Lambda Learner: Fast Incremental Learning on Data Streams

Rohan Ramanath, Konstantin Salomatin, Jeffrey D. Gee, Kirill Talanin

Onkar Dalal, Gungor Polatkan, Sara Smoot, Deepak Kumar

{rramanath,ksalomatin,jgee,odalal,gpolatkan,ssmoot,dekumar}@linkedin.com

LinkedIn Corporation, USA

ABSTRACT

One of the most well-established applications of machine learning is in deciding what content to show website visitors. When observation data comes from high-velocity, user-generated data streams, machine learning methods perform a balancing act between model complexity, training time, and computational costs. Furthermore, when model freshness is critical, the training of models becomes time-constrained. Parallelized batch offline training, although horizontally scalable, is often not time-considerate or cost-effective. In this paper, we propose Lambda Learner, a new framework for training models by incremental updates in response to mini-batches from data streams. We show that the resulting model of our framework closely estimates a periodically updated model trained on offline data and outperforms it when model updates are time-sensitive. We provide theoretical proof that the incremental learning updates improve the loss-function over a stale batch model. We present a large-scale deployment on the sponsored content platform for a large social network, serving hundreds of millions of users across different channels (e.g., desktop, mobile). We address challenges and complexities from both algorithms and infrastructure perspectives, illustrate the system details for computation, storage, stream processing training data and open-source the system.

CCS CONCEPTS

• **Theory of computation** → *Streaming models*; • **Information systems** → *Personalization*; • **Computing methodologies** → *Online learning settings*.

KEYWORDS

Mixed-effect Models; Incremental Online Learning; Recommender Systems; Personalization; Reinforcement Learning

ACM Reference Format:

Rohan Ramanath, Konstantin Salomatin, Jeffrey D. Gee, Kirill Talanin and Onkar Dalal, Gungor Polatkan, Sara Smoot, Deepak Kumar. 2021. Lambda Learner: Fast Incremental Learning on Data Streams. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3447548.3467172>

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.

KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467172>

1 INTRODUCTION

Content in a professional social network is distributed across the diverse member experiences the site has to offer, from advertising and recruiting, to job and knowledge seeking. Each experience depends on a suite of machine learned models, which rank and curate content to create quality experiences, typically measured by engagement. Since quality experiences are intrinsically member-specific, personalized models are more performant in most settings.

A common challenge for recommender systems is to achieve *memorization* and *generalization*. Most state-of-the-art models [6, 23] address generalization via embedding based methods [7, 20] or other nonlinear models and capture memorization via interaction features, ids of the user/item or a crossing network. The memorization piece is important because it is difficult to learn effective low-dimensional representations for users and items when the underlying user-item matrix is sparse and high-rank. Since the features used in the memorization part of the model could be both transient and ever growing, being able to update this part of the model is critical to performance. For example, a new advertising campaign can reach thousands of impressions in a few minutes. Likewise, a job posting can reach millions of job seekers in a similarly short time window. In these cases, recommendation models need to learn from the information as quickly as possible. Within the context of computational advertising, we specifically observe the following challenges when estimating viewer engagement:

- **Cold start:** Over the course of a day, a model needs to make decisions on advertisers and advertisements not observed in the training set. These can be brand-new advertisers and advertisements, or advertisers who have not marketed through the social network for extended periods of time. Under these scenarios, the memorization piece of the model must quickly adapt to new data.
- **Data shifts:** Viewers' most recent engagements are the best interest indicators, as interest can dramatically shift over short periods of time. For example, a viewer who might be heavily engaging with advertisements about cloud services on the site may be less-inclined to engage immediately after signing up with some of the cloud services.

Although the prior work [6, 7, 20, 23] acknowledges the need to re-train the overall model frequently (Cheng et al. [6] recommends warm start based fine-tuning and empirically checking performance as a sanity check), there is no model agnostic solution with theoretical guarantees of performance improvement. One class of models that have been successful for personalized recommendations & ranking is Generalized Additive Mixed Effect models (GAME)[26]. This consists of a typically large fixed-effects model (*generalization*) that is trained on the whole dataset to improve the model's performance on previously unseen user-item pairs, and a series of

simpler linear random-effects models (*memorization*) trained on data corresponding to each entity (e.g. user or article or ad) for more granular personalization. The fixed effects part of the model can be arbitrarily complex (e.g. wide & deep, boosted trees, ensembles, or linear [26]^{*}) depending on the targeted task. The entity specific random-effect models (detailed in Section 2.2) are a powerful way to personalize both user and item (content/ad) patterns.

The most general solution—frequent offline model retraining and deployment—often proves insufficient for the following reasons: (1) to model newly introduced content or content with short-lived significance, computation time becomes a valuable commodity, and limiting it causes a computationally-constrained learning problem [3], i.e. the computational constraints introduced to retrain the model disallow us from meeting self-imposed update frequencies; (2) the cold-start problem can cause long-term repercussions for the site, i.e. poor recurring recommendations may influence future engagement for new or returning dormant members. Additionally, the model performance slowly degrades over time, demonstrating the need to update the model continuously.

In this paper, we propose a framework to incrementally train the memorization part of the model as a booster over the generalization part. We incrementally update the memorization model between full batch offline updates of the generalization model to balance training performance against model stability. We show that for a scenario that suffers from the aforementioned cold-start, incremental training enables quicker adaptation to new and short-lived data, and reduces the rate of model degradation. The key concept is to prepare mini-batches of data from an incoming stream of logged data to train an incremental update to the memorization model. We approximate the loss function on previously observed data by a local quadratic approximation around the previous optimal value and combine it with the loss function on the most recent mini-batch. This allows us to incrementally update the model without relying on all historical data yet do better than just relying on the new incoming data. This results in a regularized learning problem with a weighted penalty. In the Bayesian interpretation, the local quadratic approximation is equivalent to using the posterior distribution of the previous model as prior for training the update.

1.1 Related Work

Similar problems have been studied from various angles. Gepperth and Hammer [11] survey incremental learning under streaming data, giving an overview of popular approaches, foundations, and applications. Ditzler et al. [10] provide a survey on non-stationary scenarios in which modeling frameworks fail to address evolving or drifting dynamics (e.g., seasonality or periodicity effects, the user behavioral changes or preferences). Bayesian, stochastic, tree-based, and instance-based learning approaches have been explored in the online scenario with varying degrees of success on both stationary and non-stationary data [17]. For potentially non-stationary data, it has been shown that Stochastic Gradient Descent (SGD) on a linear model best balanced performance and growth in model complexity [17]. Bottou [3] demonstrated that SGD in very large datasets out-performed more modern approaches due to the speed of convergence (to a predefined “expected risk”). It was also found that

second-order SGD (2SGD) in a single-pass, under convexity and metric regularity assumptions was able to approximate iterative, full-batch SGD [3].

From an optimization point of view, McMahan [18] investigates different strategies for online convex optimization and finds that several such algorithms are actually highly related, being distinguishable based on how they handle $L1$ regularization. d’Alché Buc and Ralaivola [8] focus on a local strategy to avoid repeated learning of all parameters. The proposed algorithm selects a local subset to run incremental learning. Wu et al. [24] study model performance degradation due to missing past data and recommends using a data bias elimination mechanism to make incremental learning methods more effective.

