Surrogate Objectives for Batch Policy Optimization in One-step Decision Making

Minmin Chen* **Ramki Gummadi*** **Chris Harris*** **Dale Schuurmans*†**
*Google †University of Alberta

Abstract

We investigate batch policy optimization for cost-sensitive classification and contextual bandits—two related tasks that obviate exploration but require generalizing from observed rewards to action selections in unseen contexts. When rewards are fully observed, we show that the expected reward objective exhibits suboptimal plateaus and exponentially many local optima in the worst case. To overcome the poor landscape, we develop a convex surrogate that is calibrated with respect to entropy regularized expected reward. We then consider the partially observed case, where rewards are recorded for only a subset of actions. Here we generalize the surrogate to partially observed data, and uncover novel objectives for batch contextual bandit training. We find that surrogate objectives remain provably sound in this setting and empirically demonstrate state-of-the-art performance.

1 Introduction

Cost-sensitive classification [1] and batch contextual bandits [34–36] are two problems that share the goal of inferring, given a batch of training data, a *policy* that chooses high reward actions in potentially unseen contexts. The problems differ in the assumed completeness of the data: in cost-sensitive classification, rewards are given (or inferable [8]) for every action, whereas in batch contextual bandits, rewards are only observed for a small subset of actions (typically one). The batch contextual bandit problem is more prevalent in practice, since massive data logs routinely record contexts encountered, actions taken in response, and the outcomes that resulted [18]. Rarely, if ever, are counterfactual outcomes recorded for actions that might have been taken instead [3]. Nevertheless, we find it helpful to reconsider cost-sensitive classification, since a core learning challenge is orthogonal to reward incompleteness: both tasks create difficult optimization landscapes.

There is an extensive literature on cost-sensitive classification. Problems with two actions have been particularly well studied [5, 8], and subsequent work has sought to reduce multiple-action learning to learning binary decisions [1, 19]. A reduction strategy has also been used to convert simply trained stochastic policies to cost-sensitive variants via post-processing [24]. Unfortunately, such reductions do not compose well with current policy learning methods, which are gradient based and best formulated as optimizing a single policy model over a well formed optimization objective.

In this paper we investigate cost-sensitive classification with stochastic policy representations, to ensure the developments are compatible with current deep learning methods. Our first result is negative: for natural policy representations, the expected reward objective generates a poor optimization landscape that exhibits plateaus and potentially an exponential number of local maxima. In response, we develop surrogate objectives for training [28]. Supervised learning research has observed that solution quality can be ensured by using surrogates that satisfy "calibration" with respect to a difficult to optimize target loss [32, 37, 42]. This idea has also recently been applied to cost-sensitive classification [26]. We extend this approach to stochastic policies and deep models by considering expected reward augmented with entropy regularization. This allows a convex surrogate to be developed that improves trainability while approximating expected cost to controllable accuracy.

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

We then consider batch contextual bandits, where rewards are observed only for a subset (typically one) of the available actions in each training context. Current work has focused on direct maximization of expected reward, using importance correction to provide unbiased (or nearly unbiased) estimates of target gradients [14, 16, 41, 43]. Unfortunately, as we illustrate, such an objective creates an extremely difficult optimization landscape, even if variance can be reduced to *zero* [6, 9]. Alternatively, we extend the calibrated surrogate to partially observed rewards, through the introduction of imputed estimates. We prove soundness of the approach and demonstrate empirical performance benefits.

2 Cost-sensitive Classification

We first consider cost-sensitive classification. For simplicity assume a finite set of actions $A = \{1, ..., K\}$, and thus we are given training data $\mathcal{D} = \{(x_i, \boldsymbol{r}_i)\}_{i=1}^T$, where $\boldsymbol{r}_i \in \mathbb{R}^K$ is a vector that specifies the reward for each action in context x_i. The goal is to infer a mapping $h : X \to A$ that specifies a high reward action $a \in A$ for a given context $x \in X$. **Notation:** We let Δ^K denote the K dimensional simplex, $\mathbf{1}$ the vector of all 1s, and $\mathbf{1}_a$ the vector of 0s except for 1 in coordinate a.

Much of the literature on cost-sensitive classification has focused on deterministic classifiers h, but we consider stochastic policies $\boldsymbol{\pi} : X \to \Delta^K$. Any deterministic classifier h can be equivalently expressed by $\boldsymbol{\pi}(x) = \mathbf{1}_{h(x)}$. We seek a policy that maximizes expected reward, or equivalently minimizes expected cost. If we assume the data source is i.i.d. with a joint distribution $p(x, \boldsymbol{r})$, the *true risk* of a policy $\boldsymbol{\pi}$ and its *empirical risk* on data set $\mathcal{D}$ can be defined respectively by

$$\mathcal{R}(\boldsymbol{\pi}) \;=\; -\mathbb{E}[\boldsymbol{\pi}(x) \cdot \boldsymbol{r}] \quad \text{and} \quad \hat{\mathcal{R}}(\boldsymbol{\pi}, \mathcal{D}) \;=\; -\tfrac{1}{T} \textstyle\sum_{(x_i, \boldsymbol{r}_i) \in \mathcal{D}} \boldsymbol{\pi}(x_i) \cdot \boldsymbol{r}_i. \tag{1}$$

Since expected cost is the target, one might presume that directly minimizing empirical risk would be a reasonable approach; unfortunately, this proves problematic [13]. In practice, it is nearly universal to train an unconstrained model $\boldsymbol{q} : X \to \mathbb{R}^K$ that is converted to a policy via a "softmax" transfer; that is, policies are normally represented with the composition $\boldsymbol{\pi}(x) = \boldsymbol{f}(\boldsymbol{q}(x))$ where the model output $\boldsymbol{q}(x)$ is converted to a probability vector via

$$\boldsymbol{f}(\boldsymbol{q}) = e^{\boldsymbol{q} - F(\boldsymbol{q})} \quad \text{with} \quad F(\boldsymbol{q}) = \log(\mathbf{1} \cdot e^{\boldsymbol{q}}). \tag{2}$$

The true and empirical risk can then be re-expressed in terms of $\boldsymbol{q}$ by

$$\mathcal{R}(\boldsymbol{f} \circ \boldsymbol{q}) \;=\; -\mathbb{E}[\boldsymbol{f}(\boldsymbol{q}(x)) \cdot \boldsymbol{r}] \quad \text{and} \quad \hat{\mathcal{R}}(\boldsymbol{f} \circ \boldsymbol{q}, \mathcal{D}) \;=\; -\tfrac{1}{T} \textstyle\sum_{(x_i, \boldsymbol{r}_i) \in \mathcal{D}} \boldsymbol{f}(\boldsymbol{q}(x_i)) \cdot \boldsymbol{r}_i. \tag{3}$$

Unfortunately, the dot product $\boldsymbol{r} \cdot \boldsymbol{f}(\boldsymbol{q}(x))$ creates significant difficulty, as this interacts poorly with the softmax transfer $\boldsymbol{f}$. A well known consequence is that the expected cost plateaus whenever the corresponding policy probabilities are nearly deterministic. A potentially greater challenge, however, is that the softmax transfer can also induce exponentially many local optima.

Theorem 1 *Even for a single context x, a deterministic reward vector $\boldsymbol{r}$, and a linear model $\boldsymbol{q}(x) = W\boldsymbol{\phi}(x)$, the function $\boldsymbol{r} \cdot \boldsymbol{f}(\boldsymbol{q}(x))$ can have a number of local maxima in W that is exponential in the number of actions K and the number of features in $\boldsymbol{\phi}$.* ***(All proofs given in the appendix.***[1]***)***

It is therefore unsurprising that empirical risk minimization with stochastic policies is not considered viable in the cost-sensitive classification literature. Nevertheless, it remains the dominant approach for batch contextual bandits. We seek to bridge the apparent disconnect between these two settings.

2.1 Calibrated Strongly Convex Surrogate

A key idea in cost-sensitive classification has been the development of convex surrogate objectives that exhibit "calibration" with respect to the target risk [2, 32]. We require additional definitions. Let $\mathcal{Q}$ denote the set of measurable functions $X \to \mathbb{R}^K$, and define the minimum risk by $\mathcal{R}^* = \inf_{\boldsymbol{q} \in \mathcal{Q}} \mathcal{R}(\boldsymbol{f} \circ \boldsymbol{q})$. Note that the minimum is generally achieved at a deterministic policy, which cannot be represented by $\boldsymbol{q} \in \mathcal{Q}$; however, the infimum can be arbitrarily well approximated within $\mathcal{Q}$. It will be convenient to expand the risk definition through a notion of pointwise risk: define the *local risk* as $\mathcal{R}(\boldsymbol{\pi}, \boldsymbol{r}, x) = -\boldsymbol{r} \cdot \boldsymbol{\pi}(x)$, which is related to the true risk via $\mathcal{R}(\boldsymbol{\pi}) = \mathbb{E}[\mathcal{R}(\boldsymbol{\pi}, \boldsymbol{r}, x)]$, with the expectation taken over pairs $(x, \boldsymbol{r}) \sim p(x, \boldsymbol{r})$. For each $(x, \boldsymbol{r})$ define the minimal risk by

$$\mathcal{R}^*(\boldsymbol{r}, x) \;=\; \inf_{\boldsymbol{\pi} \in \mathcal{P}} \mathcal{R}(\boldsymbol{\pi}, \boldsymbol{r}, x) \;=\; \inf_{\boldsymbol{q} \in \mathcal{Q}} \mathcal{R}(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x). \tag{4}$$

[1] Appendix and code available at `https://www.cs.ualberta.ca/~dale/neurips19/supplement`

Consider a surrogate loss function $L : (\mathcal{Q}, \mathbb{R}^K, X) \to \mathbb{R}$ and let $L^*(\boldsymbol{r}, x) = \inf_{\boldsymbol{q}\in\mathcal{Q}} L(\boldsymbol{q}, \boldsymbol{r}, x)$. We say that a surrogate L is *calibrated* with respect to the target risk $\mathcal{R}$ if there exists a calibration function $\delta(\epsilon, x) \geq 0$ such that for all $\epsilon > 0$, all $x \in X$, all $\boldsymbol{r} \in \mathbb{R}^K$ and all $\boldsymbol{q} \in \mathcal{Q}$:

$$L(\boldsymbol{q}, \boldsymbol{r}, x) - L^*(\boldsymbol{r}, x) \;<\; \delta(\epsilon, x) \quad \text{implies} \quad \mathcal{R}(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) \;<\; \mathcal{R}^*(\boldsymbol{r}, x) + \epsilon. \tag{5}$$

Although calibrated convex surrogates have been developed for cost-sensitive classification [26], these do not consider stochastic policies. Rather than extending these constructions to stochastic policies, which is not straightforward, we develop a new surrogate for the stochastic case. Consider an entropy regularized version of the target risk [25] which we call the *smoothed risk*:

$$\mathcal{S}_\tau(\boldsymbol{\pi}, \boldsymbol{r}, x) \;=\; -\boldsymbol{r} \cdot \boldsymbol{\pi}(x) + \tau \boldsymbol{\pi}(x) \cdot \log \boldsymbol{\pi}(x) \quad \text{and} \quad \mathcal{S}_\tau(\boldsymbol{\pi}) \;=\; \mathbb{E}[\mathcal{S}_\tau(\boldsymbol{\pi}, \boldsymbol{r}, x)]. \tag{6}$$

The smoothed risk approximates the true risk, with a discrepancy that can be made arbitrarily small.

Proposition 2 *Let* $\tilde{\boldsymbol{\pi}}_\tau = \arg\min_{\boldsymbol{\pi}\in\mathcal{P}} \mathcal{S}_\tau(\boldsymbol{\pi})$. *Then* $\tilde{\boldsymbol{\pi}}_\tau(x) = \exp(\mathbb{E}[\boldsymbol{r}|x] - F(\mathbb{E}[\boldsymbol{r}|x])/\tau)$ *and* $\mathcal{R}(\tilde{\boldsymbol{\pi}}_\tau) < \mathcal{R}^* + \tau \log K$. *Hence for any* $\epsilon > 0$ *setting* $\tau < \epsilon / \log K$ *ensures* $\mathcal{R}(\tilde{\boldsymbol{\pi}}_\tau) < \mathcal{R}^* + \epsilon$.

Note that the smoothed risk is not convex in $\boldsymbol{q}$ due to the softmax transfer $\boldsymbol{\pi}(x) = \boldsymbol{f}(\boldsymbol{q}(x))$. Nevertheless, it is possible to develop a convex surrogate that is calibrated for the smoothed risk as follows. First we need a few properties of Bregman divergences in general and the KL divergence in particular. The Bregman divergence D_F, specified by the convex differentiable potential F, satisfies:

$$D_F(\boldsymbol{q} \| \boldsymbol{r}) \;=\; F(\boldsymbol{q}) - F(\boldsymbol{r}) - \boldsymbol{f}(\boldsymbol{r}) \cdot (\boldsymbol{q} - \boldsymbol{r}) \;=\; F(\boldsymbol{q}) - \boldsymbol{q} \cdot \boldsymbol{p} + F^*(\boldsymbol{p}) \;=\; D_{F^*}(\boldsymbol{p} \| \boldsymbol{\pi}), \tag{7}$$

where $\boldsymbol{f} = \nabla F$, $F^*(\boldsymbol{p})$ is the convex conjugate of F, $\boldsymbol{p} = \boldsymbol{f}(\boldsymbol{r})$ and $\boldsymbol{\pi} = \boldsymbol{f}(\boldsymbol{q})$ [27]. Clearly, D_F is convex in its first argument $\boldsymbol{q}$, but not necessarily in the second. For the KL divergence in particular we have $F(\boldsymbol{q}) = \log \boldsymbol{1} \cdot e^{\boldsymbol{q}}$, $\boldsymbol{f}(\boldsymbol{q}) = e^{\boldsymbol{q} - F(\boldsymbol{q})}$, $F^*(\boldsymbol{p}) = \boldsymbol{p} \cdot \log \boldsymbol{p}$, hence

$$D_{KL}(\boldsymbol{\pi} \| \boldsymbol{p}) = \boldsymbol{\pi} \cdot (\log \boldsymbol{\pi} - \log \boldsymbol{p}) = D_{F^*}(\boldsymbol{\pi} \| \boldsymbol{p}) = D_F(\boldsymbol{r} \| \boldsymbol{q}). \tag{8}$$

This means that the local smoothed risk (6) can be shown to be equivalent to

$$\mathcal{S}_\tau(\boldsymbol{\pi}, \boldsymbol{r}, x) \;=\; -\tau \left(\tfrac{\boldsymbol{r}}{\tau} \cdot \boldsymbol{\pi}(x) - \boldsymbol{\pi}(x) \cdot \log \boldsymbol{\pi}(x) \right) = -\tau F(\tfrac{\boldsymbol{r}}{\tau}) + \tau D_F \left(\tfrac{\boldsymbol{r}}{\tau} \| \boldsymbol{q}(x) \right). \tag{9}$$

Later, in Section 3, we will find it helpful to consider a shift v of the expected cost; i.e. $\mathcal{R}(\boldsymbol{\pi}, \boldsymbol{r} - v, x) = v - \boldsymbol{r} \cdot \boldsymbol{\pi}$, noting this does not affect the location of the minimizer in $\boldsymbol{q}$. The above characterization then allows us to formulate a convex calibrated surrogate by reversing the divergence.

