

Projected Task-Specific Layers for Multi-Task Reinforcement Learning

Josselin Somerville Roberts¹ and Julia Di²

Abstract—Multi-task reinforcement learning could enable robots to scale across a wide variety of manipulation tasks in homes and workplaces. However, generalizing from one task to another and mitigating negative task interference still remains a challenge. Addressing this challenge by successfully sharing information across tasks will depend on how well the structure underlying the tasks is captured. In this work, we introduce our new architecture, Projected Task-Specific Layers (PTSL), that leverages a common policy with dense task-specific corrections through task-specific layers to better express shared and variable task information. We then show that our model outperforms the state of the art on the MT10 and MT50 benchmarks of Meta-World consisting of 10 and 50 goal-conditioned tasks for a Sawyer arm.

I. INTRODUCTION

Complex manipulation is common in a number of desirable real-world robotic use cases—such as wiping various kitchen surfaces in a busy restaurant, routing cables in a datacenter, or screwing parts in an manufacturing assembly. While these individual tasks are different, they are often composed of similar manipulation primitives. Humans intuitively recognize and scaffold upon these primitives to implement different tasks, but it is challenging for robots, which are often only trained on individual, specific tasks, to do the same. Enabling robots to reason through multiple related tasks with multi-task reinforcement learning [1] can unlock more general-purpose robotics—allowing efficient learning across similar tasks and using their shared structure to learn a better performing policy [2]–[4].

¹Josselin Somerville Roberts is with the Department of Computer Science, Stanford University, Stanford CA, 94305, USA josselin@stanford.edu

²Julia Di is with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA juliadi@stanford.edu

Multi-task learning often, however, requires a careful choice of tasks and balanced sampling, and even then may not always improve learning. For complex manipulation, learning-based approaches may generalize over unseen tasks [5] but can still be difficult to scale successfully [6], [7]. Recent work has argued that learning methods should use shared task structures [8], [9] but most approaches still learn a single shared policy used across all tasks which may not adequately represent variations between tasks [6], [10].

Instead, we propose a new backbone architecture, the Projected Task-Specific Layers (PTSL), which combines a large, shared fully-connected policy with low-rank task-specific layers as shown in Fig. 1. After each layer, the hidden state from the shared policy and the low-rank task-specific policy are combined, making PTSL expressive of different tasks. We evaluate PTSL as a standalone backbone or on top of the Context-based Representation (CARE [4]) encoder that leverages text descriptions of the task as metadata to project a mixture of encodings.

The main contributions of this work are:

- We propose the **PTSL** architecture for deep multi-task reinforcement learning that adds low-rank task-specific layers on top of each layer of a shared backbone.
- Our results with PTSL outperforms CARE [4], the current state-of-the-art, on both the MT10 and MT50 Goal-conditioned benchmarks from Meta-World [11].
- Further, our results suggest that multi-task learning with a shared projection is more sample efficient and can improve learning on individual tasks.
- Finally, our results provide insights into the benefits of intermediate architectures sharing an embedding space between task-specific layers and a backbone.

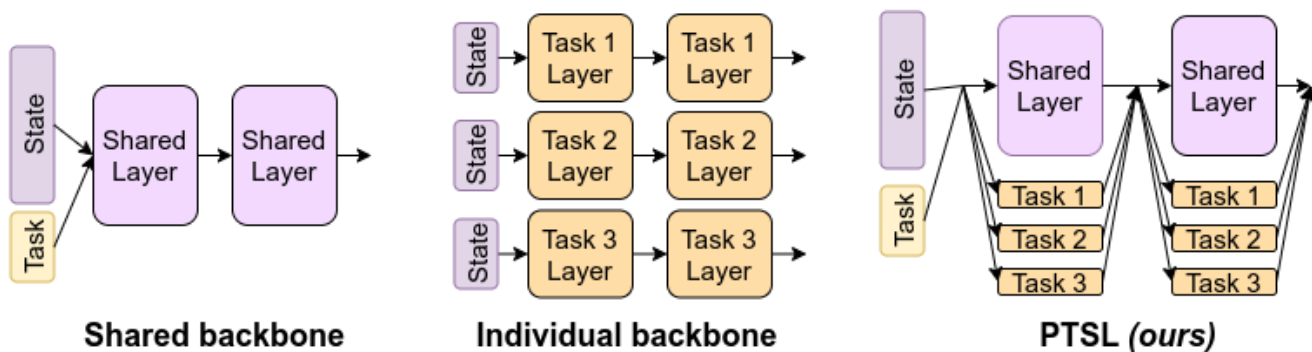


Fig. 1: Simplified diagram of different architectures for multi-task reinforcement learning: Shared backbone for all tasks (left), Individual backbone for each task (center) and Projected Task-Specific Layers (ours) (right)

