Efficient Multi-Task Learning via Iterated Single-Task Transfer

K.R. Zentner[1], Ujjwal Puri[1], Yulun Zhang[1], Ryan Julian[2], and Gaurav S. Sukhatme[1]

***Abstract*— In order to be effective general purpose machines in real world environments, robots not only will need to adapt their existing manipulation skills to new circumstances, they will need to acquire entirely new skills on-the-fly. One approach to achieving this capability is via Multi-task Reinforcement Learning (MTRL). Most recent work in MTRL trains a single policy to solve all tasks at once. In this work, we investigate the feasibility of instead training separate policies for each task, and only transferring from a task once the policy for it has finished training. We describe a method of finding near optimal sequences of transfers to perform in this setting, and use it to show that performing the optimal sequence of transfer is competitive with other MTRL methods on the MetaWorld MT10 benchmark. Lastly, we describe a method for finding nearly optimal transfer sequences during training that is able to improve on training each task from scratch.**

I. Introduction

Multi-task reinforcement learning (MTRL) methods that train a single neural network jointly on multiple tasks have many potential advantages. If the tasks have shared structure, then joint training may be able to learn to exploit that shared structure to learn faster or learn a higher performing policy [1], [2]. Furthermore, some recent work has demonstrated that training a network which uses an appropriate latent task representation can allow for generalization across the training tasks [3], [4], [5]. Despite these advantages, there are also some reasons to train separate neural networks for each task in a set. In particular, the performance of single-task policies continues to be higher than multi-task policies in many cases [2], [6]. Training separate policies for each task may also be more computationally efficient, since each optimization step of a single task network only needs to update the small number of parameters for that task. In some of the most comprehensive work on multi-task supervised learning, a mix of these approaches has been demonstrated, with each task only using and updating a small sub-network of the full multi-task network [7]. Lastly, a training separate policy for each task allows achieving a minimum performance level on *all tasks*, in contrast to several prior MTRL methods, which may achieve a high performance on average by ignoring more difficult tasks. In this work we choose to investigate the case of MTRL via training separate policies for each task. We further restrict ourselves to only transferring fully trained policies from one task to another task using Proximal Policy Optimization [8], with minimal modifications.

[1] Department of Computer Science, University of Southern California, Los Angeles, CA 90089. kzentner|ujjwalpu|yulunzha|gaurav@usc.edu GS holds concurrent appointments as a Professor at USC and as an Amazon Scholar.
[2] rjulian@google.com

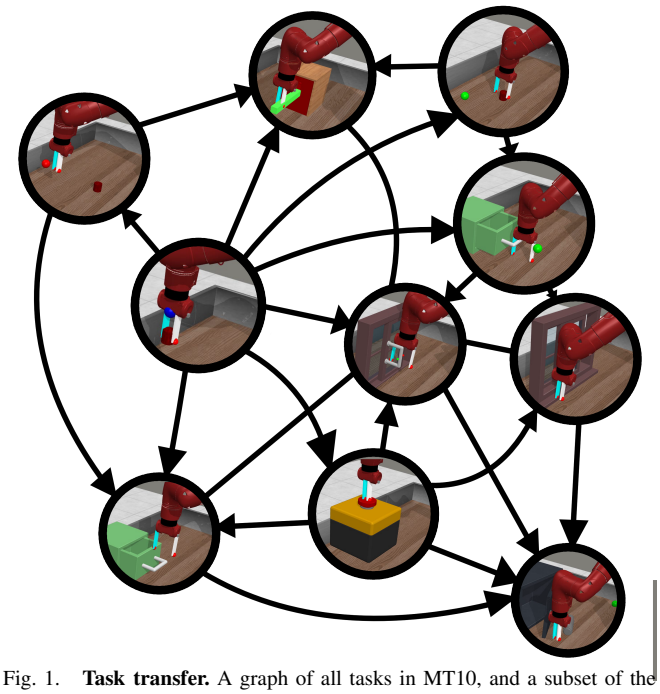

Fig. 1. **Task transfer.** A graph of all tasks in MT10, and a subset of the possible edges between them. We are interested in transfers between two tasks at a time, which correspond to edges in such a graph. Some parts of this work will also include a "random" or "from-scratch" vertex.

We contribute evidence towards answering three questions of interest in this setting:

1) How well, in principle, can this restricted form of transfer improve sample efficiency, if at all?
2) What conditions are required for this type of transfer to improve sample efficiency?
3) How do policies trained using transfer differ from those trained from scratch, if at all?

