
Transformers as Decision Makers: Provable In-Context Reinforcement Learning via Supervised Pretraining

Licong Lin
University of California, Berkeley
liconglin@berkeley.edu

Yu Bai
Salesforce AI Research
yu.bai@salesforce.com

Song Mei
University of California, Berkeley
songmei@berkeley.edu

Abstract

Large transformer models pretrained on offline reinforcement learning datasets have demonstrated remarkable in-context reinforcement learning (ICRL) capabilities, where they can make good decisions when prompted with interaction trajectories from unseen environments. However, when and how transformers can be trained to perform ICRL have not been theoretically well-understood. In particular, it is unclear which reinforcement-learning algorithms transformers can perform in context, and how distribution mismatch in offline training data affects the learned algorithms. This paper provides a theoretical framework that analyzes supervised pretraining for ICRL. This includes two recently proposed training methods — algorithm distillation and decision-pretrained transformers. First, assuming model realizability, we prove the supervised-pretrained transformer will imitate the conditional expectation of the expert algorithm given the observed trajectory. The generalization error will scale with model capacity and a distribution divergence factor between the expert and offline algorithms. Second, we show transformers with ReLU attention can efficiently approximate near-optimal online reinforcement learning algorithms like LinUCB and Thompson sampling for stochastic linear bandits, and UCB-VI for tabular Markov decision processes. This provides the first quantitative analysis of the ICRL capabilities of transformers pretrained from offline trajectories.

1 Introduction

The transformer architecture [66] for sequence modeling has become a key weapon for modern artificial intelligence, achieving success in language [14, 11, 44] and vision [16]. Motivated by these advances, the research community has actively explored how to best harness transformers for reinforcement learning (RL) [12, 29, 34, 54, 30, 33, 75]. While promising empirical results have been demonstrated, the theoretical understanding of transformers for RL remains limited.

This paper provides theoretical insights into in-context reinforcement learning (ICRL)—an emerging approach that utilizes sequence-to-sequence models like transformers to perform reinforcement learning in newly encountered environments. In ICRL, the model takes as input the current state and past interaction history with the environment (the *context*), and outputs an action. The key hypothesis in ICRL is that pretrained transformers can act as *RL algorithms*, progressively improving their policy based on past observations. Approaches such as Algorithm Distillation [30] and Decision-Pretrained Transformers [33] have demonstrated early successes, finding that *supervised pretraining* can produce

good ICRL performance. However, many concrete theoretical questions remain open about the ICRL capabilities of transformers, including but not limited to (1) what RL algorithms can transformers implement in-context; (2) what performance guarantees (e.g. regret bounds) can such transformers achieve when used iteratively as an online RL algorithm; and (3) when can supervised pretraining find such a good transformer. Specifically, this paper investigates the following open question:

How can supervised pretraining on Transformers learn in-context reinforcement learning?

In this paper, we initiate a theoretical study of the ICRL capability of transformers under supervised pretraining to address the open questions outlined above. We show that (1) Transformers can implement prevalent RL algorithms, including LinUCB and Thompson sampling for stochastic linear bandits, and UCB-VI for tabular Markov decision processes; (2) The algorithms learned by transformers achieve near-optimal regret bounds in their respective settings; (3) Supervised pretraining find such algorithms as long as the sample size scales with the covering number of transformer class and distribution ratio between expert and offline algorithms.

Summary of contributions and paper outline

- We propose a general framework for supervised pretraining approaches to ICRL (Section 2). This framework encompasses existing methods like Algorithm Distillation [30], where the expert and context algorithms are identical, as well as Decision-Pretrained Transformers [33], where the expert generates optimal actions for the MDP. It also includes approximate DPT variants where the expert estimates optimal actions from full interaction trajectories.
- We prove that the supervised-pretrained transformer will imitate the conditional expectation of the expert algorithm given the observed trajectory (Section 3 & Appendix C). The generalization error scales with both model capacity and a distribution ratio measuring divergence between the expert algorithm and the algorithm that generated offline trajectories.
- We demonstrate that transformers can effectively approximate several near-optimal reinforcement learning algorithms by taking observed trajectories as context inputs (Section 4 & Appendix E). Specifically, we show transformers can approximate LinUCB (Appendix E.1) and Thompson sampling algorithms (Appendix E.2) for stochastic linear bandit problems, and UCB-VI (Appendix E.3) for tabular Markov decision processes. Combined with generalization error bounds for pretraining and regret bounds of these RL algorithms, this provides regret bounds for supervised-pretrained transformers.
- Preliminary experiments validate that transformers can perform ICRL in our setup (Appendix I).
- Technically, we prove efficient approximation of LinUCB by showing transformers can implement accelerated gradient descent for solving ridge regression (Appendix F.4), enabling fewer attention layers than the vanilla gradient descent approach in [6]. To enable efficient Thompson sampling implementation, we prove transformers can compute matrix square roots through the Pade decomposition (Appendix G.3). These approximation results are interesting in their own right.

Related work Our work is intimately related to the lines of work on meta-reinforcement learning, in-context learning, transformers for decision-making, and the approximation theory of transformers. Due to limited space, we discuss these related works in Appendix A.

2 Framework for In-Context Reinforcement Learning

Let \mathcal{M} be the space of decision-making environments, where each environment $M \in \mathcal{M}$ shares the same number of rounds T and state-action-reward spaces $\{\mathcal{S}_t, \mathcal{A}_t, \mathcal{R}_t\}_{t \in [T]}$. Each $M = \{\mathbb{T}_M^{t-1}, \mathbb{R}_M^t\}_{t \in [T]}$ has its own transition model $\mathbb{T}_M^t : \mathcal{S}_t \times \mathcal{A}_t \rightarrow \Delta(\mathcal{S}_{t+1})$ (with $\mathcal{S}_0, \mathcal{A}_0 = \{\emptyset\}$ so $\mathbb{T}_M^0(\cdot) \in \Delta(\mathcal{S}_1)$ gives the initial state distribution) and reward functions $\mathbb{R}_M^t : \mathcal{S}_t \times \mathcal{A}_t \rightarrow \Delta(\mathcal{R}_t)$. We equip \mathcal{M} with a distribution $\Lambda \in \Delta(\mathcal{M})$, the environment prior. While this setting is general, we will later instantiate it on T rounds of bandits or K episodes of H -step MDPs with $T = KH$.

Distributions of offline trajectories We denote a partial interaction trajectory, consisting of observed state-action-reward tuples, by $D_t = \{(s_1, a_1, r_1), \dots, (s_t, a_t, r_t)\} \in \mathcal{T}_t = \prod_{s \leq t} (\mathcal{S}_s \times \mathcal{A}_s \times \mathcal{R}_s)$ and write $D = D_T$ for short. An algorithm Alg maps a partial trajectory $D_{t-1} \in \mathcal{T}_{t-1}$ and state $s_t \in \mathcal{S}_t$ to a distribution over the actions $\text{Alg}(\cdot | D_{t-1}, s_t) \in \Delta(\mathcal{A}_t)$. Given an environment M and algorithm Alg, the distribution over a full trajectory D_T is fully specified:

$$\mathbb{P}_M^{\text{Alg}}(D_T) = \prod_{t=1}^T \mathbb{T}_M^{t-1}(s_t | s_{t-1}, a_{t-1}) \text{Alg}(a_t | D_{t-1}, s_t) \mathbb{R}_M^t(r_t | s_t, a_t).$$

In supervised pretraining, we use a *context algorithm* Alg_0 (which we also refer to as the offline algorithm) to collect the offline trajectories D_T . For each trajectory D_T , we also assume access to expert actions $\bar{a} = (\bar{a}_t \in \mathcal{A}_t)_{t \in [T]} \sim \text{Alg}_E(\cdot | D_T, M)$, sampled from an expert algorithm $\text{Alg}_E : \mathcal{T}_T \times M \rightarrow \prod_{t \in [T]} \Delta(\mathcal{A}_t)$. This expert could omnisciently observe the full trajectory D_T and environment M to recommend actions. Let $\bar{D}_T = D_T \cup \{\bar{a}\}$ be the augmented trajectory. Then we have

$$\mathbb{P}_M^{\text{Alg}_0, \text{Alg}_E}(\bar{D}_T) = \mathbb{P}_M^{\text{Alg}_0}(D_T) \prod_{t=1}^T \text{Alg}_E^t(\bar{a}_t | D_T, M).$$

We denote $\mathbb{P}_\Lambda^{\text{Alg}_0, \text{Alg}_E}$ as the joint distribution of (M, \bar{D}_T) where $M \sim \Lambda$ and $\bar{D}_T \sim \mathbb{P}_M^{\text{Alg}_0, \text{Alg}_E}$, and $\mathbb{P}_\Lambda^{\text{Alg}_0}$ as the joint distribution of (M, D_T) where $M \sim \Lambda$ and $D_T \sim \mathbb{P}_M^{\text{Alg}_0}$.

Three special cases of expert algorithms We consider three special cases of the expert algorithm Alg_E , corresponding to three supervised pretraining setups:

- (a) *Algorithm distillation* [30]. The algorithm depends only on the partial trajectory D_{t-1} and current state s_t : $\text{Alg}_E^t(\cdot | D_T, M) = \text{Alg}_E^t(\cdot | D_{t-1}, s_t)$. For example, Alg_E could be a bandit algorithm like the Uniform Confidence Bound (UCB).
- (b) *Decision pretrained transformer (DPT)* [33]. The algorithm depends on the environment M and the current state s_t : $\text{Alg}_E^t(\cdot | D_T, M) = \text{Alg}_E^t(\cdot | s_t, M)$. For example, Alg_E could output the optimal action a_t^* in state s_t for environment M .
- (c) *Approximate DPT*. The algorithm depends on the full trajectory D_T but not the environment M : $\text{Alg}_E^t(\cdot | D_T, M) = \text{Alg}_E^t(\cdot | D_T)$. For example, Alg_E could estimate the optimal action \hat{a}_t^* from the entire trajectory D_T .

For any expert algorithm Alg_E , we define its reduced algorithm where the t -th step is $\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t) := \mathbb{E}_\Lambda^{\text{Alg}_0}[\text{Alg}_E^t(\cdot | D_T, M) | D_{t-1}, s_t]$. The expectation on the right is over $\mathbb{P}_\Lambda^{\text{Alg}_0}(D_T, M | D_{t-1}, s_t) = \Lambda(M) \cdot \mathbb{P}_M^{\text{Alg}_0}(D_T) / \mathbb{P}_M^{\text{Alg}_0}(D_{t-1}, s_t)$. Note that the reduced expert algorithm $\overline{\text{Alg}}_E$ generally depends on the context algorithm Alg_0 . However, for cases (a) and (b), $\overline{\text{Alg}}_E$ is independent of the context algorithm Alg_0 . Furthermore, in case (a), we have $\overline{\text{Alg}}_E^t = \text{Alg}_E^t$.

3 Statistical analysis of supervised pretraining

In supervised pretraining, we are given n i.i.d offline trajectories $\{D_T^i = (s_1^i, a_1^i, r_1^i, \dots, s_T^i, a_T^i, r_T^i)\}_{i=1}^n \sim_{iid} \mathbb{P}_\Lambda^{\text{Alg}_0}$ from the interaction of $M^i \sim_{iid} \Lambda$ with an offline algorithm Alg_0 . Given an expert algorithm Alg_E , we augment each trajectory D_T^i by $\{\bar{a}_t^i \sim_{iid} \text{Alg}_E(\cdot | D_{t-1}^i, s_t^i)\}_{t \in [T]}$. Supervised pretraining maximizes the log-likelihood over the algorithm class $\{\text{Alg}_\theta\}_{\theta \in \Theta}$

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \log \text{Alg}_\theta(\bar{a}_t^i | D_{t-1}^i, s_t^i). \quad (1)$$

This section discusses the statistical properties of the algorithm learned via supervised pretraining.

3.1 Main result

Our main result demonstrates that the algorithm maximizing the supervised pretraining loss will imitate $\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t) = \mathbb{E}_{M \sim \Lambda, D_T \sim \text{Alg}_0}[\text{Alg}_E^t(\cdot | D_T, M) | D_{t-1}, s_t]$, the conditional expectation of the expert algorithm Alg_E given the observed trajectory. The imitation error bound will scale with the covering number of the algorithm class and a distribution ratio factor, defined in Appendix C.

We aim to bound the performance gap between $\text{Alg}_{\hat{\theta}}$ and Alg_E in terms of expected cumulative rewards, where the expected cumulative reward is defined as

$$\mathfrak{R}_{\Lambda, \text{Alg}}(T) := \mathbb{E}_{M \sim \Lambda}[\mathfrak{R}_{M, \text{Alg}}(T)], \quad \mathfrak{R}_{M, \text{Alg}}(T) = \mathbb{E}_{D_T \sim \mathbb{P}_M^{\text{Alg}_0}}[\sum_{t=1}^T r_t].$$

An intermediate step of the result is controlling the expected Hellinger distance between two algorithms, where for distributions p, q , we have $D_H^2(p, q) = \int (\sqrt{p(x)} - \sqrt{q(x)})^2 dx$.

Theorem 1 (Performance gap between expected cumulative rewards). *Let Assumption A hold and let $\hat{\theta}$ be a solution to Eq. (1). Take $\mathcal{R} = \mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_0}$ as defined in Definition 5, and $\mathcal{N}_\Theta = \mathcal{N}_\Theta((nT)^{-2})$*

as defined in Definition 4. Then for some universal constant $c > 0$, with probability $\geq 1 - \delta$, we have

$$\mathbb{E}_{D_T \sim \mathbb{P}_\Lambda^{\text{Alg}_E}} \left[\sum_{t=1}^T D_H(\text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t), \overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t)) \right] \leq cT\sqrt{\mathcal{R}} \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right). \quad (2)$$

Further assume that $|r_t| \leq 1$ almost surely. Then with probability at least $1 - \delta$, the difference of the expected cumulative rewards between $\text{Alg}_{\hat{\theta}}$ and $\overline{\text{Alg}}_E$ satisfies

$$\left| \mathfrak{R}_{\Lambda, \text{Alg}_{\hat{\theta}}}(T) - \mathfrak{R}_{\Lambda, \overline{\text{Alg}}_E}(T) \right| \leq cT^2\sqrt{\mathcal{R}} \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right). \quad (3)$$

The proof of Theorem 1 is contained in Section D.1. We remark that when the expectation on the left-hand-side of (2) is with respect to the measure $\mathbb{P}_\Lambda^{\text{Alg}_0}$, standard MLE analysis will provide a bound without the distribution ratio factor $\mathcal{R} = \mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_0}$ in the right-hand side. The distribution ratio factor arises from the distribution shift between trajectories generated by the expert algorithm Alg_E versus the context algorithm Alg_0 . In addition, it should be noted that the result in Theorem 1 holds generally provided Assumption A is satisfied, which does not necessitate that the algorithm class be induced by transformers.

Implications for special cases Theorem 1 directly implies statistical guarantees for Algorithm Distillation [30], Decision-Pretrained Transformers [33], and its approximate version (Appendix C.1).

4 Approximation by transformers

As another main contribution of this paper, we demonstrate the capability of transformers to implement various near-optimal reinforcement learning algorithms in bandits and MDPs. Due to the space limit, the formal settings and theorem statements are deferred to Appendix E.

Theorem 2 (Approximating the soft LinUCB; Informal version of Theorem 12). *For a soft version of the LinUCB algorithm [13] for linear bandits, $\text{Alg}_{\text{sLinUCB}(\tau)}$, there exists a transformer $\text{TF}_\theta^{\text{R}}(\cdot)$ with $D \leq O(dA)$, $L = \tilde{O}(\sqrt{T})$, $M \leq 4A$, $D' = \tilde{O}(d + A\sqrt{Td}/(\tau\varepsilon))$, $\|\theta\| = \tilde{O}(A + T\sqrt{d}/(\tau\varepsilon^{1/4}))$,*

such that taking Alg_θ as defined in Eq. (7), we have

$$\left| \log \text{Alg}_{\text{sLinUCB}(\tau)}(\mathbf{a}_{t,k} | D_{t-1}, s_t) - \log \text{Alg}_\theta(\mathbf{a}_{t,k} | D_{t-1}, s_t) \right| \leq \varepsilon, \quad \forall t \in [T], k \in [A].$$

Theorem 3 (Regret of LinUCB and ICRL; Informal version of Theorem 13). *With probability $\geq 1 - \delta$, the learned algorithm $\text{Alg}_{\hat{\theta}}$, a solution to Eq. (1), achieves regret bound*

$$\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\hat{\theta}}}(T) \right] \leq \mathcal{O} \left(d\sqrt{T} \log(T) + T^2 \sqrt{\frac{\log(\mathcal{N}_\Theta \cdot T/\delta)}{n}} \right),$$

where $\log \mathcal{N}_\Theta \leq \tilde{O}(L^2 D(MD + D') \log n) \leq \tilde{O}(T^{3.5} d^2 A^3 \log n)$.

Besides LinUCB, we also provide similar transformer approximation results and regret bounds for the following algorithms, where near-optimal regret can be achieved by polynomial-sized transformers:

- Thompson sampling for linear bandit (Appendix E.2).
- UCB-VI [4] for MDPs (Appendix E.3).

5 Conclusions

This paper theoretically investigates the ICRL capability of supervised-pretrained transformers. We demonstrate how transformers can efficiently implement prevalent RL algorithms including LinUCB, Thompson sampling, and UCB-VI, achieving near-optimal regrets in respective settings. We also provide sample complexity guarantees for the supervised pretraining approach to learning these algorithms. The generalization error scales with the covering number of the transformer class as well as the distribution ratio between the expert and offline algorithms. Simulations validate our theoretical findings. Finally, we discuss the limitations of our results and provide additional discussions in Appendix A.1.

References

- [1] A. Agarwal, N. Jiang, S. M. Kakade, and W. Sun. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep.*, 32, 2019.
- [2] K. Ahn, X. Cheng, H. Daneshmand, and S. Sra. Transformers learn to implement preconditioned gradient descent for in-context learning. *arXiv preprint arXiv:2306.00297*, 2023.
- [3] E. Akyürek, D. Schuurmans, J. Andreas, T. Ma, and D. Zhou. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*, 2022.
- [4] M. G. Azar, I. Osband, and R. Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- [5] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [6] Y. Bai, F. Chen, H. Wang, C. Xiong, and S. Mei. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. *arXiv preprint arXiv:2306.04637*, 2023.
- [7] Y. Bengio, S. Bengio, and J. Cloutier. *Learning a synaptic learning rule*. Citeseer, 1990.
- [8] S. Bhattamishra, A. Patel, and N. Goyal. On the computational power of transformers and its implications in sequence modeling. *arXiv preprint arXiv:2006.09286*, 2020.
- [9] D. Brandfonbrener, A. Bietti, J. Buckman, R. Laroché, and J. Bruna. When does return-conditioned supervised learning work for offline reinforcement learning? *Advances in Neural Information Processing Systems*, 35:1542–1553, 2022.
- [10] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [11] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [12] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [13] W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- [14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [15] R. Dorfman, I. Shenfeld, and A. Tamar. Offline meta reinforcement learning—identifiability challenges and effective data collection strategies. *Advances in Neural Information Processing Systems*, 34:4607–4618, 2021.
- [16] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [17] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel. RI^2 : Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [18] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [19] D. J. Foster, S. M. Kakade, J. Qian, and A. Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.

- [20] H. Fu, T. Guo, Y. Bai, and S. Mei. What can a single attention layer learn? a study through the random features lens. *arXiv preprint arXiv:2307.11353*, 2023.
- [21] J. Fu, S. Levine, and P. Abbeel. One-shot learning of manipulation skills with online dynamics adaptation and neural network priors. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4019–4026. IEEE, 2016.
- [22] S. Garg, D. Tsipras, P. S. Liang, and G. Valiant. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35: 30583–30598, 2022.
- [23] S. A. Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- [24] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, and S. Levine. Meta-reinforcement learning of structured exploration strategies. *Advances in neural information processing systems*, 31, 2018.
- [25] S. Hochreiter, A. S. Younger, and P. R. Conwell. Learning to learn using gradient descent. In *Artificial Neural Networks—ICANN 2001: International Conference Vienna, Austria, August 21–25, 2001 Proceedings 11*, pages 87–94. Springer, 2001.
- [26] J. Hron, Y. Bahri, J. Sohl-Dickstein, and R. Novak. Infinite attention: Nngp and ntk for deep attention networks. In *International Conference on Machine Learning*, pages 4376–4386. PMLR, 2020.
- [27] J. Humplik, A. Galashov, L. Hasenclever, P. A. Ortega, Y. W. Teh, and N. Heess. Meta reinforcement learning as task inference. *arXiv preprint arXiv:1905.06424*, 2019.
- [28] S. Ishii, W. Yoshida, and J. Yoshimoto. Control of exploitation–exploration meta-parameter in reinforcement learning. *Neural networks*, 15(4-6):665–687, 2002.
- [29] M. Janner, Q. Li, and S. Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- [30] M. Laskin, L. Wang, J. Oh, E. Parisotto, S. Spencer, R. Steigerwald, D. Strouse, S. Hansen, A. Filos, E. Brooks, et al. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022.
- [31] T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [32] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of statistics*, pages 1302–1338, 2000.
- [33] J. N. Lee, A. Xie, A. Pacchiano, Y. Chandak, C. Finn, O. Nachum, and E. Brunskill. Supervised pretraining can learn in-context reinforcement learning. *arXiv preprint arXiv:2306.14892*, 2023.
- [34] K.-H. Lee, O. Nachum, M. S. Yang, L. Lee, D. Freeman, S. Guadarrama, I. Fischer, W. Xu, E. Jang, H. Michalewski, et al. Multi-game decision transformers. *Advances in Neural Information Processing Systems*, 35:27921–27936, 2022.
- [35] L. Li, R. Yang, and D. Luo. Focal: Efficient fully-offline meta-reinforcement learning via distance metric learning and behavior regularization. *arXiv preprint arXiv:2010.01112*, 2020.
- [36] Y. Li, M. E. Ildiz, D. Papailiopoulos, and S. Oymak. Transformers as algorithms: Generalization and implicit model selection in in-context learning. *arXiv preprint arXiv:2301.07067*, 2023.
- [37] B. Liu, J. T. Ash, S. Goel, A. Krishnamurthy, and C. Zhang. Transformers learn shortcuts to automata. *arXiv preprint arXiv:2210.10749*, 2022.
- [38] Y. Y. Lu. A padé approximation method for square roots of symmetric positive definite matrices. *SIAM journal on matrix analysis and applications*, 19(3):833–845, 1998.
- [39] E. Mitchell, R. Rafailov, X. B. Peng, S. Levine, and C. Finn. Offline meta-reinforcement learning with advantage weighting. In *International Conference on Machine Learning*, pages 7780–7791. PMLR, 2021.

- [40] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint arXiv:1803.11347*, 2018.
- [41] D. K. Naik and R. J. Mammone. Meta-neural networks that learn by learning. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 1, pages 437–442. IEEE, 1992.
- [42] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- [43] A. Nichol, J. Achiam, and J. Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [44] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [45] K. Paster, S. McIlraith, and J. Ba. You can’t count on luck: Why decision transformers and rvs fail in stochastic environments. *Advances in Neural Information Processing Systems*, 35: 38966–38979, 2022.
- [46] J. Pérez, J. Marinković, and P. Barceló. On the turing completeness of modern neural network architectures. *arXiv preprint arXiv:1901.03429*, 2019.
- [47] V. H. Pong, A. V. Nair, L. M. Smith, C. Huang, and S. Levine. Offline meta-reinforcement learning with online self-supervision. In *International Conference on Machine Learning*, pages 17811–17829. PMLR, 2022.
- [48] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [49] N. Rajaraman, L. Yang, J. Jiao, and K. Ramchandran. Toward the fundamental limits of imitation learning. *Advances in Neural Information Processing Systems*, 33:2914–2924, 2020.
- [50] N. Rajaraman, Y. Han, L. F. Yang, K. Ramchandran, and J. Jiao. Provably breaking the quadratic error compounding barrier in imitation learning, optimally. *arXiv preprint arXiv:2102.12948*, 2021.
- [51] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pages 5331–5340. PMLR, 2019.
- [52] P. Rashidinejad, B. Zhu, C. Ma, J. Jiao, and S. Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Advances in Neural Information Processing Systems*, 34:11702–11716, 2021.
- [53] A. Raventós, M. Paul, F. Chen, and S. Ganguli. Pretraining task diversity and the emergence of non-bayesian in-context learning for regression. *arXiv preprint arXiv:2306.15063*, 2023.
- [54] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, et al. A generalist agent. *arXiv preprint arXiv:2205.06175*, 2022.
- [55] S. Ross and D. Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668. JMLR Workshop and Conference Proceedings, 2010.
- [56] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [57] J. Rothfuss, D. Lee, I. Clavera, T. Asfour, and P. Abbeel. Prompt: Proximal meta-policy search. *arXiv preprint arXiv:1810.06784*, 2018.

- [58] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen, et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.
- [59] T. Schaul and J. Schmidhuber. Metalearning. *Scholarpedia*, 5(6):4650, 2010.
- [60] J. Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.
- [61] J. Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992.
- [62] N. M. Shafiqullah, Z. Cui, A. A. Altanzaya, and L. Pinto. Behavior transformers: Cloning k modes with one stone. *Advances in neural information processing systems*, 35:22955–22968, 2022.
- [63] K. Shen, J. Guo, X. Tan, S. Tang, R. Wang, and J. Bian. A study on relu and softmax in transformer. *arXiv preprint arXiv:2302.06461*, 2023.
- [64] M. Strupl, F. Faccio, D. R. Ashley, J. Schmidhuber, and R. K. Srivastava. Upside-down reinforcement learning can diverge in stochastic environments with episodic resets. *arXiv preprint arXiv:2205.06595*, 2022.
- [65] S. Thrun and L. Pratt. *Learning to learn*. Springer Science & Business Media, 2012.
- [66] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [67] J. Von Oswald, E. Niklasson, E. Randazzo, J. Sacramento, A. Mordvintsev, A. Zhmoginov, and M. Vladymyrov. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pages 35151–35174. PMLR, 2023.
- [68] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [69] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumar, and M. Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- [70] L. Wang, Q. Cai, Z. Yang, and Z. Wang. On the global optimality of model-agnostic meta-learning. In *International conference on machine learning*, pages 9837–9846. PMLR, 2020.
- [71] C. Wei, Y. Chen, and T. Ma. Statistically meaningful approximation: a case study on approximating turing machines with transformers. *Advances in Neural Information Processing Systems*, 35:12071–12083, 2022.
- [72] M. Wortsman, J. Lee, J. Gilmer, and S. Kornblith. Replacing softmax with relu in vision transformers. *arXiv preprint arXiv:2309.08586*, 2023.
- [73] S. M. Xie, A. Raghunathan, P. Liang, and T. Ma. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021.
- [74] M. Yang, D. Schuurmans, P. Abbeel, and O. Nachum. Dichotomy of control: Separating what you can control from what you cannot. *arXiv preprint arXiv:2210.13435*, 2022.
- [75] S. Yang, O. Nachum, Y. Du, J. Wei, P. Abbeel, and D. Schuurmans. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023.
- [76] S. Yao, B. Peng, C. Papadimitriou, and K. Narasimhan. Self-attention networks can process bounded hierarchical languages. *arXiv preprint arXiv:2105.11115*, 2021.
- [77] C. Yun, S. Bhojanapalli, A. S. Rawat, S. J. Reddi, and S. Kumar. Are transformers universal approximators of sequence-to-sequence functions? *arXiv preprint arXiv:1912.10077*, 2019.

- [78] R. Zhang, S. Frei, and P. L. Bartlett. Trained transformers learn linear models in-context. *arXiv preprint arXiv:2306.09927*, 2023.
- [79] Y. Zhang, A. Backurs, S. Bubeck, R. Eldan, S. Gunasekar, and T. Wagner. Unveiling transformers with lego: a synthetic reasoning task. *arXiv preprint arXiv:2206.04301*, 2022.
- [80] L. Zintgraf, K. Shiarlis, M. Igl, S. Schulze, Y. Gal, K. Hofmann, and S. Whiteson. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Framework for In-Context Reinforcement Learning | 2 |
| 3 | Statistical analysis of supervised pretraining | 3 |
| 3.1 | Main result | 3 |
| 4 | Approximation by transformers | 4 |
| 5 | Conclusions | 4 |
| A | Related work | 11 |
| A.1 | Limitation and discussion | 11 |
| B | Technical preliminaries | 12 |
| C | Details for supervised pretraining results | 15 |
| C.1 | Implications in special cases | 16 |
| D | Proofs in Section 3 | 16 |
| D.1 | Proof of Theorem 1 | 17 |
| D.2 | Proof of Proposition 10 | 18 |
| D.3 | An auxiliary lemma | 18 |
| E | Approximation by transformers | 20 |
| E.1 | LinUCB for linear bandits | 20 |
| E.2 | Thompson sampling for linear bandit | 21 |
| E.3 | UCB-VI for Tabular MDPs | 22 |
| F | Soft LinUCB for linear stochastic bandit | 23 |
| F.1 | Embedding and extraction mappings | 23 |
| F.2 | LinUCB and soft LinUCB | 23 |
| F.3 | Approximation of the ridge estimator | 24 |
| F.4 | Proof of Theorem 12 | 26 |
| F.5 | Proof of Theorem 13 | 30 |
| G | Thompson sampling for stochastic linear bandit | 32 |
| G.1 | Thompson sampling algorithm | 32 |
| G.2 | Definitions and assumptions | 33 |
| G.3 | Proof of Theorem 19 (and hence Theorem 14) | 34 |
| G.4 | Proof of Theorem 15 | 41 |
| G.5 | Proof of Lemma 20 | 43 |
| H | Learning in-context RL in markov decision processes | 45 |
| H.1 | Embedding and extraction mappings | 45 |
| H.2 | UCB-VI and soft UCB-VI | 46 |
| H.3 | Proof of Theorem 16 | 46 |
| H.4 | Proof of Theorem 17 | 53 |
| I | Experiments | 55 |
| J | Experimental details | 56 |
| J.1 | Implementation details | 56 |
| J.2 | Additional experiments and plots | 57 |
| J.3 | The effect of distribution ratio | 57 |

A Related work

Meta-learning and meta-reinforcement learning In-context reinforcement learning can be cast into the framework of meta-learning and meta-reinforcement learning [60, 61, 7, 41, 28, 59, 65]. More recently, a line of work focuses on meta-learn certain shared structures such as the dynamics of the shared tasks [21, 40], a task context identifier [51, 27, 80], exploration strategies [24], or the initialization of the network policy [18, 25, 43, 57]. Theories for this last approach of model-agnostic meta-learning have been explored by [70].

Our work focuses on a more agnostic approach to learning the learning algorithm itself [69, 17, 15, 39, 35, 47, 30, 33]. Among these works, [69, 17] focus on the online meta-RL setting with the training objective to be the total reward. Furthermore, [15, 39, 35, 47] focus on offline meta-RL, but their training objectives differ from the cross entropy loss used here, requiring explicit handling of distribution shift. The supervised pretraining approach we consider is most similar to the algorithm distillation methods of [30] and the decision-pretrained transformers of [33]. We provide quantitative sample complexity guarantees and transformer constructions absent from previous work.

In-context learning The in-context learning (ICL) capability of pretrained transformers has gained significant attention since being demonstrated on GPT-3 [11]. Recent work investigates why and how pretrained transformers perform ICL [22, 36, 67, 3, 73, 6, 78, 2, 53]. In particular, [73] propose a Bayesian framework explaining how ICL works. [22] show transformers can be trained from scratch to perform ICL of simple function classes. [67, 3, 6] demonstrate transformers can implement in-context learning algorithms via in-context gradient descent, with [6] showing transformers can perform in-context algorithm selection. [78] studied training dynamics of a single attention layer for in-context learning of linear functions. Our work focuses on the related but distinct capability of in-context decision-making for pretrained transformers.

Transformers for decision making Besides the ICRL approach, recent work has proposed goal-conditioned supervised learning (GCSL) for using transformers to make decisions [12, 29, 34, 54, 10, 62, 75]. In particular, Decision Transformer (DT) [12, 29] uses transformers to autoregressively model action sequences from offline data, conditioned on the achieved return. During inference, one queries the model with a desired high return. Limitations and modifications of GCSL have been studied in [74, 45, 64, 9]. A key distinction between GCSL and ICRL is that GCSL treats the transformer as a policy, whereas ICRL treats it as an algorithm for improving the policy based on observed trajectories.

Expressivity of transformers The transformer architecture, introduced by [66], has revolutionized natural language processing and is used in most recently developed large language models like BERT and GPT [14, 11]. The expressivity of transformers has been extensively studied [77, 46, 26, 76, 8, 79, 37, 71, 20, 6, 3, 67]. Transformers have been shown to efficiently approximate various algorithms including automata [37], Turing machines [71], and gradient descent [6, 3, 67]. Our work provides efficient transformer constructions that implement accelerated gradient descent and matrix square root algorithms, complementing existing expressivity results.

Statistical theories of imitation learning Our generalization error analysis adapts classical analysis of maximum-likelihood estimator [23]. The error compounding analysis for imitation learning appeared in early works [56, 55]. More recent theoretical analyses of imitation learning also appear in [49, 50, 52].

A.1 Limitation and discussion

In this section, we discuss some limitations of our work and some potential future directions.

Distribution ratio In Theorem 1, our regret bound of the learned algorithms $\text{Alg}_{\hat{\theta}}$ depends on the distribution ratio $\mathcal{R}_{\text{Alg}_E, \text{Alg}_0}$. While in cases like algorithm distillation [30] the distribution ratio equals one since the offline algorithm matches the expert algorithm, in the worst case, the ratio can exponentially depend on T or even become arbitrarily large. To control the distribution ratio in practice, one approach is to augment the offline trajectory dataset with a portion of trajectories generated by an expert algorithm or no-regret algorithms resembling the expert algorithm. On the other hand, further research could investigate structural assumptions of decision-making problems that avoid pessimistic dependence on the distribution ratio in regret bounds.

Guarantee of pretrained transformer Our statistical result (Theorem 1) only guarantees that the pretrained transformer learns an “algorithm” matching the expert algorithm under the pre-training distribution, even though our approximation results (Theorem 12, 14, 16) show the existence of a

transformer approximating the expert algorithm over the entire input space. In our early experiments, we noticed the learned transformers do not generalize well on out-of-distribution instances, such as with shifted reward distributions or increased number of runs T . Similar phenomena occur in other in-context learning problems (e.g. [22]). Understanding the actual algorithm implemented by the pretrained transformer through theoretical and empirical analysis is an interesting question for future work.

Alternative pretraining methods Our theoretical results study pretraining the transformer by maximizing the log-likelihood of i.i.d. offline trajectories as in Eq. (1). This aligns with standard supervised pretraining of large language models. However, alternative pretraining methods may also be effective. For instance, one could replace the log-probability in Eq. (1) with an ℓ_2 loss for continuous action spaces, consider other objectives like cumulative reward [17], or explore goal-conditioned reinforcement learning [12] for in-context RL. While our work focuses on log-likelihood pretraining, theoretical investigation of alternative methods is an interesting direction for future work.

Possibility of surpassing the expert algorithm by online training Our work considers offline pretraining by imitating the expert algorithm (i.e., $\overline{\text{Alg}}_E$), which can only learn a transformer matching the expert’s performance at best. However, through online training, where the transformer interacts with the environment, the learned transformer may surpass existing experts by training to improve itself rather than imitating a specific algorithm. Investigating whether online training enables surpassing expert algorithms is an interesting direction for future work.

Implications for practice While the focus of our work is primarily theoretical, our results lead to several practical implications for in-context reinforcement learning. One key implication is the importance of training labels (i.e., expert actions \bar{a}). When the expert algorithm depends solely on past observations, we can learn Alg_E (see Theorem 13). In contrast, when Alg_E is the optimal action a^* (involving knowledge of the underlying MDP), we can learn the posterior average of this algorithm given past observations. This corresponds to the Thompson sampling algorithm, as in Decision-Pretrained Transformers (see Theorem 15).

Furthermore, as discussed previously, the distribution ratio between the offline and expert algorithms may impact the generalization of the learned algorithm. Both our theory (see Theorem 12) and simulations (see Figure 4) show that a small distribution ratio between the offline algorithm Alg_0 and the expert algorithm $\overline{\text{Alg}}_E$ is essential, otherwise the online performance of the learned algorithm may substantially degrade. This suggests that incorporating trajectories generated purely from the expert (“on-policy ICRL”) into the offline dataset is advantageous, when feasible.

B Technical preliminaries

In this work, we will apply the following standard concentration inequality (see e.g. Lemma A.4 in [19]).

