
A Statistical Guarantee for Representation Transfer in Multitask Imitation Learning

Bryan Chan*

Department of Computing Science
University of Alberta
Edmonton, Canada
bryan.chan@ualberta.ca

Karime Pereida[†] & James Bergstra[†]

AI Platform
Ocado Technology
Toronto, Canada
{k.pereidaperez, james.bergstra}@ocado.com

Abstract

Transferring representation for multitask imitation learning has the potential to provide improved sample efficiency on learning new tasks, when compared to learning from scratch. In this work, we provide a statistical guarantee indicating that we can indeed achieve improved sample efficiency on the target task when a representation is trained using sufficiently diverse source tasks. Our theoretical results can be readily extended to account for commonly used neural network architectures with realistic assumptions. We conduct empirical analyses that align with our theoretical findings on four simulated environments—in particular leveraging more data from source tasks can improve sample efficiency on learning in the new task.

1 Introduction

Imitation learning (IL) is a common approach to learn sequential decision making agents—it involves imitating the expert through matching distributions induced by the expert demonstrations [Osa et al., 2018]. However, current methods require thousands of demonstrations even in simple tasks [Mandlekar et al., 2022, Jang et al., 2021, Ablett et al., 2023]. Acquiring large amounts of data can be expensive and even infeasible in domains including robotic and healthcare applications. To address this challenge, empirical research has proposed transferring part of the agent trained from one or more tasks to a target task with the goal of improving sample efficiency on the target task [Brohan et al., 2023, Hansen et al., 2022, Jang et al., 2021, Li et al., 2022]. In this paper we show how much sample efficiency on the target task improves, compared to training an agent from scratch, when we transfer a pretrained representation to the target task via multitask imitation learning (MTIL).

We have three main contributions: (1) Although Arora et al. [2020] has investigated the benefits of representation transfer, their result does not relate the target task and the source tasks. We cannot guarantee that transferring the representation to a particular target task can yield any benefit. Inspired by Tripuraneni et al. [2020], we provide an analysis that relates the source and target tasks via the notion of task diversity. (2) This line of work [Arora et al., 2020, Tripuraneni et al., 2020, Maurer et al., 2016] relies on Gaussian complexity, in this paper we instead use Rademacher complexity and provide a tighter bound by a log factor than using Gaussian complexity. Our result is due to the objective of behavioural cloning, where the method aims to minimize the Kullback–Leibler (KL) divergence between the expert and the learner [Xu et al., 2020]. The consequence is that we can connect our result with deep-learning theory, where commonly used neural networks are directly quantified with Rademacher complexity [Bartlett et al., 2021]. (3) Based on our theory, we further

*Work done during internship at Ocado Technology.

[†]Equal advising.

conduct experiments to demonstrate that transferring representations from source tasks to target tasks is a valid approach—the agent performs better as we increase the number of tasks and data.

2 Preliminaries

Sequential decision making problems can be formulated as Markov Decision Processes (MDPs). An infinite-horizon MDP is a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, r, P, \rho, \gamma \rangle$, where \mathcal{S} is a finite state space, \mathcal{A} is a finite action space, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the bounded reward function, $P \in \Delta_{\mathcal{S} \times \mathcal{A}}^{\mathcal{S}}$ is the transition distribution over the states for each state-action pair, $\rho \in \Delta^{\mathcal{S}}$ is the initial state distribution over the states, and $\gamma \in [0, 1)$ is the discount factor. The agent interacts with the environment through a policy $\pi \in \Pi$, which we assume to be stationary and Markovian (i.e. $\Pi = \Delta_{\mathcal{S}}^{\mathcal{A}}$.)

The interconnection between the policy π and the environment \mathcal{M} induces a random infinite-length trajectory $S_0, A_0, S_1, A_1, \dots$, where $S_0 \sim \rho$, $A_h \sim \pi(\cdot | S_h)$, and $S_{h+1} \sim P(\cdot | S_h, A_h)$. The corresponding discounted stationary state(-action) distribution for policy π , which describes the “frequency” of visiting state s (and action a) under π , can be written as $\nu_{\pi}(s) = (1 - \gamma) \sum_{h=0}^{\infty} \gamma^h \mathbb{P}(S_h = s; \pi)$ (and $\mu_{\pi}(s, a) = (1 - \gamma) \sum_{h=0}^{\infty} \gamma^h \mathbb{P}(S_h = s, A_h = a; \pi)$.)

For any policy $\pi \in \Pi$, the corresponding value function $v^{\pi} : \mathcal{S} \rightarrow \mathbb{R}$ is defined as $v^{\pi}(s) = \mathbb{E}_{\pi, \rho, P} [\sum_{h=0}^{\infty} \gamma^h r(S_h, A_h) | S_0 = s]$. The optimal policy π^* is such that $\pi^*(s) = \arg \max_{\pi \in \Pi} v^{\pi}(s)$, for any $s \in \mathcal{S}$ (randomly breaking ties in the case of two or more maxima.) In general, the goal is to obtain an ε -optimal policy (i.e. $v^{\pi}(s) \geq v^{\pi^*}(s) - \varepsilon$, for any $s \in \mathcal{S}$.)

In imitation learning, the learner has no access to the reward function r and is given demonstrations from the optimal (expert) policy π^* instead. The demonstrations form a set of N state-action pairs $\{(s_n, a_n)\}_{n=1}^N$, where $(s_n, a_n) \stackrel{\text{i.i.d.}}{\sim} \mu_{\pi^*}$. The goal is to obtain an ε -optimal policy using the demonstrations. Behavioural cloning treats this problem as a supervised learning problem and aims to minimize the risk [Pomerleau, 1988]. The risk is defined as $\ell(\pi) = \mathbb{E}_{(s,a) \sim \mu_{\pi^*}} [\ell(\pi(s), a)]$, where $\pi \in \Pi$ and $\ell(\cdot, \cdot) : \Delta^{\mathcal{A}} \times \mathcal{A} \rightarrow \mathbb{R}^+$ is a loss function. This paper considers the log loss $\ell(\pi(s), a) = -\log \pi(a|s)$, which is a surrogate of the 0-1 loss. The log loss is also equivalent to the KL-divergence between the expert and the learner when the expert is deterministic.

3 Multitask Imitation Learning (MTIL) with Representation Transfer

We consider the transfer-learning setting where we are given demonstrations of T source tasks. We define each task t as an MDP with different transition distributions and reward functions, but with the same state and action spaces. When the context is unclear, we use subscript, e.g. π_t , to denote the task-specific objects. Our goal is to leverage the demonstrations of T source tasks to learn an ε -optimal policy on a new target task τ with better sample efficiency, when compared to learning from scratch. To achieve this, we first learn a shared representation from the source tasks and transfer it to the target task. During the transfer, we fix the representation and only learn the task-specific mapping, similar to a finetuning procedure [Howard and Ruder, 2018, Liu et al., 2022].

Formally, we consider softmax parameterized policies of the form $\pi^{f, \phi}(s) = \text{softmax}((f \circ \phi)(s))$, where $f \in \mathcal{F}$ is the task-specific mapping and $\phi \in \Phi$ is the representation. We perform a two-phase procedure for transfer learning: (1) learn a representation $\hat{\phi}$ from the source tasks, and (2) learn a task-specific mapping \hat{f} such that $\pi^{\hat{f}, \hat{\phi}}$ performs well on the target task. We call the phases respectively the training phase and the transfer phase.

In the training phase, for each task t of the T source tasks, we are given N state-action pairs $\{(s_{t,n}, a_{t,n})\}_{n=1}^N$. Let $\mathbf{f} = (f_1, \dots, f_T)$, where $f_t \in \mathcal{F}$ is the task-specific mapping for task t , for $t \in [T]$, which we write $\mathbf{f} \in \mathcal{F}^{\otimes T}$ for conciseness. Then, we define the **empirical training risk** as

$$\hat{R}_{train}(\mathbf{f}, \phi) := \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N \ell(\pi^{f_t, \phi}(s_{t,n}), a_{t,n}), \quad (1)$$

and the corresponding minimizer of equation 1 is $\hat{\phi} = \arg \min_{\phi \in \Phi} \min_{\mathbf{f} \in \mathcal{F}^{\otimes T}} \hat{R}_{train}(\mathbf{f}, \phi)$.

