

Scalable Privacy-Compliant Virality Prediction on Twitter^{*}

Damian Konrad Kowalczyk^{1,2}[0000-0002-5612-0859]
and Jan Larsen²[0000-0003-1880-1810]

¹ Microsoft Development Center Copenhagen, Kanalvej 7
2800 Kgs. Lyngby, Denmark dakowalc@microsoft.com

² DTU Compute, Matematiktorvet 303B
2800 Kgs. Lyngby, Denmark {[@dank](mailto:dank), [@janla](mailto:janla)}@dtu.dk

Abstract. The digital town hall of Twitter becomes a preferred medium of communication for individuals and organizations across the globe. Some of them reach audiences of millions, while others struggle to get noticed. Given the impact of social media, the question remains more relevant than ever: how to model the dynamics of attention in Twitter. Researchers around the world turn to machine learning to predict the most influential tweets and authors, navigating the volume, velocity, and variety of social big data, with many compromises. In this paper, we revisit content popularity prediction on Twitter. We argue that strict alignment of data acquisition, storage and analysis algorithms is necessary to avoid the common trade-offs between scalability, accuracy and privacy compliance. We propose a new framework for the rapid acquisition of large-scale datasets, high accuracy supervisory signal and multilanguage sentiment prediction while respecting every privacy request applicable. We then apply a novel gradient boosting framework to achieve state-of-the-art results in virality ranking, already before including tweet’s visual or propagation features. Our Gradient Boosted Regression Tree is the first to offer explainable, strong ranking performance on benchmark datasets. Since the analysis focused on features available early, the model is immediately applicable to incoming tweets in 18 languages.

Keywords: Twitter · virality · privacy · sentiment · explainability · scalability · popularity

1 Introduction and motivation

”The role of the social and professional networks in the spread and acceptance of innovations, knowledge, business practices, products, behavior, rumors, and memes, is a much-studied problem in social sciences, marketing and economics. Online environments like Twitter, offer an unprecedented opportunity to track such phenomena.” [2]

^{*} Supported by Microsoft Development Center Copenhagen and the Danish Innovation Fund Case No.5189-00089B

The knowledge discovery process, however, is becoming even more tangled with the arrival of social big data. 700 million tweets have been posted on the day of writing this introduction. The volume, velocity, and variety of mostly unstructured information even from a single social network are evolving at an extremely fast pace. From an engineering and data science perspective, near real-time analysis via online services and algorithms scalable in-memory are required, and demand substantial computational resources. Scientific endeavors to date offer progress toward specific subtasks of social network analysis (SNA) yet data collection and privacy compliance remain among the biggest challenges in extracting knowledge [3]. Arguably the most significant among them is privacy [34]. The social nature of nodes in these networks makes data subjective to many privacy concerns and laws. The new European General Data Protection Regulation (GDPR and ISO/IEC 27001) in force since May 25th, 2018 makes SNA and black-box approaches (like deep neural networks) more difficult to use in business, requiring the results to be retraceable (explainable) on demand [17]. In machine learning, explainable (compliant) real-time analysis is often at odds with predictive accuracy. In social popularity prediction, some of the best results today are achieved using deep neural networks, difficult to interpret [37] or data modalities time-consuming to acquire [12]. Modeling popularity relies on a precise count of responses (subject to privacy requests, i.e., retweets in virality prediction) which exposes them further. Accuracy in such studies depends on processing documents no longer available, while privacy compliance requires removing them. Ensuring accurate and explainable analysis via quality of the data and methods, while respecting user privacy, remain conflicting goals and open research issues individually. In this work we argue that significant advancement in SNA requires avoiding such trade-offs and addressing all the above issues simultaneously. We draw inspiration from multiple disciplines, to challenge state of the art in content virality prediction on Twitter. We propose a framework which to the best of our knowledge, is the first one that satisfies the properties of model preserving and privacy-compliant simultaneously. We use it to train a scalable and explainable model, and are the first to achieve strong [9] ranking performance on benchmark datasets.

2 Related work

2.1 Social big data analysis before GDPR

Social big data has become essential for various distributed services, applications, and systems [31], enabling event detection [10], sentiment analysis [11], popularity prediction [38], natural language processing, finding influential bloggers, personalized recommendation [14], online advertising, viral marketing, opinion leader detection etc. Computational and storage requirements of such applications have led to cloud scale reinvention of data storage and processing technologies. New tools are constantly emerging to replace the conventional non-effective ones, and a hybrid of techniques [20,15] is now a requirement to extract value from the social big data. [35] proposes a solution based on Hadoop technology

and a Naive Bayes classification for sentiment analysis of tweets. The sentiment analysis is performed in MapReduce layer and results stored in distributed NO-SQL data-base. [18] uses Lucene indexing with full-text searching ability on top of Hadoop for spectral clustering, to detect Twitter communities during the Hurricane Sandy disaster. In our work we pursue close alignment of data acquisition and analysis algorithms, with the strict constraints of storage and time, to accommodate both user-generated content (UGC) and privacy requests, arriving at high volume and velocity. Instead of perturbing or anonymizing the data, sensitive or deleted information is permanently eliminated from storage and subsequent analysis.

