
Large-scale Stochastic Optimization of NDCG Surrogates for Deep Learning with Provable Convergence

Zi-Hao Qiu^{*1} Quanqi Hu^{*2} Yongjian Zhong² Lijun Zhang¹ Tianbao Yang²

Abstract

NDCG, namely Normalized Discounted Cumulative Gain, is a widely used ranking metric in information retrieval and machine learning. However, efficient and provable stochastic methods for maximizing NDCG are still lacking, especially for deep models. In this paper, we propose a principled approach to optimize NDCG and its top- K variant. First, we formulate a novel compositional optimization problem for optimizing the NDCG surrogate, and a novel bilevel compositional optimization problem for optimizing the top- K NDCG surrogate. Then, we develop efficient stochastic algorithms with provable convergence guarantees for the non-convex objectives. Different from existing NDCG optimization methods, the per-iteration complexity of our algorithms scales with the mini-batch size instead of the number of total items. To improve the effectiveness for deep learning, we further propose practical strategies by using initial warm-up and stop gradient operator. Experimental results on multiple datasets demonstrate that our methods outperform prior ranking approaches in terms of NDCG. To the best of our knowledge, this is the first time that stochastic algorithms are proposed to optimize NDCG with a provable convergence guarantee. Our proposed methods are implemented in the LibAUC library at <https://libauc.org/>.

1. Introduction

NDCG is a performance metric of primary interest for learning to rank in information retrieval (Liu, 2011), and is also adopted in many machine learning tasks where ranking is of foremost importance (Liu & Yang, 2008; Bhatia et al., 2015). In the following, we use the terminologies from information retrieval to describe NDCG and our methods. The goal is to rank the relevant items higher than irrelevant items for any given query. For a query q and a list of n items, the ranking model assigns a score for each item, and then we obtain an ordered list by sorting these scores in descending order. The NDCG score for q can be computed by:

$$\text{NDCG}_q = \frac{1}{Z_q} \sum_{i=1}^n \frac{2^{y_i} - 1}{\log_2(1 + r(i))}, \quad (1)$$

where y_i denotes the relevance score of the i -th item, $r(i)$ denotes the rank of the i -th item in the ordered list, and Z_q is a normalization factor that is the Discounted Cumulative Gain (DCG) score (Järvelin & Kekäläinen, 2002) of the optimal ranking for q . The top- K variant of NDCG can be defined similarly by summing over items whose ranks are in the top K positions of the ordered list. In many real-world applications, e.g., recommender systems, we want to recommend a small set of K items from a large collection of items (Cremonesi et al., 2010), thus top- K NDCG is a popular metric in these applications.

There are several challenges for optimizing NDCG and its top- K variant. First, computing the rank of each item among all n items is expensive. Second, the rank operator is non-differentiable in terms of model parameters. To tackle non-differentiability, surrogate functions have been proposed in the literature for approximating NDCG and its top- K variant (Taylor et al., 2008; Qin et al., 2010; Swezey et al., 2021; Pobrotyn & Bialobrzeski, 2021). However, to the best of our knowledge, *the computational challenge* of computing the gradient of (1) that involves sorting n items has never been addressed. All existing gradient-based methods have a complexity of $O(nd)$ per-iteration, where d is the number of model parameters, which is prohibitive for deep learning tasks with big n and big d . A naive approach is to update the model parameters by the gradient of the NDCG surrogate over a mini-batch of samples, however, since the

^{*}Equal contribution ¹National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China ²the University of Iowa, Iowa City, USA. Correspondence to: Lijun Zhang <zhanglj@lamda.nju.edu.cn>, Tianbao Yang <tianbaoyang@uiowa.edu>.

surrogate for NDCG is complicated and non-convex, an unbiased stochastic gradient is not readily computed, which makes existing methods lack theoretical guarantee.

In this paper, we propose the first stochastic algorithms with a per-iteration complexity of $O(Bd)$, where B is the mini-batch size, for optimizing the surrogates for NDCG and its top- K variant, and establishing their convergence guarantees. For optimizing the NDCG surrogate, we first formulate a novel *finite-sum coupled compositional optimization (FCCO)* problem. Then, we develop an efficient stochastic algorithm inspired by a recent work on average precision maximization (Qi et al., 2021). We establish an iteration complexity of $O(\frac{1}{\epsilon^4})$ for finding an ϵ -level stationary solution, which is better than that proved by Qi et al. (2021), i.e., $O(\frac{1}{\epsilon^8})$. To tackle the challenge of optimizing the top- K NDCG surrogate that involves a selection operator, we propose a novel *bilevel optimization* problem, which contains many lower level problems for top- K selection of all queries. Then we smooth the non-smooth functions in the selection operator, and propose an efficient algorithm with the iteration complexity of $O(\frac{1}{\epsilon^4})$. The algorithm is based on recent advances of stochastic bilevel optimization (Guo et al., 2021a), but with unique features to tackle the compositional upper level problem and a mini-batch of randomly sampled lower level problems per iteration for optimizing the top- K NDCG surrogate.

To improve the effectiveness of optimizing the NDCG surrogates, we also study two practical strategies. First, we propose initial warm-up to find a good initial solution. Second, we use stop gradient operator to simplify the optimization of the top- K NDCG surrogate. We conduct comprehensive experiments on two tasks, learning to rank and recommender systems. Empirical results demonstrate that the proposed algorithms can consistently outperform prior approaches in terms of NDCG, and show the effectiveness of two proposed strategies.

We summarize our contributions below:

- We formulate the optimization of the NDCG surrogate as a finite-sum coupled compositional optimization problem, and propose a novel stochastic algorithm with provable convergence guarantees.
- We propose a novel bilevel compositional optimization formulation for optimizing the top- K NDCG surrogate. Then we develop a novel stochastic algorithm and establish its convergence rate.
- To improve the effectiveness for deep learning, we also study practical strategies by using initial warm-up and stop gradient operator. Experimental results on multiple datasets demonstrate the effectiveness of our algorithms and strategies.

2. Related Work

Listwise LTR approaches. Learning to rank (LTR) is an extensively studied area (Liu, 2011), and we only review the listwise LTR approaches that are closely related to this work. The listwise methods can be classified into three groups. The first group uses ranking metrics to dynamically re-weight instances during training. For example, LambdaRank algorithms (Burgess et al., 2005a; Burgess, 2010) define a weight ΔNDCG , which is the NDCG difference when a pair of items is swapped in the current list, and use it to re-weight the pair during training. Although algorithms in this group take NDCG into account, the underlying losses of them remain unknown and their theoretical relations to NDCG are difficult to analyze. The second group defines loss functions over the entire item lists to optimize the agreement between predictions and ground truth rankings. For example, ListNet (Cao et al., 2007) minimizes cross-entropy between predicted and ground truth top-one probability distributions. ListMLE (Xia et al., 2008) aims to maximize the likelihood of the ground truth list given the predicted results. However, optimizing these loss functions might not necessarily maximize NDCG. In addition, efficient stochastic algorithms for optimizing these losses are still lacking. The third group directly optimizes ranking metrics, and most of works focus on the widely used NDCG, as reviewed below.

NDCG Optimization. Some earlier works employ traditional optimization techniques, e.g., genetic algorithm (Yeh et al., 2007), boosting (Xu & Li, 2007; Valizadegan et al., 2009), and SVM framework (Chakrabarti et al., 2008). However, these methods are not scalable to big data. A popular class of approaches is to approximate ranks in NDCG with smooth functions and then optimize the resulting surrogates. For example, SoftRank (Taylor et al., 2008) tries to use rank distributions to smooth NDCG, however, it suffers from a high computational complexity of $O(n^3)$. ApproxNDCG (Qin et al., 2010) approximates the indicator function in the computation of ranks, and the top- K selector in the computation of top- K variant by a generalized sigmoid function. Recently, PiRank (Swezey et al., 2021) and NeuralNDCG (Pobrotyn & Bialobrzewski, 2021) are proposed to smooth NDCG by approximating non-continuous sorting operator based on NeuralSort (Grover et al., 2019). However, these methods mainly focus on how to approximate NDCG with differentiable functions, and remain computationally expensive as their per-iteration complexity is $O(nd)$. Moreover, little attention has been paid to the convergence guarantee for the stochastic optimization of these surrogates. In contrast, this is the first work to develop stochastic algorithms with provable convergence guarantee for optimizing the surrogates for NDCG and its top- K variant.

Stochastic Compositional Optimization. Optimization of a two-level compositional function in the form of

$\mathbb{E}_\xi[f(\mathbb{E}_\zeta[g(\mathbf{w}; \zeta)]; \xi)]$ where ξ and ζ are independent random variables, or its finite-sum variant has been studied extensively (Wang et al., 2017; Balasubramanian et al., 2020; Chen et al., 2021b). In this paper, we formulate the surrogate function of NDCG into a similar but more complicated two-level compositional function of the form $\mathbb{E}_\xi[f(\mathbb{E}_\zeta[g(\mathbf{w}; \zeta, \xi)])]$ where ξ and ζ are independent and ξ has a finite support inspired by (Qi et al., 2021). The key difference between our compositional function and the ones considered in previous work is that the inner function $g(\mathbf{w}; \zeta, \xi)$ also depends on the random variable ξ of the outer level. Our algorithm is developed based on that of Qi et al. (2021) for average precision maximization, but establishes an improved complexity of $O(\frac{1}{\epsilon^4})$ for finding an ϵ -stationary solution. It is also notable that our algorithm and convergence result for optimizing NDCG is similar to that in a concurrent work (Wang & Yang, 2022) dedicated to FCCO. However, the key difference is that our convergence analysis for optimizing NDCD follows that for optimizing top- K NDCG in a novel bi-level optimization framework.

Stochastic Bilevel Optimization. Stochastic bilevel optimization (SBO) has a long history in the literature (Colson et al., 2007; Kunisch & Pock, 2013; Liu et al., 2020). Recent works on SBO focus on algorithms with provable convergence rates (Ghadimi & Wang, 2018; Ji et al., 2020; Hong et al., 2020; Chen et al., 2021a). However, most of these studies do not explicitly consider the challenge for dealing with SBO with many lower level problems. Guo et al. (2021a) consider SBO with many lower level problems and develop a stochastic algorithm with convergence guarantee. However, their algorithm is not applicable to our problem for optimizing the compositional top- K NDCG surrogate and a mini-batch of randomly sampled lower level problems in each iteration, and is not practical as it requires evaluating the stochastic gradients twice per-iteration at two different points. In this paper, we propose a novel stochastic algorithm for optimizing the top- K NDCG surrogate, which contains many lower level problems, and establish its iteration complexity of $O(\frac{1}{\epsilon^4})$.

3. Preliminaries

In this section, we provide some preliminaries and notations. Let \mathcal{Q} denote the query set of size N , and $q \in \mathcal{Q}$ denote a query. \mathcal{S}_q denotes a set of N_q items (e.g., documents, movies) to be ranked for q . For each $\mathbf{x}_i^q \in \mathcal{S}_q$, let $y_i^q \in \mathbb{R}^+$ denote its relevance score, which measures the relevance between query q and item x_i^q . Let $\mathcal{S}_q^+ \subseteq \mathcal{S}_q$ denote a set of N_q^+ items *relevant* to q , whose relevance scores are *non-zero*. Denoted by $\mathcal{S} = \{(q, \mathbf{x}_i^q), q \in \mathcal{Q}, \mathbf{x}_i^q \in \mathcal{S}_q^+\}$ all relevant query-item (Q-I) pairs. Let $h_q(\mathbf{x}; \mathbf{w})$ denote the predictive function for \mathbf{x} with respect to the query q , whose parameters are denoted by $\mathbf{w} \in \mathbb{R}^d$ (e.g., a deep neural network). Let

$\mathbb{I}(\cdot)$ denote the indicator function, which outputs 1 if its input is true and 0 otherwise. Let

$$r(\mathbf{w}; \mathbf{x}, \mathcal{S}_q) = \sum_{\mathbf{x}' \in \mathcal{S}_q} \mathbb{I}(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) \geq 0)$$

denote the rank of \mathbf{x} with respect to the set \mathcal{S}_q , where we simply ignore the tie.

According to the definition in (1), the averaged NDCG over all queries can be expressed by

$$\text{NDCG} := \frac{1}{N} \sum_{q=1}^N \frac{1}{Z_q} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \frac{2^{y_i^q} - 1}{\log_2(r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)},$$

where Z_q is the maximum DCG of a perfect ranking of items in \mathcal{S}_q , which can be pre-computed. Note that \mathbf{x}_i^q are summed over \mathcal{S}_q^+ instead of \mathcal{S}_q , because only relevant items have non-zero relevance scores and contribute to NDCG.

An important variant of NDCG is its top- K variant, which is defined over the items $\mathbf{x}_i^q \in \mathcal{S}_q$ whose prediction scores are in the top- K positions, i.e.,

$$\frac{1}{N} \sum_{q=1}^N \frac{1}{Z_q^K} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \mathbb{I}(\mathbf{x}_i^q \in \mathcal{S}_q[K]) \frac{2^{y_i^q} - 1}{\log_2(r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)},$$

where $\mathcal{S}_q[K]$ denotes the top- K items whose prediction scores are in the top- K positions among all items in \mathcal{S}_q , and Z_q^K denotes the top- K DCG score of the perfect ranking.

4. Optimizing a Smooth NDCG Surrogate

To address the non-differentiability of the rank function $r(\mathbf{w}; \mathbf{x}, \mathcal{S}_q)$, we approximate it by a continuous and differentiable surrogate function

$$\bar{g}(\mathbf{w}; \mathbf{x}, \mathcal{S}_q) = \sum_{\mathbf{x}' \in \mathcal{S}_q} \ell(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w})),$$

where $\ell(\cdot)$ is a surrogate loss function of $\mathbb{I}(\cdot \geq 0)$. In this paper, we use a convex and non-decreasing smooth surrogate loss, e.g., squared hinge loss $\ell(x) = \max(0, x + c)^2$, where c is a margin parameter. Other choices are possible with pros and cons discussed in the literature (Wu et al., 2009; Qin et al., 2010). Below, we abuse the notation $\ell(\mathbf{w}; \mathbf{x}', \mathbf{x}, q) = \ell(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}))$.

Using the surrogate function, we cast NDCG maximization into:

$$\max_{\mathbf{w} \in \mathbb{R}^d} L(\mathbf{w}) := \frac{1}{|\mathcal{S}|} \sum_{q=1}^N \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \frac{2^{y_i^q} - 1}{Z_q \log_2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}. \quad (2)$$

The following lemma justifies the maximization over $L(\mathbf{w})$ for NDCG maximization:

Lemma 4.1. *When $\ell(\mathbf{w}; \mathbf{x}', \mathbf{x}, q) \geq \mathbb{I}(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) \geq 0)$, then $L(\mathbf{w})$ is a lower bound of NDCG.*

The key challenge in designing an efficient algorithm for solving the above problem lies at (i) computing $\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$

Algorithm 1 Stochastic Optimization of NDCG: SONG

Require: $\eta, \gamma_0, \beta_1, u^{(1)} = 0$
Ensure: \mathbf{w}_T

- 1: **for** $t = 1, \dots, T$ **do**
- 2: Draw some relevant Q-I pairs $\mathcal{B} = \{(q, \mathbf{x}_i^q)\} \subset \mathcal{S}$
- 3: For each sampled q draw a batch of items $\mathcal{B}_q \subset \mathcal{S}_q$
- 4: **for** each sampled Q-I pair $(q, \mathbf{x}_i^q) \in \mathcal{B}$ **do**
- 5: Let $\hat{g}_{q,i}(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_q|} \sum_{\mathbf{x}' \in \mathcal{B}_q} \ell(\mathbf{w}_t; \mathbf{x}', \mathbf{x}_i^q, q)$
- 6: Compute $u_{q,i}^{(t+1)} = (1 - \gamma_0)u_{q,i}^{(t)} + \gamma_0 \hat{g}_{q,i}(\mathbf{w}_t)$
- 7: Compute $p_{q,i} = \nabla f_{q,i}(u_{q,i}^{(t)})$
- 8: **end for**
- 9: Compute the stochastic gradient estimator $G(\mathbf{w}_t)$ by

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{B}} p_{q,i} \nabla \hat{g}_{q,i}(\mathbf{w}_t)$$

- 10: Compute $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1)G(\mathbf{w}_t)$
- 11: update $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{m}_{t+1}$
- 12: **end for**

and its gradient is expensive when $N_q = |\mathcal{S}_q|$ is very large; and (ii) an unbiased stochastic gradient of the objective function is not readily available. To highlight the second challenge, let us consider the gradient of the function

$$\phi(\mathbf{w}) = \frac{1}{\log_2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}, \text{ which is given by}$$

$$\nabla \phi(\mathbf{w}) = \frac{-\log_2(e) \cdot \nabla \bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)}{\log_2^2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1) \cdot (\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}.$$

We can estimate $\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ by its unbiased estimator using a mini-batch of B_q items $\mathbf{x}' \in \mathcal{B}_q \subset \mathcal{S}_q$, i.e., $\frac{N_q}{B_q} \sum_{\mathbf{x}' \in \mathcal{B}_q} \ell(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}_i^q; \mathbf{w}))$. However, directly plug this unbiased estimator of $\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ into the above expression will produce a biased estimator of $\nabla \phi(\mathbf{w})$ due to the non-linear function of \bar{g} . The optimization error will be large if the mini-batch size B_q is small (Hu et al., 2020).

To address this challenge, we cast the problem into the following equivalent minimization form:

$$\min_{\mathbf{w} \in \mathbb{R}^d} F(\mathbf{w}) := \frac{1}{|\mathcal{S}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{S}} f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)), \quad (3)$$

where $g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) = \frac{1}{N_q} \bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ and $f_{q,i}(g) = \frac{1}{Z_q} \frac{1 - 2^{y_i^q}}{\log_2(N_q g + 1)}$. It is a special case of a family of **finite-sum coupled compositional stochastic optimization** problems, which was first studied by Qi et al. (2021) for maximizing average precision. Inspired by their method, we develop a stochastic algorithm for solving (3). The complete procedure is provided in Algorithm 1, which is named as Stochastic Optimization of NDCG (SONG).

To motivate the proposed method, we first derive the gradient of $F(\mathbf{w})$ by the chain rule, which is given by

$$\nabla F(\mathbf{w}) = \frac{1}{|\mathcal{S}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{S}} \nabla f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)) \nabla g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q).$$

The major cost for computing $\nabla F(\mathbf{w})$ lies at computing $g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ and its gradient, which involves all items in \mathcal{S}_q . To this end, we approximate these quantities by stochastic samples. The gradient $\nabla g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ can be simply approximated by the stochastic gradient $\nabla \hat{g}_{q,i}(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_q|} \sum_{\mathbf{x}' \in \mathcal{B}_q} \nabla \ell(\mathbf{w}_t; \mathbf{x}', \mathbf{x}_i^q, q)$, where \mathcal{B}_q is sampled from \mathcal{S}_q . Note that $\nabla f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q))$ is non-linear with $g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$, thus we need a better way to estimate $g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ to control the approximation error and provide convergence guarantee. We borrow a technique from Qi et al. (2021) by using a moving average estimator to keep track of $g(\mathbf{w}_t; \mathbf{x}_i^q, \mathcal{S}_q)$ for each $\mathbf{x}_i^q \in \mathcal{S}_q^+$. To this end, we maintain a scalar $u_{q,i}$ for each relevant query-item pair (q, \mathbf{x}_i^q) and update it by a linear combination of historical one $u_{q,i}^{(t)}$ and an unbiased estimator of $g(\mathbf{w}_t; \mathbf{x}_i^q, \mathcal{S}_q)$ denoted by $\hat{g}_{q,i}(\mathbf{w}_t)$ in Step 5 and 6, where $\gamma_0 \in (0, 1)$ is a parameter. Intuitively, when t increases, \mathbf{w}_{t-1} is getting closer to \mathbf{w}_t , hence the previous value of the estimator, i.e., $u_{q,i}^{(t)}$ is useful for estimating $g_{q,i}(\mathbf{w}_t)$. With these stochastic estimators, we can compute the gradient of the objective in (3) with controllable approximation error in Step 9. We implement the momentum update for \mathbf{w}_{t+1} in Step 10 and 11, where $\beta_1 \in (0, 1)$ is the momentum parameter. The momentum update can be also replaced by the Adam-style update (Guo et al., 2021b), where the step size η is replaced by an adaptive step size. We can establish the same convergence rate for the Adam-style update.

