# Scalable Multi-Objective and Meta Reinforcement Learning via Gradient Estimation

**Zhenshuo Zhang, Minxuan Duan, Youran Ye, Hongyang R. Zhang**

Northeastern University, Boston, Massachusetts
{zhang.zhens, duan.mi, ye.you, ho.zhang}@northeastern.edu

## Abstract

We study the problem of efficiently estimating policies that simultaneously optimize multiple objectives in reinforcement learning (RL). Given $n$ objectives (or tasks), we seek the optimal partition of these objectives into $k \ll n$ groups, where each group comprises related objectives that can be trained together. This problem arises in applications such as robotics, control, and preference optimization in language models, where learning a single policy for all $n$ objectives is suboptimal as $n$ grows. We introduce a two-stage procedure—meta-training followed by fine-tuning—to address this problem. We first learn a meta-policy for all objectives using multitask learning. Then, we adapt the meta-policy to multiple randomly sampled subsets of objectives. The adaptation step leverages a first-order approximation property of well-trained policy networks, which is empirically verified to be accurate within a 2% error margin across various RL environments. The resulting algorithm, PolicyGradEx, efficiently estimates an aggregate task-affinity score matrix given a policy evaluation algorithm. Based on the estimated affinity score matrix, we cluster the $n$ objectives into $k$ groups by maximizing the intra-cluster affinity scores. Experiments on three robotic control and the Meta-World benchmarks demonstrate that our approach outperforms state-of-the-art baselines by 16% on average, while delivering up to $26\times$ faster speedup relative to performing full training to obtain the clusters. Ablation studies validate each component of our approach. For instance, compared with random grouping and gradient-similarity-based grouping, our loss-based clustering yields an improvement of 19%. Finally, we analyze the generalization error of policy networks by measuring the Hessian trace of the loss surface, which gives non-vacuous measures relative to the observed generalization errors.

## Introduction

Reinforcement learning (RL) is a technique for sequential decision-making in interactive environments, enabling agents to learn from feedback signals. A central challenge in RL is to develop agents that generalize across a wide range of tasks rather than solving a single task in isolation. For instance, a robot needs to master multiple related skills that share an underlying structure (Yu et al. 2020b). Prior work has studied this problem for two closely related settings. The first, multitask reinforcement learning, seeks to find a single policy that performs well across a given set of objectives (or tasks) (Sodhani, Zhang, and Pineau 2021; Joshi et al. 2025). The second, meta-reinforcement learning, aims to learn a meta-initialization that can rapidly adapt to unseen tasks (Finn, Abbeel, and Levine 2017; Rakelly et al. 2019). In both settings, however, the computational cost of evaluating and adapting a shared policy over many competing objectives grows quickly with $n$.

This paper studies the problem of designing a (scalable) multi-objective optimizer for reinforcement learning, with extensions to policy meta-adaptation. Learning from a diverse set of tasks has been shown to improve generalization (Thrun and Pratt 1998; Cobbe et al. 2019; Yin et al. 2020). However, the main difficulty is modeling task relationships and how different tasks transfer information to one another when trained together in a complex network (Wu, Zhang, and Ré 2020). Existing approaches have considered measuring the similarity between gradient vectors across tasks to quantify task relationships (Yu et al. 2020a). Notably, this gradient-similarity-based measure corresponds to pairwise task-affinity scores, whereas in practice it is natural to model how multiple policies interact when trained together (Li, Nguyen, and Zhang 2023). Exhaustively searching over all subsets of $\{1, 2, \ldots, n\}$ is computationally infeasible: evaluating each subset typically requires a full training procedure, and the number of subsets grows exponentially in $n$ (Caruana 1997). Greedy stepwise selection requires evaluation on $O(n^2)$ subsets, which is again impractical for large $n$ (Li et al. 2023; Li, Sharma, and Zhang 2024).

To overcome these challenges, we introduce a scalable algorithm for estimating policy performance on arbitrary task subsets without full training. Our algorithm starts by learning a single meta-policy across all tasks, then uses a first-order surrogate model—derived from a Taylor approximation around the meta-policy to efficiently estimate the outcome of fine-tuning on any subset. This approach is analogous to a random ensemble (Zhang et al. 2025), which allows us to compute a pairwise task affinity matrix and partition the $n$ objectives into $k$ subgroups using a convex clustering procedure. The resulting groups can be trained separately using any multitask or meta-RL optimizers. This overall approach is illustrated in Figure 1.

We validate our approach through extensive experiments on both Meta-World (Yu et al. 2020b) and robotic control

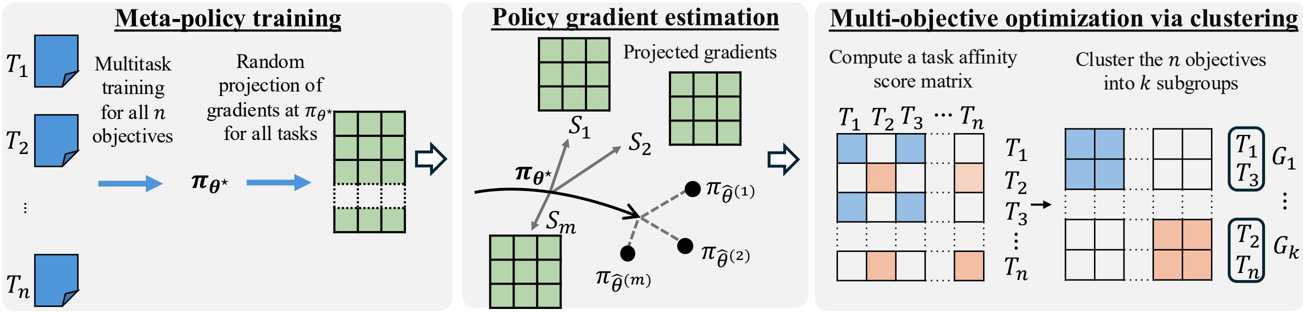


Figure 1: An overview of our approach. *Left:* Run multitask training on all the tasks, $T_1, T_2, \ldots, T_n$, and obtain to a meta-initialization policy $\pi_{\theta^\star}$. Store the projected gradients of a surrogate loss with meta-policy $\theta^\star$ for every transition from $t = 1, 2, \ldots, N$. *Middle:* Estimate policy adaptation performance on $m$ task subsets using projected gradients as features in logistic regression. *Right:* Compute an $n \times n$ task affinity score matrix based on the estimated loss values. Lastly, run a clustering algorithm to group similar objectives, resulting in $k$ subgroups $G_1, \ldots, G_k$, each of which share the same policy within group.

