

Post-Training with Policy Gradients: Optimality and the Base Model Barrier

Alireza Mousavi-Hosseini¹ Murat A. Erdogdu²

August 13, 2026

Abstract

We study post-training linear autoregressive models with outcome and process rewards. Given a context $\mathbf{x}$, the model must predict the response $\mathbf{y} \in \mathcal{Y}^N$, a sequence of length N that satisfies a γ margin condition, an extension of the standard separability to sequences. We prove that on test samples where the base model achieves a non-trivial likelihood α , a variant of policy gradient (PG) can achieve likelihood $1 - \varepsilon$ with an essentially minimax optimal number of reward queries $\tilde{O}((\alpha^{-1} + \varepsilon^{-1})/\gamma^2)$. However, a barrier arises for going beyond the support of the base model. We prove that the overall expected error after post-training with outcome rewards is governed by a property of the base model called the *Likelihood Quantile* (LQ), and that variants of PG, while minimax optimal, may require a number of reward queries exponential in N to go beyond this support, regardless of the pre-training algorithm. To overcome this barrier, we study post-training with a process reward model, and demonstrate how PG variants in this setting avoid the curse of dimensionality in N via dependence on a token-level LQ. Along the way, we prove that under the margin condition, SGD with adaptive learning rate (LR) achieves a near optimal test error for statistical learning, and PG with adaptive LR achieves a near optimal number of mistakes for online learning while being computationally efficient whenever possible, both of which may be of independent interest.

1 Introduction

Outcome-based reinforcement learning (RL) has been a promising approach to enable large language models (LLMs) to go beyond their pre-training data. RL has been particularly successful in domains such as math and coding where one is able to *verify* the correctness of the model’s response, without requiring a complete ground-truth solution [GYZ⁺25, TDG⁺25]. However, the extent to which RL can create new knowledge that is missing from the base model is unclear. While some works argue that RL with outcome rewards can genuinely enhance the base model [WLZ⁺26], others have observed that the effect of RL post-training has mostly been to sharpen the base distribution, and the fine-tuned model is not able to generate responses that are outside the support of the base model [YCL⁺25, WXL⁺25, SLX⁺25]. Indeed, explicitly sampling from the sharpened base distribution can match or outperform RL [KD26], which suggests that RL cannot successfully improve upon the base model unless it is already highly capable.

On the theoretical side, prior works have attempted to characterize the benefits of RL post-training compared to supervised fine-tuning (SFT) [SRLK25] or on top of pre-training [TMUK26, KWLS25]. While providing positive results for RL, these works leave out the role of the base model. [FMR25] demonstrates that a notion of *coverage* is needed for computationally efficient exploration with RL. However, their definition does not depend on the final desired likelihood, the probability that the model generates a correct response to a context. [CHG⁺25] studies coverage profile which characterizes the probability of success with best-of- N sampling, and consider how log-likelihood maximization improves this coverage. However, the consequences of this definition for RL post-training, which is used to achieve the best-of- N performance using a single generation, is less clear. [CJRX25] studies outcome-based online RL, with a focus on sample complexity,

¹University of Toronto, Vector Institute. mousavi@cs.toronto.edu.

²University of Toronto, Vector Institute, and Numurho. erdogdu@cs.toronto.edu




Figure 1: The evolution of the model’s likelihood to generate the correct response over different contexts throughout PG. **(Left)** The average likelihood over samples with initial likelihood ≈ 0 under the base model. On-policy PG with PRM is able to improve this average while with ORM stays at 0. **(Center)** The comparison between the expected test error with ORM where the error plateaus at a threshold, and with PRM where the error continues to decrease. **(Right)** The likelihood of individual samples throughout PG with ORM, where the color denotes initial likelihood. Experiment details are presented in Section 6.

while the computational efficiency of their algorithm is unclear. Importantly, the works above leave open the study of practical algorithms for post-training, specifically PG variants such as proximal policy optimization (PPO) [SWD⁺17] and group relative policy optimization (GRPO) [SWZ⁺24]. In this work, we aim to answer the following question in a theoretically tractable setting:

Q1. *How do the number of reward queries and policy gradient steps in post-training depend on the base model for on- and off-support samples?*

We answer the above question by contrasting the on- and off-support cases. For sequences of length N , the probability that the uniform policy generates a correct response is exponentially small in N . If the base policy from pre-training achieves a probability of correct response (likelihood) that is non-trivial, i.e. at most polynomially small in N , we consider this sample to be on the base model support, and the complexities above remain polynomial. However, if the base policy is not significantly better than a uniform distribution on a sample, then the sample is off-support and improving its likelihood may require exponentially many iterations. This phenomenon is demonstrated in Figure 1. While this gives us a per-sample answer to the above question, a natural next step is to investigate the expected error over the entire distribution:

Q2. *Can RL post-training achieve a significantly smaller expected test error compared to the base model while being computationally efficient?*

We answer the above question for outcome and process rewards. For outcome reward models, we show that in the worst-case, any post-training algorithm may require exponentially many reward queries to significantly improve upon the base model obtained by SGD. However, this exponential dependence on the sequence length can be alleviated by using a process reward model, where after generating each token, the learner can query the reward to verify the correctness of the response so far.

1.1 Our Contributions

Our main contributions are as follows:

- In Section 3.1, we study post-training with outcome rewards and access to a base model q , with the goal of predicting a response $\mathbf{y}^*(\mathbf{x}) = (y_1^*(\mathbf{x}), \dots, y_N^*(\mathbf{x}))$, where $y_i \in \{1, \dots, k\}$, to an input context $\mathbf{x}$, which satisfies a γ margin assumption. We show that conditioned on the base model achieving a non-trivial likelihood α on a sample $(\mathbf{x}, \mathbf{y}^*(\mathbf{x}))$, a variant of policy-gradient can boost this likelihood and get expected error ε with $\tilde{O}(1/(\alpha\gamma^2\varepsilon))$ iterations. As a consequence, we show that PG with a uniform behavior policy can achieve the minimax optimal mistake bound $\tilde{O}(k^N/\gamma^2)$ in online learning.
- In Section 3.2, we show that for the overall expected error to go below ε , the number of reward queries depends on a property of the base model, which we refer to as the *Likelihood Quantile* (LQ). This creates a

fundamental *barrier of the base model*; if the base model is pre-trained with SGD, PG requires exponentially many reward queries to go below the SGD error rate.

- In Section 4, we demonstrate how using process rewards instead of outcome rewards can alleviate the above problem, with the number of reward queries instead depending linearly on N and another property of the base model, called the *Token-Level Likelihood Quantile*, which is linear in k in the worst-case.
- In Section 5, we establish the minimax optimality of the policy gradient variants explored in Section 3. We also rule out the possibility of significant algorithmic improvements beyond SGD in pre-training for better LQ; the minimax optimal LQ is of order k^{-N} unless enough labeled samples are given for an SGD pre-trained model to have very low test error of order k^{-N} .

1.2 Additional Related Works

Online Linear Classification with Bandit Feedback. The problem of post-training a policy with outcome rewards can be cast as an instance of a contextual bandit problem. For the case of $N = 1$, there has been a rich literature on learning a linear multiclass classifier with bandit feedback. The seminal work of [KSST08] introduced the Banditron algorithm, a modification of the classical Perceptron algorithm, that only achieves a mistake bound $\tilde{O}(k^2/\gamma^2)$. [DH13] shows that a mistake bound of $\tilde{O}(k/\gamma^2)$ is information-theoretically optimal, and presents an inefficient algorithm to achieve it. [BPS⁺19] provides a computationally efficient algorithm that only achieves the $\tilde{O}(k/\gamma^2)$ rate under a stronger notion of margin. Others have considered more general data distributions and achieved a relative mistake bound of order $\sqrt{T}$, see e.g. [vdHFCB21] and references therein.

In contrast, we introduce a simple online algorithm as a variant of PG with a uniform behavior policy that achieves the minimax optimal mistake bound $\tilde{O}(k^N/\gamma^2)$, thus achieving the desirable $\tilde{O}(k/\gamma^2)$ when $N = 1$ while also being computationally efficient per iteration. Thus, we answer the open question of achieving such an efficient online learning algorithm raised in [BPS⁺19].

SGD for Separable Data. The study of SGD on logistic loss for linearly separable data first gained attention through its connection to the max-margin SVM and its implicit bias [SHN⁺18, JT19]. Extensions to multiclass settings have been considered in [RSSW24, SK25]. Our analysis for both SGD and PG is more similar to [Sha21]. In contrast, we significantly go beyond the binary setting by considering autoregressive models and adapting the margin assumption to sequences of labels, where each label has $k \geq 2$ choices.

Policy Gradient Analysis. A comprehensive study of theoretical properties of policy gradient methods in tabular settings was carried out by [AKLM21]. However, their guarantees for PG with softmax parameterization are only asymptotic. Further works have made this guarantee non-asymptotic and extended beyond the tabular setting [ZKZB20, MXSS20, ZNY⁺21, LZBY20]. We remark however that none of these results are directly applicable in our setting, since their generality prevents establishing tight bounds. We instead carry out an analysis tailored to the 0-1 reward structure for post-training an autoregressive model.

Notation. For quantities a and b , we use $a \lesssim b$ or $a = \mathcal{O}(b)$ to denote $a \leq Cb$ for some absolute constant $C > 0$, and define $a \gtrsim b$ and $a = \Omega(b)$ similarly. We use $\tilde{O}$ and $\tilde{\Omega}$ to hide polylogarithmic factors. We use $f(\varepsilon) = o(1)$ to denote $f(\varepsilon) = \mathcal{O}(1/\log(1/\varepsilon))$, a stricter version of the usual definition $\lim_{\varepsilon \downarrow 0} f(\varepsilon) = 0$. We use the notation $\mathbf{y}_{1:i} = (y_j)_{j=1}^i$, and $\mathbf{y}_{1:0}$ is an empty sequence. We use $\text{Clip}(x, a, b) := ((x \vee a) \wedge b)$ for $x, a, b \in \mathbb{R}$. For a set $\mathcal{Y}$, we define $\mathcal{Y}^+ := \bigcup_{i=1}^N \mathcal{Y}^i$. We use $\mathbf{e}_y$ to denote the y^{th} standard basis of $\mathbb{R}^k$.

2 Autoregressive Linear Models and Pre-Training

Given access to a verifier, post-training can be performed using outcome rewards. This approach yields a contextual bandit formulation: given a prompt $\mathbf{x} \in \mathcal{X}$, the model $p_{\mathbf{W}}(\cdot | \mathbf{x})$ generates a response $\mathbf{y} = (y_1, \dots, y_N) \in \mathcal{Y}^N$ of length N , and receives a reward $r(\mathbf{x}, \mathbf{y})$ from an outcome reward model (ORM). For

simplicity, we fix the length of the response, while our arguments can easily be adapted to dynamic lengths. We assume $|\mathcal{Y}| = k$. For many verifiers this reward is binary, i.e. $r(\mathbf{x}, \mathbf{y}) = 1$ if $\mathbf{y}$ is a correct answer to $\mathbf{x}$, and $r(\mathbf{x}, \mathbf{y}) = 0$ otherwise. There are also settings where a process reward model (PRM) can provide intermediate signal to the learner. We study this formulation in the subsequent section.

We specify the generative model $p_{\mathbf{w}}(\cdot | \mathbf{x})$ using the decomposition $p_{\mathbf{w}}(\mathbf{y} | \mathbf{x}) = \prod_{i=1}^N p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1})$, where

$$p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}) = \frac{\exp(\langle \mathbf{w}, \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y_i) \rangle)}{\sum_{y' \in \mathcal{Y}} \exp(\langle \mathbf{w}, \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y') \rangle)}.$$

Here, $\phi: \mathcal{X} \times \mathcal{Y}^+ \rightarrow \mathbb{R}^D$ is the feature map. Current models typically implement ϕ using a deep Transformer [VSP⁺17]. The above formulation simplifies the analysis by only focusing on the final linear layer while freezing the feature map ϕ .

We assume that each prompt $\mathbf{x}$ has a unique correct response $\mathbf{y}^*(\mathbf{x})$ that satisfies the following assumption.

Assumption 1. Suppose there exists $\mathbf{w}^* \in \mathbb{R}^D$ such that $\|\mathbf{w}^*\|_2 = 1$ and for all $i \in [N]$, $y \neq y_i^*(\mathbf{x})$,

$$\langle \mathbf{w}^*, \phi(\mathbf{x}, \mathbf{y}_{1:i}^*(\mathbf{x})) \rangle \geq \langle \mathbf{w}^*, \phi(\mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x}), y) \rangle + \gamma,$$

a.s. over $\mathbf{x}$ for some $\gamma > 0$. We further assume $\|\phi(\mathbf{x}, \mathbf{y}_{1:i}(\mathbf{x}))\|_2 \leq 1$ for all $i \in [N]$, $\mathbf{x}$, and $\mathbf{y}$.

When $N = 1$, letting $\phi(\mathbf{x}, y) = \text{Vec}(\mathbf{e}_y \mathbf{x}^\top)$ and $\mathbf{w} = \text{Vec}(\mathbf{W})$ with $\mathbf{W} \in \mathbb{R}^{k \times d}$, the above reduces to a standard γ -margin assumption for k -class linear classification, and the autoregressive model reduces to the usual logistic model.

Notice that the above only requires separability at the token level while assuming all previous tokens are correct, as opposed to sequence level separability which would be harder to satisfy. With Assumption 1, it might be tempting to only train on a single token to recover $\mathbf{w}^*$, thus remove dependence on N from the complexity of supervised training. However, such a direction may not be unique, and a direction that separates some tokens might not be useful for the rest.

Before moving to the analysis of post-training, we first recall results about pre-training. During pre-training, we have access to i.i.d. labeled examples $(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))_{t \geq 0}$, and we can run SGD iterates as

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}), \quad (\text{SGD})$$

where η_t is a potentially adaptive learning rate. [CHG⁺25] demonstrated that with a constant LR, SGD suffers from a convergence rate that depends linearly on the sequence length N . They mitigate this issue by analyzing an Adagrad-type [DHS11] adaptive LR on SGD with a sufficiently large mini-batch size. Under Assumption 1, we can study the simpler variant above that operates in a one-pass mode over the pre-training data and does not require batching, hence can also be applied in online settings.

Proposition 1. Suppose Assumption 1 holds, and let $(\mathbf{w}_t)_{t=0}^{T-1}$ denote the iterates of SGD with learning rate schedule (η_t) . Then,

1. (Constant LR) If $\eta_t = 1/(2N)$, then

$$1 - \mathbb{E} \left[p_{\bar{\mathbf{w}}_T^{\text{SGD}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \right] \lesssim \frac{N \log(Tk\gamma^2)^2}{\gamma^2 T}$$

where $\bar{\mathbf{w}}_T^{\text{SGD}} := \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{w}_t$ is the averaged iterate.

2. (Adaptive LR) If

$$\eta_t = (2 + 4\|\nabla \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}_t) | \mathbf{x}_t)\|_2)^{-1},$$

then

$$1 - \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \lesssim \frac{\log(TNk\gamma^2)^2}{\gamma^2 T}.$$

In other words, if we draw $\tau \sim \text{Unif}(\{0, \dots, T-1\})$, then

$$1 - \mathbb{E} [p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \lesssim \frac{\log(TNk\gamma^2)^2}{\gamma^2 T}.$$

A few remarks are in order. First, it is well-known that the VC dimension of the class of γ -margin half-spaces in dimension $D \geq 1/\gamma^2$ is at least $\Omega(1/\gamma^2)$ (e.g. can be proved by considering an orthonormal set of points). Therefore, by standard minimax lower bounds on sample complexity using VC dimension in realizable settings [SSBD14, Theorem 6.8], we require at least $\Omega(1/(\gamma^2\varepsilon))$ samples to achieve a test error ε in the worst-case for $N = 1$ and $k = 2$, which can thus be extended to all N and k . This lower bound can alternatively be derived using the simpler technique we employ in Theorem 10. Interestingly, while for $N = \mathcal{O}(1)$ SGD with constant or adaptive LR (almost) achieves the minimax optimal sample complexity, for $N \gg 1$, it is only the adaptive LR that does so, with nearly the same rate as the $N = 1, k = 2$ case.

Moreover, while we state Proposition 1 and other results in the paper in a conventional statistical learning setup where data is i.i.d., proofs are obtained by an online analysis and an online-to-batch conversion (see Section A for details). As a consequence, an online learner that samples responses $\mathbf{y}^{(t)} \sim p_{\mathbf{w}_t}(\cdot | \mathbf{x}^{(t)})$ would make at most $\tilde{\mathcal{O}}(1/\gamma^2)$ mistakes in expectation over online (potentially adversarially chosen) data, which nearly matches the classical Perceptron algorithm and is minimax optimal, while covering the broader setting of autoregressive linear models.

3 Policy Gradient with Outcome Reward (Bandit Feedback)

For post-training, we only have access to i.i.d. prompts $(\mathbf{x}_t)_{t \geq 0}$ as well as a reward model, but not the labels directly. In this section, we assume access to a reward model for the final response generated, i.e. $r : \mathcal{X} \times \mathcal{Y}^N \rightarrow \{0, 1\}$ where

$$r(\mathbf{x}, \mathbf{y}) = \mathbb{1}[\mathbf{y} = \mathbf{y}^*(\mathbf{x})].$$

This stage is typically performed using some policy gradient algorithm such as PPO or GRPO. Here, we study the most basic policy gradient algorithm, also referred to as REINFORCE [Wil92]. The PG update with outcome rewards is given by

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t r_t \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)}), \quad (\text{PG-OR})$$

where $r_t := r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$, $\mathbf{y}_t \sim q_t(\cdot | \mathbf{x}^{(t)})$ and q_t is the behavior policy at round t which can be adaptively chosen using past samples. Unless otherwise stated, we assume $\mathbf{w}_0 = \mathbf{0}$ for simplicity. While in practice the initialization is provided by the base model, in our analysis the benefit of the base model comes from the behavior policy, therefore we can simplify the analysis by initializing at zero. Two remarks are in order.

- For (PG-OR) to be policy gradient in the sense of following the population reward gradient, we have to introduce the importance weights $p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})/q_t(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})$ unless we are fully on-policy, i.e. $q_t = p_{\mathbf{w}_t}$. We defer the study of this fully on-policy PG to Section A.5 since our analysis for it yields a qualitatively similar behavior to subsequent PG variants we study, yet with a suboptimal convergence rate. When $q_t \neq p_{\mathbf{w}_t}$, practical algorithms such as PPO clip the importance weight to be close to 1. Our analysis can also accommodate clipped importance weights; in Section A.6 we study the following updates

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t r_t \rho_t \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})$$

where $\rho_t = \text{Clip}\left(\frac{p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})}{q_t(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})}, 1/\zeta, \zeta\right)$ for some $\zeta \geq 1$. In summary, we can show that all convergence rates will be scaled by a factor of ζ^2 . In the main text, we focus the presentation on the case where the importance weights are removed, i.e. $\zeta = 1$, as it optimizes the convergence rate and achieves minimax optimality, which suggests that following the (unbiased estimate of) population reward gradient in this setting is not strictly necessary.

- Note that (PG-OR) considers a simple advantage estimator with no baseline. Once again, this is motivated by the minimax optimality of the updates. No advantage estimator achieves a better rate with only Assumption 1.

3.1 Conditional Convergence of PG

We begin with the following conditional guarantee. The proof of all results in this section are stated in Section A.

Theorem 2 (PG-OR with Constant LR). *Let $(\mathbf{w}_t)$ denote the (PG-OR) iterates. Suppose Assumption 1 holds. For simplicity, assume $\eta = 1/(2N)$ and $\|\mathbf{w}_0\|_2 \leq 1/\gamma \cdot \log(NkT\gamma^2)$. Let $\mathbf{x}$ denote a test sample. For any $\alpha \in (0, 1]$, define the event $\mathcal{E}_\alpha := \{\min_t q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha\}$ and $\pi_\alpha := \mathbb{P}(\mathcal{E}_\alpha)$. Then, $\bar{\mathbf{w}}_T^{\text{PG}} = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{w}_t$ satisfies*

$$\mathbb{E}[p_{\bar{\mathbf{w}}_T^{\text{PG}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha] \geq 1 - \tilde{\mathcal{O}}\left(\frac{N}{\pi_\alpha \alpha \gamma^2 T}\right). \quad (3.1)$$

The conditioning above highlights that the test performance on an individual sample after post-training depends on how likely the behavior policy q_t is to generate the correct response for that sample. In the bound above, α denotes this likelihood. If we set $q_t = \text{Unif}$ for all t , then $\pi_\alpha = 1$ for $\alpha = k^{-N}$ and $\pi_\alpha = 0$ for any $\alpha > k^{-N}$. The way a base model enters the above bound is through the function $\alpha \mapsto \pi_\alpha$. The slower it decays, the faster the convergence will be for samples that have an initial likelihood at least α .

