
Convergence of Uncertainty Sampling for Active Learning

Anant Raj^{1,2} Francis Bach¹

Abstract

Uncertainty sampling in active learning is heavily used in practice to reduce the annotation cost. However, there has been no wide consensus on the function to be used for uncertainty estimation in binary classification tasks and convergence guarantees of the corresponding active learning algorithms are not well understood. The situation is even more challenging for multi-category classification. In this work, we propose an efficient uncertainty estimator for binary classification which we also extend to multiples classes, and provide a non-asymptotic rate of convergence for our uncertainty sampling based active learning algorithm in both cases under no-noise conditions (*i.e.*, linearly separable data). We also extend our analysis to the noisy case and provide theoretical guarantees for our algorithm under the influence of noise in the tasks of binary and multi-class classifications.

1. Introduction

Over the last decade, machine learning algorithms have achieved a lot of success on various tasks in computer vision, natural language processing, and speech recognition. This success has been led by various factors which include improvement in computing architectures and improved machine learning algorithms. Moreover, the rapid growth in the number of large labeled public datasets is also one of the most important factors which contributed to the rise of machine learning. However, in many practical scenarios, labeled data are hard to obtain, as it requires a lot of time and human efforts to label a dataset. Hence, it is a time consuming as well as economically expensive procedure. For this reason, there have been many efforts to build machine learning algorithms which require a significantly lower number

of labeled samples to train. One major direction of research in this area is to devise efficient *active learning* algorithms.

Active learning algorithms propose efficient labeling schemes to reduce the number of labels required in order to train a classifier, resulting in minimal annotation cost while still maintaining high performance. Active learning methods can be categorized in two major categories : (i) stream-based active learning where samples from the data generating distribution are sequentially presented to the active learner, and (ii) pool-based active learning where there exists a very small number of labeled samples and the rest of the samples have no label. We note that any streaming based active learning algorithm can be converted to a pool-based active learning algorithm and vice versa (Sabato & Hess, 2016). For both of these categories, there exists an acquisition function which characterizes which informative samples should be labelled. The most popular way of determining if a sample is informative or not is by estimating uncertainty. Other than uncertainty sampling based approaches, the other major approaches for performing active learning are query-by-committee (Seung et al., 1992), expected model change (Settles et al., 2007), expected error reduction (Roy & McCallum, 2001), expected variance reduction (Wang et al., 2015), among others.

The main focus of this paper is *uncertainty sampling* based active learning algorithms. Despite of it being one of the most used active learning algorithms in practice, little is known about its theoretical properties. Specifically, there has been no common consensus on the uncertainty estimation approach used to perform active learning which also provably converges. Also, most of the active learning algorithms studied previously in theory as well as in practice are for binary classification and are not trivial to extend to multi-category classification problems at least theoretically. In this paper, we investigate these questions from the lens of optimization through stochastic gradient descent, and propose a sampling function for which the uncertainty sampling based active learning algorithm provably converges. We make the following contributions in this work :

- (i) We propose a family of theoretically motivated functions to estimate uncertainty for margin based predictions for binary and multi-category classification. Our approach can also be used with any non-linear feature

¹Inria, Ecole Normale Supérieure, PSL Research University, Paris, France ²Department of Electrical and Computer Engineering, Coordinated Science Laboratory University of Illinois at Urbana-Champaign, USA. Correspondence to: Anant Raj <anant.raj@inria.fr>.

representation (kernel methods, random features and neural network features) for classification.

- (ii) We show that the active learning algorithm based on the proposed sampling scheme converges to the optimal predictor in the separable regime for binary classification, and can be easily extended to multi-category classification while performing only first order update in every iteration ($O(d)$ computation cost, where d is the feature dimension) for both the cases. We provide a non-asymptotic rate of convergence of order $O(1/n)$, where n is number of iterations of the algorithms which is also number of unlabeled samples seen by the algorithm for both binary and multi-class classification problems.
- (iii) We extend our analysis to the inseparable regime for both cases (binary and multi-class) and show that the probability of mis-classification is bounded by $O(1/n) + O(\eta)$ where n is number of iteration performed by the algorithms and η is the noise parameter.
- (iv) We perform experimental evaluations for our algorithm on classification tasks.

1.1. Related Work

There has been a vast amount of work done in the field of active learning, and we only give a brief overview here. For binary classification, online active learning has been studied under the name of selective sampling under adversarial assumptions (Cesa-Bianchi et al., 2009; Dekel et al., 2010; Orabona & Cesa-Bianchi, 2011; Cavallanti et al., 2011). However, these methods are hard to extended to multi-class classification and are computationally expensive (e.g., with second order updates). Agarwal (2013) proposed a selective sampling scheme for multi-class classification for generalized linear models, but still each step of the algorithm is computationally expensive to perform. However, our proposed algorithm only performs the computation of $O(d)$ in each step while making stochastic gradient descent or online gradient descent like update, where d is the dimension of the feature vector. Settles (2012) provides an excellent survey of empirical studies in active learning.

Apart from empirical studies, there has been a lot of works in active learning on the theoretical front. Disagreement based active learning has been an active area of research in machine learning. The primary idea is as follows. A set of possible empirical risk minimizers are maintained with time and a label is queried if two minimizers disagree on the predicted label of that sample. An excellent survey of disagreement based active learning is provided by Hanneke (2014).

Our work can also be related to the vast line of work done in margin based active learning (Balcan et al., 2007; Das-

gupta et al., 2005; Balcan & Long, 2013; Wang & Singh, 2016). However for most of works in this area, the gain and convergence of algorithms have only been shown under strong distributional assumptions, while we do not make any such assumption on the data for our uncertainty sampling based active learning algorithm, and still manage to show a non-asymptotic rate of convergence for our algorithm.

Uncertainty sampling based machine learning algorithms have a long history. They were first proposed by Lewis & Gale (1994) who experimentally show that a probabilistic model with uncertainty sampling can improve the performance of text classification by up to 500 fold. Later, Schohn & Cohn (2000) applied uncertainty sampling to SVM classification and showed improved performance. Since then, it has been widely used for performing active learning (Yang et al., 2015; Zhu et al., 2008; Lughofer & Pratama, 2017; Yang & Loog, 2016; Wang et al., 2017). However, none of the above mentioned works focuses on the theoretical understanding of uncertainty sampling. Recently, Musmann & Liang (2018) showed that threshold based uncertainty sampling on a convex (e.g., logistic) loss can be interpreted as performing a preconditioned stochastic gradient step on the population zero-one loss. However, a proper convergence analysis was not proposed by Musmann & Liang (2018), in part due to the underlying non-convexity of their formulation.

2. Background

2.1. Uncertainty Sampling

Because of the ease of application, uncertainty sampling remains one of the most popular approaches used to perform active learning. Uncertainty sampling relies on the idea of querying the data point about which the current predictor is most uncertain. In simpler terms, uncertainty sampling usually identifies those points which are close to the decision boundary of the current model. However, the most important task here is to compute the uncertainty of prediction. There have been several approaches proposed to measure the uncertainty of a prediction. Here below, we discuss few of them that are widely used in practice (Monarch, 2021). Let us assume that a probabilistic model generates predictions in the form of probability distributions $p_\theta(\cdot|x)$ on \mathcal{Y} for $x \in \mathcal{X}$ and model parameter θ .

Margin of confidence sampling (Monarch, 2021; Nguyen et al., 2021). An intuitive way to estimate the uncertainty is by computing the margin in the confidence of top two predictions. Mathematically, sampling probability of querying a label $p_u(x, \theta) = \sigma(p_\theta(y_1^*|x) - p_\theta(y_2^*|x))$ where $\sigma : \mathbb{R} \rightarrow [0, 1]$, and y_1^* and y_2^* correspond to the two top most predictions for x given the model parameter θ .

Least confidence sampling (Monarch, 2021; Nguyen et al., 2021). Least confidence sampling considers the difference between 100% confidence and the most confident prediction to compute sampling probability of query a label. That means $p_u(x, \theta) \propto 1 - p_\theta(y_1^*|x)$ where y_1^* corresponds to the top most prediction for x given the model parameter θ .

Entropy-based sampling (Monarch, 2021; Nguyen et al., 2021). Entropy is an information theoretically motivated way to compute the uncertainty and widely used to estimate the uncertainty. Sampling probability of querying a label can be written as: $p_u(x, \theta) \propto \sum_{y \in \mathcal{Y}} p_\theta(y|x) \log p_\theta(y|x)$.

In this paper, we use the margin of confidence sampling scheme to estimate uncertainty in the prediction. Details of the sampling function σ will be provided in Section 3, where we discuss convergence of the algorithm. However, before discussing the theoretical results (Section 3), in the next section we discuss the relation between the hinge loss and corresponding test accuracy in binary and multi-class classification.

2.2. Max-margin linear classification

In this paper, we consider the simplest possible set-up of linear classification, with inputs $x \in \mathbb{R}^d$, and linear prediction functions. We note that by replacing x by some feature function $\Phi(x)$ we can deal with non-linear problems, the feature map being explicit or implicit through kernel methods (Hofmann et al., 2008).

Binary classification. With two classes, we consider $y \in \{-1, 1\}$ and a prediction function $x \mapsto \theta^\top x$ parameterized by $\theta \in \mathbb{R}^d$. We then classify according to the sign of $\theta^\top x$.

The associated error rate can be computed as

$$\mathbb{P}(y\theta^\top x \leq 0),$$

but is a non-convex function of θ . Among the many convex surrogates, we will consider the classical hinge loss:

$$\hat{\ell}(x, y, \theta) = \max\{0, 1 - y\theta^\top x\},$$

and its square, leading to the traditional support vector machine. The regular hinge loss is non differentiable, while the squared hinge loss is smooth in θ . In this paper, we will consider an algorithm based on the squared hinge loss, but obtain a guarantee for the non-squared one. Given that the two losses lead to guarantees on the misclassification error, as

$$\begin{aligned} \mathbb{P}(y\theta^\top x \leq 0) &\leq \mathbb{E}(\hat{\ell}(x, y, \theta)) \\ \mathbb{P}(y\theta^\top x \leq 0) &\leq \mathbb{E}(\hat{\ell}(x, y, \theta)^2), \end{aligned}$$

this allows to get the desired bounds.

