K-12 Guidelines for Artificial Intelligence: What Students Should Know

David S. Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn

Supported by NSF DRL-1846073.

Outline

1. Why is this the right time to begin teaching AI in K-12?
2. Overview of the AI4K12 Initiative
3. Exploring the “Big Ideas” in AI
    - What are the big ideas?
    - Grade band learning progression
4. Resources for teaching the big ideas in K-12
5. Discussion and Q&A

Why is this the right time to be teaching AI in K-12?

1. AI is playing an increasingly prominent role in society:
    - Intelligent assistants
    - Self-driving cars
    - Autonomous robots in the workplace (and someday the home)
2. Informed citizens need to understand the basics of AI as our society faces important public policy decisions surrounding AI technologies.
3. AI technologies will cause job loss in some areas, and gains in other areas.
4. There is a growing need for AI-literate workers. Students should be encouraged from a young age to consider STEM careers.

A basic understanding

In the near future, perhaps sooner than we think, **virtually everyone will need a basic understanding of** the technologies that underpin **machine learning and artificial intelligence.**

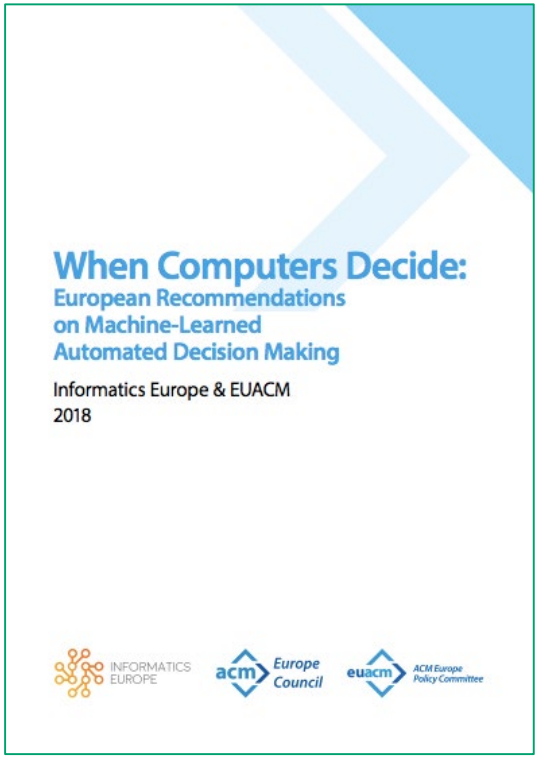

ACM Europe / Informatics Europe

For some, or for all?

Children need to be adequately prepared for working with, and using, AI.

For a proportion, this will mean a **thorough education in AI-related subjects**, requiring adequate resourcing of the computing curriculum and support for teachers.

For all children, the **basic knowledge and understanding** necessary to navigate an AI driven world will be essential. In particular, we recommend that the ethical design and use of technology becomes an integral part of the curriculum.

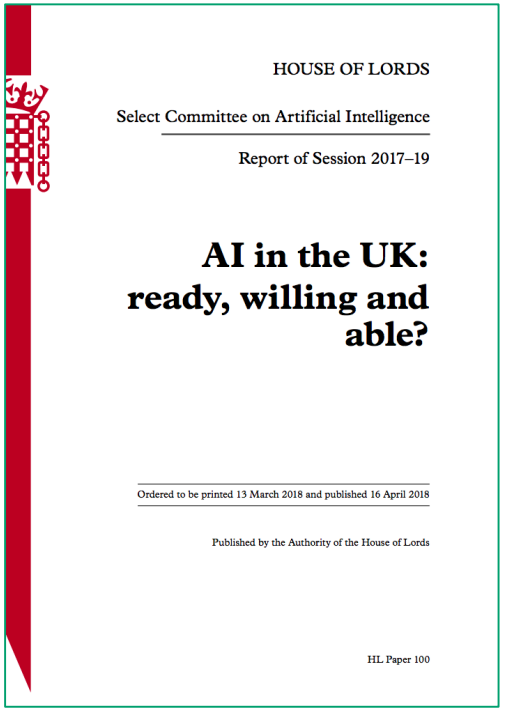
HOUSE OF LORDS

Select Committee on Artificial Intelligence

Report of Session 2017–19

AI in the UK: ready, willing and able?

Ordered to be printed 13 March 2018 and published 16 April 2018

Published by the Authority of the House of Lords

HL Paper 100

AI4K12.org Mission

- Develop national guidelines for teaching AI in K-12
  - Modeled after the CSTA standards for computing education.
  - Four grade bands: K-2, 3-5, 6-8, and 9-12
  - What should students know?
  - What should students be able to do?
- Develop a curated AI resource directory for K-12 teachers
- Foster a community of K-12 AI resource developers

The AI4K12 Initiative, a joint project of:

AAAI (Association for the Advancement of Artificial Intelligence)

CSTA (Computer Science Teachers Association)

With funding from National Science Foundation ITEST Program (DRL-1846073)

Carnegie Mellon University
School of Computer Science

Steering Committee

Dave Touretzky
Carnegie Mellon
AI for K-12 Working Group Chair

Christina Gardner-McCune
University of Florida
AI For K-12 Working Group Co-Chair

Fred Martin
UMass Lowell
CSTA Chair of Board of Directors

Deborah Seehorn
Co-Chair of CSTA Standards Committee

K-12 Teacher Working Group Members

Grades K-2
Vicky Sedgwick (Lead)
Susan Amsler-Akacem
Dr. April DeGennaro
Charlotte Dungan

Grades 3-5
Kelly Powers (Lead)
Dr. Marlo Barnett
Dr. Phillip Eaglin
Brian Stamford

Grades 6-8
Padmaja Bandaru (Lead)
Minsoo Park
Juan Palomares
Josh Caldwell
Sheena Vaidyanathan

Grades 9-12
Dianne O'Grady-Cunniff (Lead)
Jared Amalong
Dr. Smadar Bergman
Kate Lockwood

Academia/Industry Working Group Members

Hal Abelson
MIT

Cynthia Breazeal
MIT

Matt Dawson
Google

Emily Reid
AI4ALL

Matthijs Spaan
TU Delft
AAAI

AI4K12.org Advisory Board

Miles Berry, Roehampton University, UK

Amy Eguchi, UC San Diego, CA

Laura Schmidt, Milwaukee School of Engineering University, WI

Maitreyee Joshi, Microsoft, Seattle, WA

Irene Lee, MIT, Cambridge, MA

Dahua Lin, Chinese University of Hong Kong, China

Hari Raghavan, IBM, New York, NY

Joseph South, ISTE, Portland, OR

Tom Vander Ark, Getting Smart, Federal Way, WA

Joyce D. Williams, National Geospatial-Intelligence Agency, VA

K-12 Computing Education

Worldwide, we are making progress on integrating **computing** into K-12:

- **Israel** started working on National standards in 1995
- **United Kingdom:**
  - Computing At School
  - **First country in the European Union** to **mandate computer science classes for all children** between the ages of 5 and 16.

- **18 European Countries**
  France, Spain, Switzerland, Slovakia, Finland, Poland, Portugal, Scandinavia, Italy, Estonia, Bulgaria, Cyprus, Czech Republic, Denmark, Greece, Ireland, Lithuania

- In progress: **US, India, China, Japan, Australia**
  - In the US: CSTA Computing Standards, CSForAll, Code.org

K-12 AI Education

We are not as far along when it comes to AI

- Many countries are trying:
  China, UK, Thailand, Korea, and EU Countries

- The 2017 CSTA Computing Standards contain just two sentences about AI.
  - Both are for the 11-12 grade band. Nothing for younger students.

3B-AP-08	11-12	Describe how artificial intelligence drives many software and physical systems.	Algorithms & Programming	Algorithms	Communicating
3B-AP-09	11-12	Implement an artificial intelligence algorithm to play a game against a human opponent or solve a problem.	Algorithms & Programming	Algorithms	Creating

Five Big Ideas in AI

- Organizing framework for the K-12 guidelines.
- 5 big ideas are enough to cover the richness of the field, but small enough to be manageable by teachers.
- CSTA experience shows 5 is a good number.
- Not necessarily the way AI practitioners view their field, but appropriate for the needs of the K-12 audience.

