

Delayed Feedback Modeling with Influence Functions

Chenlu Ding^{1*}, Jiancan Wu^{1,2†}, Yancheng Yuan^{3†}, Cunchun Li¹, Xiang Wang¹, Dingxian Wang⁴, Frank Yang⁴, Andrew Rabinovich⁴

¹University of Science and Technology of China

²Shanghai Key Laboratory of Data Science

³The Hong Kong Polytechnic University

⁴Upwork

dingchenlu200103@gmail.com, wujcan@gmail.com, yancheng.yuan@polyu.edu.hk, xiangwang1223@gmail.com,

Abstract

In online advertising under the cost-per-conversion (CPA) model, accurate conversion rate (CVR) prediction is crucial. A major challenge is delayed feedback, where conversions may occur long after user interactions, leading to incomplete recent data and biased model training. Existing solutions partially mitigate this issue but often rely on auxiliary models, making them computationally inefficient and less adaptive to user interest shifts. We propose IF-DFM, an Influence Function-empowered for Delayed Feedback Modeling which estimates the impact of newly arrived and delayed conversions on model parameters, enabling efficient updates without full retraining. By reformulating the inverse Hessian-vector product as an optimization problem, IF-DFM achieves a favorable trade-off between scalability and effectiveness. Experiments on benchmark datasets show that IF-DFM outperforms prior methods in both accuracy and adaptability.

Code — <https://github.com/oceanoceanna/IF-DFM>

Introduction

In online advertising, predicting conversion rate (CVR) is vital for revenue optimization under the cost-per-conversion (CPA) model (Ma et al. 2018; Lu et al. 2017; Lee et al. 2012), where advertisers pay only after users complete specific actions. Unlike immediate clicks, conversions often occur hours or even weeks after the ad click, leading to the *delayed feedback issue* (Chapelle 2014; Ktena et al. 2019; Yang et al. 2021). To address this, platforms typically define a fixed time window to wait for conversions and periodically update the model (Dai et al. 2023). Clicks followed by conversions within the window are treated as positives, while others—including those with longer delays—are marked negative (Guo et al. 2022; Ktena et al. 2019; Li et al. 2021). This creates a trade-off between label accuracy and model timeliness.

Many efforts have been made to mitigate the delayed feedback issue in CVR, which can be roughly categorized into offline and online approaches. Offline methods

(Yoshikawa and Imai 2018; Saito, Morishita, and Yasui 2020; Yasui et al. 2020) rely on historical data to train models, usually with an extra component to model the distribution of the delay time (Chapelle 2014) or adjust for the correct labels (Wang et al. 2023). They generally work under the i.i.d assumption that future data distributions will remain consistent with historical patterns. On the other hand, online methods (Chen et al. 2022; Gu et al. 2021; Li et al. 2021) attempt to update the model by ingesting newly observed data and correcting labels in near real-time or waiting short intervals. Such updates are conducted on the duplicated data with corrected labels (Yang et al. 2021) or using a surrogate loss, to approach the oracle model (Dai et al. 2023; Ktena et al. 2019). The framework of offline CVR methods, online CVR methods can be found in Figure 1.

Despite effectiveness, current approaches suffer from two significant limitations:

- **Inadequate Adaptation to Evolving User Interest.** The dynamic nature of user interests challenges the assumption of an unchanged data distribution, which may not hold in environments characterized by evolving user behavior (Dai et al. 2023). Offline methods, operating under the i.i.d assumption, struggle to capture users’ emerging behaviors and preference shifts. While online methods update with new data promptly, their reliance on duplicated samples for label correction can cause confusion and limit the effective use of accurate labels.
- **Reliance on Auxiliary Models.** The mainstream approaches typically incorporate auxiliary components to estimate potential label reversals for observed negatives or to model the probabilities of fake negatives as latent variables. However, developing these additional components can be as complex as constructing the primary CVR model. Moreover, their effectiveness is limited by the quality and quantity of historical data, resulting in computational inefficiencies and adding additional complexities to the CVR prediction task.

We explore a new paradigm for mitigating delayed feedback that naturally adapts to new data without auxiliary models. While retraining with correctly labeled data is straightforward, it is impractical in large-scale CVR settings

*Work done at The Hong Kong Polytechnic University.

†Corresponding author.

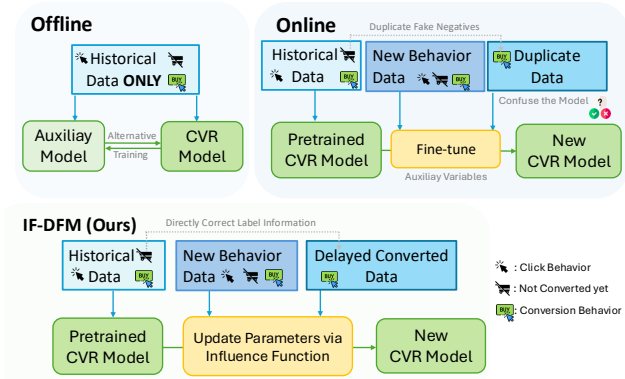


Figure 1: The framework of offline CVR methods, online CVR methods, and IF-DFM.

due to high computational costs. To address this, we leverage influence functions—a tool from robust statistics gaining traction in machine learning (Koh and Liang 2017). Influence functions estimate parameter changes caused by data perturbations (e.g., removing a sample (Zhang et al. 2023) or editing features (Wu et al. 2023)) by up-weighting the sample’s loss, resulting in inverse Hessian-vector products. Conceptually, incorporating the influence function to mitigate the delayed feedback problem offers a notable benefit: it allows for the direct modification of model parameters based on the estimated impact of mislabeled or newly injected data. However, given the significant computational and space complexities of inverting the Hessian matrix, integrating the influence function also presents its unique challenge: *How to efficiently get model parameter changes for CVR tasks?*

