A Proof for Algorithm 1

We start by stating a lemma below involving our assumption on $\mu(\cdot), \beta, x$ that implies the instantaneous one-sided loss is bounded by $C_{T,\delta} := LBM + \gamma + c + \phi\sqrt{\log(2T/\delta)}$ with probability $1 - \delta/T$.

Lemma 1. *Given $\|x\|_2 \le B$, $\|\beta\|_2 \le M$, ϵ zero-mean subgaussian with parameter ϕ and $\mu(\cdot)$ has Lipschitz constant L with $\mu(0) \le \gamma$, we have the following tail bound on the empirical instantaneous one-sided loss*

$$u(x, a) := |\mu(x^\top \beta^*) + \epsilon - c| \cdot \mathbb{1}\Big\{\mathbb{1}\{\mu(x^\top \beta^*) + \epsilon > c\} \neq a\Big\},$$

as

$$\mathbb{P}(u(x, a) \ge LBM + \gamma + c + \tau) \le 2 \exp\Big(-\frac{\tau^2}{\phi^2}\Big)$$

for all $\tau > 0$.

Proof. By Cauchy-Schwartz and Lipschitz assumption, we have $|\mu(x^\top \beta^*) - c| \le LBM + \gamma + c$. Now by triangle inequality since $|u| \le LBM + \gamma + c + |\epsilon|$ and ϵ is subgaussian with parameter ϕ,

$$\mathbb{P}(|u| \ge LBM + \gamma + c + \tau) \le \mathbb{P}(|\epsilon| \ge \tau) \le 2 \exp\Big(-\frac{\tau^2}{\phi^2}\Big)$$

for all $\tau > 0$. ◻

All the analysis that follow will condition on this event, where we have that the one-sided loss is bounded by $C_{T,\delta}$ for all $t \le T$ with probability exceeding $1 - \delta$.

Before proceeding, we give a lemma below that characterizes the optimal solution to the expected one-sided loss minimization problem on the population level.

Lemma 2. *The optimal strategy for the expected one-sided loss minimization problem satisfies*

$$\min_{\pi\in\Pi} \mathbb{E}_{\mathcal{P}}[u(x, \pi(x))] = \mathbb{E}_{\mathcal{P}}\Big[|y - c| \cdot \mathbb{1}\Big\{\mathbb{1}\{y > c\} \neq \mathbb{1}\{\mu(x^\top \beta^*) > c\}\Big\}\Big]$$

for strategy class $\Pi = \{\pi^\beta : \|\beta\|_2 \le M\}$, where $\pi^\beta(x) := \mathbb{1}\{\mu(x^\top\beta) > c\}$ and expectation is taken over data that follows $y = \mu(x^\top \beta^*) + \epsilon$. In other words, $a = \mathbb{1}\{\mu(x^\top \beta^*) > c\}$ is the optimal strategy for the objective at population level.

Proof. We can rewrite the objective in terms of β as

$$\begin{aligned}\min_{\pi\in\Pi} \mathbb{E}_{\mathcal{P}}[u(x,\pi(x))] &= \min_{\beta:\|\beta\|_2\le M} \mathbb{E}_{\mathcal{P}}\Big[(y - c) \cdot \Big(\mathbb{1}\{y > c\} - \mathbb{1}\{\mu(x^\top\beta) > c\}\Big)\Big] \\ &= \min_{\beta:\|\beta\|_2\le M} \mathbb{E}_{\mathcal{P}}\Big[(\mu(x^\top\beta^*) + \epsilon - c) \cdot \Big(\mathbb{1}\{\mu(x^\top\beta^*) + \epsilon > c\} - \mathbb{1}\{\mu(x^\top\beta) > c\}\Big)\Big].\end{aligned}$$

Therefore it suffices to show that

$$\beta^* = \operatorname*{argmax}_{\beta:\|\beta\|_2\le M} \mathbb{E}_{\mathcal{P}}\Big[(\mu(x^\top\beta^*) + \epsilon - c) \cdot \mathbb{1}\{\mu(x^\top\beta) > c\}\Big].$$

As ϵ is zero-mean and independent of x by assumption, we have

$$\mathbb{E}_{\mathcal{P}}\Big[(\mu(x^\top\beta^*) + \epsilon - c) \cdot \mathbb{1}\{\mu(x^\top\beta) > c\}\Big] = \mathbb{E}_{\mathcal{P}}\Big[(\mu(x^\top\beta^*) - c) \cdot \mathbb{1}\{\mu(x^\top\beta) > c\}\Big],$$

and the claim above immediately follows. ◻

With this in hand, we are ready to show the one-sided loss bound for the offline learner in Algorithm 1.

Proof of Proposition 1. To construct an ϵ-cover $\hat{\Pi}$ in the pseudo-metric $\rho(\pi, \hat{\pi}) = \mathbb{P}(\pi(x) \neq \hat{\pi}(x))$ for i.i.d feature-utility pairs $(x_t, u_t) \sim \mathcal{P}$, since the linear threshold functions in $\mathbb{R}^d$ has a VC dimension of $d+1$, a standard argument with Sauer's lemma (see e.g. (Beygelzimer et al., 2011)) concludes that for sequences $x_1, \cdots, x_T$ drawn i.i.d, with probability $1-\delta/2$ over a random subset of size K,

$$\min_{\pi \in \hat{\Pi}} \mathbb{E}_{\mathcal{P}(u,x)}\Big[\sum_{t=1}^{T} u_t(x_t, \pi(x_t))\Big] \leq \min_{\pi \in \Pi} \mathbb{E}_{\mathcal{P}(u,x)}\Big[\sum_{t=1}^{T} u_t(x_t, \pi(x_t))\Big] + \frac{C_{T,\delta}T}{K}\Big(2(d+1)\log\Big(\frac{eT}{(d+1)}\Big) + \log\Big(\frac{2}{\delta}\Big)\Big)$$

for $|\hat{\Pi}| = (\frac{eT}{d+1})^{(d+1)}$. Now running the passive algorithm for the discretized strategy class $\hat{\Pi}$ on the exploration data we collected in the first phase (consists of $K+S$ rounds), we have for a fixed straegy $\hat{\pi} \in \hat{\Pi}$, since the $K+S$ terms are i.i.d and unbiased,

$$\mathbb{E}_{\mathcal{P}(u,x)}\Big[\sum_{t=1}^{K+S} u_t(x_t, \hat{\pi}(x_t))\Big] = (K+S) \cdot \mathbb{E}_{\mathcal{P}(u,x)}[u(x, \hat{\pi}(x))]\,.$$

Now via a Hoeffding's inequality for bounded random variables and a union bound over $|\hat{\Pi}|$, with probability at least $1-\delta$, simultaneously for all $\hat{\pi} \in \hat{\Pi}$,

$$\Big|\frac{1}{K+S}\sum_{t=1}^{K+S} u_t(x_t, \hat{\pi}(x_t)) - \mathbb{E}_{\mathcal{P}(u,x)}[u(x, \hat{\pi}(x))]\Big| \leq \sqrt{\frac{C_{T,\delta}^2}{2(K+S)}\log\Big(\frac{2|\hat{\Pi}|}{\delta}\Big)}\,.$$

Therefore applying the inequality twice with $\hat{\pi}_K := \hat{\pi}_K^{\hat{\beta}^*}$ and $\hat{\pi}^* := \min_{\pi \in \hat{\Pi}} \mathbb{E}_{\mathcal{P}}[u(x, \pi(x))]$, and using the optimality of $\hat{\pi}_K$ as the empirical minimizer, we have for each round,

$$\mathbb{E}_{\mathcal{P}(u,x)}[u(x, \hat{\pi}_K(x))] \leq \mathbb{E}_{\mathcal{P}(u,x)}[u(x, \hat{\pi}^*(x))] + 2\sqrt{\frac{C_{T,\delta}^2}{2(K+S)}\Big(\log\Big(\frac{4}{\delta}\Big) + (d+1)\log\Big(\frac{eT}{d+1}\Big)\Big)}$$

with probability at least $1-\delta/2$. Now summing up over T rounds, putting together the inequalities established and minimizing over K and S, gives the final utility bound as

$$\begin{aligned}\sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[u(x, a_t)] &\leq \min_{\pi \in \Pi}\sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[u(x, \pi(x))] + T\sqrt{\frac{2C_{T,\delta}^2}{(K+S)}\Big(\log\Big(\frac{4}{\delta}\Big) + (d+1)\log\Big(\frac{eT}{d+1}\Big)\Big)} \\ &+ C_{T,\delta}(K+S) + \frac{C_{T,\delta}T}{K}\Big(2(d+1)\log\Big(\frac{eT}{(d+1)}\Big) + \log\Big(\frac{2}{\delta}\Big)\Big) \\ &= \min_{\pi \in \Pi}\sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[u(x, \pi(x))] + \mathcal{O}\Big(C_{T,\delta}T^{2/3} d\log\Big(\frac{T}{d\delta}\Big)\Big)\,,\end{aligned}$$

with probability at least $1-\delta$. This in turn gives the one-sided loss bound

$$\begin{aligned}&\sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[u(x, a_t)] - \min_{\pi \in \Pi}\sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[u(x, \pi(x))] \\ &= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[|y-c| \cdot \mathbb{1}\Big\{\mathbb{1}\{y > c\} \neq a_t\Big\} - |y-c| \cdot \mathbb{1}\Big\{\mathbb{1}\{y > c\} \neq \mathbb{1}\{\mu(x^\top\beta^*) > c\}\Big\}\Big] \\ &= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[(y-c) \cdot \Big(\mathbb{1}\{\mu(x^\top\beta^*) > c\} - a_t\Big)\Big] \\ &= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[(\mu(x^\top\beta^*) - c) \cdot \Big(\mathbb{1}\{\mu(x^\top\beta^*) > c\} - a_t\Big)\Big] + \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[\epsilon \cdot \Big(\mathbb{1}\{\mu(x^\top\beta^*) > c\} - a_t\Big)\Big] \\ &= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[|\mu(x^\top\beta^*) - c| \cdot \mathbb{1}\Big\{\mathbb{1}\{\mu(x^\top\beta^*) > c\} \neq a_t\Big\}\Big] + \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}\Big[\epsilon \cdot \Big(\mathbb{1}\{\mu(x^\top\beta^*) > c\} - a_t\Big)\Big]\end{aligned}$$

$$= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{P}}[r_t] = \mathcal{O}\Big(C_{T,\delta} T^{2/3} d \log\Big(\frac{T}{d\delta}\Big)\Big)$$

with the same probability, where we used Lemma 2 in the first equality and the fact that ϵ is zero-mean and independent of x for the last step. □

B Proof for Section 5

The following Gaussian anti-concentration bound is used throughout the proof.

