

Gen-HypRL : Generative Policy learning Framework for Multi-Task Reinforcement Learning

Sanket Kalwar^{1*}, Jayaram Reddy^{1*}, Brojeshwar Bhowmick², Arun Singh³, K. Madhava Krishna¹

{sankethkalwar@gmail.com , ramreddyai1010@gmail.com}

¹Robotics Research Center, IIIT Hyderabad

²TCS Research, India

³University of Tartu, Estonia




Figure 1: **Overview:** We present Gen-HypRL, a framework for directly generating policies π using conditioning information (Example: Behavior embeddings for MetaWorld and Velocity inputs for Half-Cheetah). These generated policies can be directly deployed to successfully accomplish the task. On the right side are the directly deployed policies generated by Gen-HypRL on the MetaWorld tasks.

Abstract

A key challenge in building generalist agents is enabling them to perform multiple tasks while simultaneously adapting to variations across the tasks efficiently, particularly in a zero-shot manner. Multi-task Reinforcement Learning (MTRL) is a paradigm that enables agents to learn a single policy that can be deployed to perform multiple tasks in a given environment. A straightforward approach like parameter sharing introduces challenges such as conflicting gradients and determining the optimal way to distribute shared parameters across tasks. In this work, we introduce Gen-HypRL, a framework for training hypernetworks in MTRL that consists of HypLa-tent, an adversarial autoencoder that generates diverse task-conditioned latent policy parameters, and HypFormer, a single-layer transformer that performs soft-weighted aggregation on these priors towards expert policy parameters. Our approach not only outperforms previous hypernetwork based methods but also performs comparably to the existing state-of-the-art methods in MTRL on MetaWorld benchmark. Additionally, experiments on MuJoCo continuous control tasks demonstrate the framework’s strong zero-shot learning capabilities, allowing it to generalize to unseen in-distribution tasks without additional fine-tuning. Our framework also achieves performance comparable to state-of-the-art offline meta-RL

*Equal contribution.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

methods.

Project Code: <https://gen-hyp-rl.github.io/>

INTRODUCTION

Robots deployed in factories and particularly in domestic environments such as houses are expected to perform wide variety of tasks. Each of these tasks may vary in reward functions, dynamics or both and its required to adapt to these variations in the tasks. Developing a framework which can perform multiple tasks while generalizing to variations in these tasks is crucial especially in the environments where changes can arise unexpectedly.

Multi Task reinforcement learning (MTRL) is a paradigm which aims to learn a single policy that can perform multiple tasks. Natural way to solve MTRL is to have shared parameters which captures the common representations such as skills or the objects being manipulated among various tasks (D’Eramo et al. 2024; Yang et al. 2020). Sharing the parameters across various tasks can lead to conflicts in gradients if the tasks are not aligned (Standley et al. 2020; Kendall, Gal, and Cipolla 2018; Chen et al. 2018). This can lead to under-performance on certain tasks. Over the years, many works have tackled this challenge by developing methods that manipulate task-specific gradients to enable efficient learning across multiple tasks (Désidéri 2012; Sener

and Koltun 2018; Yu et al. 2020a; Liu et al. 2021a,b; Navon et al. 2022), and is still an active area of research. On the other hand, CARE(Sodhani, Zhang, and Pineau 2021) proposes learning diverse representations—skills, behaviors, or objects through a mixture of encoders, combined using attention mechanism based on context such as language. MOORE(Hendawy, Peters, and D’Eramo 2023) further enhances representation diversity with Gram-Schmidt orthogonalization. PACO (Sun et al. 2022) learns a policy subspace where task-specific policies are composed by interpolating the learned parameters. However, scaling to many tasks increases the learnable parameters.