In this paper, we propose an influence function-empowered framework for delayed feedback modeling (IF-DFM). The core idea is to leverage influence functions to estimate the impact of newly injected data and directly update model parameters without retraining. For a deployed CVR model, IF-DFM computes the parameter changes caused by correcting mislabeled samples, approximating the difference between the current model and the one trained on corrected data. This avoids sample duplication and eliminates the need for auxiliary models. Additionally, IF-DFM supports the integration of newly arrived behavior data (e.g., post-deployment clicks), enabling the model to adapt to evolving user interactions and improve prediction accuracy. To reduce computation, we reformulate the inverse Hessian-vector product as a finite-sum quadratic problem, which can be efficiently solved using stochastic optimization methods like SGD (Robbins and Monro 1951) and its variants (Duchi, Hazan, and Singer 2011; Kingma and Ba 2015), avoiding full-batch gradients. We conduct extensive experiments on two benchmark datasets, demonstrating the superiority of the proposed method. Our main contributions can be summarized as follows:

- We propose a new paradigm IF-DFM for mitigating the delayed feedback problem. It utilizes the influence function to estimate the impact of fake label correction and




Figure 2: An illustration of the delayed feedback problem in CVR tasks.

new behavior data integration, and directly modify model parameters without retraining.

- We propose to compute the influence function by formulating the inverse Hessian-vector-product calculation as an optimization problem, achieving a balance between computational efficiency and effectiveness.
- Extensive experiments on benchmark datasets demonstrate the superiority of IF-DFM.

Preliminaries

Problem Formulation

The fundamental task of CVR prediction is to estimate the rate of conversion from the user clicks into purchases, based on historical data. Throughout this paper, we represent the i -th training sample as $z_i = (x_i, y_i^O)$, where x_i encompasses the features, y_i^O is the observed label. Furthermore, in CVR prediction, each sample is also associated with a true label y_i^T , which is inaccessible within the training dataset. The feature set x_i typically comprises temporal components—click timestamp c_i , payment timestamp p_i , and the elapsed time e_i from the click to a model-training-start timestamp T —along with user/item characteristics. Here, the default value of p_i is -1 if no conversion has occurred before T . The training set accumulated up to T can be described as $\mathcal{Z} = \{z_i \mid i = 1, 2, \dots, n\}$, where n is the sample size. Due to the delayed feedback, where the ground truth y_i^T is absent, the training process could be misled by fake negative samples.

The relationship between sample labels and the timeline is illustrated in Figure 2. If a conversion occurs before T (i.e., $p_i < T$), the sample is labeled as $y_i^O = y_i^T = 1$, representing a true positive. If the conversion is unobserved by T (i.e., $p_i > T$), it is labeled as $y_i^O = 0$, which may not match y_i^T , resulting in a fake negative due to delayed conversion. Thus, delayed feedback arises solely from negatives that may later become positives. The model is evaluated at timestamp T' , using only the test samples with their true labels.

Vanilla and Retrain Versions of Loss

The CVR model, $f(x, \theta)$ parameterized by θ , operates as a binary classifier, predicting whether a click action x will result in a conversion. Training such models typically employs binary cross-entropy loss (Dai et al. 2023; Wang et al. 2023). We begin with the **Vanilla** version (Chapelle 2014;

Yasui et al. 2020), where all unconverted samples are treated as negatives:

$$\mathcal{L}_V(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{BCE}(z_i, \theta), \quad (1)$$

where for any sample $z = (x, y)$ and model parameter θ , $\mathcal{L}_{BCE}(z, \theta) = -[y \log f(x, \theta) + (1 - y) \log(1 - f(x, \theta))]$ is the commonly used binary cross-entropy loss.

The vanilla loss is computed over observed labels without correction, serving as a lower bound for CVR performance. In contrast, the **Oracle** version assumes access to true labels y_i^T , providing an upper bound:

$$\mathcal{L}_O(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{BCE}(z_i^T, \theta), \quad (2)$$

where $z_i^T = (x_i, y_i^T)$. Since many conversions occur after model deployment, true labels are typically unavailable during training. To address this, we define a practical **Retrain** version:

$$\mathcal{L}_R(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{BCE}(\tilde{z}_i, \theta), \quad (3)$$

where $\tilde{z}_i = (x_i, \tilde{y}_i)$ and \tilde{y}_i is defined as:

$$\tilde{y}_i = \begin{cases} y_i^O & \text{the sample was not converted before testing,} \\ y_i^T & \text{the sample converted before testing.} \end{cases}$$

That is, retrain loss assumes access to all true conversions observed up to the evaluation time, representing the best achievable practical performance.

Influence Function

Influence function (Hampel 1974), which is a classical tool in robust statistics for measuring the changes of the model parameters to small changes in the weights of training samples, has been recently introduced in the machine learning community for understanding black-box predictions and beyond (Koh and Liang 2017). Given a training data set $\mathcal{Z} = \{z_1, z_2, \dots, z_n\}$, we consider the following empirical risk minimization problem:

$$\hat{\theta} \in \arg \min_{\theta} R(\theta) := \frac{1}{n} \sum_{i=1}^n \mathcal{L}(z_i, \theta), \quad (4)$$

where $\mathcal{L}(\cdot, \theta)$ is the loss function (e.g., Equations (1) (2)), θ is the model parameter.

We are interested in estimating the model parameter change $\Delta\theta(\epsilon) = \hat{\theta}_{new}(\epsilon) - \hat{\theta}$ if some training sample z_j is slightly reweighted by ϵ , where

$$\hat{\theta}_{new}(\epsilon) \in \arg \min_{\theta} \hat{\mathcal{L}}(z_j; \theta, \epsilon) = R(\theta) + \epsilon \mathcal{L}(z_j, \theta). \quad (5)$$

When $\epsilon \approx 0$, the influence function provides an elegant tool for estimating $\Delta\theta(\epsilon)$ without solving (5) as

$$\begin{aligned} \Delta\theta(\epsilon) &\approx -\epsilon H_{\hat{\theta}}^{-1} \nabla_{\theta} \mathcal{L}(z_j, \hat{\theta}) \\ \text{and } \left. \frac{d\Delta\theta(\epsilon)}{d\epsilon} \right|_{\epsilon=0} &= -H_{\hat{\theta}}^{-1} \nabla_{\theta} \mathcal{L}(z_j, \hat{\theta}), \end{aligned} \quad (6)$$

where $H_{\hat{\theta}} = \nabla_{\theta}^2 R(\hat{\theta})$ is the Hessian matrix of $R(\cdot)$ at $\theta = \hat{\theta}$. A detailed derivation of the expression in (6) can be found in (Koh and Liang 2017).