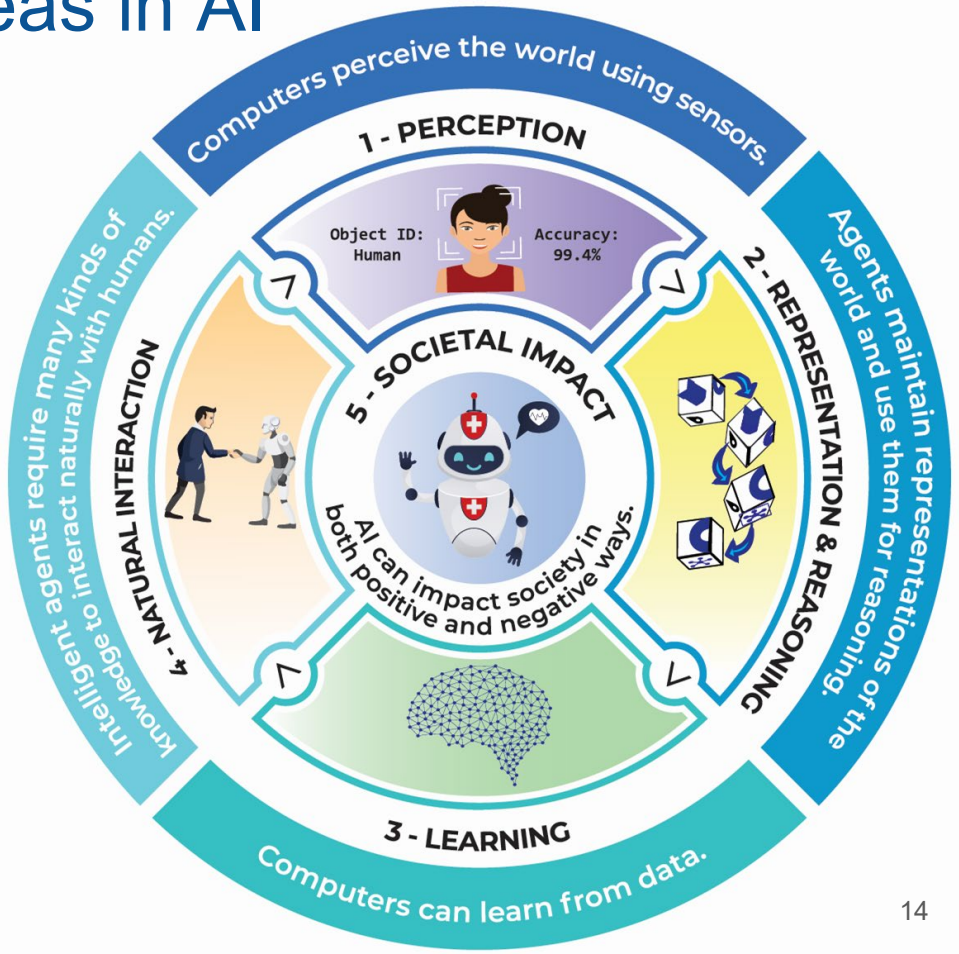

Big Idea #1: Perception

Computers perceive the world using sensors.

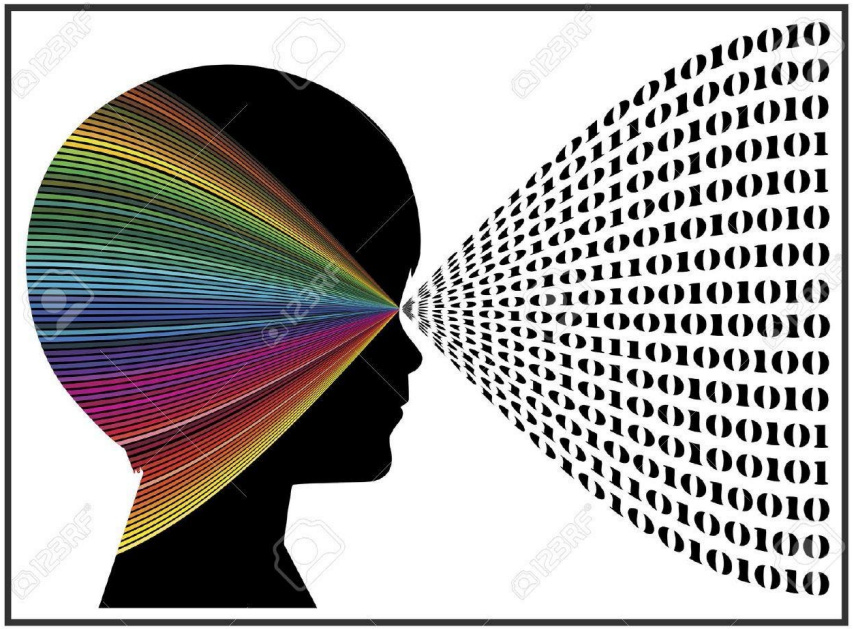

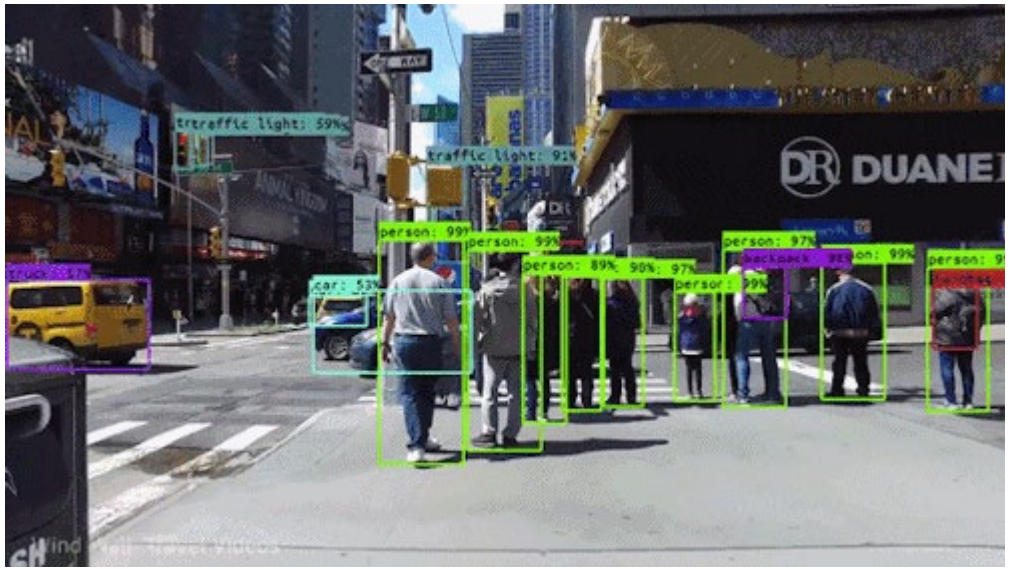

Perception is the extraction of *meaning* from sensory signals.