benchmarks (Towers et al. 2025). We find that the first-order approximation applied to the outputs of a policy network yields accurate estimates—typically within 2% error margin—while being up to 26× faster than full training. Our algorithm consistently outperforms multitask and meta-RL baselines (Yang et al. 2020; Sodhani, Zhang, and Pineau 2021; Sun et al. 2022) by 16% on average, and surpasses baseline grouping strategies (such as random grouping or gradient-similarity based grouping) by 23% and 16%, respectively. These results highlight the effectiveness and scalability of our task-affinity estimation approach.

Lastly, we conduct a theoretical analysis of the generalization errors of policy networks. The main technical approach is to compute (and measure) the trace of the Hessian of the loss surface (Ju, Li, and Zhang 2022; Ju et al. 2023; Zhang, Li, and Ju 2024). See Theorem 1 for the precise statement. When evaluated on the above RL environments, we find that the Hessian-based generalization bounds match the scale of empirical generalization errors (see Figure 2).

In summary, the contributions of this paper are threefold. (1) We propose a first-order gradient estimation algorithm to estimate policy performance on any given task subset efficiently. Based on this gradient estimation, we design a scalable algorithm that computes an $n \times n$ task affinity score matrix and partitions $n$ objectives into $k$ clusters via convex relaxation. (2) We conduct extensive experiments to validate the efficiency of our approach for several multi-objective and meta reinforcement learning benchmarks. The code for reproducing our results can be accessed at: https://github.com/VirtuosoResearch/PolicyGradEx. (3) Finally, we show non-vacuous bounds on the generalization error of the loss surface of policy networks via a PAC-Bayes analysis, and report empirical estimates of the Hessian trace.

## Problem Setup

We consider a Markov decision process in which an agent must learn a set of $n$ tasks $\mathcal{T} = \{T_1, T_2, \ldots, T_n\}$. Each task $T_i$ is modeled as a Markov decision process (MDP) $T_i = (\mathcal{S}, \mathcal{A}, P, P_0, r_i, \gamma)$, where all tasks share the same state space $\mathcal{S}$, action space $\mathcal{A}$, transition model $P$, and initial-state distribution $P_0$, discount factor $\gamma$, but differ in their reward functions $r_i$. For a policy $\pi_\theta$ with parameters $\theta$, we define its task-specific expected reward as

$$R_i(\theta) = \mathbb{E}_{\pi_\theta}\left[\sum_{t=0}^{\infty} \gamma^t\, r_i(s_t, a_t)\right].$$

In multi-objective RL, the goal is to learn a single policy that performs well across all tasks. A natural aggregate objective is the average reward

$$R_{\mathcal{T}}(\theta) = \frac{1}{n}\sum_{i=1}^{n} R_i(\theta).$$

More generally, for any subset $S \subseteq T$, we define

$$R_S(\theta) = \frac{1}{|S|}\sum_{T_i \in S} R_i(\theta).$$

Pairwise task interactions can be characterized through this notion. For example, $R_{\{i,j\}}(\theta)$ captures the joint performance on tasks $T_i$ and $T_j$, providing a meaningful measure of their affinity. A key challenge in multitask learning is *negative transfer*, where conflicting gradients from two tasks degrade optimization efficiency and final performance (Wu, Zhang, and Ré 2020). This occurs when

$$R_{\{i,j\}}(\theta) < \left(R_i(\theta) + R_j(\theta)\right)/2,$$

indicating that training the pair together is detrimental compared to training each task separately.

Our goal is to model the underlying relationship between the $n$ tasks, which can be used to group similar tasks or select a small representative subset for meta-learning. However, searching for the optimal partition requires evaluating $R_S(\theta)$ for multiple subsets $S \subseteq \mathcal{T}$, since the number of possible subsets is $2^n$. Greedy stepwise selection again requires evaluating $\mathcal{O}(n^2)$ subsets, where each evaluation typically involves a full RL training or adaptation procedure (Li et al. 2023, 2024). This is again intractable for large $n$.

To address this challenge, we introduce **an efficient gradient-estimation algorithm to approximate the loss of any subset trained on a policy network without repeated training**. This enables us to scale up multi-objective RL to a large number of objectives $n$, as we will describe next.

# Our Approach

This section describes our proposed algorithm for multi-objective reinforcement learning. The algorithm consists of two main stages. We first find a meta-policy $\theta^\star$ trained using all tasks. Then, compute a surrogate model to estimate the adapted performance of $\theta^\star$ on a subset of $\{1, 2, \ldots, n\}$. This surrogate modeling framework enables efficient estimation of subset combinations without repeated training. In the second stage, we estimate the fine-tuning performance $R_{S_i}$ on $m$ randomly chosen subsets $S_1, S_2, \ldots, S_m$ of $\mathcal{T}$, for every $1 \le i \le m$. For each subset, we solve a weighted logistic regression problem to estimate the adaptation parameters $\hat{\theta}_{S_i}$ and then estimate an approximate loss. Importantly, we run this second step using precomputed gradients and functional values at $\theta^\star$ from the first stage. The key technique that enables this gradient estimation is **a first-order approximation property of the policy reward** that we also empirically verify for various RL settings.