The key to applying the influence function for estimating $\Delta\theta$ lies in efficiently computing the right-hand side of (6), which is challenging in practice due to high-dimensional parameters and large-scale data. As a core contribution, we propose a novel stochastic algorithm to tackle this computational bottleneck and unlock the practical potential of influence functions.

Method

Delayed feedback in CVR prediction presents challenges from both label reversal and new data integration. We propose IF-DFM, a unified framework that addresses both issues without retraining, and introduce a stochastic algorithm to reduce the computational cost of influence function estimation.

IF-DFM

In this section, we elaborate on the proposed IF-DFM framework, which models both label reversals and newly arrived data as data perturbations.

Label Reversal In CVR prediction, delays between click and conversion behaviors often lead to label reversals, where some fake negatives become positives after training. This phenomenon can be naturally modeled as data perturbation. Specifically, if a label reversal occurs for a training sample z_i , it can be represented as perturbing z_i to $z_i^{\delta} = (x_i, y_i^O + \delta)$. Since observed positives are correctly labeled, the perturbation for CVR is $\delta = 1$. Without loss of generality, we consider the case where a label reversal occurs for z_i , and study the perturbed empirical risk minimization by reweighting this sample:

$$\hat{\theta}_{\delta}(\epsilon) \in \arg \min_{\theta} \left(\mathcal{L}_V(\theta) + \epsilon \mathcal{L}_{BCE}(z_i^{\delta}, \theta) - \epsilon \mathcal{L}_{BCE}(z_i, \theta) \right). \quad (7)$$

Note that the right-hand side of (7) coincides with the retrain loss (3) if we take $\epsilon = 1/n$. When $\epsilon \approx 0$ (which is true if $\epsilon = 1/n$ and n is sufficiently large), following the derivation of influence function (e.g., see (Koh and Liang 2017)), we can obtain that

$$\begin{aligned} \left. \frac{d\hat{\theta}_{\delta}}{d\epsilon} \right|_{\epsilon=0} &= - \left(\nabla_{\theta}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1} \left(\nabla_{\theta} \mathcal{L}_{BCE}(z_i^{\delta}, \hat{\theta}) \right. \\ &\quad \left. - \nabla_{\theta} \mathcal{L}_{BCE}(z_i, \hat{\theta}) \right), \end{aligned} \quad (8)$$

where $\hat{\theta}$ is a minimizer of (1) (which is available if the Vanilla model is trained).

Therefore, after aggregating the gradients of all label-reversed samples, we can estimate the changes of the model parameter from (8) by taking $\epsilon = 1/n$ as

$$\begin{aligned} \Delta\theta_{\text{delay}} &\approx \frac{1}{n} \left(\nabla_{\theta}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1} \left(\sum_{j \in J} \nabla_{\theta} \mathcal{L}_{BCE}(z_j, \hat{\theta}) \right. \\ &\quad \left. - \sum_{j \in J} \nabla_{\theta} \mathcal{L}_{BCE}(z_j^{\delta}, \hat{\theta}) \right), \end{aligned} \quad (9)$$

where J is the index set of samples that clicked before T and converted during the $[T, T']$ interval (*i.e.*, samples with label reversal).

Newly Arrived Behavioral Data Integration Conventionally, the influence functions estimate how model parameters or test loss change when a sample is removed or altered. In our IF-DFM framework, we extend this to incorporate new behavioral data arriving during the interval $[T, T']$. To this end, we consider the following perturbed loss function for a new sample $z = (x, y)$ and $0 \leq \epsilon < 1$:

$$\hat{\theta}_z(\epsilon) \in \arg \min_{\theta} ((1 - \epsilon)\mathcal{L}_V(\theta) + \epsilon\mathcal{L}_{BCE}(z, \theta)). \quad (10)$$

Note that we can recover the retrain loss for integrating the additional new sample z by taking $\epsilon = 1/(n + 1)$. By careful calculations, we can find an approximation to $\Delta_{z,\epsilon} = \hat{\theta}_z(\epsilon) - \hat{\theta}$ for sufficiently small $\epsilon > 0$ as

$$\Delta_{z,\epsilon} \approx -\frac{\epsilon}{1 - \epsilon} \left(\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1} \nabla_{\theta} \mathcal{L}_{BCE}(z, \hat{\theta}). \quad (11)$$

Therefore, for sufficiently large n , we obtain an approximation to the change of model parameter for integrating the new data sample z as

$$\begin{aligned} \Delta\theta_{add} &= \hat{\theta}_z \left(\frac{1}{n+1} \right) - \hat{\theta} \\ &\approx -\frac{1}{n} \left(\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1} \nabla_{\theta} \mathcal{L}_{BCE}(z, \hat{\theta}). \end{aligned} \quad (12)$$

Considering both the label reversed data and the newly arrived data, the equation for the total parameter change can be given as

$$\Delta\theta_{total} = \Delta\theta_{delay} + \Delta\theta_{add}. \quad (13)$$

More specifically, we need to estimate

$$\begin{aligned} \Delta\theta_{total} &\approx \left(\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1} \left(\frac{1}{n} \sum_{j \in J} \nabla_{\theta} \mathcal{L}_V(z_j, \hat{\theta}) \right. \\ &\quad \left. - \frac{1}{n} \sum_{j \in J} \nabla_{\theta} \mathcal{L}_V(z_j^{\delta}, \hat{\theta}) - \frac{1}{n} \sum_{k \in K} \nabla_{\theta} \mathcal{L}_V(z_k, \hat{\theta}) \right), \end{aligned} \quad (14)$$

where J and K are the index sets for the label reversed data samples and newly arrived data samples, respectively.