Similar to SGD, we can remove N in the convergence rate above by using adaptive learning rates.

Theorem 3 (PG-OR with Adaptive LR). *Consider the same setting as Theorem 2, except we use the adaptive learning rate $\eta_t = (4 + 2\|\nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})\|_2)^{-1}$. Then, in expectation over the randomness in training, in the test prompt $\mathbf{x}$, and in $\tau \sim \text{Unif}(\{0, \dots, T-1\})$, we have*

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha] \geq 1 - \tilde{\mathcal{O}}\left(\frac{1}{\pi_\alpha \alpha \gamma^2 T}\right), \quad (3.2)$$

where we recall $\mathcal{E}_\alpha := \{\min_t q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha\}$ and $\pi_\alpha := \mathbb{P}(\mathcal{E}_\alpha)$ for any $\alpha \in (0, 1]$.

Remark. [Online Learning Regret Bound] As discussed earlier, while we stated the guarantee of Theorem 3 for i.i.d. data, we can utilize the same proof technique for an arbitrary sequence $\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)})$, where an adversary can choose $\mathbf{x}^{(t)}$ and $\mathbf{y}^*(\mathbf{x}^{(t)})$ for all t , with the only constraint being Assumption 1 (see Proposition 19 and its proof). In particular, if q_t is simply the uniform policy, we have $\pi_\alpha = 1$ for $\alpha = k^{-N}$, thus we can remove conditioning and write

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})] \geq 1 - \tilde{\mathcal{O}}\left(\frac{k^N}{\gamma^2 T}\right).$$

Therefore, the online learning algorithm that draws the response $\mathbf{y}^{(t)}$ to $\mathbf{x}^{(t)}$ according to $p_{\mathbf{w}_t}$ will have an expected number of mistakes $\tilde{\mathcal{O}}(k^N/\gamma^2)$ over T rounds¹. Indeed, this number of mistakes is minimax optimal (up to log terms) as established in Theorem 11. To our knowledge, this is the first online learning algorithm in this setting that is computationally efficient per iteration, while nearly achieving the minimax optimal number of mistakes. Prior works achieved this number of expected mistakes under a stronger notion of separability and only for $N = 1$ [BPS⁺19].

3.2 Unconditional Convergence of PG

In the previous section, we conditioned the expected error on the initial likelihoods of test samples. Our focus now is to use those results to bound the expected test error over the entire distribution. For a policy q , we define the Likelihood Quantile (LQ) function $\mathcal{Q}_q : [0, 1] \rightarrow [0, 1]$ as

$$\mathcal{Q}_q(\varepsilon) := \sup \{\alpha \in [0, 1] : \mathbb{P}_{\mathbf{x}}(q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq \alpha) \leq \varepsilon\}.$$

Specifically, $\mathcal{Q}_q(\varepsilon)$ provides the ε -quantile of the random variable $q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})$, which is the likelihood of the ground-truth data $(\mathbf{x}, \mathbf{y}^*(\mathbf{x}))$ under the base model q . As can be seen from Figure 2b, the more concentrated around 1 the random variable $q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})$ is (corresponding to longer pre-training), the larger $\mathcal{Q}_q(\varepsilon)$ will be for a fixed ε . With $\pi_\alpha := \mathbb{P}_{\mathbf{x}}(q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha)$, the mapping $\alpha \mapsto 1 - \pi_\alpha$ is the CDF of the

¹The prediction query can also be used for exploration: first sample $\mathbf{y}^{(t)} \sim p_{\mathbf{w}_t}(\cdot | \mathbf{x}^{(t)})$ and use it for the update if it receives positive reward; otherwise, make one additional query from the uniform policy. Even when mistakes are counted on both queries, their expected number is at most twice the prediction error, so the same $\tilde{\mathcal{O}}(k^N/\gamma^2)$ mistake bound holds.

likelihood over test samples, and $\mathcal{Q}_q$ is the generalized inverse of this function; it satisfies $\mathcal{Q}_q(1 - \pi_\alpha) \geq \alpha$. Through this observation, LQ is connected to other notions studied in the literature; $1 - \pi_\alpha$ is referred to as the α -coverage profile in [CHG⁺25]. We refer to references therein for other definitions in the literature.

We can examine the lower bounds on $\mathcal{Q}_q$ when $q = p_{\tilde{w}_n^{\text{SGD}}}$ is obtained by SGD pre-training with $T^{\text{SGD}} = n$ samples:

- With the constant LR $\eta \asymp 1$: $\mathcal{Q}_q(\varepsilon) \geq e^{-\tilde{O}(N/(n\gamma^2\varepsilon))}$.
- With the adaptive LR $\eta_t \asymp (1 + \|\nabla \log p_{w_t}\|_2)^{-1}$: $\mathcal{Q}_q(\varepsilon) \geq (1 - \tilde{O}((n\gamma^2\varepsilon)^{-1})) \vee 0$.

Further, for the uninformative uniform policy q , we have $\mathcal{Q}_q(\varepsilon) = k^{-N}$ for all $\varepsilon < 1$. We can now state unconditional bounds on the accuracy of the post-trained model.

Corollary 4 (Test Error of PG-OR). *Suppose Assumption 1 holds. Given a base policy q_0 , let $q = \frac{1}{2}q_0 + \frac{1}{2}\text{Unif}$. Consider (PG-OR) with $q_t = q$ for all t and with the adaptive learning rate of Theorem 3. Let $\tau \sim \text{Unif}(\{0, \dots, T-1\})$. Then, for all $\varepsilon \in (0, 1)$, with*

$$T = \tilde{O}\left(\frac{\min(\mathcal{Q}_{q_0}((1 - o(1))\varepsilon)^{-1}, k^N)}{\gamma^2\varepsilon}\right)$$

iterations, we have $\mathbb{E}[p_{w_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$.

Remark. We distinguish between the reciprocal and inverse in our notation; $\mathcal{Q}_{q_0}(\varepsilon)^{-1} = 1/\mathcal{Q}_{q_0}(\varepsilon)$ while $\mathcal{Q}_{q_0}^{-1}(\cdot)$ denotes the inverse function of $\mathcal{Q}_{q_0}(\cdot)$. By monotonicity of $\mathcal{Q}_{q_0}$, we have $\varepsilon \geq 1 - \pi_{k^{-N}} = \mathcal{Q}_{q_0}^{-1}(k^{-N})$ iff $\mathcal{Q}_{q_0}(\varepsilon) \geq k^{-N}$. Therefore, we can express the above bound as

$$T = \begin{cases} \tilde{O}\left(\frac{\mathcal{Q}_{q_0}((1 - o(1))\varepsilon)^{-1}}{\gamma^2\varepsilon}\right) & \text{if } \varepsilon \geq \mathcal{Q}_{q_0}^{-1}(k^{-N}) \\ \tilde{O}\left(\frac{k^N}{\gamma^2\varepsilon}\right) & \text{otherwise.} \end{cases} \quad (3.3)$$

To understand the base model barrier, consider the case where q_0 is given by SGD with adaptive LR on n labeled samples. Then, $\mathcal{Q}_{q_0}^{-1}(k^{-N}) = \tilde{\Omega}(1/(n\gamma^2))$. This suggests that to go below expected error $\tilde{O}(1/(\gamma^2n))$, which is already achieved by the SGD-trained base model, policy gradient will require exponentially many iterations. One may wonder if this is due to a suboptimal bound on $\mathcal{Q}_{q_0}$ for SGD, or due to a fundamental suboptimality of SGD itself.

In Section 5, we rule out both cases by showing a minimax bound of $\mathcal{Q}_q(\varepsilon) \leq k^{-N}$ when $n \lesssim 1/(\gamma^2\varepsilon)$, among all algorithms that output q with n labeled i.i.d. samples. Note that when $n \gg 1/(\gamma^2\varepsilon)$, the SGD pre-trained model already achieves error ε . This phenomenon can also be observed in practice. Figure 2a shows that while for sufficiently large ε the post-trained Qwen3-8B model has larger LQ than Qwen3-8B-Base, for ε below a threshold, both models fail completely and have LQ of zero. This suggests that post-training was unable to extend the model’s support beyond regions where the base model already exhibited non-trivial LQ.

The barrier above is not specific to policy gradient and its variants. We demonstrate in Theorem 11 that (PG-OR) with the behavior policy described above is minimax optimal (up to log factors) among algorithms that have access to q_0 and only make one reward query per context. However, it is possible to improve both sample/iteration complexity and number of reward queries by considering a modification that makes multiple queries per round. Specifically, by choosing q_t according to Algorithm 1, we have the following.

Corollary 5. *Suppose Assumption 1 holds. Given a base policy q_0 and for any $\varepsilon \in (0, 1)$, consider (PG-OR) with adaptive learning rate of Theorem 3 and q_t given by Algorithm 1 with $m = \min(\lceil \mathcal{Q}_{q_0}((1 - o(1))\varepsilon)^{-1} \rceil, k^N)$.*

Algorithm 1 Best-of- m exploration for PG-OR.

Input: $\mathbf{x}^{(t)}$, p_{w_t} , q_0 , m
 $\mathbf{y} \sim p_{w_t}(\cdot | \mathbf{x}^{(t)})$
 $j = 1$
while $j \leq m$ **and** $r(\mathbf{x}^{(t)}, \mathbf{y}) = 0$ **do**
 $j = j + 1$, $\mathbf{y} \sim (\frac{1}{2}q_0 + \frac{1}{2}\text{Unif})(\cdot | \mathbf{x}^{(t)})$
end while
return $\mathbf{y}$




Figure 2: (a) Estimated LQ of Qwen3-8B and Qwen3-8B-Base on Math500 and (b) true LQ of our linear autoregressive model trained with Adagrad with varying number of steps on a synthetic task. As the model trains for longer, likelihoods increase and $\mathcal{Q}_q(\epsilon)$ gets closer to 1 for fixed ϵ . Details provided in Section 6.

Algorithm 2 Best-of- m exploration for PG-PR.

Input: $(\mathbf{x}^{(t)}, p_{\mathbf{w}_t}, q_0, m)$
 $\mathbf{y} \sim p_{\mathbf{w}_t}(\cdot | \mathbf{x}^{(t)})$
if $r^*(\mathbf{x}^{(t)}, \mathbf{y}) = 1$ **then**
 return $\mathbf{y}$
end if
for $i = 1, \dots, N$ **do**
 $j = 1, y_{i,j} \sim (\frac{1}{2}q_0 + \frac{1}{2}\text{Unif})(\cdot | \mathbf{x}^{(t)}, \mathbf{y}_{1:i-1})$
 while $j \leq m$ **and** $r^*(\mathbf{x}^{(t)}, \mathbf{y}_{1:i-1}, y_{i,j}) = 0$ **do**
 $j = j + 1,$
 $y_{i,j} \sim (\frac{1}{2}q_0 + \frac{1}{2}\text{Unif})(\cdot | \mathbf{x}^{(t)}, \mathbf{y}_{1:i-1})$
 end while
 $y_i = y_{i,j}$
end for
return $\mathbf{y}$

Let T and Q denote the total number of iterations (also samples) and reward queries (also queries to q_0) respectively, and $\tau \sim \text{Unif}(\{0, \dots, T-1\})$. Then, with $Q = \tilde{\mathcal{O}}((m + \epsilon^{-1})/\gamma^2)$ reward queries and $T = \tilde{\mathcal{O}}(1/(\gamma^2\epsilon))$ iterations, we have $\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \epsilon$.

Remark. Compared to Corollary 4, the number of PG iterations and context samples only scale with $\tilde{\mathcal{O}}(1/(\gamma^2\epsilon))$. On the other hand, the number of reward queries (and sampling queries from q_0) still blows up exponentially for $\epsilon < \mathcal{Q}_{q_0}^{-1}(k^{-N})$, thus still facing the base model barrier.

Remark. As we prove in Theorem 10, both the number of samples and reward queries above are minimax optimal (up to log factors) among learners with access to q_0 .

Unlike Corollary 4, the behavior policy in Corollary 5 depends on the current policy $p_{\mathbf{w}_t}$, taking advantage of online exploration. If we remove line 1 from Algorithm 1, then the expected number of reward queries will always be exponential in N . When $\epsilon < \mathcal{Q}_{q_0}^{-1}(k^{-N})$, we could choose $m = \infty$ and obtain the same bounds in expectation. In this case, PG and SGD coincide, and we only rely on Algorithm 1 to recover the ground-truth label of $\mathbf{x}^{(t)}$.

4 Policy Gradient with Process Rewards

In the previous section, we demonstrated the hardness of post-training when we only have access to outcome rewards. An alternative is to provide the model with reward feedback throughout the generation process [SNF⁺25, CYW⁺25]. While no longer a contextual bandit problem, the policy gradient iterates can be written as

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{w}_t - \eta_t \nabla_{\mathbf{w}} \ell_t(\mathbf{w}_t), \\ \ell_t(\mathbf{w}_t) &:= - \sum_{i=1}^N A_i^{(t)} \log p_{\mathbf{w}_t}(y_i^{(t)} | \mathbf{x}, \mathbf{y}_{1:i-1}^{(t)}) \end{aligned} \quad (\text{PG-PR})$$

where $\mathbf{y}^{(t)} \sim q_t(\cdot | \mathbf{x}_t)$ is sampled from some behavior policy, and $A_i^{(t)}$ is the advantage of taking action $y_i^{(t)}$ at step i of rollout t . Suppose we have a 0-1 reward function $r^* : \mathcal{X} \times \mathcal{Y}^+ \rightarrow \{0, 1\}$ that satisfies

$$r^*(\mathbf{x}, \mathbf{y}_{1:i}) = \mathbb{1}[\mathbf{y}_{1:i} = \mathbf{y}_{1:i}^*(\mathbf{x})], \quad \forall i \in [N]. \quad (4.1)$$

The most basic estimator for advantage would be the (undiscounted) return from taking action $y_i^{(t)}$, i.e. $A_i^{(t)} = \sum_{j \geq i} r^*(\mathbf{x}^{(t)}, \mathbf{y}_{1:j}^{(t)})$. However, we can achieve optimal rates by using a simpler advantage $A_i^{(t)} = r^*(\mathbf{x}, \mathbf{y}_{1:i}^{(t)})$, which is what we consider in the following. This form of reward allows us to efficiently explore

sequences with a base model that is only accurate for next token prediction (and not predicting the entire sequence). In particular, we can define the Token-Level Likelihood Quantile function $\mathcal{Q}_q^{\text{TL}} : [0, 1] \rightarrow [0, 1]$ as

$$\mathcal{Q}_q^{\text{TL}}(\varepsilon) := \sup \{ \alpha \in [0, 1] : \mathbb{P}_{\mathbf{x}} \left(\min_{i \in [N]} q(y_i^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x})) \leq \alpha \right) \leq \varepsilon \}$$

Unlike LQ, for a uniform base policy q , we have $\mathcal{Q}_q^{\text{TL}}(\varepsilon) = k^{-1}$ for all $\varepsilon < 1$, notably independent of N . This definition leads to the following guarantee. The proofs of this section are deferred to Section B.

Theorem 6 (PG-PR with Adaptive Learning Rate). *Suppose Assumption 1 holds. Given a base policy q_0 and for any $\varepsilon \in (0, 1)$, let q_t be the output of Algorithm 2 with $m = \lceil 2(\log N + 1) \cdot \min(\mathcal{Q}_{q_0}^{\text{TL}}((1 - o(1))\varepsilon)^{-1}, k) \rceil$. Let T and Q denote the total number of post-training iterations (also samples) and reward (also sampling from q_0) queries respectively, and $\tau \sim \text{Unif}(\{0, \dots, T - 1\})$. Consider (PG-PR) iterates with the adaptive learning rate $\eta_t = (4 + 2\|\nabla_{\mathbf{w}} \ell_t(\mathbf{w}_t)\|)^{-1}$. Then, with $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$ iterations and*

$$Q = \tilde{\mathcal{O}} \left(\frac{N \min(\mathcal{Q}_{q_0}^{\text{TL}}((1 - o(1))\varepsilon)^{-1}, k)}{\gamma^2} + \frac{1}{\gamma^2 \varepsilon} \right)$$

reward queries we have $\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$.

We immediately observe that compared to Corollary 5 which required $\tilde{\mathcal{O}}((k^N + \varepsilon^{-1})/\gamma^2)$ reward queries in the worst-case, the above theorem requires only $\tilde{\mathcal{O}}((Nk + \varepsilon^{-1})/\gamma^2)$ reward queries. Further, the requirement on the base model is much milder for this theorem. q_0 only needs to generate the next correct token conditioned on a partially correct response. One can observe that for any base model q_0 , we have $\mathcal{Q}_{q_0}^{\text{TL}}(\varepsilon) \geq \mathcal{Q}_{q_0}(\varepsilon)$. In certain settings, the benefit of process reward models will be significant, and they will enable PG to go beyond the support of the base model and achieve significantly smaller expected test error. One such example is provided below.

Corollary 7. *Suppose Assumption 1 holds, and additionally, the feature map satisfies Assumption 2. Let $(\mathbf{w}_t^{\text{SGD}})_{t=0}^{n-1}$ denote the SGD iterates with constant learning rate $\eta = 1/(2N)$, and define $q_0 = p_{\bar{\mathbf{w}}_n^{\text{SGD}}}$ where $\bar{\mathbf{w}}_n^{\text{SGD}} = \frac{1}{n} \sum_{t=0}^{n-1} \mathbf{w}_t^{\text{SGD}}$. Let $(\mathbf{w}_t^{\text{PG}})_{t=0}^{T-1}$ denote the (PG-PR) iterates in the setting of Theorem 6, and let $\tau \sim \text{Unif}(\{0, \dots, T - 1\})$.*

1. *To achieve $\mathbb{E}[p_{\bar{\mathbf{w}}_n^{\text{SGD}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$, we have a supervised pre-training sample complexity of $n = \tilde{\mathcal{O}}(N/(\gamma^2 \varepsilon))$.*
2. *To achieve $\mathbb{E}[p_{\mathbf{w}^{\text{PG}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$, we have a post-training iteration complexity $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$. Further, the sufficient number of reward queries is given by*

$$Q = \begin{cases} \tilde{\mathcal{O}}((N + \varepsilon^{-1})/\gamma^2) & \text{if } n = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon)) \\ \tilde{\mathcal{O}}((Nk + \varepsilon^{-1})/\gamma^2) & \text{otherwise} \end{cases}.$$

The above clearly demonstrates (PG-PR) with an SGD-trained base policy as an optimal combination of PG and SGD. While learning with (constant LR) SGD alone would require a number of supervised samples scaling linearly with N , and doing PG alone would require a number of reward queries scaling with Nk ; their combination removes the N factor in the supervised sample complexity and the k factor in the number of reward queries. We remark however that this still does not outperform SGD with adaptive LR with the same number of supervised samples. We leave such an attempt to future work.

5 Lower Bounds

In this section, we argue about the tightness and optimality of the policy gradient algorithms studied for post-training, and of SGD with adaptive LR for pre-training. To establish post-training optimality, it is natural to define classes of algorithms with access to a base model and a reward function. Our lower bounds cover both online learning settings where the learner receives only one reward feedback per iteration and its

goal is to minimize the number of mistakes, and the statistical setting where the learner has access to the reward model and can make multiple queries on the same context, and its goal is to have high accuracy on a new i.i.d. sample.

We are specifically interested in the role the base model plays in the lower bounds. Hence, we formalize the definition of online and statistical learners with access to a base model as an oracle, and we believe these definitions might be useful in subsequent works on this topic.