From the representation learning point of view, Yang et al. [25] study incremental learning for deep models, focusing on flexible model evolution and updating models such that previous knowledge is not forgotten. They propose using a hidden attention layer. Song et al. [21] consider the problem from collaborative member behavior and propose an item clustering-based system. Similarly, Dhurandhar and Petrik [9] approach the problem of incremental feature learning for GLMs where a snapshot of batch data is available. Though such approaches might fit more static content libraries, they do not suit more dynamic environments, such as an advertising platform in which decent relevance accuracy should be accomplished quickly once content is added to the ecosystem. Online learning for recommending content from large and dynamic pools is studied by Agarwal et al. [1]. A factor model is proposed to fit content-specific factors quickly using online regression. Though the online learning portion has similarities to our work, the problem does not consider large co-variate space. Lastly, a family of reinforcement learning algorithms could be used for incrementally updating recommender systems (by changing the objective function), Chen et al. [5] propose a model applied to the candidate generation phase. Our work is strictly focused on the ranking phase, assuming independence of the candidate selection generation methodology, and fixing the objective function used to rank the ads.

1.2 Our Contributions

In this paper, our contributions are in four fold: 1) We propose a new framework for incremental learning for personalization at scale. We perform model adjustments over the most recent data observations and adapt to the latest trends, additionally training new models for content unobserved in the training set. 2) We provide theoretical guarantees for improvements in loss-function for the incrementally updated model over the stale batch model. 3) We conduct experiments showing that incrementally updated models in a single pass over a data stream can closely approximate full batch offline models. 4) Finally, we present a large-scale deployment and impact on the ads platform at a large social network, address infrastructure challenges, and open-source the system[†].

1.3 Terminology and Notation

Throughout this document we will commonly reference the following notation. We refer to a training dataset $\Omega^{(t)}$, as the mini-batch/slice of ordered “data instances” observed in time interval

^{*}GLMix is an instantiation of GAME with a linear fixed-effect model

[†]<https://github.com/linkedin/lambda-learner>

Figure 1: An illustration of the incremental learning process

t . Each $\Omega^{(t)}$ consists of data instances represented by the tuple $(\mathbf{x}_n, y_n)$, where $\mathbf{x}_n$ is an observed feature vector describing instance n , and y_n is a corresponding observed label. We denote the cumulative training dataset observed until time t using $\bar{\Omega}^{(t)} = \bar{\Omega}^{(t-1)} \cup \Omega^{(t)}$. We use $\beta^{*(t)}$ to refer to a parameterized model trained on the full dataset $\bar{\Omega}^{(t)}$, and $\beta_{inc}^{(t)}$ to refer to the incremental model trained using the new Lambda Learner framework.

In the GAME framework, the model $\beta = [\beta_f, \mathbf{B}]$ consists of two components: a fixed-effects model, β_f , and a set of random-effect models, $\mathbf{B} = \{B_r\}_{r \in \mathcal{R}}$, where $\mathcal{R}$ denotes the set of random-effect types, see section 2.2 for details. Coefficients of a single random-effect type r are denoted by $B_r = \{\beta_{rl}\}_{l=1}^{N_r}$, where N_r is the total number of random-effects with random-effect type, r (e.g. r may correspond to “viewers”, “articles”, “ads” or “jobs”; l is a specific viewer, article, ad or job ID, and N_r is the cardinality of viewers, articles, ads or jobs). Lastly, for each data instance $\mathbf{x}_n, y_n$, we use $i(r, n)$ to represent the index value of random-effect type r for instance n .

2 BACKGROUND

In this section, we review a high-level foundation of personalization through random effect models and incremental learning.

2.1 Learning on Data Streams

Given a stream of observation data $\Omega^{(0)}, \Omega^{(1)}, \dots, \Omega^{(t)}, \dots$, and a base convex loss function f (e.g. log-logistic), we define a series of loss functions over the mini-batch of datastream $\Omega^{(t)}$ as

$$f_t(\beta) = \sum_{n \in \Omega^{(t)}} f(\mathbf{x}_n, y_n, \beta), \quad (1)$$

and a time-discounted cumulative loss over the total data $\bar{\Omega}^{(t)}$

$$F_t(\beta) = \sum_{k=0}^{t-1} \delta^{t-k} f_k(\beta) = \delta F_{t-1}(\beta) + f_t(\beta),$$

where, $\delta < 1$ is a time discount or forgetting factor to emphasize the fresher data.

2.2 Generalized Additive Mixed Effect Models

2.2.1 Generalized Linear Mixed Models. Zhang et al [26] proposed Generalized Linear Mixed Models (GLMix) as an approach for personalizing traditional Generalized Linear Models (GLM’s) through the addition of linear, “random-effect” models, keyed to discretely valued random-effect variables. Example random-effect variables range from unique entity identifiers (e.g., advertisement, job, or user) to contextual information such as the consumption device

(e.g., desktop, mobile application, etc). In practice, the coefficients fit conditionally on the values of these variables and serve a supplementary role in providing personalization atop the GLM.

We solve for the fixed-effect coefficients (unconditional part of the model), β_f , and the random-effect coefficients, $\mathbf{B}$, through the optimization of negative log likelihood objective

$$\underset{\beta=[\beta_f, \mathbf{B}]}{\text{minimize}} \sum_{n \in \Omega} -\log p(y_n | s_n) - \log p(\beta_f) - \sum_{r \in \mathcal{R}} \sum_{l=1}^{N_r} \log p(\beta_{rl}) \quad (2)$$

where the first summation term is data likelihood and the score s_n is a sum of fixed-effect and random-effect scores for the n -th observation:

$$s_n = \mathbf{x}_n^T \beta_f + \sum_{r \in \mathcal{R}} \mathbf{z}_{rn}^T \beta_{r, i(r, n)}, \quad (3)$$

$\mathbf{x}_n$ is the feature vector for the fixed-effect of observation n , $\mathbf{z}_{rn}$ is the feature vector for random-effect type r of n , and $\beta_{r, i(r, n)}$ is the coefficients vector conditional on the random-effect variable value $i(r, n)$. For example, in a user-item recommendation model with per-user and per-item random effects: $r \in \{\text{“user”, “item”}\}$ and $i(\text{“user”, } n)$ is simply the id of the user in the observation n , and $i(\text{“item”, } n)$ - id of the item. The latter terms $\log p(\beta_f)$ and $\log p(\beta_{rl})$ in (2) represent the prior of the model coefficients, in particular, a Gaussian prior $\mathcal{N}(0, \lambda I)$.

2.2.2 Generalized Additive Mixed Effect. Although the GLMix formulations in (2) and (3) describe fixed-effect and random-effects as linear models, they are not limited to such. By substituting s_n for:

$$s_n = g_f(\mathbf{x}_n, \beta_f) + \sum_{r \in \mathcal{R}} g_r(\mathbf{z}_{rn}, \beta_{r, i(r, n)}), \quad (4)$$

we define the Generalized Additive Mixed Effect (GAME) model class, where g_f and g_r are arbitrary scoring functions. For example, g_f can be a deep neural network or a bag of decision trees. GLMix is a purely linear example of GAME.

2.3 Incremental Learning

Incremental Learning (IL) describes the general set of algorithms attempting to solve (often) large-scale machine learning through immediate model updates in response to new data instances (Figure 1), without additional passes over pre-existing data. These scenarios occur when training models on large streams of data, where there is a limited time window for models to adjust, constraining computational cost of the learning process. Model update rates can be controlled over batch sizes of new observations (incremental batch learning), with a batch size of 1 corresponding to online learning.

3 LAMBDA LEARNING

In this section, we present our incremental learning framework. Our approach provides an efficient mechanism for updating the random-effects model on fresh data flows. One key aspect is the parallelization of random-effect model training, which is proved possible when data instances are grouped by random-effect types [26].

3.1 Formulation for Incremental Updates

For every batch of data $\Omega^{(t)}$, we want to solve the problem

$$\underset{\beta}{\text{minimize}} \quad F_t(\beta) - \log p(\beta), \quad (5)$$

where $p(\beta)$ is the prior distribution of the model parameters β . Assuming a Gaussian prior $p(\beta) \sim \mathcal{N}(0, \lambda I)$, as in the case of log-logistic loss, the output of the problem is a Gaussian posterior distribution $p_t(\beta) \sim \mathcal{N}(\beta^{*(t)}, \Sigma^{*(t)})$ (using Laplace approximation).