Addressing problem (14) poses significant computational challenges. When the Hessian matrix $\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta})$ is positive definite, we can apply the conjugate gradient (CG) method to estimate $\Delta\theta_{total}$, where only Hessian-vector products are required. For $\theta \in \mathbb{R}^p$ (where p is usually huge in applications), the CG method requires at most p iterations to calculate the right-hand side of (14) and it usually requires much fewer iterations to obtain an approximate solution if the condition number of $\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta})$ is close to 1 (Nocedal and Wright 1999). However, the CG method typically requires full-batch gradient and Hessian matrix computations, making it impractical for large-scale datasets. To address the computational challenges of the CG method, an approximation of $\left(\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \right)^{-1}$ via its truncated power series has

been widely used to estimate $\Delta\theta_{total}$ (Koh and Liang 2017). However, its accuracy is often compromised. Finding a more precise and efficient solution for (14) remains a challenge. Next, we will address this challenge by designing an efficient and scalable stochastic algorithm.

An Efficient and Scalable Method for Calculating the parameter changes $\Delta\theta_{total}$

In this section, we will address the computational challenges for calculating $\Delta\theta_{total}$. Our core idea is to compute $\Delta\theta_{total}$ by converting (14) into an equivalent finite-sum optimization problem. Consequently, we can apply popular scalable algorithms, such as the stochastic gradient descent (Robbins and Monro 1951) and its variants (Duchi, Hazan, and Singer 2011; Kingma and Ba 2015; Johnson and Zhang 2013), to calculate $\Delta\theta_{total}$ efficiently. The definition of $\Delta\theta_{total}$ implies that it is a solution to the linear system

$$\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \Delta = b \quad (15)$$

with

$$b = \frac{1}{n} \sum_{j \in J} \nabla_{\theta} \mathcal{L}_V(z_j, \hat{\theta}) - \frac{1}{n} \sum_{j \in J} \nabla_{\theta} \mathcal{L}_V(z_j^{\delta}, \hat{\theta}) \quad (16)$$

$$- \frac{1}{n} \sum_{k \in K} \nabla_{\theta} \mathcal{L}_V(z_k, \hat{\theta}). \quad (17)$$

Since $\hat{\theta}$ is a minimizer of (1), it satisfies the second-order necessary optimality condition (Nocedal and Wright 1999), implying that the Hessian $\nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta})$ is symmetric and positive semidefinite. As a result, $\Delta\theta_{total}$ corresponds to the solution of a convex quadratic optimization problem:

$$\min_{\Delta} F(\Delta) := \frac{1}{2} \Delta^{\top} \nabla_{\hat{\theta}}^2 \mathcal{L}_V(\hat{\theta}) \Delta - \langle b, \Delta \rangle. \quad (18)$$

It follows from the finite-sum structure of \mathcal{L}_V that

$$F(\Delta) = \frac{1}{n} \sum_{i=1}^n f_i(\Delta), \quad (19)$$

where

$$f_i(\Delta) = \frac{1}{2} \Delta^{\top} \nabla_{\hat{\theta}}^2 \mathcal{L}_{BCE}(z_i, \hat{\theta}) \Delta - \langle b, \Delta \rangle. \quad (20)$$

Therefore, we have converted the calculation of $\Delta\theta_{total}$ equivalently to minimizing an optimization problem with a finite-sum objective function. Note that the function value and the gradient of f_i can be efficiently evaluated via the built-in auto-differentiation and the Hessian-vector-product modules implemented by popular frameworks, such as PyTorch. Consequently, we can apply efficient and scalable optimization algorithms to calculate $\Delta\theta_{total}$. In particular, we choose ADAM (Kingma and Ba 2015) in this paper. We summarize the details in Algorithm 1.

Experiment

Our experiments aim to investigate the following research questions: **RQ1**: How does our method's performance compare to state-of-the-art methods in CVR prediction? **RQ2**: Can our method adapt to dynamic changes in user interests over time? **RQ3**: How efficient is our method in calculating parameter changes?

Algorithm 1: An efficient algorithm to calculate $\Delta\theta_{total}$.

- 1: **Input:** Training data \mathcal{D} ; Label reversed data $z_j \rightarrow z_j^\delta$;
Newly arrived data z_k ; Trained vanilla model $\hat{\theta}$.
 - 2: **Output:** The model parameter change $\Delta\theta_{total}$.
 - 3: **Initialization:** Construct b in (16) and the function f_i .
 - 4: Call ADAM to minimize (19) and obtain Δ^* .
 - 5: Return $\Delta\theta_{total} = \Delta^*$.
-

Experimental Settings

Datasets. To assess our method, we conduct experiments on two large-scale datasets: Criteo and Taobao. The statistics of these datasets are summarized in Table 1.

| Name | # clicks | CVR | Name | # clicks | CVR |
|--------|-----------|--------|--------|-----------|--------|
| Criteo | 4,019,339 | 0.2227 | Taobao | 8,544,800 | 0.0100 |

Table 1: Table 1: Statistics of Criteo and Taobao datasets.

Criteo¹: A public dataset with click logs from Criteo’s live traffic, containing timestamps, conversion labels, and both categorical and continuous features. We use 3 million samples spanning 14 days for training.

Taobao²: A dataset of user interactions on Taobao, including behaviors such as page views and purchases, along with timestamps and associated features.

Following common settings in previous offline research (Wang et al. 2023; Yasui et al. 2020), we employ a temporal partitioning approach to divide each dataset into training, validation, and test sets. Specifically, we use data from the interval $[T' - d_{test}, T']$ for validation and data from $[T', T' + d_{test}]$ for testing, where d_{test} represents the size of the validation/test set. The interval $[T, T']$ covers a period of c days. We will explore the value of c in the following sections. Such a setting enables us to assess the CVR model’s ability to capture evolving user interests.