Definition 8 (Online Learner with Bandit Feedback). *An online learner with access to base model $q(\cdot | \mathbf{x})$ is a potentially randomized algorithm that at every iteration receives $\mathbf{x}_t$, predicts $\mathbf{y}_t$ (using its history and the base model), receives reward $r(\mathbf{x}_t, \mathbf{y}_t)$, and proceeds to the next round.*

Definition 9 (Statistical Learner with Bandit Feedback). *A statistical learner with access to base model $q(\cdot | \mathbf{x})$ is a potentially randomized algorithm that receives n samples $\{\mathbf{x}_t\}_{t=1}^n$, and on sample t makes m_t reward queries $\{r(\mathbf{x}_t, \mathbf{y}_t^{(i)})\}_{i=1}^{m_t}$ where $\mathbf{y}_t^{(i)}$ can be chosen adaptively using access to the base model. The algorithm then constructs the training set $S = \{(\mathbf{x}_t, \mathbf{y}_t^{(i)}, r(\mathbf{x}_t, \mathbf{y}_t^{(i)})) : t \in [n], i \in [m_t]\}$ and returns a prediction rule $\hat{\mathbf{y}}(S) : \mathcal{X} \rightarrow \mathcal{Y}^N$.*

The performance of an online learner is measured by the number of mistakes it makes over T rounds, i.e. $\sum_{t=0}^{T-1} \mathbb{1}[\mathbf{y}_t \neq \mathbf{y}^*(\mathbf{x}_t)]$, while the performance of a statistical learner is measured by its expected error over a new test sample, i.e. $\mathbb{P}_{\mathbf{x}}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}))$. Note that (PG-OR) is an example of both algorithms. It can be used to make online predictions as $\mathbf{y}_t \sim p_{\mathbf{w}_t}(\cdot | \mathbf{x}_t)$ (assuming the behavior policy does not make additional reward queries), and it can also be used to make predictions on new test samples at the end of training via $\hat{\mathbf{y}} \sim p_{\bar{\mathbf{w}}}(\cdot | \mathbf{x})$ where $\bar{\mathbf{w}}$ is the averaged iterate.

With the above definitions, we are ready to state the lower bound for statistical learners, which establishes the near optimality of the PG variant with adaptive learning rate presented in Corollary 5. The proofs of this section are deferred to Section C.

Theorem 10 (Lower Bound for Statistical Learner under ORM). *Let $(\mathbf{x}_t)_{t \geq 0}$ be an i.i.d. sequence drawn from some distribution $\mathcal{D}$. Consider the statistical learner of Definition 9, where r is the outcome reward. For every $\gamma \in (0, N^{-1})$, $\alpha \in [k^{-N}, 0.5]$, $\varepsilon^*, \varepsilon \leq c$ for a sufficiently small absolute constant $c > 0$, and $D \geq k/\gamma^2$, there exists a base policy q satisfying*

$$\mathcal{Q}_q(\varepsilon) = \begin{cases} \alpha & \text{if } \varepsilon \geq \varepsilon^* \\ k^{-N} & \text{if } \varepsilon < \varepsilon^* \end{cases} \quad (5.1)$$

and a feature map ϕ , such that for every learner with access to q , there exist a distribution $\mathcal{D}$ and $\mathbf{w}^*$ that satisfy Assumption 1 with margin γ and to achieve expected test error less than ε the learner requires at least $Q = \tilde{\Omega}(\mathcal{Q}_q((1 + o(1))\varepsilon)^{-1}/\gamma^2)$ reward queries and at least $T = \Omega(1/(\gamma^2\varepsilon))$ samples.

We highlight a contrast with the result of [JRX25]. They show that learning a reward model using trajectory samples can be done efficiently with a polynomial dependence on N . Our lower bound however, states that the number of reward queries, either true reward or queries to a learned reward model, may have to scale exponentially with N in the worst-case, which provides a computational, rather than a statistical, barrier.

Next, we establish the near optimality of the online PG variant with adaptive learning rate of Corollary 4 among online learners.

Theorem 11 (Lower Bound for Online Learner under ORM). *Consider the same setting as Theorem 10, except with the online learner of Definition 8. There exists a base policy q satisfying (5.1) and a feature map ϕ such that for every learner, there is a distribution such that to achieve expected test error ε on a new sample the learner needs at least $T = \tilde{\Omega}(\mathcal{Q}_q((1 + o(1))\varepsilon)^{-1}/(\gamma^2\varepsilon))$ reward queries (equivalently samples). Further, over $T \gtrsim k^N/(\gamma^2\varepsilon^*N)$ rounds the expected number of mistakes is at least $\Omega(k^N/(\gamma^2N))$.*

5.1 Lower Bounds on the Base Model

So far we established that under an ORM, the optimal number of reward queries to achieve expected error ε scales as $\mathcal{Q}_q(\varepsilon)^{-1}/\gamma^2$. One may hope that there is a supervised pre-training algorithm that with a modest

number of samples n can achieve an LQ that is polynomially (rather than exponentially) small in N . The following lower bound shows that such an algorithm does not exist in general.

Theorem 12 (Lower Bound for Pre-Training). *Let $S_n := \{(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))\}_{t=1}^n$ be n i.i.d. samples drawn from some distribution $\mathcal{D}$. Consider a learner that receives S_n and outputs a learned weight $\hat{\mathbf{w}}(S_n)$ and a corresponding autoregressive base policy $q = p_{\hat{\mathbf{w}}(S_n)}$. For every $\gamma \in (0, 0.5]$, $D \geq k/\gamma^2$, and $\varepsilon < c$ for some absolute constant $c > 0$, there exists a distribution $\mathcal{D}$ that satisfies Assumption 1, we have $\mathcal{Q}_q(\varepsilon) \leq k^{-N}$ for all $n \lesssim 1/(\gamma^2\varepsilon)$.*

The above lower bound implies that to have $\mathcal{Q}_q(\varepsilon) > k^{-N}$, any pre-training algorithm that outputs an autoregressive linear model needs $n \gtrsim 1/(\gamma^2\varepsilon)$. On the other hand, SGD (with adaptive LR) has an expected error of $\tilde{\mathcal{O}}(1/(\gamma^2n))$ over n samples. Thus with $n \gg 1/(\gamma^2\varepsilon)$ samples, SGD can already achieve an expected error of order ε (up to log terms), and there will be no significant improvement coming from PG in terms of expected error. This argument shows that the barrier of the base model is not an artifact of our analysis or specific to pre-training with SGD, it is a fundamental property of post-training with outcome rewards.

6 Experiments

To validate intuitions from our theory, we perform experiments on synthetic as well as a mathematical reasoning tasks.

For the mathematical reasoning task, we use two models from the Qwen3 family [YLY+25], namely Qwen3-8B-Base which is a pre-trained base model, and Qwen3-8B which has undergone post-training for improved reasoning capabilities. These models are evaluated on the Math-500 benchmark, a dataset consisting of 500 problems from the MATH dataset [LKB+24]. To estimate the likelihood of generating a correct answer, we generate 64 completions for each problem, and estimate the likelihood as the ratio of correct over all completions. We use temperature 0.7, top-p 0.95, and top-k 50 for generation. The resulting likelihood quantile functions are plotted in Figure 2a.

In the synthetic task, $\mathbf{x}$ is drawn according to either a mixture model or uniformly from the hypercube $\{\pm 1\}^d$. For the mixture model, to generate $\mathbf{x}$, we sample uniformly from the standard basis vectors, add an isotropic Gaussian noise with per-dimension standard deviation $0.05/\sqrt{d}$ whose norm is clipped at 0.05, and project the resulting vector on $\mathbb{S}^{d-1}(\sqrt{d})$.

To generate $\mathbf{y}^*(\mathbf{x})$, we first fix $\mathbf{W}_1^* \in \mathbb{R}^{k \times d}$ and $\mathbf{W}_2^* \in \mathbb{R}^{k \times k}$ once by drawing their entries i.i.d. from $\mathcal{N}(0, 1)$. Then $y_1^*(\mathbf{x}) = \arg \max(\mathbf{W}_1^* \mathbf{x})$ and $y_i^* = \arg \max(\mathbf{W}_1^* \mathbf{x} + \mathbf{W}_2^* \mathbf{e}_{y_{i-1}^*(\mathbf{x})})$. Since y_{i+1} can be predicted from $\mathbf{x}$ and y_i , we can use the feature map $\phi(\mathbf{x}, \mathbf{y}_{1:i}) = \text{Concat}(\text{Vec}(\mathbf{e}_{y_i} \mathbf{x}^\top), \text{Vec}(\mathbf{e}_{y_i} \mathbf{e}_{y_{i-1}}^\top))$ if $i > 1$ and $\phi(\mathbf{x}, y) = \text{Concat}(\text{Vec}(\mathbf{e}_y \mathbf{x}^\top), \mathbf{0}_{k^2})$ otherwise. Thus $\phi(\cdot, \cdot) \in \mathbb{R}^{dk+k^2}$.

In Figure 1, we use the mixture dataset described above with $N = 128$ and $d = k = 32$. To obtain the base model, we train a linear autoregressive model initialized from zero using Adagrad on negative log-likelihood for 1000 steps with a learning rate of 0.1 and a batch size of 256. We then use this model to initialize on-policy PG with either outcome rewards of Section 3 or process rewards of Section 4. For PG, we use Adagrad with 4000 steps, 0.1 learning rate, and a batch size of 1024. Each batch is a fresh draw of i.i.d. samples during both pre- and post-training.

For Figure 1, we define off-support samples as mixture centers that have a likelihood under the base model smaller than 10^{-12} . 4 out of 32 centers satisfy this constraint. We plot the average likelihood of generating a correct response over these 4 samples throughout PG training. PG with ORM is unable to boost the likelihood over any of these 4 samples, while the average likelihood improves under PRM. We further plot the expected error over 1024 i.i.d. test samples throughout PG. We also plot the likelihood trajectory of 16 out of 32 mixture centers under ORM, where we can observe samples that start near 0 stay near 0. This is in line with Theorem 3. For these samples, the initial likelihood is $\alpha \asymp k^{-N}$, and improving over the base model takes exponentially many iterations. However, samples with larger α show progress throughout PG.

In Figure 2b, we illustrate the evolution of the LQ function over the supervised training trajectory. For this experiment, we use the uniform over hypercube distribution for $\mathbf{x}$, and set $N = 128$, $d = 32$, and $k = 10$. We use the Adagrad optimizer with learning rate 1.0 and batch size 256.

The results can be reproduced using <https://github.com/mousaviah/rlvr-base-model-barrier>.

7 Conclusion

In this paper, we studied variants of policy gradient for post-training an autoregressive linear model, where the responses satisfy a separability assumption. We proved that with an outcome reward model, PG can efficiently increase the likelihood of samples inside the support of the base model. To go outside this support and achieve test error significantly smaller than the base model, the number of reward queries depends on a property of the base model we call the likelihood quantile, and can be exponential in the sequence length. This issue can be alleviated using a process reward model, where we introduce a token-level likelihood quantile and show that it scales more favorably with sequence length. We further demonstrated the tightness of our analysis, as well as fundamental statistical limitations on LQ, through minimax lower bounds.

While we assume access to accurate process rewards, it is interesting to find algorithms that can learn such a reward with statistical and computational efficiency whenever possible. Another open question is to extend the results of this paper to settings with noisy and non-separable responses, and develop optimal post-training algorithms accordingly.

Acknowledgments

The authors would like to thank Denny Wu, Juno Kim, and anonymous reviewers for helpful discussions and feedback. MAE was partially supported by the NSERC Grant [2019-06167], the CIFAR AI Chairs program, the CIFAR Catalyst grant, and the Ontario Early Researcher Award.

References

- [AKLM21] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76, 2021.
- [BPS⁺19] Alina Beygelzimer, David Pal, Balazs Szorenyi, Devanathan Thiruvengatathari, Chen-Yu Wei, and Chicheng Zhang. Bandit multiclass linear classification: Efficient algorithms for the separable case. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 624–633. PMLR, 2019.
- [CHG⁺25] Fan Chen, Audrey Huang, Noah Golowich, Sadhika Malladi, Adam Block, Jordan T Ash, Akshay Krishnamurthy, and Dylan J Foster. The coverage principle: How pre-training enables post-training. *arXiv preprint arXiv:2510.15020*, 2025.
- [CJRX25] Fan Chen, Zeyu Jia, Alexander Rakhlin, and Tengyang Xie. Outcome-based online reinforcement learning: Algorithms and fundamental limits. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [CYW⁺25] Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Yuchen Zhang, Jiacheng Chen, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, et al. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*, 2025.
- [DH13] Amit Daniely and Tom Helbertal. The price of bandit information in multiclass online classification. In *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 93–104. PMLR, 2013.
- [DHS11] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

- [FMR25] Dylan J Foster, Zakaria Mhammedi, and Dhruv Rohatgi. Is a good foundation necessary for efficient reinforcement learning? The computational role of the base model in exploration. In Nika Haghtalab and Ankur Moitra, editors, *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 2026–2142. PMLR, 2025.
- [GYZ⁺25] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [JRX25] Zeyu Jia, Alexander Rakhlin, and Tengyang Xie. Do we need to verify step by step? Rethinking process supervision from a theoretical perspective. In *Forty-second International Conference on Machine Learning*, 2025.
- [JT19] Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1772–1798. PMLR, 25–28 Jun 2019.
- [KD26] Aayush Karan and Yilun Du. Reasoning with sampling: Your base model is smarter than you think. In *The Fourteenth International Conference on Learning Representations*, 2026.
- [KSST08] Sham M Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. Efficient bandit algorithms for online multiclass prediction. In *Proceedings of the 25th international conference on Machine learning*, pages 440–447, 2008.
- [KWLS25] Juno Kim, Denny Wu, Jason D. Lee, and Taiji Suzuki. Metastable dynamics of chain-of-thought reasoning: Provable benefits of search, RL and distillation. In *Forty-second International Conference on Machine Learning*, 2025.
- [LKB⁺24] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2024.
- [LZBY20] Yanli Liu, Kaiqing Zhang, Tamer Basar, and Wotao Yin. An improved analysis of (variance-reduced) policy gradient and natural policy gradient methods. *Advances in Neural Information Processing Systems*, 33:7624–7636, 2020.
- [MXSS20] Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International conference on machine learning*, pages 6820–6829. PMLR, 2020.
- [RSSW24] Hrithik Ravi, Clayton Scott, Daniel Soudry, and Yutong Wang. The implicit bias of gradient descent on separable multiclass data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [Sha21] Ohad Shamir. Gradient methods never overfit on separable data. *Journal of Machine Learning Research*, 22(85):1–20, 2021.
- [SHN⁺18] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70):1–57, 2018.
- [SK25] Matan Schliserman and Tomer Koren. Multiclass loss geometry matters for generalization of gradient descent in separable classification. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

- [SLX⁺25] Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, and Luke Zettlemoyer. Spurious rewards: Rethinking training signals in rlvr. *arXiv preprint arXiv:2506.10947*, 2025.
- [SNF⁺25] Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang Geng, Jacob Eisenstein, Rishabh Agarwal, Alekh Agarwal, Jonathan Berant, and Aviral Kumar. Rewarding progress: Scaling automated process verifiers for LLM reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [SRLK25] Amrith Setlur, Nived Rajaraman, Sergey Levine, and Aviral Kumar. Scaling test-time compute without verification or RL is suboptimal. In *Forty-second International Conference on Machine Learning*, 2025.
- [SSBD14] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [SWD⁺17] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [SWZ⁺24] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- [TDG⁺25] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.
- [TMUK26] Nikolaos Tsilivis, Eran Malach, Karen Ullrich, and Julia Kempe. How reinforcement learning after next-token prediction facilitates learning. In *The Fourteenth International Conference on Learning Representations*, 2026.
- [vdHFCB21] Dirk van der Hoeven, Federico Fusco, and Nicolò Cesa-Bianchi. Beyond bandit feedback in online multiclass classification. *Advances in neural information processing systems*, 34:13280–13291, 2021.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [Wil92] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- [WLZ⁺26] Xumeng Wen, Zihan Liu, Shun Zheng, Shengyu Ye, Zhirong Wu, Yang Wang, Zhijian Xu, Xiao Liang, Junjie Li, Ziming Miao, Jiang Bian, and Mao Yang. Reinforcement learning with verifiable rewards implicitly incentivizes correct reasoning in base LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026.
- [WXL⁺25] Fang Wu, Weihao Xuan, Ximing Lu, Zaid Harchaoui, and Yejin Choi. The invisible leash: Why rlvr may not escape its origin. *arXiv preprint arXiv:2507.14843*, 2025.
- [YCL⁺25] Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in LLMs beyond the base model? In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [YLY⁺25] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang,

Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

- [ZKZB20] Kaiqing Zhang, Alec Koppel, Hao Zhu, and Tamer Basar. Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612, 2020.
- [ZNY⁺21] Junyu Zhang, Chengzhuo Ni, Zheng Yu, Csaba Szepesvari, and Mengdi Wang. On the convergence and sample efficiency of variance-reduced policy gradient method. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

A Omitted Results and Proofs of Section 3

Our core argument for PG upper bounds is that PG can be seen as online gradient descent on a convex loss. Therefore, we state and use the following general fact for online gradient descent.

Proposition 13. *Let $(\ell_t)_{t=0}^{T-1}$ be a sequence of convex and non-negative functions on $\mathbb{R}^D$. Consider online gradient descent updates $(\mathbf{w}_t)_{t=0}^T$, initialized from $\mathbf{w}_0$ and defined as*

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \nabla \ell_t$$

where $(\eta_t)_{t=0}^{T-1}$ is the learning rate sequence. Then:

1. If $\|\nabla^2 \ell_t\|_2 \leq L$ for some $L > 0$ and all t , let $\eta_t = \eta$ where $\eta \in (0, L^{-1})$. Then, for all $\mathbf{v} \in \mathbb{R}^D$ we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \ell_t(\mathbf{w}_t) \leq \frac{1}{1 - \eta L} \left(\frac{1}{T} \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2\eta T} \right). \quad (\text{A.1})$$

2. If $\|\nabla \ell_t\|_2 \leq C\ell_t$ for some $C > 0$ and all t , let $\eta_t = \eta/(\lambda + \|\nabla \ell_t\|_2)$ where $\lambda > 0$ and $\eta \in (0, 2C^{-1})$. Then, for all $\mathbf{v} \in \mathbb{R}^D$ we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{\ell_t(\mathbf{w}_t)}{1 + C\lambda^{-1}\ell_t(\mathbf{w}_t)} \leq \frac{1}{1 - C\eta/2} \left(\frac{1}{T} \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{\lambda\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2\eta T} \right). \quad (\text{A.2})$$

While the constant learning rate case is well-known in the literature, the adaptive learning rate case, inspired by [CHG⁺25], is particularly relevant for gradient descent to maximize log-likelihood on softmax-linear policies where the smoothness constant can grow with the sequence length.

To prove this proposition, we recall the following lemmas.

Lemma 14. *Suppose $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies $\|\nabla^2 f\|_2 \leq L$. Then, $\|\nabla f(x)\|_2^2 \leq 2L(f(x) - \inf f)$ for all $x \in \mathbb{R}^d$.*

Proof. By smoothness, for any pair $\mathbf{x}, \mathbf{x}'$ we have $f(\mathbf{x}') \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{x}' - \mathbf{x}\|_2^2$. The proof is completed by letting $\mathbf{x}' = \mathbf{x} - \frac{1}{L} \nabla f(\mathbf{x})$. $\square$

Lemma 15. *Suppose $(\mathbf{u}_t)_{t=0}^{T-1}$ is an arbitrary sequence of vectors, and starting from $\mathbf{w}_0$, define $\mathbf{w}_{t+1} := \mathbf{w}_t - \mathbf{u}_t$ for all $t \leq T-1$. Then, for any $\mathbf{v}$,*

$$\sum_{t=0}^{T-1} \langle \mathbf{u}_t, \mathbf{w}_t - \mathbf{v} \rangle \leq \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} + \frac{\sum_{t=0}^{T-1} \|\mathbf{u}_t\|_2^2}{2}.$$

Proof. For any $t \geq 0$,

$$\|\mathbf{w}_{t+1} - \mathbf{v}\|^2 = \|\mathbf{w}_t - \mathbf{v}\|^2 + \|\mathbf{u}_t\|^2 - 2\langle \mathbf{u}_t, \mathbf{w}_t - \mathbf{v} \rangle.$$

By rearranging the terms,

$$2 \sum_{t=0}^{T-1} \langle \mathbf{u}_t, \mathbf{w}_t - \mathbf{v} \rangle = \sum_{t=0}^{T-1} \|\mathbf{w}_t - \mathbf{v}\|^2 - \|\mathbf{w}_{t+1} - \mathbf{v}\|^2 + \sum_{t=0}^{T-1} \|\mathbf{u}_t\|^2 \leq \|\mathbf{w}_0 - \mathbf{v}\|^2 + \sum_{t=0}^{T-1} \|\mathbf{u}_t\|^2.$$

$\square$

We can now state the proof of the proposition.