We also have several remarks about SONG: (i) the total per-iteration complexity of SONG is $O(Bd + B^2)$. The details can be found in Appendix A. For a large model size $d \gg B$, we have the per-iteration complexity of $O(Bd)$, which is similar to the standard cost of deep learning and is independent of the length of \mathcal{S}_q for each query; and (ii) the additional memory cost is the size of $u_{q,i}$, i.e., the number of all relevant Q-I pairs. It is worth to mention that in many real-world datasets the number of relevant Q-I pairs are much fewer than all Q-I pairs (i.e., data is sparse) (Yuan et al., 2014; Yin et al., 2020; Singh, 2020). Thus the additional memory cost is acceptable in most cases.

Next, we establish the convergence guarantee of SONG in the following theorem.

Theorem 4.2. *Under appropriate conditions and proper settings of parameters $\gamma_0, \gamma_1, \eta = O(\epsilon^2)$, $\beta_1 = 1 - \gamma_1$, Algorithm 1 ensures that after $T = O(\frac{1}{\epsilon^4})$ iterations we can find an ϵ -stationary solution of $F(\mathbf{w})$, i.e., $\mathbb{E}[\|\nabla F(\mathbf{w}_\tau)\|^2] \leq \epsilon^2$ for a randomly selected $\tau \in \{1, \dots, T\}$.*

Remark: The above theorem indicates that SONG has the same $O(\frac{1}{\epsilon^4})$ iteration complexity as the standard SGD for solving standard non-convex losses (Ghadimi & Lan, 2013). We refer the interested readers to Appendix E for the proof, where we also exhibit the settings for γ_0, β_1, η and the conditions. The conditions are imposed mainly for ensuring

$f_{q,i}(g)$ and $g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ are smooth and Lipschitz continuous. It is worth mentioning that the above complexity is better than that proved by Qi et al. (2021), i.e., $O(1/\epsilon^5)$. In addition, we do not have any requirement on the batch size, i.e., $|\mathcal{B}|, |\mathcal{B}_q|$, which can be as small as 1. However, we can enjoy parallel speed-up for a large batch size.

5. Optimizing a Smooth Top- K NDCG Surrogate

In this section, we propose an efficient stochastic algorithm to optimize the top- K variant of NDCG. By using the smooth surrogate loss $\ell(\cdot)$ for approximating the rank function, we have the following objective for top- K NDCG:

$$\frac{1}{N} \sum_{q=1}^N \frac{1}{Z_q^K} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \mathbb{I}(\mathbf{x}_i^q \in \mathcal{S}_q[K]) \frac{2^{y_i^q} - 1}{\log_2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)},$$

where $\mathcal{S}_q[K]$ denotes the set of top- K items in \mathcal{S}_q whose prediction scores are in the top- K positions. Compared with optimizing the NDCG surrogate in (3), there is another level of complexity, i.e., the selection of top- K items from \mathcal{S}_q , which is non-differentiable. In the literature, Qin et al. (2010) and Wu et al. (2009) use the relationship $\mathbb{I}(\mathbf{x}_i^q \in \mathcal{S}_q[K]) = \mathbb{I}(K - r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) \geq 0)$ and approximate it by $\psi(K - \bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q))$, where ψ is a continuous surrogate of the indicator function. However, there are two levels of approximation error, one lies at approximating $r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ by $\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ and the other one lies at approximating $\mathbb{I}(\cdot \geq 0)$ by $\psi(\cdot)$. To reduce the error for selecting $\mathbf{x}_i^q \in \mathcal{S}_q[K]$, we propose a more effective method, which relies on the following lemma:

Lemma 5.1. *Let $\lambda_q(\mathbf{w}) = \arg \min_{\lambda} (K + \varepsilon)\lambda + \sum_{\mathbf{x}' \in \mathcal{S}_q} (h_q(\mathbf{x}'; \mathbf{w}) - \lambda)_+$, where $\varepsilon \in (0, 1)$, then $\lambda_q(\mathbf{w})$ is the $(K + 1)$ -th largest value among $h_q(\mathbf{x}', \mathbf{w}), \forall \mathbf{x}' \in \mathcal{S}_q$, and hence $\mathbf{x}_i^q \in \mathcal{S}_q[K]$ is equivalent to $h_q(\mathbf{x}_i^q; \mathbf{w}) > \lambda_q(\mathbf{w})$.*

Remark: We can show that the optimal solution $\lambda_q(\mathbf{w})$ can be served as the threshold for selecting top- K items in \mathcal{S}_q .

As a result, the problem can be converted into

$$\min_{|\mathcal{S}|} \frac{1}{|\mathcal{S}|} \sum_{q=1}^N \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \frac{\mathbb{I}(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}) > 0)(1 - 2^{y_i^q})}{Z_q^K \log_2(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}$$

$$s.t., \lambda_q(\mathbf{w}) = \arg \min_{\lambda} \frac{K + \varepsilon}{N_q} \lambda + \frac{1}{N_q} \sum_{\mathbf{x}' \in \mathcal{S}_q} (h_q(\mathbf{x}'; \mathbf{w}) - \lambda)_+.$$

However, there are still several challenges that prevent us developing a provable algorithm. In particular, the selection operator $\mathbb{I}(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}) > 0)$ is a non-smooth function of \mathbf{w} due to (i) the indicator function $\mathbb{I}(\cdot)$ is non-continuous and non-differentiable; and (ii) $\lambda_q(\mathbf{w})$ is a non-smooth function of \mathbf{w} because the lower optimization problem is non-smooth and non-strongly convex.

To address the above challenges, we first approximate $\mathbb{I}(\cdot > 0)$ by a smooth and Lipschitz continuous function $\psi(\cdot)$. The choice of ψ can be justified by the following lemma:

Lemma 5.2. *If $\psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w})) \leq C \mathbb{I}(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}) > 0)$ holds for some constant $C > 0$ and $\ell(\mathbf{w}; \mathbf{x}', \mathbf{x}, q) \geq \mathbb{I}(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) > 0)$, then the function $\frac{1}{N} \sum_{q=1}^N \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \frac{\psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}))(2^{y_i^q} - 1)}{C Z_q^K \log_2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}$ is a lower bound of the top- K NDCG.*

Remark: When $h_q(\mathbf{x}; \mathbf{w})$ is bounded, it is not hard to find a smooth and Lipschitz continuous function $\psi(\cdot)$ satisfying the above condition. A simple choice is $\psi(s) = \max(s, 0)^2$.

Next, we smooth $\lambda(\mathbf{w})$. The idea is to make the objective function in the lower level problem smooth and strongly convex, while not affecting the optimal solution $\lambda(\mathbf{w})$ too much. To this end, we replace the lower level problem by

$$\hat{\lambda}_q(\mathbf{w}) = \arg \min_{\lambda} L_q(\lambda; \mathbf{w}) := \frac{K + \varepsilon}{N_q} \lambda + \frac{\tau_2}{2} \lambda^2$$

$$+ \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} \tau_1 \ln(1 + \exp((h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)/\tau_1)).$$

The following lemma justifies the above smoothing.

Lemma 5.3. *Assuming $h_q(\mathbf{x}, \mathbf{w}) \in (0, c_h]$, if $\tau_1 = \tau_2 = \varepsilon$ for some $\varepsilon \ll 1$, then we have $|\hat{\lambda}_q(\mathbf{w}) - \lambda_q(\mathbf{w})| \leq O(\varepsilon)$ for any \mathbf{w} . In addition, $L_q(\lambda; \mathbf{w})$ is a smooth and strongly convex function in terms of λ for any \mathbf{w} .*

As a result, we propose to solve the following optimization problem for top- K NDCG maximization:

$$\min_{|\mathcal{S}|} \frac{1}{|\mathcal{S}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{S}} \psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \hat{\lambda}_q(\mathbf{w})) f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q))$$

$$s.t., \hat{\lambda}_q(\mathbf{w}) = \arg \min_{\lambda} L_q(\lambda; \mathbf{w}), \forall q \in \mathcal{Q}, \quad (4)$$

where we employ $f_{q,i}(g)$ to denote $\frac{1}{Z_q^K} \frac{1 - 2^{y_i^q}}{\log_2(N_q g + 1)}$.

Our bilevel formulation is more advantageous than previous NDCG@ K formulation. First, our formulation only approximates $r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ by $\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ **once** in the denominator, while previous one approximates $r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)$ **twice** (one in $\psi(K - r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q))$ and one in the denominator). In addition, $\psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}))$ is arguably better than $\psi(K - \bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q))$ for approximating $\mathbb{I}(K - r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) \geq 0)$ due to Lemma 5.1.

Although (4) is a bilevel optimization problem, existing stochastic algorithms for bilevel optimization are not applicable to solving the above problem. That is because there are several differences from the standard bilevel optimization problem studied in the literature. First, an unbiased stochastic gradient of the objective function is not readily computed as we explained before. Second, there are multiple lower level problems in (4), whose solutions cannot be updated at the same time for all $q \in \mathcal{Q}$ when N is large. To

address these challenges, we develop a tailored stochastic algorithm for solving (4).

The proposed algorithm is presented in Algorithm 2, to which we refer as K-SONG. To motivate K-SONG, we first consider the gradient of the objective function denoted by $F_K(\mathbf{w})$, which can be computed as

$$\begin{aligned} \nabla F_K(\mathbf{w}) &= \frac{1}{|\mathcal{S}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{S}} \left(\psi'(h_q(\mathbf{x}_i^q; \mathbf{w}) - \hat{\lambda}_q(\mathbf{w})) \right. \\ &\quad \left. (\nabla h_q(\mathbf{x}_i^q; \mathbf{w}) - \nabla_{\mathbf{w}} \hat{\lambda}_q(\mathbf{w})) \right) f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)) \\ &\quad + \psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \hat{\lambda}_q(\mathbf{w})) \nabla g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) f'_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)). \end{aligned} \quad (5)$$

Similar to SONG, we can estimate $g(\mathbf{w}_t; \mathbf{x}_i^q, \mathcal{S}_q)$ by $u_{q,i}^{(t)}$. An inherent challenge of bilevel optimization is to estimate the implicit gradient $\nabla_{\mathbf{w}} \hat{\lambda}(\mathbf{w})$. According to the optimality condition of $\hat{\lambda}(\mathbf{w})$ (Ghadimi & Wang, 2018), we can derive $\nabla_{\mathbf{w}} \hat{\lambda}_q(\mathbf{w}) = -\nabla_{\lambda, \mathbf{w}}^2 L_q(\hat{\lambda}_q(\mathbf{w}); \mathbf{w}) (\nabla_{\lambda}^2 L_q(\hat{\lambda}_q(\mathbf{w}); \mathbf{w}))^{-1}$. To estimate $\nabla_{\lambda, \mathbf{w}}^2 L_q(\hat{\lambda}(\mathbf{w}); \mathbf{w})$ at the t -th iteration, we use the current estimate $\lambda_{q,t}$ in place of $\hat{\lambda}_q(\mathbf{w}_t)$ and use $L_q(\hat{\lambda}, \mathbf{w}; \mathcal{B}_q)$ that is defined by a mini-batch samples of \mathcal{B}_q in place of $L_q(\hat{\lambda}; \mathbf{w})$, i.e.,

$$\begin{aligned} L_q(\lambda, \mathbf{w}; \mathcal{B}_q) &= \frac{K}{N_q} \lambda + \frac{\tau_2}{2} \lambda^2 \\ &\quad + \frac{1}{|\mathcal{B}_q|} \sum_{\mathbf{x}_i \in \mathcal{B}_q} \tau_1 \ln(1 + \exp((h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)/\tau_1)). \end{aligned} \quad (6)$$

The issue of estimating $(\nabla_{\lambda}^2 L_q(\hat{\lambda}_q(\mathbf{w}); \mathbf{w}))^{-1}$ is more tricky. In the literature (Ghadimi & Wang, 2018), a common method is to use von Neuman series with stochastic samples to estimate it. However, such method requires multiple samples in the order of $O(1/\tau_2)$, which is a large number when τ_2 is small. To address this issue, we follow a similar strategy of Guo et al. (2021a) to estimate $\nabla_{\lambda}^2 L_q(\hat{\lambda}_q(\mathbf{w}); \mathbf{w})$ directly by using mini-batch samples. In the proposed algorithm, we use a moving average estimator denoted by s_q as shown in Step 10. Finally, we have the following stochastic gradient estimator:

$$\begin{aligned} G(\mathbf{w}_t) &= \frac{1}{|\mathcal{B}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{B}} p_{q,i} \nabla \hat{g}_{q,i}(\mathbf{w}_t) \\ &\quad + \psi'(h_q(\mathbf{x}_i^q; \mathbf{w}_t) - \lambda_{q,t}) \left[\nabla_{\mathbf{w}} h_q(\mathbf{x}_i^q; \mathbf{w}_t) \right. \\ &\quad \left. - \nabla_{\lambda, \mathbf{w}}^2 L_q(\mathbf{w}_t, \lambda_{q,t}^t; \mathcal{B}_t) s_{q,t}^{-1} \right] f(u_{q,i}^{(t)}) \end{aligned} \quad (7)$$

where $p_{q,i}$ is computed in Step 7 in K-SONG.

We follow a similar strategy as Guo et al. (2021b) to update $\lambda_{q,t+1}$ by a simple stochastic gradient update, shown in Step 11. It is notable that different from Guo et al. (2021a), we update $\lambda_{q,t+1}$ for a mini-batch of randomly sampled queries q , which makes the analysis more challenging.

Algorithm 2 Stochastic Optimization of top- K NDCG: K-SONG

Require: $\eta_0, \eta_1, \gamma_0, \gamma'_0, \beta_1, u^{(1)} = 0, \lambda = 0$

Ensure: \mathbf{w}_T

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Draw some relevant Q-I pairs $\mathcal{B} = \{(q, \mathbf{x}_i^q)\} \subset \mathcal{S}$
 - 3: For each sampled q draw a batch of items $\mathcal{B}_q \subset \mathcal{S}_q$
 - 4: **for** each sampled Q-I pair $(q, \mathbf{x}_i^q) \in \mathcal{B}$ **do**
 - 5: Let $\hat{g}_{q,i}(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_q|} \sum_{\mathbf{x}' \in \mathcal{B}_q} \ell(\mathbf{w}_t; \mathbf{x}', \mathbf{x}_i^q, q)$
 - 6: Let $u_{q,i}^{(t+1)} = (1 - \gamma_0) u_{q,i}^{(t)} + \gamma_0 \hat{g}_{q,i}(\mathbf{w}_t)$
 - 7: Let $p_{q,i} = \psi(h_q(\mathbf{x}_i^q; \mathbf{w}_t) - \lambda_{q,t}) \nabla f_{q,i}(u_{q,i}^t)$
 - 8: **end for**
 - 9: **for** each sampled query $q \in \mathcal{B}$ **do**
 - 10: Let $s_{q,t+1} = (1 - \gamma'_0) s_{q,t} + \gamma'_0 \nabla_{\lambda}^2 L_q(\lambda_{q,t}; \mathbf{w}_t; \mathcal{B}_q)$
 - 11: Let $\lambda_{q,t+1} = \lambda_{q,t} - \eta_0 \nabla_{\lambda} L_q(\lambda_{q,t}; \mathbf{w}_t; \mathcal{B}_q)$
 - 12: **end for**
 - 13: Compute a stochastic gradient $G(\mathbf{w}_t)$ according to (7) or (8)
 - 14: Compute $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1) G(\mathbf{w}_t)$
 - 15: Update $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_1 \mathbf{m}_{t+1}$
 - 16: **end for**
-

Finally, we present the convergence guarantee of K-SONG.

Theorem 5.4. *Under appropriate conditions and proper settings of parameters $\gamma_0, \gamma'_0, \eta_0 = \mathcal{O}(|\mathcal{B}_q| \epsilon^2)$, $\gamma_1 = \mathcal{O}(\min\{|\mathcal{B}|, |\mathcal{B}_q|\} \epsilon^2)$, $\beta_1 = 1 - \gamma_1$, $\eta_1 = \mathcal{O}\left(\min\left\{\frac{|\mathcal{B}||\mathcal{B}_q| \epsilon^2}{|\mathcal{S}|}, \min\{|\mathcal{B}|, |\mathcal{B}_q|\} \epsilon^2\right\}\right)$, Algorithm 2 ensures that after $T = \mathcal{O}\left(\max\left\{\frac{|\mathcal{S}|}{|\mathcal{B}||\mathcal{B}_q| \epsilon^4}, \frac{1}{\min\{|\mathcal{B}|, |\mathcal{B}_q|\} \epsilon^4}\right\}\right)$ iterations we can find an ϵ -stationary solution of $F_K(\mathbf{w})$, i.e., $\mathbb{E}[\|\nabla F_K(\mathbf{w}_T)\|^2] \leq \epsilon^2$ for a randomly selected $\tau \in \{1, \dots, T\}$.*

Remark: The above theorem indicates that K-SONG also has the iteration complexity of $O(\frac{1}{\epsilon^4})$ in terms of ϵ . We refer the interested readers to Appendix E for details.

6. Practical Strategies

In this section, we present two practical strategies for improving the effectiveness of SONG/K-SONG.

Initial Warm-up. A potential problem of optimizing NDCG is that it may not lead to a good local minimum if a bad initial solution is given. To address this issue, we use warm-up to find a good initial solution by solving a well-behaved objective. Similar strategies have been used in the literature (Yuan et al., 2020; Qi et al., 2021), however, their objectives are not suitable for ranking. Here we choose the listwise cross-entropy loss (Cao et al., 2007), i.e.,

$$\min_{\mathbf{w}} \frac{1}{N} \sum_{q=1}^N \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} -\ln \left(\frac{\exp(h_q(\mathbf{x}_i^q; \mathbf{w}))}{\sum_{\mathbf{x}_j^q \in \mathcal{S}_q} \exp(h_q(\mathbf{x}_j^q; \mathbf{w}))} \right),$$

which is the cross-entropy between predicted and ground truth top-one probability distributions. The objective can be

formulated as a similar finite-sum coupled compositional problem as NDCG, and a similar algorithm to SONG can be used to solve it. We present the formulation and detailed algorithm in Appendix B.

Stop Gradient for the top- K Selector. Given a good initial solution, we justify that the second term in (7) is close to 0 under a reasonable condition, and present the details in Appendix C. Thus, the gradient of the top- K selector $\psi(h(\mathbf{x}_i^q, \mathbf{w}) - \hat{\lambda}_q(\mathbf{w}))$ is not essential. We can apply the stop gradient operator on the top- K selector, and compute the gradient estimator by

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{B}} p_{q,i} \nabla \hat{g}_{q,i}(\mathbf{w}_t), \quad (8)$$

which **simplifies** K-SONG by avoiding maintaining and updating $s_{q,t}$. We refer to the K-SONG using the gradient in (7) as theoretical K-SONG, and the K-SONG using the gradient in (8) as practical K-SONG.

