Theoretical Investigations and Practical Enhancements on Tail Task Risk Minimization in Meta Learning

Yiqin Lv[1], Qi Wang[1], Dong Liang[1], Zheng Xie[1]

[1] National University of Defense Technology, Changsha, China

Overview

1. Introduction
2. Preliminaries
3. Theoretical Investigations
4. Empirical Findings
5. Conclusion

Introduction

Introduction

◻ Meta Learning

Leverages previous experience as priors to quickly adapt to unseen tasks ☺.

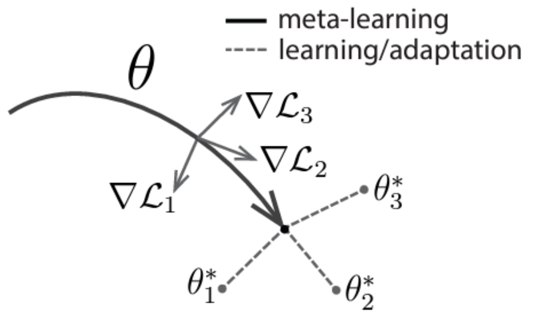

Figure 1. Diagram of our model-agnostic meta-learning algorithm (MAML), which optimizes for a representation θ that can quickly adapt to new tasks.

Source: MAML[1]

◻ Robustness Concern

Worst fast adaptation can be catastrophic in risk-sensitive scenarios, e.g., autonomous driving ☹.

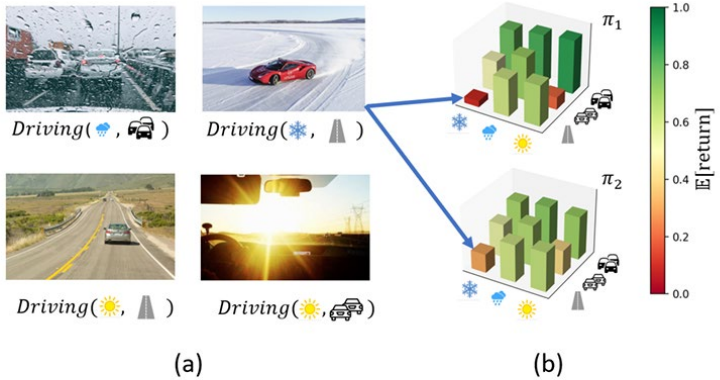

Source: RoML[2]

It is desirable to watch adaptation differences across tasks when deploying meta learning models.

Introduction

Previous Works

DR-MAML [3] increases task distributional robustness via employing the tail risk minimization principle for meta learning.

Two-stage optimization strategy

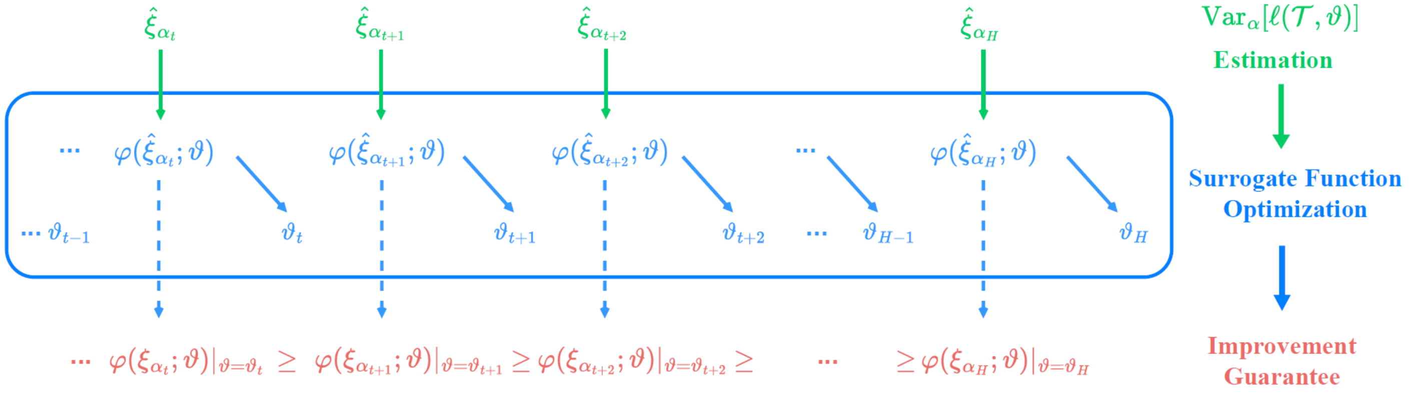

(i) Estimate the risk quantile VaR_α with the crude Monte Carlo method in the task space.

(ii) Update the meta learning model parameters from the screened subset of tasks.

Introduction

☐ Existing Limitations

- **Theoretically**

  (i) There constitutes no notion of solutions.

  (ii) Lacks an algorithmic understanding of the two-stage optimization strategy.

  (iii) The analysis on generalization capability is ignored in the tail risk of tasks.

- **Empirically**

  The use of the crude Monte Carlo might be less efficient in quantile estimates and suffers from a higher approximation error of the VaR_α, degrading the adaptation robustness.

*We propose translating the **two-stage optimization strategy** for distributionally robust meta learning into **a max-min optimization problem** ☺.*

Preliminaries

Preliminaries

□ Notations

Task distribution $p(\tau)$ defined in task space Ω_τ; the set of all tasks $\mathcal{T}$; Meta dataset $\mathfrak{D}_\tau$, *e.g.*, $\mathfrak{D}_\tau = \{(x_i, y_i)\}_{i=1}^{m} = \mathfrak{D}_\tau^S \cup \mathfrak{D}_\tau^Q$ in few-shot regression problems; Parameter space Θ

Meta risk function $\ell: \mathfrak{D}_\tau \times \Theta \mapsto \mathbb{R}^+$ evaluating fast adaptation performance;

Cumulative distribution of the meta risk function

$$F_\ell(l;\theta) := \mathbb{P}\left(\{\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S;\theta\right) \leq l; \tau \in \mathcal{T}, l \in \mathbb{R}^+\}\right)$$

Preliminaries

◻ Notations

Value-at-risk (VaR_α)

$$\mathrm{VaR}_\alpha[\ell(\mathcal{T},\theta)] = \inf_{l\in\mathbb{R}^+}\{l|F_\ell(l;\theta) \geq \alpha, \tau \in \mathcal{T}\}$$

Conditional value-at-risk (CVaR_α)

$$\mathrm{CVaR}_\alpha = \mathbb{E}_{p(\tau)}[\ell|\ell \geq \mathrm{VaR}_\alpha]$$

Normalized cumulative distribution $F_\ell^\alpha(l;\theta)$;

Tail risk task subspace $\Omega_{\alpha,\tau}$;