II. RELATED WORK

A. Reinforcement Learning for Robotics

Reinforcement learning (RL) has recently shown success in various domains such as Atari [12], Go [13], and Starcraft [14]. In robotics, RL has been well studied in self-driving cars [15], manipulation [16]–[18], and locomotion [19]. Multi-task reinforcement learning is a subfield of RL for teaching a single agent to solve multiple tasks. Whereas single-task RL optimises for one reward function, multi-task learning optimises for multiple objectives [2].

One of the main challenges of multi-task learning is *task interference* or *negative transfer*: learning a new task can deteriorate the performance of a previously learned task. This has been documented in previous literature [3], [20] and several approaches have been proposed for mitigation. One approach is to train single policies for each task and use knowledge transfer [21] or policy distillation [22], but this requires separate networks or a large number of network parameters. Other previous approaches [23], [24] have addressed this issue by modifying the training algorithm or having a model architecture that can support different sub-policies [4], [25], but these methods are often slow and scale poorly to the number of tasks. Thus, in this work we choose to address this issue by focusing not on the optimization method but the underlying architecture.

B. Soft Actor-Critic

The Soft Actor-Critic (SAC) [26] algorithm optimizes the maximum-entropy RL objective using off-policy data to learn. SAC has been demonstrated to perform better [11] than other algorithms such as Proximal Policy Optimization (PPO) [27] or Trust Region Policy Optimization (TRPO) [28]. Accordingly, we choose to use the multi-task adaptation of SAC with disentangled alphas (which refers to having separate alpha coefficients for every task learned by the policy [29]) to focus on the architecture of the agent.

C. Multi-task Architectures

Two common and opposing architectures for multi-task learning are the multi-headed actor and the shared actor architecture [11]. The multi-headed actor consists of a single network with one head per task, but does not scale well due to the sheer number of parameters. Meanwhile, the shared actor consists of having a single network for all tasks, which does not allow for task-specific corrections and often leads to poor performance as the number of tasks increases [30].

Within these paradigms, a number of approaches exist. Mixture of Experts (MOE) [31], [32] methods have independent experts with a learned gating network to output weights. Similar work has been done with Hard routing [33] which consists of having a routing network that selects the expert for each task. However, good expert assignments can be nontrivial. A third improvement on the multi-headed architecture is called Soft Modularization [25], where each "step" of the network is composed of multiple linear layers and a routing network decides how to route the linear layers of one step to the ones of the next step.

Another approach is to have a shared network with a preprocessor that produces a vector to represent the state, often denoted as an encoder. CARE (Contextual Attention-based Representation learning) proposes a mixture of encoder architecture [4]. The idea is to have several independent encoders that encode the task in a feature vector. Then, another component produces attention scores based only on the task of combining these embeddings. This embedding is then fed to a fully linear network. This process is done for the actor, critic, and value functions. However, while a task-dependent encoder is useful, a shared policy like in CARE [4] may not be the best approach because some tasks may require small variations in the shared policy.

III. METHOD

In this section, we introduce our **Projected Task-Specific Layers (PTSL)** architecture, which can be used on top of methods such as CARE [4]. PTSL reconciles encoding approaches and routing approaches for multi-tasking.

A. Projected Attention Layers

Encoding the state in a task-specific way, like in CARE, may not be enough for all tasks. For example, in manipulation, knowing how to push and pull may not benefit the actions of picking and placing objects. Yet it is still beneficial to have a shared policy because some skills can be shared.