2.2 Content popularity prediction

Social network influence can be defined as the ability of a user to spread information in the network [32], with the retweet count assumed as a measure of a tweets popularity. One common challenge for content-based popularity prediction is the 140-character constraint imposed by Twitter, making it difficult to identify and extract predictive features [5]. [36] showed that carefully crafted wording of the message could help propagate the tweets better, but there’s much more to UGC than the caption. [19,37] demonstrate social-oriented features were the best performers to predict image popularity on Twitter. [25] utilized textual, visual, and social cues to predict the image popularity on Flickr. [37] proposed a joint-embedding neural network combining the same cues to rival state-of-the-art methods. Recurrent and Deep Neural Networks advance feature extraction from high-dimensional unstructured data (i.e., image attachments), however due to low explainability also introduce a major drawback for critical decision-making processes (with recent advances by [33]). In this study, we prioritize explainable methods in application to structured data. [32,23,7] demonstrate relationships between the number of followers of Twitter users and their influence on information spreading. Ranking users by the number of followers is found to perform similarly to PageRank [23]. [32] models the probability to be retweeted by a power law function. [29] have used an explainable Random Forrest classifier to predict a range of the logarithm of the retweets volume. He demonstrates the predictive value of user features (e.g., count of followers), network features, and the popularity of hashtags included. [4] provide a comparison of learning methods and features, regarding retweet prediction accuracy and feature importance. They find Random Forests to achieve the best performance in binary classification of retweetability and highlight the value of author features: number of times the user is listed by other users, number of followers and the average number of tweets posted per day. [28] uses recursive partitioning trees to achieve 0.682 classification accuracy on a large topical dataset, albeit using features unavailable early (favorites count) or anymore (local publication time) challenging both scalability and reproducibility. [16] investigated the features of tweets contributing to retweetability and is the first to explore the impact of negative sentiment in diffusion of news on Twitter. We follow [16] to consider affect in our model. Substantial gains are seen when including network features extracted from the

content graph formed by retweets, or relationship graph formed by "friendships". The document level subgraphs to inform prediction are often acquired via real-time monitoring of the diffusion process. [39] predicted the popularity of a tweet through the time-series path of its retweets, using a Bayesian probabilistic model. [37] uses preconditioned recurrent neural network to model the temporal diffusion, and shows SOTA ranking performance of 0.366 on benchmark datasets. [1] used temporal evolution patterns to predict the popularity of online UGC. [8] use temporal and structural features to predict the cascades of photo shares on Facebook. [41] model the retweeting cascades as a self-exciting point process. [12] argues that determining the topic of interest of a user based on his past tweets might boost predictive accuracy. [30] studied retweet network propagation trends using conditional random fields, demonstrating gains in accuracy when considering social relationships and retweet history. Access to subgraphs on the author or even document level is however strictly limited by social networks, thus leveraging tweets (early) performance, authors relationships, preferences or retweet history is prohibitive for a scalable, near real-time prediction on a single tweet.

In this study we seek to maximize virality ranking performance. We follow [37] to approach the problem as Poisson regression, and [16] to consider tweet sentiment in prediction. However, in the contrast to prior work, we don't sacrifice scalability or privacy compliance, nor rely on available retweet count for ground truth.

3 Solution overview

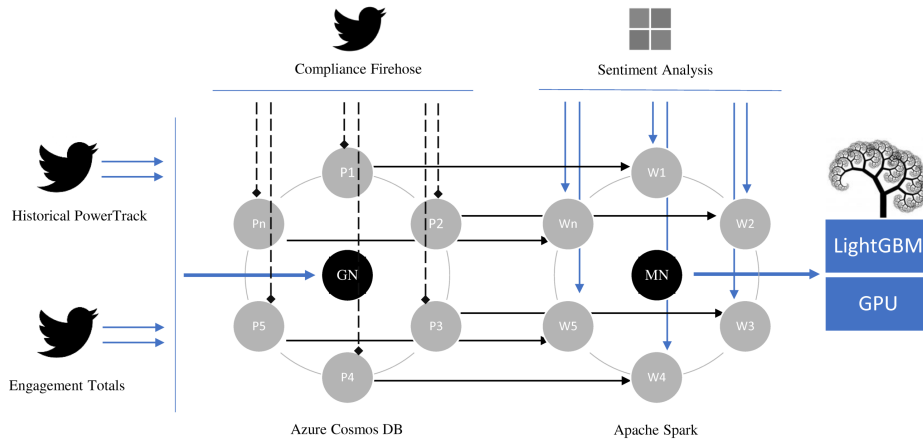


Fig. 1. Solution overview, including data acquisition, storage and analysis components. Cosmos DB gateway node GN orchestrates indexing of Twitter's historical data to partitions P, for simultaneous feature extraction by Spark worker nodes W, before aggregation by master node MN for GPU accelerated predictive analysis.

