Improved Regret Bounds for Non-Convex Online-Within-Online Meta Learning

Jiechao Guan, Hui Xiong

AI Thrust, The Hong Kong University of Science and Technology (Guangzhou), China

{jiechaoguan, xionghui}@hkust-gz.edu.cn

March 10, 2024

Presentation Outline

Basic Notations

- In online learning, an action space Θ is a compact subset of $\mathbb{R}^d$.
- A loss function $\ell : \Theta \mapsto \mathbb{R}_{\geq 0}$ is called uniformly L-Lipschitz over Θ w.r.t. the norm $\|\cdot\|$, if for any $x, y \in \Theta$, $|\ell(x) - \ell(y)| \leq L\|x - y\|$.
- Let $\mathcal{P}(\Theta)$ be the set of all probability distributions over action space Θ. For any $\rho \in \mathcal{P}(\Theta)$, we use $\langle \rho, \ell \rangle = \mathbb{E}_{\theta \sim \rho} \ell(\theta)$ for brevity.

Non-Convex Online Learning

- At each round $i \in [m]$, the online learner selects an action $\theta_i \in \Theta$, receives a loss function $\ell_i(\cdot) : \Theta \mapsto \mathbb{R}_{\geq 0}$, and suffers loss $\ell_i(\theta_i)$.
- The quantity used to measure performance of online learner is the *regret* $R_m \triangleq \sum_{i=1}^{m} \ell_i(\theta) - \ell_i(\theta^*)$, where $\theta^* \in \arg\min_{\theta \in \Theta} \sum_{i=1}^{m} \ell_i(\theta)$.
- In non-convex online learning, at each round $i \in []$ we need to add randomness into the algorithm, otherwise there exists setting where using any algorithm obtains vacuous regret $R_m = O(m)$.
- We choose action drawn from distribution $\rho_i \in \mathcal{P}(\Theta)$. The regret is $R_m \triangleq \sum_{i=1}^{m} \mathbb{E}_{\theta \sim \rho_i} \ell_i(\theta) - \ell_i(\theta^*)$, $\theta^* \in \arg\min_{\theta \in \Theta} \sum_{i=1}^{m} \mathbb{E}_{\theta \sim \rho_i} \ell_i(\theta)$.

Online-Within-Online (OWO) Meta Learning

- In non-convex OWO meta learning setting, the online meta-learner will encounter T tasks with each composed of functions $\{\ell_{ti}\}_{i\in[m]}$.

- Concretely, at i-th round of the t-th task, the online meta-learner selects an action $\theta_{ti} \in \Theta$, and then suffers the loss $\ell_{ti}(\theta_{ti})$.

- The quantity to measure the performance of online meta-learner is the following task-averaged regret [1, 2]:

$$\bar{R}_{T,m} = \frac{1}{T}\sum_{t=1}^{T} R_{t,m} = \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{m} \mathbb{E}_{\theta\sim\rho_{ti}}\ell_{ti}(\theta) - \ell_{ti}(\theta_t^*).$$

Regret Bound Decomposition Framework for Non-Convex OWO Meta Learning

- Non-convex algorithms: online mirror descent (OMD), exponential weighted aggregation (EWA), follow the perturbed leader (FTPL).

- The regret upper bound for many non-convex algorithms (e.g. EWA) at task $t \in [T]$ has the form: $U_t(\rho_{t1}, \lambda_t) = mb\lambda_t + \frac{V(\rho_{t1}, \rho_t^*)^2}{\lambda_t}$, where $b > 0$, λ_t is the step size, $\rho_{t1} \in \mathcal{P}(\Theta)$ the initialization distribution over Θ , $V(\rho_{t1}, \rho_t^*)^2$ the non-negative function of ρ_{t1} and ρ_t^*.

- Let $\lambda = v/\sqrt{mb}$ for $v > 0$, then $U_t(\rho, v) = (v + \frac{f_t(\rho)}{v})\sqrt{mb}$. The regret bound decomposition framework consists of two processes: one is to minimize $\{f_s(\rho) = V(\rho, \rho_s^*)^2\}_{s\in[t-1]}$ to determine the initialization distribution ρ_{t1} for task t, another is to minimize $\{h_s(v) = v + f_s(\rho_{s1})/v\}_{s\in[t-1]}$ to determine the step size v_t.

