

Disentangled Self-Attentive Neural Networks for Click-Through Rate Prediction

Yichen Xu^{1,*}, Yanqiao Zhu^{2,3,*}, Feng Yu⁴, Qiang Liu^{2,3}, and Shu Wu^{2,3,5,†}

¹School of Computer Science, Beijing University of Posts and Telecommunications

²Center for Research on Intelligent Perception and Computing, Institute of Automation, Chinese Academy of Sciences

³School of Artificial Intelligence, University of Chinese Academy of Sciences ⁴Alibaba Group

⁵Artificial Intelligence Research, Chinese Academy of Sciences

linyus@bupt.edu.cn, yanqiao.zhu@cripac.ia.ac.cn, yf271406@alibaba-inc.com, {qiang.liu, shu.wu}@nlpr.ia.ac.cn

ABSTRACT

Click-Through Rate (CTR) prediction, whose aim is to predict the probability of whether a user will click on an item, is an essential task for many online applications. Due to the nature of data sparsity and high dimensionality of CTR prediction, a key to making effective prediction is to model high-order feature interaction. An efficient way to do this is to perform inner product of feature embeddings with self-attentive neural networks. To better model complex feature interaction, in this paper we propose a novel Disentangled Self-attentive Network (DESTINE) framework for CTR prediction that explicitly decouples the computation of unary feature importance from pairwise interaction. Specifically, the unary term models the general importance of one feature on all other features, whereas the pairwise interaction term contributes to learning the pure impact for each feature pair. We conduct extensive experiments using two real-world benchmark datasets. The results show that DESTINE not only maintains computational efficiency but achieves consistent improvements over state-of-the-art baselines.

KEYWORDS

Click-through rate prediction; high-order feature interaction; disentangled self-attention

ACM Reference Format:

Yichen Xu, Yanqiao Zhu, Feng Yu, Qiang Liu, and Shu Wu. 2021. Disentangled Self-Attentive Neural Networks for Click-Through Rate Prediction. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, QLD, Australia. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3459637.3482088>

1 INTRODUCTION

Click-Through Rate (CTR) prediction, seeking to predict the probability that a user will interact with a candidate item, is essential for many online applications, such as computational advertising

*The first two authors made equal contribution to this work.

†To whom correspondence should be addressed.

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 the author(s) 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.
CIKM '21, November 1–5, 2021, Virtual Event, QLD, Australia

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8446-9/21/11...\$15.00
<https://doi.org/10.1145/3459637.3482088>




Figure 1: Our proposed disentangled self-attentive networks for CTR prediction that decouple the learning of the pairwise term and the unary term.

[11] and recommender systems [2]. One major challenge of making accurate prediction is that the data used in CTR tasks usually involve numerous categorical features, e.g., categories of ads, users' device models, etc., and thus are high-dimensional and extremely sparse, distinct from continuous numerical features such as images. With such high-dimensional and sparse features as input, one complex model would be inevitably prone to overfitting. Therefore, a successful solution to extracting useful information from these high-dimensional data is to model combinatorial interaction among feature fields based on embedding lookup techniques, also known as *cross features*. For example, for movie CTR prediction, one informative feature based on third-order feature interaction could be {Age, Gender, Genre}, considering that young men tend to prefer action movies. However, it is not possible to enumerate all combinatorial feature interaction due to the exponential complexity [2, 14, 20]. How to automatically model high-order feature interaction thereby attracts a lot of interests.

Recent development in CTR prediction has witnessed a transition from simple linear models [13] to more sophisticated methods that model arbitrary-order interaction among these sparse categorical features to make effective prediction. To model implicit feature interaction, the pioneering work Factorization Machines (FM) [12] proposes to model second-order feature interaction via inner-product of embedding vectors. Following this line, many other methods extend second-order FM to model higher-order interaction [1, 4, 10]. However, these methods suffer from high computational complexity, limiting its practical application in real world.

One widely-used solution to modeling high-order feature interaction is to compute inner product of feature embeddings [7, 16], which resembles self-attentive neural networks [18] in deep learning literature. Specifically, the dot-product attention scores between every feature pairs can be regarded as the importance of each

feature pair. Then, we can compute the second-order feature interaction as weighted sum over feature embeddings. To model arbitrary-order feature interaction, we can stack multiple layers of self-attentive networks with residual connections.