3.1 Data acquisition

We use Twitters Historical APIs to acquire datasets of tweets for training and validation against other studies. In contrast to sampling Twitters x-hose, predominant in prior work, we apply Twitters PowerTrack search rules, to formulate and collect entire datasets retroactively. The documents are then stored in a globally distributed NO-SQL database, hosted by Microsoft Azure. The data remains online, exposed to every privacy request applicable.

3.2 Privacy compliant storage

Data analyzed in this study is publicly available during collection. Exactly how much of it remains public, changes rapidly afterwards. Account removal, suspension, or deleting of a single tweet render affected content unavailable for analysis in a privacy-compliant way. Users exercise their right to be forgotten at an unprecedented rate. We consume an average of 4,000 of such requests per second via Twitters Compliance Firehose API and apply to our storage simultaneously with analysis. For perspective, the average rate of new tweets published today is 8,000/s. To support this velocity and rapid feature extraction for dependent analysis we choose Azure Cosmos DB as the persistent data store.

3.3 High accuracy labels

In the contrast to prior work, we do not rely on available retweet count for training supervision. Twitter’s Engagement Totals API is called during data collection, to retrieve the number of retweets and favorites ever registered for the tweet (including those deleted shortly after). This enables our data collection effort to focus on unique content only, reducing the document volume required for the task (and proportional compliance responsibility) by more than half, while ensuring 100% accuracy of the supervisory signal.

3.4 Sentiment analysis

To compute document sentiment, we adopt Text Analytics API from Microsoft Cognitive Services [27], a collection of readily consumable ML algorithms in the cloud. At the time of this study, the service supports 18 languages: English, Spanish, Portuguese, French, German, Italian, Dutch, Norwegian, Swedish, Polish, Danish, Finnish, Russian, Greek, Turkish, Arabic, Japanese and Chinese. The service is for-profit and continuously improving (changing) over time, which might challenge reproduction. To address this, we share the score of each document.

3.5 Compute

We conduct an in-memory analysis of entries no longer personally identifiable. This prevents fragmentation of sensitive data outside of the central store exposed to user privacy requests. Instead of anonymizing the datasets, sensitive

or deleted information is eliminated from storage and future analysis as soon as the request from the user is processed by the social media platform. We dedicate an Apache Spark cluster to data preprocessing and analysis. Spark is efficient at iterative computations and is thus well-suited for the development of large-scale machine learning applications [26]. Communication performance between Spark and our privacy-compliant Cosmos DB enables feature extraction at rates exceeding 65,000 tweets per second. The resulting in-memory dataset is then aggregated by the Spark master node, equipped with Tesla K80 GPUs (Graphics Processing Units) for predictive analysis and model tuning. We choose LightGBM framework to train our Gradient Boosted Regression Tree and explain the choice in the following section.

4 Data collection

We use the new framework to build multiple datasets across different time periods for training and evaluation of our models (Table 1)

Table 1. Datasets acquired

Dataset	Timeframe	Months	Language	w/images only	Total	Unique (acquired)	Never retweeted
MBI [6]	2013.02-2013.03	2	English	TRUE	2,724,764	1,319,288	1,042,411
T2015 [37]	2015.11-2016.04	6	English	TRUE	9,025,826	2,804,153	2,106,475
T2016 [37]	2016.10-2015.12	3	English	TRUE	8,469,016	2,736,600	2,088,377
T16-BIO	2015.06-2017.06	12	Multi (18x)	FALSE	27,032,417	14,788,552	12,809,021
T2017-BIO	2017.01-2018-02	14	Multi (18x)	FALSE	19,850,448	9,719,264	8,774,009

Benchmark datasets We acquire three benchmark datasets MBI, T2015 and T2016 (with a total of 6,860,041 unique tweets) to enable comparison with the work of [25,22,6,37]. The datasets match the same filters, as applied before (e.g., timeframe, language or presence of image attachment) yet result in higher volume. We follow [37,6] to split the tweets into 70% training, 10% validation, and 20% test sets respectively.

Twitter 2017 For the general multilanguage model, we have collected 10 million unique tweets and used 9.7M of them for predictive analysis, after applying privacy requests. The dataset has been downsampled from the entire Twitter 2017 volume to 18 languages supported by the sentiment scoring service, then using Twitter PowerTracks sample and bio operators, to manage the volume without sacrificing our models generalization capability over the full year.