In the transfer phase, we are given M state-action pairs $\{(s_m, a_m)\}_{m=1}^M$ for a target task τ . With the same loss function ℓ as equation 1, we define the **empirical test risk** as

$$\hat{R}_{test}(f_\tau, \phi) := \frac{1}{M} \sum_{m=1}^M \ell(\pi^{f_\tau, \phi}(s_m), a_m), \quad (2)$$

where $f_\tau \in \mathcal{F}$ is the task-specific mapping for task τ . We obtain a task-specific mapping that minimizes equation 2 based on the representation $\hat{\phi}$ obtained from the training phase. That is, $\hat{f}_\tau = \arg \min_{f \in \mathcal{F}} \hat{R}_{test}(f, \hat{\phi})$.

We perform empirical risk minimization (ERM) on both stages—the empirical risks allow us to quantify the generalization error of the learner. Observing that the log loss corresponds to the KL-divergence, we convert the generalization error to policy error. Finally, this allows us to establish the sample complexity bound of achieving ε -optimal policy based on the diversity of the source tasks. The diversity is measured with a positive constant σ , where larger σ corresponds to higher diversity. We now state our main theorem and defer the analysis to appendix B:

Theorem 1 (*Transfer Imitation Learning Policy Error Bound.*) *Let π_τ^* be the optimal policy for the target task τ . Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1 and let \hat{f}_τ be the ERM of \hat{R}_{test} defined in equation 2 by fixing $\hat{\phi}$. Let $\sigma > 0$. Suppose the source tasks are σ -diverse. With a deterministic expert policy π_τ^* and under some assumptions, we have that with probability $1 - 2\delta$,*

$$\text{Policy Error} = \|v^{\pi_\tau^*} - v^{\text{softmax}(\hat{f}_\tau \circ \hat{\phi})}\|_\infty \leq \frac{2\sqrt{2}}{(1-\gamma)^2} \sqrt{\varepsilon_{gen} + 2\zeta}, \quad (3)$$

where $\varepsilon_{gen} = \mathcal{O}\left(1/\sqrt{\sigma^2 NT}, 1/\sqrt{M}, \mathfrak{R}_{NT}(\Phi)/\sigma\right)$ is the generalization error, $\mathfrak{R}_{NT}(\Phi)$ is the Rademacher complexity, and $\zeta \in (0, 1)$ is a constant related to the policy realizability.

Theorem 1 indicates that we can obtain ε -optimal policy through the transfer-learning procedure. Notably, we can trade-off the number of target data M at the cost of the number of source tasks T and number of training data per task N . In contrast, behavioural cloning can only leverage the target data to train both the representation and the task-specific mapping. Consequently, if the representation class Φ is expressive, it is beneficial to pretrain the representation using the existing source data.

Remark 1 In practice, our policies are neural networks [Fujimoto et al., 2018, Haarnoja et al., 2018, Yarats et al., 2020]—we can consider the last layer as the task-specific mapping f and the remaining layers as the representation ϕ . These neural networks include multilayer perceptrons and convolutional neural networks with Lipschitz activation functions (e.g. ReLU, tanh, sigmoid, max-pooling, etc.). Their Rademacher complexities can be upper bounded based on the amount of training data, number of layers, number of hidden units, and their parameter norms [Neysshabur et al., 2015, Sokolic et al., 2016, Golowich et al., 2018, Truong, 2022]. Due to lemma 4 of Bartlett and Mendelson [2002], our result provides a tighter bound than the Gaussian complexity used in existing works by $\mathcal{O}(\ln NT)$ [Arora et al., 2020, Tripuraneni et al., 2020, Maurer et al., 2016].

4 Experiments

We aim to investigate the following questions regarding MTIL: (i) Can we achieve better imitation performance than behavioural cloning (BC) with less target data?, (ii) How does the number of source tasks or number of source data impact imitation performance?, (iii) How does the number of target data influence on imitation performance compared to the number of source tasks and source data?

We implement our MTIL procedure with multitask behavioural cloning (MTBC). MTBC is first pretrained with N source data from T source tasks, then finetuned with M target data (see appendix D for details), whereas BC is trained with only M target data. The analyses are done on four simulated environments: frozen lake, pendulum, cheetah, and walker. We convert the continuous action spaces into discrete action spaces described in appendix E.1. The detailed descriptions of the implementation and the varying environmental parameters are found in appendix E. For the analysis below, we determined that each environment requires $|\mathcal{D}|$ demonstrations to achieve near-expert performance when training BC from scratch. Table 1 specifies $|\mathcal{D}|$ for each environment.

Table 1: The number of demonstrations $|\mathcal{D}|$ for each environment.

Frozen Lake	Pendulum	Cheetah	Walker
500	1000	1000	500

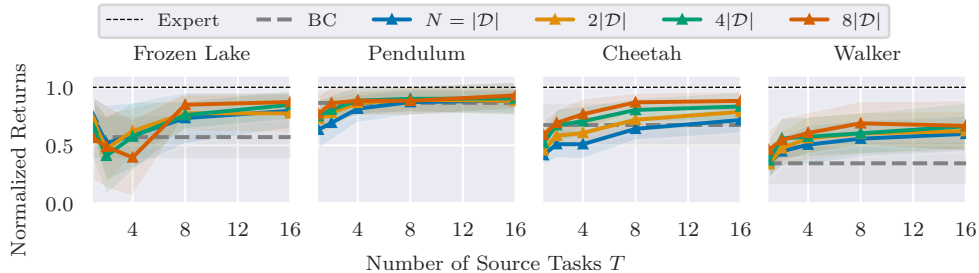


Figure 1: The performance of MTBC as we vary N and T . Each solid line colour corresponds to a particular N . The solid line corresponds to the mean and the shaded region is 1 standard error from the mean.

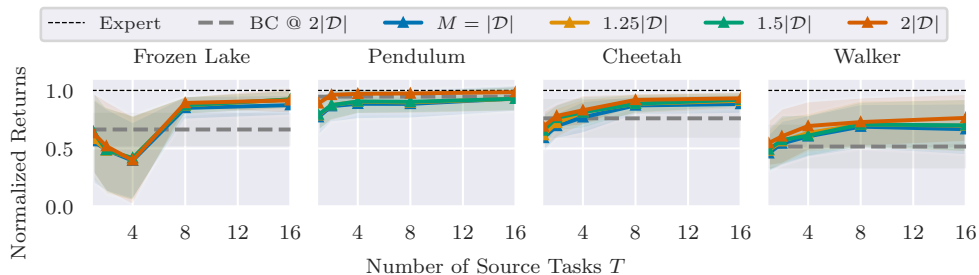


Figure 2: The performance of MTBC as we vary M and T . Each solid line colour corresponds to a particular M . The solid line corresponds to the mean and the shaded region is 1 standard error from the mean.

We first investigate questions (1) and (2). We fix the number of target data $M = |\mathcal{D}|$ while varying the number of source data N and the number of source tasks T . Figure 1 indicates that as we increase N and T , MTBC generally improves and can outperform BC with the same number of target data. For cheetah and walker, which are more complex environments, MTBC outperforms BC while MTBC and BC have comparable performance in simpler environments.

To investigate question (3), we fix the amount of source data $N = 8|\mathcal{D}|$ while varying the amount of target data M by multiples of $|\mathcal{D}|$. For comparison we also include BC trained with $2|\mathcal{D}|$ target data. Figure 2 shows that increasing the number of target data only marginally improves imitation performance. This result indicates that higher returns are mainly achieved by increasing the amount of source data N and number of source tasks T . The latter is due to the learned representation being more expressive than the task-specific mapping.

5 Conclusion and Future Directions

We have theoretically shown a statistical bound to recover the expert policy when transferring representation from a set of source tasks to a target task. This bound can be directly extended to commonly used neural networks. We have also conducted experiments to further support our theoretical analysis. In particular our experiments show that, when we pretrain a representation increasing amount source tasks and source data, one can reduce the amount of target data to imitate the expert policy, compared to training from scratch. However, we still require a practical method to quantify task diversity prior to training. Another limitation in our work is assuming that there exists one shared representation. This assumption can be relaxed by considering the meta-learning formulation where we allow the representation per task to be in a neighbourhood of some reference representation [Collins et al., 2022].

