Follow the Prophet: Accurate Online Conversion Rate Prediction in the Face of Delayed Feedback

Haoming Li^{1,2}, Feiyang Pan^{1,2}, Xiang Ao^{*1,2}, Zhao Yang^{1,2},
Min Lu³, Junwei Pan³, Dapeng Liu³, Lei Xiao³, Qing He^{1,2}

¹ Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS),
Institute of Computing Technology, CAS, Beijing 100190, China

² University of Chinese Academy of Sciences, Beijing 100049, China

³ Tencent, Shenzhen, China

{lihaoming19s,panfeiyang,aoxiang,yangzhao20s,heqing}@ict.ac.cn,{alfredolv,jonaspan,rocliu,shawnxiao}@tencent.com

ABSTRACT

The delayed feedback problem is one of the imperative challenges in online advertising, which is caused by the highly diversified feedback delay of a conversion varying from a few minutes to several days. It is hard to design an appropriate online learning system under these non-identical delay for different types of ads and users. In this paper, we propose to tackle the delayed feedback problem in online advertising by “Following the Prophet” (FTP for short). The key insight is that, if the feedback came instantly for all the logged samples, we could get a model without delayed feedback, namely the “prophet”. Although the prophet cannot be obtained during online learning, we show that we could predict the prophet’s predictions by an aggregation policy on top of a set of multi-task predictions, where each task captures the feedback patterns of different periods. We propose the objective and optimization approach for the policy, and use the logged data to imitate the prophet. Extensive experiments on three real-world advertising datasets show that our method outperforms the previous state-of-the-art baselines.

CCS CONCEPTS

• Computing methodologies → Machine learning.

KEYWORDS

online advertising, conversion rate prediction, delayed feedback

ACM Reference Format:

Haoming Li, Feiyang Pan, Xiang Ao, Zhao Yang, Min Lu, Junwei Pan, Dapeng Liu, Lei Xiao, Qing He. 2021. Follow the Prophet: Accurate Online Conversion Rate Prediction in the Face of Delayed Feedback. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463045>

*Correspondence to: Xiang Ao <aoxiang@ict.ac.cn>.

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.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3463045>

1 INTRODUCTION

Conversion rate (CVR) prediction is one of the most essential tasks in online advertising, which predicts the probability of whether a user will accomplish certain goals after viewing an ad, such as order placement [1, 15, 17, 20]. For such a task, the delayed feedback problem is one of the imperative challenges. After an ad is shown to or clicked by a user, the system can receive the conversion feedback with a delay ranging from a few minutes to several days, if there is one [10]. For those ad impressions that do not have feedback yet, it is hard for the system to tell whether they will convert in the future or not. In this case, labelling these samples as negative may introduce bias to CVR prediction models. A customary practice is to wait for a period before assigning negative labels [5, 9, 14], but it makes the model hard to catch up with the most recent data.

In online advertising, since there are always a number of samples whose labels are undetermined, it is difficult to use a standard objective function (e.g., binary cross-entropy) to train the CVR prediction model. So, we consider two questions: 1) How to design an appropriate learning objective in online advertising under delayed feedback, and 2) How to learn a CVR predictor that optimizes such objective.

In this paper, we propose a general methodology for the delayed feedback problem, named “Follow the Prophet” (FTP). Our insight is, we should first construct a prophet that knows how to use the data properly so as to define the learning objective, and then learn the actual model by imitating the prophet.

The “prophet” is a model trained over an idealized dataset, which is constructed to eliminate the issue of delayed feedback. We show that such a model as well as the constructed dataset can be easily obtained using a subset of historical data with complete feedback.

Second, we propose to use multiple tasks and objective functions to characterize the conversion feedback under different delayed periods. In such case, the “prophet” could be the optimization goal for cherry-picking the proper task for each sample. Hence, our objective is to aggregate these tasks to imitate the idealized prophet.

Finally, to accomplish that, we leverage an aggregation policy on top of the predictions of each task to yield an appropriate predictions for each sample. This aggregation policy is optimized in a well-designed training pipeline to minimize the error with an artificial prophet. This ultimately achieves our objective, “behaving like a prophet”.

Compared to previous works, our method does not rely on specific assumptions, neither the delay distribution nor the true labels

distribution. Instead, it is optimized in an objective-oriented way, by “following the prophet.” Our solution is naturally suitable for online training and easy to deploy in real-world advertising systems.