Theorem 3 *For an arbitrary baseline* v *and* $\tau > 0$*, let*

$$L(\boldsymbol{q}, \boldsymbol{r}, x) = \tau D_F \left(\boldsymbol{q}(x) + \tfrac{v}{\tau} \left\| \tfrac{\boldsymbol{r}}{\tau} \right) + \tfrac{\tau}{4} \left\| \boldsymbol{q}(x) - \tfrac{\boldsymbol{r} - v}{\tau} \right\|^2. \tag{10}$$

Then, for any fixed v, L *is strongly convex in* $\boldsymbol{q}$ *and calibrated with respect to the smoothed (shifted) risk* $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x) = \mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) - v$ *with calibration function* $\delta(\epsilon, x) = \epsilon \; \forall x$.

Therefore, any desired level of accuracy in minimizing empirical smoothed risk can be achieved by approximately minimizing the surrogate loss L to appropriate accuracy.

2.2 Experimental Evaluation

To first assess the overall approach, we evaluate how well optimizing the surrogate (10) minimizes true risk, using a separate test set for evaluation. As baselines, we compare to directly minimizing empirical risk $\hat{\mathcal{R}}(\boldsymbol{\pi})$ (1), and the standard supervised objectives, log-likelihood, $-\mathbb{E}_{\boldsymbol{p}}[\log \boldsymbol{\pi}]$, and squared error, $\|\boldsymbol{q}(x) - \frac{\boldsymbol{r} - v}{\tau}\|^2$. Empirically, we found it beneficial to relax (10) to a tunable combination between the components and empirical risk. We refer to such a tuned loss as "Composite" in all experimental results. Since the surrogate objective is a combination of the "reversed KL" objective $D_{F^*}(\boldsymbol{p} \| \boldsymbol{\pi})$ (7) and the squared error, we also evaluate $D_{F^*}(\boldsymbol{p} \| \boldsymbol{\pi})$ alone to isolate its effect.

MNIST We first consider MNIST data, training a fully connected model with one hidden layer of 512 ReLU units. The original training data was partitioned into the first 55K examples for training and the last 5K examples for validation. We use the validation data to select hyperparameters, including learning rate, mini-batch size, and combination weights (details in appendix). The policy was trained by minimizing each objective using SGD with momentum fixed at 0.9 [33] for 100 epochs.

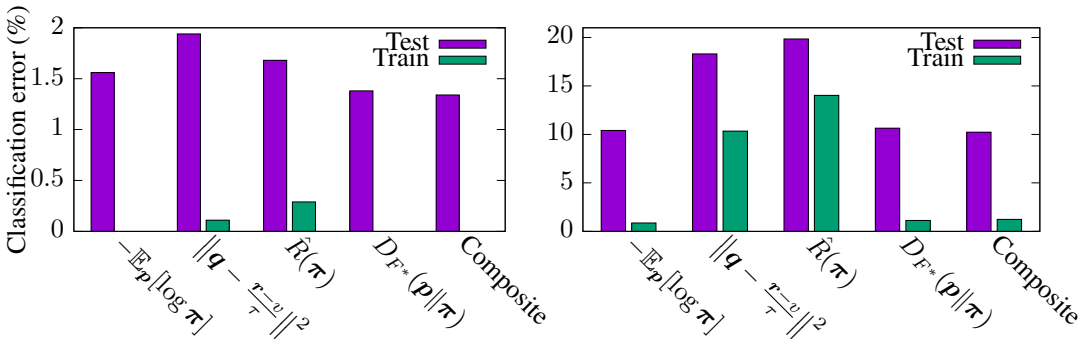

(a) MNIST (b) CIFAR10

Figure 1: Training with full reward feedback across all actions (see appendix for additional results).

CIFAR-10 Next we considered the CIFAR-10 data set [15] and trained a Resnet-20 architecture [12], using the standard 50K training, 10K validation split. We set any unspecified model hyperparameters to the defaults for resnet in the open source tensor2tensor library [39] and tuned learning rate and the composite loss combination weights on validation data. All objectives were trained using the Momentum optimizer with cosine decay learning rate for 250 epochs (details in appendix).

The results in Figure 1 confirm that directly minimizing $\hat{\mathcal{R}}(\boldsymbol{\pi})$ (1), is not always competitive: it yields the highest training error on both MNIST and CIFAR10, as well as poor test error. For MNIST, it is striking that generalization can still be improved with respect to the standard log-likelihood baseline. For CIFAR-10, as shown in Figure 1b, minimizing the reverse KL $D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$ achieved 10.7% test error, which significantly improves directly optimizing empirical risk $\hat{\mathcal{R}}(\boldsymbol{\pi})$, which obtained 19.9% test error. The reverse KL was competitive even against the baseline log-likelihood, which achieved 10.4% test error. The results for squared error are worse than log-likelihood, while the $D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$ objective performs better in both data sets. This suggests that the generalization improvements are coming from better minimization of $D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$, while the squared error term is helping improve the optimization landscape. To further investigate whether $\hat{\mathcal{R}}(\boldsymbol{\pi})$ suffers from a difficult optimization landscape, we ran a much longer training experiment (see appendix), finding that every method except squared loss is eventually able to achieve about 6.5% test error, but at significant cost.

3 Batch Contextual Bandits

We now extend these developments to the contextual bandit case. To focus on the most challenging and practical scenario, we assume a single action has been observed in each context. Therefore, the training data consists of tuples $\mathcal{D} = \{(x_i, a_i, r_i, \beta_i)\}$, where $x_i \in X$ is a context, $a_i \in \{1, ..., K\}$ is an action, $r_i \in \mathbb{R}$ is a reward, and β_i is the proposal probability of a_i. For simplicity, we assume a stationary behaviour (logging) policy $\boldsymbol{\beta} : X \to \Delta^K$ was used to select the actions, hence $\beta_i = \beta(a_i|x_i)$. Although β might not be known [20], estimating it from $\mathcal{D}$ has proved effective [41, 43, 4]. We continue to assume contexts and rewards are generated i.i.d. from a joint distribution $p(x, \boldsymbol{r})$, but the distribution of rewards $\boldsymbol{r}(x) \sim p(\boldsymbol{r}|x)$ and actions $a \sim \beta(a|x)$ are conditionally independent given the context x [43]. Other more elaborate models for missing data are possible, but require committing to stronger assumptions about the data generation and behavior process [3, 21].

As before, the goal is to infer a policy $\boldsymbol{\pi} : X \to \Delta^K$ that maximizes expected reward. Here we define the *true risk* of a policy $\boldsymbol{\pi}$ as in (1), but the empirical risk, also defined in (1), is no longer directly observable because it requires rewards for all actions. The standard solution is to formulate an unbiased estimate of the full empirical risk (which is itself an unbiased estimate of the true risk), then use this as a policy optimization objective. In fact, the current literature is dominated by such an approach, where an unbiased (or nearly unbiased) estimate of the empirical risk (1) is first formulated via *importance correction* then used as a training objective [14, 16, 17, 29, 34–36]. Unfortunately, importance correction introduces significant variance in gradient estimates, even using standard variance reduction techniques. Also, as identified in Section 2, *even if variance could be completely eliminated*, the underlying optimization landscape presents difficulties.

3.1 Reward Estimation

Before focusing on policy optimization, we first need to address the problem of estimating rewards from incomplete data. We here adopt a simple approach of imputing missing values with a model $\boldsymbol{q} : X \to \mathbb{R}^K$. That is, for a context x, observed action a and observed reward r_a, we estimate the full reward vector $\boldsymbol{r}$ by

$$\hat{\boldsymbol{r}}(x) \quad = \quad \tau \boldsymbol{q}(x) + \mathbf{1}_a \lambda(x,a)(r_a - \tau \boldsymbol{q}(x)_a), \tag{11}$$

with parameters $\lambda(x,a)$ and τ. This construction allows the local risk of any policy $\boldsymbol{\pi}$ to be estimated by $\mathcal{R}(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) = -\boldsymbol{\pi}(x) \cdot \hat{\boldsymbol{r}}(x)$. Although (11) seems simplistic, it is able to express most estimators in the literature by suitable choices of τ and $\lambda(x,a)$. For example, choosing $\tau = 0$ and $\lambda(x,a) = \beta(a|x)^{-1}$ yields importance weighting $\mathcal{R}(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) = \frac{\boldsymbol{\pi}(x)_a}{\beta(a|x)} r_a$; choosing $\tau = 1$ and $\lambda(x,a) = 0$ yields the "direct method" $\mathcal{R}(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) = \boldsymbol{\pi}(x) \cdot \boldsymbol{q}(x)$; and choosing $\tau = 1$ and $\lambda(x,a) = \beta(a|x)^{-1}$ yields the "doubly robust" estimate $\mathcal{R}(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) = \boldsymbol{\pi}(x) \cdot \boldsymbol{q}(x) + \frac{\boldsymbol{\pi}(x)_a}{\beta(a|x)}(r_a - \boldsymbol{q}(x)_a)$ [6]. The "switch" estimator [40] can be expressed for a given threshold $\theta > 0$ by setting $\tau = 1$ and $\lambda(x,a) = 0$ if $\boldsymbol{\pi}(x)_a / \theta > \beta(a|x)$, otherwise $\tau = 0$ and $\lambda(x,a) = \beta(a|x)^{-1}$. The "switch" estimator generalizes the trimmed importance estimator [3], and is argued in [40] to be superior to the "magic" estimator [38]. The "self-normalized" importance estimator [36] can also be expressed by setting $\tau = 0$ and $\lambda(x,a) = \beta(a|x)^{-1} / \sum_{(x_i,a_i,r_i) \in \mathcal{D}} \boldsymbol{\pi}(x_i)_{a_i} \beta(a_i|x_i)^{-1}$.

For any fixed $\boldsymbol{q}(x)$ and τ it is easy to show that $\lambda(x,a) = \beta(a|x)^{-1}$ implies $\mathbb{E}_{a \sim \beta(\cdot|x)}[\hat{\boldsymbol{r}}(x)] = \boldsymbol{r}(x)$.

A key question is the provenance of $\boldsymbol{q}$. The standard approach is to recover $\boldsymbol{q}$ by regressing to the observed rewards $\boldsymbol{q} = \arg\min_{\boldsymbol{q} \in \mathcal{H}} \sum_{(x_i,a_i,r_i) \in \mathcal{D}} (r_i - \boldsymbol{q}(x_i)_{a_i})^2$ for a class of models $\mathcal{H}$. Note that this is equivalent to conducting a policy evaluation step for the policy β. Stochastic contextual bandits are a restricted case of reinforcement learning where every policy has the same action value function, $\mathbb{E}[\boldsymbol{r}(x)]$. Hence, a single policy evaluation, yielding $\boldsymbol{q}$, can in principle be used to evaluate any policy, since evaluating $\boldsymbol{\pi}$ instead of β does not change action values but only introduces covariate shift.

3.2 Policy Optimization

For policy optimization, one could adopt the least squares estimate $\boldsymbol{q}$ and optimize a separate policy, but if $\boldsymbol{\pi}$ uses the same architecture the optimum is simply $\boldsymbol{\pi} = \boldsymbol{f} \circ \boldsymbol{q}$ (under $\mathcal{S}$). In Section 2, we saw that least squares estimation of $\boldsymbol{q}$ did not perform well, nor do we expect so here. We would like to gain the advantages realized in Section 2, but an actor-critic approach obviates policy optimization.

Instead, to couple the value estimator to policy optimization, we consider a unified approach where the actor and the critic *are the same model*. That is, we use the policy transformation $\boldsymbol{\pi} = \boldsymbol{f} \circ \boldsymbol{q}$ from Section 2, but now explicitly treat the logits as action value estimates. A unified actor-critic model has been considered previously [22]. In the partially observed case, we propose to replace the observed reward vector with the estimate $\hat{\boldsymbol{r}}$ derived from $\boldsymbol{q}$, allowing any loss to be applied. Although such an approach seems naive, we find that maintaining this form of strict mutual consistency between the value estimates and policy, combined with the estimator $\hat{\boldsymbol{r}}$ and surrogate losses, leads to effective empirical performance. Moreover, we will find that this approach is theoretically justified.

3.3 Calibrated Surrogate

Given (x, a, r_a), define the optimal imputed local risk and the suboptimality gap respectively by

$$\mathcal{S}^*_\tau(\hat{\boldsymbol{r}}, x) \quad = \quad \inf_{\boldsymbol{q} \in \mathcal{Q}} \mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \hat{\boldsymbol{r}}, x) \quad \text{and} \quad \mathcal{G}_\tau(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) \;=\; \mathcal{S}_\tau(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) - \mathcal{S}^*_\tau(\hat{\boldsymbol{r}}, x). \tag{12}$$

Equality (9) can then be used to show the divergence $D_F\left(\frac{\hat{\boldsymbol{r}}}{\tau} \| \boldsymbol{q}\right)$ characterizes the suboptimality gap:

Proposition 4 *For any $\boldsymbol{q}$, $\tau > 0$ and observation (x, a, r_a): $\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau} \middle\| \boldsymbol{q}(x)\right) = \mathcal{G}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \hat{\boldsymbol{r}}, x)$.*

If we consider the imputed form of the surrogate objective $L(\boldsymbol{q}, \hat{\boldsymbol{r}}, x)$ defined in Theorem 3 we then find that the surrogate remains calibrated for the imputed smoothed risk.

Theorem 5 *For any model $\boldsymbol{q}$, $\tau > 0$, observation (x, a, r_a), and baseline v:*

$$L(\boldsymbol{q}, \hat{\boldsymbol{r}}, x) \quad \geq \quad \tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau} \middle\| \boldsymbol{q}(x) + \tfrac{v}{\tau}\right) \;=\; \mathcal{G}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \hat{\boldsymbol{r}}, x) \;\geq\; 0. \tag{13}$$

Moreover, L is calibrated with respect to $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \hat{\boldsymbol{r}} - v, x)$ with calibration function $\delta(x, \epsilon) = \epsilon$.