Density function $p_\alpha(\tau;\theta)$

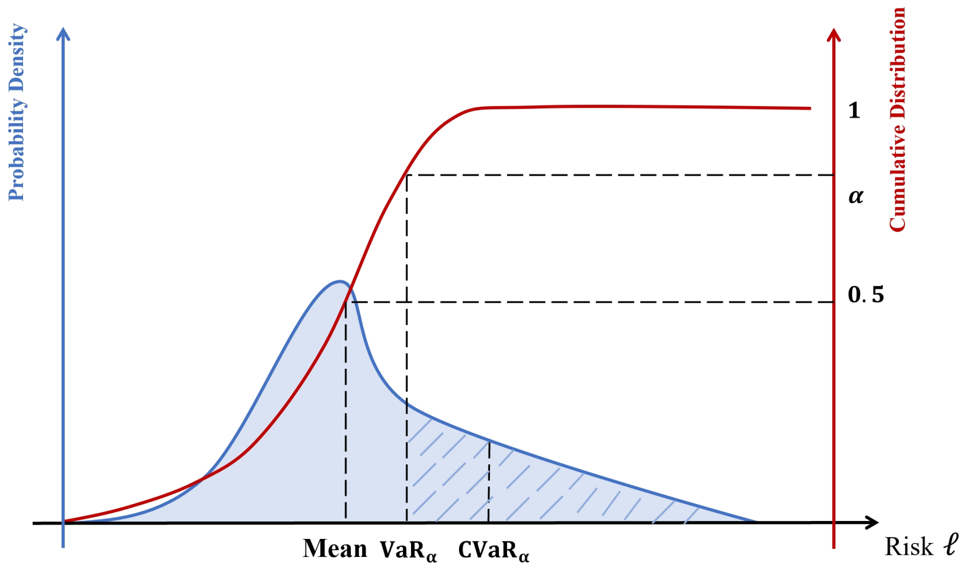

Risk Minimization Principles

- **Expected Risk Minimization.** It minimizes meta risk based on the sampling chance of tasks from the original task distribution:

$$\min_{\theta\in\Theta}\mathcal{E}(\theta) := \mathbb{E}_{p(\tau)}\left[\ell\left(\mathfrak{D}_{\tau}^{Q},\mathfrak{D}_{\tau}^{S};\theta\right)\right].$$

- **Worst-case Risk Minimization.** Noticing that the worst fast adaptation can be disastrous in some risk sensitive scenarios, [4] proposes to conduct the worst-case optimization in meta learning:

$$\min_{\theta\in\Theta}\max_{\tau\in\mathcal{T}}\mathcal{E}_{\mathrm{W}}(\theta) := \ell\left(\mathfrak{D}_{\tau}^{Q},\mathfrak{D}_{\tau}^{S};\theta\right).$$

Risk Minimization Principles

- **Expected Tail Risk Minimization (CVaR_α).** To balance the average performance and the worst-case performance, [3] minimizes the expected tail risk, or equivalently CVaR_α risk measure:

$$\min_{\theta\in\Theta,\xi\in\mathbb{R}} \mathcal{E}_\alpha(\theta,\xi) := \frac{1}{1-\alpha}\int_\alpha^1 v_\beta \, d\beta = \xi + \frac{1}{1-\alpha}\mathbb{E}_{p(\tau)}\left[\left[\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S; \theta\right) - \xi\right]^+\right],$$

$v_\beta := F_\ell^{-1}(\beta)$ denotes the quantile statistics

$\left[\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S; \theta\right) - \xi\right]^+ := \max\{\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S; \theta\right) - \xi, 0\}$ is the hinge risk.

Examples

Example 1 (DR-MAML).

Given $p(\tau)$ and vanilla MAML [1], the distributionally robust MAML within CVaR_α *can be written as a bi-level optimization problem:*

$$\min_{\substack{\theta\in\Theta\\ \xi\in\mathbb{R}}} \xi + \frac{1}{1-\alpha}\mathbb{E}_{p(\tau)}\left[\left[\ell\left(\mathfrak{D}_\tau^Q;\theta-\lambda\nabla_\theta\ell(\mathfrak{D}_\tau^S;\theta)\right)-\xi\right]^+\right],$$

where the gradient update w.r.t. the support set $\nabla_\theta\ell(\mathfrak{D}_\tau^S;\theta)$ indicates the inner loop with a learning rate λ. The outer loop executes the gradient updates and seeks the robust meta initialization in the parameter space.

Two-Stage Optimization Strategies

The pipelines of DR-MAML (**Example 1):**

Stage-I includes the fast adaptation *w.r.t.* individual task in Eq. (1) and the quantile estimate in Eq. (2)

Stage-II applies the sub-gradient updates to the model parameters in Eq. (3)/(4).

$$\theta_t^{\tau_i} = \theta_t^{meta} - \lambda_1 \nabla_\theta \ell\big(\mathfrak{D}_{\tau_i}^S; \theta\big),\ i = 1, \dots, \mathcal{B} \tag{1}$$

$$\hat{\xi} = \hat{F}_{MC-\mathcal{B}}^{-1}(\alpha), \tag{2}$$

$$\delta(\tau_i) = 1\big[\ell\big(\mathfrak{D}_{\tau_i}^Q; \theta_t^{\tau_i}\big) \geq \hat{\xi}\big],\ i = 1, \dots, \mathcal{B} \tag{3}$$

$$\theta_{t+1}^{meta} \leftarrow \theta_t^{meta} - \lambda_2 \left[\sum_{i=1}^{\mathcal{B}} \nabla_\theta \left[\delta(\tau_i) \cdot \ell\big(\mathfrak{D}_{\tau_i}^Q; \theta_t^{\tau_i}\big)\right]\right]. \tag{4}$$

These two stages repeat until reaching the convergence required iterations.