Intuitively, the dot product between each pair of feature embedding vectors encodes *pairwise* semantics of feature interaction, which, however, neglects the modeling of *general* influence of each feature field. To explicitly model such *unary* semantics, we propose to decouple a unary term from the vanilla self-attention network that computes the *general* impact of one certain feature to all other features. We term the resulting framework `DisentangleSelf-attentiveNetwork`, DESTINE for brevity. In particular, to better model feature interaction, DESTINE consists of two independent computational blocks illustrated in Figure 1: a *whitened* pairwise term that models *specific* interaction between two features and a unary term for the *general* influence of one feature to all others.

For the CTR prediction problem, we first embed input features into low-dimensional spaces and then compute high-order feature interaction by stacking multiple disentangled self-attentive layers. Finally, the embeddings resulting from the last interaction layer are used to estimate the click behavior. Extensive experiments on two real-world datasets demonstrate that our proposed DESTINE not only achieves state-of-the-art performance but also retains high computational efficiency. Our code is made publicly available at <https://github.com/CRIPAC-DIG/DESTINE>.

2 THE PROPOSED DESTINE APPROACH

2.1 Problem Definition

Suppose the training dataset $\mathbb{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ contains N samples, where each sample \mathbf{x}_i consists of M fields of users' and items' features and its associated label $y_i \in \{0, 1\}$ represents that user's behavior (e.g., whether to click an item). The problem of click-through rate prediction is to predict \hat{y}_i , given a feature vector \mathbf{x}_i , for accurately estimating whether a user will interact with an item.

2.2 Learning Decoupled Feature Interaction

The proposed DESTINE consists of three key components: (a) the embedding layer, (b) the interaction layer, and (c) the output layer. At first, the input features are fed into the embedding layer, which transforms input features into dense, low-dimensional embedding vectors. Then, these feature embeddings are fed into several stacked interaction layers, which model high-order interaction. After that, we feed the embeddings from the last interaction layer into the output layer to estimate the click behavior.

For each sparse input feature \mathbf{x}_i , we transform it into dense embeddings $\mathbf{e}_i \in \mathbb{R}^d$ via embedding lookup. Once we obtained a compact representation for each feature, we use the scaled dot-product attention scheme to model high-order feature interaction among feature fields. Specifically, we formulate each feature interaction as a (*key*, *value*) pair and learn the importance of each feature interaction by multiplying each feature embedding, such that important key-value pairs get higher attention scores. Formally, we first transform each feature embedding into a new embedding space \mathbb{R}^d as follows

$$\mathbf{q}_m = \mathbf{W}_q \mathbf{e}_m, \quad (1)$$

$$\mathbf{k}_n = \mathbf{W}_k \mathbf{e}_n, \quad (2)$$

where the *query* and *key* transformation are parameterized by two linear transformation matrices $\mathbf{W}_q, \mathbf{W}_k \in \mathbb{R}^{d' \times d}$, respectively.

Then, we compute the correlation $\alpha(\mathbf{e}_m, \mathbf{e}_n)$ between feature m and feature n . Previous work [22] in visual representation learning demonstrates that the importance score of feature m over feature n could be decomposed into two terms: a pairwise term to model pure specific interaction and a unary term to model general impact over all feature fields. We take summation of these two terms:

$$\alpha(\mathbf{e}_m, \mathbf{e}_n) = \alpha_p(\mathbf{e}_m, \mathbf{e}_n) + \alpha_u(\mathbf{e}_m, \mathbf{e}_n). \quad (3)$$

For the pairwise term, we perform whitening [3] on the key and the query vector to model pure interaction among features, which makes the two interacting features less correlated with each other:

$$\alpha_p(\mathbf{e}_m, \mathbf{e}_n) = \sigma \left((\mathbf{q}_m - \boldsymbol{\mu}_q)^\top (\mathbf{k}_n - \boldsymbol{\mu}_k) \right), \quad (4)$$

where $\sigma(\cdot)$ is the softmax function; $\boldsymbol{\mu}_q = \frac{1}{M} \sum_{i=1}^M \mathbf{W}_q \mathbf{e}_i$ and $\boldsymbol{\mu}_k = \frac{1}{M} \sum_{j=1}^M \mathbf{W}_k \mathbf{e}_j$ takes average of the key and the query vectors, respectively.

