIGASTRIC', 'PAIN GASTRIC', 'PAIN STOMACH', 'PERSISTENT VOMITING', 'POSTPRANDIAL EMESIS', 'POSTPRANDIAL NAUSEA', 'SECRETORY DIARRHEA', 'SECRETORY DIARRHOEA', 'SICKNESS/NAUSEA', 'SOFT STOOLS', 'STOMACH ACHE', 'STOMACH CRAMPS', 'STOMACH DULL PAIN OF', 'STOMACH PAIN', 'STOOLS LOOSE', 'STOOLS WATERY', 'ULCER TYPE PAIN', 'UPPER ABDOMINAL DISCOMFORT', 'UPPER ABDOMINAL PAIN', 'UPSET GASTROINTESTINAL', 'URGENT DIARRHEA', 'URGENT DIARRHOEA', 'VOMITED', 'VOMITING', 'VOMITING AGGRAVATED', 'VOMITING ALONE', 'VOMITING NOS', 'VOMITING REFLEX', 'WATERY DIARRHEA', 'WATERY DIARRHOEA'

- diltiazem:'ABNORMAL ECG', 'ABNORMAL EKG', 'ACUTE HYPOTENSION', 'ACUTE MYOPATHY', 'ARREST CARDIAC', 'ARTERIAL HYPOTENSION', 'ASYSTOLE', 'ASYSTOLIC', 'BLOOD MYOGLOBIN INCREASED', 'BLOOD PRESSURE LOW', 'BRADYCARDIA', 'BRADYCARDIA NOS', 'BRADYCARDIA NOS (EXCL FOETAL)', 'CK INCREASED', 'CPK INCREASE', 'CPK INCREASED', 'CPK-MM INCREASED', 'CARDIAC ARREST', 'CARDIAC ARREST TRANSIENT', 'CARDIAC DEATH', 'CARDIAC SYNCOPE', 'CARDIAC TELEMETRY ABNORMAL', 'CARDIO-RESPIRATORY ARREST', 'CARDIOPULMONARY ARREST', 'CHANGE IN ECG', 'CONSCIOUSNESS AWAKING LOSS', 'CONSCIOUSNESS LOSS', 'CONSCIOUSNESS LOSS OF', 'CREATINE KINASE HIGH', 'CREATINE KINASE INCREASED', 'CREATINE PHOSPHOKINASE INCREASED', 'CREATINE PHOSPHOKINASE SERUM INC', 'CREATINE PHOSPHOKINASE SERUM INCREASED', 'CREATININE ABNORMAL NOS', 'CREATININE BLOOD INCREASED', 'CREATININE SERUM INCREASED', 'DEATH OCCURRING IN LESS THAN 24 HOURS FROM ONSET OF SYMPTOMS, NOT OTHERWISE EXPLAINED', 'DEATH SUDDEN', 'DEATH SUDDEN (NOS)', 'DISORDER ECG/EKG (NOS)', 'ECG EKG ABNORMAL (NOS)', 'ECG ABNORMAL', 'ECG ABNORMAL NOS', 'ECG ABNORMAL NON-SPECIFIC', 'ECG ABNORMAL SPECIFIC', 'ECG ABNORMALITIES NON-SPECIFIC', 'ECG PLUS VOLTAGE MARKED', 'ECG/EKG CHANGES NON-SPECIFIC', 'EKG ABNORMAL', 'EKG ABNORMAL NON-SPECIFIC', 'EKG/ECG ABNORMALITIES NON-SPECIFIC', 'ELECTROCARDIOGRAM ABNORMAL', 'ELECTROCARDIOGRAM ABNORMAL