Hypernetworks (Ha, Dai, and Le 2016) have gained increasing attention over the years for their soft parameter sharing capabilities and have been applied across diverse domains (Mahabadi et al. 2021; Von Oswald et al. 2019; Ruiz et al. 2024; Alaluf et al. 2022; Zhmoginov, Sandler, and Vladymyrov 2022) but are relatively less explored in the context of Reinforcement learning (Rezaei-Shoshtari et al. 2023; Liang et al. 2024; Zhao et al. 2020; Beck et al. 2023). Hypernetworks generate the weights for a target network, enabling it to adapt to specific tasks or contexts. This capability can be used to generate weights for the related tasks just by conditioning on task-specific information. Leveraging their soft weight-sharing property, we aim to explore the potential of hypernetworks in MTRL and assess their zero-shot generalization capabilities.

To summarize, our key contributions are:

1. We propose **Gen-HypRL**, a framework for training hypernetworks for MTRL. It consists of:
 - **HypLatent**, an adversarial autoencoder which learns to generate **diverse** latent policy parameters conditioned on the task information.
 - **HypFormer** is a single-layer transformer network that performs **soft-weighted aggregation** on the prior generated by *HypLatent*, refining it towards expert policy parameters.
2. We regularize hypernetworks by reconstructing behavior embeddings. Here, the hypernetwork generates task-specific weights while a discriminator ensures consistency with the trajectory distribution. This method improves generalization by encouraging coherent and adaptable embeddings across tasks.
3. Experiments in MuJoCo control tasks shows that our framework has impressive zero-shot generalization capabilities on unseen tasks that are in distribution.

METHOD

We propose a three-stage pipeline as shown in Fig 2. In the first stage, an autoencoder (Kingma 2013) maps policy parameters to the latent space, whose latent features are then used in the second stage, *HypLatent* which is based on Adversarial autoencoder (Makhzani et al. 2015). *HypLatent* seeks to approximate the autoencoder’s latent distribution through adversarial training. The generator aims to map behavior embeddings and noise sampled from a normal distribution, where the inclusion of noise enables the generator

to learn diverse latent features. These diverse features are then utilized in the third stage, *HypFormer*. *HypFormer* uses the diverse latent features generated by *HypLatent* for each behavior embedding that performs soft-weighted aggregation and refines them towards expert policy parameters. A transformer encoder applies self-attention across both the latent features and behavior embeddings, grounding the latent features through behavior embedding interactions. Ultimately, the latent features are guided by learned residuals. The subsequent sections provide detailed discussions on the autoencoder training, behavior embeddings, *HypLatent*, and *HypFormer*.

Stage 1: Autoencoder and Behaviour Embedding:

Autoencoder: Takes trained SAC policies of multiple tasks as input, mapping them to a latent space. These policies are randomly sampled from various tasks during training and passed through autoencoder, which learns a latent space that captures the distribution of task-specific policies. It consists of encoder network E_θ which takes the policy parameters X as input and outputs latent representation of the policy parameters Z , and a decoder network D_θ which takes the output latent Z from the encoder E_θ and reconstructs the policy parameters $\hat{X} = D_\theta(E_\theta(X))$. The objective function for training autoencoder is shown in equation (1).

$$L_{autoencoder} = \frac{\lambda}{M} \sum_{i=1}^M (\hat{X}^i - X^i)^2 \quad (1)$$

where M is the number of training samples for SAC policies, and λ is the scaling factor set to $1e3$. This will learn the distribution of policy parameters effectively and the encoder’s output is used in Stage-2 by *HypLatent* during training to generate diverse samples of latents.

Behaviour Embedding: These embeddings are used to provide conditional information to *HypLatent*, enabling the generation of diverse, task-specific policies. This approach is adapted from MakeAnAgent (MAA) (Liang et al. 2024). Consider the trajectory of N steps. We now define the n -step trajectory $\tau^n = (s_1, a_1, a_2, a_3, \dots, s_{n-2}, a_{n-2}, a_{n-1}, a_n)$, and post-success states $\hat{\tau} = (s_K, s_{K+1}, s_{K+2}, \dots, s_{K+M})$, which are collected after the success step K until M fixed steps. The objective is to maximize the mutual information $I(\hat{\tau}; \tau^n)$ between post-success states and the n -step trajectory. We train behaviour embeddings using the contrastive loss defined in Equation (2):

$$L_{behaviour} = -\frac{1}{N} \sum_{i=1}^N \log \frac{h_i^T W v_i}{\sum_{j=1}^N h_j^T W v_j} \quad (2)$$

where, $h_i = \phi(\tau_i^n)$ and $v_i = \psi(\hat{\tau}_i)$ along with learned similarity weights W between h_i and v_i which finally forms behaviour embedding $\tau_e = (h_i, v_i)$.