## Surrogate Modeling

Our goal is to estimate the output of a policy update starting from the meta-initialization $\theta^\star$. To do so efficiently, we linearize the policy gradient objective around $\theta^\star$ and reformulate the one-step update as a weighted logistic regression problem.

We begin with the standard policy gradient objective used in algorithms such as Proximal Policy Optimization (PPO) (Schulman et al. 2017):

$$J(\theta) = \mathbb{E}_t \left[ r_t(\theta) \hat{A}_t \right], \text{ where } r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta^\star}(a_t|s_t)}$$

is the probability ratio and $\hat{A}_t$ is the estimated reward for the state-action pair $(s_t, a_t)$ at time $t$. We perform a first-order Taylor expansion of the log-probability $\log \pi_\theta(a_t|s_t)$ around the initial parameters $\theta^\star$. Let $\Delta\theta = \theta - \theta^\star$ and define the gradient feature vector

$$g_t = \nabla \log \pi_\theta(a_t|s_t)\big|_{\theta=\theta^\star}.$$

The first-order Taylor expansion of the above log-probability-ratio is given by:

$$\begin{aligned} \log r_t(\theta) &= \log \pi_\theta(a_t|s_t) - \log \pi_{\theta^\star}(a_t|s_t) \\ &= g_t^\top \Delta\theta + \epsilon, \end{aligned} \tag{1}$$

where $\epsilon$ refers to the approximation error of the expansion. Provided that $\Delta\theta$ is small relative to $\theta^\star$, we have another approximation as (with an error of $(1 - r_t)^2/2$)

$$r_t(\theta) \approx 1 + \log r_t(\theta). \tag{2}$$

By applying equations (1) and (2) back to $J(\theta)$, we have:

$$J(\theta) \approx \mathbb{E}_t \left[ \hat{A}_t (1 + g_t^\top \Delta\theta) \right].$$

Since $\mathbb{E}_t[\hat{A}_t]$ does not depend on $\Delta\theta$, maximizing $J(\theta)$ reduces to maximizing the following which is linear in $\theta$:

$$\mathbb{E}_t[\hat{A}_t g_t^\top (\theta - \theta^\star)].$$

**Surrogate loss.** Next, we turn the above reward maximization on $J(\theta)$ into a weighted binary classification problem. For every $(s_t, a_t, \hat{A}_t)$, for any $t \in \{1, 2, \ldots, N\}$, define:

---

Algorithm 1: Policy gradient estimation (POLICYGRADEX)

**Input:** $m$ random subsets from $\mathcal{T}$, $S_1, S_2, \ldots, S_m$
**Require:** Initial meta-policy parameter $\theta^\star \in \mathbb{R}^p$; Number of episodes $N$; Projection dimension $d$

```
π_θ* ← Setup the meta-policy for all tasks in T
S_1, S_2, ..., S_m ← m random subsets of T with size α
P ← A p by d isotropic Gaussian random projection matrix
for t = 1, ..., N do
    for T_i ∈ T do
        g_{i,t} ← P^⊤ ∇π_θ*(a_t|s_t)
        w_{i,t} ← |Â_t|
        y_{i,t} ← sign(Â_t)
    end for
end for
for j = 1, ..., m do
    Δθ̂_d ← argmin_{Δθ∈R^d} L̂_{S_j}(θ* + PΔθ)
    θ̂^(j) ← θ* + PΔθ̂_d
    f̂(S_j) ← −L̂_{S_j}(θ̂^(j))
end for
Return f̂(S_1), ..., f̂(S_m)
```

---

- Target binary label $y_t = \text{sign}(\hat{A}_t) \in \{-1, +1\}$, indicating if the action was better or worse than zero.
- Classifier score $z_t = g_t^\top \Delta\theta$, linear score for the update.
- Sample weight $w_t = |\hat{A}_t|$, magnitude of the reward.

This yields a per-sample surrogate loss as follows:

$$\ell(g_t, y_t, w_t; \Delta\theta) = w_t \cdot \log(1 + (-y_t(g_t^\top \Delta\theta))).$$

For a given task subset $S$, let $\mathcal{D}_S$ denote the set of trajectory samples $(g_{i,t}, y_{i,t}, w_{i,t})$ from all the tasks $T_i \in S$, across all the samples (with an extra index $i$). The average loss for the subset $S$ is defined by the average loss over all the samples from all the tasks within the subset $S$:

$$\hat{L}_S(\theta) = \frac{1}{|\mathcal{D}_S|} \sum_{(g,y,w) \in \mathcal{D}_S} \ell(g, y, w; \theta). \tag{3}$$

Minimizing this loss over $\Delta\theta = \theta - \theta^\star$ yields the estimated adaptation for the subset $S$ from the meta-policy $\theta^\star$.

**Random projection.** Since the gradient vectors $g_{i,t} \in \mathbb{R}^p$ may be high-dimensional (with $p$ in the order of millions), we apply random projections to reduce the dimension down to a few hundred (in practice, $d = 400$ suffices). This projection provably preserves the pairwise similarity between all pairs of gradient vectors via the Johnson–Lindenstrauss Lemma (Johnson and Lindenstrauss 1984). We sample a random projection matrix $P \in \mathbb{R}^{p \times d}$, where $d \ll p$ and every entry of $P$ is independently drawn from a Gaussian $\mathcal{N}(0, d^{-1})$. We project the high-dimensional gradients to a $d$-dimensional space as

$$\tilde{g}_{i,t} = P^\top g_{i,t}.$$

| Distance | MT10 | CartPole | Highway | Lunarlander |
|---|---|---|---|---|
| 0.1% | $0.01_{\pm 0.01}\%$ | $0.12_{\pm 0.14}\%$ | $0.02_{\pm 0.02}\%$ | $0.06_{\pm 0.01}\%$ |
| 0.5% | $0.43_{\pm 0.73}\%$ | $0.73_{\pm 0.10}\%$ | $0.11_{\pm 0.09}\%$ | $0.03_{\pm 0.02}\%$ |
| 1.0% | $0.32_{\pm 0.56}\%$ | $0.98_{\pm 0.65}\%$ | $2.04_{\pm 0.58}\%$ | $0.48_{\pm 0.01}\%$ |

Table 1: We empirically find that the approximation error of $\epsilon$ is negligible for several interactive RL environments. We report the mean and standard deviation of the approximation error across 10 random subsets. We defer a detailed description of the setup to the experiments section.