Proof. [Proof of Proposition 13] By convexity of ℓ_t , for any $\mathbf{v} \in \mathbb{R}^D$ we have $\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{v}) \leq \langle \nabla \ell_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{v} \rangle$. Therefore

$$\sum_{t=0}^{T-1} \eta_t (\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{v})) \leq \sum_{t=0}^{T-1} \langle \eta_t \nabla \ell_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{v} \rangle$$

$$\leq \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} + \frac{\sum_{t=0}^{T-1} \eta_t^2 \|\nabla \ell_t(\mathbf{w}_t)\|_2^2}{2}. \quad (\text{A.3})$$

Now, suppose $\|\nabla^2 \ell_t\|_2 \leq L$. By Lemma 14 and non-negativity of ℓ_t , we have $\|\nabla \ell_t\|_2^2 \leq 2L\ell_t$, consequently

$$\sum_{t=0}^{T-1} \eta_t (\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{v})) \leq \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} + \sum_{t=0}^{T-1} \eta_t^2 L \ell_t(\mathbf{w}_t).$$

Recall that in the first case we choose $\eta_t = \eta < 1/L$, therefore by rearranging the above

$$(1 - \eta L) \sum_{t=0}^{T-1} \ell_t(\mathbf{w}_t) \leq \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2\eta},$$

which finishes the proof of the first case.

For the second case, we have $\eta_t = \eta/(\lambda + \|\nabla \ell_t(\mathbf{w}_t)\|_2)$. Thus $\eta_t \|\nabla \ell_t(\mathbf{w}_t)\|_2 \leq \eta$. Combining this fact with (A.3) yields

$$\begin{aligned} \sum_{t=0}^{T-1} \eta_t (\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{v})) &\leq \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} + \frac{\eta \sum_{t=0}^{T-1} \eta_t \|\nabla \ell_t(\mathbf{w}_t)\|_2}{2} \\ &\leq \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} + \frac{C\eta \sum_{t=0}^{T-1} \eta_t \ell_t(\mathbf{w}_t)}{2}. \end{aligned}$$

By rearranging the terms,

$$\begin{aligned} \left(1 - \frac{C\eta}{2}\right) \sum_{t=0}^{T-1} \eta_t \ell_t(\mathbf{w}_t) &\leq \sum_{t=0}^{T-1} \eta_t \ell_t(\mathbf{v}) + \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2} \\ &\leq \sum_{t=0}^{T-1} \frac{\eta \ell_t(\mathbf{v})}{\lambda} + \frac{\|\mathbf{v} - \mathbf{w}_0\|_2^2}{2}. \end{aligned}$$

Furthermore, notice that

$$\eta_t \ell_t(\mathbf{w}_t) = \frac{\eta \ell_t(\mathbf{w}_t)}{\lambda + \|\nabla \ell_t(\mathbf{w}_t)\|_2} \geq \frac{\eta \ell_t(\mathbf{w}_t)}{\lambda + C\ell_t(\mathbf{w}_t)}.$$

As a result

$$\left(1 - \frac{C\eta}{2}\right) \sum_{t=0}^{T-1} \frac{\ell_t(\mathbf{w}_t)}{1 + C\lambda^{-1}\ell_t(\mathbf{w}_t)} \leq \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{2\eta},$$

which concludes the proof of the second case. $\square$

With this general theorem in hand, we can prove the results of Section 3. Note that Proposition 1 is a special case of Theorem 2 and Theorem 3, where $q_t(\cdot | \mathbf{x}) = \delta_{\mathbf{y}^*(\mathbf{x})}$ is the Dirac measure. As a result, we only need to prove the two theorems.

A.1 Proof of Theorem 2

This theorem uses the first case of Proposition 13. Thus, we need to estimate the smoothness of log-likelihood, achieved by the following lemma.

Lemma 16. *For any $\mathbf{w} \in \mathbb{R}^D$, $\mathbf{x} \in \mathbb{R}^d$, and $\mathbf{y} \in \mathcal{Y}^N$, we have*

$$\nabla_{\mathbf{w}}^2 \log p_{\mathbf{w}}(\mathbf{y} | \mathbf{x}) = \sum_{i=1}^N -\mathbb{E}_{Y_i} [\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i) \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)^\top] + \mathbb{E}_{Y_i} [\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)] \mathbb{E}_{Y_i} [\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)^\top]$$

where $Y_i \sim p_{\mathbf{w}}(\cdot | \mathbf{x}, \mathbf{y}_{1:i-1})$. As a result, if $\|\phi(\mathbf{x}, \mathbf{y}_{1:i})\|_2 \leq R$ for all $\mathbf{x}, \mathbf{y}, i$, then $\|\nabla^2 \log p_{\mathbf{w}}(\mathbf{y} | \mathbf{x})\|_2 \leq NR^2$ for all $\mathbf{x}, \mathbf{y}$.

Proof. We have

$$\nabla_{\mathbf{w}}^2 \log p_{\mathbf{w}}(\mathbf{y} | \mathbf{x}) = \sum_{i=1}^N \nabla_{\mathbf{w}}^2 \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}).$$

Further,

$$\nabla_{\mathbf{w}} \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}) = \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y_i) - \mathbb{E}_{Y_i}[\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)].$$

Taking a second derivative yields

$$\begin{aligned} \nabla_{\mathbf{w}}^2 \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}) &= - \sum_{y=1}^k \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y) \nabla p_{\mathbf{w}}(y | \mathbf{x}, \mathbf{y}_{1:i-1}) \\ &= - \sum_{y=1}^k p_{\mathbf{w}}(y | \mathbf{x}, \mathbf{y}_{1:i-1}) \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y) \nabla \log p_{\mathbf{w}}(y | \mathbf{x}, \mathbf{y}_{1:i-1}) \\ &= - \mathbb{E}_{Y_i}[\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i) \nabla \log p_{\mathbf{w}}(Y_i | \mathbf{x}, \mathbf{y}_{1:i-1})] \\ &= - \mathbb{E}_{Y_i}[\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i) \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)^\top] + \mathbb{E}_{Y_i}[\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)] \mathbb{E}_{Y_i}[\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, Y_i)^\top], \end{aligned}$$

which completes the proof. $\square$

To handle settings that involve both i.i.d. and online settings, we state a more general version of Theorem 2 as a corollary of Proposition 13, and then show how Theorem 2 for i.i.d. data follows from this corollary.

Proposition 17. *Let $(\mathbf{w}_t)$ denote the (PG-OR) iterates. Suppose the sequence $(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))$ is chosen arbitrarily (in particular not necessarily i.i.d.), with the only constraint being Assumption 1. For simplicity, assume $\eta_t = 1/(2N)$ for all t and $\|\mathbf{w}_0\|_2 \leq 1/\gamma \cdot \log(kT\gamma^2)$. Then,*

$$\mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \right] \leq \frac{10N \log(kT\gamma^2)^2}{\gamma^2 T}, \quad (\text{A.4})$$

where the expectation is over the randomness in sampling $\mathbf{y}^{(t)} \sim q_t(\cdot | \mathbf{x}^{(t)})$.

Proof. Note that since $r(\mathbf{x}_t, \mathbf{y}_t) = \mathbb{1}[\mathbf{y}_t = \mathbf{y}^*(\mathbf{x}_t)]$, we can rewrite (PG-OR) as

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}).$$

For simplicity, throughout the proof we will use the shorthand $\mathbf{y}^{*(t)} := \mathbf{y}^*(\mathbf{x}^{(t)})$. Define

$$\ell_t(\mathbf{w}) := -r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{w}}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}).$$

Then ℓ_t is convex and $\|\nabla^2 \ell_t\|_2 \leq N$ by Lemma 16 and the fact that $r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \leq 1$. Therefore, for $\eta_t = 1/(2N)$, we can invoke the first case of Proposition 13 to write

$$\frac{1}{T} \sum_{t=0}^{T-1} -r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \leq \frac{2}{T} \sum_{t=0}^{T-1} -r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{v}}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) + \frac{2N \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T} \quad (\text{A.5})$$

for any $\mathbf{v} \in \mathbb{R}^D$. Next, we take expectation over the randomness of $\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}$ from both sides. In the above, the only random quantities are $(\mathbf{y}^{(t)})_{t=0}^{T-1}$ and $(\mathbf{w}_t)_{t=0}^{T-1}$, as we are considering an arbitrary sequence of context $(\mathbf{x}^{(t)})_{t=0}^{T-1}$. Note that $\mathbf{w}_t$ is independent of the draw of $\mathbf{y}^{(t)}$, therefore

$$\begin{aligned} \mathbb{E} \left[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \right] &= \mathbb{E} \left[\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \mathbb{E} \left[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) | \mathbf{w}_t \right] \right] \\ &= \mathbb{E} \left[q_t(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \right]. \end{aligned}$$

In the above, we used the fact that $\mathbb{E}[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) | \mathbf{w}_t] = q_t(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)})$. By plugging this into (A.5), we have

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \right] \leq \frac{2}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \log p_{\mathbf{v}}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) + \frac{2N \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}$$

$$\leq -\frac{2}{T} \sum_{t=0}^{T-1} \log p_{\mathbf{v}}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) + \frac{2N\|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}.$$

Let $\mathbf{v}^* = R_v \mathbf{w}^*$ where $\mathbf{w}^*$ is the unit-norm vector of Assumption 1, and r is to be chosen later. Then, for any $\mathbf{x}$,

$$\begin{aligned} -\log p_{\mathbf{v}}(\mathbf{y}^* | \mathbf{x}) &= -\sum_{i=1}^N \log p_{\mathbf{v}}(y_i^* | \mathbf{x}, \mathbf{y}_{1:i-1}^*) \\ &= \sum_{i=1}^N \log \left(1 + \sum_{y' \neq y_i^*} e^{R_v \langle \mathbf{w}^*, \phi(\mathbf{x}, \mathbf{y}_{1:i-1}^*, y') - \phi(\mathbf{x}, \mathbf{y}_{1:i}^*) \rangle} \right) \\ &\leq N k e^{-R_v \gamma}. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \right] &\leq 2N k e^{-R_v \gamma} + \frac{4N\|\mathbf{v}\|_2^2 + 4N\|\mathbf{w}_0\|_2^2}{T} \\ &\leq 2N k e^{-R_v \gamma} + \frac{4N R_v^2 + 4N\|\mathbf{w}_0\|_2^2}{T}. \end{aligned}$$

We can balance the first two terms by choosing $R_v = \frac{1}{\gamma} \log(kT\gamma^2)$, and obtain

$$\begin{aligned} \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t \log p_{\mathbf{w}_t}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) \right] &\leq \frac{6N}{\gamma^2 T} \log^2(kT\gamma^2) + \frac{2N\|\mathbf{w}_0\|_2^2}{T} \\ &\leq \frac{10N \log^2(kT\gamma^2)}{\gamma^2 T}, \end{aligned}$$

which completes the proof. $\square$

To finish the proof of Theorem 2, we can employ a simple online-to-batch conversion.

Proof. [Proof of Theorem 2] By taking expectation from both sides of (A.4), we have

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^*(\mathbf{x}_t) | \mathbf{x}_t) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}_t) | \mathbf{x}_t) \right] \leq \frac{10N \log(kT\gamma^2)^2}{\gamma^2 T}.$$

Since $\mathbf{x}_t$ is independent of $\mathbf{w}_t$ and q_t , we can replace $\mathbf{x}_t$ with the test sample $\mathbf{x}$ that has the same law as $\mathbf{x}_t$ and is independent from all $(\mathbf{w}_t)$ and (q_t) while preserving the expectation above, and obtain

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \right] \leq \frac{10N \log(kT\gamma^2)^2}{\gamma^2 T}.$$

Let

$$\mathcal{L} := \frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}).$$

Using the law of total expectation and the fact that $\mathcal{L} \geq 0$, we have $\mathbb{E}[\mathcal{L} | \mathcal{E}_\alpha] \mathbb{P}(\mathcal{E}_\alpha) \leq \mathbb{E}[\mathcal{L}]$. As a result,

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha \right] \leq \frac{10N \log(kT\gamma^2)^2}{\pi_\alpha \gamma^2 T}.$$

Note that on $\mathcal{E}_\alpha$, we have $q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha$ for all t , thus

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} -\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha \right] \leq \frac{10N \log(kT\gamma^2)^2}{\pi_\alpha \alpha \gamma^2 T}.$$

Using the convexity of negative log-likelihood and the Jensen inequality, we obtain

$$\mathbb{E} \left[-\log p_{\tilde{\mathbf{w}}_T^{\text{PG}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha \right] \leq \frac{10N \log(kT\gamma^2)^2}{\pi_\alpha \alpha \gamma^2 T}.$$

Finally, we can use the inequality $1 - x \leq -\log x$ for $x > 0$ to complete the proof. $\square$

A.2 Proof of Theorem 3

For this part, we are going to use the second case of Proposition 13. To do so, we will take advantage of the structure of the log-likelihood gradient characterized by the following lemma.

Lemma 18. *Let $R = \sup_{\mathbf{w}, \mathbf{x}, \mathbf{y}, i} \|\phi(\mathbf{x}, \mathbf{y}_{1:i})\|_2$. Then, for any $\mathbf{w} \in \mathbb{R}^D$, $\mathbf{x} \in \mathbb{R}^d$, $\mathbf{y} \in \mathcal{Y}^N$, and $i \in [N]$, we have*

$$\|\nabla_{\mathbf{w}} \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1})\|_2 \leq -2R \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}).$$

Therefore, by the triangle inequality, we also have

$$\|\nabla_{\mathbf{w}} \log p_{\mathbf{w}}(\mathbf{y} | \mathbf{x})\|_2 \leq -2R \log p_{\mathbf{w}}(\mathbf{y} | \mathbf{x}).$$

Proof. Let us define $p_y = p_{\mathbf{w}}(y | \mathbf{x}, \mathbf{y}_{1:i-1})$. Then,

$$\begin{aligned} \|\nabla_{\mathbf{w}} \log p_{\mathbf{w}}(y_i | \mathbf{x}, \mathbf{y}_{1:i-1})\|_2 &= \left\| \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y_i) - \sum_{y=1}^k \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y) p_y \right\|_2 \\ &\leq (1 - p_{y_i}) \|\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y_i)\|_2 + \sum_{y \neq y_i} p_y \|\phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y)\|_2 \\ &\leq 2R(1 - p_{y_i}) \\ &\leq -2R \log p_{y_i}, \end{aligned}$$

where the last inequality used the fact that $1 - x \leq -\log x$ for $x > 0$. $\square$

Similar to the previous section, we first prove a more general online guarantee for PG with adaptive learning rate.

Proposition 19. *Let $(\mathbf{w}_t)$ denote the (PG-OR) iterates. Suppose the sequence $(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))$ is chosen arbitrarily (in particular not necessarily i.i.d.), with the only constraint being Assumption 1. Assume $\eta_t = 1/(2(\lambda + \|\nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})\|_2))$ for all t and some $\lambda > 0$, and $\|\mathbf{w}_0\|_2 \leq 1/\gamma \cdot \log(NkT\gamma^2/\lambda)$. Then,*

$$\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \min\left(\frac{\lambda}{2}, -\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})\right) \right] \leq \frac{20\lambda \log(NkT\gamma^2/\lambda)^2}{\gamma^2 T}. \quad (\text{A.6})$$

Proof. We can make a similar observation as in Proposition 17 about the (PG-OR) update, and rewrite it as

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{w}_t + \eta_t r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \\ &= \mathbf{w}_t - \eta_t \nabla \ell_t(\mathbf{w}_t) \end{aligned}$$

for $\ell_t(\mathbf{w}_t) := -r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$, where we recall $\mathbf{y}^{*(t)} := \mathbf{y}^*(\mathbf{x}^{(t)})$. Note that ℓ_t is convex, and by Lemma 18, it satisfies $\|\nabla \ell_t\|_2 \leq 2\ell_t$. For simplicity, define $r_t := r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ and $\tilde{\ell}_t(\mathbf{w}) := -\log p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$. Recall the learning rate $\eta_t = (\lambda + \|\nabla \tilde{\ell}_t(\mathbf{w}_t)\|_2)^{-1}/2$, and note that η_t is only important when $r_t = 1$, in which case $\tilde{\ell}_t = \ell_t$, thus $\eta_t = (\lambda + \|\nabla \ell_t(\mathbf{w}_t)\|_2)^{-1}/2$. We can therefore invoke the second case of Proposition 13 and write

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{r_t \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\lambda^{-1} \tilde{\ell}_t(\mathbf{w}_t)} \leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{r_t \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\lambda^{-1} r_t \tilde{\ell}_t(\mathbf{w}_t)} \leq \frac{2}{T} \sum_{t=0}^{T-1} r_t \tilde{\ell}_t(\mathbf{v}) + \frac{2\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}$$

for any $\mathbf{v} \in \mathbb{R}^D$. Similar to the proof of Proposition 17, we can take expectation over the randomness of $\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}$ from both sides above, and use the independence of $\mathbf{y}^{(t)}$ and $\mathbf{w}_t$ as well as the fact that $\mathbb{E}[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) | \mathbf{w}_t] = q_t(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}) =: q_t$ to write

$$\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{q_t \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\lambda^{-1} \tilde{\ell}_t(\mathbf{w}_t)} \right] \leq \frac{2}{T} \sum_{t=0}^{T-1} q_t \tilde{\ell}_t(\mathbf{v}) + \frac{2\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}$$

Note that we have $\frac{\tilde{\ell}_t}{1 + 2\lambda^{-1} \tilde{\ell}_t} = \frac{(\lambda/2) \tilde{\ell}_t}{\lambda/2 + \tilde{\ell}_t} \geq \frac{1}{2} \min(\lambda/2, \tilde{\ell}_t)$. Thus,

$$\begin{aligned} \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t \min(\frac{\lambda}{2}, \tilde{\ell}_t) \right] &\leq \frac{4}{T} \sum_{t=0}^{T-1} q_t \tilde{\ell}_t(\mathbf{v}) + \frac{4\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T} \\ &\leq \frac{4}{T} \sum_{t=0}^{T-1} \tilde{\ell}_t(\mathbf{v}) + \frac{4\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}. \end{aligned}$$

Recall from the proof of Proposition 17 that for $\mathbf{v} = R_v \mathbf{w}^*$ we have $\tilde{\ell}_t(\mathbf{v}) \leq Nk e^{-R_v \gamma}$. Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t \min(\frac{\lambda}{2}, \tilde{\ell}_t(\mathbf{w}_t)) \right] &\leq 4Nk e^{-R_v \gamma^2} + \frac{4\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T} \\ &\leq 4Nk e^{-R_v \gamma^2} + \frac{4\lambda \|\mathbf{v}\|_2^2 + 4\lambda \|\mathbf{w}_0\|_2^2}{T} \\ &\leq 4Nk e^{-R_v \gamma^2} + \frac{8\lambda R_v^2 + 8\lambda \|\mathbf{w}_0\|_2^2}{T}. \end{aligned}$$

By choosing $R_v = 1/\gamma \cdot \log(NkT\gamma^2/\lambda)$, we obtain

$$\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t \min(\frac{\lambda}{2}, \tilde{\ell}_t(\mathbf{w}_t)) \right] \leq \frac{20\lambda \log(NkT\gamma^2/\lambda)^2}{\gamma^2 T},$$

which completes the proof. $\square$

Similar to the previous section, we can use an online-to-batch conversion to prove Theorem 3.

Proof. [Proof of Theorem 3] Note that by taking expectation from both sides of (A.6) and replacing $\mathbf{x}_t$ with $\mathbf{x}$ (as they are both independent of $\mathbf{w}_t$, q_t and have the same law) we obtain

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \min(\frac{\lambda}{2}, -\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})) \right] \lesssim \frac{\lambda \log(NkT\gamma^2/\lambda)^2}{\gamma^2 T}.$$

Using the inequality $1 - x \leq -\log x$ for all x , we obtain

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \min(\frac{\lambda}{2}, 1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})) \right] \lesssim \frac{\lambda \log(NkT\gamma^2/\lambda)^2}{\gamma^2 T}.$$

Now by choosing $\lambda = 2$, we can remove min and write

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) (1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})) \right] \lesssim \frac{\log(NkT\gamma^2)^2}{\gamma^2 T}.$$

Similar to the proof of Theorem 2, we can use the non-negativity of the LHS above along with the law of total expectation to write

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) (1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})) | \mathcal{E}_\alpha \right] \lesssim \frac{\log(NkT\gamma^2)^2}{\pi_\alpha \gamma^2 T}.$$

Observing that $q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha$ on π_α concludes the proof. $\square$

A.3 Proof of Corollary 4

Let $\tilde{\mathcal{E}}_{\tilde{\alpha}} := \{\mathbf{x} : q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \tilde{\alpha}\}$. Then, for any $\tilde{\alpha} \in (0, 1]$,

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] = \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}) + \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}^c] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}^c). \quad (\text{A.7})$$

We will consider two approaches to choose $\tilde{\alpha}$ for bounding this expression.