Lemma 4. *For any sequence of random variables $(X_t)_{t \leq T}$ adapted to a filtration $\{\mathcal{F}_t\}_{t=1}^T$, we have with probability at least $1 - \delta$ that*

$$\sum_{s=1}^t X_s \leq \sum_{s=1}^t \log \mathbb{E}[\exp(X_s) \mid \mathcal{F}_{s-1}] + \log(1/\delta), \quad \text{for all } t \in [T].$$

Lemma 5. *Adopt the notations in Definition 4. Then for any $\theta \in \Theta$, there exists $\theta_0 \in \Theta_0$ such that $\|\text{Alg}_{\theta_0}(\cdot \mid D_{t-1}, s_t) - \text{Alg}_{\theta}(\cdot \mid D_{t-1}, s_t)\|_1 \leq 2\rho$.*

Proof of Lemma 5. For any $\theta \in \Theta$, let $\theta_0 \in \Theta_0$ be such that $\|\log \text{Alg}_{\theta_0}(\cdot \mid D_{t-1}, s_t) - \log \text{Alg}_{\theta}(\cdot \mid D_{t-1}, s_t)\|_{\infty} \leq \rho$ for all D_{t-1}, s_t and $t \in [T]$. Then

$$\begin{aligned} & \|\text{Alg}_{\theta_0}(\cdot \mid D_{t-1}, s_t) - \text{Alg}_{\theta}(\cdot \mid D_{t-1}, s_t)\|_1 \\ &= \sum_{a \in \mathcal{A}_t} |\text{Alg}_{\theta_0}(a \mid D_{t-1}, s_t) - \text{Alg}_{\theta}(a \mid D_{t-1}, s_t)| \\ &\leq \sum_{a \in \mathcal{A}_t} e^{\max\{\log \text{Alg}_{\theta_0}(\cdot \mid D_{t-1}, s_t), \log \text{Alg}_{\theta}(\cdot \mid D_{t-1}, s_t)\}} \\ &\quad \cdot |\log \text{Alg}_{\theta_0}(\cdot \mid D_{t-1}, s_t) - \log \text{Alg}_{\theta}(\cdot \mid D_{t-1}, s_t)| \end{aligned}$$

$$\begin{aligned}
&\leq \rho \sum_{a \in \mathcal{A}_t} e^{\max\{\log \text{Alg}_{\theta_0}(\cdot|D_{t-1}, s_t), \log \text{Alg}_{\theta}(\cdot|D_{t-1}, s_t)\}} \\
&\leq \rho \sum_{a \in \mathcal{A}_t} (\text{Alg}_{\theta_0}(\cdot|D_{t-1}, s_t) + \text{Alg}_{\theta}(\cdot|D_{t-1}, s_t)) \leq 2\rho,
\end{aligned}$$

where the second line uses a Taylor expansion of e^x , the fourth line uses the assumption on θ_0 , the last line uses $e^{\max\{x,y\}} \leq e^x + e^y$ and the fact that $\text{Alg}_{\theta_0}(\cdot|D_{t-1}, s_t), \text{Alg}_{\theta}(\cdot|D_{t-1}, s_t)$ are probability functions. \square

We have the following upper bound on the covering number of the transformer class $\{\text{TF}_{\theta}^{\text{R}} : \theta \in \Theta_{D,L,M,D',B}\}$.

Lemma 6. *For the space of transformers $\{\text{TF}_{\theta}^{\text{R}} : \theta \in \bar{\Theta}_{D,L,M,D',B}\}$ with*

$$\bar{\Theta}_{D,L,M,D',B} := \left\{ \theta = (\theta_{\text{attn}}^{[L]}, \theta_{\text{mlp}}^{[L]}) : \max_{\ell \in [L]} M^{(\ell)} \leq M, \max_{\ell \in [L]} D'^{(\ell)} \leq D', \|\theta\| \leq B \right\},$$

where $M^{(\ell)}, D'^{(\ell)}$ denote the number of heads and hidden neurons in the ℓ -th layer respectively, the covering number of the set of induced algorithms $\{\text{Alg}_{\theta}, \theta \in \bar{\Theta}_{D,L,M,D',B}\}$ (c.f. Eq. (7)) satisfies

$$\log \mathcal{N}_{\bar{\Theta}_{D,L,M,D',B}}(\rho) \leq cL^2 D(MD + D') \log \left(2 + \frac{\max\{B, L, R\}}{\rho} \right)$$

for some universal constant $c > 0$.

Remark of Lemma 6. Note that the transformer classes $\Theta_{D,L,M,D',B}, \bar{\Theta}_{D,L,M,D',B}$ have the same expressivity as one can augment any $\text{TF}_{\theta} \in \bar{\Theta}_{D,L,M,D',B}$ such that the resulting $\text{TF}_{\theta, \text{aug}} \in \Theta_{D,L,M,D',B}$ by adding heads or hidden neurons with fixed zero weights. Therefore, the same bound in Lemma 6 follows for $\Theta_{D,L,M,D',B}$, and throughout the paper we do not distinguish $\Theta_{D,L,M,D',B}$ and $\bar{\Theta}_{D,L,M,D',B}$ and use them interchangeably. We also use $M^{(\ell)}, D'^{(\ell)}$ to represent the number of heads and hidden neurons in the ℓ -th layer of transformers, respectively.

Proof of Lemma 6. We start with introducing Proposition J.1 in [6].

Proposition 7 (Proposition J.1 in [6]). *The function TF^{R} is $(LB_H^L B_{\Theta})$ -Lipschitz w.r.t. $\theta \in \Theta_{D,L,M,D',B}$ for any fixed input \mathbf{H} . Namely, for any $\theta_1, \theta_2 \in \Theta_{D,L,M,D',B}$, we have*

$$\|\text{TF}_{\theta_1}^{\text{R}}(\mathbf{H}) - \text{TF}_{\theta_2}^{\text{R}}(\mathbf{H})\|_{2,\infty} \leq LB_H^L B_{\Theta} \|\theta_1 - \theta_2\|,$$

where $\|\mathbf{A}\|_{2,\infty} := \sup_{t \in [T]} \|\mathbf{A}_{\cdot t}\|_2$ for any matrix $\mathbf{A} \in \mathbb{R}^{K \times T}$, and $B_{\Theta} := BR(1 + BR^2 + B^3R^2), B_H := (1 + B^2)(1 + B^2R^3)$.

As in the Proof of Theorem 20 in [6], we can verify using Example 5.8 in [68] that the δ -covering number

$$\log N(\delta; B_{\|\cdot\|}(r), \|\cdot\|) \leq L(3MD^2 + 2DD') \log(1 + 2r/\delta), \quad (5)$$

where $B_{\|\cdot\|}(r)$ denotes any ball of radius r under the norm $\|\cdot\|$. Moreover, we have the following continuity result on the log-softmax function

Lemma 8 (Continuity of log-softmax). *For any $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$, we have*

$$\left\| \log \left(\frac{e^{\mathbf{u}}}{\|e^{\mathbf{u}}\|_1} \right) - \log \left(\frac{e^{\mathbf{v}}}{\|e^{\mathbf{v}}\|_1} \right) \right\|_{\infty} \leq 2 \|\mathbf{u} - \mathbf{v}\|_{\infty}$$

We defer the proof of Lemma 8 to the end of this section.

Note that $\text{Alg}_{\theta}(\cdot|D_{t-1}, s_t)$ corresponds to K entries in one column of $\mathbf{H}^{(L)}$ applied through the softmax function. Therefore, combining Proposition 7, Lemma 8 and Eq. (5), we conclude that for any $r > 0$, there exists a subset $\Theta_0 \in \Theta_{D,L,M,D',B}$ with size $L(3MD^2 + 2DD') \log(1 + 2r/\delta)$ such that for any $\theta \in \Theta_{D,L,M,D',B}$, there exists $\theta_0 \in \Theta_0$ with

$$\left\| \log \text{Alg}_{\theta}(\cdot|D_{t-1}, s_t) - \log \text{Alg}_{\theta_0}(\cdot|D_{t-1}, s_t) \right\|_{\infty} \leq 2LB_H^L B_{\Theta} \delta$$

for all D_T . Substituting $r = B$ and letting $\delta = \rho/(2LB_H^L B_{\Theta})$ yields the upper bound on $\mathcal{N}_{\Theta_{D,L,M,D',B}}(\rho)$ in Lemma 6.

Proof of Lemma 8. Define $\mathbf{w} := \mathbf{u} - \mathbf{v}$. Then

$$\begin{aligned}
& \left\| \log \left(\frac{e^{\mathbf{u}}}{\|e^{\mathbf{u}}\|_1} \right) - \log \left(\frac{e^{\mathbf{v}}}{\|e^{\mathbf{v}}\|_1} \right) \right\|_{\infty} \\
& \leq \|\mathbf{u} - \mathbf{v}\|_{\infty} + |\log \|e^{\mathbf{u}}\|_1 - \log \|e^{\mathbf{v}}\|_1| \\
& = \|\mathbf{u} - \mathbf{v}\|_{\infty} + \int_0^1 \left\langle \frac{e^{\mathbf{v}+t\mathbf{w}}}{\|e^{\mathbf{v}+t\mathbf{w}}\|_1}, \mathbf{w} \right\rangle dt \\
& \leq \|\mathbf{u} - \mathbf{v}\|_{\infty} + \int_0^1 \left\| \frac{e^{\mathbf{v}+t\mathbf{w}}}{\|e^{\mathbf{v}+t\mathbf{w}}\|_1} \right\|_1 \cdot \|\mathbf{w}\|_{\infty} dt \\
& = 2 \|\mathbf{u} - \mathbf{v}\|_{\infty},
\end{aligned}$$

where the third line uses the Newton-Leibniz formula. \square

\square

We present the following standard results on the convergence of GD and AGD. We refer the reader to [42] for the proof of these results.

Proposition 9 (Convergence guarantee of GD and AGD). *Suppose $L(\mathbf{w})$ is a α -strongly convex and β -smooth function on \mathbb{R}^d . Denote the condition number $\kappa := \beta/\alpha$ and $\mathbf{w}^* := \arg \min_{\mathbf{w}} L(\mathbf{w})$.*

(a). *The gradient descent iterates $\mathbf{w}_{\text{GD}}^{t+1} := \mathbf{w}_{\text{GD}}^t - \eta \nabla L(\mathbf{w}_{\text{GD}}^t)$ with stepsize $\eta = 1/\beta$ and initial point $\mathbf{w}_{\text{GD}}^0 = \mathbf{0}_d$ satisfies*

$$\begin{aligned}
\|\mathbf{w}_{\text{GD}}^t - \mathbf{w}^*\|_2^2 & \leq \exp\left(-\frac{t}{\kappa}\right) \|\mathbf{w}_{\text{GD}}^0 - \mathbf{w}^*\|_2^2, \\
L(\mathbf{w}_{\text{GD}}^t) - L(\mathbf{w}^*) & \leq \frac{\beta}{2} \exp\left(-\frac{t}{\kappa}\right) \|\mathbf{w}_{\text{GD}}^0 - \mathbf{w}^*\|_2^2.
\end{aligned}$$

(b). *The accelerated gradient descent (AGD, [42]) iterates $\mathbf{w}_{\text{AGD}}^{t+1} := \mathbf{v}_{\text{AGD}}^t - \frac{1}{\beta} L(\mathbf{v}_{\text{AGD}}^t)$, $\mathbf{v}_{\text{AGD}}^{t+1} := \mathbf{w}_{\text{AGD}}^{t+1} + \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}(\mathbf{w}_{\text{AGD}}^{t+1} - \mathbf{w}_{\text{AGD}}^t)$ with $\mathbf{w}_{\text{AGD}}^0 = \mathbf{v}_{\text{AGD}}^0 = \mathbf{0}_d$ satisfies*

$$\begin{aligned}
\|\mathbf{w}_{\text{AGD}}^t - \mathbf{w}^*\|_2^2 & \leq (1 + \kappa) \exp\left(-\frac{t}{\sqrt{\kappa}}\right) \|\mathbf{w}_{\text{AGD}}^0 - \mathbf{w}^*\|_2^2, \\
L(\mathbf{w}_{\text{AGD}}^t) - L(\mathbf{w}^*) & \leq \frac{\alpha + \beta}{2} \exp\left(-\frac{t}{\sqrt{\kappa}}\right) \|\mathbf{w}_{\text{AGD}}^0 - \mathbf{w}^*\|_2^2.
\end{aligned}$$

Transformer architecture We consider a sequence of N input vectors $\{\mathbf{h}_i\}_{i=1}^N \subset \mathbb{R}^D$, compactly written as an input matrix $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_N] \in \mathbb{R}^{D \times N}$, where each \mathbf{h}_i is a column of \mathbf{H} (also a *token*). Throughout this paper, we define $\sigma(t) := \text{ReLU}(t) = \max\{t, 0\}$ as the standard relu activation function.

Definition 1 (Masked attention layer). *A masked attention layer with M heads is denoted as $\text{Attn}_{\theta}(\cdot)$ with parameters $\theta = \{(\mathbf{V}_m, \mathbf{Q}_m, \mathbf{K}_m)\}_{m \in [M]} \subset \mathbb{R}^{D \times D}$. On any input sequence $\mathbf{H} \in \mathbb{R}^{D \times N}$, we have $\bar{\mathbf{H}} = \text{Attn}_{\theta}(\mathbf{H}) = [\bar{\mathbf{h}}_1, \dots, \bar{\mathbf{h}}_N] \in \mathbb{R}^{D \times N}$, where*

$$\bar{\mathbf{h}}_i = [\text{Attn}_{\theta}(\mathbf{H})]_i = \mathbf{h}_i + \sum_{m=1}^M \frac{1}{i} \sum_{j=1}^i \sigma(\langle \mathbf{Q}_m \mathbf{h}_i, \mathbf{K}_m \mathbf{h}_j \rangle) \cdot \mathbf{V}_m \mathbf{h}_j \in \mathbb{R}^D.$$

We remark that the use of ReLU attention layers is for technical reasons. In practice, both ReLU attention and softmax attention layers should perform well. Indeed, several studies have shown that ReLU transformers achieve comparable performance to softmax transformers across a variety of tasks [72, 63, 6].

Definition 2 (MLP layer). *An MLP layer with hidden dimension D' is denoted as $\text{MLP}_{\theta}(\cdot)$ with parameters $\theta = (\mathbf{W}_1, \mathbf{W}_2) \in \mathbb{R}^{D' \times D} \times \mathbb{R}^{D \times D'}$. On any input sequence $\mathbf{H} \in \mathbb{R}^{D \times N}$, we have $\bar{\mathbf{H}} = \text{MLP}_{\theta}(\mathbf{H}) = [\bar{\mathbf{h}}_1, \dots, \bar{\mathbf{h}}_N] \in \mathbb{R}^{D \times N}$, where*

$$\bar{\mathbf{h}}_i = \mathbf{h}_i + \mathbf{W}_2 \cdot \sigma(\mathbf{W}_1 \mathbf{h}_i) \in \mathbb{R}^D.$$

We next define L -layer decoder-based transformers. Each layer consists of a masked attention layer (see Definition 1) followed by an MLP layer (see Definition 2) and a clip operation.

Definition 3 (Decoder-based Transformer). *An L -layer decoder-based transformer, denoted as $\text{TF}_\theta^R(\cdot)$, is a composition of L masked attention layers, each followed by an MLP layer and a clip operation: $\text{TF}_\theta^R(\mathbf{H}) = \mathbf{H}^{(L)} \in \mathbb{R}^{D \times N}$, where $\mathbf{H}^{(L)}$ is defined iteratively by taking $\mathbf{H}^{(0)} = \text{clip}_R(\mathbf{H}) \in \mathbb{R}^{D \times N}$, and for $\ell \in [L]$,*

$$\mathbf{H}^{(\ell)} = \text{clip}_R\left(\text{MLP}_{\theta_{\text{mlp}}^{(\ell)}}\left(\text{Attn}_{\theta_{\text{mattn}}^{(\ell)}}\left(\mathbf{H}^{(\ell-1)}\right)\right)\right) \in \mathbb{R}^{D \times N}, \quad \text{clip}_R(\mathbf{H}) = [\text{proj}_{\|\mathbf{h}\|_2 \leq R}(\mathbf{h}_i)]_i.$$

Above, the parameter $\theta = (\theta_{\text{mattn}}^{(1:L)}, \theta_{\text{mlp}}^{(1:L)})$ consists of $\theta_{\text{mattn}}^{(\ell)} = \{(\mathbf{V}_m^{(\ell)}, \mathbf{Q}_m^{(\ell)}, \mathbf{K}_m^{(\ell)})\}_{m \in [M]} \subset \mathbb{R}^{D \times D}$ and $\theta_{\text{mlp}}^{(\ell)} = (\mathbf{W}_1^{(\ell)}, \mathbf{W}_2^{(\ell)}) \in \mathbb{R}^{D' \times D} \times \mathbb{R}^{D \times D'}$. We define the parameter class of transformers as $\Theta_{D,L,M,D',B} := \{\theta = (\theta_{\text{attn}}^{(1:L)}, \theta_{\text{mlp}}^{(1:L)}) : \|\theta\| \leq B\}$, where the norm of a transformer TF_θ^R is denoted as

$$\|\theta\| := \max_{\ell \in [L]} \left\{ \max_{m \in [M]} \left\{ \|\mathbf{Q}_m^{(\ell)}\|_{\text{op}}, \|\mathbf{K}_m^{(\ell)}\|_{\text{op}} \right\} + \sum_{m=1}^M \|\mathbf{V}_m^{(\ell)}\|_{\text{op}} + \|\mathbf{W}_1^{(\ell)}\|_{\text{op}} + \|\mathbf{W}_2^{(\ell)}\|_{\text{op}} \right\}. \quad (6)$$

We introduced clipped operations in transformers for technical reasons. For brevity, we will write $\text{TF}_\theta = \text{TF}_\theta^R$ when there is no ambiguity. We will set the clipping value R to be sufficiently large so that the clip operator does not take effect in any of our approximation results.

Algorithm induced by Transformers We equip the transformer with an embedding mapping $\mathbf{h} : \cup_{t \in [T]} \mathcal{S}_t \cup \cup_{t \in [T]} (\mathcal{A}_t \times \mathcal{R}_t) \rightarrow \mathbb{R}^D$. This assigns any state $s_t \in \mathcal{S}_t$ a D -dimensional embedding vector $\mathbf{h}(s_t) = \mathbf{h}(s_t) \in \mathbb{R}^D$, and any action-reward pair $(a_t, r_t) \in \mathcal{A}_t \times \mathcal{R}_t$ a D -dimensional embedding $\mathbf{h}(a_t, r_t) \in \mathbb{R}^D$. The embedding function \mathbf{h} should encode the time step t of the state, action, and reward. With abuse of notation, we denote $\mathbf{h}(D_{t-1}, s_t) = [\mathbf{h}(s_1), \mathbf{h}(a_1, r_1), \dots, \mathbf{h}(a_{t-1}, r_{t-1}), \mathbf{h}(s_t)]$. We define a concatenation operator $\text{cat} : \mathbb{R}^{D \times *} \rightarrow \mathbb{R}^{D \times *}$ that concatenates its inputs $\text{cat}(\mathbf{h}_1, \dots, \mathbf{h}_N) = [\mathbf{h}_1, \dots, \mathbf{h}_N]$ in most examples, but it could also insert special tokens at certain positions (in MDPs we add an additional token at the end of each episode). For a partial trajectory and current state (D_{t-1}, s_t) , we input $\mathbf{H} = \text{cat}(\mathbf{h}(s_1), \mathbf{h}(a_1, r_1), \dots, \mathbf{h}(a_{t-1}, r_{t-1}), \mathbf{h}(s_t)) \in \mathbb{R}^{D \times *}$ into the transformer. This produces output $\bar{\mathbf{H}} = \text{TF}_\theta^R(\mathbf{H}) = [\bar{\mathbf{h}}_1, \bar{\mathbf{h}}_2, \dots, \bar{\mathbf{h}}_{-2}, \bar{\mathbf{h}}_{-1}]$ with the same shape as \mathbf{H} . To extract a distribution over the action space \mathcal{A}_t with $|\mathcal{A}_t| = A$ actions, we assume a fixed linear extraction mapping $\mathbf{A} \in \mathbb{R}^{A \times D}$. The induced algorithm is then defined as: $\text{Alg}_\theta(\cdot | D_{t-1}, s_t) = \text{softmax}(\mathbf{A} \cdot \bar{\mathbf{h}}_{-1})$. The overall algorithm induced by the transformer is:

$$\text{Alg}_\theta(\cdot | D_{t-1}, s_t) = \text{softmax}(\mathbf{A} \cdot \text{TF}_\theta^R(\text{cat}(\mathbf{h}(D_{t-1}, s_t)))_{-1}). \quad (7)$$

We will always choose a proper concatenation operator cat in examples, so that in the pretraining phase, all the algorithm outputs $\{\text{Alg}_\theta(\cdot | D_{t-1}, s_t)\}_{t \leq T}$ along the trajectory can be computed in a single forward propagation.

C Details for supervised pretraining results

Definition 4 (Covering number). *For a class of algorithms $\{\text{Alg}_\theta, \theta \in \Theta\}$, we say $\Theta_0 \subseteq \Theta$ is an ρ -cover of Θ , if Θ_0 is a finite set such that for any $\theta \in \Theta$, there exists $\theta_0 \in \Theta_0$ such that*

$$\|\log \text{Alg}_{\theta_0}(\cdot | D_{t-1}, s_t) - \log \text{Alg}_\theta(\cdot | D_{t-1}, s_t)\|_\infty \leq \rho, \quad \text{for all } D_{t-1}, s_t, t \in [T].$$

The covering number $\mathcal{N}_\Theta(\rho)$ is the minimal cardinality of Θ_0 such that Θ_0 is a ρ -cover of Θ .

Definition 5 (Distribution ratio). *We define the distribution ratio of two algorithms $\text{Alg}_1, \text{Alg}_2$ by*

$$\mathcal{R}_{\text{Alg}_1, \text{Alg}_2} := \mathbb{E}_{M \sim \Lambda, D_T \sim \mathbb{P}_M^{\text{Alg}_1}} \left[\prod_{s=1}^T \frac{\text{Alg}_1(a_s | D_{s-1}, s_s)}{\text{Alg}_2(a_s | D_{s-1}, s_s)} \right] = 1 + \chi^2\left(\mathbb{P}_\Lambda^{\text{Alg}_1}, \mathbb{P}_\Lambda^{\text{Alg}_2}\right).$$

Our main result requires the realizability assumption of algorithm class $\{\text{Alg}_\theta\}_{\theta \in \Theta}$ with respect to the conditional expectation of the expert algorithm.

Assumption A (Approximate realizability). *There exists $\theta^* \in \Theta$ and $\varepsilon_{\text{real}} > 0$ such that for all $t \in [T]$,*

$$\log \mathbb{E}_{M \sim \Lambda, \bar{D}_T \sim \mathbb{P}_M^{\text{Alg}_0, \text{Alg}_E}} \left[\frac{\overline{\text{Alg}}_E(\bar{a}_t | D_{t-1}, s_t)}{\text{Alg}_{\theta^*}(\bar{a}_t | D_{t-1}, s_t)} \right] \leq \varepsilon_{\text{real}}. \quad (8)$$

C.1 Implications in special cases

Algorithm Distillation When we set $\text{Alg}_E = \text{Alg}_0$, the supervised pretraining approach corresponds to the Algorithm Distillation method introduced in [30]. In this case, it suffices to set $\bar{a}^i = a^i$ for every pretraining trajectory, eliminating the need to sample additional expert actions. The conditional expectation of the expert algorithm is given by $\overline{\text{Alg}}_E = \text{Alg}_0$, and the distribution ratio $\mathcal{R}_{\text{Alg}_E, \text{Alg}_0} = 1$. Under these conditions, Theorem 1 ensures that $\text{Alg}_{\hat{\theta}}$ imitates Alg_0 with a reward difference bounded by

$$\left| \mathfrak{R}_{\Lambda, \text{Alg}_{\hat{\theta}}}(T) - \mathfrak{R}_{\Lambda, \text{Alg}_0}(T) \right| \leq cT^2 \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right).$$

If the context algorithm Alg_0 does not perform well, we cannot expect the learned algorithm $\text{Alg}_{\hat{\theta}}$ to have good performance, regardless of the number of offline trajectories.

Decision Pretrained Transformer When we set $\text{Alg}_E^t = \text{Alg}_E^t(s_t, M) = a_t^*$ to be the optimal action at time t , the supervised pretraining approach corresponds to Decision-Pretrained Transformers (DPT) proposed in [33]. In this case, the conditional expectation of the expert algorithm $\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t) = \mathbb{E}[\text{Alg}_E(\cdot | s_t, M) | D_{t-1}, s_t] = \text{Alg}_{\text{TS}}(\cdot | D_{t-1}, s_t)$ is the Thompson sampling algorithm [33, Theorem 1], which samples from the posterior distribution of the optimal action a_t^* given by $\mathbb{P}(a_t^*(M) | D_{t-1}, s_t) \propto \Lambda(M) \cdot \mathbb{P}_M^{\text{Alg}_0}(D_{t-1}, s_t)$. This implies that learning from optimal actions effectively learns to imitate Thompson sampling. Furthermore, the context algorithm is not required to perform well for the learned algorithm to be consistent with Thompson sampling. However, a high-quality context algorithm Alg_0 may help reduce the distribution ratio \mathcal{R} , thereby learning Thompson sampling with fewer samples.

Approximate DPT In practical scenarios, the learner may not have access to the optimal action a_t^* of the environment M during pretraining. Instead, they might rely on an estimated optimal action $\hat{a}_t^* \sim \text{Alg}_E^t(\cdot | D_T)$, derived from the entire trajectory D_T . We can offer a guarantee analogous to Theorem 1, provided the distribution of the estimated action closely aligns with its posterior distribution:

$$\mathbb{E}_{D_T \sim \mathbb{P}_\Lambda^{\text{Alg}_0}} \text{KL}(\text{Alg}_E^t(\cdot | D_T) \parallel \mathbb{P}_{\text{TS}, t}(\cdot | D_T)) \leq \varepsilon_{\text{approx}}, \quad \forall t \in [T]. \quad (9)$$

Here, $\mathbb{P}_{\text{TS}, t}(\cdot | D_T)$ represents the posterior distribution of the optimal action $a_t^* = a_t^*(M)$ at time t , given the observation D_T , where $(M, D_T) \sim \mathbb{P}_\Lambda^{\text{Alg}_0}$.

Proposition 10. *Let Assumption A hold and let $\hat{\theta}$ be the solution to Eq. (1). Take $\mathcal{R} = \mathcal{R}_{\text{Alg}_{\text{TS}}, \text{Alg}_0}$ as defined in Definition 5, and $\mathcal{N}_\Theta = \mathcal{N}_\Theta((nT)^{-2})$ as defined in Definition 4. Assume that for each trajectory, an estimated optimal action is provided $\hat{a}_t^* \sim \text{Alg}_E^t(\cdot | D_T)$ at each time $t \in [T]$ satisfying Eq. (9). Assume that the rewards $|r_t| \leq 1$ almost surely. Then for some universal constant $c > 0$, with probability at least $1 - \delta$, the difference of the expected cumulative rewards between $\text{Alg}_{\hat{\theta}}$ and Alg_{TS} satisfies*

$$\left| \mathfrak{R}_{\Lambda, \text{Alg}_{\hat{\theta}}}(T) - \mathfrak{R}_{\Lambda, \text{Alg}_{\text{TS}}}(T) \right| \leq c\sqrt{\mathcal{R}} \cdot T^2 \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} + \sqrt{\varepsilon_{\text{approx}}} \right).$$

The proof of Proposition 10 is contained in Appendix D.2.

D Proofs in Section 3

In this section, $c > 0$ denotes universal constants that may differ across equations.

D.1 Proof of Theorem 1

Proof of Eq. (2) Note that we have

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \overline{\text{Alg}}_E, s_t} \sqrt{D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \\
&= \sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \text{Alg}_0, s_t} \left[\left(\prod_{s=1}^{t-1} \frac{\overline{\text{Alg}}_E(a_s|D_{s-1}, s_s)}{\text{Alg}_0(a_s|D_{s-1}, s_s)} \right) \cdot \sqrt{D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \right] \\
&\leq \sum_{t=1}^T \sqrt{\mathbb{E}_{\Lambda, \text{Alg}_0} \left(\prod_{s=1}^{t-1} \frac{\overline{\text{Alg}}_E(a_s|D_{s-1}, s_s)}{\text{Alg}_0(a_s|D_{s-1}, s_s)} \right)^2 \cdot \mathbb{E}_{\Lambda, \text{Alg}_0} D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \\
&\leq \sqrt{\mathbb{E}_{\Lambda, \text{Alg}_0} \left(\prod_{s=1}^T \frac{\overline{\text{Alg}}_E(a_s|D_{s-1}, s_s)}{\text{Alg}_0(a_s|D_{s-1}, s_s)} \right)^2 \cdot \sum_{t=1}^T \sqrt{\mathbb{E}_{\Lambda, \text{Alg}_0} D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))},
\end{aligned}$$

where the second line follows from a change of distribution argument, the third line follows from Cauchy-Schwartz inequality, and the fourth line uses the fact that

$$\begin{aligned}
& \mathbb{E}_{x, y \sim \mathbb{P}_1, \mathbb{P}_2} \left(\frac{\mathbb{Q}_1(x)\mathbb{Q}_2(y|x)}{\mathbb{P}_1(x)\mathbb{P}_2(y|x)} \right)^2 = \int \frac{\mathbb{Q}_1(x)^2 \mathbb{Q}_2^2(y|x)}{\mathbb{P}_1(x)\mathbb{P}_2(y|x)} d\mu(x, y) \\
&= \int \frac{\mathbb{Q}_1(x)^2}{\mathbb{P}_1(x)} \left(\int \frac{\mathbb{Q}_2^2(y|x)}{\mathbb{P}_2(y|x)} d\mu(y|x) \right) d\mu(x) \geq \int \frac{\mathbb{Q}_1(x)^2}{\mathbb{P}_1(x)} d\mu(x) = \mathbb{E}_{x \sim \mathbb{P}_1} \left(\frac{\mathbb{Q}_1(x)}{\mathbb{P}_1(x)} \right)^2,
\end{aligned}$$

for any probability densities $\{\mathbb{Q}_i, \mathbb{P}_i\}_{i=1,2}$ with respect to some base measure μ .

Continuing the calculation of the above lines of bounds, we have

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \overline{\text{Alg}}_E, s_t} \sqrt{D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \\
&\leq \sqrt{T} \sqrt{\mathbb{E}_{M \sim \Lambda, a_{1:T-1} \sim \text{Alg}_0, s_t} \left(\prod_{s=1}^T \frac{\overline{\text{Alg}}_E(a_s|D_{s-1}, s_s)}{\text{Alg}_0(a_s|D_{s-1}, s_s)} \right)^2} \\
&\quad \cdot \sqrt{\sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \text{Alg}_0, s_t} D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \\
&= \sqrt{T} \sqrt{\mathbb{E}_{M \sim \Lambda, a_{1:T-1} \sim \overline{\text{Alg}}_E, s_t} \left[\prod_{s=1}^T \frac{\overline{\text{Alg}}_E(a_s|D_{s-1}, s_s)}{\text{Alg}_0(a_s|D_{s-1}, s_s)} \right]} \\
&\quad \cdot \sqrt{\sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \text{Alg}_0, s_t} D_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot|D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot|D_{t-1}, s_t))} \\
&\leq cT \sqrt{\mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_0}} \sqrt{\frac{\log \mathcal{N}_{\Theta}(1/(nT)^2) + \log(T/\delta)}{n}} + \varepsilon_{\text{real}} \\
&\leq cT \sqrt{\mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_0}} \left(\sqrt{\frac{\log[\mathcal{N}_{\Theta}(1/(nT)^2)T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right),
\end{aligned}$$

where the first inequality follows from the Cauchy-Schwartz inequality, the first equality is due to a change of distribution argument, the second inequality uses Lemma 11. This completes the proof of Eq. (2).

Proof of Eq. (3) For any bounded function f such that $|f(D_T)| \leq F$ for some $F > 0$, we have

$$\begin{aligned}
& \left| \mathbb{E}_{M \sim \Lambda, a \sim \overline{\text{Alg}}_E} [f(D_T)] - \mathbb{E}_{M \sim \Lambda, a \sim \text{Alg}_{\hat{\theta}}} [f(D_T)] \right| \\
&= \left| \sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t} \sim \overline{\text{Alg}}_E, a_{t+1:T} \sim \text{Alg}_{\hat{\theta}}} [f(D_T)] - \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \overline{\text{Alg}}_E, a_{t:T} \sim \text{Alg}_{\hat{\theta}}} [f(D_T)] \right|
\end{aligned}$$

$$\leq 2F \sum_{t=1}^T \mathbb{E}_{M \sim \Lambda, a_{1:t-1} \sim \overline{\text{Alg}}_E, s_t} \text{D}_{\text{TV}}(\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t), \text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t)),$$

where the first equality uses the performance difference lemma, the last line follows from the variational representation of the total variation distance

$$\text{D}_{\text{TV}}(\mathbb{P}, \mathbb{Q}) = \sup_{\|f\|_{\infty}=1} \mathbb{E}_{\mathbb{P}}[f(X)]/2 - \mathbb{E}_{\mathbb{Q}}[f(X)]/2,$$

and

$$\text{D}_{\text{TV}}(\mathbb{P}_1(x)\mathbb{P}_2(y|x)\mathbb{P}_3(z|y), \mathbb{P}_1(x)\mathbb{P}_4(y|x)\mathbb{P}_3(z|y)) = \mathbb{E}_{x \sim \mathbb{P}_1} \text{D}_{\text{TV}}(\mathbb{P}_2(y|x), \mathbb{P}_4(y|x)) \quad (10)$$

for probability densities $\{\mathbb{P}_i\}_{i=1,2,3,4}$ with respect to some base measure μ . Since $\text{D}_{\text{TV}}(\mathbb{P}, \mathbb{Q}) \leq \sqrt{\text{D}_{\text{H}}^2(\mathbb{P}, \mathbb{Q})}$ for any distributions \mathbb{P}, \mathbb{Q} , it follows from Eq. (2) that

$$\begin{aligned} & \left| \mathbb{E}_{M \sim \Lambda, a \sim \overline{\text{Alg}}_E} [f(D_T)] - \mathbb{E}_{M \sim \Lambda, a \sim \text{Alg}_{\hat{\theta}}} [f(D_T)] \right| \\ & \leq cF \sqrt{\mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_{\hat{\theta}}}} \cdot T \left(\sqrt{\frac{\log [\mathcal{N}_{\Theta}(1/(nT)^2)T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right) \end{aligned}$$

with probability at least $1 - \delta$ for some universal constant $c > 0$. Letting $f(D_T) = \sum_{t=1}^T r_t$ and noting that $|f(D_T)| \leq T$ concludes the proof of Theorem 1.

D.2 Proof of Proposition 10

By the jointly convexity of $\text{KL}(\mathbb{P} \parallel \mathbb{Q})$ with respect to (\mathbb{P}, \mathbb{Q}) and the fact that $\text{D}_{\text{H}}^2(\mathbb{P}, \mathbb{Q}) \leq \text{KL}(\mathbb{P} \parallel \mathbb{Q})$, we have

$$\begin{aligned} & \mathbb{E}_{D_{t-1}, s_t \sim \mathbb{P}_{\Lambda}^{\text{Alg}_0}} \text{D}_{\text{H}}^2(\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t), \mathbb{P}_{\text{TS}}(\cdot | D_{t-1}, s_t)) \\ & \leq \mathbb{E}_{D_{t-1}, s_t \sim \mathbb{P}_{\Lambda}^{\text{Alg}_0}} \text{KL}(\overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t) \parallel \mathbb{P}_{\text{TS}}(\cdot | D_{t-1}, s_t)) \\ & \leq \mathbb{E}_{D_T \sim \mathbb{P}_{\Lambda}^{\text{Alg}_0}} \text{KL}(\hat{a}_t^* \parallel \mathbb{P}_{\text{TS}, t}(\cdot | D_T)) \leq \varepsilon_{\text{approx}}. \end{aligned}$$

Therefore, applying Lemma 11 gives

$$\begin{aligned} & \mathbb{E}_{M \sim \Lambda, D_T \sim \mathbb{P}_M^{\text{Alg}_0}} \left[\sum_{t=1}^T \text{D}_{\text{H}}^2(\text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t), \text{Alg}_{\text{TS}}(\cdot | D_{t-1}, s_t)) \right] \\ & \leq 2 \mathbb{E}_{M \sim \Lambda, D_T \sim \mathbb{P}_M^{\text{Alg}_0}} \left[\sum_{t=1}^T \text{D}_{\text{H}}^2(\text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t), \overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t)) + \text{D}_{\text{H}}^2(\text{Alg}_E(\cdot | D_{t-1}, s_t), \text{Alg}_{\text{TS}}(\cdot | D_{t-1}, s_t)) \right] \\ & \leq c \left(\frac{T \log [\mathcal{N}_{\Theta}(1/(nT)^2)T/\delta]}{n} + T(\varepsilon_{\text{real}} + \varepsilon_{\text{approx}}) \right) \end{aligned}$$

with probability at least $1 - \delta$. Proposition 10 follows from similar arguments as in the proof of Theorem 1 with $\varepsilon_{\text{real}}$ replaced by $\varepsilon_{\text{real}} + \varepsilon_{\text{approx}}$.