Then, we solve a logistic regression problem in the $d$-dimensional space with $\theta_d \in \mathbb{R}^d$:

$$\widehat{\theta}_d \leftarrow \arg\min_{\theta_d \in \mathbb{R}^d} \frac{1}{|\mathcal{D}_{S_j}|} \sum_{(g,y,w)\in\mathcal{D}_{S_j}} \ell(\tilde{g}, y, w; \theta_d), \tag{4}$$

for any $j = 1, 2, \ldots, m$. Finally, we turn the solution from dimension $d$ back to the original dimension $p$ as

$$\widehat{\theta}^{(j)} = \theta^\star + P\widehat{\theta}_d.$$

The complete description of this surrogate modeling procedure is provided in Algorithm 1.

## Evaluation of First-Order Policy Approximation

A critical assumption is the accuracy of the approximation in equation (1). Here we observe that the approximation error remains negligible for $\theta$ around $\theta^\star$. We evaluate this approximation by measuring the relative residual sum of squares (RSS) error across multiple RL environments. Specifically, we compute:

$$\frac{\left(\log \pi_\theta(a_t|s_t) - \log \pi_{\theta^\star}(a_t|s_t) - g_t^\top \Delta\theta\right)^2}{(\log \pi_\theta(a_t|s_t))^2},$$

where $\Delta\theta = \theta - \theta^\star$ and $g_t = \nabla \log \pi_\theta(a_t|s_t)\big|_{\theta=\theta^\star}$.

We begin by verifying that adapted policies remain close to the meta-initialization $\theta^\star$. Using ten randomly sampled subsets, we evaluate the relative distance $\frac{\|\theta-\theta^\star\|_F}{\|\theta^\star\|_F}$, confirming that adaptation stays within a local neighborhood where the approximation is valid.

We report the RSS error measured on MT10 from the MetaWorld benchmark, CartPole, LunarLander, and Highway in Table 1. In these environments, the policy, a four-layer MLP, is trained to perform various control tasks by selecting an action based on its output given the states. We sample 2048 steps per task to obtain the gradients. The results are averaged over 10 randomly sampled subsets of size 5. For a proper initialization $\theta^\star$, the approximation error is less than $2\%$ when the updated policy $\theta$ remains close to the initial policy. This verifies that the first-order expansion provides a sufficiently accurate local model for surrogate-based policy estimation. The approximation error increases from less than $2\%$ to up to $10\%$ when the policy's parameter distance from initialization grows to around $5\%$, defining the boundaries of our method's applicability.

Algorithm 2: Clustering related objectives into subgroups using the task affinity score matrix

**Input:** $n$ tasks $\mathcal{T}$; number of desired clusters $k$
**Require:** Number of subsets $m$ with size $\alpha$; Regularization parameter $\lambda$; Number of episodes $N$; Projected dimension $d$
**Output:** A disjoint partition of $\{1, 2, \ldots, n\}$ into $k$ groups

$\theta^\star \leftarrow$ Train a meta-initialization policy with $\mathcal{T}$
$S_1, \ldots, S_m \leftarrow$ Sample $m$ subsets of size $\alpha$ from $\mathcal{T}$
$\hat{f}(S_1), \hat{f}(S_2), \ldots, \hat{f}(S_m) \leftarrow$ Apply PolicyGradEx with $(S_1, S_2, \ldots, S_m; \theta^\star; N, d)$
$U \leftarrow$ An $n \times n$ affinity score matrix via equation (5)
$X \leftarrow$ Solve convex relaxation program (7) (in the Appendix) with $U$ and regularization parameter $\lambda$
$G_1, G_2, \ldots, G_k \leftarrow$ Round $X$ into $k$ subgroups

## Task Affinity Grouping

We next quantify the pairwise relationships between tasks by constructing a task affinity matrix $U \in \mathbb{R}^{n\times n}$, where each entry $U_{i,j}$ captures the collaborative behavior between tasks $T_i$ and $T_j$. Larger values indicate that jointly fine-tuning on the two tasks produces higher surrogate performance.

A naive approach is to repeatedly train a model for all the $\binom{n}{2}$ subsets, which is impractical. Instead, we sample $m$ random subsets $\{S_1, \ldots, S_m\}$ and use the surrogate model to compute their estimated performance $\hat{f}(S_i) = -\hat{L}_{S_i}(\hat{\theta}_i)$, for all $i = 1, 2, \ldots, m$ based on Algorithm 1. We then compute an $n \times n$ task affinity score matrix as follows:

$$U_{i,j} = \frac{1}{n_{i,j}} \sum_{1\le l\le m: \{T_i, T_j\}\in S_l} \hat{f}(S_l), \tag{5}$$

for all $i$ and $j$ in $\{1, 2, \ldots, n\}$, where $n_{i,j}$ is the number of sampled subsets containing both $T_i$ and $T_j$. As a remark, provided that $m = O(n^2)$, then $n_{i,j}$ must be nonzero with high probability, for all $i$ and $j$ in $\{1, 2, \ldots, n\}$.