We conduct extensive experiments on three real-world advertising conversion rate prediction datasets, including two public datasets and one large-scale in-house dataset. We use a streaming simulation for evaluation to reflect the performance of each method in the online advertising system. Experimental results show that our method outperforms the other state-of-the-arts in previous works.

2 RELATED WORKS

Conversion prediction plays a critical role in online advertising [1, 15, 17, 20], and there are several works on resolving the feedback delay issue. Chapelle [5] proposed to apply survival time analysis with a delayed feedback model (DFM) to estimate the delay under an assumption of exponential distribution. Yoshikawa and Imai [26] extended DFM to a non-parametric model (NoDeF). Yasui et al. [25] regarded the delayed feedback as a data shift and proposed to use an importance weighting approach (FSIW) to handle different distribution between test data and observed data.

One of the disadvantages of these methods is that they are hard to apply in gradient-descent-based online training [2, 16, 23]. Ktena et al. [14] proposed to use a “fake negative” data pipeline in online streaming: all samples are assigned with a fake negative label without delay. If the system receives a conversion feedback later, they just *duplicate* that sample with a positive label. To estimate conversion rate under this form of data, they proposed to use positive-unlabeled loss (PU) [6, 13, 21], fake negative weighted loss (FNW) and fake negative calibration approach (FNC) [14].

3 PROBLEM FORMULATION

First, a CVR prediction problem without considering the delayed feedback can be formulated as probabilistic predictions of binary classification over a dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$, where each instance has input features (including ad and user attributes) x and a binary label y indicating whether there is a conversion [15, 19].

Next, we formulate an online CVR prediction problem with delayed feedback by adding a few notations. During online learning, at time τ , the system maintains a dataset:

$$\mathcal{D}_\tau = \{(s_i, t_i, x_i, y_{i,\tau-s_i})\}_{i=1}^{N_\tau} \quad (1)$$

where each instance has s the timestamp when it is initially logged, and t the timestamp when receiving its conversion feedback. $t := \infty$ if there is no feedback until τ , and the observed label $y_{\tau-s}$ with a period of $d = \tau - s$ is defined as:

$$y_d = y_{\tau-s} = \begin{cases} 1, & \text{if there is a conversion before } \tau, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Moreover, we define $d_{\max}$ as the maximum feedback delay, i.e., an instance initialized at time s will be considered having no conversion if there is no feedback until $s + d_{\max}$:

$$y_{d'} \equiv y_{d_{\max}}, \forall d' > d_{\max} \quad (3)$$

Finally, the conversion rate can be defined as:

$$\Pr(\text{conversion}|x) := \Pr(y_{d_{\max}} = 1|x) \quad (4)$$

4 FOLLOW THE PROPHET

The outline of our model and data pipeline for FTP is shown in Fig. 1. In the following sections we will illustrate three essential components in FTP: the prophet that can give prophecies, the multiple tasks for capturing conversion feedback patterns in different periods, and the policy optimized towards the prophet to aggregate these tasks.

4.1 Construct the prophet

As we describe in the introduction, a “prophet” is an ideal model that can eliminate the issue of delayed feedback. The issue is that, when a new instance is logged, we only know its features x , but does not have the feedback to determine the label.

So formally, at any time τ , the $y_{\tau-s}$ in logging data $\mathcal{D}_\tau$ is not the label of CVR prediction problem (which need to be $y_{d_{\max}}$ in Eq. 4).

An ideal CVR prediction model could be trained if we could get a set of ideal data with all the feedback, which would be:

$$\mathcal{D}_\tau^* = \{(s_i, t_i, x_i, y_{d_{\max}})\}_{i=1}^{N_\tau} \quad (5)$$

We define that, at any time τ :

A *prophet* is a well-trained CVR model based on $\mathcal{D}_\tau^*$.

A *prophecy* is the prediction of the *prophet*.

Obviously, we cannot obtain $\mathcal{D}_\tau^*$ if the time now is τ . However, according to Eq. 3, on a subset of $\mathcal{D}_\tau$ with the instances subject to $s < \tau - d_{\max}$, we have:

$$\mathcal{D}_{\tau, d_{\max}} := \mathcal{D}_\tau|_{s_i < \tau - d_{\max}} \equiv \mathcal{D}_{\tau - d_{\max}}^* \quad (6)$$

$\mathcal{D}_{\tau, d_{\max}}$ contains instances with complete feedback until τ , since by the definition of $d_{\max}$, all the feedback in this subset should already be received. Therefore, it is possible to construct a prophet until $\tau - d_{\max}$ with logging data subset $\mathcal{D}_{\tau, d_{\max}}$.