Big Idea #1 – Examples of Perception

- Speech recognition
- Computer vision:
  - object recognition
  - face recognition
  - license plate readers
  - scene understanding
- Other forms of perception e.g., music recognition, or interpreting sonar, radar, or LIDAR data

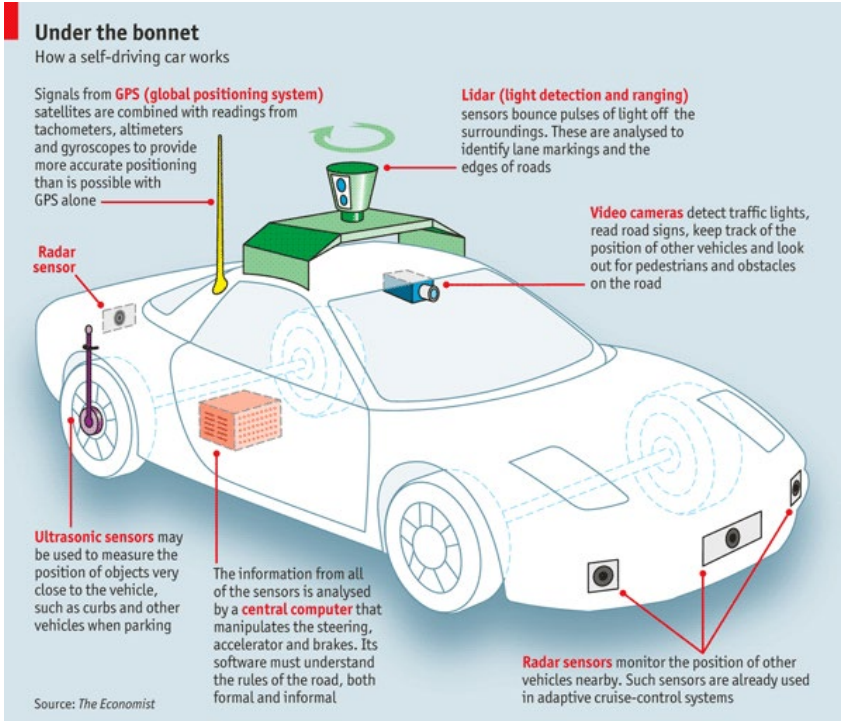

Big Idea #1 – Major Concepts

- Human senses vs. computer sensors
- Going from sensing to perception
- Types of perception: vision, speech recognition etc.
- How perception works: algorithms
- Limitations of computer perception
- Intelligent vs. non-intelligent machines

Big Idea #1 – What should students be able to do?

Grades K-2:

- Identify sensors on computers, robots, and intelligent appliances.
- Interact with intelligent agents such as Alexa or Siri.

Grades 6-8:

- Explain how sensor limitations affect computer perception.
- Explain that perception systems may draw on multiple algorithms as well as multiple sensors.
- Build an application using multiple sensors and types of perception (possibly with Scratch plugins, or Calypso).

Grades 3-5:

- Describe how sensor inputs are converted to analog or digital signals.
- Demonstrate a limitation of computer perception.
- Build an application using perception (possibly with Scratch plugins, or Calypso).

Grades 9-12:

- Describe the domain knowledge underlying different forms of computer perception.
- Demonstrate speech recognition difficulty in handling homophones and other types of ambiguity.

Big Idea #2: Representation and Reasoning

Agents maintain representations of the world, and use them for reasoning.

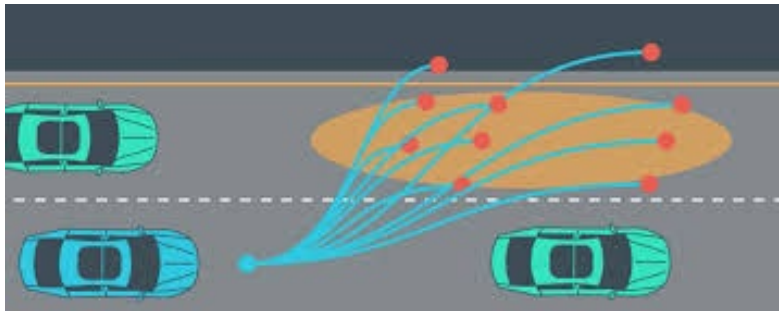

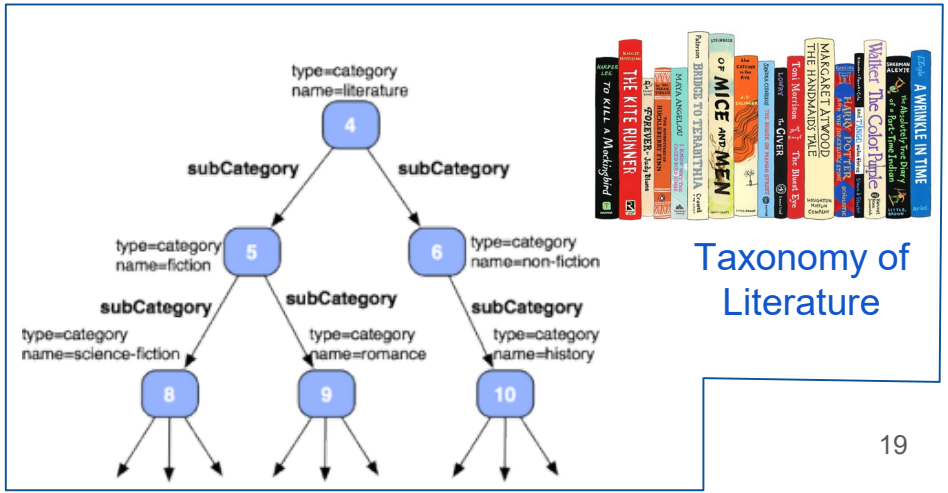

Big Idea #2 – Examples of Representation & Reasoning

- Path planning for self-driving cars
  - Map of the city; description of the scene around car; goal: find best path to the destination
- Internet Search
  - Representations - content of web pages
  - Reasoning - choosing which pages to return in response to a query
- Playing chess, checkers, backgammon, go
  - Representations - board state
  - Reasoning - finding the best move
- Designing a school bus route
  - Representations - list of students and where they live; list of buses and seating capacities
  - Reasoning - assigning students to buses and deciding the order in which the stops are made

Big Idea #2 – Major Concepts

- Types of representations
- Types of reasoning algorithms
- Representation supports reasoning: algorithms operate on representations
- Families of algorithms and the work they do
- Limitations of common reasoning algorithms

Big Idea #2 – What should students be able to do?

Grades K-2:

- Construct a model of something and compare it to the thing being modeled
- Use a decision tree to make a decision

Grades 3-5:

- Create/design a representation of an (animal) classification system using a tree structure.
- Describe how AI representations support reasoning to answer questions

Grades 6-8:

- Design a graph model of their home or locations in their community and apply reasoning to determine the shortest path to key locations on their map
- Create/design a representation of an (animal) classification system using a tree structure.

Grades 9-12:

- Draw a search tree for tic-tac-toe
- Describe the differences between types of search algorithms

Big Idea #3: Learning

Computers can learn from data.

Big Idea #3 – Examples of Machine Learning

Personal experience:

- Training your phone to recognize your face.
- Netflix learning what movies you like.

Industrial applications:

- Training a speech recognition system.
- Training a machine translation system: Google can translate between more than 100 different languages.

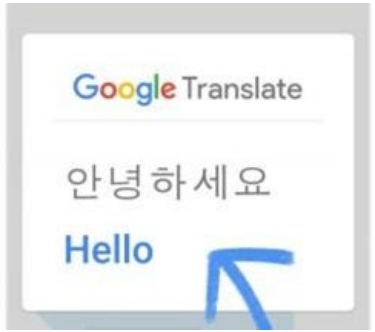

- Image search: training a classifier to recognize pictures of people, animals, vehicles, etc.

Big Idea #3 – Major Concepts in Machine Learning

- What is learning?
- Approaches to machine learning
- Types of learning algorithms
- Fundamentals of neural networks
- Types of neural network architecture
- How training data influences learning
- Limitations of machine learning

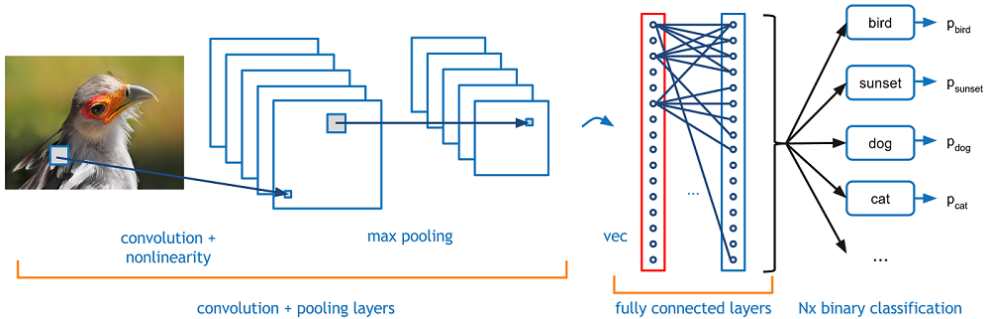

Big Idea #3 – What should students be able to do?