References

- Trevor Ablett, Bryan Chan, and Jonathan Kelly. Learning from guided play: Improving exploration for adversarial imitation learning with simple auxiliary tasks. *IEEE Robotics and Automation Letters*, 8(3):1263–1270, 2023.
- Sanjeev Arora, Simon S. Du, Sham M. Kakade, Yiping Luo, and Nikunj Saunshi. Provable representation learning for imitation learning via bi-level optimization. In *International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 367–376. PMLR, 2020.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta Numer.*, 30:87–201, 2021.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, et al. RT-1: robotics transformer for real-world control at scale. In *Robotics: Science and Systems*, 2023.
- Liam Collins, Aryan Mokhtari, Sewoong Oh, and Sanjay Shakkottai. MAML and ANIL provably learn representations. In *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 4238–4310. PMLR, 2022.
- Robert Dadashi, Léonard Hussenot, Damien Vincent, Sertan Girgin, Anton Raichuk, Matthieu Geist, and Olivier Pietquin. Continuous control with action quantization from demonstrations. In *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 4537–4557. PMLR, 2022.
- Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. Implementation matters in deep policy gradients: A case study on PPO and TRPO. In *International Conference on Learning Representations*, 2020.
- Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1582–1591. PMLR, 2018.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *Conference on Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 297–299. PMLR, 2018.
- Taosha Guo, Abed AlRahman Al Makdah, Vishaal Krishnan, and Fabio Pasqualetti. Imitation and transfer learning for LQG control. *IEEE Control Systems Letters*, 7:2149–2154, 2023.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1856–1865. PMLR, 2018.
- Nicklas Hansen, Zhecheng Yuan, Yanjie Ze, Tongzhou Mu, Aravind Rajeswaran, Hao Su, Huazhe Xu, and Xiaolong Wang. On pre-training for visuo-motor control: Revisiting a learning-from-scratch baseline. In *Conference on Robot Learning 2022 Workshop on Pre-training Robot Learning*, 2022.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In *ACL (1)*, pages 328–339. Association for Computational Linguistics, 2018.
- Chloe Ching-Yun Hsu, Celestine Mendler-Dünnner, and Moritz Hardt. Revisiting design choices in proximal policy optimization. *arXiv preprint arXiv:2009.10897*, 2020.

- Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. BC-Z: zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 991–1002. PMLR, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Vladimir Koltchinskii and Dmitry Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, Jacob Andreas, Igor Mordatch, Antonio Torralba, and Yuke Zhu. Pre-trained language models for interactive decision-making. In *Advances in Neural Information Processing Systems*, volume 35, pages 31199–31212, 2022.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. In *Advances in Neural Information Processing Systems*, volume 35, pages 1950–1965, 2022.
- Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning*, pages 1678–1690. PMLR, 2022.
- Andreas Maurer. A vector-contraction inequality for Rademacher complexities. In *International Conference on Algorithmic Learning Theory*, volume 9925 of *Lecture Notes in Computer Science*, pages 3–17, 2016.
- Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 1376–1401, 2015.
- Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2): 1–179, 2018.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744, 2022.
- Dean A. Pomerleau. *Alvinn: An autonomous land vehicle in a neural network*. In *Advances in Neural Information Processing Systems*, volume 1, 1988.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Tim Seyde, Igor Gilitschenski, Wilko Schwarting, Bartolomeo Stellato, Martin Riedmiller, Markus Wulfmeier, and Daniela Rus. Is Bang-Bang control all you need? solving continuous control with Bernoulli policies. In *Advances in Neural Information Processing Systems*, volume 34, pages 27209–27221, 2021.
- Tim Seyde, Peter Werner, Wilko Schwarting, Igor Gilitschenski, Martin Riedmiller, Daniela Rus, and Markus Wulfmeier. Solving continuous control via Q-learning. In *International Conference on Learning Representations*, 2022.
- Jure Sokolic, R Gyries, Guillermo Sapiro, and Miguel RD Rodrigues. Lessons from the Rademacher complexity for deep learning. *International Conference on Learning Representations*, 2016.

- Mark Towers, Jordan K. Terry, Ariel Kwiatkowski, John U. Balis, Gianluca de Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Arjun KG, Markus Krimmel, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Andrew Tan Jin Shen, and Omar G. Younis. Gymnasium, March 2023. URL <https://zenodo.org/record/8127025>.
- Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. In *Advances in Neural Information Processing Systems*, volume 33, pages 7852–7862, 2020.
- Lan V Truong. On Rademacher complexity-based generalization bounds for deep learning. *arXiv preprint arXiv:2208.04284*, 2022.
- Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. dm_control: Software and tasks for continuous control. *Software Impacts*, 6:100022, 2020. ISSN 2665-9638. doi: <https://doi.org/10.1016/j.simpa.2020.100022>. URL <https://www.sciencedirect.com/science/article/pii/S2665963820300099>.
- Martin J Wainwright. *High-dimensional Statistics: A Non-asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2021.
- Tian Xu, Ziniu Li, and Yang Yu. Error bounds of imitating policies and environments. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 15737–15749, 2020.
- Denis Yarats, Ilya Kostrikov, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*, 2020.
- Thomas T. C. K. Zhang, Katie Kang, Bruce D. Lee, Claire J. Tomlin, Sergey Levine, Stephen Tu, and Nikolai Matni. Multi-task imitation learning for linear dynamical systems. In *LADC*, volume 211 of *Proceedings of Machine Learning Research*, pages 586–599. PMLR, 2023.

A Related Work

Our work is heavily inspired by the theoretical analyses on transferring shared representations across tasks [Arora et al., 2020, Tripuraneni et al., 2020, Maurer et al., 2016]. Existing works use Gaussian complexity as the function class complexity measure when analyzing the sample complexity bound. We use Rademacher complexity instead, which allows us to directly connect our analysis with deep-learning theory [Bartlett et al., 2021, Truong, 2022]. Another line of works focuses on continuous-control problem [Guo et al., 2023, Zhang et al., 2023], where they analyze structure-specific systems. Our work is similar in that we focus on connecting the source tasks and the target task via task diversity, and providing concrete sample complexity bounds based on the task diversity. However, we remain operating in the infinite-horizon finite MDP setting with non-linear representation class.

B Theoretical Analysis

Our goal is to analyze the sample complexity of obtaining an ε -optimal policy. The analysis contains the following two steps: (1) establishing a sample complexity bound for achieving near optimal risk when transferring the representation to a new target task, and (2) bounding the policy error of the transferred policy through the transfer risk. The **transfer risk** is the excess risk on the target task τ :

$$R_{transfer}(\hat{f}_\tau, \hat{\phi}) = R_{test}(\hat{f}_\tau, \hat{\phi}) - R_{test}(f_\tau^*, \phi^*), \quad (4)$$

where $R_{phase}(\cdot) = \mathbb{E}\hat{R}_{phase}(\cdot)$, for $phase \in \{train, test\}$ is the expectation of the corresponding empirical risk. The expectation taken is over the randomness of the state-action pairs.

The first step is heavily inspired by Tripuraneni et al. [2020] but we only leverage Rademacher complexity to quantify the sample complexity bound. Our result differs in that we use only Rademacher complexity throughout the analysis due to our assumptions. The second step leverages the fact that our loss function is the KL-divergence that allows us to upper bound the policy error through the transfer risk. In general, the results may include few extra constant terms including C_Φ , $C_{\mathcal{F}}$, and B that follow from our assumptions. We provide the detailed assumptions and proofs in appendix C.

First, we compare the training error between the learned representation and true representation. This error is quantified by their similarity—Tripuraneni et al. [2020] proposed to consider the task-average prediction difference:

Definition 1 (*Task-average Representation Difference [Tripuraneni et al., 2020].*) Fix a function class \mathcal{F} , T functions $\mathbf{f} = (f_1, \dots, f_T) \in \mathcal{F}^{\otimes T}$, a loss function $\ell(\cdot, \cdot)$, and data $(s_t, a_t) \sim \mu_{\pi_t^*}$. The task-average representation difference between $\phi, \phi' \in \Phi$ is defined as

$$\bar{d}_{\mathcal{F}}(\phi'; \phi, \mathbf{f}) := \frac{1}{T} \sum_{t=1}^T \inf_{f' \in \mathcal{F}} \mathbb{E}_{(s_t, a_t)} [\ell((f' \circ \phi')(s_t), a_t) - \ell((f_t \circ \phi)(s_t), a_t)]. \quad (5)$$

The error between the learned representation $\hat{\phi}$ and the true representation ϕ^* can be upper bounded by the task-average difference—the bound is dependent on the number of source tasks T , the amount of training data per task N , and the Rademacher complexity of the representation class.