Theoretical Investigations

The Sketch of our Study

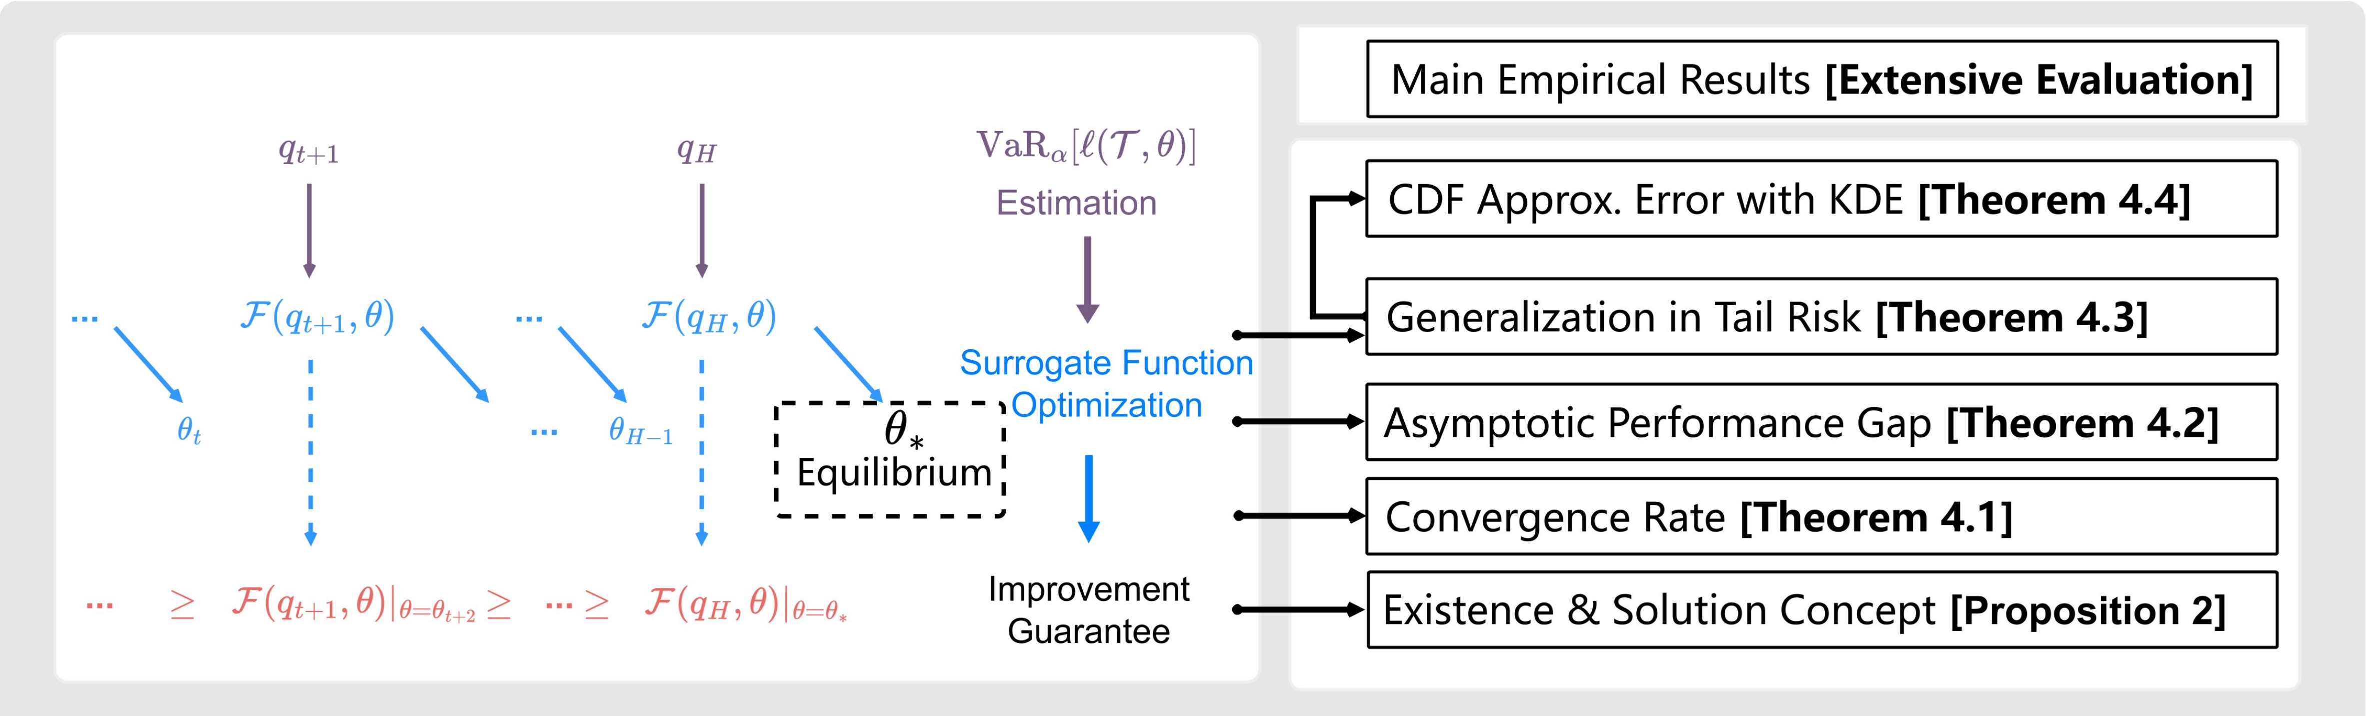

Stackelberg Game for Meta Learning Robustification

On the left side is the two-stage distributionally robust strategy.

The contributed theoretical understanding is right-down, with the right-up the empirical improvement.

Max-Min Optimization

In the two-stage optimization strategy, the default is the minimization of the risk measure *w.r.t.* the parameter space after the maximization of the risk measure *w.r.t.* the task subspace.

- **Max-Min Optimization.** With the pre-assigned decision-making orders, the studied problem can be characterized as:

$$\max_{q(\tau)\in\mathcal{Q}_\alpha}\min_{\theta\in\Theta}\mathcal{F}(q,\theta):=\mathbb{E}_{q(\tau)}\left[\ell\left(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta\right)\right],$$

where $\mathcal{Q}_\alpha:=\{q(\tau)|\mathcal{T}_q\subseteq\mathcal{T},\int_{\tau\in\mathcal{T}_q}p\,(\tau)d\tau=1-\alpha\}$ constitutes a collection of uncertainty sets over task subspace $\mathcal{T}_q$, and $q(\tau)$ is the normalized probability density over the task subspace.

Stackelberg Game

- **Stackelberg Game:** the example pipelines can be understood as approximately solving a stochastic two-player zero-sum Stackelberg game

$$\mathcal{SG} := \langle \mathcal{P}_L, \mathcal{P}_F; \{q \in \mathcal{Q}_\alpha\}, \{\theta \in \Theta\}; \mathcal{F}(q,\theta) \rangle.$$

- **Best Responses:**

$$\mathcal{SG}:\ q_t = \underbrace{\arg\max_{q\in\mathcal{Q}_\alpha} \mathbb{E}_q\left[\ell(\mathfrak{D}^Q_\tau, \mathfrak{D}^S_\tau; \theta_t)\right]}_{Leader\ Player},$$

$$\theta_{t+1} = \underbrace{\arg\min_{\theta\in\Theta} \mathbb{E}_{q_t}\left[\ell(\mathfrak{D}^Q_\tau, \mathfrak{D}^S_\tau; \theta)\right]}_{Follower\ Player},$$