D.3 An auxiliary lemma

Lemma 11 (General guarantee for supervised pretraining). *Suppose Assumption A holds. Then the solution to Eq. (1) achieves*

$$\mathbb{E}_{D_T \sim \mathbb{P}_{\Lambda}^{\text{Alg}_0}} \left[\sum_{t=1}^T \text{D}_{\text{H}}^2(\text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t), \overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t)) \right] \leq c \frac{T \log [\mathcal{N}_{\Theta}(1/(nT)^2)T/\delta]}{n} + T \varepsilon_{\text{real}}.$$

with probability at least $1 - \delta$ for some universal constant $c > 0$.

Proof of Lemma 11.

Define

$$\mathcal{L}_{nt}(\theta) := \sum_{i=1}^n \log \text{Alg}_{\theta}(\bar{a}_t^i | D_{t-1}^i, s_t^i), \quad \text{and} \quad \mathcal{L}_{nt}(\text{expert}) := \sum_{i=1}^n \log \overline{\text{Alg}}_E(\bar{a}_t^i | D_{t-1}^i, s_t^i),$$

and let $\mathcal{L}_n(\theta) = \sum_{t=1}^T \mathcal{L}_{nt}(\theta)$, $\mathcal{L}_n(\text{expert}) = \sum_{t=1}^T \mathcal{L}_{nt}(\text{expert})$. We claim that with probability at least $1 - \delta$

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{D_T} \left[D_{\text{H}}^2(\text{Alg}_{\theta}(\cdot | D_{t-1}, s_t), \overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t)) \right] \\ & \leq \frac{\mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta)}{n} + 2 \frac{T \log \mathcal{N}_{\Theta}(1/(nT)^2)}{n} + 2 \frac{T \log(T/\delta)}{n} + \frac{4}{n} \end{aligned} \quad (11)$$

for all $\theta \in \Theta$, $i \in [T]$, where D_T follows distribution $\mathbb{P}_M^{\text{Alg}_0}(\cdot)$, $M \sim \Lambda$. For now, we assume this claim holds. Moreover, it follows from Lemma 4 and the fact $\mathcal{L}_n(\hat{\theta}) \geq \mathcal{L}_n(\theta^*)$ that

$$\begin{aligned} \frac{\mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\hat{\theta})}{n} & \leq \frac{\mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta^*)}{n} = \sum_{t=1}^T \frac{\mathcal{L}_{nt}(\text{expert}) - \mathcal{L}_{nt}(\theta^*)}{n} \\ & \leq \frac{T \log(T/\delta)}{n} + \sum_{t=1}^T \log \mathbb{E}_{D_T} \left[\frac{\overline{\text{Alg}}_E(\bar{a}_t | D_{t-1}, s_t)}{\text{Alg}_{\theta^*}(\bar{a}_t | D_{t-1}, s_t)} \right] \\ & \leq \frac{T \log(T/\delta)}{n} + T \varepsilon_{\text{real}} \end{aligned} \quad (12)$$

with probability at least $1 - \delta$.

Choosing $\theta = \hat{\theta}$ in Eq. (11) and combining it with Eq. (12) and a union bound, we obtain

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{D_T} \left[D_{\text{H}}^2(\text{Alg}_{\hat{\theta}}(\cdot | D_{t-1}, s_t), \overline{\text{Alg}}_E(\cdot | D_{t-1}, s_t)) \right] \\ & \leq T \varepsilon_{\text{real}} + 2 \left(\frac{T \log \mathcal{N}_{\Theta}(1/(nT)^2) + 2T \log(2T/\delta) + 2}{n} \right) \\ & \leq T \varepsilon_{\text{real}} + cT \left(\frac{\log \mathcal{N}_{\Theta}(1/(nT)^2) + \log(T/\delta)}{n} \right) \end{aligned}$$

with probability at least $1 - \delta$ for some universal constant $c > 0$. This completes the proof.

Proof of Eq. (11) Let Θ_0 be a $1/(nT)^2$ -covering set of Θ with covering number $n_{\text{cov}} = |\Theta_0|$. For $k \in [n_{\text{cov}}]$, $t \in [T]$, $i \in [n]$, define

$$\ell_{kt}^i = \log \frac{\overline{\text{Alg}}_E(\bar{a}_t^i | D_{t-1}^i, s_t^i)}{\text{Alg}_{\theta_k}(\bar{a}_t^i | D_{t-1}^i, s_t^i)},$$

where (D_T^i, \bar{a}^i) are the trajectory and expert actions collected in the i -th instance. Using Lemma 4 with $X_s = -\ell_{kt}^s$ and a union bound over (k, t) , conditioned on the trajectories (D_T^1, \dots, D_T^n) , we have

$$\frac{1}{2} \sum_{i=1}^n \ell_{kt}^i + \log(n_{\text{cov}} T / \delta) \geq \sum_{i=1}^n -\log \mathbb{E} \left[\exp \left(-\frac{\ell_{kt}^i}{2} \right) \right]$$

for all $k \in [n_{\text{cov}}]$, $t \in [T]$ with probability at least $1 - \delta$. Note that

$$\begin{aligned} \mathbb{E} \left[\exp \left(-\frac{\ell_{kt}^i}{2} \right) \middle| D_{t-1}^i, s_t^i \right] & = \mathbb{E}_{\mathbb{D}} \left[\sqrt{\frac{\text{Alg}_{\theta_k}(\bar{a}_t^i | D_{t-1}^i, s_t^i)}{\overline{\text{Alg}}_E(\bar{a}_t^i | D_{t-1}^i, s_t^i)}} \middle| D_{t-1}^i, s_t^i \right] \\ & = \sum_{a \in \mathcal{A}_t} \sqrt{\text{Alg}_{\theta_k}(a | D_{t-1}^i, s_t^i) \overline{\text{Alg}}_E(a | D_{t-1}^i, s_t^i)}, \end{aligned}$$

where the last inequality uses the assumption that the actions \bar{a}^i are generated using the expert $\overline{\text{Alg}}_E(\cdot | D_{t-1}^i, s_t^i)$. Therefore, for any $\theta \in \Theta$ covered by θ_k , we have

$$-\log \mathbb{E} \left[\exp \left(-\frac{\ell_{kt}^i}{2} \right) \right]$$

$$\begin{aligned}
&\geq 1 - \mathbb{E}_{D^i} \left[\sum_{a \in \mathcal{A}_t} \sqrt{\text{Alg}_{\theta_k}(a|D_{t-1}^i, s_t^i) \overline{\text{Alg}_E}(a|D_{t-1}^i, s_t^i)} \right] \\
&= 1 - \mathbb{E}_{D^i} \left[\sum_{a \in \mathcal{A}_t} \sqrt{\text{Alg}_{\theta}(a|D_{t-1}^i, s_t^i) \overline{\text{Alg}_E}(a|D_{t-1}^i, s_t^i)} \right] \\
&\quad - \mathbb{E}_{D^i} \left[\sum_{a \in \mathcal{A}_t} \sqrt{\overline{\text{Alg}_E}(a|D_{t-1}^i, s_t^i)} \left(\sqrt{\text{Alg}_{\theta_k}(a|D_{t-1}^i, s_t^i)} - \sqrt{\text{Alg}_{\theta}(a|D_{t-1}^i, s_t^i)} \right) \right] \\
&\geq \frac{1}{2} \mathbb{E}_{D^i} \left[D_{\text{H}}^2(\overline{\text{Alg}_E}(\cdot|D_{t-1}^i, s_t^i), \text{Alg}_{\theta}(\cdot|D_{t-1}^i, s_t^i)) \right] \\
&\quad - \mathbb{E}_{D^i} \left[\sum_{a \in \mathcal{A}} \left(\sqrt{\text{Alg}_{\theta}(\cdot|D_{t-1}^i, s_t^i)} - \sqrt{\text{Alg}_{\theta_k}(\cdot|D_{t-1}^i, s_t^i)} \right)^2 \right]^{1/2} \\
&\geq \frac{1}{2} \mathbb{E}_{D^i} \left[D_{\text{H}}^2(\overline{\text{Alg}_E}(\cdot|D_{t-1}^i, s_t^i), \text{Alg}_{\theta}(\cdot|D_{t-1}^i, s_t^i)) \right] - \|\text{Alg}_{\theta}(\cdot|D_{t-1}^i, s_t^i) - \text{Alg}_{\theta_k}(\cdot|D_{t-1}^i, s_t^i)\|_1^{1/2} \\
&\geq \frac{1}{2} \mathbb{E}_{D^i} \left[D_{\text{H}}^2(\overline{\text{Alg}_E}(\cdot|D_{t-1}^i, s_t^i), \text{Alg}_{\theta}(\cdot|D_{t-1}^i, s_t^i)) \right] - \frac{\sqrt{2}}{nT}
\end{aligned}$$

for all $i \in [n]$, $t \in [T]$, where the first inequality uses $-\log x \geq 1 - x$, the second inequality follows from Cauchy-Schwartz inequality, the third inequality uses $(\sqrt{x} - \sqrt{y})^2 \leq |x - y|$ for $x, y \geq 0$, the last inequality uses the fact that θ is covered by θ_k and Lemma 5. Since any $\theta \in \Theta$ is covered by θ_k for some $k \in [n_{\text{cov}}]$, and for this k summing over $t \in [T]$ gives

$$\sum_{i=1}^n \sum_{t=1}^T \ell_{kt}^i = \mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta_k) \leq \mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta) + \frac{1}{nT} \leq \mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta) + 1.$$

Therefore, with probability at least $1 - \delta$, we have

$$\begin{aligned}
&\frac{1}{2} \left(\mathcal{L}_n(\text{expert}) - \mathcal{L}_n(\theta) + 1 \right) + T \log(n_{\text{cov}}T/\delta) + \sqrt{2} \\
&\geq \frac{n}{2} \sum_{t=1}^T \mathbb{E}_{D_T} \left[D_{\text{H}}^2(\text{Alg}_{\theta}(\cdot|D_{t-1}, s_t), \text{Alg}_{\text{expert}}(\cdot|D_{t-1}, s_t)) \right]
\end{aligned}$$

for all $\theta \in \Theta$, where D_T follows $\mathbb{P}_{\Lambda}^{\text{Alg}_0}$. Multiplying both sides by $2/n$ and letting $n_{\text{cov}} = \mathcal{N}_{\Theta}(1/(nT)^2)$ yields Eq. (11). \square

E Approximation by transformers

In this section, we demonstrate the capability of transformers to implement prevalent reinforcement learning algorithms that produce near-optimal regret bounds. Specifically, we illustrate the implementation of LinUCB for stochastic linear bandits in Section E.1, Thompson sampling for stochastic linear bandits in Section E.2, and UCB-VI for tabular Markov decision process in Section E.3.

E.1 LinUCB for linear bandits

A stochastic linear bandit environment is defined by $M = (\mathbf{w}^*, \mathcal{E}, \mathbb{A}_1, \dots, \mathbb{A}_T)$. For each time step $t \in [T]$, the learner chooses an action $a_t \in \mathbb{R}^d$ from a set of actions $\mathbb{A}_t = \{\mathbf{a}_{t,1}, \dots, \mathbf{a}_{t,A}\}$, which consists of A actions and may vary over time. Upon this action selection, the learner receives a reward $r_t = \langle a_t, \mathbf{w}^* \rangle + \varepsilon_t$. Here, $\{\varepsilon_t\} \sim_{\text{iid}} \mathcal{E}$ are zero-mean noise variables, and $\mathbf{w}^* \in \mathbb{R}^d$ represents an unknown parameter vector. Stochastic linear bandit can be cast into our general framework by setting $s_t = \mathbb{A}_t$ and adopting a deterministic transition where s_t transits to s_{t+1} deterministically regardless of the chosen action.

We assume the context algorithm Alg_0 is the soft LinUCB [13]. Specifically, for each time step $t \in [T]$, the learner estimates the parameter \mathbf{w}^* using linear ridge regression $\mathbf{w}_{\text{ridge},\lambda}^t := \arg \min_{\mathbf{w} \in \mathbb{R}^d} \sum_{j=1}^{t-1} (r_j - \langle \mathbf{a}_j, \mathbf{w} \rangle)^2 + \lambda \|\mathbf{w}\|_2^2$. Subsequently, the learner calculates the upper confidence bounds for the reward of each action as $v_{tk}^* := \langle \mathbf{a}_{t,k}, \mathbf{w}_{\text{ridge},\lambda}^t \rangle + \alpha \cdot (\mathbf{a}_{t,k}^\top (\lambda \mathbf{I}_d + \sum_{j=1}^{t-1} a_j a_j^\top)^{-1} \mathbf{a}_{t,k})^{1/2}$. Finally, the learner selects an action a_t according to probability

$\{p_{t,j}^*\}_{j \in [A]} = \text{softmax}(\{v_{t,j}^*/\tau\}_{j \in [A]})$ for some sufficiently small $\tau > 0$. Note that the soft LinUCB $\text{Alg}_{\text{sLinUCB}(\tau)}$ recovers LinUCB as $\tau \rightarrow 0$.

We further assume the existence of constants $\sigma, b_a, B_a, B_w > 0$ such that the following conditions hold: $|\varepsilon_t| \leq \sigma$, $b_a \leq \|\mathbf{a}_{t,k}\|_2 \leq B_a$, and $\|\mathbf{w}^*\|_2 \leq B_w$ for all $t \in [T], k \in [A]$. Given these, the confidence parameter is defined as: $\alpha = \sqrt{\lambda}B_w + \sigma\sqrt{2\log(2B_aB_wT) + d\log((d\lambda + TB_a^2)/(d\lambda))} = \tilde{O}(\sqrt{d})$. The following result shows that the soft LinUCB algorithm can be efficiently approximated by transformers, for which the proof is contained in Appendix F.4.

Theorem 12 (Approximating the soft LinUCB). *Consider the embedding mapping \mathbf{h} , extraction mapping \mathbf{A} , and concatenation operator cat as in Appendix F.1. For any small $\varepsilon, \tau > 0$, there exists a transformer $\text{TF}_{\theta}^{\text{R}}(\cdot)$ with $\log R = \tilde{O}(1)$,*

$$D \leq O(dA), L = \tilde{O}(\sqrt{T}), M \leq 4A, D' = \tilde{O}(d + A\sqrt{Td}/(\tau\varepsilon)), \|\theta\| = \tilde{O}(A + T\sqrt{d}/(\tau\varepsilon^{1/4})), \quad (13)$$

such that taking Alg_{θ} as defined in Eq. (7), we have

$$\left| \log \text{Alg}_{\text{sLinUCB}(\tau)}(\mathbf{a}_{t,k}|D_{t-1}, s_t) - \log \text{Alg}_{\theta}(\mathbf{a}_{t,k}|D_{t-1}, s_t) \right| \leq \varepsilon, \quad \forall t \in [T], k \in [A].$$

Here $O(\cdot)$ hides some absolute constant, and $\tilde{O}(\cdot)$ additionally hides polynomial terms in $(\sigma, b_a^{-1}, B_a, B_w, \lambda^{\pm 1})$, and poly-logarithmic terms in $(T, A, d, 1/\varepsilon, 1/\tau)$.

A key component in proving Theorem 12 is demonstrating that the transformer can approximate the accelerated gradient descent algorithm for solving linear ridge regression (Lemma 18), a result of independent interest. Leveraging Theorem 12, we can derive the following regret bound for the algorithm obtained via Algorithm Distillation, with the proof provided in Appendix F.5.

Theorem 13 (Regret of LinUCB and ICRL). *Let $\Theta = \Theta_{D,L,M,D',B}$ be the class of transformers satisfying Eq. (13) with $\varepsilon = 1/T^3$ and $\tau = 1/\log(4TAB_a(B_w + 2\alpha/\sqrt{\lambda}))/\sqrt{4T} = \tilde{O}(T^{-1/2})$, and choose the clip value $\log R = \tilde{O}(1)$. Let both the context algorithm Alg_0 and the expert algorithm Alg_E coincide with the soft LinUCB algorithm $\text{Alg}_{\text{sLinUCB}(\tau)}$ with parameter τ during supervised pretraining. Then with probability at least $1 - \delta$, the learned algorithm $\text{Alg}_{\hat{\theta}}$, a solution to Eq. (1), entails the regret bound*

$$\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\hat{\theta}}}(T) \right] \leq \mathcal{O} \left(d\sqrt{T} \log(T) + T^2 \sqrt{\frac{\log(\mathcal{N}_{\Theta} \cdot T/\delta)}{n}} \right),$$

where $\log \mathcal{N}_{\Theta} \leq \tilde{O}(L^2D(MD + D') \log n) \leq \tilde{O}(T^{3.5}d^2A^3 \log n)$. Here \mathcal{O} hides polynomial terms in $(\sigma, b_a^{-1}, B_a, B_w, \lambda^{\pm 1})$, and \tilde{O} additionally hides poly-logarithmic terms in $(T, A, d, 1/\varepsilon, 1/\tau)$.

E.2 Thompson sampling for linear bandit

We continue to examine the stochastic linear bandit framework of Section E.1, now assuming a Gaussian prior $\mathbf{w}^* \sim \mathcal{N}(0, \lambda \mathbf{I}_d)$ and Gaussian noises $\{\varepsilon_t\}_{t \geq 0} \sim_{\text{iid}} \mathcal{N}(0, r)$. Additionally, we assume existence of (b_a, B_a) such that $b_a \leq \|\mathbf{a}_{t,k}\|_2 \leq B_a$. In this model, Thompson sampling also utilizes linear ridge regression. Subsequently, we establish that transformers trained under the DPT methodology can learn Thompson sampling algorithms. We state the informal theorem in Theorem 14 below, where its formal statement and proof are contained in Appendix G.

Theorem 14 (Approximating the Thompson sampling, Informal). *Consider the embedding mapping \mathbf{h} , extraction mapping \mathbf{A} , and concatenation operator cat as in F.1. Under Assumption B, C, for sufficiently small ε , there exists a transformer $\text{TF}_{\theta}^{\text{R}}(\cdot)$ with $\log R = \tilde{O}(1)$,*

$$D = \tilde{O}(AT^{1/4}d), \quad L = \tilde{O}(\sqrt{T}), \quad M = \tilde{O}(AT^{1/4}), \quad (14)$$

$$\|\theta\| = \tilde{O}(T + AT^{1/4} + \sqrt{A}), \quad D' = \tilde{O}(AT^{1/4}d),$$

such that taking Alg_{θ} as defined in Eq. (7), with probability at least $1 - \delta_0$ over $(M, D_T) \sim \mathbb{P}_{\Lambda}^{\text{Alg}}$ for any Alg , we have

$$\log \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) - \log \text{Alg}_{\theta}(\mathbf{a}_{t,k}|D_{t-1}, s_t) \leq \varepsilon, \quad \forall t \in [T], k \in [A].$$

Here, $\tilde{O}(\cdot)$ hides polynomial terms in $(M_0, C_0, \lambda^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$, and poly-logarithmic terms in $(T, A, d, 1/\varepsilon, 1/\delta_0)$, where (M_0, C_0) are parameters in Assumption B and C.

Central to proving Theorem 14 is establishing that the transformer can approximate matrix square roots via Pade decomposition (Appendix G.3), a result of independent interest. Theorem 14 thereby implies the subsequent regret bound for transformers trained under DPT.

Theorem 15 (Regret of Thompson sampling and ICRL). *Follow the assumptions of Theorem 14. Let $\Theta = \Theta_{D,L,M,D',B}$ be the class of transformers satisfying Eq. (14) with $\varepsilon = 1/(\mathcal{R}T^3)$, $\delta_0 = \delta/(2n)$, and choose the clip value $\log R = \tilde{\mathcal{O}}(1)$. Assume the trajectories are collected by some context algorithm Alg_0 , and we choose the expert algorithm $\text{Alg}_E(s_t, M) = a_t^* = \arg \max_{a \in \mathbb{A}_t} \langle \mathbf{a}, \mathbf{w}^* \rangle$ to be the optimal action of the bandit instance M for each trajectory. Then with probability at least $1 - \delta$, the learned algorithm $\text{Alg}_{\hat{\theta}}$, a solution to Eq. (1), entails regret bound*

$$\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\hat{\theta}}}(T) \right] \leq \mathcal{O} \left(d\sqrt{T} \log(Td) + \sqrt{\mathcal{R}} \cdot T^2 \sqrt{\frac{\log(\mathcal{N}_{\Theta} T / \delta)}{n}} \right),$$

where $\mathcal{R} = \mathcal{R}_{\text{Alg}_{\text{TS}}, \text{Alg}_0}$, and $\log \mathcal{N}_{\Theta} \leq \tilde{\mathcal{O}}(L^2 D(MD + D') \log n) \leq \tilde{\mathcal{O}}(T^{5/4} A^2 d(M_0 + A\sqrt{T}d) \log n)$. Here \mathcal{O} hides polynomial terms in $(\lambda^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$, and $\tilde{\mathcal{O}}$ additionally hides poly-logarithmic terms in $(M_0, C_0, T, A, d, 1/\varepsilon, 1/\delta_0)$.

E.3 UCB-VI for Tabular MDPs

A finite-horizon tabular MDP is specified by $M = (\mathcal{S}, \mathcal{A}, H, \{P_h\}_{h \in [H]}, \{R_h\}_{h \in [H]}, \mu_1)$, with H being the time horizon, \mathcal{S} the state space of size S , \mathcal{A} the action space of size A , and $\mu_1 \in \Delta(\mathcal{S})$ defining the initial state distribution. At each time step $h \in [H]$, $P_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ denotes the state transition dynamics and $R_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ gives the reward function. A policy $\pi := \{\pi_h : (\mathcal{S} \times \mathcal{A} \times \mathbb{R})^{h-1} \times \mathcal{S} \rightarrow \Delta(\mathcal{A})\}_{h \in [H]}$ maps history and state to a distribution over actions. The value of policy π interacting with environment M is defined as the expected cumulative reward $V_M(\pi) = \mathbb{E}_{M, \pi}[\sum_{h=1}^H R_h(s_h, a_h)]$. A policy π^* is said to be optimal if $\pi^* = \arg \max_{\pi \in \Delta(\Pi)} V_M(\pi)$.

We let the context algorithm Alg_0 interact with an MDP instance M to generate K episodes, each consisting of H horizon sequences $(s_{k,h}, a_{k,h}, r_{k,h})_{k \in [K], h \in [H]}$. These can be reindexed into a single trajectory $D_T = \{(s_t, a_t, r_t)\}_{t \in [T]}$ with $t = H(k-1) + h$ and $T = KH$. The Bayes regret of any algorithm Alg gives $\mathbb{E}_{M \sim \Lambda}[KV_M(\pi^*) - \mathfrak{R}_{M, \text{Alg}}(T)]$.

Near minimax-optimal regret for tabular MDPs can be attained through the UCB-VI algorithm [4]. We demonstrate that transformers are capable of approximating the soft UCB-VI algorithm $\text{Alg}_{\text{sUCBVI}(\tau)}$, a slight modification of UCB-VI formalized in Appendix H.1.

Theorem 16 (Approximating the soft UCB-VI). *Consider the embedding mapping \mathfrak{h} , extraction mapping \mathfrak{A} , and concatenation operator cat as in Appendix H.1. There exists a transformer $\text{TF}_{\hat{\theta}}^{\text{R}}(\cdot)$ with $\log R = \tilde{\mathcal{O}}(1)$,*

$$\begin{aligned} D &= \mathcal{O}(HS^2A), \quad L = 2H + 8, \quad M = \mathcal{O}(HS^2A), \\ D' &= \mathcal{O}(K^2HS^2A), \quad \|\theta\| \leq \tilde{\mathcal{O}}(K^2HS^2A + K^3 + 1/\tau), \end{aligned} \quad (15)$$

such that $\text{Alg}_{\text{sUCBVI}(\tau)}(a|D_{t-1}, s_t) = \text{Alg}_{\hat{\theta}}(a|D_{t-1}, s_t)$ for all $t \in [T], a \in \mathcal{A}$. Here $\mathcal{O}(\cdot)$ hides universal constants and $\tilde{\mathcal{O}}(\cdot)$ hides poly-logarithmic terms in $(H, K, S, A, 1/\tau)$.

Leveraging Theorem 16, we can derive the following regret bound for the algorithm obtained via Algorithm Distillation.

Theorem 17 (Regret of UCB-VI and ICRL). *Let $\Theta = \Theta_{D,L,M,D',B}$ be the class of transformers satisfying Eq. (15) with $\tau = 1/K$, and choose the clip value $\log R = \tilde{\mathcal{O}}(1)$. Let both the context algorithm Alg_0 and the expert algorithm Alg_E coincide with the soft UCB-VI algorithm $\text{Alg}_{\text{sUCBVI}(\tau)}$ during supervised pretraining. Then with probability at least $1 - \delta$, the learned algorithm $\text{Alg}_{\hat{\theta}}$, a solution to Eq. (1), entails regret bound*

$$\mathbb{E}_{M \sim \Lambda}[KV_M(\pi^*) - \mathfrak{R}_{M, \text{Alg}_{\hat{\theta}}}(T)] \leq \tilde{\mathcal{O}} \left(H^2 \sqrt{SAK} + H^3 S^2 A + T^2 \sqrt{\frac{\log(\mathcal{N}_{\Theta} T / \delta)}{n}} \right),$$

where $\log \mathcal{N}_{\Theta} \leq \tilde{\mathcal{O}}(L^2 D(MD + D') \log n) = \tilde{\mathcal{O}}(H^4 S^4 A^3 (K^2 + HS^2 A) \log n)$, and $\tilde{\mathcal{O}}(\cdot)$ hides poly-logarithmic terms in (H, K, S, A) .

F Soft LinUCB for linear stochastic bandit

Throughout this section, we use $c > 0$ to denote universal constants whose values may vary from line to line. Moreover, for notational simplicity, we use $O(\cdot)$ to hide universal constants, $\mathcal{O}(\cdot)$ to hide polynomial terms in the problem parameters $(\sigma, b_a^{-1}, B_a, B_w, \lambda^{\pm 1})$, and $\tilde{\mathcal{O}}(\cdot)$ to hide both poly-logarithmic terms in $(T, A, d, 1/\varepsilon, 1/\tau)$ and polynomial terms in $(\sigma, b_a^{-1}, B_a, B_w, \lambda^{\pm 1})$. We also use the bold font $\mathbf{a}_t \in \mathbb{R}^d$ to denote the selected action vector a_t at time $t \in [T]$.

This section is organized as follows. Section F.1 discusses the embedding and extraction formats of transformers for the stochastic linear bandit environment. Section F.2 describes the LinUCB and the soft LinUCB algorithms. Section F.3 introduces and proves a lemma on approximating the linear ridge regression estimator, which is important for proving Theorem 12. We prove Theorem 12 in Section F.4 and prove Theorem 13 in Section F.5.

F.1 Embedding and extraction mappings

Consider the embedding in which for each $t \in [T]$, we have two tokens $\mathbf{h}_{2t-1}, \mathbf{h}_{2t} \in \mathbb{R}^D$ such that

$$\mathbf{h}_{2t-1} = \begin{bmatrix} \mathbf{0}_{d+1} \\ \mathbb{A}_t \\ \mathbf{0}_A \\ \mathbf{0} \\ \mathbf{pos}_{2t-1} \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2t-1}^a \\ \mathbf{h}_{2t-1}^b \\ \mathbf{h}_{2t-1}^c \\ \mathbf{h}_{2t-1}^d \end{bmatrix}, \quad \mathbf{h}_{2t} = \begin{bmatrix} \mathbf{a}_t \\ r_t \\ \mathbf{0}_{Ad} \\ \mathbf{0}_A \\ \mathbf{0} \\ \mathbf{pos}_{2t} \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2t}^a \\ \mathbf{h}_{2t}^b \\ \mathbf{h}_{2t}^c \\ \mathbf{h}_{2t}^d \end{bmatrix},$$

where $\mathbf{h}_{2t-1}^b = \mathbb{A}_t = [\mathbf{a}_{t,1}^\top \ \dots \ \mathbf{a}_{t,A}^\top]^\top$ denotes the action set at time t , $\mathbf{h}_{2t}^a = [\mathbf{a}_t^\top \ r_t]^\top$ denotes the action and the observed reward at time t , \mathbf{h}_{2t-1}^c is used to store the (unnormalized) policy at time t , $\mathbf{0}$ in \mathbb{R}^d denotes an additional zero vector with dimension $O(dA)$, and $\mathbf{pos}_i := (i, i^2, 1)^\top$ for $i \in [2T]$ is the positional embedding. Note that the token dimension $D = O(dA)$. In addition, we define the token matrix $\mathbf{H}_t := [\mathbf{h}_1, \dots, \mathbf{h}_{2t}] \in \mathbb{R}^{D \times 2t}$ for all $t \in [T]$.

Offline pretraining During pretraining, the transformer TF_θ takes in $\mathbf{H}_T^{\text{pre}} := \mathbf{H}_T$ as the input token matrix and generates $\mathbf{H}_T^{\text{post}} := \text{TF}_\theta(\mathbf{H}_T^{\text{pre}})$ as the output. For each step $t \in [T]$, we define the induced policy $\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2t-1}^{\text{post},c})}{\|\exp(\mathbf{h}_{2t-1}^{\text{post},c})\|_1} \in \Delta^A$, whose i -th entry is the probability of selecting action $\mathbf{a}_{t,i}$ given (D_{t-1}, s_t) . We then find the transformer $\hat{\theta} \in \Theta$ by solving Eq. (1). Due to the decoder structure of transformer TF_θ , the $2t - 1$ -th token only has access to the first $2t - 1$ tokens. Therefore the induced policy is determined by the historical data (D_{t-1}, s_t) and does not depend on future observations.

Rollout At each time $t \in [T]$, given the action set \mathbb{A}_t (i.e., current state s_t) and the previous data D_{t-1} , we first construct the token matrix $\mathbf{H}_{\text{roll},t}^{\text{pre}} = [\mathbf{H}_{t-1}, \mathbf{h}_{2t-1}] \in \mathbb{R}^{D \times (2t-1)}$. The transformer then takes $\mathbf{H}_{\text{roll},t}^{\text{pre}}$ as the input and generates $\mathbf{H}_{\text{roll},t}^{\text{post}} = [\mathbf{H}_{t-1}^{\text{post}}, \mathbf{h}_{2t-1}^{\text{post}}] = \text{TF}_\theta(\mathbf{H}_{\text{roll},t}^{\text{pre}})$. Next, the agent selects an action $\mathbf{a}_t \in \mathbb{A}_t$ according to the induced policy $\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2t-1}^{\text{post},c})}{\|\exp(\mathbf{h}_{2t-1}^{\text{post},c})\|_1} \in \Delta^A$ and observes the reward r_t .

Embedding and extraction mappings To integrate the above construction into the framework described in Section 2, we have the embedding vectors $\mathbf{h}(s_t) := \mathbf{h}_{2t-1}, \mathbf{h}(a_t, r_t) := \mathbf{h}_{2t}$, the concatenation operator $\text{cat}(\mathbf{h}_1, \dots, \mathbf{h}_N) := [\mathbf{h}_1, \dots, \mathbf{h}_N]$, the input token matrix $\mathbf{H} = \mathbf{H}_{\text{roll},t}^{\text{pre}} := \text{cat}(\mathbf{h}(s_1), \mathbf{h}(a_1, r_1), \dots, \mathbf{h}(a_{t-1}, r_{t-1}), \mathbf{h}(s_t)) \in \mathbb{R}^{D \times (2t-1)}$, the output token matrix $\bar{\mathbf{H}} = \mathbf{H}_{\text{roll},t}^{\text{post}}$, and the linear extraction map \mathbf{A} satisfies $\mathbf{A} \cdot \bar{\mathbf{h}}_{-1} = \mathbf{A} \cdot \bar{\mathbf{h}}_{2t-1}^{\text{post}} = \mathbf{h}_{2t-1}^{\text{post},c}$.

F.2 LinUCB and soft LinUCB

Let T be the total time and $\lambda, \alpha > 0$ be some prespecified values. At each time $t \in [T]$, LinUCB consists of the following steps:

1. Computes the ridge estimator $\mathbf{w}_{\text{ridge},\lambda}^t = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2t} \sum_{j=1}^{t-1} (r_j - \langle \mathbf{a}_j, \mathbf{w} \rangle)^2 + \frac{\lambda}{2t} \|\mathbf{w}\|_2^2$.
2. For each action $k \in [A]$, computes $v_{tk}^* := \langle \mathbf{a}_{t,k}, \mathbf{w}_{\text{ridge},\lambda}^t \rangle + \alpha \sqrt{\mathbf{a}_{t,k}^\top \mathbf{A}_t^{-1} \mathbf{a}_{t,k}}$, where $\mathbf{A}_t = \lambda \mathbf{I}_d + \sum_{j=1}^{t-1} \mathbf{a}_j \mathbf{a}_j^\top$.

3. Selects the action $\mathbf{a}_{t,j}$ with $j := \arg \max_{k \in [A]} v_{tk}^*$.

Unless stated otherwise, in step 2 above we choose $\alpha = \alpha(\delta)$ with $\delta = 1/(2B_a B_w T)$ and

$$\alpha(\delta) := \sqrt{\lambda B_w + \sigma \sqrt{2 \log(1/\delta) + d \log((d\lambda + TB_a^2)/(d\lambda))}} = \mathcal{O}(\sqrt{d \log T}) = \tilde{\mathcal{O}}(\sqrt{d}).$$

In this work, to facilitate the analysis of supervised pretraining, we consider soft LinUCB (denoted by sLinUCB(τ)), which replaces step 3 in LinUCB with

3' Selects the action $\mathbf{a}_{t,j}$ with probability $\frac{\exp(v_{tj}^*/\tau)}{\|\exp(v_{tj}^*/\tau)\|_1}$ for $j \in [A]$.

Note that soft LinUCB recovers the standard LinUCB as $\tau \rightarrow 0$.

F.3 Approximation of the ridge estimator

In this section, we present a lemma on how transformers can approximately implement the ridge regression estimator in-context.

Throughout the proof, for $t \in [2T]$, we let $\mathbf{h}_t^{(L)}$ denote the i -th token in the output token matrix obtained after passing through an L -layer transformer. We also define $\text{read}_{\mathbf{w}_{\text{ridge}}} : \mathbb{R}^D \mapsto \mathbb{R}^d$ be the operator that gives the values of d coordinates in the token vector that are used to store the estimation of the ridge estimate.

Lemma 18 (Approximation of the ridge estimator). *For any small $\varepsilon > 0$, there exists an attention-only (i.e., no MLP layers) transformer $\text{TF}_\theta(\cdot)$ with*

$$L = \left\lceil \frac{4T(B_a^2 + \lambda)}{\lambda} \log(TB_a(B_a B_w + \sigma)/(\lambda\varepsilon)) \right\rceil = \tilde{\mathcal{O}}(T), \quad \max_{\ell \in [L]} M^{(\ell)} \leq 3, \quad \|\boldsymbol{\theta}\| \leq \sqrt{2} + \frac{\lambda + 2}{B_a^2 + \lambda} = \mathcal{O}(1)$$

such that $\|\text{read}_{\mathbf{w}_{\text{ridge}}}(\mathbf{h}_{2t-1}^{(L)}) - \mathbf{w}_{\text{ridge},\lambda}^t\|_2 \leq \varepsilon$ for all $t \in [T]$.

Moreover, there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = \left\lceil 2\sqrt{2T} \sqrt{\frac{B_a^2 + \lambda}{\lambda}} \log\left(\frac{(2T(B_a^2 + \lambda) + \lambda)TB_a(B_a B_w + \sigma)}{\lambda^2 \varepsilon}\right) \right\rceil = \tilde{\mathcal{O}}(\sqrt{T}), \quad \max_{\ell \in [L]} M^{(\ell)} \leq 4,$$

$$\max_{\ell \in [L]} D^{(\ell)} \leq 4d, \quad \|\boldsymbol{\theta}\| \leq 10 + \frac{\lambda + 2}{B_a^2 + \lambda} = \mathcal{O}(1)$$

such that $\|\text{read}_{\mathbf{w}_{\text{ridge}}}(h_{2t-1}^{(L)}) - \mathbf{w}_{\text{ridge},\lambda}^t\|_2 \leq \varepsilon$ for all $t \in [T]$.

Results similar to Lemma 18 have been shown in [6] under a different scenario. However, we remark that the second part of Lemma 18 has a weaker requirement on the number of layers as we prove that transformers can implement accelerated gradient descent (AGD, [42]) in-context.