PTSL permits small task-specific variations in the policy and was inspired by Projected Attention Layers (PAL) [34], a method that permits a transformer to have small task-specific variations for Natural Language Processing (NLP). For each task, the input goes through a series of transformer layers. Each layer is made of two parts: one large shared attention layer and one small task-specific attention layer. The latter is computed by projecting down the input to a smaller dimension, applying an attention layer trained specifically for the task, and then projecting it up to the original dimension and adding it to the shared attention layer output. This process is repeated for each layer of the transformer.

The intuition is that a general backbone can contain most of the knowledge (akin to grammar rules for a NLP Transformer), but small task-specific variations may be added as some tasks may require different treatment (e.g. classification versus summarization). PALs have been shown to be very efficient in NLP tasks and we believe that this approach can be applied to robotic manipulation as well.

B. Projected Task-Specific Layers

We propose **Projected Task-Specific Layers (PTSL)**, which adapts the PAL architecture to a linear layer setup instead of a transformer, with a few modifications. This architecture can replace the backbone of other methods to obtain complex agents such as **CARE + PTSL**.

Similarly to PAL, PTSL is made of a shared backbone and low-rank task-specific layers (one low-rank layer for each task, on top of each shared layer) as detailed in Fig. 2. The backbone is a linear layer that is shared between all tasks. The task-specific layers are linear layers that are specific to

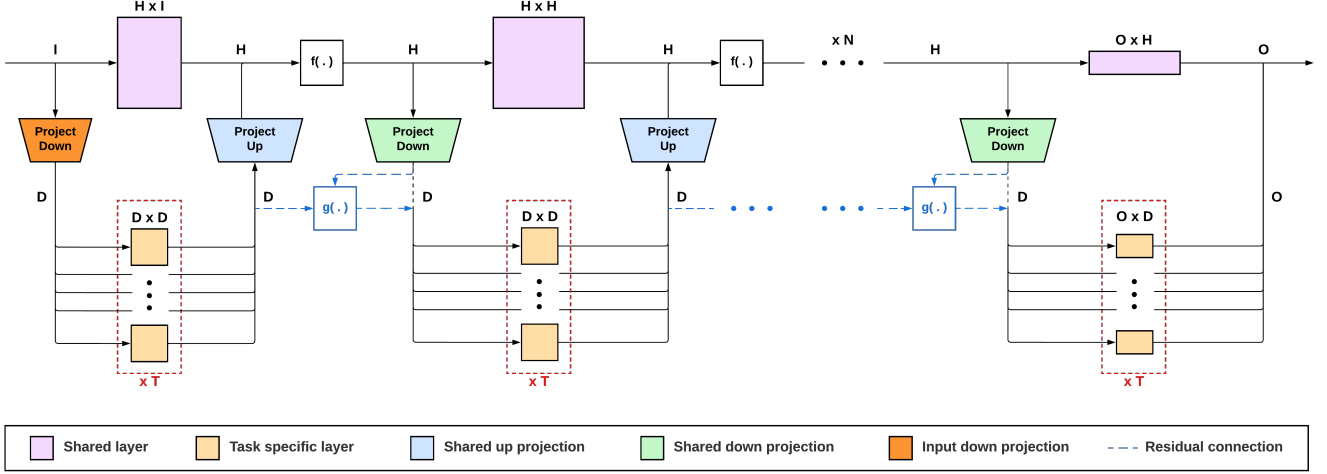


Fig. 2: PTSL Architecture (**ours**), explained in Section III-B for details. The dotted red lines represent residual connections (not always present). See Section III-C.3 for details. Projection modules that are reused are represented with the same color. See Section III-C.2 for details.

each task. Where our approach differs from PAL is regarding the projections to the task-specific dimension.

C. Problem formulation and Preliminaries

In a multi-task setting, we have a set of tasks $\mathcal{T} = \{T_1, T_2, \dots, T_T\}$ and a set of environments $\mathcal{E} = \{E_1, E_2, \dots, E_T\}$ where each task T_i is associated with an environment E_i . Each environment E_i is a Markov Decision Process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, \mathcal{P} is the transition probability function, \mathcal{R} is the reward function, and γ is the discount factor. The goal of the agent is to learn a policy π_θ that maximizes the expected return $J(\pi_\theta)$ where θ are the parameters of the policy.

