An Ensemble-based Approach to Click-Through Rate Prediction for Promoted Listings at Etsy

Kamelia Aryafar Etsy 117 Adams Street Brooklyn, NY karyafar@etsy.com Devin Guillory Etsy 20 California Street San Francisco, CA dguillory@etsy.com Liangjie Hong Etsy 117 Adams Street Brooklyn, NY lhong@etsy.com

ABSTRACT

Etsy¹ is a global marketplace where people across the world connect to make, buy and sell unique goods. Sellers at Etsy can promote their product listings via advertising campaigns similar to traditional sponsored search ads. Click-Through Rate (CTR) prediction is an integral part of online search advertising systems where it is utilized as an input to auctions which determine the final ranking of promoted listings to a particular user for each query. In this paper, we provide a holistic view of Etsy's promoted listings' CTR prediction system and propose an ensemble learning approach which is based on historical or behavioral signals for older listings as well as content-based features for new listings. We obtain representations from texts and images by utilizing state-of-the-art deep learning techniques and employ multimodal learning to combine these different signals. We compare the system to non-trivial baselines on a large-scale real world dataset from Etsy, demonstrating the effectiveness of the model and strong correlations between offline experiments and online performance. The paper is also the first technical overview to this kind of product in e-commerce context.

CCS CONCEPTS

•Information systems \rightarrow Learning to rank; Evaluation of retrieval results; •Computing methodologies \rightarrow Neural networks; Ensemble methods;

KEYWORDS

Click-Through Rate Prediction, Logistic Regression, Multimodal Learning, Deep Learning

ACM Reference format:

Kamelia Aryafar, Devin Guillory, and Liangjie Hong. 2017. An Ensemblebased Approach to Click-Through Rate Prediction for Promoted Listings at Etsy. In *Proceedings of ADKDD'17, Halifax, NS, Canada, August 14, 2017,* 6 pages.

DOI: 10.1145/3124749.3124758

¹http://www.etsy.com

ADKDD'17, Halifax, NS, Canada

© 2017 ACM. 978-1-4503-5194-2/17/08...\$15.00 DOI: 10.1145/3124749.3124758



Figure 1: Promoted and organic search results for a query layering necklace is shown where promoted listings are presented in the first row and organic results followed.

1 INTRODUCTION

Etsy, founded in 2005, is a global marketplace where people around the world connect to make, buy and sell unique goods: handmade items, vintage goods, and craft supplies. Users come to Etsy to search for and buy listings other users offer for sale. Currently, Etsy has more than 45M items to sell with nearly 2M active sellers and 30M active buyers across the globe². As millions of people search for items on the site every day and increasingly more sellers compete for more customers, Etsy has started to offer *Promoted Listings Program*³ to help sellers get their items in front of these interested shoppers by placing some of their listings higher up in related search results. An example of promoted listings, blending with organic results, are shown in Figure 1.

In a nutshell, promoted listings at Etsy work as follows. A seller, who is willing to participate in the program, would specify a total budget that he/she wants to spend during the whole campaign and Etsy, as the platform, would choose queries or keywords that the campaign runs on and how much bid price for each query. In a simplified setting, for each query q and each relevant promoted listing l, the platform computes $b_{l,q}$, the bidding price of the listing l to the query q, and the expected Click-Through Rate (CTR) $\theta_{l,q}$. The score $b_{l,q} \times \theta_{l,q}$ is used in a generalized second price auction⁴

²https://www.etsy.com/about/

4https://en.wikipedia.org/wiki/Generalized_second-price_auction

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

³https://www.etsy.com/advertising#promoted_listings

ADKDD'17, August 14, 2017, Halifax, NS, Canada

with other qualified listings. The final pricing and the position of a promoted listing is determined by the auction. The whole process is similar to traditional sponsored search with two distinctions:

- Sellers can only specific the overall budget without the control on each bidding price.
- Sellers cannot choose which queries they want to bid on.

