Parameter-Free Online Linear Optimization with Side Information via Universal Coin Betting

Jongha (Jon) Ryu[1], Alankrita Bhatt[1], and Young-Han Kim[1,2]

[1]University of California, San Diego
[2]Gauss Labs, Inc

AISTATS 2022

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$
  - Receives a vector $\mathbf{g}_t \in V$ s.t. $\|\mathbf{g}_t\| \leq L$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$
  - Receives a vector $\mathbf{g}_t \in V$ s.t. $\|\mathbf{g}_t\| \le L$
  - Gains reward $\langle \mathbf{g}_t, \mathbf{w}_t \rangle$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$
  - Receives a vector $\mathbf{g}_t \in V$ s.t. $\|\mathbf{g}_t\| \le L$
  - Gains reward $\langle \mathbf{g}_t, \mathbf{w}_t \rangle$
- **Goal**: maximize the cumulative reward $\sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{w}_t \rangle$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$
  - Receives a vector $\mathbf{g}_t \in V$ s.t. $\|\mathbf{g}_t\| \le L$
  - Gains reward $\langle \mathbf{g}_t, \mathbf{w}_t \rangle$
- **Goal**: maximize the cumulative reward $\sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{w}_t \rangle$
- The standard performance measure: regret with respect to the best static competitor in hindsight

$$\mathrm{Reg}(\mathbf{u}; \mathbf{g}^T) := \sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{u} \rangle - \sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{w}_t \rangle \text{ for } \mathbf{u} \in V$$

Problem Setting: Online Linear Optimization (OLO)

- V: Hilbert space with norm $\|\cdot\|$
- In each round $t = 1, 2, \ldots$
  - Learner picks $\mathbf{w}_t \in V$
  - Receives a vector $\mathbf{g}_t \in V$ s.t. $\|\mathbf{g}_t\| \le L$
  - Gains reward $\langle \mathbf{g}_t, \mathbf{w}_t \rangle$
- **Goal**: maximize the cumulative reward $\sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{w}_t \rangle$
- The standard performance measure: regret with respect to the best static competitor in hindsight

$$\mathrm{Reg}(\mathbf{u}; \mathbf{g}^T) := \sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{u} \rangle - \sum_{t=1}^{T} \langle \mathbf{g}_t, \mathbf{w}_t \rangle \text{ for } \mathbf{u} \in V$$

- **Issues**: (1) learning rate tuning; (2) static competitors are weak

(1) Parameter-Free OLO via Universal Gambling

- To attain optimal rates, naive approaches require $\|\mathbf{u}\|$

(1) Parameter-Free OLO via Universal Gambling

- To attain optimal rates, naive approaches require $\|\mathbf{u}\|$
- **Q.** Can we attain optimal regret w/o such a priori knowledge?

(1) Parameter-Free OLO via Universal Gambling

- To attain optimal rates, naive approaches require $\|\mathbf{u}\|$
- **Q.** Can we attain optimal regret w/o such a priori knowledge?
- **A.** Orabona and Pál [2016] showed that universal coin betting algorithm can be converted to a near-optimal-regret parameter-free OLO algorithm!

(1) Parameter-Free OLO via Universal Gambling

- To attain optimal rates, naive approaches require $\|\mathbf{u}\|$
- **Q.** Can we attain optimal regret w/o such a priori knowledge?
- **A.** Orabona and Pál [2016] showed that universal coin betting algorithm can be converted to a near-optimal-regret parameter-free OLO algorithm!
- Key tool: Fenchel duality

(1) Parameter-Free OLO via Universal Gambling

- To attain optimal rates, naive approaches require $\|\mathbf{u}\|$
- **Q.** Can we attain optimal regret w/o such a priori knowledge?
- **A.** Orabona and Pál [2016] showed that universal coin betting algorithm can be converted to a near-optimal-regret parameter-free OLO algorithm!
- Key tool: Fenchel duality
- **Note**: Other parameter-free algorithms exist

(2) OLO with Side Information

- Static competitors $\{\mathbf{u} \colon \mathbf{u} \in V\}$ are weak

(2) OLO with Side Information

- Static competitors $\{\mathbf{u}\colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero

(2) OLO with Side Information

- Static competitors $\{\mathbf{u} \colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero
  - In general, $\langle \sum_{t=1}^{T} \mathbf{g}_t, \mathbf{u} \rangle$ can be large iff $\| \sum_{t=1}^{T} \mathbf{g}_t \|$ is large

(2) OLO with Side Information

- Static competitors $\{\mathbf{u}\colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero
  - In general, $\langle \sum_{t=1}^{T} \mathbf{g}_t, \mathbf{u} \rangle$ can be large iff $\| \sum_{t=1}^{T} \mathbf{g}_t \|$ is large
- **Q.** Can we leverage a possible structure in $(\mathbf{g}_t)_{t \geq 1}$?