Lemma 3. *For $X \sim \mathcal{N}(\mu, \sigma^2)$, we have the lower bound on Gaussian density as:*

$$\mathbb{P}\left(X - \mu < -t\right) \geq \frac{1}{\sqrt{2\pi}} \frac{\sigma t}{t^2 + \sigma^2} e^{-\frac{t^2}{2\sigma^2}} .$$

B.1 Proof of Theorem 1

Proof of Theorem 1. The optimality condition for MLE fit $\hat{\beta}$ gives that $\sum_{i=1}^{n} y_i \cdot x_i = \sum_{i=1}^{n} \mu(x_i^\top \hat{\beta}) \cdot x_i$, which implies

$$\sum_{i=1}^{n} x_i \cdot \epsilon_i = \sum_{i=1}^{n} x_i \cdot \left[\mu(x_i^\top \hat{\beta}) - \mu(x_i^\top \beta^*)\right] .$$

In the direction v, where v is orthogonal to every other vector drawn from P, we have for n' the number of times v has appeared in the n samples used for warm-starting the greedy learner,

$$\sum_{i=1}^{n'} v \cdot \epsilon_i = \sum_{i=1}^{n'} v \cdot \left[\mu(v^\top \hat{\beta}) - \mu(v^\top \beta^*)\right]$$

and therefore assuming $\epsilon_i \sim \mathcal{N}(0,1)$, we have $\mu(v^\top \hat{\beta}) \sim \mathcal{N}(\mu(v^\top \beta^*), \frac{1}{n'})$. If $\mu(v^\top \beta^*) = c + \tau$ for $0 < \tau \leq 1/\sqrt{n'}$, by anticoncentration property of the Gaussian distribution (Lemma 3):

$$\mathbb{P}\left(\mu(v^\top \hat{\beta}) < c\right) \geq \frac{1}{\sqrt{2\pi}} \frac{\tau/\sqrt{n'}}{1/n' + \tau^2} \exp\left(-\frac{\tau^2}{2/n'}\right) \geq 1/10 .$$

Hence, with constant probability the greedy procedure will reject v at the next round. From then on-wards, the greedy learner will reject all instances of v and will incur a loss of τ every time it encounters it in any subsequent round. □

B.2 Linear One-Sided Loss for Empirical One-Sided Loss Minimization

Let $(x)_+ = \max(0, x)$. For observed pairs of (x_i, y_i) denoted by set $\mathcal{S}$, the empirical risk minimization of our one-sided loss can be rewritten as

$$\arg\min_{\beta} \sum_{i \in \mathcal{S}} |y_i - c| \cdot \frac{\big((x_i^\top \beta - c)(c - y_i)\big)_+}{(x_i^\top \beta - c)(c - y_i)},$$

from which it's obvious that it's not convex in β. However, in the case $x \in \{e_i\}_{i=1}^{d}$ (or any other orthogonal system), the coordinate decouples (after rotation) and the problem becomes solving d problems in 1D as

$$\arg\min_{\beta_i} \sum_{j \in \mathcal{S}: x_j = e_i} \frac{\big((\beta_i - c)(c - y_j)\big)_+}{|\beta_i - c|}$$

and the following procedure would find the optimal solution to the problem: (1) for all y_i such that $y_i \leq c$, compute $l = \sum_{i: y_i \leq c} (c - y_i)$; (2) similarly compute $u = \sum_{i: y_i \geq c} (y_i - c)$; (3) compare the two quantities if $l < u$, any $\beta_i > c$ would be a global optimum for the problem with objective function value l, otherwise any $\beta_i < c$

would be a global optimum for the problem with objective function value u. Therefore for group i ($x = e_i$) with response $y_i \sim \mathcal{N}(c + \tau_i, \sigma_i^2)$, where we initialize with t observed samples, again using Lemma 3,

$$\begin{aligned}
\mathbb{P}(u < l) &= \mathbb{P}\Big(\sum_{t: y_i^t \geq c} (y_i^t - c) < \sum_{t: y_i^t \leq c} (c - y_i^t) \Big) \\
&= \mathbb{P}\Big(\sum_t (y_i^t - c) < 0 \Big) \\
&= \mathbb{P}\Big(\mathcal{N}(t\tau_i, t\sigma_i^2) < 0 \Big) \\
&\geq \frac{1}{2\pi} \frac{\sqrt{t}\sigma_i \tau_i}{t\tau_i^2 + \sigma_i^2} \exp\Big(-\frac{t\tau_i^2}{2\sigma_i^2} \Big) > 1/5
\end{aligned}$$

for $\tau_i = 1/\sqrt{t}$ and $\sigma_i = 1$, after which no observations will be made on group i as $\hat{\beta}_i < c$ and linear one-sided loss will be incurred with constant probability.

B.3 Simulation

We use the example provided in Theorem 1 with $n = 10000$ and $d = 20$. The x axis shows the number of rounds (t) (i.e., number of batches) and the y axis shows the average one-sided loss R_t/t. Batch size is chosen to be 1 for linear regression and 100 for logistic regression. The average one-sided loss fails to decrease for the greedy method but our method (Algorithm 2) exhibits vanishing one-sided loss.

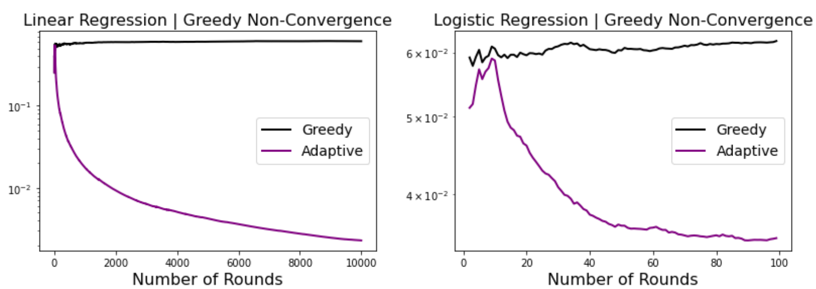

Figure 3: Examples when greedy fails to converge where th y-axis is the average one-sided loss up to the round indicated in the x-axis.

C Proof for Algorithm 2

Lemma 4. *Suppose that Assumption 1 holds. Let $B > 0$. Then there exists $\eta > 0$ such that $\mu'(x^\top \beta) \geq \eta$ for all $x, \beta \in \mathbb{R}^d$ satisfying $\|x\|_2 \leq B$ and $\|\beta\|_2 \leq M$.*

Proof. By Assumption 1, we have μ' is continuous and positive everywhere. Define the interval $\mathcal{I} := [-B \cdot M, B \cdot M]$. We have by Cauchy-Schwarz that $x^\top \beta \in \mathcal{I}$. Since $\mathcal{I}$ is closed and bounded, by Heine–Borel $\mathcal{I}$ is compact in $\mathbb{R}$. Since the image of a continuous function on a compact set is also compact, it follows that there exists $\eta > 0$ such that $\mu'(x) \geq \eta$ for all $x \in \mathcal{I}$, as desired. □

Much of the analysis in the lemma below is built upon (Filippi et al., 2010), generalized to our setting.

Lemma 5 (Instantaneous One-Sided Loss)**.** *For all $1 \leq t \leq T$ and some $0 < \delta < \min(1, d/e)$, with $A_0 := \sum_{i=1}^K x_i^0 x_i^{0\top} \succeq \lambda_0 \cdot I_d$ and $A_t := A_0 + \sum_{i=1}^t X_i^\top X_i$, we have under the assumption stated in Theorem 2 that*

$$r_t \leq 4\frac{L}{\eta} \kappa C_{T,\delta} \sqrt{2d \log t} \sqrt{\log(2dT/\delta)} \cdot \sum_{i=1}^N \sqrt{\bar{x}_i^{t\top} A_{t-1}^{-1} \bar{x}_i^t}$$

with probability at least $1 - \delta$ for $\kappa = \sqrt{3 + 2\log(1 + 2NB^2/\lambda_0)}$, where $\bar{x}_i^t$ is either x_i^t or 0_d depending on whether we choose to observe the context.

Proof. We recast the problem as picking 1 out of 2^N choices (induced by all possible binary decisions on each of the N samples in the batch) in each round with linear reward function. To this end, for each feature vector $x \in \mathbb{R}^d$, we encode the algorithm's two choices as $(d+1)$-dimensional vectors $[0; x]$ (selecting it) and $[1; 0_d]$ (not selecting it). Let us denote $c' = \mu^{-1}(c)$, where c is the cutoff. Then for each $i \in [N]$ at round t, OPT chooses x_i^t to predict positively and observe y_i^t if $\mu([0; x_i^t] \cdot \tilde{\beta}^*) = \mu(0 \cdot c' + x_i^{t\top} \beta^*)$ exceeds $\mu([1; 0_d] \cdot \tilde{\beta}^*) = \mu(1 \cdot c' + 0_d^\top \beta^*)$ where $\tilde{\beta}^* := [c'; \beta^*]$. We can then define the following notation X_t representing Algorithm 2's choices at round t:

$$X_t := \underset{X \in \mathbb{R}^{N \times (d+1)}:\ x_i \in \{[0; x_i^t], [1; 0_d]\}\ \forall i \in [N]}{\operatorname{argmax}} 1^\top \mu(X \tilde{\beta}_t) + \rho_t(\delta/2T) \cdot \sum_{i=1}^N \sqrt{x_i^{[2:d+1]\top} A_{t-1}^{-1} x_i^{[2:d+1]}} \tag{4}$$

where $\tilde{\beta}_t := [c'; \beta_t] \in \mathbb{R}^{d+1}$ with β_t as the MLE fit on (x_i, y_i) pairs observed so far. Note that this is the same as the X_t one would get from Algorithm 2 up to padding of 0 at front for the observed contexts and appending vector $[1, 0_d]$ for the unobserved ones.