In a multi-task setting, the agent has to learn a policy π_{θ_i} for each task T_i . The agent is evaluated on its ability to learn a policy for each task T_i so therefore we are interested in the average return over all tasks. Since we evaluate our method on Meta-World, we will focus on the discrete signal that is the success rate of the agent on each task. Because we consider a multi-task setting with no unseen tasks during training, we do not consider the problem of generalization to unseen tasks for this work.

1) *Notation:* We introduce the following notation:

- I the input dimension,
- O the output dimension,
- H the hidden dimension
- D the task-specific dimension,
- T the number of tasks,
- N the number of hidden layers (meaning that we have $N + 1$ layers in total).

In addition, we note the input x and x_i as the input of the i -th layer (e.g., $x = x_0$). Additionally, we define:

- SH^i the shared i -th linear layer.
- TS_j^i the i -th task specific layer for the j -th task.

- P_{down}^i the i -th projection layer that projects from H to D (except for P_{down}^0 that projects from I to D).
- P_{up}^i the i -th projection layer that projects from D to H (except for P_{up}^N that projects from D to O).

2) *Projection to the task-specific dimension:* In a transformer setup, we often have $I = H = O$, which allows PAL to have a single shared down projection P_{down}^* and a single shared up projection P_{up}^* . Under a maximum number of parameters constraint, it is better to have a single shared projection than a projection for each layer [34].

This is less relevant in our case as I and O are not necessarily equal to H and the gain of sharing the projection is less important as transformers such as BERT have 12 layers while we only have 3 hidden layers (meaning 4 layers in total). Thus, we experiment with both options: a single shared projection and a projection for each layer. In the case of a single shared projection, we will note P_{down}^* and P_{up}^* the shared projection layers. Note that we still have two independent projections P_{down}^0 and P_{up}^N that are not shared due to the difference in input and output dimensions.

In our case, we investigate whether having these two specific projections is worthwhile. Indeed if I and O are small, it may be more beneficial to skip the projection so that TS_j^0 goes from I to D and TS_j^N goes from D to O . We can compute the number of parameters for each case:

- **First Task-specific layer TS_j^0 :** with a projection the number of parameters is $I \times D + T \cdot (D \times D + D)$ while without a projection it is $T \cdot (I \times D + D)$.
- **Last Task-specific layer TS_j^N :** with a projection the number of parameters is $D \times O + T \cdot (D \times D + D)$ while without a projection it is $T \cdot (D \times O + O)$.

For Meta-World, it was more beneficial to have an individual down projection P_{down}^0 but no up projection P_{up}^N as $I = 104$ (output of the CARE encoder) and $O = 8$.

3) *Residuals across task-specific layers*: We also add a residual connection between the task-specific layers, as illustrated by the dotted red in Figure 2. Let us note the function $g : (\mathbb{R}^D, \mathbb{R}^D) \rightarrow \mathbb{R}^D$ that combines the projection of our input $P_{\text{down}}^i(x_i)$ with the output of the previous task-specific layer $\text{TS}_j^i(P_{\text{down}}^{i-1}(x_{i-1}))$ to obtain the input of the current task-specific layer TS_j^i .

In this paper we will consider four different functions g :

- **No residual**: $g(x_i^{\text{proj}}, y_{i-1}^{\text{task}}) = x_i^{\text{proj}}$.
- **Addition**: $g(x_i^{\text{proj}}, y_{i-1}^{\text{task}}) = x_i^{\text{proj}} + y_{i-1}^{\text{task}}$.
- **Learnable sum**: $g(x_i^{\text{proj}}, y_{i-1}^{\text{task}}) = \alpha x_i^{\text{proj}} + \beta y_{i-1}^{\text{task}}$ with $\alpha, \beta \in \mathbb{R}$ learnable parameters.
- **Learnable projection**: two vectors are concatenated to a vector of \mathbb{R}^{2D} that is projected to \mathbb{R}^D . $g(x_i^{\text{proj}}, y_{i-1}^{\text{task}}) = P_g \times \text{Concat}(x_i^{\text{proj}}, y_{i-1}^{\text{task}})$ with $P_g \in \mathbb{R}^{D \times 2D}$, a learnable projection shared across tasks and layers.