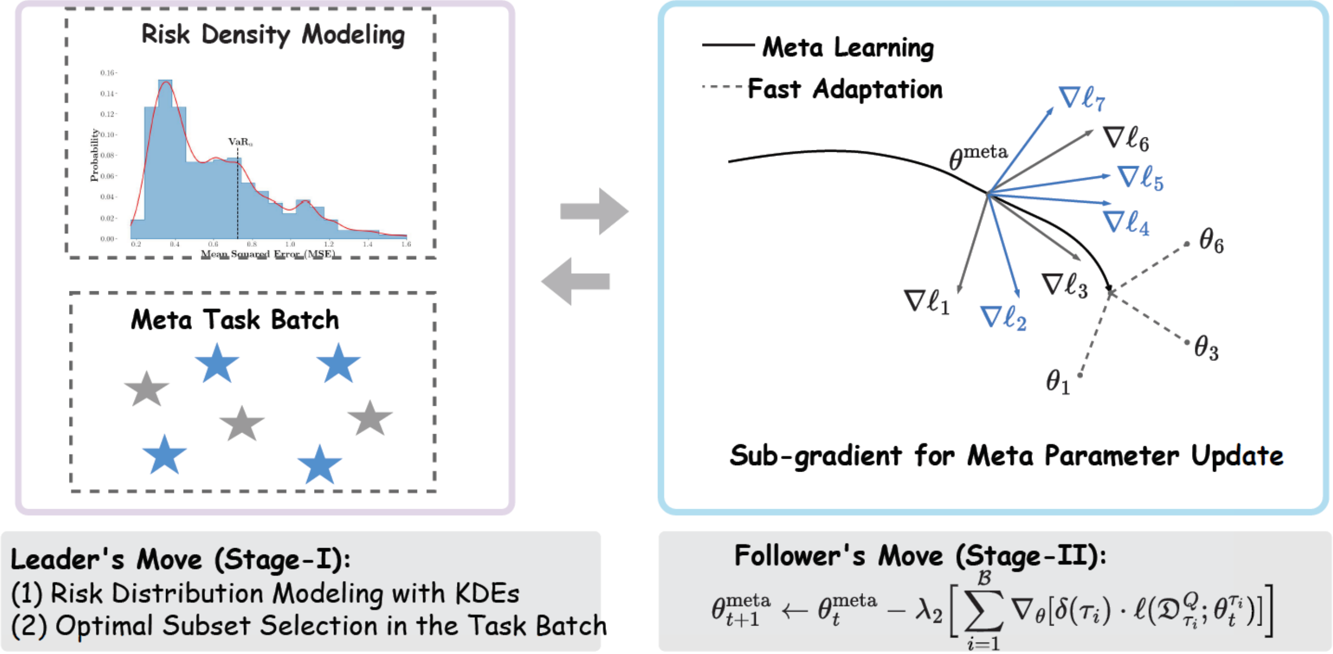

Solution Concept

☐ **Definition** (Local Stackelberg Equilibrium).

Let $(q_*, \theta_*) \in \mathcal{Q}_\alpha \times \Theta$ be the solution. With the leader $q_* \in \mathcal{Q}_\alpha$ and the follower $\theta_* \in \Theta$, (q_*, θ_*) is called a *local Stackelberg equilibrium* for the leader if the following inequalities hold,

$$\inf_{\theta \in \mathcal{S}_{\Theta'}(q_*)} \mathcal{F}(q_*, \theta) \geq \inf_{\theta \in \mathcal{S}_{\Theta'}(q)} \mathcal{F}(q, \theta),$$

where $\mathcal{S}_{\Theta'}(q) := \{\bar{\theta} \in \Theta' | \mathcal{F}(q, \bar{\theta}) \leq \mathcal{F}(q, \theta), \forall \theta \in \Theta'\}$.

☐ **Interpretation of the Obtained Equilibrium $(\boldsymbol{q_*}, \boldsymbol{\theta_*})$.**

Given the follower's decision θ_* and the induced task risk distribution $F_\ell(l; \theta_*)$, the leader cannot further raise a proposal of a task subset with a probability $1 - \alpha$ to degrade the tailed expected performance.

Convergence Rate

☐ Assumption 1

1. The meta risk function $\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S; \theta\right)$ is β_τ-Lipschitz continuous *w.r.t.* θ;
2. The cumulative distribution $F_\ell(l;\theta)$ is β_ℓ-Lipschitz continuous *w.r.t.* l, and the normalized density function $p_\alpha(\tau;\theta)$ is β_θ-Lipschitz continuous *w.r.t.* θ;
3. For arbitrary valid $\theta \in \Theta$ and corresponding $p_\alpha(\tau;\theta)$, $\ell\left(\mathfrak{D}_\tau^Q, \mathfrak{D}_\tau^S; \theta\right)$ is bounded:

$$\sup_{\tau \in \Omega_{\alpha,\tau}} \ell\left(\mathfrak{D}_{\tau_i}^Q, \mathfrak{D}_{\tau_i}^S; \theta\right) \leq \mathcal{L}_{\max}.$$

☐ Assumption 2

The implicit function $h(\cdot)$ is β_h-Lipschitz continuous w.r.t. $\theta \in \Theta$, and $\nabla_\theta \mathcal{F}(q, \theta)$ is β_q-Lipschitz continuous w.r.t. $q \in \mathcal{Q}_\alpha$.

Convergence Rate

□ **Theorem 1** (Convergence Rate for the Second Player).

Let the iteration sequence in optimization be: $\cdots \mapsto \{q_{t-1}, \theta_t\} \mapsto \{q_t, \theta_{t+1}\} \mapsto \cdots \mapsto \{q_*, \theta_*\}$*, with the converged equilibirum* (q_*, θ_*)*. Under the* ***Assumption*** *2 and suppose that* $\left\|I - \lambda \nabla^2_{\theta\theta} \mathcal{F}(q_*, \theta_*)\right\|_2 < 1 - \lambda \beta_q \beta_h$*, we can have* $\lim_{t\to\infty} \frac{\|\theta_{t+1} - \theta_*\|_2}{\|\theta_t - \theta_*\|_2} \leq 1$*, and the iteration converges with the rate* $\left(\left\|I - \lambda \nabla^2_{\theta\theta} \mathcal{F}(q_*, \theta_*)\right\|_2 + \lambda \beta_q \beta_h\right)$*.*

Convergence Rate

Theorem 2 (Asymptotic Performance Gap in Tail Task Risk).