For any feasible context matrix $X \in \mathbb{R}^{N \times (d+1)}$ at round t, using the Lipschitz assumption on $\mu(\cdot)$, and denote $\tilde{\beta}^* = [c'; \beta^*] \in \mathbb{R}^{d+1}$, we have

$$\begin{aligned}
|1^\top \mu(X\tilde{\beta}^*) - 1^\top \mu(X\tilde{\beta}_t)| &\leq L\|X(\tilde{\beta}^* - \tilde{\beta}_t)\|_1 \\
&= L\|\bar{X}\beta^* - \bar{X}\beta_t\|_1
\end{aligned}$$

where we used that both β's have constant c' in the first coordinate, so the problem is reduced to looking at the last d coordinates, defined as $\bar{X}^*$ and $\bar{X}_t$ respectively. Now let $g_t(\beta) := \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta)$, Mean Value Theorem gives that

$$\begin{aligned}
g_t(\beta^*) - g_t(\beta_t) &= \int_0^1 \nabla g_t(s\beta^* + (1-s)\beta_t)\, ds \cdot (\beta^* - \beta_t) \\
&= \int_0^1 \sum_{i=1}^{t-1} \bar{X}_i^\top \mu'\Big(\bar{X}_i(s\beta^* + (1-s)\beta_t)\Big) \bar{X}_i\, ds \cdot (\beta^* - \beta_t) \\
&=: G_t \cdot (\beta^* - \beta_t)
\end{aligned}$$

where G_t satisfies $G_t \succeq \eta A_{t-1} \succ 0$ since the middle term $\mu'(\cdot)$ can be seen as a diagonal matrix with entries $\geq \eta$ by Lemma 4. Therefore we have for some $v \in \{\pm 1\}^N$ by (1) Cauchy-Schwarz (as $G_t^{-1} \succ 0$); (2) triangle inequality; (3) β_t is optimal for the projection problem,

$$\begin{aligned}
|1^\top \mu(X\tilde{\beta}^*) - 1^\top \mu(X\tilde{\beta}_t)| &\leq L\|X(\tilde{\beta}^* - \tilde{\beta}_t)\|_1 \\
&= L\|\bar{X} G_t^{-1}(g_t(\beta^*) - g_t(\beta_t))\|_1 \\
&\leq L\|g_t(\beta^*) - g_t(\beta_t)\|_{G_t^{-1}} \|\bar{X}^\top v\|_{G_t^{-1}} \\
&\leq \frac{L}{\eta}\|g_t(\beta^*) - g_t(\beta_t)\|_{A_{t-1}^{-1}} \|\bar{X}^\top v\|_{A_{t-1}^{-1}} \\
&\leq \frac{L}{\eta}\|g_t(\beta^*) - g_t(\beta_t)\|_{A_{t-1}^{-1}} \cdot \sum_{i=1}^N \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&= \frac{L}{\eta}\Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta^*) - \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta_t) \Big\|_{A_{t-1}^{-1}} \cdot \sum_{i=1}^N \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&\leq \frac{L}{\eta}\Big(\Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta^*) - \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \hat{\beta}_t) \Big\|_{A_{t-1}^{-1}} \\
&\qquad + \Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta_t) - \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \hat{\beta}_t) \Big\|_{A_{t-1}^{-1}} \Big) \cdot \sum_{i=1}^N \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&= \frac{L}{\eta}\Big(\Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i \beta^*) - \sum_{i=1}^{t-1} \bar{X}_i^\top y_i \Big\|_{A_{t-1}^{-1}}
\end{aligned}$$

$$
\begin{aligned}
&+\Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i\beta_t) - \sum_{i=1}^{t-1} \bar{X}_i^\top \mu(\bar{X}_i\hat{\beta}_t) \Big\|_{A_{t-1}^{-1}} \Big) \cdot \sum_{i=1}^{N} \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&\le \frac{2L}{\eta} \Big\| \sum_{i=1}^{t-1} \bar{X}_i^\top (\mu(\bar{X}_i\beta^*) - y_i) \Big\|_{A_{t-1}^{-1}} \cdot \sum_{i=1}^{N} \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&\le \frac{2L}{\eta} \kappa C_{T,\delta} \sqrt{2d \log t} \sqrt{\log(d/\delta)} \cdot \sum_{i=1}^{N} \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i} \\
&=: \zeta_t^{\bar{X}}(\delta) =: \rho_t(\delta) \cdot \sum_{i=1}^{N} \sqrt{\bar{x}_i^\top A_{t-1}^{-1} \bar{x}_i}
\end{aligned}
$$

where we used Lemma 1 from (Filippi et al., 2010) on vector-valued martingales for bounding the first term in the last step, which holds with probability at least $1-\delta$ for $0 < \delta < \min(1, d/e)$ and $\kappa := \sqrt{3 + 2\log(1 + 2NB^2/\lambda_0)}$. We used the fact that (1) the N context vectors $\{\bar{x}_i\}$ are independent of each other in each round; (2) each of the error term $\mu(\bar{X}_i\beta^*) - y_i$ has sub-gaussian tail by assumption.

Leveraging this, for the instantaneous one-sided loss, we get with probability at least $1-\delta/T$ (denote $X_t^* \in \mathbb{R}^{N\times(d+1)}$ as the context chosen by the best action at time t)

$$
\begin{aligned}
r_t &= 1^\top(\mu(X_t^*\tilde{\beta}^*) - \mu(X_t\tilde{\beta}^*)) \\
&= 1^\top(\mu(X_t^*\tilde{\beta}^*) - \mu(X_t^*\tilde{\beta}_t)) + 1^\top(\mu(X_t^*\tilde{\beta}_t) - \mu(X_t\tilde{\beta}_t)) + 1^\top(\mu(X_t\tilde{\beta}_t) - \mu(X_t\tilde{\beta}^*)) \\
&\le \zeta_t^{\bar{X}_t^*}(\delta/2T) + \zeta_t^{\bar{X}_t}(\delta/2T) + 1^\top(\mu(X_t^*\tilde{\beta}_t) - \mu(X_t\tilde{\beta}_t)) \\
&= \zeta_t^{\bar{X}_t^*}(\delta/2T) + \zeta_t^{\bar{X}_t}(\delta/2T) + 1^\top\mu(X_t^*\tilde{\beta}_t) + \zeta_t^{\bar{X}_t^*}(\delta/2T) - 1^\top\mu(X_t\tilde{\beta}_t) - \zeta_t^{\bar{X}_t^*}(\delta/2T) \\
&\le \zeta_t^{\bar{X}_t^*}(\delta/2T) + \zeta_t^{\bar{X}_t}(\delta/2T) + 1^\top\mu(X_t\tilde{\beta}_t) + \zeta_t^{\bar{X}_t}(\delta/2T) - 1^\top\mu(X_t\tilde{\beta}_t) - \zeta_t^{\bar{X}_t^*}(\delta/2T) \\
&= 2\zeta_t^{\bar{X}_t}(\delta/2T)
\end{aligned}
$$

where we used the optimality of X_t for (4) in the last inequality. Union bounding over T time steps yields the claim. ☐

Below we state a helper lemma for bounding the second term $\sum_{i=1}^N \|\bar{x}_i^t\|_{A_{t-1}^{-1}}$ from the previous lemma.

Lemma 6 (Helper Lemma)**.** *For all $T \ge 1$, let $A_0 = \sum_{i=1}^K x_i^0 x_i^{0\top} \succeq \lambda_0 \cdot I_d$ and $A_t = A_0 + \sum_{i=1}^t X_i^\top X_i$, for $s = \min(N, d)$, under the assumption $\|x_i^t\|_2 \le B \; \forall i, t$,*

$$
\sum_{t=1}^{T} \min\Big\{ \sum_{i=1}^{N} \|\bar{x}_i^t\|^2_{A_{t-1}^{-1}}, N \Big\} \le 2dNs \log\Big(\frac{\lambda_0 + B^2NT}{d}\Big) - 2dNs\log(\lambda_0)\,,
$$

where $\bar{x}_i^t$ is either x_i^t or 0_d depending on whether we choose to observe the context.