To investigate these questions, we conduct a series of experiments using the MetaWorld MT10 [9] benchmark, which provides 10 related robotic control tasks each containing internal variation.

II. Related Work

a) Reinforcement learning for robotics: Reinforcement learning (RL) for learning robotic capabilities has been well studied [10], [11], [12], [13]. The recent resurgence of interest in neural networks for use in supervised learning domains has resulted in a resurgence of interest in neural networks for RL [14], [15].

b) Transfer, continual, and lifelong learning for robotics: Transfer learning is a heavily-studied problem outside the robotics domain [16], [17], [18], [19]. Many

approaches have been proposed for rapid transfer of robot policies to new domains, including residual policy learning [20], simultaneously learning across multiple goals and tasks [21], methods which use model-based RL [22], [23], [24], [25], and goal-conditioned RL [26], [27]. Similarly, work in robotic meta-learning focuses on learning representations which can be quickly adapted to new dynamics [28], [29] and objects [30], [31], [32], but has thus far been less successful for task-task transfer [9].

Like [33], our approach relies on rapidly adapting policies for an already-acquired task into a policy for a new task. Much like [34], and [35], we use experiments to analyze different transfer techniques from a geometric perspective on the task-task adaptation problem. As in prior work [36], [37], we observe that the selection of pre-training tasks is essential for preparing RL agents for rapid adaptation.

c) Reusable skill libraries for efficient learning and transfer: Learning reusable skill libraries is a classic approach [38] for efficient acquisition and transfer of robot motion policies. Prior to the popularity of DRL-based methods, Associative Skill Memories [39] and Probabilistic Movement Primitives [40], [41] were proposed for acquiring a set of reusable skills for robotic manipulation. In addition to manipulation [42], [43], [44], DRL-based skill decomposition methods are popular today for learning and adaptation in locomotion and whole-body humanoid control [45], [46], [47], [48]. [5] proposed learning reusable libraries of robotic manipulation skills in simulation using RL and learned latent spaces, and [1] showed these skill latent spaces could be used for efficient simulation-to-real transfer and rapid hierarchical task acquisition with real robots.

III. Setting

We are interested in acquiring policies for each task in a finite task space $\mathcal{T}$. As is common in multi-task RL, we presume all tasks in $\mathcal{T}$ share a single continuous state space $\mathcal{S}$ and continuous action space $\mathcal{A}$, and the MTRL problem is defined by the tuple $(\mathcal{T}, \mathcal{S}, \mathcal{A})$. Each task $\tau \in \mathcal{T}$ is an infinite-horizon Markov decision process (MDP) defined by the tuple $\tau = (\mathcal{S}, \mathcal{A}, p_\tau(s, a, s'), r_\tau(s, a, s'))$. As tasks are differentiated only by their reward functions r_τ and state transition dynamics p_τ, we may abbreviate this definition to simply $\tau = (r_\tau, p_\tau)$. Critically, we will assume that the order in which tasks are acquired can be controlled.

There are many possible ways of measuring the effectiveness of a learning procedure in multi-task learning. In this work, we focus on achieving a required level of performance S^* for each task in $\mathcal{T}$ using as few environment steps as possible. Notably, this is strictly *more difficult* than achieving the same performance on average across all tasks. Further, we limit ourselves to repeatedly transferring a fully-trained policy from one task (henceforth the "base" task) to another task (henceforth the "target" task). We will describe such transfer as "efficient" if it takes fewer environment steps to train a policy for the target task then training a policy "from scratch" (i.e. using a randomly initialized policy).

IV. Finding a Near Optimal Transfer Ordering

In Algorithm 1 we define a transfer procedure, which attempts to transfer from a policy π_τ to a target task τ' using PPO [8]:

Algorithm 1 Transfer Procedure

Input: Base policy π_{base}, target task $\tau' \in \mathcal{T}$,
target success rate S^*,
cutoff performance for each training epoch M_i
$D_0 \leftarrow \texttt{run_policy}(\pi_{base}, \tau')$
$V_0(s) \leftarrow \texttt{train_value_function}(D_0)$
$\pi_{\tau'} \leftarrow \texttt{clone}(\pi_{base})$
for $i \in (1, \ldots)$ **do**
  $D_i \leftarrow \texttt{ppo_step}(\pi_{\tau'}, V_{i-1}, \tau')$
  if $\texttt{performance}(D_i) \geq S^*$ **then**
    return $\pi_{\tau'}$
  else if $\texttt{performance}(D_i) < M_i$ **then**
    return $Reject$
  end if
end for

Algorithm 1 is a slightly modified PPO used as a building block in Algorithm 2. If the value function is not pre-trained, PPO rapidly destroys the policy before it can be transferred. On the benchmark that we are experimenting with, PPO often fails to train a working policy, even during single-task learning. Therefore, we have also added a rejecting rule, which rejects a training run if it begins to fail. Specifically, for each task we have generated a smooth curve "cutoff curve" M_i from a successful training run, and reject the run if it falls behind this curve by more than 10% of the training period. This allows us to reliably complete enough runs with PPO to acquire all of the tasks.