However, as the data grows, it becomes unfeasible to store in memory, and too computationally expensive to completely iterate over to compute F_t , ∇F_t , and/or $\nabla^2 F_t$. Therefore, we maintain an approximation of the loss function over the past observations, F_{t-1} , and combine it with the loss function over the latest batch f_t . In particular, we decompose (5) into a loss on past data and a loss on an incremental batch,

$$\underset{\beta}{\text{minimize}} \quad \delta F_{t-1}(\beta) + f_t(\beta) - \log p(\beta), \quad (6)$$

We substitute the first term of the objective function though approximation using a quadratic function centered around previous optimum $\beta^{*(t-1)}$,

$$\underset{\beta}{\text{minimize}} \quad \frac{\delta}{2} \|\beta - \beta^{*(t-1)}\|_{\nabla^2 F_{t-1}(\beta^{*(t-1)})}^2 + f_t(\beta) + \frac{\lambda}{2} \|\beta\|^2. \quad (7)$$

Note that optimality of $\beta^{*(t-1)}$ implies $\nabla F_{t-1}(\beta^{*(t-1)}) = 0$ (hence the linear term is dropped).

In the Bayesian framework, since $\left(\Sigma^{*(t-1)}\right)^{-1} = \nabla^2 F_{t-1}(\beta^{*(t-1)})$, this approximate problem is equivalent to training the mini-batch using the posterior distribution, $\mathcal{N}(\beta^{*(t-1)}, \Sigma^{*(t-1)})$, from the previous step as a prior distribution for $\beta^{(t)}$.

3.2 Asynchronous Updates

Additionally, in the GAME framework, we can use asynchronous updates for the random-effect models to keep them fresh in production systems. We control the rate at which a random-effect model is updated by a mini batch-size parameter. Once a sufficiently-sized batch of data $\Omega_{\bar{r}\bar{l}}^{(t)} = \{n \in \Omega^{(t)} | i(\bar{r}, n) = \bar{l}\}$, for random-effect $\bar{r}$ and index $\bar{l}$, has been observed, we update the corresponding random-effect by solving:

$$\begin{aligned} \underset{\beta_{\bar{r}\bar{l}}^{(t)}}{\text{minimize}} \quad & \sum_{n \in \Omega_{\bar{r}\bar{l}}^{(t)}} -\log \left(y_n | \zeta_n^{(t-1)} + \mathbf{z}_{\bar{r}n}^T \beta_{\bar{r}\bar{l}}^{(t)} \right) - \delta \log p(\beta_{\bar{r}\bar{l}}^{(t)}) \\ \zeta_n^{(t-1)} = & \mathbf{x}_n^T \beta_f^{*(0)} + \sum_{s \in \mathcal{R}, s \neq \bar{r}} \mathbf{z}_{sn}^T \beta_{s, i(s, n)}^{(t-1)}, \end{aligned} \quad (8)$$

and the posterior distribution for $\beta_{\bar{r}\bar{l}}^{(t-1)}$, i.e., the Gaussian distribution $\mathcal{N}(\beta_{\bar{r}\bar{l}}^{*(t-1)}, \Sigma_{\bar{r}\bar{l}}^{*(t-1)})$ is used as the prior distribution $p(\beta_{\bar{r}\bar{l}}^{(t)})$.

We keep the fixed-effect model, $\beta_f^{*(0)}$, stationary throughout the incremental training process and use the updated model $\beta^{*(t)} = [\beta_f^{*(0)}, \{\beta_r^{*(t)}\}_{r \in \mathcal{R}}]$ to serve traffic after time t . The fixed-effect model ($\beta_f^{*(0)}$), typically an ensemble of embedding based models and boosted forests, are continuously retrained offline.

3.3 Algorithm Details

Algorithm 1 describes the complete algorithm. We start with large volumes of historical engagement data to train, in offline fashion, a fixed-effects (ensemble) model as the initial batch of random-effects models. This produces fixed-effects coefficients (*generalization* part of the model), $\beta_f^{*(0)}$, as well as a starting point $\beta_r^{*(0)} = \beta_{r, inc}^{*(0)}$ for incrementally learning the random-effect (*memorization* part of the model) coefficients. This is also used to set the initial Hessian matrix, $H_r^{(0)} = \nabla^2 F_0(\beta_r^{*(0)})$ as well as the time discount $\delta < 1$. Every time the fixed-effects models are retrained offline, all of the above parameters are asynchronously updated from outside the system.

Now, as the system processes the next mini-batch of data, it uses the latest coefficients from the fixed-effect model to compute an offset for each training example. Note the same stationary $\beta_f^{*(0)}$ for all t . The offset provides the connection of the fixed-effect models with the random-effects that are to be updated while allowing different model families to be used for generalization and memorization. We then train the random-effect coefficients ($\beta_{r, inc}^{*(t)}$) using the current batch of data while constraining the solution to be “near” the previous optimum (line 13 of Algorithm 1). In practice, the entire training loop for each random-effect executes in under a second. The new coefficients are used to update the running Hessian estimate and the posterior covariance. The model coefficients, $\beta^{(t)} = [\beta_f^{*(0)}, \beta_r^{(t)}]$, are then pushed to stores for online serving. Since we have access to the posterior covariance, we rely on $\beta_f^{*(0)}$ and perform Thompson Sampling from the updated distribution for $\beta_r^{(t)}$ to compute scores online.

Connection to Reinforcement Learning: The ad ranking problem is often cast as an instance of a contextual multi-armed bandit problem [16]. Each of the possible ads is represented by an arm ($r \in \mathcal{R}$). The action space (arms) is dynamic as ads get introduced into the system when they are created and retired as the budget runs out. The context for opportunity t is encoded as a feature vector $\mathbf{x}_t$. The policy uses the context to make a decision at every opportunity (a member request) to serve an ad. Once a decision is made an observation is collected, and the corresponding reward computed; e.g., member clicked on the ad. The policy is then revised with the data collected for this opportunity. An effective way to solve this bandit problem is via Thompson sampling [4, 22]. Since our algorithm produces the posterior covariance as an output, the system behaves like an online contextual bandit. As new arms get introduced, the system is able to estimate to value of each arm much faster because of the fast feedback loop. Additionally, reducing the number of high variance opportunities for new ads improves the auction efficiency and overall revenue.

4 THEORY

In this section, we give theoretical guarantees of why incremental training improves model performance. To start, we prove that the first incremental step at $t = 1$ is better and then extend it to subsequent steps $t = 2, 3, \dots$

4.1 First incremental update ($t = 1$)

Starting with an exact solution $\beta^{*(0)}$ trained on the initial batch $\Omega^{(0)}$, we make an incremental update based on the first mini-batch of observations $\Omega^{(1)}$.

THEOREM 4.1. *For a convex loss function f , the first incremental step update $\beta_{inc}^{*(1)}$ obtained by solving problem (7) for $t = 1$, i.e.,*

$$\beta_{inc}^{*(1)} = \underset{\beta}{\operatorname{argmin}} \quad f_1(\beta) + \frac{\delta}{2} \|\beta - \beta^{*(0)}\|_{\nabla^2 F_0(\beta^{*(0)})}^2, \quad (9)$$

improves the loss on $\bar{\Omega}^{(1)}$, i.e. $F_1(\beta)$, compared to the previous batch update $\beta^{(0)}$. In other words,*

$$F_1(\beta^{*(0)}) - F_1(\beta_{inc}^{*(1)}) \geq \frac{\lambda}{2} \|\beta^{*(0)} - \beta_{inc}^{*(1)}\|_2^2 + \mathcal{O}(\gamma^3), \quad (10)$$

where, $\gamma = \|\beta_{inc}^{(1)} - \beta^{*(0)}\|_2$.*

4.2 Subsequent incremental updates ($t = 2, \dots$)

We begin by defining a series of incremental objective functions (and respective optimas as incremental updates) similar to (??),

$$F_t^{inc}(\beta) = f_t(\beta) + \frac{1}{2} \|\beta - \beta_{inc}^{*(t-1)}\|_{\nabla^2 F_{t-1}^{inc}(\beta_{inc}^{*(t-1)})}^2, \quad (11)$$

and show that these incremental updates trained on mini-batches are better using the stale model parameters for any t .