For both aspects, Etsy as a platform would operate on seller's behalf. Etsy uses Cost-Per-Click (CPC) model, meaning that the site charges sellers budget when a buyer clicks on the promoted listing. A similar program exists in eBay⁵ while it uses Cost-Per-Action (CPA) model. As Etsy is using CPC model to operate promoted listings, in order to optimize the platform's revenue, it implies that we need more clicks for each promoted listing and each such clicked listing pays more. In other words, we would like to have higher $b_{l,q}$ and $\theta_{l,q}$ for each clicked listing *l* to the query *q*.

In this paper, we discuss the methodologies and systems to drive CTR θ , given a fixed bidding strategy which computes *b*. To our knowledge, our paper is the first study to systematically discuss how a promoted listings system can be built with practical considerations. In particular, we focus on areas to address these research and engineering questions:

- What is the architecture to balance the effectiveness and simplicity of the system, given the current scale of Etsy?
- What are the effective CTR prediction algorithm and features as well as modeling techniques?
- As an e-commerce site where images are ubiquitous and vital to the user experience, what is the strategy to incorporate such information into the algorithm?
- Is there any correlation between offline evaluation metrics and online model performance, such that we can constantly improve our models?

Although some of these questions might have been tackled in prior work (e.g., [2, 8, 10, 12]), we provide a comprehensive view of these issues in this paper. In addition to the holistic view, in terms of modeling techniques, we propose an ensemble learning approach which leverages historical or users' behavioral signals for existing listings as well as content-based features for new listings. In order to learn meaningful representations from features, we utilize feature hashing and deep learning techniques to obtain representations from texts and images and employ a multimodal learning process to combine these different signals. To demonstrate the effectiveness of the system, we compare the proposed system to non-trivial baselines on a large-scale real world dataset from Etsy in the offline setting as well as online experimental results. During this process, we also establish a strong correlation between offline evaluation metrics and online performance ones, which serves as a guidance for future experimental plans.

The paper is organized as follows. We start a brief review of related work in Section 2, followed by the discussion of our methodology in Section 3. We demonstrate the effectiveness of the system in Section 4 and conclude the paper in Section 5.

In this section, we briefly review some related work to online advertising and multimodal modeling with respect to learning to rank.

Ads' CTR Prediction: Several industrial companies shared their practical lessons and methodologies to build large-scale ads systems. Early authors from Microsoft proposed online Bayesian probit regression [6] to tackle the CTR prediction problem. Arguably the first comprehensive review is from Google's ads system [12] where the paper not only talks about different aspects of machine learning algorithms (e.g., models, features and etc.) to improve ads clickthrough-rate modeling but also layouts what kind of components need to be built for a robust industrial system (e.g., monitoring and alerts). The paper also populated Follow-The-Regularized-Leader (FTRL) as a strong ads CTR prediction baseline. In [8], authors from Facebook described their ads system with the novelty of utilizing Gradient Boosted Decision Trees (GBDT) as feature transformers, as well as online model updates and data joining. Twitter's engineers and researchers [10] advanced the field by utilizing both pointwise and pairwise learning to CTR prediction problem, with detailed feature analysis. LinkedIn's authors [2] proposed to utilize ADMM to fit large-scale Bayesian logistic regression and also provided detailed discussion on how to cache model coefficients and features. Chapelle et al. [4] described Yahoo's CTR prediction algorithms with the emphasis on a wide range of techniques to improve the accuracy of the model including subsampling, feature hashing and multi-task formalism.

Multimodal Learning to Rank: There exists many works to utilize multimodal data, especially images, for the scenario of learning to rank. The basic idea behind multimodal learning is to learn different representations from different data types and combine their predictive strengths in another layer (e.g., [3, 9]). Specifically in learning to rank, Lynch et al. [11] utilized deep neural networks to extract images features and learned a pairwise SVM model to rank organic search results to queries. Recently proposed ResNet [7] has been widely treated as a generic image feature generation tool, which is used in the current work.

3 SYSTEMS AND METHODOLOGIES

On high level, our promoted listings system matches a user query with an initial set of promoted listing results based on the textual relevance. The final ranking of the listings is determined by a generalized second price auction mechanism that takes multiple factors such as bids, budgets, pacing and CTR of those items into account. In this paper, we focus on the CTR prediction of the system.