Proof. We have from the Matrix Determinant Lemma that

$$
\begin{aligned}
\det(A_t) = \det(A_{t-1} + X_t^\top X_t) &= \det(A_{t-1})\det(I_N + X_t(A_{t-1})^{-1}X_t^\top) \\
&= \det(A_{t-1}) \prod_{i=1}^{s} (1 + \lambda_i^t) \\
&= \det(A_0) \prod_{j=1}^{t}\prod_{i=1}^{s} (1 + \lambda_i^j),
\end{aligned}
$$

where we denoted the nonzero eigenvalues of the PSD matrix $X_t(A_{t-1})^{-1}X_t^\top$ as $\{\lambda_i^t\}_{i=1}^s$. Since $\{\lambda_i^t\}_{i=1}^s$ are eigenvalues of $(A_{t-1})^{-1}$ restricted to $\operatorname{span}\{x_i^t\}_{i=1}^N$, we have

$$
\sum_{i=1}^{N} x_i^{t\top}(A_{t-1})^{-1}x_i^t \le N \cdot \max_i(\lambda_i^t) \le N \sum_{i=1}^{s} \lambda_i^t\,. \tag{5}
$$

Now using that $x \leq 2\log(1+x)$ for $x \in [0,1]$,

$$\sum_{j=1}^{t}\sum_{i=1}^{s}\min\{1/s, \lambda_i^t\} \leq 2\sum_{j=1}^{t}\sum_{i=1}^{s}\log(1+\lambda_i^t) = 2(\log\det(A_t) - \log\det(A_0))\,.$$

Since $\mathrm{Trace}(A_t) = \lambda_0 + \sum_{j=1}^t \mathrm{Trace}(X_j^\top X_j) \leq \lambda_0 + B^2 Nt$ if all covariates are bounded as $\|x_i^t\|_2 \leq B$. Therefore from AM-GM inequality, since the determinant is the product of the eigenvalues, we have

$$\sum_{j=1}^{t}\sum_{i=1}^{s}\min\{1/s, \lambda_i^t\} \leq 2d\log\left(\frac{\lambda_0 + B^2 Nt}{d}\right) - 2d\log(\lambda_0)\,.$$

Implying from (5) that

$$\begin{aligned}\sum_{t=1}^{T}\min\left\{\sum_{i=1}^{N}\bar{x}_i^{t\top}A_{t-1}^{-1}\bar{x}_i^t, N\right\} &\leq N\sum_{t=1}^{T}\min\left\{\sum_{i=1}^{s}\lambda_i^t, 1\right\} \leq Ns\sum_{t=1}^{T}\sum_{i=1}^{s}\min\{\lambda_i^t, 1/s\}\\ &\leq 2dNs\log\left(\frac{\lambda_0 + B^2 NT}{d}\right) - 2dNs\log(\lambda_0)\,,\end{aligned}$$

as claimed. $\square$

We are now ready to put things together to give the final one-sided loss bound for our algorithm.

Proof of Theorem 2. For the cumulative one-sided loss, using Lemma 5 and Lemma 6 above, and the conditional event that the instantaneous one-sided loss is bounded by $C_{T,\delta}$, with probability at least $1-\delta$,

$$\begin{aligned}R_T &\leq C_{T,\delta}\cdot K + \sum_{t=1}^{T} r_t\\ &\leq C_{T,\delta}\cdot K + \sum_{t=1}^{T}\min\left\{2\rho_t(\delta/2T)\cdot\sum_{i=1}^{N}\sqrt{\bar{x}_i^{t\top}A_{t-1}^{-1}\bar{x}_i^t}, NC_{T,\delta}\right\}\\ &\leq C_{T,\delta}\cdot K + 2\rho_T(\delta/2T)\cdot\sum_{t=1}^{T}\min\left\{\sum_{i=1}^{N}\sqrt{\bar{x}_i^{t\top}A_{t-1}^{-1}\bar{x}_i^t}, N\right\}\\ &\leq C_{T,\delta}\cdot K + 2\rho_T(\delta/2T)\sqrt{TN}\cdot\sqrt{\sum_{t=1}^{T}\min\left\{\sum_{i=1}^{N}\bar{x}_i^{t\top}A_{t-1}^{-1}\bar{x}_i^t, N\right\}}\\ &\leq C_{T,\delta}\cdot K + 2\rho_T(\delta/2T)\sqrt{TN}\cdot\sqrt{2dNs\log\left(\frac{\lambda_0 + B^2 NT}{d}\right) - 2dNs\log(\lambda_0)}\end{aligned}$$

where we used the fact that $NC_{T,\delta} \leq 2N\rho_T(\delta/2T)$ and Cauchy-Schwarz. Plugging in the definition of $\rho_T(\delta/2T) = \frac{2L}{\eta}\kappa C_{T,\delta}\sqrt{2d\log T}\sqrt{\log(2dT/\delta)}$ finishes the proof. $\square$

D Iterative Method

Proof of Proposition 2. We begin by showing iterate contraction. Looking at the condition for potential gradient update, assuming that we have an upper bound as $\mathbb{E}[\|\beta_t - \beta^*\|_2] \leq d_t$ at iteration t, on the event that the distance to OPT satistifes $\|\beta^* - \beta_t\|_2 \leq \delta\cdot d_t$, which happens with probability at least $1-\frac{1}{\delta}$ by Markov's inequality,

$$\begin{aligned}\mathbb{1}\{\mu(x_t^\top\beta_t) + s_t \geq c\} &= \mathbb{1}\left\{\mu(x_t^\top\beta^*) + s_t \geq c + \mu(x_t^\top\beta^*) - \mu(x_t^\top\beta_t)\right\}\\ &\geq \mathbb{1}\left\{\mu(x_t^\top\beta^*) + s_t \geq c + L\|x_t\|_2\cdot\|\beta^* - \beta_t\|_2\right\}\\ &\geq \mathbb{1}\left\{\mu(x_t^\top\beta^*) \geq c\right\}\end{aligned}$$

where we used that the exploration bonus $s_t = L \cdot \delta \cdot d_t \|x_t\|_2$ and the Lipschitz condition of $\mu'(\cdot)$. Now since $\mu(x_t^\top \beta^*) - \mathbb{E}[\mu(x^\top \beta^*)]$ is $CL\|\beta^*\|\sigma$-subgaussian for some numerical constant C, i.e.,

$$\mathbb{P}_x \left\{ \mu(x_t^\top \beta^*) \geq c \right\} = \mathbb{P}_x \left\{ \mu(x_t^\top \beta^*) \geq \mathbb{E}_x[\mu(x^\top \beta^*)] - \zeta \right\} \geq 1 - e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}} ,$$

therefore with probability at least $1 - e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}} - \delta^{-1}$, we accept x_t and are presented with the corresponding response y_t. It remains to work out the update for d_t such that $\mathbb{E}[\|\beta_t - \beta^*\|_2] \leq d_t$ holds at all iterations (so that we can set s_t appropriately).

For this, we have if $\left| y_t - \mu(x_t^\top \beta_t) \right| \leq \alpha + B$, thanks to the projection that maintains $\mu'(\cdot) \geq \gamma$ and the norm bound assumption on noise ϵ_t,

$$\gamma |x_t^\top (\beta_t - \beta^*)| - B \leq \left| \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta_t) \right| - B \leq \left| y_t - \mu(x_t^\top \beta_t) \right| \leq \alpha + B \,, \tag{6}$$

implying that we are already accurate enough on this sample as $|x_t^\top (\beta_t - \beta^*)| \leq \frac{\alpha + 2B}{\gamma}$.

Otherwise if $\left| y_t - \mu(x_t^\top \beta_t) \right| > \alpha + B$, we have

$$L |x_t^\top (\beta_t - \beta^*)| + B > \left| \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta_t) \right| + B > \left| y_t - \mu(x_t^\top \beta_t) \right| > \alpha + B$$

therefore $|x_t^\top (\beta_t - \beta^*)| > \frac{\alpha}{L}$, and making a gradient update gives the contraction

$$\mathbb{E}[\|\beta_{t+1} - \beta^*\|_2^2 \,|\, \beta_t] \leq \|\beta_t - \beta^*\|_2^2 + \left[\eta^2 \cdot \mu'(z)^2 \|x_t\|_2^2 - 2\eta \cdot \mu'(z) \right] (x_t^\top (\beta_t - \beta^*))^2$$

at step t. Taking stepsize $\eta = \frac{1}{L\|x_t\|_2^2} < \frac{1}{\mu'(z) \cdot \|x_t\|_2^2}$, we have the distance to OPT progress recursion using (3) as

$$\mathbb{E}[\|\beta_{t+1} - \beta^*\|^2 \,|\, \beta_t] \leq \|\beta_t - \beta^*\|_2^2 - \frac{\alpha^2}{\|x_t\|_2^2 L^2} \,.$$

Using law of total expectation and putting everything together, at iteration t of the algorithm, with probability at least $1 - \rho := 1 - e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}} - \delta^{-1}$, either

$$\mathbb{E}[\|\beta_{t+1} - \beta^*\|^2] \leq \mathbb{E}[\|\beta_t - \beta^*\|_2^2] - \frac{\alpha^2}{\|x_t\|_2^2 L^2} \,, \tag{7}$$

or in light of (6),

$$\left| \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta_t) \right| \leq L \cdot |x_t^\top (\beta^* - \beta_t)| \leq \frac{L}{\gamma} (\alpha + 2B) \,,$$

otherwise with the remaining probability ρ the distance to OPT stays the same since we don't accept the sample. Equation (7) therefore prescribes that we should update d_t as

$$\mathbb{E}\left[\|\beta_t - \beta^*\| \right] \leq \sqrt{\mathbb{E}\left[\|\beta_t - \beta^*\|^2 \right]} \leq \left(\mathbb{E}[\|\beta_{t-1} - \beta^*\|_2^2] - \frac{\alpha^2}{\|x_{t-1}\|_2^2 L^2} \right)^{1/2} .$$

This concludes the first part of the claim. Turning to misclassification error, since again by sub-gaussianity,

$$\mathbb{P}_x\{\mu(x_t^\top \beta^*) \geq c\} = \mathbb{P}_x \left\{ \mu(x_t^\top \beta^*) \geq \mathbb{E}_x[\mu(x^\top \beta^*)] - \zeta \right\} \geq 1 - e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}} ,$$

we focus on this side of the error (i.e., false-negatives). In this case, we don't make a mistake w.r.t the oracle predictor $\mathbb{1}\{\mu(x_t^\top \beta^*) \geq c\}$ at step t if $\mu(x_t^\top \beta_t) + s_t \geq c$, in which case we predict "accept" and is revealed y_t. To bound the probability, we observe that on this event,

$$\begin{aligned}
\mathbb{P}(\mu(x_t^\top \beta_t) + s_t \geq c) &= \mathbb{P}(\mu(x_t^\top \beta_t) + \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta^*) + s_t \geq c) \\
&\geq \mathbb{P}(s_t \geq \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta_t)) \\
&= \mathbb{P}(L \cdot \delta \cdot d_t \|x_t\|_2 \geq \mu(x_t^\top \beta^*) - \mu(x_t^\top \beta_t)) \\
&\geq \mathbb{P}(\|\beta_t - \beta^*\|_2 \leq \delta \cdot d_t) \\
&\geq 1 - \frac{\mathbb{E}[\|\beta_t - \beta^*\|_2]}{\delta \cdot d_t} \geq 1 - \frac{1}{\delta} \,,
\end{aligned}$$

where we used the definition of s_t and Markov's inequality. Consequently, picking $\delta^{-1} = \rho - e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}}$ in setting s_t and assuming $\zeta \geq \sqrt{2L\|\beta^*\|^2\sigma^2\log(\rho^{-1})}$,

$$\begin{aligned}\mathbb{P}(\text{making a mistake at step } t) &\leq \mathbb{P}(\mu(x_t^\top\beta^*) < c) + \mathbb{P}(\mu(x_t^\top\beta_t) + s_t < c \,|\, \mu(x_t^\top\beta^*) \geq c) \\ &\leq e^{-\frac{\zeta^2}{2L^2\|\beta^*\|^2\sigma^2}} + \frac{1}{\delta} \leq \rho\end{aligned}$$

for any $\rho \in (0, 1)$. □

E Additional Experiment Results

We include additional tables and plots in this section to further support our findings, starting on next page.