Multi-class classification. Here, we consider $y \in \{1, \dots, k\}$. In multi-class classification the model parameter θ is a vector in \mathbb{R}^{dk} which consists of predictors $\theta(i) \in \mathbb{R}^d$ for all $i \in \{1, 2, \dots, k\}$. Hence, we denote model parameter θ as collection of k predictors, *i.e.*, $\theta = [\theta(1); \theta(2); \dots; \theta(k)]$. We consider the multi-class SVM formulation of Crammer & Singer (2001). Following the structured SVM notation from Tsochantaridis et al. (2005), let us assume that $\phi(x, y)$ represents the feature map corresponding to the sample pair (x, y) . In multi-class classification with k classes, $\phi(x, y) \in \mathbb{R}^{dk}$ consists of k blocks of d -dimensional vector and if we allow ourselves to denote each d -dimensional block with $\phi(x, y)(i)$ for $i \in \{1, 2, \dots, k\}$, then $\phi(x, y) = [\phi(x, y)(1); \phi(x, y)(2); \dots; \phi(x, y)(k)]$ where $\phi(x, y)(i) = 0$ for all $i \neq y$ and $\phi(x, y)(y) = x$. Define a loss function $\Delta : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$. The general hinge loss for maximum margin training function can be written as:

$$\hat{\ell}(x_t, y_t, \theta) = \max_{y \in \mathcal{Y}} [\Delta(y, y_n) - \theta^\top (\phi(x_n, y_n) - \phi(x_n, y))].$$

In the case of multi-class classification, generally $\Delta(y, y') = 1$ if $y \neq y'$, otherwise $\Delta(y, y') = 0$. Hence, the multi-class hinge loss can be written as,

$$\hat{\ell}(x, y, \theta) = \max [0, 1 - \theta^\top (\phi(x, y) - \phi(x, y^*(\theta, x, y)))] , \quad (1)$$

where $y^*(\theta, x, y) = \arg \max_{z \in \mathcal{Y} \setminus y} \theta^\top \phi(x, z)$.

Predicting a label for the data point $x \in \mathbb{R}^d$ by the predictor θ is done by computing $\arg \max_{z \in \mathcal{Y}} \theta^\top \phi(x, z)$. Similar to the binary case, the hinge loss for multiclass classification is non differentiable, while the square hinge loss is smooth and differentiable in the model parameter θ . For multi-class hinge loss as well, the misclassification error can be bounded by expected loss for both the losses. That is for a given sample pair (x, y) and model parameter θ , we have,

$$\begin{aligned} \mathbb{P}(\theta^\top \phi(x, y) - \theta^\top \phi(x, z^*) \leq 0) &\leq \mathbb{E}(\hat{\ell}(x, y, \theta)) \\ \mathbb{P}(\theta^\top \phi(x, y) - \theta^\top \phi(x, z^*) \leq 0) &\leq \mathbb{E}(\hat{\ell}(x, y, \theta)^2), \end{aligned}$$

where $z^* = \arg \max_{z \in \mathcal{Y}} \theta^\top \phi(x, z)$.

3. Convergent Uncertainty Sampling for Classification

In this work, we consider the streaming data setting. However, the algorithm can also be applied for non-streaming data setting. In the next three sections, we discuss the convergence results for binary and multi-class classification.

Let us consider n *i.i.d.* samples (x_i, y_i) jointly sampled from \mathcal{P} such that $x_i \in \mathbb{R}^d$, $i = 1, \dots, n$, and $y_i \in \mathcal{Y}$ where

Algorithm 1 Uncertainty Sampling in Binary Classification

Input: learning rate γ , Streaming (x_i, y_i) for $i \in [n]$, initial model parameter θ_1 , parameter μ and sampling function σ .

Output: average iterate $\bar{\theta}_{n+1}$.

for $t \leftarrow 1$ **to** n **do**

Compute probability $p_u(x_t, \theta_t) = \sigma(\theta_t, x_t)$.
 Sample Bernoulli random variable z_t with $p = p_u(x_t, \theta_t)$.
 Compute $\hat{\ell}_t(x_t, y_t, \theta_t) \leftarrow \max(0, 1 - y_t(\theta_t^\top x_t))$.
 Update $\theta_{t+1} \leftarrow \theta_t + \gamma z_t (y_t x_t \hat{\ell}_t(x_t, y_t, \theta_t))$.
 Update $\bar{\theta}_{t+1} \leftarrow \left(1 - \frac{1}{t+1}\right) \bar{\theta}_{t+1} + \frac{1}{t+1} \theta_{t+1}$.

$\mathcal{Y} = \{-1, 1\}$ for binary classification and $\mathcal{Y} = \{1, 2, \dots, k\}$ for multi-class classification.

Before going into the details of theoretical results, we first discuss the intuition of the uncertainty sampling algorithm. As discussed previously, the main idea behind the uncertainty sampling based active learning algorithm is to query the labels for those data points about which the predictor is most uncertain about. In the streaming data setting, every time we see a new data point, we compute its uncertainty by computing the prediction score. Then, we convert this score into the density function with the help of the given function σ , which maps prediction score to probability of querying a label. This probability score is used to generate a Bernoulli random variable which decides if the algorithm queries the label of the presented sample or not, and then performs a stochastic gradient step (for the squared hinge loss) if the label is accessed. The exact expression for the function σ is provided in Section 3.1 (for binary classification) and in Section 3.2 for multiclass classification. The pseudo-codes of the algorithms are presented in Algorithm 1 (binary classification) and Algorithm 2 (multi-class classification). It is important to note that in Algorithm 1 and in Algorithm 2, when the sampled Bernoulli random variable z_t is zero for any t , then the *for* loop breaks for that iteration on its own as the θ_t does not get updated.

3.1. Binary Classification

In the separable case for binary classification, we assume that there exists an optimal classifier $\theta_\star \in \mathbb{R}^d$ such that for all $x \in \mathcal{X}$ and its corresponding label $y \in \mathcal{Y}$ where $\mathcal{Y} = \{-1, 1\}$,

$$y(\theta_\star^\top x) \geq \rho^\star > 1. \quad (2)$$

The above assumption is a standard assumption made in the analysis of maximum margin classifier in the realizable case (Balcan et al., 2007; Dasgupta et al., 2005). We minimize the square hinge loss to obtain our predictor. The expression

Algorithm 2 Uncertainty Sampling in Multi-Class Classification

Input: learning rate γ , Streaming (x_i, y_i) for $i \in [n]$, initial model parameter θ_1 , parameter μ , number of classes k , and sampling function σ .

Output: average iterate $\bar{\theta}_{n+1}$.

for $t \leftarrow 1$ **to** n **do**

Compute score $s_t(j) = \theta(j)^\top x_t$ for all $j \in [k]$.
 Compute probability $p_u(x_t, \theta_t) = \sigma(\theta_t, x_t)$.
 Sample Bernoulli random variable z_t with $p = p_u(x_t, \theta_t)$.
 $y_t^\star \leftarrow \arg \max_{j \in \mathcal{Y} \setminus y_t} s_t(j)$
 $\hat{\ell}_t(x_t, y_t, \theta_t) \leftarrow \max(0, 1 - \theta_t(y_t)^\top x_t + \theta_t(y_t^\star)^\top x_t)$.
 Update $\theta_{t+1}(y_t) \leftarrow \theta_t(y_t) + \gamma z_t x_t \hat{\ell}_t(x_t, y_t, \theta_t)$.
 Update $\theta_{t+1}(y_t^\star) \leftarrow \theta_t(y_t^\star) - \gamma z_t x_t \hat{\ell}_t(x_t, y_t, \theta_t)$.
 Update $\theta_{t+1} \leftarrow [\theta_{t+1}(1); \theta_{t+1}(2); \dots; \theta_{t+1}(k)]$.
 Update $\bar{\theta}_{t+1} \leftarrow \left(1 - \frac{1}{t+1}\right) \bar{\theta}_{t+1} + \frac{1}{t+1} \theta_{t+1}$.

for square hinge loss can be written as (with an extra factor of $1/2$),

$$\ell(x, y, \theta) = \frac{1}{2} \max[0, 1 - y\theta^\top x]^2. \quad (3)$$

We have the following update rule to update the model parameter θ for uncertainty sampling based active learning :

$$\theta_{t+1} = \theta_t + \gamma z_t (y_t x_t) [1 - y_t(\theta_t^\top x_t)]_+, \quad (4)$$

where z_t is a Bernoulli random variable for fixed θ_t and x_t such that $p(z_t = 1 | x_t, \theta_t) = \sigma(\theta_t, x_t)$ where $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ is an even function.

When $\sigma = 1$, i.e., querying label for every sample, then this is exactly stochastic gradient descent for the squared hinge loss, which is known to converge with rate $O(1/t)$ in the separable situation (Vaswani et al., 2019). In the next result we show that the algorithm proposed in this paper (Algorithm 1) converges.

Theorem 3.1. Consider a set of n i.i.d. samples (x_i, y_i) from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{-1, 1\}$ for all $i = 1, \dots, n$ then, if there exists a θ_\star for which $y(\theta_\star^\top x) \geq \rho^\star$ for all (x, y) pair in \mathcal{P} , we have the following convergence guarantee for Algorithm 1:

$$\mathbb{E}(1 - y\bar{\theta}_n^\top x)_+ \leq \frac{R^2 \max\left\{1, \frac{1}{\mu}\right\} \|\theta_1 - \theta_\star\|^2}{\min\left\{\frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu}\right\}^2 n}, \quad (5)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu|\theta^\top x|}$, step size $\gamma = \frac{\min\left\{\frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu}\right\}}{R^2 \max\left\{1, \frac{1}{\mu}\right\}}$ and $\|x\| \leq R$ for all x in the domain \mathcal{X} .

Proof sketch. See the complete proof in the Appendix. Using the update given in Equation (4) and taking expectations only with respect to z_t considering x_t, y_t, θ_t fixed, we get

$$\mathbb{E} \|\theta_{t+1} - \theta_\star\|^2 = \|\theta_t - \theta_\star\|^2 + 2\gamma\sigma(\theta_t, x_t)[1 - y_t(\theta_t^\top x_t)]_+ y_t(\theta_t^\top x_t - y_t\theta_\star^\top x_t) + \gamma^2\sigma(\theta_t, x_t)R^2[1 - y_t(\theta_t^\top x_t)]_+^2,$$

where R is the upper bound on $\|x\|$ for all $x \in \mathcal{X}$. Our next goal is to find the function $\sigma(\theta, x)$ which satisfies the following properties for $y_t\theta_t^\top x_t < 1$,

$$\sigma(\theta_t, x_t)(1 - y_t\theta_t^\top x_t)_+ \leq c_1(1 - y_t\theta_t^\top x_t)_+ \quad (6)$$

$$\sigma(\theta_t, x_t)(y_t\theta_t^\top x_t - y_t\theta_\star^\top x_t) \leq -c_2, \quad (7)$$

for some positive constants c_1 and c_2 . We show in Lemma A.1 that choosing

$$\sigma(\theta, x) = \frac{1}{1 + \mu|\theta^\top x|}, \quad (8)$$

for $\mu > 0$ satisfies the conditions in Equation (6) and Equation (7) for $c_1 \geq \max\left\{1, \frac{1}{\mu}\right\}$ and $c_2 \leq \min\left\{\frac{\rho^\star - 1}{1 + \mu}, \frac{1}{\mu}\right\}$. Finally, after taking expectations, applying Jensen's inequality, and for the optimal choice of step size γ , we get

$$\mathbb{E}(1 - y\theta_t^\top x)_+ \leq \frac{R^2 \max\left\{1, \frac{1}{\mu}\right\} \|\theta_1 - \theta_\star\|^2}{\min\left\{\frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu}\right\}^2 n}.$$

□

On the mistake bound. As discussed in Section 2.2, the probability of misclassification is bounded by the expected classification loss (non squared hinge loss). Hence, for an independently sampled pair (x, y) ,

$$\mathbb{P}(y\theta_n^\top x \leq 0) \leq \frac{R^2 \max\left\{1, \frac{1}{\mu}\right\} \|\theta_1 - \theta_\star\|^2}{\min\left\{\frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu}\right\}^2 n}. \quad (9)$$

Discussion. It is important to note that the conditions mentioned in Equation (6) and Equation (7) are required only when $y_t\theta_t^\top x_t < 1$ as the gradient is zero when $y_t\theta_t^\top x_t \geq 1$, and hence the model parameter is not updated. Ideally, the label should not be queried when $y_t\theta_t^\top x_t \geq 1$, however there is no way to compute it beforehand. Hence, our sampling scheme provides a good trade-off for sampling. The value of μ should be decided by experimental evaluation as for $\mu \rightarrow \infty$, the upper bound for our algorithm diverges. It is also to keep in mind that the choice of μ depends on the sampling budget. Less budget will require to do more aggressive sampling and hence will require to choose much larger μ and vice versa.