Under the ***Assumption*** *1 and given a batch of tasks* $\{\tau_i\}_{i=1}^{\mathcal{B}}$ *we can have*

$$CVaR_\alpha\left(\theta_T^{meta}\right) - CVaR_\alpha(\theta_*) \leq \beta_\tau \parallel \theta_T^{meta} - \theta_* \parallel + \frac{VaR_\alpha^*}{1-\alpha}\left(\mathbb{P}(\mathcal{T}_1) - \mathbb{P}(\mathcal{T}_2)\right),$$

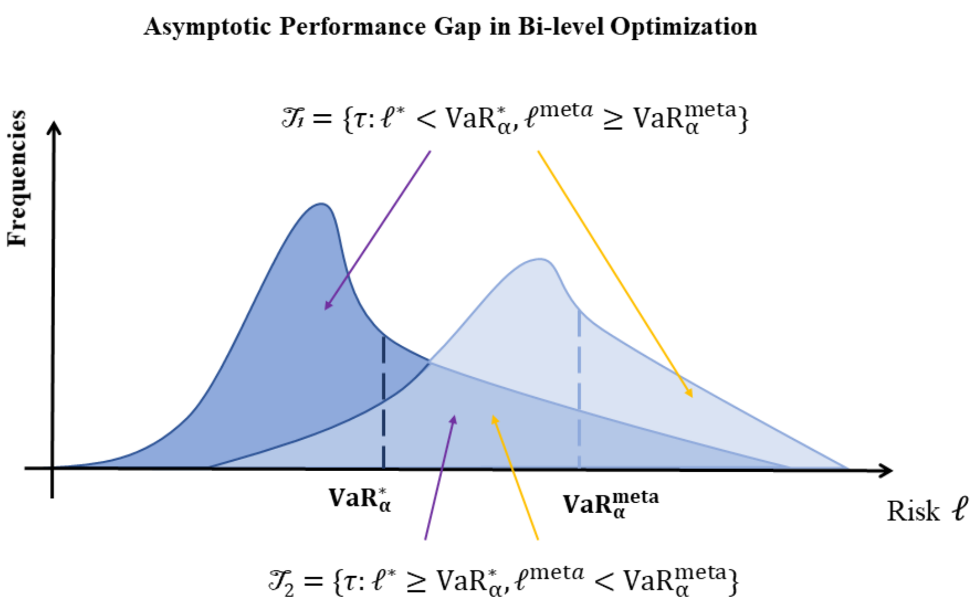

For sufficiently large T, the first term can be bounded by a small number due to the convergence, and the second term vanishes.

Generalization Bound

Theorem 3 (Generalization Bound in the Tail Risk Cases).

Given a collection of task samples $\{\tau_i\}_{i=1}^{\mathcal{B}}$ *and corresponding meta datasets, we can derive the following generalization bound in the presence of tail risk:*

$$R(\theta_*) \leq \hat{R}(\theta_*) + \sqrt{\frac{2\left(\frac{\alpha}{1-\alpha}\mathcal{L}_{\max}^2 + \mathbb{V}_{\tau_i \sim p_\alpha(\tau)}\left[\ell\left(\mathfrak{D}_{\tau_i}^Q, \mathfrak{D}_{\tau_i}^S; \theta_*\right)\right]\right)\ln\left(\frac{1}{\epsilon}\right)}{\mathcal{B}}}$$

$$+\frac{1}{3(1-\alpha)}\frac{\mathcal{L}_{\max}}{\mathcal{B}}\left(2\ln\left(\frac{1}{\epsilon}\right) + 3\alpha\mathcal{B}\right),$$

where the inequality holds with probability at least $1-\epsilon$ *and* $\epsilon \in (0,1)$, $\mathbb{V}[\cdot]$ *denotes the variance operation, and* $\mathcal{L}_{\max}$ *is from the* ***Assumption*** *1.*

Practical Enhancements

KDE can handle arbitrary complex distributions compared to crude Monte Carlo (MC) methods

$$F_{\ell\text{-KDE}}(l;\theta) = \int_{-\infty}^{l} \frac{1}{\mathcal{B}h_\ell} \sum_{i=1}^{\mathcal{B}} K\left(\frac{t-\ell\left(\mathfrak{D}_{\tau_i}^{Q}, \mathfrak{D}_{\tau_i}^{S};\theta\right)}{h_\ell}\right) dt,$$

- **Theorem 4**. *Let* $F_{\ell\text{-KDE}}^{-1}(\alpha;\theta) = \mathrm{VaR}_\alpha^{\mathrm{KDE}}[\ell(\mathcal{T},\theta)]$ *and* $F_\ell^{-1}(\alpha;\theta) = \mathrm{VaR}_\alpha[\ell(\mathcal{T},\theta)]$. *Suppose that* $K(x)$ *is lower bounded by a constant,* $\forall x$. *For any* $\epsilon > 0$, *with probability at least* $1-\epsilon$, *we can have the following bound:* $\sup_{\theta\in\Theta}\left(F_{\ell\text{-KDE}}^{-1}(\alpha;\theta) - F_\ell^{-1}(\alpha;\theta)\right) \le \mathcal{O}\left(\frac{h_\ell}{\sqrt{\mathcal{B}*\log\mathcal{B}}}\right)$.

- **Remark 1.** The crude Monte Carlo used in typically incurs an error of approximately $\mathcal{O}\left(\frac{1}{\sqrt{\mathcal{B}}}\right)$ in estimating quantiles . In contrast, that of KDE is no more than $\mathcal{O}\left(\frac{h_\ell}{\sqrt{\mathcal{B}*\log\mathcal{B}}}\right)$ from **Theorem** 4.

Empirical Findings

Benchmarks & Baselines

- **Benchmarks.**

- **Baselines.** MAML mainly works as the base meta learner, and we consider vanilla MAML , TR-MAML , DRO-MAML and DR-MAML.

- **Evaluations.** Expected/empirical risk minimization (Average), worst-case risk minimization (Worst), and tail risk minimization (CVaR$_\alpha$).

Sinusoid Regression

- **Problem Setup.** The goal of the sinusoid regression is to quickly fit an underlying function $f(x) = A\sin(x - B)$ from K randomly sampled data points, and tasks are specified by (A, B).