Proof of Lemma 18. Note that $\lambda \mathbf{I}_d \preceq \mathbf{A}_t \preceq (TB_a^2 + \lambda) \mathbf{I}_d$. Therefore the optimization problem

$$\mathbf{w}_{\text{ridge},\lambda}^t = \arg \min_{\mathbf{w} \in \mathbb{R}^d} L(\mathbf{w}) := \arg \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2(2t-1)} \sum_{j=1}^{t-1} (r_j - \langle \mathbf{a}_j, \mathbf{w} \rangle)^2 + \frac{\lambda}{2(2t-1)} \|\mathbf{w}\|_2^2$$

is $\lambda/(2t-1)$ -strongly convex and $(B_a^2 + \lambda)$ -smooth and the condition number $\kappa \leq 2T(B_a^2 + \lambda)/\lambda$. Moreover, by the definition of $\mathbf{w}_{\text{ridge},\lambda}^t$ we have

$$\begin{aligned} \|\mathbf{w}_{\text{ridge},\lambda}^t\|_2 &= \|(\lambda \mathbf{I}_d + \sum_{j=1}^{t-1} \mathbf{a}_j \mathbf{a}_j^\top)^{-1} (\sum_{j=1}^{t-1} \mathbf{a}_j r_j)\|_2 \leq \|(\lambda \mathbf{I}_d + \sum_{j=1}^{t-1} \mathbf{a}_j \mathbf{a}_j^\top)^{-1}\|_2 \cdot \|\sum_{j=1}^{t-1} \mathbf{a}_j r_j\|_2 \\ &\leq \frac{TB_a(B_a B_w + \sigma)}{\lambda} \end{aligned}$$

for all $t \in [T]$.

Proof of part 1 By Proposition 9, we see that $L = \lceil 4T(B_a^2 + \lambda) \log(TB_a(B_a B_w + \sigma)/(\lambda\varepsilon))/\lambda \rceil$ steps of gradient descent with stepsize $\eta = 1/(B_a^2 + \lambda)$ starting from $\mathbf{w}_{\text{GD}}^0 = \mathbf{0}_d$ finds \mathbf{w}_{GD}^L such that $\|\mathbf{w}_{\text{GD}}^L - \mathbf{w}_{\text{ridge},\lambda}^t\|_2 \leq \varepsilon$.

Now we prove that one attention-only layer can implement one step of gradient descent

$$\mathbf{w}_{\text{GD}}^{\ell+1} := \mathbf{w}_{\text{GD}}^\ell - \frac{\eta}{2t-1} \sum_{j=1}^{t-1} (\langle \mathbf{a}_j, \mathbf{w}_{\text{GD}}^\ell \rangle - r_j) \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \mathbf{w}_{\text{GD}}^\ell.$$

We encode the algorithm using the last token (i.e., the $2t-1$ -th token). Denote the first d entries of \mathbf{h}_{2t-1}^d by $\widehat{\mathbf{w}}$ and define $\text{read}_{\mathbf{w}_{\text{ridge}}}(\mathbf{h}_{2t-1}) = \widehat{\mathbf{w}}$. Starting from $\widehat{\mathbf{w}}^0 = \mathbf{0}_d$, for each layer $\ell \in [L]$, we let the number of heads $M^{(\ell)} = 3$ and choose $\mathbf{Q}_{1,2,3}^{(\ell)}, \mathbf{K}_{1,2,3}^{(\ell)}, \mathbf{V}_{1,2,3}^{(\ell)}$ such that for even tokens \mathbf{h}_{2j} with $j \leq t-1$ and odd tokens \mathbf{h}_{2j-1} with $j \leq t$

$$\begin{aligned} \mathbf{Q}_1^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} \widehat{\mathbf{w}}^{\ell-1} \\ 1 \end{bmatrix}, \quad \mathbf{K}_1^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = \begin{bmatrix} \mathbf{a}_j \\ -r_j \end{bmatrix}, \quad \mathbf{V}_1^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = -\eta \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_1^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \mathbf{0}, \quad \mathbf{V}_1^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \mathbf{0} \\ \mathbf{Q}_2^{(\ell)} &= -\mathbf{Q}_1^{(\ell)}, \quad \mathbf{K}_2^{(\ell)} = \mathbf{K}_1^{(\ell)}, \quad \mathbf{V}_2^{(\ell)} = -\mathbf{V}_1^{(\ell)}, \\ \mathbf{Q}_3^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} 1 \\ -(2t-1) \\ 1 \end{bmatrix}, \quad \mathbf{K}_3^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = \begin{bmatrix} 1 \\ 1 \\ 2j \end{bmatrix}, \quad \mathbf{K}_3^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \begin{bmatrix} 1 \\ 1 \\ 2j-1 \end{bmatrix}, \quad \mathbf{V}_3^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} = -\eta\lambda \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{w}}^{\ell-1} \\ \mathbf{0} \end{bmatrix}. \end{aligned}$$

Summing up the three heads and noting that $t = \sigma(t) - \sigma(-t)$, we see that the $\widehat{\mathbf{w}}$ part of \mathbf{h}_{2t-1} (i.e., $\text{read}_{\mathbf{w}_{\text{ridge}}}(\mathbf{h}_{2t-1})$) has the update

$$\begin{aligned} \widehat{\mathbf{w}}^l &= \widehat{\mathbf{w}}^{l-1} - \frac{\eta}{2t-1} \sum_{j=1}^t [\sigma(\langle \mathbf{a}_j, \widehat{\mathbf{w}}^{l-1} \rangle - r_j) - \sigma(r_j - \langle \mathbf{a}_j, \widehat{\mathbf{w}}^{l-1} \rangle)] \mathbf{a}_j \\ &\quad - \frac{\eta\lambda}{2t-1} \left[\sum_{j=1}^{t-1} (\sigma(1+2j-2t) \mathbf{V}_3^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} + \sigma(1+2j-2t+1) \mathbf{V}_3^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)}) + \mathbf{V}_3^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} \right] \\ &= \widehat{\mathbf{w}}^{l-1} - \frac{\eta}{2t-1} \sum_{j=1}^t [\langle \mathbf{a}_j, \widehat{\mathbf{w}}^{l-1} \rangle - r_j] \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \mathbf{V}_3^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} \\ &= \widehat{\mathbf{w}}^{l-1} - \frac{\eta}{2t-1} \sum_{j=1}^{t-1} [\langle \mathbf{a}_j, \widehat{\mathbf{w}}^{l-1} \rangle - r_j] \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \widehat{\mathbf{w}}^{l-1}, \end{aligned}$$

which is one step of gradient descent with stepsize η . Moreover, it is easy to see that one can choose the matrices such that $\max_{m \in [3]} \|\mathbf{Q}_m^{(\ell)}\|_{\text{op}} = \max_{m \in [3]} \|\mathbf{K}_m^{(\ell)}\|_{\text{op}} = \sqrt{2}$ and $\|\mathbf{V}_1^{(\ell)}\|_{\text{op}} = \|\mathbf{V}_2^{(\ell)}\|_{\text{op}} = \eta, \|\mathbf{V}_3^{(\ell)}\|_{\text{op}} = \lambda\eta$. Therefore the norm of the transformer $\|\boldsymbol{\theta}\| \leq \sqrt{2} + (\lambda+2)/(B_a^2 + \lambda)$.

Proof of part 2 Similarly, Proposition 9 shows $L = \lceil 2\sqrt{2T(B_a^2 + \lambda)}/\lambda \log((1 + \kappa)TB_a(B_a B_w + \sigma)/(\lambda\varepsilon)) \rceil$ steps of accelerated gradient descent gives $\|\mathbf{w}_{\text{AGD}}^L - \mathbf{w}_{\text{ridge},\lambda}^t\|_2 \leq \varepsilon$.

Again, we encode the algorithm using the last token (i.e., the $2t-1$ -th token). Denote the first $d, d+1 \sim 2d, 2d+1 \sim 3d$ entries of \mathbf{h}_{2t-1}^d by $\widehat{\mathbf{w}}_a, \widehat{\mathbf{w}}_b, \widehat{\mathbf{v}}$ respectively. Starting from $\widehat{\mathbf{w}}_a^0 = \widehat{\mathbf{w}}_b^0 = \widehat{\mathbf{v}}^0 = \mathbf{0}_d$, AGD updates the parameters as follows:

$$\widehat{\mathbf{w}}_a^\ell = \widehat{\mathbf{w}}_a^{\ell-1} + (\widehat{\mathbf{v}}^{\ell-1} - \widehat{\mathbf{w}}_a^{\ell-1}) - \eta \nabla L(\widehat{\mathbf{v}}^{\ell-1}), \quad (16a)$$

$$\widehat{\mathbf{v}}^\ell = \widehat{\mathbf{v}}^{\ell-1} + [\widehat{\mathbf{w}}_a^\ell + \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} (\widehat{\mathbf{w}}_a^\ell - \widehat{\mathbf{w}}_b^{\ell-1}) - \widehat{\mathbf{v}}^{\ell-1}], \quad (16b)$$

$$\widehat{\mathbf{w}}_b^\ell = \widehat{\mathbf{w}}_b^{\ell-1} + (\widehat{\mathbf{w}}_a^\ell - \widehat{\mathbf{w}}_b^{\ell-1}). \quad (16c)$$

We show that one attention layer and one MLP layer can implement one step of AGD as above. Namely, Eq. (16a) can be obtained using the same attention layer we constructed for gradient descent with $\widehat{\mathbf{v}}$ replacing $\widehat{\mathbf{w}}$, and an extra head with

$$\mathbf{Q}_4^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} = \begin{bmatrix} 2t-1 \\ -(2t-1)^2 \\ 1 \end{bmatrix}, \quad \mathbf{K}_4^{(\ell)} \mathbf{h}_i^{(\ell-1)} = \begin{bmatrix} 1 \\ 1 \\ i^2 \end{bmatrix}, \quad \mathbf{V}_4^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} = \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{v}}^{\ell-1} - \widehat{\mathbf{w}}_a^{\ell-1} \\ \mathbf{0} \end{bmatrix}$$

for $i \leq 2t - 1$ that gives $\widehat{\mathbf{v}}^{\ell-1} - \widehat{\mathbf{w}}_a^{\ell-1}$. Denote the output tokens of the attention layer by $\tilde{\mathbf{h}}$. Eq. (16b), (16c) can be implemented using one layer of MLP. Concretely, we choose $\mathbf{W}_1^{(\ell)}, \mathbf{W}_2^{(\ell)}$ such that

$$\mathbf{W}_1^{(\ell)} \tilde{\mathbf{h}}_{2t-1}^{(\ell-1)} = \begin{bmatrix} \mathbf{w}_a^\ell + \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}(\widehat{\mathbf{w}}_a^\ell - \widehat{\mathbf{w}}_b^{\ell-1}) - \widehat{\mathbf{v}}^{t-1} \\ -\mathbf{w}_a^\ell - \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}(\widehat{\mathbf{w}}_a^{\ell-1} - \widehat{\mathbf{w}}_b^{\ell-1}) + \widehat{\mathbf{v}}^{t-1} \\ \mathbf{w}_a^\ell - \mathbf{w}_b^{\ell-1} \\ -\mathbf{w}_a^\ell + \mathbf{w}_b^{\ell-1} \end{bmatrix}, \quad \mathbf{W}_2^{(\ell)} \sigma(\mathbf{W}_1^{(\ell)} \tilde{\mathbf{h}}_{2t-1}^{(\ell-1)}) = \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{w}}_b^\ell \\ \widehat{\mathbf{v}}^\ell \\ \mathbf{0} \end{bmatrix}.$$

Since $t = \sigma(t) - \sigma(-t)$ for $t \in \mathbb{R}$, it is readily verified that one can choose the linear maps such that $\|\mathbf{W}_1^{(\ell)}\|_{\text{op}} \leq 4\sqrt{2}$, $\|\mathbf{W}_2^{(\ell)}\|_{\text{op}} = \sqrt{2}$. Combining this with the attention layer for Eq. (16a) and noting that $\|\mathbf{V}_4^{(\ell)}\|_{\text{op}} = \sqrt{2}$, we verify that the transformer we constructed has norm $\|\boldsymbol{\theta}\| \leq 10 + (\lambda + 2)/(B_a^2 + \lambda)$. This completes the proof of Lemma 18. \square

F.4 Proof of Theorem 12

We construct a transformer that implements the following steps at each time $t \in [T]$ starting with $\mathbf{h}_{2t-1}^x = \mathbf{h}_{2t-1}^{\text{pre},x}$ for $x \in \{a, b, c, d\}$

$$\mathbf{h}_{2t-1} = \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},a} \\ \mathbf{h}_{2t-1}^{\text{pre},b} \\ \mathbf{h}_{2t-1}^{\text{pre},c} \\ \mathbf{h}_{2t-1}^{\text{pre},d} \end{bmatrix} \xrightarrow{\text{step 1}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mathbf{w}}_{\text{ridge}} \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 2}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mathbf{w}}_{\text{ridge}} \\ \star \\ \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,1} \\ \vdots \\ \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,A} \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 3}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mathbf{w}}_{\text{ridge}} \\ \star \\ \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,1} \\ \vdots \\ \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,A} \\ \widehat{v}_{t1}/\tau \\ \vdots \\ \widehat{v}_{tA}/\tau \\ \mathbf{0} \\ \text{pos} \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{post},a} \\ \mathbf{h}_{2t-1}^{\text{post},b} \\ \mathbf{h}_{2t-1}^{\text{post},c} \\ \mathbf{h}_{2t-1}^{\text{post},d} \end{bmatrix}, \quad (17)$$

where $\text{pos} := [t, i^2, 1]^\top$; $\widehat{\mathbf{w}}_{\text{ridge}}$ is an approximation to the ridge estimator $\mathbf{w}_{\text{ridge},\lambda}^t$; $\widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k}$ are approximations to $\mathbf{A}_t^{-1} \mathbf{a}_{t,k}$; \widehat{v}_{tk} are approximations to $v_{tk} := \langle \widehat{\mathbf{w}}_{\text{ridge}}, \mathbf{a}_{t,k} \rangle + \alpha \sqrt{\langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle}$, which are also approximations to

$$v_{tk}^* := \langle \mathbf{w}_{\text{ridge},\lambda}^t, \mathbf{a}_{t,k} \rangle + \alpha \sqrt{\langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle}$$

for $k \in [A]$. After passing through the transformer, we obtain the policy

$$\text{Alg}_{\boldsymbol{\theta}}(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2t-1}^{\text{post},c})}{\|\exp(\mathbf{h}_{2t-1}^{\text{post},c})\|_1} \in \Delta^A.$$

We claim the following results which we will prove later.

Step 1 For any $\varepsilon > 0$, there exists a transformer $\text{TF}_{\boldsymbol{\theta}}(\cdot)$ with

$$L = \left\lceil 2\sqrt{2T} \sqrt{\frac{B_a^2 + \lambda}{\lambda}} \log \left(\frac{(2T(B_a^2 + \lambda) + \lambda)TB_a(B_a B_w + \sigma)}{\lambda^2 \varepsilon} \right) \right\rceil = \tilde{\mathcal{O}}(\sqrt{T}),$$

$$\max_{\ell \in [L]} M^{(\ell)} \leq 4, \quad \max_{\ell \in [L]} D^{(\ell)} \leq 4d, \quad \|\boldsymbol{\theta}\| \leq 10 + \frac{\lambda + 2}{B_a^2 + \lambda} = \mathcal{O}(1)$$

that implements step 1 in (17) with $\|\widehat{\mathbf{w}}_{\text{ridge}} - \mathbf{w}_{\text{ridge},\lambda}^t\|_2 \leq \varepsilon$.

Step 2 For any $\varepsilon > 0$, there exists a transformer $\text{TF}_{\boldsymbol{\theta}}(\cdot)$ with

$$L = \left\lceil 2\sqrt{2T} \sqrt{\frac{B_a^2 + \lambda}{\lambda}} \log \left(\frac{(2T(B_a^2 + \lambda) + \lambda)B_a}{\lambda^2 \varepsilon} \right) \right\rceil = \tilde{\mathcal{O}}(\sqrt{T}), \quad \max_{\ell \in [L]} M^{(\ell)} \leq 4A,$$

$$\max_{\ell \in [L]} D^{(\ell)} \leq 4dA, \quad \|\boldsymbol{\theta}\| \leq 10 + A\left(\frac{\lambda + 3}{B_a^2 + \lambda} + \sqrt{2}\right) = \mathcal{O}(A)$$

that implements step 2 in (17) with $\|\widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} - \mathbf{A}_t^{-1} \mathbf{a}_{t,k}\|_2 \leq \varepsilon$ for $k \in [A]$.

Step 3 Suppose that the approximation error in Step 2 satisfies $\varepsilon_2 \leq b_a^2/[2(B_a^2 + \lambda)TB_a]$. For any $\varepsilon > 0$, there exists a one-layer transformer $\text{TF}_\theta(\cdot)$ with

$$L = 2, \quad \max_{\ell \in [L]} M^{(\ell)} \leq 4A, \quad \max_{\ell \in [L]} D^{(\ell)} \leq \mathcal{O}(A\sqrt{T\alpha/(\tau\varepsilon)}), \quad \|\boldsymbol{\theta}\| \leq \mathcal{O}(A+T(\alpha/(\tau\varepsilon))^{1/4} + \alpha/\tau)$$

that implements step 3 in (17) with $|\widehat{v}_{tk}/\tau - v_{tk}/\tau| \leq \varepsilon$ for $k \in [A]$.

Denote the errors ε appear in each step by $\varepsilon_1, \varepsilon_2, \varepsilon_3$, respectively. Define for all $k \in [A]$ that

$$v_{tk}^* := \langle \mathbf{w}_{\text{ridge},\lambda}^t, \mathbf{a}_{t,k} \rangle + \alpha \sqrt{\langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle},$$

which are the actual values used to compare across different actions in LinUCB. Then for all $k \in [A]$, we have the approximation error

$$\begin{aligned} \left| \frac{v_{tk}^*}{\tau} - \frac{\widehat{v}_{tk}}{\tau} \right| &\leq \left| \frac{v_{tk}^*}{\tau} - \frac{v_{tk}}{\tau} \right| + \left| \frac{v_{tk}}{\tau} - \frac{\widehat{v}_{tk}}{\tau} \right| \\ &\leq \frac{1}{\tau} |\langle \mathbf{w}_{\text{ridge},\lambda}^t - \widehat{\mathbf{w}}_{\text{ridge}}, \mathbf{a}_{t,k} \rangle| + \frac{1}{\tau} \left| \alpha \sqrt{\langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle} - \alpha \sqrt{\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} \rangle} \right| + \varepsilon_3 \\ &\leq \frac{B_a \varepsilon_1}{\tau} + \frac{\alpha B_a \varepsilon_2}{2\tau \min \left\{ \sqrt{\langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle}, \sqrt{\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} \rangle} \right\}} + \varepsilon_3 \\ &\leq \frac{B_a \varepsilon_1}{\tau} + \frac{\sqrt{T(B_a^2 + \lambda)} \alpha B_a \varepsilon_2}{b_a \tau} + \varepsilon_3, \end{aligned}$$

where the last line uses Eq. (18). For a targeted approximation error ε , choosing $\varepsilon_1 = \varepsilon\tau/(12B_a)$, $\varepsilon_2 = \min\{b_a\tau\varepsilon/(12\sqrt{T(B_a^2 + \lambda)}\alpha B_a), b_a^2/[2(B_a^2 + \lambda)TB_a]\}$ and $\varepsilon_3 = \varepsilon/12$, we obtain $|v_{tk}^*/\tau - \widehat{v}_{tk}/\tau| \leq \varepsilon/2$ for all $k \in [A]$.

From the proof of each step, we can verify that the token dimension D can be chosen to be of order $\mathcal{O}(dA)$. Moreover, due to the convergence guarantee for each iteration of AGD in Proposition 9, it can be verified that there exists some sufficiently large value $R > 0$ with $\log R = \widetilde{\mathcal{O}}(1)$ such that we have $\|\mathbf{h}_i^{(\ell)}\|_2 \leq R$ for all layer $\ell \in [L]$ and all token $i \in [2T]$ in our TF construction. Therefore, TF_θ^R and TF_θ^∞ generate the same output for all the token matrices we consider, and w.l.o.g. we may assume in the proof of each step that the transformers we consider are those without truncation (i.e., $\text{TF}_\theta = \text{TF}_\theta^\infty$).

Finally, combining Step 1–3 with $\alpha = \widetilde{\mathcal{O}}(\sqrt{d})$ and applying Lemma 8 completes the proof of Theorem 12.

Proof of Step 1 We use the first d entries of \mathbf{h}_{2t-1}^d to represent $\widehat{\mathbf{w}}_{\text{ridge}}$ and the $d+1 \sim 3d$ entries (denoted by \star) to record intermediate results for computing $\widehat{\mathbf{w}}_{\text{ridge}}$. Step 1 follows immediately from the second part of Lemma 18.

Proof of Step 2 Note that

$$\mathbf{A}_t^{-1} \mathbf{a}_{t,k} = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2(2t-1)} \mathbf{x}^\top \mathbf{A}_t \mathbf{x} - \frac{1}{(2t-1)} \langle \mathbf{x}, \mathbf{a}_{t,k} \rangle =: \arg \min_{\mathbf{x} \in \mathbb{R}^d} L_k(\mathbf{x})$$

is the global minimizer of a $\lambda/(2t-1)$ -strongly convex and $(B_a^2 + \lambda)$ -smooth quadratic function with the condition number $\kappa \leq 2T(B_a^2 + \lambda)/\lambda$. Moreover, we have

$$\|\mathbf{A}_t^{-1} \mathbf{a}_{t,k}\|_2 \leq \|\mathbf{A}_t^{-1}\|_{\text{op}} \|\mathbf{a}_{t,k}\|_2 \leq B_a/\lambda.$$

It follows from Proposition 9 that $L = \lceil 2\sqrt{2T(B_a^2 + \lambda)/\lambda} \log((1 + \kappa)B_a/(\lambda\varepsilon)) \rceil$ steps of accelerated gradient descent finds $\widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k}$ with $\|\widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} - \mathbf{A}_t^{-1} \mathbf{a}_{t,k}\|_2 \leq \varepsilon$.

Similar to the proof of Lemma 18, we can construct a transformer such that each (self-attention+MLP) layer implements one step of the accelerated gradient descent (AGD) for all $k \in [A]$. Denote the $(k+2)d+1 \sim (k+3)d$, $(A+1+2k)d+1 \sim (A+2+2k)d$, $(A+2+2k)d+1 \sim (A+3+2k)d$ entries of \mathbf{h}_{2t-1}^d by $\widehat{\mathbf{w}}_{a,tk}$, $\widehat{\mathbf{w}}_{b,k}$, $\widehat{\mathbf{v}}_k$ for $k \in [A]$. Note that in the input vector $\mathbf{h}_{2t-1}^{\text{pre},d}$ we have $\widehat{\mathbf{w}}_{a,tk}^0$, $\widehat{\mathbf{w}}_{b,k}^0$, $\widehat{\mathbf{v}}_k^0 = \mathbf{0}_d$.

For each layer $\ell \in [L]$ and $k \in [A]$, we choose $\mathbf{Q}_{k1}^{(\ell)}$, $\mathbf{K}_{k1}^{(\ell)}$, $\mathbf{V}_{k1}^{(\ell)}$ such that for even tokens \mathbf{h}_{2j} with $j \leq t-1$ and odd tokens \mathbf{h}_{2j-1} with $j \leq t$

$$\begin{aligned} \mathbf{Q}_{k1}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} \widehat{\mathbf{v}}_k^{\ell-1} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = \begin{bmatrix} \mathbf{a}_j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \mathbf{0}, \quad \mathbf{V}_{k1}^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = -\eta \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_j \end{bmatrix}, \quad \mathbf{V}_{k1}^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \mathbf{0} \\ \mathbf{Q}_{k2}^{(\ell)} &= -\mathbf{Q}_{k1}^{(\ell)}, \quad \mathbf{K}_{k2}^{(\ell)} = \mathbf{K}_{k1}^{(\ell)}, \quad \mathbf{V}_{k2}^{(\ell)} = -\mathbf{V}_{k1}^{(\ell)}, \\ \mathbf{Q}_{k3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} 1 \\ 1-2t \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)} = \begin{bmatrix} 1 \\ 1 \\ 2j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)} = \begin{bmatrix} 1 \\ 1 \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} = \eta \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_{j,k} - \lambda \widehat{\mathbf{v}}_k^{\ell-1} \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{Q}_{k4}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} 2t-1 \\ -(2t-1)^2 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k4}^{(\ell)} \mathbf{h}_{2j}^{(\ell)} = \begin{bmatrix} 1 \\ 1 \\ (2j)^2 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k4}^{(\ell)} \mathbf{h}_{2j-1}^{(\ell)} = \begin{bmatrix} 1 \\ 1 \\ (2j-1)^2 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k4}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} = \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{v}}_k^{\ell-1} - \widehat{\mathbf{w}}_{a,k}^{\ell-1} \\ \mathbf{0} \end{bmatrix}, \end{aligned}$$

where $\eta = 1/(B_a^2 + \lambda)$ and the values $\mathbf{V}_{kt}^{(\ell)} \mathbf{h}_{2j}^{(\ell-1)}$, $\mathbf{V}_{kt}^{(\ell)} \mathbf{h}_{2j-1}^{(\ell-1)}$, $t = 1, 2, 3, 4$ are supported on the entries corresponding to $\widehat{\mathbf{w}}_{a,k}$. Summing up the $M = 4A$ heads and noting that $t = \sigma(t) - \sigma(-t)$, we see that the $\widehat{\mathbf{w}}_{a,k}$ part of \mathbf{h}_t has the update

$$\begin{aligned} \widehat{\mathbf{w}}_{a,k}^\ell &= \widehat{\mathbf{w}}_{a,k}^{\ell-1} - \frac{\eta}{2t-1} \sum_{j=1}^{t-1} [\sigma(\langle \mathbf{a}_j, \widehat{\mathbf{v}}_k^{\ell-1} \rangle) - \sigma(-\langle \mathbf{a}_j, \widehat{\mathbf{v}}_k^{\ell-1} \rangle)] \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \mathbf{V}_{k3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} + \mathbf{V}_{k4}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} \\ &= \widehat{\mathbf{w}}_{a,k}^{\ell-1} - \frac{\eta}{2t-1} \sum_{j=1}^{t-1} \langle \mathbf{a}_j, \widehat{\mathbf{v}}_k^{\ell-1} \rangle \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \mathbf{V}_{k3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} + \mathbf{V}_{k4}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} \\ &= \widehat{\mathbf{v}}_k^{\ell-1} - \frac{\eta}{2t-1} \sum_{j=1}^{t-1} \langle \mathbf{a}_j, \widehat{\mathbf{v}}_k^{\ell-1} \rangle \mathbf{a}_j - \frac{\eta\lambda}{2t-1} \widehat{\mathbf{v}}_k^{\ell-1} + \frac{\eta}{2t-1} \mathbf{a}_{t,k} \\ &= \widehat{\mathbf{v}}_k^{\ell-1} - \eta \nabla L(\mathbf{v}_k^{\ell-1}), \end{aligned}$$

which is one step of gradient descent with step size η (c.f. Eq. 16a). Moreover, it can be verified that one can choose the matrices such that $\max_{k \in [A], m \in [4]} \|\mathbf{Q}_{km}^{(\ell)}\|_{\text{op}} = \max_{k \in [A], m \in [4]} \|\mathbf{K}_{km}^{(\ell)}\|_{\text{op}} \leq \sqrt{2}$ and $\max_{k \in [A]} \|\mathbf{V}_{k1}^{(\ell)}\|_{\text{op}} = \max_{k \in [A]} \|\mathbf{V}_{k2}^{(\ell)}\|_{\text{op}} = \eta$, $\max_{k \in [A]} \|\mathbf{V}_{k3}^{(\ell)}\|_{\text{op}} \leq (\lambda + 1)\eta$, $\max_{k \in [A]} \|\mathbf{V}_{k4}^{(\ell)}\|_{\text{op}} \leq \sqrt{2}$. Therefore, the norm of the attention layer

$$\|\boldsymbol{\theta}\| \leq \sqrt{2}(A+1) + A(\lambda+3)/(B_a^2 + \lambda).$$

Following the construction as in the proof of Lemma 18, we can choose $\mathbf{W}_1^{(\ell)}$, $\mathbf{W}_2^{(\ell)}$ that implement Eq. (16b), (16b) for all $k \in [A]$ simultaneously and we also have $\|\mathbf{W}_1^{(\ell)}\|_{\text{op}} \leq 4\sqrt{2}$, $\mathbf{W}_2^{(\ell)} = \sqrt{2}$ with $D'^{(\ell)} = 4dA$. It follows from combining the bounds for the weight matrices that

$$\|\boldsymbol{\theta}\| \leq \sqrt{2}(A+1) + A\left(\frac{\lambda+3}{B_a^2 + \lambda} + \sqrt{2}\right) + \sqrt{2} + 4\sqrt{2} \leq 10 + A\left(\frac{\lambda+3}{B_a^2 + \lambda} + \sqrt{2}\right) = \mathcal{O}(A).$$

Proof of Step 3 Denote the i -th token of the output of step 2 (i.e., the input of step 3) by $\mathbf{h}_i^{(0)}$. We use the $(3A+3)d+1 \sim (3A+3)d+A$ entries of \mathbf{h}_{2t-1}^d to record $\widehat{v}_{t1}/\tau, \dots, \widehat{v}_{tA}/\tau$ and use the $(3A+3)d+A+1 \sim (3A+3)d+2A$ entries to store additional information (denoted by $v_{a,t1}, \dots, v_{a,tA}$) for computing $\widehat{v}_{t1}/\tau, \dots, \widehat{v}_{tA}/\tau$. Concretely, for all $k \in [A]$, we choose

$\mathbf{Q}_{k1,k2,k3,k4}^{(\ell)}, \mathbf{K}_{k1,k2,k3,k4}^{(\ell)}, \mathbf{V}_{k1,k2,k3,k4}^{(\ell)}$ such that for even tokens \mathbf{h}_{2j} with $j \leq t-1$ and odd tokens \mathbf{h}_{2j-1} with $j \leq t$

$$\begin{aligned} \mathbf{Q}_{k1}^{(1)} \mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} \widehat{\mathbf{w}}_{\text{ridge}} \\ 2t-1 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j-1}^{(0)} = \begin{bmatrix} \mathbf{a}_{j,k} \\ -\mathbf{B} \\ \mathbf{B}(2j-1) \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0}_d \\ -\mathbf{B} \\ 2\mathbf{B}j \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{V}_{k1}^{(1)} \mathbf{h}_{2j-1}^{(0)} &= \begin{bmatrix} \mathbf{0} \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0} \\ 2j \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{Q}_{k2}^{(1)} \mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} -\widehat{\mathbf{w}}_{\text{ridge}} \\ 2t-1 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k2}^{(1)} = \mathbf{K}_{k1}^{(1)}, \quad \mathbf{V}_{k2}^{(1)} = -\mathbf{V}_{k1}^{(1)}, \\ \mathbf{Q}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}} \\ 2t-1 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j-1}^{(0)} = \begin{bmatrix} \mathbf{a}_{j,k} \\ -\mathbf{B} \\ \mathbf{B}(2j-1) \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0}_d \\ -\mathbf{B} \\ 2\mathbf{B}j \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{V}_{k3}^{(1)} \mathbf{h}_{2j-1}^{(0)} &= \begin{bmatrix} \mathbf{0} \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k3}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0} \\ 2j \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{Q}_{k4}^{(1)} \mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} -\widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}} \\ 2t-1 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k4}^{(1)} = \mathbf{K}_{k3}^{(1)}, \quad \mathbf{V}_{k4}^{(1)} = -\mathbf{V}_{k3}^{(1)}, \end{aligned}$$

where $\mathbf{B} := TB_a^2(B_a B_w + \sigma)/\lambda + 2B_a^2/\lambda$; $\mathbf{V}_{k1}^{(1)} \mathbf{h}_c^{(0)}, \mathbf{V}_{k2}^{(1)} \mathbf{h}_c^{(0)}$ ($c = 2j-1, 2j$) are supported on the $[(3A+3)d+k]$ -th entry of \mathbf{h}_c^d ; $\mathbf{V}_{k3}^{(1)} \mathbf{h}_c^{(0)}, \mathbf{V}_{k4}^{(1)} \mathbf{h}_c^{(0)}$ ($c = 2j-1, 2j$) are supported on the $[(3A+3)d+A+k]$ -th entry of \mathbf{h}_c^d .

Since $\langle \widehat{\mathbf{w}}_{\text{ridge}}, \mathbf{a}_{j,k} \rangle \leq \|\widehat{\mathbf{w}}_{\text{ridge}}\|_2 \|\mathbf{a}_{j,k}\|_2 \leq \mathbf{B}$, it follows that

$$\left\langle \mathbf{Q}_{k1}^{(1)} \mathbf{h}_{2t-1}^{(0)}, \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j-1}^{(0)} \right\rangle = \langle \widehat{\mathbf{w}}_{\text{ridge}}, \mathbf{a}_{j,k} \rangle + (2j-1 - (2t-1))\mathbf{B} \leq 0$$

for $j < i$. Likewise $\left\langle \mathbf{Q}_{k1}^{(1)} \mathbf{h}_{2t-1}^{(0)}, \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)} \right\rangle \leq 0$ for $j < i$. Since we assume the error $\varepsilon_2 \leq b_a^2/[2(B_a^2 T + \lambda)B_a]$ in Step 2, $b_a \leq \|\mathbf{a}_{t,k}\|_2 \leq B_a$ and $\lambda \mathbf{I}_d \preceq \mathbf{A}_t \preceq (B_a^2 T + \lambda)\mathbf{I}_d$, it follows that

$$\begin{aligned} \left\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}} \right\rangle &\geq \langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle - \|\mathbf{a}_{t,k}\|_2 \|\mathbf{A}_t^{-1} \mathbf{a}_{t,k} - \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}}\|_2 \\ &\geq \frac{b_a^2}{2(B_a^2 T + \lambda)} \geq \frac{b_a^2}{2T(B_a^2 + \lambda)} =: \frac{1}{T} \cdot l_1, \end{aligned} \quad (18)$$

$$\left\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}} \right\rangle \leq \langle \mathbf{a}_{t,k}, \mathbf{A}_t^{-1} \mathbf{a}_{t,k} \rangle + \|\mathbf{a}_{t,k}\|_2 \|\mathbf{A}_t^{-1} \mathbf{a}_{t,k} - \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}}\|_2 \leq \frac{2B_a^2}{\lambda} =: l_2. \quad (19)$$

Therefore, $\left\langle \mathbf{Q}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)}, \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j-1}^{(0)} \right\rangle = \left\langle \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}}, \mathbf{a}_{j,k} \right\rangle + (2j-1 - (2t-1))\mathbf{B} \geq 0$ iff $j = i$. Likewise $\left\langle \mathbf{Q}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)}, \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j}^{(0)} \right\rangle \leq 0$ for $j < i$. Similar results hold for the $k2, k4$ -th heads. By some basic algebra and noting that $t = \sigma(t) - \sigma(-t)$ for $t \in \mathbb{R}$, we see that the attention layer updates the position for $\widehat{v}_{tk}/\tau, \widehat{v}_{a,tk}$ with the values $\langle \widehat{\mathbf{w}}_{\text{ridge}}, \mathbf{a}_{t,k} \rangle, \left\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}_t^{-1} \mathbf{a}_{t,k}} \right\rangle$ for all $k \in [A]$, respectively. Moreover, it can be verified that one can choose the matrices such that

$$\max_{k \in [A], m \in [4]} \|\mathbf{Q}_{km}^{(1)}\|_{\text{op}} = \max_{k \in [A], m \in [4]} \|\mathbf{V}_{km}^{(1)}\|_{\text{op}} = 1, \quad \max_{k \in [A], m \in [4]} \|\mathbf{K}_{km}^{(1)}\|_{\text{op}} \leq \mathbf{B}.$$

Now, to compute the value of \widehat{v}_{tk}/τ in step 3 in (17), what remains is to approximately compute $\alpha\sqrt{\widehat{v}_{a,tk}}$, add the result to the position for \widehat{v}_{tk}/τ , and multiplied it by $1/\tau$.

Since $\widehat{v}_{a,tk} = \langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} \rangle \in [l_1/T, l_2]$, to approximately compute $\sqrt{\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} \rangle}$, it suffices to approximate $f(x) = \sqrt{x}$, $x \in [l_1/T, l_2]$. For any level of approximation error $\varepsilon_{\text{approx}} > 0$, let $(x_1, x_2, \dots, x_N) \in [l_1/T, l_2]$ satisfy

$$x_1 = l_1/T, \quad x_N = l_2, \quad 0 \leq \sqrt{x_{j+1}} - \sqrt{x_j} \leq \varepsilon_{\text{approx}}, \quad \text{for } j \in [N-1].$$

Define the function

$$\tilde{f}(x) := \sqrt{x_1} + \sum_{j=1}^{N-1} \sigma(x - x_j) \frac{1}{\sqrt{x_{j+1}} + \sqrt{x_j}}.$$

Note that $\tilde{f}(x)$ is a piecewise linear function on $[l_1/T, l_2]$ with $\tilde{f}(x_i) = \sqrt{x_i}$ for $i \in [N]$. By some basic algebra, it can be shown that for $\varepsilon_{\text{approx}} < l_1/T$, the difference between $f(x)$ and $\tilde{f}(x)$

$$\max_{c \in [x_j, x_{j+1}]} |\tilde{f}(c) - f(c)| = \max_{t \in [0,1]} \left| \sqrt{x_j} + \frac{t}{\sqrt{x_j+1} + \sqrt{x_j}} - \sqrt{x_j + t(x_{j+1} - x_j)} \right| \leq \varepsilon_{\text{approx}}$$

when $\sqrt{x_{j+1}} - \sqrt{x_j} < c\sqrt{\varepsilon_{\text{approx}} l_1/T}$ for some universal constant $c > 0$ and all $j \in [N-1]$.