Metrics. We adopt three standard metrics to evaluate CVR performance following (Wang et al. 2023; Dai et al. 2023): **AUC** (Area Under the ROC Curve), which measures the ranking quality between positive and negative samples; **PRAUC** (Precision-Recall AUC), which evaluates performance across precision-recall thresholds; and **Log Loss (LL)**, which assesses prediction accuracy via probabilistic error (lower is better). We also report the **Relative Improvement (RI)** over the Vanilla model, normalized by the performance gap between Retrain and Vanilla. For a method f and metric M , RI is computed as:

$$RI_M = \frac{M(f) - M(\text{Vanilla})}{M(\text{Retrain}) - M(\text{Vanilla})}. \quad (21)$$

This yields RI-AUC, RI-PRAUC, and RI-LL, reflecting how much each method narrows the gap to Retrain.

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

²<https://tianchi.aliyun.com/dataset/dataDetail?dataId=649&userId=1&lang=en-us>

Baselines. We compare our method with 13 state-of-the-art baselines, covering both offline (e.g., Vanilla, Retrain, DFM (Chapelle 2014), FSIW (Yasui et al. 2020), nnDF (Yasui and Kato 2022), ULC (Wang et al. 2023)) and online approaches (e.g., Pretrain, Retrain-online, FNC (Ktena et al. 2019), FNW (Ktena et al. 2019), DDFM (Dai et al. 2023), ES-DFM (Yang et al. 2021), DEFUSE (Chen et al. 2022)). For fair comparison, we follow (Wang et al. 2023) and adopt four widely-used CVR backbones: MLP, DeepFM (Guo et al. 2017), AutoInt (Song et al. 2019), and DCNV2 (Wang et al. 2021). Detailed descriptions of all baselines are provided in Appendix.

Implementation Details. In the offline setting, IF-DFM leverages the influence function to estimate $\Delta\theta_{delay}$ and adjust the Vanilla model’s parameters. In the online setting, we use the offline Vanilla model as the pretrained model. For methods requiring auxiliary models or variables (e.g., DDFM), we follow their pretraining methods. The model parameters are modified based on the approximate $\Delta\theta_{total}$. More implementation details can be found in Appendix.

Performance Comparison


The offline results on the Criteo dataset across different backbones are summarized in Table 2, from which we draw the following observations:

- Retrain consistently outperforms Vanilla, confirming that delayed feedback negatively impacts model performance. Both FSIW and DFM improve upon Vanilla, with FSIW achieving better results due to its pre-trained auxiliary models. In contrast, nnDF underperforms due to its reliance on global information and the need for full data during updates.
- ULC shows limited performance, as it depends on a label correction model that requires high-quality historical data and struggles to adapt to shifting user interests.
- IF-DFM significantly outperforms all baselines across metrics and backbones, demonstrating its effectiveness in handling delayed feedback. It achieves average improvements of 0.55% in AUC, 1.29% in PRAUC, and 4.99% in LL. Prior work from Google (Cheng et al. 2016; Wang et al. 2017) and Microsoft (Ling et al. 2017) indicates that even a 0.1% AUC gain can lead to substantial real-world benefits.
- In terms of relative gains, IF-DFM approaches the performance of Retrain, with improvements of 85.16% (AUC), 82.07% (PRAUC), and 85.89% (LL), underscoring its effectiveness.

Compared to Criteo, the Taobao dataset has a lower conversion rate. As shown in Figure 3, IF-DFM surpasses the best baseline in most cases, with AUC and PRAUC gains of 0.28% and 0.94%, respectively. FSIW benefits from pre-trained importance estimators, giving it short-term advantages. However, IF-DFM, by leveraging real label information, better approximates retraining and, as shown in subsequent experiments, adapts more effectively to evolving user interests.

| Backbone | Method | AUC \uparrow | PRAUC \uparrow | LL \downarrow | RI-AUC \uparrow | RI-PRAUC \uparrow | RI-LL \uparrow |
|----------|----------------------|-----------------|------------------|-----------------|-------------------|---------------------|------------------|
| MLP | Vanilla | 0.8353 | 0.6398 | 0.4187 | 0.000 | 0.000 | 0.000 |
| | Retrain | 0.8419 | 0.6513 | 0.3930 | 1.000 | 1.000 | 1.000 |
| | DFM | 0.8355 | 0.6409 | 0.4152 | 0.0303 | 0.0957 | 0.1362 |
| | FSIW | 0.8369 | 0.6433 | 0.4020 | 0.2424 | 0.3043 | 0.6498 |
| | nnDF | 0.6887 | 0.4029 | 0.5639 | -22.21 | -20.60 | -5.670 |
| | ULC | 0.8343 | 0.6412 | 0.3994 | -0.1515 | 0.1217 | 0.7510 |
| | IF-DFM (Ours) | 0.8411** | 0.6491** | 0.3958** | 0.8788** | 0.8087** | 0.9662** |
| DeepFM | Vanilla | 0.8368 | 0.6411 | 0.4168 | 0.000 | 0.000 | 0.000 |
| | Retrain | 0.8418 | 0.6509 | 0.3926 | 1.000 | 1.000 | 1.000 |
| | DFM | 0.8381 | 0.6434 | 0.4065 | 0.260 | 0.2347 | 0.4256 |
| | FSIW | 0.8378 | 0.6436 | 0.4023 | 0.200 | 0.2551 | 0.5992 |
| | nnDF | 0.7025 | 0.4146 | 0.5651 | -26.86 | -23.11 | -6.128 |
| | ULC | 0.8376 | 0.6446 | 0.3982 | 0.160 | 0.3571 | 0.7686 |
| | IF-DFM (Ours) | 0.8412** | 0.6492** | 0.3962** | 0.8800** | 0.8265** | 0.8512** |
| AutoInt | Vanilla | 0.8361 | 0.6401 | 0.4174 | 0.000 | 0.000 | 0.000 |
| | Retrain | 0.8415 | 0.6503 | 0.3929 | 1.000 | 1.000 | 1.000 |
| | DFM | 0.8374 | 0.6426 | 0.4091 | 0.2407 | 0.2451 | 0.3388 |
| | FSIW | 0.8376 | 0.6442 | 0.4026 | 0.2778 | 0.4020 | 0.6041 |
| | nnDF | 0.6820 | 0.3770 | 0.5789 | -28.54 | -25.79 | -6.592 |
| | ULC | 0.8374 | 0.6443 | 0.3987 | 0.2407 | 0.4118 | 0.7633 |
| | IF-DFM (Ours) | 0.8404** | 0.6475** | 0.3965** | 0.7963** | 0.7255** | 0.8530** |
| DCNV2 | Vanilla | 0.8370 | 0.6418 | 0.4174 | 0.000 | 0.000 | 0.000 |
| | Retrain | 0.8417 | 0.6508 | 0.3927 | 1.000 | 1.000 | 1.000 |
| | DFM | 0.8372 | 0.6422 | 0.4087 | 0.0426 | 0.044 | 0.3522 |
| | FSIW | 0.8388 | 0.6456 | 0.3997 | 0.3830 | 0.422 | 0.6045 |
| | nnDF | 0.6850 | 0.3878 | 0.5728 | -32.34 | -28.22 | -6.291 |
| | ULC | 0.8368 | 0.6441 | 0.3989 | -0.0426 | 0.2556 | 0.7490 |
| | IF-DFM (Ours) | 0.8410** | 0.6501** | 0.3985** | 0.8511** | 0.9222** | 0.7652** |