Grades K-2:

- Learn from patterns in data with “unplugged” activities
- Use a classifier that recognizes drawings. Use Google Autodraw or Cognimates Train Doodle to investigate how training sets work to identify images and discuss how the program knows what they are drawing

Grades 3-5:

- Describe and compare the three different machine learning approaches: supervised, unsupervised and reinforcement learning.
- Modify an interactive machine learning project by training its model..
- Describe how algorithms and machine learning can exhibit biases.

Grades 6-8:

- Identify bias in a training data set and extend the training set to address the bias
- Hand-simulate the training of a simple neural network

Grades 9-12:

- Train a neural net (1-3 layers) *TensorFlow Playground*
- Trace and experiment with a simple ML algorithm

Big Idea #4: Natural Interaction

Intelligent agents require many types of knowledge to interact naturally with humans.

Humans are among the hardest things for AI agents to understand.

Big Idea #4 – Examples of Natural Interaction

- Intelligent assistants such as Alexa and Siri
- Conversational agents (chatbots)
- Intelligent tutoring systems that provide for adaptive education
- Gesture and facial expression recognition

Pressing questions about the capabilities of AI for natural interaction:

- Can computers exhibit common sense reasoning comparable to people?
- Can a computer ever be conscious or self-aware?

Big Idea #4 – Major Concepts in Natural Interaction

- Natural language understanding
- Affective computing
- Common sense reasoning
- Consciousness and philosophy of mind
- Natural interaction applications
- Human-robot interaction
- Limitations of AI for natural interaction

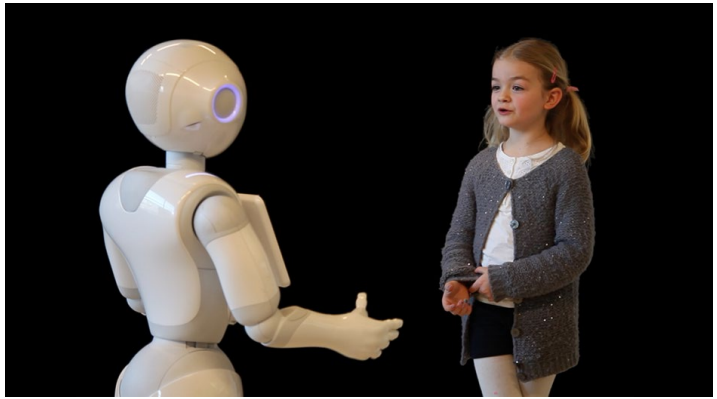

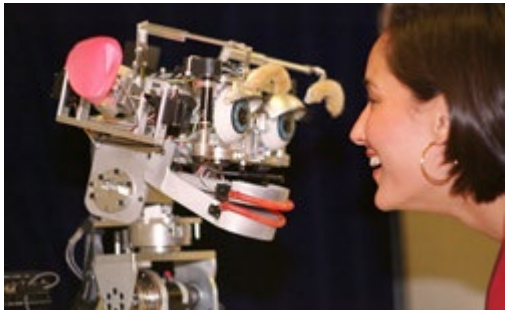

Big Idea #4 – What should students be able to do?

Grades K-2:

- Identify words in stories that have positive and negative connotations.
- Recognize and label facial expressions into appropriate emotions (happiness, sadness, anger) and explain why they are labeled the way they are
- Experiment with software that recognizes emotions in facial expressions

Grades 3-5:

- Identify how humans combine multiple inputs (tone, facial expressions, posture, etc) in order to understand communication.
- Describe some tasks where AI outperforms humans, and tasks where it does not

Grades 6-8:

- Construct a simple chatbot
- Explain and give examples of how language can be ambiguous
- Reason about the nature of intelligence, and identify approaches to determining whether an agent is or is not intelligent.

Grades 9-12:

- Demonstrate how sentence parsers handle ambiguity
- Explore the Google Knowledge Graph
- Identify and debate the issues of AI and consciousness

Big Idea #5: Societal Impact

"Artificial Intelligence can impact society in both positive and negative ways."

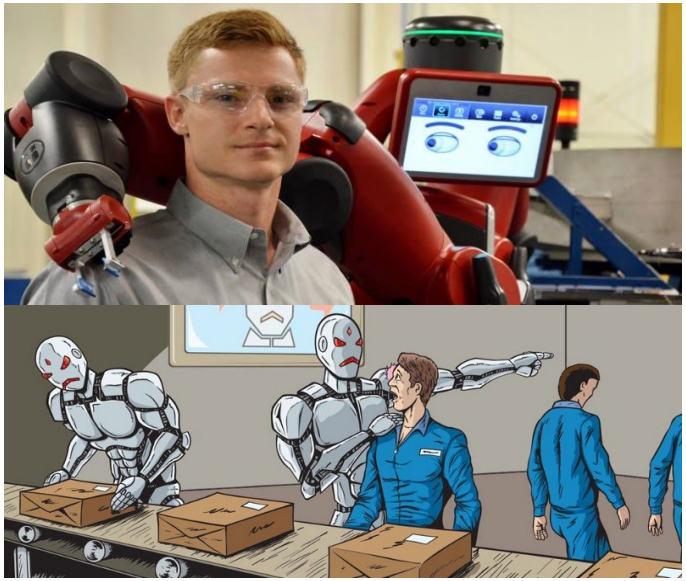

Big Idea #5 – Examples of Societal Impact of AI

- Ethics: What sorts of applications are desirable/permissible?
  - Transparency and accountability of AI systems
  - Competing definitions of “fairness”
  - Values tradeoffs, e.g., privacy vs. security; who should own your data?

- Effects: What are the likely impacts of AI technology on society?
  - Robot servants, rescuers, and companions
  - Economic disruption; changes in the nature of work
  - Effects of unintended biases

Big Idea #5 – Major Concepts around Societal Impact

- AI technologies are changing business, healthcare, education, and government
- Use of AI is an economic driver that makes new services possible and businesses more efficient
- Humans make numerous technical and ethical decisions when developing AI applications
- AI technologies impact communities and people in different ways
- Ethical standards are needed for AI systems that make decisions about people
- AI and robotics will change the way people work, create new jobs, and eliminate some jobs

Big Idea #5 – What should students be able to do?

Grades K-2:

- Identify common AI applications encountered in their daily lives
- Discuss whether common uses of AI technology are a good or bad thing

Grades 3-5:

- Explore how behavior is influenced by bias and how it affects decision making
- Describe ways that AI systems can be designed for inclusivity

Grades 6-8:

- Explain potential sources of bias in AI decision making
- Understand tradeoffs in the design of AI systems and how decisions can have unintended consequences in the function of a system

Grades 9-12:

- Critically explore the positive and negative impacts of an AI system
- Design an AI system to address social issues (or explain how AI could be used to address a social issue)

Five Big Ideas in Artificial Intelligence

1. Perception

Computers perceive the world using sensors. Perception is the process of extracting meaning from sensory signals. Making computers "see" and "hear" well enough for practical use is one of the most significant achievements of AI to date.

2. Representation & Reasoning

Agents maintain representations of the world and use them for reasoning. Representation is one of the fundamental problems of intelligence, both natural and artificial. Computers construct representations using data structures, and these representations support reasoning algorithms that derive new information from what is already known. While AI agents can reason about very complex problems, they do not think the way a human does.

3. Learning

Computers can learn from data. Machine learning is a kind of statistical inference that finds patterns in data. Many areas of AI have progressed significantly in recent years thanks to learning algorithms that create new representations. For the approach to succeed, tremendous amounts of data are required. This "training data" must usually be supplied by people, but is sometimes acquired by the machine itself.

4. Natural Interaction

Humans are among the hardest things for AI agents to understand. Intelligent agents require many kinds of knowlege to interact naturally with humans. Agents must be able to converse in human languages, recognize facial expressions and emotions, and draw upon knowledge of culture and social conventions to infer intentions from observed behavior. Today's AI systems can use language to a limited extent, but lack the general reasoning and conversational capabilities of even a child.

5. Societal Impact

AI can impact society in both positive and negative ways. AI technologies are changing the ways we work, travel, communicate, and care for each other. But we must be mindful of the harms that can potentially occur. For example, biases in the data used to train an AI system could lead to some people being less well served than others. Thus, it is important to discuss the impacts that AI is having on our society and develop criteria for the ethical design and deployment of AI-based systems.