3.1 System Overview

Out CTR prediction system consists of three main components:

- (1) data collection and training instances creation
- (2) model training and deployment
- (3) inference and CTR prediction scores serving

Firstly, real-time events, including which promoted listings are served and how users interact with them, are processed through Kakfa⁶ pipeline to our Hadoop Distributed File System (HDFS) regularly. We then extract listings' information with historical user

⁶https://kafka.apache.org/

² RELATED WORK

⁵http://pages.ebay.com/seller-center/stores/promoted-listings/benefits.html



Figure 2: The high-level CTR prediction system overview, which is described in Section 3.1 and historical and multimodal listings embedding is covered in Section 3.3, model training in Section 3.2, offline evaluations and predictions in Section 4.1 and Section 4.2

behavior data from HDFS. During this process, we assign positive and negative labels to listings where each clicked listing is considered as a positive data instance while each non-clicked listing is considered as negative. The listings information are transformed into a feature vector containing both historical and multimodal embeddings described in Section 3.3. All listings features excluding the image representations are extracted via multiple MapReduce Hadoop jobs. The deep visual semantic features are extracted in a batch mode on a set of single boxes with GPUs using Torch ⁷. The visual features are then transferred to HDFS where they are joined with all other listing features to create a labeled training instance.

We utilize Vowpal Wabbit⁸ to train our CTR prediction models where all training processes happen on a single box. The details about CTR prediction model is described in Section 3.2. The model is then deployed to multiple boxes as an API endpoint for inference. The new model is validated through a set of offline evaluation metrics discussed in Section 4.1 on a holdout dataset of promoted listings search logs.

At inference time a new CTR prediction score is generated for every active listing available for sale. These listings are tagged with the feature representations similar to training side via multiple MapReduce Hadoop jobs to create testing instances. The model is applied on each testing instance via the API endpoint in a MapReduce job where a CTR prediction score is returned. The CTR prediction scores are calibrated to control for CPC by keeping the CTR mean and standard deviation the same across multiple experimental variants. The calibrated CTR prediction are the final input to the auction mechanism deciding the final ranking of promoted listings. Figure 2 provides an overview of this architecture.

3.2 CTR Prediction

In a nutshell, given a query q issued by a user, CTR prediction is to determine how likely a listing l is going to be clicked by the user under the context of q. In other words, the system is to model the conditional probability P(c | q, l) where c is the binary response, click or not.

In practice, we form a feature vector $x_{q,l} \in \mathbb{R}^d$ to represent the overall feature representation of query q and listing l. In later part of the paper, we will discuss how $x_{q,l}$ is constructed in detail. As

logistic regression is a widely-used *de-facto* machine learning model for CTR prediction in industry, due to it's scalability and speed, we also follow the same formalism, modeling $P(c | q, l) = \sigma(\boldsymbol{w} \cdot \boldsymbol{x}_{q,l})$ where $\boldsymbol{w} \in \mathbb{R}^d$ is a coefficient vector and $\sigma(a) = 1/(1 + \exp(-a))$ is a sigmoid function. Many approaches exist to obtain optimal \boldsymbol{w} . Here, we utilize FTRL-Proximal algorithm as the primary learning tool for both single and multimodal feature representations discussed in Section 3.3:

$$\mathbf{w}_{t+1} = \arg\min_{\mathbf{w}} \left(\mathbf{g}_{1:t} \cdot \mathbf{w} + \frac{1}{2} \sum_{s=1}^{t} \sigma_{s} ||\mathbf{w} - \mathbf{w}_{s}||_{2}^{2} + \lambda_{1} ||\mathbf{w}||_{1} \right)$$

where \mathbf{w}_{t+1} is the optimal coefficient vector at iteration t + 1, $\mathbf{g}_{1:t} = \sum_{s=1}^{t} \mathbf{g}_s$ (\mathbf{g}_s is the gradients at iteration s) and σ_s is the learning-rate schedule and λ_1 is the regularization parameter. The detailed description of the algorithm is discussed in [12].