Therefore, there exists a function $\tilde{f}(x)$ with $N = \mathcal{O}(\sqrt{T}/\varepsilon_{\text{approx}})$ that satisfies

$$\max_{[l_1/T, l_2]} |\tilde{f}(x) - f(x)| \leq \varepsilon_{\text{approx}}.$$

As a consequence, we verify that one can implement $\tilde{f}(\widehat{v}_{a,tk})$ for all $k \in [A]$ simultaneously by constructing a two-layer MLP with

$$\|\mathbf{W}_1^{(1)}\|_{\text{op}} \leq \mathcal{O}(\sqrt{N}), \quad \|\mathbf{W}_2^{(1)}\|_{\text{op}} \leq \mathcal{O}(\sqrt{TN}), \quad D' \leq AN.$$

Choose $\varepsilon_{\text{approx}} = \tau\varepsilon/\alpha$. Substituting the expressions for $N, \varepsilon_{\text{approx}}$ into the upper bounds on the norms, we obtain

$$\|\mathbf{W}_1^{(1)}\|_{\text{op}} \leq \mathcal{O}((\alpha T/(\tau\varepsilon))^{1/4}), \quad \|\mathbf{W}_2^{(1)}\|_{\text{op}} \leq \mathcal{O}(T^{3/4}(\alpha/(\tau\varepsilon))^{1/4}), \quad D' \leq \mathcal{O}(A(\alpha T/(\tau\varepsilon))^{1/2}).$$

Lastly, we can construct another two-layer MLP with weights $\mathbf{W}_1^{(2)}, \mathbf{W}_2^{(2)}$ such that it implements the summation and multiplication updates

$$\widehat{\mathbf{v}} \leftarrow \widehat{\mathbf{v}} + \mathbf{W}_2^{(2)} \sigma(\mathbf{W}_1^{(2)} \mathbf{h}_{2t-1}^{(1)}) \approx \widehat{\mathbf{v}} + \left(\frac{1}{\tau} - 1\right) \widehat{\mathbf{v}} + \frac{\alpha}{\tau} \left[\sqrt{\widehat{v}_{a,t1}}, \dots, \sqrt{\widehat{v}_{a,tA}} \right]^\top = \frac{\widehat{v}_{tk}}{\tau}$$

with $\|\widehat{v}_{tk}/\tau - v_{tk}/\tau\| \leq \varepsilon$ for all $k \in [A]$. We verify that the weight matrices can be chosen with

$$\|\mathbf{W}_1^{(2)}\|_{\text{op}} \leq \mathcal{O}(1), \quad \|\mathbf{W}_2^{(2)}\|_{\text{op}} \leq \mathcal{O}(\alpha/\tau)$$

and $D' \leq \mathcal{O}(A)$.

Therefore the norm of the transformer that implements step 3 satisfies

$$\|\boldsymbol{\theta}\| \leq \mathcal{O}(B + 1 + 4A + T^{3/4}(\alpha/(\tau\varepsilon))^{1/4} + \alpha/\tau) = \mathcal{O}(A + T(\alpha/(\tau\varepsilon))^{1/4} + \alpha/\tau).$$

This concludes the proof of Step 3.

F.5 Proof of Theorem 13

By Theorem 1 and Theorem 12 with $\varepsilon_{\text{real}} = \varepsilon = 1/T^3$, it suffices to show soft LinUCB with parameter τ has the regret guarantee

$$\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\text{LinUCB}(\tau)}}(T) \right] \leq \mathcal{O}(d\sqrt{T} \log(T)).$$

This follows directly from a regret analysis similar to that for LinUCB (see e.g. [13] or Theorem 19.2 in [31]). Concretely, note that $v_{tk}^* = \langle \mathbf{w}_{\text{ridge}, \lambda}^t, \mathbf{a}_{t,k} \rangle + \alpha \sqrt{\langle \mathbf{a}_{t,k}, \widehat{\mathbf{A}}_t^{-1} \mathbf{a}_{t,k} \rangle}$ is the solution to the optimization problem

$$\text{maximize} \quad \langle \mathbf{w}, \mathbf{a}_{t,k} \rangle$$

subject to $\mathbf{w} \in \mathbb{C}_t := \{\mathbf{w} | (\mathbf{w} - \mathbf{w}_{\text{ridge},\lambda}^t)^\top \mathbf{A}_t (\mathbf{w} - \mathbf{w}_{\text{ridge},\lambda}^t) \leq \alpha^2\}$,

where we recall $\alpha = \alpha(\delta_0)$ with $\delta_0 = 1/(2B_a B_w T)$ and

$$\alpha = \alpha(\delta_0) := \sqrt{\lambda B_w + \sigma \sqrt{2 \log(1/\delta_0) + d \log((d\lambda + TB_a^2)/(d\lambda))}}. \quad (20)$$

Moreover, standard analysis as in the proof of Theorem 19.2 in [31] shows with probability over $1 - 1/(2B_a B_w T)$ we have $\mathbf{w}^* \in \mathbb{C}_t$ for all $t \in [T]$. Denote this event by \mathcal{E}_0 . Moreover, let $p_{t,k}$ denote the probability of soft LinUCB selecting the action $\mathbf{a}_{t,k}$ at time t for all $k \in [A]$. For any $\varepsilon > 0$, let $\mathbb{S}_t(\varepsilon) := \{k \in [A] : v_{tk}^* - \max_{j \in [A]} v_{tj}^* \leq \varepsilon\}$.

Therefore, on the event \mathcal{E}_0 at time t we have

$$\begin{aligned} \max_j v_{tj}^* - \sum_{k=1}^A p_{t,k} v_{tk}^* &= \sum_{k \in \mathbb{S}_t(\varepsilon_0)} p_{t,k} (\max_j v_{tj}^* - v_{tk}^*) + \sum_{k \notin \mathbb{S}_t(\varepsilon_0)} p_{t,k} (\max_j v_{tj}^* - v_{tk}^*) \\ &\leq \varepsilon_0 + \sum_{k \notin \mathbb{S}_t(\varepsilon_0)} \exp\left(-\frac{\varepsilon_0}{\tau}\right) (\max_j v_{tj}^* - v_{tk}^*) \\ &\leq \varepsilon_0 + 2A \exp\left(-\frac{\varepsilon_0}{\tau}\right) B_a (B_w + 2\alpha/\sqrt{\lambda}), \end{aligned}$$

where the second line uses

$$p_{t,k} \leq \exp\left(-\frac{\varepsilon_0}{\tau}\right) \cdot \max p_{t,k} \leq \exp\left(-\frac{\varepsilon_0}{\tau}\right),$$

and the last line follows from that $|v_{tj}^*| \leq B_a (B_w + 2\alpha/\sqrt{\lambda})$ on the event \mathcal{E}_0 . Choosing $\varepsilon_0 = \varepsilon_1/2 := 1/\sqrt{4T}$ and noting that $\tau = \varepsilon_0/\log(4TAB_a(B_w + 2\alpha/\sqrt{\lambda})) = \tilde{\mathcal{O}}(1/\sqrt{T})$, we obtain

$$\max_j v_{tj}^* - \sum_{k=1}^A p_{t,k} v_{tk}^* \leq \varepsilon_1.$$

Now, on the event \mathcal{E}_0 , we have

$$\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,j} \rangle \leq \max_{j \in [A]} v_{tj}^* \leq \sum_{k=1}^A p_{t,k} v_{tk}^* + \varepsilon_1 = \sum_{k=1}^A p_{t,k} \langle \tilde{\mathbf{w}}_k, \mathbf{a}_{t,k} \rangle + \varepsilon_1$$

for some $\tilde{\mathbf{w}}_k \in \mathbb{C}_t, k \in [A]$. Therefore, on \mathcal{E}_0 for each $t \in [T]$

$$\begin{aligned} \max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,j} \rangle - \sum_{k=1}^A p_{t,k} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle &\leq \varepsilon_1 + \sum_{k=1}^A p_{t,k} \langle \tilde{\mathbf{w}}_k - \mathbf{w}^*, \mathbf{a}_{t,k} \rangle \\ &\leq \varepsilon_1 + \sum_{k=1}^A p_{t,k} \|\tilde{\mathbf{w}}_k - \mathbf{w}^*\|_{\mathbf{A}_t} \cdot \|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}} \leq \varepsilon_1 + 2\alpha \mathbb{E}_{k \sim \mathbf{p}_t} \|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}}. \end{aligned}$$

Moreover, note that $\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,j} \rangle - \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle \leq 2B_w B_a$ and $\|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}} \leq B_a/\sqrt{\lambda}$. Summing over $t \in [T]$ and using the tower property of martingales, we obtain

$$\begin{aligned} &\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\text{LinUCB}(\tau)}}(T) \right] \\ &= \mathbb{E} \left[\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,j} \rangle - \sum_{k=1}^A p_{t,k} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle \right] \\ &\leq \mathbb{E} \left[2 \sum_{t=1}^T \alpha \mathbb{E}_{k \sim \mathbf{p}_t} \|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}} + \varepsilon_1 T + 2B_w B_a T \cdot \mathbf{1}_{\{\mathcal{E}_0^c\}} \right] \\ &\leq \mathbb{E} \left[2 \sum_{t=1}^T \alpha \left(\frac{B_a}{\sqrt{\lambda}} \wedge \|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}} \right) + \varepsilon_1 T + 2B_w B_a T \cdot \mathbf{1}_{\{\mathcal{E}_0^c\}} \right] \end{aligned}$$

$$\begin{aligned}
&\leq 2\mathbb{E}\left[\alpha\sqrt{T}\left(\frac{B_a}{\sqrt{\lambda}}+1\right)\sqrt{\sum_{t=1}^T(1\wedge\|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}}^2)+\varepsilon_1T}\right]+2B_wB_aT\mathbb{P}(\mathcal{E}_0^c) \\
&\leq\sqrt{8d(B_a/\sqrt{\lambda}+1)^2T\alpha^2\log((d\lambda+TB_a^2)/(d\lambda))+\varepsilon_1T}+1,
\end{aligned}$$

where the fourth line uses the fact that

$$\frac{B_a}{\sqrt{\lambda}}\wedge\|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}}\leq\left(\frac{B_a}{\sqrt{\lambda}}+1\right)\cdot(1\wedge\|\mathbf{a}_{t,k}\|_{\mathbf{A}_t^{-1}}^2)$$

and Cauchy-Schwartz inequality, the last line follows from Lemma 19.4 of [31]. Plugging in $\varepsilon_1 = 1/\sqrt{T}$ and Eq. (20) gives the upper bound on expected regret

$$\mathbb{E}_{M\sim\Lambda}\left[\sum_{t=1}^T\max_k\langle\mathbf{a}_{t,k},\mathbf{w}^*\rangle-\mathfrak{R}_{M,\text{Alg}_{\text{LinUCB}}(\tau)}(T)\right]\leq\mathcal{O}(d\sqrt{T}\log(T))$$

for soft LinUCB with parameter τ .

Moreover, the second part of Theorem 13 (i.e., the upper bound on $\log\mathcal{N}_\Theta$) follows directly from Lemma 6 and Eq. (13).

G Thompson sampling for stochastic linear bandit

Throughout this section, we use $c > 0$ to denote universal constants whose values may vary from line to line. Moreover, for notational simplicity, we use $\mathcal{O}(\cdot)$ to hide universal constants, $\tilde{\mathcal{O}}(\cdot)$ to hide polynomial terms in the problem parameters $(\lambda^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$, and $\tilde{\mathcal{O}}(\cdot)$ to hide both poly-logarithmic terms in $(M_0, C_0, T, A, d, 1/\varepsilon, 1/\delta_0)$ and polynomial terms in $(\lambda^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$. We also use the bold font $\mathbf{a}_t \in \mathbb{R}^d$ to denote the selected action vector \mathbf{a}_t at time $t \in [T]$.

This section is organized as follows. Section G.1 describes the Thompson sampling algorithm for stochastic linear bandits. Section G.2 introduces some additional definitions, assumptions, and the formal version of Theorem 14 as in Theorem 19. We prove Theorem 19 in Section G.3 and prove Theorem 15 in Section G.4. Lastly, the proof of Lemma 20 used in the proof of Theorem 19 is provided in Section G.5.

G.1 Thompson sampling algorithm

Consider the stochastic linear bandit setup as in Section E.1, but instead we assume a Gaussian prior distribution $\mathbf{w}^* \sim \mathcal{N}(0, \lambda\mathbf{I}_d)$ and Gaussian noises $\{\varepsilon_t\}_{t \geq 0} \sim_{iid} \mathcal{N}(0, r)$. Furthermore, we assume there exist (b_a, B_a) such that $b_a \leq \|\mathbf{a}_{t,k}\|_2 \leq B_a$. At each time $t \in [T]$, Thompson sampling consists of the following steps:

1. Computes

$$\boldsymbol{\mu}_t := \left(\frac{r}{\lambda}\mathbf{I}_d + \sum_{j=1}^{t-1}\mathbf{a}_j\mathbf{a}_j^\top\right)^{-1}\sum_{j=1}^{t-1}\mathbf{a}_j y_j, \quad \boldsymbol{\Sigma}_t := \frac{r}{\lambda}\mathbf{I}_d + \sum_{j=1}^{t-1}\mathbf{a}_j\mathbf{a}_j^\top.$$

2. Selects the action $\mathbf{a}_t = \mathbf{a}_{t,k}$ with probability

$$\mathbb{P}_{\tilde{\mathbf{w}}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, r\boldsymbol{\Sigma}_t^{-1})}\left(k = \arg\max_{j \in [A]}\langle\mathbf{a}_{t,j}, \tilde{\mathbf{w}}_t\rangle\right).$$

Note that Thompson sampling is equivalent to the posterior sampling procedure in our stochastic linear bandit setup, i.e., we select an action with probability that equals the posterior probability of the action being optimal. We allow λ to be either some constant independent of T, d , or has the form $\lambda = \lambda_0/d$ for some constant $\lambda_0 > 0$. The latter case is considered so that the bandit parameter vector \mathbf{w}^* has ℓ_2 norm of order $\tilde{\mathcal{O}}(1)$ with high probability. In this case, we use $\mathcal{O}(\cdot)$ to hide polynomial terms in the problem parameters $(\lambda_0^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$, and $\tilde{\mathcal{O}}(\cdot)$ to hide both poly-logarithmic terms in $(M_0, C_0, T, A, d, 1/\varepsilon, 1/\delta_0)$ and polynomial terms in $(\lambda_0^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$.

G.2 Definitions and assumptions

For any actions $\mathbf{a}_{t,1}, \dots, \mathbf{a}_{t,A} \in \mathbb{R}^d$, we define

$$f_k(\mathbf{a}_{t,1}, \dots, \mathbf{a}_{t,A}; \mu_t, r\Sigma_t^{-1}) := \log \mathbb{P}_{\tilde{\mathbf{w}}_t \sim \mathcal{N}(\mu_t, r\Sigma_t^{-1})} \left(k = \arg \max_{j \in [A]} \langle \mathbf{a}_{t,j}, \tilde{\mathbf{w}}_t \rangle \right).$$

For any $k \in [A]$, $\mathbf{x}_1, \dots, \mathbf{x}_A \in \mathbb{R}^d$, $y_1, \dots, y_A \in \mathbb{R}$, we introduce

$$g_k(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_A) := \log \mathbb{P}_{\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_d)} \left(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle + y_k - y_j \geq 0, \text{ for all } j \in [A] \right).$$

It can be verified that

$$f_k(\mathbf{a}_1, \dots, \mathbf{a}_A; \mu_t, r\Sigma_t^{-1}) = g_k(\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,1}, \dots, \sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,A}, \langle \mu_t, \mathbf{a}_{t,1} \rangle, \dots, \langle \mu_t, \mathbf{a}_{t,A} \rangle).$$

For any $\eta_1 \in [0, 1]$, we also define the truncated log-probabilities

$$f_{k,\eta_1}(\mathbf{a}_{t,1}, \dots, \mathbf{a}_{t,A}; \mu_t, r\Sigma_t^{-1}) := \log \left[\mathbb{P} \left(k = \arg \max_{j \in [A]} \langle \mathbf{a}_{t,j}, \tilde{\mathbf{w}}_t \rangle \right) \vee \eta_1 \right],$$

$$g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_A) := \log \left[\mathbb{P} \left(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle + y_k - y_j \geq 0, \text{ for all } j \in [A] \right) \vee \eta_1 \right].$$

Define in addition the region $D_\eta := \{\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_A : \|\mathbf{x}_i - \mathbf{x}_j\|_2 \geq \eta, \text{ for all } i \neq j\}$. We verify that on the set D_η , the function g_{k,η_1} is Lipschitz continuous in any of its arguments (see Lemma 20 for more).

We adopt the following definition in [6].

Definition 6 (Approximability by sum of relus). *A function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is $(\varepsilon, R, M_0, C_0)$ -approximable by sum of relus, if there exists a “ (M_0, C_0) -sum of relus” function*

$$f_{M_0, C_0}(\mathbf{z}) = \sum_{m=1}^{M_0} c_m \sigma(\mathbf{w}_m^\top [\mathbf{z}; 1]) \quad \text{with} \quad \sum_{m=1}^{M_0} |c_m| \leq C_0, \max_{m \in [M_0]} \|\mathbf{w}_m\|_1 \leq 1, \mathbf{w}_m \in \mathbb{R}^{d+1}, c_m \in \mathbb{R},$$

such that $\sup_{\mathbf{z} \in [-R, R]^d} |g(\mathbf{z}) - f_{M_0, C_0}(\mathbf{z})| \leq \varepsilon$.

Assumption B (Approximation of log-posterior probability). *There exist $M_0, C_0 > 0$ depending on $(1/\varepsilon, 1/\eta_1, 1/\eta_2, R_\delta, A)$ such that for any $\varepsilon > 0, \eta_1 \in (0, 1), \eta_2 > 0, \delta \in (0, 1/2)$ and $k \in [A]$, $g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_A)$ is $(\varepsilon, R_\delta, M_0, C_0)$ -approximable by sum of relus on D_{η_2} with $R_\delta := 2B_a \sqrt{\lambda}(1 + 2\sqrt{\log(2/\delta)} + \sqrt{d}) = \tilde{O}(\sqrt{\lambda d})$.*

Assumption B states that the (truncated) log-policy of Thompson sampling can be approximated via a two-layer MLP on a compact set with $\tilde{O}(\sqrt{d})$ -radius when $\lambda = \tilde{O}(1)$ (or with $\tilde{O}(1)$ -radius when $\lambda = \lambda_0/d = \tilde{O}(1/d)$).

Assumption C (Difference between the actions). *There exists some $\eta > 0$ such that for all instances M and any time $t \in [T]$, we have $\|\mathbf{a}_{t,j} - \mathbf{a}_{t,k}\|_2 \geq \eta$ for all $1 \leq j < k \leq A$.*

With the definitions and assumptions at hand, we now present the formal statement of Theorem 14 as in Theorem 19.

Theorem 19 (Approximating the Thompson sampling, Formal statement of Theorem 14). *For any $0 < \delta_0 < 1/2$, consider the same embedding mapping \mathbf{h} and extraction mapping \mathbf{A} as for soft LinUCB in F.1, and consider the standard concatenation operator cat . Under Assumption B, C, for $\varepsilon < (\eta \wedge 1)/4$, there exists a transformer $\text{TF}_\theta^R(\cdot)$ with $\log R = \tilde{O}(1)$,*

$$D = \tilde{O}(T^{1/4}Ad), L = \tilde{O}(\sqrt{T}), M = \tilde{O}(AT^{1/4}), D' = \tilde{O}(A(T^{1/4}d + M_0)), \\ \|\theta\| \leq \tilde{O}(T + AT^{1/4} + \sqrt{M_0A} + C_0), \quad (21)$$

such that with probability at least $1 - \delta_0$ over $(M, D_T) \sim \mathbb{P}_\Lambda^{\text{Alg}}$ for any Alg, we have

$$\log \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k} | D_{t-1}, s_t) - \log \text{Alg}_\theta(\mathbf{a}_{t,k} | D_{t-1}, s_t) \leq \varepsilon, \quad \text{for all } t \in [T], k \in [A].$$

Here M_0, C_0 are the values defined in Assumption B with $\eta_1 = \varepsilon/(4A), \eta_2 = \eta, \delta = \delta_0$, and $\tilde{O}(\cdot)$ hides polynomial terms in $(\lambda^{\pm 1}, r^{\pm 1}, b_a^{-1}, B_a)$ and poly-logarithmic terms in $(M_0, C_0, T, A, d, 1/\delta_0, 1/\varepsilon)$.

G.3 Proof of Theorem 19 (and hence Theorem 14)

We construct a transformer implementing the following steps at each time $t \in [T]$ starting with $\mathbf{h}_{2t-1}^* = \mathbf{h}_{2t-1}^{\text{pre}, \star}$ for $\star \in \{a, b, c, d\}$

$$\begin{aligned}
 \mathbf{h}_{2t-1} &= \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre}, a} \\ \mathbf{h}_{2t-1}^{\text{pre}, b} \\ \mathbf{h}_{2t-1}^{\text{pre}, c} \\ \mathbf{h}_{2t-1}^{\text{pre}, d} \end{bmatrix} \xrightarrow{\text{step 1}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre}, \{a, b, c\}} \\ \widehat{\mu}_t \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 2}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre}, \{a, b, c\}} \\ \widehat{\mu}_t \\ \Sigma_t^{1/2} \mathbf{a}_{t,1} \\ \vdots \\ \Sigma_t^{1/2} \mathbf{a}_{t,A} \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 3}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre}, \{a, b, c\}} \\ \widehat{\mu}_t \\ \sqrt{r} \Sigma_t^{-1/2} \mathbf{a}_{t,1} \\ \vdots \\ \sqrt{r} \Sigma_t^{-1/2} \mathbf{a}_{t,A} \\ \langle \widehat{\mu}_t, \mathbf{a}_{t,1} \rangle \\ \vdots \\ \langle \widehat{\mu}_t, \mathbf{a}_{t,A} \rangle \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \quad (22) \\
 &\xrightarrow{\text{step 4}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre}, \{a, b\}} \\ \widehat{v}_{t1} \\ \vdots \\ \widehat{v}_{tA} \\ \mathbf{h}_{2t-1}^d \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{post}, a} \\ \mathbf{h}_{2t-1}^{\text{post}, b} \\ \mathbf{h}_{2t-1}^{\text{post}, c} \\ \mathbf{h}_{2t-1}^{\text{post}, d} \end{bmatrix},
 \end{aligned}$$

where $\text{pos} := [t, t^2, 1]^\top$; μ_t, Σ_t are the mean and covariance of the distribution we sample $\tilde{\mathbf{w}}$ from; \widehat{v}_{tk} are approximations to $v_{tk}^* := \log \mathbb{P}(j = \arg \max_{k \in [A]} \langle \mathbf{a}_{t,k}, \tilde{\mathbf{w}}_t \rangle)$. In addition, we use $\mathbf{h}^*, \star \in \{a, b, c, d\}$ to denote the corresponding parts of a token vector \mathbf{h} . After passing through the transformer, we obtain the policy

$$\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2t-1}^{\text{post}, c})}{\|\exp(\mathbf{h}_{2t-1}^{\text{post}, c})\|_1} = \frac{\exp(\widehat{\mathbf{v}}_t)}{\|\exp(\widehat{\mathbf{v}}_t)\|_1} \in \Delta^A.$$

In step 1—3 of (22), we use transformer to approximately generate the arguments

$$(\sqrt{r} \Sigma_t^{-1/2} \mathbf{a}_{t,1}, \dots, \sqrt{r} \Sigma_t^{-1/2} \mathbf{a}_{t,A}, \langle \mu_t, \mathbf{a}_{t,1} \rangle, \dots, \langle \mu_t, \mathbf{a}_{t,A} \rangle)$$

of the function g_k (or g_{k, η_2}), and in step 4 of (22), we use transformer to approximate the truncated log-probability g_{k, η_1} for some $\eta_1 \in (0, 1)$ by exploiting Assumption B, C.

For any $0 < \delta_0 < 1/2$, define $B_w := \sqrt{\lambda}(\sqrt{d} + 2\sqrt{\log(2/\delta_0)})$ and the event

$$\mathcal{E}_{\delta_0} := \left\{ \max_{t \in [T]} |\varepsilon_t| \leq \sqrt{2r \log(2T/\delta_0)} \right\} \cup \{ \|\mathbf{w}^*\|_2 \leq B_w \}.$$

Then by a standard tail bound for gaussian variables $\{\varepsilon_t\}_{t=1}^T$, a union bound over $t \in [T]$, and Eq. (4.3) in [32], we have

$$\mathbb{P}(\mathcal{E}_{\delta_0}) \geq 1 - \delta_0.$$

We claim the following results which we will prove later.

Step 1 Under the high probability event \mathcal{E}_{δ_0} , for any $\varepsilon > 0$, there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = \left\lceil 2\sqrt{2T} \sqrt{\frac{B_a^2 + \tilde{\lambda}}{\tilde{\lambda}}} \log \left(\frac{(2T(B_a^2 + \tilde{\lambda}) + \tilde{\lambda})TB_a(B_a B_w + \sqrt{2r \log(2T/\delta_0)})}{\tilde{\lambda}^2 \varepsilon} \right) \right\rceil = \tilde{\mathcal{O}}(\sqrt{T}),$$

$$\max_{\ell \in [L]} M^{(\ell)} \leq 4, \quad \max_{\ell \in [L]} D^{(\ell)} \leq 4d, \quad \|\theta\| \leq 10 + (\tilde{\lambda} + 2)/(B_a^2 + \tilde{\lambda}),$$

where $\tilde{\lambda} := r/\lambda$ that implements step 1 in (22) with $\|\widehat{\mu}_t - \mu_t\|_2 \leq \varepsilon$.

Step 2 For any $\varepsilon > 0$, there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = \tilde{\mathcal{O}}(\sqrt{T}), \quad \max_{\ell \in [L]} M^{(\ell)} = \tilde{\mathcal{O}}(AT^{1/4}), \quad \max_{\ell \in [L]} D'^{(\ell)} = \tilde{\mathcal{O}}(T^{1/4}Ad), \quad \|\theta\| \leq \tilde{\mathcal{O}}(T+AT^{1/4})$$

that implements step 2 in (22) such that $\|\widehat{\Sigma}_t^{1/2} \mathbf{a}_{t,k} - \Sigma_t^{1/2} \mathbf{a}_{t,k}\|_2 \leq \varepsilon$ for all $k \in [A]$.

Step 3 Under the high probability event \mathcal{E}_{δ_0} , for any $\varepsilon > 0$, assume Step 1, 2 above are implemented with the approximation error less than $\varepsilon/B_a, \varepsilon\tilde{\lambda}/\sqrt{4r}$ respectively, then there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = \lceil 2 + 2\sqrt{2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda} \log((1 + \tilde{\kappa})4\sqrt{r}\sqrt{T(B_a^2 + \tilde{\lambda})B_a/\varepsilon})} \rceil = \tilde{\mathcal{O}}(\sqrt{T}), \quad \max_{\ell \in [L]} M^{(\ell)} = 4A, \\ \max_{\ell \in [L]} D'^{(\ell)} \leq 4Ad, \quad \|\theta\| \leq \tilde{\mathcal{O}}(T + A)$$

that implements step 3 in (22) with $\|\sqrt{r}\widehat{\Sigma}_t^{-1/2} \mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1/2} \mathbf{a}_{t,k}\|_2 \leq \varepsilon, |\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle - \langle \mu_t, \mathbf{a}_{t,k} \rangle| \leq \varepsilon$ for all $k \in [A]$.

Step 4 Under Assumption B, C and the high probability event \mathcal{E}_{δ_0} , suppose the approximation error ε_3 in Step 3 satisfies $\varepsilon_3 \leq R_{\delta_0}/2 = \tilde{\mathcal{O}}(\sqrt{\lambda d})$, and suppose the vector

$$\mathbf{v} := (\sqrt{r}\widehat{\Sigma}_t^{-1/2} \mathbf{a}_{t,1}, \dots, \sqrt{r}\widehat{\Sigma}_t^{-1/2} \mathbf{a}_{t,A}, \langle \widehat{\mu}_t, \mathbf{a}_{t,1} \rangle \dots \langle \widehat{\mu}_t, \mathbf{a}_{t,A} \rangle)$$

lies in $D_{\eta/2}$, for any $\varepsilon > 0$ there exists an MLP-only transformer $\text{TF}_\theta(\cdot)$ with

$$L = 1, \quad D' = M_0A, \quad \|\mathbf{W}_1\|_{\text{op}} \leq \sqrt{M_0A}, \quad \|\mathbf{W}_2\|_{\text{op}} \leq C_0$$

that implements step 4 in (22) such that $|\widehat{v}_{tk} - g_{k,\eta_1}(\mathbf{v})| \leq \varepsilon$ for all $k \in [A]$ and $\eta_1 = c\varepsilon/A$ for some universal constant $c > 0$.¹

To complete the proof, we in addition present the following lemma.

Lemma 20. For any $\eta_1 \in (0, 1), \eta > 0$, $g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_A)$ is 1/2-Holder continuous in its arguments on D_η , namely,

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_j, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{2A}{\eta_1} \left(\sqrt{\frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta}} + \frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta} \right), \\ & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_j, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y'_j, \dots, y_A)| \\ & \leq \frac{2A|y_j - y'_j|}{\eta\eta_1} \end{aligned}$$

for any

$$\begin{aligned} & (\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_A, y_1, \dots, y_A), (\mathbf{x}_1, \dots, \mathbf{x}'_j, \dots, \mathbf{x}_A, y_1, \dots, y_A) \in D_\eta, \\ & (\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_j, \dots, y_A), (\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y'_j, \dots, y_A) \in D_\eta \end{aligned}$$

for all $k, j \in [A]$.

See the proof in Section G.5.

Now, we complete the proof by combining Step 1—4 and using Lemma 20.

Let $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4$ denote the approximation errors ε appearing in Step 1, 2, 3, 4, respectively. W.l.o.g., we assume $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4 < 1/4 \wedge \eta/4$. Define the vector

$$\mathbf{v}^* := (\sqrt{r}\Sigma_t^{-1/2} \mathbf{a}_{t,1}, \dots, \sqrt{r}\Sigma_t^{-1/2} \mathbf{a}_{t,A}, \langle \mu_t, \mathbf{a}_{t,1} \rangle \dots \langle \mu_t, \mathbf{a}_{t,A} \rangle).$$

By Assumption C and a triangular inequality, we have $\mathbf{v}, \mathbf{v}^* \in D_{\eta/2}$. By the Lipschitz continuity of $f(x) = \exp(x)$ on $(-\infty, 1.5]$, we have

$$|\exp(\widehat{v}_{tk}) - \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t)| \leq |\exp(\widehat{v}_{tk}) - \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) \vee \eta_1| + \eta_1$$

¹Note M_0, C_0 in the formula implicitly depend on $1/\varepsilon$.

$$\begin{aligned}
&\leq e^{3/2}(|\widehat{v}_{tk} - g_{k,\eta_1}(\mathbf{v})| + |g_{k,\eta_1}(\mathbf{v}) - g_{k,\eta_1}(\mathbf{v}^*)|) + \eta_1, \\
&\leq e^{3/2}\left(\varepsilon_4 + \frac{2A^2}{\eta_1}\left(\sqrt{\frac{2\varepsilon_3}{\eta}} + \frac{2\varepsilon_3}{\eta}\right) + \frac{2A^2\varepsilon_3}{\eta\eta_1}\right) + \eta_1 =: \varepsilon_5,
\end{aligned}$$

where the second inequality uses

$$g_{k,\eta_1}(\mathbf{v}^*) = \log[\text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) \vee \eta_1],$$

and the third inequality uses Lemma 20 and Step 4. Therefore,

$$\left|\sum_{k=1}^A \exp(\widehat{v}_{t,k}) - 1\right| \leq A\varepsilon_5.$$

and the constructed transformer TF_θ satisfies (assume $A\varepsilon_5 < 1$)

$$\begin{aligned}
&\log \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) - \log \text{Alg}_\theta(\mathbf{a}_{t,k}|D_{t-1}, s_t) \\
&\leq (\log[\text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) \vee \eta_1] - \widehat{v}_{t,k}) + \log\left(\sum_{k=1}^A \exp(\widehat{v}_{t,k})\right) \\
&\leq |\widehat{v}_{tk} - g_{k,\eta_1}(\mathbf{v})| + |g_{k,\eta_1}(\mathbf{v}) - g_{k,\eta_1}(\mathbf{v}^*)| + A\varepsilon_5 \\
&\leq (A+1)\varepsilon_5,
\end{aligned}$$

where third line uses $\log(1+x) < x$. Finally, for the prespecified $\varepsilon > 0$, choosing $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \eta_1$ such that $\varepsilon_5 \leq \varepsilon/(2A)$ gives

$$\log \text{Alg}_{\text{TS}}(\mathbf{a}_{t,k}|D_{t-1}, s_t) - \log \text{Alg}_\theta(\mathbf{a}_{t,k}|D_{t-1}, s_t) \leq \varepsilon.$$

This can be done via choosing $\varepsilon_1 = c_1\varepsilon^4, \varepsilon_2 = c_2\varepsilon^4, \varepsilon_3 = c_3\varepsilon^4, \varepsilon_4 = c_4\varepsilon, \eta_1 = \varepsilon/(4A)$, where c_i ($i = 1, 2, 3, 4$) hide values that could depend polynomially on $(A, 1/\eta)$, such that $\varepsilon_5 \leq \varepsilon/(4A)$.

Combining the construction in Step 1—4 yields Theorem 19.

Similar to the proof of Theorem 12, from the proof of each step, we can verify that the token dimension D can be chosen to be of order $\tilde{O}(T^{1/4}Ad)$ (see the proof of Step 2b for details). Moreover, due to the convergence guarantee for each iteration of AGD in Proposition 9, we can be verified that there exists some sufficiently large value $R > 0$ with $\log R = \tilde{O}(1)$ such that we have $\|\mathbf{h}_i^{(\ell)}\|_2 \leq R$ for all layer $\ell \in [L]$ and all token $i \in [2T]$ in our TF construction. Therefore, TF_θ^R and TF_θ^∞ yield identical outputs for all token matrices considered, and hence we do not distinguish them in the proof of each step.

Proof of Step 1 Note that μ_t is a ridge estimator of \mathbf{w}_* with parameter $\tilde{\lambda} = r/\lambda$ and the noise $\sup_t |\varepsilon_t| \leq \sqrt{2r \log(T/\delta_0)}$. Step 1 follows immediately from the second part of Lemma 18.