7. Experiments

In this section, we evaluate our algorithms through comprehensive experiments on two different domains: learning to rank and recommender systems. Experimental results show that our algorithms can outperform prior ranking methods in terms of NDCG. We also conduct experiments to demonstrate the convergence speed of training and verify our algorithmic designs, including the moving average estimator and the bilevel formulation for K-SONG. In addition, we examine the effectiveness of initial warm-up and stop gradient operator. We implement our proposed methods in the LibAUC¹ library. To show the advantages of our library, we compare SONG and K-SONG in LibAUC with several listwise ranking approaches implemented in TensorFlow Ranking² library. The code to reproduce the results in this paper is available at <https://github.com/zhqiu/NDCG-Optimization>.

We compare our algorithms, SONG and K-SONG, against the following methods that optimize different loss functions. **RankNet** (Burges et al., 2005b) is a commonly used pairwise loss. **ListNet** (Cao et al., 2007) and **ListMLE** (Xia et al., 2008) are two listwise losses that optimize the agreement between predictions and ground truth rankings. **LambdaRank** (Burges et al., 2005a) is a listwise loss that takes NDCG into account, but not directly optimizes NDCG. **ApproxNDCG** (Qin et al., 2010) and **NeuralNDCG** (Pobrotyn & Bialobrzewski, 2021) are two losses that optimize the NDCG surrogates directly. Similar to NeuralNDCG, Pi-Rank (Swezey et al., 2021) also employs NeuralSort (Grover et al., 2019) to approximate NDCG, so we do not compare with it. We do not compare with SoftRank (Taylor et al.,

2008), as its $O(n^3)$ complexity is prohibitive.

For all methods, we sample a batch of queries, and a few (e.g., 10) relevant items and some irrelevant items for each query per iteration. For K-SONG, we report its practical version results unless specified otherwise. We use the Adam-style update for all methods and set the momentum parameters to their default values (Kingma & Ba, 2015). The hyperparameters of all losses are fine-tuned using grid search with training/validation splits mentioned below. Due to the limited space, we present the detailed implementation and datasets information in Appendix D.1 and D.2, respectively. To further show the effectiveness of our methods, we conduct more experiments on multi-label classification and provide the results in Appendix D.4.

7.1. Learning to Rank

Data. Learning to rank (LTR) algorithms aim to rank a set of candidate items for a given search query. We consider two datasets: MSLR-WEB30K (Qin & Liu, 2013) and Yahoo! LTR dataset (Chapelle & Chang, 2011), which are the largest public LTR datasets from commercial search engines. Both datasets contain query-document pairs represented by real-valued feature vectors, and have associated relevance scores on the scale from 0 to 4. Following Ai et al. (2019), we use the training/validation/test sets in the Fold1 of MSLR-WEB30K dataset for evaluation. The Yahoo! LTR dataset splits the queries arbitrarily and uses 19,944 for training, 2,994 for validation and 6,983 for testing.

Setup. For the backbone network, we adopt the Context-Aware Ranker (Pobrotyn et al., 2020), a ranking model based on the Transformer. For all methods, we first pre-train a model by initial warm-up. Then we re-initialize the last layer and train the model by different methods as mentioned before. In both stages, we set the initial learning rate and batch size to 0.001 and 64, respectively. We train the networks for 100 epochs, decaying the learning rate by 0.1 after 50 epochs. We tune γ_0 and K in our algorithms from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ and $\{10, 20, 50\}$, respectively.

Results. We evaluate all methods and calculate $\text{NDCG}@k$ ($k \in [1, 3, 5]$) on the test data. We provide partial results in Table 1, and full results in Table 5 in Appendix D.3. We notice that, in general, methods that directly optimize the NDCG surrogates achieve higher performance. Similar conclusions have been reached in other studies (Qin et al., 2010; Pobrotyn & Bialobrzewski, 2021). We also observe that our SONG and K-SONG can consistently outperform all baselines on both datasets. These results clearly show that our methods are effective for LTR tasks.

7.2. Recommender Systems

Data. Recommender systems (RS) are widely used in IT

¹<https://libauc.org/>

²<https://www.tensorflow.org/ranking>

Table 1. The test NDCG on four datasets. We report the average NDCG@3 for two LTR datasets, the average NDCG@20 for two RS datasets, and standard deviation (within brackets) over 3 runs with different random seeds. Full results are in Appendix D.3.

METHOD	NDCG@3		NDCG@20	
	MSLE WEB30K	YAHOO! LTR	MOVIELENS20M	NETFLIX PRIZE
RANKNET	0.5105±0.0004	0.7150±0.0004	0.0744±0.0013	0.0489±0.0003
LISTNET	0.5058±0.0001	0.7151±0.0004	0.0875±0.0004	0.0700±0.0002
LISTMLE	0.5074±0.0002	0.7146±0.0006	0.0799±0.0001	0.0508±0.0004
LAMBDA RANK	0.5118±0.0003	0.7155±0.0002	0.0913±0.0002	0.0693±0.0002
APPROXNDCG	0.5114±0.0005	0.7152±0.0007	0.0938±0.0003	0.0592±0.0009
NEURALNDCG	0.5101±0.0005	0.7139±0.0001	0.0901±0.0003	0.0718±0.0003
SONG	0.5136±0.0006	0.7187±0.0004	0.0969±0.0002	0.0749±0.0002
K-SONG	0.5147±0.0006	0.7191±0.0004	0.0973±0.0003	0.0743±0.0003

industry (Lu et al., 2015). We use two large-scale movie recommendation datasets: MovieLens20M (Harper & Konstan, 2015) and Netflix Prize dataset (Bennett et al., 2007). Both datasets contain large numbers of users and movies, which are represented with integer IDs. All users have rated several movies, with ratings range from 1 to 5. To create training/validation/test sets, we use the most recent rated item of each user for testing, the second recent item for validation, and the remaining items for training, which is widely-used in the literature (He et al., 2018; Wang et al., 2020a). When evaluating models, we need to collect irrelevant (unrated) items and rank them with the relevant (rated) item to compute NDCG metrics. During training, inspired by Wang et al. (2019a), we randomly sample 1000 unrated items to save time. When testing, however, we adopt the all ranking protocol (Wang et al., 2019b; He et al., 2020) — all unrated items are used for evaluation.

Setup. We choose NeuMF (He et al., 2017) as the backbone network, which is commonly used in RS tasks. For all methods, models are first pre-trained by our initial warm-up method for 20 epochs with the learning rate 0.001 and a batch size of 256. Then the last layer is randomly re-initialized and the network is fine-tuned by different methods. At the fine-tuning stage, the initial learning rate and weight decay are set to 0.0004 and 1e-7, respectively. We train the models for 120 epochs with the learning rate multiplied by 0.25 at 60 epochs. The hyper-parameters of all methods are individually tuned for fair comparison, e.g., we tune γ_0 in SONG and K-SONG from {0.1, 0.2, 0.3, 0.4, 0.5}, and K in K-SONG in a range {50, 100, 300, 500}.

Results. We evaluate all methods and calculate NDCG@ k ($k \in [10, 20, 50]$) on the test data. Part of the results are reported in Table 1, and full results are in Table 6 in Appendix D.3. First, SONG outperforms all baselines on both datasets. Specifically, SONG achieves 3.30% and 4.32% improvements on NDCG@20 over the best baseline on MovieLens20M and Netflix Prize, respectively. Besides, K-SONG performs better than SONG in most cases. These results clearly demonstrate that our algorithms are effective for optimizing NDCG and its top- K variant. It is worth to




Figure 1. Comparison of convergence of different methods in terms of validation NDCG@5 scores on two RS datasets.




Figure 2. Ablation study on two variants of SONG.

mention that the improvements from our methods on RS datasets are higher than that on LTR datasets. The reason is that RS datasets have about 20,000 items per query, while most queries in LTR datasets have less than 1,000 items (detailed statistics in Appendix D.2). These results validate that our methods are more advantageous for large-scale data.

7.3. More Studies

Convergence Speed. We plot the convergence curves for optimizing NDCG on two RS datasets in Figure 1. All convergence curves for four datasets are shown in Figure 9 in Appendix D.3. We can observe that our proposed SONG and K-SONG converge much faster than other methods.

Ablation Studies. We now study the effects of the moving average estimators in our methods and initial warm-up. We present the experimental results of two RS datasets in Figure 2 and more results in Figure 10 in Appendix D.3. First, we can observe that maintaining the moving average estimators enables our algorithm perform better. To further study the effect of γ_0 , we provide more results and analysis in Appendix D.3. Second, we consistently observe that initial warm-up can bring the model to a good initialization state and improve the final performance of the model.

Comparison with Full-Items Training. We compare three




Figure 3. Comparison of full-items and mini-batch training.




Figure 4. Comparison of theoretical and practical K-SONG.

different training methods: full-items gradient descent that uses *all items* in S_q to computing $g(\mathbf{w}; \mathbf{x}_i^q, S_q)$ and its gradient, biased mini-batch gradient descent (i.e., set $\gamma_0 = 1.0$ in our algorithms), and our algorithms (i.e., with γ_0 tuned). We compare these methods for NDCG maximization and present the results in Figure 3. We can see that our methods converge to that of full-items gradient descent, which proves the effectiveness of our algorithms. We also provide the negative loglikelihood loss curves of three different training methods for warm-up in Figure 12 in Appendix D.3, and similar conclusions can be reached.

Theoretical and Practical K-SONG. To verify the effectiveness of stop gradient operator, we present the comparison of theoretical K-SONG and practical K-SONG in Figure 4. We observe that practical K-SONG and theoretical K-SONG achieve similar performance on both datasets, which indicates that the proposed stop gradient operator is effective in simplifying theoretical K-SONG.

The advantage of the bi-level formulation. To demonstrate the advantage of our bi-level formulation for optimizing the top- K NDCG surrogate, we implement previous NDCG@ K formulation by modifying our Algorithm 1 for optimizing the NDCG@ K objective with $\psi(K - \bar{g}(\mathbf{w}, \mathbf{x}))$ in place of $\mathbb{I}(K \geq r(\mathbf{w}; \mathbf{x}))$. We compare these two formulations and present the results in Figure 5, and we can see that our bi-level formulation is more advantageous.

Comparison with TensorFlow Ranking. We implement our SONG and K-SONG into the LibAUC library and compare it with TensorFlow Ranking library (Pasumarthi et al., 2019) (TFR), which is an open-source library for neural learning to rank implemented by Google. Specifically, we compare our implementations of SONG and K-SONG with four listwise ranking methods implemented in TFR, including ListNet, ListMLE, ApproxNDCG, and Gumbel-ApproxNDCG. For all methods, models are trained for 120 epochs on MovieLens20M with the learning rate 0.001 and a batch size of 256. For SONG and K-SONG, we first train




Figure 5. Comparison of our bilevel NDCG@ K formulation and previous NDCG@ K formulation.




Figure 6. Comparison of convergence (left) and training time (right) between LibAUC (ours) and TensorFlow Ranking library.

the models by initial warm-up for the first 20 epochs, and then keep training the models by SONG or K-SONG for 100 epochs. We present the comparison of convergence and training time per epoch in Figure 6. We notice that our implementation of SONG and K-SONG in the LibAUC library with initial warm-up converge faster than the algorithms implemented in the TFR library by Google, and the training time our methods is competitive if not better than that of TFR library, which indicates the advantages of our implementations in LibAUC. In addition, our algorithm for optimizing ListNet in LibAUC is better than that implemented in TFR due to that our algorithm has rigorous convergence guarantee and that in TFR is a mini-batch based heuristic method.

8. Conclusion

In this work, we propose stochastic methods to optimize NDCG and its top- K variant that have been widely used in various ranking tasks. The optimization problems of NDCG and top- K NDCG are casted into a novel compositional optimization problem and a novel bilevel optimization problem, respectively. We design efficient stochastic algorithms with provable convergence guarantee to compute the solutions. We also study initial warm-up and stop gradient operator to improve the effectiveness for deep learning. Extensive experimental results on multiple domains demonstrate that our methods can achieve promising results.

Acknowledgements

Q. Hu, Y. Zhong and T. Yang were partially supported by NSF Grant 2110545 and NSF Career Award 1844403. Z. Qiu and L. Zhang were partially supported by NSFC (62122037, 61921006). Part work of Z. Qiu was done when he was visiting T. Yang’s lab virtually.

References

- Ai, Q., Wang, X., Bruch, S., Golbandi, N., Bendersky, M., and Najork, M. Learning groupwise multivariate scoring functions using deep neural networks. In *Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval*, pp. 85–92, 2019.
- Balasubramanian, K., Ghadimi, S., and Nguyen, A. Stochastic multi-level composition optimization algorithms with level-independent convergence rates. *arXiv preprint arXiv:2008.10526*, 2020.
- Bennett, J., Lanning, S., et al. The netflix prize. In *Proceedings of KDD Cup and Workshop*, volume 2007, pp. 35, 2007.
- Bhatia, K., Jain, H., Kar, P., Varma, M., and Jain, P. Sparse local embeddings for extreme multi-label classification. In *Advances in Neural Information Processing Systems*, volume 29, pp. 730–738, 2015.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*, pp. 89–96, 2005a.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*, pp. 89–96, 2005b.
- Burges, C. J. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11(23-581):81, 2010.
- Cao, Z., Qin, T., Liu, T., Tsai, M., and Li, H. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th International Conference on Machine Learning*, pp. 129–136, 2007.
- Chakrabarti, S., Khanna, R., Sawant, U., and Bhattacharyya, C. Structured learning for non-smooth ranking losses. In *Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 88–96, 2008.
- Chapelle, O. and Chang, Y. Yahoo! learning to rank challenge overview. In *Proceedings of the Learning to Rank Challenge*, pp. 1–24. PMLR, 2011.
- Chen, T., Sun, Y., and Yin, W. A single-timescale stochastic bilevel optimization method. *arXiv preprint arXiv:2102.04671*, 2021a.
- Chen, T., Sun, Y., and Yin, W. Solving stochastic compositional optimization is nearly as easy as solving stochastic optimization. *IEEE Transactions on Signal Processing*, 69:4937–4948, 2021b.
- Colson, B., Marcotte, P., and Savard, G. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, 2007.
- Cremonesi, P., Koren, Y., and Turrin, R. Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the 4th ACM Conference on Recommender Systems*, pp. 39–46, 2010.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Ghadimi, S. and Wang, M. Approximation methods for bilevel programming. *arXiv preprint arXiv:1802.02246*, 2018.
- Grover, A., Wang, E., Zweig, A., and Ermon, S. Stochastic optimization of sorting networks via continuous relaxations. In *the 7th International Conference on Learning Representations*, 2019.
- Guo, Z., Hu, Q., Zhang, L., and Yang, T. Randomized stochastic variance-reduced methods for multi-task stochastic bilevel optimization. *arXiv preprint arXiv:2105.02266*, 2021a.
- Guo, Z., Xu, Y., Yin, W., Jin, R., and Yang, T. On stochastic moving-average estimators for non-convex optimization. *arXiv preprint arXiv:2104.14840*, 2021b.
- Harper, F. M. and Konstan, J. A. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4):1–19, 2015.
- He, X., Liao, L., Zhang, H., Nie, L., Hu, X., and Chua, T.-S. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 173–182, 2017.
- He, X., He, Z., Song, J., Liu, Z., Jiang, Y.-G., and Chua, T.-S. Nais: Neural attentive item similarity model for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 30(12):2354–2366, 2018.
- He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., and Wang, M. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 639–648, 2020.
- Hong, M., Wai, H.-T., Wang, Z., and Yang, Z. A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic. *arXiv preprint arXiv:2007.05170*, 2020.

- Hu, Y., Zhang, S., Chen, X., and He, N. Biased stochastic first-order methods for conditional stochastic optimization and applications in meta learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- Järvelin, K. and Kekäläinen, J. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002.
- Ji, K., Yang, J., and Liang, Y. Provably faster algorithms for bilevel optimization and applications to meta-learning. *arXiv preprint arXiv:2010.07962*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *the 3rd International Conference on Learning Representations*, 2015.
- Kunisch, K. and Pock, T. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, 2013.
- Lin, T., Jin, C., and Jordan, M. I. On gradient descent ascent for nonconvex-concave minimax problems. *arXiv preprint arXiv:1906.00331*, 2019.
- Liu, J., Chang, W.-C., Wu, Y., and Yang, Y. Deep learning for extreme multi-label text classification. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*, pp. 115–124, 2017.
- Liu, N. N. and Yang, Q. Eigenrank: a ranking-oriented approach to collaborative filtering. In *Proceedings of the 31st International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 83–90, 2008.
- Liu, R., Mu, P., Yuan, X., Zeng, S., and Zhang, J. A generic first-order algorithmic framework for bi-level programming beyond lower-level singleton. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 6305–6315, 2020.
- Liu, T.-Y. *Learning to Rank for Information Retrieval*. Springer, 2011.
- Lu, J., Wu, D., Mao, M., Wang, W., and Zhang, G. Recommender system application developments: a survey. *Decision Support Systems*, 74:12–32, 2015.
- Mencia, E. L. and Fürnkranz, J. Efficient pairwise multilabel classification for large-scale problems in the legal domain. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 50–65. Springer, 2008.
- Nesterov, Y. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005.
- Pasumarthi, R. K., Bruch, S., Wang, X., Li, C., Bender-sky, M., Najork, M., Pfeifer, J., Golbandi, N., Anil, R., and Wolf, S. Tf-ranking: Scalable tensorflow library for learning-to-rank. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2970–2978, 2019.
- Pobrotyn, P. and Białobrzeski, R. Neuralndcg: Direct optimisation of a ranking metric via differentiable relaxation of sorting. *arXiv preprint arXiv:2102.07831*, 2021.
- Pobrotyn, P., Bartczak, T., Synowiec, M., Białobrzeski, R., and Bojar, J. Context-aware learning to rank with self-attention. *arXiv preprint arXiv:2005.10084*, 2020.
- Qi, Q., Luo, Y., Xu, Z., Ji, S., and Yang, T. Stochastic optimization of area under precision-recall curve for deep learning with provable convergence. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- Qin, T. and Liu, T.-Y. Introducing letor 4.0 datasets. *arXiv preprint arXiv:1306.2597*, 2013.
- Qin, T., Liu, T.-Y., and Li, H. A general approximation framework for direct optimization of information retrieval measures. *Information Retrieval*, 13(4):375–397, 2010.
- Singh, M. Scalability and sparsity issues in recommender datasets: a survey. *Knowledge and Information Systems*, 62(1):1–43, 2020.
- Swezey, R., Grover, A., Charron, B., and Ermon, S. Pirank: Scalable learning to rank via differentiable sorting. *Advances in Neural Information Processing Systems*, 34, 2021.
- Taylor, M., Guiver, J., Robertson, S., and Minka, T. Soft-rank: optimizing non-smooth rank metrics. In *Proceedings of the 2008 International Conference on Web Search and Web Data Mining*, pp. 77–86, 2008.
- Valizadegan, H., Jin, R., Zhang, R., and Mao, J. Learning to rank by optimizing ndcg measure. In *Advances in Neural Information Processing Systems*, volume 22, pp. 1883–1891, 2009.
- Wang, B. and Yang, T. Finite-sum coupled compositional stochastic optimization: Theory and applications. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.
- Wang, C., Zhang, M., Ma, W., Liu, Y., and Ma, S. Modeling item-specific temporal dynamics of repeat consumption for recommender systems. In *Proceedings of the 28th International Conference on World Wide Web*, pp. 1977–1987, 2019a.