(NOS)', 'ELECTROCARDIOGRAM ABNORMAL NOS', 'ELECTROCARDIOGRAM ABNORMAL NON-SP', 'ELECTROCARDIOGRAM ABNORMAL NON-SPECIFIC', 'ELECTROCARDIOGRAM ABNORMAL SPECIFIC', 'ELECTROCARDIOGRAM AMBULATORY ABNORMAL', 'ELECTROCARDIOGRAM CHANGE', 'ELECTROCARDIOGRAM CHANGE NOS', 'GENERALISED MUSCLE ACHES', 'GENERALIZED MUSCLE ACHES', 'HEART ARREST', 'HEART RATE ABNORMAL', 'HEART RATE DECREASED', 'HEART RATE LOW', 'HOLTER MONITORING ABNORMAL', 'HYPOTENSION', 'HYPOTENSION NOS', 'HYPOTENSION AGGRAVATED', 'HYPOTENSION ASYMPTOMATIC', 'HYPOTENSION PAROXYSM', 'HYPOTENSION SYMPTOMATIC', 'HYPOTENSION, UNSPECIFIED', 'HYPOTENSIVE', 'HYPOTENSIVE EPISODE', 'IATROGENIC HYPOTENSION', 'LOC', 'LOCALISED MUSCLE PAIN', 'LOCALIZED MUSCLE PAIN', 'LOSS OF CONSCIOUSNESS', 'LOSS OF CONSCIOUSNESS NEC', 'LOSS OF CONSIOUSNESS', 'LOST CONSCIOUSNESS', 'LOW BP', 'LOW BLOOD PRESSURE', 'LOW PULSE RATE', 'MUSCLE ACHE', 'MUSCLE PAIN', 'MUSCLE SORENESS', 'MUSCLE TENDERNESS ANY SITE', 'MUSCULAR PAIN', 'MUSCULAR PAINS', 'MYALGIA', 'MYALGIA AGGRAVATED', 'MYALGIA OF LOWER EXTREMITIES', 'MYOGLOBIN BLOOD INCREASED', 'MYOGLOBIN BLOOD PRESENT', 'MYOGLOBIN URINE INCREASED', 'MYOGLOBIN URINE PRESENT', 'MYOGLOBINAEMIA', 'MYOGLOBINEMIA', 'MYOGLOBINURIA', 'MYONECROSIS', 'MYOPATHY', 'MYOPATHY AGGRAVATED', 'MYOPATHY

TOXIC', 'MYOPATHY, UNSPECIFIED', 'NONSPECIFIC ABNORMAL ELECTROCARDIOGRAM (ECG) (EKG)', 'ORTHOSTATIC COLLAPSE', 'OTHER MYOPATHIES', 'PALPITATION', 'PALPITATIONS', 'PALPITATIONS AGGRAVATED', 'PHOSPHOKINASE CREATINE SERUM INCREASED', 'PLASMA CREATINE PHOSPHOKINASE ABNORMAL', 'PLASMA CREATINE PHOSPHOKINASE INCREASED', 'POLYMYALGIA', 'POLYMYALGIA AGGRAVATED', 'POLYMYALGIA WORSENED', 'PROXIMAL MYOPATHY', 'PROXIMAL MYOPATHY AGGRAVATED', 'PULSE DECREASED', 'PULSE RATE DECREASE', 'PULSE RATE DECREASE MARKED', 'PULSE RATE DECREASED', 'PULSE RATE FELL', 'PULSE RATE LOW', 'RHABDOMYOLYSIS', 'SERUM CREATINE PHOSPHOKINASE ABNORMAL', 'SERUM CREATINE PHOSPHOKINASE INCREASED', 'SERUM CREATININE ABNORMAL', 'SERUM CREATININE INCREASED', 'SLOW PULSE', 'STANDSTILL CARDIAC', 'STANDSTILL CARDIAC', 'SUDDEN CARDIAC DEATH', 'SUDDEN DEATH', 'SUDDEN DEATH NOS', 'SUDDEN DEATH UNEXPLAINED', 'SUDDEN DEATH, CAUSE UNKNOWN', 'SYNCOPAL ATTACK', 'SYNCOPE', 'SYNCOPE AGGRAVATED', 'SYNCOPE CONVULSIVE', 'SYNCOPE EXERTIONAL', 'SYNCOPE HYPOTENSIVE', 'SYNCOPE POSTURAL', 'TENDERNESS MUSCLE', 'TOXIC MYOPATHY', 'TRANSIENT SYSTOLIC HYPOTENSION', 'UNCONSCIOUS', 'UNCONSCIOUSNESS'