This result suggests a simple algorithmic approach for policy optimization: given the data $\mathcal{D} = \{(x_i, a_i, r_i, \beta_i)\}$, minimize the imputed empirical surrogate objective with respect to the model $\boldsymbol{q}$:

$$\min_{\boldsymbol{q} \in \mathcal{Q}} \hat{L}(\boldsymbol{q}, \mathcal{D}) \quad \text{where} \quad \hat{L}(\boldsymbol{q}, \mathcal{D}) = \tfrac{1}{T} \textstyle\sum_{(x_i, a_i, r_i, \beta_i) \in \mathcal{D}} L(\boldsymbol{q}, \hat{\boldsymbol{r}}, x_i). \tag{14}$$

That is, we combine the estimate $\hat{\boldsymbol{r}}$ from Section 3.1, (11), with the surrogate L from Section 2, (10).

3.4 Analysis

The expected smoothed risk quantities we seek to control are defined by:

$$\mathcal{S}_\tau(\boldsymbol{\pi}) \;=\; \mathbb{E}[\mathcal{S}_\tau(\boldsymbol{\pi}, \boldsymbol{r}, x)], \quad \mathcal{S}_\tau^* \;=\; \inf_{\boldsymbol{q} \in \mathcal{Q}} \mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}) \quad \text{and} \quad \mathcal{G}_\tau(\boldsymbol{\pi}) \;=\; \mathcal{S}_\tau(\boldsymbol{\pi}) - \mathcal{S}_\tau^*. \tag{15}$$

For the purposes of analysis, we assume training data consists of tuples drawn from $(x, a, r_a) \sim p(x, \boldsymbol{r})\beta(a|x)$, and that the estimate $\hat{\boldsymbol{r}}$ is *unbiased*; i.e., $\mathbb{E}[\hat{\boldsymbol{r}}|x] = \mathbb{E}[\boldsymbol{r}|x]$, using $\lambda(x, a) = \beta(a|x)^{-1}$.

First, observe that, in expectation, the surrogate objective upper bounds the divergence in Theorem 5, which, in turn, by Jensen's inequality, bounds the suboptimality gap in the expected smoothed risk.

Theorem 6 *For any model $\boldsymbol{q}$, any $\hat{\boldsymbol{r}}$ such that $\mathbb{E}[\hat{\boldsymbol{r}}|x] = \mathbb{E}[\boldsymbol{r}|x]$, and any baseline v:*

$$\mathbb{E}[L(\boldsymbol{q}, \hat{\boldsymbol{r}}, x)] \;\;\geq\;\; \mathbb{E}\left[\tau D_F\left(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau} \middle\| \boldsymbol{q}(x) + \tfrac{v}{\tau}\right)\right] \;\geq\; \mathcal{G}_\tau(\boldsymbol{f} \circ \boldsymbol{q}) \;\geq\; 0. \tag{16}$$

Therefore, minimizing (14), in expectation, minimizes the true smoothed risk (15).

This result can be made stronger by observing that, under mild assumptions, the empirical divergence $\hat{D}(\boldsymbol{q}, \mathcal{D}) = \frac{1}{T}\sum_i D_F\big(\frac{\hat{\boldsymbol{r}}(x_i)}{\tau} \big\| \boldsymbol{q}(x_i)\big)$ also concentrates to its expectation, uniformly over $\boldsymbol{q} \in \mathcal{H}$, for a well behaved model class $\mathcal{H}$. ***In the appendix***, we specify the conditions on $\mathcal{H}$, β, and $p(x, \boldsymbol{r})$ that, in addition to $\hat{\boldsymbol{r}}$ being unbiased (i.e. $\mathbb{E}[\hat{\boldsymbol{r}}|x] = \mathbb{E}[\boldsymbol{r}|x]$), ensure finite sample concentration. We refer to a collection $\mathcal{H}$, β, $p(x, \boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ that satisfies these conditions as "well behaved".

Lemma 7 *Assume $\mathcal{H}$, β, $p(x, \boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then for any $\tau, \delta > 0$ there exists a constant C such that with probability at least $1 - \delta$:*

$$\mathbb{E}\left[D_F\left(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau} \big\| \boldsymbol{q}(x)\right)\right] \;\;\leq\;\; \hat{D}_F(\boldsymbol{q}, \mathcal{D}) + \tfrac{C}{\sqrt{T}} \quad \forall \boldsymbol{q} \in \mathcal{H}. \tag{17}$$

Combining Theorem 6 with Lemma 7 it can be shown that for finite sample size T, with high probability, the empirical surrogate (14) is approximately calibrated with respect to smoothed risk.

Theorem 8 *Assume $\mathcal{H}$, β, $p(x, \boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then for any v and $\tau, \delta > 0$, there exists a C such that with probability at least $1 - \delta$: if $\hat{L}(\boldsymbol{q}, \mathcal{D}) < \frac{\tau C}{\sqrt{T}}$ for $\boldsymbol{q} \in \mathcal{H}$ then $\mathcal{G}_\tau(\boldsymbol{f} \circ \boldsymbol{q}) \leq \frac{2\tau C}{\sqrt{T}}$.*

That is, if $\hat{L}(\boldsymbol{q}, \mathcal{D})$ can be sufficiently minimized within $\mathcal{H}$, the suboptimality gap achieved by $\boldsymbol{q}$ will be near-optimal with high probability, with bound diminishing to zero for large sample size.

3.5 Discussion

If we let $\tau = 0$ in the definition of $\hat{\boldsymbol{r}}$, (11), then $\hat{\boldsymbol{r}}$ exhibits no dependence on $\boldsymbol{q}$, making $L(\boldsymbol{q}, \hat{\boldsymbol{r}}, x)$ convex in $\boldsymbol{q}$. However, we have found that empirical results are improved by choosing $\tau > 0$, since this compels the logits $\boldsymbol{q}$ to also model observed rewards. In addition, even though using an unbiased $\hat{\boldsymbol{r}}$ enables the theory above, achieving unbiasedness via importance correction increases variance, degrades the quality of the reward estimate, and yields inferior results. In our experiments, we considered τ to be a hyperparameter, and also considered different choices for λ, including $\lambda(x, a) = \beta(a|x)^{-1}$ and $\lambda(x, a) = 1$. We also introduced tunable combination weights between the Bregman divergence and the squared error terms in (14), similar to the relaxation in Section 2.2. In all cases, we chose hyperparameters from validation data only.

Note that the approach developed in this paper differs fundamentally from recent trust-region and proximal methods in reinforcement learning [30, 31], which still directly optimize expected return, possibly with entropy regularization [23]. These methods use proximal constraints/regularization to improve the stability of optimization, but apply a "surrogate" as a local not a global modification of the objective. By contrast, we are changing the entire optimization objective globally, not locally, and train to maximize a target that is different from expected return.

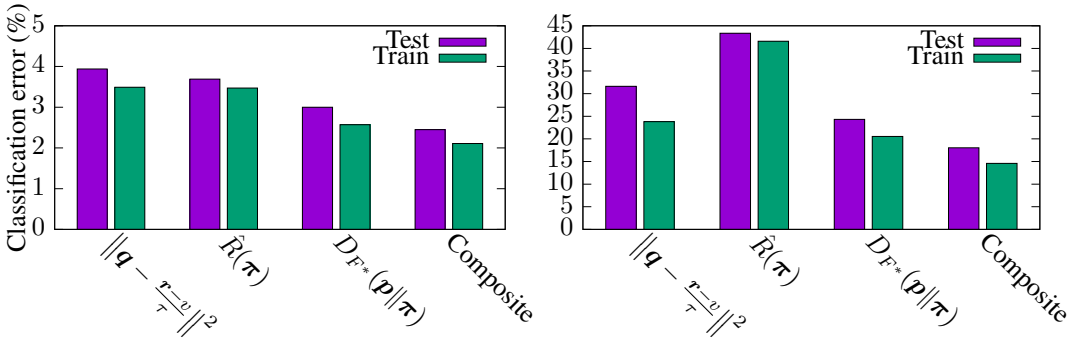

(a) MNIST (b) CIFAR10

Figure 2: Training with partial (single action) reward feedback (see appendix for additional results).

Unlike previous entropy regularized approaches [11, 22], which generally consider split actor-critic models, we achieve success with a single model that serves as both.

Another subtlety with optimizing importance corrected objectives, such as $\mathcal{R}(\boldsymbol{\pi}, \hat{\boldsymbol{r}}, x) = \frac{\boldsymbol{\pi}(x)_a}{\beta(a|x)} r_a$, is that this does not account for the policy's data coverage [20]; that is, a policy might minimize such an objective by moving mass $\boldsymbol{\pi}(x_i)_{a_i}$ away from the training observations (x_i, a_i, r_i), leading to a phenomenon known as "propensity overfitting" [34–36]. This effect can be countered by adding coverage-dependent confidence intervals to the estimates [34–36], or constraining [10] or regularizing [20] toward the logging policy choices. Although such regularization is helpful, it is orthogonal to the aim of the current investigation, as any objective can be augmented in this way.

3.6 Experimental Evaluation

As is standard in the field [34–36], we form a partially observed version of a supervised learning task by sampling actions from a behaviour policy $\boldsymbol{\pi}_0$, assigning a reward of 1 when the action chosen matches the correct label and a reward of 0 when it does not. Reward on all counterfactual actions is therefore missing. For the MNIST and CIFAR-10 experiments, we used the same architecture, optimizer and model configurations used in the fully observed label experiments. For the empirical risk estimator $\hat{\mathcal{R}}(\boldsymbol{\pi})$ we used importance correction $\lambda(x, a) = \beta(a|x)^{-1}$; however $\lambda(x, a) = 1$ proved to be more effective for $D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$, which is equivalent to replacing the counterfactual rewards with the model estimate.

MNIST We evaluate the results when data is collected by the uniform behavior policy, $\boldsymbol{\pi}_0(x) = \mathbf{1} \cdot \frac{1}{10}$. The hyperparameters for all objectives were re-optimized on validation data, using the same optimization algorithm as before. Details are given in the appendix.

CIFAR-10 Here we also evaluate the results when data is collected by the uniform behavior policy, $\boldsymbol{\pi}_0(x) = \mathbf{1} \cdot \frac{1}{10}$. However, in addition, we also evaluate the proposed objectives using data released in a recently published benchmark on CIFAR-10 [14]. (Note that the behavior policy itself was not released in this benchmark; instead different sized training sets of size 50k, 100k, 150k, and 250k were generated using this policy.) We used this alternative data to produce each column in Table 1. For the CIFAR-10 experiments we simply set $\tau = 1$. Additional details are given in the appendix.

Figure 2 shows the results for training on MNIST and CIFAR-10 given data collected by the random behavior policy. In both cases, the composite objective yields improvements over optimizing $\hat{R}(\boldsymbol{\pi})$ directly. To investigate whether this difficulty is due to plateaus, we again conduct significantly longer training in the appendix, finding that the Composite and $D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$ objectives remain advantageous. Table 1 then shows results on CIFAR-10 using the alternative behavior data from [14]. This data appears to be more condusive to optimizing $\hat{R}(\boldsymbol{\pi})$ directly, although even in this scenario the composite objective is still competitive, significantly improving the results reported in [14].

Criteo We also test the proposed surrogate objective on the Criteo data set [18], a large-scale test-bed for evaluating batch contextual bandit methods [3, 34]. Here again the behavior policy was not

Examples	50k	100k	150k	250k
$\hat{\mathcal{R}}(\boldsymbol{\pi})$	8.71	7.51	6.92	6.66
$\left\|\boldsymbol{q} - \frac{\boldsymbol{r}-v}{\tau}\right\|^2$	21.85	18.26	14.65	12.86
$D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$	16.00	9.85	8.68	8.75
Composite	8.34	6.92	6.57	6.36

Table 1: CIFAR-10: Test error % for the bandit feedback data sets from [14] with increasing number of training examples.

Objectives	$\hat{\mathcal{R}}(\pi) \times 10^4$
Random	43.68 ± 2.11
Behavior	53.55
DRO $\hat{\mathcal{R}}(\boldsymbol{\pi})$ [7]	53.07 ± 2.27
POEM [34]	51.89 ± 1.73 [2]
$\hat{\mathcal{R}}(\boldsymbol{\pi})$	51.72 ± 1.42
$\left\|\boldsymbol{q} - \frac{\boldsymbol{r}-v}{\tau}\right\|^2$	52.00 ± 1.28
$D_{F^*}(\boldsymbol{p}\|\boldsymbol{\pi})$	52.30 ± 0.83
Composite	55.09 ± 2.86

Table 2: Criteo: Importance sampling estimated reward on test. Error bars are 99% confidence intervals under normal distribution.

released, but only its generated data. Following [18], we use only banners with a single slot (i.e., where only a single item is chosen) in our learning and evaluation. These banners are randomly split into training, validation and test sets, each containing 7 million records, using the script provided by [18]. There are 35 features used to describe the context and candidates actions (2 continuous and the rest categorical). We encode the discrete features using one-hot encoding, and build linear models using different learning losses. For evaluation, we report the importance sampling based estimates of reward (user clicks on banner) $\hat{R}(\boldsymbol{\pi})$ on the test set, as in [18].

We compare the proposed surrogates with several state-of-the-art methods on this data set. Hyperparameters of the different methods were tuned on validation data, and all objectives optimized by SGD with momentum and batch size between 1K and 5K; more details regarding the experiment setup and hyperparameter choices are given in the appendix. All methods use the same input encoding to map inputs x to $\boldsymbol{\phi}(x)$. In particular, we evaluated the following:

Random: Choose a candidate banner (i.e. an action) uniformly at random to display.
Behavior: Simply report observed reward on the test set acting according to the logging policy $\boldsymbol{\beta}$.
Squared $\left\|\boldsymbol{q} - \frac{\boldsymbol{r}-v}{\tau}\right\|^{\mathbf{2}}$: We set $v = 0$ is effective since expected reward is close to 0.
$\hat{\mathbf{R}}(\boldsymbol{\pi})$: Directly optimize $\hat{R}(\boldsymbol{\pi})$ using importance correction; i.e. $\lambda(x, a) = \beta(a|x)^{-1}$, $\tau = 0$ in (11).
DRO $\hat{\mathbf{R}}(\boldsymbol{\pi})$: Optimize the doubly robust estimator [7]; i.e., $\lambda(x, a) = \beta(a|x)^{-1}$ and $\tau > 0$ in (11).
POEM: Combines importance corrected empirical risk estimation, $\hat{R}(\boldsymbol{\pi})$, with a regularization that penalizes the variance of the estimated $\hat{R}(\boldsymbol{\pi})$ [34]. We tuned the additional regularization factor λ. (To keep a fair comparison with the other methods, we did not impose capping on the importance weight here, which was additionally tuned in [34].)
$\mathbf{D}_{\mathbf{F}^*}(\boldsymbol{p}\|\boldsymbol{\pi})$: The imputation strategy for the reward uses $\lambda(x, a) = \beta(a|x)^{-1}$ and we tune $\tau > 0$.
Composite: We combine the $\hat{R}(\boldsymbol{\pi})$ objective with $\mathbf{D}_{\mathbf{F}^*}(\boldsymbol{p}\|\boldsymbol{\pi})$. In addition to the scaling factor τ, we tuned the combination weights.

