MULTITASK COOPERATIVE ONLINE LEARNING

Juliette Achddou

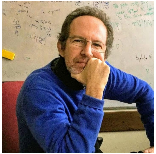

Nicolò Cesa-Bianchi

Pierre Laforgue

European Lighthouse on Secure and Safe AI

Funded by the European Union

Motivation and Objective

Motivation: many real-world applications (recommendation, health monitoring) require to provide a personalized service to multiple different clients.

Central idea: communicating information (models, privatized gradients) between agents in a decentralized fashion along the training process.

Informal objective: designing an algorithm with small multitask regret when neighbors in the communication network have similar tasks.

Relevant related works: cooperative online learning [Cesa-Bianchi et al., 2020], multitask online learning [Cesa-Bianchi et al., 2021], distributed online optimization [Hosseini et al., 2013].

Model

N **agents**, organized in communication network described by an undirected **graph** G. A hidden sequence of **convex loss functions** $\ell_1, \ell_2, \ldots$ chosen adversarially. For $t = 1, 2, \ldots$

1. agent $i_t \in [N]$ is activated
2. i_t may *fetch* information from its neighbors
3. i_t predicts $x_t \in \mathcal{X}$
4. i_t pays $\ell_t(x_t)$ and observes $g_t \in \partial \ell_t(x_t)$
5. i_t may *send* information to its neighbors

We aim at minimizing for any comparator $U \in \mathcal{U}$ and horizon T the **multitask regret**

$$R_T(U) = \sum_{i=1}^{N} \sum_{t\,:\,i_t=i} \Big(\ell_t(x_t) - \ell_t\big([U]_{i:}\big) \Big) .$$

Algorithm 2 `COOL-CN`

Requires: Base algorithm `AlgoClique`, right stochastic matrix W

for $t = 1, 2, \dots$ **do**

  Active agent i_t:

    fetches $\left[Y_t^{(j)}\right]_{i_t:}$ from each $j \in \mathcal{N}_{i_t}$

    predicts $x_t = \sum_{j \in \mathcal{N}_{i_t}} w_{i_t j} \left[Y_t^{(j)}\right]_{i_t:}$

    pays $\ell_t(x_t)$ and observes $g_t \in \partial \ell_t(x_t)$

    sends $\left(i_t, w_{i_t j}\, g_t\right)$ to each neighbor $j \in \mathcal{N}_{i_t}$

  for $j \in \mathcal{N}_{i_t}$ **do**

    Agent j feeds the linear loss $\langle w_{i_t j}\, g_t, \cdot \rangle$ to their local instance of `AlgoClique`, and obtains $Y_{t+1}^{(j)}$

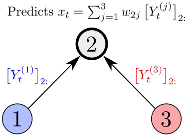

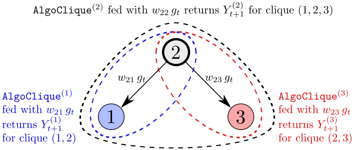

Lemma

The regret of `COOL-CN` *satisfies*

$$R_T(U) \leq \sum_{j=1}^{N} R_T^{\text{clique-}j}\left(U^{(j)}\right) .$$

where $R_T^{\text{clique-}j}$ *is the regret suffered by* `AlgoClique` *on the linear losses* $\langle w_{i_t j}\, g_t, \cdot \rangle$ *over the rounds* $t \leq T$ *such that* $i_t \in \mathcal{N}_j$, *and* $U^{(j)}$ *contains the rows* $U_{i:}$ *for* $i \in \mathcal{N}_j$.

What choice for `Algoclique`?

A valid choice of `AlgoClique` is `MT-FTRL` Cesa-Bianchi et al. [2021], in its variance adaptive version. It is an algorithm designed for the case without communication constraints (i.e., G is a clique), which satisfies

$$R_T^{\text{clique-}j}(U^{(j)}) = \tilde{\mathcal{O}}\left(\sqrt{1+\sigma_j^2(N_j-1)}\sqrt{T}\right) ,$$

where $\sigma_j^2 = \frac{1}{2N_j(N_j-1)}\sum_{i,i'\in\mathcal{N}_j}\left\|[U]_{i:} - U_{i':}\right\|_2^2$ is a measure of the local variance.

Upper bounds for adversarial activations.