We then cluster tasks by solving a convex optimization problem that maximizes intra-cluster affinity, following a trace-regularized relaxation (Awasthi et al. 2015). This step is very fast since it runs on an $n$ by $n$ matrix with $n$ being at most several hundred, which can typically be solved in just a few seconds. Moreover, this procedure can be repeated with different values of $\lambda$ to determine the optimal number of clusters $k$ for downstream performance. For brevity, we defer a complete description of the convex relaxation program to the Appendix. The whole procedure is described in Algorithm 2.

# Experiments

We conduct a comprehensive set of experiments to validate PolicyGradEx and its application across various multi-objective reinforcement learning benchmarks. Our evaluation is designed to answer two key questions: (1) Does our surrogate model accurately and efficiently approximate the outcome of full policy training on various task subsets? (2) Do task groups identified by our algorithm lead to superior performance in downstream multitask and meta-RL evaluations compared to state-of-the-art and heuristic baselines?

| RL Environment | Meta-World | CartPole | Highway | LunarLander |
|---|---|---|---|---|
| Multi-task training (Yu et al. 2020b) | $71.3_{\pm 1.2}\%$ | $145.9_{\pm 9.0}$ | $140.0_{\pm 4.6}$ | $53.8_{\pm 14.6}$ |
| Soft modularization (Yang et al. 2020) | $82.0_{\pm 1.1}\%$ | $139.3_{\pm 9.5}$ | $141.3_{\pm 3.5}$ | $66.1_{\pm 13.0}$ |
| PaCo (Sun et al. 2022) | $73.1_{\pm 1.1}\%$ | $144.5_{\pm 5.2}$ | $136.6_{\pm 6.4}$ | $62.6_{\pm 11.4}$ |
| CARE (Sodhani, Zhang, and Pineau 2021) | $84.0_{\pm 1.8}\%$ | / | / | / |
| Randomly assign each task into $k$ groups | $58.2_{\pm 6.2}\%$ | $144.1_{\pm 10.1}$ | $143.4_{\pm 5.4}$ | $73.1_{\pm 11.7}$ |
| Gradient-similarity-based grouping (Yu et al. 2020a) | $69.6_{\pm 1.9}\%$ | $142.0_{\pm 8.2}$ | $135.6_{\pm 7.6}$ | $80.8_{\pm 6.9}$ |
| **Algorithm 2 (This paper)** | $\mathbf{94.0}_{\pm \mathbf{2.8}}\%$ | $\mathbf{159.2}_{\pm 3.8}$ | $\mathbf{153.5}_{\pm 7.8}$ | $\mathbf{82.8}_{\pm 6.9}$ |

Table 2: Comparison of our approach with several baseline methods. We report the average success rate on the Meta-World benchmark and the rewards (after adaptation) for three robotic control environments. We note that CARE does not apply to the control tasks in the meta-RL setting. We report the mean and standard deviation from five runs.

Our finding shows that our surrogate modeling approach achieves high accuracy in estimating multitask RL performance, with over $0.73$ normalized mutual information relative to the ground-truth, while reducing the number of floating-point operations (FLOPs) by up to $26\times$. In downstream evaluations, PolicyGradEx outperforms existing multitask optimizers by $19\%$ in multitask RL benchmarks and achieves a $13\%$ improvement in the meta-RL setting. Furthermore, we compare our task affinity grouping algorithm to random grouping and gradient-similarity based grouping, showing that PolicyGradEx achieves a $19\%$ improvement over both. Ablation analysis validates the use of random projections and the choice of $k$ in the algorithm.

## Experimental Setup

*Environments.* We evaluate our method on two types of benchmarks. First, we use MT10 from the Meta-World benchmark (Yu et al. 2020b), which consists of 10 diverse robotic manipulation tasks. While these tasks share common state and action spaces, their reward functions and dynamics vary (e.g., a goal position in one task may correspond to an object's position in another). Second, we use three classic control environments from Gymnasium (Towers et al. 2025) and MO-Gymnasium (Felten et al. 2023): CartPole, Highway, and LunarLander. For each, we generate a distribution of 10 source tasks by altering key physical parameters (e.g., pole length for CartPole, traffic density for Highway, and gravity for LunarLander).

*Baselines.* We compare our method against two categories of baselines. (1) Multi-objective RL baselines: These represent standard and state-of-the-art multitask approaches for multitasks training: a single policy trained on all tasks; Soft Modularization (Yang et al. 2020), which employs a routing-based selection mechanism; PaCo (Sun et al. 2022), an ensemble method that composes policies from a shared parameter subspace; and CARE (Sodhani, Zhang, and Pineau 2021), which performs inference using contextual information to adapt policies. (2) Grouping baselines: These methods use the same number of groups as our approach but employ different heuristic grouping procedures. This allows us to isolate the benefit of our affinity modeling. We include random grouping (where tasks are randomly assigned to $k$ groups) and gradient-similarity-based grouping (where tasks are clustered based on the cosine similarity of their respective policy gradients).

| # MLP layers | Meta-World | LunarLander | Speedup |
|---|---|---|---|
| 2 | 0.76 | 0.73 | $21\times$ |
| 4 | 0.76 | 0.73 | $24\times$ |
| 8 | 0.76 | 0.73 | $26\times$ |

Table 3: Normalized Mutual Information (NMI) between estimated and actual clusters, measured on two environments trained with a multi-layer perceptron (MLP). In the last column, the speedup is measured as the ratio of the FLOP count between full training and PolicyGradEx.

*Implementations.* For the MT10 benchmark, we group the $n = 10$ tasks into $k = 3$ subsets and train a separate policy for each subgroup using soft modularization. We measure performance using the average success rate per task.

For the control environments, we form the meta-training set by randomly selecting one from each task group provided by PolicyGradEx. We then use MAML (Finn, Abbeel, and Levine 2017) to train a meta-policy on the assigned tasks.