- lovastatin, simvastatin, atorvastatin:'ACUTE MYOPATHY', 'BLOOD CREATINE PHOSPHOKINASE ABNORMAL', 'BLOOD CREATINE PHOSPHOKINASE ABNORMAL NOS', 'BLOOD CREATINE PHOSPHOKINASE INCREASED', 'BLOOD CREATINE PHOSPHOKINASE MM INCREASED', 'BLOOD CREATININE ABNORMAL', 'BLOOD CREATININE INCREASED', 'BLOOD MYOGLOBIN INCREASED', 'CK INCREASED', 'CPK INCREASE', 'CPK INCREASED', 'CPK-MM INCREASED', 'CREATINE KINASE HIGH', 'CREATINE KINASE INCREASED', 'CREATINE PHOSPHOKINASE INCREASED', 'CRE-

ATINE PHOSPHOKINASE SERUM INC', 'CREATINE PHOSPHOKINASE SERUM INCREASED', 'CREATININE ABNORMAL NOS', 'CREATININE BLOOD INCREASED', 'CREATININE HIGH', 'CREATININE INCREASED', 'CREATININE SERUM INCREASED', 'GENERALISED MUSCLE ACHES', 'GENERALIZED MUSCLE ACHES', 'INCREASED SERUM CREATININE', 'LOCALISED MUSCLE PAIN', 'LOCALIZED MUSCLE PAIN', 'MUSCLE ACHE', 'MUSCLE DISSOLUTION', 'MUSCLE NECROSIS', 'MUSCLE PAIN', 'MUSCLE SORENESS', 'MUSCLE TENDERNESS ANY SITE', 'MUSCULAR PAIN', 'MUSCULAR PAINS', 'MYALGIA', 'MYALGIA AGGRAVATED', 'MYALGIA OF LOWER EXTREMITIES', 'MYOGLOBIN BLOOD INCREASED', 'MYOGLOBIN BLOOD PRESENT', 'MYOGLOBIN URINE INCREASED', 'MYOGLOBIN URINE PRESENT', 'MYOGLOBINAEMIA', 'MYOGLOBINEMIA', 'MYOGLOBINURIA', 'MYONECROSIS', 'MYOPATHY', 'MYOPATHY AGGRAVATED', 'MYOPATHY TOXIC', 'MYOPATHY, UNSPECIFIED', 'OTHER MYOPATHIES', 'PAIN MUSCLE', 'PHOSPHOKINASE CREATINE SERUM INCREASED', 'PLASMA CREATINE PHOSPHOKINASE ABNORMAL', 'PLASMA CREATINE PHOSPHOKINASE INCREASED', 'PLASMA CREATININE ABNORMAL', 'PLASMA CREATININE INCREASED', 'POLYMYALGIA', 'POLYMYALGIA AGGRAVATED', 'POLYMYALGIA WORSENED', 'PROXIMAL MYOPATHY', 'PROXIMAL MYOPATHY AGGRAVATED', 'RAISED SERUM CREATININE', 'RHABDOMYOLYSIS', 'SERUM CREATINE PHOSPHOKINASE ABNORMAL', 'SERUM CREATINE PHOSPHOKINASE INCREASED', 'SERUM CREATININE ABNORMAL', 'SERUM CREATININE INCREASED', 'SYMPTOMATIC INFLAMMATORY MYOPATHY', 'SYMPTOMATIC INFLAMMATORY MYOPATHY IN DISEASES CLASSIFIED ELSEWHERE', 'TENDERNESS MUSCLE', 'TOXIC MYOPATHY', 'URINE MYOGLOBIN INCREASED'