The first question we address is if such a simple transfer procedure can efficiently transfer policies from one task to another. We evaluate this by training a base policy for each task in MT10, and then running a transfer using that policy to each target task in MT10. For each run we measure the number of environment steps needed to reach a 90% success rate. The results are shown in Figure 2. Notably, some tasks can be learned more quickly by transferring from another task than learning that task from scratch (as shown in the last row of Figure 2). This shows that, potentially, this type of transfer could improve sample efficiency.

To evaluate this further, we assume that policies acquired through transfer will transfer as well as policies learned from scratch. Under that assumption, the most efficient way of acquiring all tasks is to minimize the total environment steps needed to acquire each task. This can be found by computing a Directed Minimum Spanning Tree (DMST), rooted at the "from scratch" node, where the edges of the tree correspond to a series of transfers. See Figure 4 for an example. Henceforth, we will refer to such trees, or an arbitrary ordering of their edges consistent with traversal as a "curriculum". A detailed description of this method is given in Algorithm 2.

Algorithm 2 DMST-BASED OPTIMAL TRANSFER

1: **Input:** task space $\mathcal{T}$, target performance S^*
2: $V \leftarrow \mathcal{T} \cup scratch$
3: $E \leftarrow \{\}$
4: **for** $\tau \in \mathcal{T}$ **do**
5: $\quad \pi_\tau, C_\tau \leftarrow \text{PPO}(\tau, \pi_{random})$
6: $\quad E \leftarrow (scratch \rightarrow \tau, C_\tau)$
7: $\quad$ **for** $\tau' \in \mathcal{T}$ **do**
8: $\quad\quad \cdot, C_{\tau \rightarrow \tau'} \leftarrow \text{PPO}(\tau', \pi_\tau)$
9: $\quad\quad E \leftarrow (\tau \rightarrow \tau', C_{\tau \rightarrow \tau'}) \cup E$
10: $\quad$ **end for**
11: **end for**
12: $T_{optimal} \leftarrow \text{DMST}(V, E)$
13: $\pi_\tau \leftarrow \pi_{random}$
14: **for** $(\tau \rightarrow \tau') \in \texttt{traverse}(T_{optimal})$ **do**
15: $\quad$ **while** τ' not solved **do**
16: $\quad\quad \pi_{\tau'} \leftarrow \text{TRANSFERPROCEDURE}(\pi_\tau, \tau', S^*)$
17: $\quad\quad$ **if** $(\tau \rightarrow \tau')$ rejected **then**
18: $\quad\quad\quad E \leftarrow E \setminus (\tau \rightarrow \tau')$
19: $\quad\quad\quad T_{optimal} \leftarrow \text{DMST}(V, E)$
20: $\quad\quad$ **end if**
21: $\quad$ **end while**
22: **end for**

Algorithm 2: The for loop on line 4 computes Figure 2, by running PPO to transfer along each potential DMST "edge", $\tau \rightarrow \tau'$, and computing the number of environment steps $C_{\tau \rightarrow \tau'}$ to reach S^*.

Given this way of computing an optimal curriculum, we can also compute a pessimal curriculum and random curricula. The results of actually running such curricula are shown in Figure 3. The plot shows that the number of environment steps predicted by each type of curriculum match closely with the environment steps actually used to train using that curriculum. This indicates that our assumption that policies training from transfer will transfer similarly to policies trained from scratch is reasonable. Figure 3 also shows that the optimal curriculum is, both in prediction and when run, more sample efficient than learning from scratch. The rest of our work in this direction analyses the behavior of our transfer procedure with a eye towards approximating this optimal curriculum.