- Wang, C., Zhang, M., Ma, W., Liu, Y., and Ma, S. Make it a chorus: knowledge-and time-aware item modeling for sequential recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 109–118, 2020a.
- Wang, C., Zhang, M., Ma, W., Liu, Y., and Ma, S. Make it a chorus: knowledge-and time-aware item modeling for sequential recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 109–118, 2020b.
- Wang, M., Fang, E. X., and Liu, H. Stochastic compositional gradient descent: algorithms for minimizing compositions of expected-value functions. *Mathematical Programming*, 161(1-2):419–449, 2017.
- Wang, X., He, X., Wang, M., Feng, F., and Chua, T.-S. Neural graph collaborative filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 165–174, 2019b.
- Wu, M., Chang, Y., Zheng, Z., and Zha, H. Smoothing dcg for learning to rank: A novel approach using smoothed hinge functions. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, pp. 1923–1926, 2009.
- Xia, F., Liu, T.-Y., Wang, J., Zhang, W., and Li, H. List-wise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th International Conference on Machine Learning*, pp. 1192–1199, 2008.
- Xu, J. and Li, H. Adarank: a boosting algorithm for information retrieval. In *Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 391–398, 2007.
- Yang, T. and Lin, Q. Rsg: Beating subgradient method without smoothness and strong convexity. *Journal of Machine Learning Research*, 19(6):1–33, 2018.
- Yeh, J.-Y., Lin, J.-Y., Ke, H.-R., and Yang, W.-P. Learning to rank for information retrieval using genetic programming. In *Proceedings of SIGIR 2007 Workshop on Learning to Rank for Information Retrieval*, 2007.
- Yin, H., Wang, Q., Zheng, K., Li, Z., and Zhou, X. Overcoming data sparsity in group recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- Yuan, T., Cheng, J., Zhang, X., Qiu, S., and Lu, H. Recommendation by mining multiple user behaviors with group sparsity. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- Yuan, Z., Yan, Y., Sonka, M., and Yang, T. Robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification. *arXiv preprint arXiv:2012.03173*, 2020.
- Zubiaga, A. Enhancing navigation on wikipedia with social tags. *arXiv preprint arXiv:1202.5469*, 2012.

A. Per-iteration Complexity

For complexity analysis, let \mathcal{Q}_t denote the sampled queries at the t -th iteration and \mathcal{B}_q^+ denote the sampled relevant items for each sampled query. In terms of the per-iteration complexity of SONG, we need to conduct forward propagation for computing $h_q(\mathbf{x}_i^q, \mathbf{w}), \forall \mathbf{x}_i^q \in \mathcal{B}_q^+ \cup \mathcal{B}_q$ and back-propagation for computing $\nabla h_q(\mathbf{x}_i^q, \mathbf{w}), \forall \mathbf{x}_i^q \in \mathcal{B}_q^+ \cup \mathcal{B}_q$. The complexity for these forward propagations and back-propagations is $\sum_{q \in \mathcal{Q}_t} (|\mathcal{B}_q^+| + |\mathcal{B}_q|)d \leq O(Bd)$, where $B = \sum_{q \in \mathcal{Q}_t} (|\mathcal{B}_q^+| + |\mathcal{B}_q|)$ is the total mini-batch size. With these computed, the cost for computing $\hat{g}_{q,i}(\mathbf{w}_t)$ and $\nabla \hat{g}_{q,i}(\mathbf{w}_t)$ for all $q, \mathbf{x}_i^q \in \mathcal{B}_q^+$ is $\sum_{q \in \mathcal{Q}_t} |\mathcal{B}_q^+| |\mathcal{B}_q| \leq O(B^2)$. Hence, the total complexity per iteration is $O(Bd + B^2)$. For a large model size $d \gg B$, we have the per-iteration complexity of $O(Bd)$, which is similar to the standard cost of deep learning per-iteration and is independent of the length of \mathcal{S}_q for each query.

B. Initial Warm-up

The listwise cross-entropy loss can be reformulated as follows:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{N} \sum_{q=1}^N \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} -\ln \left(\frac{\exp(h_q(\mathbf{x}_i^q; \mathbf{w}))}{\sum_{\mathbf{x}_j^q \in \mathcal{S}_q} h_q(\mathbf{x}_j^q; \mathbf{w})} \right) \\ & = \frac{1}{N} \sum_{q=1}^N \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q^+} \ln \left(\sum_{\mathbf{x}_j^q \in \mathcal{S}_q} \exp(h_q(\mathbf{x}_j^q) - h_q(\mathbf{x}_i^q)) \right). \end{aligned}$$

The above objective has the same structure of the NDCG surrogate, i.e., it is an instance of finite-sum coupled compositional stochastic optimization problem. Hence, we can use a similar algorithm to SONG to solve the above problem. We present the details in Algorithm 3.

Algorithm 3 Stochastic Optimization of Listwise CE loss: SOLC

Require: $\eta, \beta_0, \beta_1, u^{(1)} = 0$

Ensure: \mathbf{w}_T

for $t = 1, \dots, T$ **do**

draw a set of queries denoted by \mathcal{Q}_t

For each query draw a batches of examples $\{\mathcal{B}_q^+, \mathcal{B}_q\}$, where \mathcal{B}_q^+ denote a set of sampled relevant documents for q and \mathcal{B}_q denote a set of sampled documents from \mathcal{S}_q

for $\mathbf{x}_i^q \in \mathcal{B}_q^+$ for each $q \in \mathcal{Q}_t$ **do**

$$u_{q,i}^{(t+1)} = (1 - \gamma_0)u_{q,i}^{(t)} + \gamma_0 \frac{1}{|\mathcal{B}_q|} \sum_{\mathbf{x}' \in \mathcal{B}_q} \exp(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}))$$

Compute $p_{q,i} = 1/u_{q,i}^{t+1}$

end for

Compute gradient

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{Q}_t|} \frac{1}{|\mathcal{B}_q^+|} \frac{1}{|\mathcal{B}_q|} \sum_{q \in \mathcal{Q}_t} \sum_{\mathbf{x}_i^q \in \mathcal{B}_q^+} \sum_{\mathbf{x}_j^q \in \mathcal{B}_q} p_{q,i} \nabla_{\mathbf{w}} (h_q(\mathbf{x}_j^q; \mathbf{w}_t) - h_q(\mathbf{x}_i^q; \mathbf{w}_t))$$

Compute $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1)G(\mathbf{w}_t)$



Update $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{m}_{t+1}$

end for

C. Justification of Stop Gradient Operator

Below, we provide a justification by showing that the second term in (7) is close to 0 under a reasonable condition. For simplicity of notation, we let $\psi_i(\mathbf{w}, \hat{\lambda}_q(\mathbf{w})) = \psi(h(\mathbf{x}_i^q, \mathbf{w}) - \hat{\lambda}_q(\mathbf{w}))$. Its gradient is given by

$$\nabla_{\mathbf{w}} \psi_i = \psi'_i(\mathbf{w}, \hat{\lambda}_q(\mathbf{w})) \left(\nabla_{\mathbf{w}} h(\mathbf{x}_i^q, \mathbf{w}) - \nabla_{\mathbf{w}}^2 L_q(\mathbf{w}, \hat{\lambda}_q(\mathbf{w})) [\nabla_{\lambda}^2 L_q(\mathbf{w}, \hat{\lambda}_q(\mathbf{w}))]^{-1} \right).$$


 Figure 7. Curves of $[\cdot]_+$, $\phi(\cdot)$, and $\kappa(\cdot)$.

 Figure 8. An example of $\psi(\cdot)$

For the purpose of justification, we can approximate $\phi(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda) = \tau_1 \log(1 + \exp((h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)/\tau_1))$ by a smoothed hinge loss function, $\kappa(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda) = \max_{\alpha} \alpha(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda) - \tau_1 \alpha^2/2$, which is equivalent to

$$\kappa(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda) = \begin{cases} 0, & h_q(\mathbf{x}_i; \mathbf{w}) - \lambda \leq 0 \\ \frac{(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)^2}{2\tau_1}, & 0 < h_q(\mathbf{x}_i; \mathbf{w}) - \lambda \leq \tau_1 \\ h_q(\mathbf{x}_i; \mathbf{w}) - \lambda - \frac{\tau_1}{2}, & h_q(\mathbf{x}_i; \mathbf{w}) - \lambda > \tau_1 \end{cases}$$

Please refer to Figure 7 for the curves of $[\cdot]_+$ and $\phi(\cdot)$ and $\kappa(\cdot)$. Below, we assume $L_q(\mathbf{w}, \lambda)$ is defined by using $\kappa(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)$ in place of $\phi(h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)$.

For any \mathbf{w} , let us consider a subset $\mathcal{C}_q = \{\mathbf{x}_i^q \in \mathcal{S}_q^+ : h_{\mathbf{w}}(\mathbf{x}_i^q) - \hat{\lambda}_q(\mathbf{w}) \in (0, \tau_1)\}$. It is not difficult to show that

$$\begin{aligned} \nabla_{\mathbf{w}\lambda}^2 L_q(\mathbf{w}, \hat{\lambda}_q(\mathbf{w})) &= \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} \frac{-\partial_{\mathbf{w}} h(\mathbf{x}_i^q; \mathbf{w})}{\tau_1} \\ \nabla_{\lambda}^2 L_q(\mathbf{w}, \hat{\lambda}_q(\mathbf{w})) &= \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} \frac{1}{\tau_1} + \tau_2 \approx \frac{1}{N_q} \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} \frac{1}{\tau_1} \end{aligned}$$

for sufficiently small τ_1, τ_2 . Then we have

$$\frac{\nabla_{\mathbf{w}\lambda}^2 L_q(\mathbf{w}, \lambda(\mathbf{w}))}{\nabla_{\lambda}^2 L_q(\mathbf{w}, \lambda(\mathbf{w}))} = \frac{1}{|\mathcal{C}_q|} \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} -\partial_{\mathbf{w}} h(\mathbf{x}_i^q).$$

Assume that ψ is chosen such that $\psi'_i(\mathbf{w}, \lambda_q(\mathbf{w})) \approx 0$ if $h_{\mathbf{w}}(\mathbf{x}_j^q) - \lambda_q(\mathbf{w}) \notin [0, \tau_1]$, and $\psi'_i(\mathbf{w}, \lambda_q(\mathbf{w})) \approx c_1$ and $f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)) \approx c_2$ if $h_{\mathbf{w}}(\mathbf{x}_j^q) - \lambda_q(\mathbf{w}) \in [0, \tau_1]$, then we have

$$\begin{aligned} \sum_{\mathbf{x}_i^q \in \mathcal{S}_q} \nabla_{\mathbf{w}} \psi_i f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)) &\approx \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} \psi'_i(\mathbf{w}, \lambda_q(\mathbf{w})) \cdot \left(\nabla_{\mathbf{w}} h(\mathbf{x}_i^q; \mathbf{w}) - \frac{1}{|\mathcal{C}_q|} \sum_{\mathbf{x}_j^q \in \mathcal{C}_q} \nabla_{\mathbf{w}} h(\mathbf{x}_j^q; \mathbf{w}) \right) f_{q,i}(g(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q)) \\ &\approx c_1 c_2 \sum_{\mathbf{x}_i^q \in \mathcal{C}_q} \left(\nabla_{\mathbf{w}} h_{\mathbf{w}}(\mathbf{x}_i^q; \mathbf{w}) + \frac{1}{|\mathcal{C}_q|} \sum_{\mathbf{x}_j^q \in \mathcal{C}_q} -\nabla_{\mathbf{w}} h(\mathbf{x}_j^q; \mathbf{w}) \right) \\ &= 0 \end{aligned}$$

As a result, when τ_1 is small enough the condition $\psi'_i(\mathbf{w}, \lambda_q(\mathbf{w})) \approx 0$ if $h_{\mathbf{w}}(\mathbf{x}_j^q) - \lambda_q(\mathbf{w}) \notin [0, \tau_1]$, and $\psi'_i(\mathbf{w}, \lambda_q(\mathbf{w})) \approx c$ if $h_{\mathbf{w}}(\mathbf{x}_j^q) - \lambda_q(\mathbf{w}) \in [0, \tau_1]$ is well justified. An example of such $\psi(\cdot)$ is provided in the Figure 8. As a result, with initial warm-up, we can compute the gradient estimator by

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}|} \sum_{(q, \mathbf{x}_i^q) \in \mathcal{B}} p_{q,i} \nabla \hat{g}_{q,i}(\mathbf{w}_t),$$

which simplifies K-SONG by avoiding maintaining and updating $s_{q,t}$.

Table 2. Statistics of Learning to Rank Datasets.

DATASET	MSLR-WEB30K	YAHOO! LTR DATASET
QUERY	30,000	29,921
Q-D PAIR	3,771,125	709,877
MAX Q-D PAIR PER QUERY	1,245	135
MIN Q-D PAIR PER QUERY	1	1

Table 3. Statistics of Recommender Systems Datasets.

DATASET	# USERS	# ITEMS	# INTERACTIONS	SPARSITY
MOVIELENS20M	138,493	26,744	20,000,263	99.46%
NETFLIX PRIZE DATASET	236,117	17,770	89,973,534	97.86%

D. Experiments

D.1. Details of Implementation

For the experiments on two LTR datasets, we adopt allRank framework³ (Pobrotyn et al., 2020). We implement some baseline methods based on their code. For the recommender systems experiments, we use ReChorus framework⁴ (Wang et al., 2020b), which is a general PyTorch framework for Top-K recommendation. We also follow the scripts in ReChorus to preprocess the datasets. The hyper-parameters for SONG and K-SONG are presented in Table 4.

We train our models on one Tesla V100 GPU with 32GB memory. The training on the Context-Aware Ranker model takes about 2~3 hours for convergence, while the training of the NeuMF model takes about 8~12 hours for convergence.

D.2. Details of Data

MSLR-WEB30K⁵ and Yahoo! LTR dataset⁶ are the largest public LTR datasets from commercial English search engines. We provide the statistics of these two datasets in Table 2. In MSLR-WEB30K dataset, there are 5 folds containing the same data, and each fold randomly splits to training, validation, and test sets. Due to privacy concerns, these datasets do not disclose any text information and only provide feature vectors for each query-document pair. For these two LTR datasets, we standardize the features, log-transforming selected ones, before feeding them to the learning algorithms. Since the lengths of search results lists in the datasets are unequal, we truncate or pad samples to the length of 40 and 100 for Yahoo! LTR dataset and MSLR-WEB30K when training, respectively, but use the full list for evaluation.

MovieLens20M⁷ contains 20 million ratings applied to 27,000 movies by 138,000 users, and all users have rated at least 20 movies. Netflix Prize dataset⁸ consists of about 100,000,000 ratings for 17,770 movies given by 480,189 users. We filter the Netflix Prize dataset by retaining users with at least 100 interactions to cater sufficient information for modeling. In both datasets, users and movies are represented with integer IDs, while ratings range from 1 to 5. The statistics of these two datasets are shown in Table 3.

D.3. Additional Experimental Results

Convergence Speed. We present the training curves on four different datasets (MovieLens20M, Netflix Prize dataset, MSLR WEB30K, and Yahoo! LTR dataset) in Figure 9.

Ablation Studies. We provide the full ablation studies on four datasets in Figure 10.

The Effect of Varying γ_0 . We adjust γ_0 in our algorithms from $\{0.1, 0.3, 0.5, 0.7, 1.0\}$, and report the training curves of

³<https://github.com/allegro/allRank>

⁴<https://github.com/THUwangcy/ReChorus>

⁵<https://www.microsoft.com/en-us/research/project/mslr/>

⁶<https://webscope.sandbox.yahoo.com>

⁷<https://grouplens.org/datasets/movielens/20m/>

⁸<https://www.kaggle.com/netflix-inc/netflix-prize-data>

Table 4. Hyper-parameters for SONG and K-SONG.

	MOVIELENS20M	NETFLIX PRIZE	MSLR-WEB30K	YAHOO! LTR
γ_0	0.1	0.3	0.3	0.3
K	300	300	10	10

Table 5. The test NDCG on two Learning to Rank datasets. We report the average NDCG@ k ($k \in [1, 3, 5]$) and standard deviation (within brackets) over 3 runs with different random seeds.

METHOD	MSLR WEB30K			YAHOO! LTR DATASET		
	NDCG@1	NDCG@3	NDCG@5	NDCG@1	NDCG@3	NDCG@5
RANKNET	0.5138±0.0008	0.5105±0.0004	0.5159±0.0003	0.7066±0.0006	0.7150±0.0004	0.7368±0.0005
LISTNET	0.5105±0.0001	0.5058±0.0001	0.5146±0.0002	0.7066±0.0002	0.7151±0.0004	0.7352±0.0004
LISTMLE	0.5153±0.0012	0.5074±0.0002	0.5136±0.0005	0.7067±0.0008	0.7146±0.0006	0.7353±0.0007
LAMBDA RANK	0.5173±0.0014	0.5118±0.0003	0.5187±0.0003	0.7084±0.0003	0.7155±0.0002	0.7352±0.0004
APPROXNDCG	0.5204±0.0007	0.5114±0.0005	0.5179±0.0006	0.7085±0.0009	0.7152±0.0007	0.7350±0.0006
NEURALNDCG	0.5160±0.0006	0.5101±0.0005	0.5155±0.0002	0.7076±0.0003	0.7139±0.0001	0.7349±0.0003
SONG	0.5265±0.0005	0.5136±0.0006	0.5206±0.0003	0.7131±0.0002	0.7187±0.0004	0.7390±0.0002
K-SONG	0.5271±0.0006	0.5147±0.0006	0.5204±0.0003	0.7128±0.0004	0.7191±0.0004	0.7394±0.0008




Figure 9. Comparison of convergence of different methods in terms of validation NDCG@5 scores.




Figure 10. Ablation study on two variants of SONG on four different datasets.

warm-up and SONG in Figure 11. We observe that $\gamma_0 = 0.1$ achieves the best performance in most cases. Setting $\gamma_0 = 1.0$ is equivalent to update the model with a biased stochastic gradient, which leads to the worst performance. These results signify the importance of moving average estimators in our methods.




Figure 11. The effect of varying γ_0 for warm-up (left two) and SONG (right two).

Comparison with Full-Items Training. We provide the negative loglikelihood loss curves of three different training methods for warm-up in Figure 12.

Table 6. The test NDCG on two movie recommendation datasets. We report the average NDCG@k (k ∈ [10, 20, 50]) and standard deviation (within brackets) over 3 runs with different random seeds.