Table 2 reports the estimated reward obtained by each method on test data. Here we can see that training with the proposed surrogate performs competitively against previous state-of-the-art-methods.

4 Conclusion

We investigated alternative objectives for policy optimization in cost-sensitive classification and contextual bandits. The formulations developed are directly applicable to deep learning and improve the underlying optimization landscape. The empirical results in both the cost-sensitive classification and batch contextual bandit scenarios replicate or surpass training with state-of-the-art baseline objectives, merely through the optimization of the non-standard loss functions. There remain several opportunities for further development of surrogate training objectives for sequential decision making tasks (i.e. in planning and reinforcement learning).

[2] The number reported here is lower than reported in the original paper [34]. One hypothesis is that this is due to the removal of the cap on importance weights. A similar result to what we obtain was reported in [20].

References

[1] Naoki Abe, Bianca Zadrozny, and John Langford. An iterative method for multi-class cost-sensitive learning. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 3–11, 2004.

[2] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473), 2006.

[3] Léon Bottou, Jonas Peters, Joaquin Quiñonero Candela, Denis Xavier Charles, Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Y. Simard, and Ed Snelson. Counterfactual reasoning and learning systems: the example of computational advertising. *Journal of Machine Learning Research*, 14(1):3207–3260, 2013.

[4] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H Chi. Top-k off-policy correction for a reinforce recommender system. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 456–464, 2019.

[5] Jacek P. Dmochowski, Paul Sajda, and Lucas C. Parra. Maximum likelihood in cost-sensitive learning: Model specification, approximations, and upper bounds. *Journal of Machine Learning Research*, 11:3313–3332, 2010.

[6] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):458–511, 2014.

[7] Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1097–1104, 2011.

[8] Charles Elkan. The foundations of cost-sensitive learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 973–978, 2001.

[9] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1446–1455, 2018.

[10] Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2052–2062, 2019.

[11] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1856–1865, 2018.

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.

[13] Klaus-Uwe Höffgen, Hans Ulrich Simon, and Kevin S. Van Horn. Robust trainability of single neurons. *Journal of Computer and System Sciences*, 50(1):114–125, 1995.

[14] Thorsten Joachims, Adith Swaminathan, and Maarten de Rijke. Deep learning with logged bandit feedback. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.

[15] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

[16] Carolin Lawrence and Stefan Riezler. Improving a neural semantic parser by counterfactual learning from human bandit feedback. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1820–1830, 2018.

[17] Carolin Lawrence, Artem Sokolov, and Stefan Riezler. Counterfactual learning from bandit feedback under deterministic logging : A case study in statistical machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2566–2576, 2017.

[18] Damien Lefortier, Adith Swaminathan, Xiaotao Gu, Thorsten Joachims, and Maarten de Rijke. Large-scale validation of counterfactual learning methods: A test-bed. *CoRR*, abs/1612.00367, 2016.

[19] Hsuan-Tien Lin. Reduction from cost-sensitive multiclass classification to one-versus-one binary classification. In *Proceedings of the Asian Conference on Machine Learning (ACML)*, 2014.

[20] Yifei Ma, Yu-Xiang Wang, and Balakrishnan (Murali) Narayanaswamy. Imitation-regularized offline learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.

[21] Karthika Mohan, Judea Pearl, and Jin Tian. Graphical models for inference with missing data. In *Advances in Neural Information Processing Systems 26*, pages 1277–1285, 2013.

[22] Ofir Nachum, Mohammad Norouzi, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems 31*, pages 2772–2782, 2017.

[23] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Trust-PCL: An off-policy trust region method for continuous control. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.

[24] Deirdre B. O'Brien, Maya R. Gupta, and Robert M. Gray. Cost-sensitive multi-class classification from probability estimates. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 712–719, 2008.

[25] Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton. Regularizing neural networks by penalizing confident output distributions. *CoRR*, 1701.06548, 2017.

[26] Bernardo Ávila Pires, Csaba Szepesvári, and Mohammad Ghavamzadeh. Cost-sensitive multiclass classification risk bounds. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1391–1399, 2013.

[27] Mark D. Reid and Robert C. Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12:731–817, 2011.

[28] Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto, Michele Piana, and Alessandro Verri. Are loss functions all the same? *Neural Computation*, 16(5):1063–107, 2004.

[29] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1670–1679, 2016.

[30] John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1889–1897, 2015.

[31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.

[32] Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.

[33] Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1139–1147, 2013.

[34] Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research*, 16:1731–1755, 2015.

[35] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 814–823, 2015.

[36] Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. In *Advances in Neural Information Processing Systems 28*, pages 3231–3239, 2015.

[37] Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8:1007–1025, 2007.

[38] Philip S. Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2139–2148, 2016.

[39] Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N. Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. Tensor2tensor for neural machine translation. *CoRR*, abs/1803.07416, 2018.

[40] Yu-Xiang Wang, Alekh Agarwal, and Miroslav Dudík. Optimal and adaptive off-policy evaluation in contextual bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3589–3597, 2017.

[41] Yuan Xie, Boyi Liu, Qiang Liu, Zhaoran Wang, Yuan Zhou, and Jian Peng. Off-policy evaluation and learning from logged bandit feedback: Error reduction via surrogate policy. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

[42] Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5, 2004.

[43] Zhengyuan Zhou, Susan Athey, and Stefan Wager. Offline multi-action policy learning: Generalization and optimization. *CoRR*, abs/1810.04778, 2018.

Supplement

Surrogate Objectives for Batch Policy Optimization in One-step Decision Making

Minmin Chen* **Ramki Gummadi*** **Chris Harris*** **Dale Schuurmans*†**
*Google †University of Alberta

1 Definitions

Throughout this appendix we use the same notation and definitions from the main body of the paper. In particular, for a vector $\boldsymbol{q} \in \mathbb{R}^K$ let

$$\begin{aligned} \boldsymbol{\pi}(\boldsymbol{q}) &= \boldsymbol{f}(\boldsymbol{q}) && (1)\\ \boldsymbol{f}(\boldsymbol{q}) &= e^{\boldsymbol{q}-F(\boldsymbol{q})} && (2)\\ F(\boldsymbol{q}) &= \log(\mathbf{1}\cdot e^{\boldsymbol{q}}). && (3) \end{aligned}$$

We also use the same risk definitions as the main body, in particular:

local risk

$$\begin{aligned} \mathcal{R}(\boldsymbol{\pi},\boldsymbol{r},x) &= -\boldsymbol{r}\cdot\boldsymbol{\pi}(x) && (4)\\ \mathcal{R}^*(\boldsymbol{r},x) &= \inf_{\boldsymbol{q}\in\mathcal{Q}} \mathcal{R}(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r},x), && (5) \end{aligned}$$

expected risk

$$\mathcal{R}(\boldsymbol{\pi}) = -\mathbb{E}[\boldsymbol{\pi}(x)\cdot\boldsymbol{r}], \quad (6)$$

local smoothed risk

$$\begin{aligned} \mathcal{S}_\tau(\boldsymbol{\pi},\boldsymbol{r},x) &= -\boldsymbol{r}\cdot\boldsymbol{\pi}(x) + \tau\boldsymbol{\pi}\cdot\log\boldsymbol{\pi}(x) && (7)\\ \mathcal{S}^*_\tau(\boldsymbol{r},x) &= \inf_{\boldsymbol{q}\mathcal{Q}} \mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r},x) && (8)\\ \mathcal{G}_\tau(\boldsymbol{\pi},\boldsymbol{r},x) &= \mathcal{S}_\tau(\boldsymbol{\pi},\boldsymbol{r},x) - \mathcal{S}^*_\tau(\boldsymbol{r},x), && (9) \end{aligned}$$

expected smoothed risk

$$\begin{aligned} \mathcal{S}_\tau(\boldsymbol{\pi}) &= \mathbb{E}[\mathcal{S}_\tau(\boldsymbol{\pi},\boldsymbol{r},x)] && (10)\\ \mathcal{S}^*_\tau &= \inf_{\boldsymbol{q}\mathcal{Q}} \mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q}) && (11)\\ \mathcal{G}_\tau(\boldsymbol{\pi}) &= \mathcal{S}_\tau(\boldsymbol{\pi}) - \mathcal{S}^*_\tau. && (12) \end{aligned}$$

2 Proofs for Section 2: Cost-sensitive Classification

Theorem 1 *Even for a single context x, a deterministic reward vector $\boldsymbol{r}$, and a linear model $\boldsymbol{q}(x) = W\boldsymbol{\phi}(x)$, the function $\boldsymbol{r}\cdot\boldsymbol{f}(\boldsymbol{q}(x))$ can have a number of local maxima in W that is exponential in the number of actions K and the number of features in $\boldsymbol{\phi}$.*

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

Proof: To demonstrate the possibility of separated local maxima, start by considering a concrete construction with 5 actions and 1 feature. In particular, let

$$\boldsymbol{r}_1 \;=\; \begin{bmatrix} 1 \\ 2 \\ -1 \\ 2 \\ 1 \end{bmatrix} \quad \text{and} \quad \Phi_1 \;=\; \begin{bmatrix} 2 \\ 1 \\ 0 \\ -1 \\ -2 \end{bmatrix} \tag{13}$$

hence $\boldsymbol{q} = \Phi_1 w$ for a scalar parameter w. Note that in this case the policy is given by

$$\boldsymbol{\pi} \;=\; \boldsymbol{f}(\Phi_1 w) \;=\; \frac{1}{d(w)} \begin{bmatrix} e^{2w} \\ e^{w} \\ 1 \\ e^{-w} \\ e^{-2w} \end{bmatrix}, \tag{14}$$

$$\text{where} \quad d(w) \;=\; e^{2w} + e^{w} + 1 + e^{-w} + e^{-2w} \;=\; 2\cosh(2w) + 2\cosh(w) + 1. \tag{15}$$

Therefore, the value function is given by

$$v(w) \;=\; \boldsymbol{r}_1^\top \boldsymbol{\pi} \;=\; \frac{2\cosh(2w) + 4\cosh(w) - 1}{2\cosh(2w) + 2\cosh(w) + 1} \;=\; \frac{n(w)}{d(w)}, \tag{16}$$

$$\text{where} \quad n(w) \;=\; 2\cosh(2w) + 4\cosh(w) - 1. \tag{17}$$

To determine the critical points of the value function, consider the derivative

$$\frac{dv}{dw} \;=\; \frac{2\sinh(w)(8\cosh(w) - 4\cosh^2(w) + 1)}{d(w)^2}. \tag{18}$$

Recall that $\cosh(w) \geq 1$, hence $d(w) \geq 5$, and therefore the zeros for $\frac{dv}{dw}$ occur whenever the numerator is zero. This implies there are exactly three critical points, at $w = 0$ and $w = \pm\,\mathrm{acosh}(1 + \frac{\sqrt{5}}{2})$ ($\approx \pm 1.3826$). One can also determine that $n(w) \geq d(w)$, hence $v(w) \geq 1$. Finally, observe that since $\cosh(w)$ is an even function, so must be $n(w)$, $d(w)$, and $v(w)$. The function $v(w)$ is plotted in Figure 1.

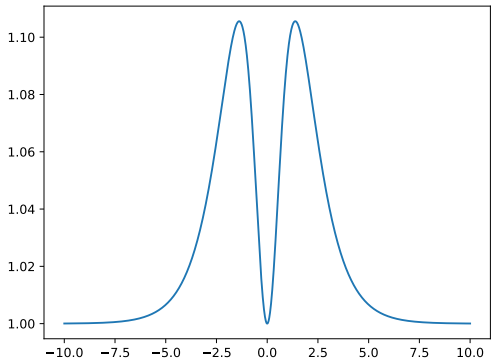

Figure 1: Plot of the value function $v(w)$.

We can now use this simple example as a widget for creating a combinatorial explosion of local maxima. We achieve this by tiling the previous construction as follows. Let t denote the number of tiles. Expand the previous construction to $5t$ actions and t features by replicating $\boldsymbol{r}_1$ and Φ_1, each t times, in the following manner:

$$\boldsymbol{r}_t \;=\; \mathbf{1} \otimes \boldsymbol{r}_1 \;=\; \begin{bmatrix} \boldsymbol{r}_1 \\ \vdots \\ \boldsymbol{r}_1 \end{bmatrix} \quad \text{and} \quad \Phi_t \;=\; I \otimes \Phi_1 \;=\; \begin{bmatrix} \Phi_1 & & \\ & \ddots & \\ & & \Phi_1 \end{bmatrix}, \tag{19}$$

hence $\boldsymbol{r}_t$ is a $5t \times 1$ vector and Φ_t is a $5t \times t$ matrix. A policy over $5t$ actions can then be parameterized by a t dimensional weight vector $\boldsymbol{w}$ via

$$\boldsymbol{\pi} \;=\; \boldsymbol{f}(\Phi_t \boldsymbol{w}) \;=\; \frac{1}{d(\boldsymbol{w})} \begin{bmatrix} e^{2w_1} \\ e^{w_1} \\ 1 \\ e^{-w_1} \\ e^{-2w_1} \\ \\ \vdots \\ \\ e^{2w_t} \\ e^{w_t} \\ 1 \\ e^{-w_t} \\ e^{-2w_t} \end{bmatrix}, \quad \text{where} \quad d(\boldsymbol{w}) \;=\; \sum_{i=1}^{t} d(w_i). \tag{20}$$

The value function in this case then becomes

$$v(\boldsymbol{w}) \;=\; \boldsymbol{r}_t^\top \boldsymbol{\pi} \;=\; \frac{\sum_{i=1}^{t} n(w_i)}{\sum_{i=1}^{t} d(w_i)}, \quad \text{where} \quad n(\boldsymbol{w}) = \sum_{i=1}^{t} n(w_i). \tag{21}$$

To determine the locations of the critical points, consider the partial derivative of v with respect to a single parameter, say w_i:

$$\frac{\partial v}{\partial w_i} \;=\; \frac{n'(w_i) d(\boldsymbol{w}) - d'(w_i) n(\boldsymbol{w})}{d(\boldsymbol{w})^2} \;=\; \frac{n'(w_i) - d'(w_i) v(\boldsymbol{w})}{d(\boldsymbol{w})}. \tag{22}$$

As before, since $d(w_i) \geq 1$ for all i, hence $d(\boldsymbol{w}) \geq t$, we know the zeros of $\frac{\partial v}{\partial w_i}$ are determined by w_i such that $n'(w_i) = d'(w_i) v(\boldsymbol{w})$. One root value for w_i will always be $w_i = 0$, regardless of the other values of w_j, $j \neq i$, since the individual numerator and denominator functions each satisfy $n'(0) = d'(0) = 0$ respectively. It remains only to show that there are always two other roots for w_i, symmetrically placed around but separated from 0, regardless of the values for the other w_j, $j \neq i$.