I.2.3.1 A Python Program to Map MedDRA PTs to LLTs

```
## map-pt-to-llt.py
##
## get a list of all 'LLT' terms for a given 'PT' term
## Requires  data from the MedDRA ascii files

pt_f = open('pt.asc','r')
pt_buf = pt_f.read()
pt_f.close()
pt_l = pt_buf.split('\r\n')
pt_d = {}
for pt in pt_l:
    pt_att = pt.split('$')
    if not len(pt_att) > 1:
        break
    pt_code = pt_att[0]
    pt_name = pt_att[1]
```

```
    pt_d[pt_name] = pt_code

ll_f = open('llt.asc','r')
llt_buf = ll_f.read()
ll_f.close()
llt_l = llt_buf.split('\r\n')
llt_d = {}
for llt in llt_l:
    llt_att = llt.split('$')
    if not len(llt_att) > 1:
        break
    (llt_name, pt_code) = (llt_att[1], llt_att[2])
    if not llt_d.has_key(pt_code):
        llt_d[pt_code] = [llt_name]
    else:
        llt_d[pt_code].append(llt_name)

## example query, returns a list of all LLT's for a PT term
#llt_d[pt_d['Bradycardia']]

## rhabdo PTs expanded to LLTs
rhabdoPTs = ['Muscle necrosis', 'Myoglobin blood increased', 'Myoglobin blood present',\
   'Myoglobin urine present', 'Myoglobinaemia', 'Myoglobinuria', 'Myopathy', \
   'Myopathy toxic', 'Rhabdomyolysis', 'Blood creatine phosphokinase abnormal', \
   'Blood creatine phosphokinase increased',\
   'Blood creatine phosphokinase MM increased', 'Blood creatinine abnormal',\
   'Blood creatinine increased',\
   'Myalgia', 'Myalgia intercostal']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in rhabdoPTs])


## diltiazem PTs expanded to LLTs
diltPTs = ['Cardiac arrest', 'Bradycardia', 'Cardiac arrest', 'Cardiac death', \
   'Cardiac telemetry abnormal',\
   'Cardio-respiratory arrest', 'Electrocardiogram abnormal', \
   'Electrocardiogram ambulatory abnormal', \
   'Electrocardiogram change', 'Electrocardiogram repolarisation abnormality', \
   'Gallop rhythm present',\
   'Heart rate abnormal', 'Heart rate decreased', 'Loss of consciousness', \
```

```
    'Palpitations', 'Sudden cardiac death',\
   'Sudden death', 'Syncope', 'Hypotension']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in diltPTs])

## nefazodone PTs expanded to LLTs
nefazPTs = ['Nausea', 'Vomiting', 'Somnolence', 'Hepatocellular damage', 'Serotonin syndrome']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in nefazPTs])

## Macrolide PTs from label
macroPTs_label = ['Nausea', 'Vomiting', 'Abdominal discomfort', 'Abdominal pain lower',\
   'Abdominal pain upper', 'Diarrhoea']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in macroPTs_label])

## macrolide arrhithmia-related PTs
macroArrythPTs = ['Cardiac arrest', 'Bradycardia', 'Cardiac arrest', \
    'Cardiac death', 'Cardiac telemetry abnormal',\
    'Cardio-respiratory arrest', 'Electrocardiogram abnormal', \
    'Electrocardiogram ambulatory abnormal', 'Electrocardiogram change',\
    'Electrocardiogram repolarisation abnormality', 'Gallop rhythm present',\
    'Heart rate abnormal', 'Heart rate decreased',\
    'Loss of consciousness', 'Palpitations', 'Sudden cardiac death',\
     'Sudden death', 'Syncope']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in macroArrythPTs])

## macrolide QT PTs
macroQT_PTs = ['Long QT syndrome', 'Torsade de pointes', \
   'Ventricular tachycardia', 'Electrocardiogram QT prolonged']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in macroQT_PTs])

## macrolide liver damage
macroLiverPT = ['Hepatocellular damage']
print "\n".join(["\n".join(llt_d[pt_d[t]]) for t in macroLiverPT])
```

Appendix J

A SAMPLE AERS REPORT RETURNED FROM OUR QUERIES

1 : clarithromycin-nefazodone-interaction REPORT: 4270220-82001st (safetyreportid: 4270220-8)

First received: 2003-12-17 Most recent info: 2003-05-30 qualification: physician

PATIENT INFO

---

age: 41 (years) gender: female weight(kg): 70.3 death date: \N

REACTION

---

seriousness list: seriousnesshospitalization,seriousnessother

Reaction (MedDRA):