THEOREM 4.2. *For a convex loss function f , for $t = 2, 3, \dots$, the series of incremental updates $\beta_{inc}^{*(t)}$ obtained by solving problem (11) for $t = 2, 3, \dots$, improves the loss compared to the stale batch update $\beta^{*(0)}$. In other words,*

$$F_t(\beta^{*(0)}) - F_t(\beta_{inc}^{*(t)}) \geq \frac{\lambda}{2} \|\beta^{*(0)} - \beta_{inc}^{*(t)}\|_2^2 + \mathcal{O}(\bar{\gamma}^3), \quad (12)$$

where $\bar{\gamma} = \max \left\{ \|\beta_{inc}^{(k)} - \beta_{inc}^{*(t)}\|_2, \|\beta_{inc}^{*(k)} - \beta^{*(0)}\|_2 \right\}$, for $0 \leq k \leq t$.*

Note: The optimality gap between the loss-function values depend on the distances between $\beta^{*(0)}$ and $\beta_{inc}^{*(t)}$, which in turn depend on the drift in distributions of the data set between $\Omega^{(0)}$ and subsequent $\Omega^{(t)}$'s.

5 SYSTEM OVERVIEW

This section describes a baseline system capable of offline retraining and online scoring, and how we augment it to support high-frequency nearline updates. Any system using the GAME paradigm can adopt our solution and benefit from improved model freshness, and resulting lifts to publisher ROI and content engagement.

5.1 Motivation

We would ideally like to update a time-sensitive model as soon as new observations are available, with zero delay. A typical system, where offline batch updates occur with period T , is insufficient. The data size itself increases with T , since a new batch of data

accumulates while the model is trained on the previous batch. In practice, we observe T can be 8-24 hours due to the time needed to prepare a new training data batch and retrain on this batch.

Such offline batch retraining is common practice for GAME (fixed + random effects) models (7). The promise of *Lambda Learning* is to reduce the time to adapt to new observations by processing data incrementally and learning on sequences of mini-batches nearline. This is a significant step towards the ideal of immediate model updates.

5.2 Overview of sponsored content ad system

Figure 2 illustrates a baseline system capable of offline retraining and online scoring, and an extension, shown in green, which enables high-frequency nearline model updates. The components we introduce are: feature tracking, nearline training data generation, fixed-effect model scoring, and nearline random-effects model training (*lambda learning*).

Our solution operates in the context of a large sponsored content advertising system (A5 in Fig. 2). A machine learned response prediction model plays a central role, scoring content by estimating user engagement. The base production system can be logically separated into two phases: (1) a cold start phase where the fixed-effects part of the GAME model is trained (A in Fig. 2), and (2) a warm start phase where the random-effects coefficients are trained, boosting the scores of the fixed-effects model (B in Fig. 2). The fixed-effects part is typically a large non-linear model, trained on user attributes, activity, and contextual cues, and retrained periodically. The random-effects capture information specific to the ad, campaign, and advertiser. Random-effect models are small and linear, allowing modeling control at different levels of granularity. Random-effects of a given effect-type can be trained independently, allowing us to train only those for which there is new data (e.g. new ad content). They are also time-sensitive (especially for previously unseen advertisers/creatives), and are most effective when retrained frequently. See [2] for more background on this paradigm.

5.3 Prerequisite system components

The baseline system is capable of offline retraining and online serving:

5.3.1 Offline Batch Model Training for $\Omega^{(0)}$. This phase involves training an initial GAME model, $\beta^{*(0)} = [\beta_f^{*(0)}, \{\beta_r^{*(0)}\}_{r \in \mathcal{R}}]$. In our case, we train a GAME model $\Omega^{(0)}$ over a training window of D days (typically 14-30 days, depending on naturally occurring periodicity). Since $\beta_f^{(0)}$ represents a generalized fixed-effects model learned across a complete dataset, it is tolerant to minor data shifts and can be updated less frequently (in our case, once every d days). However, since new items are constantly added to the system and item-specific behaviors are fast-changing, $\{\beta_r\}_{r \in \mathcal{R}}$ must be updated frequently, every h hours (e.g. 6-8 hours) to capture these changes.

Step 1: For each training cadence of d days, we start fresh and train $\beta^{*(0)} = [\beta_f^{*(0)}, \{\beta_r^{*(0)}\}_{r \in \mathcal{R}}]$ on the observations, $\Omega^{(0)}$, collected over training window T by solving Equation (2).

Figure 2: System design for incremental nearline model updates. Nearline components are highlighted in green.

5.3.2 *Online model serving (A4 in Fig. 2).* The random-effects are trained on varying volumes of training data, resulting in different covariance estimates for the β_r 's. We use Thompson Sampling to incorporate these covariance estimates in online inference, enabling exploration of low-mean, high-variance ads [2].

Step 2: For each of the random-effects coefficients, we use the posterior mean $\beta_r^{*(0)}$, and the variance $\Sigma_r^{*(0)}$ to sample a β_r for online inference and serving.

5.4 System components for Lambda Learning

To extend the baseline system to support frequent nearline model updates, we add four components: feature tracking, nearline training data generation, fixed-effect model scoring, and nearline random-effect model training (*lambda learning*). These communicate via Kafka topics partitioned by random-effect id. Horizontal scaling is achieved by increasing the number of partitions and machines running these components.

5.4.1 *Feature Tracking (C1 in Fig. 2).* Common practice in model training pipelines is to join labeled data instances (“observations”) with offline stores of uniquely keyed features (e.g. member features, ad features, etc.) to produce training data, Ω . The trained models are then applied on-demand in an online serving layer, where “feature producer” functions produce a real-time representation of $\mathbf{x}_n$ prior to being scored with the model. However, we observe drawbacks to this approach: (1) differences between offline/online feature production (e.g. due to bugs, versioning, unavailability, etc), and (2) time-sensitivity of certain feature sets. Training pipelines that rely on the present value of time-sensitive features (e.g. activity-based features) can cause overfitting, and, in extreme instances, introduce label information into training data.

As a solution, we introduce an online feature tracking system which records feature values used in online scoring. This is implemented using Kafka [15], an distributed messaging system, which allows services to asynchronously share data in near real-time. This solves the issue of feature time-sensitivity, since feature values are computed at the same time as the observations from which we create training labels. For online services that already score and fetch features on-demand, tracking these features takes little extra work. For efficiency, we track only time-sensitive features, joining on static features in the next component.

5.4.2 *Training Data Generation (C2 in Fig. 2).* Kafka provides topics such as impressions, clicks, and tracked features. We use Samza [19] to listen to these topics and produce featurized training examples, which are published to a new Kafka topic. This data is used by subsequent nearline training, but is also ETLED offline for use in offline training. Online/offline feature parity is guaranteed. Features that are fairly static, but dense (e.g. embeddings of static entities) are good candidates to be introduced at this stage since it not cost efficient to serialize highly redundant, dense features.