Proof of Step 2 By the boundedness assumption of the actions, we have

$$\tilde{\lambda} \leq \sigma_{\min}(\boldsymbol{\Sigma}_t) \leq \sigma_{\max}(\boldsymbol{\Sigma}_t) \leq T(B_a^2 + \tilde{\lambda}).$$

Define the condition number $\tilde{\kappa} = T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}$ and $\mu := \sqrt{T\tilde{\lambda}(B_a^2 + \tilde{\lambda})}$. Using the Pade decomposition for the square root function in Theorem 3.1 and the discussion afterward in [38], we have

$$\begin{aligned}
\boldsymbol{\Sigma}_t^{1/2} &= (\tilde{\lambda}\mathbf{I}_d + \sum_{j=1}^{t-1} \mathbf{a}_j \mathbf{a}_j^\top)^{1/2} = \sqrt{\mu} \left(\mathbf{I}_d + \frac{(\boldsymbol{\Sigma}_t - \mu\mathbf{I}_d)}{\mu} \right)^{1/2} \\
&= \sqrt{\mu} \left[\mathbf{I}_d + \sum_{j=1}^m \left(\mathbf{I}_d + \frac{b_j^{(m)}(\boldsymbol{\Sigma}_t - \mu\mathbf{I}_d)}{\mu} \right)^{-1} \frac{a_j^{(m)}(\boldsymbol{\Sigma}_t - \mu\mathbf{I}_d)}{\mu} \right] + \mathbf{E}_m
\end{aligned}$$

for any $m \geq 0$, where

$$a_j^{(m)} = \frac{2}{2m+1} \sin^2 \frac{j\pi}{2m+1}, \quad b_j^{(m)} = \cos^2 \frac{j\pi}{2m+1},$$

and the error term \mathbf{E}_m satisfies

$$\begin{aligned}
& \|\mathbf{E}_m\|_{\text{op}} \\
& \leq \max \left\{ 2\sqrt{T(B_a^2 + \tilde{\lambda})} \left[1 + \left(\frac{\sqrt{T(B_a^2 + \tilde{\lambda})} + \sqrt{\mu}}{\sqrt{T(B_a^2 + \tilde{\lambda})} - \sqrt{\mu}} \right)^{2m+1} \right]^{-1}, 2\sqrt{\tilde{\lambda}} \left[\left(\frac{\sqrt{\mu} + \sqrt{\tilde{\lambda}}}{\sqrt{\mu} - \sqrt{\tilde{\lambda}}} \right)^{2m+1} - 1 \right]^{-1} \right\} \\
& = \max \left\{ 2\sqrt{T(B_a^2 + \tilde{\lambda})} [1 + (\frac{\tilde{\kappa}^{1/4} + 1}{\tilde{\kappa}^{1/4} - 1})^{2m+1}]^{-1}, 2\sqrt{\tilde{\lambda}} [(\frac{\tilde{\kappa}^{1/4} + 1}{\tilde{\kappa}^{1/4} - 1})^{2m+1} - 1]^{-1} \right\} \\
& \leq 2 \max \left\{ \sqrt{T(B_a^2 + \tilde{\lambda})} (1 + \frac{2}{\tilde{\kappa}^{1/4}})^{-2m-1}, \sqrt{\tilde{\lambda}} \left[(1 + \frac{2}{\tilde{\kappa}^{1/4}})^{2m+1} - 1 \right]^{-1} \right\}.
\end{aligned}$$

Since $(1 + 2/\tilde{\kappa}^{1/4})^{\tilde{\kappa}^{1/4}/2+1} > e$, it follows that choosing

$$m = \left(\frac{\tilde{\kappa}^{1/4}}{4} + 1 \right) \max \left\{ \left\lceil \log \left(\frac{2\sqrt{T(B_a^2 + \tilde{\lambda})}}{\varepsilon} \right) \right\rceil, \left\lceil \log \left(\frac{2\sqrt{\tilde{\lambda}}}{\varepsilon} + 1 \right) \right\rceil \right\} = \tilde{\mathcal{O}}(T^{1/4}).$$

gives $\|\mathbf{E}_m\|_{\text{op}} \leq \varepsilon$ for any $0 < \varepsilon < 1$.

Thus, using Pade decomposition, we can write

$$\Sigma_t^{1/2} \mathbf{a}_{t,k} = \sqrt{\mu} \left[\mathbf{a}_{t,k} + \sum_{j=1}^m \left(\mu \mathbf{I}_d + b_j^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \right)^{-1} a_j^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k} \right] + \mathbf{E}_m^k$$

with $\|\mathbf{E}_m^k\|_2 \leq \varepsilon$ for all $k \in [A]$ and some $m = \tilde{\mathcal{O}}(T^{1/4})$. Next, we show that there exists a transformer that can implement the following intermediate steps that give Step 2.

$$\begin{aligned}
& \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mu}_t \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 2a}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mu}_t \\ \mathbf{0}_{dA} \\ (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,1} \\ \vdots \\ (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,A} \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \\
& \xrightarrow{\text{step 2b}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mu}_t \\ \mathbf{0}_{dA} \\ (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,1} \\ \vdots \\ (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,A} \\ \left(\mu \mathbf{I}_d + b_1^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \right)^{-1} a_1^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,1} \\ \vdots \\ \left(\mu \mathbf{I}_d + b_m^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \right)^{-1} a_m^{(m)} (\Sigma_t - \mu \mathbf{I}_d) \mathbf{a}_{t,A} \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix} \xrightarrow{\text{step 2c}} \begin{bmatrix} \mathbf{h}_{2t-1}^{\text{pre},\{a,b,c\}} \\ \widehat{\mu}_t \\ \Sigma_t^{1/2} \mathbf{a}_{t,1} \\ \vdots \\ \Sigma_t^{1/2} \mathbf{a}_{t,A} \\ \star \\ \mathbf{0} \\ \text{pos} \end{bmatrix}, \quad (23)
\end{aligned}$$

where \star denotes additional terms in \mathbf{h}^d that are not of concern to our analysis.

Step 2a There exists an attention-only transformer $\text{TF}_\theta(\cdot)$ with

$$L = 2, \quad \max M^{(\ell)} = 3A, \quad \|\theta\| \leq T + 2 + \mu \leq \mathcal{O}(T)$$

that implements step 2a in (23).

Step 2b Denote $(\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k}$ by \mathbf{q}_k and $(\mu \mathbf{I}_d + b_m^{(m)} (\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d))$ by \mathbf{M}_m . For any $\varepsilon > 0$, there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = 2\sqrt{2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda} \log((1 + \tilde{\kappa})B_a T(B_a^2 + \tilde{\lambda})/(\tilde{\lambda}\varepsilon))} = \tilde{\mathcal{O}}(\sqrt{T}),$$

$$\max_{\ell \in [L]} M^{(\ell)} = 4Am = \tilde{\mathcal{O}}(T^{1/4}A), \quad \max_{\ell \in [L]} D'^{(\ell)} \leq \mathcal{O}(Adm) = \tilde{\mathcal{O}}(T^{1/4}Ad), \quad \|\boldsymbol{\theta}\| \leq \mathcal{O}(Am) \leq \tilde{\mathcal{O}}(T^{1/4}A)$$

approximately implements step 2b in (23) such that the output component $\widehat{\mathbf{M}_m^{-1} \mathbf{q}_k}$ satisfies $\|a_j^{(m)} \widehat{\mathbf{M}_m^{-1} \mathbf{q}_k} - a_j^{(m)} \mathbf{M}_j^{-1} \mathbf{q}_k\|_2 \leq \varepsilon$ for all $j \in [m]$ and $k \in [A]$.

Step 2c There exists an MLP-only transformer $\text{TF}_\theta(\cdot)$ with

$$L = 1, D' = 2Ad(m+1) = \tilde{\mathcal{O}}(T^{1/4}Ad), \quad \|\mathbf{W}_1\|_{\text{op}} = \sqrt{2}, \quad \|\mathbf{W}_2\|_{\text{op}} \leq \sqrt{\mu}(1+m) = \tilde{\mathcal{O}}(T^{3/4})$$

that implements step 2c in (23).

Combining the intermediate steps with the approximation error in Step 2b chosen as ε/m gives Step 2 as desired.

Proof of Step 2a For all $k \in [A]$, we choose $\mathbf{Q}_{k1}^{(1)}, \mathbf{K}_{k1,k2,k3}^{(1)}, \mathbf{V}_{k1,k2,k3}^{(1)}$ such that for even token indices $2j$ with $j \leq t-1$ and odd token indices with $j \leq t$

$$\mathbf{Q}_{k1}^{(1)} \mathbf{h}_{2t-1}^{(0)} = \begin{bmatrix} \mathbf{a}_{t,k} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{a}_j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_j \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)} \mathbf{h}_{2j-1}^{(0)} = [\mathbf{0}], \quad \mathbf{V}_{k1}^{(1)} \mathbf{h}_{2j-1}^{(0)} = [\mathbf{0}]$$

$$\mathbf{Q}_{k2}^{(1)} = -\mathbf{Q}_{k1}^{(1)}, \quad \mathbf{K}_{k2}^{(1)} = \mathbf{K}_{k1}^{(1)}, \quad \mathbf{V}_{k2}^{(1)} = -\mathbf{V}_{k1}^{(1)},$$

$$\mathbf{Q}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)} = \begin{bmatrix} 1 \\ -(2t-1) \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j-1}^{(0)} = \begin{bmatrix} 1 \\ 1 \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j}^{(0)} = \begin{bmatrix} 1 \\ 1 \\ 2j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)} = \begin{bmatrix} \mathbf{0} \\ (\tilde{\lambda} - \mu) \mathbf{a}_{t,k} \\ \mathbf{0} \end{bmatrix},$$

where for each $k \in [A]$, $\mathbf{V}_{k1}^{(1)} \mathbf{h}_{2j}^{(0)}, \mathbf{V}_{k3}^{(1)} \mathbf{h}_{2t-1}^{(0)}$ are supported on the same d entries of \mathbf{h}^d . It is readily verified that summing over the attention heads and $k \in [A]$ gives the updates

$$\mathbf{0}_d \mapsto \frac{(\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k}}{2t-1}$$

for all $k \in [A]$. We assume the updated vectors are supported on some Ad coordinates of \mathbf{h}_{2t-1}^d . Moreover, one can choose the matrices such that $\|\mathbf{Q}_{k1,k2,k3}^{(1)}\|_{\text{op}} \leq 1, \|\mathbf{K}_{k1,k2,k3}^{(1)}\|_{\text{op}} \leq 1, \|\mathbf{V}_{k1,k2,k3}^{(1)}\|_{\text{op}} \leq \max\{1, |\tilde{\lambda} - \mu|\} \leq 1 + \mu$. Therefore the norm of the first layer of the attention-only transformer $\|\boldsymbol{\theta}^{(1)}\| \leq 2 + \mu$.

The second layer is used to multiply the updated vectors by a factor of $2t-1$, namely, to perform the map

$$\frac{(\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k}}{2t-1} \mapsto (\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k}$$

for all $k \in [A]$, where the output vectors are supported on coordinates different from the input vectors (therefore we need $2Ad$ coordinates for embedding in step 2a). This can be achieved by choosing $\|\mathbf{Q}^{(2)}\|_{\text{op}} \leq T, \|\mathbf{K}^{(2)}\|_{\text{op}} \leq T, \|\mathbf{V}^{(2)}\|_{\text{op}} \leq 1$ such that

$$\mathbf{Q}_1^{(2)} \mathbf{h}_{2t-1}^{(1)} = \begin{bmatrix} (2t-1)^2 \\ -T(2t-1)^2 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_1^{(2)} \mathbf{h}_{2j-1}^{(1)} = \begin{bmatrix} 1 \\ 1 \\ T(2j-1)^2 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k3}^{(1)} \mathbf{h}_{2j}^{(1)} = \begin{bmatrix} 1 \\ 1 \\ T(2j)^2 \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{V}_1^{(2)} \mathbf{h}_{2t-1}^{(1)} = \begin{bmatrix} \mathbf{0} \\ \frac{(\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,1}}{(2t-1)} \\ \vdots \\ \frac{(\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,A}}{(2t-1)} \\ \mathbf{0} \end{bmatrix}.$$

Therefore $\|\boldsymbol{\theta}^{(2)}\| \leq T + 1$ and hence the two layer transformer we constructed has norm $\|\boldsymbol{\theta}\| \leq T + 2 + \mu$.

Proof of step 2b The construction is similar to the construction in Step 2 of the proof of Theorem 12. Hence we only provide a sketch of proof here. Note that

$$a_j^{(m)} \mathbf{M}_j^{-1} \mathbf{q}_k = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2(2t-1)} \mathbf{x}^\top \mathbf{M}_j \mathbf{x} - \frac{1}{2t-1} \langle \mathbf{x}, a_j^{(m)} \mathbf{q}_k \rangle =: \arg \min_{\mathbf{x} \in \mathbb{R}^d} L_{k,j}(\mathbf{x})$$

is the global minimizer of a $\tilde{\lambda}/(2t-1)$ -convex and $(B_a^2 + \tilde{\lambda})$ -smooth function with the conditional number $\tilde{\kappa} \leq 2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}$. Since

$$\|a_j^{(m)} \mathbf{M}_j^{-1} \mathbf{q}_k\|_2 \leq |a_j^{(m)}| \|\mathbf{M}_j^{-1}\|_{\text{op}} \|\mathbf{q}_k\|_2 \leq B_a T (B_a^2 + \tilde{\lambda}) / \tilde{\lambda}.$$

Therefore by Proposition 9 we have $L = \lceil 2\sqrt{2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}} \log((1 + \tilde{\kappa})B_a T (B_a^2 + \tilde{\lambda})/(\tilde{\lambda}\varepsilon)) \rceil$

steps of accelerated gradient descent with step size $\eta = 1/(B_a^2 + \tilde{\lambda})$ gives $\|a_j^{(m)} \widehat{\mathbf{M}}_j^{-1} \mathbf{q}_k - a_j^{(m)} \mathbf{M}_j^{-1} \mathbf{q}_k\|_2 \leq \varepsilon$. Now, it remains to construct a transformer that can implement the accelerated gradient descent steps. Here we only provide the construction of the gradient $\nabla L_{k,j}(\mathbf{x})$ at the l -th iteration $\mathbf{x} = \mathbf{x}_{k,j}^{\ell-1} \in \mathbb{R}^d$ which belongs to the output after $\ell-1$ transformer layers. The full construction of AGD steps follows from similar techniques as in Step 2 of the proof of Theorem 12. Concretely, for each layer $\ell \in [L]$ and $k \in [A], j \in [m]$, we choose $\mathbf{Q}_{kj1,kj2,kj3}^{(\ell)}, \mathbf{K}_{kj1,kj2,kj3}^{(\ell)}, \mathbf{V}_{kj1,kj2,kj3}^{(\ell)}$ such that for even token indices $2j$ with $s \leq t-1$ and odd token indices with $s \leq t$

$$\begin{aligned} \mathbf{Q}_{kj1}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} \mathbf{x}_{k,j}^{\ell-1} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{kj1}^{(\ell)} \mathbf{h}_{2s}^{(\ell-1)} = \begin{bmatrix} \mathbf{a}_s \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{kj1}^{(\ell)} \mathbf{h}_{2s-1}^{(\ell-1)} = \mathbf{0}, \\ \mathbf{V}_{kj1}^{(\ell)} \mathbf{h}_{2s}^{(\ell-1)} &= -\eta \begin{bmatrix} \mathbf{0} \\ b_j^{(m)} \mathbf{a}_s \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{kj1}^{(\ell)} \mathbf{h}_{2s-1}^{(\ell-1)} = \mathbf{0}, \\ \mathbf{Q}_{kj2}^{(\ell)} &= -\mathbf{Q}_{kj1}^{(\ell)}, \quad \mathbf{K}_{kj2}^{(\ell)} = \mathbf{K}_{kj1}^{(\ell)}, \quad \mathbf{V}_{kj2}^{(\ell)} = -\mathbf{V}_{kj1}^{(\ell)}, \\ \mathbf{Q}_{kj3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \begin{bmatrix} 1 \\ -(2t-1) \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{kj3}^{(\ell)} \mathbf{h}_{2s-1}^{(\ell)} = \begin{bmatrix} 1 \\ 1 \\ (2s-1) \\ \mathbf{0} \end{bmatrix}, \\ \mathbf{V}_{kj3}^{(\ell)} \mathbf{h}_{2t-1}^{(\ell-1)} &= \eta \begin{bmatrix} \mathbf{0} \\ a_j^{(m)} \mathbf{q}_k - [(1 - b_j^{(m)})\mu + b_j^{(m)}\tilde{\lambda}] \mathbf{x}_{k,j}^{\ell-1} \\ \mathbf{0} \end{bmatrix}. \end{aligned}$$

Similarly, it can be verified that the constructed attention layer generates

$$-\eta \cdot \nabla L_{k,j}(\mathbf{x}_{k,j}^{\ell-1}) = -\frac{\eta}{2t-1} [[(1 - b_j^{(m)})\mu \mathbf{I}_d + b_j^{(m)}\tilde{\lambda}] + b_j^{(m)} \sum_{s=1}^{t-1} \mathbf{a}_s \mathbf{a}_s^\top] \mathbf{x}_{k,j}^{\ell-1} + \frac{\eta a_j^{(m)} \mathbf{q}_k}{2t-1}.$$

Therefore, a construction similar to Step 2 of the proof of Theorem 12 yields Step 2b. Moreover, note that for the construction to exist we need the embedding dimension $D = O(Adm) = \tilde{O}(T^{1/4}Ad)$ and the number of hidden neurons $D' = O(Adm) = \tilde{O}(T^{1/4}Ad)$.

Proof of Step 2c Note that step 2c is a linear transformation from $\mathbf{a}_{t,k}, a_j^{(m)} \widehat{\mathbf{M}}_j^{-1} \mathbf{q}_k, k \in [A], j \in [m]$ to $\sqrt{\mu} \left[\mathbf{a}_{t,k} + \sum_{j=1}^m \left(\mu \mathbf{I}_d + b_j^{(m)} (\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \right)^{-1} a_j^{(m)} (\boldsymbol{\Sigma}_t - \mu \mathbf{I}_d) \mathbf{a}_{t,k} \right]$ and we have the fact $x = \sigma(x) - \sigma(-x)$. One can thus choose $\mathbf{W}_1 = \begin{bmatrix} \mathbf{I}_{A(m+1)d} & -\mathbf{I}_{A(m+1)d} & \mathbf{0} \end{bmatrix}$ with $D' = 2A(m+1)d$ and \mathbf{W}_2 with $\|\mathbf{W}_2\|_{\text{op}} \leq \sqrt{\mu}(1+m)$ that implements the linear map.

Proof of Step 3 Similar to the proof of Step 2b, given $\widehat{\boldsymbol{\Sigma}}_t^{1/2} \mathbf{a}_{t,k}$ we can approximate $\boldsymbol{\Sigma}_t^{-1} \widehat{\boldsymbol{\Sigma}}_t^{1/2} \mathbf{a}_{t,k} \approx \boldsymbol{\Sigma}_t^{-1/2} \mathbf{a}_{t,k}$ using accelerated gradient descent. Concretely, note that

$$\sqrt{r} \boldsymbol{\Sigma}_t^{-1} \widehat{\boldsymbol{\Sigma}}_t^{1/2} \mathbf{a}_{t,k} = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2(2t-1)} \mathbf{x}^\top \boldsymbol{\Sigma}_t \mathbf{x} - \frac{1}{2t-1} \langle \mathbf{x}, \sqrt{r} \widehat{\boldsymbol{\Sigma}}_t^{1/2} \mathbf{a}_{t,k} \rangle =: \arg \min_{\mathbf{x} \in \mathbb{R}^d} L_k(\mathbf{x})$$

is the global minimizer of a $\tilde{\lambda}/(2t-1)$ -convex and $(B_a^2 + \tilde{\lambda})$ -smooth function with the conditional number $\tilde{\kappa} \leq 2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}$. Since

$$\|\sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 \leq \sqrt{r}\|\mathbf{M}_j^{-1}\|_{\text{op}}\|\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 \leq 2\sqrt{r}\sqrt{T(B_a^2 + \tilde{\lambda})}B_a,$$

where the last inequality uses the assumption in Step 3. Therefore for any $\varepsilon_0 > 0$, it follows from Proposition 9 that $L = \lceil 2\sqrt{2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}}\log((1 + \tilde{\kappa})2\sqrt{r}\sqrt{T(B_a^2 + \tilde{\lambda})}B_a/\varepsilon_0) \rceil$ steps of accelerated gradient descent with step size $\eta = 1/(B_a^2 + \tilde{\lambda})$ gives

$$\|\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 \leq \varepsilon_0.$$

Following the construction in Step 2 of the proof of Theorem 12 it can be verified that there exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = \lceil 2\sqrt{2T(B_a^2 + \tilde{\lambda})/\tilde{\lambda}}\log((1 + \tilde{\kappa})2\sqrt{r}\sqrt{T(B_a^2 + \tilde{\lambda})}B_a/\varepsilon_0) \rceil = \tilde{\mathcal{O}}(\sqrt{T}), \quad \max_{\ell \in [L]} M^{(\ell)} = 4A,$$

$$\max_{\ell \in [L]} D^{(\ell)} = 4Ad, \quad \|\theta\| \leq \mathcal{O}(A)$$

that implements the AGD steps. Therefore, the approximation error

$$\begin{aligned} & \|\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 \\ & \leq \|\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 + \|\sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1}\Sigma_t^{1/2}\mathbf{a}_{t,k}\|_2 \\ & \leq \varepsilon_0 + \sqrt{r}\|\Sigma_t^{-1}\|_{\text{op}}\|\Sigma_t^{1/2}\mathbf{a}_{t,k} - \Sigma_t^{1/2}\mathbf{a}_{t,k}\|_2 \\ & \leq \varepsilon_0 + \frac{\sqrt{r}}{\tilde{\lambda}}\|\Sigma_t^{1/2}\mathbf{a}_{t,k} - \Sigma_t^{1/2}\mathbf{a}_{t,k}\|_2 \leq \varepsilon_0 + \varepsilon/2, \end{aligned}$$

where the last inequality uses the assumption on the approximation error in Step 2. Letting $\varepsilon_0 = \varepsilon/2$ yields $\|\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k} - \sqrt{r}\Sigma_t^{-1}\widehat{\Sigma}_t^{1/2}\mathbf{a}_{t,k}\|_2 \leq \varepsilon$.

In addition to the calculation of $\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k}$, we construct a two-layer attention-only transformer that computes $\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle$. Namely, we choose $\mathbf{Q}_{k1}^{(1)}, \mathbf{K}_{k1,k2}^{(1)}, \mathbf{V}_{k1,k2}^{(1)}$ such that

$$\begin{aligned} \mathbf{Q}_{k1}^{(1)}\mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} \widehat{\mu}_t \\ -(2t-1) \\ \mathbf{B} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)}\mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{B} \\ 2j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k1}^{(1)}\mathbf{h}_{2j-1}^{(0)} = \begin{bmatrix} \mathbf{a}_{j,k} \\ \mathbf{B} \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k1}^{(1)}\mathbf{h}_{2t-1}^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_k \\ \mathbf{0} \end{bmatrix} \\ \mathbf{Q}_{k2}^{(1)}\mathbf{h}_{2t-1}^{(0)} &= \begin{bmatrix} -\widehat{\mu}_t \\ -(2t-1) \\ \mathbf{B} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k2}^{(1)}\mathbf{h}_{2j}^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{B} \\ 2j \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_{k2}^{(1)}\mathbf{h}_{2j-1}^{(0)} = \begin{bmatrix} \mathbf{a}_{j,k} \\ \mathbf{B} \\ 2j-1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_{k2}^{(1)}\mathbf{h}_{2t-1}^{(0)} = -\begin{bmatrix} \mathbf{0} \\ \mathbf{e}_k \\ \mathbf{0} \end{bmatrix}, \end{aligned}$$

where $\mathbf{B} = 2TB_a^2(B_aB_w + \sqrt{2r\log(T/\delta_0)})/\tilde{\lambda} = \tilde{\mathcal{O}}(T)$ is an upper bound of $\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle$ for all $k \in [A]$ under the event \mathcal{E}_{δ_0} , and $\mathbf{e}_k = (0, 0, \dots, 1, 0, \dots, 0) \in \mathbb{R}^A$ is the one-hot vector supported on the k -th entry. Summing up the attention heads gives the update

$$\mathbf{0} \mapsto \frac{\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle}{2t-1}.$$

Note that one can choose the matrices such that

$$\|\mathbf{Q}_{k1,k2}^{(1)}\|_{\text{op}} \leq \mathbf{B}, \quad \|\mathbf{K}_{k1,k2}^{(1)}\|_{\text{op}} \leq \mathbf{B}, \quad \|\mathbf{V}_{k1,k2}^{(1)}\|_{\text{op}} \leq 1.$$

Thus the norm of the attention layer $\|\theta^{(1)}\| \leq \mathbf{B} + 2$.

Finally, as in the proof of Step 2a we can construct a single-layer single-head attention-only transformer with $\|\theta^{(2)}\| \leq T + 1$ that performs the multiplication

$$\frac{\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle}{2t-1} \mapsto \langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle.$$

To estimate the approximation error, note that $\|\mathbf{a}_{t,k}\|_2 \leq B_a$ and $\|\widehat{\mu}_t - \mu_t\|_2 \leq \varepsilon/B_a$ by our assumption in Step 3, it follows immediately that $|\langle \widehat{\mu}_t, \mathbf{a}_{t,k} \rangle - \langle \mu_t, \mathbf{a}_{t,k} \rangle| \leq \varepsilon$ for all $k \in [A]$. Combining the construction of the transformer layers above gives Step 3.

Proof of Step 4 By Assumption B, $g_{k,\eta_1}(\mathbf{v})$ are $(\varepsilon, R_{\delta_0}, \mathbf{M}_0, \mathbf{C}_0)$ -approximable by sum of relus for some $\mathbf{M}_0, \mathbf{C}_0$ depend polynomially on $(1/\varepsilon, 1/\eta_1, 1/\eta, 1/\delta_0, A)$. Since

$$\|\sqrt{r}\Sigma_t^{-1/2}\mathbf{a}_{t,k}\|_2 \leq \sqrt{r}\|\Sigma_t^{-1/2}\|_2\|\mathbf{a}_{t,k}\|_2 \leq \sqrt{\lambda}B_a, \quad |\langle \mu_t, \mathbf{a}_{t,k} \rangle| \leq \|\mu_t\|_2\|\mathbf{a}_{t,k}\|_2 = B_w B_a$$

and $R_{\delta_0} = 2(B_w B_a + \sqrt{\lambda}B_a)$, it follows from the assumption $\varepsilon_3 \leq R_{\delta_0}/2$ and a triangular inequality that $\|\mathbf{v}\|_\infty \leq R_\delta$. Therefore, using Assumption B and stacking up the approximation functions for each coordinate $k \in [A]$ we construct a two-layer MLP with $\|\mathbf{W}_1\|_{\text{op}} \leq \sqrt{A}\mathbf{M}_0, \|\mathbf{W}_1\|_{\text{op}} \leq \mathbf{C}_0, D' = \mathbf{M}_0 A$ such that

$$\mathbf{W}_2 \sigma(\mathbf{W}_1 \mathbf{h}_{2t-1}^{(1)}) = \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{v}}_{t1} \\ \vdots \\ \widehat{\mathbf{v}}_{tA} \\ \mathbf{0} \end{bmatrix}$$

where $(\widehat{\mathbf{v}}_{t1}, \dots, \widehat{\mathbf{v}}_{tA})$ is supported on \mathbf{h}_{2t-1}^c and $|\widehat{v}_{tk} - g_{k,\eta_1}(\mathbf{v})| \leq \varepsilon$ for all $k \in [A]$.

G.4 Proof of Theorem 15

Denote the transformer constructed in Theorem 19 by TF_θ . From the proof of Theorem 19, we have

$$\log \frac{\text{Alg}_{\text{TS}}(\bar{a}_t | D_{t-1}, s_t)}{\text{Alg}_\theta(\bar{a}_t | D_{t-1}, s_t)} \leq \varepsilon$$

under the event

$$\mathcal{E}_{\delta_0} := \left\{ \max_{t \in [T]} |\varepsilon_t| \leq \sqrt{2r \log(2T/\delta_0)} \right\} \cup \left\{ \|\mathbf{w}^*\|_2 \leq B_w \right\} \quad \text{for all } t \in [T]$$

with probability at least $1 - \delta_0$, where $B_w := \sqrt{\lambda}(\sqrt{d} + 2\sqrt{\log(2/\delta_0)})$. Note that due to the unboundedness of the noise ε_t and parameter vector \mathbf{w}^* , Assumption A may not be satisfied. However, setting $\delta_0 = \delta/(2n)$ and applying a union bound gives

$$\log \frac{\text{Alg}_{\text{TS}}(\bar{a}_t^i | D_{t-1}^i, s_t^i)}{\text{Alg}_\theta(\bar{a}_t^i | D_{t-1}^i, s_t^i)} \leq \varepsilon, \quad \text{for } t \in [T], i \in [n]. \quad (24)$$

with probability at least $1 - \delta/2$. From the proof of Theorem 1 we see that Assumption A is only used in Eq. (12) in the proof of Lemma 11. Moreover, it can be verified that the same result holds with Assumption A replaced by the condition in Eq. (24). Therefore, we have

$$\begin{aligned} \left| \mathfrak{R}_{\Lambda, \text{Alg}_{\widehat{\theta}}}(T) - \mathfrak{R}_{\Lambda, \text{Alg}_{\text{TS}}}(T) \right| &\leq cT^2 \sqrt{\mathcal{R}} \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} + \sqrt{\varepsilon_{\text{real}}} \right) \\ &\leq cT^2 \sqrt{\mathcal{R}} \left(\sqrt{\frac{\log[\mathcal{N}_\Theta \cdot T/\delta]}{n}} \right) + \sqrt{T} \end{aligned}$$

with probability at least $1 - \delta$ as in Theorem 1, where the second inequality follows as in our setting $\varepsilon_{\text{real}} = \varepsilon = 1/(T^3 \mathcal{R})$. Now, it suffices to show Thompson sampling has the expected regret with

$$\mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\text{TS}}}(T) \right] = \mathcal{O}(d\sqrt{T} \log(Td)).$$

The proof follows similar arguments as in Theorem 36.4 in [31]. Define

$$\tilde{\lambda} := r/\lambda, \quad \beta := \sqrt{r} \left(\sqrt{2rd \log(4d/\delta_{\text{TS}})} + \sqrt{2 \log(2/\delta_{\text{TS}}) + d \log(1 + TB_a^2/\tilde{\lambda}d)} \right),$$

where δ_{TS} will be specified later, and recall $\Sigma_t = \tilde{\lambda} \mathbf{I}_d + \sum_{j=1}^{t-1} \mathbf{a}_j \mathbf{a}_j^\top$. Since $\|\mathbf{w}^*\|_2 \leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})}$ with probability at least $1 - \delta_{\text{TS}}/2$ by a union bound, it follows from Theorem 20.5 in [31] that

$$\mathbb{P}(\|\mathbf{w}^*\|_2 \leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})}, \quad \text{and} \quad \|\mu_t - \mathbf{w}^*\|_{\Sigma_t} \geq \beta, \text{ for some } i \in [T]) \leq \delta_{\text{TS}},$$

where the probability is taken over the both randomness of the noise and of the bandit instance M .

Let \mathcal{E} be the event where $\|\mu_t - \mathbf{w}^*\|_{\Sigma_t} \leq \beta$ for all $i \in [T]$, and let \mathcal{E}_0 be the event where $\{\|\mathbf{w}^*\|_2 \leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})}\}$. Then $\mathbb{P}(\mathcal{E} \cap \mathcal{E}_0) \geq 1 - \delta_{\text{TS}}$ and the expected regret

$$\begin{aligned}
& \mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\text{TS}}}(T) \right] \\
&= \mathbb{E} \left[\sum_{t=1}^T \max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle \right] \\
&= \mathbb{E} \left[\sum_{t=1}^T (\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \right] + \mathbb{E} \left[\sum_{t=1}^T \max_{j \in [A]} (\langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{(\mathcal{E} \cap \mathcal{E}_0)^c} \right] \\
&\leq \mathbb{E} \left[\sum_{t=1}^T (\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \right] + \mathbb{E} [2B_a T \|\mathbf{w}^*\|_2 \mathbf{1}_{(\mathcal{E} \cap \mathcal{E}_0)^c}] \\
&\leq \mathbb{E} \left[\sum_{t=1}^T (\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \right] + 2B_a T \mathbb{E} [\|\mathbf{w}^*\|_2 \mathbf{1}_{(\mathcal{E} \cap \mathcal{E}_0)^c}].
\end{aligned}$$

Since

$$\begin{aligned}
\mathbb{E} [\|\mathbf{w}^*\|_2 \mathbf{1}_{(\mathcal{E} \cap \mathcal{E}_0)^c}] &\leq \mathbb{P}((\mathcal{E} \cap \mathcal{E}_0)^c) \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})} + \int_{\sqrt{2\lambda d \log(4d/\delta_{\text{TS}})}}^{\infty} \mathbb{P}(\|\mathbf{w}^*\|_2 \geq t) dt \\
&\leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})} \delta_{\text{TS}} + d^{3/2} \int_{\sqrt{2\lambda \log(4d/\delta_{\text{TS}})}}^{\infty} \mathbb{P}(|w_1| \geq t) dt \\
&\leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})} \delta_{\text{TS}} + 2\sqrt{2} d^{3/2} \lambda^{1/2} \int_{\sqrt{\log(4d/\delta_{\text{TS}})}}^{\infty} \exp(-t^2) dt \\
&\leq \sqrt{2\lambda d \log(4d/\delta_{\text{TS}})} \delta_{\text{TS}} + \sqrt{2} d^{3/2} \lambda^{1/2} \int_{\log(4d/\delta_{\text{TS}})}^{\infty} \frac{1}{t^{1/2}} \exp(-t) dt \\
&\leq 2\sqrt{2\lambda d \log(4d/\delta_{\text{TS}})} \delta_{\text{TS}},
\end{aligned}$$

where the second line follows from a union bound over $[d]$, and the third line uses properties of subgaussian variables. Therefore, choosing $\delta_{\text{TS}} = 1/[T\sqrt{d}]$ gives

$$\begin{aligned}
& \mathbb{E}_{M \sim \Lambda} \left[\sum_{t=1}^T \max_k \langle \mathbf{a}_{t,k}, \mathbf{w}^* \rangle - \mathfrak{R}_{M, \text{Alg}_{\text{TS}}}(T) \right] \\
&\leq \mathbb{E} \left[\sum_{t=1}^T (\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \right] + 6B_a \sqrt{\lambda \log(4d^2 T)}. \tag{25}
\end{aligned}$$

Now define the event $\mathcal{E}_t := \{\|\mu_t - \mathbf{w}^*\|_{\Sigma_t} \leq \beta\}$, then we have $\mathcal{E}_t \in \mathcal{F}_{t-1}$ and $\cap_{t=1}^T \mathcal{E}_t = \mathcal{E}$. Also, we define the upper confidence bound $U_t(\mathbf{a}) := \langle \mu_t, \mathbf{a} \rangle + \beta \|\mathbf{a}\|_{\Sigma_t^{-1}}$, which does not depend on the true parameter \mathbf{w}^* . Let $(\mathcal{F}_t)_{t \geq 0}$ denote the filtration generated by the data collected up to time t and the random parameter vector \mathbf{w}^* .

Let \mathbf{a}_t^* denote the optimal action at time t . Due to the construction of Thompson sampling, we have the distribution of \mathbf{a}_t^* and \mathbf{a}_t are the same conditioned on \mathcal{F}_{t-1} . Therefore, $\mathbb{E}[U_t(\mathbf{a}_t^*) | \mathcal{F}_{t-1}] = \mathbb{E}[U_t(\mathbf{a}_t) | \mathcal{F}_{t-1}]$ and

$$\begin{aligned}
\mathbb{E} \left[(\langle \mathbf{w}^*, \mathbf{a}_t^* \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \mid \mathcal{F}_{t-1} \right] &\leq \mathbb{E} \left[(\langle \mathbf{w}^*, \mathbf{a}_t^* \rangle - U_t(\mathbf{a}_t^*) + U_t(\mathbf{a}_t) - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E}_t} \mid \mathcal{F}_{t-1} \right] \\
&\leq \mathbb{E} \left[(U_t(\mathbf{a}_t) - \langle \mathbf{w}^*, \mathbf{a}_t \rangle) \mathbf{1}_{\mathcal{E}_t} \mid \mathcal{F}_{t-1} \right] \\
&\leq \mathbb{E} \left[(\|\mathbf{a}_t\|_{\Sigma_t^{-1}} \|\mu_t - \mathbf{w}^*\|_{\Sigma_t} + \beta \|\mathbf{a}_t\|_{\Sigma_t^{-1}}) \mathbf{1}_{\mathcal{E}_t} \mid \mathcal{F}_{t-1} \right] \\
&\leq 2\beta \mathbb{E} [\|\mathbf{a}_t\|_{\Sigma_t^{-1}} | \mathcal{F}_{t-1}].
\end{aligned}$$

Moreover, we have

$$\|\mathbf{a}_t\|_{\Sigma_t^{-1}} \leq B_a/\sqrt{\tilde{\lambda}}.$$

Combining the last two displays, we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(\max_{j \in [A]} \langle \mathbf{w}^*, \mathbf{a}_{t,k} \rangle - \langle \mathbf{w}^*, \mathbf{a}_t \rangle \right) \mathbf{1}_{\mathcal{E} \cap \mathcal{E}_0} \right] \\ & \leq \sum_{t=1}^T \mathbb{E} \left[2\beta \|\mathbf{a}_t\|_{\Sigma_t^{-1}} \wedge (B_a/\sqrt{\tilde{\lambda}}) \right] \\ & \leq 2(\beta \vee (B_a/\sqrt{\tilde{\lambda}})) \mathbb{E} \left[\sum_{t=1}^T (\|\mathbf{a}_t\|_{\Sigma_t^{-1}} \wedge 1) \right] \\ & \leq 2(\beta \vee (B_a/\sqrt{\tilde{\lambda}})) \sqrt{T} \sqrt{\mathbb{E} \left[\sum_{t=1}^T (\|\mathbf{a}_t\|_{\Sigma_t^{-1}}^2 \wedge 1) \right]} \\ & \leq 2(\beta \vee (B_a/\sqrt{\tilde{\lambda}})) \sqrt{T} \sqrt{2d \log(1 + TB_a^2/(\tilde{\lambda}d))} \\ & = \mathcal{O}(d\sqrt{T} \log(Td)), \end{aligned}$$

where the fourth line uses Cauchy-Schwartz inequality and the fifth line follows from Lemma 19.4 in [31]. Combining the last display with Eq. (25) completes the proof of first part of Theorem 15. Moreover, the second part of Theorem 15 (i.e., the upper bound on $\log \mathcal{N}_{\Theta}$) follows directly from Lemma 6 and Eq. (14).