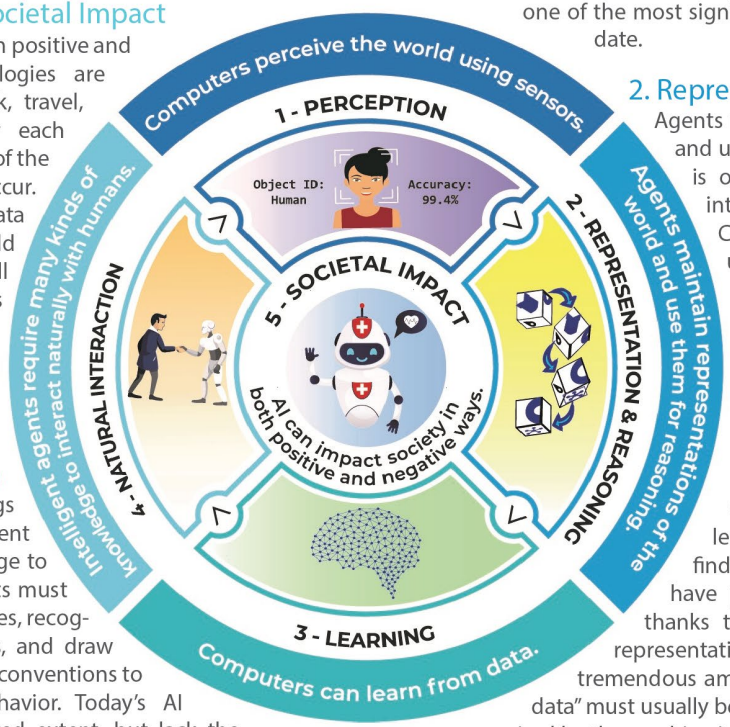

The AI for K-12 Initiative is a joint project of the Association for the Advancement of Artificial Intelligence (AAAI) and the Computer Science Teachers Association (CSTA), funded by National Science Foundation award DRL-1846073

AI4K12.org

11x14 Poster available for free download at AI4K12.org

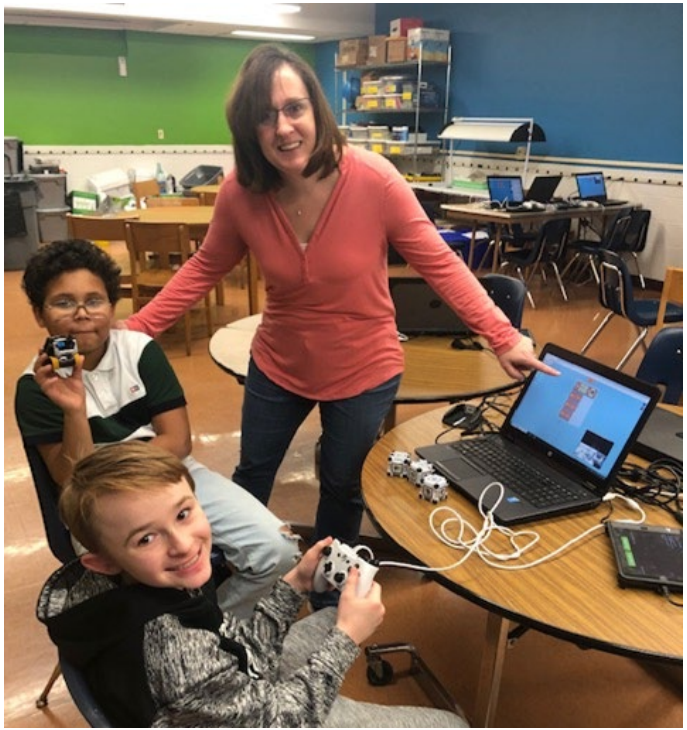

Teaching AI in K-12 Classrooms

Guidelines for supporting K-12 students

1. **Use transparent AI demonstrations** *that help students see what is going on inside the black box: it's not magic!*
2. **Help students build mental models** of *what is happening under the hood in AI applications.*
3. **Encourage students to develop AI applications** *using AI services.*

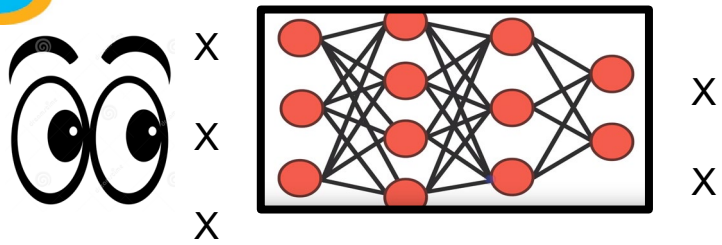

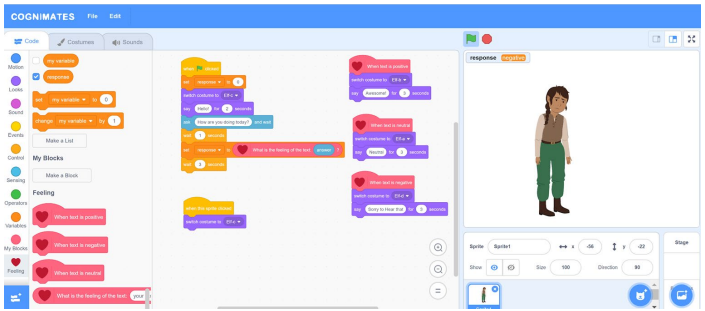

Student Activity Considerations

- **Experiment with AI agents** to investigate their behavior
- **Hand simulate** AI algorithms
- Encourage students to **build their own AI applications**
- **Explore case studies of AI-related societal issues** from multiple perspectives

These activities promote understanding of:

- How AI works
- Limitations of AI
- Systems thinking (AI systems are built from smaller components)
- Sources of bias in AI
- Societal impacts of AI systems

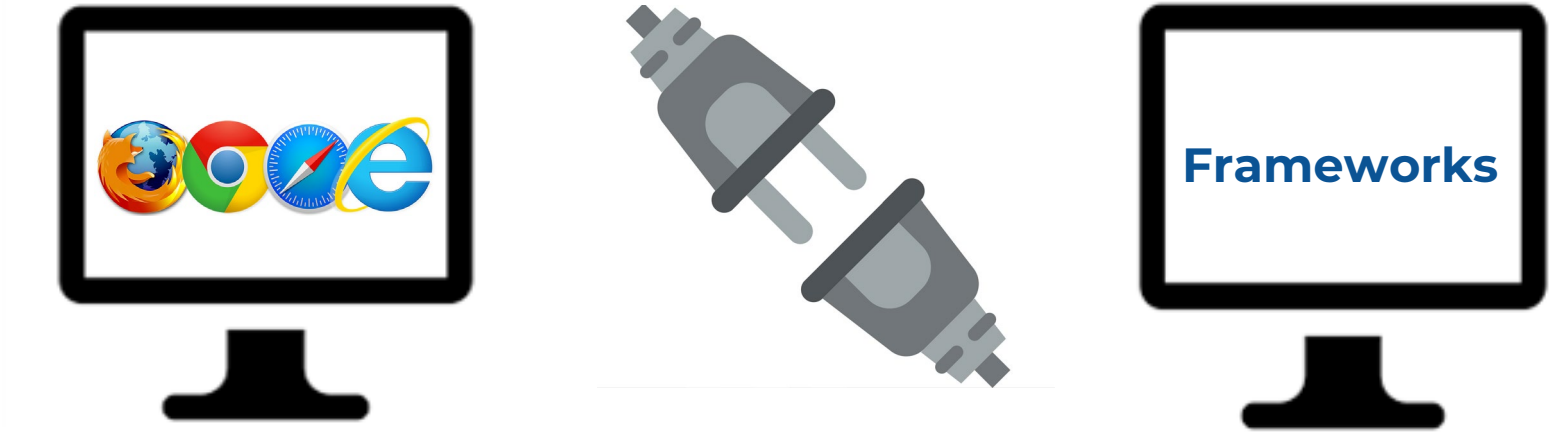

Overview of the Resource Library:
AI Tools & Resources for K-12

Google's Quick, Draw!

https://quickdraw.withgoogle.com/

Can a neural network learn to recognize doodling?

