Online Constrained Meta-Learning: Provable Guarantees for Generalization

Siyuan Xu & Minghui Zhu

School of Electrical Engineering and Computer Science
The Pennsylvania State University

Neural Information Processing Systems
December, 2023

Online constrained meta-learning

Online meta-learning: learn to perform tasks while accumulating the learning experience for future tasks

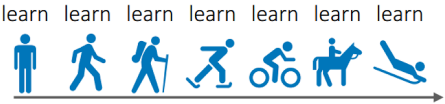

Constrained learning tasks:

Safe Learning

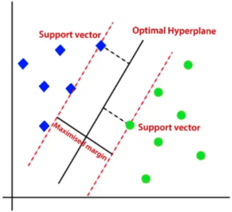

Supported Vector Machine

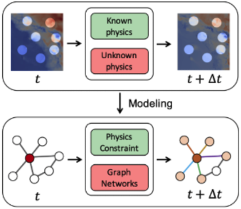

Physics-informed Learning

Constrained meta-learning: learn to constrained learning

Sequential constrained learning tasks

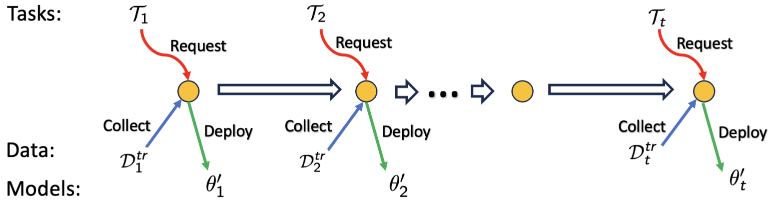

- The optimal solution θ_t^* for task $\mathcal{T}_t$

$$\theta_t^* = \underset{\theta \in \Theta}{\operatorname{argmin}} \ \mathbb{E}_{z \sim \mathcal{D}_{0,t}} [\ell_0(\theta, z)] \ \text{s.t.} \ \mathbb{E}_{z \sim \mathcal{D}_{i,t}} [\ell_i(\theta, z)] \leq c_{i,t}, \ \ i = 1, \ldots, m.$$

- Challenges:
  - The data distribution $\mathcal{D}_t = \{\mathcal{D}_{0,t}, \mathcal{D}_{1,t}, \ldots, \mathcal{D}_{m,t}\}$ is unknown, and only $\mathcal{D}_t^{tr} = \{\mathcal{D}_{0,t}^{tr}, \cdots, \mathcal{D}_{m,t}^{tr}\}$ can be (i.i.d.) sampled from the distribution $\mathcal{D}_t$.
  - Safety constraints should be satisfied.

Online constrained meta-learning framework

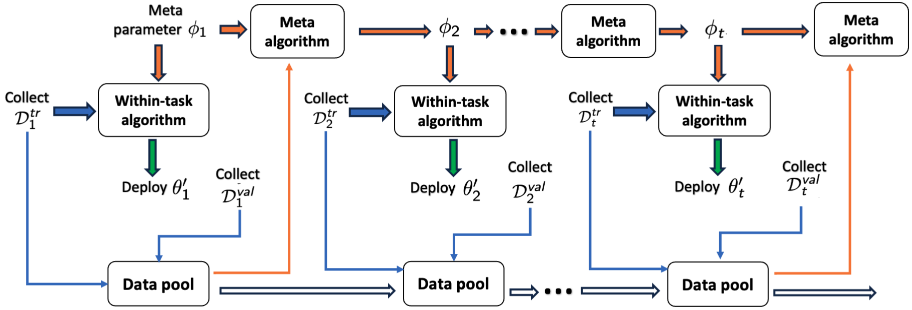

In each round:

- Adapt task-specific parameter θ_t' from the meta-parameter ϕ_t by a within-task algorithm
- Deploy model with parameter θ_t' for task $\mathcal{T}_t$
- Update the meta-parameter ϕ_t by a meta-algorithm