METHOD	MOVIELENS20M			NETFLIX PRIZE DATASET		
	NDCG@10	NDCG@20	NDCG@50	NDCG@10	NDCG@20	NDCG@50
RANKNET	0.0538±0.0011	0.0744±0.0013	0.1086±0.0013	0.0362±0.0002	0.0489±0.0003	0.0730±0.0003
LISTNET	0.0660±0.0003	0.0875±0.0004	0.1227±0.0003	0.0532±0.0002	0.0700±0.0002	0.0992±0.0002
LISTMLE	0.0588±0.0001	0.0799±0.0001	0.1137±0.0001	0.0376±0.0003	0.0508±0.0004	0.0753±0.0001
LAMBDA RANK	0.0697±0.0001	0.0913±0.0002	0.1259±0.0001	0.0531±0.0002	0.0693±0.0002	0.0976±0.0003
APPROXNDCG	0.0735±0.0005	0.0938±0.0003	0.1284±0.0002	0.0434±0.0005	0.0592±0.0009	0.0873±0.0012
NEURALNDCG	0.0692±0.0003	0.0901±0.0003	0.1232±0.0007	0.0554±0.0002	0.0718±0.0003	0.1003±0.0002
SONG	0.0748±0.0002	0.0969±0.0002	0.1326±0.0001	0.0571±0.0002	0.0749±0.0002	0.1050±0.0003
K-SONG	0.0747±0.0002	0.0973±0.0003	0.1340±0.0001	0.0573±0.0003	0.0743±0.0003	0.1042±0.0001




Figure 12. Comparison of full-items and mini-batch training on SONG (left two) and warm-up (right two).

D.4. Experiments on Multi-label Classification

To further verify the effectiveness of our methods, we also conduct experiments on multi-label datasets. Similar to Learning to Rank task, we treat each instance as a query and each label as an item. We adopt XML-CNN⁹ (Liu et al., 2017) as our base model. We download data from The Extreme Classification Repository¹⁰ and conduct experiments on two datasets: EUR-Lex (Mencia & Fürnkranz, 2008) and Wiki10-31K (Zubiaga, 2012). The statistics of these two datasets are presented in Table 7. In our experiments, we use raw data to classify.

Table 7. Statistics of Multi-label Datasets.

Dataset	Labels	Training Samples	Testing Samples	Avg. Points per Label	Avg. Labels per Points
EURLex-4K	3,993	15,539	3,809	25.73	5.31
Wiki10-31K	309,38	14,146	6,616	8.52	18.64

In extreme multi-label (XML) classification, label spaces usually are large; however, each instance only has very few relevant labels. Therefore, we adopt NDCG@k as our evaluation metric, which is also a common way in evaluating XML methods.

Upon XML-CNN, we compare our method with other NDCG optimization methods: ApproxNDCG (Qin et al., 2010) and NeuralNDCG (Pobrotyn & Bialobrzewski, 2021). The results are summarized in Table 8.

Table 8. Results in NDCG@k; bold indicates the best performance among all methods

Datasets	Metrics	Baseline	ApproxNDCG	NeuralNDCG	SONG	K-SONG
EUR-Lex	NDCG@3	67.15	66.59	67.68	67.84	68.11
	NDCG@5	61.13	60.23	61.86	61.32	61.74
Wiki10-31K	NDCG@3	71.26	71.49	71.52	72.90	74.01
	NDCG@5	63.23	62.43	62.85	65.10	66.15

⁹<https://github.com/siddsax/XML-CNN>

¹⁰<http://manikvarma.org/downloads/XC/XMLRepository.html#Bi13>

E. Convergence Analysis

E.1. Analysis of SONG

For simplicity, we rewrite problem (3) as the following compositional optimization problem,

$$\min_{\mathbf{w}} \quad \frac{1}{n} \sum_{i \in \mathcal{S}} f_i(g_i(\mathbf{w})). \quad (9)$$

One may reorder the set of queries \mathcal{S} so that each pair (q, \mathbf{x}_i^q) has a single index. We abuse the notation \mathcal{S} denoting the set of the new indexing. Then the equivalence between problem (3) and (9) is established. Furthermore, SONG can be rewritten as Algorithm 4 accordingly. In fact, problem (9) can be seen as a special case of problem (10), where ψ_i 's are constant functions. Hence, Theorem 4.2 naturally follows from Theorem 5.4, of which the proof will be presented in the following section.

Algorithm 4

Require: $\mathbf{w}_0, \mathbf{m}_0, u^0, \gamma_0, \beta_1, \eta$

Ensure: \mathbf{w}_T

for $t = 0, 1, \dots, T - 1$ **do**

 Draw batch of queries $\mathcal{B}_1^t \in \{1, \dots, n\}$

 Draw batch of items $\mathcal{B}_{2,i}^t$ for each $i \in \mathcal{B}_1^t$

 Compute $u_i^{t+1} = \begin{cases} (1 - \gamma_0)u_i^t + \gamma_0 g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t \\ u_i^t & \text{o.w.} \end{cases}$

 Compute stochastic gradient estimator $G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \nabla g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) \nabla f_i(u_i^t)$

$\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1)G(\mathbf{w}_t)$

$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{m}_{t+1}$

end for

E.2. Analysis of K-SONG

In this section, we present a convergence analysis for K-SONG. Similarly to the analysis of SONG, we reorder the set of queries \mathcal{S} and generalize problem (4) into the following compositional bilevel optimization problem,

$$\begin{aligned} \min_{\mathbf{w}} \quad & F(\mathbf{w}) := \frac{1}{n} \sum_{i \in \mathcal{S}} \psi_i(\mathbf{w}, \lambda_i(\mathbf{w})) f_i(g_i(\mathbf{w})) \\ \text{s.t.} \quad & \lambda_i(\mathbf{w}) = \arg \min_{\lambda} L_i(\mathbf{w}, \lambda). \end{aligned} \quad (10)$$

This allows us to rewrite K-SONG into Algorithm 5 accordingly.

Notations: Throughout this convergence analysis section, all subscript i represents the block of variable or function corresponding to the i th query. The following notations will be used,

$$\delta_{\lambda,t} := \|\lambda(\mathbf{w}_t) - \lambda^t\|^2, \quad \delta_{g,t} := \|g(\mathbf{w}_t) - u^t\|^2, \quad \delta_{L\lambda\lambda,t} := \|\nabla_{\lambda\lambda}^2 L(\mathbf{w}_t, \lambda(\mathbf{w}_t)) - s^t\|^2$$

We make the following assumptions regarding problem (10).

Assumption E.1.

- Functions ψ_i, f_i, g_i, L_i are L_ψ, L_f, L_g, L_L -smooth respectively for all i .
- Functions ψ_i, f_i, g_i are C_ψ, C_f, C_g -Lipschitz continuous respectively for all i . Function L_i is μ_L -strongly convex for all i .
- $\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}, \lambda), \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}, \lambda)$ are $L_{L\mathbf{w}\lambda}, L_{L\lambda\lambda}$ -Lipschitz continuous respectively with respect to (\mathbf{w}, λ) for all i .
- ψ_i and f_i are bounded by B_ψ and B_f respectively, i.e. $\|\psi_i(\mathbf{w}, \lambda)\| \leq B_\psi$ and $\|f_i(g)\| \leq B_f$ for all $\mathbf{w}, \lambda, i, g$.
- $\|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}, \lambda)\|^2 \leq C_{L\mathbf{w}\lambda}^2, \gamma I \preceq \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}, \lambda; \mathcal{B}) \preceq L_L I$ for all i
- Unbiased stochastic oracles $g_i, \nabla g_i, \nabla_\lambda L_i, \nabla_{\lambda\lambda}^2 L_i, \nabla_{\mathbf{w}\lambda}^2 L_i$ have bounded variance σ^2 .

Algorithm 5
Require: $\mathbf{w}_0, \mathbf{m}_0, \lambda^0, u^0, s^0, \gamma_0, \gamma'_0, \beta_1, \eta_0, \eta_1$
Ensure: \mathbf{w}_T
for $t = 0, 1, \dots, T - 1$ **do**

 Draw batch of queries $\mathcal{B}_1^t \in \{1, \dots, n\}$

 Draw batch of items $\mathcal{B}_{2,i}^t$ for each $i \in \mathcal{B}_1^t$

 Compute $u_i^{t+1} = \begin{cases} (1 - \gamma_0)u_i^t + \gamma_0 g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t \\ u_i^t & \text{o.w.} \end{cases}$

 Compute $\lambda_i^{t+1} = \begin{cases} \lambda_i^t - \eta_0 \nabla_{\lambda} L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t \\ \lambda_i^t & \text{o.w.} \end{cases}$

 Compute $s_i^{t+1} = \begin{cases} (1 - \gamma'_0)s_i^t + \gamma'_0 \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t \\ s_i^t & \text{o.w.} \end{cases}$

 Compute stochastic gradient estimator $G(\mathbf{w}_t)$ according to (11)

 $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1)G(\mathbf{w}_t)$
 $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_1 \mathbf{m}_{t+1}$
end for

Now we show that problem (4) satisfies Assumption E.1. Here we consider the squared hinge loss $\ell(h_q(\mathbf{x}'; \mathbf{w}), h_q(\mathbf{x}; \mathbf{w})) = \max\{0, h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) + c\}^2$ where c is a margin parameter. Suppose the score function and its gradients $h_q(\mathbf{x}; \mathbf{w}), \nabla_{\mathbf{w}} h_q(\mathbf{x}; \mathbf{w}), \nabla_{\mathbf{w}}^2 h_q(\mathbf{x}; \mathbf{w})$ are bounded by finite constants $c_h, c_{h'}, c_{h''}$ respectively. As an average of squared hinge loss, function $g_i(\mathbf{w})$ in problem (10) has bounded gradients $\nabla g_i(\mathbf{w}) \leq 8c_h c_{h'}$ and $\nabla^2 g_i(\mathbf{w}) \leq 8c_{h'}^2 + 8c_h c_{h''}$ for each $i \in \mathcal{S}$. Hence g_i is Lipschitz continuous and smooth. Moreover, with $m > 2c_h$, there exists $c_\ell > 0$ such that $\ell(h_q(\mathbf{x}_1; \mathbf{w}) - h_q(\mathbf{x}_2; \mathbf{w})) \geq c_\ell$ for all $\mathbf{x}_1, \mathbf{x}_2$. Function $f_i(g) = f_{q,i}(g) = \frac{1}{Z_q} \frac{1 - 2^{g_i}}{\log_2(N_q g + 1)}$ is thus bounded, Lipschitz continuous and smooth for each $i = (q, \mathbf{x}_i) \in \mathcal{S}$. For function $\psi_i = \psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda)$, we consider the logistic loss, then ψ_i is naturally bounded, Lipschitz continuous and smooth. The smoothness and strong convexity of L_i are proved in Lemma 5.3. In fact, the strong convexity of L_i implies the lower boundedness $\gamma = \tau_2$ of $\nabla_{\lambda\lambda} L_i(\mathbf{w}, \lambda; \mathcal{B})$. To show the Lipschitz continuity of $\nabla_{\mathbf{w}\lambda}^2 L_q(\lambda; \mathbf{w})$ and $\nabla_{\lambda\lambda}^2 L_q(\lambda; \mathbf{w})$ one may simply take the third gradients of $L_q(\lambda; \mathbf{w})$ and use the fact $\exp(\frac{\lambda - h_q(\mathbf{x}_i; \mathbf{w})}{\tau_1}) > 0$ and the assumption of the boundedness of $h_q(\mathbf{x}; \mathbf{w})$ and its gradients to verify.

By using the implicit function theorem, the stochastic gradient estimator of $\nabla F(\mathbf{w}_t)$ in Algorithm 5 is given by

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \left[\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) [s_i^t]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t) \right] f_i(u_i^t) \quad (11)$$

$$+ \psi_i(\mathbf{w}_t, \lambda_i^t) \nabla g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) \nabla f_i(u_i^t)$$

Note that the parameter τ_λ in the update of λ^{t+1} exists only for theoretical analysis reason. In practical, $\tau_\lambda \eta_0$ can be treated as one parameter. Moreover, we define the gradient approximation at iteration t

$$\nabla F(\mathbf{w}_t, \lambda^t, u^t) = \frac{1}{n} \sum_{i \in \mathcal{S}} \left[\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t) \right] f_i(u_i^t)$$

$$+ \psi_i(\mathbf{w}_t, \lambda_i^t) \nabla g_i(\mathbf{w}_t) \nabla f_i(u_i^t)$$

Now we present the formal statement of Theorem 5.4 regarding to problem (10).

Theorem E.2. Let $F(\mathbf{w}_0) - F(\mathbf{w}^*) \leq \Delta_F$. Under Assumption E.1 and consider Algorithm 5, with $\eta_0 \leq \min\left\{\frac{\mu_L}{L_F^2}, \frac{2n}{|\mathcal{B}_1^t| \mu_L}, \frac{\mu_L \epsilon^2}{48C_5 \sigma^2}\right\}$, $\gamma_0 \leq \left\{\frac{1}{2}, \frac{\epsilon^2}{96C_6 \sigma^2}\right\}$, $\gamma'_0 \leq \left\{1, \frac{\epsilon^2}{96C_7 \sigma^2}\right\}$, $\gamma_1 \leq \frac{\epsilon^2}{12(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2)}$, $\beta_1 = 1 - \gamma_1$, $\eta_1^2 \leq \min\left\{\frac{\gamma_1^2}{64L_F^2}, \frac{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2}{128n^2 C_5 C_\lambda^2}, \frac{|\mathcal{B}_1^t|^2 \gamma_0^2}{128n^2 C_6 C_g^2}, \frac{|\mathcal{B}_1^t|^2 \gamma_0^2}{512n^2 C_7 L_{\lambda\lambda}^2 (1 + C_\lambda^2)}\right\}$, $T \geq \left\{\frac{30\Delta_F}{\eta_1 \epsilon^2}, \frac{15\mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2]}{\gamma_1 \epsilon^2}, \frac{30C_5 \delta_{\lambda,0}}{|\mathcal{B}_1^t| \eta_0 \mu_L \epsilon^2}, \frac{30C_6 \delta_{g,0}}{|\mathcal{B}_1^t| \gamma_0 \epsilon^2}, \frac{60C_7 \delta_{L\lambda\lambda,0}}{|\mathcal{B}_1^t| \gamma_0^2 \epsilon^2}\right\}$, we have

$$\mathbb{E}[\|\nabla F(\mathbf{w}_\tau)\|^2] \leq \epsilon^2, \quad \mathbb{E}[\|\nabla F(\mathbf{w}_\tau) - \mathbf{m}_{\tau+1}\|^2] < 2\epsilon^2,$$

where τ is randomly sampled from $\{0, \dots, T\}$, C_5, C_6, C_7, C_8 are constants defined in the proof, and L_F is the Lipschitz continuity constant of $\nabla F(\mathbf{w})$.

E.3. Proof of Theorem E.2

To prove Theorem E.2, we first present some required Lemmas.

Lemma E.3. *Under Assumption E.1, $F(\mathbf{w})$ is L_F -smooth for some constant $L_F \in \mathbb{R}$.*

Lemma E.4. *Consider the update $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_1 \mathbf{m}_{t+1}$. Then under Assumption E.1, with $\eta_1 L_F \leq \frac{1}{2}$, we have*

$$F(\mathbf{w}_{t+1}) \leq F(\mathbf{w}_t) + \frac{\eta_1 \omega}{2} \|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2 - \frac{\eta_1}{2} \|\nabla F(\mathbf{w}_t)\|^2 - \frac{\eta_1}{4} \|\mathbf{m}_{t+1}\|^2.$$

Lemma E.5 (Lemma 4.3 Lin et al. (2019)). *Under Assumption E.1, $\lambda_i(\mathbf{w})$ is C_λ -Lipschitz continuous with $C_\lambda = L_L/\mu_L$ for all $i = 1, \dots, n$.*

Lemma E.6. *Consider the updates in Algorithm 5, under Assumption E.1, with $\eta_0 \leq \min\{\mu_L/L_L^2, \frac{2n}{|\mathcal{B}_1^t| \mu_L}\}$ we have*

$$\sum_{t=0}^T \mathbb{E}[\delta_{\lambda,t}] \leq \frac{2n}{|\mathcal{B}_1^t| \eta_0 \mu_L} \delta_{\lambda,0} + \frac{4n\eta_0 T \sigma^2}{\mu_L} + \frac{8n^3 C_\lambda^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \quad (12)$$

Lemma E.7. *Consider Algorithm 5, under Assumption E.1, with $\gamma_0 < 1/2$ we have*

$$\sum_{t=0}^T \mathbb{E}[\delta_{g,t}] \leq \frac{2n}{|\mathcal{B}_1^t| \gamma_0} \delta_{g,0} + 8n\gamma_0 \sigma^2 T + \frac{8n^3 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \quad (13)$$

Lemma E.8. *Consider Algorithm 5, under Assumption E.1, with $\gamma'_0 \leq 1$ we have*

$$\sum_{t=0}^T \mathbb{E}[\delta_{L\lambda\lambda,t}] \leq \frac{4n}{|\mathcal{B}_1^t| \gamma'_0} \delta_{L\lambda\lambda,0} + 32L_{L\lambda\lambda}^2 \sum_{t=0}^{T-1} \mathbb{E}[\delta_{\lambda,t}] + 8n\gamma'_0 T \sigma^2 + \frac{32n^3 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2]$$

Proof of Theorem E.2. First, recall and define the following definitions

$$\begin{aligned} \nabla F(\mathbf{w}_t) &:= \frac{1}{n} \sum_{i \in \mathcal{S}} \left[\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) \right] f_i(g_i(\mathbf{w}_t)) \\ &\quad + \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) \nabla g_i(\mathbf{w}_t) \nabla f_i(g_i(\mathbf{w}_t)) \end{aligned}$$

$$\begin{aligned} \nabla F(\mathbf{w}_t, \lambda^t) &:= \frac{1}{n} \sum_{i \in \mathcal{S}} \nabla F_i(\mathbf{w}_t, \lambda^t) \\ &:= \frac{1}{n} \sum_{i \in \mathcal{S}} \left[\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) [s_i^t]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t) \right] f_i(u_i^t) \\ &\quad + \psi_i(\mathbf{w}_t, \lambda_i^t) \nabla g_i(\mathbf{w}_t) \nabla f_i(u_i^t) \end{aligned}$$

$$G(\mathbf{w}_t) = \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \left[\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) [s_i^t]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t) \right] f_i(u_i^t) + \psi_i(\mathbf{w}_t, \lambda_i^t) \nabla g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) \nabla f_i(u_i^t)$$