After retrieving latent policy parameter Z from autoencoder and behaviour embedding τ_e , both are used in the stage-2 for training *HypLatent* network.

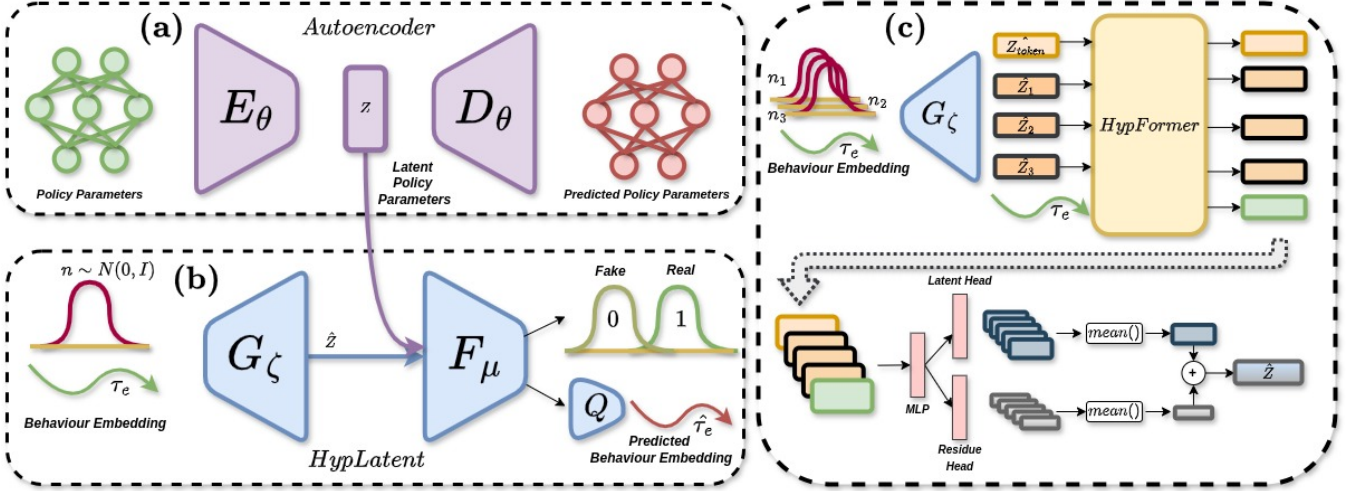


Figure 2: **Overview:** The figure in a section (a) consists of an *Autoencoder* which learns to map policy parameters to latent space, more details is discussed at . Section (b) shows *HypLatent* architecture in which there is a Generator G_ζ and a Discriminator F_μ , objective of a generator is to learn latent policy parameter similar to the *Autoencoder's* encoder output, given a behaviour embedding τ_e and noise $n \sim N(0, I)$, discriminator assesses whether the samples generated by the generator belong to the true latent policy parameter distribution, also there is a auxillary network Q which reconstructs the trajectory embedding given the output features of the discriminator, refer for in detail explanation. We also show *HypFormer* in the section (c) where we use the trained generator to produce diverse latent policy parameters. For a given behavior embedding τ_e , multiple noise samples $n_1, n_2, n_3 \sim N(0, I)$ are used to generate multiple latent policy parameters. These are then processed by *HypFormer*, which predicts the ground truth latent policy parameters using the latent head and residue head. For a detailed explanation, see Section .