IV. EXPERIMENTS

In this section, we evaluate PTSL in the Meta-World multi-task RL environment and compare against baselines. We also conduct ablation studies to verify the effectiveness of our method. To support the community, our code is made publicly available at <https://github.com/JosselinSomervilleRoberts/PTSL> which is a fork of the CARE [4] repository that can be found at <https://github.com/facebookresearch/mtrl>. In addition to the code, the repository contains all the commands to reproduce our experiments, our run results, and some implementation and training tips. All trainings were performed on a NVIDIA RTX A6000 and averaged across n runs.

The first goal of our experimental evaluation is to assess if the PSTL architecture improves the performance of a multi-task agent without increasing the number of parameters. This comparison is done in two different settings: short horizon to evaluate the sample efficiency of PTSL and long horizon to assess convergence to an efficient policy. We compare our method to current state-of-the-art architecture baselines: CARE [4], Soft Modularization [25], and MT-SAC. The choice of baselines is described in detail in Section IV-B.

We also assess which variant of PTSL yields the best results, and evaluate whether having a shared or independent projection is better and what residual function should be used. The residual functions considered are no residual, the addition, the learnable sum, and the learnable projection (See Section III-C.3 for the definitions).

A. Benchmark

We use Meta-World’s MT10 and MT50 Goal-conditioned tasks as our benchmarks. Meta-World is a multi-task RL benchmark containing 50 robotic manipulation tasks performed by a simulated Sawyer robot. MT10 and MT50 are two evaluation protocols based on Meta-World, where MT10 contains 10 tasks (shown in Fig. 3), and MT50 contains all 50 tasks. The state space is 12-dimensional and consists of tuples of 3D Cartesian end-effector position, 3D Cartesian positions of one or two objects, and the goal position. All our tasks in MT10 and MT50 are goal-conditioned tasks.




Fig. 3: The MT10 benchmark from Meta-World contains 10 tasks: reach, push, pick and place, open door, open drawer, close drawer, press button top-down, insert peg side, open window, and open box.

B. Baselines

We compare our method to the following baselines that are implemented in the CARE and our repository:

MT-SAC: a shared backbone [26] baseline with disentangled alphas [29] that has a simple shared fully connected backbone. There are $N = 3$ hidden layers with a hidden dimension of $H = 400$ resulting in $P = 1,641,222$ parameters for both MT10 and MT50.

Soft Modularization: a soft modularization baseline that learns different policies for each task using a routing network. With also $N = 3$ hidden layers with a hidden dimension of $H = 400$ resulting in $P = 485,766$ parameters for MT10 and $P = 485,806$ for MT50.

CARE: the CARE [4] baseline¹. Also $N = 3$ hidden layers with a hidden dimension of $H = 400$. There are $A = 4$ experts in the Mixture of Experts with a hidden dimension of $H_A = 50$ resulting in $P = 1,871,534$ parameters for both MT10 and MT50.

CARE + PTSL (ours): The proposed method on top of CARE. To keep the comparison fair, we use nearly the same number of parameters and layers as CARE, with $P = 1,871,532$ parameters and $N = 3$ layers. To obtain these numbers, we had to reduce H to 326 and set D to 50 for MT10 (with shared projections) and we kept the same encoder parameters as CARE. For MT50, we reduced D to 32 and H to 274 to reach $P = 1,869,908$ parameters.

CARE + PTSL (shallow) (ours): Variation of PTSL with fewer layers on top of CARE, $N = 2$ hidden layers, $H = 400$, and $D = 50$ resulting in $P = 1,556,454$ parameters. All other parameters are identical to the previous CARE + PTSL architecture for MT10. We show that with fewer layers and less parameters, we obtain similar performance on a short horizon. This method is only tested for MT10.

PTSL only (ours): The proposed PTSL as a simple standalone shared backbone with disentangled alphas. Once

¹Although we used the implementation provided by CARE with the same parameters and the same version of Meta-World, we were unable to reproduce their results even after averaging over 10 or 20 seeds. This is a known issue (see <https://github.com/facebookresearch/mtrl/issues>). The results provided here may appear different from what was published in the original paper. In the rest of this work, for the purpose of comparison, all the reported CARE values will be the results that we were able to reproduce on the same horizon and not the results published in the original paper.