Considering the update $\mathbf{m}_{t+1} = (1 - \gamma_1)\mathbf{m}_t + \gamma_1 G(\mathbf{w}_t)$ in Algorithm 5, where $\gamma_1 = 1 - \beta_1$, we have

$$\begin{aligned}
 & \mathbb{E}_t[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] \\
 &= \mathbb{E}_t[\|\nabla F(\mathbf{w}_t) - (1 - \gamma_1)\mathbf{m}_t - \gamma_1 G(\mathbf{w}_t)\|^2] \\
 &= \mathbb{E}_t[\|(1 - \gamma_1)(\nabla F(\mathbf{w}_{t-1}) - \mathbf{m}_t) + (1 - \gamma_1)(\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_{t-1})) + \gamma_1(\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_t, \lambda^t)) \\
 &\quad + \gamma_1(\nabla F(\mathbf{w}_t, \lambda^t) - G(\mathbf{w}_t))\|^2] \\
 &\stackrel{(a)}{=} \mathbb{E}_t[\|(1 - \gamma_1)(\nabla F(\mathbf{w}_{t-1}) - \mathbf{m}_t) + (1 - \gamma_1)(\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_{t-1})) + \gamma_1(\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_t, \lambda^t))\|^2 \\
 &\quad + \|\gamma_1(\nabla F(\mathbf{w}_t, \lambda^t) - G(\mathbf{w}_t))\|^2] \\
 &\stackrel{(b)}{\leq} (1 + \gamma_1)(1 - \gamma_1)^2 \|\nabla F(\mathbf{w}_{t-1}) - \mathbf{m}_t\|^2 + 2(1 + \frac{1}{\gamma_1}) \left[\|\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_{t-1})\|^2 + \gamma_1^2 \|\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_t, \lambda^t)\|^2 \right] \\
 &\quad + \gamma_1^2 \mathbb{E}_t[\|\nabla F(\mathbf{w}_t, \lambda^t) - G(\mathbf{w}_t)\|^2] \\
 &\leq (1 - \gamma_1) \|\nabla F(\mathbf{w}_{t-1}) - \mathbf{m}_t\|^2 + 2(1 + \frac{1}{\gamma_1}) \left[L_F^2 \|\mathbf{w}_t - \mathbf{w}_{t-1}\|^2 + \underbrace{\gamma_1^2 \|\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_t, \lambda^t)\|^2}_{\textcircled{a}} \right] \\
 &\quad + \gamma_1^2 \underbrace{\mathbb{E}_t[\|\nabla F(\mathbf{w}_t, \lambda^t) - G(\mathbf{w}_t)\|^2]}_{\textcircled{b}}
 \end{aligned} \tag{14}$$

where the (a) follows from $\mathbb{E}_t[\widehat{\nabla} F(\mathbf{w}_t, \lambda^t)] = G(\mathbf{w}_t)$, and (b) is due to $\|a+b\|^2 \leq (1+\beta)\|a\|^2 + (1+\frac{1}{\beta})\|b\|^2$. Furthermore, one may bound the last two terms in (14) as following

$$\begin{aligned}
 \textcircled{a} &= \mathbb{E}_t[\|\nabla F(\mathbf{w}_t) - \nabla F(\mathbf{w}_t, \lambda^t)\|^2] \\
 &\leq \frac{1}{n} \sum_{i \in \mathcal{S}} 6 \|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) [f_i(g_i(\mathbf{w}_t)) - f_i(u_i^t)]\|^2 + 6 \|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_{\mathbf{w}} \psi_i(\mathbf{w}_t, \lambda_i^t)\|^2 \|f_i(u_i^t)\|^2 \\
 &\quad + 12 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) f_i(g_i(\mathbf{w}_t))\|^2 \\
 &\quad + 12 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)]^{-1} \|\nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t)\|^2 \|f_i(g_i(\mathbf{w}_t))\|^2 \\
 &\quad + 12 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)]^{-1} \|\nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t)\|^2 \|f_i(g_i(\mathbf{w}_t)) - f_i(u_i^t)\|^2 \\
 &\quad + 12 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)]^{-1} - [s_i^t]^{-1} \|\nabla_{\lambda} \psi_i(\mathbf{w}_t, \lambda_i^t)\|^2 \|f_i(u_i^t)\|^2 \\
 &\quad + 6 \|\psi_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \psi_i(\mathbf{w}_t, \lambda_i^t)\| \|\nabla g_i(\mathbf{w}_t)\| \|\nabla f_i(g_i(\mathbf{w}_t))\|^2 \\
 &\quad + 6 \|\psi_i(\mathbf{w}_t, \lambda_i^t)\| \|\nabla g_i(\mathbf{w}_t)\| \|\nabla f_i(g_i(\mathbf{w}_t)) - \nabla f_i(u_i^t)\|^2 \\
 &\leq \left(\frac{6C_{\psi}^2 C_f^2}{n} + \frac{12C_{L\mathbf{w}\lambda}^2 C_{\psi}^2 C_f^2}{\mu_L^2 n} + \frac{6B_{\psi}^2 C_g^2 L_f^2}{n} \right) \|g(\mathbf{w}_t) - u^t\|^2 + \frac{12C_{L\mathbf{w}\lambda}^2 C_{\psi}^2 B_f^2}{\mu_L^2 \gamma^2 n} \|\nabla_{\lambda\lambda}^2 L(\mathbf{w}_t, \lambda(\mathbf{w}_t)) - s^t\|^2 \\
 &\quad + \left(\frac{6L_{\psi}^2 B_f^2}{n} + \frac{12L_{L\mathbf{w}\lambda}^2 C_{\psi}^2 B_f^2}{\mu_L^2 n} + \frac{12C_{L\mathbf{w}\lambda}^2 L_{\psi}^2 B_f^2}{\mu_L^2 n} + \frac{6C_g^2 C_f^2}{n} \right) \|\lambda(\mathbf{w}_t) - \lambda^t\|^2 \\
 &=: \frac{C_6}{4n} \delta_{g,t} + \frac{C_7}{4n} \delta_{L\lambda\lambda,t} + \frac{\tilde{C}_5}{4n} \delta_{\lambda,t}
 \end{aligned}$$

with some properly chosen constants \tilde{C}_5, C_6, C_7 .

On the other hand, part ⑥ can be bounded by some constant,

$$\begin{aligned}
 \textcircled{6} &= \mathbb{E}_t[\|\nabla F(\mathbf{w}_t, \lambda^t) - G(\mathbf{w}_t)\|^2] \\
 &\leq \mathbb{E}_t \left[2 \left\| \frac{1}{n} \sum_{i \in \mathcal{S}} \nabla F_i(\mathbf{w}_t, \lambda^t) - \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \nabla F_i(\mathbf{w}_t, \lambda^t) \right\|^2 + 2 \left\| \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \nabla F_i(\mathbf{w}_t, \lambda^t) - \frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} G_i(\mathbf{w}_t) \right\|^2 \right] \\
 &\leq \frac{12C_\psi^2 B_f^2 + \frac{12C_{L\mathbf{w}\lambda}^2 C_\psi^2 B_f^2}{\gamma^2} + 12B_\psi^2 C_g^2 C_f^2}{|\mathcal{B}_1^t|} \\
 &\quad + 2\mathbb{E}_t \left[\frac{1}{|\mathcal{B}_1^t|} \sum_{i \in \mathcal{B}_1^t} \left\| [\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t)] [s_i^t]^{-1} \nabla_\lambda \psi_i(\mathbf{w}_t, \lambda_i^t) f_i(u_i^t) \right\|^2 \right. \\
 &\quad \left. + \left\| \psi_i(\mathbf{w}_t, \lambda_i^t) [\nabla g_i(\mathbf{w}_t) - \nabla g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t)] \nabla f_i(u_i^t) \right\|^2 \right] \\
 &\leq \frac{12C_\psi^2 B_f^2 + \frac{12C_{L\mathbf{w}\lambda}^2 C_\psi^2 B_f^2}{\gamma^2} + 12B_\psi^2 C_g^2 C_f^2}{|\mathcal{B}_1^t|} + \frac{C_\psi^2 B_f^2 \sigma^2}{\gamma^2} + B_\psi^2 C_f^2 \sigma^2 =: \frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2
 \end{aligned}$$

Thus, with the natural assumption $\gamma_1 \leq 1$, we have

$$\begin{aligned}
 &\mathbb{E}_t[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] \\
 &\leq (1 - \gamma_1) \|\nabla F(\mathbf{w}_{t-1}) - \mathbf{m}_t\|^2 + \frac{4}{\gamma_1} \left[L_F^2 \eta_1^2 \|\mathbf{m}_{t-1}\|^2 + \gamma_1^2 \frac{\tilde{C}_5}{4n} \delta_{\lambda,t} + \gamma_1^2 \frac{C_6}{4n} \delta_{g,t} + \gamma_1^2 \frac{C_7}{4n} \delta_{L\lambda\lambda,t} \right] + \gamma_1^2 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right)
 \end{aligned} \tag{15}$$

Take expectation over all randomness and summation over $t = 1, \dots, T$ to get

$$\begin{aligned}
 \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] &\leq \frac{1}{\gamma_1} \mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2] + \frac{4L_F^2 \eta_1^2}{\gamma_1^2} \sum_{t=1}^T \mathbb{E}[\|\mathbf{m}_t\|^2] + \frac{\tilde{C}_5}{n} \sum_{t=1}^T \mathbb{E}[\delta_{\lambda,t}] \\
 &\quad + \frac{C_6}{n} \sum_{t=1}^T \mathbb{E}[\delta_{g,t}] + \frac{C_7}{n} \sum_{t=1}^T \mathbb{E}[\delta_{L\lambda\lambda,t}] + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right) T
 \end{aligned} \tag{16}$$

Recall that from Lemma E.6 Lemma E.7, and Lemma E.8 we have bounds for $\sum_{t=0}^T \mathbb{E}[\delta_{\lambda,t}]$, $\sum_{t=0}^T \mathbb{E}[\delta_{g,t}]$ and $\sum_{t=0}^T \mathbb{E}[\delta_{L\lambda\lambda,t}]$,

$$\sum_{t=0}^T \mathbb{E}[\delta_{\lambda,t}] \leq \frac{2n}{|\mathcal{B}_1^t| \eta_0 \mu_L} \delta_{\lambda,0} + \frac{4n\eta_0 T \sigma^2}{\mu_L} + \frac{8n^3 C_\lambda^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \tag{17}$$

$$\sum_{t=0}^T \mathbb{E}[\delta_{g,t}] \leq \frac{2n}{|\mathcal{B}_1^t| \gamma_0} \delta_{g,0} + 8n\gamma_0 \sigma^2 T + \frac{8n^3 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \tag{18}$$

$$\sum_{t=0}^T \mathbb{E}[\delta_{L\lambda\lambda,t}] \leq \frac{4n}{|\mathcal{B}_1^t| \gamma_0'} \delta_{L\lambda\lambda,0} + 32L_{L\lambda\lambda}^2 \sum_{t=0}^{T-1} \mathbb{E}[\delta_{\lambda,t}] + 8n\gamma_0' T \sigma^2 + \frac{32n^3 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \tag{19}$$

By plugging (17), (18) and (19) into inequality (16), we obtain

$$\begin{aligned}
 & \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] \\
 & \leq \frac{1}{\gamma_1} \mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2] + \frac{4L_F^2 \eta_1^2}{\gamma_1^2} \sum_{t=1}^T \mathbb{E}[\|\mathbf{m}_t\|^2] + \frac{C_5}{n} \sum_{t=0}^T \mathbb{E}[\delta_{\lambda,t}] + \frac{C_6}{n} \sum_{t=0}^T \mathbb{E}[\delta_{g,t}] \\
 & \quad + \frac{C_7}{n} \left[\frac{4n}{|\mathcal{B}_1^t| \gamma_0'} \delta_{L\lambda\lambda,0} + 8n\gamma_0' T \sigma^2 + \frac{32n^3 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \right] + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right) T \\
 & \leq \frac{1}{\gamma_1} \mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2] + \frac{2C_5}{|\mathcal{B}_1^t| \eta_0 \mu_L} \delta_{\lambda,0} + \frac{4C_5 \eta_0 T \sigma^2}{\mu_L} + \frac{2C_6}{|\mathcal{B}_1^t| \gamma_0} \delta_{g,0} + 8C_6 \gamma_0 \sigma^2 T \\
 & \quad + \frac{4C_7}{|\mathcal{B}_1^t| \gamma_0'} \delta_{L\lambda\lambda,0} + 8C_7 \gamma_0' T \sigma^2 + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right) T \\
 & \quad + \left[\frac{4L_F^2 \eta_1^2}{\gamma_1^2} + \frac{8n^2 C_5 C_\lambda^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} + \frac{8n^2 C_6 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} + \frac{32n^2 C_7 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} \right] \sum_{t=1}^T \mathbb{E}[\|\mathbf{m}_t\|^2]
 \end{aligned} \tag{20}$$

where $C_5 := 32L_{L\lambda\lambda}^2 C_7 + \tilde{C}_5$.

Recall Lemma E.4, we have

$$F(\mathbf{w}_{t+1}) \leq F(\mathbf{w}_t) + \frac{\eta_w}{2} \|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2 - \frac{\eta_1}{2} \|\nabla F(\mathbf{w}_t)\|^2 - \frac{\eta_1}{4} \|\mathbf{m}_{t+1}\|^2.$$

Combing with (20), we obtain

$$\begin{aligned}
 & \frac{1}{T+1} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{x}_t)\|^2] \\
 & \leq \frac{2\mathbb{E}[F(\mathbf{w}_0) - F(\mathbf{w}_{T+1})]}{\eta_1 T} + \frac{1}{T} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] - \frac{1}{2T} \sum_{t=0}^T \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \\
 & \leq \frac{2[F(\mathbf{w}_0) - F(\mathbf{w}^*)]}{\eta_1 T} + \frac{1}{T} \left[\frac{\mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2]}{\gamma_1} + \frac{2C_5 \delta_{\lambda,0}}{|\mathcal{B}_1^t| \eta_0 \mu_L} + \frac{2C_6 \delta_{g,0}}{|\mathcal{B}_1^t| \gamma_0} + \frac{4C_7 \delta_{L\lambda\lambda,0}}{|\mathcal{B}_1^t| \gamma_0'} \right] \\
 & \quad + \frac{4C_5 \eta_0 \sigma^2}{\mu_L} + 8C_6 \gamma_0 \sigma^2 + 8C_7 \gamma_0' \sigma^2 + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right) \\
 & \quad + \frac{1}{T} \left[\frac{4L_F^2 \eta_1^2}{\gamma_1^2} + \frac{8n^2 C_5 C_\lambda^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} + \frac{8n^2 C_6 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} + \frac{32n^2 C_7 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} - \frac{1}{2} \right] \sum_{t=1}^T \mathbb{E}[\|\mathbf{m}_t\|^2]
 \end{aligned} \tag{21}$$

By setting

$$\eta_1^2 \leq \min \left\{ \frac{\gamma_1^2}{64L_F^2}, \frac{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2}{128n^2 C_5 C_\lambda^2}, \frac{|\mathcal{B}_1^t|^2 \gamma_0^2}{128n^2 C_6 C_g^2}, \frac{|\mathcal{B}_1^t|^2 \gamma_0'^2}{512n^2 C_7 L_{L\lambda\lambda}^2 (1 + C_\lambda^2)} \right\}$$

we have

$$\frac{4L_F^2 \eta_1^2}{\gamma_1^2} + \frac{8n^2 C_5 C_\lambda^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} + \frac{8n^2 C_6 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} + \frac{32n^2 C_7 L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0'^2} - \frac{1}{4} \leq 0$$

which implies that the last term of the RHS of inequality (21) are less or equal to zero. Hence

$$\begin{aligned}
 & \frac{1}{T+1} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{x}_t)\|^2] \\
 & \leq \frac{2[F(\mathbf{w}_0) - F(\mathbf{w}^*)]}{\eta_1 T} + \frac{1}{T} \left[\frac{\mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2]}{\gamma_1} + \frac{2C_5 \delta_{\lambda,0}}{|\mathcal{B}_1^t| \eta_0 \mu_L} + \frac{2C_6 \delta_{g,0}}{|\mathcal{B}_1^t| \gamma_0} + \frac{4C_7 \delta_{L\lambda\lambda,0}}{|\mathcal{B}_1^t| \gamma_0'} \right] \\
 & \quad + \frac{4C_5 \eta_0 \sigma^2}{\mu_L} + 8C_6 \gamma_0 \sigma^2 + 8C_7 \gamma_0' \sigma^2 + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9 \sigma^2 \right)
 \end{aligned} \tag{22}$$

With

$$\eta_0 \leq \frac{\mu_L \epsilon^2}{48C_5\sigma^2}, \gamma_0 \leq \frac{\epsilon^2}{96C_6\sigma^2}, \gamma'_0 \leq \frac{\epsilon^2}{96C_7\sigma^2}, \gamma_1 \leq \frac{\epsilon^2}{12(\frac{C_8}{|\mathcal{B}_1^t|} + C_9\sigma^2)},$$

$$T \geq \left\{ \frac{30[F(\mathbf{w}_0) - F(\mathbf{w}^*)]}{\eta_1\epsilon^2}, \frac{15\mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2]}{\gamma_1\epsilon^2}, \frac{30C_5\delta_{\lambda,0}}{|\mathcal{B}_1^t|\eta_0\mu_L\epsilon^2}, \frac{30C_6\delta_{g,0}}{|\mathcal{B}_1^t|\gamma_0\epsilon^2}, \frac{60C_7\delta_{L\lambda\lambda,0}}{|\mathcal{B}_1^t|\gamma'_0\epsilon^2} \right\}$$

we have

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{x}_t)\|^2] \leq \frac{1}{3}\epsilon^2 + \frac{1}{3}\epsilon^2 < \epsilon^2.$$

Furthermore, to show the second part of the theorem, following from inequality (20), we have

$$\begin{aligned} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] &\leq \frac{1}{\gamma_1} \mathbb{E}[\|\nabla F(\mathbf{w}_0) - \mathbf{m}_1\|^2] + \frac{2C_5}{|\mathcal{B}_1^t|\eta_0\mu_L} \delta_{\lambda,0} + \frac{4C_5\eta_0 T\sigma^2}{\mu_L} + \frac{2C_6}{|\mathcal{B}_1^t|\gamma_0} \delta_{g,0} \\ &\quad + 8C_6\gamma_0\sigma^2 T + \frac{4C_7}{|\mathcal{B}_1^t|\gamma'_0} \delta_{L\lambda\lambda,0} + 8C_7\gamma'_0 T\sigma^2 + \gamma_1 \left(\frac{C_8}{|\mathcal{B}_1^t|} + C_9\sigma^2 \right) T \\ &\quad + \frac{1}{2} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla F(\mathbf{w}_t)\|^2 + \|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2]. \end{aligned}$$

With parameters set above, it follows that

$$\frac{1}{T+1} \sum_{t=0}^T \mathbb{E}[\|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2] < 2\epsilon^2.$$

□

E.4. Proofs of Lemmas

E.4.1. PROOF OF LEMMA 4.1

Proof. Given $\ell(\mathbf{w}; \mathbf{x}', \mathbf{x}, q) \geq \mathbb{I}(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) \geq 0)$, we have

$$\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) \geq r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q),$$

for each (q, \mathbf{x}_i^q) , which immediately follows the desired conclusion. □

E.4.2. PROOF OF LEMMA 5.1

Proof. To show the equivalence in the Lemma, it suffices to show that $\lambda_q(\mathbf{w})$ is the $(K+1)$ -th largest value in the set $\{h_q(\mathbf{x}'; \mathbf{w}) | \mathbf{x}' \in \mathcal{S}_q\}$. Let $\{\theta_1, \theta_2, \dots, \theta_{N_q}\}$ denote a sequence of values defined by sorting $\{h_q(\mathbf{x}'; \mathbf{w}) | \mathbf{x}' \in \mathcal{S}_q\}$ in descending order, i.e., $\theta_1 \geq \theta_2 \geq \dots \geq \theta_{N_q}$. θ_k denote the k -th largest value.

Recall

$$\lambda_q(\mathbf{w}) = \arg \min_{\lambda} (K + \varepsilon)\lambda + \sum_{\mathbf{x}' \in \mathcal{S}_q} (h_q(\mathbf{x}'; \mathbf{w}) - \lambda)_+,$$

where $\varepsilon \in (0, 1)$. Define function $\Lambda_q(\lambda) := (K + \varepsilon)\lambda + \sum_{i=1}^{N_q} (\theta_i - \lambda)_+$, then it follows that $\lambda_q(\mathbf{w}) = \arg \min_{\lambda} \Lambda_q(\lambda)$. Take the derivative of $\Lambda_q(\lambda)$, we have

$$\nabla_{\lambda} \Lambda_q(\lambda) = K + \varepsilon - \sum_{i=1}^{N_q} d(\theta_i - \lambda), \text{ where } d(\theta_i - \lambda) = \begin{cases} 1, & \theta_i > \lambda \\ \varepsilon' \in [0, 1], & \theta_i = \lambda \\ 0, & \theta_i < \lambda \end{cases}$$

First, we assume $\theta_K > \theta_{K+1}$. One may consider this problem in three cases.