General Regret Bound for Non-Convex OWO Meta Learning

Theorem 2.1

Assume the upper regret bound for task $t \in [T]$ *has the form* $U_t(\rho, v) = (v + \frac{f_t(\rho)}{v})\sqrt{mb} + g(m)$. *Assume we have a sub-algorithm that achieves* $F_T(\rho)$ *regret w.r.t. any* $\rho \in \mathcal{P}(\Theta)$ *by setting distributions* ρ_{t1} *on* $f_t(\rho) = V(\rho, \rho_t^*)^2$, *and another sub-algorithm that achieves non-increasing* $H_T(v)$ *regret w.r.t. any* $v > 0$ *by playing actions* $v_t > 0$ *on* $h_t(v) = v + \frac{f_t(\rho_{t1})}{v}$ *for all* $t \in [T]$. *Then, running the OWO meta learning algorithm with the step size* $v_t/\sqrt{mb}$ *and initialization* ρ_{t1} *at each task* t, *for* $\rho^* = \arg\min_{\rho \in \mathcal{P}(\Theta)} \sum_{t=1}^{T} f_t(\rho)$ *the optimal initialization and* V *the task-similarity, we get the task-averaged regret upper bound:*

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{m}\langle \rho_{ti}-\rho_t^*, \ell_{ti}\rangle \leq \Big(\frac{H_T(V)}{T}+\min\{\frac{F_T(\rho^*)}{VT}, 2\sqrt{\frac{F_T(\rho^*)}{T}}\}+2V\Big)\sqrt{mb}+g(m).$$

General Regret Bound for Learning Step Size in Non-Convex OWO Meta Learning

For learning the step size v_t, we consistently use Follow-The-Leader (FTL) algorithm to achieve the following logarithmic regret bound.

Proposition 1

Assume that FTL algorithm runs on the sequence of functions $\{h_t(v) = v + \frac{f_t(\rho_{t1})}{v}\}_{t\in[T]}$ *over the domain* $[0, D]$*, where* $D^2 \geq \max_{t\in[T]} f_t(\rho_{t1})$*, then we have the regret bound:*

$$\sum_{t=1}^{T} h_t(v_t) - \min_{v\in[0,D]} \sum_{t=1}^{T} h_t(v) \leq \frac{D^3}{4} \sum_{t=1}^{T} \frac{\left|1 - \sum_{s=1}^{t-1} f_t(\rho_{t1}) / \sum_{s=1}^{t-1} f_s(\rho_{s1})\right|^2}{\sum_{s=1}^{t} f_s(\rho_{s1})}.$$

Furthermore, if for all $t \in [T]$, $f_t(\rho_{t1}) \in [B^2, D^2]$ *with* $D \geq B > 0$*, then we have the logarithmic regret upper bound:* $\sum_{t=1}^{T} h_t(v_t) - \min_{v\in[0,D]} \sum_{t=1}^{T} h_t(v) \leq \frac{D^7}{4B^6}(\log T + 1)$.

Comparisons between Different Regret Bounds for OWO Meta Learning

Table 1: Different task-averaged regret bounds for OWO meta learning algorithms under different assumptions of loss functions $\{\ell_{ti}\}_{t,i=1}^{T,m}$. T is the number of tasks, and m is the number of iterations per task. Concretely, the **task-averaged regret upper bound** = $\big($**Bound I** + **Bound II** + $V\big)\sqrt{m}$, where **Bound I** is the regret upper bound for learning the initialization, **Bound II** is the regret upper bound for learning the step size, V represents the task similarity among different tasks, $\alpha \in (0, 1/2)$. The explicit form of these regret bounds and task similarities are given in Table B.1 of the Appendix.