3.3 Feature Representations

The feature representation x plays a vital role in the effectiveness of the model. Here, we focus on the set of listing features that are important to the performance. The listing features used in CTR prediction can be divided into two sets of features: *historical* features based on promoted listing search logs that record how users interact with each listing and *content-based* features which are extracted from the information presented in each listing's page.

Historical Features: This type of features track how a particular listing performs in terms of CTR and other behavior metrics in the past, which is a usually strong baseline to the problem. Here, we describe how historical smoothed CTR works while other types of historical features can be easily extended from this discussion.

For a listing l_i , we assume that click events are random variables, drawn from a Binomial distribution with parameter θ_i , being the probability to be clicked by users. The naïve estimator, which is also the Maximum Likelihood Estimator (MLE), $\hat{\theta}_i = \frac{c_i}{v_i}$ where c_i is the number of clicks for the item *i* and v_i is the number of impressions for the item *i*. This estimator is only reliable for listings with significant enough data. However, this is usually not the case for new listings with much less or zero impressions. This motivates us to use a prior distribution to smooth the estimation, similar to the method described in Section 3.1 of [1]. Essentially, we put a Beta(α, β) prior on θ_i and the mean of the posterior distribution (which is also a Beta distribution due to conjugacy) becomes:

$$\frac{c_i + \alpha}{v_i + \alpha + \beta}$$

where α can be set as global average of clicks and β can be set as global average of impressions (We understand that there exists empirical Bayes method to estimate both parameters from data). We denote this estimator as *smoothed* CTR, which has a smaller variance and is more stable compared to the MLE estimator. In practice, we recompute average number of clicks and impressions for α and β for every time period and compute the exponential smoothing over these numbers to make sure the prior numbers are stable. The smoothed CTR is computed over the training set per day and is used as the only feature for a logistic regression model, denoted as *historical* model discussed in Section 4.2. From our experiments, we found that this model serves a strong baseline and smoothed CTR also contributes significantly in the final predictive model.

⁷http://torch.ch/ ⁸http://hunch.net/~vw/

In addition to historical CTR, we compute similar historical features for other user behaviors, including listing favorites, purchases and etc.

Content-based Features: Etsy listings are composed of text information such as tags, title and description, numerical information such as price and listing id and at least one main image representing the item for sale. The goal here is to learn an effective multimodal feature representation F_l for a listing item l, given its title and tag words, a numerical listing id, a numerical price, and a listing image.

For text, we obtain a feature vector by embedding unigram and bigrams of tags and title in a text vector space, utilizing the hashing trick [14]. As raw text can grow unbounded and most raw features only appear handful of times, the feature hashing technique can help us have a finite and meaningful representation of texts. We denote the text embedding as T_l .

For images, we embed each image into a representation denoted by I_l . Here, I_l is obtained by adopting a transfer learning method [13] where the internal layers of a pre-trained convolutional neural network act as generic feature extractors on other datasets with or without fine tuning depending on the specs of the dataset. The image features are transfer-learned from a pre-trained deep residual neural network (ResNet101 [7]) that is pre-trained on ImageNet [5]. We utilize the last fully connected layer of the residual neural network to obtain a 2048 dimensional representation of listing l in the image space.

The multimodal embedding of the item l is then obtained as $F_l = [T_l, I_l]$, which is a simple concatenation of both text embedding features and the image embedding features. The final multimodal feature representation is used to train the logistic regression (dented as *content-based* model) discussed in Section 4.2.

3.4 Ensemble Learning

As we discussed in 3.3, different features might capture different aspects of listings such that it may not perform well by simply concatenating them together. From our preliminary experiments, we found that historical features and listing embedding features are effective in slightly different ways. Historical features show strong performance for listings with enough data while perform poorly for the ones with little data or no data (behaving like prior distribution in those cases). On the other hand, content-based features demonstrate robust performance for listings with little data. Based on these observations, we propose an ensemble learning approach to combine these different features. To be specific, we train separate CTR prediction models purely based on historical features and content-based features. Then, we combine these individual models with a higher level model, again another logistic regression model. This higher level model determines how to balance signals (in this case, prediction scores) from models based on historical features and models based on content-based features. In practice, we train multiple such higher level models with different regions of data (e.g., listings with enough historical data or not), which will be discussed in Section 4.