- If $\lambda > \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) \leq K$, so we have $\nabla_{\lambda} \Lambda_q(\lambda) \geq K + \varepsilon - K = \varepsilon > 0$.
- If $\lambda < \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) \geq K + 1$, so we have $\nabla_{\lambda} \Lambda_q(\lambda) \leq K + \varepsilon - K - 1 = \varepsilon - 1 < 0$.
- If $\lambda = \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) = K + \varepsilon'$, so we have $\nabla_{\lambda} \Lambda_q(\lambda) = K + \varepsilon - K - \varepsilon' = \varepsilon - \varepsilon'$. Thus we will have $\nabla_{\lambda} \Lambda_q(\lambda) = 0$ by setting $\varepsilon' = \varepsilon$. Hence $\lambda_q(\mathbf{w}) = \theta_{K+1}$.

Second, if $\theta_K = \theta_{K+1}$. One may consider this problem in three cases.

- If $\lambda > \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) \leq K - 1$, so we have $\nabla_\lambda \Lambda_q(\lambda) \geq K + \varepsilon - K + 1 = 1 + \varepsilon > 0$.
- If $\lambda < \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) \geq K + 1$, so we have $\nabla_\lambda \Lambda_q(\lambda) \leq K + \varepsilon - K - 1 < 0$.
- If $\lambda = \theta_K = \theta_{K+1}$, then $\sum_{i=1}^{N_q} d(\theta_i - \lambda) = K - 1 + 2\varepsilon'$, so we have $\nabla_\lambda \Lambda_q(\lambda) = K + \varepsilon - K + 1 - 2\varepsilon' = 1 + \varepsilon - 2\varepsilon'$. Thus we will have $\nabla_\lambda \Lambda_q(\lambda) = 0$ by setting $\varepsilon' = (1 + \varepsilon)/2$. Hence $\lambda_q(\mathbf{w}) = \theta_{K+1}$.

In summary, $\theta_{K+1} = \lambda_q(\mathbf{w}) = \arg \min_\lambda \Lambda_q(\lambda)$. The proof is finished. \square

E.4.3. PROOF OF LEMMA 5.2

Proof. Given the condition $\psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w})) \leq C\mathbb{I}(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}) > 0)$ and $\ell(\mathbf{w}; \mathbf{x}', \mathbf{x}, q) \geq \mathbb{I}(h_q(\mathbf{x}'; \mathbf{w}) - h_q(\mathbf{x}; \mathbf{w}) > 0)$, we have

$$\frac{\psi(h_q(\mathbf{x}_i^q; \mathbf{w}) - \lambda_q(\mathbf{w}))(2^{y_i^q} - 1)}{CZ_q^K \log_2(\bar{g}(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)} \leq \frac{\mathbb{I}(\mathbf{x}_i^q \in \mathcal{S}_q[K])(2^{y_i^q} - 1)}{Z_q^K \log_2(r(\mathbf{w}; \mathbf{x}_i^q, \mathcal{S}_q) + 1)}$$

for each (q, \mathbf{x}_i^q) . The desired result follows. \square

E.4.4. PROOF OF LEMMA 5.3

Proof. Recall

$$L_q(\lambda; \mathbf{w}) = \frac{K}{N_q} \lambda + \frac{\tau_2}{2} \lambda^2 + \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} \tau_1 \ln(1 + \exp((h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)/\tau_1)).$$

Define

$$\begin{aligned} \tilde{L}_q(\lambda; \mathbf{w}) &= \frac{K}{N_q} \lambda + \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} (h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)_+ \\ \hat{L}_q(\lambda; \mathbf{w}) &= \frac{K}{N_q} \lambda + \frac{\tau_2}{2} \lambda^2 + \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} (h_q(\mathbf{x}_i; \mathbf{w}) - \lambda)_+. \end{aligned}$$

For simplicity, we denote $\lambda_* = \arg \min_\lambda L_q(\lambda; \mathbf{w})$, $\tilde{\lambda}_* = \arg \min_\lambda \tilde{L}_q(\lambda; \mathbf{w})$, $\hat{\lambda}_* = \arg \min_\lambda \hat{L}_q(\lambda; \mathbf{w})$. Note that it is obvious to see that when $\lambda \geq 2c_h$, function $\tilde{L}_q(\lambda; \mathbf{w})$ is monotonically increasing, and monotonically decreasing when $\lambda \leq 0$. Thus the optimal point is bounded, i.e. $\tilde{\lambda}_* \in [0, 2c_h]$. Similarly, we have $\nabla_\lambda L_q(\lambda; \mathbf{w}) < 0$ when $\lambda \leq 0$ and $\nabla_\lambda L_q(\lambda; \mathbf{w}) \geq 0$ when $\lambda \geq c_h + \tau_1 \ln N_{max}$ where $N_{max} = \max_q N_q$. This allows us to show that the optimal point λ_* is also bounded, i.e. $\lambda_* \in [0, c_h + \tau_1 \ln N_{max}]$. By applying Lemma 8 in Yang & Lin (2018) to $\tilde{L}_q(\lambda; \mathbf{w})$, we know that there exists a constant $c_1 > 0$ such that for all λ we have

$$|\lambda - \lambda_q(\mathbf{w})| \leq c_1 (\tilde{L}_q(\lambda; \mathbf{w}) - \tilde{L}_q(\lambda_q(\mathbf{w}); \mathbf{w})). \quad (23)$$

It is trivial to show $\tau_1 \ln(1 + \exp(x/\tau_1)) \geq x_+ \forall x \in \mathbb{R}$ and $\tau_1 \ln(1 + \exp(x/\tau_1)) - x_+ \leq (\ln 2)\tau_1$. Then it follows easily that

$$\hat{L}_q(\lambda; \mathbf{w}) \leq L_q(\lambda; \mathbf{w}) \leq \hat{L}_q(\lambda; \mathbf{w}) + c_2 \tau_1 \quad (24)$$

where $c_2 = \ln 2$. Then with inequality (24) and the optimality of λ_* , we have

$$\begin{aligned} \tilde{L}_q(\lambda_*; \mathbf{w}) &= \hat{L}_q(\lambda_*; \mathbf{w}) - \frac{\tau_2}{2} \lambda_*^2 \leq L_q(\lambda_*; \mathbf{w}) - \frac{\tau_2}{2} \lambda_*^2 \leq L_q(\tilde{\lambda}_*; \mathbf{w}) - \frac{\tau_2}{2} \lambda_*^2 \\ &\leq \hat{L}_q(\tilde{\lambda}_*; \mathbf{w}) + c_2 \tau_1 - \frac{\tau_2}{2} \lambda_*^2 = \tilde{L}_q(\tilde{\lambda}_*; \mathbf{w}) + \frac{\tau_2}{2} \tilde{\lambda}_*^2 + c_2 \tau_1 - \frac{\tau_2}{2} \lambda_*^2 \end{aligned}$$

which follows that

$$|\tilde{L}_q(\lambda_*; \mathbf{w}) - \tilde{L}_q(\tilde{\lambda}_*; \mathbf{w})| \leq \frac{\tau_2}{2} \tilde{\lambda}_*^2 + c_2 \tau_1 - \frac{\tau_2}{2} \lambda_*^2 \quad (25)$$

Combining inequalities (23), (25) and the boundedness of λ_* , $\tilde{\lambda}_*$, and setting $\tau_1 = \tau_2 = \varepsilon$, we obtain

$$|\lambda_q(\mathbf{w}) - \hat{\lambda}_q(\mathbf{w})| \leq c_1 \left(\frac{\tau_2}{2} \tilde{\lambda}_*^2 + c_2 \tau_1 - \frac{\tau_2}{2} \lambda_*^2 \right) = \mathcal{O}(\varepsilon)$$

To show the smoothness of $L_q(\lambda; \mathbf{w})$, we first show

$$\tau_1 \ln(1 + \exp(x/\tau_1)) = \max_{\alpha \in [0, 1]} x\alpha - \tau_1 [\alpha \ln(\alpha) + (1 - \alpha) \ln(1 - \alpha)] =: A(\alpha) \quad (26)$$

Note that the solution to $A'(\alpha) = x - \tau_1[\ln(\alpha) - \ln(1 - \alpha)] = 0$ is $\alpha^* = 1 - (1 + \exp(x/\tau_1))^{-1}$. Then $A(\alpha^*) = \tau_1 \ln(1 + \exp(x/\tau_1))$, which implies (26). Moreover, $A(\alpha)$ is strong concave because

$$(A(\alpha) + \tau_1 \alpha^2)'' = -\tau_1 \left(\frac{1}{\alpha} + \frac{1}{1 - \alpha} \right) + 2\tau_1 < 0$$

It follows that

$$L_q(\lambda; \mathbf{w}) = \frac{K}{N_q} \lambda + \frac{\tau_2}{2} \lambda^2 + \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} \max_{\alpha \in (0,1)} (h_q(\mathbf{x}_i; \mathbf{w}) - \lambda) \alpha - \tau_1 [\alpha \ln(\alpha) + (1 - \alpha) \ln(1 - \alpha)].$$

Then by Theorem 1 in Nesterov (2005), $L_q(\lambda; \mathbf{w})$ is smooth. The strong convexity of $L_q(\lambda; \mathbf{w})$ follows from the convexity of $L_q(\lambda; \mathbf{w}) - \frac{\tau_2}{2} \lambda^2$, which can be proved by checking the non-negativity of its second derivative

$$\nabla^2(L_q(\lambda; \mathbf{w}) - \frac{\tau_2}{2} \lambda^2) = \frac{1}{N_q} \sum_{\mathbf{x}_i \in \mathcal{S}_q} \frac{\frac{1}{\tau_1} \exp((\lambda - h_q(\mathbf{x}_i; \mathbf{w}))/\tau_1)}{[1 + \exp((\lambda - h_q(\mathbf{x}_i; \mathbf{w}))/\tau_1)]^2} \geq 0$$

□

E.4.5. PROOF OF LEMMA E.3

Proof. Take arbitrary $\mathbf{w}_1, \mathbf{w}_2 \in \mathbb{R}^d$, we have

$$\begin{aligned} \|\nabla F(\mathbf{w}_1) - \nabla F(\mathbf{w}_2)\| &\leq \frac{1}{n} \sum_{i \in \mathcal{S}} \|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) f_i(g_i(\mathbf{w}_1)) - \nabla_{\mathbf{w}} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2))\| \\ &\quad + \frac{1}{n} \sum_{i \in \mathcal{S}} \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2)) \\ &\quad \quad - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) f_i(g_i(\mathbf{w}_1))\| \\ &\quad + \frac{1}{n} \sum_{i \in \mathcal{S}} \|\psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) \nabla g_i(\mathbf{w}_1) \nabla f_i(g_i(\mathbf{w}_1)) - \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) \nabla g_i(\mathbf{w}_2) \nabla f_i(g_i(\mathbf{w}_2))\| \end{aligned}$$

For each i we have

$$\begin{aligned} &\|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) f_i(g_i(\mathbf{w}_1)) - \nabla_{\mathbf{w}} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2))\|^2 \\ &\leq \|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) [f_i(g_i(\mathbf{w}_1)) - f_i(g_i(\mathbf{w}_2))]\|^2 + \|\nabla_{\mathbf{w}} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) - \nabla_{\mathbf{w}} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))\|^2 \|f_i(g_i(\mathbf{w}_2))\|^2 \\ &\leq C_{\psi}^2 C_f^2 \|g_i(\mathbf{w}_1) - g_i(\mathbf{w}_2)\|^2 + L_{\psi}^2 [\|\mathbf{w}_1 - \mathbf{w}_2\|^2 + \|\lambda_i(\mathbf{w}_1) - \lambda_i(\mathbf{w}_2)\|^2] B_f^2 \\ &\leq (C_{\psi}^2 C_f^2 C_g^2 + B_f^2 L_{\psi}^2 (1 + C_{\lambda})) \|\mathbf{w}_1 - \mathbf{w}_2\|^2 \\ &=: C_1^2 \|\mathbf{w}_1 - \mathbf{w}_2\|^2 \end{aligned}$$

and

$$\begin{aligned} &\|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2)) \\ &\quad - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) f_i(g_i(\mathbf{w}_1))\| \\ &\leq 4 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) - \nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2))\|^2 \\ &\quad + 4 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))\| \left[[\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))]^{-1} - [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))]^{-1} \right] \nabla_{\lambda} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) f_i(g_i(\mathbf{w}_2))\|^2 \\ &\quad + 4 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))]^{-1} \|\nabla_{\lambda} \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) - \nabla_{\lambda} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))\| f_i(g_i(\mathbf{w}_2))\|^2 \\ &\quad + 4 \|\nabla_{\mathbf{w}\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))\| [\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1))]^{-1} \nabla_{\lambda} \psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) \|f_i(g_i(\mathbf{w}_2)) - f_i(g_i(\mathbf{w}_1))\|^2 \\ &\leq \left[\left(\frac{4L_{L\mathbf{w}\lambda}^2 C_{\psi}^2 B_f^2}{\mu_L^2} + \frac{4C_{L\mathbf{w}\lambda}^2 L_{L\lambda}^2 C_{\psi}^2 B_f^2}{\mu_L^4} + \frac{4C_{L\mathbf{w}\lambda}^2 L_{\psi}^2 B_f^2}{\mu_L^2} \right) (1 + C_{\lambda}^2) + \frac{4C_{L\mathbf{w}\lambda}^2 C_{\psi}^2 C_f^2 C_g^2}{\mu_L^2} \right] \|\mathbf{w}_1 - \mathbf{w}_2\|^2 \\ &=: C_2^2 \|\mathbf{w}_1 - \mathbf{w}_2\|^2 \end{aligned}$$

and

$$\begin{aligned}
 & \|\psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) \nabla g_i(\mathbf{w}_1) \nabla f_i(g_i(\mathbf{w}_1)) - \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) \nabla g_i(\mathbf{w}_2) \nabla f_i(g_i(\mathbf{w}_2))\|^2 \\
 & \leq 3\|\psi_i(\mathbf{w}_1, \lambda_i(\mathbf{w}_1)) - \psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))\| \|\nabla g_i(\mathbf{w}_1) \nabla f_i(g_i(\mathbf{w}_1))\|^2 \\
 & \quad + 3\|\psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2))\| \|\nabla g_i(\mathbf{w}_1) - \nabla g_i(\mathbf{w}_2)\| \|\nabla f_i(g_i(\mathbf{w}_1))\|^2 \\
 & \quad + 3\|\psi_i(\mathbf{w}_2, \lambda_i(\mathbf{w}_2)) \nabla g_i(\mathbf{w}_2)\| \|\nabla f_i(g_i(\mathbf{w}_1)) - \nabla f_i(g_i(\mathbf{w}_2))\|^2 \\
 & \leq [3C_\psi^2 C_g^2 C_f^2 (1 + C_\lambda^2) + 3B_\psi^2 L_g^2 C_f^2 + 3B_\ell^2 C_g^2 L_f^2 C_g^2] \|\mathbf{w}_1 - \mathbf{w}_2\|^2 \\
 & =: C_3^2 \|\mathbf{w}_1 - \mathbf{w}_2\|^2.
 \end{aligned}$$

Hence

$$\|\nabla F(\mathbf{w}_1) - \nabla F(\mathbf{w}_2)\| \leq \frac{1}{n} \sum_{i \in S} (C_1 + C_2 + C_3) \|\mathbf{w}_1 - \mathbf{w}_2\| = L_F \|\mathbf{w}_1 - \mathbf{w}_2\|,$$

where $L_F := C_1 + C_2 + C_3$. □

E.4.6. PROOF OF LEMMA E.4

Proof. By L_F -smoothness of $F(\mathbf{w})$, with $\eta_1 \leq \frac{1}{2L_F}$, we have

$$\begin{aligned}
 F(\mathbf{w}_{t+1}) & \leq F(\mathbf{w}_t) + \nabla F(\mathbf{w}_t)^T (\mathbf{w}_{t+1} - \mathbf{w}_t) + \frac{L_F}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2 \\
 & = F(\mathbf{w}_t) - \eta_1 \nabla F(\mathbf{w}_t)^T \mathbf{m}_{t+1} + \frac{L_F}{2} \eta_1^2 \|\mathbf{m}_{t+1}\|^2 \\
 & = F(\mathbf{w}_t) + \frac{\eta_1}{2} \|\nabla F(\mathbf{w}_t) - \mathbf{m}_{t+1}\|^2 - \frac{\eta_1}{2} \|\nabla F(\mathbf{w}_t)\|^2 + \left(\frac{L_F}{2} \eta_1^2 - \frac{\eta_1}{2} \right) \|\mathbf{m}_{t+1}\|^2.
 \end{aligned}$$
□

E.4.7. PROOF OF LEMMA E.6

Proof. Recall and define the following notations

$$\lambda_i^{t+1} = \begin{cases} \lambda_i^t - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t, \\ \lambda_i^t & \text{o.w.} \end{cases}, \quad \tilde{\lambda}_i^t := \lambda_i^t - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t)$$

Then

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{\lambda}_i^t - \lambda_i(\mathbf{w}_t)\|^2] \\
 & = \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\lambda_i^t - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) - \lambda_i(\mathbf{w}_t)\|^2] \\
 & = \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\lambda_i^t - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) - \lambda_i(\mathbf{w}_t) + \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) + \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t) - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t)\|^2] \\
 & = \|\lambda_i^t - \lambda_i(\mathbf{w}_t) + \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t)\|^2 + \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t) - \eta_0 \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t)\|^2] \\
 & \leq \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \eta_0^2 \|\nabla_\lambda L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t)\|^2 \\
 & \quad + 2\eta_0 \langle \lambda_i^t - \lambda_i(\mathbf{w}_t), \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t)) - \nabla_\lambda L_i(\mathbf{w}_t, \lambda_i^t) \rangle + \eta_0^2 \sigma^2 \\
 & \stackrel{(a)}{\leq} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \eta_0^2 L_L^2 \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 - 2\eta_0 \mu_L \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \eta_0^2 \sigma^2 \\
 & \stackrel{(b)}{\leq} (1 - \eta_0 \mu_L) \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \eta_0^2 \sigma^2
 \end{aligned} \tag{27}$$

where inequality (a) uses the strong monotonicity of $L_i(\mathbf{w}_t, \cdot)$ as it is assumed to be μ_L -strongly convex, and (b) uses the assumption $\eta_0 \leq \mu_L / L_L^2$.