Denoting the expected number of samples labelled in t steps as $\#_t$, we get:

$$\#_n = \sum_{t=0}^{n-1} \sigma(\theta_t, x_t) = \sum_{t=0}^{n-1} \frac{1}{1 + \mu|\theta_t^\top x_t|},$$

which can be significantly less than n if the absolute value of $\theta_t^\top x_t$ is large for most of t , i.e., most of the points are far from the decision boundary. Also, small margin (ρ^\star) will only affect the result of convergence up to a constant factor as the data can be re-normalized to get appropriate margin.

3.2. Extension to Multi-class Classification

In the separable multi-class case, we assume that there exists a set of optimal half-spaces $\theta_\star(i) \in \mathbb{R}^d$ for $i \in \{1, 2, \dots, k\}$ corresponding to each class such that for all $x \in \mathcal{X}$ and its corresponding label $y \in \mathcal{Y}$ where $\mathcal{Y} = \{1, 2, \dots, k\}$,

$$(\theta_\star(i) - \theta_\star(j))^\top x \geq \rho^\star > 1 \text{ for all } i \neq j \in [k]. \quad (10)$$

Note that the loss function given in Equation (1) is the same as that of used to compute the loss in Algorithm 2. We will optimize the multi-class square hinge loss and not the hinge loss for the reason discussed in Section 2.2. Let us introduce the following notation,

$$\delta_x(y, y') = \phi(x, y) - \phi(x, y').$$

Hence, the expression for the square hinge loss is

$$\ell(x, y, \theta) = \frac{1}{2} \hat{\ell}^2(x, y, \theta) = \frac{1}{2} [1 - \theta^\top \delta_x(y, y^\star(\theta, x, y))]_+^2.$$

The gradient of $\ell(x, y, \theta)$ with respect to θ can be written as,

$$\nabla \ell(x, y, \theta) = -\hat{\ell}(x, y, \theta) \delta_x(y, y^\star(\theta, x, y)).$$

We consider the projected stochastic gradient descent update to update the model parameter θ for uncertainty sampling based active learning in multi-class classification:

$$\theta_{t+1} = \Pi_{\|\theta\| \leq B} \left[\theta_t + \gamma z_t \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \hat{\ell}(x_t, y_t, \theta_t) \right], \quad (11)$$

where z_t is a Bernoulli random variable for fixed θ_t and x_t such that $p(z_t = 1 | x_t, \theta_t) = \sigma(\theta_t, x_t)$ where $\sigma : \mathbb{R} \rightarrow [0, 1]$ is an even function. $\Pi_{\|\theta\| \leq B}$ denotes the projection operator which projects θ_t for all t in the ball of radius B centered around the origin. In the theorem below, we show the convergence of the algorithm proposed in Algorithm 2.

Theorem 3.2. Consider a set of n i.i.d. samples (x_i, y_i) from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{1, 2, \dots, k\}$ for all $i = 1, \dots, n$. Then, if there exists a set of d -dimensional

optimal half-spaces $\theta_* = \{\theta_*(1), \theta_*(2), \dots, \theta_*(k)\}$ corresponding to each class for which $\theta_*^\top \delta_x(y, y^*(\theta_*, x, y)) \geq \rho^*$ for all (x, y) pair in \mathcal{P} , we have the following convergence guarantee for Algorithm 2 with the projected gradient descent update in Equation (11),

$$\mathbb{E} \hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2(1+BR)\|\theta_1 - \theta_*\|^2}{\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}^2 n}, \quad (12)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1+\mu|\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|}$, step size $\gamma = \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} \frac{1}{2R^2(1+BR)}$ and $\|\delta_x(i, j)\| \leq R$, for all x in the domain \mathcal{X} , and for all $i \neq j \in [k]$ and $\|\theta\| \leq B$.

Proof sketch. See the complete proof given in the Appendix. Using the update given in Equation (11) and taking expectations only with respect to z_t considering x_t, y_t, θ_t fixed, we get the following

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t)\hat{\ell}(x_t, y_t, \theta_t) \\ &\quad \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \theta^{*\top} \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))\right) \\ &\quad + \gamma^2\sigma(\theta_t, x_t)R^2\hat{\ell}^2(x_t, y_t, \theta_t), \end{aligned}$$

where R is the upper bound on $\|\delta_x(i, j)\|$ for all $x \in \mathcal{X}$ and $i, j \in \{1, 2, \dots, k\}$. Similar to the case of binary classification, we would need to find the function $\sigma(\theta, x)$ which satisfy the following properties for $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) < 1$,

$$\sigma(\theta_t, x_t)\hat{\ell}^2(x_t, y_t, \theta_t) \leq c_1\hat{\ell}(x_t, y_t, \theta_t) \quad (13)$$

$$\sigma(\theta_t, x_t) \left(\theta^{*\top} \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))\right) \geq c_2, \quad (14)$$

for some positive constants c_1 and c_2 . Let us now define top two predictions by the classifier θ for sample pair (x, y) are as follows,

$$y_{(1)}^* = \arg \max_{z \in \mathcal{Y}} \theta^\top \phi(x, z) \quad (15)$$

$$y_{(2)}^* = \arg \max_{z \in \mathcal{Y} \setminus y_{(1)}^*} \theta^\top \phi(x, z). \quad (16)$$

Then, we show in Lemma B.1 that choosing

$$\begin{aligned} \sigma(\theta_t, x_t) &= \frac{1}{1 + \mu \left| \theta_t^\top \delta_{x_t}(y_{t(1)}^*, y_{t(2)}^*) \right|} \\ &= \frac{1}{1 + \mu \left| \theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t \right|}, \end{aligned}$$

where μ is a positive constants, satisfies the conditions in Equation (13) and Equation (14) for $c_1 \geq (1+BR)$ and for

$c_2 \leq \frac{\rho^*-1}{1+\mu}$. Finally after taking expectations and applying Jensen's inequality, we get the following

$$\mathbb{E} \hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2(1+BR)(1+\mu)^2 \|\theta_0 - \theta_*\|^2}{n(\rho^*-1)^2},$$

for all $\mu \geq 0$. \square

On the mistake bound. Similar to the case of binary classification we discussed in Section 2.2, the probability of misclassification for the multiclass case is bounded by the expected classification loss (multi-hinge loss). Hence, for an independently sampled pair (x, y) and for $\mu \geq 0$,

$$\begin{aligned} \mathbb{P}(\bar{\theta}_n^\top \delta_x(y, y^*(\bar{\theta}_n, x, y)) \leq 0) &\leq \frac{R^2(1+BR)(1+\mu)^2 \|\theta_0 - \theta_*\|^2}{n(\rho^*-1)^2} \\ \Rightarrow \mathbb{P}(\bar{\theta}_n(y)^\top x - \bar{\theta}_n(y^*(\bar{\theta}_n, x, y))^\top x \leq 0) &\leq \frac{R^2(1+BR)(1+\mu)^2 \|\theta_0 - \theta_*\|^2}{n(\rho^*-1)^2}. \end{aligned}$$

Discussion. It is directly not clear from the bound how to choose μ to have a direct gain of applying active learning method. However similar to the binary classification case, the querying of a label is only required when $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) < 1$ as the gradient is 0 when $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \geq 1$. In that case, choosing larger μ will query less number of labels when $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \geq 1$, however, then it will also start to discard informative samples. A good μ should be chosen based on experimental evidence. Nevertheless, the algorithm converges for all choices of μ .

Denoting the expected number of samples labeled in t steps as $\#_t$, we get,

$$\#_n = \sum_{t=0}^{n-1} \sigma(\theta_t, x_t) = \sum_{t=0}^{n-1} \frac{1}{1 + \mu \left| \theta_t^\top \delta_{x_t}(y_{t(1)}^*, y_{t(2)}^*) \right|},$$

which can be significantly less than n if the absolute value of $\theta_t^\top \delta_{x_t}(y_{t(1)}^*, y_{t(2)}^*)$ is large for most of the time instance t , i.e., most of the points are far from the decision boundary. Comment that we made about small margin in binary classification remains valid in this case as well. It can also be noted that our analysis can easily be extended to structured output spaces under max-margin loss. The formulation for structured output spaces under max-margin loss function would resemble the multi-class formulation.

Discussion on extension for other loss functions. Our proof technique relies on choosing specific sampling function σ which can be related to the gradient term in SGD like update for square hinge loss (square max-margin loss). For other loss functions, we believe that it is possible to

choose sampling function σ and relate it to the (stochastic) gradient of the loss function, however, relating the gradient to the classification error is a non-trivial task for general loss function which is not the case for square max-margin loss.

3.3. Towards Uncertainty Sampling for Inseparable Data

In this section, we discuss uncertainty sampling for active learning in a more realistic scenario, that is, the inseparable case. We assume the existence of mild noise in the data. That means there exists a θ_* such that the following conditions about the classification noise hold in the case of binary and multi-class prediction problem respectively for a given sample pair (x, y) ,

$$\mathbb{P}(y\theta_*^\top x | (x, y) \leq \rho^*) \leq \eta \text{ (binary classification)}, \quad (17)$$

$$\mathbb{P}(\theta_*^\top \delta_x(y, y^*(\theta, x, y)) | (x, y) \leq \rho^*) \leq \eta \text{ (multi-class classification)}. \quad (18)$$

The assumption made about the noise in above equations are relatively stronger than noise conditions often assumed in the statistical learning theory literature (Massart & Nédélec, 2006; Tsybakov, 2004). However, extending our analysis under Tsybakov's noise condition (Tsybakov, 2004) is beyond the scope of this paper and could be considered in a subsequent work. Before moving to present the main results in the noisy case, we mention below the projected stochastic gradient descent update for the binary classification, the update looks as follows,

$$\theta_{t+1} = \Pi_{\|\theta\| \leq B} \left[\theta_t + \gamma z_t (y_t x_t) [1 - y_t (\theta_t^\top x_t)]_+ \right]. \quad (19)$$

Here below, we present our convergence result for inseparable data.