4.2 Multi-task learning for delayed feedback

Recall that, without considering the delayed feedback, we can use the standard binary cross-entropy objective function to optimize a single CVR prediction model. However, when under delayed feedback, this kind of optimization approach fails to properly handle samples without feedback yet.

Therefore, instead of optimize the CVR prediction directly, we propose to set up multiple tasks and objective functions to characterize different delayed periods:

$$\text{Task } k: \quad \hat{y}_k = f_k(x) = \Pr(y_{d_k} = 1|x), \quad k = 1, 2, \dots, K$$

We assume that $0 < d_1 < d_2 < \dots < d_K = d_{\max}$.

Similar to Eq. 6, for each task k , the available logging subset is:

$$\mathcal{D}_{\tau, d_k} := \mathcal{D}_\tau|_{s_i < \tau - d_k} \quad (7)$$

The loss function for each task k is the binary cross-entropy

$$\mathcal{L}_k = -y_{d_k} \log(f_k(x)) - (1 - y_{d_k}) \log(1 - f_k(x)) \quad (8)$$

By limited the training data to a subset subject to $s < \tau - d_k$, the label y_{d_k} is always available at time τ . So, all the samples are entirely labeled during training. This is one of the major differences between our method and previous works.

Another advantage is that it is easy to build data pipelines for online training, as shown in Fig. 1(a). Each task can fetch new

Figure 1: Design of the FTP Model and Data Pipeline.

available training data continuously from its own data pipeline, and perform incremental updates.

Various forms of DNN-based multi-task networks can be applied to the model [3, 22], including the widely used shared-bottom structure [4], cross-lingual parameter sharing [7], MoE [8, 11, 24], MMoE [18] and MT-FwFM [20], etc.

4.3 Aggregation policy to imitate the prophet

Now, we introduce the aggregation policy to imitate the prophet, which combines these tasks to make the final prediction.

Our aggregation policy is a model that outputs a set of weights (by a softmax activation) for the tasks, denoted by $(g_1(x), \dots, g_K(x))$ for a given sample with input x . While making prediction (c.f. Fig. 1(b)), we use the weighted average of the tasks, i.e.,

$$\hat{y}_{\text{FTP}} = \sum_{k=1}^K g_k(x) f_k(x) \quad (9)$$

4.3.1 Policy loss. The optimization approach for the policy model is exactly as the name of our method: follow the prophet, as shown in Fig. 1(c).

Since we have all the models, we can have the predictions of the tasks $\hat{y}_k = f_k(x)$, $k = 1, \dots, K$, as well as the prediction of the prophet $y^* = f^*(x)$. Then, we can find the task k^* which has the closest prediction to the prophecy,

$$k^* = \arg \min_k |\hat{y}^* - \hat{y}_k|. \quad (10)$$

We want the policy to assign the largest weight to the task k^* , so the loss on the sample is an ordinary cross-entropy

$$\mathcal{L}_{\text{policy}} = - \sum_{k=1}^K \mathbb{1}(k = k^*) \log(g_k(x)) \quad (11)$$

4.3.2 Training data pipeline. In order to compute Eq. 10, we save the tasks' prediction into the online logs, and we denote this extended online logs as:

$$\mathcal{D}'_{\tau} = \{(s_i, x_i, y_{\tau-s_i}, \hat{y}_{\text{tasks}, i})\}_{i=1}^{N_{\tau}} \quad (12)$$

where $\hat{y}_{\text{tasks}, i} = (\hat{y}_{1, i}, \hat{y}_{2, i}, \dots, \hat{y}_{K, i})$ is the *online* predictions of multiple tasks.

As we discuss in the previous sections, the prophet can only be constructed on $\mathcal{D}_{\tau, d_{\max}}$, so the training of aggregation policy should also based on the subset of the extended logs $\mathcal{D}'_{\tau, d_{\max}}$.

5 EXPERIMENTS

5.1 Streaming simulation set-up

We focus on the online system and use a simulation of online streaming while performing offline evaluation. The original dataset is $\mathcal{D}^*$. A simulated "current" timestamp τ is used to control the evaluation procedure. During the training process, only $\mathcal{D}_{\tau}$ is available. Every time after the model update, it will be used to predict and evaluate on the instances in next hour. Finally, the timestamp τ moves to one hour later, and we repeat the steps above.