Dataset	c	greedy	ϵ-grdy	os-ϵ-grdy	noise	os-noise	margin	ours
Adult	50%	239.45	236.34	211.74	230.77	165.77	162.31	**144.92**
	55%	175.88	175.46	170.21	175.26	143.64	135.9	**118.7**
	60%	140.18	138.37	138.13	138.53	126.76	125.37	**114.53**
	65%	129.29	128.39	128.26	126.8	123.48	120.94	**115.62**
	70%	134.74	134.18	133.8	131.66	132.39	132.67	**129.81**
	75%	145.99	145.14	146.09	**144.38**	146.26	145.28	145.86
	80%	188.48	186.69	186.03	185.47	**184.68**	185.44	186.2
	85%	246.93	244.71	244.93	243.14	243.77	**242.94**	243.16
	90%	318.33	295.72	**290.51**	279.49	293.3	294.4	293.14
	95%	179.24	146.21	158.73	**131.2**	148.85	131.75	152.95
Bank	50%	164.23	162.67	117.86	136.0	88.49	86.26	**74.64**
	55%	142.01	138.29	107.39	125.15	83.59	82.46	**71.68**
	60%	141.04	139.42	110.72	131.61	90.86	89.29	**81.02**
	65%	146.56	140.44	121.98	135.96	100.98	96.59	**94.46**
	70%	207.6	197.0	185.9	198.66	153.3	150.75	**137.24**
	75%	166.72	166.08	166.81	162.88	153.38	150.07	**142.5**
	80%	148.93	147.06	148.75	148.18	142.68	137.37	**134.17**
	85%	145.63	125.29	122.99	122.59	122.28	130.96	**119.19**
	90%	104.98	102.89	104.59	103.0	102.1	**101.1**	104.19
	95%	119.77	119.06	**116.17**	119.42	119.48	119.63	119.73
COMPAS	50%	41.56	36.67	36.93	36.93	28.09	28.12	**26.01**
	55%	41.71	38.18	39.22	37.72	33.47	31.85	**31.17**
	60%	44.83	44.02	42.78	42.2	34.84	36.56	**34.48**
	65%	47.33	40.04	40.06	38.23	35.44	33.7	**32.52**
	70%	41.66	39.16	39.61	39.87	38.03	36.98	**34.07**
	75%	46.14	40.49	41.12	37.84	35.11	34.27	**33.89**
	80%	45.8	45.25	44.88	44.58	43.97	41.11	**40.63**
	85%	58.59	54.62	50.54	50.52	43.56	46.24	**41.93**
	90%	33.42	32.1	33.66	**28.79**	31.71	31.91	30.47
	95%	20.38	20.34	20.37	20.31	20.38	**19.68**	20.35
Crime	50%	15.77	15.77	15.5	15.66	14.93	14.73	**13.95**
	55%	14.64	14.52	14.64	14.27	14.46	14.46	**14.07**
	60%	17.44	17.42	17.31	17.12	16.99	16.35	**15.95**
	65%	18.64	18.59	18.72	18.27	18.52	18.53	**18.15**
	70%	22.0	21.75	21.99	20.33	20.63	20.1	**19.19**
	75%	21.87	21.87	21.74	21.33	21.05	21.4	**20.75**
	80%	23.03	22.94	23.38	22.46	22.61	22.58	**21.87**
	85%	23.75	23.46	23.82	23.75	23.65	23.51	**22.68**
	90%	22.43	22.19	22.34	21.65	22.13	21.3	**20.88**
	95%	12.34	12.34	12.43	12.34	**12.21**	12.34	12.34
German	50%	14.7	14.51	14.12	13.62	11.12	10.52	**9.63**
	55%	12.43	12.3	12.24	12.42	11.06	10.98	**9.43**
	60%	15.16	14.48	14.2	14.09	13.83	13.32	**11.68**
	65%	17.02	16.39	16.52	15.82	13.75	13.86	**12.48**
	70%	15.89	15.53	15.93	15.41	14.09	14.52	**13.07**
	75%	15.44	15.26	15.13	14.86	14.49	14.95	**14.2**
	80%	12.8	12.69	12.87	12.61	12.68	12.63	**12.45**
	85%	11.55	11.45	11.23	11.38	11.27	11.23	**10.98**
	90%	10.09	9.96	10.14	**9.07**	9.84	9.97	9.93
	95%	8.23	8.13	8.23	**7.59**	7.97	8.18	8.22

Table 3: **Linear Regression Results**.

Dataset	c	greedy	ϵ-grdy	os-ϵ-grdy	noise	os-noise	margin	ours
Blood	50%	2.06	2.06	2.06	2.06	1.92	1.72	**1.52**
	55%	1.4	1.4	1.4	**1.36**	1.39	1.39	1.38
	60%	3.63	2.72	3.11	1.94	1.91	1.96	**1.87**
	65%	3.28	2.74	2.07	2.81	2.02	1.69	**1.59**
	70%	3.7	2.78	3.04	**2.38**	3.13	3.06	2.65
	75%	5.03	3.91	4.16	3.29	4.08	3.99	**3.13**
	80%	4.16	3.32	4.12	3.07	**3.06**	3.92	3.58
	85%	4.1	3.73	3.58	**3.28**	3.98	4.05	3.67
	90%	5.09	4.58	5.11	**3.97**	4.26	4.51	4.66
	95%	2.64	2.59	2.68	2.61	2.57	2.56	**2.55**
Diabetes	50%	4.17	4.16	4.23	3.94	3.81	3.95	**3.61**
	55%	4.93	4.93	4.97	4.88	4.79	4.9	**4.74**
	60%	6.01	6.01	5.92	5.83	5.75	5.97	**5.65**
	65%	5.45	5.45	5.48	5.29	5.32	5.3	**5.25**
	70%	6.05	5.56	6.14	6.05	5.6	5.39	**5.33**
	75%	7.61	7.57	6.74	6.69	6.21	6.35	**5.5**
	80%	8.18	7.9	8.24	7.64	8.01	7.06	**6.65**
	85%	6.84	6.84	6.84	6.84	6.75	6.69	**6.64**
	90%	5.78	5.73	5.86	5.47	5.65	5.53	**5.51**
	95%	4.52	4.52	4.56	4.36	4.37	4.29	**4.27**
EEG Eye	50%	256.47	200.04	175.8	173.52	106.26	**96.85**	119.7
	55%	227.08	191.27	169.19	177.03	118.03	**109.69**	128.09
	60%	196.1	169.52	163.42	155.87	121.5	121.09	**119.67**
	65%	162.28	159.8	154.73	148.5	133.16	130.01	**129.27**
	70%	175.71	167.94	168.73	157.68	167.52	160.76	**155.79**
	75%	157.93	147.06	154.61	146.3	**123.48**	124.6	136.1
	80%	164.19	140.15	139.47	133.81	142.57	**125.39**	149.71
	85%	143.94	125.09	118.29	**115.81**	117.24	131.08	136.51
	90%	121.78	116.0	121.3	**104.72**	115.71	115.27	117.57
	95%	149.06	**142.67**	145.99	139.5	151.72	148.77	149.09
Australian	50%	3.74	3.74	3.77	3.63	3.0	2.79	**2.65**
	55%	3.38	3.38	3.38	3.26	2.96	3.19	**2.69**
	60%	5.0	4.97	5.0	4.33	**3.73**	3.99	3.75
	65%	4.69	4.57	4.69	4.3	3.88	**3.84**	3.9
	70%	6.77	6.77	6.77	6.66	5.09	5.26	**4.65**
	75%	5.78	5.77	5.78	5.57	4.99	4.8	**4.77**
	80%	5.43	5.43	5.43	5.21	5.27	5.25	**4.85**
	85%	5.09	5.09	5.09	**5.01**	5.03	4.98	5.06
	90%	4.14	4.14	4.14	**3.93**	4.08	4.07	4.11
	95%	2.15	2.14	2.27	**2.08**	2.15	2.15	2.15
Churn	50%	46.98	43.65	30.65	36.64	21.24	18.83	**14.89**
	55%	57.49	51.72	47.36	46.52	30.75	**23.76**	24.39
	60%	61.44	56.94	50.43	55.59	35.65	32.82	**29.55**
	65%	40.83	38.89	37.29	37.67	29.44	29.33	**26.14**
	70%	49.99	47.84	47.91	49.89	41.18	36.17	**35.27**
	75%	58.96	56.91	58.34	56.99	55.42	53.88	**49.48**
	80%	52.66	50.43	51.85	51.23	48.62	48.74	**48.25**
	85%	52.41	50.66	49.02	**45.95**	50.76	49.67	49.89
	90%	60.33	**59.5**	60.37	59.98	59.78	59.82	59.97
	95%	56.36	54.33	54.7	**53.09**	54.29	55.91	56.39

Table 4: **Linear Regression Results (continued)**.