Stage 2: *HypLatent* Generative Adversarial Network

We use the *HypLatent* to learn the latent manifold Z from the autoencoder's encoder output. This is conditioned on the behavior embedding τ_e and a sampled noise $n \sim N(0, I)$. The network consists of generator G_ζ which is conditioned on τ_e along with sampled noise n to generate \hat{Z} , i.e $\hat{Z} = G_\zeta(\tau_e, n)$, and a discriminator F_μ tries to predict whether the generated \hat{Z} belong to real latent policy parameter manifold, i.e $p_Z = F_\mu(\hat{Z})$. For training G_ζ we minimize the objective function V_ζ below,

$$V_\zeta(G_\zeta, F_\mu) = \nabla_\zeta \frac{1}{N} \sum_{i=1}^N \log(1 - F_\mu(\hat{Z}^i)) \quad (3)$$

And, maximize the objective V_μ for F_μ ,

$$V_\mu(G_\zeta, F_\mu) = \nabla_\mu \frac{1}{N} \sum_{i=1}^N [\log(F_\mu(Z^i)) + \log(1 - F_\mu(\hat{Z}^i))] \quad (4)$$

For further grounding the generated latent \hat{Z} , we regularize the generator and discriminator by reconstructing the input τ_e back from the discriminator auxiliary head Q . We

do this by minimizing L_{reg} in Equation (5).

$$L_{reg} = \frac{\lambda_{reg}}{N} \sum_{i=1}^N (Q(F_\mu(\hat{Z}^i)) - \tau_e^i)^2 \quad (5)$$

Now, we use the generator G_ζ to generate latent policy parameter \hat{Z} . Using G_ζ alone will generate diverse latent policy parameters, resulting in various behaviors, meaning different ways of performing the same task but won't guarantee high success rate on an average. *HypFormer* tackles this by using **soft-weighted aggregation** to refine the generated latent policy parameters toward the expert policy parameters, as detailed in the section below.

Stage 3: *HypFormer*

In this stage, we use the generator G_ζ to generate latent policy parameters $S_{\hat{Z}} = (\hat{Z}_1, \hat{Z}_2, \hat{Z}_3, \dots, \hat{Z}_L)$ by conditioning it on the trajectory embedding τ_e and sampled noise $(n_1, n_2, \dots, n_L) \sim N(0, I)$. $S_{\hat{Z}}$ is then input to *HypFormer* along with the learnable token \hat{Z}_{token} and the behaviour embedding τ_e . Now, *HypFormer* takes the combination $(\hat{Z}_{token}, S_{\hat{Z}}, \tau_e)$ as input tokens, applies self-attention to produce enhanced latent policy parameters for each trajectory τ_e . This self-attention helps us to perform soft-weighted aggregation of latent policy parameters. Applying MSE loss to align *HypFormer* predictions with ground truth resulted in unstable training. To address this, the enhanced tokens are processed by a shared MLP, which

then splits into two branches: the latent Head and the residue Head. Latent Head learns to predict ground truth latent policy parameters, while the residue Head predicts the residue between Latent Head prediction and ground truth latent policy parameters, after which residue tokens and predicted policy tokens are averaged to get per trajectory single residue token ∇_{pred}^i and a single latent policy parameter token Z_{pred}^i . As we have ground truth latent policy parameter token Z_{gt}^i along with the residue token ∇_{gt}^i for each trajectory τ_e^i , We apply cosine similarity loss between Z_{gt}^i and Z_{pred}^i .

$$L_{sim} = \frac{1}{N} \cdot \sum_{i=1}^N \left(1 - \frac{Z_{pred}^i \cdot Z_{gt}^i}{\max(\|Z_{pred}^i\|_2, \|Z_{gt}^i\|_2, \epsilon)} \right) \quad (6)$$

For residue token prediction we minimize L_{res} as follows,

$$L_{res} = \frac{1}{N} \sum_{i=1}^N (\nabla_{pred}^i - \nabla_{gt}^i)^2 \quad (7)$$

Finally to ensure that residue head and latent head output's are consistent with each other consistency loss is applied on $\hat{Z}_{pred}^i = Z_{pred}^i + \nabla_{pred}^i$,

$$L_c = \frac{1}{N} \sum_{i=1}^N (\hat{Z}_{pred}^i - Z_{gt}^i)^2 \quad (8)$$

So, final objective to minimize is as follows,

$$L_{HypFormer} = \lambda_{sim} L_{sim} + \lambda_{res} L_{res} + \lambda_c L_c \quad (9)$$

The size of L is typically a power of 2; in our case, it is set to 2^3 during training. N is the number of trajectories in the batch while training. λ_{sim} , λ_{res} , λ_c are all set to $1e3$.