5.2 Datasets

Criteo Conversion Logs. ¹ (Criteo-2014). This is a widely used public dataset for the delayed feedback problem [5]. It includes 60 days data, with $d_{\max}$ as 30 days. The last 20 days are used for evaluation.

Tencent Advertising Algorithm Competition 2017. ² (Tencent-2017). This dataset includes 14 days of click and conversion logs. According to the discussion about this competition, we set $d_{\max}$ to 2 days, and data of last 2 days are excluded. The last remaining 5 days are used for evaluation.

Tencent Ads Conversion Logs. (Tencent Ads). To verify the effectiveness of our method on production data, we collect 3 months of conversion logs from Tencent Ads, with about 200 million instances after sampling. We use the instances in the last month for evaluation.

5.3 Compared methods

Prophet*: An ideal CVR model constructed on $\mathcal{D}^*$ for evaluation.

Waiting(d): A method that only updates on $\mathcal{D}_{\tau, d}$ ($s < \tau - d$).

¹<https://labs.criteo.com/2013/12/conversion-logs-dataset/>

²<https://algo.qq.com/?lang=en>

Figure 2: Results. The methods “(d)” means “Waiting(d)” for short.

FSIW: Feedback shift importance weighting [25]. Note that this method is not originally designed for streaming settings.

PU, FNW, FNC: Using a fake negative data pipeline with a PU loss / FNW loss / FN Calibration.[14].

FTP: Our method. We set the tasks to (1 day, 7 days, 14 days, 21 days, 30 days) on Criteo-2014, (1 hour, 6 hours, 24 hours, 48 hours) on Tencent-2017, and (4 hours, 1 day, 5 days) on Tencent Ads.

We use a DNN model as the base model for all the methods. For the two public datasets, all features are transformed into embeddings, and concatenated together to feed into two dense layers with 128 units in each. The activation function is Leaky ReLU. L2 regularization is set to 10^{-6} on Criteo-2014 and 10^{-7} on Tencent-2017. The models are update by the Adam optimizer [12]. In our FTP method, all the tasks and the policy network share the embedding layers and the first dense layer at the bottom³.

5.4 Results

We conduct 5 random runs on the two public datasets, and the results are shown in Fig. 2 (a,b). Results on in-house Tencent Ads data is shown in Fig. 2 (c), and the relative metrics compared to prophet* are reported.

From Fig. 2, we observe that our FTP outperforms all other methods stably on all datasets. The performance of FSIW is not satisfactory, because it is not design for streaming settings, and fail to make use of recent feedback directly. PU loss is hard to converge on Criteo-2014 and Tencent Ads in our experiments. A similar issue is also mentioned in [14].

Over all the datasets, we find that the *Waiting* method with a well-tuned waiting period d can also have considerable performance. There isn’t much mentioned about this baseline in previous works. Our FTP is the only method that outperforms the best-tuned *Waiting* baseline.

³Our code on the public datasets: <https://github.com/kmdict/FollowTheProphet>

d_k	Feedback%	Best%
1d	60%	9%
7d	80%	18%
14d	90%	24%
21d	95%	24%
30d	100%	25%

(a) Criteo-2014

d_k	Feedback%	Best%
1h	73%	23%
6h	83%	23%
24h	95%	26%
48h	100%	28%

(b) Tencent-2017

Table 1: Feedback rate for $d < d_k$ and frequency of task k being the best task.

To further analyze the success of FTP, we show statistics about different tasks in FTP in Table. 1. “Feedback%” indicates how many percentages of feedback can be received within the specific task’s period. “Best%” counts up how many times that each task becomes the “best” task, with a prediction closest to the prophecy.

We highlight two characteristics that may lead to the success of FTP. First, all the tasks have considerable opportunity to become the best task, which indicates that they can capture different patterns under different delay periods. Second, although there is already 60% of feedback can be received within 1 day on Criteo-2014, it obviously does *not* mean the task with a delay of 1 day should always be the best task among them. FTP does not learn from the unknown delay time distribution directly, but use a more objective-oriented way: following the prophet.

6 CONCLUSION

In this paper, we focus on the online CVR prediction and propose “Following the Prophet” (FTP) method to tackle the delayed feedback problem. We construct an aggregation policy on top of multi-task predictions, where each task captures the feedback pattern during different periods. The policy is optimized towards an ideal “prophet” based on historical logged data. Extensive experiments on three real-world advertising datasets show that our method outperforms previous state-of-the-arts.