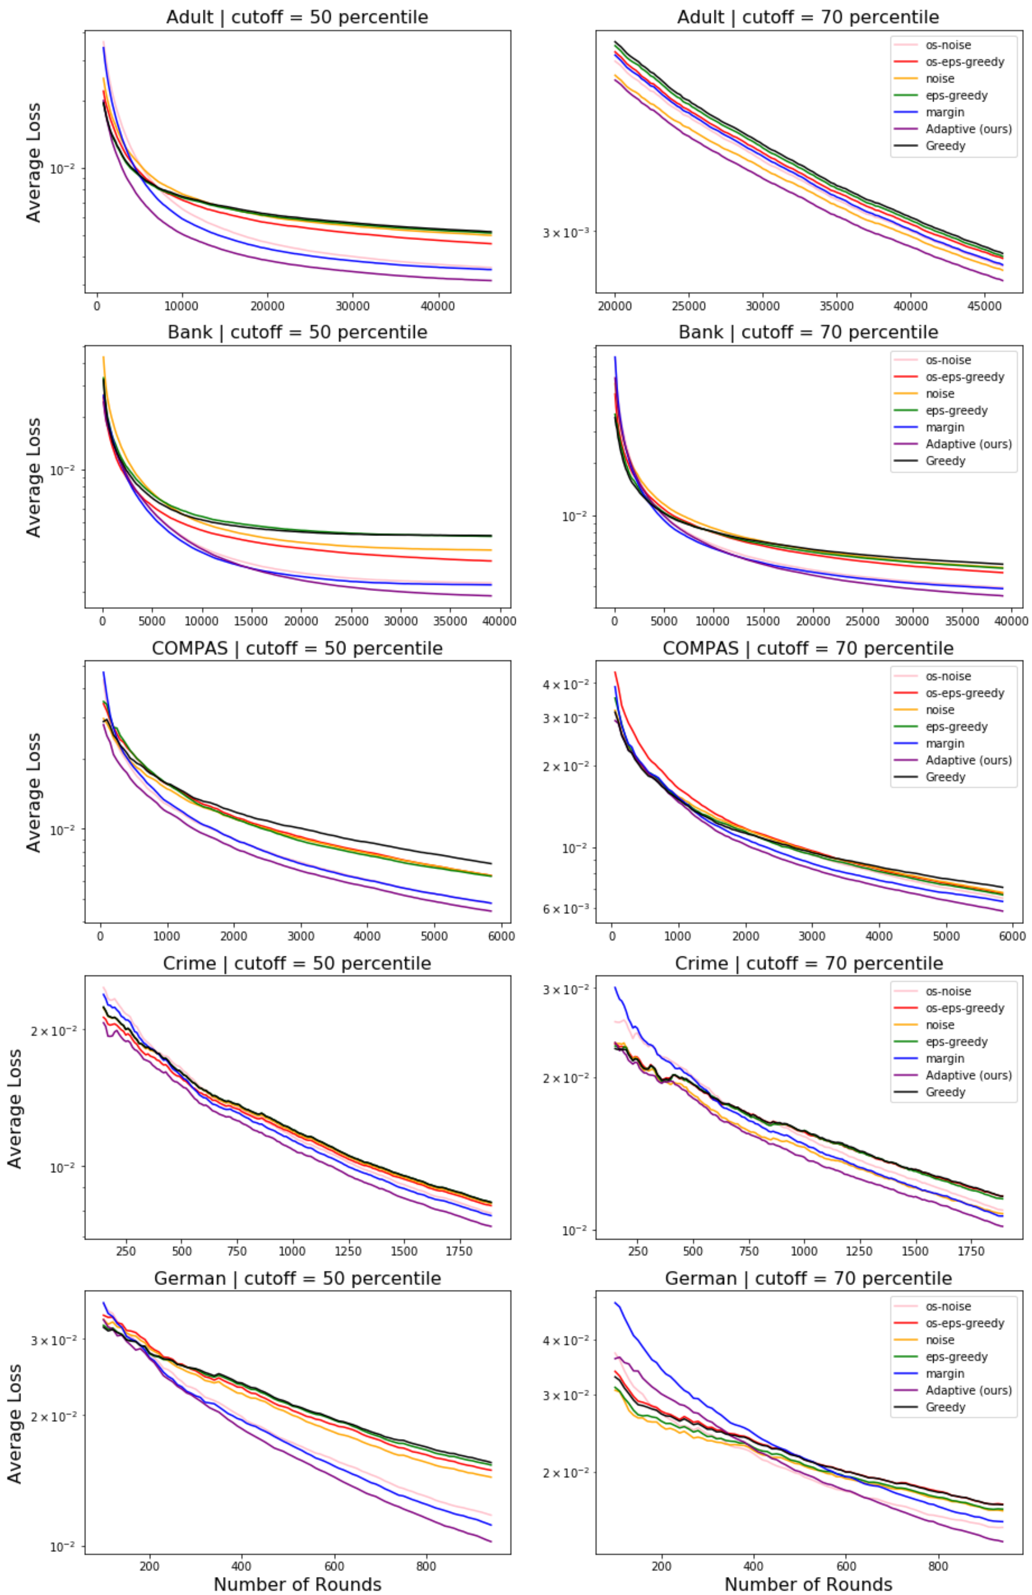

Figure 4: **Linear Regression plots**.

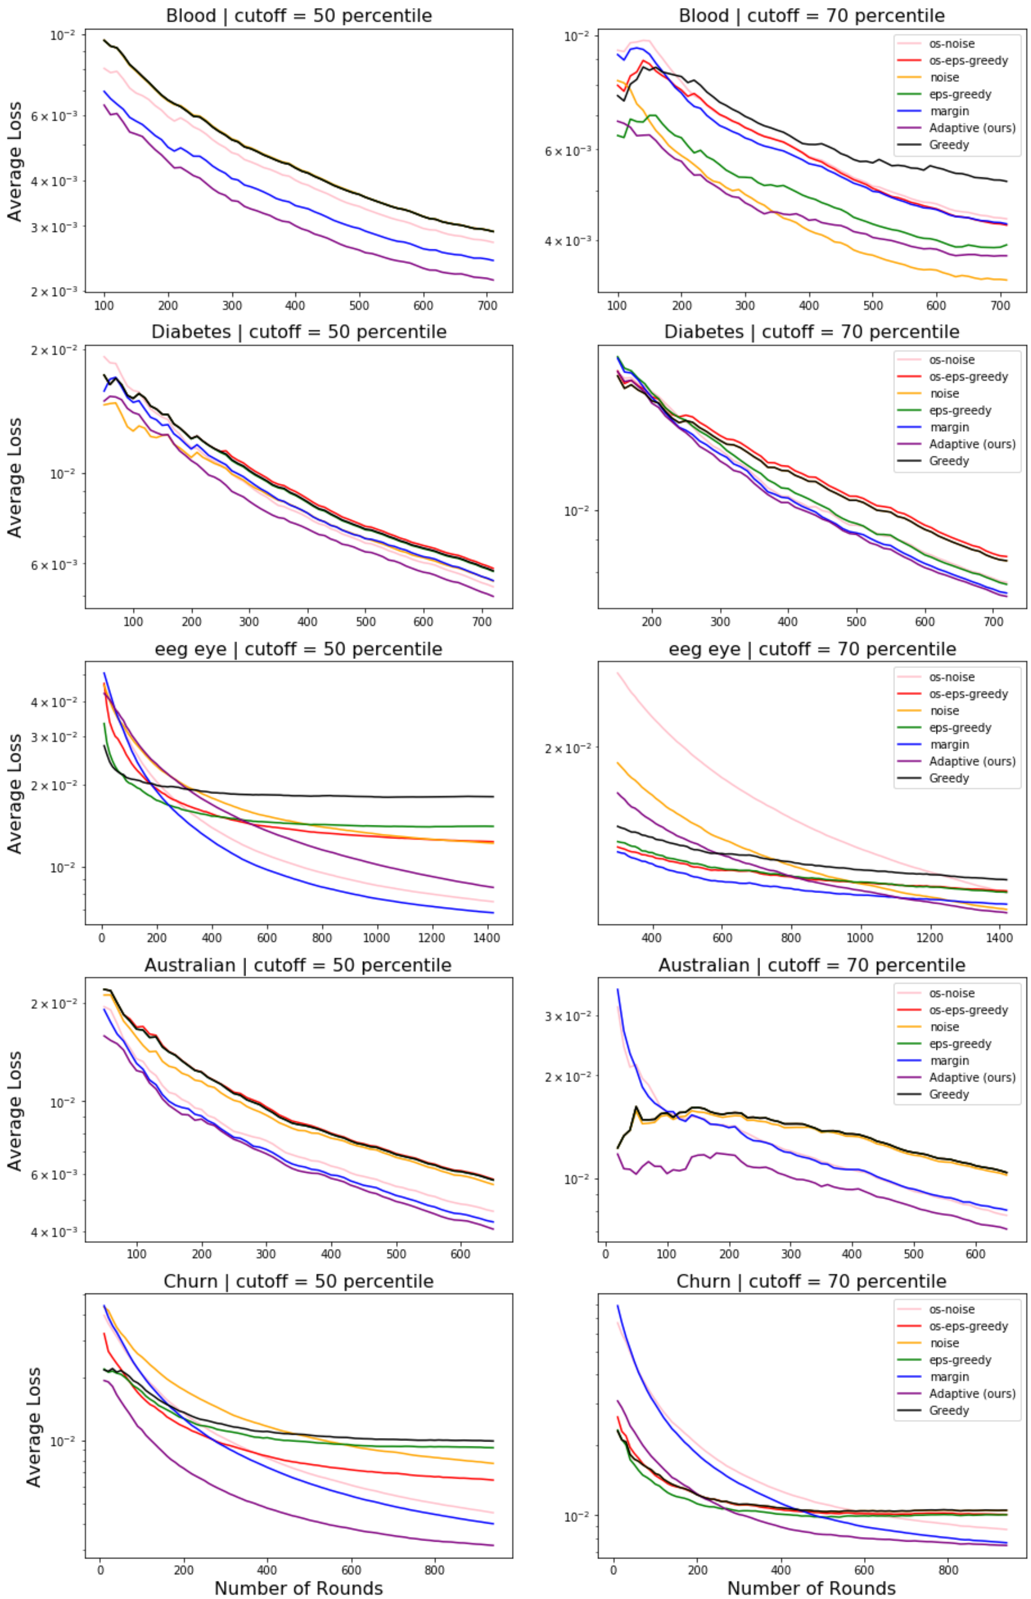

Figure 5: **Linear Regression plots (continued)**.