The final performance is measured by the average reward after 200 adaptation steps on 50 unseen target tasks. All experiments are conducted using PyTorch on an Ubuntu server with an Intel Xeon E5-2623 CPU and an NVIDIA Quadro RTX 6000 GPU.

## Results for Loss Function Estimation

We first evaluate both the approximation accuracy and computational cost of PolicyGradEx. We establish ground-truth subset clusters based on task-affinity scores derived from rewards obtained by training policies on different task subsets. We then use Normalized Mutual Information (NMI) to quantify the accuracy of the subsets estimated by PolicyGradEx against this ground truth.

As shown in Table 3, on both MetaWorld and LunarLander, PolicyGradEx identifies subset clusters that achieve over $0.73$ similarity to those obtained from full-policy training, while reducing FLOPs by a factor of $26\times$. By contrast, the NMI under random clustering is approximately $0.2$. Moreover, the relative error between the estimated task affinity matrix and the ground-truth matrix is at most $0.2$.

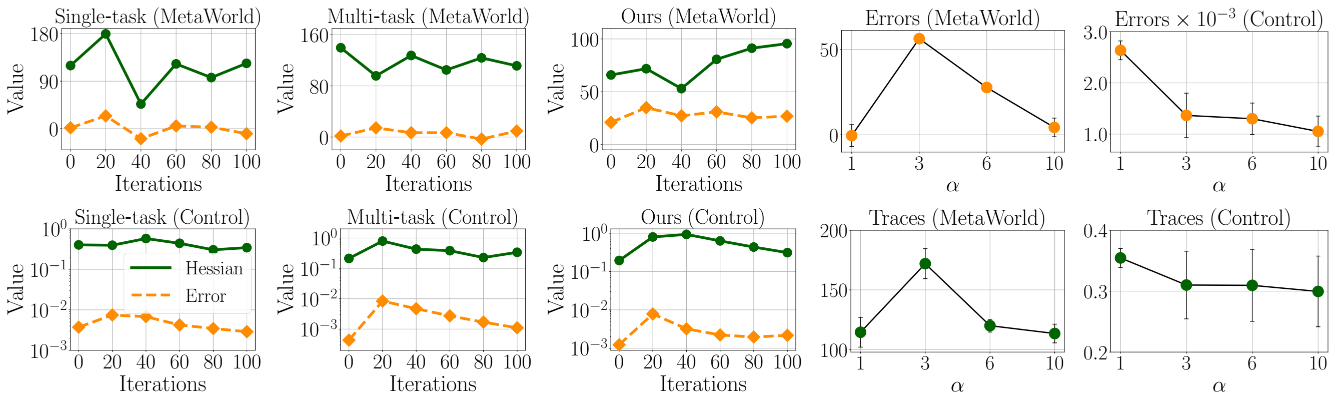


(a) Illustration of Hessian traces and generalization errors during training

(b) Varying subset size $\alpha$ in Algorithm 2

Figure 2: Illustrating the Hessian trace measurements and empirical generalization errors with respect to the policy network. Figure 2a: Showing that the Hessian trace is comparable in scale to the observed generalization errors, tested on Meta-World and a control task. Figure 2b: Showing that in a Meta-World environment, the generalization error reaches the highest when the subset size $\alpha = 3$, suggesting negative transfer among a small set of tasks. In the meta-RL control task, the generalization performance monotonically improves with the addition of more tasks in each group.

## Results for Multi-Objective RL

We report the evaluation results in Table 2. On the Meta-World benchmark, our approach improves the average success rate by 21% compared to multitask optimizers. It also achieves a 62% improvement over training with random groups and a 35% improvement over training with gradient-similarity-based groups. This observation suggests that random grouping and gradient similarity-based grouping in Meta-World tasks can result in negative transfer. In contrast, PolicyGradEx selects task groups with positive transferability, thereby improving the final success rate.

Next, we evaluate the ability of PolicyGradEx to select a representative task subset for meta-learning in the three control RL environments. Using MAML as the meta-learner, our approach improves the final adapted reward by 7% compared to multitask optimizers.

Finally, we compare with two baseline meta-training settings: training on all available source tasks and training on randomly grouped tasks. Our approach achieves a 13% improvement over both baselines and a 9% improvement compared to gradient-similarity-based grouping.

## Generalization Error Measurements via Hessians

Next, we analyze the generalization behavior. by evaluating a sharpness measure based on the Hessian trace of the test loss for the policy network and the policy gradient loss. A smaller trace value suggests a flatter loss landscape. We implement Hutchinson's estimator and leverage a faster version of this estimator (Meyer et al. 2021) to estimate the Hessian trace of RL models. We compute the policy gradient loss from the target tasks and measure the difference between the training loss and test loss as the generalization error. Results for the largest eigenvalue of the Hessian are similar.

First, we evaluate the Hessian trace and the policy generalization error between single-task training, multitask training on ten tasks, and training on a task group selected by PolicyGradEx. We begin training from the initial policy for 100 iterations, each with 2048 steps.

Figure 2a shows the dynamics of the trace of the Hessian and the generalization errors during training. We find that they both follow a qualitatively similar trend. Next, as shown in Figure 2b, we find that single-task training yields a very low generalization error in the Meta-World benchmark, indicating that a single policy can solve an individual task. However, when the task count reaches three, the generalization error sharply increases, suggesting that a single policy struggles to perform well across multiple tasks. As the number of tasks increases, the generalization error and the Hessian trace decrease, indicating the positive transfer among tasks. In the control environments, we observe that both the generalization error and the Hessian trace decrease as $\alpha$ increases.

## Ablation Analysis

Given that each task is solvable in the Meta-World benchmark, our goal is to determine the minimum number of groups $k$ required to maximize the success rate. We vary $k$ from 1 to 4, observing average success rates of 94.0% when $k = 3$, compared to 89.5% when $k = 2$ and 95.1% when $k = 4$. Thus, we report the final results with three groups. For the control environments, we vary $k$ from 1 to 5. We find that the success rate stabilizes after $k$ reaches 3, so we report results using three groups.