Help teach it by adding your drawings to the world's largest doodlin
data set, shared publicly to help with machine learning research.

Let's Draw!

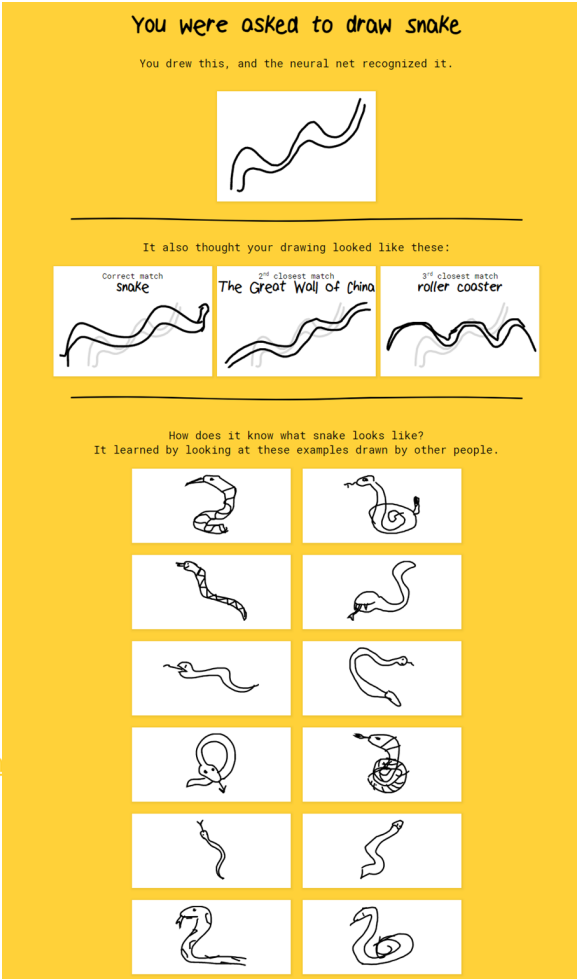

TEACHABLE MACHINE

Built with TensorFlow

- Teach a machine using your camera.
- Live, in the browser.
- No coding required.

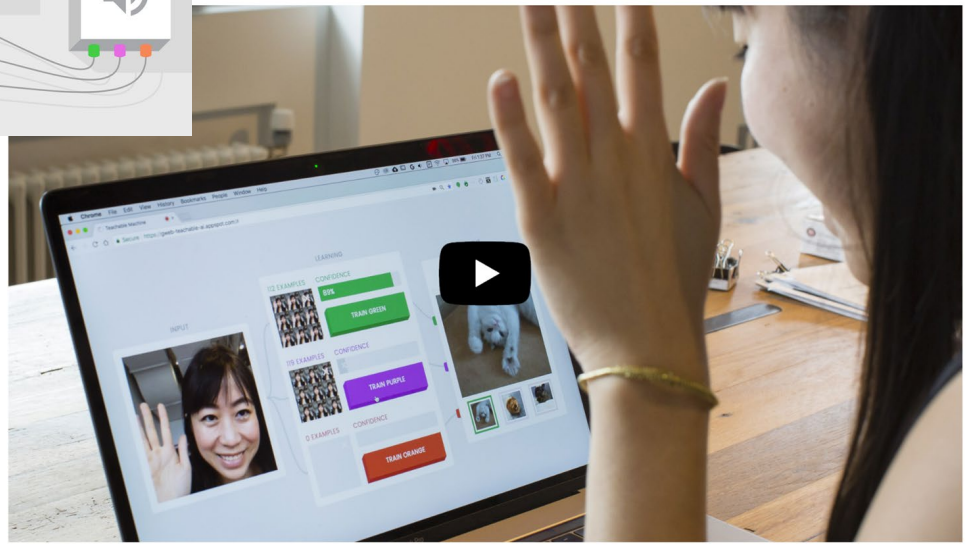

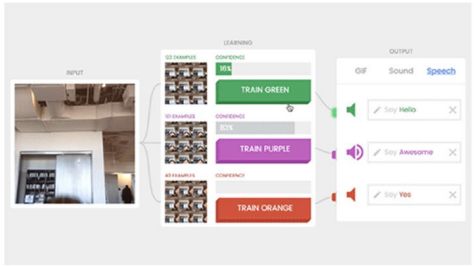

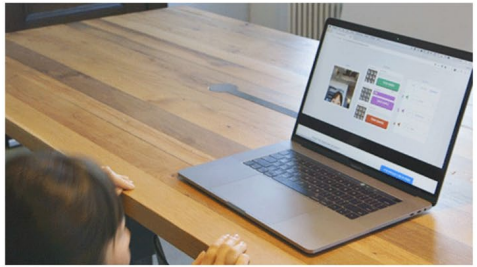

https://experiments.withgoogle.com/teachable-machine

Machine Learning for Kids

https://machinelearningforkids.co.uk

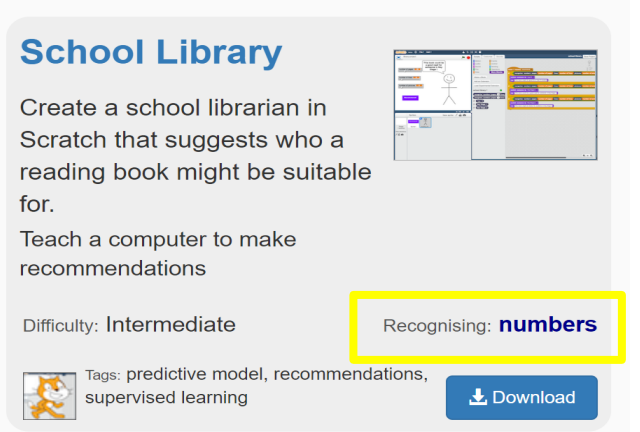

1. Collect examples of things you want to be able to recognise
2. Use the examples to train a computer to be able to recognise them
3. Make a game in Scratch that uses the computer's ability to recognise them

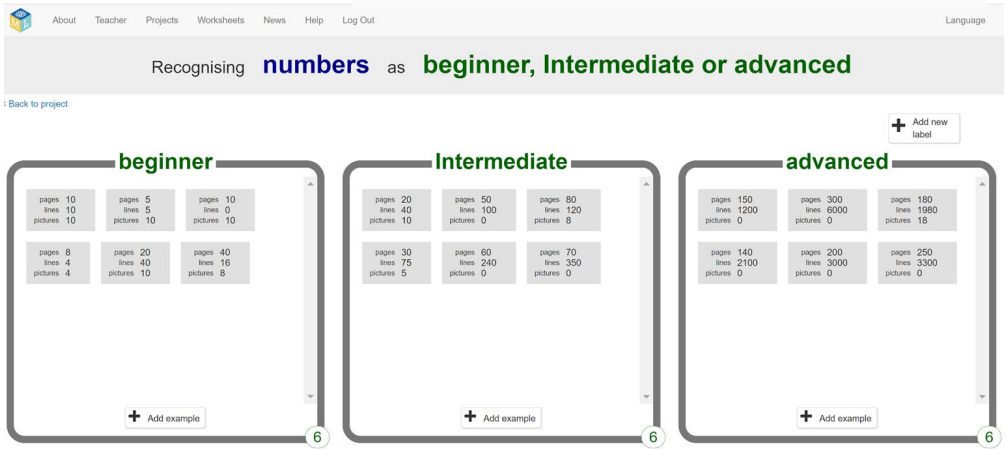

It is being built by Dale Lane using APIs from IBM Watson Developer Cloud.