4.1 Sentiment score and all-time totals

Retweet counts, favorite counts, and sentiment scores were collected for ca. 30 million unique tweets, simultaneously with applying privacy requests. It is worth noting that 85% of unique tweets acquired had never been retweeted.

4.2 Feature selection

Multiple features have been extracted from the rich Twitter metadata, to capture what is being said (content), by who (author), when (temporal) and how (sentiment). Table 2 describes selected features and their Pearson correlation coefficient with the logarithm of retweet count in T2017-BIO. Only the information available at the time of acquisition or immediately after is considered, to maximize the scalability of the solution. Specifically, we do not consider the early performance of the tweet (i.e., retweet or favorite counts received) or image-based features at this point.

Some authors (e.g., celebrities) receive more attention than others despite low activity. We calculate the two author ratio features in an attempt to isolate such examples. Number of attachments (like hashtags, mentions, URLs, images, symbols and videos) compete for viewers attention with the original 140-character body of the tweet, and their total count is also considered. Finally, we log-transform selected author features (e.g. author’s favorite and listed counts) due to power-law distribution [5].

Table 2. Feature summary

Modality	Feature	Type	Pearson
(A) Author	followersCount	ordinal	0.205920
	friendsCount	ordinal	0.082779
	accountAgeDays	ordinal	0.020379
	statusesCount	ordinal	-0.001455
	actorFavoritesCount	ordinal	0.029914
	actorListedCount	ordinal	0.221067
	actorVerified	categorical	0.202722
(C) Content	attachmentsTotal	ordinal	0.085333
	mentionCount	ordinal	-0.006590
	hashtagsCount	ordinal	0.104335
	mediaCount	ordinal	0.147623
	urlCount	ordinal	0.082549
	isQuote	categorical	0.061915
(L) Language	languageIndex	categorical	0.005199
	sentimentValue	continuous	0.059863
(T) Temporal	postedHour	ordinal	0.016639
	postedDay	ordinal	-0.000963
	postedMonth	ordinal	-0.004129
	postedDayTime	categorical	0.016639
	postedWeekDay	categorical	-0.001002

5 Methodology

We consider the problem of predicting the scale of retweet cascade for a given tweet based on data modalities available immediately after its delivery. The

author features are used together with the content, language, and temporal to predict the number of future retweets. In this study, we assume the future retweet count r of a tweet follows Poisson distribution:

$$P(R = r | \lambda) = \frac{e^{-\lambda} \lambda^{-r}}{r!} \quad (1)$$

where the latent variable $\lambda \in R^+$ defines the mean and variance of the distribution, and maximize the Poisson log-likelihood given a collection of N training tuples of tweets t_i and their retweet counts $r_{gt,i}$

$$\theta^* = \arg \min_{\theta} \frac{1}{N} \sum [r_{gt,i} \ln \lambda(t_i) + \lambda(t_i)] \quad (2)$$

where θ contains all parameters of the proposed model.

5.1 Gradient Boosted Regression Tree

GBRT is a tree ensemble algorithm which builds one regression tree at a time by fitting the residual of the trees that preceded it. With our twice-differentiable loss function, denoted as:

$$L_{\text{Poisson}}(r_{gt}, t) = r_{gt} \ln \lambda(t) + \lambda(t) \quad (3)$$

GBRT minimizes the loss function (regularization term omitted for simplicity):

$$L = \sum_{i=1}^N L_{\text{Poisson}}(r_{gt,i}, F(t_i)) \quad (4)$$

with a function estimation $F(t)$ represented in an additive form:

$$F(t) = \sum_{m=1}^T f_m(t) \quad (5)$$

where each $F_m(t)$ is a regression tree and T is the number of trees. GBRT learns these regression trees in an incremental way: at m -stage, fixing the previous $m - 1$ trees when learning the m -th trees. To construct the m -th tree, GBRT minimizes the following loss:

$$L_m = \sum_{t=1}^N L_{\text{Poisson}}(r_{gt,i}, F_{m-1}(t_i) + f_m(t_i)) \quad (6)$$

where $F_{m-1}(t) = \sum_k^{m-1} f_k(t)$.

The optimization problem (6) can be solved by Taylor expansion of the loss function:

$$L_m \approx \bar{L}_m = \sum_{i=0}^N [L_{\text{Poisson}}(r_{gt,i}, F_{m-1}(t_i)) + \nabla_i f_m(t_i) + \frac{\nabla_i^2}{2} f_m^2(t_i)] \quad (7)$$

with the gradient and Hessian defined as:

$$\begin{aligned}\nabla_i &= \frac{\partial L_{\text{Poisson}}(r_{gt,i}, F(t_i))}{\partial F(t_i)} \mid F(t_i) = F_{m-1}(t_i) \\ \nabla_i^2 &= \frac{\partial L_{\text{Poisson}}^2(r_{gt,i}, F(t_i))}{\partial^2 F(t_i)} \mid F(t_i) = F_{m-1}(t_i)\end{aligned}\tag{8}$$