Existing Works	**Task-Averaged Regret**	**Assumptions of ℓ_{ti}**	**Bound I**	**Bound II**
Khodak et al. (2019)	$\frac{1}{T}\sum_{t,i=1}^{T,m} \ell_{ti}(\theta_{ti}) - \ell_{ti}(\theta_t^*)$	Convex & Uniformly Lipschitz	$O(\frac{\log T}{T})$	$O(\frac{\log T}{\sqrt{T}})$
Balcan et al. (2021)	$\frac{1}{T}\sum_{t,i=1}^{T,m} \langle \rho_{ti}, \ell_{ti} \rangle - \ell_{ti}(\theta_t^*)$	Bounded & Piecewise Lipschitz	$O(\frac{m^{d/2}}{T^{1/4}})$	$O(\frac{(\log m)\log T}{\sqrt{T}})$
Our Theorem 2	$\frac{1}{T}\sum_{t,i=1}^{T,m} \langle \rho_{ti}, \ell_{ti} \rangle - \ell_{ti}(\theta_t^*)$	Bounded & Piecewise Lipschitz	$O(\frac{1}{T^{1/2-\alpha}})$	$O(\frac{(\log T)^{9/2}}{T})$
Our Theorem 3	$\frac{1}{T}\sum_{t,i=1}^{T,m} \langle \rho_{ti} - \rho_t^*, \ell_{ti} \rangle$	Bounded	$O(\frac{\log T}{T})$	$O(\frac{\log T}{T})$

Regret Bound Decomposition Framework for Non-Convex Piecewise Lipschitz Functions

Definition 2.2

(**Piecewise Lipschitzness** [3]) The sequence of random loss functions $\{\ell_i\}_{i=1}^m$ is piecewise L-Lipschitz ($L > 0$) that are β-dispersed, if $\forall m, \forall \epsilon \geq m^{-\beta}$, in expectation over the randomness of the loss functions, we have $\mathbb{E}[\max_{\|\theta-\theta'\|_2 \leq \epsilon} \left|\{i \in [m] \mid \ell_i(\theta) - \ell_i(\theta') > L\|\theta - \theta'\|_2\}\right|] \leq \tilde{O}(\epsilon m)$.

- With M-bounded piecewise Lipschitz functions, the regret bound for EWA algorithm with the initialization ρ_1 and the step size λ:

$$\sum_{i=1}^{m} \mathbb{E}_{\theta\sim\rho_i}\ell_i(\theta) - \sum_{i=1}^{n} \ell_i(\theta^*) \leq \lambda M^2 m + \frac{V(\rho_1, \theta^*)^2}{\lambda} + \tilde{O}((L+1)m^{1-\beta}).$$

- Choose Follow-The-Regularized-Leader (FTRL) and FTL algorithms to learn initialization ρ_{t1} and step size v_t respectively for task t.

Algorithm 1 Non-convex OWO meta learning algorithm for bounded piecewise Lipschitz functions.

1: **Input:** step size η for FTRL, mixture parameter $\gamma \in (0,1]$, domain upper bound D; initialized distribution $\rho_{11} : \Theta \mapsto \mathbb{R}_{\geq 0}$ and initialized step size $\lambda_1 = \sqrt{(D^2 - \log\gamma)/(mM^2)}$ for EWA.
2: **for** task $t \in [T]$ **do**
3: **for** round $i \in [m]$ **do**
4: Set $P_{ti} = \int_\Theta \rho_{ti}(\theta)\mathrm{d}\theta$ and sample θ_{ti} with probability $p_{ti}(\theta_{ti}) = \rho_{ti}(\theta_{ti})/P_{ti}$
5: Suffer loss $\ell_{ti}(\theta_{ti})$, observe $\ell_{ti}(\cdot)$, and update $\rho_{t,i+1}(\theta) = e^{-\lambda_t \ell_{ti}(\theta)}\rho_{ti}(\theta)$ // EWA step
6: Sample θ_t^* with probability p_{tm} and obtain task-t optimum $\theta_t^* \in \Theta$
7: Set $\mathbf{1}_{\mathcal{B}(\theta_t^*, m^{-\beta})}$ to be the function that is **1** in the ball $\mathcal{B}(\theta_t^*, m^{-\beta})$ and otherwise 0
8: Update $\rho_{t+1,1}$ to $\boldsymbol{\rho}_{t+1}(t) = \arg\min_{\|\boldsymbol{\rho}\|_1 = 1, \boldsymbol{\rho} \geq \gamma\hat{\boldsymbol{\nu}}(t)} \mathrm{KL}(\boldsymbol{\rho}\|\hat{\boldsymbol{\nu}}(t)) - \eta\sum_{s \leq t} \log\langle \boldsymbol{\rho}_s^*(t), \boldsymbol{\rho}\rangle$, $\lambda_{t+1} = \sqrt{\frac{-\sum_{s\leq t}\log\langle \boldsymbol{\rho}_s^*(s), \boldsymbol{\rho}_s(s)\rangle}{tmM^2}}$ // meta-update step