Two important knobs that are fine-tuned by analysing traffic patterns for data quality are:

- **Trade-off: training cost and accuracy:** Since offline data pipelines have access to “delayed” clicks when preparing training data, all clicks are accounted for. In the streaming case, the use of windowing means “late” click may be missed. Longer windows mean higher cost, but lower rate of misses.
- **Partition for throughput:** Topics are partitioned based on random-effect id. This avoids network shuffle, and ensures all data for a given advertiser is processed on the same node.

5.4.3 *Fixed Effect Model Scoring (C3 in Fig. 2).* This Samza processor listens to a stream of training data and scores each incoming instance with the fixed-effects model parameterized by $\beta_f^{(0)}$. If the fixed-effects model is an ensemble, the entire ensemble is scored just like it would be online to compute a score for every impression. The score serves as the offset for random-effects training as explained in Algorithm 1.

Step 3: Compute the score from the fixed-effects model, $g_f(\mathbf{x}_n, \beta_f^{(0)})$.

5.4.4 *Lambda Learning (C4 in Fig. 2).* The random-effects coefficients produced by offline training are served online via a key-value store. The last step, Lambda Learning, involves iterative retraining of these using the nearline data. This component groups stream data into mini-batches based on random-effect id, fetches the corresponding random-effects coefficients from the store, updates them, and writes them back, ready for online scoring.

In practice, the velocity of incoming data streams can fluctuate throughout a day. Therefore, mini-batch updates can be triggered by meeting one of many conditions, such impression count (high velocity periods), time since last update (low velocity periods), or coefficient variance (new ids).

Step 4: For each mini-batch of data $\Omega^{(t)}$, we solve the inner loop of Algorithm 1 to update $\beta_r^{*(t)}$ and $\Sigma_r^{*(t)}$.

Several subtleties arise in implementing the read-train-write (RTW) part of this component:

- **Concurrency:** While different random-effects can be trained in parallel, RTWs for a given random-effect must be sequential. We enforce this by making each RTW hold a lock unique to the random-effect being retrained.
- **Weak store consistency:** Derived data key-value stores may offer weak consistency guarantees to meet strict latency or throughput SLAs. In particular, the store we use doesn't guarantee read-your-write semantics. We've found that placing a TTL cache between the processor and the store solves this issue.
- **Replay after batch updates:** The store supports both batch writes and nearline streaming writes. When a batch write occurs it is based on hours-old data. Replaying the training data from the last few hours on top of the new batch coefficients ensures coefficients reflect the most recent data.

6 EXPERIMENTS

We evaluate our work across two datasets, focusing on the performance comparison between the proposed incremental-online updates, and a traditional iterative, batched-offline approach. Our experiments, though conducted offline using static datasets, consider the online data stream from which they were sourced. The purpose of the experiments is to measure the long-term stability of incremental updates, as well as the immediate impact of including near real-time data in our models. Our hypothesis is that the incremental approach's increased access to recent data will increase its predictive power. This data advantage is proportional to the time-costs associated with iterative-offline training, and therefore, we believe incremental updates will be most effective as offline training time increases. However, we also hypothesize that prolonged incremental updating without iterative-offline training can compound the error associated with the incremental loss-function.

6.1 Datasets

We use two datasets: a public “movie lens 20M” dataset [13] for reproducibility and a proprietary sponsored updates dataset which is derived from the system we have built and used to train the ads response prediction models. All of our findings are consistent between the two datasets, so we present the detailed discussion of the results on the ads dataset. The corresponding results on the movie-lens dataset are in Appendix C.

6.1.1 Social network ad click data. We collected a small number of weeks of ad click and impression data from a large social network. This time frame is sufficient to simulate life-cycle of an offline model and incremental lambda-learner model. We are omitting details about the features and transformations used for confidentiality reasons. However, the dataset includes sufficient information about the advertisements themselves and the audience of the impressions.

6.1.2 Movie Lens. Movie lens 20M [13] is a benchmark dataset with 20 million ratings given to 27,000 movies by 138,000 users. We

Figure 3: Illustrates how models are evaluated. Timeline $[0, t)$ represents a time-ordered dataset. Δ represents the size of an evaluation set after each model update, Ω_{eval} , and τ is an estimated delay incurred by full-batch training. β_{NU} , β_{RWBU} and β_{IBU} represent the parameterized models produced using full-batch training, and β_{LL} represents parameterized models produced using incremental online updates (shown in waves). Incremental updates are executed every Δ , and total number of incremental updates is given by $T/\Delta - 1$

use this dataset for reproducibility of results and apply transformations to make the setup more analogous to that of our real-world “sponsored updates” dataset: 1) We binarize responses, using scores greater or equal to 4.0 as positive labels and scores less than 4.0 as negative labels, 2) We use ALS matrix factorization [14] of rank 30 to produce latent movie features used in the model, 3) As the dataset spans 20 years and is not as time sensitive as a typical web recommendation, we speed up time and scale 20 years into 14 days.

6.2 Evaluation Framework

We chronologically order each of the datasets and split them into two parts. The first part of the data, Ω_{cold} , is used to train an initial, baseline model, $\beta^{(0)}$. The latter part, Ω_{warm} , is used to simulate model updates and for evaluation. Let us assume that Ω_{warm} starts at time $t = 0$ and ends at time $t = T$. We break Ω_{warm} into T/Δ increments of equal duration, Δ . The models, described below, are evaluated sequentially on all the increments. After an increment is used for evaluation, it becomes available for training (with a delay τ , representing the cost of batch training and deployment).

Our work is motivated by the computational costs associated with offline iterative training of random-effect models. We compare the iterative approach against batch updates with different training costs, τ . When $\tau = 0$, offline model has no associated cost, and updates happen instantaneously. As $\tau \rightarrow \infty$, offline training becomes impossible, and model updates are infeasible. Considering this, we compare four model update variants (illustrated in Figure 3):

- **No Updates (NU):** Under this strategy, we assume $\tau = \infty$, and no model updates are performed. This means that all increments (starting at 0, Δ , 2Δ ...), are evaluated using the model, $\beta^{(0)}$. This serves as a strict lower-baseline for other variants.
- **Ideal Batch Update (IBU):** To approximate the error of incremental model updates against an upper bound, we compare against an offline batch model trained without compute cost ($\tau = 0$).
- **Real-World Batch Update ($RWBU_\tau$):** We evaluate full-batch updates over different training costs τ . Each increment, $i\Delta$, is evaluated using, $\beta^{(i\Delta-\tau)}$, which is trained offline on all data instances preceding the Ω_{eval} by more than τ .

(a) Change in ROC-AUC, as model update frequency is reduced. The x-axis denote the models being updated every Δ hours - in the incremental using ONLY data in the last Δ hours and in the batch using the entire history including the data in the last Δ hours.

(b) The best LL variant performance is stable and the batch performance degrades when the area under ROC metrics are computed every 15 minutes for a period of 72 hours. 95% confidence intervals included.

Figure 4: Model evaluation using lambda learning on ad click production data. Note: y-axis on both charts is scaled to obfuscate absolute metric values, highest value corresponds to 1.0. Historical weight was fixed, $\delta = 0.95$.

- **Lambda Learner (LL_{Δ}):** We compare against model updates based on the proposed method in Section 3. For LL, data availability is made identical to that of IBU. However, models are updated every Δ . In practice, Δ should be small compared to the batch training cost τ , as comparable or larger values imply updates as infrequent as RWBU.

Updated models are evaluated on each Δ , using the area-under-ROC (AUC) [12] as the metric. We compute AUC on the full Ω_{warm} dataset (i.e. all T/Δ increments).

6.3 Incremental vs. Batch Loss

First, we quantify the drop in performance caused by the approximation of the incremental loss function. We measure this drop at different incremental update frequencies (i.e. different values of Δ , shown in Figure 4a). We report the aggregate metric for all

model updates over Ω_{warm} to remove the effects of seasonality, intra/inter-day trends, and any other idiosyncrasies that may cause inaccuracies over a single window.