Dataset	c	greedy	ϵ-grdy	os-ϵ-grdy	noise	os-noise	margin	ours
Adult	50%	43.48	43.55	43.48	43.35	43.41	43.38	**42.63**
	55%	59.78	59.9	59.79	59.61	**59.58**	59.64	59.61
	60%	77.25	77.44	77.26	76.98	**76.48**	76.58	77.05
	65%	95.62	95.86	95.76	95.43	94.95	94.75	**93.13**
	70%	102.86	102.86	102.9	102.6	102.81	102.47	**100.06**
	75%	123.34	123.55	123.11	120.7	120.08	119.94	**118.38**
	80%	117.39	117.44	117.56	115.95	116.39	115.5	**112.05**
	85%	94.83	94.94	94.83	94.71	93.4	92.53	**90.86**
	90%	77.83	78.05	78.36	77.12	76.61	77.24	**75.22**
	95%	32.12	32.48	32.8	31.86	31.14	31.47	**30.81**
Bank	50%	23.22	23.26	**23.18**	23.3	23.33	23.2	23.23
	55%	34.55	34.55	34.53	**34.48**	34.6	34.52	34.58
	60%	49.57	49.56	49.58	49.52	49.86	**49.5**	49.56
	65%	71.09	**70.76**	71.05	71.16	71.33	71.01	71.01
	70%	85.72	85.94	85.67	85.51	**85.26**	85.27	85.75
	75%	115.67	115.52	115.71	**114.63**	115.42	115.19	115.79
	80%	149.07	148.31	149.12	147.86	147.34	148.19	**147.1**
	85%	162.05	161.33	162.16	159.92	159.59	160.8	**158.34**
	90%	135.85	135.48	135.76	134.89	132.68	131.4	**121.53**
	95%	98.26	98.07	98.15	96.57	97.01	96.52	**89.9**
COMPAS	50%	44.47	43.88	44.15	43.07	42.11	42.64	**40.34**
	55%	47.87	47.45	47.85	47.19	46.98	46.61	**45.51**
	60%	44.38	43.91	44.36	43.95	41.94	41.62	**40.62**
	65%	42.54	41.98	41.59	40.69	39.94	40.58	**39.63**
	70%	43.7	43.59	**43.41**	43.66	43.83	43.7	43.7
	75%	35.85	35.8	35.17	**33.77**	35.51	35.51	35.07
	80%	38.79	38.74	38.85	**36.71**	38.29	37.16	36.79
	85%	27.51	27.51	27.47	27.28	**27.14**	27.19	27.28
	90%	21.98	21.98	21.98	21.55	**20.16**	20.69	21.04
	95%	16.72	**15.93**	16.46	16.63	16.63	16.71	16.24
Crime	50%	11.04	10.83	11.04	10.85	10.33	10.44	**9.42**
	55%	12.85	12.63	12.85	12.14	12.36	12.14	**11.45**
	60%	16.75	16.75	16.79	16.81	16.21	15.77	**15.04**
	65%	26.51	26.51	26.01	25.21	23.3	23.27	**22.0**
	70%	26.05	25.93	26.13	25.94	25.84	25.55	**24.46**
	75%	29.17	28.65	29.22	27.99	27.92	27.06	**26.34**
	80%	31.01	31.01	31.01	29.68	30.6	30.4	**30.08**
	85%	22.33	22.33	22.33	22.11	21.82	21.33	**21.3**
	90%	13.3	13.3	13.3	**12.62**	12.74	12.88	12.67
	95%	2.96	2.96	**2.95**	2.96	3.03	2.96	2.96
German	50%	35.71	35.21	33.55	33.35	24.19	23.19	**20.33**
	55%	37.24	34.7	35.09	37.01	29.42	26.65	**23.74**
	60%	42.12	39.95	39.12	37.53	31.77	29.35	**25.19**
	65%	35.27	35.14	34.63	33.39	31.23	30.55	**28.16**
	70%	42.55	41.14	42.18	40.98	40.64	40.3	**37.12**
	75%	31.49	31.41	31.49	31.26	31.02	30.77	**29.9**
	80%	31.83	31.67	31.83	30.76	30.0	29.94	**29.11**
	85%	29.61	29.61	29.61	**29.08**	29.26	29.31	29.45
	90%	24.73	24.73	**24.54**	24.41	24.61	24.62	24.7
	95%	18.19	18.19	18.19	17.78	17.88	17.88	**17.7**

Table 5: **Logistic Regression Results**.

Dataset	c	greedy	ϵ-grdy	os-ϵ-grdy	noise	os-noise	margin	ours
Blood	50%	5.05	5.05	4.87	4.83	4.71	4.53	**4.24**
	55%	6.44	6.44	6.16	6.46	6.17	6.15	**5.91**
	60%	6.97	6.97	6.97	6.87	6.19	5.5	**5.45**
	65%	11.34	11.34	11.34	11.59	9.02	8.72	**8.18**
	70%	13.04	13.04	13.03	13.04	10.84	12.14	**9.69**
	75%	6.13	6.13	6.13	5.76	5.44	4.54	**4.31**
	80%	8.92	8.88	8.92	8.77	9.05	8.79	**8.03**
	85%	2.63	2.63	2.63	**2.54**	2.62	2.63	2.59
	90%	6.98	6.98	6.98	6.88	6.98	6.98	**5.89**
	95%	5.63	5.63	5.63	5.25	5.43	5.55	**5.22**
Diabetes	50%	28.23	28.23	27.75	27.22	26.67	26.18	**25.16**
	55%	25.18	25.18	25.17	24.89	24.35	24.58	**24.28**
	60%	26.51	25.85	26.17	25.12	25.25	25.1	**24.78**
	65%	29.47	29.32	29.14	29.05	28.91	28.86	**28.66**
	70%	29.36	28.0	27.79	28.0	**27.4**	27.9	28.11
	75%	26.99	26.96	26.52	**25.52**	26.45	26.47	26.42
	80%	25.9	24.65	25.58	**24.9**	25.71	25.64	25.86
	85%	21.11	**20.94**	21.01	21.36	21.12	21.11	21.07
	90%	24.23	23.82	24.23	**23.72**	24.12	24.01	24.12
	95%	12.34	12.25	12.33	12.31	12.34	12.31	**12.16**
EEG Eye	50%	239.33	238.92	239.09	236.65	200.61	201.51	**187.28**
	55%	239.03	238.71	238.66	239.85	217.58	217.15	**206.65**
	60%	227.14	227.13	226.59	223.53	218.24	219.47	**211.05**
	65%	222.98	218.47	220.71	218.1	211.26	210.72	**199.73**
	70%	209.48	207.89	208.83	206.63	204.94	205.4	**199.04**
	75%	194.56	193.68	194.44	193.3	194.25	193.04	**189.83**
	80%	208.11	207.76	207.95	208.23	208.84	207.93	**202.14**
	85%	186.23	186.23	186.25	184.02	185.49	186.12	**178.63**
	90%	182.61	181.96	179.12	177.37	180.55	180.83	**176.03**
	95%	160.11	160.15	159.98	159.79	159.99	159.62	**156.21**
Australian	50%	21.88	21.88	21.87	21.21	21.76	20.81	**20.38**
	55%	22.7	22.61	22.7	22.3	21.62	21.91	**20.2**
	60%	21.23	21.15	21.0	21.09	20.58	21.23	**20.05**
	65%	16.82	16.72	16.65	15.98	15.94	15.76	**15.07**
	70%	17.47	17.29	17.46	**16.49**	17.24	17.46	17.43
	75%	11.02	11.02	11.02	**10.63**	11.27	11.02	11.02
	80%	8.28	8.28	8.09	**8.02**	8.06	8.14	8.17
	85%	8.01	7.95	8.01	**7.62**	8.08	8.01	8.01
	90%	5.79	5.79	5.79	**5.55**	5.92	5.79	5.78
	95%	2.88	2.88	2.88	**2.86**	2.92	2.88	2.88
Churn	50%	61.04	57.74	54.13	53.85	39.46	38.88	**34.89**
	55%	60.84	56.7	52.18	56.13	47.21	45.4	**42.94**
	60%	66.42	59.53	59.76	57.13	48.36	47.35	**41.68**
	65%	70.57	65.32	66.35	62.78	62.02	58.32	**53.09**
	70%	122.96	117.49	116.04	112.36	94.61	88.3	**82.23**
	75%	81.49	80.11	81.49	79.56	77.21	74.51	**72.74**
	80%	86.62	86.12	82.48	84.84	84.25	82.92	**81.61**
	85%	99.19	**93.62**	97.05	96.59	94.6	96.04	95.61
	90%	93.81	93.76	92.39	**90.29**	90.69	92.64	93.03
	95%	76.27	76.01	72.63	72.69	73.09	72.27	**70.87**

Table 6: **Logistic Regression Results (continued)**.