Regarding the unary term, we introduce another query transformation matrix $\mathbf{W}'_q \in \mathbb{R}^{d' \times d}$ for modeling significant features:

$$\alpha_u(\mathbf{e}_m, \mathbf{e}_n) = \sigma \left((\boldsymbol{\mu}'_q)^\top \mathbf{k}_n \right), \quad (5)$$

where $\boldsymbol{\mu}'_q$ is the mean vector from another transformation to model the general impact on key vectors, i.e. $\boldsymbol{\mu}'_q = \frac{1}{M} \sum_{i=1}^M \mathbf{W}'_q \mathbf{e}_i$.

After computing the attention score for each feature interaction (m, n), we transform each candidate feature to a new embedding space by a *value* transformation, parameterized by a linear transformation matrix $\mathbf{W}_v \in \mathbb{R}^{d' \times d}$, as given in the sequel,

$$\mathbf{v}_k = \mathbf{W}_v \mathbf{e}_k. \quad (6)$$

At last, we update the final representation for feature field m by linearly combining all features.

Extension to multihead self-attention. To allow the model to learn distinct feature interaction in different subspaces, we use multiple attention heads. Specifically, we compute the representation of feature field m under an attention head h by

$$\mathbf{z}_m^{(h)} = \sum_{k=1}^M \alpha^{(h)}(\mathbf{e}_m, \mathbf{e}_k) \cdot \mathbf{v}_k^{(h)}, \quad (7)$$

where $\alpha^{(h)}$ is the attention score computed via Eq. (3). Note that, each attention head h keeps its distinct weight parameters $\mathbf{W}_k^{(h)}$, $\mathbf{W}'_q^{(h)}$, $(\mathbf{W}'_q)^{(h)}$, and $\mathbf{W}_v^{(h)}$.

Then, we obtain \mathbf{z}_m , the overall hidden representation for feature m , by concatenating the representation of all attention heads as

$$\mathbf{z}_m = \left[\mathbf{z}_m^{(1)}; \mathbf{z}_m^{(2)}; \dots; \mathbf{z}_m^{(H)} \right], \quad (8)$$

where H is the number of attention heads. Additionally, following previous work [16, 23], we incorporate raw, individual features (first-order features) via residual connections [5], which is formulated as,

$$\hat{\mathbf{z}}_m = \varphi(\mathbf{z}_m + \mathbf{W}_r \mathbf{e}_m), \quad (9)$$

where $\mathbf{W}_r \in \mathbb{R}^{d' \times H \times d}$ is a linear projection matrix to avoid dimension mismatch and $\varphi(\cdot) = \max(0, \cdot)$ is the ReLU activation function.

With features \hat{z}_m received from the last interaction layer, we concatenate all M features and utilize a simple logistic regression model on top of them to predict user behavior, formulated as

$$\hat{y} = \sigma(\mathbf{w}^\top [\hat{z}_1; \hat{z}_2; \dots; \hat{z}_M] + b), \quad (10)$$

where $\mathbf{w} \in \mathbb{R}^{d'HM}$ is a linear projection vector, b is a bias term, and $\sigma(x) = 1/(1 + e^{-x})$ is the sigmoid function.

At last, we employ the binary cross-entropy function, which is widely-used in CTR prediction models, as the loss function,

$$\mathcal{L} = -\frac{1}{N} \sum_{(y, \hat{y}) \in \mathbb{D}} (y \log \hat{y} + (1 - y) \log(1 - \hat{y})), \quad (11)$$

where y and \hat{y} is the ground truth and the predicted click, respectively. We use gradient descent algorithms to update model weights.

2.3 Complexity Analysis

For the proposed model, we only introduce a new learnable parameter $\mathbf{W}'_q \in \mathbb{R}^{d \times d'}$ for each attention head, leading to additional space complexity of $O(dd'H)$ for each interaction layer. In addition, the time complexity of DESTINE is $O(H(d'dM + d'M + d'M^2 + d'M^2)) = O(MHd'(2M + d + 1))$, compared to that of the original self-attention module of $O(MHd'(M + d))$. Note that d , d' , and H are usually small, demonstrating that our model is memory-efficient and computation-friendly.

3 EXPERIMENTS

In this section, we present empirical analysis to answer the following three questions.

RQ1. Does the proposed DESTINE method outperform existing state-of-the-art baseline methods that model feature interaction for CTR prediction?

RQ2. Many models integrate implicit feature interaction via deep neural networks (DNN); how does the proposed model with implicit feature interaction compare with them?

RQ3. How do other model variants perform compared to the proposed DESTINE?