- ABDOMINAL DISCOMFORT
- ANXIETY
- BRONCHITIS
- COORDINATION ABNORMAL
- CRYING
- DELUSION
- DEPRESSION
- DIFFICULTY IN WALKING
- DRUG ABUSER
- DRUG DEPENDENCE
- DRUG INEFFECTIVE
- DRUG WITHDRAWAL SYNDROME
- HALLUCINATION, AUDITORY
- HYPERHIDROSIS
- INSOMNIA
- MALAISE
- MUSCLE CRAMP
- NAUSEA
- PALPITATIONS
- PANIC REACTION
- PNEUMONIA
- RASH
- RIGORS
- ROAD TRAFFIC ACCIDENT
- SUICIDAL IDEATION

TINNITUS
VISUAL DISTURBANCE
VOMITING

MEDICATIONS

Medicinal product: OXYCONTIN
indication: PAIN
route: \N

Dosage: 20 MG

Start date: 1997-12-18 End date: 1999-03-01 Tx duration: \N (units unkown)

characterization: suspect

Medicinal product: OXYCONTIN
indication: PAIN
route: \N

Dosage: 40 MG

Start date: 1999-03-23 End date: 2000-05-01 Tx duration: \N (units unkown)

characterization: suspect

Medicinal product: OXYCONTIN
indication: PAIN
route: \N

Dosage: 80 MG, TID

Start date: 2000-05-01 End date: 2001-03-19 Tx duration: \N (units unkown)

characterization: suspect

Medicinal product: OXYCODONE HCL
indication: PAIN

route: \N

Dosage: 5 MG

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: suspect

Medicinal product: ALBUTEROL

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: AMBIEN

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: AMERGE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: BENZONATATE

indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: BIAXIN
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: BUSPAR
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: CEPHALEXIN
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: CLONAZEPAM
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: VALIUM
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: DOXEPIN HYDROCHLORIDE
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: EFFEXOR
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: FLOVENT
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: FLOXIN ^R.W. JOHNSON^
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: HYDROCODONE W/ACETAMINOPHEN
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: IMITREX ^GLACO-WELLCOME^
indication: \N
route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: LEVAQUIN

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: MEDROXYPROGESTERONE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: NEURONTIN

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: PAXIL

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: PREDNISONE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: PROMETHAZINE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: PROTUSS-DM

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: ROXICET

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: ROXICODONE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: SERZONE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: SONATA

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: TRAZODONE HCL

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: ZOLOFT

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: REMERON

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: WELLBUTRIN

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: DOXYCYCLINE

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Medicinal product: KETOPROFEN

indication: \N

route: \N

Dosage: \N

Start date: \N End date: \N Tx duration: \N (units unkown)