We train our GBRT by minimizing \bar{L}_m which is equivalent to minimizing:

$$\min_{f \in F} \sum_{i=1}^N \frac{\nabla_i^2}{2} (f_m(t_i) + \frac{\nabla_i}{\nabla_i^2})^2\tag{9}$$

This approach is vulnerable to overdispersion and power-law distribution, characterizing the retweet count. In extreme cases where Hessian is nearly zero (9) approaches positive infinity. To safeguard the optimization, we cap each trees weight estimation at 1.5 and follow [5] to use total retweet count as ground-truth after log-transformation:

$$r_{gt} = \ln(r_{total} + 1)\tag{10}$$

5.2 Gradient Boosting Framework

LightGBM [21] implementation of GBDT is chosen for the task, due to distinctive techniques applicable. Experiments on multiple public datasets show that Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) can accelerate the training process by over 20 times while achieving almost the same accuracy [21]. Most of all, LightGBM implements a novel histogram-based algorithm to approximately find the best splits which is highly scalable on GPUs [40]. The framework allows us to explore substantially larger hyperparameter space during cross-validation. Finally, LightGBM offers good accuracy with integer-encoded categorical features by applying [13] to find the optimal split over categories. This often performs better than one-hot encoding and enables treating more features as categorical while avoiding dimensionality explosion.

6 Experiments

We exercise gradient boosted Poisson regression in experiments organized by datasets, to tune and compare our approach against recent state-of-the-art methods, before attempting to generalize the prediction across topics and cultures in the multilingual extended timeframe study.

6.1 Evaluation metrics

We compute the Spearman Rho ranking coefficient, to measure our models ability to rank the content by expected popularity. Interpretation of this coefficient

is domain specific, with guidelines for social/behavioral sciences proposed by [9]. SpearmanR from SciPy version 1.4.0 is used to ensure tie handling. We did not find this concern expressed in prior work. The p-value for all reported Spearman results is $p < 0.001$

Relative and absolute measures of fit: R^2 , and RMSE are chosen for optimization, to penalize large error higher (i.e. when underestimating highly viral content or vice-versa). The mean-absolute-percentage-error (MAPE) is computed due to popularity in previous studies [37], but not considered for tuning. We dispute MAPEs value relative to above when fitting asymmetric, zero-inflated distribution of the dependent variable (like retweet count). It is undefined for the majority of examples (Table 1), which never receive a retweet and penalizes errors for least retweeted higher.

6.2 Validation on benchmark datasets

We begin with evaluation of our multimodal GBRT against previous state-of-the-art methods. For a fair comparison, we use Poisson regression on the joint author, content and temporal features (ACT), before including sentiment (ACTL). Table 4 demonstrates that our proposed model achieves substantially higher ranking performance, compared to other content-based methods, already before considering image and propagation modalities. Using more advanced feature representations, sentiment score and high accuracy ground-truth, we outperform the state-of-the-art by more than 37% on multiple datasets.

Table 3. Method performance on benchmark datasets.

Method	SpearmanR			MAPE		
	MBI	T2015	T2016	MBI	T2015	T2016
McParlene [25,37]	0.188	0.269	0.257	0.093	0.121	0.137
Khosla [22,37]	0.185	0.273	0.254	0.097	0.103	0.124
Cappallo [6,37]	0.189	0.265	0.258	0.089	0.095	0.119
Mazloom [24,37]	0.190	0.287	0.262	0.073	0.097	0.117
Wang [37]	0.229	0.358	0.350	0.057	0.084	0.103
Ours (ACT)	0.322	0.498	0.503	0.247	0.266	0.256
Ours (all)	0.323	0.499	0.504	0.247	0.266	0.255
	R^2			RMSE		
	MBI	T2015	T2016	MBI	T2015	T2016
Ours (ACT)	0.303	0.417	0.391	0.444	0.553	0.555

6.3 Multilingual, extended timeframe experiments

We apply our method to the new T2017-BIO dataset to generalize popularity prediction across languages and time. Tweet $t(A, C, T, L)$ includes content descriptions C, language descriptions L and is rst issued by author A, at the

time T. Table 4 summarizes contributions of these modalities individually and in combination. The baseline model is trained on a single feature, most popular in literature: the count of authors followers, notified about the tweet.