Theorem 2 (*Learned Representation Risk Bound.*) Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1. Let ϕ^* be the true representation and $\mathbf{f}^* = (f_1^*, \dots, f_T^*) \in \mathcal{F}^{\otimes T}$ be the T true task-specific mappings. Under some assumptions, we have that with probability at least $1 - \delta$,

$$\bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*) \leq 8\sqrt{2}C_{\mathcal{F}}\mathfrak{R}_{NT}(\Phi) + 2B\sqrt{\frac{\log(2/\delta)}{2NT}}.$$

We now consider the test error when using the learned representation on the target task. This requires a different notion of similarity—Tripuraneni et al. [2020] proposed to consider the worst-case task-specific mapping in \mathcal{F} .

Definition 2 (*Worst-case Representation Difference [Tripuraneni et al., 2020].*) Fix a task τ , a function class \mathcal{F} , a loss function $\ell(\cdot, \cdot)$, and data $(s, a) \sim \mu_{\pi_\tau^*}$. The worst-case representation difference between $\phi, \phi' \in \Phi$ is defined as

$$d_{\tau, \mathcal{F}}(\phi'; \phi) := \sup_{f \in \mathcal{F}} \inf_{f' \in \mathcal{F}} \mathbb{E}_{(s, a)} [\ell((f' \circ \phi')(s), a) - \ell((f \circ \phi)(s), a)]. \quad (6)$$

We can then use definition 2 to upper bound the generalization error for the transfer phase ERM estimator:

Theorem 3 (*Transfer Risk Bound.*) Let \hat{f}_τ be the ERM of \hat{R}_{test} defined in equation 2 with some fixed $\hat{\phi} \in \Phi$. Under some assumptions, then with probability $1 - \delta$,

$$R_{transfer}(\hat{f}_\tau, \hat{\phi}) \leq 8C_{\mathcal{F}}C_{\Phi}\sqrt{\frac{|\mathcal{A}|}{M}} + 2B\sqrt{\frac{\log(2/\delta)}{2M}} + d_{\tau, \mathcal{F}}(\hat{\phi}; \phi^*). \quad (7)$$

We connect the training error of the learned representation $\hat{\phi}$ with the transfer risk via the notion of task diversity similar to the one defined by Tripuraneni et al. [2020]:

Definition 3 (σ -diversity.) Let $\sigma > 0$ and fix a task τ . Fix a function class \mathcal{F} , T functions $\mathbf{f} = (f_1, \dots, f_T) \in \mathcal{F}^{\otimes T}$. The T tasks are σ -diverse for representation ϕ , if for all $\phi' \in \Phi$, we have that $d_{\tau, \mathcal{F}}(\phi'; \phi) \leq \bar{d}_{\mathcal{F}}(\phi'; \phi, \mathbf{f})/\sigma$.

Intuitively, if the T source tasks differ too much from the new task, then inequality only holds with small σ . In other words, there is little diversity. Another perspective is that any ϕ' is overfitted to the T tasks and is unable to generalize to the new task. Consequently, combining theorems 2 and 3, we get the following sample complexity bound:

Corollary 1 (*Learned Representation Transfer Risk Bound.*) Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1 and let \hat{f}_τ be the ERM of \hat{R}_{test} defined in equation 2 by fixing $\hat{\phi}$. Suppose the source tasks are σ -diverse. Under some assumptions, we have that with probability $1 - 2\delta$, the transfer risk $R_{transfer}(\hat{f}_\tau, \hat{\phi})$ is upper bounded by:

$$\mathcal{O} \left(C_{\mathcal{F}} C_{\Phi} \sqrt{\frac{|\mathcal{A}|}{M}} + B \sqrt{\frac{\log(2/\delta)}{M}} + \frac{1}{\sigma} \left(C_{\mathcal{F}} \mathfrak{R}_{NT}(\Phi) + B \sqrt{\frac{\log(2/\delta)}{NT}} \right) \right). \quad (8)$$

Corollary 1 tells us that so long as we can bound the task-average representation difference, we can achieve ε -error with sufficiently large number of demonstration data in the target task. Furthermore, if the source tasks are not diverse, then we will need more samples from them to reduce the error. Finally, since the log loss is the KL-divergence with deterministic experts, we can use the result from Xu et al. [2020] to obtain the policy error of the policy induced by the transfer-learning procedure:

Theorem 4 (*Policy Error Bound [Xu et al., 2020].*) Given any two policies π, π' with $\mathbb{E}_{s \sim \nu_\pi} [D_{KL}(\pi(s) \parallel \pi'(s))] < \varepsilon$, we have that $\|v^\pi - v^{\pi'}\|_\infty \leq \frac{2\sqrt{2}}{(1-\gamma)^2} \sqrt{\varepsilon}$.

Thus, by theorem 1, for specific representation class, we can upper bound their Rademacher complexities and retrieve a concrete sample complexity bound to achieve ε -optimal policies with high probability.

C Detailed Proofs

In this section we provide the proofs for the theorems in appendix B. We first list out our assumptions (appendix C.1), then a few known lemmas and definitions used (appendix C.2), and finally the proofs of our theorems (appendices C.3–C.6).

C.1 Assumptions

In the analysis, we offload the softmax function to the loss, then we can define a new loss function that can be used for analysis. That is, let $x \in \mathbb{R}^D$, then we define the log-softmax loss to be

$$\ell(x, a) = -\log(\text{softmax}_a(x)), \quad (9)$$

where $\text{softmax}_a(x)$ corresponds to the a 'th component of $\text{softmax}(x)$. We further make the following assumptions for the analysis.

Assumption 1 (*Bounded Representation.*) The representation $\phi \in \Phi$ is bounded in ℓ_2 -norm: $\Phi \subseteq \{\phi : \mathcal{S} \rightarrow \mathbb{R}^D \mid \|\phi(s)\|_2 \leq C_\Phi, \forall s \in \mathcal{S}\}$.

Assumption 2 (*Linear Task-specific Mapping and Bounded Parameters.*) The task-specific mapping is linear and is bounded: $\mathcal{F} = \{f : \mathbb{R}^D \rightarrow \mathbb{R}^{|\mathcal{A}|} \mid f = Wx, W \in \mathbb{R}^{|\mathcal{A}| \times D}, \|W\|_F \leq C_{\mathcal{F}}, x \in \mathbb{R}^D\}$, where $\|\cdot\|_F$ is the Frobenius norm

Assumption 3 (*Deterministic Expert Policies.*) The expert policy π_τ^* for each task τ is deterministic and can be written as $\pi_\tau^*(s) = \mathbf{1}_{a=\arg \max_{a' \in \mathcal{A}} (f_\tau^* \circ \phi^*)(s)}$ for all $s \in \mathcal{S}$.

Assumption 4 (*Shared Representation.*) There is a representation ϕ^* such that for every task τ , there exists a task-specific mapping f_τ^* such that the discounted state-action stationary distribution is $\mu_{\pi_\tau^*}(s, a)$.

Assumption 5 (*Realizability.*) The true shared representation ϕ^* is contained in Φ . Additionally, for some fixed $\zeta < 1/2$, we have that for all tasks τ , there exists a task-specific mapping $f_\tau \in \mathcal{F}$ such that for all $s \in \mathcal{S}$,

$$\pi^{f_\tau, \phi^*}(a^* | s) \geq 1 - \zeta,$$

where $a^* = \arg \max_{a \in \mathcal{A}} \pi_\tau^*(s)$.

Note that for any infinite-horizon MDPs, there always exists a deterministic optimal policy, thus assumption 3 is often reasonable to obtain. With the recent successes in foundation models in various applications [Bommasani et al., 2021, Wei et al., 2021, Ouyang et al., 2022], both assumptions 4 and 5 may be reasonable. We note that our results also apply when replacing assumptions 3 and 5 with the standard realizability assumption on stochastic expert policies. Finally, assumptions 1 and 2 are standard regularity conditions in statistical learning theory—we further note that consequently the composed mapping $f \circ \phi$ is bounded: $\sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} |(f \circ \phi)(s)| \leq C_S$, for any $f \in \mathcal{F}$ and $\phi \in \Phi$.