Bounds for Learning Initialization and Step Size

Proposition 2

Assume that FTRL algorithm with initialization $\hat{\boldsymbol{\nu}}$ *and regularization* $\mathrm{KL}(\boldsymbol{\rho}||\hat{\boldsymbol{\nu}})$ *runs on functions* $\{f_t(\boldsymbol{\rho}_t)\}_{t\in[T]}$ *over the set* $\Delta = \{\boldsymbol{\rho} : \|\boldsymbol{\rho}\|_1 = 1, \boldsymbol{\rho} \geq \gamma\hat{\boldsymbol{\nu}}\}$*, then* $\sum_{t=1}^{T} f_t(\boldsymbol{\rho}_t) - f_t(\boldsymbol{\rho}^*) \leq 2GL\sqrt{T}/\gamma$*, where* $\boldsymbol{\rho}^* \in \arg\min_{\boldsymbol{\rho}\in\Delta} \sum_{t=1}^{T} f_t(\boldsymbol{\rho})$*,* $G^2 \geq \mathrm{KL}(\boldsymbol{\rho}^*||\hat{\boldsymbol{\nu}})$*,* $L^2 = \frac{1}{T}\sum_{t=1}^{T} \frac{\mathrm{vol}(\Theta)^2}{\mathrm{vol}(\Theta_t)^2}$*.*

Proposition 3

Assume that FTL algorithm runs on the sequence of functions $\{h_t(v) = v + \frac{f_t(\rho_{t1})}{v}\}_{t\in[T]}$ *with* $f_t(\rho_{t1}) = -\log\langle\boldsymbol{\rho}_t^*, \boldsymbol{\rho}_{t1}\rangle$ *over the domain* $[0, D]$*, where* $D^2 = \max_{t\in[T]} \log\frac{\mathrm{vol}(\Theta)}{\gamma\mathrm{vol}(\Theta_t)} \geq f_t(\rho_{t1})$*, for any* $t \in [T]$*. Denote* $B^2 = \min_{t\in[T]} \log\frac{\mathrm{vol}(\Theta)}{\gamma\mathrm{vol}(\Theta_t)+(1-\gamma)\mathrm{vol}(\Theta)}$*. Then if we set* $\gamma = m^{d\beta}/T^{\alpha} \in (0, 1]$ *with* $\alpha \in (0, 1/2)$*, we have the following regret upper bound:*

$$\sum_{t=1}^{T} h_t(v_t) - \min_{v\in[0,D]} \sum_{t=1}^{T} h_t(v) \leq \frac{(\alpha\log T + \log(\mathrm{vol}(\Theta)/V_d))^{7/2}(\log T + 1)}{4B^6}.$$

Regret Bounds for OWO Meta Learning Algorithms with Piecewise Lipschitz Functions

Theorem 2.3

Under the conditions of Theorem 2.1, for any task $t \in [T]$, let $\{\ell_{ti} : \Theta \mapsto [0, M]\}_{i \in [m]}$ be a sequence of piecewise L-Lipschitz functions that are β-dispersed. Set $\alpha \in (0, \frac{1}{2})$, let $g(m) = \tilde{O}((L+1)m^{1-\beta})$, then using FTRL algorithm in Proposition 2 and FTL algorithm in Proposition 3 respectively to learn the initialization and step size in Algorithm 1 obtains the regret upper bound:

$$\bar{R}_{T,m} \leq g(m) + \sqrt{m} M \Big\{ \frac{(\alpha \log T + \log(\mathrm{vol}(\Theta)/V_d))^{7/2} (\log T + 1)}{4TB^6} + \min\left\{ \frac{\mathrm{vol}(\Theta)(\mathrm{KL}(\boldsymbol{\rho}^* \| \hat{\boldsymbol{\nu}}) + 1)}{V_d V T^{\frac{1}{2}-\alpha}}, 2\sqrt{\frac{\mathrm{vol}(\Theta)(\mathrm{KL}(\boldsymbol{\rho}^* \| \hat{\boldsymbol{\nu}}) + 1)}{V_d T^{\frac{1}{2}-\alpha}}} \right\} + 2V \Big\}.$$

Regret Bound Decomposition Framework for Non-Convex Non-Lipschitz Functions

- EWA has the following regret upper bound [4]:

$$\sum_{i=1}^{m} \langle \rho_i - \rho, \ell_i \rangle \leq \lambda M^2 m + \frac{\mathrm{KL}(\rho || \rho_1)}{\lambda}, \quad \forall \rho \in \mathcal{P}(\Theta).$$

- Use FTL algorithm to learn initialization ρ_{t1} and step size v_t in task t.

Algorithm 2 Non-convex OWO meta learning algorithm for bounded non-Lipschitz functions.

1: **Input:** initialized distribution $\rho_{11} \in \mathcal{P}(\Theta)$ and learning rate $\lambda_1 > 0$.
2: **for** task $t \in [T]$ **do**
3: **for** round $i \in [m]$ **do**
4: $\rho_{ti} = \arg\min_{\rho \in \mathcal{P}(\Theta)} \mathrm{KL}(\rho || \rho_{t1}) + \lambda_t \sum_{j=1}^{i-1} \langle \ell_{tj}, \rho \rangle$ // EWA step
5: Suffer loss $\langle \rho_{ti}, \ell_{ti} \rangle$ and observe $\ell_{ti}(\cdot)$
6: Update $\rho_{t+1,1} = \frac{1}{t} \sum_{s=1}^{t} \rho_{sm}$, $\lambda_{t+1} = \sqrt{\frac{\sum_{s=1}^{t} \mathrm{KL}(\rho_{sm} || \rho_{s1})}{tmM^2}}$ // meta-update step

Bounds for Learning Initialization and Step Size

Proposition 4

Given a sequence of distributions $\{\rho_t^*\}_{t\in[T]}$, *assume that FTL algorithm runs on the sequence of* $\{\mathrm{KL}(\rho_t^*||\rho)\}_{t\in[T]}$ *to determine* ρ, *i.e.* $\rho_{t1} = \arg\min_{\rho\in\mathcal{P}(\Theta)} \sum_{s=1}^{t-1} \mathrm{KL}(\rho_s^*||\rho)$, *and further assume* $G^2 \geq \max_{t\in[T]} \chi^2(\rho_t^*||\rho_{t1})$, *then we can obtain the following regret upper bound:*

$$\sum_{t=1}^{T} \mathrm{KL}(\rho_t^*||\rho_{t1}) - \min_{\rho\in\mathcal{P}(\Theta)} \sum_{t=1}^{T} \mathrm{KL}(\rho_t^*||\rho) \leq \sum_{t=1}^{T} \frac{\chi^2(\rho_t^*||\rho_{t1})}{t} \leq G^2(\log T + 1).$$

Proposition 5

Given a sequence of functions $\{f_t(\rho_{t1}) = \mathrm{KL}(\rho_t^*||\rho_{t1})\}_{t\in[T]}$, *assume there exist* $D^2 \geq \max_{t\in[T]} \mathrm{KL}(\rho_t^*||\rho_{t1}), B^2 \leq \min_{t\in[T]} \mathrm{KL}(\rho_t^*||\rho_{t1})$ *with* $B > 0$. *Assume FTL algorithm runs on functions* $\{h_t(v) = v + \frac{f_t(\rho_{t1})}{v}\}_{t\in[T]}$ *over domain* $[0, D]$. *Then*

$$\sum_{t=1}^{T} h_t(v_t) - \min_{v\in[0,D]} \sum_{t=1}^{T} h_t(v) \leq \frac{D^7}{4B^6}(\log T + 1).$$