Theorem 3.3. *Consider a set of n i.i.d. samples (x_i, y_i) from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{-1, 1\}$ for all $i = 1, \dots, n$. Then under the assumption in Equation (17) for all (x, y) pair in \mathcal{P} and for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu |\theta^\top x|}$, $\mu > 0$, we have the following convergence guarantee for Algorithm 1:*

1. If the noise parameter

$$\eta < \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{\max \left\{ R \|\theta_*\|, \frac{1 + R \|\theta_*\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}$$

and iterates in Algorithm 1 are updated via stochastic gradient descent update in equation (4), then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R \|\theta_*\|}{1 + \mu}, R \|\theta_*\| \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$,

we have

$$\mathbb{E}(1 - y\bar{\theta}_n^\top x)_+ \leq \frac{R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \|\theta_1 - \theta_*\|^2}{\Gamma^2 n},$$

such that $\|x\| \leq R$ for all $x \in \mathcal{X}$ and where $\Gamma = \left[(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R \|\theta_*\|}{1 + \mu}, R \|\theta_*\| \right\} \right]$.

2. If the noise parameter η satisfies

$$\eta \geq \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{\max \left\{ R \|\theta_*\|, \frac{1 + R \|\theta_*\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}},$$

and iterates in Algorithm 1 are updated via projected stochastic gradient descent update in equation (19),

then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$, we have

$$\mathbb{E}(1 - y\bar{\theta}_n^\top x)_+ \leq \frac{\left(R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \right) \|\theta_1 - \theta_*\|^2}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}^2 n} + O(\eta),$$

such that $\|x\| \leq R$ for all $x \in \mathcal{X}$.

A similar result also holds for the case of multi-class classification. Consider the definition of the max-margin loss function for multiclass prediction as in equation (1), i.e.,

$$\hat{\ell}(x, y, \theta) = \max [0, 1 - \theta^\top (\phi(x, y) - \phi(x, y^*(\theta, x, y)))] ,$$

where $y^*(\theta, x, y) = \arg \max_{z \in \mathcal{Y} \setminus y} \theta^\top \phi(x, z)$.

Then we have the following convergence result.

Theorem 3.4. *Consider a set of n i.i.d. samples (x_i, y_i) from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{1, 2, \dots, k\}$ for all $i = 1, \dots, n$. Then under the assumption in Equation (18) for all (x, y) pair in \mathcal{P} and for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu |\theta(y_1^\top x - \theta(y_2^\top x))|}$, $\mu > 0$, we have the following convergence guarantee for Algorithm 2 with the projected stochastic gradient descent update (Equation (11)):*

1. If $\eta < \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{1 + \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} + R \|\theta_*\|}$, then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta (1 + B R)}{R^2 (1 + B R)}$

$$\mathbb{E} \hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2 (1 + B R) \|\theta_1 - \theta_*\|^2}{\Gamma^2 n},$$

where $\Gamma = \left[(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta (1 + R \|\theta_*\|) \right]$ and $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$.

2. If $\eta \geq \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{1 + \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} + R \|\theta_*\|}$, then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{R^2 (1 + B R)}$

$$\mathbb{E} \hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2 (1 + B R) \|\theta_1 - \theta_*\|^2}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}^2 n} + O(\eta),$$

such that $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$.

Discussion. In the above two results, we see that our uncertainty sampling algorithm is robust to small noise. The two cases of Theorems 3.3 and 3.4 are about existence of low and high noise in the data. When the noise level is small then it does not affect the rate of convergence as the other error term (optimization error) is dominant. However, when the noise is higher then errors due to noise are significant and that is the dominant term. Hence, in the case of significantly large noise, our algorithm has to pay for an extra error cost of the order $O(\eta)$.

4. Experiments

In this section, we perform an experimental evaluation for our proposed uncertainty sampling based active learning algorithm. The experiments are performed on both synthetic as well as real world data. For the experimental evaluation, we only consider binary classification setting. It is clear from the previous sections that we utilize the property of max-margin loss to obtain the convergence. Max-margin loss has similar properties for binary as well as multiclass classification. Hence, we omit the experimental evaluation for multiclass classification.

Synthetic Data (Binary Classification). We generate n number of data points $x_i \in \mathbb{R}^d$ for $i \in \{1, 2, \dots, n\}$ having dimension $d = 200$ from Gaussian distribution centered at 0 with covariance matrix Σ which is a diagonal matrix. Similarly, we generate a 200-dimensional prediction vector θ sampled from the normal distribution. We compute the prediction vector y_i for x_i as follows, $y_i = \text{sign}(\theta^\top x_i)$. Separately from the training set, another set of 15,000 data points are generated which are used to evaluate the test performance. The data generating distribution for training and test samples are kept the same.

To generate Fig. 1a, we fix $n = 200,000$, $\mu = 4$ and the i -th singular value of Σ is $1/i$. We run one pass of vanilla SGD algorithm and our uncertainty sampling based active learning algorithm. As clear from our plot, out of 200,000 samples, our algorithm asks labels for only 40,000 sample. Hence, for the comparison, we compare the performance on test set of first 40,000 iteration of vanilla SGD with our algorithm.

In Fig. 1b, we compare the test performance for various μ . It is obvious that larger value of μ implies a more aggressive sampling scheme. Hence, it is expected that for a fixed budget, larger value of μ will provide better performance which is also reflected in the plot 1b.

In Fig. 1c, we plot the test accuracy with respect to the number of labels queries given with fixed budget of unlabeled

data. As can be seen from the plot, vanilla SGD asks for the label of every data point. As we increase μ , the number of queried labels goes down. However, the performance on test data remains almost the same surprisingly for uncertainty sampling based algorithm with increasing μ and for vanilla SGD at the end of one epoch.

Real World Datasets. Next, we evaluate our algorithm on real world datasets. We consider *Covertype* and *letter-binary* datasets to evaluate the test performance of uncertainty sampling based active learning for binary classification. Normalized binary version of datasets are downloaded from manikvarma.org/code/LDKL/download.html. Since, the decision boundary for these datasets are non-linear, we use randomized Fourier features (Rahimi et al., 2007) to perform the task of classification. For both the datasets, we use 500 random fourier feature representation. We perform vanilla SGD and uncertainty sampling based active learning. We plot the result in Fig. 2a for *covertype* and in Fig. 2b for *letter-binary*. We observe that uncertainty sampling based active learning algorithm has low prediction error for a fixed budget.

In the last Fig. 2c, we plot the fraction of labels queried for various value of μ when number of unlabeled synthetically generated samples are fixed to 20,000 and 40,000. The plot shows that, our algorithm tends to query similar fraction of labels irrespective of number of available unlabeled data points.

5. Conclusion

In this paper, we show that our uncertainty sampling based active learning method converges to the optimal predictor for linear models under the proposed sampling scheme for binary as well as for multi-class classification. We also extend our analysis for noisy case under restrictive noise assumptions (eqs. (17) and (18)). As a future research direction, we would like to analyze uncertainty sampling algorithm based active learning algorithm under more generalized noise assumptions of Tsybakov (2004), and would like to consider more aggressive sampling schemes.

Acknowledgment

Anant Raj is supported by the a Marie Sklodowska-Curie Fellowship (project NN-OVEROPT 101030817). We also acknowledge support from the French government under the management of the Agence Nationale de la Recherche as part of the ‘‘Investissements d’avenir’’ program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), as well as from the European Research Council (grant SEQUOIA 724063).




Figure 1. Experimental results for binary classification on synthetic data.




Figure 2. Experimental results for binary classification.

References

Agarwal, A. Selective sampling algorithms for cost-sensitive multiclass prediction. In *International Conference on Machine Learning*, pp. 1220–1228, 2013.

Balcan, M.-F. and Long, P. Active and passive learning of linear separators under log-concave distributions. In *Conference on Learning Theory*, pp. 288–316, 2013.

Balcan, M.-F., Broder, A., and Zhang, T. Margin based active learning. In *International Conference on Computational Learning Theory*, pp. 35–50, 2007.

Cavallanti, G., Cesa-Bianchi, N., and Gentile, C. Learning noisy linear classifiers via adaptive and selective sampling. *Machine Learning*, 83(1):71–102, 2011.

Cesa-Bianchi, N., Gentile, C., and Orabona, F. Robust bounds for classification via selective sampling. In *Proceedings of the International Conference on Machine Learning*, pp. 121–128, 2009.

Crammer, K. and Singer, Y. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2(Dec):265–292, 2001.

Dasgupta, S., Kalai, A. T., and Monteleoni, C. Analysis of perceptron-based active learning. In *International*

conference on computational learning theory, pp. 249–263, 2005.

Dekel, O., Gentile, C., and Sridharan, K. Robust selective sampling from single and multiple teachers. In *COLT*, pp. 346–358, 2010.

Hanneke, S. Theory of disagreement-based active learning. *Foundations and Trends® in Machine Learning*, 7(2-3): 131–309, 2014.

Hofmann, T., Schölkopf, B., and Smola, A. J. Kernel methods in machine learning. *The Annals of Statistics*, 36(3): 1171–1220, 2008.

Lewis, D. D. and Gale, W. A. A sequential algorithm for training text classifiers. In *SIGIR '94*, pp. 3–12. Springer, 1994.

Lughofer, E. and Pratama, M. Online active learning in data stream regression using uncertainty sampling based on evolving generalized fuzzy models. *IEEE Transactions on Fuzzy Systems*, 26(1):292–309, 2017.

Massart, P. and Nédélec, É. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.