C.2 Useful Definitions and Results

Definition 4 (*Rademacher Complexity*.) For a vector-valued function $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathbb{R}^K\}$, and N data points $X = (x_1, \dots, x_N)$, where $x_n \in \mathcal{X}$ for $n \in [N]$, the empirical Rademacher complexity of \mathcal{F} is defined as

$$\hat{\mathfrak{R}}_X(\mathcal{F}) = \mathbb{E}_\varepsilon \left[\sup_{f \in \mathcal{F}} \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \varepsilon_{n,k} f_k(x_n) \right],$$

where $\varepsilon_{n,k}$ are sampled i.i.d. from the Rademacher random variable and $f_k(\cdot)$ is the k 'th component of $f(\cdot)$. Fix a data distribution \mathcal{D}_X over \mathcal{X} . The corresponding Rademacher complexity of \mathcal{F} is defined as $\mathfrak{R}_M(\mathcal{F}) = \mathbb{E}_X [\hat{\mathfrak{R}}_X(\mathcal{F})]$, where the expectation is taken over the distribution \mathcal{D}_X . In the case of $K = 1$, we recover the scalar Rademacher complexity.

Intuitively, the Rademacher complexity of \mathcal{F} measures the expressiveness of \mathcal{F} over all datasets \mathbf{X} .

Proposition 1 With assumptions 1 and 2, the log-softmax loss, defined in equation 9, is bounded and Lipschitz (continuous) in its first argument.

Proof: For the first claim, we first fix $(s, a) \in \mathcal{S} \times \mathcal{A}$. Let $\phi_s = \phi(s) \in \mathbb{R}^D$ and $x_s = W\phi_s$. Then, we have that

$$-\log(\text{softmax}_a(x_s)) = -(W\phi_s)(a) + \log \sum_{a' \in \mathcal{A}} \exp[(W\phi_s)(a')].$$

Let $a^* = \arg \max_{a \in \mathcal{A}} W\phi_s$ and overload the notation $(W\phi_s)(a) = \phi_s^\top W^\top a$, where a is a one-hot vector with non-zero at the a 'th entry. Then, the second term can be upper bounded:

$$\begin{aligned} \log \sum_{a' \in \mathcal{A}} \exp[\phi_s^\top W^\top a'] &\leq \log \sum_{a' \in \mathcal{A}} \exp[\phi_s^\top W^\top a^*] \\ &= \log|\mathcal{A}| + \phi_s^\top W^\top a^*. \end{aligned}$$

Thus,

$$\begin{aligned} -\log(\text{softmax}_a(x_s)) &\leq -\phi_s^\top W^\top a + \log|\mathcal{A}| + \phi_s^\top W^\top a^* \\ &\leq \log|\mathcal{A}| + 2\phi_s^\top W^\top a^*. \end{aligned}$$

Taking the supremum norm over $\mathcal{S} \times \mathcal{A}$, we have that

$$\begin{aligned} \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left| -\log(\text{softmax}_a(x_s)) \right| &\leq \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left| \log|\mathcal{A}| + 2\phi_s^\top W^\top a^* \right| \\ &\leq \log|\mathcal{A}| + 2 \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left| \phi_s^\top W^\top a^* \right| \\ &\leq \log|\mathcal{A}| + 2C_\Phi C_{\mathcal{F}}, \end{aligned}$$

where the last inequality follows upper bounding the second term through Hölder's inequality, setting both norms to be the 2-norm. This verifies the boundedness of the log-softmax loss.

For the second claim, we can bound the gradient of the log-softmax loss with respect to the first argument and apply mean-value theorem. Let us first consider the partial derivatives of $\ell(x, a)$. Let x_i be the i 'th component of x . For any $x \in \mathbb{R}^D$, we have that

$$\frac{\partial}{\partial x_i} \ell(x, a) \Big|_{i=a} = -\frac{\sum_{a' \neq i} \exp x^\top a'}{\sum_{a' \in \mathcal{A}} \exp x^\top a'} \quad \frac{\partial}{\partial x_i} \ell(x, a) \Big|_{i \neq a} = \frac{\exp x^\top i}{\sum_{a' \in \mathcal{A}} \exp x^\top a'},$$

where we overload the notation and represent $a' \in \mathcal{A}$ as a one-hot vector. Then, consider the ℓ_2 -norm of the gradient, we have that:

$$\begin{aligned} \|\nabla_x \ell(x, a)\|_2^2 &= \left(-\frac{\sum_{a' \neq a} \exp x^\top a'}{\sum_{a' \in \mathcal{A}} \exp x^\top a'} \right)^2 + \frac{\sum_{a' \neq a} (\exp x^\top a')^2}{(\sum_{a' \in \mathcal{A}} \exp x^\top a')^2} \\ &\leq 2 \left(\frac{\sum_{a' \neq a} \exp x^\top a'}{\sum_{a' \in \mathcal{A}} \exp x^\top a'} \right)^2 \\ &= 2, \end{aligned}$$

where the first inequality follows from Jensen's inequality. Consequently, by mean-value theorem, we have that $|\ell(x, a) - \ell(y, a)| \leq \sqrt{2} \|x - y\|_2$, for any $x, y \in \mathbb{R}^D$, verifying the Lipschitzness of the log-softmax loss in its first argument. \square

Proposition 2 *Let ℓ be the log-softmax loss defined in equation 9 and $f \in \mathcal{F}$. Under assumptions 1 and 2, the function $h_a(\phi(s)) = \ell(f(\phi(s)), a)$ is Lipschitz in $\phi(s)$, for any $s \in \mathcal{S}$.*

Proof: We first note that $h_a(\phi_s)$ can be written as $\log \sum_{b \in \mathcal{A}} \exp(b^\top W \phi_s) - a^\top W \phi_s$, where $\phi_s = \phi(s) \in \mathbb{R}^D$, and we overload the notation and write a, b as the one-hot vectors.

For the first term, the ℓ_2 -norm of its gradient with respect to ϕ_s can be upper bounded by $C_{\mathcal{F}}$:

$$\begin{aligned} \left\| \nabla_{\phi_s} \log \sum_{b \in \mathcal{A}} \exp(b^\top W \phi_s) \right\|_2 &= \left\| \frac{\sum_b \exp(b^\top W \phi_s) W^\top b}{\sum_b \exp(b^\top W \phi_s)} \right\|_2 \\ &= \|W^\top \text{softmax}(W \phi_s)\|_2 \\ &\leq \|W\|_F \|\text{softmax}(W \phi_s)\|_2 \\ &\leq C_{\mathcal{F}}. \end{aligned}$$

For the second term, the ℓ_2 -norm of its gradient with respect to ϕ_s can be upper bounded by $C_{\mathcal{F}}$:

$$\|\nabla_{\phi_s} a^\top W \phi_s\|_2 = \|W^\top a\|_2 \leq \|W\|_F \|a\|_2 \leq C_{\mathcal{F}}.$$

Since the first term is convex and the second term is linear, $h_a(\phi_s)$ is $2C_{\mathcal{F}}$ -Lipschitz. \square

Definition 5 (*Bounded Difference Property.*) *The function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ satisfies the bounded difference inequality with positive constants (L_1, \dots, L_N) if, for each $n \in [N]$,*

$$\sup_{x_1, \dots, x_N, x'_n \in \mathbb{R}} |f(x_1, \dots, x_n, \dots, x_N) - f(x_1, \dots, x'_n, \dots, x_N)| \leq L_n. \quad (10)$$

Lemma 1 (*McDiarmid's Inequality/Bounded Difference Inequality.*) *Fix a data distribution $\mathcal{D}_{\mathcal{X}}$. Suppose f satisfies the bounded difference property defined in equation 10, with positive constants L_1, \dots, L_N and that $X = (X_1, \dots, X_N)$ is drawn independently from $\mathcal{D}_{\mathcal{X}}$. Then*

$$\mathbb{P}[|f(X) - \mathbb{E}f(X)| \geq t] \leq 2 \exp\left(\frac{-2t^2}{\sum_{n=1}^N L_n^2}\right), \forall t \geq 0.$$

Proof: We refer the readers to corollary 2.21 of Wainwright [2019] for the proof. \square