Figure 4a compares the performance of LL against the theoretical upper bound (IBU, $\tau = 0$), lower bound (NU) and some real world scenarios (RWBU, $\tau \in \{8, 16, 24\}$). For sake of completeness we show the performance values of all variants for each setting of Δ , however, in an online deployment, it is hard to observe RWBU performance for values of $\Delta \leq 4$ hours. This is because the time taken to finish a full retrain of the batch algorithm (e.g. $\tau = 8$ hours) is likely to exceed this update interval. The performance of the IBU degrades as Δ increases providing evidence for the need to have more frequent model updates. This behavior is typical in systems with dynamic content or when there are temporal patterns of data access. LL demonstrates consistent performance for all update frequencies except $\Delta = 0.5$ hours. We believe that two opposite trends balance each other and result in this stable performance: firstly, less frequent updates increase the delay, a detrimental factor; secondly, with less frequent updates there is less dependence on objective function approximation, a beneficial factor.

The model training and deployment delay (τ) has strong impact on the performance of the batch update algorithm. LL starts to outperform the batch update with the delay of 8 hours (RWBU, $\tau = 8h$) and significantly outperforms the batch update with higher values of delay ($\tau = 16h$ and $\tau = 24h$).

The results and observations discussed in this section for the proprietary sponsored updates dataset are consistent with the results on the public Movie Lens data. We report these results in the Appendix C for reproducibility.

6.4 Time Decay Comparison

In this section we investigate how the performance of competing algorithms decays with time. We expect the batch algorithm to decay between consecutive updates, while LL may accumulate error with every update. The magnitude of this decay determines whether “frequent enough” RWBU updates are sufficient for optimal performance, or if incremental updating can improve the performance between RWBU updates. In particular, we look at the performance of the non-updating batch offline model (NU) and LL model updated every hour ($\Delta = 1$ hour) within 3 days. To remove seasonality factors such as time-of-day and day-of-week we averaged over multiple runs with random start times.

The NU model is trained using data up to time $t = 0$ in Figure 4b and the model performance clearly degrades with time. LL, on the other hand, shows no significant degradation.

6.5 Effect of Forgetting Factor (δ)

The higher the frequency of an incremental model update, the smaller the batch of data used to update the random-effect models. With this, we find that the optimal δ -values were inversely proportional to the batch size. Recall, as $\delta \rightarrow 0$ the updates focus more on the incremental data. Therefore, the smaller the batch size, the more we want to rely on the prior information learned and less on the incoming batch of data. On the contrary, if we have a large batch of data, then we have enough information to influence the Hessian in a meaningful way. We observed this trend on both datasets. Figure 5

shows the trend of best performing δ for different update intervals, Δ . All AUC values were scaled with respect to $(\mu_{auc}, \sigma_{auc})$, given a Δ value. It can be seen that when models are updated in shorter increments (smaller batches), then higher δ values shows better performance. However, as intervals grow longer (bigger batches), incorporating the new data becomes more effective.

Figure 5: Heatmaps showing model performance against different δ -values and update intervals (Δ). Gradients are scaled against AUC values within a Δ value.

6.6 Effect of Hessian Approximation

On both datasets, we observe that model performance was not sensitive to using the full Hessian or its diagonal approximation to compute the prior. This could occur if the Hessian is well-conditioned. In production, we use the diagonal approximation for efficiency.

7 CONCLUSION

Using the proposed nearline architecture, we close the feedback loop in the order of minutes. This reduces the number of high variance opportunities for new ads and allows Thompson sampling to converge faster. As a result new advertisers saw improvements in member engagement (1.76%), ROI (reduced cost-per-click by 0.55%) and the platform saw an improvement in revenue (2.59%). Existing advertisers also saw modest improvements across these key metrics resulting in an overall site-wide improvement in member engagement and platform revenue without hurting advertiser ROI.

It is a well-established industry practice to use the most recent user activity data to train machine learning models. Intuitively, in not doing so, one runs the risk of omitting time-sensitive information (e.g. trends, topics) that can better inform a recommendation system. That said, there are systemic delays (e.g. extract, transform, load data) and the training cost of machine learning algorithm (e.g. wide & deep models) that limit the recency of data that can be used to update the model. To this end, our proposed *Lambda Learning* framework trains the model on an incremental loss function using data prepared in near realtime. While the overall model is being periodically refreshed (batch retrain), *Lambda Learner* continuously learns on an incremental objective and updates the model. This setup, along with the nearline system, improves the performance of the model serving production traffic. As shown in our experiments, the forgetting factor δ can be tuned to perform well on incremental batches of varying sizes. This also means that if we were to use custom δ_i for each random-effect (tuned based on the volume of data we've seen in the past), we can improve the model further.

REFERENCES

- [1] Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. 2010. Fast online learning through offline initialization for time-sensitive recommendation. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 703–712. <https://doi.org/10.1145/1835804.1835894>

- [2] Deepak Agarwal, Bo Long, Jonathan Traupman, Doris Xin, and Liang Zhang. 2014. Laser: A scalable response prediction platform for online advertising. In *Proceedings of the 7th ACM international conference on Web search and data mining*, 173–182.
- [3] Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*. Springer, 177–186.
- [4] Olivier Chapelle and Lihong Li. 2011. An Empirical Evaluation of Thompson Sampling. *Advances in neural information processing systems* 24 (2011), 2249–2257.
- [5] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H Chi. 2019. Top-k off-policy correction for a REINFORCE recommender system. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 456–464.
- [6] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ipsir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, 7–10.
- [7] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, 191–198.
- [8] Florence d'Alché Buc and Liva Ralaivola. 2002. Incremental Learning Algorithms for Classification and Regression: local strategies. 627 (08 2002), 320–329. <https://doi.org/10.1063/1.1503700>
- [9] Amit Dhurandhar and Marek Petrik. 2014. Efficient and accurate methods for updating generalized linear models with multiple feature additions. *The Journal of Machine Learning Research* 15, 1 (2014), 2607–2627.
- [10] Gregory Ditzler, Manuel Roveri, Cesare Alippi, and Robi Polikar. 2015. Learning in Nonstationary Environments: A Survey. *Computational Intelligence Magazine, IEEE* 10 (11 2015), 12–25. <https://doi.org/10.1109/MCI.2015.2471196>
- [11] Alexander Gepperth and Barbara Hammer. 2016. Incremental learning algorithms and applications. In *ESANN*.
- [12] James A Hanley and Barbara J McNeil. 1982. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* 143, 1 (1982), 29–36.
- [13] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TIIS)* 5, 4 (2015), 1–19.
- [14] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (Aug. 2009), 30–37. <https://doi.org/10.1109/MC.2009.263>
- [15] Jay Kreps, Neha Narkhede, Jun Rao, et al. 2011. Kafka: A Distributed Messaging System for Log Processing. *Proceedings of the NetDB* 11, 1–7.
- [16] John Langford and Tong Zhang. 2007. Epoch-Greedy algorithm for multi-armed bandits with side information. *Advances in Neural Information Processing Systems (NIPS 2007)* 20 (2007), 1.
- [17] Viktor Losing, Barbara Hammer, and Heiko Wersing. 2018. Incremental on-line learning: A review and comparison of state of the art algorithms. *Neurocomputing* 275 (2018), 1261–1274.
- [18] H. McMahan. 2011. Follow-the-Regularized-Leader and Mirror Descent: Equivalence Theorems and L1 Regularization. *Journal of Machine Learning Research - Proceedings Track* 15 (01 2011), 525–533.
- [19] Shadi A Noghabi, Kartik Paramasivam, Yi Pan, Navina Ramesh, Jon Bringham, Indranil Gupta, and Roy H Campbell. 2017. Samza: Stateful Scalable Stream Processing at LinkedIn. *Proceedings of the VLDB Endowment* 10, 12 (2017), 1634–1645.
- [20] Steffen Rendle. 2010. Factorization machines. In *2010 IEEE International Conference on Data Mining*. IEEE, 995–1000.
- [21] Linqi Song, Cem Tekin, and Mihaela Schaar. 2015. Online Learning in Large-Scale Contextual Recommender Systems. *IEEE Transactions on Services Computing* 9 (01 2015), 1–1. <https://doi.org/10.1109/TSC.2014.2365795>
- [22] Liang Tang, Romer Rosales, Ajit Singh, and Deepak Agarwal. 2013. Automatic ad format selection via contextual bandits. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, 1587–1594.
- [23] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. 2017. Deep & cross network for ad click predictions. In *Proceedings of the ADKDD'17*, 1–7.
- [24] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. 2019. Large Scale Incremental Learning.
- [25] Yang Yang, Da-Wei Zhou, De-Chuan Zhan, Hui Xiong, and Yuan Jiang. 2019. Adaptive Deep Models for Incremental Learning: Considering Capacity Scalability and Sustainability. 74–82. <https://doi.org/10.1145/3292500.3330865>
- [26] XianXing Zhang, Yitong Zhou, Yiming Ma, Bee-Chung Chen, Liang Zhang, and Deepak Agarwal. 2016. Glmix: Generalized linear mixed models for large-scale response prediction. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 363–372.