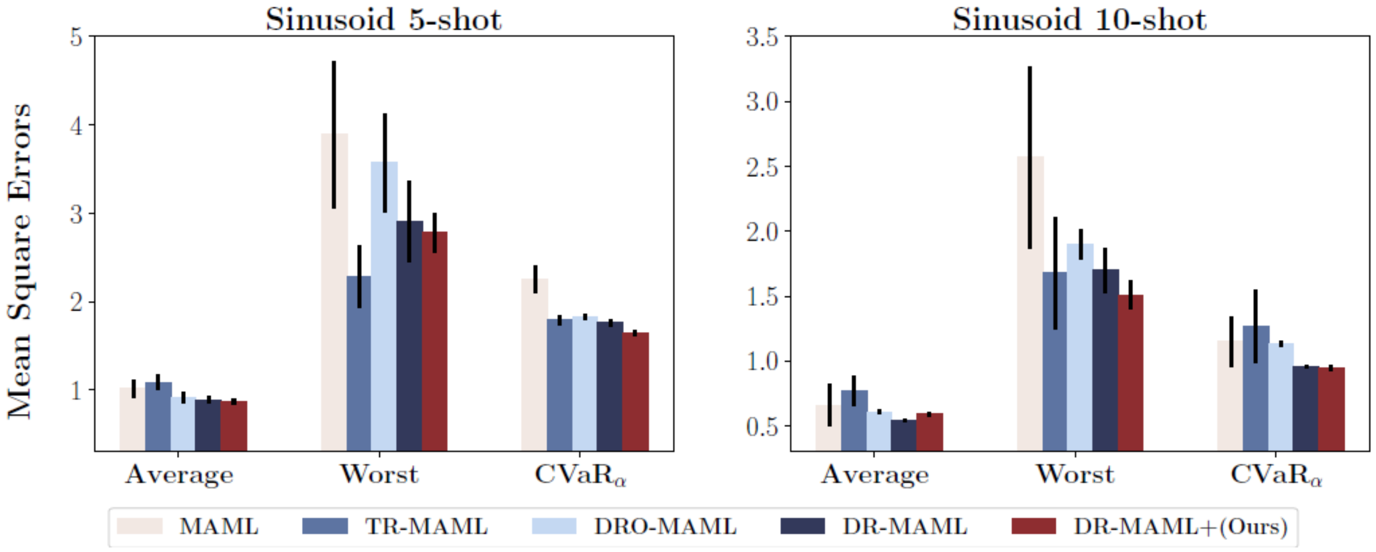

- **Result Analysis.** (1) DR-MAML+ consistently outperforms all baselines across average and CVaR$_\alpha$ indicators in the **5-shot** case. DR-MAML+ exhibits more robustness in challenging task distributions, *e.g.*, **5-shot** case.

  (2) The standard error associated with our method is significantly smaller than others, underscoring the stability of DR-MAML+.

System Identification

- **Problem Setup.** The system identification corresponds to learning a dynamics model from a few collected transitions in physics systems.

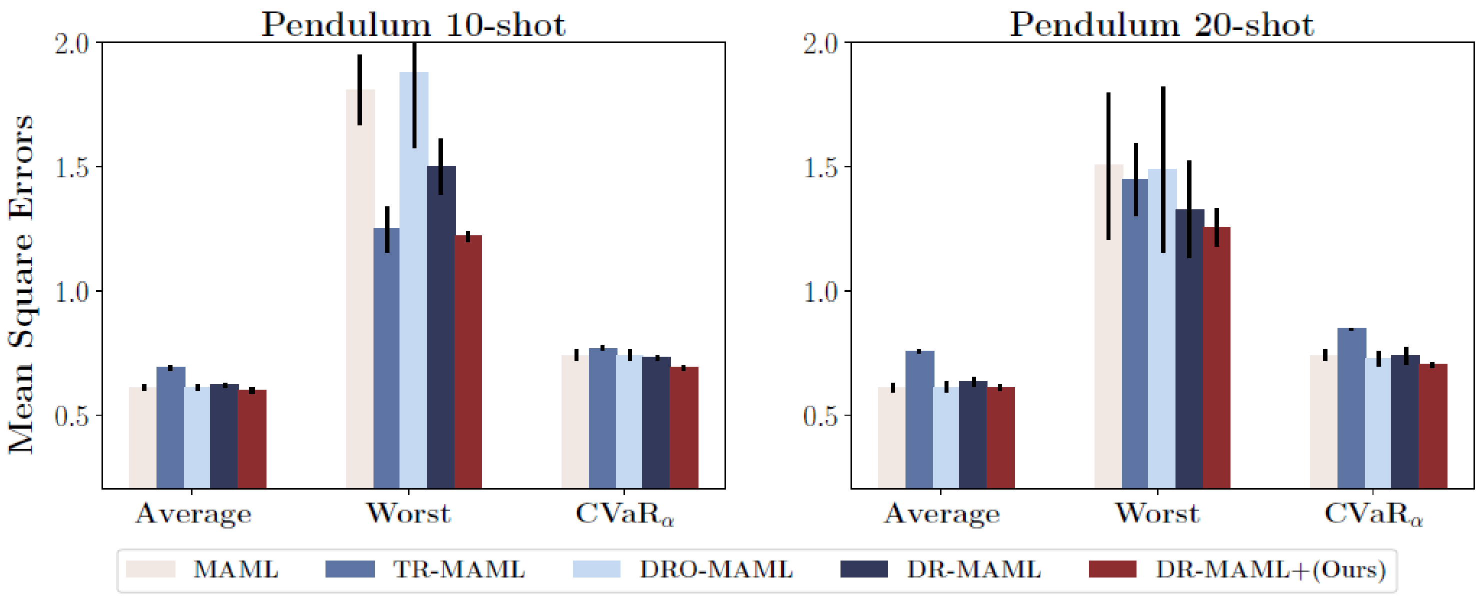

- **Result Analysis.** (1) There is no significant difference between **10-shot** and **20-shot** cases. DR-MAML+ dominates the performance across all indicators in both cases.

  (2) TR-MAML behaves well in the worst-case but sacrifices too much average performance.

  (3) DR-MAML+ exhibits an advantage over DR-MAML regarding CVaR$_\alpha$.

Few-Shot Image Classification

- **Problem Setup.** The task is a 5-way 1-shot classification problem. And 64 classes are selected for constructing meta-training tasks, with the remaining 32 classes for meta-testing.

Table 1: **Average** `5-way 1-shot` **classification accuracies in *mini*-ImageNet with reported standard deviations (3 runs).** With $\alpha = 0.5$, the best results are in bold.

Method	Eight Meta-Training Tasks			Four Meta-Testing Tasks		
	Average	Worst	CVaR_α	Average	Worst	CVaR_α
MAML [23]	70.1±2.2	48.0±4.5	63.2±2.6	46.6±0.4	44.7±0.7	44.6±0.7
TR-MAML [37]	63.2±1.3	60.7±1.6	62.1±1.2	48.5±0.6	45.9±0.8	46.6±0.5
DRO-MAML [60]	67.0±0.2	56.6±0.4	61.6±0.2	49.1±0.2	46.6±0.1	47.2±0.2
DR-MAML [1]	70.2±0.2	63.4±0.2	67.2±0.1	49.4±0.1	47.1±0.1	47.5±0.1
DR-MAML+(Ours)	**70.4**±**0.1**	**63.8**±**0.2**	**67.5**±**0.1**	**49.9**±**0.1**	**47.2**±**0.1**	**48.1**±**0.1**

- **Result Analysis.** (1) Methods within a two-stage distributionally robust strategy, namely DR-MAML and DR-MAML+, show superiority to others across all indicators in both training and testing scenarios.

  (2) DR-MAML+ and DR-MAML are comparable in most scenarios, and we attribute this to the small batch size in training, which weakens KDE' quantile approximation advantage.