- Monarch, R. M. *Human-in-the-Loop Machine Learning: Active learning and annotation for human-centered AI*. Simon and Schuster, 2021.
- Mussmann, S. and Liang, P. S. Uncertainty sampling is pre-conditioned stochastic gradient descent on zero-one loss. In *Advances in Neural Information Processing Systems*, pp. 6955–6964, 2018.
- Nguyen, V.-L., Shaker, M. H., and Hüllermeier, E. How to measure uncertainty in uncertainty sampling for active learning. *Machine Learning*, pp. 1–34, 2021.
- Orabona, F. and Cesa-Bianchi, N. Better algorithms for selective sampling. In *International Conference on Machine Learning*, pp. 433–440. Omnipress, 2011.
- Rahimi, A., Recht, B., et al. Random features for large-scale kernel machines. In *NIPS*, volume 3, pp. 5, 2007.
- Roy, N. and McCallum, A. Toward optimal active learning through monte carlo estimation of error reduction. *ICML*, 2:441–448, 2001.
- Sabato, S. and Hess, T. Interactive algorithms: from pool to stream. In *Conference on Learning Theory*, pp. 1419–1439, 2016.
- Schohn, G. and Cohn, D. Less is more: Active learning with support vector machines. In *ICML*, volume 2, pp. 6, 2000.
- Settles, B. *Active learning: Synthesis lectures on artificial intelligence and machine learning*. Long Island, NY: Morgan & Clay Pool, 2012.
- Settles, B., Craven, M., and Ray, S. Multiple-instance active learning. *Advances in Neural Information Processing Systems*, 20:1289–1296, 2007.
- Seung, H. S., Opper, M., and Sompolinsky, H. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pp. 287–294, 1992.
- Tsochantaridis, I., Joachims, T., Hofmann, T., Altun, Y., and Singer, Y. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6(9), 2005.
- Tsybakov, A. B. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- Vaswani, S., Bach, F., and Schmidt, M. Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In *International Conference on Artificial Intelligence and Statistics*, pp. 1195–1204, 2019.
- Wang, G., Hwang, J.-N., Rose, C., and Wallace, F. Uncertainty sampling based active learning with diversity constraint by sparse selection. In *International Workshop on Multimedia Signal Processing*, pp. 1–6. IEEE, 2017.
- Wang, R., Chow, C.-Y., and Kwong, S. Ambiguity-based multiclass active learning. *IEEE Transactions on Fuzzy Systems*, 24(1):242–248, 2015.
- Wang, Y. and Singh, A. Noise-adaptive margin-based active learning and lower bounds under tsybakov noise condition. In *AAAI Conference on Artificial Intelligence*, 2016.
- Yang, Y. and Loog, M. Active learning using uncertainty information. In *International Conference on Pattern Recognition (ICPR)*, pp. 2646–2651, 2016.
- Yang, Y., Ma, Z., Nie, F., Chang, X., and Hauptmann, A. G. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127, 2015.
- Zhu, J., Wang, H., Yao, T., and Tsou, B. K. Active learning with sampling by uncertainty and density for word sense disambiguation and text classification. In *Proceedings of the International Conference on Computational Linguistics*, pp. 1137–1144, 2008.

Appendix

A. Binary Separable Classification

Theorem (Restatement of Theorem 3.1). *Consider a set of n i.i.d samples (x_i, y_i) jointly sampled from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{-1, 1\}$ for all $i = 1, \dots, n$ then under the assumption that there exists a θ_* for which $y(\theta_*^\top x) \geq \rho^*$ for all (x, y) pair in \mathcal{P} , the following convergence guarantee exists for Algorithm 1,*

$$\mathbb{E}(1 - y\theta_t^\top x)_+ \leq \frac{R^2 \max\left\{1, \frac{1}{\mu}\right\} \|\theta_1 - \theta_*\|^2}{\min\left\{\frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu}\right\}^2 n}, \quad (20)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu|\theta^\top x|}$, step size $\gamma = \frac{\min\left\{\frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu}\right\}}{R^2 \max\left\{1, \frac{1}{\mu}\right\}}$ and $\|x\| \leq R$ for all x in the domain \mathcal{X} .

Proof. We minimize the square hinge loss which can be written as,

$$\ell(x, y, \theta) = \frac{1}{2} \max[0, 1 - y\theta^\top x]^2.$$

We have the following update rule for:

$$\theta_{t+1} = \theta_t + \gamma z_t(y_t x_t) [1 - y_t(\theta_t^\top x_t)]_+$$

where z_t is a bernoulli random variable for fixed θ_t and x_t such that $p(z_t = 1 | x_t, \theta_t) = \sigma(\theta_t, x_t)$ where $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ is an even function. Following the update, we have

$$\begin{aligned} \|\theta_{t+1} - \theta^*\|^2 &= \left\| \theta_t - \theta^* + \gamma z_t(y_t x_t) [1 - y_t(\theta_t^\top x_t)]_+ \right\|^2 \\ &= \|\theta_t - \theta^*\|^2 + 2\gamma z_t [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) + \gamma^2 z_t^2 (y_t x_t)^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \\ &= \|\theta_t - \theta^*\|^2 + 2\gamma z_t [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) + \gamma^2 z_t^2 (y_t x_t)^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \end{aligned} \quad (21)$$

In the last equation we used the fact that $z_t \in \{0, 1\}$, hence $z_t^2 = z_t$. Taking expectations on both sides only with respect to z_t considering (x_t, y_t, θ_t) fixed, we have

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 &= \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) \\ &\quad + \gamma^2\sigma(\theta_t, x_t) \|y_t x_t\|^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \\ &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) \\ &\quad + \gamma^2\sigma(\theta_t, x_t) R^2 [1 - y_t(\theta_t^\top x_t)]_+^2. \end{aligned} \quad (22)$$

Now in the above equation, we want to find function $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ for $y_t(\theta_t^\top x_t) < 1$ for all t , such that the following holds,

$$\begin{aligned} \sigma(\theta_t, x_t) (1 - y_t \theta_t^\top x_t)_+ &\leq c_1 (1 - y_t \theta_t^\top x_t)_+ \\ \sigma(\theta_t, x_t) (1 - y_t \theta_t^\top x_t)_+ (y_t \theta_t^\top x_t - y_t \theta^{*\top} x_t) &\leq -c_2 (1 - y_t \theta_t^\top x_t)_+ \end{aligned}$$

for some positive constants c_1 and c_2 . This is then equivalent to, for all $y_t \theta_t^\top x_t < 1$ and $y_t \theta_*^\top x_t \geq \rho^*$:

$$\sigma(\theta_t, x_t) (1 - y_t \theta_t^\top x_t)_+ \leq c_1 \quad (23)$$

$$\sigma(\theta_t, x_t)(\rho^* - y_t \theta_t^\top x_t) \geq c_2 \quad (24)$$

IF we choose, $\sigma(\theta, x) = \frac{1}{1+\mu|\theta^\top x|}$ for some constant $\mu > 0$, then from Lemma A.1, we have

$$\frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{(\rho^* - y\theta^\top x)} \leq \sigma(\theta, x) \leq \frac{\max \left\{ 1, \frac{1}{\mu} \right\}}{(1 - y\theta^\top x)_+}.$$

Hence, we have

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 - 2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} (1 - y_t \theta_t^\top x_t)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} (1 - y_t \theta_t^\top x_t)_+ \\ \Rightarrow (1 - y_t \theta_t^\top x_t)_+ &\leq \frac{1}{2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \left[\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 \right] \end{aligned}$$

Taking expectation on both the sides, we have

$$(1 - y\theta_t^\top x)_+ \leq \frac{1}{2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \left[\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 \right] \quad (25)$$

Summing the above equation for $t = 1$ to n , choosing $\gamma = \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$ and applying Jensen's inequality give us,

$$\mathbb{E}(1 - y\bar{\theta}_n^\top x)_+ \leq \frac{R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \|\theta_1 - \theta_*\|^2}{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}^2 n}. \quad (26)$$

□

Lemma A.1. *If there exist a function $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ of the form,*

$$\sigma(\theta, x) = \frac{1}{1 + \mu|\theta^\top x|} \quad (27)$$

where $\theta \in \mathbb{R}^d$ and $x \in \mathbb{R}^d$, then under the assumptions of Theorem 3.1 we have

$$\frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{(\rho^* - y\theta^\top x)} \leq \sigma(\theta, x) \leq \frac{\max \left\{ 1, \frac{1}{\mu} \right\}}{(1 - y\theta^\top x)_+} \quad (28)$$

for all $y(\theta^\top x) \leq 1$ where $y \in \{-1, 1\}$.

Proof. We want to find an even function $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$, such that for all $y(\theta^\top x) \leq 1$ where $y \in \{-1, 1\}$, the following conditions hold

$$\begin{aligned} \sigma(\theta, x)(1 - y\theta^\top x)_+ &\leq c_1 \\ \sigma(\theta, x)(y\theta^\top x - \rho^*) &\leq -c_2 \end{aligned}$$

for some positive constants c_1 and c_2 and $\frac{c_1}{c_2}$ is as small as possible. Replacing $\sigma(\theta, x)$ with $\frac{1}{1+\mu|\theta^\top x|}$ and using the fact that $y\theta^\top x \leq 1$, we get to find constants c_1 and c_2 such that

$$\begin{aligned} c_1 &\geq \frac{1 - y\theta^\top x}{1 + \mu|\theta^\top x|}, \\ c_2 &\leq \frac{\rho^* - y\theta^\top x}{1 + \mu|\theta^\top x|}. \end{aligned}$$

When $\mu \geq \frac{1}{\rho^*}$, $\frac{\rho^*-u}{1+\mu|u|}$ is decreasing function for $u > 0$ and increasing function for $u \leq 0$. When $\mu \leq \frac{1}{\rho^*}$, $\frac{\rho^*-u}{1+\mu|u|}$ is a decreasing function everywhere. Checking at $y\theta^\top x = 0$, $y\theta^\top x = 1$ and $y\theta^\top x \rightarrow -\infty$, we have

$$c_1 \geq \max \left\{ 1, \frac{1}{\mu} \right\} \text{ and } c_2 \leq \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}.$$

Hence, finally we have

$$\frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^*-1}{1+\mu} \right\}}{(\rho^* - y\theta^\top x)} \leq \sigma(\theta, x) \leq \frac{\max \left\{ 1, \frac{1}{\mu} \right\}}{(1 - y\theta^\top x)_+}. \quad (29)$$

□

B. Separable Multi-class Classification

Theorem (Restatement of Theorem 3.2). *Consider a set of n i.i.d samples (x_i, y_i) jointly sampled from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{1, 2, \dots, k\}$ for all $i = 1, \dots, n$ then under the assumption that there exists a set of d -dimensional optimal half-spaces $\theta_\star = \{\theta_\star(1), \theta_\star(2), \dots, \theta_\star(k)\}$ corresponding to each class for which $\theta_\star^\top \delta_x(y, y^\star(\theta_\star, x, y)) \geq \rho^*$ for all (x, y) pair in \mathcal{P} , the following convergence guarantee exists for Algorithm 2 under projected gradient descent update in equation (11),*

$$\mathbb{E} \hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2(1 + BR) \|\theta_1 - \theta_\star\|^2}{\min \left\{ \frac{1}{\mu}, \frac{\rho^*-1}{1+\mu} \right\}^2 n}, \quad (30)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu |\theta_t(y_{t(1)}^\top x_t - \theta_t(y_{t(2)}^\top x_t)|}$, step size $\gamma = \min \left\{ \frac{1}{\mu}, \frac{\rho^*-1}{1+\mu} \right\} \frac{1}{R^2(1+BR)}$ and $\|\delta_x(i, j)\| \leq R$ for all x in the domain \mathcal{X} , and for all $i \neq j \in [k]$ and $\|\theta\| \leq B$.