Theorem 5 (*Rademacher Complexity Bound.*) *Fix a data distribution $\mathcal{D}_{\mathcal{X}}$ and parameter $\delta \in (0, 1)$. Suppose $\mathcal{F} \subseteq \{f : \mathcal{X} \rightarrow [0, B]\}$ and $X = (X_1, \dots, X_N)$ is drawn i.i.d. from $\mathcal{D}_{\mathcal{X}}$. Then with probability at least $1 - \delta$ over the draw of X , for any function $f \in \mathcal{F}$,*

$$\left| \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}} [f(x)] - \frac{1}{n} \sum_{n=1}^N f(X_n) \right| \leq 2\mathfrak{R}_N(\mathcal{F}) + B \sqrt{\frac{\log(2/\delta)}{2n}}. \quad (11)$$

Proof: This result follows closely to the proof of Theorem 10 in Koltchinskii and Panchenko [2002], where the only modifications come from applying B -boundedness of f when applying McDiarmid's inequality (i.e. $L_n \leq \frac{B}{N}, \forall n \in [N]$.) \square

Theorem 6 (Vector-contraction Inequality [Maurer, 2016].) Let $x_1, \dots, x_N \in \mathcal{X}$, $\mathcal{F} \subseteq \{f : \mathcal{X} \rightarrow \mathbb{R}^D\}$ be a class of functions, and $h_n : \mathbb{R}^D \rightarrow \mathbb{R}$ to be L -Lipschitz, for all $n \in [N]$. Then

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{n=1}^N \varepsilon_n h_n(f(x_n)) \right] \leq \sqrt{2} L \mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{d=1}^D \sum_{n=1}^N \varepsilon_{d,n} f_d(x_n) \right],$$

where $\varepsilon_n, \varepsilon_{d,n}$ are i.i.d. sequences of Rademacher variables, and $f_d(x_n)$ is the d 'th component of $f(x_n)$.

Proof: We refer the readers to theorem 3 and corollary 4 of Maurer [2016] for the proof. \square

C.3 Proof of Theorem 2

Theorem 2 states the following:

Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1. Let ϕ^* be the true representation and $\mathbf{f}^* = (f_1^*, \dots, f_T^*) \in \mathcal{F}^{\otimes T}$ be the T true task-specific mappings. If the assumptions 1 to 5 hold, then with probability $1 - \delta$,

$$\bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*) \leq 8\sqrt{2} C_{\mathcal{F}} \mathfrak{R}_{NT}(\Phi) + 2B \sqrt{\frac{\log(2/\delta)}{2NT}}.$$

Proof: The proof follows closely from the analysis of Tripuraneni et al. [2020]. First, recall that $\hat{\phi}, \hat{\mathbf{f}}$ are respectively the ERMs of equation 1 and equation 2. For any T task-specific mappings $\mathbf{f}' = (f'_1, \dots, f'_T)$ and a representation ϕ' . Let $\mathbf{f}^* = (f_1^*, \dots, f_T^*)$ be the true task-specific mappings and ϕ^* be the true shared representation. We define the centered training risk and its empirical counterpart respectively as:

$$\begin{aligned} L(\mathbf{f}', \phi', \mathbf{f}^*, \phi^*) &= \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{(s_t, a_t)} \left[\ell(\pi^{f'_t, \phi'}(s_t), a_t) - \ell(\pi^{f_t^*, \phi^*}(s_t), a_t) \right], \\ \hat{L}(\mathbf{f}', \phi', \mathbf{f}^*, \phi^*) &= \frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N \left(\ell(\pi^{f'_t, \phi'}(s_{t,n}), a_{t,n}) - \mathbb{E}_{(s_t, a_t)} \left[\ell(\pi^{f_t^*, \phi^*}(s_t), a_t) \right] \right). \end{aligned}$$

Define $\tilde{\mathbf{f}} = \frac{1}{T} \sum_{t=1}^T \arg \inf_{f_t \in \mathcal{F}} \mathbb{E}_{s_t, a_t} \left[\ell(\pi^{f_t, \hat{\phi}}(s_t), a_t) - \ell(\pi^{f_t^*, \phi^*}(s_t), a_t) \right]$ to be the minimizer of the centered training risk by fixing $\hat{\phi}$. Then, we have that $L(\tilde{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) = \bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*)$. Now, we aim to upper bound $\bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*)$ through the difference in the centered training risk:

$$\begin{aligned} \bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*) &= L(\tilde{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - L(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) \\ &= \underbrace{L(\tilde{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - L(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*)}_{\leq 0} + L(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - L(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*), \end{aligned}$$

where the first difference is non-positive by definition of $\tilde{\mathbf{f}}$. Thus, it remains to bound the second difference, which can be done via standard risk decomposition. First, recall that:

$$\begin{aligned} \hat{R}_{train}(\mathbf{f}, \phi) &= \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N \ell(\pi^{f_t, \phi}(s_{t,n}), a_{t,n}), \\ R_{train}(\mathbf{f}, \phi) &= \mathbb{E} \left[\hat{R}_{train}(\mathbf{f}, \phi) \right]. \end{aligned}$$

Then, we have that

$$\begin{aligned}
L(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - L(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) &= L(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - \hat{L}(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) \\
&\quad + \hat{L}(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - \hat{L}(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) \\
&\quad + \hat{L}(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) - L(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) \\
&\stackrel{(i)}{\leq} L(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) - \hat{L}(\hat{\mathbf{f}}, \hat{\phi}, \mathbf{f}^*, \phi^*) \\
&\quad + \hat{L}(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) - L(\mathbf{f}^*, \phi^*, \mathbf{f}^*, \phi^*) \\
&\leq 2 \sup_{\mathbf{f} \in \mathcal{F}^{\otimes T}, \phi \in \Phi} |R_{train}(\mathbf{f}, \phi) - R_{train}(\mathbf{f}^*, \phi^*)| \\
&\stackrel{(ii)}{\leq} 2 \left(2\mathfrak{R}_{NT}(\ell \circ \mathcal{F} \circ \Phi) + B \sqrt{\frac{\log(2/\delta)}{2NT}} \right) \text{ w.p. } 1 - \delta,
\end{aligned}$$

where in (i) the second difference is non-positive due to assumptions 3 and 5; (ii) is a result of theorem 5, where we used proposition 1.

We now bound the Rademacher complexity term $\mathfrak{R}_{NT}(\ell \circ \mathcal{F} \circ \Phi)$. Note that $\ell \circ \mathcal{F}$ is $2C_{\mathcal{F}}$ -Lipschitz by proposition 2, thus by theorem 6, we have that $\mathfrak{R}_{NT}(\ell \circ \mathcal{F} \circ \Phi) \leq 2\sqrt{2}C_{\mathcal{F}}\mathfrak{R}_{NT}(\Phi)$. Substituting this upper bound into the above, we get that with probability at least $1 - \delta$,

$$\bar{d}_{\mathcal{F}}(\hat{\phi}; \phi^*, \mathbf{f}^*) \leq 8\sqrt{2}C_{\mathcal{F}}\mathfrak{R}_{NT}(\Phi) + 2B\sqrt{\frac{\log(2/\delta)}{2NT}}.$$

□

C.4 Proof of Theorem 3

Theorem 3 states the following:

let \hat{f}_{τ} be the ERM of \hat{R}_{test} defined in equation 2 with some fixed $\hat{\phi} \in \Phi$. If the assumptions 1 to 5 hold, then with probability $1 - \delta$,

$$R_{transfer}(\hat{f}_{\tau}, \hat{\phi}) \leq 8C_{\mathcal{F}}C_{\Phi}\sqrt{\frac{|\mathcal{A}|}{M}} + 2B\sqrt{\frac{\log(2/\delta)}{2M}} + d_{\tau, \mathcal{F}}(\hat{\phi}; \phi^*).$$

Proof: The proof follows closely from the analysis of Tripuraneni et al. [2020]. First recall that:

$$\begin{aligned}
R_{transfer}(\hat{f}_{\tau}, \hat{\phi}) &= R_{test}(\hat{f}_{\tau}, \hat{\phi}) - R_{test}(f_{\tau}^*, \phi^*), \\
R_{test}(f, \phi) &= \mathbb{E} \left[\hat{R}_{test}(f, \phi) \right], \\
\hat{R}_{test}(f, \phi) &= \frac{1}{M} \sum_{m=1}^M \ell((f \circ \phi)(s_m), a_m), \\
\hat{f} &= \arg \min_{f \in \mathcal{F}} \hat{R}_{test}(f, \hat{\phi}).
\end{aligned}$$