EXPERIMENTAL SETUP

In our experiments, we aim to evaluate and answer the following: 1.) Performance of our method in Multi-task Reinforcement Learning (MTRL) on seen tasks. 2.) Does our method zero-shot generalize to related but unseen tasks which are in-distribution? 3.) How well do the representations learned by the hypernetwork perform on entirely unseen tasks, specifically in terms of out-of-distribution task performance in MTRL?

Datasets and Environments details

We evaluate MTRL performance on MetaWorld (Yu et al. 2020b) and zero-shot generalization capabilities on MuJoCo control tasks.

1.) **MetaWorld:** The MAA (MakeAnAgent) splits are utilized for training and evaluating MTRL performance on both seen and unseen tasks. Dataset is sourced from MAA.

2.) **Cheetah-Vel:** Following (Mitchell et al. 2021; Xu et al. 2022), the task involves achieving target velocities sampled from [0-3], with 35 velocities in the training set and 5 in the test set. Ablation studies are conducted on training with 10, 20, and 25 random velocities, while the test set remains unchanged, demonstrating the framework's generalization and

sample efficiency. We collect a dataset of 800 expert policies for each training velocity by training SAC (Haarnoja et al. 2018). The first policy is saved at 150k training steps, followed by checkpoints taken every 500 training steps thereafter.

RESULTS AND ANALYSIS

We train the *HypLatent* using a 1D UNet architecture (Ronneberger, Fischer, and Brox 2015) for the generator and a 3-layer MLP for the discriminator. For regularization, an auxiliary network consisting of a 2-layer MLP is used. The noise sample size is set to 128 for MetaWorld tasks and 4 for Half-Cheetah. The behavior embedding size is 128 for MetaWorld, while for Half-Cheetah, the conditioning information is represented by a single scalar value (velocity).

We utilize a single-layer Transformer (Vaswani 2017) encoder without positional encoding, with a token size of 256 and 128 heads in the multi-head attention layer.

RESULTS AND ANALYSIS

We qualitatively verify Gen-HypRL performance on the MetaWorld and MuJoCo continuous control task (Half-Cheetah) as shown in the Fig.1 and Fig.3. Policy generated by Gen-HypRL when deployed exhibits correct behaviour qualitatively.

Quantitative Analysis

We quantitatively evaluate our framework on the MetaWorld and Cheetah-vel datasets to demonstrate its effectiveness in both Multi-task reinforcement learning (MTRL) and zero-shot generalization in RL. For the MetaWorld dataset, the success rate is measured by the proportion of the task completed given a trajectory which is applicable to both seen and unseen tasks. Specifically, we assess our method's performance in MTRL by evaluating it on the seen MetaWorld tasks and then testing the zero-shot generalization capability of the learned representations on the unseen MetaWorld tasks. We show the comparison of our method directly with MAA as it is the closest method to our approach. In Table 1, it is evident that our Gen-HypRL is better than MAA by 34.3% directly and Gen-HypRL without *HypFormer* is better by 15% on MetaWorld Seen test tasks. Table 2 shows that our top-5 and top-10 generated policies achieve 100% success rate on seen test environments. For more broad comparison, we have compared against CARE (Sodhani, Zhang, and Pineau 2021), Decision Transformer DT (Chen et al. 2021). Table 3 shows the success rate on completely unseen tasks of MetaWorld.