Moreover, consider the randomness on the query sampling \mathcal{B}_1^t , we have

$$\mathbb{E}_t [\|\lambda_i^{t+1} - \lambda_i(\mathbf{w}_t)\|^2] = \frac{|\mathcal{B}_1^t|}{n} \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{\lambda}_i^t - \lambda_i(\mathbf{w}_t)\|^2] + \frac{n - |\mathcal{B}_1^t|}{n} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2$$

which follows

$$\mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{\lambda}_i^t - \lambda_i(\mathbf{w}_t)\|^2] = \frac{n}{|\mathcal{B}_1^t|} \mathbb{E}_t [\|\lambda_i^{t+1} - \lambda_i(\mathbf{w}_t)\|^2] - \frac{n - |\mathcal{B}_1^t|}{|\mathcal{B}_1^t|} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 \tag{28}$$

Combining inequalities (27) and (28), we obtain

$$\begin{aligned}\mathbb{E}_t[\|\lambda_i^{t+1} - \lambda_i(\mathbf{w}_t)\|^2] &\leq \frac{n - |\mathcal{B}_1^t|}{n} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{|\mathcal{B}_1^t|}{n} (1 - \eta_0 \mu_L) \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{|\mathcal{B}_1^t|}{n} \eta_0^2 \sigma^2 \\ &\leq (1 - \frac{|\mathcal{B}_1^t| \eta_0 \mu_L}{n}) \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{|\mathcal{B}_1^t| \eta_0^2 \sigma^2}{n}\end{aligned}\quad (29)$$

Thus

$$\begin{aligned}\mathbb{E}_t[\|\lambda_i^{t+1} - \lambda_i(\mathbf{w}_{t+1})\|^2] &\leq (1 + \frac{|\mathcal{B}_1^t| \eta_0 \mu_L}{2n}) \mathbb{E}_t[\|\lambda_i^{t+1} - \lambda_i(\mathbf{w}_t)\|^2] + (1 + \frac{2n}{|\mathcal{B}_1^t| \eta_0 \mu_L}) \mathbb{E}_t[\|\lambda_i(\mathbf{w}_{t+1}) - \lambda_i(\mathbf{w}_t)\|^2] \\ &\leq (1 - \frac{|\mathcal{B}_1^t| \eta_0 \mu_L}{2n}) \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{2|\mathcal{B}_1^t| \eta_0^2 \sigma^2}{n} + \frac{4nC_\lambda^2}{|\mathcal{B}_1^t| \eta_0 \mu_L} \mathbb{E}_t[\|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2]\end{aligned}\quad (30)$$

where we use the assumption $\eta_0 \leq \frac{2n}{|\mathcal{B}_1^t| \mu_L}$ i.e. $\frac{|\mathcal{B}_1^t| \eta_0 \mu_L}{2n} \leq 1$.

Taking summation over all queries and expectation over all randomness, we have

$$\mathbb{E}[\|\lambda^{t+1} - \lambda(\mathbf{w}_{t+1})\|^2] \leq (1 - \frac{|\mathcal{B}_1^t| \eta_0 \mu_L}{2n}) \mathbb{E}[\|\lambda^t - \lambda(\mathbf{w}_t)\|^2] + 2|\mathcal{B}_1^t| \eta_0^2 \sigma^2 + \frac{4n^2 C_\lambda^2}{|\mathcal{B}_1^t| \eta_0 \mu_L} \mathbb{E}[\|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2]\quad (31)$$

Taking summation over $t = 0, \dots, T-1$, we have

$$\sum_{t=0}^T \mathbb{E}[\|\lambda^t - \lambda(\mathbf{w}_t)\|^2] \leq \frac{2n}{|\mathcal{B}_1^t| \eta_0 \mu_L} \|\lambda^0 - \lambda(\mathbf{w}_0)\|^2 + \frac{4n\eta_0 T \sigma^2}{\mu_L} + \frac{8n^3 C_\lambda^2}{|\mathcal{B}_1^t|^2 \eta_0^2 \mu_L^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2]\quad (32)$$

□

E.4.8. PROOF OF LEMMA E.7

Proof. Consider

$$\begin{aligned}\|u^{t+1} - g(\mathbf{w}_t)\|^2 &= \|u^{t+1} - u^t + u^t - g(\mathbf{w}_t)\|^2 \\ &= \|u^{t+1} - u^t\|^2 + \|u^t - g(\mathbf{w}_t)\|^2 + 2\langle u^{t+1} - u^t, u^t - g(\mathbf{w}_t) \rangle \\ &= \|u^{t+1} - u^t\|^2 + \|u^t - g(\mathbf{w}_t)\|^2 + 2 \sum_{i \in \mathcal{B}_1^t} \langle u_i^{t+1} - u_i^t, u_i^t - g_i(\mathbf{w}_t) \rangle \\ &= \|u^{t+1} - u^t\|^2 + \|u^t - g(\mathbf{w}_t)\|^2 + 2 \underbrace{\sum_{i \in \mathcal{B}_1^t} \langle u_i^{t+1} - u_i^t, u_i^t - g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) \rangle}_{A_5} \\ &\quad + 2 \underbrace{\sum_{i \in \mathcal{B}_1^t} \langle u_i^{t+1} - u_i^t, g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) - g_i(\mathbf{w}_t) \rangle}_{A_6}\end{aligned}\quad (33)$$

With $u_i^t - u_i^{t+1} = \gamma_0(u_i^t - g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t)) \forall i \in \mathcal{B}_1^t$ and the inequality $2(b-a, a-c) = \|b-c\|^2 - \|a-b\|^2 - \|a-c\|^2$, we have

$$\begin{aligned}A_5 &= 2 \sum_{i \in \mathcal{B}_1^t} \langle u_i^{t+1} - g_i(\mathbf{w}_t), u_i^t - g_i(\mathbf{w}_t, \mathcal{B}_{2,i}^t) \rangle + 2 \sum_{i \in \mathcal{B}_1^t} \langle g_i(\mathbf{w}_t) - u_i^t, u_i^t - g_i(\mathbf{w}_t, \mathcal{B}_{2,i}^t) \rangle \\ &= \frac{2}{\gamma_0} \sum_{i \in \mathcal{B}_1^t} \langle u_i^{t+1} - g_i(\mathbf{w}_t), u_i^t - u_i^{t+1} \rangle + 2 \sum_{i \in \mathcal{B}_1^t} \langle g_i(\mathbf{w}_t) - u_i^t, u_i^t - g_i(\mathbf{w}_t, \mathcal{B}_{2,i}^t) \rangle \\ &= \frac{1}{\gamma_0} \sum_{i \in \mathcal{B}_1^t} [\|u_i^t - g_i(\mathbf{w}_t)\|^2 - \|u_i^{t+1} - g_i(\mathbf{w}_t)\|^2 - \|u_i^{t+1} - u_i^t\|^2] \\ &\quad + 2 \sum_{i \in \mathcal{B}_1^t} \langle g_i(\mathbf{w}_t) - u_i^t, u_i^t - g_i(\mathbf{w}_t, \mathcal{B}_{2,i}^t) \rangle \\ &= \frac{1}{\gamma_0} \|u^t - g(\mathbf{w}_t)\|^2 - \frac{1}{\gamma_0} \|u^{t+1} - g(\mathbf{w}_t)\|^2 - \frac{1}{\gamma_0} \|u^{t+1} - u^t\|^2 \\ &\quad + 2 \sum_{i \in \mathcal{B}_1^t} \langle g_i(\mathbf{w}_t) - u_i^t, u_i^t - g_i(\mathbf{w}_t, \mathcal{B}_{2,i}^t) \rangle\end{aligned}\quad (34)$$

where the last equality is due to the fact $\|u_i^t - g_i(\mathbf{w}_t)\|^2 = \|u_i^{t+1} - g_i(\mathbf{w}_t)\|^2$ and $\|u_i^{t+1} - u_i^t\|^2 = 0$ for all $i \notin \mathcal{B}_1^t$. Taking expectation over the randomness at iteration t we have

$$\begin{aligned} \mathbb{E}_t[A_5] &\leq \frac{1}{\gamma_0} \|u^t - g(\mathbf{w}_t)\|^2 - \frac{1}{\gamma_0} \mathbb{E}_t[\|u^{t+1} - g(\mathbf{w}_t)\|^2] - \frac{1}{\gamma_0} \mathbb{E}_t[\|u^{t+1} - u^t\|^2] \\ &\quad - 2\mathbb{E}_{\mathcal{B}_1^t} \left[\sum_{i \in \mathcal{B}_1^t} \|u_i^t - g_i(\mathbf{w}_t)\|^2 \right] \\ &= \frac{1}{\gamma_0} \|u^t - g(\mathbf{w}_t)\|^2 - \frac{1}{\gamma_0} \mathbb{E}_t[\|u^{t+1} - g(\mathbf{w}_t)\|^2] - \frac{1}{\gamma_0} \mathbb{E}_t[\|u^{t+1} - u^t\|^2] \\ &\quad - 2\frac{|\mathcal{B}_1^t|}{n} \|u^t - g(\mathbf{w}_t)\|^2 \end{aligned} \quad (35)$$

On the other hand, with the assumption $\gamma_0 < 1/2$, we have

$$\begin{aligned} A_6 &\leq \left(\frac{1}{\gamma_0} - 1\right) \sum_{i \in \mathcal{B}_1^t} \|u_i^{t+1} - u_i^t\|^2 + \frac{1}{\frac{1}{\gamma_0} - 1} \sum_{i \in \mathcal{B}_1^t} \|g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) - g_i(\mathbf{w}_t)\|^2 \\ &\leq \left(\frac{1}{\gamma_0} - 1\right) \sum_{i \in \mathcal{B}_1^t} \|u_i^{t+1} - u_i^t\|^2 + 2\gamma_0 \sum_{i \in \mathcal{B}_1^t} \|g_i(\mathbf{w}_t; \mathcal{B}_{2,i}^t) - g_i(\mathbf{w}_t)\|^2 \\ &\leq \left(\frac{1}{\gamma_0} - 1\right) \|u^{t+1} - u^t\|^2 + 2\gamma_0 |\mathcal{B}_1^t| \sigma^2 \end{aligned} \quad (36)$$

Then by plugging (34), (35), (36) back into (33), we obtain

$$\begin{aligned} &\mathbb{E}[\|u^{t+1} - g(\mathbf{w}_t)\|^2] \\ &\leq \mathbb{E}[\|u^{t+1} - u^t\|^2] + \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] + \frac{1}{\gamma_0} \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] - \frac{1}{\gamma_0} \mathbb{E}[\|u^{t+1} - g(\mathbf{w}_t)\|^2] \\ &\quad - \frac{1}{\gamma_0} \mathbb{E}[\|u^{t+1} - u^t\|^2] - 2\frac{|\mathcal{B}_1^t|}{n} \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] + \left(\frac{1}{\gamma_0} - 1\right) \mathbb{E}[\|u^{t+1} - u^t\|^2] + 2\gamma_0 |\mathcal{B}_1^t| \sigma^2 \\ &= \left(1 + \frac{1}{\gamma_0} - 2\frac{|\mathcal{B}_1^t|}{n}\right) \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] - \frac{1}{\gamma_0} \mathbb{E}[\|u^{t+1} - g(\mathbf{w}_t)\|^2] + 2\gamma_0 |\mathcal{B}_1^t| \sigma^2 \end{aligned}$$

Note that $\frac{(1 + \frac{1}{\gamma_0} - 2\frac{|\mathcal{B}_1^t|}{n})}{1 + \frac{1}{\gamma_0}} = 1 - \frac{2|\mathcal{B}_1^t|\gamma_0}{(1 + \gamma_0)n} \leq 1 - \frac{|\mathcal{B}_1^t|\gamma_0}{n}$ and $(1 + \frac{a}{2})(1 - a) \leq 1 - \frac{a}{2}$. It follows

$$\mathbb{E}[\|u^{t+1} - g(\mathbf{w}_t)\|^2] \leq \left(1 - \frac{|\mathcal{B}_1^t|\gamma_0}{n}\right) \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] + 2\gamma_0^2 |\mathcal{B}_1^t| \sigma^2$$

Moreover, we have

$$\begin{aligned} &\mathbb{E}[\|u^{t+1} - g(\mathbf{w}_{t+1})\|^2] \\ &\leq \left(1 + \frac{|\mathcal{B}_1^t|\gamma_0}{2n}\right) \mathbb{E}[\|u^{t+1} - g(\mathbf{w}_t)\|^2] + \left(1 + \frac{2n}{|\mathcal{B}_1^t|\gamma_0}\right) \mathbb{E}[\|g(\mathbf{w}_t) - g(\mathbf{w}_{t+1})\|^2] \\ &\leq \left(1 + \frac{|\mathcal{B}_1^t|\gamma_0}{2n}\right) \left[\left(1 - \frac{|\mathcal{B}_1^t|\gamma_0}{n}\right) \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] + 2\gamma_0^2 |\mathcal{B}_1^t| \sigma^2 \right] + \left(1 + \frac{2n}{|\mathcal{B}_1^t|\gamma_0}\right) C_g^2 n \mathbb{E}[\|\mathbf{w}_t - \mathbf{w}_{t+1}\|^2] \\ &\leq \left(1 - \frac{|\mathcal{B}_1^t|\gamma_0}{2n}\right) \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] + 4\gamma_0^2 |\mathcal{B}_1^t| \sigma^2 + \frac{4n^2 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|\gamma_0} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2] \end{aligned}$$

Take summation over $t = 0, \dots, T-1$ to get

$$\sum_{t=0}^{T-1} \mathbb{E}[\|u^t - g(\mathbf{w}_t)\|^2] \leq \frac{2n}{|\mathcal{B}_1^t|\gamma_0} \mathbb{E}[\|u^0 - g(\mathbf{w}_0)\|^2] + 8n\gamma_0\sigma^2 T + \frac{8n^3 C_g^2 \eta_1^2}{|\mathcal{B}_1^t|^2 \gamma_0^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{m}_{t+1}\|^2]$$

□

E.4.9. PROOF OF LEMMA E.8

Proof. Recall and define the following notations

$$s_i^{t+1} = \begin{cases} (1 - \gamma'_0)s_i^t + \gamma'_0 \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) & \text{if } i \in \mathcal{B}_1^t \\ s_i^t & \text{o.w.} \end{cases}, \quad \tilde{s}_i^t = (1 - \gamma'_0)s_i^t + \gamma'_0 \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t)$$

Consider

$$\begin{aligned} & \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{s}_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] \\ &= \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|(1 - \gamma'_0)s_i^t + \gamma'_0 \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] \\ &= \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|(1 - \gamma'_0)[s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))] + \gamma'_0[\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)] \\ &\quad + \gamma'_0[\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] \\ &= \|(1 - \gamma'_0)[s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))] + \gamma'_0[\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 \\ &\quad + \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\gamma'_0[\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t; \mathcal{B}_{2,i}^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t)]\|^2] \\ &\leq (1 + \frac{\gamma'_0}{2})(1 - \gamma'_0)^2 \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + (1 + \frac{2}{\gamma'_0})\gamma_0'^2 \|\nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i^t) - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + \gamma_0'^2 \sigma^2 \\ &\leq (1 - \frac{\gamma'_0}{2}) \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + 4\gamma_0' L_{L\lambda\lambda}^2 \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \gamma_0'^2 \sigma^2 \end{aligned} \tag{37}$$

Note that for the randomness of query sampling we have

$$\mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] = \frac{|\mathcal{B}_1^t|}{n} \mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{s}_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] + \frac{n - |\mathcal{B}_1^t|}{n} \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2$$

which follows that

$$\mathbb{E}_{\mathcal{B}_{2,i}^t} [\|\tilde{s}_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] = \frac{n}{|\mathcal{B}_1^t|} \mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] - \frac{n - |\mathcal{B}_1^t|}{|\mathcal{B}_1^t|} \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 \tag{38}$$

Then by plugging the equality (38) into inequality (37), we obtain

$$\begin{aligned} & \frac{n}{|\mathcal{B}_1^t|} \mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] - \frac{n - |\mathcal{B}_1^t|}{|\mathcal{B}_1^t|} \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 \\ &\leq (1 - \frac{\gamma'_0}{2}) \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + 4\gamma_0' L_{L\lambda\lambda}^2 \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \gamma_0'^2 \sigma^2 \end{aligned}$$

It follows

$$\begin{aligned} & \mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] \\ &\leq (1 - \frac{|\mathcal{B}_1^t| \gamma'_0}{2n}) \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + \frac{4|\mathcal{B}_1^t| \gamma_0' L_{L\lambda\lambda}^2}{n} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{|\mathcal{B}_1^t| \gamma_0'^2 \sigma^2}{n} \end{aligned}$$

Furthermore,

$$\begin{aligned} & \mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_{t+1}, \lambda_i(\mathbf{w}_{t+1}))\|^2] \\ &\leq (1 + \frac{|\mathcal{B}_1^t| \gamma'_0}{4n}) \mathbb{E}_t [\|s_i^{t+1} - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2] + (1 + \frac{4n}{|\mathcal{B}_1^t| \gamma'_0}) L_{L\lambda\lambda}^2 (1 + C_\lambda^2) \mathbb{E}_t [\|\mathbf{w}_t - \mathbf{w}_{t+1}\|^2] \\ &\leq (1 - \frac{|\mathcal{B}_1^t| \gamma'_0}{4n}) \|s_i^t - \nabla_{\lambda\lambda}^2 L_i(\mathbf{w}_t, \lambda_i(\mathbf{w}_t))\|^2 + \frac{8|\mathcal{B}_1^t| \gamma_0' L_{L\lambda\lambda}^2}{n} \|\lambda_i^t - \lambda_i(\mathbf{w}_t)\|^2 + \frac{2|\mathcal{B}_1^t| \gamma_0'^2 \sigma^2}{n} \\ &\quad + \frac{8n L_{L\lambda\lambda}^2 (1 + C_\lambda^2)}{|\mathcal{B}_1^t| \gamma'_0} \mathbb{E}_t [\|\mathbf{w}_t - \mathbf{w}_{t+1}\|^2] \end{aligned}$$

where we use the assumption $\gamma'_0 \leq 1 \leq \frac{4n}{|\mathcal{B}_1^t|}$ i.e. $\frac{4n}{|\mathcal{B}_1^t| \gamma'_0} \geq 1$.

Taking expectation over all randomness and taking summation over all queries, we have

$$\begin{aligned} & \mathbb{E} [\|s^{t+1} - \nabla_{\lambda\lambda}^2 L(\mathbf{w}_{t+1}, \lambda(\mathbf{w}_{t+1}))\|^2] \\ &\leq (1 - \frac{|\mathcal{B}_1^t| \gamma'_0}{4n}) \mathbb{E} [\|s^t - \nabla_{\lambda\lambda}^2 L(\mathbf{w}_t, \lambda(\mathbf{w}_t))\|^2] + \frac{8|\mathcal{B}_1^t| \gamma_0' L_{L\lambda\lambda}^2}{n} \mathbb{E} [\|\lambda^t - \lambda(\mathbf{w}_t)\|^2] + 2|\mathcal{B}_1^t| \gamma_0'^2 \sigma^2 \\ &\quad + \frac{8n^2 L_{L\lambda\lambda}^2 (1 + C_\lambda^2)}{|\mathcal{B}_1^t| \gamma'_0} \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}_{t+1}\|^2] \end{aligned}$$

Taking summation over $t = 0, \dots, T - 1$, we obtain

$$\begin{aligned}
 & \sum_{t=0}^T \mathbb{E}[\|s^t - \nabla_{\lambda\lambda}^2 L(\mathbf{w}_t, \lambda(\mathbf{w}_t))\|^2] \\
 & \leq \frac{4n}{|\mathcal{B}_1^t| \gamma_0'} \|s^0 - \nabla_{\lambda\lambda}^2 L(\mathbf{w}_0, \lambda(\mathbf{w}_0))\|^2 + 32L_{L\lambda\lambda}^2 \sum_{t=0}^{T-1} \mathbb{E}[\|\lambda^t - \lambda(\mathbf{w}_t)\|^2] + 8n\gamma_0' T \sigma^2 \\
 & \quad + \frac{32n^3 L_{L\lambda\lambda}^2 (1 + C_\lambda^2)}{|\mathcal{B}_1^t|^2 \gamma_0'^2} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{w}_t - \mathbf{w}_{t+1}\|^2]
 \end{aligned}$$

□