First, note that

$$\mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}^c) = \mathbb{P}\left(\frac{1}{2}q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) + \frac{k^{-N}}{2} < \tilde{\alpha}\right) \leq \mathbb{P}(q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq 2\tilde{\alpha}) = 1 - \pi_{2\tilde{\alpha}}$$

where we recall $\pi_\alpha := \mathbb{P}(q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha)$. Then, we can rewrite (A.7) as

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}) + 1 - \pi_{2\tilde{\alpha}}. \quad (\text{A.8})$$

By choosing $\tilde{\alpha} = \frac{1}{2}\mathcal{Q}_{q_0}\left((1 - \frac{1}{\log(1/\varepsilon)})\varepsilon\right)$, we have $1 - \pi_{2\tilde{\alpha}} \leq \varepsilon(1 - \frac{1}{\log(1/\varepsilon)})$. Combined with the bound from Theorem 3 for the first term, we obtain

$$\begin{aligned} \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] &\leq \tilde{\mathcal{O}}\left(\frac{\tilde{\alpha}^{-1}}{\gamma^2 T}\right) + 1 - \pi_{2\tilde{\alpha}} \\ &\leq \tilde{\mathcal{O}}\left(\frac{\mathcal{Q}_{q_0}((1 - o(1))\varepsilon)^{-1}}{\gamma^2 T}\right) + \varepsilon(1 - \frac{1}{\log(1/\varepsilon)}), \end{aligned}$$

where $o(1)$ is hiding the $1/\log(1/\varepsilon)$ factor. By letting $T = \tilde{\mathcal{O}}(\tilde{\alpha}^{-1}/(\gamma^2 \varepsilon))$, we can derive the first term below $\varepsilon/\log(1/\varepsilon)$, and the total expectation below ε . This gives us a sufficient number of iterations

$$T = \tilde{\mathcal{O}}\left(\frac{\mathcal{Q}_{q_0}((1 - o(1))\varepsilon)^{-1}}{\gamma^2 \varepsilon}\right).$$

For the second approach, note that for $\tilde{\alpha} = \frac{k^{-N}}{2}$ we have $\mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}) = 1$. Thus, we can directly use Theorem 3 to write $\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}(k^N/(\gamma^2 T))$. Hence, another sufficient number of iterations is given by $T = \tilde{\mathcal{O}}(k^N/(\gamma^2 \varepsilon))$. We conclude the proof by considering the minimum between the two approaches. $\square$

A.4 Proof of Corollary 5

Recall the definitions

$$\begin{aligned} \mathcal{E}_\alpha &:= \{\mathbf{x} : q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha\}, \\ \tilde{\mathcal{E}}_{\tilde{\alpha}} &:= \{\mathbf{x} : q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \tilde{\alpha}\}. \end{aligned}$$

From Theorem 3,

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}) \leq \tilde{\mathcal{O}}\left(\frac{1}{\tilde{\alpha}\gamma^2 T}\right).$$

By inserting the above bound into (A.7) we obtain

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}\left(\frac{1}{\tilde{\alpha}\gamma^2 T}\right) + \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}^c). \quad (\text{A.9})$$

Let

$$m := \min\left(\left\lceil \mathcal{Q}_{q_0}\left((1 - \frac{1}{\log(1/\varepsilon)})\varepsilon\right)^{-1}, k^N \right\rceil\right).$$

We break down the argument based on which term attains the minimum.

- Suppose the minimum is attained by the first term. Since $q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - (1 - q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})/2)^m$, with $\alpha = m^{-1}$ we have $\tilde{\alpha} = 1 - e^{-1/2}$. Further,

$$\mathbb{P}(\mathcal{E}_\alpha^c) = 1 - \pi_\alpha \leq 1 - \pi_{m^{-1}} = \mathcal{Q}_{q_0}^{-1}(m^{-1}) \leq (1 - \frac{1}{\log(1/\varepsilon)})\varepsilon.$$

Since $\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) \leq \mathbb{P}(\mathcal{E}_\alpha^c)$, from (A.9) we have

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}\left(\frac{1}{\gamma^2 T}\right) + \varepsilon(1 - \frac{1}{\log(1/\varepsilon)}).$$

Thus $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$ is sufficient to ensure the first term is less than $\varepsilon/\log(1/\varepsilon)$ and the overall expectation is less than ε .

- Suppose instead that the minimum is instead attained by k^N . Note that

$$q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - (1 - k^{-N}/2)^m \geq 1 - e^{-1/2}.$$

Hence, this time for $\tilde{\alpha} = 1 - e^{-1/2}$ we have $\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) = 0$. As a result, (A.9) reads

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}\left(\frac{1}{\gamma^2 T}\right),$$

and taking $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$ is once again sufficient.

For the number of reward queries, we observe that every time $\mathbf{y}^{(t)} = \mathbf{y}^*(\mathbf{x}^{(t)})$ for $\mathbf{y}^{(t)} \sim p_{\mathbf{w}_t}$, Algorithm 1 makes one reward query, otherwise it makes at most $1 + m$ queries. The (PG-OR) also makes one reward query. Therefore

$$\begin{aligned} Q &= 2T + m \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\hat{\mathbf{y}}^{(t)} \neq \mathbf{y}^*(\mathbf{x}^{(t)})]\right] \\ &= 2T + m \sum_{t=1}^T \mathbb{E}\left[1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})\right] \\ &= 2T + m \sum_{t=1}^T \mathbb{E}[1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] && \text{(Using independence of } \mathbf{x}^{(t)} \text{ from } \mathbf{w}_t) \\ &\leq 2T + mT\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) + m\tilde{\mathcal{O}}(\gamma^{-2}) && \text{(Using (A.9) and } \tilde{\alpha} = 1 - e^{-1/2}). \end{aligned}$$

Recall that $\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) \leq \varepsilon(1 - o(1))$ and $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$. Therefore we can combine the middle and last terms on the RHS above, and write

$$Q \leq 2T + m\tilde{\mathcal{O}}(\gamma^{-2}) \leq \tilde{\mathcal{O}}((m + \varepsilon^{-1})/\gamma^2),$$

finishing the proof. $\square$

A.5 Fully On-Policy Policy Gradient

In this section, we will consider (PG-OR) when $q_t = p_{\mathbf{w}_t}$ for all t starting from a base model $q_0 = p_{\mathbf{w}_0}$. We first state this result and then explain the intuition behind it, and how it is qualitatively related to the observations in Section 3. We then proceed to prove the result.

Proposition 20 (Fully Online PG). *Consider the (PG-OR) updates under Assumption 1, with a base policy $q_0 = p_{w_0}$ and assume $\|w_0\|_2^2 \lesssim \log(Tk\gamma^2)/\gamma$ for simplicity. Let $\tau \sim \text{Unif}\{0, \dots, T-1\}$. For any $\delta \in (0, 1)$, with a proper choice of learning rate of order $\eta = \tilde{\Theta}(\gamma^2\delta/N)$, after*

$$T = \tilde{O}\left(\frac{N}{\gamma^4\delta^2\varepsilon} \cdot \max\left(1, \frac{1 - \mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]}{\delta}\right)\right)$$

we have $p_{w_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - \varepsilon$ with probability at least $\mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \delta$.

We can observe that with a fixed probability of $\mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \delta$, online PG can boost the likelihood to be $1 - \varepsilon$, with a rate that although suboptimal, still depends only linearly on the sequence length and polynomially on other parameters. Crucially, the role of the base model here is to determine the probability of success, which can at most be $\mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]$. While the above proposition does not characterize which samples at test-time win this lottery ticket, Figure 1 hints that these are samples where the initial likelihood is non-trivial. Further, we expect the above rate to be an artifact of our analysis and possibly improvable.

To prove the above proposition, we make a number of observations.

Lemma 21. *Recall $\sup_{w, \mathbf{x}, \mathbf{y}, i} \|\phi(\mathbf{x}, \mathbf{y}_{1:i})\|_2 \leq 1$. We then have $\sup_{w, \mathbf{x}, \mathbf{y}} \|\nabla_w^2 p_w(\mathbf{y} | \mathbf{x})\|_2 \lesssim 1$.*

Proof. Define the shorthand notations

$$p := p_w(\mathbf{y} | \mathbf{x}), \quad p_y^i := p_w(y | \mathbf{x}, \mathbf{y}_{1:i-1}), \quad p^i := p_w(y_i | \mathbf{x}, \mathbf{y}_{1:i-1}), \quad Y_i \sim p_w(\cdot | \mathbf{x}, \mathbf{y}_{1:i-1}).$$

We have

$$\nabla^2 p = p \nabla \log p \nabla \log p^\top + p \nabla^2 \log p.$$

Therefore,

$$\begin{aligned} \|\nabla^2 p\|_2 &\leq p \|\nabla \log p\|_2^2 + p \|\nabla^2 \log p\|_2 \\ &\leq p \left(\sum_{i=1}^N \|\nabla \log p^i\|_2 \right)^2 + p \sum_{i=1}^N \|\nabla^2 \log p^i\|_2 \end{aligned} \tag{A.10}$$

Further, from Lemma 18 with $R = 1$, recall that $\|\nabla \log p^i\|_2 \leq -2 \log p^i$. Also, from the proof of Lemma 16, recall that

$$\begin{aligned} \|\nabla^2 \log p^i\|_2 &= \left\| \mathbb{E}_{Y_i} [\phi_{Y_i} \phi_{Y_i}^\top] - \mathbb{E}_{Y_i} [\phi_{Y_i}] \mathbb{E}[\phi_{Y_i}]^\top \right\|_2 \\ &\leq \mathbb{E}_{Y_i} \left[\|\phi_{Y_i}\|_2^2 \right] - \|\mathbb{E}_{Y_i} [\phi_{Y_i}]\|_2^2 && \text{(Since } \|\text{Cov}_{Y_i}(\phi_{Y_i})\|_2 \leq \text{Tr}(\text{Cov}_{Y_i}(\phi_{Y_i}))\text{)} \\ &= \mathbb{E}_{Y_i} \left[\|\phi_{Y_i} - \mathbb{E}_{Y_i} [\phi_{Y_i}]\|_2^2 \right] \\ &\leq 2 \mathbb{E}_{Y_i} [\|\phi_{Y_i} - \mathbb{E}_{Y_i} [\phi_{Y_i}]\|_2] && \text{(Since } \|\phi_{Y_i} - \mathbb{E}_{Y_i} [\phi_{Y_i}]\|_2 \leq 2\text{)} \\ &\leq 2 \|\phi_{y_i} - \mathbb{E}_{Y_i} [\phi_{Y_i}]\|_2 + 2 \sum_{y \neq y_i} p_y^i \|\phi_y - \mathbb{E}[Y_i] \phi_{Y_i}\|_2 \\ &\leq 4(1 - p_i) + 4 \sum_{y \neq y_i} p_y^i && \text{(From the proof of Lemma 18)} \\ &\leq 8(1 - p_i) \\ &\leq -8 \log p_i, \end{aligned}$$

where $\phi_y := \phi(\mathbf{x}, \mathbf{y}_{1:i-1}, y)$. Plugging these estimates back into (A.10) yields

$$\|\nabla^2 p\|_2 \lesssim p \left(\left(\sum_{i=1}^N -\log p_i \right)^2 - \sum_{i=1}^N \log p_i \right) = e^{\log p} ((-\log p)^2 - \log p).$$

Using the fact that the function $x \mapsto e^{-x}(x^2 + x)$ is bounded over $x \geq 0$ completes the proof. $\square$

To apply Proposition 17, we need to lower bound q_τ where $\tau \sim \text{Unif}\{0, \dots, T-1\}$. In the case of on-policy PG, this means lower bounding $p_{\mathbf{w}_\tau}$ itself, which is the objective of Proposition 17. We thus need another tool to lower bound $p_{\mathbf{w}_\tau}$, and use this to achieve an improved conditional lower bound on $p_{\mathbf{w}_\tau}$ using 17. We establish this lower bound by using the fact that on-policy PG is effectively doing stochastic gradient ascent on the function $\mathbf{w} \mapsto \mathbb{E}[p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]$. Therefore, this quantity should not decrease significantly from its initial value along the PG dynamics if the learning rate is sufficiently small. This fact is established by the following lemma.

Lemma 22. *Let $(\mathbf{w}_t)$ denote the on-policy (PG-OR) iterates where $q_t = p_{\mathbf{w}_t}$, initialized from the base policy $p_{\mathbf{w}_0} = q_0$. Suppose Assumption 1 holds and $\|\mathbf{w}_0\|_2^2 \lesssim \log(Tk\gamma^2)/\gamma^2$ for simplicity. Define $\tau \sim \text{Unif}(\{0, \dots, T-1\})$. Then, for any constant learning rate $\eta \leq 1/(2N)$, we have*

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq \mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \frac{C\eta N \log(Tk\gamma^2)^2}{\gamma^2},$$

where $C > 0$ is an absolute constant.

Proof. Let $F(\mathbf{w}) := -\mathbb{E}[p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]$ and recall $\ell_t(\mathbf{w}) := -r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \log p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$. Also define $p_t := p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$ and $r_t := r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ for brevity. Recall that using the 0-1 structure of the reward, we can write (PG-OR) as

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \nabla \ell_t(\mathbf{w}_t).$$

Then

$$\mathbb{E}[\nabla \ell_t(\mathbf{w}_t) | \mathbf{w}_t] = -\mathbb{E}[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) | \mathbf{w}_t] = \nabla F(\mathbf{w}_t),$$

as expected from policy gradient, where we used $\mathbb{E}[r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) | \mathbf{x}^{(t)}, \mathbf{w}_t] = p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$. Further, Lemma 21 proves that $\|\nabla^2 F(\mathbf{w})\|_2 \leq L$ where $L = \mathcal{O}(1)$ for all $\mathbf{w}$. Hence, using the smoothness of F , we can write

$$\begin{aligned} F(\mathbf{w}_{t+1}) &\leq F(\mathbf{w}_t) - \langle \nabla F(\mathbf{w}_t), \eta \nabla \ell_t(\mathbf{w}_t) \rangle + \frac{L\eta^2}{2} \|\nabla \ell_t(\mathbf{w}_t)\|_2^2 \\ &= F(\mathbf{w}_t) - \langle \nabla F(\mathbf{w}_t), \eta \nabla \ell_t(\mathbf{w}_t) \rangle + \frac{L\eta^2 r_t}{2} \|\nabla \log p_t\|_2^2, \end{aligned}$$

where we used $r^2 = r$. By taking expectation from both sides, we obtain

$$\begin{aligned} \mathbb{E}[F(\mathbf{w}_{t+1})] &\leq \mathbb{E}[F(\mathbf{w}_t)] - \eta \mathbb{E}[\|\nabla F(\mathbf{w}_t)\|_2^2] + \frac{L\eta^2}{2} \mathbb{E}[p_t \|\nabla \log p_t\|_2^2] \\ &\leq \mathbb{E}[F(\mathbf{w}_t)] + \frac{L\eta^2 N}{2} \mathbb{E}[-p_t \log p_t] \quad (\text{By Lemma 16}) \\ &\leq \mathbb{E}[F(\mathbf{w}_0)] + \frac{L\eta^2 N}{2} \sum_{t=0}^{T-1} \mathbb{E}[-p_t \log p_t]. \end{aligned}$$

Therefore, using $\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] = \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]$, we have

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq \mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \frac{L\eta^2 N}{2} \sum_{t=0}^{T-1} \mathbb{E}[-p_t \log p_t]. \quad (\text{A.11})$$

To bound $\sum_{t=0}^{T-1} \mathbb{E}[-p_t \log p_t]$, we can follow the exact same path as the proof of Proposition 17, while keeping $\eta \leq 1/(2N)$ in all equations, which yields

$$\sum_{t=0}^{T-1} \mathbb{E}[-p_t \log p_t] \lesssim \frac{\log(Nk\gamma^2\eta T)^2}{\gamma^2\eta} + \frac{\|\mathbf{w}_0\|_2^2}{\eta T}. \quad (\text{A.12})$$

Using the bound $\|\mathbf{w}_0\|_2 \lesssim \log(Tk\gamma^2)/\gamma$, we obtain

$$\sum_{t=0}^{T-1} \mathbb{E}[-p_t \log p_t] \lesssim \frac{\log(Tk\gamma^2)^2}{\gamma^2 \eta}.$$

Substituting this bound into (A.11) results in

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq \mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \frac{CN\eta \log(Tk\gamma^2)^2}{\gamma^2},$$

for some absolute constant $C > 0$, where we used the fact that $L = \mathcal{O}(1)$. $\square$

With the tools above, we can state the proof of the proposition.

Proof. [Proof of Proposition 20] Let $\alpha \in (0, 1)$ to be fixed later. For convenience, we use the notation $p_{\mathbf{w}_\tau} := p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})$ where $(\mathbf{x}, \mathbf{y}^*(\mathbf{x}))$ is an independently drawn test sample. We rely on the following decomposition

$$\begin{aligned} \mathbb{P}(1 - p_{\mathbf{w}_\tau} \geq \varepsilon) &= \mathbb{P}(1 - p_{\mathbf{w}_\tau} \geq \varepsilon | p_{\mathbf{w}_\tau} \geq \alpha) \mathbb{P}(p_{\mathbf{w}_\tau} \geq \alpha) + \mathbb{P}(1 - p_{\mathbf{w}_\tau} \geq \varepsilon | p_{\mathbf{w}_\tau} < \alpha) \mathbb{P}(p_{\mathbf{w}_\tau} < \alpha) \\ &\leq \mathbb{P}(1 - p_{\mathbf{w}_\tau} \geq \varepsilon | p_{\mathbf{w}_\tau} \geq \alpha) \mathbb{P}(p_{\mathbf{w}_\tau} \geq \alpha) + \mathbb{P}(p_{\mathbf{w}_\tau} < \alpha) \\ &\leq \frac{1}{\varepsilon} \mathbb{E}[1 - p_{\mathbf{w}_\tau} | p_{\mathbf{w}_\tau} \geq \alpha] \mathbb{P}(p_{\mathbf{w}_\tau} \geq \alpha) + \frac{\mathbb{E}[1 - p_{\mathbf{w}_\tau}]}{1 - \alpha} \quad (\text{By Markov's Inequality}) \\ &\leq \frac{1}{\alpha \varepsilon} \mathbb{E}[p_{\mathbf{w}_\tau}(1 - p_{\mathbf{w}_\tau}) | p_{\mathbf{w}_\tau} \geq \alpha] \mathbb{P}(p_{\mathbf{w}_\tau} \geq \alpha) + \frac{\mathbb{E}[1 - p_{\mathbf{w}_\tau}]}{1 - \alpha} \\ &\leq \frac{1}{\alpha \varepsilon} \mathbb{E}[-p_{\mathbf{w}_\tau} \log p_{\mathbf{w}_\tau} | p_{\mathbf{w}_\tau} \geq \alpha] \mathbb{P}(p_{\mathbf{w}_\tau} \geq \alpha) + \frac{\mathbb{E}[1 - p_{\mathbf{w}_\tau}]}{1 - \alpha} \\ &\leq \frac{1}{\alpha \varepsilon} \mathbb{E}[-p_{\mathbf{w}_\tau} \log p_{\mathbf{w}_\tau}] + \frac{\mathbb{E}[1 - p_{\mathbf{w}_\tau}]}{1 - \alpha} \quad (\text{By the law of total expectation}) \\ &\leq \frac{C \log(Tk\gamma^2)^2}{\alpha \varepsilon \gamma^2 \eta T} + \frac{\mathbb{E}[1 - p_{\mathbf{w}_\tau}]}{1 - \alpha} \quad (\text{Using (A.12)}) \\ &\leq \frac{C \log(Tk\gamma^2)^2}{\alpha \varepsilon \gamma^2 \eta T} + \frac{1 - \mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]}{1 - \alpha} + \frac{C\eta N \log(Tk\gamma^2)^2}{\gamma^2(1 - \alpha)}, \quad (\text{By Lemma 22}) \end{aligned}$$

where $C > 0$ is an absolute constant. For a fixed α , the optimal (constant) learning rate minimizing the above bound is given by $\eta = \sqrt{(1 - \alpha)/(\alpha NT\varepsilon)}$. We will have to choose T sufficiently large so that this learning rate is also smaller than $1/(2N)$, needed to obtain (A.12). By using this optimal η , we obtain

$$\mathbb{P}(1 - p_{\mathbf{w}_\tau} \geq \varepsilon) \leq 1 - \mathbb{E}[q_0] + \frac{\alpha(1 - \mathbb{E}[q_0])}{1 - \alpha} + \tilde{\mathcal{O}}\left(\sqrt{\frac{N}{\alpha(1 - \alpha)\gamma^4 T\varepsilon}}\right),$$

where we used $q_0 := q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})$ for brevity. Next, we must choose α .