Fig. 4: Training curves of different methods on all benchmarks. For MT10, PTSL converges faster than baselines, and for MT50, we see a gain in sample efficiency. The bolded line represents the mean over $n = 10$ runs for the short horizon and $n = 4$ for the long horizon. The shaded area represents the standard error.

again we set $N = 3$ and we chose H and D to match the number of parameters of CARE. This means that for MT10, $H = 367$ and $D = 50$ resulting in $P = 1,869,833$ parameters. For MT50, $H = 325$ and $D = 32$ resulting in $P = 1,871,187$ parameters.

C. Comparative evaluation

Success — 200K ($n = 10$)	MT10-Cond.	MT50-Cond.
Multitask SAC [11]	0.389 ± 0.029	0.212 ± 0.018
Soft Mod. [25]	0.274 ± 0.033	0.152 ± 0.022
CARE [4]	0.360 ± 0.030	0.320 ± 0.027
CARE + PTSL (Ours)	0.511 ± 0.034	0.370 ± 0.020
CARE + PTSL (Shallow) (Ours)	0.503 ± 0.051	-

TABLE I: Success rate of the baselines on the short horizon (200k steps per task) for MT10 and MT50 Goal-conditioned benchmark. Results are averaged over $n = 10$ seeds for each method. We report the mean and standard error.

Success — MT10-Cond.	After 1M steps	Best
Multitask SAC [11]	0.706 ± 0.050	0.737 ± 0.055
Soft Mod. [25]	0.533 ± 0.039	0.554 ± 0.050
CARE [4]	0.648 ± 0.060	0.683 ± 0.066
CARE + PTSL (Ours)	0.742 ± 0.067	0.772 ± 0.053
PTSL only (Ours)	0.697 ± 0.043	0.721 ± 0.050

TABLE II: Success rate of the baselines on MT10 Goal-conditioned on the long horizon (1M steps per task). Results are reported both at the end of the 1M steps and at the best average value. The Results are averaged over $n = 4$ seeds for each method. We report the mean and standard error.

Figure 4a shows the average success rate on the 10 tasks of the MT10 Goal-conditioned benchmark from Meta-World [11] for MT-SAC, Soft Modularization, CARE, and CARE + PTSL (both deep and shallow architectures). Success rate is noisy since it is a binary variable, so we averaged the results across multiple seeds (the number of seeds is noted as n , set to 10 for short horizon and 4 for long horizon).

We consider 1 million steps as our long horizon, and 200 thousand steps as our short horizon. As noted earlier, all methods are trained using SAC with disentangled alphas.

Success — MT50-Cond.	After 1M steps	Best
Multitask SAC [11]	0.466 ± 0.013	0.489 ± 0.016
Soft Mod. [25]	0.154 ± 0.010	0.206 ± 0.018
CARE [4]	0.388 ± 0.028	0.495 ± 0.024
CARE + PTSL (Ours)	0.354 ± 0.015	0.427 ± 0.020
PTSL only (Ours)	0.610 ± 0.021	0.614 ± 0.020

TABLE III: Success rate of the baselines on MT50 Goal-conditioned on the long horizon (1M steps per task). Results are reported both at the end of the 1M steps and at the best average value. Results are averaged over $n = 4$ seeds for each method. We report the mean and standard error.

As explained in Section IV-B, CARE does not perform as well as described in the original paper [4]. In particular, we noticed decreasing performances on long horizons (See Figure 4d). To remove uncertainty coming from CARE in our experiments, we trained two PTSL agents on long horizons, one with CARE and one without.

Table I and Fig. 4a show that our method outperforms all baselines on the short horizon on MT10. For comparison, in the Meta-World paper [11], it takes around 1.5M steps for the Multitask SAC agent to reach the accuracy that our CARE + PTSL agent reaches within 200K steps. This suggests that our method is highly sample-efficient. Furthermore, we show that using a shallow PTSL network yields very similar results on MT10 after a short horizon, suggesting that lighter PTSL architectures can still generalize.

Table I and Figure 4c also show that PTSL is beneficial as CARE + PTSL outperforms all baselines with more tasks (MT50) even if this means reducing the size of the hidden layers (to keep the same number of parameters).