Proof. We have multi-class hinge loss,

$$\hat{\ell}(x_t, y_t, \theta) = [1 - \theta^\top \delta_{x_t}(y_t, y^\star(\theta, x_t, y_t))]_+, \text{ where } y^\star(\theta, x_t, y_t) = \arg \max_{y \in \mathcal{Y} \setminus y_t} \theta^\top \phi(x_t, y),$$

and $\phi(x, y) \in \mathbb{R}^{dk}$ represents the feature map corresponding to the sample (x, y) . We have assumed that

$$\theta_\star^\top \delta_{x_t}(y_t, y^\star(\theta, x_t, y_t)) \geq \rho^* \geq 1 \text{ for all } t.$$

Let us consider the smooth loss function,

$$\ell(x_t, y_t, \theta) = [\hat{\ell}(x_t, y_t, \theta)]^2 = [1 - \theta^\top \delta_{x_t}(y_t, y^\star(\theta, x_t, y_t))]_+^2.$$

We perform the following projected stochastic update

$$\theta_{t+1} = \Pi_{\|\theta\| \leq B} \left[\theta_t + \gamma z_t \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) [1 - \theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))]_+ \right]$$

where z_t is a bernoulli random variable such that $p(z_t = 1 | x_t, \theta_t) = \sigma(\theta_t, \top x_t)$ such that $\sigma : \mathbb{R}^{dk} \times \mathbb{R}^{dk} \rightarrow [0, 1]$ is an even function. Now,

$$\begin{aligned} \|\theta_{t+1} - \theta_\star\|^2 &\leq \left\| \theta_t - \theta_\star + \gamma z_t \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) [1 - \theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))]_+ \right\|^2 \\ &= \|\theta_t - \theta_\star\|^2 + 2\gamma z_t \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) - \theta_\star^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \right) \\ &\quad + \gamma^2 z_t^2 \|\delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))\|^2 \hat{\ell}^2(x_t, y_t, \theta_t) \\ &= \|\theta_t - \theta_\star\|^2 + 2\gamma z_t \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) - \theta_\star^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \right) \\ &\quad + \gamma^2 z_t \|\delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))\|^2 \hat{\ell}^2(x_t, y_t, \theta_t). \end{aligned} \quad (31)$$

In the last equation we used the fact that $z_t \in \{0, 1\}$, hence $z_t^2 = z_t$. Taking expectations on both sides only with respect to z_t considering (x_t, y_t, θ_t) fixed, we have

$$\begin{aligned}
 \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t)\hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \theta^{*\top} \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \right) \\
 &\quad + \gamma^2\sigma(\theta_t, x_t)\|\delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))\|^2\hat{\ell}^2(x_t, y_t, \theta_t) \\
 &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t)\hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \theta^{*\top} \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \right) \\
 &\quad + \gamma^2\sigma(\theta_t, x_t)R^2\hat{\ell}^2(x_t, y_t, \theta_t) \\
 &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t, x_t)\hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \rho^* \right) \\
 &\quad + \gamma^2\sigma(\theta_t, x_t)R^2\hat{\ell}^2(x_t, y_t, \theta_t)
 \end{aligned} \tag{32}$$

Now, we want to choose such a function σ for $\hat{\ell}(x_t, y_t, \theta_t) > 0$ such that

$$\begin{aligned}
 \sigma(\theta_t, x_t)\hat{\ell}^2(x_t, y_t, \theta_t) &\leq c_1\hat{\ell}(x_t, y_t, \theta_t) \\
 \sigma(\theta_t, x_t) \left(\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \right) &\geq c_2
 \end{aligned}$$

for some positive constants c_1 and c_2 . If we choose,

$$\sigma(\theta, x) = \frac{1}{1 + \mu \left| \theta(y_{(1)}^*)^\top x - \theta(y_{(2)}^*)^\top x \right|},$$

the from Lemma B.1 we have,

$$\begin{aligned}
 \sigma(\theta_t, x_t) &\leq \frac{1 + BR}{\hat{\ell}(x_t, y_t, \theta_t)}, \\
 \sigma(\theta_t, x_t) &\geq \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \frac{1}{\left(\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \right)}.
 \end{aligned} \tag{33}$$

Putting the inequalities in equation (33) back in the equation (32), we get the following,

$$\begin{aligned}
 \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 - 2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \hat{\ell}(x_t, y_t, \theta_t) + \gamma^2 R^2(1 + BR)\hat{\ell}(x_t, y_t, \theta_t) \\
 \Rightarrow \hat{\ell}(x_t, y_t, \theta_t) &\leq \frac{1}{2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2(1 + BR)} \left[\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 \right]
 \end{aligned}$$

Taking expectation on both the sides, we have

$$\hat{\ell}(x, y, \theta_t) \leq \frac{1}{2\gamma \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2(1 + BR)} \left[\mathbb{E}\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 \right] \tag{34}$$

Summing the above equation for $t = 1$ to n , choosing $\gamma = \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \frac{1}{R^2(1 + BR)}$ and applying Jensen's inequality give us,

$$\mathbb{E}\hat{\ell}(x, y, \bar{\theta}_n) \leq \frac{R^2(1 + BR)\|\theta_1 - \theta^*\|^2}{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}^2 n}. \tag{35}$$

□

Lemma B.1. Consider the setting and notations of Theorem 3.2. If there exist a function $\sigma : \mathbb{R}^{dk} \times \mathbb{R}^{dk} \rightarrow \mathbb{R}$ of the form,

$$\sigma(\theta, x) = \frac{1}{1 + \mu \left| \theta(y_{(1)}^*)^\top x - \theta(y_{(2)}^*)^\top x \right|} \tag{36}$$

where $\theta \in \mathbb{R}^{dk}$ and $x \in \mathbb{R}^{dk}$, then under the assumptions of Theorem 3.2 we have

$$\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \frac{1}{(\rho^* - \theta^\top \delta_x(y, y^*(\theta, x, y)))} \leq \sigma(\theta, x) \leq \frac{1 + BR}{\hat{\ell}(x, y, \theta)} \quad (37)$$

for all $\theta^\top \delta_x(y, y^*(\theta, x, y)) \leq 1$ where $y \in \{1, 2, \dots, k\}$.

Proof. We want to find an even function $\sigma : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$, such that for all $y(\theta^\top x) \leq 1$ where $y \in \{-1, 1\}$, the following conditions hold

$$\begin{aligned} \sigma(\theta_t, x_t) \hat{\ell}^2(x_t, y_t, \theta_t) &\leq c_1 \hat{\ell}(x_t, y_t, \theta_t) \\ \sigma(\theta_t, x_t) (\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))) &\geq c_2 \end{aligned}$$

for some positive constants c_1 and c_2 and $\frac{c_1}{c_2}$ is as small as possible. Replacing $\sigma(\theta, x)$ with $\frac{1}{1 + \mu |\theta(y_{t(1)}^*)^\top x - \theta(y_{t(2)}^*)^\top x|}$ and using the fact that $\hat{\ell}(x_t, y_t, \theta_t) y > 0$, we get to find constants c_1 and c_2 such that

$$c_1 \geq \frac{\hat{\ell}(x_t, y_t, \theta_t)}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|} \quad (38)$$

$$c_2 \leq \frac{\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|} \quad (39)$$

Under bounded θ assumption we have,

$$\frac{\hat{\ell}(x_t, y_t, \theta_t)}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|} \leq 1 + BR. \quad (40)$$

In the above equation, we have used cauchy-shwartz inequality, $\|\theta\| \leq B$ and $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$. Let us now consider to get the inequality for c_2 . We here observe that

$$|\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t| \leq |\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))|$$

Hence,

$$\frac{\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))}{1 + \mu |\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))|} \leq \frac{\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|}$$

When $\mu \geq \frac{1}{\rho^*}$, $\frac{\rho^* - u}{1 + \mu |u|}$ is decreasing function for $u > 0$ and increasing function for $u \leq 0$. When $\mu \leq \frac{1}{\rho^*}$, $\frac{\rho^* - u}{1 + \mu |u|}$ is a decreasing function everywhere. Hence, checking the value at $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) = 1$ and $\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \rightarrow -\infty$ gives us the following,

$$c_2 \leq \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \leq \frac{\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|}.$$

Hence, we have

$$\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \frac{1}{(\rho^* - \theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)))} \leq \sigma(\theta_t, x_t) \leq \frac{1 + BR}{\hat{\ell}(x_t, y_t, \theta_t)}. \quad (41)$$

□

C. Binary Classification in the Presence of Noise

Lemma C.1. *Under the assumption in equation (17), the iterates in Algorithm 1 updated via stochastic gradient descent (equation (4)) or projected stochastic gradient descent (equation (19)) satisfy*

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \mathbb{E} \|\theta_t - \theta^*\|^2 - 2\gamma(1-\eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1+\mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R\|\theta_\star\|}{1+\mu}, R\|\theta_\star\| \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \end{aligned} \quad (42)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1+\mu|\theta^\top x|}$, $\mu > 0$, and positive step size γ such that for all $\|x\| \leq R$ for all $x \in \mathcal{X}$.

Proof.

$$\begin{aligned} \|\theta_{t+1} - \theta^*\|^2 &= \left\| \theta_t - \theta^* + \gamma z_t(y_t x_t) [1 - y_t(\theta_t^\top x_t)]_+ \right\|^2 \\ &= \|\theta_t - \theta^*\|^2 + 2\gamma z_t [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) + \gamma^2 z_t^2 (y_t x_t)^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \\ &= \|\theta_t - \theta^*\|^2 + 2\gamma z_t [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) + \gamma^2 z_t (y_t x_t)^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \end{aligned} \quad (43)$$

In the last equation we used the fact that $z_t \in \{0, 1\}$, hence $z_t^2 = z_t$. Taking expectations on both sides only with respect to z_t considering (x_t, y_t, θ_t) fixed, we have

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t^\top x_t) [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) \\ &\quad + \gamma^2 \sigma(\theta_t^\top x_t) (y_t x_t)^2 [1 - y_t(\theta_t^\top x_t)]_+^2 \\ &\leq \|\theta_t - \theta^*\|^2 + 2\gamma\sigma(\theta_t^\top x_t) [1 - y_t(\theta_t^\top x_t)]_+ \left(y_t(\theta_t^\top x_t) - y_t(\theta^{*\top} x_t) \right) \\ &\quad + \gamma^2 \sigma(\theta_t^\top x_t) R^2 [1 - y_t(\theta_t^\top x_t)]_+^2. \end{aligned} \quad (44)$$

In the last line, we used the fact that $\|x\| \leq R$ for all $x \in \mathcal{X}$. Now finally we take the expectation with respect to the noise in y_t , using the assumption made in Equation (17) and conditioning on θ_t, x_t and y_t . We get the following,