3.1 Experimental Setup

We use two large-scale datasets Avazu and Criteo for evaluation. **Avazu** contains click logs of mobile ads in 10 days, which is composed of 23 categorical features including domains, categories, connection types, etc. **Criteo** comprises traffic logs of display ads over 7 days. Each sample contains 13 numerical and 26 categorical feature fields. The detailed statistics is summarized in Table 1.

For fair comparison, following most existing studies [9, 16, 24], we randomly sample 80% data as the training set, 10% as the test set, and the remaining 10% as the validation set. We report the performance in terms of two widely-used metrics AUC and logloss. Please kindly note that considering a large scale of user base, performance improvements of AUC at 1%-level are considered as practically significant for industrial deployment [2, 4, 16, 19, 20, 24].

3.2 Performance of Feature Interaction Models

3.2.1 Baselines. Representative baseline CTR prediction models can be grouped into three lines, according to the order of feature

Table 1: Dataset statistics

Dataset	# Instances	# Fields	# Features	Positives
Avazu	40,428,967	23	1,544,488	17%
Criteo	45,840,617	39	998,960	26%

Table 2: Performance (AUC and logloss) and training time (mins) of methods that model feature interaction. The best performance is highlighted in boldface.

Model	Criteo			Avazu		
	AUC	Logloss	Time	AUC	Logloss	Time
LR	0.7820	0.4695	535.2	0.7560	0.3964	342.6
FM	0.7836	0.4700	391.3	0.7706	0.3856	480.2
AFM	0.7938	0.4584	468.3	0.7718	0.3854	130.7
DeepCrossing	0.8009	0.4513	—	0.7643	0.3889	—
CrossNet	0.7907	0.4591	216.7	0.7667	0.3868	56.3
CIN	0.8009	0.4517	219.0	0.7758	0.3829	179.6
HOFM	0.8005	0.4508	696.2	0.7701	0.3854	903.0
AutoInt	0.8061	0.4455	375.9	0.7752	0.3824	112.6
DESTINE	0.8087	0.4425	477.3	0.7831	0.3789	104.9

Note: For fair comparison, the training time of DeepCrossing is not listed since its implementation is based on distributed multi-GPU environments.

interaction they model: (a) first-order method LR [13], (b) second-order methods FM [12] and AFM [21], and (c) higher-order methods DeepCrossing [14], CrossNet [19], CIN [10], HOFM [1], and AutoInt [16]. In this section, we only include counterparts that directly model feature interaction; for full models that involve a DNN component to model implicit interaction, such as Deep&Cross and AutoInt+, the results are presented later in Section 3.3.

3.2.2 Implementation details. We set the number of attention heads to $H = 2$; the hidden dimension of embeddings is set to $d' = 32$ and the size of attentive embeddings is set to $d = 64$. The model is trained using the Adam optimizer [8] with a learning rate of 0.001. We find that the weight for ℓ_2 regularization slightly impacts the model performance, and thus we carefully tuned it in $[5 \times 10^{-3}, 5 \times 10^{-4}, 5 \times 10^{-5}, 5 \times 10^{-6}]$. Moreover, we set the dropout [17] rate to 0.2 to avoid overfitting.

3.2.3 Results and analysis. The overall performance is summarized in Table 2. We also report the total training time (mins) of each model. All baseline performance is referenced from their original papers. The training time of baselines is measured using their official implementations. Overall, from the table, it is evident that the proposed DESTINE achieves the best performance on all datasets. Moreover, DESTINE also enjoys another merit of low computational complexity as it takes relatively low time to finish training.

We also make other observations. Firstly, methods that model higher-order feature interaction generally gain larger performance improvements, which verifies the necessity of modeling high-order interaction in CTR prediction. Secondly, although several high-order methods such as DeepCrossing utilize a feedforward neural network to learn feature interaction, their performance is even inferior to second-order methods, which implies that it is not sufficient to learn useful feature interaction in an implicit way. Our method,

Table 3: Performance of models integrated with DNNs that implicitly model feature interaction. The averaged changes in the last column are relative changes of performance compared to the corresponding base models.