V. CONDITIONS FOR EFFICIENT TRANSFERS

We have shown that, under ideal circumstances, a simple transfer procedure can improve sample efficiency if applied to perform a carefully selected sequence of transfers. This makes predicting that set of transfers potentially valuable, and naturally leads to the question of what properties of base and target task lead to efficient transfers. Several prior works have attempted to predict transfer between tasks by measuring task similarity [49], [6], including using the same benchmark we have chosen [4]. However, one very distinct feature of transfer we observe in this setting is that it is very asymmetrical, as shown in Figure 2. This implies that attempting to predict the contents of Figure 2 using sym-

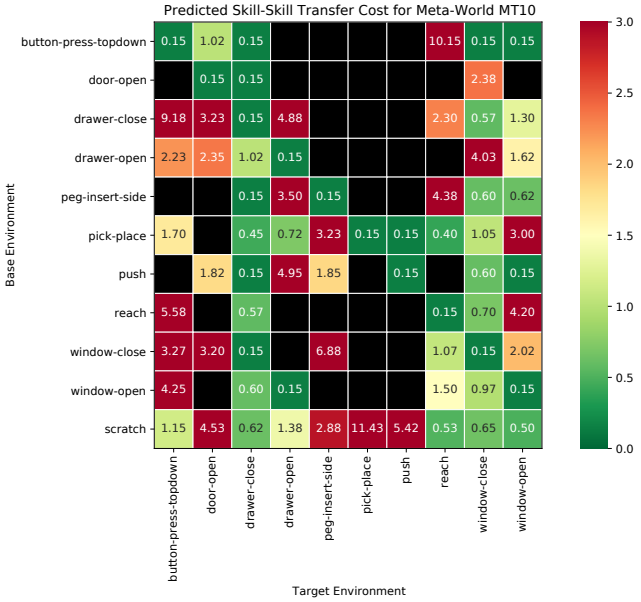

Fig. 2. This figure shows the total number of environment steps in millions to reach 90% success rate on a target task, starting with a policy trained on a base task, or trained from scratch. Black entries did not reach 90% success after 12 million timesteps. All entries ran for at least 150,000 timesteps. Each column corresponds to all possible ways of attempting to train a policy to solve a target task. Note the bottom row, which contains the time require to train each policy starting from a randomly initialized policy. The diagonal confirms that each policy solves the task it was trained on within 1 more epoch of training.

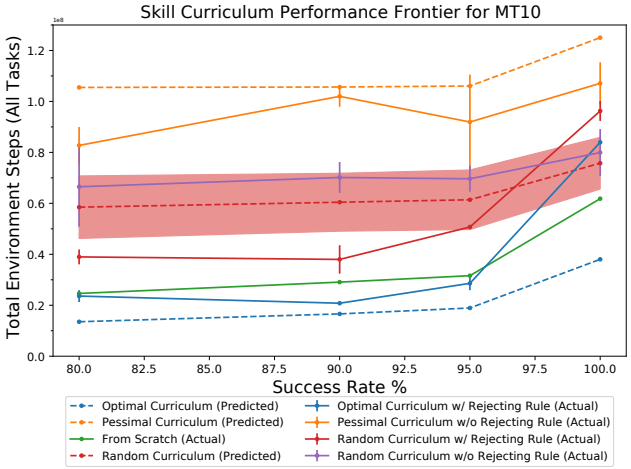

Fig. 3. Comparison of the performance-sample efficiency frontier for several curricula for Meta-World MT10. The high level of agreement between the actual and predicted curricula indicate that policies trained using transfer learning transfer just as well as policies trained from scratch. This implies that the method we describe for approximating the optimal transfer in this setting is nearly optimal. Note the solid blue line, which matches or improves on the efficiency of the from scratch curve for a range of target success rates $80\% \leq S^* \leq 95\%$. This shows a potential improvement in sample efficiency greater than has been achieved in other methods, such as [2]. Error bars show one standard deviation in number of total environment steps require to solve all tasks. However, runs using the rejecting rule exhibit very low variance between runs. Pessimal curricula with the rejecting rule and optimal curricula without the rejecting rule perform similarly to random curricula.

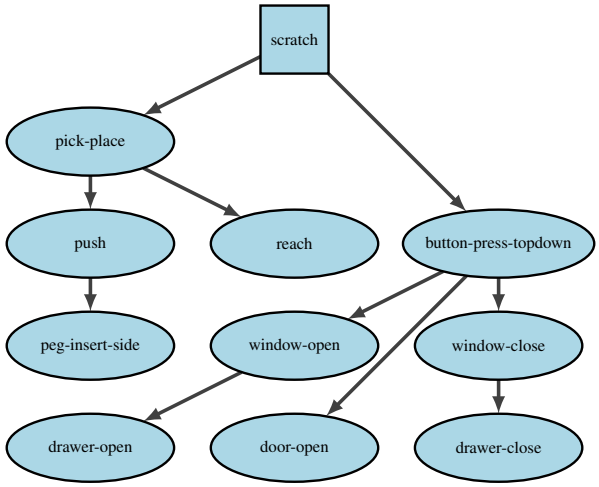

Fig. 4. Predicted optimal curriculum DMST for MT10 computed from Figure 2 as described in Algorithm 2. Note that the curriculum begins by learning the *most general* tasks first, then refines behavior through iterated transfer. Additional curricula may be computed by Algorithm 2 if any transfer edges are rejected as specified by Algorithm 1.