Table 2: Experimental results on Criteo. Superscript ** indicates $p \leq 0.05$ for the t-test comparing IF-DFM with the best baseline.



(a) AUC Results


(b) PRAUC Results

Figure 3: Offline experimental results on Taobao dataset.

Effectiveness of Adapting to Evolving User Interest

User interest drift is a critical challenge in CVR prediction, as outdated models may fail to capture evolving preferences, leading to degraded performance. Table 3 reports the results of various methods under different values of c , which denotes the time gap between training and test sets. From these results, we can draw the following conclusions:

- ULC and DFM degrade notably over time. At $c = 14$, both perform worse than Vanilla. FSIW shows some resilience by adaptively reweighting samples based on feedback, but its effectiveness is limited as the temporal gap grows. ULC, constrained by its use of counterfactual deadlines, fails to generalize well.
- The performance gap between IF-DFM and baselines



(a) AUC Results

(b) PRAUC Results

Figure 4: Online experimental results on Criteo dataset.

grows with increasing c , highlighting its robustness to temporal drift. By estimating the influence of label reversed behavior data, IF-DFM effectively approximates retraining without incurring its full cost.

In the online setting, IF-DFM estimates parameter updates from both fake negatives and newly arrived data. We also evaluate a variant, IF-DFM-w/o-add, which omits updates from new data ($\Delta\theta_{add}$), using only delayed conversions ($\Delta\theta_{delay}$). The results are shown in Figure 4. We can find that:

- Compared to pretraining, online CVR models can mitigate the delayed feedback problem to some extent. This can be attributed to their fine-tuning pipeline using new data and delayed conversion data.

| | | AUC \uparrow | PRAUC \uparrow | LL \downarrow |
|----------|-------------|----------------|------------------|-----------------|
| $c = 5$ | Vanilla | 0.8427 | 0.6479 | 0.4085 |
| | Retrain | 0.8464 | 0.6541 | 0.3879 |
| | nnDF | 0.6840 | 0.3979 | 0.5565 |
| | DFM | 0.8412 | 0.6469 | 0.4064 |
| | FSIW | 0.8441 | 0.6508 | 0.3898 |
| | ULC | 0.8445 | 0.6513 | 0.3894 |
| | Ours | 0.8452 | 0.6514 | 0.3883 |
| $c = 14$ | Vanilla | 0.8369 | 0.6385 | 0.4114 |
| | Retrain | 0.8453 | 0.6529 | 0.3860 |
| | nnDF | 0.7041 | 0.4116 | 0.5521 |
| | DFM | 0.8351 | 0.6370 | 0.4076 |
| | FSIW | 0.8385 | 0.6421 | 0.3965 |
| | ULC | 0.8345 | 0.6385 | 0.3954 |
| | Ours | 0.8420 | 0.6463 | 0.3938 |

Table 3: Offline experimental results on the Criteo dataset for different values of c , with MLP as the backbone.

| | Vanilla | ULC | DFM | FSIW | Vanilla + Ours |
|-----------------|---------|--------|--------|--------|----------------|
| Time (s) | 1351.4 | 2467.4 | 1979.8 | 1018.5 | 1351.4 + 14.8 |

Table 4: Running time of different methods under the offline setting, with MLP as the backbone and $c = 10$.

- IF-DFM outperforms IF-DFM-w/o-add, with AUC and PRAUC gains of 0.36% and 0.47%, respectively, indicating the benefit of incorporating newly arrived data.
- IF-DFM exceeds all baselines and closely approaches retraining performance, demonstrating its ability to effectively integrate delayed and new label information for improved adaptability.

Efficiency of Calculating Parameter Changes

Ensuring rapid updates of model parameters is crucial for maintaining the freshness of a CVR model. Due to space constraints, we report runtime results only for methods with the MLP backbone on the Criteo dataset, as shown in Table 4. We have the following observations:

- Compared to Vanilla, ULC requires the training of additional auxiliary models and the alternating training of two models, nearly doubling the training time. This indicates that relying on auxiliary models incurs extra time costs, making it inefficient and computationally expensive.
- IF-DFM updates model parameters without retraining, taking only 14.8 seconds—just 1.1% of Vanilla’s training time—demonstrating its high efficiency in addressing delayed feedback.

Related work

Delayed Feedback

Delayed feedback is a central issue in CVR modeling. Offline methods assume *i.i.d.* data, using historical samples to predict future conversions. DFM (Chapelle 2014) models delay with an exponential distribution, while

NoDeF (Yoshikawa and Imai 2018) adopts a non-parametric approach. Methods like FSIW (Yasui et al. 2020), nnDF (Yasui and Kato 2022), and ULC (Wang et al. 2023) aim to construct unbiased estimators of oracle loss through reweighting or label correction. However, most require training auxiliary models, hindering adaptability to evolving user interests. Online methods update models with real-time feedback. ES-DFM (Yang et al. 2021) uses elapsed-time sampling for delayed positives, DEFUSE (Chen et al. 2022) employs a bi-distribution strategy, and DDFM (Dai et al. 2023) uses dual estimators for different data phases. Yet, duplicating delayed conversions may introduce ambiguity, and many approaches still depend on auxiliary components.