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 - 2\gamma(1-\eta)\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ [(y_t(\theta_t^\top x_t) - \rho^*)] \\ &\quad + 2\gamma\eta\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ [y_t(\theta_t^\top x_t) + |\theta^{*\top} x_t|] + \gamma^2 R^2 \sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+^2 \\ &\leq \|\theta_t - \theta^*\|^2 - 2\gamma(1-\eta)\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ [(y_t(\theta_t^\top x_t) - \rho^*)] \\ &\quad + 2\gamma\eta\sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+ [y_t(\theta_t^\top x_t) + R\|\theta^*\|] + \gamma^2 R^2 \sigma(\theta_t, x_t) [1 - y_t(\theta_t^\top x_t)]_+^2. \end{aligned} \quad (45)$$

Similar to the noiseless case, we choose

$$\sigma(\theta, x) = \frac{1}{1 + \mu|\theta^\top x|},$$

for $\mu > 0$ and after applying result from lemma A.1, we have for $y_t\theta_t^\top x_t < 1$

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \|\theta_t - \theta^*\|^2 - 2\gamma(1-\eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1+\mu} \right\} (1 - y_t\theta_t^\top x_t)_+ \\ &\quad + 2\gamma\eta \frac{y_t\theta_t^\top x_t + |\theta_\star^\top x_t|}{1 + \mu|\theta_t^\top x_t|} (1 - y_t\theta_t^\top x_t)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} (1 - y_t\theta_t^\top x_t)_+ \\ &\leq \|\theta_t - \theta^*\|^2 - 2\gamma(1-\eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1+\mu} \right\} (1 - y_t\theta_t^\top x_t)_+ \\ &\quad + 2\gamma\eta \frac{y_t\theta_t^\top x_t + R\|\theta_\star\|}{1 + \mu|\theta_t^\top x_t|} (1 - y_t\theta_t^\top x_t)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} (1 - y_t\theta_t^\top x_t)_+ \end{aligned} \quad (46)$$

It is easier to see that $\frac{y_t \theta_t^\top x_t + R \|\theta_\star\|}{1 + \mu \|\theta_t^\top x_t\|} \leq \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\}$. Hence,

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^\star\|^2 &\leq \|\theta_t - \theta^\star\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} (1 - y_t \theta_t^\top x_t)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\} (1 - y_t \theta_t^\top x_t)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} (1 - y_t \theta_t^\top x_t)_+ \end{aligned} \quad (47)$$

Taking expectation on both sides we have,

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^\star\|^2 &\leq \mathbb{E} \|\theta_t - \theta^\star\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} \mathbb{E}(1 - y \theta_t^\top x)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\} \mathbb{E}(1 - y \theta_t^\top x)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \mathbb{E}(1 - y \theta_t^\top x)_+. \end{aligned} \quad (48)$$

□

Theorem (Restatement of Theorem 3.3). Consider a set of n i.i.d samples (x_i, y_i) jointly sampled from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{-1, 1\}$ for all $i = 1, \dots, n$ then under the assumption in equation (17) for all (x, y) pair in \mathcal{P} and for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu \|\theta^\top x\|}$, $\mu > 0$, the following convergence guarantee exists for Algorithm 1:

1. If the noise parameter $\eta < \frac{\min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}}{\max \left\{ R \|\theta_\star\| + \frac{1 + R \|\theta_\star\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}}$ and iterates in Algorithm 1 are updated via stochastic gradient descent update in equation (4), then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$

$$\mathbb{E}(1 - y \bar{\theta}_n^\top x)_+ \leq \frac{R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \|\theta_1 - \theta_\star\|^2}{\left[(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\} \right]^2 n}, \quad (49)$$

such that $\|x\| \leq R$ for all $x \in \mathcal{X}$.

2. If the noise parameter $\eta \geq \frac{\min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}}{\max \left\{ R \|\theta_\star\| + \frac{1 + R \|\theta_\star\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}}$ and iterates in Algorithm 1 are updated via projected stochastic gradient descent update in equation (19), then for step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$

$$\mathbb{E}(1 - y \theta_t^\top x)_+ \leq \frac{\left(R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \right) \|\theta_1 - \theta^\star\|^2}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\}^2 n} + O(\eta), \quad (50)$$

such that $\|x\| \leq R$ for all $x \in \mathcal{X}$.

Proof. From Lemma C.1, we have

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^\star\|^2 &\leq \mathbb{E} \|\theta_t - \theta^\star\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} \mathbb{E}(1 - y \theta_t^\top x)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R \|\theta_\star\|}{1 + \mu}, R \|\theta_\star\| \right\} \mathbb{E}(1 - y \theta_t^\top x)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \mathbb{E}(1 - y \theta_t^\top x)_+ \end{aligned}$$

Now we consider the two cases.

i When

$$\eta < \frac{\min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}}{\max \left\{ R \|\theta_\star\| + \frac{1 + R \|\theta_\star\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu} + \frac{\rho^\star - 1}{1 + \mu} \right\}},$$

then, $(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} > \eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\}$. Hence,

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \mathbb{E} \|\theta_t - \theta^*\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \\ \Rightarrow \mathbb{E}(1 - y\theta_t^\top x)_+ &\leq \frac{\mathbb{E} \|\theta_t - \theta^*\|^2 - \mathbb{E} \|\theta_{t+1} - \theta^*\|^2}{2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - 2\gamma\eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \end{aligned}$$

Now in the above equation, summing for all t from 1 to n and applying Jensen's inequality after choosing the optimal step size $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$, we get

$$\mathbb{E}(1 - y\theta_n^\top x)_+ \leq \frac{R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \|\theta_1 - \theta_*\|^2}{\left[(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\} \right]^2 n}. \quad (51)$$

ii When

$$\eta \geq \frac{\min \left\{ \frac{1}{\mu} + \frac{\rho^* - 1}{1 + \mu} \right\}}{\max \left\{ R\|\theta_*\| + \frac{1 + R\|\theta_*\|}{1 + \mu} \right\} + \min \left\{ \frac{1}{\mu} + \frac{\rho^* - 1}{1 + \mu} \right\}},$$

then $(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} < \eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\}$. Hence,

$$\begin{aligned} \mathbb{E} \|\theta_{t+1} - \theta^*\|^2 &\leq \mathbb{E} \|\theta_t - \theta^*\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \\ &\quad + 2\gamma\eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ + \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \mathbb{E}(1 - y\theta_t^\top x)_+ \\ \Rightarrow \mathbb{E}(1 - y\theta_t^\top x)_+ &\leq \frac{\|\theta_t - \theta^*\|^2 - \mathbb{E} \|\theta_{t+1} - \theta^*\|^2}{2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \\ &\quad + \frac{2\gamma\eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\}}{2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \mathbb{E}(1 - y\theta_t^\top x)_+ \end{aligned}$$

We update via projected stochastic gradient descent. Hence,

$$(1 - y\theta_t^\top x)_+ \leq 1 + BR$$

for all (x, y) pair. This gives,

$$\begin{aligned} \mathbb{E}(1 - y\theta_t^\top x)_+ &\leq \frac{\mathbb{E} \|\theta_t - \theta^*\|^2 - \mathbb{E} \|\theta_{t+1} - \theta^*\|^2}{2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}} \\ &\quad + \frac{2\gamma\eta \max \left\{ \frac{1 + R\|\theta_*\|}{1 + \mu}, R\|\theta_*\| \right\} (1 + BR)}{2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \gamma^2 R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}. \end{aligned}$$

Choosing optimal $\gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{R^2 \max \left\{ 1, \frac{1}{\mu} \right\}}$ we have

$$\mathbb{E}(1 - y\theta_t^\top x)_+ \leq \frac{\left(R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \right) [\mathbb{E} \|\theta_t - \theta^*\|^2 - \mathbb{E} \|\theta_{t+1} - \theta^*\|^2]}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}^2}$$

$$+ \left(\frac{2(1 + BR) \max \left\{ \frac{1+R\|\theta_\star\|}{1+\mu}, R\|\theta_\star\| \right\}}{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1+\mu} \right\}} \right) \eta.$$

Summing the above equation for $t = 1$ to n and applying Jensen's inequality give us,

$$\mathbb{E}(1 - y\bar{\theta}_n^\top x)_+ \leq \frac{\left(R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \right) \|\theta_1 - \theta^\star\|^2}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1+\mu} \right\}^2 n} + \underbrace{\left(\frac{2(1 + BR) \max \left\{ \frac{1+R\|\theta_\star\|}{1+\mu}, R\|\theta_\star\| \right\}}{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1+\mu} \right\}} \right)}_{:\approx \text{constant}} \eta \quad (52)$$

$$= \frac{\left(R^2 \max \left\{ 1, \frac{1}{\mu} \right\} \right) \|\theta_1 - \theta^\star\|^2}{(1 - \eta)^2 \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1+\mu} \right\}^2 n} + O(\eta). \quad (53)$$

□

D. Multi-class Classification in the Presence of Noise

Lemma D.1. *Under the assumption in equation (18), the iterates in Algorithm 2 updated via projected stochastic gradient descent (equation (11)) satisfy*

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta^\star\|^2 &\leq \mathbb{E}\|\theta_t - \theta^\star\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^\star - 1}{1 + \mu} \right\} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\ &\quad + 2\gamma\eta(1 + R\|\theta_\star\|)\mathbb{E}[\hat{\ell}(x, y, \theta_t)] + \gamma^2(1 + BR)R^2\mathbb{E}[\hat{\ell}(x, y, \theta_t)] \end{aligned} \quad (54)$$

for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu|\theta(y_1^\top x - \theta(y_2^\top x))|}$, $\mu > 0$, and positive step size γ such that for all $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$.