Consider the minimizer of the test risk given the estimated representation $\hat{\phi}$: $\hat{f}^* = \arg \min_{f \in \mathcal{F}} R_{test}(f, \hat{\phi})$. Then, we have that

$$\begin{aligned}
&R_{test}(\hat{f}_{\tau}, \hat{\phi}) - R_{test}(f_{\tau}^*, \phi^*) \\
&= \underbrace{R_{test}(\hat{f}_{\tau}, \hat{\phi}) - R_{test}(\hat{f}_{\tau}^*, \hat{\phi})}_{(a)} + \underbrace{R_{test}(\hat{f}_{\tau}^*, \hat{\phi}) - R_{test}(f_{\tau}^*, \phi^*)}_{(b)}.
\end{aligned}$$

We bound the term (a) via standard risk decomposition:

$$\begin{aligned}
& R_{test}(\hat{f}_\tau, \hat{\phi}) - R_{test}(\hat{f}_\tau^*, \hat{\phi}) \\
&= R_{test}(\hat{f}_\tau, \hat{\phi}) - \hat{R}_{test}(\hat{f}_\tau, \hat{\phi}) + \hat{R}_{test}(\hat{f}_\tau, \hat{\phi}) - \hat{R}_{test}(\hat{f}_\tau^*, \hat{\phi}) + \hat{R}_{test}(\hat{f}_\tau^*, \hat{\phi}) - R_{test}(\hat{f}_\tau^*, \hat{\phi}) \\
&\stackrel{(i)}{\leq} R_{test}(\hat{f}_\tau, \hat{\phi}) - \hat{R}_{test}(\hat{f}_\tau, \hat{\phi}) + \hat{R}_{test}(\hat{f}_\tau^*, \hat{\phi}) - R_{test}(\hat{f}_\tau^*, \hat{\phi}) \\
&\leq 2 \sup_{f \in \mathcal{F}} |R_{test}(f, \hat{\phi}) - \hat{R}_{test}(f, \hat{\phi})| \\
&\stackrel{(ii)}{\leq} 2 \left(2\mathfrak{R}_M(\ell \circ \mathcal{F}) + B \sqrt{\frac{\log(2/\delta)}{2M}} \right) \text{ w.p. } 1 - \delta,
\end{aligned}$$

where (i) follows from the fact that \hat{f} is the empirical test risk minimizer by definition—clearly $\hat{R}_{test}(\hat{f}_\tau, \hat{\phi}) - \hat{R}_{test}(\hat{f}_\tau^*, \hat{\phi}) \leq 0$; (ii) is a result of theorem 5.

We now bound the Rademacher complexity term $\mathfrak{R}_M(\ell \circ \mathcal{F})$. Note that ℓ is $\sqrt{2}$ -Lipschitz in its first argument for every $a \in \mathcal{A}$ from proposition 1. By theorem 6, we have that $\mathfrak{R}_M(\ell \circ \mathcal{F}) \leq 2\mathfrak{R}_M(\mathcal{F})$. Thus, substituting this upper bound into the above, we get that with probability at least $1 - \delta$,

$$R_{test}(\hat{f}_\tau, \hat{\phi}) - R_{test}(\hat{f}_\tau^*, \hat{\phi}) \leq 8\mathfrak{R}_M(\mathcal{F}) + 2B \sqrt{\frac{\log(2/\delta)}{2M}}.$$

Now, by assumption 2, \mathcal{F} is the class of linear functions with bounded Frobenius norm. Consider the empirical Rademacher complexity $\hat{\mathfrak{R}}_S(\mathcal{F})$:

$$\begin{aligned}
\hat{\mathfrak{R}}_S(\mathcal{F}) &= \frac{1}{M} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{a=1}^{|\mathcal{A}|} \sum_{m=1}^M \varepsilon_{am} f_a(\phi_m) \right] \\
&\stackrel{(i)}{\leq} \frac{1}{M} C_{\mathcal{F}} \sqrt{|\mathcal{A}| \sum_{m=1}^M \|\phi_m\|_2^2} \\
&\stackrel{(ii)}{\leq} \frac{1}{M} C_{\mathcal{F}} \sqrt{\frac{|\mathcal{A}| C_{\Phi}^2}{M}} \\
&\leq C_{\mathcal{F}} C_{\Phi} \sqrt{\frac{|\mathcal{A}|}{M}},
\end{aligned}$$

where (i) follows from section 4.2 of Maurer [2016] and applying the Frobenius norm to the inequality; (ii) follows from the fact that $\phi_m = \phi(s_m)$, $s_m \sim \nu_\pi$, thus $\|\phi_m\|_2 \leq C_{\Phi}$ by assumption 1. Consequently, taking the expectation over X we get $\mathfrak{R}_M(\mathcal{F}) \leq C_{\mathcal{F}} C_{\Phi} \sqrt{|\mathcal{A}|/M}$.

We can bound the term (b) as follows:

$$\begin{aligned}
& R_{test}(\hat{f}_\tau^*, \hat{\phi}) - R_{test}(f_\tau^*, \phi^*) \\
&= \inf_{f \in \mathcal{F}} R_{test}(f, \hat{\phi}) - \inf_{f' \in \mathcal{F}} R_{test}(f', \phi^*) \\
&\leq \sup_{f' \in \mathcal{F}} \inf_{f \in \mathcal{F}} R_{test}(\hat{f}_\tau^*, \hat{\phi}) - R_{test}(f', \phi^*) \\
&= d_{\mathcal{F}}(\hat{\phi}; \phi^*).
\end{aligned}$$

Finally, by comining both (a) and (b) terms, we conclude that with probability at least $1 - \delta$,

$$R_{transfer}(\hat{f}_\tau, \hat{\phi}) \leq 8C_{\mathcal{F}} C_{\Phi} \sqrt{\frac{|\mathcal{A}|}{M}} + 2B \sqrt{\frac{\log(2/\delta)}{2M}} + d_{\tau, \mathcal{F}}(\hat{\phi}; \phi^*).$$

□

C.5 Proof of Corollary 1

Corollary 1 states the following:

Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1 and let \hat{f}_τ be the ERM of \hat{R}_{test} defined in equation 2 by fixing $\hat{\phi}$. Suppose the source tasks are σ -diverse. If the assumptions 1 to 5 hold, then with probability $1 - 2\delta$, $R_{transfer}(\hat{f}_\tau, \hat{\phi})$ is upper bounded by:

$$\mathcal{O} \left(C_{\mathcal{F}} C_{\Phi} \sqrt{\frac{|\mathcal{A}|}{M}} + B \sqrt{\frac{\log(2/\delta)}{M}} + \frac{1}{\sigma} \left(C_{\mathcal{F}} \mathfrak{R}_{NT}(\Phi) + B \sqrt{\frac{\log(2/\delta)}{NT}} \right) \right).$$

Proof: Set the probability of bad events of theorems 2 and 3 to be δ each. We obtain the desired results by taking complement of the union bound over the bad events and merging the terms. \square

C.6 Proof of Theorem 1

Theorem 1 states the following:

Let $\hat{\phi}$ be the ERM of \hat{R}_{train} defined in equation 1 and let \hat{f}_τ be the ERM of \hat{R}_{test} defined in equation 2 by fixing $\hat{\phi}$. Suppose the source tasks are σ -diverse. Under assumptions 1 to 5, we have that with probability $1 - 2\delta$,

$$\|v^{\pi_\tau^*} - v^{\text{softmax}(\hat{f}_\tau \circ \hat{\phi})}\|_\infty \leq \frac{2\sqrt{2}}{(1-\gamma)^2} \sqrt{\varepsilon_{gen} + 2\zeta}, \quad (12)$$

where ε_{gen} is the RHS of equation 8.