- If $1 - \mathbb{E}[q_0] \ll \delta$, then to achieve probability of error $1 - \mathbb{E}[q_0] + \delta$, the optimal α is $1/2$, and $T = \tilde{\mathcal{O}}(N/(\gamma^4 \varepsilon \delta^2))$ iterations are sufficient. Note that the learning rate will satisfy $\eta \asymp \sqrt{1/(NT\varepsilon)} = \tilde{\Theta}(\gamma^2 \delta/N) \ll 1/N$.
- If $1 - \mathbb{E}[q_0] \gg \delta$, then a suitable choice of α is $\alpha = \delta/(2(1 - \mathbb{E}[q_0]))$, and $T = \tilde{\mathcal{O}}(N(1 - \mathbb{E}[q_0])/(\gamma^4 \varepsilon \delta^3))$ iterations are sufficient. Note that the learning rate will satisfy $\eta \asymp \sqrt{\delta/((1 - \mathbb{E}[q_0])NT\varepsilon)} = \Theta(\gamma^2 \delta/N) \ll 1/N$.

Overall, we can conclude that with a learning rate of $\eta = \tilde{\Theta}(\gamma^2 \delta/N)$ and with

$$T = \tilde{\mathcal{O}}\left(\frac{N}{\gamma^4 \varepsilon \delta^2} \cdot \max\left(1, \frac{1 - \mathbb{E}[q_0]}{\delta}\right)\right)$$

iterations, we can achieve $p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - \varepsilon$ with probability at least $\mathbb{E}[q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] - \delta$. $\square$

A.6 Policy Gradient with Clipped Importance Weights

In this section, we study the effect of clipped importance weights on the convergence of policy gradients. Namely, we study updates of the form

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t r(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \rho_t \nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)}), \quad (\text{A.13})$$

where $\mathbf{y}^{(t)} \sim q_t(\cdot | \mathbf{x}^{(t)})$ and

$$\rho_t = \text{Clip}\left(\frac{p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})}{q_t(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})}, \zeta^{-1}, \zeta\right), \quad \zeta \geq 1.$$

We begin by adapting the conditional convergence results of (PG-OR) to use clipped importance weights, as outlined below.

Theorem 23 (PG-OR with Clipped Importance Weights). *Let $(\mathbf{w}_t)_{t=0}^{T-1}$ denote the iterates of PG with outcome rewards and clipped importance weights as in (A.13). Suppose Assumption 1 holds. For simplicity, assume $\|\mathbf{w}_0\|_2 \lesssim 1/\gamma \cdot \log(NTk\gamma^2)$. Let $\tau \sim \text{Unif}(\{0, \dots, T-1\})$ and $\mathbf{x}$ denote an i.i.d. test sample. For any $\alpha \in (0, 1]$, define the event $\mathcal{E}_\alpha := \{\min_t q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha\}$ and $\pi_\alpha := \mathbb{P}(\mathcal{E}_\alpha)$. Then*

1. *For the constant learning rate $\eta_t = 1/(2\zeta N)$ we have*

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha] \geq 1 - \tilde{\mathcal{O}}\left(\frac{\zeta^2 N}{\pi_\alpha \alpha \gamma^2 T}\right).$$

Further, the same bound holds for the averaged iterate $\bar{\mathbf{w}}_T^{\text{PG}} = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{w}_t$.

2. *For the adaptive learning rate $\eta_t = (4\zeta + 2\rho_t \|\nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})\|_2)^{-1}$ we have*

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha] \geq 1 - \tilde{\mathcal{O}}\left(\frac{\zeta^2}{\pi_\alpha \alpha \gamma^2 T}\right).$$

Proof. The proof of both statements follow by closely tracking the proof of Theorem 2 and Theorem 3 respectively. Since they are similar, we only state the proof of the second case.

We can recreate the online bound of Proposition 17 by retracing its proof steps and for learning rate $\eta_t = 1/(2(\zeta\lambda + \rho_t \|\nabla \log p_{\mathbf{w}_t}(\mathbf{y}^{(t)} | \mathbf{x}^{(t)})\|_2))$ obtain

$$\mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \min\left(\frac{\lambda}{2}, -\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})\right)\right] \lesssim \frac{\zeta^2 \log(NTk\gamma^2)^2}{\gamma^2 T}. \quad (\text{A.14})$$

Specifically, this time we define $\ell_t(\mathbf{w}) := -r_t \rho_t \log p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$. Importantly, a consequence of this definition is that while ρ_t depends on the weight $\mathbf{w}_t$, this weight is frozen and there is no gradient from ρ_t appearing in $\nabla \ell_t(\mathbf{w})$. Namely $\ell_t(\mathbf{w})$ is a still convex function of $\mathbf{w}$, and satisfies $\|\nabla \ell_t\|_2 \leq 2\ell_t$ thanks to Lemma 18. Further, notice that (A.13) is equivalent to $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \nabla \ell_t$ due to the 0-1 reward structure, and recall that we use $\tilde{\ell}_t(\mathbf{w}) := -\log p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})$ to denote negative log-likelihood. Therefore, we can invoke the second case of Proposition 13 with the learning rate given above, which yields

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{r_t \rho_t \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\zeta^{-1} \lambda^{-1} \rho_t r_t \tilde{\ell}_t(\mathbf{w}_t)} \leq \frac{2}{T} \sum_{t=0}^{T-1} r_t \rho_t \tilde{\ell}_t(\mathbf{v}) + \frac{2\zeta\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}.$$

Using the bound $\rho_t \in [\zeta^{-1}, \zeta]$ and $r_t \leq 1$, we can lower bound the RHS above and write

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{r_t \zeta^{-1} \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\lambda^{-1} \tilde{\ell}_t(\mathbf{w}_t)} \leq \frac{2}{T} \sum_{t=0}^{T-1} r_t \rho_t \tilde{\ell}_t(\mathbf{v}) + \frac{2\zeta\lambda \|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}.$$

Next, we take expectation from both sides and use the fact that $\mathbb{E}[r_t | \mathbf{x}^{(t)}, \mathbf{w}_t] = q_t(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) =: q_t$ and that

$$\mathbb{E}[r_t \rho_t | \mathbf{x}^{(t)}, \mathbf{w}_t] \leq q_t \max(p_{\mathbf{w}_t}/q_t, \zeta^{-1}) \leq 1.$$

We can therefore write

$$\frac{\zeta^{-1}}{T} \mathbb{E} \left[\sum_{t=0}^{T-1} \frac{q_t \tilde{\ell}_t(\mathbf{w}_t)}{1 + 2\lambda^{-1} \tilde{\ell}_t(\mathbf{w}_t)} \right] \leq \frac{2}{T} \sum_{t=0}^{T-1} \tilde{\ell}_t(\mathbf{v}) + \frac{2\zeta\lambda\|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}.$$

By following the remaining steps from the proof of Proposition 19, we arrive at (A.14). By letting $\lambda = 2$, using the inequality $1 - x \leq -\log x$, and employing the same online-to-batch conversion of replacing $\mathbf{x}^{(t)}$ with an independent copy $\mathbf{x}$, we can show

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})(1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})) \right] \lesssim \frac{\zeta^2 \log(NTk\gamma^2)^2}{\gamma^2 T}$$

or equivalently

$$\mathbb{E}[q_\tau(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})(1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}))] \lesssim \frac{\zeta^2 \log(NTk\gamma^2)^2}{\gamma^2 T}.$$

Combining the above bound with the fact that

$$\alpha \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_\alpha] \mathbb{P}(\mathcal{E}_\alpha) \leq \mathbb{E}[q_\tau(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})(1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}))]$$

finishes the proof. $\square$

Using the above theorem, we can also adapt the unconditional post-training guarantees of Corollary 4 and 5 to use clipped importance weights. We state this result without presenting its proof, which can simply be obtained by going through the proofs of Corollary 4 (in Section A.3) and the proof of Corollary 5 (in Section A.4) respectively.

Corollary 24. *Suppose Assumption 1 holds and a base policy q_0 is given. Let $(\mathbf{w}_t)_{t=0}^{T-1}$ denote the iterates of PG with outcome rewards and clipped importance weights as in (A.13), with the adaptive learning rate of Theorem 23. Let T denote the number of PG iterations, and Q denote the number of reward queries. Let $\tau \sim \text{Unif}(\{0, \dots, T-1\})$.*

1. *Suppose $q_t = \frac{1}{2}q_0 + \frac{1}{2}\text{Unif}$ for all t . Then, for any $\varepsilon \in (0, 1)$, with $Q = T = \tilde{O}\left(\frac{\zeta^2 \min(Q_{q_0}((1-o(1))\varepsilon)^{-1}, k^N)}{\gamma^2 \varepsilon}\right)$ iterations and reward queries, we achieve $\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$.*
2. *Suppose $q_t = q$ where q is the output of Algorithm 1 with $m = \min(\lceil Q_{q_0}((1-o(1))\varepsilon)^{-1} \rceil, k^N)$. Then, for any $\varepsilon \in (0, 1)$, with $Q = \tilde{O}(\zeta^2(m + \varepsilon^{-1})/\gamma^2)$ reward queries and $T = \tilde{O}(\zeta^2/(\gamma^2 \varepsilon))$ iterations, we achieve $\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \geq 1 - \varepsilon$.*

B Proofs of Section 4

Here we will state the proof of the results that appeared in Section 4.

B.1 Proof of Theorem 6

To prove Theorem 6, similar to the previous section, we first prove a general statement that handles any input sequence.

Proposition 25. *Let $(\mathbf{w}_t)$ denote the (PG-PR) iterates with $\mathbf{y}^{(t)} \sim q_t(\cdot | \mathbf{x}^{(t)})$. Suppose the sequence $(\mathbf{x}^{(t)}, \mathbf{y}^*(\mathbf{x}^{(t)}))$ is chosen arbitrarily (in particular not necessarily i.i.d.), with the only constraint being Assumption 1. Assume $\eta_t = 1/(2(\lambda + \|\nabla \ell_t(\mathbf{w}_t)\|_2))$ for all t and some $\lambda > 0$, and $\|\mathbf{w}_0\|_2 \leq 1/\gamma \cdot \log(NkT\gamma^2/\lambda)$. Then, for any $i \in [N]$, we have*

$$\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}_{1:i}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)}) \min\left(\frac{\lambda}{2}, -\log p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)})_{1:i} | \mathbf{x}^{(t)})\right) \right] \leq \frac{20\lambda \log(NkT\gamma^2/\lambda)^2}{\gamma^2 T}. \quad (\text{B.1})$$

Proof. The proof is similar to that of Proposition 19. Namely, using the fact that $A_t^{(j)} \neq 0$ if and only if $\mathbf{y}_{1:j}^{(t)} = \mathbf{y}_{1:j}^{*(t)}$ (where we recall the shorthand notation $\mathbf{y}^{*(t)} := \mathbf{y}^*(\mathbf{x}^{(t)})$). Then, we can write (PG-PR) as

$$\begin{aligned}\mathbf{w}_{t+1} &= \mathbf{w}_t + \eta_t \sum_{j=1}^N A_j^{(t)} \nabla \log p_{\mathbf{w}_t}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)}) \\ &= \mathbf{w}_t - \eta_t \nabla \ell_t(\mathbf{w}_t)\end{aligned}$$

where we recall $\ell_t(\mathbf{w}_t) := -\sum_{j=1}^N A_j^{(t)} \log p_{\mathbf{w}_t}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)})$. Once again, by non-negativity of the advantages $A_j^{(t)} \geq 0$, ℓ_t is convex. Therefore, we attempt at applying the second case of Proposition 13. For that, we need to establish a bound on $\|\nabla \ell_t\|_2$, which we do so by the triangle inequality,

$$\|\nabla \ell_t(\mathbf{w})\|_2 \leq \sum_{j=1}^N A_j^{(t)} \left\| \nabla \log p_{\mathbf{w}}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)}) \right\|_2 \leq -2 \sum_{j=1}^N A_j^{(t)} \log p_{\mathbf{w}}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)}) = 2\ell_t(\mathbf{w}),$$

where the second inequality follows from Lemma 18. Using the above and the learning rate schedule $\eta_t = 1/(2(\lambda + \|\nabla \ell_t(\mathbf{w}_t)\|_2))$ we can invoke the second case of Proposition 13 to obtain

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{\ell_t(\mathbf{w}_t)}{1 + 2\lambda^{-1}\ell_t(\mathbf{w}_t)} \leq \frac{2}{T} \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{2\lambda\|\mathbf{v} - \mathbf{w}_0\|_2^2}{T} \quad (\text{B.2})$$

for any $\mathbf{v} \in \mathbb{R}^D$. Further, we use the following definitions

$$\begin{aligned}\ell_t^i(\mathbf{w}) &:= -\sum_{j=1}^i A_j^{(t)} \log p_{\mathbf{w}}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)}), \\ \tilde{\ell}_t^i(\mathbf{w}) &:= -\log p_{\mathbf{w}}(\mathbf{y}_{1:i}^{*(t)} | \mathbf{x}^{(t)}), \\ \tilde{\ell}_t(\mathbf{w}) &:= -\log p_{\mathbf{w}}(\mathbf{y}^{*(t)} | \mathbf{x}^{(t)}).\end{aligned}$$

Since $\ell_t^i \leq \ell_t$ and $\tilde{\ell}_t^i \leq \tilde{\ell}_t$, we have

$$\begin{aligned}\frac{1}{T} \sum_{t=0}^{T-1} \frac{\ell_t^i(\mathbf{w}_t)}{1 + 2\lambda^{-1}\tilde{\ell}_t^i(\mathbf{w}_t)} &\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{\ell_t^i(\mathbf{w}_t)}{1 + 2\lambda^{-1}\tilde{\ell}_t^i(\mathbf{w}_t)} \\ &\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{\ell_t(\mathbf{w}_t)}{1 + 2\lambda^{-1}\ell_t(\mathbf{w}_t)} \\ &\leq \frac{2}{T} \sum_{t=0}^{T-1} \ell_t(\mathbf{v}) + \frac{2\lambda\|\mathbf{v} - \mathbf{w}_0\|_2^2}{T},\end{aligned} \quad (\text{B.3})$$

where the last inequality follows from (B.2). Next, we take expectation over the randomness in the draw of $\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}$. Note that $\mathbf{w}_t$ is independent of $\mathbf{y}^{(t)}$, and $\mathbb{E}[A_j^{(t)} | \mathbf{w}_t] = q_t(\mathbf{y}_{1:j}^{(t)} | \mathbf{x}^{(t)})$ for all $j \in [N]$. Therefore,

$$\begin{aligned}\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} [\ell_t^i(\mathbf{w}_t)] &= \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\sum_{j=1}^i -q_t(\mathbf{y}_{1:j}^{*(t)} | \mathbf{x}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}_j^{*(t)} | \mathbf{x}^{(t)}, \mathbf{y}_{1:j-1}^{*(t)}) \right] \\ &\geq \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[-q_t(\mathbf{y}_{1:i}^{*(t)} | \mathbf{x}^{(t)}) \log p_{\mathbf{w}_t}(\mathbf{y}_{1:i}^{*(t)} | \mathbf{x}^{(t)}) \right] \\ &= \mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[q_t(\mathbf{y}_{1:i}^{*(t)} | \mathbf{x}^{(t)}) \tilde{\ell}_t^i(\mathbf{w}_t) \right].\end{aligned} \quad (\text{B.4})$$

By combining (B.3), (B.4), and the fact that $\mathbb{E}[\ell_t(\mathbf{v})] \leq \tilde{\ell}_t(\mathbf{v})$, we obtain

$$\mathbb{E}_{\mathbf{y}^{(0)}, \dots, \mathbf{y}^{(T-1)}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \frac{q_t(\mathbf{y}_{1:i}^{*(t)} | \mathbf{x}^{(t)}) \tilde{\ell}_t^i(\mathbf{w}_t)}{1 + 2\lambda^{-1}\tilde{\ell}_t^i(\mathbf{w}_t)} \right] \leq \frac{2}{T} \sum_{t=0}^{T-1} \tilde{\ell}_t(\mathbf{v}) + \frac{2\lambda\|\mathbf{v} - \mathbf{w}_0\|_2^2}{T}.$$

The remainder of the proof follows exactly as in that of Proposition 19. Namely, we let $\mathbf{v} = R_v \mathbf{w}^*$ for which we have $\tilde{\ell}_t(\mathbf{v}) \leq N k e^{-R_v \gamma}$. We choose $R_v = 1/\gamma \cdot \log(N k T \gamma^2 / \lambda)$ and notice that $\frac{\tilde{\ell}_t^i}{1+2\lambda^{-1}\tilde{\ell}_t^i} \geq \frac{1}{2} \min(\lambda/2, \tilde{\ell}_t^i)$. Combined with the bound on $\|\mathbf{w}_0\|_2$, the proof is complete. $\square$

Note that the above guarantee is more general than that of Proposition 19. While it can recover the same bound for $i = N$, it can cover cases where $i < N$, where it shows that (PG-PR) can boost the accuracy on the first i tokens if the behavior policy assigns a non-trivial probability to the first i ground-truth tokens. As done before, we can use an online-to-batch conversion to turn the above bound into a conditional guarantee on test error.

Proposition 26. *Let $(\mathbf{w}_t)$ denote the (PG-PR) iterates with adaptive learning rate $\eta_t = (4 + 2\|\nabla \ell_t(\mathbf{w}_t)\|_2)^{-1}$. Suppose Assumption 1 holds, and $\|\mathbf{w}_0\|_2 \leq 1/\gamma \cdot \log(N k T \gamma^2 / 2)$. For any $i \in [N]$ and any $\alpha \in (0, 1]$, define the event $\mathcal{E}_{i,\alpha} := \{\inf_t q_t(\mathbf{y}^*(\mathbf{x})_{1:i} | \mathbf{x}) \geq \alpha\}$ and $\pi_{i,\alpha} := \mathbb{P}(\mathcal{E}_{i,\alpha})$. , in expectation over the training randomness, over the test sample $\mathbf{x}$, and over the choice $\tau \sim \text{Unif}(\{0, \dots, T-1\})$, we have*

$$\mathbb{E}[p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x})_{1:i} | \mathbf{x}) | \mathcal{E}_{i,\alpha}] \geq 1 - \tilde{\mathcal{O}}\left(\frac{1}{\pi_{i,\alpha} \alpha \gamma^2 T}\right). \quad (\text{B.5})$$

Proof. Similar to the proof of Theorem 3, we take expectation from both sides of (B.1), and use the independence of $\mathbf{x}^{(t)}$ from $\mathbf{w}_t$ to replace it with an independent test sample $\mathbf{x}$, and write

$$\mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}) \min\left(\frac{\lambda}{2}, -\log p_{\mathbf{w}_t}(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x})\right)\right] \lesssim \frac{\lambda \log(N k T \gamma^2 / \lambda)^2}{\gamma^2 T}.$$

Using $1 - x \leq -\log x$, we can write

$$\mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} q_t(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}) \min\left(\frac{\lambda}{2}, 1 - p_{\mathbf{w}_t}(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x})\right)\right] \lesssim \frac{\lambda \log(N k T \gamma^2 / \lambda)^2}{\gamma^2 T}.$$

By choosing $\lambda = 2$, we can remove the minimum and write

$$\mathbb{E}[q_\tau(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}) (1 - p_{\mathbf{w}_\tau}(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}))] \lesssim \frac{\log(N k T \gamma^2)^2}{\gamma^2 T},$$

where we replaced the average with expectation over τ . Next, we can use the law of total expectation and note that on $\mathcal{E}_{i,\alpha}$ we have $q_\tau(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}) \geq \alpha$ for all t , which yields

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}_{1:i}^*(\mathbf{x}) | \mathbf{x}) | \mathcal{E}_{i,\alpha}] \mathbb{P}(\mathcal{E}_{i,\alpha}) \lesssim \frac{\log(N k T \gamma^2)^2}{\gamma^2 T}$$

which completes the proof. $\square$

With this proposition, we can present the proof of the main theorem of this section.