Regret Bounds for OWO Meta Learning Algorithms with Non-Lipschitz Functions

Theorem 2.4

Under the conditions of Theorem 2.1, for any task $t \in [T]$, *let* $\{\ell_{ti} : \Theta \mapsto [0, M]\}_{i\in[m]}$ *be a sequence of M-bounded functions. Let* $G^2 \geq \max_{t\in[T]} \chi^2(\rho_t^*||\rho_{t1})$, $D^2 \geq \max_{t\in[T]} \mathrm{KL}(\rho_t^*||\rho_{t1})$ *and* $\min_{t\in[T]} \mathrm{KL}(\rho_t^*||\rho_{t1}) \geq B^2 > 0$, *then using FTL algorithm to respectively learn the initialization and step size of EWA algorithm in Algorithm 2 attains the task-averaged regret upper bound:*

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{m}\langle \rho_{ti} - \rho_t^*, \ell_{ti}\rangle \leq \Big(\frac{D^7(\log T + 1)}{4TB^6} + \min\{\frac{G^2(\log T+1)}{VT}, 2G\sqrt{\frac{\log T + 1}{T}}\} + 2V\Big)\sqrt{m}M.$$

Transfer Risk Bounds for Batch Meta Learning

Theorem 3.1

Assume that for each task $t \in [T]$*, there exist* $G^2 \geq \max_{t\in[T]} \chi^2(\rho_t^*||\rho_{t1})$*, and* $\min_{t\in[T]} \mathrm{KL}(\rho_t^*||\rho_{t1}) \geq B^2 > 0$*. Assume that the novel task consists of loss functions* $\{\ell_i\}_{i\in[m]} \overset{i.i.d.}{\sim} \mu$*,* $\mu \overset{i.i.d.}{\sim} \tau$*, and for any optimal distribution* ρ^* *over task* μ*, there exists* $H > 0$ *such that* $\mathrm{KL}(\rho^*||\frac{1}{T}\sum_{t=1}^T \rho_{t1}) \leq H$*. Then we use* $(\frac{1}{T}\sum_{t=1}^T \rho_{t1}, \sqrt{\sum_{t=1}^T \mathrm{KL}(\rho_t^*||\rho_{t1})/(TmM^2)})$ *to run EWA algorithm for novel task* $\mu \sim \tau$ *with loss functions* $\{\ell_i\}_{i\in[m]}$ *to output probability distributions* $\{\rho_i\}_{i\in[m]}$*. Then let* $\bar{\rho} = \frac{1}{m}\sum_{i=1}^m \rho_i$*, for the optimal distribution* ρ^* *over task* μ *that does not dependent on* $\{\ell_i\}_{i\in[m]}$*, with probability* $1-\delta$ *over the draw of probability distributions* $\{\mu_t\}_{t=1}^T$*:*

$$\mathbb{E}_{\mu\sim\tau}\mathbb{E}_{\{\ell_i\}_{i=1}^m\sim\mu^m}\mathbb{E}_{\ell\sim\mu}\mathbb{E}_{\theta\sim\bar{\rho}}\ell(\theta) \leq \mathbb{E}_{\mu\sim\tau}\mathbb{E}_{\ell\sim\mu}\mathbb{E}_{\theta\sim\rho^*}\ell(\theta) + M\left(\sqrt{\frac{4G^2\log T}{Tm}} + \frac{3V^2}{\sqrt{m}B} + \frac{H}{B}\sqrt{\frac{\log 1/\delta}{2Tm}}\right).$$

PAC-Bayes Generalization Bounds for Statistical Multi-Task Learning

Proposition 6

Let $\bar{R}_{T,m} = \sum_{t,i=1}^{T,m} \langle \rho_{ti} - \mathcal{A}_t(S_t), \frac{\ell_{ti}}{T} \rangle$, *then* $\sum_{t=1}^{T} \overline{\mathrm{gen}}(\mathcal{A}_t, S_t) = \frac{\bar{R}_{T,m}}{m} - \frac{1}{T} \sum_{t=1}^{T} M_{\Pi_{tm}}$.