Table 4. Quantitative evaluation of A: actor, C: content, T: temporal, and L: language features. SpearmanR, R squared: higher is better. RMSE, MAPE: lower is better

Features	SpearmanR	R^2	RMSE	MAPE
A	0.310	0.317	0.359	0.133
C	0.211	0.055	0.422	0.160
T	0.062	0.001	0.432	0.171
L	0.164	0.017	0.430	0.167
AC	0.356	0.396	0.337	0.121
AT	0.311	0.316	0.359	0.132
AL	0.324	0.320	0.358	0.130
CT	0.220	0.059	0.421	0.159
CL	0.269	0.076	0.417	0.154
TL	0.170	0.019	0.430	0.166
ATL	0.324	0.320	0.358	0.130
ACT	0.357	0.395	0.338	0.120
ACL	0.369	0.399	0.336	0.119
ACTL	0.369	0.402	0.336	0.118
Baseline	0.180	0.091	0.414	0.160

7 Discussion

When prioritizing social posts by expected popularity, model’s ranking performance might precede metrics of overall fit. Interpretation of Spearman and R^2 metrics is domain specific. For social/behavioral sciences, reaching 0.5 indicates strong correlation [9]. The final study aimed to explore generalizability of our method over an extended time-frame and 18 languages. The relative insignificance of the Temporal modality (Table 4) suggests low correlation between the time of posting and the content popularity, thereby challenging the common intuition, that posting at the time of audiences activity helps propagating the content. We also find that content-based features alone have higher value for expected popularity ranking than the number of followers. How many people like you appears less important than what you have to say.

Non-linear advanced ML algorithms like deep neural networks and gradient boosted decision trees are among the most successful methods used today. The fact is often attributed to the inherent capability of discovering non-linear relationships between groups of features. It was not necessary in our study to compute e.g., all cross-products to rival state-of-the-art, and at times we have noticed a higher cumulative contribution of combined modalities over their individual gains (Table 4). The size of the audience immediately exposed to the tweet, measured as the count of the authors followers, remains the single strongest predictor




Fig. 2. Feature level importance

of tweet popularity when considered in isolation (Figure 2). The number of times an author has been listed by others, followed others or favorited other content are also among significant features, open to interpretation. Number of friends is arguably related to the diversity of content the author is exposed to. We expect the count of tweets favorited over time (i.e. age of account) to differentiate active from passive consumers. Assuming the authors influence is measured by her capacity to spread information in the social network [32], could the diversity of content actively consumed over time maximize authors influence? We propose this hypothesis for computational social science.

8 Conclusions and future work

In this paper, we have studied the problem of predicting tweet popularity under scalability, explainability and privacy compliance constraints. Our method estimates the potential reach of a tweet i.e. size of retweet cascades based on modalities available immediately after document creation. We prove it is possible to rival state-of-the-art results without compromising on explainability, scalability or privacy compliance. Our Gradient Boosted Regression Tree, combining available modalities with sentiment score and high accuracy ground-truth achieves state-of-the-art results on multiple datasets and is the first to achieve strong [9] virality ranking performance. In the final round of experiments, we apply our method to generalize prediction across extended time-frame in 18 languages and explain the contribution of each modality.

Training the final model on NVidia Tesla K80 took 10 minutes. Computing predictions for the 2 million unique tweets in the validation set, took another 45 seconds. Thats over 44,000 tweets scored per second, with a single GPU. Assuming incoming tweets are already vectorized, the ACT model deployed on Tesla K80 can cope with 5 (five) times todays Twitter volume and velocity. [37] take up to 72 additional hours (after data collection) to acquire propagation features

for the prediction. During that time, our model will have predicted popularity for up to 11 billion tweets.

8.1 Applications

Our model is ready for production with immediate application to social media monitoring. The proposed framework is extendable to other data modalities (e.g. visual) and other methods (e.g. deep neural networks). Our privacy compliant storage solution is immediately applicable to data collection and analysis from other social networks exposing privacy signal (e.g. Tumblr and WordPress, with privacy requests available as compliance interactions from DataSift). Our solution to focus analysis on temporary in-memory samples, created ad-hoc for every iteration, from a single central persistent storage to receive compliance requests, is applicable to any social network sourced data. Our solution to rely on dedicated APIs for high accuracy labels, instead of error prone counting or crawling used in prior work, is immediately applicable to Instagram, Tumblr and Facebook Pages. Our explainable GBRT approach is immediately applicable to Instagram and Tumblr.

8.2 Acknowledgements

This project is supported by Microsoft Development Center Copenhagen and the Danish Innovation Fund, Case No. 5189-00089B. We would like to thank Charlotte Mark, Lars Kai Hansen, Joerg Derungs, Petter Stengard and Uffe Kjall. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