Proof. [Proof of Theorem 6.] The proof is roughly similar to the argument presented in Section A.4. This time, we define

$$\mathcal{E}_\alpha := \left\{ \min_{j \in [N]} q_0(\mathbf{y}_j^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:j-1}^*(\mathbf{x})) \geq \alpha \right\},$$

$$\tilde{\mathcal{E}}_{\tilde{\alpha}} := \{q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \tilde{\alpha}\},$$

for $\alpha, \tilde{\alpha} \in (0, 1]$. Using the law of total expectation and Proposition 26, we have

$$\begin{aligned} \mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] &= \mathbb{E}\left[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}\right] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}) + \mathbb{E}\left[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \tilde{\mathcal{E}}_{\tilde{\alpha}}^c\right] \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}^c) \\ &\leq \tilde{\mathcal{O}}\left(\frac{1}{\tilde{\alpha} \gamma^2 T}\right) + \mathbb{P}(\tilde{\mathcal{E}}_{\tilde{\alpha}}^c). \end{aligned} \quad (\text{B.6})$$

Let

$$m := \lceil 2(\log N + 1) \cdot m^* \rceil, \quad m^* := \min\left(\mathcal{Q}_{q_0}^{\text{TL}}\left((1 - \frac{1}{\log(1/\varepsilon)})\varepsilon\right)^{-1}, k\right).$$

Once again, we proceed by considering each case in the minimum above separately.

- Suppose the minimum is attained by the first term. Note that

$$q(y_j^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:j-1}^*(\mathbf{x})) \geq 1 - (1 - q_0(y_j^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:j-1}^*(\mathbf{x}))/2)^m$$

for all j . Therefore, on $\mathcal{E}_\alpha$ we have $q(y_j^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:j-1}^*(\mathbf{x})) \geq 1 - e^{-\alpha m/2}$ and

$$q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - Ne^{-\alpha m/2} \geq 1 - N^{-\alpha m^* (\log N + 1)}.$$

By letting $\alpha = 1/m^*$ we obtain $q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - e^{-1}$. Thus, if we let $\tilde{\alpha} := 1 - e^{-1}$, we have

$$\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) \leq \mathbb{P}(\mathcal{E}_\alpha^c) = \mathcal{Q}_{q_0}^{-1}(1/m^*) \leq (1 - \frac{1}{\log(1/\varepsilon)})\varepsilon.$$

Going back to (B.6), we obtain

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}\left(\frac{1}{\gamma^2 T}\right) + (1 - \frac{1}{\log(1/\varepsilon)})\varepsilon.$$

Choosing $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$ is sufficient to ensure the first term remains bounded by $\varepsilon/\log(1/\varepsilon)$, and the overall expected error remains bounded by ε .

- Now suppose the minimum is attained by the second term. Then,

$$q(y_j^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:j-1}^*(\mathbf{x})) \geq 1 - (1 - 1/(2k))^m \geq 1 - e^{-m/(2k)}.$$

And $q(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq 1 - Ne^{-m/(2k)} \geq 1 - Ne^{-\log N - 1}$. Thus, for $\tilde{\alpha} = 1 - e^{-1}$ we have $\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) = 0$. Hence (B.6) reads

$$\mathbb{E}[1 - p_{\mathbf{w}_\tau}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \tilde{\mathcal{O}}\left(\frac{1}{\gamma^2 T}\right),$$

and choosing $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$ is sufficient.

The argument for the number of reward queries is also similar to that of Section A.4, adapted to Algorithm 2. Note that if $\mathbf{y}^{(t)} = \mathbf{y}^*(\mathbf{x}^{(t)})$ for $\mathbf{y}^{(t)} = p_{\mathbf{w}_t}$, then Algorithm 2 only makes one reward query, and (PG-PR) does not need any additional reward queries. Otherwise, Algorithm 2 makes at most $1 + mN$ reward queries, and (PG-PR) makes at most N reward queries. Thus, we can write the total number of reward queries as

$$\begin{aligned} Q &= T + (m+1)N \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\hat{\mathbf{y}}^{(t)} \neq \mathbf{y}^*(\mathbf{x}^{(t)})]\right] \\ &= T + (m+1)N \sum_{t=1}^T \mathbb{E}\left[1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)})\right] \\ &= T + (m+1)N \sum_{t=1}^T \mathbb{E}[1 - p_{\mathbf{w}_t}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] && \text{(Using independent of } \mathbf{x}^{(t)} \text{ from } \mathbf{w}_t) \\ &= T + (m+1)NT \mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) + (m+1)N\tilde{\mathcal{O}}(\gamma^{-2}) && \text{(Using (B.6) and } \tilde{\alpha} = 1 - e^{-1}). \end{aligned}$$

Using $\mathbb{P}(\tilde{\mathcal{E}}_\alpha^c) \leq (1 - o(1))\varepsilon$ and $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$, we have

$$Q = T + (m+1)N\tilde{\mathcal{O}}(\gamma^{-2}) = \tilde{\mathcal{O}}\left(\frac{mN + \varepsilon^{-1}}{\gamma^2}\right).$$

Plugging in m completes the proof. $\square$

B.2 Proof of Corollary 7

We first state the precise assumption on the feature map.

Assumption 2. *There exist k functions $(\psi_i)_{i=1}^k$ where $\psi_i : \mathcal{X} \rightarrow \mathbb{R}^D$, such that for all $i \in [N]$, $\phi(\mathbf{x}, \mathbf{y}_{1:i}^*(\mathbf{x})) = \psi_1(x)$, and $(\phi(\mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x}), 1), \dots, \phi(\mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x}), k))$ is a permutation of $(\psi_1(\mathbf{x}), \dots, \psi_k(\mathbf{x}))$. The permutation itself is allowed to depend on the position i .*

A consequence of the above assumption is that for all $\mathbf{w}$, $\mathbf{x}$, and $i \in [N]$,

$$p_{\mathbf{w}}(y_i^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x})) = p_{\mathbf{w}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})^{1/N}. \quad (\text{B.7})$$

We can now state the proof of the corollary.

Proof. [Proof of Corollary 7] The first case of the corollary directly follows from Proposition 1.

By going back to Proposition 17 in the SGD setting, we have $\mathbb{E}[-\log p_{\bar{\mathbf{w}}_n^{\text{SGD}}}(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})] \leq \mathcal{O}(N \log(nk\gamma^2)/(\gamma^2 n))$, where we used the Jensen inequality. Recall $q_0 = p_{\bar{\mathbf{w}}_n^{\text{SGD}}}$. Define $Z(\mathbf{x}) := \min_{i \in N} q_0(y_i^*(\mathbf{x}) | \mathbf{x}, \mathbf{y}_{1:i-1}^*(\mathbf{x}))$. By (B.7), we have $Z(\mathbf{x}) = q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})^{1/N}$. Consequently we have

$$\begin{aligned} \mathbb{P}(Z(\mathbf{x}) < \alpha) &= \mathbb{P}(q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})^{1/N} < \alpha) = \mathbb{P}\left(-\frac{1}{N} \log q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \geq \log(1/\alpha)\right) \\ &\leq \frac{\mathbb{E}[-\log q_0(\mathbf{y}^*(\mathbf{x}) | \mathbf{x})]}{N \log(1/\alpha)} \quad (\text{By Markov's Inequality}) \\ &\leq \frac{\log(nk\gamma^2)}{n\gamma^2 \log(1/\alpha)}. \end{aligned}$$

Recall that $\mathcal{Q}_{q_0}^{\text{TL}}(\varepsilon) = \sup\{\alpha \in [0, 1] : \mathbb{P}(Z \leq \alpha) \leq \varepsilon\}$. Thus we have

$$\mathcal{Q}_{q_0}^{\text{TL}}(\varepsilon) \geq \exp\left(-\frac{\log(n\gamma^2 k)}{n\gamma^2 \varepsilon}\right).$$

Recall from Theorem 6 that while the number of iterations is always given by $T = \tilde{\mathcal{O}}(1/(\gamma^2 \varepsilon))$, the number of reward queries is given by

$$Q = \tilde{\mathcal{O}}\left(\frac{N \min(\mathcal{Q}_{q_0}^{\text{TL}}(\varepsilon)^{-1}, k) + \varepsilon^{-1}}{\gamma^2}\right).$$

When $n \geq \frac{\log(k/\varepsilon)}{\gamma^2 \varepsilon}$, one can verify that $\mathcal{Q}_{q_0}^{\text{TL}}(\varepsilon) \geq e^{-3}$. Thus, we can choose this $\mathcal{Q}_{q_0}^{\text{TL}}$ in the minimum above. Otherwise, we can choose k . This finishes the proof. $\square$

C Proofs of Section 5

Our minimax lower bounds rely on the following fact.

Lemma 27. *Suppose y is a random variable drawn uniformly from $\{1, \dots, m\}$. To estimate y , an algorithm is allowed $l \leq m$ rounds of guesses, where at round t the algorithm produces a guess $\hat{y}_t$ and receives the answer $\mathbb{1}[\hat{y}_t = y]$. Let $\hat{y}_l$ denote the final guess of the algorithm. For any such algorithm we have*

$$\mathbb{P}(\hat{y}_l \neq y) \geq \frac{m-l}{m}.$$

Proof. Define $A_l = \cap_{i=1}^l \{\hat{y}_i \neq y\}$. We will prove by induction that $\mathbb{P}(A_l) \geq (m-l)/m$ for $l \leq m$. Note that $\hat{y}_1$ is independent of y , therefore $\mathbb{P}(A_1) = (m-1)/m$. For $l \geq 2$, we have

$$\mathbb{P}(A_l) = \mathbb{P}(\hat{y}_l \neq y | A_{l-1}) \mathbb{P}(A_{l-1}). \quad (\text{C.1})$$

Further,

$$\begin{aligned}
\mathbb{P}(\hat{y}_l \neq y \mid A_{l-1}) &= \mathbb{P}(\hat{y}_l \neq y \mid A_{l-1}, \hat{y}_l \in \{\hat{y}_i\}_{i=1}^{l-1}) \mathbb{P}(\hat{y}_l \in \{\hat{y}_i\}_{i=1}^{l-1}) \\
&\quad + \mathbb{P}(\hat{y}_l \neq y \mid A_{l-1}, \hat{y}_l \notin \{\hat{y}_i\}_{i=1}^{l-1}) \mathbb{P}(\hat{y}_l \notin \{\hat{y}_i\}_{i=1}^{l-1}) \\
&\geq \mathbb{P}(\hat{y}_l \neq y \mid A_{l-1}, \hat{y}_l \notin \{\hat{y}_i\}_{i=1}^{l-1}) \\
&\geq \frac{m-l}{m-(l-1)},
\end{aligned}$$

where we used the fact that $\mathbb{P}(\hat{y}_l \neq y \mid A_{l-1}, \hat{y}_l \in \{\hat{y}_i\}_{i=1}^{l-1}) = 1$. Plugging back into (C.1) and using the induction hypothesis $\mathbb{P}(A_{l-1}) \geq (m-(l-1))/m$, we obtain

$$\mathbb{P}(A_l) \geq \frac{m-l}{m}.$$

This completes the proof of the induction. The proof of the lemma is completed by noticing that $\mathbb{P}(\hat{y}_l \neq y) \geq \mathbb{P}(A_l)$. $\square$

Remark. Note that the above lower bound is sharp. An optimal algorithm draws $\hat{y}_l$ uniformly from $\{1, \dots, m\} \setminus \{\hat{y}_1, \dots, \hat{y}_{l-1}\}$.

C.1 Proof of Theorem 10

We first specify the precise log dependencies hidden by the big-O notation in the statement, namely

$$Q = \Omega \left(\frac{\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1}}{\gamma^2 \cdot \log(1/\varepsilon) \cdot \log_k \left\{ \mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1} \right\}} \right).$$

The proof proceeds by randomizing the data distribution and providing bounds in expectation with respect to this randomness, which implies the existence of a worst-case distribution with the same bound.

Setup and Definitions. Define $I := 1/\gamma^2$ and $m := 1/\alpha$. For simplicity, we assume I and m are integers, while general real numbers can be similarly handled. Define four sets $(\mathcal{X}_i)_{i=1}^4$ where $\mathcal{X}_i = \{\mathbf{x}^{i,1}, \dots, \mathbf{x}^{i,I_i}\}$, for $(I_i)_{i=1}^4$ to be determined later. We define the data distribution $\mathcal{D}$ as follows: let

$$p_1 := (1 - \varepsilon^*)(1 - \delta), \quad p_2 := (1 - \varepsilon^*)\delta, \quad p_3 := \varepsilon^*(1 - \delta), \quad p_4 := \varepsilon^*\delta$$

for some $\delta \in (0, 1)$ to be fixed later. With probability p_i , we draw $\mathbf{x}$ uniformly from $\mathcal{X}_i$. The ground-truth label $\mathbf{y}^*(\mathbf{x}^{i,r})$ of input example $\mathbf{x}^{i,r}$ is drawn independently and uniformly at random from $\mathbb{Y}_i$. Suppose for simplicity that m is a perfect power of k , the following argument can be easily extended to general m . Let $L := \log_k m$. Define

$$\begin{aligned}
\mathbb{Y}_1 &= \{(y_1, \dots, y_L, 1, \dots, 1) : y_1, \dots, y_L \in \mathcal{Y}^L\} \subseteq \mathcal{Y}^N, \\
\mathbb{Y}_2 &= \mathbb{Y}_4 = \{(1, 1, \dots, 1), (2, 1, \dots, 1)\} \subseteq \mathcal{Y}^N, \\
\mathbb{Y}_3 &= \mathcal{Y}^N
\end{aligned}$$

In words, $\mathbb{Y}_1$ is made of all sequences of length L by appending $N - L$ 1s to the end of each sequence. In particular, $|\mathbb{Y}_1| = m$.

We will use the notation $(m_i)_{i=1}^4$ where $m_i = |\mathbb{Y}_i|$. With this definition, the label for $\mathbf{x} \in \mathcal{X}_i$ is drawn uniformly from m_i sequences. Note that $m_i \geq 2$ for all i . Throughout the proof, we will use T for the number of labeled samples.

We define $q(\cdot \mid \mathbf{x})$ to be $\text{Unif}(\mathbb{Y}_1)$ for $\mathbf{x} \in \mathcal{X}_1 \cup \mathcal{X}_2$ and $\text{Unif}(\mathcal{Y}^N)$ for $\mathbf{x} \in \mathcal{X}_3 \cup \mathcal{X}_4$. Since $\mathbb{Y}_2 \subseteq \mathbb{Y}_1$, we have $q(\mathbf{y}^*(\mathbf{x}) \mid \mathbf{x}) = m^{-1}$ if $\mathbf{x} \in \mathcal{X}_1 \cup \mathcal{X}_2$ and $q(\mathbf{y}^*(\mathbf{x}) \mid \mathbf{x}) = k^{-N}$ otherwise. One can then verify that q satisfies

$$\mathcal{Q}_q(\varepsilon) = \begin{cases} k^{-N} & \text{if } \varepsilon < \varepsilon^* \\ m^{-1} & \text{if } \varepsilon \geq \varepsilon^*. \end{cases}$$

Let $L_1 = L = \log_k m$, $L_2 = L_4 = 1$ and $L_3 = N$. Define the following set of standard basis vectors of $\mathbb{R}^D$,

$$\{\mathbf{e}_{i,r,j,a} : i \in [4], r \in [I_i], j \in [L_i], a \in [k]\},$$

Add an extra vector to this orthonormal collection, denoted by $\mathbf{u}$. Note that this construction requires $D \geq 1 + k \sum_{i=1}^4 I_i L_i$.

For all $i \in [4]$ and $r \in [I_i]$, define the feature map as

$$\phi(\mathbf{x}^{i,r}, \mathbf{y}_{1:j}) := \begin{cases} \mathbf{e}_{i,r,j,y_j}, & j \leq L_i \\ \mathbf{u}, & j > L_i, y_j = 1. \\ \mathbf{0}, & j > L_i, y_j \neq 1 \end{cases}$$

Define

$$\mathbf{w}^* := \gamma \mathbf{u} + \gamma \sum_{i=1}^4 \sum_{r=1}^{I_i} \sum_{j=1}^{L_i} \mathbf{e}_{i,r,j,y_j^*(\mathbf{x}^{i,r})}.$$

It is straightforward to verify that

$$\langle \mathbf{w}^*, \phi(\mathbf{x}^{i,r}, \mathbf{y}_{1:j}^*(\mathbf{x}^{i,r})) \rangle = \gamma, \quad \langle \mathbf{w}^*, \phi(\mathbf{x}^{i,r}, \mathbf{y}_{1:j-1}^*(\mathbf{x}^{i,r}), y') \rangle = 0, \quad \forall i, r, j, y' \neq y_j^*(\mathbf{x}^{i,r}).$$

Additionally $\|\mathbf{w}^*\|_2^2 = \gamma^2(1 + \sum_{i=1}^4 I_i L_i)$. We define $I_i = (I - 1)/(4L_i)$ so that $\|\mathbf{w}^*\|_2 = 1$. Therefore, $\mathbf{w}^*$ satisfies Assumption 1. Once again, we assume I_i are integers (which is W.L.O.G.) and assume $I_i \geq 1$. The latter has the implication that $\gamma^2 \leq 1/N$. This choice of I_i also gives the sufficient condition on D , namely $D \geq k/\gamma^2$.

Lower Bound on Number of Samples. Let $m(\mathbf{x})$ denote the number of reward queries made on context $\mathbf{x}$. Let $\mathbf{x} \sim \mathcal{D}$ independent of S . We have

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) = \sum_{i=1}^4 \mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) \mathbb{P}(\mathbf{x} \in \mathcal{X}_i). \quad (\text{C.2})$$

Moreover,

$$\begin{aligned} \mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) &= \mathbb{E}[\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid S, \mathbf{x} \in \mathcal{X}_i) \mid \mathbf{x} \in \mathcal{X}_i] \\ &\geq \mathbb{E}\left[\mathbb{1}[\mathbf{x} \notin S] \cdot \frac{m_i - 1}{m_i} + \mathbb{1}[\mathbf{x} \in S] \cdot \frac{m_i - 1 - m(\mathbf{x})}{m_i} \vee 0 \mid \mathbf{x} \in \mathcal{X}_i\right] \\ &= \mathbb{P}(\mathbf{x} \notin S \mid \mathbf{x} \in \mathcal{X}_i) \cdot \frac{m_i - 1}{m_i} + \mathbb{E}\left[\frac{\mathbb{1}[\mathbf{x} \in S](m_i - 1 - m(\mathbf{x}))}{m_i} \vee 0 \mid \mathbf{x} \in \mathcal{X}_i\right]. \end{aligned} \quad (\text{C.3})$$

To obtain a lower bound on the number of samples, we remove the second term from (C.3) to obtain

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) \geq \frac{m_i - 1}{m_i} \mathbb{P}(\mathbf{x} \notin S \mid \mathbf{x} \in \mathcal{X}_i)$$

where the second inequality follows from Lemma 27. Note that if $\mathbf{x} \in \mathcal{X}_i$, then the probability that it is drawn once is $\frac{p_i}{I_i}$, and by a union bound, the probability that it is drawn at least once is at most $\frac{T p_i}{I_i}$. Therefore $\mathbb{P}(\mathbf{x} \notin S \mid \mathbf{x} \in \mathcal{X}_i) \geq (1 - \frac{p_i T}{I_i}) \vee 0$. Thus

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq \sum_{i=1}^4 p_i \cdot \frac{m_i - 1}{m_i} \cdot \left(1 - \frac{p_i T}{I_i}\right) \vee 0.$$

Since each term is non-negative, we can remove $i = 1$ and $i = 3$ terms, from the lower bound, and write

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq \frac{p_2}{2} \left(1 - \frac{p_2 T}{I_2}\right) + \frac{p_4}{2} \left(1 - \frac{p_4 T}{I_4}\right) \geq \frac{p_2 + p_4}{2} \left(1 - \frac{(p_2 \vee p_4) T}{I_2}\right),$$

where we used $m_2 = m_4 = 2$ and $I_2 = I_4$. For the probability of error to be less than ε we need

$$T \geq \frac{I_2}{p_2 \vee p_4} \left(1 - \frac{2\varepsilon}{p_2 + p_4}\right).$$

Recall $p_2 = (1 - \varepsilon^*)\delta$ and $p_4 = \varepsilon^*\delta$. Assume $\varepsilon < 1/4$ and let $\delta = 4\varepsilon$. Then, the above implies

$$T \geq \frac{I_2}{2\delta((1 - \varepsilon^*) \vee \varepsilon^*)} \geq \frac{I_2}{8\varepsilon} \gtrsim \frac{1}{\gamma^2\varepsilon},$$

which completes the proof of the sample lower bound.