A few useful properties of $v(\boldsymbol{w})$ will allow us to show this. First, since the individual $n(w_i)$ and $d(w_i)$ functions are even, the function $v(\boldsymbol{w})$ must also be even along any coordinate w_i. Second, since $n(w_i) \geq d(w_i)$ for all i, and moreover $n(w_i) > d(w_i)$ if $w_i \neq 0$, we have $v(\boldsymbol{w}) > 1$ if $w_i \neq 0$. Third, since it always holds that $n(w_i) < 2d(w_i)$, we also have $\sum_i n(w_i) < 2 \sum_i d(w_i)$, hence $v(\boldsymbol{w}) < 2$. Therefore, even though the value of $v(\boldsymbol{w})$ will determine the exact location of the symmetric nonzero roots in (22), we can establish the existence of these nonzero roots simply by assuming $v(\boldsymbol{w})$ takes on any arbitrary value in the range $1 < v < 2$, as we now show.

Consider the zeros of $n'(w_i) - d'(w_i) v$ where v is any quantity such that $1 < v < 2$. From the definitions we know that $n'(w_i) = 4\sinh(2w_i) + 4\sinh(w_i)$ and $d'(w_i) = 4\sinh(2w_i) + 2\sinh(w_i)$, hence we seek the values of w_i such that

$$2(1 - v)\sinh(2w_i) \;=\; (v - 2)\sinh(w_i) \tag{23}$$

As noted, one solution is $w_i = 0$ but we particularly seek the nonzero roots, so consider $w_i \neq 0$, hence $\sinh(w_i) \neq 0$. Under this assumption (23) reduces to

$$4(1 - v)\cosh(w_i) \;=\; v - 2, \tag{24}$$

which has a putative solution pair $w_i = \pm \operatorname{acosh}(\frac{2-v}{4(v-1)})$. This solution pair exists (and is nonzero) if $1 < \frac{2-v}{4(v-1)} < \infty$, which is guaranteed by $1 < v < 2$.

To summarize: in characterizing the landscape of $v(\boldsymbol{w})$, we know the function is continuous, smooth, and sandwiched between $1 \leq v(\boldsymbol{w}) < 2$ for all $\boldsymbol{w}$. Along each coordinate axis, w_i, regardless of the values of the other weight parameters, w_j, $j \neq i$, there are exactly three critical points: one at zero, and two others symmetrically placed around but distinct from zero (attaining equal value, since v is an even function along each coordinate). Since the point $w_i = 0$ is a local minimum along the coordinate, the other two critical points must be local maxima. The overall weight vector $\boldsymbol{w}$ is only at a critical point if each of its coordinates are at a critical point. Therefore, in total there are 3^t critical points, of which 2^t are local maxima (i.e. each coordinate is at a local maximum).

■

Comment Clearly, the above construction creates 2^t local maxima that all have the same expected value. Intuitively, a small perturbation of one of the modes in the initial construction can preserve the number of local maxima while elevating a single such maximum to global dominance.

Proposition 2 *Let* $\tilde{\boldsymbol{\pi}}_\tau = \arg\min_{\boldsymbol{\pi}\in\mathcal{P}} \mathcal{S}_\tau(\boldsymbol{\pi})$. *Then* $\tilde{\boldsymbol{\pi}}_\tau(x) = \exp(\mathbb{E}[\boldsymbol{r}|x] - F(\mathbb{E}[\boldsymbol{r}|x])/\tau)$ *and* $\mathcal{R}(\tilde{\boldsymbol{\pi}}_\tau) < \mathcal{R}^* + \tau \log K$. *Hence for any* $\epsilon > 0$ *setting* $\tau < \epsilon/\log K$ *ensures* $\mathcal{R}(\tilde{\boldsymbol{\pi}}_\tau) < \mathcal{R}^* + \epsilon$.

Proof: Let $\Delta(\boldsymbol{z})$ denote putting a vector $\boldsymbol{z}$ on the main diagonal of a square matrix. First, it is easy to prove that $\tilde{\boldsymbol{\pi}}_\tau(x) = \exp(\mathbb{E}[\boldsymbol{r}|x] - F(\mathbb{E}[\boldsymbol{r}|x])/\tau)$ is optimal. Consider a fixed x and note:

$$\frac{d\mathcal{S}_\tau(\boldsymbol{\pi}, \boldsymbol{r}, x)}{d\boldsymbol{q}(x)} = \left(\Delta(\boldsymbol{\pi}(x)) - \boldsymbol{\pi}(x)\boldsymbol{\pi}(x)^\top\right)(\tau\boldsymbol{q} - \boldsymbol{r}), \tag{25}$$

thus, $\boldsymbol{q}(x) = \boldsymbol{r}/\tau - \mathbf{1}v/\tau$ determines an equilibrium point in $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x)$ for any constant v. Since (25) is linear in $\boldsymbol{r}$, taking an expectation in $\boldsymbol{r}$ still yields equilibria of the form $\boldsymbol{q}(x) = \mathbb{E}\left[\boldsymbol{r}/\tau|x\right]$ setting $v = 0$. Thus, the optimal policy conditioned on x can be written as $\tilde{\boldsymbol{\pi}}_\tau(x) = \exp\left(\mathbb{E}[\boldsymbol{r}|x]/\tau - F(\mathbb{E}[\boldsymbol{r}|x]/\tau)\right) = \exp\left(\bar{\boldsymbol{r}}/\tau - F(\bar{\boldsymbol{r}}/\tau)\right)$, where we let $\bar{\boldsymbol{r}} = \mathbb{E}[\boldsymbol{r}|x]$.

For the second part of the claim, for any fixed x and $\boldsymbol{r}$ we consider the gap between the exact optimum and the approximate optimum produced by $\tilde{\boldsymbol{\pi}}_\tau(x) = \boldsymbol{f}(\boldsymbol{q}(x)) = \boldsymbol{f}(\frac{\bar{\boldsymbol{r}}}{\tau})$:

$$\text{Gap} = \max_a r_a - \boldsymbol{f}(\tfrac{\bar{\boldsymbol{r}}}{\tau}) \cdot \bar{\boldsymbol{r}}. \tag{26}$$

We can bound this gap by lower bounding the expected reward achieved by the policy at x:

$$\begin{aligned}
\boldsymbol{f}(\tfrac{\bar{\boldsymbol{r}}}{\tau}) \cdot \bar{\boldsymbol{r}} &= \tau \boldsymbol{f}(\tfrac{\bar{\boldsymbol{r}}}{\tau}) \cdot \tfrac{\bar{\boldsymbol{r}}}{\tau} && (27)\\
&= \tau F(\tfrac{\bar{\boldsymbol{r}}}{\tau}) + \tau F^*(\boldsymbol{f}(\tfrac{\bar{\boldsymbol{r}}}{\tau})) && (28)\\
&\geq \max_a r_a - \tau \log K, && (29)
\end{aligned}$$

where the second step uses the fact that the Young-Fenchel inequality is tight at a dual pair $\tilde{\boldsymbol{\pi}}$ and $\boldsymbol{q}$ [2, §3.3.2], and the last step using the fact that F^* is negative entropy and the maximum entropy of any distribution over K actions is $\log K$. From this it is easy to conclude that whenever $\tau \leq \epsilon/\log K$ we must have Gap $\leq \epsilon$. The result then follows by noting that this inequality holds pointwise for all x, hence also in expectation over x. ■

Theorem 3 *For an arbitrary baseline* v *and* $\tau > 0$*, let*

$$L(\boldsymbol{q}, \boldsymbol{r}, x) = \tau D_F\left(\boldsymbol{q}(x) + \tfrac{v}{\tau} \,\middle\|\, \tfrac{\boldsymbol{r}}{\tau}\right) + \tfrac{\tau}{4}\left\|\boldsymbol{q}(x) - \tfrac{\boldsymbol{r}-v}{\tau}\right\|^2, \tag{30}$$

Then, for any fixed v*,* L *is strongly convex in* $\boldsymbol{q}$ *and calibrated with respect to the smoothed (shifted) risk* $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x) = \mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) - v$ *with calibration function* $\delta(\epsilon, x) = \epsilon\ \forall x$.

Proof: Strong convexity is immediate from the inclusion of the squared loss. We need to establish two additional properties. First, that the global minimizer of (30) is also a global minimizer of the local smoothed risk $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x)$ (7). Second, that the surrogate objective is an *upper bound* on the suboptimality of the local smoothed risk.

For the equilibrium condition, note that $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x)$ must satisfy (25), hence, again, we have $\boldsymbol{q}(x) = \boldsymbol{r}/\tau - \mathbf{1}v/\tau$ is an equilibrium point for any fixed v. By comparison, taking the gradient of the surrogate with respect to $\boldsymbol{q}(x)$ yields

$$\frac{dL(\boldsymbol{q}, \boldsymbol{r}, x)}{d\boldsymbol{q}(x)} = \tau(\tilde{\boldsymbol{\pi}}(x) - \boldsymbol{p}(x)) + \tfrac{\tau}{2}\left(\boldsymbol{q}(x) - \tfrac{\boldsymbol{r}}{\tau} + \tfrac{v}{\tau}\right), \tag{31}$$

where $\tilde{\boldsymbol{\pi}}(x) = \boldsymbol{f}(\boldsymbol{q}(x) + \frac{v}{\tau})$ for any fixed v and $\boldsymbol{p} = \boldsymbol{f}(\frac{\boldsymbol{r}}{\tau})$. Thus, an equilibrium point for $L(\boldsymbol{q}, \boldsymbol{r}, x)$ is also given by $\boldsymbol{q}(x) = \frac{\boldsymbol{r}}{\tau} - \frac{v}{\tau}$. Moreover, any such point must be a *unique* global minimizer for $L(\boldsymbol{q}, \boldsymbol{r}, x)$ by strong convexity. Since this choice of $\boldsymbol{q}(x)$ uniquely determines $\boldsymbol{\pi}(x)$, it characterizes the equilibria of $\mathcal{S}(\boldsymbol{\pi}, \boldsymbol{r} - v, x)$ and therefore also the global minimizer.

For the second part, first let $\mathcal{S}^*(\boldsymbol{r} - v, x) = \inf_{\boldsymbol{q}\in\mathcal{Q}} \mathcal{S}(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x)$, and note that since we know

$$\begin{aligned}
\mathcal{S}(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r} - v, x) &= -\tau F(\tfrac{\boldsymbol{r}-v}{\tau}) + \tau D_F\left(\tfrac{\boldsymbol{r}-v}{\tau} \,\middle\|\, \boldsymbol{q}(x)\right) && (32)\\
&= v - \tau F(\tfrac{\boldsymbol{r}}{\tau}) + \tau D_F\left(\tfrac{\boldsymbol{r}-v}{\tau} \,\middle\|\, \boldsymbol{q}(x)\right), && (33)
\end{aligned}$$

it follows that $\mathcal{S}^*(\boldsymbol{r}-v,x) = v - \tau F(\frac{\boldsymbol{r}}{\tau})$, which is achieved at $\boldsymbol{q} = \boldsymbol{r}/\tau - \mathbf{1}v/\tau$. Finally, to establish that $L(\boldsymbol{q},\boldsymbol{r},x) \geq \mathcal{S}(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r}-v,x) - \mathcal{S}^*(\boldsymbol{r}-v,x)$ we consider a second order Taylor analysis along the lines of [6], which uses two Taylor expansions of $F(\boldsymbol{q})$. Using the same derivation, it can be shown that

$$\begin{aligned} D_F\left(\tfrac{\boldsymbol{r}}{\tau}\middle\|\boldsymbol{q}(x)+\tfrac{v}{\tau}\right) &= D_F\left(\boldsymbol{q}(x)+\tfrac{v}{\tau}\middle\|\tfrac{\boldsymbol{r}}{\tau}\right) \\ &\quad +\frac{1}{4}\left(\boldsymbol{q}(x)-\tfrac{\boldsymbol{r}}{\tau}+\tfrac{v}{\tau}\right)^\top\left(H_F(\boldsymbol{b})-H_F(\boldsymbol{a})\right)\left(\boldsymbol{q}(x)-\tfrac{\boldsymbol{r}}{\tau}+\tfrac{v}{\tau}\right), \quad (34) \end{aligned}$$

where H_F denotes the Hessian of F, $\boldsymbol{a} = (1-\frac{\eta}{2})\frac{\boldsymbol{r}}{\tau} + \frac{\eta}{2}(\boldsymbol{q}(x)+\frac{v}{\tau})$ for some $0 \leq \eta \leq 1$, and $\boldsymbol{b} = (1-\frac{\rho}{2})(\boldsymbol{q}(x)+\frac{v}{\tau}) + \frac{\rho}{2}\frac{\boldsymbol{r}}{\tau}$ for some $0 \leq \rho \leq 1$. Since the Hessian has the form

$$H_F(\boldsymbol{a}) = \Delta(\boldsymbol{f}(\boldsymbol{a})) - \boldsymbol{f}(\boldsymbol{a})\boldsymbol{f}(\boldsymbol{a})^\top \quad (35)$$

for all $\boldsymbol{a}$, we know that $I \succeq H_F(\boldsymbol{a}) \succeq 0$ and $I \succeq H_F(\boldsymbol{b}) \succeq 0$, hence $I \succeq H_F(\boldsymbol{b}) - H_F(\boldsymbol{a})$. Therefore, from (34) it follows that

$$D_F\left(\tfrac{\boldsymbol{r}}{\tau}\middle\|\boldsymbol{q}(x)+\tfrac{v}{\tau}\right) \leq D_F\left(\boldsymbol{q}(x)+\tfrac{v}{\tau}\middle\|\tfrac{\boldsymbol{r}}{\tau}\right) + \tfrac{1}{4}\left\|\boldsymbol{q}(x)-\tfrac{\boldsymbol{r}}{\tau}+\tfrac{v}{\tau}\right\|^2. \quad (36)$$

Therefore,

$$\begin{aligned} \mathcal{S}(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r}-v,x) &= \tau D_F\left(\tfrac{\boldsymbol{r}}{\tau}\middle\|\boldsymbol{q}(x)+\tfrac{v}{\tau}\right) + v - \tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) & (37)\\ &\leq \tau D_F\left(\boldsymbol{q}(x)+\tfrac{v}{\tau}\middle\|\tfrac{\boldsymbol{r}}{\tau}\right) + \tfrac{\tau}{4}\left\|\boldsymbol{q}(x)-\tfrac{\boldsymbol{r}}{\tau}+\tfrac{v}{\tau}\right\|^2 + v - \tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) & (38)\\ &= L(\boldsymbol{q},\boldsymbol{r},c) + v - \tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right). & (39) \end{aligned}$$

We conclude that if $L(\boldsymbol{q},\boldsymbol{r},x) \leq \epsilon$ then $\mathcal{S}(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r}-v,x) - \mathcal{S}^*(\boldsymbol{r}-v,x) \leq \epsilon$ and the result follows. ■

3 Proofs for Section 3: Batch Contextual Bandits

Note that throughout this section, as in the main body of the paper, we use $\hat{\boldsymbol{r}}(x)$ to denote the imputed reward estimator

$$\hat{\boldsymbol{r}}(x) = \tau\boldsymbol{q}(x) + \mathbf{1}_a\lambda(x,a)(r_a - \tau\boldsymbol{q}(x)_a). \quad (40)$$

This simplified notation allows us to simply write $\hat{\boldsymbol{r}}$ in place of $\boldsymbol{r}$ in the expressions below. However, this notation also masks the dependence of $\hat{\boldsymbol{r}}(x)$ on the model output $\boldsymbol{q}$ and the observation (x,a,r_a). That is, to be more explicit, the full dependence of $\hat{\boldsymbol{r}}$ can be fully expressed as $\hat{\boldsymbol{r}}(x,a,r_a,\boldsymbol{q}(x))$.