Model	Criteo		Avazu		Avg. Changes	
	AUC	Logloss	AUC	Logloss	AUC	Logloss
Wide&Deep	0.8026	0.4494	0.7749	0.3824	+0.0292	-0.0213
DeepFM	0.8066	0.4449	0.7751	0.3829	+0.0142	-0.0113
Deep&Cross	0.8067	0.4447	0.7731	0.3836	+0.0200	-0.0164
xDeepFM	0.8070	0.4447	0.7770	0.3823	+0.0068	-0.0096
AutoInt+	0.8083	0.4434	0.7774	0.3811	+0.0023	-0.0020
DeepIM	0.8044	0.4472	0.7828	0.3809	+0.0165	-0.0138
AutoCTR	0.8104	0.4413	0.7791	0.3800	-	-
DESTINE+	0.8118	0.4398	0.7851	0.3779	+0.0026	-0.0019

on the contrary, explicitly models useful feature interaction via attentive networks, so that achieves better performance. Thirdly, compared to AutoInt that leverages a vanilla attentive net, our proposed DESTINE disentangles a pairwise and a unary term from computing attention scores, which further boosts performance. In summary, the results validate the effectiveness of DESTINE.

3.3 Performance of Deep Models Integrated

3.3.1 Evaluation protocols. Many existing CTR prediction models alternatively integrate a Deep Neural Network (DNN) component to learn implicit feature interaction. For each DNN layer, we use a simple feedforward network, followed by a batch normalization layer [6]. We stack two DNN layers for all datasets. The resulting representation \tilde{z}_m for feature m is concatenated with \hat{z}_m and is further trained with another linear layer to obtain the final embedding. For investigating whether implicit feature interaction improves the performance, we include the following baselines: Wide&Deep [2], DeepFM [4], Deep&Cross [19], xDeepFM [10], AutoInt+ [16], DeepIM [23], and AutoCTR [15]. Following the naming convention, we term our hybrid model DESTINE+ that jointly learns feature interaction using decoupled attentive networks and DNNs.

3.3.2 Results and analysis. The results are summarized in Table 3. We observe that DESTINE+ achieves new state-of-the-art performance over existing baselines, which demonstrates that integrating with implicit feature interaction with DNN, DESTINE is able to learn feature interaction more effectively. Moreover, from the last two columns, we observe that the average improvements brought by the DNN component are limited. This indicates that the base model DESTINE has already achieved relatively high performance, which once again verifies the efficacy of DESTINE. We also note that the relative improvements on the Criteo dataset are greater than that on the Avazu. This could be explained by that the number of feature fields on Criteo is much more than that on Avazu; in other words, Criteo is “wider” than Avazu, where DNN training is much easier than the proposed disentangled self-attentive model.

3.4 Performance of Model Variants

For diagnosing the proposed scheme, we further investigate four variants of the proposed disentangled self-attention module:




Figure 2: Performance of ablations and other model variants.

- **Pairwise** (DESTINE-P). We only use the pairwise term in Eq. (4).
- **Unary** (DESTINE-U). We only use the unary term in Eq. (5).
- **Multiplication** (DESTINE-M). We use multiplication rather than addition for combining the pairwise and the unary term, as described below,

$$\alpha_d(\mathbf{e}_m, \mathbf{e}_n) = \sigma\left((\mathbf{q}_m - \boldsymbol{\mu}_q)^\top (\mathbf{k}_n - \boldsymbol{\mu}_k)\right) \cdot \sigma\left((\boldsymbol{\mu}'_q)^\top \mathbf{k}_n\right). \quad (12)$$

- **Shared transformation** (DESTINE-S). We consider preserving the shared key transformation \mathbf{W}_k in the unary term, as formulated below,

$$\alpha_s(\mathbf{e}_m, \mathbf{e}_n) = \sigma\left((\mathbf{q}_m - \boldsymbol{\mu}_q)^\top (\mathbf{k}_n - \boldsymbol{\mu}_k)\right) + \sigma\left(\boldsymbol{\mu}_q^\top \mathbf{k}_n\right). \quad (13)$$

The results are summarized in Figure 2. We can see that both the pairwise and the unary term benefits model performance. However, the way to combine the two terms plays an important role. DESTINE-M performs poorly despite its utilization of both terms and is even outperformed by DESTINE-U and DESTINE-P, where only one of the two terms is used. This verifies our justification that coupled gradient in DESTINE-M will deteriorate the performance. The performance DESTINE-S is also inferior to DESTINE, since the shared \mathbf{W}_q transformation still leads to the coupling of gradients during learning. The outstanding performance of DESTINE compared to all other variants justifies the design of our model.

4 CONCLUSION

In this paper, we present a disentangled self-attention network DESTINE for click-through rate prediction, which consists of two terms for pairwise and unary semantics. Specifically, the unary term models the general impact of one feature on all others, whereas the remaining whitened pairwise term models pure feature interaction among each feature pair. Extensive experiments on two real-world datasets demonstrate the effectiveness of the proposed method.