Influence Function

Influence functions estimate the effect of individual training points on model outputs. Introduced by Koh *et al.* (Koh and Liang 2017), they have since been applied to interpretability (Chhabra et al. 2024; Basu, You, and Feizi 2020), fairness (Li and Liu 2022; Wang, Wu, and He 2024), data management (Hara, Nitanda, and Maehara 2019), unlearning (Chen et al. 2023; Liu et al. 2023), and recommendation systems (Cheng et al. 2018; Guo et al. 2024; Li et al. 2023). Extensions include robustness analysis (Koh et al. 2019) and data debugging via TracIn (Pruthi et al. 2020). IF4URec (Yu et al. 2020) improves recommender systems through sample reweighting. However, as Basu *et al.* (Basu, Pope, and Feizi 2021) highlight, scalability remains a challenge in deep learning, and although FastIF (Guo et al. 2021) improves efficiency to some extent, it does not touch the core of the inverse Hessian-vector product computation, limiting applicability to large models.

Conclusion

We propose IF-DFM, an influence function-based framework for modeling delayed feedback. IF-DFM efficiently estimates parameter updates caused by delayed conversions and new data, approximating retraining effects without duplicating samples or relying on auxiliary models. We further develop a scalable stochastic optimization algorithm to accelerate this computation. Our approach is applicable in both online and offline contexts. Extensive experiments validate that IF-DFM effectively mitigates the delayed feedback issue while incorporating new data to adjust to evolving user interests.

Limitations and Future Work

While IF-DFM effectively addresses delayed feedback in CVR tasks, it relies on gradient information and assumes access to the full training set (Koh and Liang 2017). We also plan to extend our study by deploying the method in real-world settings with A/B testing. In future work, we aim to incorporate influence functions into various recommendation methods, particularly LLM-based models (Yang et al. 2023; Bao et al. 2025), to enhance adaptive and trustworthy model behavior in diverse real-world scenarios.

Acknowledgements

This research is supported by the National Science and Technology Major Project (2023ZD0121102), National Natural Science Foundation of China (U25A20445, 62302321), Fundamental Research Funds for the Central Universities (WK2100250065). The work of Yancheng Yuan is supported by the Research Center for Intelligent Operations Research and The Hong Kong Polytechnic University under grant (P0045485). The work is also collaborated with Zhang Yang etc. from National University of Singapore. This research is also supported by the advanced computing resources provided by the Supercomputing Center of the USTC.

References

- Bao, K.; Zhang, J.; Wang, W.; Zhang, Y.; Yang, Z.; Luo, Y.; Chen, C.; Feng, F.; and Tian, Q. 2025. A bi-step grounding paradigm for large language models in recommendation systems. *ACM Transactions on Recommender Systems*, 3(4): 1–27.
- Basu, S.; Pope, P.; and Feizi, S. 2021. Influence Functions in Deep Learning Are Fragile. In *ICLR*.
- Basu, S.; You, X.; and Feizi, S. 2020. On Second-Order Group Influence Functions for Black-Box Predictions. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 715–724.
- Chapelle, O. 2014. Modeling delayed feedback in display advertising. In *KDD*, 1097–1105.
- Chen, R.; Yang, J.; Xiong, H.; Bai, J.; Hu, T.; Hao, J.; Feng, Y.; Zhou, J. T.; Wu, J.; and Liu, Z. 2023. Fast Model DeBias with Machine Unlearning. In *NeurIPS*.
- Chen, Y.; Jin, J.; Zhao, H.; Wang, P.; Liu, G.; Xu, J.; and Zheng, B. 2022. Asymptotically Unbiased Estimation for Delayed Feedback Modeling via Label Correction. In *WWW*, 369–379.
- Cheng, H.-T.; Koc, L.; Harmsen, J.; Shaked, T.; Chandra, T.; Aradhye, H.; Anderson, G.; Corrado, G.; Chai, W.; Ispir, M.; et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, 7–10.
- Cheng, W.; Shen, Y.; Zhu, Y.; and Huang, L. 2018. Explaining latent factor models for recommendation with influence functions. *arXiv preprint arXiv:1811.08120*.
- Chhabra, A.; Li, P.; Mohapatra, P.; and Liu, H. 2024. "What Data Benefits My Classifier?" Enhancing Model Performance and Interpretability through Influence-Based Data Selection. In *ICLR*.
- Dai, S.; Zhou, Y.; Xu, J.; and Wen, J. 2023. Dually Enhanced Delayed Feedback Modeling for Streaming Conversion Rate Prediction. In *CIKM*, 390–399.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7).
- Gu, S.; Sheng, X.; Fan, Y.; Zhou, G.; and Zhu, X. 2021. Real Negatives Matter: Continuous Training with Real Negatives for Delayed Feedback Modeling. In *KDD*, 2890–2898.
- Guo, H.; Rajani, N.; Hase, P.; Bansal, M.; and Xiong, C. 2021. FastIF: Scalable Influence Functions for Efficient Model Interpretation and Debugging. In *EMNLP (1)*, 10333–10350.
- Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *IJCAI*, 1725–1731.
- Guo, Y.; Cai, F.; Pan, Z.; Shao, T.; Chen, H.; and Zhang, X. 2024. A counterfactual explanation method based on modified group influence function for recommendation. *Complex & Intelligent Systems*, 1–13.
- Guo, Y.; Li, H.; Ao, X.; Lu, M.; Liu, D.; Xiao, L.; Jiang, J.; and He, Q. 2022. Calibrated Conversion Rate Prediction via Knowledge Distillation under Delayed Feedback in Online Advertising. In *CIKM*, 3983–3987.
- Hampel, F. R. 1974. The influence curve and its role in robust estimation. *Journal of the american statistical association*, 69(346): 383–393.
- Hara, S.; Nitanda, A.; and Maehara, T. 2019. Data Cleansing for Models Trained with SGD. In *NeurIPS*, 4215–4224.
- Johnson, R.; and Zhang, T. 2013. Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. In *NIPS*, 315–323.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- Koh, P. W.; Ang, K.; Teo, H. H. K.; and Liang, P. 2019. On the Accuracy of Influence Functions for Measuring Group Effects. In *NeurIPS*, 5255–5265.
- Koh, P. W.; and Liang, P. 2017. Understanding Black-box Predictions via Influence Functions. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, 1885–1894.
- Ktena, S. I.; Tejani, A.; Theis, L.; Myana, P. K.; Dilipkumar, D.; Huszár, F.; Yoo, S.; and Shi, W. 2019. Addressing delayed feedback for continuous training with neural networks in CTR prediction. In *RecSys*, 187–195.
- Lee, K.; Orten, B.; Dasdan, A.; and Li, W. 2012. Estimating conversion rate in display advertising from past performance data. In *KDD*, 768–776.
- Li, H.; Pan, F.; Ao, X.; Yang, Z.; Lu, M.; Pan, J.; Liu, D.; Xiao, L.; and He, Q. 2021. Follow the Prophet: Accurate Online Conversion Rate Prediction in the Face of Delayed Feedback. In *SIGIR*, 1915–1919.
- Li, P.; and Liu, H. 2022. Achieving Fairness at No Utility Cost via Data Reweighting with Influence. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 12917–12930.
- Li, Y.; Chen, C.; Zheng, X.; Zhang, Y.; Gong, B.; Wang, J.; and Chen, L. 2023. Selective and collaborative influence function for efficient recommendation unlearning. *Expert Systems with Applications*, 234: 121025.
- Ling, X.; Deng, W.; Gu, C.; Zhou, H.; Li, C.; and Sun, F. 2017. Model ensemble for click prediction in bing search ads. In *Proceedings of the 26th international conference on world wide web companion*, 689–698.