G.5 Proof of Lemma 20

For any $j \neq k$, by definition of g_{k,η_1}

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_j, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{1}{\eta_1} \left| \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \geq 0, \text{ for all } i \in [A]) \right. \\ & \quad \left. - \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \geq 0, \text{ for all } i \neq j, \langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \geq 0) \right| \\ & \leq \frac{1}{\eta_1} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \geq 0 \geq \langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle + y_k - y_j) \right. \\ & \quad \left. + \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \leq 0 \leq \langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle + y_k - y_j) \right) \\ & \leq \frac{1}{\eta_1} \left(\mathbb{P}(\langle \mathbf{x}_j - \mathbf{x}'_j, \mathbf{z} \rangle \geq \langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \geq 0) \right. \\ & \quad \left. + \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \leq 0 \leq \langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle + y_k - y_j) \right) \\ & \leq \frac{1}{\eta_1} \left(\mathbb{P}(\langle \mathbf{x}_j - \mathbf{x}'_j, \mathbf{z} \rangle \geq \langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \geq 0) + \mathbb{P}(\langle \mathbf{x}_j - \mathbf{x}'_j, \mathbf{z} \rangle \leq \langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \leq 0) \right). \end{aligned}$$

Note that conditioned on $\mathbf{x}_j, \mathbf{x}'_j$ we have

$$\mathbb{P}(|\langle \mathbf{x}_j - \mathbf{x}'_j, \mathbf{z} \rangle| \leq \|\mathbf{x}_j - \mathbf{x}'_j\|_2 \sqrt{2 \log(2/\delta_1)}) \geq 1 - \delta_1$$

for any $\delta_1 > 0$. Therefore we further have

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_j, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{1}{\eta_1} \left[\mathbb{P} \left(\langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle + y_k - y_j \in [-\|\mathbf{x}_j - \mathbf{x}'_j\|_2 \sqrt{2 \log(2/\delta_1)}, \|\mathbf{x}_j - \mathbf{x}'_j\|_2 \sqrt{2 \log(2/\delta_1)}] \right) + \delta_1 \right] \\ & \leq \frac{1}{\eta_1} \left[\sup_{A \in \mathcal{F}, \mu(A) = 2\|\mathbf{x}_j - \mathbf{x}'_j\|_2 \sqrt{2 \log(2/\delta_1)}} \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_j, \mathbf{z} \rangle \in A) + \delta_1 \right] \end{aligned}$$

$$\leq \frac{1}{\eta_1} \left(\frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2 \sqrt{2\log(2/\delta_1)}}{\sqrt{2\pi}\eta} + \delta_1 \right) \leq \frac{1}{\eta_1} \left(\frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta\delta_1} + \delta_1 \right)$$

for any $\delta_1 > 0$, where the last inequality follows from the fact that standard Gaussian has probability density less than $1/\sqrt{2\pi}$ everywhere, $\|\mathbf{x}_k - \mathbf{x}'_j\|_2 \leq \eta$ and $\log(2/\delta_1) \leq 4/\delta_1^2$. Choosing $\delta_1 = 1 \wedge \sqrt{\frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta}}$ gives

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_j, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{2}{\eta_1} \left(\sqrt{\frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta}} + \frac{2\|\mathbf{x}_j - \mathbf{x}'_j\|_2}{\eta} \right) \end{aligned}$$

Similarly, for $\mathbf{x}_k \neq \mathbf{x}'_k$, we have

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_k, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_k, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{1}{\eta_1} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \geq 0 \geq \langle \mathbf{x}'_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i, \text{ for some } i \in [A]) \right. \\ & \quad \left. + \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \leq 0 \leq \langle \mathbf{x}'_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i, \text{ for some } i \in [A]) \right) \\ & \leq \sum_{i \neq k} \frac{1}{\eta_1} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \geq 0 \geq \langle \mathbf{x}'_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i) \right. \\ & \quad \left. + \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i \leq 0 \leq \langle \mathbf{x}'_k - \mathbf{x}_i, \mathbf{z} \rangle + y_k - y_i) \right) \\ & \leq \frac{A}{\eta_1} \max_{i \neq k} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_k, \mathbf{z} \rangle \geq \langle \mathbf{x}_i - \mathbf{x}'_k, \mathbf{z} \rangle + y_i - y_k \geq 0) + \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}'_k, \mathbf{z} \rangle \leq \langle \mathbf{x}_i - \mathbf{x}'_k, \mathbf{z} \rangle + y_i - y_k \leq 0) \right). \end{aligned}$$

Following the same argument, we have

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_k, \dots, \mathbf{x}_A, y_1, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}'_k, \dots, \mathbf{x}_A, y_1, \dots, y_A)| \\ & \leq \frac{2A}{\eta_1} \left(\sqrt{\frac{2\|\mathbf{x}_k - \mathbf{x}'_k\|_2}{\eta}} + \frac{2\|\mathbf{x}_k - \mathbf{x}'_k\|_2}{\eta} \right). \end{aligned}$$

Likewise, for any $j \neq k$ we have

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_j, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y'_j, \dots, y_A)| \\ & \leq \frac{1}{\eta} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle \in [\min\{y_k - y_j, y_k - y'_j\}, \max\{y_k - y_j, y_k - y'_j\}]) \right) \\ & \leq \frac{1}{\eta_1} \sup_{A \in \mathcal{F}, \mu(A) = 2|y_j - y'_j|} \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle \in A) \\ & \leq \frac{1}{\eta_1} \frac{2|y_j - y'_j|}{\sqrt{2\pi}\eta} \leq \frac{2|y_j - y'_j|}{\eta\eta_1} \end{aligned}$$

and

$$\begin{aligned} & |g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y_k, \dots, y_A) - g_{k,\eta_1}(\mathbf{x}_1, \dots, \mathbf{x}_A, y_1, \dots, y'_k, \dots, y_A)| \\ & \leq \sum_{j \neq k} \frac{1}{\eta} \left(\mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle \in [\min\{y'_k - y_j, y_k - y_j\}, \max\{y'_k - y_j, y_k - y_j\}]) \right) \\ & \leq \frac{A}{\eta_1} \sup_{A \in \mathcal{F}, \mu(A) = 2|y_j - y'_j|} \mathbb{P}(\langle \mathbf{x}_k - \mathbf{x}_j, \mathbf{z} \rangle \in A) \\ & \leq \frac{A}{\eta_1} \frac{2|y_j - y'_j|}{\sqrt{2\pi}\eta} \leq \frac{2A|y_j - y'_j|}{\eta\eta_1}. \end{aligned}$$

H Learning in-context RL in markov decision processes

Throughout this section, we use $c > 0$ to denote universal constants whose values may vary from line to line. Moreover, for notational simplicity, we use $O(\cdot)$ to hide universal constants, $\tilde{O}(\cdot)$ to hide poly-logarithmic terms in $(H, K, S, A, 1/\tau)$.

This section is organized as follows. Section H.1 discusses the embedding and extraction formats of transformers for Markov decision processes. Section H.2 describes the UCB-VI and the soft UCB-VI algorithms. We prove Theorem 16 in Section H.3 and prove Theorem 17 in Section H.4.

H.1 Embedding and extraction mappings

To embed MDP problems into transformers, we consider an embedding similar to that for linear bandits. For each episode $k \in [K]$, we construct $2H + 1$ tokens. Concretely, for each $t \in [T]$ in the k -th episode, we write $t = H(k - 1) + h$ and construct two tokens

$$\mathbf{h}_{2(t-1)+k} = \begin{bmatrix} \mathbf{0}_{A+1} \\ s_{k,h} \\ \mathbf{0}_A \\ \mathbf{0} \\ \mathbf{pos}_{2(t-1)+k} \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2(t-1)+k}^a \\ \mathbf{h}_{2(t-1)+k}^b \\ \mathbf{h}_{2(t-1)+k}^c \\ \mathbf{h}_{2(t-1)+k}^d \end{bmatrix}, \quad \mathbf{h}_{2t-1+k} = \begin{bmatrix} a_{k,h} \\ r_{k,h} \\ \mathbf{0}_S \\ \mathbf{0}_A \\ \mathbf{0} \\ \mathbf{pos}_{2t-1+k} \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2t-1+k}^a \\ \mathbf{h}_{2t-1+k}^b \\ \mathbf{h}_{2t-1+k}^c \\ \mathbf{h}_{2t-1+k}^d \end{bmatrix},$$

where $s_{k,h}, a_{k,h}$ are represented using one-hot embedding (we let $s_{k,H+1} = \mathbf{0}_S$), $\mathbf{h}_{2(t-1)+k}^c$ is used to store the (unnormalized) policy at time t given current state $s_{k,h}$, $\mathbf{0}$ in \mathbf{h}^d denotes an additional zero vector. At the end of each episode k , we add an empty token

$$\mathbf{h}_{(2H+1)k} = \mathbf{h}_k^{\text{emp}} := [\mathbf{0} \quad \mathbf{pos}_{(2H+1)k}]^\top$$

to store intermediate calculations. We also include in the tokens the positional embedding $\mathbf{pos}_i := (k, h, v_i, i, i^2, 1)^\top$ for $i \in [2T+K]$, where $v_i := \mathbf{1}_{\{\mathbf{h}_i^a = \mathbf{0}\}}$ denote the tokens that do not embed actions and rewards. In addition, we define the token matrix $\mathbf{H}_t := [\mathbf{h}_1, \dots, \mathbf{h}_{2t-1+k}] \in \mathbb{R}^{D \times (2t-1+k)}$ for all $t \in [T]$.

Offline pretraining Similar to the bandit setting, during pretraining the transformer TF_θ takes in $\mathbf{H}_T^{\text{pre}} := \mathbf{H}_T$ as the input token matrix, and generates $\mathbf{H}_T^{\text{post}} := \text{TF}_\theta(\mathbf{H}_T^{\text{pre}})$ as the output. For each time $t \in [T]$, we define the induced policy $\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2(t-1)+k}^{\text{post},c})}{\|\exp(\mathbf{h}_{2(t-1)+k}^{\text{post},c})\|_1} \in \Delta^A$, whose i -th entry is the probability of selecting the i -th action (denoted by the one-hot vector \mathbf{e}_i) given (D_{t-1}, s_t) . We then find the transformer $\hat{\theta} \in \Theta$ by solving Eq. (1).

Rollout At each time $t \in [T]$, given the current state s_t and previous data D_{t-1} , we first construct the token matrix $\mathbf{H}_{\text{roll},t}^{\text{pre}} \in \mathbb{R}^{D \times (2(t-1)+k)}$ that consists of tokens up to the first token for time t . The transformer then takes $\mathbf{H}_{\text{roll},t}^{\text{pre}}$ as the input and generates $\mathbf{H}_{\text{roll},t}^{\text{post}} = \text{TF}_\theta(\mathbf{H}_{\text{roll},t}^{\text{pre}})$. Next, the agent selects an action $a_t \in \mathcal{A}$ following the induced policy $\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2(t-1)+k}^{\text{post},c})}{\|\exp(\mathbf{h}_{2(t-1)+k}^{\text{post},c})\|_1} \in \Delta^A$ and observes the reward r_t and next state s_{t+1} ($s_{t+1} \sim \mu_1$ if t is the last time step in an episode).

Embedding and extraction mappings To integrate the above construction into our general framework in Section 2, for $t = (k - 1)H + h$, we have the embedding vectors

$$\mathbf{h}(s_t) := \mathbf{h}_{2(t-1)+k}, \quad \mathbf{h}(a_t, r_t) := \mathbf{h}_{2t-1+k}.$$

For $N \geq 1$, write

$$\lceil (N + 1)/2 \rceil = (k_N - 1)H + h_N$$

for some $h_N \in [H]$, and define the concatenation operator

$\text{cat}(\mathbf{h}_1, \dots, \mathbf{h}_N) := [\mathbf{h}_1, \dots, \mathbf{h}_{2H}, \mathbf{h}_1^{\text{emp}}, \mathbf{h}_{2H+1}, \dots, \mathbf{h}_{4H}, \mathbf{h}_2^{\text{emp}}, \mathbf{h}_{4H+1}, \dots, \mathbf{h}_N] \in \mathbb{R}^{N+k_N-1}$, where we insert an empty token $\mathbf{h}_k^{\text{emp}}$ (i.e., a token with $\mathbf{h}^{\{a,b,c\}} = \mathbf{0}$) at the end of each episode k .

In this case, we have the input token matrix

$$\mathbf{H} = \mathbf{H}_{\text{roll},t}^{\text{pre}} := \text{cat}(\mathbf{h}(s_1), \mathbf{h}(a_1, r_1), \dots, \mathbf{h}(a_{t-1}, r_{t-1}), \mathbf{h}(s_t)) \in \mathbb{R}^{D \times [2(t-1)+k]},$$

the output token matrix $\bar{\mathbf{H}} = \mathbf{H}_{\text{roll},t}^{\text{post}}$, and the linear extraction map \mathbf{A} satisfies

$$\mathbf{A} \cdot \bar{\mathbf{h}}_{-1} = \mathbf{A} \cdot \bar{\mathbf{h}}_{2(t-1)+k}^{\text{post}} = \mathbf{h}_{2(t-1)+k}^{\text{post},c}.$$

H.2 UCB-VI and soft UCB-VI

We show that transformers with the embedding in Section H.1 can approximately implement the UCB-VI algorithm in [4]. Namely, UCB-VI implements the following steps:

for each episode $k \in [K]$ and each step $h = H, \dots, 1$

1. Compute the estimated transition matrix $\widehat{P}_h(s'|s, a) := \frac{N_h(s, a, s')}{N_h(s, a)\sqrt{V_1}}$, where $N_h(s, a, s')$ denotes the number of times the state-action-next-state tuple (s, a, s') has been visited in the first $k - 1$ episodes, and $N_h(s, a) = \sum_{s'} N_h(s, a, s')$ (we assume $N_H(s, a, s') = 0$ and let $N_H(s, a)$ be the number of times (s, a) is visited at timestep H).
2. Calculate the estimated Q-function

$$\widehat{Q}_h(s, a) = \min\{H, r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \widehat{P}_h(s' | s, a) \widehat{V}_{h+1}(s')\},$$

where the bonus $b_h(s, a) = 2H \sqrt{\frac{\log(SAT/\delta)}{N_h(s, a)\sqrt{V_1}}}$, $\widehat{V}_{H+1}(s) := 0$ for all $s \in \mathcal{S}$ and $\widehat{V}_h(s) := \max_{a \in \mathcal{A}} \widehat{Q}_h(s, a)$.

Throughout this section, we choose the small probability $\delta = 1/(KH)$.

During policy execution, at each step $h \in [H]$, UCB-VI takes the greedy action $a_h := \arg \max_a \widehat{Q}_h(s_h, a)$ and observes the reward and next state (r_h, s_{h+1}) . To facilitate pretraining, in this work we consider a soft version of UCB-VI, which takes action a_h following the softmax policy

$$\pi_h(a|s_h) = \frac{\exp(\widehat{Q}_h(s_h, a)/\tau)}{\left\| \exp(\widehat{Q}_h(s_h, a)/\tau) \right\|_1}$$

using the estimated Q-function for some sufficiently small $\tau > 0$. Note that soft UCB-VI recovers UCB-VI as $\tau \rightarrow 0$.

H.3 Proof of Theorem 16

Throughout the proof, we abuse the notations \mathbf{h}_i^\star for $\star \in \{a, b, c, d\}$ to denote the corresponding positions in the token vector \mathbf{h}_i . For any $t' \in [T]$, we let $k(t'), h(t')$ be the non-negative integers such that $t' = H(k(t') - 1) + h(t')$ and $h(t') \in [H]$. For the current time t , we use the shorthands $k = k(t), h = h(t)$. For a token index $i \in [(2H + 1)K]$, let $\bar{k}(i), \bar{h}(i)$ be the episode and time step the i -th token corresponds to (for the empty tokens we set $h = H + 1$). Given the input token matrix $\mathbf{H}_{\text{roll}, t}^{\text{pre}}$, we construct a transformer that implements the following steps on the last token. $\mathbf{h}_{2(t-1)+k}^\star = \mathbf{h}_{2(t-1)+k}^{\text{pre}, \star}$ for $\star \in \{a, b, c, d\}$

$$\begin{array}{c}
 \left[\begin{array}{l} \mathbf{h}_{2(t-1)+k}^{\text{pre}, a} \\ \mathbf{h}_{2(t-1)+k}^{\text{pre}, b} \\ \mathbf{h}_{2(t-1)+k}^{\text{pre}, c} \\ \mathbf{h}_{2(t-1)+k}^{\text{pre}, d} \end{array} \right] \xrightarrow{\text{step 1}} \left[\begin{array}{l} \mathbf{h}_{2(t-1)+k}^{\text{pre}, \{a, b, c\}} \\ N_1(s, a, s') \\ \vdots \\ N_H(s, a, s') \\ N_1(s, a) \\ \vdots \\ N_H(s, a, s') \\ N_1(s, a)r_1(s, a) \\ \vdots \\ N_H(s, a, s')r_H(s, a) \\ \star \\ \mathbf{0} \\ \text{pos}_{2(t-1)+k} \end{array} \right] \xrightarrow{\text{step 2}} \left[\begin{array}{l} \mathbf{h}_{2(t-1)+k}^{\text{pre}, \{a, b, c\}} \\ \widehat{P}_1(s, a, s') \\ \vdots \\ \widehat{P}_H(s, a, s') \\ \star \\ \mathbf{0} \\ \text{pos}_{2(t-1)+k} \end{array} \right] \xrightarrow{\text{step 3}} \left[\begin{array}{l} \mathbf{h}_{2(t-1)+k}^{\text{pre}, \{a, b, c\}} \\ \widehat{Q}_1(s, a, s') \\ \vdots \\ \widehat{Q}_H(s, a, s') \\ \widehat{V}_1(s) \\ \vdots \\ \widehat{V}_H(s) \\ \star \\ \mathbf{0} \\ \text{pos}_{2(t-1)+k} \end{array} \right]
 \end{array}$$

$$\xrightarrow{\text{step 4}} \begin{bmatrix} \mathbf{h}_{2^{(t-1)+k}}^{\text{pre},\{a,b\}} \\ \frac{Q_h(s_t, a_1)}{\tau} \\ \vdots \\ \frac{\widehat{Q}_h(s_t, a_1)}{\tau} \\ \mathbf{h}_{2^{(t-1)+k}}^d \end{bmatrix} =: \begin{bmatrix} \mathbf{h}_{2^{(t-1)+k}}^{\text{post},a} \\ \mathbf{h}_{2^{(t-1)+k}}^{\text{post},b} \\ \mathbf{h}_{2^{(t-1)+k}}^{\text{post},c} \\ \mathbf{h}_{2^{(t-1)+k}}^{\text{post},d} \end{bmatrix}, \quad (26)$$

where $N_h(s, a, s'), \widehat{P}_h(s, a, s'), \widehat{Q}_h(s, a, s') \in \mathbb{R}^{S^2 \times A}$, $N_h(s, a) \in \mathbb{R}^{S \times A}$, $\widehat{V}_h(s) \in \mathbb{R}^S$ for all $h \in [H]$, and $*$ denote additional quantities in $\mathbf{h}_{2^{(t-1)+k}}^d$. Given the current state s_t , the transformer $\text{TF}_\theta(\cdot)$ generates the policy

$$\text{Alg}_\theta(\cdot | D_{t-1}, s_t) := \frac{\exp(\mathbf{h}_{2^{(t-1)+k}}^{\text{post},c})}{\|\exp(\mathbf{h}_{2^{(t-1)+k}}^{\text{post},c})\|_1} \in \Delta^A.$$

We claim the following results which we will prove later.

Step 1 There exists an attention-only transformer $\text{TF}_\theta(\cdot)$ with

$$L = 4, \quad \max_{\ell \in [L]} M^{(\ell)} \leq O(HS^2A), \quad \|\theta\| \leq O(HK + HS^2A)$$

that implements step 1 in (22).

Step 2 There exists a one-layer transformer $\text{TF}_\theta(\cdot)$ with

$$L = 1, \quad M \leq O(HS^2A), \quad D' \leq O(K^2HS^2A), \quad \|\theta\| \leq \tilde{O}(HS^2A + K^3 + KH)$$

that implements step 2 in (22).

Step 3 There exists a transformer $\text{TF}_\theta(\cdot)$ with

$$L = 2H, \quad \max_{\ell \in [L]} M^{(\ell)} \leq 2SA, \quad \max_{\ell \in [L]} D'^{(\ell)} \leq 3SA, \quad \|\theta\| \leq O(H + SA)$$

that implements step 3 (i.e., value iteration) in (22).

Step 4 There exists an attention-only transformer $\text{TF}_\theta(\cdot)$ with

$$L = 3, \quad \max_{\ell \in [L]} M^{(\ell)} = O(HA), \quad \|\theta\| \leq O(H(K + A) + 1/\tau)$$

that implements step 4 in (22).

From the construction of Step 1—4, we verify that one can choose the constructed transformer to have the embedding dimension $D = O(HS^2A)$. Moreover, due to the boundedness of the reward function, Q -function and the fact that the bonus $b(s, a) \leq \tilde{O}(H)$, we verify that there exists some $R > 0$ with $\log R = \tilde{O}(1)$ such that $\|\mathbf{h}_i^{(\ell)}\|_2 \leq R$ for all layer $\ell \in [L]$ and all token $i \in [K(2H + 1)]$. Therefore, similar to what we do in the proof of Theorem 12, 14, we may w.l.o.g. consider transformers without truncation (i.e., $R = \infty$) in our construction of step 1—4 in (26).

Proof of Step 1 We prove this step by constructing a transformer that implements the following two steps:

Step 1a For each $t' < t$ with $t' = (k' - 1)H + h'$ for some $h' \in [H]$, we add $s_{k',h'}, (a_{k',h'}, r_{k',h'})$ from $\mathbf{h}_{2^{(t'-1)+k'}}^b$ and $\mathbf{h}_{2^{t'-1}+k'}^d$ to $\mathbf{h}_{2^{t'}+k'}^d$.

Step 1b Compute $N_h(s, a, s'), N_h(s, a)$ for $h \in [H]$ and assign them to the current token $\mathbf{h}_{2^{(t-1)+k}}^d$.

For step 1a, we can construct a two-layer attention-only transformer with $\mathbf{Q}_{1,2,3}^{(1)}, \mathbf{K}_{1,2,3}^{(1)}, \mathbf{V}_{1,2,3}^{(1)}$ such that for all $i \leq 2(t-1) + k$

$$\mathbf{Q}_1^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \bar{k}(i) + 1 - v_i \\ \mathbf{B} \\ 1 \\ i \end{bmatrix}, \quad \mathbf{K}_1^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} -\mathbf{B} \\ \bar{k}(i) \\ i + 3 \\ -1 \end{bmatrix}, \quad \mathbf{V}_1^{(1)} \mathbf{h}_{2^{(t'-1)+k'}}^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0}_{A+1} \\ s_{k',h'} \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{V}_1^{(1)} \mathbf{h}_{2t'-1+k'}^{(0)} = \begin{bmatrix} \mathbf{0} \\ a_{k',h'} \\ r_{k',h'} \\ \mathbf{0}_S \\ \mathbf{0} \end{bmatrix}$$

where we choose $B = 4$ and $\mathbf{V} \mathbf{h}^{(0)}$ are supported on some entries in $\mathbf{h}^{(0),d}$. Moreover, we choose $\mathbf{Q}_3^{(1)} = \mathbf{Q}_2^{(1)} = \mathbf{Q}_1^{(1)}$, $\mathbf{V}_2^{(1)} = \mathbf{V}_3^{(1)} = -\mathbf{V}_1^{(1)}$ and $\mathbf{K}_2^{(1)}, \mathbf{K}_3^{(1)}$ such that

$$\mathbf{K}_2^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} -B \\ \bar{k}(i) \\ i+2 \\ -1 \end{bmatrix}, \quad \mathbf{K}_3^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} -B \\ \bar{k}(i) \\ i+1 \\ -1 \end{bmatrix}$$

We verify that $\|\mathbf{Q}_\star^{(1)}\|_{\text{op}}, \|\mathbf{K}_\star^{(1)}\|_{\text{op}} = 4, \|\mathbf{V}_\star^{(1)}\|_{\text{op}} = 1$ for $\star \in [3]$. Summing up the heads, we obtain the following update on a subset of coordinates in $\mathbf{h}_{2t'+k'}^{(0),d}$:

$$\begin{aligned} \mathbf{0}_{S+A+1} &\rightarrow \mathbf{0}_{S+A+1} + \sum_{j=1}^3 \sum_{i=1}^{2t'+k'} \sigma(\langle \mathbf{Q}_j^{(1)} \mathbf{h}_{2t'+k'}^{(0)}, \mathbf{K}_j^{(1)} \mathbf{h}_i^{(0)} \rangle) \mathbf{V}_j \mathbf{h}_i^{(0)} \\ &= \frac{1}{2t'+k'} [(\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'-2}^{(0)} + 2\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'-1}^{(0)} + 3\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'}^{(0)}) \\ &\quad - (\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'-1}^{(0)} + 2\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'}^{(0)} - \mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'}^{(0)})] \\ &= \frac{1}{2t'+k'} (\mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'-2}^{(0)} + \mathbf{V}_1^{(1)} \mathbf{h}_{2t'+k'-1}^{(0)}) \\ &= \frac{1}{2t'+k'} \begin{bmatrix} a_{k',h'} \\ r_{k',h'} \\ s_{k',h'} \end{bmatrix}. \end{aligned}$$

Note that $\langle \mathbf{Q}^{(1)} \mathbf{h}_i^{(0)}, \mathbf{K}^{(1)} \mathbf{h}_j^{(0)} \rangle \leq 0$ for $i = 2t' - 1 + k'$ (i.e., all tokens that embed the action and reward) since $v_i = 0$, it follows that no update happens on the tokens in which we embed the action and reward (i.e., the corresponding part of \mathbf{h}^d remains zero). Moreover, it should be noted that no update happens on tokens with $h = 1$.

We then use another attention layer to multiply the updated vectors by a factor of $2t' + k'$, namely, to perform the map

$$\frac{1}{2t'+k'} \begin{bmatrix} a_{k',h'} \\ r_{k',h'} \\ s_{k',h'} \end{bmatrix} \mapsto \begin{bmatrix} a_{k',h'} \\ r_{k',h'} \\ s_{k',h'} \end{bmatrix},$$

where the output vector is supported on coordinates different from the input vectors. This can be achieved by choosing $\|\mathbf{Q}_1^{(2)}\|_{\text{op}} \leq (2H+1)K, \|\mathbf{K}_1^{(2)}\|_{\text{op}} \leq (2H+1)K, \|\mathbf{V}_1^{(2)}\|_{\text{op}} \leq 1$ such that

$$\mathbf{Q}_1^{(2)} \mathbf{h}_i^{(1)} = \begin{bmatrix} i^2 \\ -(2H+1)Ki^2 \\ 1 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{K}_1^{(2)} \mathbf{h}_j^{(1)} = \begin{bmatrix} 1 \\ 1 \\ (2H+1)Kj^2 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{V}_1^{(2)} \mathbf{h}_{2t'+k'}^{(1)} = \frac{1}{2t'+k'} \begin{bmatrix} \mathbf{0} \\ a_{k',h'} \\ r_{k',h'} \\ s_{k',h'} \\ \mathbf{0} \end{bmatrix}, \quad (27)$$

and noting that $\langle \mathbf{Q}_1^{(2)} \mathbf{h}_i^{(1)}, \mathbf{Q}_1^{(2)} \mathbf{h}_j^{(1)} \rangle = i$ when $j = i$ and otherwise 0.

For step **1b**, we show that it can be implemented using a two-layer attention-only transformer.

To compute $N_h(s, a, s')$, in the first layer we construct $M = 10HS^2A$ heads with the query, key, value matrices $\{\mathbf{Q}_{ijkh,s}^{(1)}\}_{s=1}^{10}, \{\mathbf{K}_{ijkh,s}^{(1)}\}_{s=1}^{10}, \{\mathbf{V}_{ijkh,s}^{(1)}\}_{s=1}^{10}$ such that for all $i \leq 2(t-1) + k$ and

$i, k \in [S], j \in [A], h \in [H]$

$$\mathbf{Q}_{ijkh,1}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} B(v_i - 1) \\ B\mathbf{e}_i \\ B\mathbf{e}_j \\ B\mathbf{e}_k \\ 1 \\ 1 \\ h \end{bmatrix}, \quad \mathbf{K}_{ijkh,1}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ s_{\bar{k}(i), \bar{h}(i)} - 1 \\ a_{\bar{k}(i), \bar{h}(i)} - 1 \\ s_{\bar{k}(i), \bar{h}(i)} \\ -3B \\ 1 - \bar{h}(i) \\ 1 \end{bmatrix}, \quad \mathbf{V}_{ijkh,1}^{(1)} \mathbf{h}_i^{(0)} = - \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_{ijk}^{N_h} \\ \mathbf{0} \end{bmatrix},$$

where we choose $B = 2H$ and \mathbf{e}_{ijk}^h denotes the one-hot vector supported on the (i, j, k) -entry in $N_h(s, a, s')$. We similarly construct

$$\mathbf{Q}_{ijkh,2}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} B(v_i - 1) \\ B\mathbf{e}_i \\ B\mathbf{e}_j \\ B\mathbf{e}_k \\ 1 \\ 1 \\ h \end{bmatrix}, \quad \mathbf{K}_{ijkh,2}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ s_{\bar{k}(i), \bar{h}(i)} - 1 \\ a_{\bar{k}(i), \bar{h}(i)} - 1 \\ s_{\bar{k}(i), \bar{h}(i)} \\ -3B \\ -\bar{h}(i) \\ 1 \end{bmatrix}, \quad \mathbf{V}_{ijkh,2}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_{ijk}^{N_h} \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{Q}_{ijkh,3}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} B(v_i - 1) \\ B\mathbf{e}_i \\ B\mathbf{e}_j \\ B\mathbf{e}_k \\ 1 \\ 1 \\ -h \end{bmatrix}, \quad \mathbf{K}_{ijkh,3}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ s_{\bar{k}(i), \bar{h}(i)} - 1 \\ a_{\bar{k}(i), \bar{h}(i)} - 1 \\ s_{\bar{k}(i), \bar{h}(i)} \\ -3B \\ \bar{h}(i) - 1 \\ 1 \end{bmatrix}, \quad \mathbf{V}_{ijkh,3}^{(1)} \mathbf{h}_i^{(0)} = - \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_{ijk}^{N_h} \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{Q}_{ijkh,4}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} B(v_i - 1) \\ B\mathbf{e}_i \\ B\mathbf{e}_j \\ B\mathbf{e}_k \\ 1 \\ 1 \\ -h \end{bmatrix}, \quad \mathbf{K}_{ijkh,4}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ s_{\bar{k}(i), \bar{h}(i)} - 1 \\ a_{\bar{k}(i), \bar{h}(i)} - 1 \\ s_{\bar{k}(i), \bar{h}(i)} \\ -3B \\ \bar{h}(i) - 2 \\ 1 \end{bmatrix}, \quad \mathbf{V}_{ijkh,4}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_{ijk}^{N_h} \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{Q}_{ijkh,5}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} B(v_i - 1) \\ B\mathbf{e}_i \\ B\mathbf{e}_j \\ B\mathbf{e}_k \\ 1 \\ 1 \\ 1 \end{bmatrix}, \quad \mathbf{K}_{ijkh,5}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ s_{\bar{k}(i), \bar{h}(i)} - 1 \\ a_{\bar{k}(i), \bar{h}(i)} - 1 \\ s_{\bar{k}(i), \bar{h}(i)} \\ -3B \\ -\bar{k}(i) \\ 1 \end{bmatrix}, \quad \mathbf{V}_{ijkh,5}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{0} \\ \mathbf{e}_{ijk}^{N_h} \\ \mathbf{0} \end{bmatrix}.$$

Summing up the first five heads, we verify that such attention updates the token with $\mathbf{h}_i^a = \mathbf{0}$ and has the form

$$\mathbf{0} \rightarrow \mathbf{0} + \frac{1}{i} \tilde{N}_h(i, j, k) \mathbf{e}_{ijk}^{N_h}$$

on \mathbf{h}_i^d , where $\tilde{N}_h(i, j, k)$ denote the number of visits to the state-action-next-state tuple (i, j, k) at time step h before token i . For $\star \in [5]$, we choose $\mathbf{V}_{ijkh, \star+5}^{(1)} = -\mathbf{V}_{ijkh, \star+5}^{(1)}$ and $\mathbf{Q}_{ijkh, \star+5}^{(1)}, \mathbf{K}_{ijkh, \star+5}^{(1)}$ be such that

$$\mathbf{Q}_{ijkh, \star+5}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{Q}_{ijkh, \star}^{(1)} \mathbf{h}_i^{(0)} \\ B \\ -\bar{k}(i) \end{bmatrix}, \quad \mathbf{K}_{ijkh, \star+5}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{K}_{ijkh, \star}^{(1)} \mathbf{h}_i^{(0)} \\ \bar{k}(i) \\ B \end{bmatrix}$$

which adds positional embedding about the current episode $\bar{k}(i)$. We verify that summing up the sixth to the tenth heads gives the update

$$\mathbf{0} \rightarrow \mathbf{0} + \frac{1}{i} (N_h(i, j, k) - \tilde{N}_h(i, j, k)) \mathbf{e}_{ijk}^{N_h}$$

on \mathbf{h}_i^d for $i \leq 2(t-1) + k$ with $\mathbf{h}_i^a = \mathbf{0}$. Therefore, combining all the heads together we have the update

$$\mathbf{0} \rightarrow \mathbf{0} + \frac{1}{i} N_h(i, j, k) \mathbf{e}_{ijk}^{N_h} \quad \text{for all } i, k \in [S], j \in [A], h \in [H]$$

on \mathbf{h}_i^d for $i \leq 2(t-1) + k$ with $\mathbf{h}_i^a = \mathbf{0}$, in particular when $i = 2(t-1) + k$. Moreover, notice that the matrices $\{\mathbf{Q}_{ijkh,s}^{(1)}\}_{s=1}^{10}$, $\{\mathbf{K}_{ijkh,s}^{(1)}\}_{s=1}^{10}$ can be constructed with the operator norm less than $10B = 10H$, and $\{\mathbf{V}_{ijkh,s}^{(1)}\}_{s=1}^{10}$ with the operator norm equals 1.

Following a similar construction, we can also compute $N_h(s, a)$, $N_h(s, a)r_h(s, a)$ for all h, s, a, s' on different supports of coordinates in \mathbf{h}_i^d via adding additional $M = O(HSA)$ heads to the attention-only layer.

Next, we construct the second attention layer to multiply the token vector by the index number i as in the proof of Step 1a. The construction is similar to that in Eq. (27) and we omit it here. Moreover, note that Step 1b can be implemented with the embedding dimension $D \leq O(HS^2A)$ as we need $O(1)$ dimensions for each quadruple (i, j, k, h) . Combining Step 1a, 1b concludes the proof of Step 1.

Proof of Step 2 After Step 1, for the current token $i = 2(t-1) + k$, we have $N_h(s, a, s')$, $r_h(s, a)$, $N_h(s, a)$, $N_h(s, a)r_h(s, a)$ lie in \mathbf{h}_i^d for all $h \in [H]$. Given these vectors that store the number of visits and rewards, note that

$$\begin{aligned} r_h(s, a) &= \frac{N_h(s, a)r_h(s, a)}{N_h(s, a) \vee 1}, \quad \text{when } N_h(s, a) \geq 1, \\ b_h(s, a) &= 2H \sqrt{\frac{\log(SAT/\delta)}{N_h(s, a) \vee 1}}, \\ \widehat{P}_h(s, a, s') &= \frac{N_h(s, a, s')}{N_h(s, a) \vee 1}. \end{aligned}$$

Therefore, we may compute \widehat{P}_h, b_h via using a transformer layer to implement the functions $f_1(x, y) = \frac{x}{y \vee 1}$, $f_2(y) = 2H \sqrt{\frac{\log(SAT/\delta)}{y \vee 1}}$, $f_3(x, y) = \frac{x}{y \vee 1} + H\mathbf{1}_{y=0}$ for $x, y \in \{0\} \cup [K]$. We demonstrate the computation of $\widehat{P}_h(s, a, s')$ (i.e., the computation of $f_1(x, y)$) here. We start with constructing an attention layer with $M = O(HS^2A)$ heads such that it implements $x \mapsto x^2$ for $x = N_h(s, a, s')$, $N_h(s, a)$. For $N_h(s, a, s')$, this can be done by choosing $\|\mathbf{Q}_{ijkh}^{(1)}\|_{\text{op}} \leq K$, $\|\mathbf{K}_{ijkh}^{(1)}\|_{\text{op}} \leq K$, $\|\mathbf{V}_{ijkh}^{(1)}\|_{\text{op}} = 1$ such that

$$\mathbf{Q}_{ijkh}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} K \\ -i \\ N_h(\mathbf{e}_i, \mathbf{e}_j, \mathbf{e}_k) \end{bmatrix}, \quad \mathbf{K}_{ijkh}^{(1)} \mathbf{h}_j^{(0)} = \begin{bmatrix} j \\ K \\ N_h(\mathbf{e}_i, \mathbf{e}_j, \mathbf{e}_k) \end{bmatrix}, \quad \mathbf{V}_{ijkh}^{(1)} \mathbf{h}_j^{(0)} = \begin{bmatrix} \mathbf{0} \\ j \\ \mathbf{0} \end{bmatrix},$$

where $\mathbf{e}_i, \mathbf{e}_k$ denote the i, j -th states and $\mathbf{e}_{s,j}$ denotes the k -th action. Similarly, we can construct HSA additional heads to compute $N_h(s, a)^2$ for all possible s, a .