A ALGORITHM

Algorithm 1 Lambda Learner - Sequential Online GAME Training

- 1: **initialize:**
- 2: Set prior distribution for $\beta^{(0)} \sim \mathcal{N}(0, \lambda I)$,
- 3: Train global effects model on the initial batch $\Omega^{(0)}$:
 $\beta^{(0)} = [\beta_f^{*(0)}, \beta_r^{*(0)}], \beta_r^{*(0)} = \beta_{r,inc}^{*(0)}$,
- 4: Set fixed-effects coefficient to $\beta_f^{*(0)}$,
- 5: Set posterior distribution for random-effects coefficients to
 $\beta_r^{(0)} \sim \mathcal{N}(\beta_{r,inc}^{*(0)}, \Sigma_{r,inc}^{*(0)})$,
- 6: Set $H_r^{(0)} = \nabla^2 F_0(\beta_r^{*(0)})$. Note, $\Sigma_{r,inc}^{*(0)} = (H_r^{(0)} + \lambda I)^{-1}$,
- 7: Set time discount $\delta < 1$.
- 8: **for** time windows $t = 1, 2, \dots$ **do**
- 9: Prepare new batch of data: $\Omega^{(t)}$,
- 10: Compute fixed-effect offset: $\zeta_f = \mathbf{x}_f^T \beta_f^{*(0)}, \forall \mathbf{x} \in \Omega^{(t)}$,
- 11: Overall regularized log-logistic regression problem:

$$\underset{\beta_r}{\text{minimize}} \quad \underbrace{\delta F_{t-1}(\beta_r)}_{\text{time discounted log-loss on past observations}} + \underbrace{f_t(\beta_r)}_{\text{log-loss on most recent observations}} + \underbrace{\frac{\lambda}{2} \|\beta_r\|_2^2}_{\text{regularization}},$$
- 12: Approximate past log-logistic loss using quadratic approximation around $\mu_r^{(t-1)}$:

$$\delta F_{t-1}(\beta_r) = \sum_{k=0}^{t-1} \delta^{t-k} f_k(\beta_r) \approx \frac{\delta}{2} \|\beta_r - \beta_{r,inc}^{*(t-1)}\|_{H_r^{(t-1)}}^2,$$
- 13: Solve incremental logistic reg. to compute the posterior mean:

$$\beta_{r,inc}^{*(t)} = \underset{\beta_r}{\text{argmin}} \quad \frac{\delta}{2} \|\beta_r - \beta_{r,inc}^{*(t-1)}\|_{H_r^{(t-1)}}^2 + f_t(\beta_r) + \frac{\lambda}{2} \|\beta_r\|_2^2,$$
- 14: Update for Hessian of approximated Log-logistic loss:

$$H_r^{(t)} = \delta H_r^{(t-1)} + \sum_{\mathbf{x} \in \beta^{(t)}} \frac{\exp(\zeta_f + \mathbf{x}_f^T \beta_{r,inc}^{*(t)})}{(1 + \exp(\zeta_f + \mathbf{x}_f^T \beta_{r,inc}^{*(t)}))^2} \mathbf{x}_r \mathbf{x}_r^T,$$
- 15: Compute the posterior covariance:

$$\Sigma_{r,inc}^{*(t)} = (H_r^{(t)} + \lambda I)^{-1},$$
- 16: Update the model coefficients:

$$\beta^{(t)} = [\beta_f^{*(0)}, \beta_r^{*(t)}].$$
- 17: Serve online recommendations using Thompson sampling from updated distribution $\beta_r^{(t)} \sim \mathcal{N}(\beta_{r,inc}^{*(t)}, \Sigma_{r,inc}^{*(t)})$.
- 18: **end for**

B OPTIMALITY GAPS FOR INCREMENTAL UPDATES

In this section, we provide proofs of the two theorems. Beginning with a Lemma

LEMMA B.1. The Hessian for $F_t^{inc}(\beta)$ at $\beta = \beta_{inc}^{*(t)}$ is given by

$$\nabla^2 F_t^{inc}(\beta_{inc}^{*(t)}) = \sum_{k=0}^{t-1} \delta^{t-k} \nabla^2 f_k(\beta_{inc}^{*(k)}). \quad (13)$$

PROOF. Follows by mathematical induction on Hessian of (11). $\square$

B.1 Proof of Theorem 4.1

PROOF. We start with the recursive definition of F_1 ,

$$F_1(\beta^{*(0)}) - F_1(\beta_{inc}^{*(1)}) = \delta \left(F_0(\beta^{*(0)}) - F_0(\beta_{inc}^{*(1)}) \right) + f_1(\beta^{*(0)}) - f_1(\beta_{inc}^{*(1)}). \quad (14)$$

Using Taylor approximation around $F_0(\beta^{*(0)})$, since $\nabla F_0(\beta^{*(0)}) = 0$ due to optimality of $\beta^{*(0)}$, the first term reduces to

$$\delta \left(F_0(\beta^{*(0)}) - F_0(\beta_{inc}^{*(1)}) \right) = -\frac{\delta}{2} \|\beta_{inc}^{*(1)} - \beta^{*(0)}\|_{\nabla^2 F_0(\beta^{*(0)})}^2 + O(\gamma^3). \quad (15)$$

Similarly, using Taylor's approximation around $f_1(\beta_{inc}^{*(1)})$, the second term reduces to

$$f_1(\beta^{*(0)}) - f_1(\beta_{inc}^{*(1)}) = \nabla f_1(\beta_{inc}^{*(1)})^T (\beta^{*(0)} - \beta_{inc}^{*(1)}) + \frac{1}{2} \|\beta_{inc}^{*(1)} - \beta^{*(0)}\|_{\nabla^2 f_1(\beta_{inc}^{*(1)})}^2 + O(\gamma^3), \quad (16)$$

Next, using the optimality of $\beta_{inc}^{*(1)}$ gives

$$\delta \nabla^2 F_0(\beta^{*(0)}) (\beta_{inc}^{*(1)} - \beta^{*(0)}) + \nabla f_1(\beta_{inc}^{*(1)}) = 0. \quad (17)$$