Proof: Our goal is to apply theorem 4. Let $\pi = \pi^*$ and $\pi' = \text{softmax}(\hat{f}_\tau \circ \hat{\phi})$, we need to bound their expected Kullback–Leibler (KL) divergence. To accomplish this, note that theorem 3 essentially means that we can bound the test risk:

$$R_{test}(\hat{f}_\tau, \hat{\phi}) \leq \varepsilon_{gen} + \varepsilon_{best}, \quad (13)$$

where ε_{gen} is the RHS of equation 8 and $\varepsilon_{best} = R_{test}(f^*, \phi^*)$. Indeed, $R_{test}(\hat{f}_\tau, \hat{\phi})$ is the expected KL divergence between $\pi_\tau^*(s)$ and $\text{softmax}(\hat{f}_\tau \circ \hat{\phi})(s)$:

$$\begin{aligned} R_{test}(\hat{f}_\tau, \hat{\phi}) &= \mathbb{E}_{(s,a) \sim \mu_{\pi_\tau^*}} \left[-\log \text{softmax}_a(\hat{f}_\tau \circ \hat{\phi})(s) \right] \\ &= \mathbb{E}_{s \sim \nu_{\pi^*}} \left[D_{KL}(\pi_\tau^*(s) \parallel \text{softmax}(\hat{f}_\tau \circ \hat{\phi})(s)) \right]. \end{aligned}$$

Finally, note that assumptions 3 and 5 imply that we can assume $f^* \in \mathcal{F}$, thus we have that:

$$\begin{aligned} \varepsilon_{best} &= R_{test}(f^*, \phi^*) \\ &\leq \min_{f \in \mathcal{F}} R_{test}(f, \phi^*) \\ &\leq \mathbb{E}_{(s,a) \sim \mu_{\pi_\tau^*}} \left[-\log(1 - \zeta) \right] \\ &\leq 2\zeta, \end{aligned}$$

where the third lines come from assumption 5 and last inequality comes from $-\log(1-x) \leq 2x$ for $x \leq 1/2$. Finally, we upper bound ε_{gen} using corollary 1 and get the desired result. \square

D Algorithm Details

In this section we provide details on the implementation of the multitask behaviour cloning (MTBC) algorithm used to obtain the results highlighted in section 4. Recall that in the training phase we aim to obtain a shared representation $\hat{\phi}$ by minimizing \hat{R}_{train} , as described in equation 1. As proposed by [Arora et al., 2020], this can be done by solving the following bi-level optimization objective:

$$\hat{\phi} = \arg \min_{\phi} \frac{1}{T} \sum_{t=1}^T \min_{\pi \in \Pi^\phi} \frac{1}{N} \sum_{n=1}^N \ell(\pi(s_{t,n}), a_{t,n}), \quad (14)$$

Algorithm 1 Multitask Behavioural Cloning (MTBC): Training Phase

- 1: **input:** Number of epochs K , T expert datasets $\{\mathcal{D}_t\}_{t=1}^T$, where $|\mathcal{D}_t| = N$, and learning rates η_ϕ and η_W .
- 2: Initialize parameters for each of the T task-specific mappings: $\{W_t^{(0)}\}_{t=1}^T$.
- 3: Initialize shared representation parameters $\phi^{(0)}$.
- 4: **for** $k = 1, \dots, K$ **do**
- 5: Compute loss:

$$\mathcal{L} \left(\phi^{(k)}, \{W_t^{(k)}\}_{t=1}^T \right) = \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N \ell \left(\pi^{W_t^{(k-1)}, \phi^{(k-1)}}(s_{t,n}), a_{t,n} \right). \quad (15)$$

- 6: Update parameters:

$$\begin{aligned} \phi^{(k)} &= \phi^{(k-1)} - \eta_\phi \nabla_\phi \mathcal{L} \left(\phi^{(k)}, \{W_t^{(k)}\}_{t=1}^T \right), \\ W_t^{(k)} &= W_t^{(k-1)} - \eta_W \nabla_{W_t} \mathcal{L} \left(\phi^{(k)}, \{W_t^{(k)}\}_{t=1}^T \right), \forall t \in [T]. \end{aligned}$$

- 7: **return** ϕ .
-

Algorithm 2 Multitask Behavioural Cloning (MTBC): Transfer Phase

- 1: **input:** Representation ϕ , number of epochs K , expert dataset \mathcal{D} , where $|\mathcal{D}| = M$, and learning rate η_W .
 - 2: Initialize parameters of the task-specific mapping: $W_\tau^{(0)}$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: Compute loss: $\mathcal{L} \left(\phi, W_\tau^{(k)} \right) = \frac{1}{M} \sum_{m=1}^M \ell \left(\pi^{W_\tau^{(k-1)}, \phi}(s_m), a_m \right)$
 - 5: Update parameters: $W_\tau^{(k)} = W_\tau^{(k-1)} - \eta_W \nabla_W \mathcal{L} \left(\phi, W_\tau^{(k)} \right)$.
 - 6: **return** $W^{(K)}$.
-

where $(s_{t,n}, a_{t,n})$ is the n 'th state-action pair in the t 'th expert dataset and $\ell : \Delta^A \times \mathcal{A} \rightarrow [0, \infty)$ is the loss function. MTBC optimizes equation 14 through a gradient-based approach on a joint objective equation 15 (see algorithm 1). During the transfer phase, MTBC fixes the representation $\hat{\phi}$ and minimizes $\hat{R}_{test}(f, \hat{\phi})$, as described in equation 2—this can be done via BC where we set the loss function ℓ to be the cross-entropy loss for discrete action space and mean-squared loss for continuous action space. In practice, since f is a linear function parameterized by W , MTBC propagates the gradient to update W (see algorithm 2).

E Implementation Details

In this section, we describe in detail the implementation of our experimental analysis.

E.1 Environments

We perform our experiments on four tasks: frozen lake, pendulum, cheetah run, and walker walk. The first two environments are based on Gymnasium [Towers et al., 2023] and the latter two are based on Deepmind control suite [Tunyasuvunakool et al., 2020]. To bridge the theory and practice, we provide a discrete action space variant for the continuous environments to validate our hypotheses. While in general discretizing action space may exclude the true optimal policy [Dadashi et al., 2022, Seyde et al., 2021, 2022], our goal is to demonstrate that MTBC can obtain the expert policy with less target data when compared with BC. Thus, we argue that not having the true optimal policy still validates our goal so long as the expert policy is non-trivial. To generate multiple source tasks with shared state and action spaces, we modify physical properties of the environment. All environmental parameters are sampled from uniform distributions of bounded ranges.




Figure 3: The set of environments used for the experiments. From left to right: frozen lake, pendulum, cheetah, and walker.

Frozen Lake The frozen lake task requires an agent to navigate through a 8×8 grid with non-deterministic transitions. The action space is the set $\{\text{LEFT}, \text{DOWN}, \text{RIGHT}, \text{UP}, \text{STAY}\}$. To vary the tasks, we modify the initial state, the goal state, and the transition function.

Pendulum The pendulum task requires an agent to swing and keep the link upright. The default action space is the torque of the revolute joint between $[-2, 2]$. The discretized action space is the set $\{0, \pm 2^{-3}, \dots, \pm 2^1\}$ (i.e. 11 actions.) This allows the agent to perform large-magnitude action for swinging the link up and to perform low-magnitude action for keeping the link upright. To vary the tasks, we modify the maximum torque applied to the joint.

Locomotion Both the cheetah-run and walker-walk tasks require an agent to move above a specified velocity. The latter task further scales the velocity based on the height of the agent. The default action space is the torque applied to each of the joint, all bounded between $[-1, 1]$. The discretized action space is the set $\{-1, 1\}$ per joint (i.e. Bang-Bang control [Seyde et al., 2021],) thus we have $2^{\dim(\mathcal{A})}$ actions after discretizing the action space—we emphasize that we use softmax policies as opposed to per-dimension Bernoulli policies. To vary the source tasks, we modify the links’ size and the joint parameters.

E.2 Generating the Experts and Demonstrations

In this section we describe how we generate the experts for all environments. We first obtain the expert policy $\pi^{default}$ for each environment using the default environmental parameters. Then, to obtain the expert policy π^{new} for each environmental variant, we initialize π^{new} using $\pi^{default}$ and resume training. This pretraining strategy speeds up training on new environment variant, as opposed to training a new policy from scratch. We use Proximal Policy Optimization (PPO) [Schulman et al., 2017] to train the expert policies for all environments. In general, all PPO policies are trained using Adam optimizer [Kingma and Ba, 2015]. We further use observation normalization and advantage normalization, as commonly done in practice [Engstrom et al., 2020, Hsu et al., 2020].

To gather the demonstrations, we execute the expert policies and deterministically select the action with highest logits. In practice, there is a time-out H for all environments even though in theory they are infinite-horizon MDPs. We gather a total of $\lceil N/H \rceil$ length H time-out episodes and trim the extra $N \bmod H$ samples, which is more practical and time efficient, and similar to how practitioners leverage data in real-life applications.