4 EXPERIMENTAL SETUP

In this section, we present the results of our CTR prediction system on a number of experiments on real-world large scale datasets from the online marketplace Etsy. We follow the common practice of *Progressive Validation* [12] where for each modeling pipeline we have the training data extracted from [t - 32, t - 2] and t is the current date. The date t - 1's data is used as validation set for model comparisons and parameter tuning. The winning model can be deployed on date t. We can run multiple such pipelines to support different variants for online A/B testing. Therefore, under this framework, we can consistently innovate on models and compare offline and online model performances.

Here, we report a set of experimental results from one date. The dataset for the date consists of more than 35 million Etsy listings. We collect the training data from [t - 32, t - 2] days of historical promoted listings search logs across all users. Each training instance is an active listing with an impression within the past 30 days. We assign positive labels to clicked listings and negative ones to listings with no clicks. Since the dataset is heavily imbalanced with non-clicked impressions dominating, we subsample the negative class where the down-sampled training dataset consists of 26.5 million data instances. The t - 1's validation set consists of 19.5 million instances that are collected over one day of promoted search logs. We report the offline evaluation metrics on this dataset which is in line with the live A/B tests, as discussed in section 4.1.

4.1 Evaluation Metrics

Meaningful offline evaluation metrics that establish correlations with online A/B test results are a key requirement of production quality iterative systems such as CTR prediction systems. Online A/B tests can be expensive in a number of ways, especially: 1) for a particular web service, the overall traffic is fixed for a given period of time and therefore usually cannot afford to run a large number of experiments, and 2) each experiment requires at least several days to weeks, sometimes months to tell the statistical significant difference between the control and the treatment group. Thus, it becomes critical to launch experiments with confidence. In other words, we want to launch models winning experiments with high probability and reduce the cost of running many online A/B tests. In order to achieve this, we would like to seek correlations between offline experimental metrics and online A/B testing metrics such that we could use offline metrics to determine which variant would likely make a successful online A/B candidate. Similar ideas have been also explored in [15].

The overall process to establish the correlation between offline metrics and online metrics is complicated. Here, we describe a simplified version. In short, we have launched a number of experiments with **known** offline effects (e.g., a particular model outperforms/underperforms the baseline production model in a number of offline metrics which are described below). Then, we observed how these models perform in online A/B tests and see among all these offline metrics, which ones are good indicator for online key business metrics (e.g., clicks, revenue and etc.). In general, we not only want to seek a sign-correlation (e.g., an offline win indicates an online win or an offline lose implies an online lose) but also want to have the right magnitude (e.g., single digit AUC win indicates single digit CTR win for A/B tests). It should be noted that through

An Ensemble-based Approach to Click-Through Rate Prediction for Promoted ADst Dbst Dbst at, Ausgust 14, 2017, Halifax, NS, Canada



Figure 3: Offline evaluation metrics: AUC and Average Impression Log Loss on the hold out data (one day of promoted listings search click logs)

multiple such studies, we have established *the Area Under the Receiver Operating Characteristic Curve* (AUC) as a reliable indicator of predictive power of our models.

In our system, we monitor AUC, *Average Impression Log Loss* and *Normalized Cross Entropy* [15] as key metrics, which are widely used in industry as offline evaluation metrics. Here, we discuss these three metrics in details.

AUC: AUC is a good metric for measuring ranking quality without accounting for calibration [8]. Empirically, we have found AUC to be a very reliable measure of relative predictive power of our models. Improvements in AUC (> 1%) have consistently resulted in significant CTR wins in multiple rounds of online A/B tests. Figure 3 shows offline AUC of multiple models. Figure 4 (bottom figure) shows AUC and clicks per request for a production model.

Log Loss: Log loss is the negative log-likelihood of the Bernoulli model and is often used as an indicator of model performance in online advertisement. Minimizing log loss indicates that P(c | q, l) should converge to the expected click rate and result in a better model 3. A model with lower average impression log loss is then preferred. Figure 3 (bottom figure) shows the average impression log loss for multiple variants compared to a baseline model.