Meta Reinforcement Learning

- **Problem Setup.** Take 2-D point robot navigation as the benchmark. The goal is to reach the target destination with the help of a few exploration transitions for fast adaptation.

Table 2: **Meta testing returns in point robot navigation (4 runs).** The chart reports average return and CVaR_α return with $\alpha = 0.5$.

Method	Average	CVaR_α
MAML [23]	-21.1 $\pm$ 0.69	-29.2 $\pm$ 1.37
DRO-MAML [60]	-20.9 $\pm$ 0.41	-29.0 $\pm$ 0.66
DR-MAML [1]	-19.6 $\pm$ 0.49	-28.9 $\pm$ 1.20
DR-MAML+(Ours)	**-19.2** $\pm$ **0.44**	**-28.4** $\pm$ **0.86**

- **Result Analysis.** (1) DR-MAML+ benefits from a more reliable quantile estimate and achieves superior performance.

  (2) The application of distributional robustness to reinforcement learning yields improvements in returns.

Extensions on Large Models

- **Benchmarks.** Tiered-ImageNet, ImageNetA, ImageNetSketch
- **Baselines.** CLIP, MaPLe, DR-MaPLe

Meta testing results on 5-way 1-shot ImageNetSketch classification (3 runs)

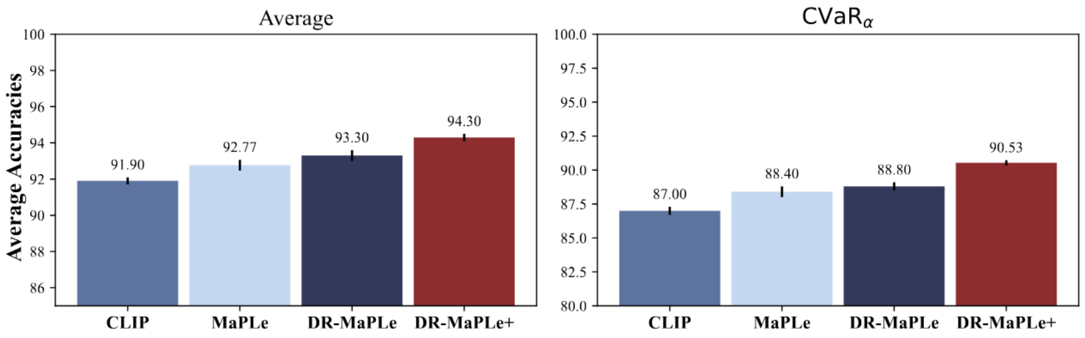

Meta testing results on 5-way 1-shot ImageNetA classification (3 runs)

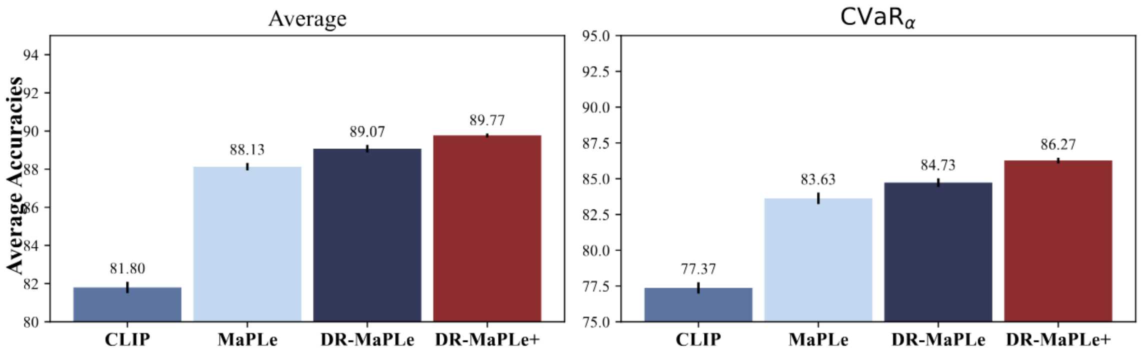

- **Result Analysis.** (1) DR-MaPLe and DR-MaPLe+ consistently outperform baselines across both average and CVaR$_\alpha$ indicators in **5-way 1-shot** cases, demonstrating the advantage of the two-stage strategy in enhancing the robustness of few-shot learning

  (2) DR-MaPLe+ achieves better results as KDE quantiles are more accurate with large batch sizes. These results examine the scalability and compatibility of our method on large models..

Assessment of Quantile Estimators

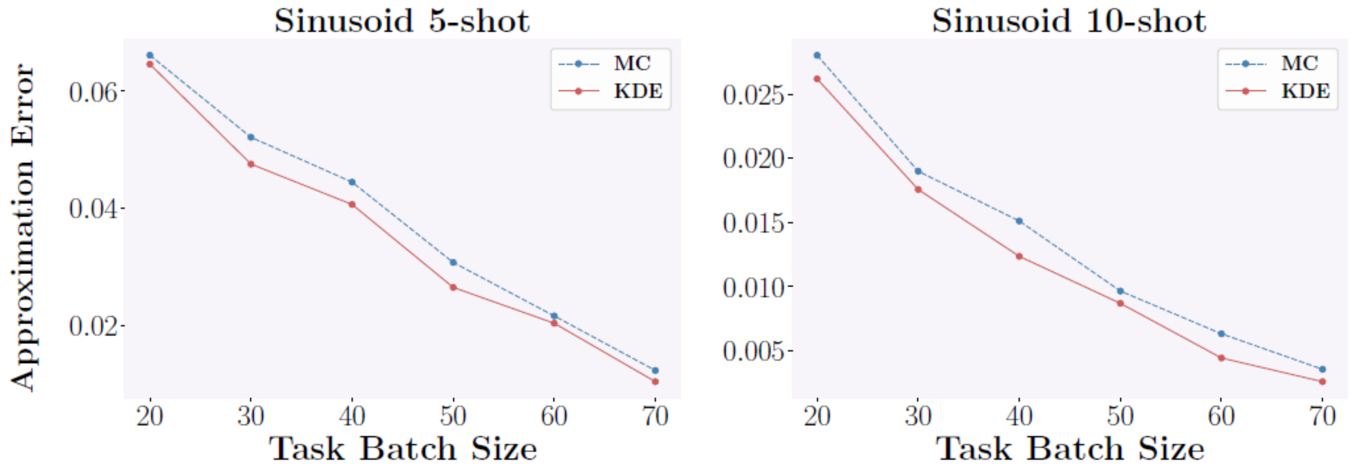

The VaR_α approximation error decreases with more tasks.

The KDE produces more accurate estimates with a sharper decreasing trend.

The above well verifies the conclusion in **Theorem** 3.

Other Investigations

■ Evaluation with Other Robust Meta Learners

Table 4: **MSEs for Sinusoid 5-shot with reported standard deviations (5 runs).** With $\alpha = 0.7$ the best results are in bold.