For the random projection dimension $d$, we vary it from 200 to 1000 and observe that values beyond 400 yield minimal gains, so we fix $d = 400$.

# Generalization Error Analysis via Hessians

In this section, we present a generalization error analysis for the RL setting described above. We present a Hessian-based technique to quantify generalization errors, applicable to any kind of policy network used in RL training. We first describe the tools that we will use. Given a data distribution of $\mathcal{D}$, let

$$L(f_W) = \mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} [\ell(f_W(x), y)]$$

denote the expected loss of an input sample $x, y$. Let $\hat{L}(f_W)$ denote the empirical loss given $n$ independent samples $(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$ sampled from $\mathcal{D}$:

$$\hat{L}(f_W) = \frac{1}{n}\sum_{i=1}^{n} \ell(f_W(x_i), y_i).$$

Let $\mathcal{W} \subseteq \mathbb{R}^d$ denote the weight space of $W$. Assume that the weighted Hessian along $x$ is bounded by a fixed constant $\mathcal{H}$ that does not grow with $d$ or $n$:

$$\mathcal{H} := \sup_{W\in\mathcal{W}} \left( W^\top \mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} \left[ [\nabla^2 \ell(f_W(x), y)]^+ \right] W \right),$$

where $[\nabla^2 \ell(f_W(x), y)]^+$ means that we truncate the negative eigen-directions of the Hessian to zero. Then, we have the following generalization error bound, which applies uniformly to the hypothesis space $\{f_W : W \in \mathcal{W}\}$.

**Theorem 1.** *Assume that the loss function $\ell$ is bounded between $0$ and $C$ for a fixed constant $C > 0$ on the data distribution $\mathcal{D}$. Suppose $\ell(f_W(\cdot), \cdot)$ is twice-differentiable in $W$ and the Hessian matrix $\nabla^2 \ell(f_W(\cdot), \cdot)$ is Lipschitz-continuous within the weight space $\mathcal{W}$. Suppose there exists a fixed bound $\mathcal{H}$ on the Hessian trace over any $W \in \mathcal{W}$. Then, for any $W$ in $\mathcal{W}$, any $\epsilon > 0$ small enough, and any small $\delta > 0$, with probability at least $1-\delta$ over the randomness of $n$ samples, we have:*

$$L(f_W) \leq (1+\epsilon)\hat{L}(f_W) + (1+\epsilon)\sqrt{\frac{C \cdot \mathcal{H}}{n}} + \epsilon, \quad (6)$$

*where $\epsilon = O(n^{-3/4} \log(\delta^{-1}))$ denotes the error term.*

The proof of Theorem 1 is based on a PAC-Bayes analysis (Ju, Li, and Zhang 2022; Zhang, Li, and Ju 2024), and the details are in the Appendix. In particular, a new aspect of this result is that we use an anisotropic perturbation in the prior and posterior distributions. We build on a linear PAC-Bayes bound (See Theorem 2 in the Appendix), and optimize the noise distribution to obtain the $\mathcal{H}$ measure.

In Figure 2, we visualize the Hessian trace on the Meta-World and the control tasks. Our results show that the Hessian trace is comparable in scale to the observed generalization errors. In the Meta-World benchmark, the generalization error increases as the number of tasks increases from one to three, suggesting negative transfer. As more tasks are added, the error decreases, indicating positive transfer.

## Related Work

**Multi-objective optimization for reinforcement learning.** Early works on multitask learning use hard parameter sharing, where a single neural network backbone is shared among all tasks, with only the final layers or heads being task-specific (Yu et al. 2020b). A canonical example is the multi-head Soft Actor-Critic (SAC) architecture, which serves as a common baseline (Yu et al. 2020b).

The multitask objective provides implicit regularization across multiple tasks trained together in a shared network (Wu, Zhang, and Ré 2020), which can be formalized in a high-dimensional regression setting (Yang et al. 2025). Training the model to learn representations useful across a diverse set of tasks helps prevent overfitting. A deeper analysis of the regularization effect behind multitask RL is an interesting question for future work.

**Data attribution and model interpretability.** Another line of work involves data modeling, which aims to trace a model's predictions back to its training data. One method is to retrain the model on a different subset of the dataset and estimate the Shapley value of each sample (Ghorbani and Zou 2019). A commonly used technique is influence functions (Koh and Liang 2017), which use Hessians to approximate the effect of removing a sample. Datamodels (Ilyas et al. 2022) finds that a linear regression method can accurately approximate the outcome of deep nets trained with a subset of samples on computer vision tasks. Li, Nguyen, and Zhang (2023) introduce a linear surrogate model for multitask learning and analyze the sample complexity of linear surrogate models. TRAK (Park et al. 2023) further demonstrates that the approximate solution of linear regression computed with projected gradients delivers comparable results to the original linear model for ImageNet, CLIP, and BERT. Our work contributes to this literature by applying attribution methods to design RL algorithms.

## Conclusion

This paper introduces an efficient algorithm for multi-objective reinforcement learning. The overall approach works by first training a meta-policy via multitask learning. Then, use a gradient-estimation algorithm to adapt this meta-policy to multiple subsets of training objectives. The key observation is that a policy can effectively adapt across different task subsets through a first-order approximation with a proper initialization. Leveraging this, we propose a surrogate model to approximate the actual training performance of the policy, which enabled us to derive task affinities and subsequently cluster similar tasks into groups. In extensive reinforcement learning tasks, our method consistently improves performance. Lastly, we measure sharpness via Hessians to analyze generalization in multi-objective RL. We hope this work inspires further studies on designing principled methods for multi-objective reinforcement learning and quantifying generalization in policy learning algorithms. Further applying our approach to broader RL settings is another promising direction for future work.