We evaluate zero-shot generalization capability on the Cheetah-vel dataset by averaging the returns across the test velocity seeds-(2, 7, 15, 23, 26) following (Mitchell et al. 2021; Xu et al. 2022). Our proposed model, Gen-HypRL, without *HypFormer*, achieves the best average return of -33.44 when trained using only 10 random velocities from the training set, outperforming state-of-the-art methods such as MACAW (Mitchell et al. 2021) and Prompt-DT




Figure 3: Qualitative results for Half-Cheetah showcasing policies generated by Gen-HypRL at two extreme, previously unseen velocities, along with their corresponding rewards.

(Xu et al. 2022), which utilize 35 seeds for training. When the number of training seeds is scaled from 10 to 35, the average return further improves to -29.6, emphasizing the robust zero-shot generalization capability of our framework.

Table 1: Gen-HypRL Success Rate(%) on the MetaWorld dataset. **Bold** numbers highlights the top achieved success-rate on the task, while the *italics* shows the 2nd best achieved success-rate.

MTRL Tasks	MAA	Gen-HypRL (w/o HypFormer)	Gen-HypRL (w/ HypFormer)
window-open	33	<i>51</i>	64
door-open	27	35	62
drawer-open	42	40	78
dial-turn	23	<i>36</i>	48
plate-slide	45	<i>66</i>	88
button-press	32	38	58
handle-press	50	<i>62</i>	82
faucet-close	45	<i>77</i>	82
Avg. Success Rate	36	<i>51</i>	70.3

Ablation Studies

We present design decisions of Gen-HypRL framework in the context of MetaWorld tasks. We demonstrate the impact of number of tokens on task performance and provide PCA analysis to further substantiate our design choices.

a) Varying number of Tokens: We train *HypFormer* using 8 tokens, which results in the best performance on the seen MTRL tasks. This is demonstrated in Figure 5, where the left image with the blue bar graph highlights this setup. As number of tokens increases, the performance on MTRL tasks decreases asymptotically. In contrast, for unseen tasks, the model achieves optimal performance with 32 tokens.

Table 2: Gen-HypRL Success Rate(%) on MetaWorld dataset. **Bold** numbers highlights the top achieved success-rate, while the *italics* shows the 2nd best achieved success-rate. Unless explicitly stated otherwise, such as for top-10 or top-5, the success rate reported in table below represents the average over 100 policies.

Methods	Seen Task	Unseen Task
CARE	<i>82.1</i>	58.5
DT	80.3	60.4
MAA	36	16.25
Gen-HypRL w/o HypFormer	51	12.13
Gen-HypRL	70.3	19.8
Gen-HypRL w/o HypFormer (top 10)	100	54.38
Gen-HypRL w HypFormer (top 10)	100	80.63
Gen-HypRL w/o HypFormer (top 5)	100	75.1
Gen-HypRL w HypFormer (top 5)	100	87.5

Increasing the number of tokens from 8 to 32 results in a corresponding improvement in performance, but beyond 32, the performance begins to degrade as number of tokens continues to increase.

b) PCA Analysis: We present a PCA analysis on the latent policy parameters generated by *HypLatent* and *HypFormer*, as illustrated in Figure 4 for three seen tasks from MetaWorld: 'Button-Press,' 'Dial-Turn,' and 'Door-Open.' In the 2D PCA plots, the latent policy parameters predicted by *HypFormer* are more closely aligned with the ground truth compared to those generated by *HypLatent*, demonstrating the effectiveness of *HypFormer*.

Table 4 demonstrates the effectiveness of *HypLatent* in the Cheetah-vel setup. In summary, *HypLatent*, together with *HypFormer*, constitutes a key component of the Gen-HypRL design framework.