characterization: concomitant

Appendix K

DIPS EVALUATIONS OF CASE REPORTS

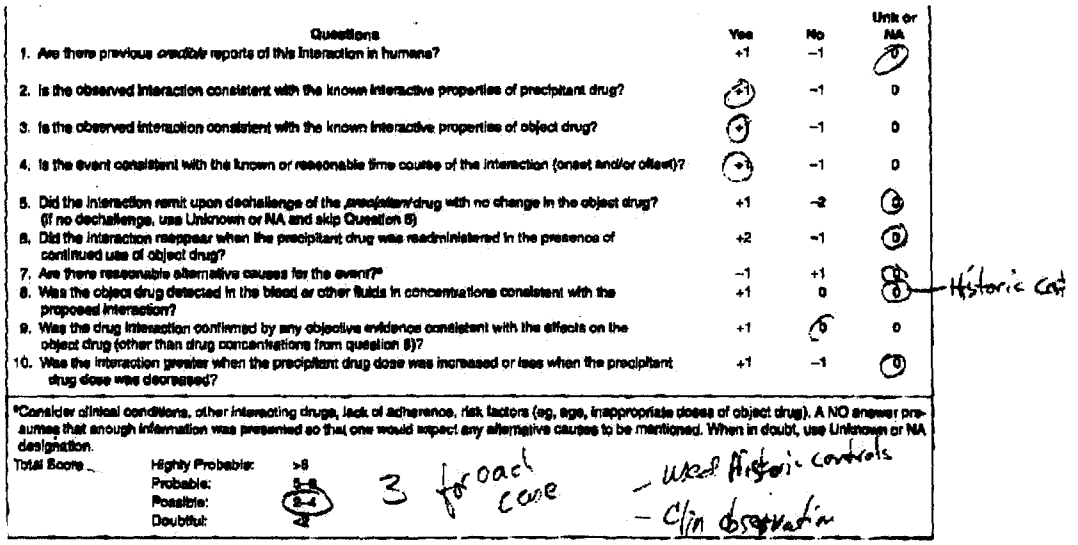

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

Historic Cont

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score _ Highly Probable: >8
Probable: 5–8
Possible: 1–4
Doubtful: ≤0

3 for each case
– used historic controls
– Clin observation

Figure K.1: DIPS evaluation of a paper involving multiple case reports published by Auclair *et al* providing evidence of an interaction between itraconazole and clarithromycin [11]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

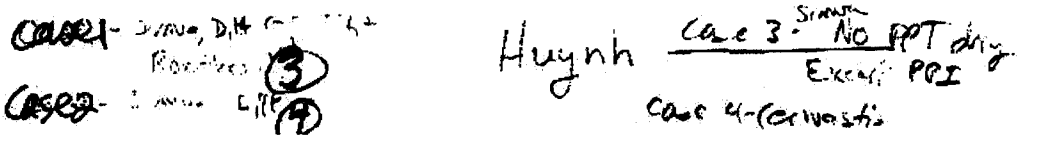

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction *remit* upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	>8
Probable:	5–8
Possible:	2–4
Doubtful:	<2

3-4

Figure K.2: DIPS evaluation of a paper involving multiple case reports published by Huynh *et al* providing evidence of an interaction between diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [89]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	(+1)	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	(+1)	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	(+1)	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	(+1)	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	(0)
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	(0)
7. Are there reasonable alternative causes for the event?[a]	−1	+1	(0)
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	(0)
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	(+1)	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	(0)

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	≥9
Probable:	(5–8)
Possible:	2–4
Doubtful:	<2

Figure K.3: DIPS evaluation of a case report published by Akram *et al* providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [5]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___ 5

Highly Probable:	>8
Probable:	5–8
Possible:	2–4
Doubtful:	≤2

Figure K.4: DIPS evaluation of a case report published by Itakura *et al* providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [91]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

All pts also on PPI's

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction recur upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

PPI?
- Rabbo
- OATP

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___ Highly Probable: >8
Probable: 5–8
Possible: 2–4
Doubtful: ≤2

All score 4

Figure K.5: DIPS evaluation of a paper involving multiple case reports published by Gladding *et al* providing evidence of seperate interactions between 1) diltiazem and atorvastatin, and 2) diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [67]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (if no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ____ . Highly Probable: >8
Probable: 5–8
Possible: 1–4
Doubtful: 0 or less

4

Figure K.6: DIPS evaluation of a case report published by Shaukat *et al* providing evidence of a DDI between fluconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [157]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score 5

Highly Probable:	≥9
Probable:	5–8
Possible:	2–4
Doubtful:	<2

Figure K.7: DIPS evaluation of a case report published by Peces and Pobes providing evidence of a DDI between diltiazem and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [140]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (if no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	>8
Probable:	5–8
Possible:	2–4
Doubtful:	<2

3

unk – pt. had CHF, vascular disease ↓ C.O ... renal output

CK Elevation mild No Myoglobin in urine.