Tables II and III as well as Figures 4b and 4d show that PTSL is able to learn a good policy on the long horizon. For MT10, CARE + PTSL performs best and reaches a top success rate of 0.772. For MT50, MT-SAC was the best-performing agent since CARE stopped learning after 400K steps. That is why while the CARE + PTSL agent does not perform exceptionally, the standalone PTSL outperforms all methods. Importantly, PTSL achieves a **score of 0.61 on MT50 Goal-Conditioned after only 1 million steps**. This surpasses the reported results from the original CARE paper

[4] of 0.54 after 2 million steps, and the Soft-Modularization paper [25] of 0.60 after 1 million steps **on MT50-Fixed** (which is easier than Goal-Conditioned).

In all settings, adding PTSL to the best-performing method improved the performance, whether it was CARE or a simple Multitask SAC, suggesting PTSL can improve any method.

One issue while training a PTSL architecture is to balance the training of the shared layers and the task-specific ones. Ideally, a shared policy should be learned first and then the task-specific layers should be learned to improve the performance. Our experiments showed that initializing the projection matrices to zero was a good way to enforce this.

D. Ablation study

Success — MT10-Cond.	200K ($n = 10$)
Independent projection	0.369 ± 0.045
Shared projection	0.511 ± 0.034

TABLE IV: Success rate of CARE + PTSL with independent versus shared projections on MT10 Goal-Conditioned. Results are averaged over $n = 4$ seeds for each method. We report the mean and standard error.

Success — MT10-Cond.	200K ($n = 10$)
No residual	0.511 ± 0.034
Learnable sum	0.410 ± 0.031
Learnable projection	0.385 ± 0.032

TABLE V: Success rate of CARE + PTSL with various residual functions. Results are averaged over $n = 4$ seeds for each method. We report the mean and standard error. The Sum residual is not indicated as it did not converge.




Fig. 5: Training curves of different methods on all benchmarks. The bolded line represents the mean over $n = 10$ runs for the short horizon and $n = 4$ for the long horizon. The shaded area represents the standard error.

The result in the previous section suggests that PTSL (with CARE or standalone) is both sample efficient and yields a good policy on long horizons. In this section, we further examine the influence of the different components of the PTSL architecture on the performance of the model.

Shared versus independent projections. In Section III-C.2, we discussed the relevance of using a shared projection in our context since we have only a small number of layers. We decided to compare the two approaches on the short horizon setting with the same parameters: $H = 326$, $D = 50$, and $N = 3$. This means that the independent projection has slightly more parameters (about 12% more). We show the results in Figure 5a.

Figure 5a shows that the shared projection is better than the independent projection, and this is the case even though the independent projection has more parameters. The shared projection is therefore more efficient than the independent projection in our context, as it helps the network to learn a consistent mapping between the shared embedding space and the task-specific embedding space. This result is in agreement with the results from Stickland et. al. [34] that show the shared projection is better than the independent projection for NLP. The shared projection is also more stable than the independent projection as it has a lower variance.

Comparison of residual functions. In Section III-C.3, we discussed the relevance of using residual functions as they have shown great results in other Deep Learning tasks [35]. To verify if residuals are relevant for PTSL, we evaluated four variants on the short horizon setting with the same number of parameters. We used $H = 326$, $D = 50$, and $N = 3$ for all methods except the *Learnable projection* that uses $H = 321$ resulting in roughly 1.871 million parameters.

Figure 5b and Table V shows that having no residuals is better on the short horizon. This means that while the model is more expressive, the additional complexity makes it less sample-efficient and is not beneficial.

V. CONCLUSION

In this work, we present the **Projected Task-Specific Layers (PTSL)**, a novel method inspired by NLP, that surpasses the state-of-the-art on the MT10 and MT50 Goal-Conditioned benchmark from Meta-World.

Here we showed that PTSL learns a high-performing policy faster than other popular methods with higher sample efficiency and without introducing more parameters. Furthermore, PTSL can be integrated *with* existing methods like CARE to improve them. Finally, we showed the benefits of sharing a low-dimensional embedding space between the shared backbone and the task-specific layers, which becomes more obvious with task diversity (MT50).

In future work, we will include transfer learning of individual layers from MT10 to MT50, hierarchical individual layers to better scale, and the implementation of a routing network for the individual layers to better share parameters.