Proposition 4 *For any $\boldsymbol{q}$, $\tau > 0$ and observation (x,a,r_a): $\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)\right) = \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)$.*

Proof: By Lemma 13 below we have $\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x) = -\tau F\left(\frac{\hat{\boldsymbol{r}}}{\tau}\right) + \tau D_F\left(\frac{\hat{\boldsymbol{r}}}{\tau}\|\boldsymbol{q}(x)\right)$. By Lemma 14 below we also know $\mathcal{S}^*_\tau(\hat{\boldsymbol{r}},x) = -\tau F\left(\frac{\hat{\boldsymbol{r}}}{\tau}\right)$. Hence

$$\begin{aligned} \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x) &= \mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x) - \mathcal{S}^*_\tau(\hat{\boldsymbol{r}},x) & (41)\\ &= \left(-\tau F\left(\tfrac{\hat{\boldsymbol{r}}}{\tau}\right) + \tau D_F\left(\tfrac{\hat{\boldsymbol{r}}}{\tau}\|\boldsymbol{q}(x)\right)\right) - \left(-\tau F\left(\tfrac{\hat{\boldsymbol{r}}}{\tau}\right)\right) & (42)\\ &= \tau D_F\left(\tfrac{\hat{\boldsymbol{r}}}{\tau}\|\boldsymbol{q}(x)\right). & (43) \end{aligned}$$

■

Theorem 5 *For any model $\boldsymbol{q}$, $\tau > 0$, observation (x,a,r_a), and baseline v:*

$$L(\boldsymbol{q},\hat{\boldsymbol{r}},x) \geq \tau D_F\left(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)+\tfrac{v}{\tau}\right) = \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x) \geq 0. \quad (44)$$

Moreover, L is calibrated with respect to $\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}}-v,x)$ with calibration function $\delta(x,\epsilon) = \epsilon$.

Proof: The middle equality in (44) is established by Proposition 4 combined with the shift invariance of D_F established in Lemma 12 below. The last inequality in (44) follows immediately from the definition of $\mathcal{G}_\tau$. The first inequality in (44) follows from the definition $L(\boldsymbol{q},\hat{\boldsymbol{r}},x) =$

$\tau D_F\left(\boldsymbol{q}(x)+\frac{v}{\tau}\middle\|\frac{\hat{\boldsymbol{r}}}{\tau}\right)+\frac{\tau}{4}\left\|\boldsymbol{q}(x)-\frac{\hat{\boldsymbol{r}}-v}{\tau}\right\|^2$ combined with the inequality (36) established in the proof of Theorem 3.

Finally, note that L is also nonnegative, yet $L(\boldsymbol{q},\hat{\boldsymbol{r}},x)=0$ at $\boldsymbol{q}(x)=\frac{\hat{\boldsymbol{r}}-v}{\tau}$, which implies this is a global minimizer of L, which also must achieve suboptimality gap $\mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}}-v,x)=0$ since L dominates $\mathcal{G}_\tau$. Hence, any desired upper bound $\epsilon>0$ on the suboptimality $\mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}}-v,x)$ is achieved by finding a $\boldsymbol{q}$ such that $L(\boldsymbol{q},\hat{\boldsymbol{r}},x)\le\epsilon$. ■

Theorem 6 *For any model $\boldsymbol{q}$, any $\hat{\boldsymbol{r}}$ such that $\mathbb{E}[\hat{\boldsymbol{r}}|x]=\mathbb{E}[\boldsymbol{r}|x]$, and any baseline v:*

$$\mathbb{E}[L(\boldsymbol{q},\hat{\boldsymbol{r}},x)] \;\ge\; \mathbb{E}\left[\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)+\frac{v}{\tau}\right)\right] \;\ge\; \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q}) \;\ge\; 0. \tag{45}$$

Proof: Assume a fixed $\boldsymbol{q}$, and note that $\hat{\boldsymbol{r}}(x)$ is a random vector derived from $\boldsymbol{q}$ and the sample $(x,a,r_a)\sim p(x,\boldsymbol{r})\beta(a|x)$ (i.e., a is independent of $\boldsymbol{r}$ given x). The last inequality in (45) is immediate from the definition of $\mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q})$. The first inequality in (45) is also immediate given Theorem 5, which establishes $L(\boldsymbol{q},\hat{\boldsymbol{r}},x)\ge\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)+\frac{v}{\tau}\right)$ pointwise for all observations (x,a,r_a). To establish the middle inequality in (45), first note that for every fixed x, the function $\mathcal{S}^*_\tau(\boldsymbol{r},x)=\inf_{\boldsymbol{q}\in\mathcal{Q}}S(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r},x)$ is a pointwise infemum of linear functions of $\boldsymbol{r}$, hence concave in $\boldsymbol{r}$ [2, §3.2.3]. Thus we obtain

$$\begin{aligned}
&\mathbb{E}\left[\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)+\frac{v}{\tau}\right)\right]\\
&= \mathbb{E}\left[\tau D_F\left(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\middle\|\boldsymbol{q}(x)\right)\right] \quad \text{by Lemma 12 below} && (46)\\
&= \mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)-\mathcal{S}^*_\tau(\hat{\boldsymbol{r}},x)\right] \quad \text{by Proposition 4} && (47)\\
&= \mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)\right]-\mathbb{E}\left[\inf_{\boldsymbol{q}\in\mathcal{Q}}\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)\right] && (48)\\
&\ge \mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)\right]-\inf_{\boldsymbol{q}\in\mathcal{Q}}\mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\hat{\boldsymbol{r}},x)\right] \quad \text{by Jensen's inequality} && (49)\\
&= \mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r},x)\right]-\inf_{\boldsymbol{q}\in\mathcal{Q}}\mathbb{E}\left[\mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q},\boldsymbol{r},x)\right] && (50)\\
&\qquad \text{by linearity of } \mathcal{S}_\tau \text{ with respect to } \boldsymbol{r} \text{ and unbiasedness of } \hat{\boldsymbol{r}}\\
&= \mathcal{S}_\tau(\boldsymbol{f}\circ\boldsymbol{q})-\mathcal{S}^*_\tau \;=\; \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{g}). && (51)
\end{aligned}$$

■

3.1 Concentration

To keep the technical presentation straightforward, we assume the domain X is a bounded subset of $\mathbb{R}^n$ for some n.

3.1.1 Well-behavedness conditions for concentration

For concentration to hold uniformly over a class of random variables, such as those defined by scalar-valued divergences between model outputs $\boldsymbol{q}(x)$ and estimated rewards $\hat{\boldsymbol{r}}(x)$, we need to impose a set of assumptions to ensure the needed quantities remain appropriately bounded. In particular, we need to assume the following about $p(x,\boldsymbol{r})$, β, $\hat{\boldsymbol{r}}$ and $\mathcal{H}$:

- There exist constants c_X and c_R such that $\|x\|_2\le c_X$ and $\|\boldsymbol{r}\|_\infty\le c_R$ for all $(x,\boldsymbol{r})$ in the support of $p(x,\boldsymbol{r})$.
- There exists a constant $\rho>0$ such that $\beta(a|x)\ge\rho$ for all $x\in X$ and $a\in A$.
- $\mathbb{E}[\hat{\boldsymbol{r}}(x)|x]=\mathbb{E}[\boldsymbol{r}|x]$ for all x; i.e., $\hat{\boldsymbol{r}}(x)$ is unbiased.
- Every $\boldsymbol{q}\in\mathcal{H}$ can be expressed as a composition of a bounded linear with a general bounded function; that is, $\mathcal{H}=\mathcal{W}\circ\mathcal{Z}$, where $\mathcal{W}=\{W:\|W\|_2\le c_W\}$ and $\mathcal{Z}=\{\boldsymbol{z}:\|\boldsymbol{z}(x)\|_2\le c_Z\ \forall x\in\text{support}(p(x,\boldsymbol{r}))\}$. This implies $\boldsymbol{q}$ can be expressed as $\boldsymbol{q}(x)=W\boldsymbol{z}(x)$. An example is a neural network with bounded weights; see Section 3.1.3. Let $c_{\mathcal{H}}=c_Wc_Z$.

We say that the collection $p(x,\boldsymbol{r})$, β, $\hat{\boldsymbol{r}}$ and $\mathcal{H}$ is "well behaved" if the above assumptions are satisfied.

The main consequence of these assumptions is that the Rademacher complexity of the class of random variables of interest will then exhibit reasonable contraction. In particular, we are interested in the scalar valued divergence $D_F\big(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big)$ obtained by a function $\boldsymbol{q}\in\mathcal{H}$ on a given sample (x,a,r_a). Consider the class of scalar-valued functions induced by composing the divergence of interest with a model $\boldsymbol{q}\in\mathcal{H}$:

$$\mathcal{F} \quad = \quad \left\{d_{a,\boldsymbol{r}} : d_{a,\boldsymbol{r}}(\boldsymbol{q}(x)) = D_F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big) \text{ where } \boldsymbol{q}\in\mathcal{H}\right\}, \tag{52}$$

where $a\in A$ and $\boldsymbol{r}\in\text{support}(p(x,\boldsymbol{r}))$. The Rademacher complexity of $\mathcal{F}$ can then be defined as

$$R_T(\mathcal{F}) \quad = \quad \frac{1}{T}\mathbb{E}\left[\sup_{\boldsymbol{q}\in\mathcal{H}}\sum_{i=1}^{T}\sigma_i d_{a_i,\boldsymbol{r}_i}(\boldsymbol{q}(x_i))\right], \tag{53}$$

where the σ_i are independent and uniformly distributed over $\{1,-1\}$ [1, 7].

The key to the well-behavedness conditions is that they allow us to establish in Lemma 11 below that there exists a constant $c_{\mathcal{F}}$ such that

$$R_T(\mathcal{F}) \quad \leq \quad \frac{c_{\mathcal{F}}}{\sqrt{T}}. \tag{54}$$

3.1.2 Main concentration results

Recall the definitions of the empirical surrogate loss and empirical divergence respectively

$$\hat{L}(\boldsymbol{q},\mathcal{D}) \quad = \quad \frac{1}{T}\sum_{(x_i,a_i,r_i,\beta_i)\in\mathcal{D}} L(\boldsymbol{q},\hat{\boldsymbol{r}},x_i) \tag{55}$$

$$\hat{D}(\boldsymbol{q},\mathcal{D}) \quad = \quad \frac{1}{T}\sum_{(x_i,a_i,r_i,\beta_i)\in\mathcal{D}} D_F\big(\tfrac{\hat{\boldsymbol{r}}(x_i)}{\tau}\big\|\boldsymbol{q}(x_i)\big). \tag{56}$$

Lemma 7 *Assume $\mathcal{H}$, β, $p(x,\boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then for any $\tau,\delta>0$ there exists a constant C such that with probability at least $1-\delta$:*

$$\mathbb{E}\left[D_F\left(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\right)\right] \quad \leq \quad \hat{D}_F(\boldsymbol{q},\mathcal{D}) + \tfrac{C}{\sqrt{T}} \quad \forall\boldsymbol{q}\in\mathcal{H}. \tag{57}$$

Proof: Assuming well-behavedness, by Lemma 9 below we know that there exists a constant c_D such that $c_D \geq D_F\big(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big)\geq 0$ for all $\boldsymbol{q}\in\mathcal{H}$ and $(x,\boldsymbol{r})$ in the support of $p(x,\boldsymbol{r})$. Using this fact, the bound [7, Theorem 26.5] can then be applied to show that with probability at least $1-\delta$, for all $\boldsymbol{q}\in\mathcal{H}$:

$$\mathbb{E}\left[D_F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big)\right] \quad \leq \quad \hat{D}_F(\boldsymbol{q},\mathcal{D}) + 2R_T(\mathcal{F}) + 4c_D\sqrt{\tfrac{2}{T}\log\tfrac{2}{\delta}}. \tag{58}$$

By Lemma 11 below we also know there exists a constant $c_{\mathcal{F}}$ such that $R_T(\mathcal{F})\leq\frac{c_{\mathcal{F}}}{\sqrt{T}}$, hence C can be chosen to be $2c_{\mathcal{F}}+4c_D\sqrt{2\log(2/\delta)}$. ■

Theorem 8 *Assume $\mathcal{H}$, β, $p(x,\boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then for any v and $\tau,\delta>0$, there exists a C such that with probability at least $1-\delta$: if $\hat{L}(\boldsymbol{q},\mathcal{D})<\frac{\tau C}{\sqrt{T}}$ for $\boldsymbol{q}\in\mathcal{H}$ then $\mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q})\leq\frac{2\tau C}{\sqrt{T}}$.*

Proof: By Theorem 5 we know $L(\boldsymbol{q},\hat{\boldsymbol{r}},x)\geq\tau D_F\big(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)+\frac{v}{\tau}\big)$ for any $\tau>0$, model $\boldsymbol{q}$, observation (x,a,r_a), and baseline v. Assuming well-behavedness, Lemma 7 above shows that for any $\tau,\delta>0$ there exists a constant C such that with probability at least $1-\delta$, for any $\boldsymbol{q}\in\mathcal{H}$:

$$\hat{L}(\boldsymbol{q},\mathcal{D}) \quad \geq \quad \tau\hat{D}(\boldsymbol{q},\mathcal{D}) \tag{59}$$

$$\geq \quad \mathbb{E}\left[\tau D_F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big)\right] - \tfrac{\tau C}{\sqrt{T}} \tag{60}$$

$$\geq \quad \mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q}) - \tfrac{\tau C}{\sqrt{T}}, \tag{61}$$

where the last inequality follows from Theorem 6. Assume there is a $\boldsymbol{q}\in\mathcal{H}$ that achieves $\hat{L}(\boldsymbol{q},\mathcal{D})\leq\frac{\tau C}{\sqrt{T}}$. Then by (61) it follows that, with probability at least $1-\delta$:

$$\mathcal{G}_\tau(\boldsymbol{f}\circ\boldsymbol{q}) \quad \leq \quad \hat{L}(\boldsymbol{q},\mathcal{D}) + \tfrac{\tau C}{\sqrt{T}} \quad \leq \quad \tfrac{2\tau C}{\sqrt{T}}. \tag{62}$$