Method	Average	Worst	CVaR$_\alpha$
CNP [15]	0.09±0.00	2.71±0.54	0.24±0.01
TR-CNP [37]	0.10±0.01	1.51±0.30	0.22±0.03
DRO-CNP [60]	0.09±0.02	2.54±1.81	0.21±0.05
DR-CNP [1]	0.09±0.01	1.62±0.45	0.20±0.02
DR-CNP+(Ours)	**0.08**±0.01	**1.47**±0.90	**0.17**±0.02

Table 5: **MSEs for Pendulum 10-shot with reported standard deviations (5 runs).** With $\alpha = 0.5$, the best results are in bold.

Method	Average	Worst	CVaR$_\alpha$
CNP [15]	0.75±0.01	1.51±0.23	0.87±0.02
TR-CNP [37]	0.76±0.00	**1.24**±0.02	0.85±0.01
DRO-CNP [60]	0.73±0.01	1.51±0.16	0.85±0.01
DR-CNP [1]	0.75±0.01	1.40±0.16	0.86±0.01
DR-CNP+(Ours)	**0.72**±0.01	1.36±0.07	**0.82**±0.00

■ Sensitivity Analysis to Confidence Level

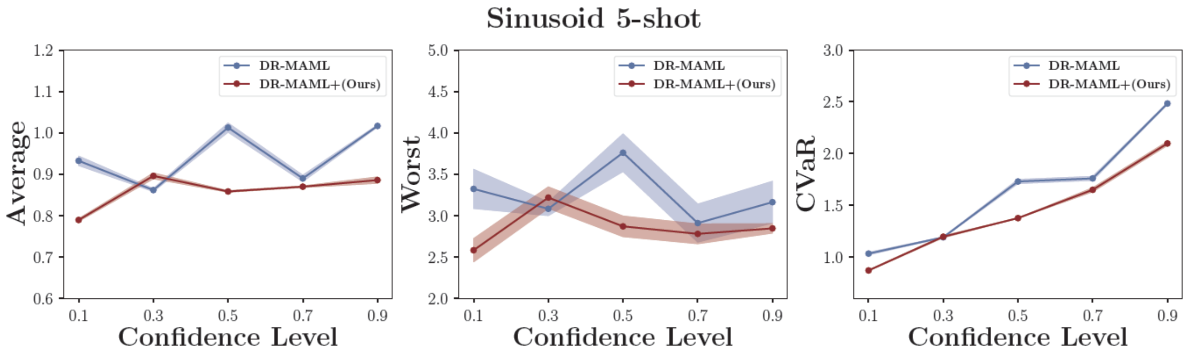

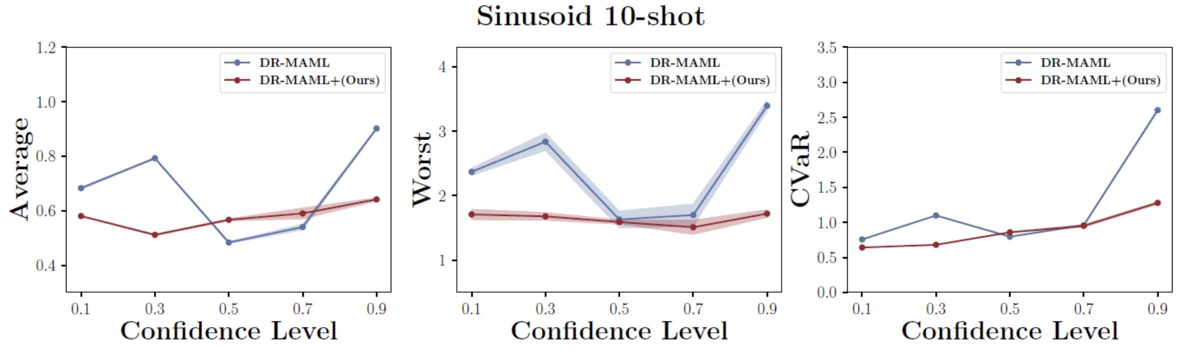

Conclusion

Technical Comparison

Principle	Meta Learner	Generalization	Convergence	Robustness Type
MAML	$\min_{\theta\in\Theta}\mathbb{E}_{p(\tau)}\left[\ell(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta)\right]$	✓	✓	$--$
DRO-MAML	$\max_{q(\tau)\in\mathcal{Q}}\min_{\theta\in\Theta}\mathbb{E}_{q(\tau)}\left[\ell(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta)\right]$	✗	✗	Uncertainty Set (Not tail risk)
TR-MAML	$\min_{\theta\in\Theta}\max_{\tau\in\mathcal{T}}\ell(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta)$	✓	✓	Worst-Case Task
DR-MAML	$\min_{\theta\in\Theta}\mathbb{E}_{p_\alpha(\tau;\theta)}\left[\ell(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta)\right]$	✗	✗	Tail Task Risk
DR-MAML+(Ours)	$\max_{q(\tau)\in\mathcal{Q}_\alpha}\min_{\theta\in\Theta}\mathbb{E}_{q(\tau)}\left[\ell(\mathfrak{D}_\tau^Q,\mathfrak{D}_\tau^S;\theta)\right]$	✓	✓	Tail Task Risk

- DRO-MAML [5] includes the uncertainty set $\mathcal{Q}$ for robust fast adaptation, there exists *no theoretical analysis*.
- TR-MAML [4] only focuses on the worst-case, which considers *a bit extreme and rarely occurred cases*.
- DR-MAML [3] *lacks generalization capability and convergence rate analysis* *w.r.t.* the meta learner.
- DR-MAML+ is a more specific instantiation of that in DR-MAML.

Conclusion

- Proposes to understand the two-stage distributionally robust strategy from optimization processes.
- Defines the convergence solution, and derives the generalization bound in the presence of tail task risk.
- Extensive evaluations demonstrate the significance of our proposal and its scalability to multimodal large models in boosting robustness.

References

[1] Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In International conference on machine learning, pages 1126–1135. PMLR, 2017.

[2] Greenberg, I., Mannor, S., Chechik, G., and Meirom, E. Train hard, fight easy: Robust meta reinforcement learning. Advances in Neural Information Processing Systems, 36, 2024.

[3] Wang, C., Lv, Y., Feng, Y., Xie, Z., and Huang, J. A simple yet effective strategy to robustify the meta learning paradigm. In Thirty-seventh Conference on Neural Information Processing Systems, 2023a.

[4] Collins, L., Mokhtari, A., and Shakkottai, S. Task-robust model-agnostic meta-learning. Advances in Neural Information Processing Systems, 33:18860-18871, 2020.

[5] Sagawa, S., Koh, P. W., Hashimoto, T. B., & Liang, P. Distributionally robust neural networks. In International Conference on Learning Representations, 2020.

Thanks for your attention  .