Figure K.8: DIPS evaluation of a case report published by Lewin *et al* providing evidence of a DDI between diltiazem and atorvastatin [113]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	(+1)	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	(+1)	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	(+1)	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	(+1)	−1	0
5. Did the interaction recur upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	(0)
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	(0)
7. Are there reasonable alternative causes for the event?*	−1	(+1)	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	(0)
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	(+1)	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	(0)

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	>8	Case 1 = 6
Probable:	(5-8)	
Possible:	2-4	Case 2 = 6
Doubtful:	<2	

Figure K.9: DIPS evaluation of two case reports published by Gilad and Lampl providing evidence of a DDI between ketoconazole and simvastatin/simvastatin acid (beta-hydroxy-simvastatin) [66]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	>8
Probable:	5-8
Possible:	2-4
Doubtful:	<2

Figure K.10: DIPS evaluation of a case report published by Spach *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [161]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of *precipitant* drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___ Highly Probable: ≥9
Probable: 5–8
Possible: 2–4
Doubtful: ≤2

Figure K.11: DIPS evaluation of a case report published by Wong *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [179]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?[a]	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

[a]Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate dose of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score 5

Highly Probable:	>8
Probable:	5–8
Possible:	2–4
Doubtful:	<2

Figure K.12: DIPS evaluation of a case report published by Ayanian *et al* providing evidence of a DDI between erythromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [28]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score		
	Highly Probable:	>8
	Probable:	5–8
	Possible:	2–4
	Doubtful:	<2

Figure K.13: DIPS evaluation of two case reports published by Stein *et al* providing evidence of a DDI between ketoconazole and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [162]. This report received a score of five on the DIPS scale giving it a DIPS rating of "probable". See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (if no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___ Highly Probable: >8
Probable: 5–8
Possible: 1–4
Doubtful: <2

Figure K.14: DIPS evaluation of a case report published by Kahri *et al* providing evidence of a DDI between fluconazole and atorvastatin [101]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

Questions	Yes	No	Unk or NA
1. Are there previous *credible* reports of this interaction in humans?	+1	−1	0
2. Is the observed interaction consistent with the known interactive properties of precipitant drug?	+1	−1	0
3. Is the observed interaction consistent with the known interactive properties of object drug?	+1	−1	0
4. Is the event consistent with the known or reasonable time course of the interaction (onset and/or offset)?	+1	−1	0
5. Did the interaction remit upon dechallenge of the *precipitant* drug with no change in the object drug? (If no dechallenge, use Unknown or NA and skip Question 6)	+1	−2	0
6. Did the interaction reappear when the precipitant drug was readministered in the presence of continued use of object drug?	+2	−1	0
7. Are there reasonable alternative causes for the event?*	−1	+1	0
8. Was the object drug detected in the blood or other fluids in concentrations consistent with the proposed interaction?	+1	0	0
9. Was the drug interaction confirmed by any objective evidence consistent with the effects on the object drug (other than drug concentrations from question 8)?	+1	0	0
10. Was the interaction greater when the precipitant drug dose was increased or less when the precipitant drug dose was decreased?	+1	−1	0

*Consider clinical conditions, other interacting drugs, lack of adherence, risk factors (eg, age, inappropriate doses of object drug). A NO answer presumes that enough information was presented so that one would expect any alternative causes to be mentioned. When in doubt, use Unknown or NA designation.

Total Score ___

Highly Probable:	>8
Probable:	5–8
Possible:	2–4
Doubtful:	<2

Case 1 = 3

Case 2 = 2

Figure K.15: DIPS evaluation of a case report published by Grunden and Fisher providing evidence of a DDI between clarithromycin and lovastatin/lovastatin acid (beta-hydroxy-lovastatin) [77]. See Table 5.14 in Chapter 5 for more details on DIPS evaluations.

VITA

AUTHOR

Richard David Boyce

EDUCATION

PhD, Biomedical and Health Informatics, University of Washington, Seattle Washington, July, 2008

Masters of Science, Biomedical and Health Informatics, University of Washington, Seattle Washington, December, 2005

Bachelor of Science, Computer Science, Central Washington University, Ellensburg Washington, June, 2003

HONORS AND AWARDS

National Library of Medicine Predoctoral Fellow in Biomedical Informatics 2003- 2006

NSF Computer Science Scholarship 2001-2003

ORGANIZATIONS

Member IEEE

Member American Medical Informatics Association

LIFE NOTES

Richard Boyce was born in Butte Montana in 1974 and is proud to be the first person in his family to finish a college degree. He and his wife, Kriss, are the busy parents of four sons and one daughter: Daymen, Paul, Zechariah, Isaiah, and FaithAnne.