metrical functions (including by projection into any metric space [1], [5]) will not be successful. Since we are not aware of any prior work which is likely to be helpful in light of the above observation, we turn towards analyzing the specific properties of Figure 2. The two most efficient tasks to transfer from, `pick-place` and `button-press-topdown`, deserve particular attention, since together they allow learning almost all other tasks very quickly, and neither of them have other tasks that can be used to efficiently acquire them in turn.

Analyzing the behavior of these two policies on the tasks they transfer to shows that they tend to solve these tasks very quickly because they contain all of the behavior necessary to solve the target task, as well as some amount of additional behavior that does not prevent success at the target task. In other words, the policies trained to solve these tasks are effective exploration policies for a large number of other tasks. In the case of `pick-place`, this is because the policy has been trained using a large number of timesteps to operate throughout the state space, and therefore rapidly generalizes to similar tasks that use a smaller portion of the state space. In contrast, the `button-press-topdown` policy has been trained with relatively few timesteps, on a task that requires relatively little precision. This results in a policy with highly stochastic behavior, which is able to succeed at several tasks due to that stochastic behavior sometimes performing the necessary manipulations.

This trend, of policies with more general behavior transferring well to tasks with more specific requirements holds in other cases among this task set, such as `push` $\rightarrow$ `peg-insert-side`, `push` $\rightarrow$ `window-open`, and `push` $\rightarrow$ `drawer-close`. If this trend continues in other benchmarks, it may suggest an alternative approach to transfer learning in MDPs. In particular, instead of transferring between tasks on the basis of similarity, transfer between tasks on the basis of one task being a more general case of the other may be worth investigating.

VI. Transfer Without Perfect Information

Although Algorithm 2 demonstrates that iterated single-task transfer has the potential for improving sample efficiency in a multi-task setting, it does not provide a practical method, since it depends on the data in Figure 2, which is more difficult to acquire than simply training each task from scratch. Furthermore, Algorithm 1 implements a "rejection rule" using data from training each task from scratch. Although the rejecting rule is necessary for efficient transfer (as shown in Figure 3, we would also like to avoid using it because it requires the data from a successful training run of the same task that it is being used to solve. In this section, we investigate one method for performing iterated single-task transfer without having access to the data from Figure 2, or even the data needed for the rejection rule. This method, described in Algorithm 3, performs inference to approximate Figure 2 while simultaneously using the inference state to direct training.

Algorithm 3 DMST-Inference Based Transfer

1: **Input:** task space $\mathcal{T}$, target performance S^*, maximum timesteps per task M
2: $V \leftarrow \mathcal{T} \cup scratch$
3: **for** $\tau \in V$ **do**
4: $\quad \pi_{BC,\tau} \leftarrow \mathrm{BC}(\tau)$
5: $\quad$ **for** $\tau' \in \mathcal{T}$ **do**
6: $\quad\quad \mu_{\tau\rightarrow\tau'} \leftarrow \frac{M}{1+2\texttt{evaluate_policy}(\tau',\pi_{BC,\tau})}$
7: $\quad\quad \theta_{\tau\rightarrow\tau'} \leftarrow M$
8: $\quad\quad i_{\tau\rightarrow\tau'} \leftarrow 0$
9: $\quad$ **end for**
10: **end for**
11: $\mathcal{T}_{solved} \leftarrow \{scratch\}$
12: **while** $\mathcal{T} \neq \mathcal{T}_{solved} \setminus scratch$ **do**
13: $\quad E_{\tau\rightarrow\tau'} \sim \Gamma(\theta_{\tau\rightarrow\tau'}; \mu_{\tau\rightarrow\tau'}) - i_{\tau\rightarrow\tau'}$
14: $\quad (\tau \rightarrow \tau') \leftarrow \texttt{select_edge}(\mathrm{DMST}(V,E), \mathcal{T}_{solved})$
15: $\quad$ **if** $i_{\tau\rightarrow\tau'} = 0$ **then**
16: $\quad\quad D_0 \leftarrow \texttt{run_policy}(\pi_\tau, \tau')$
17: $\quad\quad V_0(s) \leftarrow \texttt{train_value_function}(D_0)$
18: $\quad\quad \pi_{\tau'} \leftarrow \texttt{clone}(\pi_\tau)$
19: $\quad$ **end if**
20: $\quad i_{\tau\rightarrow\tau'} \leftarrow 1 + i_{\tau\rightarrow\tau'}$
21: $\quad D_i \leftarrow \texttt{ppo_step}(\pi_{\tau'}, V_{i-1}, \tau')$
22: $\quad \theta_{\tau\rightarrow\tau'} \leftarrow \texttt{update_theta}(\theta_{\tau\rightarrow\tau'}, D_i, S^*)$
23: $\quad \mu_{\tau\rightarrow\tau'} \leftarrow \texttt{update_mu}(\mu_{\tau\rightarrow\tau'}, D_i, S^*)$
24: $\quad$ **if** $\texttt{performance}(D_i) \geq S^*$ **then**
25: $\quad\quad \mathcal{T}_{solved} \leftarrow \mathcal{T}_{solved} \cup \tau'$
26: $\quad$ **end if**
27: **end while**