Normalized Cross Entropy: This metric is proposed in He et al. [8]. In essence normalized cross entropy is the average impression log loss normalized by the entropy of average empirical CTR of the training dataset. In each training of the model, we compute the average empirical CTR for the training set used by each variant. We then divide the average impression log loss by this empirical training CTR to obtain the normalized cross entropy. The key benefit normalized cross entropy is its robustness to empirical training CTR. Figure 4 (top figure) shows this metric for multiple variants compared to a baseline model.

4.2 Empirical Results

Here we present and discuss results of our offline evaluations on metrics discussed in Section 4.1 for a typical date. Recall that for our ensemble learning described in Section 3.4, we want final higher



Figure 4: Offline evaluation metrics: Normalized cross entropy on the hold out data (one day of promoted listing search click logs) and overlayed AUC with clicks per request per day

level models to learn how to balance different individual models in different data region. Based on some of our preliminary experiments, we found that simply dividing datasets based on the number of impressions are already a good starting point. In particular, in this experiment, the training and test sets are both divided into two partitions: listings with more than or equal to k impressions (denoted as *warm*) and listings with less than k impressions during the training period (denoted as *cold*). We denote the original training and testing dataset as *mixed* datasets. We set k = 30 empirically as the impression breakdown threshold for cold and warm datasets. Here, we only report offline experimental results but online A/B tests share similar results as AUC is a consistent indicator of the offline/online experiments correlation.

Table 1 summarizes the results of our experiments in terms of AUC, average impression log loss and normalized cross entropy. All models are compared to a baseline model: a logistic regression with FTRL trained on *mixed* training dataset using a numerical listing id as the single feature representation, which was deployed as the production model. We report the changes in metrics compared to this baseline. The **historical** model denotes the model trained on the *warm* training dataset with average smoothed CTR of Section 3as the single listing representation. The *cold*, *warm* and *mixed* metrics show the relative performance of this model in comparison with the baseline model. We can observe that the historical model outperforms the baseline in terms of evaluation metrics on *warm*, *cold* and *mixed* testing sets.

The **content-based** model is a model that utilizes a multimodal(text and image) feature representation of listings in the training dataset. The content-based model performs significantly better on the *cold* testing set as illustrated in Table 1. As expected this model performs worse in terms of offline evaluation metrics on *warm* and *mixed* training sets.

The **Ensemble** model takes the previous models predictions (historical model trained on *warm* training set and content-based model trained on *cold* to maximize it's performance) along with smoothed impressions ($\lfloor \log(impression_{count}) \rfloor$) as the input feature. This model performs significantly better in terms of AUC (> 1.9% lift on

ADKDD'17, August 14, 2017, Halifax, NS, Canada

	Historical			С	Content-based			Ensemble		
	mixed	cold	warm	mixed	cold	warm		mixed	cold	warm
AUC (%)										
	+1.56	+1.89	+1.55	-1.57	+6.39	-1.74		+1.95	+8.34	+1.83
$Log Loss (\times 10^3)$										
	-0.016	-0.048	-0.018	+0.311	-0.194	+0.335		-0.092	-0.332	-0.087
Normalized Entropy $(\times 10^3)$										
	-0.29	-1.23	-0.31	+5.67	-5.00	+6.01		-1.68	-8.55	-1.56

Table 1: Changes in AUC (%), average impression log loss $(\times 10^3)$ and normalized cross entropy $(\times 10^3)$ on the dataset is compared to a numerical listing id only model, which was deployed as the production model, across different variants. The historical model utilized historical smoothed CTR for the logistic regression while the content-based model uses the multimodal listing features discussed in Section 3.3. The ensemble model is trained with content-based and historical models scores along with $\lfloor \log(impression_{count}) \rfloor$ as an input to the ensemble. The impressions break threshold for this experiment is set at k = 30 as discussed in Section 4.2.

the mixed testing set, > +8% lift on cold testing set and > +1.8%lift on the warm testing set), log loss and normalized entropy on the mixed testing set. The model also performs significantly better compared to the baseline and previous models on *cold* and *warm* testing sets.