ACKNOWLEDGMENTS

This work is supported by National Key Research and Development Program (2018YFB1402600), National Natural Science Foundation of China (61772528), and Shandong Provincial Key Research and Development Program (2019JZZY010119).

REFERENCES

- [1] Mathieu Blondel, Akinori Fujino, Naonori Ueda, and Masakazu Ishihata. 2016. Higher-Order Factorization Machines. In *NIPS*. 3351–3359.
- [2] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhya, Glen Anderson, Greg S. Corrado, Wei Chai, Mustafa Ispir, Rohan Anil, Zakaria Haque, Lichan Hong, Vihan Jain, Xiaobing Liu, and Hemal Shah. 2016. Wide & Deep Learning for Recommender Systems. In *DLRS@RecSys*. ACM, 7–10.
- [3] Jerome H. Friedman. 1987. Exploratory Projection Pursuit. *J. Amer. Statist. Assoc.* 82, 397 (March 1987), 249–266.
- [4] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *IJCAI*. 1725–1731.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. IEEE, 770–778.
- [6] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*. 448–456.
- [7] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *ICDM*. IEEE, 197–206.
- [8] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [9] Zekun Li, Zeyu Cui, Shu Wu, Xiaoyu Zhang, and Liang Wang. 2019. Fi-GNN: Modeling Feature Interactions via Graph Neural Networks for CTR Prediction. In *CIKM*. ACM, 539–548.
- [10] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xDeepFM: Combining Explicit and Implicit Feature Interactions for Recommender Systems. In *KDD*. ACM, 1754–1763.
- [11] Qiang Liu, Feng Yu, Shu Wu, and Liang Wang. 2015. A Convolutional Click Prediction Model. In *CIKM*. ACM, 1743–1746.
- [12] Steffen Rendle. 2010. Factorization Machines. In *ICDM*. IEEE, 995–1000.
- [13] Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting Clicks: Estimating the Click-Through Rate for New Ads. In *WWW*. ACM, 521–529.
- [14] Ying Shan, T. Ryan Hoens, Jian Jiao, Haijing Wang, Dong Yu, and J C Mao. 2016. Deep Crossing: Web-Scale Modeling without Manually Crafted Combinatorial Features. In *KDD*. ACM, 255–262.
- [15] Qingquan Song, Dehua Cheng, Hanning Zhou, Jiyan Yang, Yuandong Tian, and Xia Hu. 2020. Towards Automated Neural Interaction Discovery for Click-Through Rate Prediction. In *KDD*. ACM, 945–955.
- [16] Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. 2019. AutoInt: Automatic Feature Interaction Learning via Self-Attentive Neural Networks. In *CIKM*. ACM, 1161–1170.
- [17] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks From Overfitting. *JMLR* 15, 1 (2014), 1929–1958.
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Uszkoreit Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *NIPS*. 5998–6008.
- [19] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. 2017. Deep & Cross Network for Ad Click Predictions. In *AdKDD@KDD*. ACM, 12:1–12:7.
- [20] Shu Wu, Feng Yu, Xueli Yu, Qiang Liu, Liang Wang, Tieniu Tan, Jie Shao, and Fan Huang. 2020. TFNet: Multi-Semantic Feature Interaction for CTR Prediction. In *SIGIR*. ACM, 1885–1888.
- [21] Jun Xiao, Hao Ye, Xiangnan He, Hanwang Zhang, Fei Wu, and Tat-Seng Chua. 2017. Attentional Factorization Machines: Learning the Weight of Feature Interactions via Attention Networks. In *IJCAI*. 3119–3125.
- [22] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. 2020. Disentangled Non-Local Neural Networks. In *ECCV*. Springer, 191–207.
- [23] Feng Yu, Zhaocheng Liu, Qiang Liu, Haoli Zhang, Shu Wu, and Liang Wang. 2020. Deep Interaction Machine: A Simple but Effective Model for High-order Feature Interactions. In *CIKM*. ACM, 2285–2288.
- [24] Jieming Zhu, Jinyang Liu, Shuai Yang, Qi Zhang, and Xiuqiang He. 2020. FuxiCTR: An Open Benchmark for Click-Through Rate Prediction. *arXiv.org* (Sept. 2020). arXiv:2009.05794v1 [cs.LG]