Next, we compute the exact values of $\widehat{P}(s, a, s')$ using an MLP layer. Namely, we construct $\mathbf{W}_1^{(1)} = \mathbf{W}_{12}^{(1)} \mathbf{W}_{11}^{(1)}$, $\mathbf{W}_2^{(1)} = \mathbf{W}_{23}^{(1)} \mathbf{W}_{22}^{(1)} \mathbf{W}_{21}^{(1)}$ such that for all h, s, a, s' , on the corresponding vector component we have

$$\mathbf{W}_{11}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 \\ N_h(s, a, s')^2 \\ \vdots \\ (N_h(s, a, s') - K)^2 \\ N_h(s, a)^2 \\ \vdots \\ (N_h(s, a) - K)^2 \end{bmatrix} = \begin{bmatrix} 1 \\ N_h(s, a, s')^2 \\ \vdots \\ N_h(s, a, s')^2 + K^2 - 2KN_h(s, a, s') \\ N_h(s, a)^2 \\ \vdots \\ N_h(s, a)^2 + K^2 - 2KN_h(s, a) \end{bmatrix},$$

$$\mathbf{W}_{12}^{(1)} \mathbf{W}_{11}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} 1 - N_h(s, a, s')^2 - N_h(s, a)^2 \\ \vdots \\ 1 - (N_h(s, a, s') - x)^2 - (N_h(s, a) - y)^2 \\ \vdots \\ 1 - (N_h(s, a, s') - K)^2 - (N_h(s, a) - K)^2 \end{bmatrix},$$

where $x, y \in \{0\} \cup [K]$. Moreover, we construct $\mathbf{W}_{21}^{(1)}$ so that on the entries corresponding to h, s, a, s' it implements

$$\mathbf{W}_2^{(1)} \sigma(\mathbf{W}_1^{(1)} \mathbf{h}_i^{(0)}) = \left[\sum_{x,y=0}^K \sigma(1 - (N_h(s, a, s') - x)^2 - (N_h(s, a) - y)^2) \cdot \frac{x}{y\sqrt{1}} \right] = \left[\frac{N_h(s, a, s')}{N_h(s, a)\sqrt{1}} \right]$$

It can be verified that we can find such $\mathbf{W}_1^{(1)}, \mathbf{W}_2^{(1)}$ with

$$\|\mathbf{W}_1^{(1)}\|_{\text{op}} \leq \|\mathbf{W}_{11}^{(1)}\|_{\text{op}} \|\mathbf{W}_{12}^{(1)}\|_{\text{op}} \leq O(K^2) \cdot O(K) = O(K^3),$$

$\|\mathbf{W}_2^{(1)}\|_{\text{op}} \leq O(K)$, and the number of hidden neurons $D' = O(K^2 H S^2 A)$. Similarly, we can compute $f_2(\cdot)$ (or $f_3(\cdot)$) exactly following the same construction but with a different $\mathbf{W}_2^{(1)}$ that records all possible values of $f_2(\cdot)$ (or $f_3(\cdot)$). Combining the upper bounds on the operator norm of the weight matrices, we further have $\|\boldsymbol{\theta}\| \leq \tilde{O}(H S^2 A + K^3 + KH)$.

Proof of Step 3 Given $\hat{V}_{H+1} = \hat{V}_{H+1} = \mathbf{0}$, we show there exists a transformer with

$$L = 2, \quad \max_{\ell \in [L]} M^{(\ell)} \leq 2SA, \quad \max_{\ell \in [L]} D'^{(\ell)} \leq 3SA, \quad \|\boldsymbol{\theta}\| \leq O(H + SA)$$

that implements one step of value iteration

$$\begin{aligned} \hat{Q}_h(s, a) &= \max\{\min\{H, r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s')\}, 0\}, \\ \hat{V}_h(s) &= \max_{a \in \mathcal{A}} \hat{Q}_h(s, a) \end{aligned}$$

for some $h \in [H]$. Namely, we start with constructing an attention layer with $M = 2SA$ and $\{\mathbf{Q}_{ijh,s}^{(1)}\}_{s=1}^2, \{\mathbf{K}_{ijh,s}^{(1)}\}_{s=1}^2, \{\mathbf{V}_{ijh,s}^{(1)}\}_{s=1}^2$ such that for all $i \leq 2(t-1) + k$

$$\begin{aligned} \mathbf{Q}_{ij,1}^{(1)} \mathbf{h}_i^{(0)} &= \begin{bmatrix} \mathbf{B} \\ -i \\ \hat{V}_{h+1}(\cdot) \end{bmatrix}, \quad \mathbf{K}_{ij,1}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} i \\ \mathbf{B} \\ \hat{P}_{h+1}(\cdot | s, a) \end{bmatrix}, \quad \mathbf{V}_{ij,1}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{0} \\ i \mathbf{e}_{ij}^{Q_h} \\ \mathbf{0} \end{bmatrix} \\ \mathbf{Q}_{ij,2}^{(1)} \mathbf{h}_i^{(0)} &= \begin{bmatrix} \mathbf{B} \\ -i \\ -\hat{V}_{h+1}(\cdot) \end{bmatrix}, \quad \mathbf{K}_{ij,2}^{(1)} = \mathbf{K}_{ij,1}^{(1)}, \quad \mathbf{V}_{ij,2}^{(1)} = -\mathbf{V}_{ij,1}^{(1)} \end{aligned}$$

where $\mathbf{B} = 3H$ and $\mathbf{e}_{ij}^{Q_h} \in \mathbb{R}^{SA}$ is a vector supported on some coordinates in \mathbf{h}_i^d reserved for \hat{Q}_h . Moreover, we have $\|\mathbf{Q}_{ijh,s}^{(1)}\|_{\text{op}}, \|\mathbf{K}_{ijh,s}^{(1)}\|_{\text{op}} \leq \mathbf{B}$, $\|\mathbf{V}_{ijh,s}^{(1)}\|_{\text{op}} = 1$. Since

$$\left| \left\langle \hat{V}_{h+1}(\cdot), \hat{P}_{h+1}(\cdot | s, a) \right\rangle \right| \leq \left\| \hat{V}_{h+1}(\cdot) \right\|_{\infty} \cdot \left\| \hat{P}_{h+1}(\cdot | s, a) \right\|_1 \leq H$$

as $\hat{V}_{h+1}(s) \in [0, H]$ and $\left\| \hat{P}_{h+1}(\cdot | s, a) \right\|_1 = 1$, it follows that summing up two heads gives the update for $i \leq 2(t-1) + k$

$$\begin{aligned} \mathbf{0} &\mapsto \mathbf{0} + \left[\sigma\left(\left\langle \mathbf{Q}_{ij,1}^{(1)} \mathbf{h}_i^{(0)}, \mathbf{K}_{ij,1}^{(1)} \mathbf{h}_j^{(0)} \right\rangle\right) - \sigma\left(\left\langle \mathbf{Q}_{ij,2}^{(1)} \mathbf{h}_i^{(0)}, \mathbf{K}_{ij,1}^{(1)} \mathbf{h}_j^{(0)} \right\rangle\right) \right] \mathbf{e}_{ij}^{Q_h} \\ &= \left\langle \mathbf{Q}_{ij,1}^{(1)} \mathbf{h}_i^{(0)}, \mathbf{K}_{ij,1}^{(1)} \mathbf{h}_j^{(0)} \right\rangle \mathbf{e}_{ij}^{Q_h}. \end{aligned}$$

Denote the resulting token vector by $\mathbf{h}_i^{(1)}$. Moreover, we can construct a two-layer MLP with

$$\|\mathbf{W}_1^{(1)}\|_{\text{op}} = O(H), \quad \|\mathbf{W}_2^{(1)}\|_{\text{op}} \leq 3, \quad D' = 3SA$$

such that for any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ on the corresponding coordinates

$$\mathbf{W}_1^{(1)} \mathbf{h}_i^{(1)} = \begin{bmatrix} \vdots \\ -[r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s')] \\ r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s') - H \\ r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s') \\ \vdots \end{bmatrix}$$

and

$$\begin{aligned} \mathbf{W}_2^{(1)} \sigma(\mathbf{W}_1^{(1)} \mathbf{h}_i^{(1)}) &= \sigma(-[r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s')]) \\ &\quad - \sigma(r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s') - H) \\ &\quad + \sigma(r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s')) \\ &= \max\{\min\{H, r_h(s, a) + b_h(s, a) + \sum_{s' \in \mathcal{S}} \hat{P}_h(s' | s, a) \hat{V}_{h+1}(s')\}, 0\} = \hat{Q}_h(s, a). \end{aligned}$$

Denote the resulting token vector by $\mathbf{h}_i^{(2)}$. Next, we construct a second MLP layer with

$$\|\mathbf{W}_1^{(2)}\|_{\text{op}} \leq 2, \quad \|\mathbf{W}_2^{(2)}\|_{\text{op}} \leq \sqrt{A}, \quad D' = AS$$

such that for any $s \in \mathcal{S}$ on the corresponding coordinates we have

$$\mathbf{W}_1^{(2)} \mathbf{h}_i^{(2)} = \begin{bmatrix} \vdots \\ \hat{Q}_h(s, a_1) \\ \hat{Q}_h(s, a_2) - \hat{Q}_h(s, a_1) \\ \vdots \\ \hat{Q}_h(s, a_A) - \hat{Q}_h(s, a_{A-1}) \\ \vdots \end{bmatrix},$$

where a_j denotes the j -th action, and

$$\begin{aligned} \mathbf{W}_2^{(2)} \sigma(\mathbf{W}_1^{(2)} \mathbf{h}_i^{(2)}) &= \sigma(\hat{Q}_h(s, a_1)) + \sum_{j=2}^A \sigma(\hat{Q}_h(s, a_j) - \hat{Q}_h(s, a_{j-1})) \\ &= \max_{a \in \mathcal{A}} \hat{Q}_h(s, a) = \hat{V}_h(s). \end{aligned}$$

Using the upper bounds on the operator norm of the weight matrices, we further have $\|\boldsymbol{\theta}\| \leq O(SA + H)$. Combining the steps concludes the construction in Step 3.

Proof of Step 4 we start with constructing an-attention layer with $M = 2HA$ and $\{\mathbf{Q}_{\text{jh},s}^{(1)}\}_{s=1}^2, \{\mathbf{K}_{\text{jh},s}^{(1)}\}_{s=1}^2, \{\mathbf{V}_{\text{jh},s}^{(1)}\}_{s=1}^2$ such that for all the current token $i = 2(t-1) + k$ and $j \leq i$

$$\begin{aligned} \mathbf{Q}_{\text{jh},1}^{(1)} \mathbf{h}_i^{(0)} &= \begin{bmatrix} s_{\bar{k}(i), \bar{h}(i)} \\ -i \\ \mathbf{B} \end{bmatrix}, \quad \mathbf{K}_{\text{jh},1}^{(1)} \mathbf{h}_j^{(0)} = \begin{bmatrix} \hat{Q}_h(\cdot, a_j) \\ \mathbf{B} \\ j \end{bmatrix}, \quad \mathbf{V}_{\text{jh},1}^{(1)} \mathbf{h}_i^{(0)} = \begin{bmatrix} \mathbf{0} \\ i \mathbf{e}_{\text{jh}} \\ \mathbf{0} \end{bmatrix} \\ \mathbf{Q}_{\text{jh},2}^{(1)} \mathbf{h}_i^{(0)} &= \begin{bmatrix} -s_{\bar{k}(i), \bar{h}(i)} \\ -i \\ \mathbf{B} \end{bmatrix}, \quad \mathbf{K}_{\text{jh},2}^{(1)} = \mathbf{K}_{\text{jh},1}^{(1)}, \quad \mathbf{V}_{\text{jh},2}^{(1)} = -\mathbf{V}_{\text{jh},1}^{(1)}, \end{aligned}$$

where we choose $\mathbf{B} = 2H$ and $\mathbf{V}_{\text{jh},1}^{(1)} \mathbf{h}_i^{(0)}$ is a one-hot vector supported on some entry of $\mathbf{h}_i^{(0)}$. We verify that summing up the heads gives the update

$$\mathbf{0} \mapsto \hat{Q}_h(s_{k,h}, a_j) \mathbf{e}_{\text{jh}}$$

for all $h \in [H], j \in [A]$. Moreover, we have $\|\mathbf{Q}_{jh,s}^{(1)}\|_{\text{op}} \leq 2H, \|\mathbf{K}_{jh,s}^{(1)}\|_{\text{op}} \leq 2H, \|\mathbf{V}_{jh,s}^{(1)}\|_{\text{op}} \leq 1$ for $s = 1, 2$. Through this attention-only layer, we extract the values $\widehat{Q}_h(s_{k,h}, a_j)$ for all $h \in [H]$ from the Q-function.

Similar to the proof of Step 1b, we construct a second attention-only layer with attention heads $\{\mathbf{Q}_{jh,s}^{(2)}\}_{s=1}^2, \{\mathbf{K}_{jh,s}^{(2)}\}_{s=1}^2, \{\mathbf{V}_{jh,s}^{(2)}\}_{s=1}^2$ that

$$\mathbf{Q}_{jh,1}^{(2)} \mathbf{h}_i^{(1)} = \begin{bmatrix} 1 \\ -\bar{h}(i) \\ -i \\ \mathbf{B} \end{bmatrix}, \quad \mathbf{K}_{jh,1}^{(2)} \mathbf{h}_j^{(1)} = \begin{bmatrix} \mathbf{h} \\ 1 \\ \mathbf{B} \\ j \end{bmatrix}, \quad \mathbf{V}_{jh,1}^{(2)} \mathbf{h}_i^{(1)} = - \begin{bmatrix} \mathbf{0} \\ \widehat{Q}_h(s_{k,h}, a_j) \mathbf{e}_j \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{Q}_{jh,2}^{(2)} = \mathbf{Q}_{jh,1}^{(2)}, \quad \mathbf{K}_{jh,2}^{(2)} \mathbf{h}_j^{(1)} = \begin{bmatrix} \mathbf{h} - 1 \\ 1 \\ \mathbf{B} \\ j \end{bmatrix}, \quad \mathbf{V}_{jh,2}^{(2)} = -\mathbf{V}_{jh,1}^{(2)},$$

where $\mathbf{V}_{jh,1}^{(2)} \mathbf{h}_i^{(1)}$ are supported on some entry of \mathbf{h}_i^d for $s = 1, 2$. Summing up the heads gives the update

$$\mathbf{0} \mapsto - \sum_{s=\bar{h}(i)+1}^H \frac{1}{i} \widehat{Q}_s(s_{k,h}, a_j).$$

Similarly, we can construct attention heads $\{\mathbf{Q}_{jh,s}^{(2)}\}_{s=3}^4, \{\mathbf{K}_{jh,s}^{(2)}\}_{s=3}^4, \{\mathbf{V}_{jh,s}^{(2)}\}_{s=3}^4$ that implements

$$\mathbf{0} \mapsto -\frac{1}{i} \sum_{s=1}^{\bar{h}(i)-1} \widehat{Q}_s(s_{k,h}, a_j).$$

Moreover, we construct $\mathbf{Q}_{jh,5}^{(2)}, \mathbf{K}_{jh,5}^{(2)}, \mathbf{V}_{jh,5}^{(2)}$ with

$$\mathbf{Q}_{jh,5}^{(2)} \mathbf{h}_i^{(1)} = \begin{bmatrix} 1 \\ -i \\ \mathbf{B} \end{bmatrix}, \quad \mathbf{K}_{jh,5}^{(2)} \mathbf{h}_j^{(1)} = \begin{bmatrix} 1 \\ \mathbf{B} \\ j \end{bmatrix}, \quad \mathbf{V}_{jh,5}^{(2)} \mathbf{h}_i^{(1)} = \begin{bmatrix} \mathbf{0} \\ \widehat{Q}_h(s_{k,h}, a_j) \mathbf{e}_j \\ \mathbf{0} \end{bmatrix}$$

that implements

$$\mathbf{0} \mapsto \frac{1}{i} \sum_{s=1}^H \widehat{Q}_s(s_{k,h}, a_j).$$

Therefore, summing up the $M = 5HA$ heads we obtain the update

$$\mathbf{0}_A \mapsto \frac{1}{i} \widehat{Q}_h(s_{k,h}, \cdot).$$

Note that $\|\mathbf{Q}_{jh,s}^{(1)}\|_{\text{op}} \leq 4H, \|\mathbf{K}_{jh,s}^{(1)}\|_{\text{op}} \leq 4H, \|\mathbf{V}_{jh,s}^{(1)}\|_{\text{op}} \leq 1$ for $s \in [5]$.

Finally, we apply an attention-only layer to implement the multiplication by a factor of i/τ using a similar construction as in Eq. (27) with $\|\mathbf{Q}_1^{(3)}\|_{\text{op}} = O(HK), \|\mathbf{K}_1^{(3)}\|_{\text{op}} = O(HK), \|\mathbf{V}_1^{(3)}\|_{\text{op}} = O(1/\tau)$, and assign the resulting vector $\widehat{Q}(s_{k,h}, \cdot)/\tau$ to \mathbf{h}_i^c . Combining the three attention-only layers completes Step 4.

H.4 Proof of Theorem 17

By Theorem 1 and 16, it suffices to show the regret of soft UCB-VI satisfies

$$\mathbb{E}[KV_M(\pi^*) - \mathfrak{R}_{M, \text{Alg}_{\text{UCBVVI}}(\tau)}(T)] \leq \tilde{O}(H^2 \sqrt{SAK} + H^3 S^2 A)$$

for all MDP instances M , where $\tau = 1/K$ and $\tilde{O}(\cdot)$ hides logarithmic dependencies on (H, K, S, A) .

Throughout the proof, we may drop the dependence on M for notational simplicity when there is no confusion. For each episode $k \in [K]$, let $N_h^k, \widehat{P}_h^k, \widehat{Q}_h^k, \widehat{V}_h^k, b_h^k$ denote the corresponding quantities $N_h, \widehat{P}_h, \widehat{Q}_h, \widehat{V}_h, b_h$ introduced in UCB-VI (see Section H.2).

For a policy π and time step $h \in [H]$, we define the Q-function Q_h^π and the value function V_h^π

$$Q_h^\pi(s, a) := \mathbb{E}\left[\sum_{t=h}^H r(s_t, a_t) \mid s_h = s, a_h = a, \pi\right],$$

$$V_h^\pi(s) := \mathbb{E}\left[\sum_{t=h}^H r(s_t, a) \mid s_h = s, \pi\right].$$

We use $\pi^k = (\pi_1^k, \dots, \pi_H^k)$, $\pi_{\text{sm}}^k = (\pi_{\text{sm},1}^k, \dots, \pi_{\text{sm},h}^k, \dots, \pi_{\text{sm},H}^k)$ to denote the policies given by UCB-VI and soft UCB-VI in the k -th episode, respectively. Note that we have $V_M(\pi) = \mathbb{E}_{s \sim \mu_1}[V_h^\pi(s)]$ and cumulative the regret

$$\mathbb{E}[KV_M(\pi^*) - \mathfrak{R}_{M, \text{Alg}_{\text{UCBVVI}}(\tau)}(T)] = \mathbb{E}\left[\sum_{k=1}^K [V_1^{\pi^*}(s_{k,1}) - V_1^{\pi_{\text{sm}}^k}(s_{k,1})]\right]$$

where the expectation is taken over the collected data

$$D_T = \{(s_{k,h}, a_{k,h}, r_{k,h})\}_{k \in [K], h \in [H]} \sim \mathbb{P}_M^{\text{sUCBVVI}(\tau)}.$$

For any function $f = f(s, a)$, we abuse the notation $f(s, \pi(\cdot)) := \mathbb{E}_{a \sim \pi}[f(s, a)]$. Lastly, we define

$$\varepsilon_{\text{sffmax}} = \max_{k \in [K], h \in [H], s \in \mathcal{S}} [\widehat{Q}_h^k(s, \pi_h^k(\cdot)) - \widehat{Q}_h^k(s, \pi_{\text{sm},h}^k(\cdot))].$$

We claim the following which we will prove later

$$\varepsilon_{\text{sffmax}} \leq A\tau. \quad (28)$$

The proof follows from similar arguments as in the proof of Theorem 1 in [4] (see also Theorem 7.6 in [1]). Hence we only provide a sketch of proof here. First, from the proof of Theorem 7.6 in [1], it can be shown that

$$\widehat{V}_h^k(s) \geq V_h^{\pi^*}(s)$$

for any k, h, s with probability at least $1 - \delta$. Thus with probability at least $1 - \delta$ for all $h \in [H], k \in [K]$

$$\begin{aligned} & V_h^{\pi^*}(s_{k,h}) - V_h^{\pi_{\text{sm}}^k}(s_{k,h}) \\ & \leq \widehat{V}_h^k(s_{k,h}) - V_h^{\pi_{\text{sm}}^k}(s_{k,h}) \\ & = \widehat{Q}_h^k(s_{k,h}, \pi_h^k(\cdot)) - \widehat{Q}_h^k(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) + \widehat{Q}_h^k(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) \\ & \leq \widehat{Q}_h^k(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) + \varepsilon_{\text{sffmax}} \\ & = \widehat{Q}_h^k(s_{k,h}, a_{k,h}) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, a_{k,h}) + \text{MD}_{k,h}^{(1)} + \varepsilon_{\text{sffmax}}, \end{aligned}$$

where the first equality uses $\widehat{V}_h^k(s_{k,h}) = \arg \max_a \widehat{Q}_h^k(s_{k,h}, a_{k,h}) = \widehat{Q}_h^k(s_{k,h}, \pi_h^k(\cdot))$, and in the last line

$$\text{MD}_{k,h}^{(1)} := [\widehat{Q}_h^k(s_{k,h}, \pi_{\text{sm},h}^k(\cdot)) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, \pi_{\text{sm},h}^k(\cdot))] - [\widehat{Q}_h^k(s_{k,h}, a_{k,h}) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, a_{k,h})].$$

Note that for any fixed $h \in [H]$, $\{\text{MD}_{k,h}^{(1)}\}_{k=1}^K$ is a bounded martingale difference sequence. Following the proof of Theorem 7.6 in [1], we further have

$$\begin{aligned} & V_h^{\pi^*}(s_{k,h}) - V_h^{\pi_{\text{sm}}^k}(s_{k,h}) \\ & \leq \widehat{Q}_h^k(s_{k,h}, a_{k,h}) - Q_h^{\pi_{\text{sm}}^k}(s_{k,h}, a_{k,h}) + \text{MD}_{k,h}^{(1)} + \varepsilon_{\text{sffmax}} \\ & \leq \left(1 + \frac{1}{H}\right) \left[\widehat{V}_{h+1}^k(s_{k,h+1}) - V_{h+1}^{\pi_{\text{sm}}^k}(s_{k,h+1})\right] + 2b_h^k(s_{k,h}, a_{k,h}) \\ & \quad + \frac{c_0 L_0 H^2 S}{N_h^k(s_{k,h}, a_{k,h})} + \text{MD}_{k,h}^{(2)} + \text{MD}_{k,h}^{(1)} + \varepsilon_{\text{sffmax}}, \end{aligned}$$

with probability at least $1 - c\delta$ for some universal constant $c > 0$, where $L_0 = \log(SAKH/\delta)$, $c_0 > 0$ is some universal constant and

$$\text{MD}_{k,h}^{(2)} := P_h(\cdot \mid s_{k,h}, a_{k,h}) \cdot (V_{h+1}^{\pi^*} - V_{h+1}^{\pi_{sm}^k}) - (V_{h+1}^{\pi^*}(s_{k,h+1}) - V_{h+1}^{\pi_{sm}^k}(s_{k,h+1}))$$

is a bounded martingale difference sequence for any fixed $h \in [H]$. Using the recursive formula and the fact that $(1 + 1/H)^H < e$, we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=1}^K [V_1^{\pi^*}(s_{k,1}) - V_1^{\pi_{sm}^k}(s_{k,1})] \right] \\ & \leq c\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \left[2b_h^k(s_{k,h}, a_{k,h}) + \frac{c_0 L_0 H^2 S}{N_h^k(s_{k,h}, a_{k,h})} + \text{MD}_{k,h}^{(2)} + \text{MD}_{k,h}^{(1)} \right] \right] + \mathbb{E} \left[K \sum_{h=0}^{H-1} \left(1 + \frac{1}{H}\right)^h \varepsilon_{\text{sflmax}} \right] \\ & \leq c\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \left[2b_h^k(s_{k,h}, a_{k,h}) + \frac{c_0 L_0 H^2 S}{N_h^k(s_{k,h}, a_{k,h})} + \text{MD}_{k,h}^{(2)} + \text{MD}_{k,h}^{(1)} \right] \right] + cKHA\tau \\ & \leq \tilde{\mathcal{O}}(H^2 \sqrt{SAK} + H^3 S^2 A) + cKHA\tau \\ & \leq \tilde{\mathcal{O}}(H^2 \sqrt{SAK} + H^3 S^2 A), \end{aligned}$$

where $c > 0$ is some universal constant, $\tilde{\mathcal{O}}(\cdot)$ hides logarithmic dependencies on (H, K, S, A) , and the last line follows again from the proof of Theorem 7.6 in [1], and the assumption that $\tau = 1/K$. We omit the detailed derivations here as they are similar to those in [4, 1]. Therefore, we conclude the proof of the first part of Theorem 17. Moreover, the second part of Theorem 17 (i.e., the upper bound on $\log \mathcal{N}_\Theta$) follows immediately from Lemma 6 and Eq. (15).

Proof of Eq. (28) By definition of \hat{Q}_h^k and $\pi_h^k, \pi_{sm,h}^k$, we have

$$\begin{aligned} \hat{Q}_h^k(s, \pi_h^k(\cdot)) - \hat{Q}_h^k(s, \pi_{sm,h}^k(\cdot)) &= \max_a \hat{Q}_h^k(s, a) - \sum_a \frac{\exp(\hat{Q}_h^k(s, a)/\tau)}{\sum_a \exp(\hat{Q}_h^k(s, a)/\tau)} \cdot \hat{Q}_h^k(s, a) \\ &= \sum_a \frac{\exp(\hat{Q}_h^k(s, a)/\tau)}{\sum_a \exp(\hat{Q}_h^k(s, a)/\tau)} \cdot [\max_a \hat{Q}_h^k(s, a) - \hat{Q}_h^k(s, a)] \\ &\leq \sum_a \frac{\exp(\hat{Q}_h^k(s, a)/\tau)}{\exp(\max_a \hat{Q}_h^k(s, a)/\tau)} \cdot [\max_a \hat{Q}_h^k(s, a) - \hat{Q}_h^k(s, a)] \\ &\leq A \cdot \left[\sup_{t \geq 0} t \exp(-t/\tau) \right] \leq A\tau. \end{aligned}$$

I Experiments

In this section, we perform preliminary simulations to demonstrate the ICRL capabilities of transformers and validate our theoretical findings. We remark that while similar experiments have been conducted in existing works [30, 33], our setting differs in several aspects such as imitating the entire interaction trajectory in our pretrain loss (1) as opposed to on the last (query) state only as in Lee et al. [33].

We compare pretrained transformers against empirical average, LinUCB (or UCB), and Thompson sampling. We use a GPT-2 model [22, 33] with $L = 8$ layers, $M = 4$ heads, and embedding dimension $D = 32$. We utilize ReLU attention layers, aligning with our theoretical construction. We pretrain the transformer with two setups: (1) Both context algorithm Alg_0 and expert algorithm Alg_E use LinUCB (the Algorithm Distillation approach); (2) Context algorithms Alg_0 mixes uniform policy and Thompson sampling, while expert $\text{Alg}_E = a_t^*$ provides optimal actions (DPT). See Appendix J for further experimental details.

In the first setup, we consider stochastic linear bandits with $d = 5$ and $A = 10$. At each $t \in [200]$, the agent chooses an action a_t and receives reward $r_t = \langle a_t, \mathbf{w}^* \rangle + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, 1.5^2)$. The parameter \mathbf{w}^* is from $\text{Unif}([0, 1]^d)$. The action set $\mathbb{A}_t = \mathbb{A}$ is fixed over time with actions i.i.d. from $\text{Unif}([-1, 1]^d)$. We generate 100K trajectories using $\text{Alg}_0 = \text{Alg}_E = \text{LinUCB}$ and train transformer




Figure 1: Regrets of transformer (TF), empirical average (Emp), Thompson sampling (TS) and LinUCB or UCB (LinUCB reduces to UCB for Bernoulli bandits). Left: linear bandit with $d = 5$, $A = 10$, $\sigma = 1.5$, $\text{Alg}_0 = \text{Alg}_E = \text{LinUCB}$. Right: Bernoulli bandit with $d = 5$, $\text{Alg}_0 = (\text{Alg}_{\text{unif}} + \text{Alg}_{\text{TS}})/2$ and $\text{Alg}_E = a^*$. The simulation is repeated 500 times. Shading displays the standard deviation of the regret estimates.

$\text{TF}_{\hat{\theta}}(\cdot)$ via Eq. (1). Figure 1 (left) shows regrets of the transformer (TF), empirical average (Emp), LinUCB, and Thompson sampling (TS). The transformer outperforms Thompson sampling and empirical average, and is comparable to LinUCB, agreeing with Theorem 13. The small regret gap between TF and LinUCB may stem from the limited capacity of the GPT2 model.

In the second setup, we consider multi-armed Bernoulli bandits with $d = 5$. The parameter \mathbf{w}^* is from $\text{Unif}([0, 1]^d)$. The fixed action set $\mathbb{A}_t = \mathbb{A}$ contains one-hot vectors $\{\mathbf{e}_i\}_{i=1}^d$ (multi-armed bandits). At each $t \in [200]$, the agent selects a_t receives reward $r_t \sim \text{Bern}(\langle a_t, \mathbf{w}^* \rangle)$. Let Alg_{unif} be the uniform policy. We use Alg_{unif} and Alg_{TS} as context algorithms to generate 50K trajectories each. The expert is fixed as $\text{Alg}_E = a^*$. We train transformer $\text{TF}_{\hat{\theta}}(\cdot)$ via Eq. (1). Figure 1 (right) shows regrets for the pretrained transformer (TF), empirical average (Emp), UCB, and Thompson sampling (TS). The transformer aligns with Thompson sampling, validating Theorem 15. However, TS underperforms UCB for Bernoulli bandits, as shown.

J Experimental details

This section provides implementation details of our experiments and some additional simulations. Our code is available at <https://anonymous.4open.science/r/in-context-rl>.

J.1 Implementation details

Model and embedding Our experiments use a GPT-2 model [48] with ReLU activation layers. The model has $L = 8$ attention layers, $M = 4$ attention heads, and embedding dimension $D = 32$. Following standard implementations in [66], we add Layer Normalization [5] after each attention and MLP layer to facilitate optimization. We consider the embedding and extraction mappings as described in Appendix F.1, and train transformer $\text{TF}_{\hat{\theta}}(\cdot)$ via maximizing Eq. (1).

Online algorithms We compare the regret of the algorithm induced by the transformer with empirical average, Thompson sampling, and LinUCB (or UCB for Bernoulli bandits).

(Emp) Empirical average. For time $t \leq A$, the agent selects each action once. For time $t > A$, the agent computes the average of the historical rewards for each action and selects the action with the maximal averaged historical rewards.

(TS) Thompson sampling. For linear bandits with Gaussian noises, we consider Thompson sampling introduced in Appendix G.1 with $r = \sigma = 1.5$ and $\lambda = 1$ (note that in this case TS does not correspond to posterior sampling as we assume \mathbf{w}^* follows the uniform distribution on $[0, 1]^d$). For Bernoulli bandits, we consider the standard TS sampling procedure (see, for example, Algorithm 3.2 in [58]).

(LinUCB) Linear UCB and UCB. For linear bandits, we use LinUCB (Appendix F.2) with $\lambda = 1$ and $\alpha = 2$. For multi-armed Bernoulli bandits, LinUCB reduces to UCB, which selects $a_t = \arg \max_{a \in \mathcal{A}} \{\hat{\mu}_{t,a} + \sqrt{1/N_t(a)}\}$, where $\mu_{t,a}$ is the average reward for action a up to time t , and $N_t(a)$ is the number of times action a was selected up to time t .

J.2 Additional experiments and plots

We provide additional experiments and plots in this section. In all experiments, we choose the number of samples $n = 100K$.

Additional plots of suboptimality $\langle a_t^* - a_t, \mathbf{w}^* \rangle$ over time are shown in Figure 2 for the two experiments in Section I. In both cases, the transformer is able to imitate the expected expert policy $\overline{\text{Alg}}_E$, as its suboptimality closely matches $\overline{\text{Alg}}_E$ (LinUCB and TS for the left and right panel, respectively). While the empirical average (Emp) has lower suboptimality early on, its gap does not converge to zero. In contrast, both LinUCB and Thompson sampling are near-optimal up to $\tilde{O}(1)$ factors in terms of their (long-term) regret.




Figure 2: Suboptimality of transformer (TF), empirical average (Emp), Thompson sampling (TS), and LinUCB (or UCB). Left: linear bandit with $d = 5$, $A = 10$, $\sigma = 1.5$, $\text{Alg}_0 = \text{Alg}_E = \text{LinUCB}$. Right: Bernoulli bandit with $d = 5$, $\text{Alg}_0 = (\text{Alg}_{\text{unif}} + \text{Alg}_{\text{TS}})/2$, and $\text{Alg}_E = a_t^*$. The simulation is repeated 500 times. Shading displays the standard deviation of the sub-optimality estimates.

Additional simulations were run with $\text{Alg}_0 = \text{Alg}_E = \text{UCB}$ for Bernoulli bandits, which has fewer actions ($A = 5$) than linear bandits ($A = 10$). Figure 3 shows the regret and suboptimality of UCB and the transformer overlap perfectly, with both algorithms exhibiting optimal behavior. This suggests the minor gaps between LinUCB and transformer in the left panel of Figure 1 and 2 are likely due to limited model capacity.




Figure 3: Regrets and suboptimality of transformer (TF), empirical average (Emp), Thompson sampling (TS), and UCB. Settings: Bernoulli bandit with $d = 5$, and $\text{Alg}_0 = \text{Alg}_E = \text{LinUCB}$. The simulation is repeated 500 times. Shading displays the standard deviation of the estimates.

J.3 The effect of distribution ratio

We evaluate the effect of the distribution ratio $\mathcal{R} = \mathcal{R}_{\overline{\text{Alg}}_E, \text{Alg}_0}$ (Definition 5) on transformer performance. We consider the Bernoulli bandit setting from Section I with expert $\text{Alg}_E = a^*$ giving optimal actions. The context algorithm is

$$\text{Alg}_0 = \alpha \text{Alg}_{\text{TS}} + (1 - \alpha) \text{Alg}_{\text{unif}},$$

mixing uniform policy Alg_{unif} and Thompson sampling Alg_{TS} , for $\alpha \in \{0, 0.1, 0.5, 1\}$. The case $\alpha = 0$ corresponds to the context algorithm being the i.i.d. uniform policy, and $\alpha = 1$ corresponds to the context algorithm being Thompson sampling. Note that the distribution ratio \mathcal{R} may scale as $\mathcal{O}((1/\alpha) \wedge A^{\mathcal{O}(T)})$ in the worst case.

Figure 4 evaluates the learned transformers against Thompson sampling for varying context algorithms. The left plot shows cumulative regret for all algorithms. The right plot shows the regret difference between transformers and Thompson sampling. The results indicate that an increased distribution ratio impairs transformer regret, as expected. Moreover, it is observed that the transformer, even with the uniform policy (i.e., $\alpha = 0$), is capable of imitating Thompson sampling in the early stages (Time ≤ 30), exceeding theoretical predictions. This suggests the transformer can learn Thompson sampling even when the context algorithm differs significantly from the expert algorithm.




Figure 4: Regrets and difference of regrets between transformers and Thompson sampling, for different context algorithms. Settings: Bernoulli bandit with $d = 5$, $\text{Alg}_E = a_t^*$ and $\text{Alg}_0 = \alpha \text{Alg}_{\text{TS}} + (1 - \alpha) \text{Alg}_{\text{unif}}$ with $\alpha \in \{0, 0.1, 0.5, 1\}$. The simulation is repeated 500 times. Shading displays the standard deviation of the estimates.