ACKNOWLEDGMENTS

This work is sponsored by CCF-Tencent Open Fund. The research work is also supported by the National Key Research and Development Program of China under Grant No. 2017YFB1002104, the National Natural Science Foundation of China under Grant No. 92046003, 61976204, U1811461. Xiang Ao is also supported by the Project of Youth Innovation Promotion Association CAS and Beijing Nova Program Z201100006820062.

REFERENCES

- [1] Deepak Agarwal, Rahul Agrawal, Rajiv Khanna, and Nagaraj Kota. 2010. Estimating rates of rare events with multiple hierarchies through scalable log-linear models. In *KDD*. 213–222.
- [2] Peter L Bartlett, Elad Hazan, and Alexander Rakhlin. 2007. Adaptive online gradient descent. In *NeurIPS*. 65–72.
- [3] Jonathan Baxter. 2000. A model of inductive bias learning. *JAIR* 12 (2000), 149–198.
- [4] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [5] Olivier Chapelle. 2014. Modeling delayed feedback in display advertising. In *KDD*. 1097–1105.
- [6] Marthinus Du Plessis, Gang Niu, and Masashi Sugiyama. 2015. Convex formulation for learning from positive and unlabeled data. In *ICML*. 1386–1394.
- [7] Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In *ACL*. 845–850.
- [8] David Eigen, Marc'Aurelio Ranzato, and Ilya Sutskever. 2013. Learning factored representations in a deep mixture of experts. In *ICLR*.
- [9] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *AdKDD*. 1–9.
- [10] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *SIGIR*. 549–558.
- [11] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation* 3, 1 (1991), 79–87.
- [12] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- [13] Ryuichi Kiryo, Gang Niu, Marthinus C du Plessis, and Masashi Sugiyama. 2017. Positive-unlabeled learning with non-negative risk estimator. In *NeurIPS*. 1674–1684.
- [14] Sofia Ira Ktena, Alykhan Tejani, Lucas Theis, Pranay Kumar Myana, Deepak Dilipkumar, Ferenc Huszar, Steven Yoo, and Wenzhe Shi. 2019. Addressing delayed feedback for continuous training with neural networks in CTR prediction. In *RecSys*. 187–195.
- [15] Kuang-chih Lee, Burkay Orten, Ali Dasdan, and Wentong Li. 2012. Estimating conversion rate in display advertising from past performance data. In *KDD*. 768–776.
- [16] Xun Liu, Wei Xue, Lei Xiao, and Bo Zhang. 2017. Pbodl: Parallel bayesian online deep learning for click-through rate prediction in tencent advertising system. *arXiv preprint arXiv:1707.00802* (2017).
- [17] Quan Lu, Shengjun Pan, Liang Wang, Junwei Pan, Fengdan Wan, and Hongxia Yang. 2017. A practical framework of conversion rate prediction for online display advertising. In *AdKDD*. 1–9.
- [18] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *KDD*. 1930–1939.
- [19] Feiyang Pan, Xiang Ao, Pingzhong Tang, Min Lu, Dapeng Liu, Lei Xiao, and Qing He. 2020. Field-aware Calibration: A Simple and Empirically Strong Method for Reliable Probabilistic Predictions. In *WWW*. 729–739.
- [20] Junwei Pan, Yizhi Mao, Alfonso Lobos Ruiz, Yu Sun, and Aaron Flores. 2019. Predicting different types of conversions with multi-task learning in online advertising. In *KDD*. 2689–2697.
- [21] Marthinus C du Plessis, Gang Niu, and Masashi Sugiyama. 2014. Analysis of learning from positive and unlabeled data. In *NeurIPS*. 703–711.
- [22] Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017).
- [23] Doyen Sahoo, Quang Pham, Jing Lu, and Steven CH Hoi. 2018. Online deep learning: learning deep neural networks on the fly. In *IJCAI*. 2660–2666.
- [24] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarsz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*.
- [25] Shota Yasui, Gota Morishita, Fujita Komei, and Masashi Shibata. 2020. A Feedback Shift Correction in Predicting Conversion Rates under Delayed Feedback. In *WWW*. 2740–2746.
- [26] Yuya Yoshikawa and Yusaku Imai. 2018. A nonparametric delayed feedback model for conversion rate prediction. *arXiv preprint arXiv:1802.00255* (2018).