Proof. We perform the following projected stochastic update

$$\theta_{t+1} = \Pi_{\|\theta\| \leq B} \left[\theta_t + \gamma z_t \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) [1 - \theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))]_+ \right]$$

where z_t is a bernoulli random variable such that $p(z_t = 1 | x_t, \theta_t) = \sigma(\theta_t, \top x_t)$ such that $\sigma : \mathbb{R}^{dk} \times \mathbb{R}^{dk} \rightarrow [0, 1]$ is an even function. Now,

$$\begin{aligned} \|\theta_{t+1} - \theta^\star\|^2 &\leq \left\| \theta_t - \theta^\star + \gamma z_t \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) [1 - \theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))]_+ \right\|^2 \\ &= \|\theta_t - \theta^\star\|^2 + 2\gamma z_t \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) - \theta^\star^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \right) \\ &\quad + \gamma^2 z_t^2 \|\delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))\|^2 \hat{\ell}^2(x_t, y_t, \theta_t) \\ &= \|\theta_t - \theta^\star\|^2 + 2\gamma z_t \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) - \theta^\star^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \right) \\ &\quad + \gamma^2 z_t \|\delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t))\|^2 \hat{\ell}^2(x_t, y_t, \theta_t). \end{aligned} \quad (55)$$

In the last equation we used the fact that $z_t \in \{0, 1\}$, hence $z_t^2 = z_t$. Taking expectations on both sides only with respect to z_t considering (x_t, y_t, θ_t) fixed, we have

$$\mathbb{E}\|\theta_{t+1} - \theta^\star\|^2 \leq \|\theta_t - \theta^\star\|^2 + 2\gamma\sigma(\theta_t, x_t)\hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) - \theta^\star^\top \delta_{x_t}(y_t, y^\star(\theta_t, x_t, y_t)) \right)$$

$$\begin{aligned}
 & + \gamma^2 \sigma(\theta_t, x_t) \|\delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))\|^2 \hat{\ell}^2(x_t, y_t, \theta_t) \\
 \leq & \|\theta_t - \theta^*\|^2 + 2\gamma \sigma(\theta_t, x_t) \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \theta^{*\top} \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) \right) \\
 & + \gamma^2 \sigma(\theta_t, x_t) R^2 \hat{\ell}^2(x_t, y_t, \theta_t)
 \end{aligned} \tag{56}$$

In the last line, we used the fact that $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$. Now finally we take the expectation with respect to the noise in y_t , using the assumption made in Equation (18) and conditioning on θ_t, x_t and y_t . We get the following,

$$\begin{aligned}
 \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 & \leq \|\theta_t - \theta^*\|^2 + 2\gamma(1 - \eta) \sigma(\theta_t, x_t) \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \rho^* \right) \\
 & + 2\gamma\eta \sigma(\theta_t, x_t) \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) + |\theta_\star^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t))| \right) \\
 & + \gamma^2 \sigma(\theta_t, x_t) R^2 \hat{\ell}^2(x_t, y_t, \theta_t) \\
 & \leq \|\theta_t - \theta^*\|^2 + 2\gamma(1 - \eta) \sigma(\theta_t, x_t) \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) - \rho^* \right) \\
 & + 2\gamma\eta \sigma(\theta_t, x_t) \hat{\ell}(x_t, y_t, \theta_t) \left(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) + R\|\theta_\star\| \right) \\
 & + \gamma^2 \sigma(\theta_t, x_t) R^2 \hat{\ell}^2(x_t, y_t, \theta_t)
 \end{aligned}$$

Similar to the noiseless case, we choose

$$\sigma(\theta, x) = \frac{1}{1 + \mu |\theta(y_1^*)^\top x - \theta(y_2^*)^\top x|}$$

for $\mu > 0$ and after applying result from lemma B.1, we have for $\hat{\ell}^2(x_t, y_t, \theta_t) > 0$

$$\begin{aligned}
 \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 & \leq \|\theta_t - \theta^*\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \hat{\ell}(x_t, y_t, \theta_t) \\
 & + 2\gamma\eta \hat{\ell}(x_t, y_t, \theta_t) \frac{(\theta_t^\top \delta_{x_t}(y_t, y^*(\theta_t, x_t, y_t)) + R\|\theta_\star\|)}{1 + \mu |\theta_t(y_{t(1)}^*)^\top x_t - \theta_t(y_{t(2)}^*)^\top x_t|} \\
 & + \gamma^2(1 + BR)R^2 \hat{\ell}(x_t, y_t, \theta_t) \\
 & \leq \|\theta_t - \theta^*\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \hat{\ell}(x_t, y_t, \theta_t) \\
 & + 2\gamma\eta(1 + R\|\theta_\star\|) \hat{\ell}(x_t, y_t, \theta_t) + \gamma^2(1 + BR)R^2 \hat{\ell}(x_t, y_t, \theta_t).
 \end{aligned} \tag{57}$$

Taking expectation on both sides of the above expression gives us

$$\begin{aligned}
 \mathbb{E}\|\theta_{t+1} - \theta^*\|^2 & \leq \mathbb{E}\|\theta_t - \theta^*\|^2 - 2\gamma(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\
 & + 2\gamma\eta(1 + R\|\theta_\star\|) \mathbb{E}[\hat{\ell}(x, y, \theta_t)] + \gamma^2(1 + BR)R^2 \mathbb{E}[\hat{\ell}(x, y, \theta_t)]
 \end{aligned} \tag{58}$$

□

Theorem (Restatement of Theorem 3.4). *Consider a set of n i.i.d samples (x_i, y_i) jointly sampled from \mathcal{P} such that $x_i \in \mathbb{R}^d$, and $y_i \in \{1, 2, \dots, k\}$ for all $i = 1, \dots, n$ then under the assumption in equation (18) for all (x, y) pair in \mathcal{P} and for the choice of $\sigma(\theta, x) = \frac{1}{1 + \mu |\theta(y_1^*)^\top x - \theta(y_2^*)^\top x|}$, $\mu > 0$, the following convergence guarantee exists for Algorithm 2 with projected stochastic gradient descent update (equation (11)):*

$$\begin{aligned}
 1. \text{ If the noise parameter } \eta < \frac{\min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\}}{1 + \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} + R\|\theta_\star\|}, \text{ then for step size } \gamma = \frac{(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta(1 + R\|\theta_\star\|)}{R^2(1 + BR)} \\
 \mathbb{E}[\hat{\ell}(x, y, \bar{\theta}_n)] & \leq \frac{R^2(1 + BR)\|\theta_1 - \theta_\star\|^2}{\left[(1 - \eta) \min \left\{ \frac{1}{\mu}, \frac{\rho^* - 1}{1 + \mu} \right\} - \eta(1 + R\|\theta_\star\|) \right]^2 n},
 \end{aligned} \tag{59}$$

such that $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$.

2. If the noise parameter $\eta \geq \frac{\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}{1+\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}+R\|\theta_\star\|}$, then for step size $\gamma = \frac{(1-\eta)\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}{R^2(1+BR)}$

$$\mathbb{E}[\hat{\ell}(x, y, \bar{\theta}_n)] \leq \frac{R^2(1+BR)\|\theta_1 - \theta_\star\|^2}{(1-\eta)^2 \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}^2 n} + O(\eta), \quad (60)$$

such that $\|\delta_x(i, j)\| \leq R$ for all $x \in \mathcal{X}$ and $i, j \in [k]$.

Proof. From lemma D.1, we have

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta_\star\|^2 &\leq \mathbb{E}\|\theta_t - \theta_\star\|^2 - 2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} \mathbb{E}[\hat{\ell}(x_t, y_t, \theta_t)] \\ &\quad + 2\gamma\eta(1+R\|\theta_\star\|)\mathbb{E}[\hat{\ell}(x_t, y_t, \theta_t)] + \gamma^2(1+BR)R^2\mathbb{E}[\hat{\ell}(x_t, y_t, \theta_t)] \end{aligned} \quad (61)$$

Now we consider two cases.

i When

$$\eta < \frac{\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}{1 + \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} + R\|\theta_\star\|},$$

then $(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} > \eta(1+R\|\theta_\star\|)$. Hence,

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta_\star\|^2 &\leq \mathbb{E}\|\theta_t - \theta_\star\|^2 - 2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\ &\quad + 2\gamma\eta(1+R\|\theta_\star\|)\mathbb{E}[\hat{\ell}(x, y, \theta_t)] + \gamma^2(1+BR)R^2\mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\ \Rightarrow \mathbb{E}[\hat{\ell}(x, y, \theta_t)] &\leq \frac{\mathbb{E}\|\theta_t - \theta_\star\|^2 - \mathbb{E}\|\theta_{t+1} - \theta_\star\|^2}{2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - 2\gamma\eta(1+R\|\theta_\star\|) - \gamma^2(1+BR)R^2} \end{aligned}$$

Now in the above equation, summing for all t from 1 to n and applying Jensen's inequality after choosing the optimal step size $\gamma = \frac{(1-\eta)\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \eta(1+R\|\theta_\star\|)}{R^2(1+BR)}$, we get

$$\mathbb{E}[\hat{\ell}(x, y, \bar{\theta}_n)] \leq \frac{R^2(1+BR)\|\theta_1 - \theta_\star\|^2}{\left[(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \eta(1+R\|\theta_\star\|)\right]^2 n}. \quad (62)$$

ii When

$$\eta \geq \frac{\min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}{1 + \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} + R\|\theta_\star\|},$$

then $(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} \leq \eta(1+R\|\theta_\star\|)$. Hence,

$$\begin{aligned} \mathbb{E}\|\theta_{t+1} - \theta_\star\|^2 &\leq \mathbb{E}\|\theta_t - \theta_\star\|^2 - 2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\ &\quad + 2\gamma\eta(1+R\|\theta_\star\|)\mathbb{E}[\hat{\ell}(x, y, \theta_t)] + \gamma^2(1+BR)R^2\mathbb{E}[\hat{\ell}(x, y, \theta_t)] \\ \Rightarrow \mathbb{E}[\hat{\ell}(x, y, \theta_t)] &\leq \frac{\mathbb{E}\|\theta_t - \theta_\star\|^2 - \mathbb{E}\|\theta_{t+1} - \theta_\star\|^2}{2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \gamma^2R^2(1+BR)} \\ &\quad + \frac{2\gamma\eta(1+R\|\theta_\star\|)}{2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \gamma^2R^2(1+BR)} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] \end{aligned}$$

We update via projected stochastic gradient descent. Hence,

$$\hat{\ell}(x, y, \theta_t) \leq 1 + BR$$

for all (x, y) pair. This gives,

$$\begin{aligned} \mathbb{E}[\hat{\ell}(x, y, \theta_t)] &\leq \frac{\mathbb{E}\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2}{2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \gamma^2 R^2(1+BR)} \\ &\quad + \frac{2\gamma\eta(1+R\|\theta_\star\|)}{2\gamma(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\} - \gamma^2 R^2(1+BR)}(1+BR). \end{aligned}$$

Now choosing $\gamma = \frac{(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}{R^2(1+BR)}$ gives

$$\mathbb{E}[\hat{\ell}(x, y, \theta_t)] \leq \frac{R^2(1+BR) [\mathbb{E}\|\theta_t - \theta^*\|^2 - \mathbb{E}\|\theta_{t+1} - \theta^*\|^2]}{\left[(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}\right]^2} + \left(\frac{2(1+R\|\theta_\star\|)(1+BR)}{(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}\right) \eta \quad (63)$$

Summing the above equation for $t = 1$ to n and applying Jensen's inequality give us,

$$\begin{aligned} \mathbb{E}[\hat{\ell}(x, y, \bar{\theta}_n)] &\leq \frac{R^2(1+BR)\|\theta_1 - \theta^*\|^2}{\left[(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}\right]^2 n} + \underbrace{\left(\frac{2(1+R\|\theta_\star\|)(1+BR)}{(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}}\right)}_{\approx \text{constant}} \eta \\ &= \frac{R^2(1+BR)\|\theta_1 - \theta^*\|^2}{\left[(1-\eta) \min\left\{\frac{1}{\mu}, \frac{\rho^*-1}{1+\mu}\right\}\right]^2 n} + O(\eta). \end{aligned} \quad (64)$$

□