5 CONCLUSION

In this paper we presented an overview of how promoted listings' CTR prediction system works at Etsy. In addition to the holistic view, we proposed an ensemble learning approach to leverage different signals of listings. In order to learn meaningful representations from features, we utilize feature hashing and deep learning techniques to obtain representations from texts and images and employ a multimodal learning process to combine these different signals. To demonstrate the effectiveness of the system, we compare the proposed system to non-trivial baselines on real world data from Etsy in the offline setting, which correlates online experimental results. During this process, we also established a strong correlation between offline evaluation metrics and online performance ones, which serves as a guidance for future experimental plans. This paper serves as the first study to systematically discuss how a promoted listings system can be built with practical considerations.

6 ACKNOWLEDGMENT

We gratefully acknowledge the contributions of the following: Corey Lynch, Alaa Awad, Stephen Murtagh, Rob Forgione, Jason Wain, Andrew Stanton, Aakash Sabharwal and Manju Rajashekhar.

REFERENCES

- Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. 2009. Spatio-temporal Models for Estimating Click-through Rate. In *Proceedings of the 18th International Conference on World Wide Web.* ACM, New York, NY, USA, 21–30.
- [2] Deepak Agarwal, Bo Long, Jonathan Traupman, Doris Xin, and Liang Zhang. 2014. LASER: A Scalable Response Prediction Platform for Online Advertising. In Proceedings of the 7th ACM International Conference on Web Search and Data Mining. ACM, New York, NY, USA, 173–182.
- [3] Galen Andrew, Raman Arora, Jeff A. Bilmes, and Karen Livescu. 2013. Deep Canonical Correlation Analysis. In Proceedings of the 30th International Conference on Machine Learning. 1247–1255.

- [4] Olivier Chapelle, Eren Manavoglu, and Romer Rosales. 2014. Simple and Scalable Response Prediction for Display Advertising. ACM Transactions on Intelligent Systems and Technology (TIST) 5, 4, Article 61 (Dec. 2014), 34 pages.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. 2009. ImageNet: A large-scale hierarchical image database. In 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 248–255.
- [6] Thore Graepel, Joaquin Quiñonero Candela, Thomas Borchert, and Ralf Herbrich. 2010. Web-Scale Bayesian Click-Through rate Prediction for Sponsored Search Advertising in Microsoft's Bing Search Engine. In Proceedings of the 27th International Conference on Machine Learning. 13–20.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 770–778.
- [8] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, and Joaquin Quiñonero Candela. 2014. Practical Lessons from Predicting Clicks on Ads at Facebook. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. ACM, New York, NY, USA, Article 5, 9 pages.
- [9] Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. 2014. Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models. *CoRR* abs/1411.2539 (2014).
- [10] Cheng Li, Yue Lu, Qiaozhu Mei, Dong Wang, and Sandeep Pandey. 2015. Clickthrough Prediction for Advertising in Twitter Timeline. In Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, New York, NY, USA, 1959–1968.
- [11] Corey Lynch, Kamelia Aryafar, and Josh Attenberg. 2016. Images Don'T Lie: Transferring Deep Visual Semantic Features to Large-Scale Multimodal Learning to Rank. In Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, New York, NY, USA, 541–548.
- [12] H. Brendan McMahan, Gary Holt, D. Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, Sharat Chikkerur, Dan Liu, Martin Wattenberg, Arnar Mar Hrafnkelsson, Tom Boulos, and Jeremy Kubica. 2013. Ad Click Prediction: A View from the Trenches. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, New York, NY, USA, 1222–1230.
- [13] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and Transferring Mid-level Image Representations Using Convolutional Neural Networks. In 2014 IEEE Conference on Computer Vision and Pattern Recognition. 1717–1724.
- [14] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. 2009. Feature Hashing for Large Scale Multitask Learning. In Proceedings of the 26th International Conference on Machine Learning. ACM, New York, NY, USA, 1113–1120.
- [15] Jeonghee Yi, Ye Chen, Jie Li, Swaraj Sett, and Tak W. Yan. 2013. Predictive Model Performance: Offline and Online Evaluations. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, New York, NY, USA, 1294–1302.