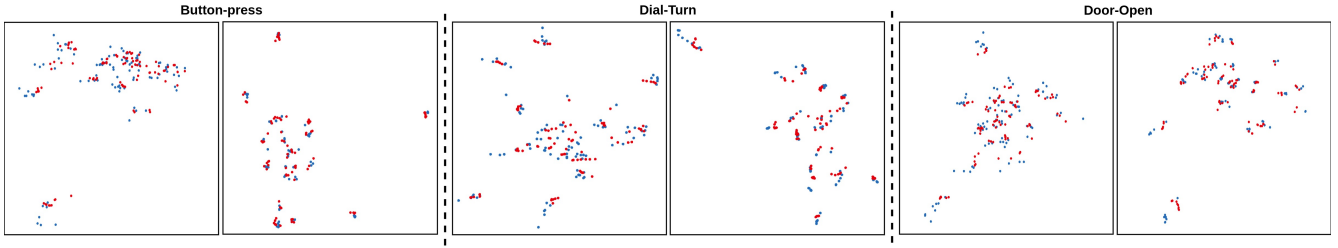


Figure 4: Visualization of the predicted latent policy parameters for three MetaWorld tasks. In each task, the image on the left represents the output of the *HypLatent* generator, while the image on the right shows the output of *HypFormer*. The red color points are the predicted latent policy parameters while the blue color points indicate the ground-truth latent policy parameters.

Table 3: Success Rate(%) on unseen tasks of MetaWorld. **Bold** numbers highlights the top achieved success-rate on the task, while the *italics* shows the 2nd best achieved success-rate.

Zero-Shot RL Tasks	MAA	Gen-HypRL (w/o HypFormer)	Gen-HypRL (w/ HypFormer)
drawer-close	55	53	80
handle-press-side	4	6	0
door-lock	13	6	6
window-close	<i>10</i>	0	12
reach-wall	13	6	<i>10</i>
coffee-button	8	3	9
button-press-wall	<i>11</i>	2	14
faucet-open	<i>16</i>	21	27
Avg.Success Rate	<i>16.25</i>	12.13	19.8

Table 4: Comparing Gen-HypRL to the offline Meta-RL works on Cheetah-Vel task.

Methods	Avg. Return
MACAW (Iter. 0)	-121.6
MACAW (Iter. 20K)	-60.5
Prompt-DT	-34.43
Gen-HypRL (w/o HypFormer) (10 seeds)	-33.44
Gen-HypRL (w/o HypFormer) (35 seeds)	-29.6




Figure 5: Performance of Gen-HypRL on MetaWorld MTRL and Zero-Shot RL tasks on varying token size from 8 to 512. Better viewed at 2x zoom.

CONCLUSION

In this paper, we propose a framework for training hypernetworks in the context of Multi-Task Reinforcement Learn-

ing (MTRL). By training a prior over the various task policies in an adversarial fashion, we encourage diversity of the generated latent policy parameters. We then use a single layer transformer architecture to guide the prior towards expert policy parameters. Our framework outperforms related hypernetwork-based baselines in MTRL and achieves performance comparable to state-of-the-art MTRL methods. Additionally, our experiments on MuJoCo control tasks demonstrate that the framework exhibits strong zero-shot generalization to unseen tasks within the same task distribution.

We believe that our framework is an important step towards tackling MTRL while retaining zero shot generalizability to in-distribution tasks.

FUTURE WORK

We currently rely on expert policies during training, making it crucial to explore performance when training with sub-optimal policies. Another promising future direction is to enhance the scalability of the framework to handle a larger number of tasks.

References

- Alaluf, Y.; Tov, O.; Mokady, R.; Gal, R.; and Bermano, A. 2022. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18511–18521.
- Beck, J.; Jackson, M. T.; Vuorio, R.; and Whiteson, S. 2023. Hypernetworks in meta-reinforcement learning. In *Conference on Robot Learning*, 1478–1487. PMLR.
- Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.
- Chen, Z.; Badrinarayanan, V.; Lee, C.-Y.; and Rabinovich, A. 2018. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, 794–803. PMLR.
- D’Eramo, C.; Tateo, D.; Bonarini, A.; Restelli, M.; and Peters, J. 2024. Sharing knowledge in multi-task deep reinforcement learning. *arXiv preprint arXiv:2401.09561*.