- Liu, J.; Xue, M.; Lou, J.; Zhang, X.; Xiong, L.; and Qin, Z. 2023. MUter: Machine Unlearning on Adversarially Trained Models. In *ICCV*, 4869–4879.
- Lu, Q.; Pan, S.; Wang, L.; Pan, J.; Wan, F.; and Yang, H. 2017. A Practical Framework of Conversion Rate Prediction for Online Display Advertising. In *ADKDD@KDD*, 9:1–9:9.
- Ma, X.; Zhao, L.; Huang, G.; Wang, Z.; Hu, Z.; Zhu, X.; and Gai, K. 2018. Entire Space Multi-Task Model: An Effective Approach for Estimating Post-Click Conversion Rate. In *SIGIR*, 1137–1140.
- Nocedal, J.; and Wright, S. J. 1999. *Numerical optimization*. Springer.
- Pruthi, G.; Liu, F.; Kale, S.; and Sundararajan, M. 2020. Estimating Training Data Influence by Tracing Gradient Descent. In *NeurIPS*.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The Annals of Mathematical Statistics*, 400–407.
- Saito, Y.; Morishita, G.; and Yasui, S. 2020. Dual Learning Algorithm for Delayed Conversions. In *SIGIR*, 1849–1852.
- Song, W.; Shi, C.; Xiao, Z.; Duan, Z.; Xu, Y.; Zhang, M.; and Tang, J. 2019. AutoInt: Automatic Feature Interaction Learning via Self-Attentive Neural Networks. In *CIKM*, 1161–1170.
- Wang, H.; Wu, Z.; and He, J. 2024. FairIF: Boosting Fairness in Deep Learning via Influence Functions with Validation Set Sensitive Attributes. In *WSDM*, 721–730.
- Wang, R.; Fu, B.; Fu, G.; and Wang, M. 2017. Deep & cross network for ad click predictions. In *Proceedings of the ADKDD'17*, 1–7.
- Wang, R.; Shivanna, R.; Cheng, D. Z.; Jain, S.; Lin, D.; Hong, L.; and Chi, E. H. 2021. DCN V2: Improved Deep & Cross Network and Practical Lessons for Web-scale Learning to Rank Systems. In *WWW*, 1785–1797.
- Wang, Y.; Sun, P.; Zhang, M.; Jia, Q.; Li, J.; and Ma, S. 2023. Unbiased Delayed Feedback Label Correction for Conversion Rate Prediction. In *KDD*, 2456–2466.
- Wu, J.; Yang, Y.; Qian, Y.; Sui, Y.; Wang, X.; and He, X. 2023. GIF: A General Graph Unlearning Strategy via Influence Function. In *WWW*, 651–661.
- Yang, J.; Li, X.; Han, S.; Zhuang, T.; Zhan, D.; Zeng, X.; and Tong, B. 2021. Capturing Delayed Feedback in Conversion Rate Prediction via Elapsed-Time Sampling. In *AAAI*, 4582–4589.
- Yang, Z.; Wu, J.; Luo, Y.; Zhang, J.; Yuan, Y.; Zhang, A.; Wang, X.; and He, X. 2023. Large language model can interpret latent space of sequential recommender. *arXiv preprint arXiv:2310.20487*.
- Yasui, S.; and Kato, M. 2022. Learning Classifiers under Delayed Feedback with a Time Window Assumption. In *KDD*, 2286–2295.
- Yasui, S.; Morishita, G.; Fujita, K.; and Shibata, M. 2020. A Feedback Shift Correction in Predicting Conversion Rates under Delayed Feedback. In *WWW*, 2740–2746.
- Yoshikawa, Y.; and Imai, Y. 2018. A Nonparametric Delayed Feedback Model for Conversion Rate Prediction. *CoRR*, abs/1802.00255.
- Yu, J.; Zhu, H.; Chang, C.; Feng, X.; Yuan, B.; He, X.; and Dong, Z. 2020. Influence Function for Unbiased Recommendation. In *SIGIR*, 1929–1932.
- Zhang, Y.; Hu, Z.; Bai, Y.; Feng, F.; Wu, J.; Wang, Q.; and He, X. 2023. Recommendation Unlearning via Influence Function. *CoRR*, abs/2307.02147.