Substituting equations (17) in (16), and combining with (15) into (14), we get

$$F_1(\beta^{*(0)}) - F_1(\beta_{inc}^{*(1)}) = \frac{1}{2} \|\beta_{inc}^{*(1)} - \beta^{*(0)}\|_{\nabla^2 F_1^{inc}(\beta_{inc}^{*(1)})}^2 + O(\gamma^3), \quad (18)$$

where $F_1^{inc}(\beta)$ is the objective function in incremental problem (9), and $\nabla^2 F_1^{inc}(\beta_{inc}^{*(1)}) = \delta \nabla^2 F_0(\beta^{*(0)}) + \nabla^2 f_1(\beta_{inc}^{*(1)})$ is the Hessian of that objective function. Therefore, for a convex loss function f , we have

$$F_1(\beta^{*(0)}) - F_1(\beta_{inc}^{*(1)}) \geq \frac{\delta m_0 + m_1}{2} \|\beta_{inc}^{*(1)} - \beta^{*(0)}\|^2 + O(\gamma^3), \quad (19)$$

where m_0 and m_1 correspond to strong convexity constants of F_0 and f_1 respectively. And, in the Algorithm 1, $\delta m_0 + m_1 \geq \lambda > 0$, thereby ensuring strict improvement. $\square$

LEMMA B.2. The Hessian for $F_t^{inc}(\beta)$ at $\beta = \beta_{inc}^{*(t)}$ is given by

$$\nabla^2 F_t^{inc}(\beta_{inc}^{*(t)}) = \sum_{k=0}^{t-1} \delta^{t-k} \nabla^2 f_k(\beta_{inc}^{*(k)}). \quad (20)$$

PROOF. Follows by mathematical induction on Hessian of (11). $\square$

B.2 Proof of Theorem 4.2

PROOF. We start with base definition of F_t ,

$$F_t(\beta^{*(0)}) - F_t(\beta_{inc}^{*(t)}) = \sum_{k=0}^{t-1} \delta^{t-k} \left(f_k(\beta^{*(0)}) - f_k(\beta_{inc}^{*(k)}) \right). \quad (21)$$

Using Taylor's approximation around $f_k(\beta_{inc}^{*(k)})$, we get

$$\begin{aligned} f_k(\beta^{*(0)}) - f_k(\beta_{inc}^{*(t)}) &= \nabla f_k(\beta_{inc}^{*(k)})^T (\beta^{*(0)} - \beta_{inc}^{*(t)}) \\ &\quad + \frac{1}{2} \|\beta^{*(0)} - \beta_{inc}^{*(k)}\|_{\nabla^2 f_k(\beta_{inc}^{*(k)})}^2 \\ &\quad - \frac{1}{2} \|\beta_{inc}^{*(t)} - \beta_{inc}^{*(k)}\|_{\nabla^2 f_k(\beta_{inc}^{*(k)})}^2 + \mathcal{O}(\bar{\gamma}^3), \end{aligned} \quad (22)$$

where $\bar{\gamma} = \max \{ \|\beta_{inc}^{*(k)} - \beta_{inc}^{*(t)}\|, \|\beta_{inc}^{*(k)} - \beta^{*(0)}\| \}$ for $0 \leq k \leq t$.

Using the optimality of $\beta_{inc}^{*(k)}$ gives

$$\delta \nabla^2 F_{k-1}^{inc}(\beta_{inc}^{*(k-1)}) (\beta_{inc}^{*(k)} - \beta_{inc}^{*(k-1)}) + \nabla f_k(\beta_{inc}^{*(k)}) = 0. \quad (23)$$

Substituting (23) in the first term of (22), and using $\nabla f_0(\beta_{inc}^{*(0)}) = \nabla F_0(\beta^{*(0)}) = 0$ and Lemma B.2, we get

$$\begin{aligned} &\sum_{k=0}^{k=t} \delta^{t-k} \nabla f_k(\beta_{inc}^{*(k)})^T (\beta^{*(0)} - \beta_{inc}^{*(t)}) \\ &= - \sum_{k=1}^{k=t} \delta^{t-k} (\beta^{*(0)} - \beta_{inc}^{*(t)})^T \delta \nabla^2 F_{k-1}^{inc}(\beta_{inc}^{*(k-1)}) (\beta_{inc}^{*(k)} - \beta_{inc}^{*(k-1)}) \\ &= - \sum_{k=1}^{k=t} \delta^{t-k} (\beta^{*(0)} - \beta_{inc}^{*(t)})^T \left(\sum_{j=0}^{j=k-1} \delta^{k-j} \nabla^2 f_j(\beta_{inc}^{*(j)}) \right) (\beta_{inc}^{*(k)} - \beta_{inc}^{*(k-1)}) \end{aligned} \quad (24)$$

Now, switching the summations over k and j , and using telescopic sum, we get

$$\begin{aligned} &\sum_{k=0}^{k=t} \delta^{t-k} \nabla f_k(\beta_{inc}^{*(k)})^T (\beta^{*(0)} - \beta_{inc}^{*(t)}) \\ &= - \sum_{j=0}^{j=t-1} \sum_{k=j+1}^{k=t} \delta^{t-j} (\beta^{*(0)} - \beta_{inc}^{*(t)})^T \nabla^2 f_j(\beta_{inc}^{*(j)}) (\beta_{inc}^{*(k)} - \beta_{inc}^{*(k-1)}) \\ &= - \sum_{j=0}^{j=t-1} \delta^{t-j} (\beta^{*(0)} - \beta_{inc}^{*(t)})^T \nabla^2 f_j(\beta_{inc}^{*(j)}) (\beta_{inc}^{*(t)} - \beta_{inc}^{*(j)}). \end{aligned} \quad (25)$$

Lastly, combining (25) and other two terms in (22) into (21), and grouping terms by $\nabla^2 f_k(\beta_{inc}^{*(k)})$, we get

$$\begin{aligned} F_t(\beta^{*(0)}) - F_t(\beta_{inc}^{*(t)}) &= \sum_{k=0}^{k=t} \frac{\delta^{t-k}}{2} \left[\|\beta^{*(0)} - \beta_{inc}^{*(k)}\|_{\nabla^2 f_k(\beta_{inc}^{*(k)})}^2 - \|\beta_{inc}^{*(t)} - \beta_{inc}^{*(k)}\|_{\nabla^2 f_k(\beta_{inc}^{*(k)})}^2 \right. \\ &\quad \left. - 2 (\beta^{*(0)} - \beta_{inc}^{*(t)})^T \nabla^2 f_k(\beta_{inc}^{*(k)}) (\beta_{inc}^{*(t)} - \beta_{inc}^{*(k)}) \right] + \mathcal{O}(\bar{\gamma}^3), \\ &= \sum_{k=0}^{k=t} \frac{\delta^{t-k}}{2} \|\beta^{*(0)} - \beta_{inc}^{*(k)}\|_{\nabla^2 f_k(\beta_{inc}^{*(k)})}^2 + \mathcal{O}(\bar{\gamma}^3), \\ &\geq \sum_{k=0}^{k=t} \frac{\delta^{t-k} m_k}{2} \|\beta^{*(0)} - \beta_{inc}^{*(k)}\|^2 + \mathcal{O}(\bar{\gamma}^3). \end{aligned} \quad (26)$$

where m_k 's are the respective strong convexity constants for f_k 's. Again, $\sum_{k=0}^{k=t} \delta^{t-k} m_k \geq \lambda > 0$ ensures strict improvement. $\square$

C RESULTS ON MOVIE LENS

We observe similar results on the publicly available Movie Lens dataset. Figure 6 shows the performance of the incremental update (LL) as compared to the theoretical upper bound (IBU) and the lower bound (NU).

Figure 6: Change in ROC-AUC, as we reduce the model update frequency. The point t , on the x-axis denotes the model being updated every t hours - in the incremental case using ONLY data in the last t hours and in the batch case using the entire history including the data in the last t hours. $\delta = 0.95$.