Theorem

For general weights w_{ij} we have

$$R_T(U) \overset{\widetilde{\mathcal{O}}}{=} \sum_{j=1}^{N} \max_{i \in \mathcal{N}_j} w_{ij} \sqrt{1 + \sigma_j^2 (N_j - 1)} \sqrt{\sum_{i \in \mathcal{N}_j} T_i}\,.$$

Setting $w_{ij} = \mathbb{I}\left\{j \in \mathcal{N}_i\right\} / N_i$, it becomes

$$R_T(U) \overset{\widetilde{\mathcal{O}}}{=} \sqrt{1 + \sigma_{\max}^2 (N_{\max} - 1)} \sqrt{\frac{N N_{\max} T}{N_{\min}^2}}\,,$$

that particularizes well, e.g., for ρ-regular graphs

$$R_T(U) \overset{\widetilde{\mathcal{O}}}{=} \sqrt{1 + \rho \sigma_{\max}^2} \sqrt{\frac{NT}{\rho + 1}}\,.$$

Upper bounds for stochastic activations.

Theorem

Using appropriate w_{ij}, with $\alpha(G)$ the independence number of G, we have

$$\mathbb{E}[R_T(U)] \overset{\widetilde{\mathcal{O}}}{=} \sqrt{1 + \sigma^2_{\max}(N_{\max} - 1)}\sqrt{\alpha(G)\, T}\,.$$

Theorem

There exists a sequence of activations and gradients such that for any algorithm

$$R_T \geq \frac{1}{3} \max\left(\sqrt{1+\sigma^2(N-1)}, \sqrt{\alpha_2(G)}\right)\sqrt{T}\,,$$

where $\alpha_2(G)$ *is the double independence number.*

Theorem

For any even number ρ, *there exists a* ρ*-regular graph and a sequence of activations and gradients such that for any algorithm*

$$R_T \geq \frac{1}{5}\sqrt{1+\rho\,\sigma^2_{\min}}\sqrt{\frac{NT}{\rho}}\,.$$

With `AlgoClique` set as `MT-FTRL`, it is possible to make the algorithm **ϵ-differentially private**, while preserving regret guarantees.

Notion of ϵ- DP:

$$\frac{P[m_1^i, \dots m_T^i | \ell_1, i_1 \dots \ell_T, i_T]}{P[m_1^i, \dots m_T^i | \ell_1', i_1 \dots \ell_T', i_T]} \leq e^{\epsilon}, \ \forall i$$

where m_t^i denotes the message sent by i to the network at time t, and $\ell_{1 \dots T}$ and $\ell'_{1 \dots T}$ **differ by one element**.

We exhibit a **cut-off privacy value** ϵ, below which communicating is useless.

Possible future works: a unifying framework for ST-COOL, MT-COOL, DOL

A unifying framework for MT-COOL, ST-COOL, DOL

	Regret	Setting
MT-COOL	$R_T(U) = \sum_{i=1}^{N} \sum_{t\,:\,i_t=i} \left(\ell_t(x_t^{(i)}) - \ell_t([U]_{i:})\right)$	▶ One activation at a time. ▶ Communication of · models · gradients
ST-COOL	$R_T(u) = \sum_{i=1}^{N} \sum_{t\,:\,i_t=i} \left(\ell_t(x_t^{(i)}) - \ell_t(u)\right)$	▶ One or more activations at a time. ▶ Communication of gradients
Decentralized OL	$R_T(u) = \sum_{i=1}^{N} \sum_{t=1}^{T} \left(\ell_{i,t}(x_t) - \ell_{i,t}(u)\right)$	▶ All agents are activated. ▶ Communication of models

- **Considering communication delays.**
  Agents can send gradients received from their neighbors to their other neighbors, but this induces a delay in the information on agent *j* that agent *i* holds.

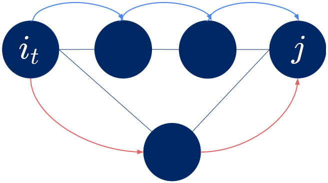

  This delay has a complicated form !
- **Considering communication constraints.**
  - Agents can only send messages every *K* rounds.
  - Alternatively: agents have a total communication budget
  - What impact on the regret?

Possible future works: different privacy notions (3)

Considering different forms of differential privacy.

- The common differential privacy measure in online learning is with respect to one individual loss function (equivalent of one data point).
- In stochastic activations, every agent will be activated $q_i T$ times, so requiring ”user-level” privacy will harm the regret.
- Is there an intermediary yet reachable objective ?

References

N. Cesa-Bianchi, T. Cesari, and C. Monteleoni. Cooperative online learning: Keeping your neighbors updated. In *Algorithmic learning theory*. PMLR, 2020.

N. Cesa-Bianchi, P. Laforgue, A. Paudice, and M. Pontil. Multitask online mirror descent. *Transactions on Machine Learning Research*, 2021.

S. Hosseini, A. Chapman, and M. Mesbahi. Online distributed optimization via dual averaging. In *IEEE Conference on Decision and Control*. IEEE, 2013.