Algorithm 3: The for loop on line 3 uses a small number of demonstrations to produce an unreliable policy $\pi_{BC,\tau}$ for each task, then evaluates it on each other task using `evaluate_policy`, to produce a value between 0 and 1.

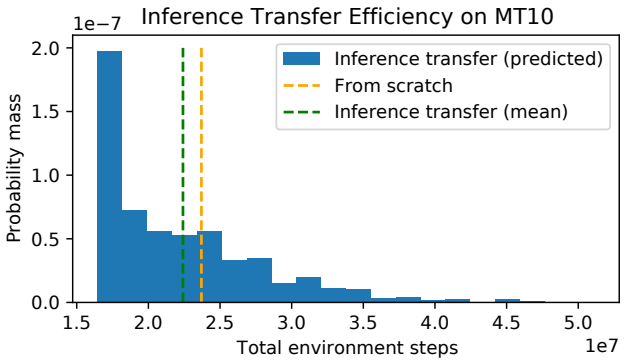

Fig. 5. Histogram showing the predicted performance of running DMST-Inference Based Transfer as described in Algorithm 3 to reach 90% success rate on MT10. Mean predicted environment steps and environment steps to train from scratch show in vertical lines.

This cross-evaluation is then used to initialize a belief state $\Gamma(\theta_{\tau\to\tau'}; \mu_{\tau\to\tau'})$ for each edge that estimates the number of environment steps to transfer along that edge. The while loop on line 12 repeatedly samples edge costs $E_{\tau\to\tau'}$ from that belief state (line 13), computes an optimal curriculum from that sample (line 14), runs a PPO step along a viable edge of that curriculum (line 21), then updates the belief state until every task has been solved (lines 22 and 23). A "viable edge" is any edge $(\tau \to \tau') \in G$ such that $\tau \in \mathcal{T}_{solved}$ and $\tau \notin \mathcal{T}_{solved}$.

Similarly to Algorithm 1, Algorithm 3 also pre-trains a value function when a new transfer is begun (lines 16 and 17). This is necessary to prevent PPO from destroying the behavior of $\pi_{tau'}$ with high-variance policy gradient steps.

The primary motivation for using the Gamma distribution is that the number of training iterations remaining on an edge is always non-negative, which matches the support of the Gamma distribution. This is especially important since many transfers, either initially or near the end of training have an expected number of remaining training iterations very near zero, which a belief distribution without a lower bound would struggle to model. Furthermore, the Gamma distribution also responds well to high-variance in the training progress, which typically corresponds to an unpredictably *long* transfer.

The specific parameterization we use is a scale parameter $\theta_{\tau\to\tau'}$ and an offset parameter $\mu_{\tau\to\tau'}$. The shape parameter k is set to 2. `update_mu` and `update_theta` predict when training is expected to reach S^* using the performance found in D_i, then use the empirical mean and variance of the last 10 predictions $\mathcal{X}$ for that transfer to update $\theta = \sqrt{\mathrm{Var}(\mathcal{X})/k}$ and $\mu = \mathrm{Mean}(\mathcal{X}) - k\theta$. Once D_i contains non-zero performance, a linear projection $\mathcal{X} = i + \frac{S^* - \texttt{performance}(D_i)}{\texttt{performance}(D_i) - \texttt{performance}(D_{i-1})}$ is used to predict the remaining training epochs. If no progress has been made so far, the prediction instead assumes that training will complete at a minimum of $\mathcal{X} = 1.5i$ iterations.