TensorFlow Playground

https://playground.tensorflow.org

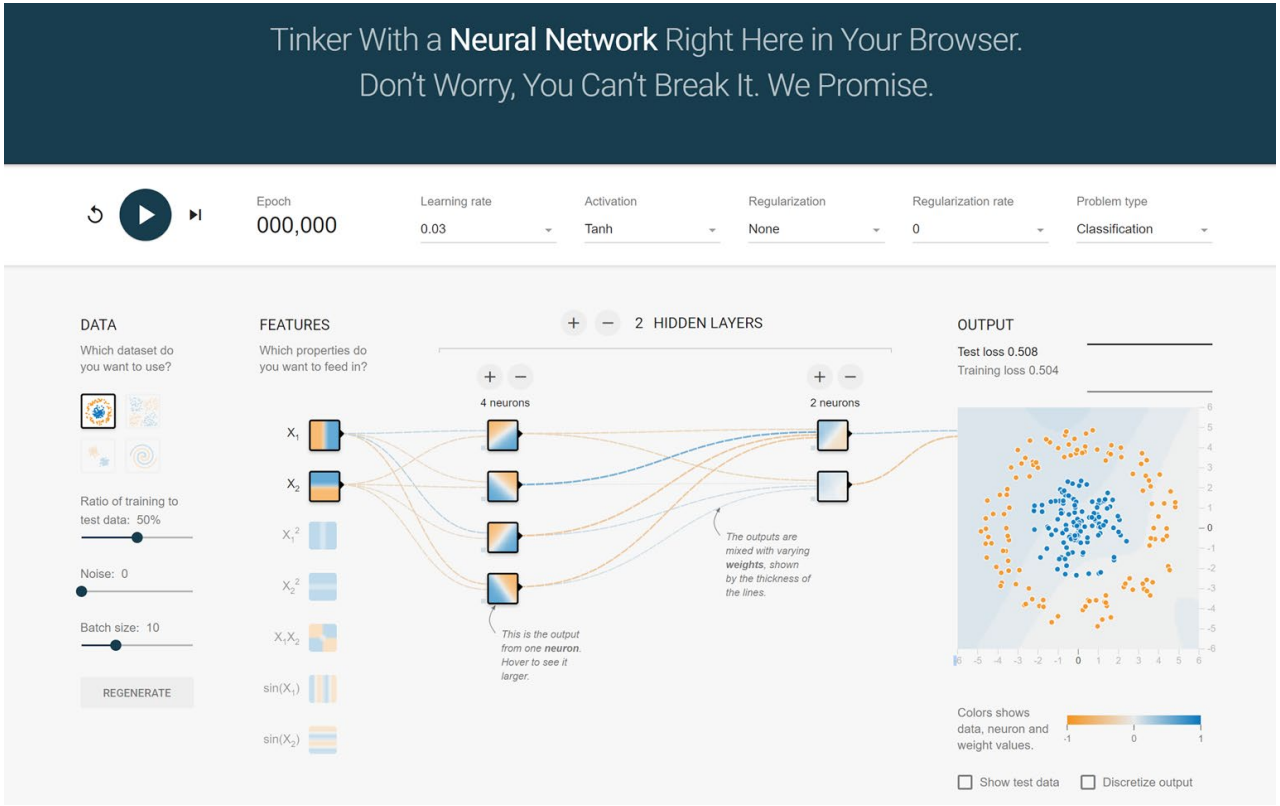

Tutorial: https://cloud.google.com/blog/products/gcp/understanding-neural-networks-with-tensorflow-playground

COGNIMATES

http://cognimates.me

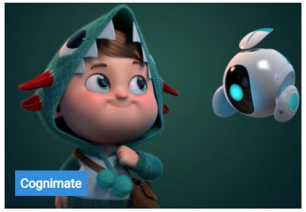

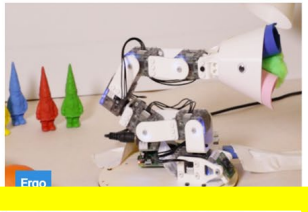

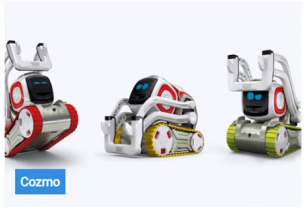

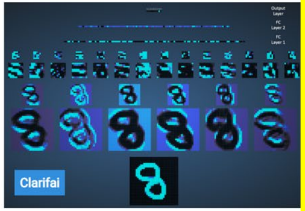

Cognimates offers AI extensions for Scratch, such as:

- speech recognition
- sentiment analysis
- visual pattern detection
- robot control

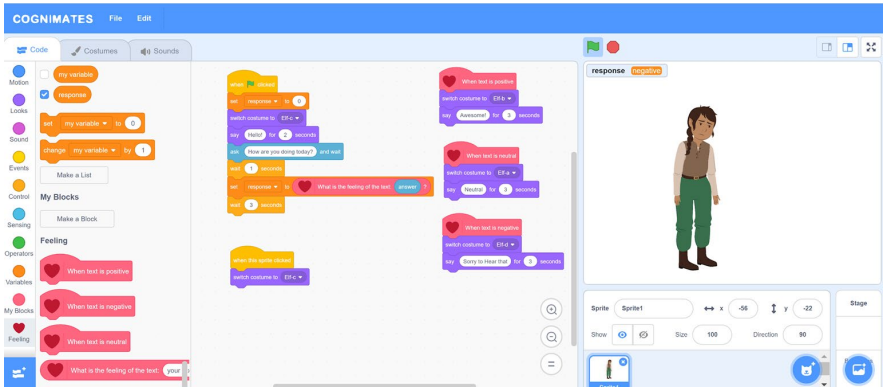

Calypso for Cozmo

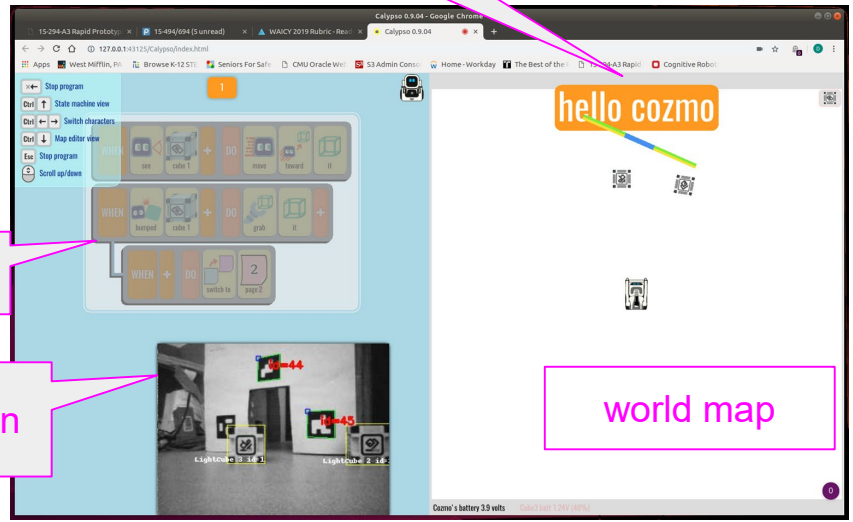

- A robot intelligence framework that Incorporates multiple AI technologies:
  - Computer vision; face recognition
  - Speech recognition and generation
  - Landmark-based navigation
  - Path planning
  - Object manipulation
- Rule-based pattern matching language inspired by Microsoft's Kodu Game Lab
- Teaches computational thinking: "Laws of Calypso", idioms, etc.
- Web site: https://Calypso.software

AI + ME (AI and Me)

edu.readyai.org/courses/aime/

“AI+ME” is an online experience intended to provide young learners with the basics of AI. The lesson takes about one hour to complete. This is the first publicly available course introducing students to the “Five Big Ideas in AI” as defined by the AI4K12 Initiative.