ACKNOWLEDGMENT

The authors would like to thank Saurabh Kumar, Yoni Gozlan, Paul-Emile Giacomelli and Brian Park for their feedback and advice during the preparation of this manuscript.

REFERENCES

- [1] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, pp. 41–75, 1997.
- [2] Y. Zhang and Q. Yang, "A survey on multi-task learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 12, pp. 5586–5609, 2021.
- [3] 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.
- [4] J. P. Shagun Sodhani, Amy Zhang, "Multi-task reinforcement learning with context-based representations," 2021.
- [5] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [6] 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," *arXiv preprint arXiv:2104.08212*, 2021.
- [7] T. Schaul, D. Borsa, J. Modayil, and R. Pascanu, "Ray interference: a source of plateaus in deep reinforcement learning," *arXiv preprint arXiv:1904.11455*, 2019.
- [8] 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.
- [9] C. D'Eramo, D. Tateo, A. Bonarini, M. Restelli, and J. Peters, "Sharing knowledge in multi-task deep reinforcement learning," in *International Conference on Learning Representations*, 2019.
- [10] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [11] Y. T., D. Quillen, H. Z., J. R., N. A., S. H., B. A., H. K., F. C., and L. S., "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning," *arXiv:1910.10897*, 2023.
- [12] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," 2013.
- [13] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [14] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser *et al.*, "Starcraft ii: A new challenge for reinforcement learning," *arXiv preprint arXiv:1708.04782*, 2017.
- [15] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. A. Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2022.
- [16] H. Nguyen and H. La, "Review of deep reinforcement learning for robot manipulation," in *2019 Third IEEE International Conference on Robotic Computing (IRC)*. IEEE, 2019, pp. 590–595.
- [17] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [19] X. B. Peng, G. Berseth, and M. Van de Panne, "Terrain-adaptive locomotion skills using deep reinforcement learning," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, pp. 1–12, 2016.
- [20] W. Zhang, L. Deng, L. Zhang, and D. Wu, "A survey on negative transfer," *IEEE/CAA Journal of Automatica Sinica*, vol. 10, no. 2, pp. 305–329, 2022.
- [21] Z. Xu, K. Wu, Z. Che, J. Tang, and J. Ye, "Knowledge transfer in multi-task deep reinforcement learning for continuous control," *Advances in Neural Information Processing Systems*, vol. 33, pp. 15 146–15 155, 2020.
- [22] A. A. Rusu, S. G. Colmenarejo, C. Gulcehre, G. Desjardins, J. Kirkpatrick, R. Pascanu, V. Mnih, K. Kavukcuoglu, and R. Hadsell, "Policy distillation," *arXiv preprint arXiv:1511.06295*, 2015.
- [23] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Gradient surgery for multi-task learning," *CoRR*, vol. abs/2001.06782, 2020. [Online]. Available: <https://arxiv.org/abs/2001.06782>
- [24] Z. Wang, Y. Tsvetkov, O. Firat, and Y. Cao, "Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models," *arXiv preprint arXiv:2010.05874*, 2020.
- [25] Y. W. Ruihan Yang, Huazhe Xu and X. Wang, "Multi-task reinforcement learning with soft modularization," 2020.
- [26] P. A. Tuomas Haarnoja, Aurick Zhou and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," 2018.
- [27] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [28] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," *CoRR*, vol. abs/1502.05477, 2015. [Online]. Available: <http://arxiv.org/abs/1502.05477>
- [29] T. garage contributors, "Garage: A toolkit for reproducible reinforcement learning research," <https://github.com/rllworkgroup/garage>, 2019.
- [30] G. Cheng, L. Dong, W. Cai, and C. Sun, "Multi-task reinforcement learning with attention-based mixture of experts," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3812–3819, 2023.
- [31] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79–87, 1991.
- [32] —, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79–87, 1991.
- [33] C. Rosenbaum, T. Klinger, and M. Riemer, "Routing networks: Adaptive selection of non-linear functions for multi-task learning," *arXiv preprint arXiv:1711.01239*, 2017.
- [34] A. Stickland and I. Murray, "Bert and pals: Projected attention layers for efficient adaptation in multi-task learning," 2019.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015.