Lower Bound on Number of Reward Queries. We use the same $\delta = 4\varepsilon$ as before. To prove the lower bound on the number of reward queries, we go back to (C.3), and note that

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) \geq \frac{m_i - 1}{m_i} - \mathbb{E}\left[\frac{\mathbb{1}[\mathbf{x} \in S]m(\mathbf{x})}{m_i} \mid \mathbf{x} \in \mathcal{X}_i\right].$$

Define $Z := \mathbb{E}[\mathbb{1}[\mathbf{x} \in S]m(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_1]$ and $Z' := \mathbb{E}[\mathbb{1}[\mathbf{x} \in S]m(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_3]$. Using (C.2), we obtain

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq (1 - \varepsilon^*)(1 - \delta) \left(\frac{m - 1 - Z}{m}\right) \vee 0 + \varepsilon^*(1 - \delta) \left(\frac{k^N - 1 - Z'}{k^N}\right) \vee 0, \quad (\text{C.4})$$

where we used the non-negativeness of each conditional probability. In particular, each term on the RHS must be bounded by the LHS, i.e.

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq (1 - \varepsilon^*)(1 - \delta) \left(\frac{m - 1 - Z}{m}\right) \quad (\text{C.5})$$

when considering only $\mathbf{x} \in \mathcal{X}_1 \cup \mathcal{X}_2$, and

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq \varepsilon^*(1 - \delta) \left(\frac{k^N - 1 - Z'}{k^N}\right) \quad (\text{C.6})$$

when considering only $\mathbf{x} \in \mathcal{X}_3$. Before proceeding, we observe that the number of reward queries is equal to $Q \geq I_1 Z + I_3 Z'$. This is because

$$\begin{aligned} Q &= \mathbb{E}\left[\sum_{i=1}^I \mathbb{1}[\mathbf{x}^i \in S]m(\mathbf{x}^i)\right] \\ &\geq \mathbb{E}\left[\sum_{\mathbf{x}^i \in \mathcal{X}_1} \mathbb{1}[\mathbf{x}^i \in S]m(\mathbf{x}^i)\right] + \mathbb{E}\left[\sum_{\mathbf{x}^i \in \mathcal{X}_3} \mathbb{1}[\mathbf{x}^i \in S]m(\mathbf{x}^i)\right] \\ &= I_1 \mathbb{E}[\mathbb{1}[\mathbf{x} \in S]m(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_1] + I_3 \mathbb{E}[\mathbb{1}[\mathbf{x} \in S]m(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_3], \end{aligned}$$

where we recall $\mathbf{x} \sim \mathcal{D}$ in the above.

Now, consider the following cases:

- First, consider the case where $(1 + 1/\log(1/\varepsilon))\varepsilon \geq \varepsilon^*$. We have

$$\varepsilon \geq (1 - \varepsilon^*)(1 - \delta) \left(\frac{m - 1 - Z}{m}\right)$$

which implies $Z \geq m(1 - \varepsilon/((1 - \varepsilon^*)(1 - \delta)) - m^{-1}) \gtrsim m$ where we used $\varepsilon/((1 - \varepsilon^*)(1 - \delta)) < 1/2$ for sufficiently small ε , and $m^{-1} < 1/2$. Therefore, the number of reward queries is $Q \geq I_1 Z \gtrsim mI/L_1 \gtrsim \alpha^{-1}/(\gamma^2 \log_k \alpha^{-1})$. Note that for the range of ε here we have $\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon) = \alpha$, therefore

$$Q \gtrsim \frac{\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1}}{\gamma^2 \log_k \{\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1}\}}$$

- Second, assume $(1 + 1/\log(1/\varepsilon))\varepsilon < \varepsilon^*$. Using (C.6), we have

$$\varepsilon \geq (1 - \delta)\varepsilon^* \left(\frac{k^N - 1 - Z'}{k^N} \right)$$

and consequently

$$\begin{aligned} Z' &\geq k^N \left(1 - k^{-N} - \frac{\varepsilon}{(1 - \delta)\varepsilon^*} \right) \geq k^N \left(1 - k^{-N} - \frac{1}{(1 - \delta)(1 + 1/\log(1/\varepsilon))} \right) \\ &= k^N \left(\frac{1}{(1 - \delta)(1 + \log(1/\varepsilon))} - \frac{\delta}{1 - \delta} - k^{-N} \right) \end{aligned}$$

Whenever $k^N \gg \log(1/\varepsilon)$, the above implies $Z' \gtrsim k^N / \log(1/\varepsilon)$, hence $Q \geq I_3 Z' \gtrsim k^N / (N\gamma^2 \log(1/\varepsilon))$. When $k^N \ll \log(1/\varepsilon)$, we can instead apply the previous bound which resulted in $Q \gtrsim \alpha^{-1} / (\gamma^2 \log_k \alpha^{-1}) \gtrsim 1/\gamma^2$. Combined with $k^N \ll \log(1/\varepsilon)$ we have $Q \gtrsim 1/\gamma^2 \gtrsim k^N / (N\gamma^2 \log(1/\varepsilon))$. Using the fact that $Q_q((1 + 1/\log(1/\varepsilon))\varepsilon) = k^{-N}$ completes the proof. $\square$

C.2 Proof of Theorem 11

We use the same setup as that of Section C.1 with the exception of having

$$\mathbb{Y}_1 = \mathbb{Y}_2 = \{(y_1, \dots, y_L, 1, \dots, 1) : y_1, \dots, y_L \in \mathcal{Y}^L\} \subseteq \mathcal{Y}^N, \quad \mathbb{Y}_3 = \mathbb{Y}_4 = \mathcal{Y}^N.$$

Note that this also implies $L_1 = L_2 = \log_k m$ and $L_3 = L_4 = N$. We use S to denote the previous samples as well the learner's prediction on them and the obtained reward. We begin with the same decomposition as in (C.2). This time, for each conditional probability, we write

$$\begin{aligned} \mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) &= \mathbb{E}[\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid S, \mathbf{x} \in \mathcal{X}_i) \mid \mathbf{x} \in \mathcal{X}_i] \\ &\geq \mathbb{E} \left[\frac{m_i - 1 - m(\mathbf{x})}{m_i} \mid \mathbf{x} \in \mathcal{X}_i \right] \vee 0, \end{aligned}$$

where $m(\mathbf{x})$ is the number of times $\mathbf{x}$ has appeared in S , and we used Lemma 27 for the inequality. Moreover, when $\mathbf{x} \in \mathcal{X}_i$, the probability that an individual draw in S is equal to $\mathbf{x}$ is p_i/I_i , hence $\mathbb{E}[m(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i] = \frac{Tp_i}{I_i}$ where we recall T is the number of rounds (hence samples). Therefore

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}_i) \geq \left(\frac{(m_i - 1)I_i - Tp_i}{m_i I_i} \right) \vee 0. \quad (\text{C.7})$$

Lower Bound on the Number of Samples. We consider the following two cases:

- Consider the case $(1 + 1/\log(1/\varepsilon))\varepsilon \geq \varepsilon^*$. Using (C.7) and (C.2), we have

$$\mathbb{P}(\hat{\mathbf{y}}(S)(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})) \geq \sum_{i=1}^4 p_i \cdot \left(\frac{(m_i - 1)I_i - Tp_i}{m_i I_i} \right) \vee 0. \quad (\text{C.8})$$

$$\geq p_2 \left(\frac{(m - 1)I_2 - Tp_2}{mI_2} \right) + p_4 \left(\frac{(k^N - 1)I_4 - Tp_4}{k^N I_4} \right) \quad (\text{C.9})$$

$$\geq (p_2 + p_4) \left(\frac{(m - 1)I_2 - T(p_2 \vee p_4)}{mI_2} \right) \quad (\text{C.10})$$

where we used the fact that $k^N \geq m$. Recall $p_2 = (1 - \varepsilon^*)\delta$, and $p_4 = \varepsilon^*\delta$. Then we can rearrange the above terms to write

$$T \geq \frac{mI_2}{\delta((1 - \varepsilon^*) \vee \varepsilon^*)} \left(1 - m^{-1} - \frac{\varepsilon}{\delta} \right) \geq \frac{mI_2}{\delta} \left(1 - m^{-1} - \frac{\varepsilon}{\delta} \right).$$

By choosing $\delta = 4\varepsilon < 1$ and using the fact that $m \geq 2$, we get $T \gtrsim mI_2/\varepsilon \gtrsim m/(\gamma^2\varepsilon \log_k m)$. Note that for $(1 + 1/\log(1/\varepsilon))\varepsilon \geq \varepsilon^*$ we have $\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon) = m^{-1}$. Hence we can write

$$T \gtrsim \frac{\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1}}{\gamma^2\varepsilon \log_k \{\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon)^{-1}\}},$$

which completes the proof of this case.

- For the case where $(1 + 1/\log(1/\varepsilon))\varepsilon < \varepsilon^*$, we use (C.8) with $i = 4$, where we recall $m_4 = k^N$ and $p_4 = \varepsilon^*\delta$. We then have

$$\varepsilon \geq \varepsilon^*\delta \left(\frac{(k^N - 1)I - 4T\delta\varepsilon^*}{k^N I} \right)$$

and consequently

$$T \geq \frac{k^N I_4}{\varepsilon^*\delta} \left(1 - k^{-N} - \frac{\varepsilon}{\varepsilon^*\delta} \right).$$

Choose $\delta = \frac{\varepsilon}{\varepsilon^*}(1 + 1/(2\log(1/\varepsilon))) < 1$. Then,

$$T \gtrsim \frac{k^N I_4}{\varepsilon} \left(\frac{1}{1 + 2\log(1/\varepsilon)} - k^{-N} \right).$$

The argument is similar to that of Section C.1. If $k^N \gg \log(1/\varepsilon)$, we have $T \gtrsim \frac{k^N}{\gamma^2 N \varepsilon \log(1/\varepsilon)}$. If $k^N \ll \log(1/\varepsilon)$, we can apply the previous bound $T \gtrsim m/(\gamma^2\varepsilon \log_k m) \gtrsim k^N/(\gamma^2\varepsilon \log(1/\varepsilon)N)$. We conclude the proof of this case by recalling that for $(1 + 1/\log(1/\varepsilon))\varepsilon < \varepsilon^*$ we have $\mathcal{Q}_q((1 + 1/\log(1/\varepsilon))\varepsilon) = k^{-N}$.

Lower Bound on the Number of Mistakes. For the lower bound on the number of mistakes, we choose $\delta = 1/2$ and unify $\mathcal{X}_1$ with $\mathcal{X}_2$ and $\mathcal{X}_3$ with $\mathcal{X}_4$. For sample t , define S_t to be the collection of previous samples as well as their corresponding predictions and rewards. Then,

$$\begin{aligned} \mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)})) &= \mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)} \in \mathcal{X}_1 \cup \mathcal{X}_2)(1 - \varepsilon^*) \\ &\quad + \mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)} \in \mathcal{X}_3 \cup \mathcal{X}_4)\varepsilon^*. \end{aligned}$$

Using Lemma 27 and arguments similar to the above paragraph, we have

$$\mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)} \in \mathcal{X}_1 \cup \mathcal{X}_2) \geq \frac{(m-1)I_2 - 2(t-1)(1-\varepsilon^*)}{mI_2} \vee 0,$$

and

$$\mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)}) | \mathbf{x}^{(t)} \in \mathcal{X}_3 \cup \mathcal{X}_4) \geq \frac{(k^N - 1)I_4 - 2(t-1)\varepsilon^*}{k^N I_4} \vee 0. \quad (\text{C.11})$$

Therefore,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)})] \right] &= \sum_{t=1}^T \mathbb{P}(\hat{\mathbf{y}}(S_t)(\mathbf{x}^{(t)}) \neq \mathbf{y}^*(\mathbf{x}^{(t)})) \\ &\geq \sum_{t=1}^T (1 - \varepsilon^*) \cdot \frac{(m-1)I_2 - 2(t-1)(1-\varepsilon^*)}{mI_2} \vee 0 + \sum_{t=1}^T \varepsilon^* \cdot \frac{(k^N - 1)I_4 - 2(t-1)\varepsilon^*}{k^N I_4} \vee 0 \end{aligned}$$

Note that for all $t \lesssim k^N I_4 / \varepsilon^*$, the terms inside the second sum above are $\Omega(\varepsilon^*)$. Therefore, the sum above for $T \gtrsim k^N I_4 / \varepsilon^* \gtrsim (k^N / (N\gamma^2\varepsilon^*))$ is at least of order $k^N I_4 \gtrsim k^N / (N\gamma^2)$, completing the proof. $\square$

C.3 Proof of Theorem 12

We first describe the setup and definitions used in this proof, which are somewhat different from those in the previous two theorems. Define $\mathcal{X}_1 = \{\mathbf{x}^1, \dots, \mathbf{x}^{I_0}\}$ and $\mathcal{X}_2 = \{\mathbf{x}^{I_0+1}, \dots, \mathbf{x}^{2I_0}\}$. A sample $\mathbf{x}$ is generated uniformly from $\mathcal{X}_1$ with probability p_1 and uniformly from $\mathcal{X}_2$ with probability p_2 , where I_0, p_1 , and p_2 are determined later. Similar to the previous proofs, we prove the existence result by randomizing the labels. For every $i \in [2I_0]$, sample i.i.d. random variables $B_i \sim \text{Unif}(\{1, 2\})$. Let $\mathbf{y}^*(\mathbf{x}^i) = (B_i, \dots, B_i)$. Let $\{\mathbf{e}_{i,1}\}_{i=1}^{2I_0} \cup \{\mathbf{e}_{i,2}\}_{i=1}^{2I_0}$ be an orthonormal set of vectors in $\mathbb{R}^D$. The feature map only depends on the last token and is given by

$$\phi(\mathbf{x}^i, \mathbf{y}_{1:j}) = \begin{cases} \mathbf{e}_{i,y_j}, & y_j \in \{1, 2\} \\ \frac{\mathbf{e}_{i,1} + \mathbf{e}_{i,2}}{2}, & y_j \notin \{1, 2\} \end{cases}. \quad (\text{C.12})$$

Let $\mathbf{w}^* = \gamma \sum_{i=1}^{2I_0} (\mathbf{e}_{i,B_i} - \mathbf{e}_{i,3-B_i})$. It is straightforward to verify $\langle \mathbf{w}^*, \phi(\mathbf{x}^i, \mathbf{y}_{1:j}^*(\mathbf{x}^i)) \rangle = \gamma$ for all i, j and $\langle \mathbf{w}^*, \phi(\mathbf{x}^i, \mathbf{y}_{1:j-1}^*(\mathbf{x}^i), y_j) \rangle \leq 0$ for all $y_j \neq y_j^*(\mathbf{x}^i)$. Further, $\|\mathbf{w}^*\|_2^2 = 4I_0\gamma^2$. By choosing $I_0 = 1/(4\gamma^2)$ (we assume I_0 is an integer W.L.O.G.) we have $\|\mathbf{w}^*\|_2^2 = 1$, therefore Assumption 1 is satisfied. Note that this construction requires $D \geq 4I_0k = k/\gamma^2$.

The above feature map implies the following property.

Lemma 28. *Let $\mathbf{1}_N := (1, \dots, 1)$ and $\mathbf{2}_N := (2, \dots, 2)$. Let the feature map be given by (C.12), and assume $\mathbf{y}^*(\mathbf{x}) = \mathbf{1}_N$ or $\mathbf{y}^*(\mathbf{x}) = \mathbf{2}_N$. Then, for any $\mathbf{x}$ and any $\mathbf{w} \in \mathbb{R}^D$, we have*

$$\min(p_{\mathbf{w}}(\mathbf{1}_N | \mathbf{x}), p_{\mathbf{w}}(\mathbf{2}_N | \mathbf{x})) \leq k^{-N}.$$

Proof. Let $a := \langle \mathbf{w}, \mathbf{e}_{i,1} \rangle$ and $b := \langle \mathbf{w}, \mathbf{e}_{i,2} \rangle$. We have $p_{\mathbf{w}}(\mathbf{1}_N | \mathbf{x}) = p_{\mathbf{w}}(1 | \mathbf{x})^N$ and $p_{\mathbf{w}}(\mathbf{2}_N | \mathbf{x}) = p_{\mathbf{w}}(2 | \mathbf{x})^N$. Further,

$$p_{\mathbf{w}}(1 | \mathbf{x}) = \frac{e^a}{e^a + e^b + (k-2)e^{(a+b)/2}} \leq \frac{e^a}{ke^{(a+b)/2}},$$

where we used the AM-GM inequality, and similarly $p_{\mathbf{w}}(2 | \mathbf{x}) \leq e^b / (ke^{(a+b)/2})$. Therefore

$$\min(p_{\mathbf{w}}(\mathbf{1}_N | \mathbf{x}), p_{\mathbf{w}}(\mathbf{2}_N | \mathbf{x}))^2 \leq p_{\mathbf{w}}(\mathbf{1}_N | \mathbf{x})p_{\mathbf{w}}(\mathbf{2}_N | \mathbf{x}) = (p_{\mathbf{w}}(1 | \mathbf{x})p_{\mathbf{w}}(2 | \mathbf{x}))^N \leq k^{-2N}.$$

□

Given the definitions and the lemma, we are ready to prove the theorem.

Proof. [Proof of Theorem 12] Let $p_1 := 1 - \delta$ and $p_2 = \delta$. Let S denote the training set, and $q(S)(\cdot | \mathbf{x})$ be the output of the algorithm. Consider the decomposition

$$\mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N}) = \sum_{i=1}^2 p_i \mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) | \mathbf{x} \in \mathcal{X}_i).$$

We choose $i = 2$, and use the non-negativity of probability to write

$$\begin{aligned} \mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N}) &\geq p_2 \mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N} | \mathbf{x} \in \mathcal{X}_2) \\ &\geq p_2 \mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N} | \mathbf{x} \in \mathcal{X}_2, \mathbf{x} \notin S) \mathbb{P}(\mathbf{x} \notin S | \mathbf{x} \in \mathcal{X}_2). \end{aligned} \quad (\text{C.13})$$

By a union bound, we have $\mathbb{P}(\mathbf{x} \notin S | \mathbf{x} \in \mathcal{X}_2) \geq 1 - \frac{p_2 n}{I_0}$. Further, note that conditioned on $\mathbf{x} \notin S$, $\mathbf{y}^*(\mathbf{x})$ is independent from $q(S)(\cdot | \mathbf{x})$, and by Lemma 28, we know $\min(q(S)(\mathbf{1}_N | \mathbf{x}), q(S)(\mathbf{2}_N | \mathbf{x})) \leq k^{-N}$. As a result

$$\mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N} | \mathbf{x} \in \mathcal{X}_2, \mathbf{x} \notin S) \geq \frac{1}{2}.$$

Therefore,

$$\mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) | \mathbf{x}) \leq k^{-N}) \geq \frac{p_2}{2} \left(1 - \frac{p_2 n}{I_0}\right)$$

$$= \frac{\delta}{2} \left(1 - \frac{\delta n}{I_0} \right).$$

By choosing $\delta = 1/2 \cdot \min(1, I_0/n)$, we have

$$\mathbb{P}(q(S)(\mathbf{y}^*(\mathbf{x}) \mid \mathbf{x}) \leq k^{-N}) \geq \frac{1}{8} \cdot \min\left(1, \frac{I_0}{n}\right) = \frac{1}{8} \cdot \min\left(1, \frac{1}{4\gamma^2 n}\right).$$

Recall that by definition of LQ, we have $\mathcal{Q}_q(\varepsilon) \leq k^{-N}$ for all $\varepsilon \leq \min(1, 1/(4\gamma^2 n))/8$ which completes the proof.

□