Online constrained meta-learning framework

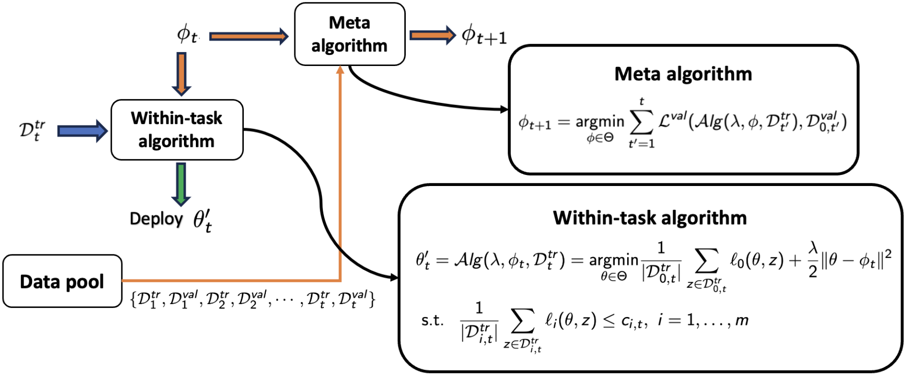

- Within-task algorithm: constrained optimization with biased regularization.
- Meta algorithm: constrained bilevel optimization[1].

[1] Xu and Zhu, "Efficient gradient approximation method for constrained bilevel optimization".

Metrics

Optimality metric

The task-averaged optimality gap (TAOG) denoted by $\bar{R}_{0,[1:T]}$:

$$\bar{R}_{0,[1:T]} = \frac{1}{T}\sum_{t=1}^{T} \max\left\{\mathbb{E}_{z\sim\mathcal{D}_{0,t}}\left[\ell_0(\theta_t', z) - \ell_0(\theta_t^*, z)\right], 0\right\}.$$

Recall that θ_t^* is the optimal parameter given the whole data distribution:

$$\theta_t^* = \operatorname*{argmin}_{\theta\in\Theta} \mathbb{E}_{z\sim\mathcal{D}_{0,t}}\left[\ell_0(\theta, z)\right] \text{ s.t. } \mathbb{E}_{z\sim\mathcal{D}_{i,t}}\left[\ell_i(\theta, z)\right] \le c_{i,t},\ i = 1, \ldots, m$$

Constraint violation metric

The task-averaged constraint violation (TACV) denoted by $\bar{R}_{i,[1:T]}$:

$$\bar{R}_{i,[1:T]} = \frac{1}{T}\sum_{t=1}^{T} \max\left\{\mathbb{E}_{z\sim\mathcal{D}_{i,t}}\left[\ell_i(\theta_t', z)\right] - c_{i,t}, 0\right\},\ i = 1, \ldots, m.$$

Task dissimilarity

The dissimilarity of tasks $\{\mathcal{T}_1, \cdots, \mathcal{T}_T\}$ is defined as

$$\mathcal{S}^*(\mathcal{T}_{1:T}) \triangleq \min_\phi \sqrt{\frac{1}{T}\sum_{t=1}^{T}\frac{1}{2}\|\theta_t^* - \phi\|^2}$$

- The dissimilarity of tasks is defined by the standard deviation of the optimal task-specific parameters.

Theoretical guarantee

Upper bounds of TAOG and TACV

Suppose that Θ is included in a compact cube with edge of length D, i.e. $\|\phi\|_\infty \leq D$ for any $\phi \in \Theta$. Choose the regularization weight $\lambda = \frac{2\sqrt{d}(\rho\mathcal{B}+L_0\sqrt{\ln|\mathcal{D}_0^{tr}|})}{\mathcal{S}^*(\mathcal{T}_{1:T})\sqrt{|\mathcal{D}_0^{tr}|}}$. For the task sequence $\{\mathcal{T}_1, \cdots, \mathcal{T}_T\}$, the following bounds hold for the TAOG and the TACV of the proposed algorithm:

$$\mathbb{E}[\bar{R}_{0,[1:T]}] \leq \mathcal{O}(\mathcal{S}^*(\mathcal{T}_{1:T})\sqrt{\frac{\ln|\mathcal{D}_0^{tr}|}{|\mathcal{D}_0^{tr}|}} + \sqrt{\frac{\ln|\mathcal{D}_+^{tr}|}{|\mathcal{D}_+^{tr}|}} + \sqrt{\frac{\ln|\mathcal{D}_0^{val}|}{|\mathcal{D}_0^{val}|}} + \frac{1}{\sqrt{T}}),$$

and

$$\mathbb{E}[\bar{R}_{i,[1:T]}] \leq \mathcal{O}\left(\sqrt{\frac{\ln|\mathcal{D}_+^{tr}|}{|\mathcal{D}_+^{tr}|}}\right), \ \forall i = 1, \ldots, m.$$

Experiment 1: Meta-imitation learning

In each round, the expert performs a demonstration of a task in a **free space**. The learner can observe a **small number** of data points from the demonstration and needs to perform the task in a new **clustered space**.