(2) OLO with Side Information

- Static competitors $\{\mathbf{u} \colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero
  - In general, $\langle \sum_{t=1}^T \mathbf{g}_t, \mathbf{u} \rangle$ can be large iff $\| \sum_{t=1}^T \mathbf{g}_t \|$ is large
- **Q.** Can we leverage a possible structure in $(\mathbf{g}_t)_{t \geq 1}$?
- **Our approach**: Assuming that we have access to a side information sequence $(h_t)_{t \geq 1}$ which may potentially capture a structure, develop a method that adapts to side information!

(2) OLO with Side Information

- Static competitors $\{\mathbf{u} \colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero
  - In general, $\langle \sum_{t=1}^{T} \mathbf{g}_t, \mathbf{u} \rangle$ can be large iff $\| \sum_{t=1}^{T} \mathbf{g}_t \|$ is large
- **Q.** Can we leverage a possible structure in $(\mathbf{g}_t)_{t \geq 1}$?
- **Our approach**: Assuming that we have access to a side information sequence $(h_t)_{t \geq 1}$ which may potentially capture a structure, develop a method that adapts to side information!
- **Example**: $h_t = \operatorname{sgn}(\langle \mathbf{g}_{t-1}, \mathbf{f} \rangle)$

(2) OLO with Side Information

- Static competitors $\{\mathbf{u} \colon \mathbf{u} \in V\}$ are weak
  - **Example**: for $\mathbf{g}, -\mathbf{g}, \mathbf{g}, -\mathbf{g}, \ldots$, the best reward with $\mathbf{u} \in V$ is zero
  - In general, $\langle \sum_{t=1}^{T} \mathbf{g}_t, \mathbf{u} \rangle$ can be large iff $\| \sum_{t=1}^{T} \mathbf{g}_t \|$ is large
- **Q.** Can we leverage a possible structure in $(\mathbf{g}_t)_{t \geq 1}$?
- **Our approach**: Assuming that we have access to a side information sequence $(h_t)_{t \geq 1}$ which may potentially capture a structure, develop a method that adapts to side information!
- **Example**: $h_t = \mathrm{sgn}(\langle \mathbf{g}_{t-1}, \mathbf{f} \rangle)$
- To capture a more complex structure, we consider tree-structured side information sequences

This Paper

- Generalize (1) the parameter-free OLO algorithm of Orabona and Pál [2016] for (2) OLO with (tree) side information

This Paper

- Generalize (1) the parameter-free OLO algorithm of Orabona and Pál [2016] for (2) OLO with (tree) side information
- **Our contribution**: develop an OLO algorithm that efficiently adapts to tree-structured side information sequences via universal gambling

This Paper

- Generalize (1) the parameter-free OLO algorithm of Orabona and Pál [2016] for (2) OLO with (tree) side information
- **Our contribution**: develop an OLO algorithm that efficiently adapts to tree-structured side information sequences via universal gambling
- **Idea**: convert the existing context-tree weighting method [Willems et al., 1995] from universal compression to an OLO algorithm, via the duality lens established by Orabona and Pál [2016]

This Paper

- Generalize (1) the parameter-free OLO algorithm of Orabona and Pál [2016] for (2) OLO with (tree) side information
- **Our contribution**: develop an OLO algorithm that efficiently adapts to tree-structured side information sequences via universal gambling
- **Idea**: convert the existing context-tree weighting method [Willems et al., 1995] from universal compression to an OLO algorithm, via the duality lens established by Orabona and Pál [2016]
- **Experiment**: online linear regression with absolute loss

Please Visit Our Poster!

- Poster session #5
- **Date/Time (PDT)**: Mar 30, 2022 (Wed) / 8:30am – 10:00am
- **Date/Time (AoE)**: Mar 30, 2022 (Wed) / 3:30am – 5:00am

References I

Francesco Orabona and Dávid Pál. Coin betting and parameter-free online learning. In *Adv. Neural Inf. Proc. Syst.*, volume 29. Curran Associates, Inc., 2016.

Frans MJ Willems, Yuri M Shtarkov, and Tjalling J Tjalkens. The context-tree weighting method: Basic properties. *IEEE Trans. Inf. Theory*, 41(3):653–664, 1995.