■

Lemma 9 *Assume $\mathcal{H}$, β, $p(x,\boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then for any $\tau > 0$ there exists a constant c_D such that $c_D \geq D_F\big(\frac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big) \geq 0$ for all $a \in A$, $\boldsymbol{q} \in \mathcal{H}$ and $(x,\boldsymbol{r})$ in the support of $p(x,\boldsymbol{r})$.*

Proof: Nonnegativity is immediate. Fix $\tau > 0$, $a \in A$, and recall the definition:

$$D_F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big) = F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big) - F\big(\boldsymbol{q}(x)\big) - \boldsymbol{f}(\boldsymbol{q}(x)) \cdot \Big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau} - \boldsymbol{q}(x)\Big) \tag{63}$$

$$= F\left(\boldsymbol{q}(x) + \mathbf{1}_a \tfrac{\boldsymbol{r}_a/\tau - \boldsymbol{q}(x)_a}{\beta(a|x)}\right) - F\big(\boldsymbol{q}(x)\big) - \boldsymbol{f}(\boldsymbol{q}(x))_a \tfrac{\boldsymbol{r}_a/\tau - \boldsymbol{q}(x)_a}{\beta(a|x)}. \tag{64}$$

We bound each term. First note that for any $\boldsymbol{q} \in \mathbb{R}^K$ we have $|F(\boldsymbol{q})| \leq \|\boldsymbol{q}\| + \log K$ [2, §3.1.5], hence

$$|F(\boldsymbol{q}(x))| \leq c_{\mathcal{H}} + \log K \tag{65}$$

$$|F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big)| \leq \|\boldsymbol{q}(x)\| + \big|\tfrac{\boldsymbol{r}_a}{\tau\beta(a|x)}\big| + \big|\tfrac{\boldsymbol{q}(x)_a}{\beta(a|x)}\big| + \log K \tag{66}$$

$$\leq \Big(1 + \tfrac{1}{\rho}\Big) c_{\mathcal{H}} + \tfrac{c_R}{\tau\rho} + \log K \tag{67}$$

$$\big|\boldsymbol{f}(\boldsymbol{q}(x))_a\big(\tfrac{\boldsymbol{r}_a}{\tau\beta(a|x)} - \tfrac{\boldsymbol{q}(x)_a}{\beta(a|x)}\big)\big| \leq \big|\tfrac{\boldsymbol{r}_a}{\tau\beta(a|x)}\big| + \big|\tfrac{\boldsymbol{q}(x)_a}{\beta(a|x)}\big| \leq \tfrac{c_R}{\tau\rho} + \tfrac{c_{\mathcal{H}}}{\rho}. \tag{68}$$

Therefore

$$D_F\big(\tfrac{\hat{\boldsymbol{r}}(x)}{\tau}\big\|\boldsymbol{q}(x)\big) \leq 2\Big(1 + \tfrac{1}{\rho}\Big) c_{\mathcal{H}} + 2\tfrac{c_R}{\tau\rho} + 2\log K. \tag{69}$$

■

Lemma 10 *For $\tau > 0$, $\boldsymbol{\beta} \geq \rho$, any $a \in A$ and any $\boldsymbol{r} \in \mathrm{support}(p(x,\boldsymbol{r}))$, the mapping $d_{a,\boldsymbol{r}}(\boldsymbol{q}) = D_F\big(\frac{\hat{\boldsymbol{r}}}{\tau}\big\|\boldsymbol{q}\big)$ is Lipchitz continuous, with Lipschitz bound at most $2\big(1 + \frac{1}{\rho}\big)$.*

Proof: For any $a \in A$ and $\boldsymbol{r} \in \mathrm{support}(p(x,\boldsymbol{r}))$, expand the definition as in (64):

$$d_{a,\boldsymbol{r}}(\boldsymbol{q}) = F\left(\boldsymbol{q} + \mathbf{1}_a \tfrac{\boldsymbol{r}_a/\tau - \boldsymbol{q}_a}{\beta_a}\right) - F\big(\boldsymbol{q}\big) - \boldsymbol{f}(\boldsymbol{q})_a \tfrac{\boldsymbol{r}_a/\tau - \boldsymbol{q}_a}{\beta_a}. \tag{70}$$

Note that $\|\nabla F(\boldsymbol{q})\| = \|\boldsymbol{f}(\boldsymbol{q})\| \leq 1$ for all $\boldsymbol{q}$, hence $F(\boldsymbol{q})$ is 1-Lipschitz. A Lipschitz bound can then be formulated for each term in (70), since the mapping $\boldsymbol{q} \mapsto \boldsymbol{q} + \mathbf{1}_a \frac{r_a/\tau - q_a}{\beta_a}$ is $\big(1 + \frac{1}{\rho}\big)$-Lipschitz, and the mapping $\boldsymbol{q} \mapsto \boldsymbol{f}(\boldsymbol{q})_a \frac{r_a/\tau - q_a}{\beta_a}$ is $\frac{1}{\rho}$-Lipschitz. Therefore, $d_{a,\boldsymbol{r}}$ is $2\big(1 + \frac{1}{\rho}\big)$-Lipschitz. ■

Lemma 11 *Assume $\mathcal{H}$, β, $p(x,\boldsymbol{r})$ and $\hat{\boldsymbol{r}}$ are "well behaved". Then there exists a constant $c_{\mathcal{F}}$ such that*

$$R_T(\mathcal{F}) \leq \frac{c_{\mathcal{F}}}{\sqrt{T}}. \tag{71}$$

Proof: To bound the Rademacher complexity of $\mathcal{F}$, it is easier to first consider the Rademacher complexity of $\mathcal{H}$ using the definition for vector-valued functions developed in [5]; define

$$R_T(\mathcal{H}) = \frac{1}{T}\mathbb{E}\left[\sup_{\boldsymbol{q}\in\mathcal{H}} \sum_{i=1}^{T}\sum_{a=1}^{K} \sigma_{ia}\boldsymbol{q}(x_i)_a\right], \tag{72}$$

where the σ_{ij} are independent and uniformly distributed over $\{1, -1\}$ [5]. As noted above, $\mathcal{F}$ can then characterized as a composition of the mappings $d_{a,\boldsymbol{r}}(\boldsymbol{q})$ specified in (70) with $\boldsymbol{q} \in \mathcal{H}$. By Lemma 10, we know that each mapping $d_{a,\boldsymbol{r}}$ is Lipschitz continuous with Lipschitz bound at most $\ell_d \stackrel{\triangle}{=} 2\big(1 + \frac{1}{\rho}\big)$. Therefore, the result of [5, Corollary 4] can be applied to establish $R_T(\mathcal{F}) \leq \sqrt{2}\ell_d R_T(\mathcal{H})$.

Then, to bound the Rademacher complexity of $\mathcal{H}$, we exploit the assumed structure $\mathcal{H} = \mathcal{W} \circ \mathcal{Z}$. Here again the result of [5, §4.2] shows that if $\mathcal{H}$ consists of mappings of the form $\boldsymbol{q}(x) = W\boldsymbol{z}(x)$, with $\|W\|_2 \leq c_W$ and $\|\boldsymbol{z}(x)\|_2 \leq c_Z$ for all $x \in \mathrm{support}(p(x,\boldsymbol{r}))$, then $R_T(\mathcal{H}) \leq \frac{\sqrt{2K}\ell_d c_W c_Z}{\sqrt{T}}$. ■

3.1.3 Feedforward neural networks

The well-behavedness conditions are sufficiently general to allow neural network representations for $\boldsymbol{q}(x)$. For example, an m-layer feedforward neural network can be written as a composition of matrix multiplications and a nonlinear transfer:

$$\boldsymbol{q}(x) \quad = \quad W^{(m)} \circ \phi \circ W^{(m-1)} \circ \phi \cdots \circ \phi \circ W^{(1)} \circ x, \tag{73}$$

where $W^{(j)}$ are the parameter matrices and ϕ is a componentwise transfer with bias:

$$\phi(\boldsymbol{z}) \quad = \quad \begin{bmatrix} \phi(z_1) \\ \vdots \\ \phi(z_K) \\ 1 \end{bmatrix}. \tag{74}$$

Standard choices for ϕ, such as ReLU, sigmoid and tanh, are Lipschitz bounded. For example, the ReLU transfer $\phi(z) = z_+$ is 1-Lipschitz. This means that if the parameter matrices $W^{(j)}$ are also bounded, i.e., $\|W^{(j)}\|_2 \leq B_j$, then $\boldsymbol{q}(x)$ in (73) is itself Lipschitz continuous with Lipschitz constant $B = \prod_{j=1}^{m} B_j$. This can be proved using a straightforward induction [8], exploiting the bounding technique for linear functions in [4].

Consider the class of functions defined by a feedforward neural network (73) with bounded parameters

$$\mathcal{H} \quad = \quad \{\boldsymbol{q} : \boldsymbol{q}(x) = W^{(m)} \circ \phi \cdots \circ \phi \circ W^{(1)} \circ x,\ \|W^{(j)}\|_2 \leq B_j,\ \phi \text{ 1-Lipschitz}\}. \tag{75}$$

This class satisfies the well-behavedness conditions for $\mathcal{H}$ stated above, since any $\boldsymbol{q} \in \mathcal{H}$ can be written as $\boldsymbol{q}(x) = W^{(m)}\boldsymbol{z}(x)$ for a function $\boldsymbol{z}(x) = \phi \circ W^{(m-1)} \circ \phi \cdots \circ \phi W^{(1)} \circ x$. Then by construction we have $\|W^{(m)}\| \leq B_m \stackrel{\triangle}{=} c_W$ and $\|\boldsymbol{z}(x)\| \leq c_X \prod_{j=1}^{m-1} B_j \stackrel{\triangle}{=} c_Z$.

3.2 Additional Lemmas

Lemma 12 *For any $\boldsymbol{q}$, $\boldsymbol{r}$ and scalar v:*

$$D_F(\boldsymbol{q} + v \| \boldsymbol{r}) \quad = \quad D_F(\boldsymbol{q} \| \boldsymbol{r}) \tag{76}$$
$$D_F(\boldsymbol{r} \| \boldsymbol{q} + v) \quad = \quad D_F(\boldsymbol{r} \| \boldsymbol{q}); \tag{77}$$

that is, D_F is shift invariant in either argument.

Proof: First, recall that by the definitions of $\boldsymbol{f}$ and F we have

$$\log \boldsymbol{f}(\boldsymbol{q} + v) \quad = \quad \boldsymbol{q} + v - F(\boldsymbol{q} + v) \ = \ \boldsymbol{q} + v - F(\boldsymbol{q}) + v \ = \ \boldsymbol{q} - F(\boldsymbol{q}) \ = \ \log \boldsymbol{f}(\boldsymbol{q}) \tag{78}$$
$$\boldsymbol{f}(\boldsymbol{q} + v) \quad = \quad \boldsymbol{f}(\boldsymbol{q}) \tag{79}$$
$$F(\boldsymbol{q} + v) \quad = \quad \log \mathbf{1} \cdot e^{\boldsymbol{q}+v} \ = \ v + \log \mathbf{1} \cdot e^{\boldsymbol{q}} \ = \ F(\boldsymbol{q}) + v. \tag{80}$$

Therefore, for the second identity (77), these identities yield

$$D_F(\boldsymbol{r} \| \boldsymbol{q} + v) \quad = \quad F(\boldsymbol{r}) - \boldsymbol{r} \cdot \boldsymbol{f}(\boldsymbol{q} + v) + F^*(\boldsymbol{f}(\boldsymbol{q} + v)) \tag{81}$$
$$= \quad F(\boldsymbol{r}) - \boldsymbol{r} \cdot \boldsymbol{f}(\boldsymbol{q}) + F^*(\boldsymbol{f}(\boldsymbol{q})) \tag{82}$$
$$= \quad D_F(\boldsymbol{r} \| \boldsymbol{q}). \tag{83}$$

For the first identity (76), note that $\boldsymbol{f}(\boldsymbol{r})$ is a probability vector for any $\boldsymbol{r}$, hence

$$D_F(\boldsymbol{q} + v \| \boldsymbol{r}) \quad = \quad F(\boldsymbol{q} + v) - (\boldsymbol{q} + v) \cdot \boldsymbol{f}(\boldsymbol{r}) + F^*(\boldsymbol{f}(\boldsymbol{r})) \tag{84}$$
$$= \quad F(\boldsymbol{q}) + v - (\boldsymbol{q} \cdot \boldsymbol{f}(\boldsymbol{r}) + v) + F^*(\boldsymbol{f}(\boldsymbol{r})) \tag{85}$$
$$= \quad F(\boldsymbol{q}) - \boldsymbol{q} \cdot \boldsymbol{f}(\boldsymbol{r}) + F^*(\boldsymbol{f}(\boldsymbol{r})) \tag{86}$$
$$= \quad D_F(\boldsymbol{q} \| \boldsymbol{r}). \tag{87}$$

■

Lemma 13 *For any x, $\boldsymbol{r}$, $\boldsymbol{q}$ and $\tau > 0$: $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) = -\tau F\left(\frac{\boldsymbol{r}}{\tau}\right) + \tau D_F\left(\frac{\boldsymbol{r}}{\tau} \| \boldsymbol{q}(x)\right)$.*

Proof: Immedate from the defintions:

$$\begin{aligned}
\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) &= -\boldsymbol{f}(\boldsymbol{q}(x)) \cdot \boldsymbol{r} + \tau \boldsymbol{f}(\boldsymbol{q}(x)) \cdot \log \boldsymbol{f}(\boldsymbol{q}(x)) && (88)\\
&= -\boldsymbol{f}(\boldsymbol{q}(x)) \cdot \boldsymbol{r} + \tau F^*(\boldsymbol{f}(\boldsymbol{q}(x))) && (89)\\
&= -\tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) + \tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) - \boldsymbol{f}(\boldsymbol{q}(x)) \cdot \boldsymbol{r} + \tau F^*(\boldsymbol{f}(\boldsymbol{q}(x))) && (90)\\
&= -\tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) + \tau D_F\left(\tfrac{\boldsymbol{r}}{\tau} \| \boldsymbol{q}(x)\right). && (91)
\end{aligned}$$

■

Lemma 14 *For any x, $\boldsymbol{r}$ and $\tau > 0$:* $\inf_{\boldsymbol{q} \in \mathcal{Q}} \mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) = -\tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right)$.