- Désidéri, J.-A. 2012. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *Comptes Rendus Mathématique*, 350(5-6): 313–318.
- Ha, D.; Dai, A.; and Le, Q. V. 2016. Hypernetworks. *arXiv preprint arXiv:1609.09106*.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Hendawy, A.; Peters, J.; and D’Eramo, C. 2023. Multi-task reinforcement learning with mixture of orthogonal experts. *arXiv preprint arXiv:2311.11385*.
- Kendall, A.; Gal, Y.; and Cipolla, R. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7482–7491.
- Kingma, D. P. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Liang, Y.; Xu, T.; Hu, K.; Jiang, G.; Huang, F.; and Xu, H. 2024. Make-An-Agent: A Generalizable Policy Network Generator with Behavior-Prompted Diffusion. *arXiv preprint arXiv:2407.10973*.
- Liu, B.; Liu, X.; Jin, X.; Stone, P.; and Liu, Q. 2021a. Conflict-averse gradient descent for multi-task learning. *Advances in Neural Information Processing Systems*, 34: 18878–18890.
- Liu, L.; Li, Y.; Kuang, Z.; Xue, J.; Chen, Y.; Yang, W.; Liao, Q.; and Zhang, W. 2021b. Towards impartial multi-task learning. *iclr*.
- Mahabadi, R. K.; Ruder, S.; Dehghani, M.; and Henderson, J. 2021. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489*.
- Makhzani, A.; Shlens, J.; Jaitly, N.; Goodfellow, I.; and Frey, B. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Mitchell, E.; Rafailov, R.; Peng, X. B.; Levine, S.; and Finn, C. 2021. Offline meta-reinforcement learning with advantage weighting. In *International Conference on Machine Learning*, 7780–7791. PMLR.
- Navon, A.; Shamsian, A.; Achituve, I.; Maron, H.; Kawaguchi, K.; Chechik, G.; and Fetaya, E. 2022. Multi-task learning as a bargaining game. *arXiv preprint arXiv:2202.01017*.
- Rezaei-Shoshtari, S.; Morissette, C.; Hogan, F. R.; Dudek, G.; and Meger, D. 2023. Hypernetworks for zero-shot transfer in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 9579–9587.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Ruiz, N.; Li, Y.; Jampani, V.; Wei, W.; Hou, T.; Pritch, Y.; Wadhwa, N.; Rubinstein, M.; and Aberman, K. 2024. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6527–6536.
- Sener, O.; and Koltun, V. 2018. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31.
- Sodhani, S.; Zhang, A.; and Pineau, J. 2021. Multi-task reinforcement learning with context-based representations. In *International Conference on Machine Learning*, 9767–9779. PMLR.
- Standley, T.; Zamir, A.; Chen, D.; Guibas, L.; Malik, J.; and Savarese, S. 2020. Which tasks should be learned together in multi-task learning? In *International conference on machine learning*, 9120–9132. PMLR.
- Sun, L.; Zhang, H.; Xu, W.; and Tomizuka, M. 2022. Paco: Parameter-compositional multi-task reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 21495–21507.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Von Oswald, J.; Henning, C.; Grewe, B. F.; and Sacramento, J. 2019. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*.
- Xu, M.; Shen, Y.; Zhang, S.; Lu, Y.; Zhao, D.; Tenenbaum, J.; and Gan, C. 2022. Prompting decision transformer for few-shot policy generalization. In *international conference on machine learning*, 24631–24645. PMLR.
- Yang, R.; Xu, H.; Wu, Y.; and Wang, X. 2020. Multi-task reinforcement learning with soft modularization. *Advances in Neural Information Processing Systems*, 33: 4767–4777.
- Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020a. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33: 5824–5836.
- Yu, T.; Quillen, D.; He, Z.; Julian, R.; Hausman, K.; Finn, C.; and Levine, S. 2020b. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, 1094–1100. PMLR.
- Zhao, D.; Kobayashi, S.; Sacramento, J.; and von Oswald, J. 2020. Meta-learning via hypernetworks. In *4th Workshop on Meta-Learning at NeurIPS 2020 (MetaLearn 2020)*. NeurIPS.
- Zhmoginov, A.; Sandler, M.; and Vladymyrov, M. 2022. Hypertransformer: Model generation for supervised and semi-supervised few-shot learning. In *International Conference on Machine Learning*, 27075–27098. PMLR.