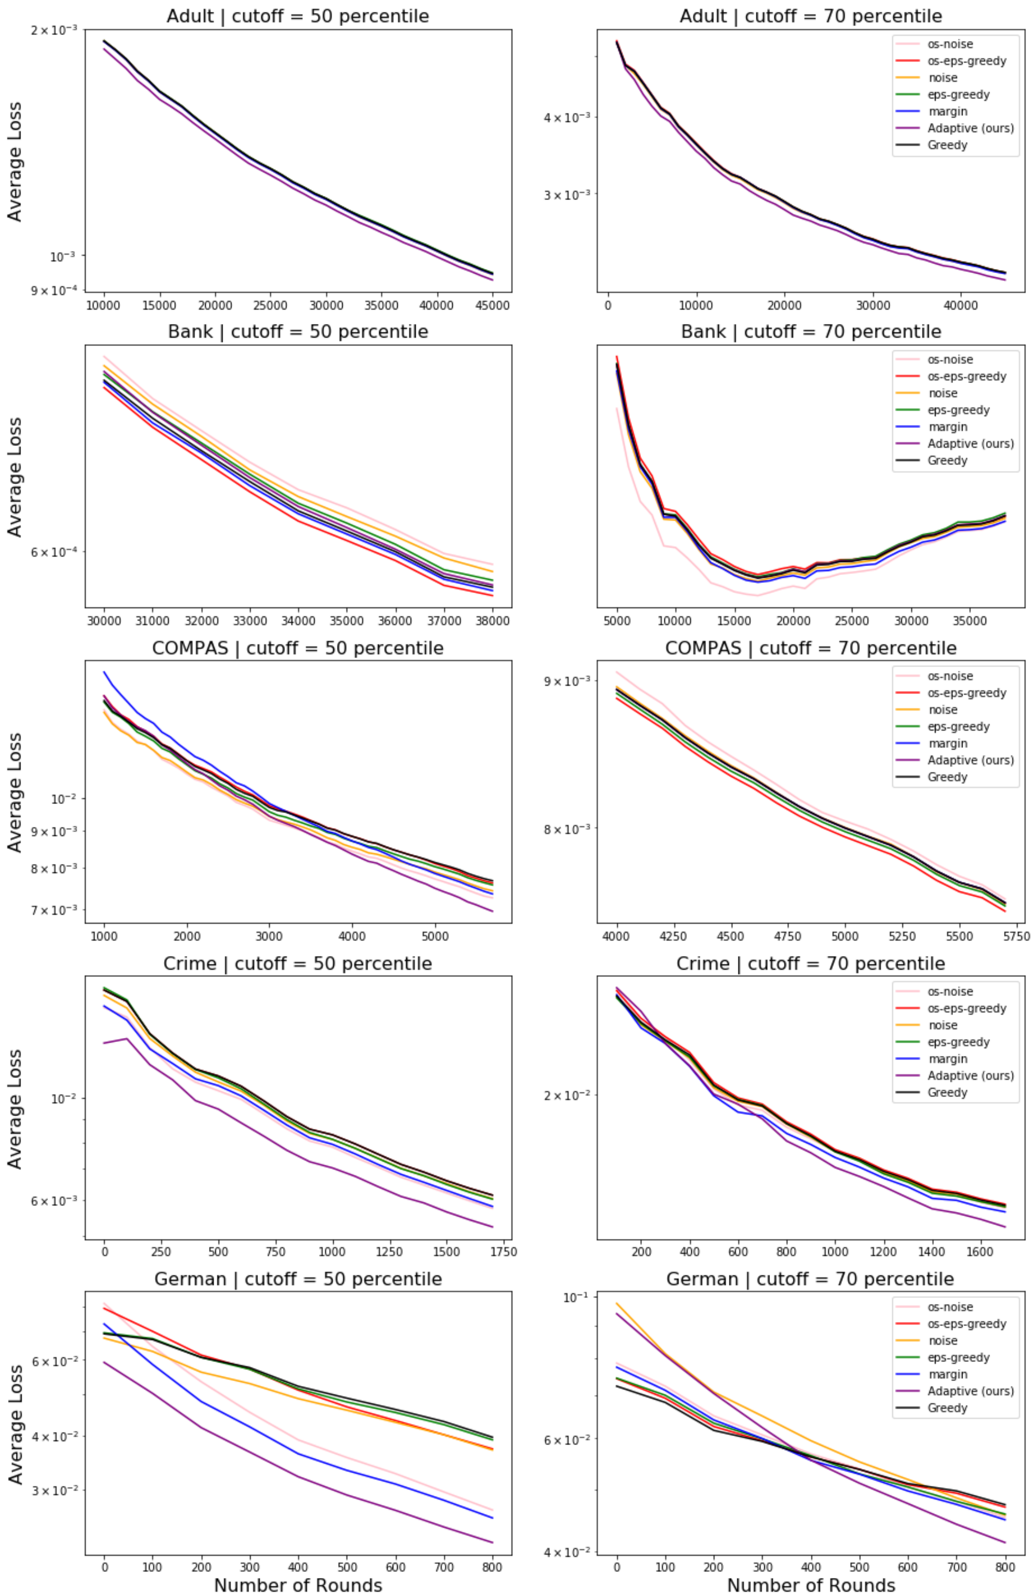

Figure 6: **Logistic Regression plots.**

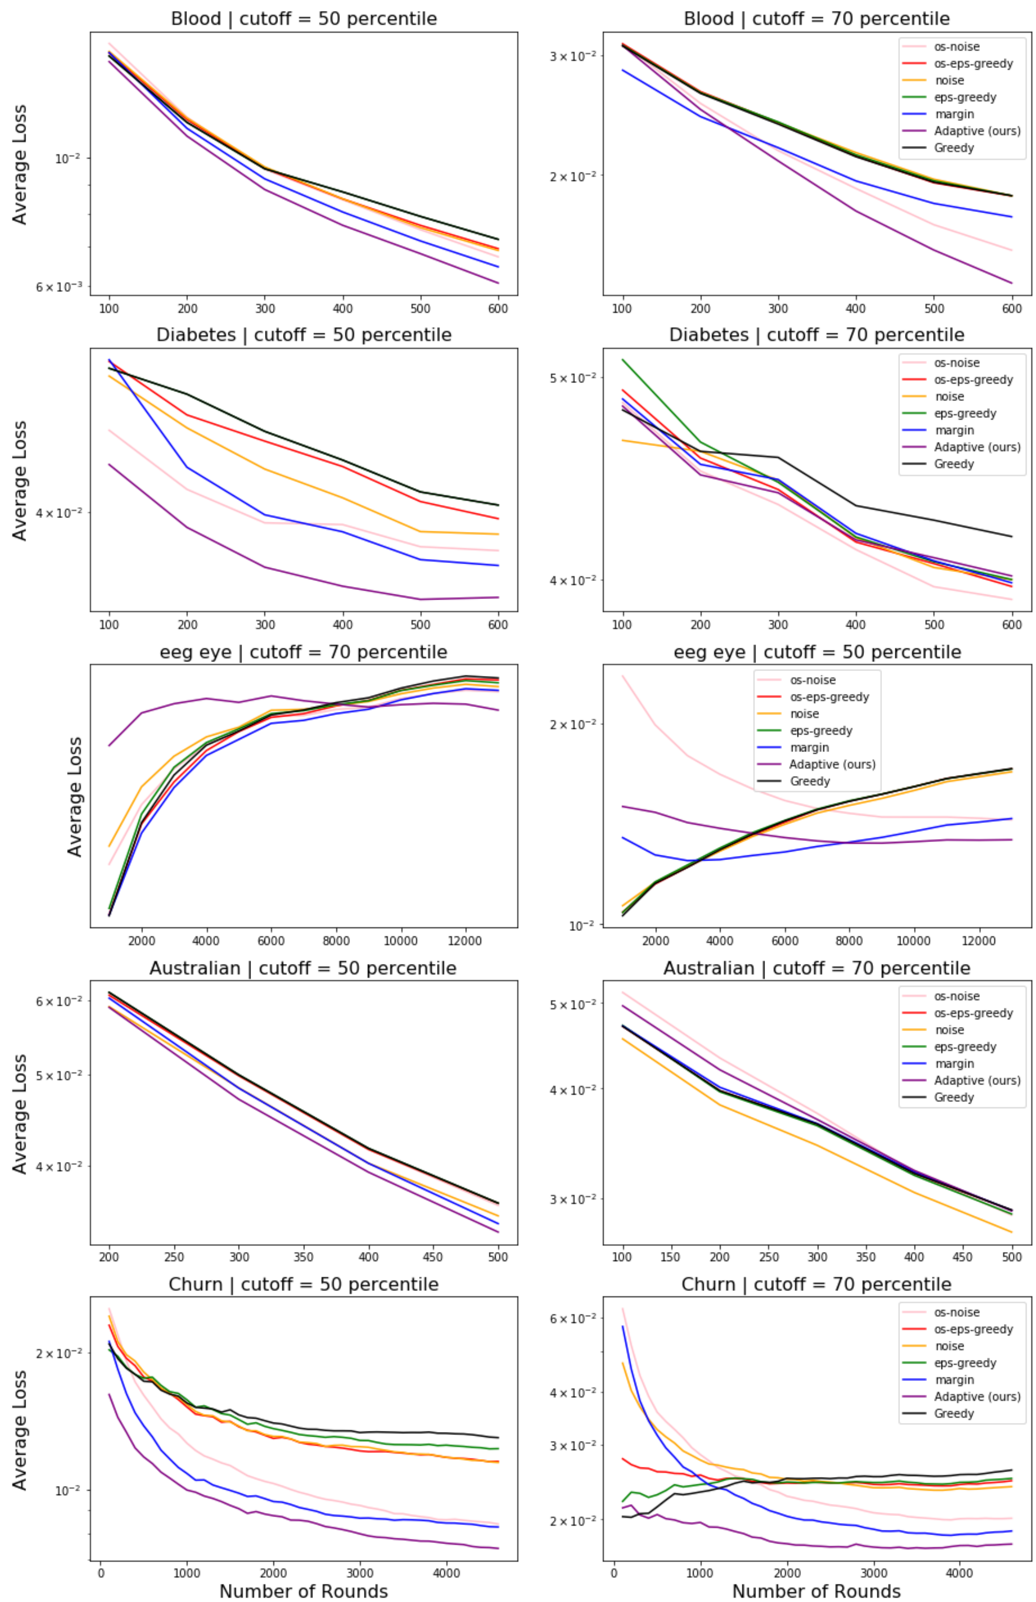

Figure 7: **Logistic Regression plots (continued).**