Target Audience: Elementary School

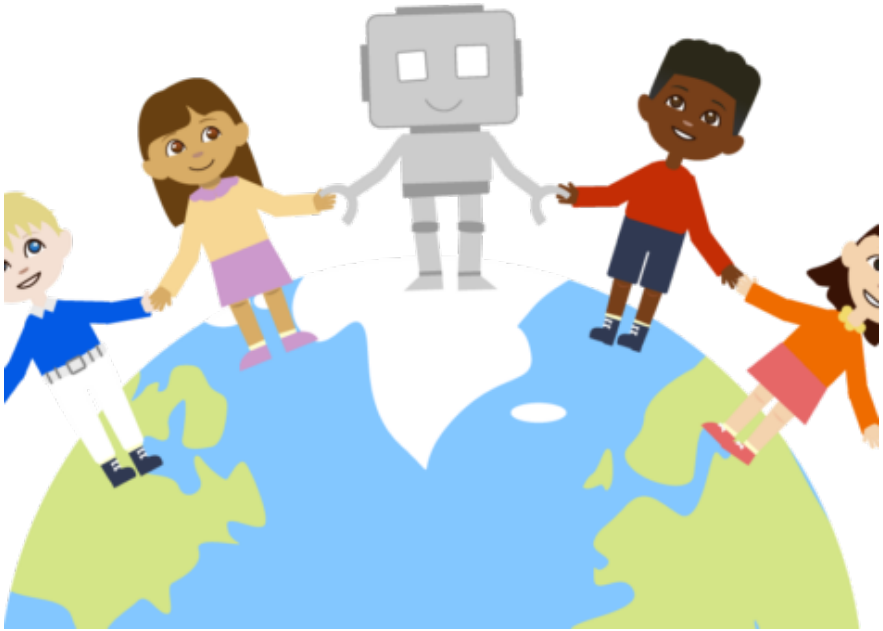

AI4All: Online Student Portal

http://ai-4-all.org/open-learning

The AI4All Open Learning platform will offer a series of online AI courses for high school students. As of summer 2019, the first course is in beta test. This courses focuses on the basics of machine learning.

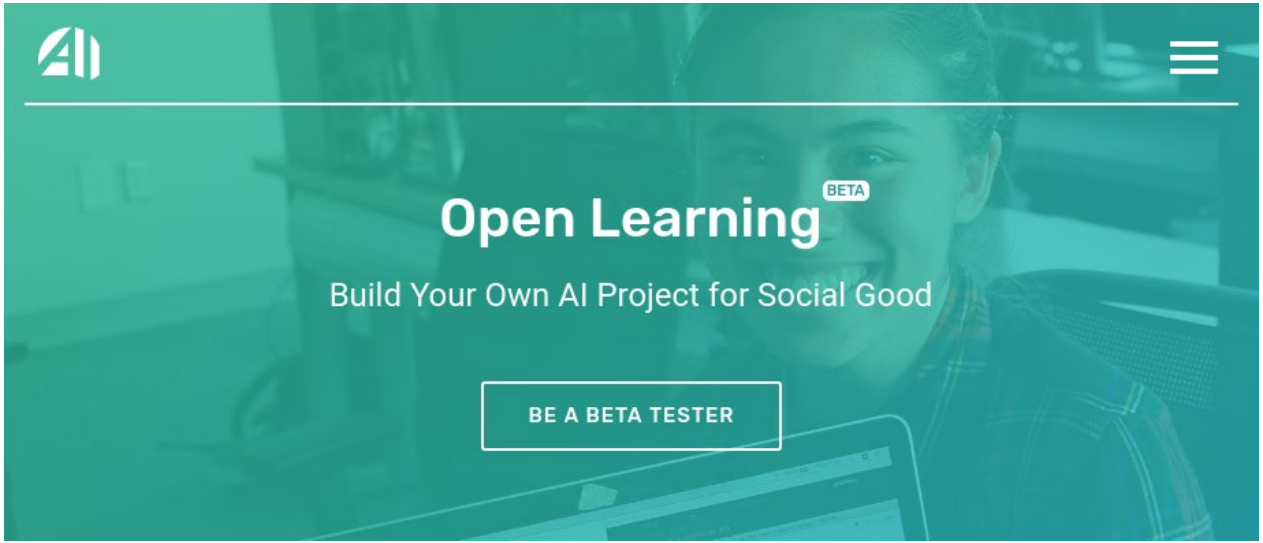

WAICY: World Artificial Intelligence Competition for Youth

2018 WAICY Stats:

5+ Time zones

200+ Students

50 Teams

(**20+** remote participation)

"S.T.E.A.M.-Powered A.I."

- 50/50 Rubric
- Winning Project

C.A.R.E.
Healthcare
Median Age: 15

East to West
Energy Transportation
Median Age: 8

AI Ninja
Home Assistant
Median Age: 7

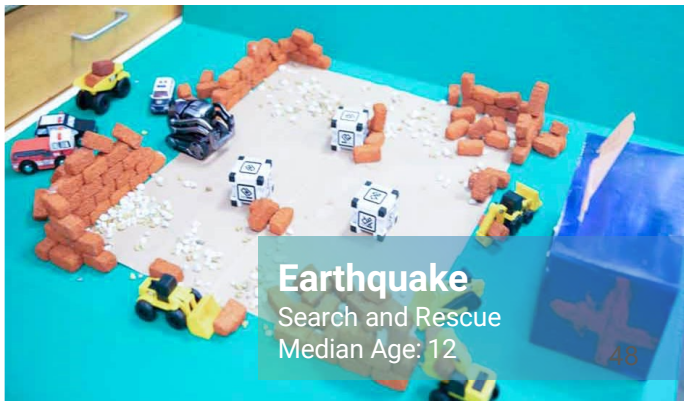

Earthquake
Search and Rescue
Median Age: 12

Professional Development Course

Artificial Intelligence Explorations and Their Practical Use in Schools

www.iste.org/learn/iste-u/artificial-intelligence

Course Dates:

Summer 2019 Session: June 3 - September 13

Enrollment period: ***Now until July 12****

ENROLL NOW

ENROLL A GROUP

Fall 2019 Session: October 14, 2019 - January 27, 2020

Enrollment period: ***Now until October 28****

Course Details:

- Grade Level: 6-12
- Course Length: 30 hours
- Cost: $224 Member / $299 Non-member
- Course Style: asynchronous with instructor

Questions?

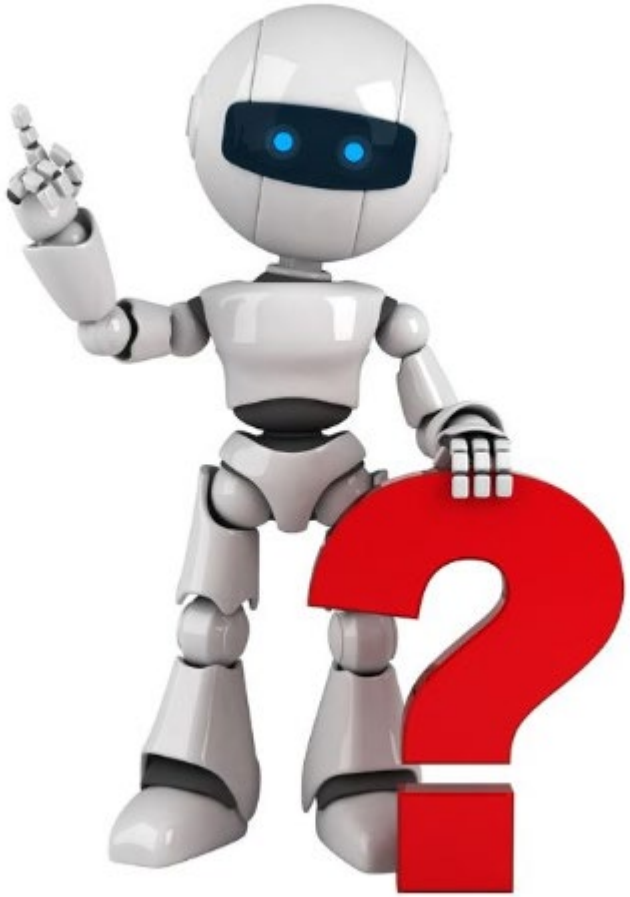

invites you to join us at the

2nd Annual AI for K-12 Symposium:

Teaching AI in K-12

November 8-9th, 2019

Westin Arlington Gateway

Arlington, VA

Part of the AAAI 2019 Fall Symposium Series

It’s time for all of us to think about AI in K-12.

Visit us:
http://AI4K12.org

Join the mailing list:
Send mail to ai4k12@aaai.org

Session Evaluation

Please take a moment to evaluate this session. Your valuable feedback helps make the overall program stronger and ensures we're meeting your learning needs. Evaluations are also used by the conference program committee to provide feedback to presenters and inform future presentations.

To provide feedback and rate the quality of this session, please use the ISTE19 app or locate the session online using the program search at isteconference.org.

Thank you!