Because Algorithm 3 predicts the costs of all edges, it is able to choose curricula based on their efficiency from future edges using the DMST. Furthermore it is able to respond to transfers that "get stuck" in a way similar to the rejection rule, without requiring data from any prior training of that task. One critical component for achieving this result is the use of "BC cross-evaluation" to create a prior for the belief distribution. Although the policies trained this way are unable to complete the tasks they are trained for reliably, they still provide essential information, because several of the most useful policies to transfer from (`pick-place` and `push`) require a significant amount of time to train. Without this prior and the use of DMST, Algorithm 3 falls back to training each task from-scratch, with a small amount of timesteps spent unsuccessfully on transfer edges.

The predicted results of running Algorithm 3 1000 times are shown in Figure 5. Notably, the inference often finds the optimal curriculum during training, as shown by the peak in the histogram at the far left. Near optimal curricula are also often found, resulting in the mean sample efficiency improving on the sample efficiency of training from-scratch. However, there is a tail of runs that execute curricula that are much worse than training from scratch.

VII. Conclusion

We have found that that performing the (near) optimal sequence of transfer is competitive with other multi-task RL methods on the MetaWorld MT10 benchmark. In general, transferring from more "general" tasks, such as `pick-place` or `button-press-topdown` can be very effective. Notably, the transfer relationships we find appear to be highly asymmetrical, indicating that traditional "task similarity" approaches to transfer are unlikely to succeed in this setting. Furthermore, policies trained via transfer in this way are indistinguishable from those trained from scratch, allowing repeated transfers. Finally, we investigated a method of performing this type of transfer without prior knowledge of the optimal sequence, opening up the possibility to base practical methods on those described in this work. In future work, we will improve on the proposed method and investigate other benchmarks to see if these patterns conclusions hold across more settings.

References

[1] R. C. Julian, E. Heiden, Z. He, H. Zhang, S. Schaal, J. Lim, G. S. Sukhatme, and K. Hausman, "Scaling simulation-to-real transfer by learning composable robot skills," in *International Symposium on Experimental Robotics*. Springer, 2018. [Online]. Available: https://ryanjulian.me/iser_2018.pdf

[2] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Gradient surgery for multi-task learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5824–5836, 2020.

[3] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," 2021.

[4] S. Sodhani, A. Zhang, and J. Pineau, "Multi-task reinforcement learning with context-based representations," in *International Conference on Machine Learning*. PMLR, 2021, pp. 9767–9779.

[5] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. Riedmiller, "Learning an embedding space for transferable robot skills," in *International Conference on Learning Representations*, 2018. [Online]. Available: https://openreview.net/forum?id=rk07ZXZRb

[6] C. Fifty, E. Amid, Z. Zhao, T. Yu, R. Anil, and C. Finn, "Efficiently identifying task groupings for multi-task learning," *Advances in Neural Information Processing Systems*, vol. 34, 2021.

[7] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," 2017.

[8] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[9] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning," in *Conference on Robot Learning (CoRL)*, 2019. [Online]. Available: https://arxiv.org/abs/1910.10897

[10] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.

[11] S. Mahadevan and J. Connell, "Automatic programming of behavior-based robots using reinforcement learning," *Artificial intelligence*, vol. 55, no. 2-3, pp. 311–365, 1992.

[12] L.-J. Lin, "Reinforcement learning for robots using neural networks," 1992.

[13] W. D. Smart and L. P. Kaelbling, "Effective reinforcement learning for mobile robots," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, vol. 4. IEEE, 2002, pp. 3404–3410.

[14] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," *arXiv preprint arXiv:1811.12560*, 2018.

[15] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.

[16] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "Decaf: A deep convolutional activation feature for generic visual recognition," in *International conference on machine learning*, 2014, pp. 647–655.

[17] 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.

[18] W. Dai, Q. Yang, G.-R. Xue, and Y. Yu, "Boosting for transfer learning," in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 193–200.

[19] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng, "Self-taught learning: transfer learning from unlabeled data," in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 759–766.

[20] T. Silver, K. Allen, J. Tenenbaum, and L. Kaelbling, "Residual policy learning," *arXiv preprint arXiv:1812.06298*, 2018.

[21] S. Ruder, "An overview of multi-task learning in deep neural networks," 2017.

[22] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2786–2793.

[23] L. Yen-Chen, M. Bauza, and P. Isola, "Experience-embedded visual foresight," in *Conference on Robot Learning*. PMLR, 2020, pp. 1015–1024.

[24] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112.