References

1. Ahmed, M., Spagna, S., Huici, F.: A Peek into the Future : Predicting the Evolution of Popularity in User Generated Content. In: Proceedings of the sixth ACM international conference on Web search and data mining (2013). <https://doi.org/10.1145/2433396.2433473>
2. Barabasi, A.L., Psfai, M.: Network science. Cambridge University Press, Cambridge (2016), <http://barabasi.com/networksciencebook/>
3. Bello-Orgaz, G., Jung, J.J., Camacho, D.: Social big data: Recent achievements and new challenges. Information Fusion (2016). <https://doi.org/10.1016/j.inffus.2015.08.005>
4. Bunyamin, H., Tunys, T.: A Comparison of Retweet Prediction Approaches: The Superiority of Random Forest Learning Method. TELKOMNIKA (Telecommunication Computing Electronics and Control) **14**(3), 1052 (sep 2016). <https://doi.org/10.12928/telkomnika.v14i3.3150>, <http://www.journal.uad.ac.id/index.php/TELKOMNIKA/article/view/3150>
5. Can, E.F., Oktay, H., Manmatha, R.: Predicting retweet count using visual cues. In: Proceedings of the 22nd ACM international conference on Conference on information & knowledge management - CIKM '13 (2013). <https://doi.org/10.1145/2505515.2507824>
6. Cappallo, S., Mensink, T., Snoek, C.G.: Latent Factors of Visual Popularity Prediction. In: Proceedings of the 5th ACM on International Conference on Multimedia Retrieval - ICMR '15 (2015). <https://doi.org/10.1145/2671188.2749405>
7. Cha, M., Haddadi, H., Benevenuto, F., Gummadi, K.P.: Measuring User Influence in Twitter: The Million Follower Fallacy. In: ICWSM 10 (2010). <https://doi.org/10.1.1.167.192>
8. Cheng, J., Adamic, L.A., Dow, P.A., Kleinberg, J., Leskovec, J.: Can Cascades be Predicted? (mar 2014). <https://doi.org/10.1145/2566486.2567997>, <http://arxiv.org/abs/1403.4608><http://dx.doi.org/10.1145/2566486.2567997>
9. Cohen, J.: Statistical Power Analysis for the Behavioral Sciences. Lawrence Erlbaum Associates (1988)
10. Dong, X., Mavroeidis, D., Calabrese, F., Frossard, P.: Multiscale event detection in social media. Data Mining and Knowledge Discovery (2015). <https://doi.org/10.1007/s10618-015-0421-2>
11. Feldman, R.: Techniques and applications for sentiment analysis. Commun. ACM **56**(4), 82–89 (Apr 2013). <https://doi.org/10.1145/2436256.2436274>, <http://doi.acm.org/10.1145/2436256.2436274>
12. Firdaus, S.N., Ding, C., Sadeghian, A.: Retweet prediction considering user's difference as an author and retweeter. In: Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2016 (2016). <https://doi.org/10.1109/ASONAM.2016.7752337>
13. Fisher, W.D.: On Grouping For Maximum Homogeneity. American Statistical Association Journal (1958), <http://www.csiss.org/SPACE/workshops/2004/SAC/files/fisher.pdf>
14. Gan, M., Jiang, R.: FLOWER: Fusing global and local associations towards personalized social recommendation. Future Generation Computer Systems (2018). <https://doi.org/10.1016/j.future.2017.02.027>
15. Gandomi, A., Haider, M.: Beyond the hype: Big data concepts, methods, and analytics. International Journal of Information Management (2015). <https://doi.org/10.1016/j.ijinfomgt.2014.10.007>