## Acknowledgments

Thanks to the anonymous referees for their constructive feedback. This work is partially funded by NSF award IIS-2412008. Any views or opinions expressed herein are solely those of the authors listed, and may differ from those expressed by NSF.

## References


Awasthi, P.; Bandeira, A. S.; Charikar, M.; Krishnaswamy, R.; Villar, S.; and Ward, R. 2015. Relax, no need to round: Integrality of clustering formulations. In *Innovations in Theoretical Computer Science*, 191–200.

Caruana, R. 1997. Multitask learning. *Machine Learning*, 28(1): 41–75.

Cobbe, K.; Klimov, O.; Hesse, C.; Kim, T.; and Schulman, J. 2019. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, 1282–1289.

Felten, F.; Alegre, L. N.; Nowé, A.; Bazzan, A. L. C.; Talbi, E. G.; Danoy, G.; and Silva, B. C. da. 2023. A Toolkit for Reliable Benchmarking and Research in Multi-Objective Reinforcement Learning. In *Advances in Neural Information Processing Systems*.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 1126–1135.

Ghorbani, A.; and Zou, J. 2019. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, 2242–2251.

Ilyas, A.; Park, S. M.; Engstrom, L.; Leclerc, G.; and Madry, A. 2022. Datamodels: Predicting predictions from training data. *International Conference on Machine Learning*.

Johnson, W. B.; and Lindenstrauss, J. 1984. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206): 1.

Joshi, V.; Xu, Z.; Liu, B.; Stone, P.; and Zhang, A. 2025. Benchmarking massively parallelized multi-task reinforcement learning for robotics tasks. *Reinforcement Learning Journal*.

Ju, H.; Li, D.; Sharma, A.; and Zhang, H. R. 2023. Generalization in graph neural networks: Improved pac-bayesian bounds on graph diffusion. In *International conference on artificial intelligence and statistics*, 6314–6341.

Ju, H.; Li, D.; and Zhang, H. R. 2022. Robust fine-tuning of deep neural networks with hessian-based generalization guarantees. In *International Conference on Machine Learning*, 10431–10461.

Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, 1885–1894.

Li, D.; Ju, H.; Sharma, A.; and Zhang, H. R. 2023. Boosting multitask learning on graphs through higher-order task affinities. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1213–1222.

Li, D.; Nguyen, H. L.; and Zhang, H. R. 2023. Identification of negative transfers in multitask learning using surrogate models. *Transactions on Machine Learning Research*.

Li, D.; Sharma, A.; and Zhang, H. R. 2024. Scalable Multitask Learning Using Gradient-based Estimation of Task Affinity. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1542–1553.

Li, D.; Zhang, Z.; Wang, L.; and Zhang, H. 2024. Scalable Fine-tuning from Multiple Data Sources: A First-Order Approximation Approach. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 5608–5623.

Meyer, R. A.; Musco, C.; Musco, C.; and Woodruff, D. P. 2021. Hutch++: Optimal stochastic trace estimation. In *Symposium on Simplicity in Algorithms (SOSA)*, 142–155. SIAM.

Park, S. M.; Tsvetkov, Y.; Költer, J. Z.; and Salakhutdinov, R. 2023. Trak: Attributing model behavior at scale. In *International Conference on Machine Learning*, 27457–27481.

Rakelly, K.; Zhou, A.; Finn, C.; Levine, S.; and Quillen, D. 2019. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International Conference on Machine Learning*, 5331–5340.

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

Sodhani, S.; Zhang, A.; and Pineau, J. 2021. Multi-task reinforcement learning with context-based representations. In *International Conference on Machine Learning*, 9767–9779.

Sun, S.; Shi, H.; Zhao, Z.; Wang, X.; and Zhou, J. 2022. PaCo: Parameter-compositional multi-task reinforcement learning. In *International Conference on Learning Representations*.

Thrun, S.; and Pratt, L. 1998. Learning to learn. Springer.

Towers, M.; Kwiatkowski, A.; Terry, J.; Balis, J. U.; De Cola, G.; Deleu, T.; Goulão, M.; Kallinteris, A.; Krimmel, M.; KG, A.; Perez-Vicente, R.; Pierré, A.; Schulhoff, S.; Tai, J. J.; Tan, H.; and Younis, O. G. 2025. Gymnasium: A standard interface for reinforcement learning environments. *Advances in Neural Information Processing Systems Datasets and Benchmarks*.

Wu, S.; Zhang, H. R.; and Ré, C. 2020. Understanding and Improving Information Transfer in Multi-Task Learning. In *International Conference on Learning Representations*.

Yang, F.; Zhang, H. R.; Wu, S.; Re, C.; and Su, W. J. 2025. Precise High-Dimensional Asymptotics for Quantifying Heterogeneous Transfers. *Journal of Machine Learning Research*, 26(113): 1–88.

Yang, R.; Xu, H.; Wu, Y.; and Wang, X. 2020. Multi-task reinforcement learning with soft modularization. In *Advances in Neural Information Processing Systems*, volume 33, 4767–4777.

Yin, M.; Tucker, G.; Zhou, M.; Levine, S.; and Finn, C. 2020. Meta-learning without memorization. *International Conference on Learning Representations*.

Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020a. Gradient surgery for multi-task learning. In *Advances in Neural Information Processing Systems*, volume 33, 5822–5833.

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.

Zhang, H. R.; Li, D.; and Ju, H. 2024. Noise Stability Optimization for Finding Flat Minima: A Hessian-based Regularization Approach. *Transactions on Machine Learning Research*.

Zhang, Z.; Zhang, Z.; Li, D.; Wang, L.; Dy, J.; and Zhang, H. R. 2025. Linear-Time Demonstration Selection for In-Context Learning via Gradient Estimation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 16470–16488.