[25] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution," in *Advances in Neural Information Processing Systems*, 2018, pp. 2450–2462.

[26] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine, "Learning to poke by poking: Experiential learning of intuitive physics," in *Advances in neural information processing systems*, 2016, pp. 5074–5082.

[27] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, F. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell, "Zero-shot visual imitation," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, June 2018. [Online]. Available: http://dx.doi.org/10.1109/CVPRW.2018.00278

[28] I. Clavera, A. Nagabandi, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," in *International Conference on Learning Representations*, 2019. [Online]. Available: https://openreview.net/forum?id=HyztsoC5Y7

[29] L. P. K. Ferran Alet, Tomas Lozano-Perez, "Modular meta-learning," in *Conference on Robot Learning (CoRL)*, 2018. [Online]. Available: http://lis.csail.mit.edu/pubs/alet-corl18.pdf

[30] S. James, M. Bloesch, and A. J. Davison, "Task-embedded control networks for few-shot imitation learning," in *Conference on Robot Learning*, 2018, pp. 783–795.

[31] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, and S. Levine, "One-shot imitation from observing humans via domain-adaptive meta-learning," in *International Conference on Learning Representations*, 2018.

[32] A. Bonardi, S. James, and A. J. Davison, "Learning one-shot imitation from humans without humans," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3533–3539, 2020.

[33] M. E. Taylor, P. Stone, and Y. Liu, "Transfer learning via inter-task mappings for temporal difference learning." *Journal of Machine Learning Research*, vol. 8, no. 9, 2007.

[34] Z. Cao, M. Kwon, and D. Sadigh, "Transfer reinforcement learning across homotopy classes," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2706–2713, 2021.

[35] S. Kumar, A. Kumar, S. Levine, and C. Finn, "One solution is not all you need: Few-shot extrapolation via structured maxent rl," *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[36] R. Luna Gutierrez and M. Leonetti, "Information-theoretic task selection for meta-reinforcement learning," in *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020)*. Leeds, 2020.

[37] L. Yen-Chen, A. Zeng, S. Song, P. Isola, and T.-Y. Lin, "Learning to see before learning to act: Visual pre-training for manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7286–7293.

[38] V. Gullapalli, J. A. Franklin, and H. Benbrahim, "Acquiring robot skills via reinforcement learning," *IEEE Control Systems Magazine*, vol. 14, no. 1, pp. 13–24, 1994.

[39] P. Pastor, M. Kalakrishnan, L. Righetti, and S. Schaal, "Towards associative skill memories," in *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*, Nov. 2012, pp. 309–315.

[40] E. Rueckert, J. Mundo, A. Paraschos, J. Peters, and G. Neumann, "Extracting low-dimensional control variables for movement primitives," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 1511–1518.

[41] X. Zhou, H. Wu, J. Rojas, Z. Xu, and S. Li, "Incremental learning robot task representation and identification," in *Nonparametric Bayesian Learning for Collaborative Robot Multimodal Introspection*. Springer, 2020, pp. 29–49.

[42] R. Yang, H. Xu, Y. Wu, and X. Wang, "Multi-task reinforcement learning with soft modularization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4767–4777, 2020.

[43] M. Wulfmeier, D. Rao, R. Hafner, T. Lampe, A. Abdolmaleki, T. Hertweck, M. Neunert, D. Tirumala, N. Siegel, N. Heess, *et al.*, "Data-efficient hindsight off-policy option learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 11 340–11 350.

[44] O. Kroemer, C. Daniel, G. Neumann, H. Van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1503–1510.

[45] X. B. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine, "Mcp: Learning composable hierarchical control with multiplicative compositional policies," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[46] L. Hasenclever, F. Pardo, R. Hadsell, N. Heess, and J. Merel, "CoMic: Complementary task learning & mimicry for reusable skills," in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 4105–4115. [Online]. Available: http://proceedings.mlr.press/v119/hasenclever20a.html

[47] J. Merel, S. Tunyasuvunakool, A. Ahuja, Y. Tassa, L. Hasenclever, V. Pham, T. Erez, G. Wayne, and N. Heess, "Catch & carry: reusable neural controllers for vision-guided whole-body tasks," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 4, pp. 39–1, 2020.

[48] T. Li, N. Lambert, R. Calandra, F. Meier, and A. Rai, "Learning generalizable locomotion skills with hierarchical reinforcement learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 413–419.

[49] N. Ferns, P. Panangaden, and D. Precup, "Metrics for finite markov decision processes." in *UAI*, vol. 4, 2004, pp. 162–169.