Proof: By Lemma 13 we know that $\mathcal{S}_\tau(\boldsymbol{f} \circ \boldsymbol{q}, \boldsymbol{r}, x) = -\tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right) + \tau D_F\left(\tfrac{\boldsymbol{r}}{\tau} \| \boldsymbol{q}(x)\right)$. Since D_F is nonnegative, yet $D_F\left(\tfrac{\boldsymbol{r}}{\tau} \| \boldsymbol{q}(x)\right) = 0$ when $\boldsymbol{q}(x) = \tfrac{\boldsymbol{r}}{\tau}$, we know the lower bound value $-\tau F\left(\tfrac{\boldsymbol{r}}{\tau}\right)$ is achieved at this point. ■

4 Additional experimental details

4.1 Additional Experiment Details: MNIST

In the MNIST experiments we trained a conventional feedforward neural network with a single hidden layer of 512 units and ReLU nonlinearities at the hidden layer. The standard training set of 60K examples was partitioned into the first 55K examples for training and the last 5K for validation. The test set of 10K examples was only used to report the final test results after all hyperparamter tuning was completed on the validation data only. All objectives were trained using the stochastic gradient descent with classical momentum set to 0.9 (i.e. Momentum(0.9)) for 100 epochs.

The hyperparameters and values considered in these experiments were:

learning rate $\in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0, 2.0\}$,
temperature $\tau \in \{0.1, 0.2, 0.5, 1.0, 2.0\}$,
offset $v \in \{0.0, 0.1, 0.2, 0.5\}$,
batch size $\in \{10, 20, 50, 100, 200, 500, 1000\}$, and
combination weights: uniform in the ten value range 0.0 to 1.0 with 0.1 increments.

4.2 Additional Experiment Details: CIFAR-10

In all the CIFAR-10 experiments, we trained a Resnet-20 model with layer sizes $(3, 4, 6, 3)$ and filter sizes $(64, 64, 128, 256, 512)$ for 12000 (then 120000; see below) iterations using a TPU with batch size of 128 * 8 = 1024, which corresponds to approximately 49 iterations per epoch for 50000 training examples, or equivalently, 250 (then 2000; see below) epochs total for each run. We used a learning rate of 0.1 with the momentum optimizer with parameter 0.9 along with Nesterov acceleration, along with batch normalization with a decay of 0.9. We also rescaled the squared loss metric by a factor of 0.01 to help stabilize learning. For the expected reward objective, we chose a baseline across $(0, 0.05, 0.1, 0.15, 0.2, 0.4, 0.6, 0.8, 1.0)$. For the composite objective, we found the best surrogate combination using a 0.05 weight on the average of the squared error and reverse imputed kl combined with the 0.95 weight on the expected reward (without any baseline) uniformly across all the bandit feedback tasks.

Although 250 epochs is already substantial training, allowing some objectives to produce good results, to better understand the relative difficulty of the different optimization landscapes we conducted longer training runs of 2000 epochs to ensure convergence was reached by all methods. The results in Table 1 in the main body of the paper were taken from the longer runs to better approximate the training set up used by [3] on the same training data.

4.3 Additional Experiment Details: Criteo

There are 35 features used to describe the context and candidates actions on the Criteo counterfactual analysis dataset. Among them, 2 are continuous and the rest are discrete categorical features. We encode the discrete features using one-hot encoding, which results in a 84017-dimensional sparse feature vector for each context x. We then build linear models using different loss functions. A

weight vector $W \in \mathbb{R}^{84017}$ is learned for each loss. Different objectives are optimized using SGD with momentum of 0.9. The table below lists the hyper-parameters we tuned for different losses. The final set of hyper-parameters for each method is chosen according to the performance on the validation set.

Hyperparameters	Values	Methods
Learning rate	$[0.01, 0.05, 0.1, 0.5, 1.0, 5.0]$	All
Batch size	$[1000, 5000]$	All
τ	$[0.01, 0.05, 0.1, 0.5, 1.0, 5.0]$	$\left\|\boldsymbol{q}(\mathbf{x}) - \frac{r-\mathbf{v}}{\tau}\right\|^2$, $\mathbf{D}_{\mathbf{F}^*}(\boldsymbol{p}\|\boldsymbol{\pi})$, **Composite**
λ	$[0.0001, 0.001, 0.01, 0.1, 1.0]$	POEM
α weight of $\mathbf{D}_{\mathbf{F}^*}(\boldsymbol{p}\|\boldsymbol{\pi})$	$[0.001, 0.01, 0.1, 1.0]$	**Composite**

5 Additional experimental results on MNIST

We repeated the experiments on MNIST 10 times to improve significance and to examine learning performance in more detail. Training with the expected reward objective was prone to getting stuck on poor plateaus in the cost sensitive misclassification (full reward feeback) case, so for that objective we repeated the experiments 20 times.

Figure 2a and Figure 2b show the average learning curves (averaged over 10 runs, 20 runs for "expected"), in terms of test misclassification error, for the various objectives. We observe that the expected reward objective is very difficult to optimize, and often gets stuck on a plateau.

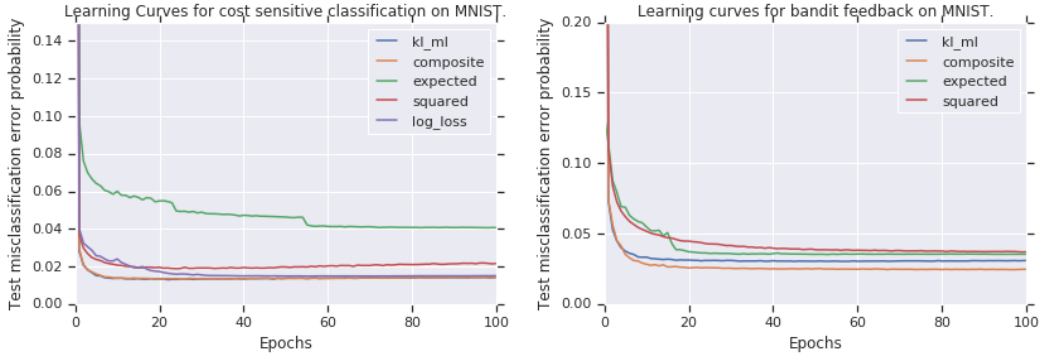

(a) Full reward feedback.

(b) Partial reward (bandit) feedback.

Figure 2: Learning curves (test misclassification error) on MNIST.

To gain a better assessment of the significance of the test misclassification results, Figure 3a and Figure 3b report the test misclassification error averaged over 10 runs (20 runs for "expected") with standard deviations illustrated. These results reinforce the observations made in the main body of the paper, except for the "expected" reward objective, which yielded poor results in the fully observed case. Note that the large error bar for training under expected reward in the cost sensitive classification setting (Figure 3a) is due to training getting stuck on a poor plateau in 6/20 runs. Removing these poor runs and recalculating the mean test misclassification error and standard deviation based on the remaining 14 runs yields the outcome given in Figure 4, which matches the findings in the main body of the paper.

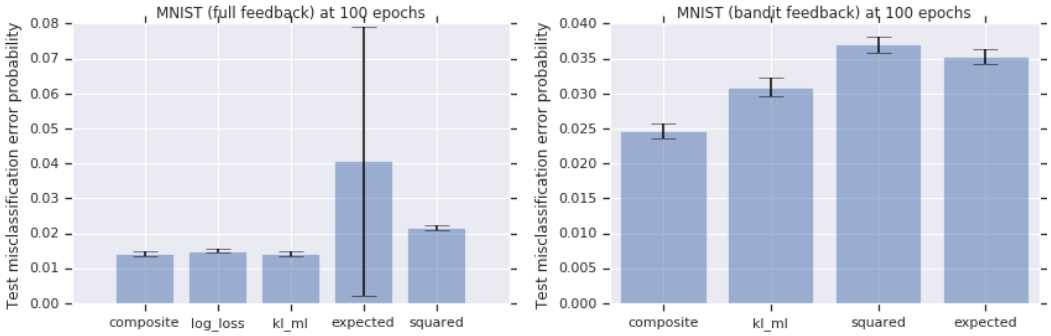

(a) Full reward feedback.

(b) Partial reward (bandit) feedback.

Figure 3: Test misclassification error on MNIST.

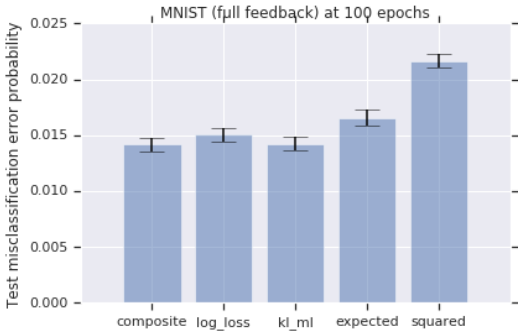

Figure 4: Test misclassification error on MNIST, full reward feedback, but for expected reward objective using 14/20 runs that escaped poor plateau.

6 Additional experimental results on CIFAR-10

We repeated the experiments on CIFAR-10 10 times to improve significance and to examine learning performance in more detail. Figure 5a and Figure 5b show the average learning curves in terms of test misclassification error, for the various objectives.

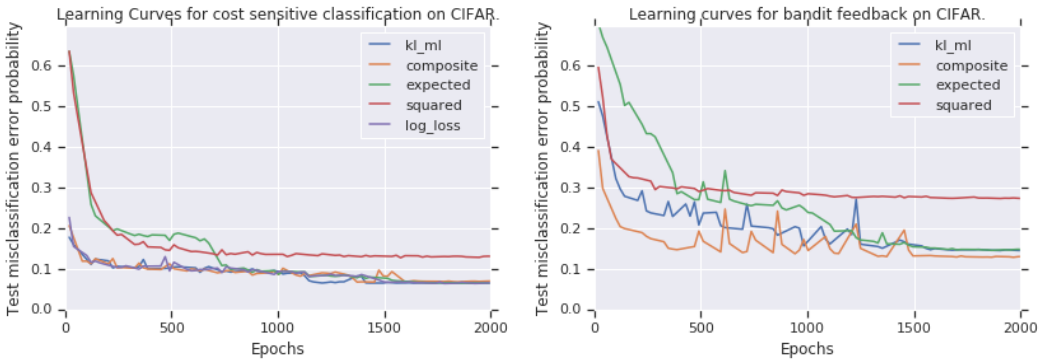

(a) Full reward feedback.

(b) Partial reward (bandit) feedback.

Figure 5: Learning curves (training misclassification error) on CIFAR-10.

As above, to gain a better assessment of the significance of the test misclassification results, Figure 6a and Figure 6b report the test misclassification error averaged over 10 runs with standard deviations illustrated. These results reinforce the observations made in the main body of the paper.

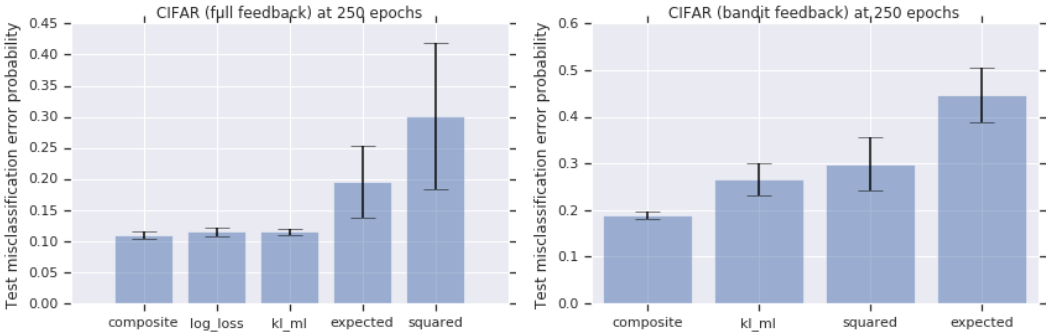

(a) Full reward feedback.

(b) Partial reward (bandit) feedback.

Figure 6: Test misclassification error on CIFAR-10.

However, as in the MNIST experiments, we find that after training for 250 epochs direct minimization of the empirical risk $\hat{\mathcal{R}}(\boldsymbol{\pi})$ is not competitive, yielding both high training and test error in both the fully observed and partially observed reward cases. To investigate whether this training difficulty was caused by plateaus that make it difficult to optimize this objective, we ran the experiments for significantly longer, for 2000 instead of 250 epochs.

In the fully observed case (i.e. cost-sensitive classification), direct empirical risk minimization is eventually able to catch up to the other objectives, achieving both small training and test misclassification error; see Figure 7a. Similarly, for the partially observed case (i.e. contextual bandit), we see a very similar phenonmenon, where direct optimization of empirical risk is able to close the performance gap with the other methods (but does not quite catch up); see Figure 7b. Thus, the hypothesis that the empirical risk objective $\hat{\mathcal{R}}(\boldsymbol{\pi})$ is indeed difficult to optimize, requiring extended training time and careful tuning to eventually reach competitive results.

We note that in both cases, the results in Figure 7a and Figure 7b significantly improve the results reported for resnet training on CIFAR-10 in [3], using a weaker exploration method in the contextual bandit case here. The main body of the paper also shows an improvement using the same logged data as [3].

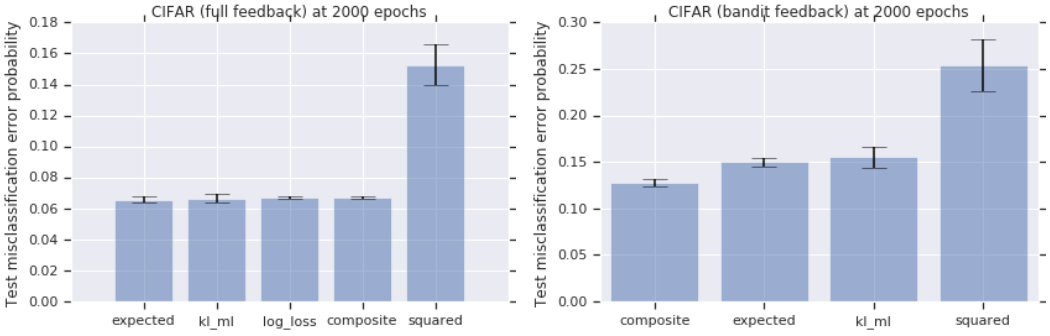

(a) Full reward feedback.

(b) Partial reward (bandit) feedback.

Figure 7: Misclassification error on CIFAR-10 data.

References

[1] Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.

[2] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge, 2004.

[3] Thorsten Joachims, Adith Swaminathan, and Maarten de Rijke. Deep learning with logged bandit feedback. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.

[4] Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in Neural Information Processing Systems 21*, pages 793–800, 2008.

[5] Andreas Maurer. A vector-contraction inequality for Rademacher complexities. In *International Conference on Algorithmic Learning Theory (ALT)*, pages 3–17, 2016.

[6] Mohammad Norouzi, Samy Bengio, Zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and Dale Schuurmans. Reward augmented maximum likelihood for neural structured prediction. In *Advances in Neural Information Processing Systems 29*, pages 1723–1731, 2016.

[7] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge, 2014.

[8] Yuchen Zhang, Jason D. Lee, Martin J. Wainwright, and Michael I. Jordan. On the learnability of fully-connected neural networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 83–91, 2017.