Use $\bar{U}_{T,m}$ to bound $\bar{R}_{T,m}$, and concentration inequality to bound $\sum_t M_{\Pi_{tm}}$, then

Theorem 3.2

Let $\mathcal{A}_t$ be the statistical learning algorithm for task $t \in [T]$. Assume that for all actions $\theta \in \Theta$, samples $z \in \mathcal{Z}$, $\ell(\theta, z) \in [0, M]$. Then, with probability at least $1 - \delta$ over the draw of $\{S_t\}_{t \in [T]}$, the multi-task generalization error of statistical learning algorithms $\{\mathcal{A}_t\}_{t \in [T]}$ satisfies:

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\theta \sim \mathcal{A}_t(S_t)} \left[\mathbb{E}_{z \sim \mu_t} \ell(\theta, z) - \frac{1}{m} \sum_{i=1}^{T} \ell(\theta, z_{ti}) \right] \leq \frac{\bar{U}_{T,m}}{m} + M \sqrt{\frac{2 \log \frac{1}{\delta}}{Tm}}.$$

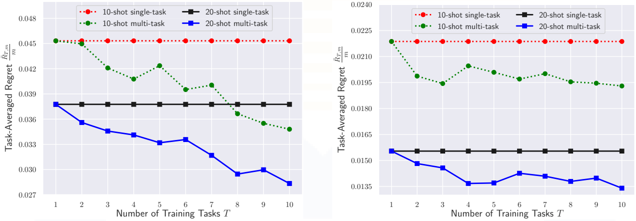

(a) Gaussian Mixture (b) Omniglot

Figure 1: Regret $\frac{\bar{R}_{T,m}}{m}$ with respect to the number T of training tasks.

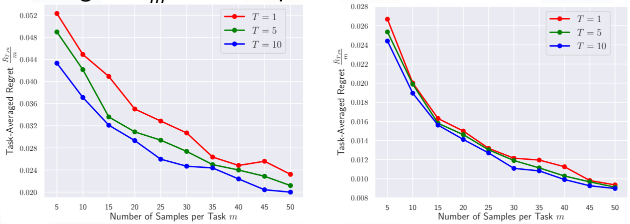

(a) Gaussian Mixture (b) Omniglot

Figure 2: Regret $\frac{\bar{R}_{T,m}}{m}$ with respect to the sample size m per training task.

Conclusions and Future Works

Our contributions for non-convex OWO meta learning are four-fold:

- We improve regret bound from $O((\frac{\sqrt{m}}{T^{1/4}} + \frac{(\log m)\log T}{\sqrt{T}} + V)\sqrt{m})$ to $O((\frac{1}{T^{1/2-\alpha}} + \frac{(\log T)^{9/2}}{T} + V)\sqrt{m})$ $(\alpha \in (0, 1/2))$ for bounded and piecewise Lipschitz functions.
- We design a new and efficient OWO meta learning algorithm of sharper regret bound $O((\frac{\log T}{T} + V)\sqrt{m})$ for bounded functions.
- We obtain a new transfer risk bound for statistical meta learning via regret analysis.
- We derive a PAC-Bayes bound for multi-task learning, shedding light on proving PAC-Bayes multi-task generalization bound with the regret bound from non-convex OWO meta learning.

Thanks!

References

[1] Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. "Adaptive Gradient-Based Meta-Learning Methods". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2019, pp. 5915–5926.

[2] Maria-Florina Balcan et al. "Learning-to-learn non-convex piecewise-Lipschitz functions". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021, pp. 15056–15069.

[3] Maria-Florina Balcan, Travis Dick, and Ellen Vitercik. "Dispersion for Data-Driven Algorithm Design, Online Learning, and Private Optimization". In: *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 2018, pp. 603–614.

[4] Pierre Alquier. "Non-Exponentially Weighted Aggregation: Regret Bounds for Unbounded Loss Functions". In: *International Conference on Machine Learning (ICML)*. 2021, pp. 207–218.