16. Hansen, L.K., Arvidsson, A., Nielsen, F.A., Colleoni, E., Etter, M.: Good friends, bad news - Affect and virality in twitter. In: Communications in Computer and Information Science (2011). https://doi.org/10.1007/978-3-642-22309-9_5
17. Holzinger, A., Biemann, C., Pattichis, C.S., Kell, D.B.: What do we need to build explainable AI systems for the medical domain? (dec 2017), <http://arxiv.org/abs/1712.09923>
18. Huang, Y., Dong, H., Yesha, Y., Zhou, S.: A Scalable System for Community Discovery in Twitter During Hurricane Sandy. In: 2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing. pp. 893–899. IEEE (may 2014). <https://doi.org/10.1109/CCGrid.2014.122>, <http://ieeexplore.ieee.org/document/6846543/>
19. Ishiguro, K., Kimura, A., Takeuchi, K.: Towards automatic image understanding and mining via social curation. In: Proceedings - IEEE International Conference on Data Mining, ICDM (2012). <https://doi.org/10.1109/ICDM.2012.37>
20. Kaisler, S., Armour, F., Espinosa, J.A., Money, W.: Big data: Issues and challenges moving forward. In: Proceedings of the Annual Hawaii International Conference on System Sciences (2013). <https://doi.org/10.1109/HICSS.2013.645>
21. Ke, G., Meng, Q., Wang, T., Chen, W., Ma, W., Liu, T.Y., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: LightGBM: A highly efficient gradient boosting decision tree. Advances in Neural Information Processing Systems (2017). <https://doi.org/10.1046/j.1365-2575.1999.00060.x>
22. Khosla, A., Das Sarma, A., Hamid, R.: What makes an image popular? In: Proceedings of the 23rd international conference on World wide web - WWW '14 (2014). <https://doi.org/10.1145/2566486.2567996>
23. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: Proceedings of the 19th international conference on World wide web - WWW '10 (2010). <https://doi.org/10.1145/1772690.1772751>
24. Mazloom, M., Rietveld, R., Rudinac, S., Worring, M., van Dolen, W.: Multimodal Popularity Prediction of Brand-related Social Media Posts. In: Proceedings of the 2016 ACM on Multimedia Conference - MM '16 (2016). <https://doi.org/10.1145/2964284.2967210>
25. McParlane, P.J., Moshfeghi, Y., Jose, J.M.: "Nobody comes here anymore, it's too crowded"; Predicting Image Popularity on Flickr. Proceedings of International Conference on Multimedia Retrieval - ICMR '14 (2014). <https://doi.org/10.1145/2578726.2578776>
26. Meng, X., Bradley, J., Yavuz, B., Sparks, E., Venkataraman, S., Liu, D., Freeman, J., Tsai, D., Amde, M., Owen, S., Xin, D., Xin, R., Franklin, M.J., Zadeh, R., Zaharia, M., Talwalkar, A.: Mllib: Machine learning in apache spark. J. Mach. Learn. Res. **17**(1), 1235–1241 (Jan 2016), <http://dl.acm.org/citation.cfm?id=2946645.2946679>
27. Microsoft: Cognitive Services APIs reference. <https://westus.dev.cognitive.microsoft.com/docs/services/TextAnalytics.V2.0/operations/56f30ceeda5650db055a3c9> (2017), accessed: 2018-09-05
28. Nesi, P., Pantaleo, G., Paoli, I., Zaza, I.: Assessing the reTweet proneness of tweets: predictive models for retweeting. Multimedia Tools and Applications (2018). <https://doi.org/10.1007/s11042-018-5865-0>
29. Palovics, R., Daroczy, B., Benczur, A.A.: Temporal prediction of retweet count. In: 4th IEEE International Conference on Cognitive Infocommunications, CogInfoCom 2013 - Proceedings (2013). <https://doi.org/10.1109/CogInfoCom.2013.6719254>

30. Peng, H.K., Zhu, J., Piao, D., Yan, R., Zhang, Y.: Retweet modeling using conditional random fields. In: Proceedings - IEEE International Conference on Data Mining, ICDM (2011). <https://doi.org/10.1109/ICDMW.2011.146>
31. Peng, S., Zhou, Y., Cao, L., Yu, S., Niu, J., Jia, W.: Influence analysis in social networks: A survey (2018). <https://doi.org/10.1016/j.jnca.2018.01.005>
32. Pezzoni, F., An, J., Passarella, A., Crowcroft, J., Conti, M.: Why do I retweet it? An information propagation model for microblogs. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2013). https://doi.org/10.1007/978-3-319-03260-3_31
33. Samek, W., Wiegand, T., Müller, K.R.: Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models (aug 2017), <http://arxiv.org/abs/1708.08296>
34. Sapountzi, A., Psannis, K.E.: Social networking data analysis tools & challenges. Future Generation Computer Systems (2018). <https://doi.org/10.1016/j.future.2016.10.019>
35. Sheela, L.J.: A Review of Sentiment Analysis in Twitter Data Using Hadoop. International Journal of Database Theory and Application (2016). <https://doi.org/10.14257/ijdta.2016.9.1.07>
36. Tan, C., Lee, L., Pang, B.: The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 175–185. Association for Computational Linguistics, Baltimore, Maryland (June 2014), <http://www.aclweb.org/anthology/P14-1017>
37. Wang, K., Bansal, M., Frahm, J.M.: Retweet wars: Tweet popularity prediction via dynamic multimodal regression. In: Proceedings - 2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018 (2018). <https://doi.org/10.1109/WACV.2018.00204>
38. Wu, B., Shen, H.: Analyzing and predicting news popularity on Twitter. International Journal of Information Management (2015). <https://doi.org/10.1016/j.ijinfomgt.2015.07.003>
39. Zaman, T.R., Herbrich, R., van Gael, J., Stern, D.: Predicting Information Spreading in Twitter. In: Workshop on Computational Social Science and the Wisdom of Crowds, NIPS 2010 (2010). <https://doi.org/10.1016/j.jclepro.2015.12.007>
40. Zhang, H., Si, S., Hsieh, C.J.: GPU-acceleration for Large-scale Tree Boosting (jun 2017), <http://arxiv.org/abs/1706.08359>
41. Zhao, Q., Erdogdu, M.A., He, H.Y., Rajaraman, A., Leskovec, J.: SEISMIC: A self-exciting point process model for predicting tweet popularity. CoRR [abs/1506.02594](https://arxiv.org/abs/1506.02594) (2015), <http://arxiv.org/abs/1506.02594>