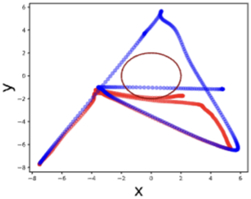

Starting from scratch

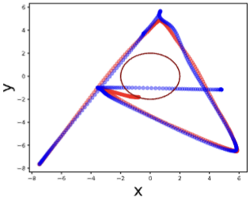

MAML with constraint penalty

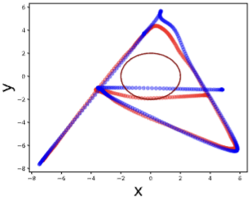

Constrained meta-learning

Experiment 1: Meta-imitation learning

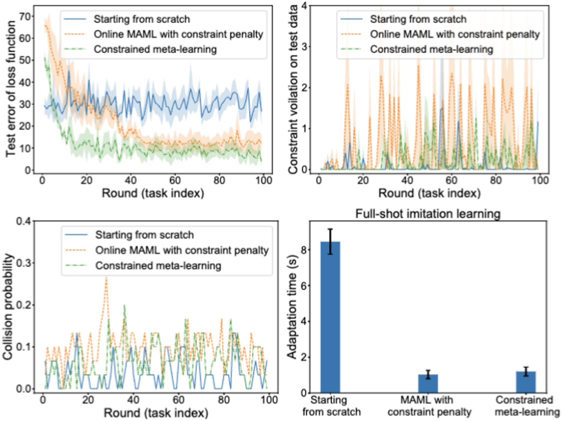

- Our algorithm outperforms the benchmarks in terms of test error and collision avoidance, and speeds up the adaptation to new tasks.

Experiment 2: Few-shot image classification with robustness

- For a new task, the training data include 25 images (5-way 5-shot setting) or 5 images (5-way 1-shot setting).
- Test on test data (clean data) and adversarial-attacked test data (PGD data).

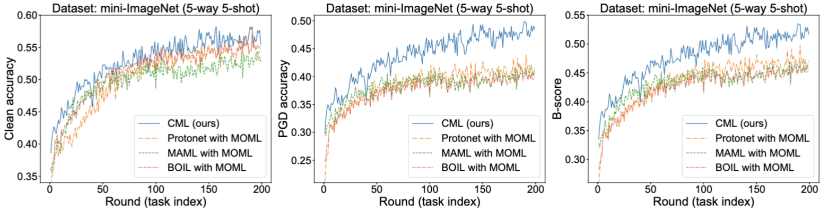

- Our method outperforms the benchmarks in terms of both test accuracy and learning speed

Experiment 2: Few-shot image classification with robustness

Table 2: Clean accuracy (abbreviated as "Clean Acc.") and PGD accuracy (abbreviated as "PGD Acc.") on the mini-ImageNet dataset for 5-way 5-shot and 5-way 1-shot learning.

	Method	Clean Acc.	PGD Acc.	B-score
1-shot	MAML + CP	40.78 ± 0.75	23.91 ± 0.67	29.83 ± 0.43
	MAML + MOML	39.23 ± 0.76	25.80 ± 0.67	31.12 ± 0.70
	ProtoNet + CP	38.65 ± 0.72	23.10 ± 0.65	28.67 ± 0.67
	ProtoNet + MOML	35.06 ± 0.70	27.24 ± 0.65	30.51 ± 0.66
	BOIL + CP	40.44 ± 0.79	25.94 ± 0.69	31.29 ± 0.75
	BOIL + MOML	41.22 ± 0.83	27.77 ± 0.75	32.98 ± 0.79
	CML (**ours**)	39.52 ± 0.80	$\mathbf{33.11 \pm 0.79}$	$\mathbf{36.03 \pm 0.79}$
5-shot	MAML + CP	56.16 ± 0.72	34.85 ± 0.72	42.91 ± 0.71
	MAML + MOML	55.66 ± 0.78	39.38 ± 0.77	45.89 ± 0.77
	ProtoNet + CP	59.11 ± 0.71	39.41 ± 0.73	46.93 ± 0.71
	ProtoNet + MOML	58.72 ± 0.74	41.59 ± 0.75	48.59 ± 0.74
	BOIL + CP	58.54 ± 0.76	34.28 ± 0.75	42.94 ± 0.78
	BOIL + MOML	60.21 ± 0.79	35.47 ± 0.78	44.37 ± 0.78
	CML (**ours**)	59.74 ± 0.75	$\mathbf{49.48 \pm 0.76}$	$\mathbf{54.01 \pm 0.74}$

- Our algorithm significantly improves the PGD accuracy than the benchmarks and keeps the clean accuracy comparable.