Luckiness in Multiscale Online Learning

Muriel F. Pérez-Ortiz and Wouter Koolen

Centrum Wiskunde & Informatica
NeurIPS 2022

October 5, 2022

Full-information online learning

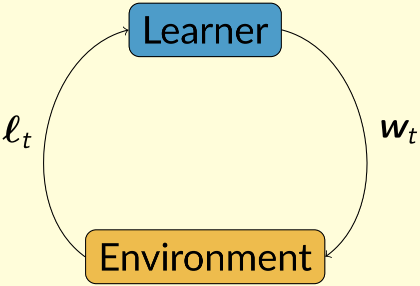

Full-information online learning

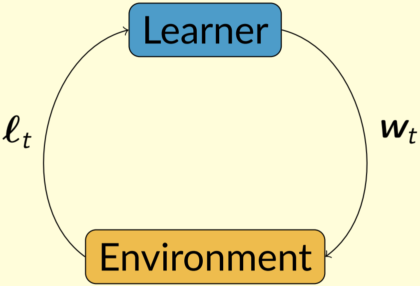

$$R_t = \sum_{s \leq t} \left\{ \langle \boldsymbol{w}_s,\ \boldsymbol{\ell}_s \rangle - \ell_{s,k^*} \right\}$$

Full-information online learning

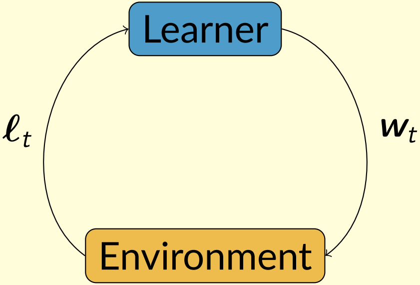

$$R_t = \sum_{s \le t} \{ \langle \boldsymbol{w}_s, \ \boldsymbol{\ell}_s \rangle - \ell_{s,k^*} \}$$

$$|\ell_k| \le \sigma$$

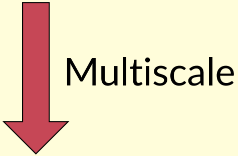

$$|\ell_k| \le \sigma_k$$

Regret Bounds of Existing Algorithms

	Worst case	Margin condition
Hedge	$\sigma_{\max}\sqrt{t \ln K}$	—

Regret Bounds of Existing Algorithms

	Worst case	Margin condition
Hedge	$\sigma_{\max}\sqrt{t \ln K}$	—
Modern alg.	$\sigma_{\max}\sqrt{t \ln K}$	$O(1)$

Regret Bounds of Existing Algorithms

	Worst case	Margin condition
Hedge	$\sigma_{\max}\sqrt{t \ln K}$	—
Modern alg.	$\sigma_{\max}\sqrt{t \ln K}$	$O(1)$
Mulstiscale alg.	$\sigma_{k^*}\sqrt{t \ln K}$	—

Regret Bounds of Existing Algorithms

	Worst case	Margin condition
Hedge	$\sigma_{\max}\sqrt{t \ln K}$	–
Modern alg.	$\sigma_{\max}\sqrt{t \ln K}$	$O(1)$
Mulstiscale alg.	$\sigma_{k^*}\sqrt{t \ln K}$	–
This work	$\sigma_{k^*}\sqrt{t \ln K}$	$O(1)$

MUSCADA: Multiscale-Adaptive Algorithm

MUSCADA: Multiscale-Adaptive Algorithm

- Regret:

$$R_t \lesssim \sigma_{k^*} \sqrt{v_t \ln(1/\pi_k)}. \quad v_t \lesssim t \; ; \quad v_t \lesssim 1 \quad \text{if lucky}$$

MUSCADA: Multiscale-Adaptive Algorithm

- Regret:

$$R_t \lesssim \sigma_{k^*} \sqrt{v_t \ln(1/\pi_k)}. \quad v_t \lesssim t \; ; \quad v_t \lesssim 1 \quad \text{if lucky}$$

- Computation

$O(K)$ operations per round

Application: two-player zero-sum games

$$\langle \boldsymbol{p}^\star,\ A\boldsymbol{q}^\star \rangle = \min_{\boldsymbol{p}} \max_{\boldsymbol{q}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle = \max_{\boldsymbol{q}} \min_{\boldsymbol{p}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle$$

Application: two-player zero-sum games

$$\langle \boldsymbol{p}^\star,\ A\boldsymbol{q}^\star \rangle = \min_{\boldsymbol{p}} \max_{\boldsymbol{q}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle = \max_{\boldsymbol{q}} \min_{\boldsymbol{p}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle$$

$$A = \begin{pmatrix} 0.1 & -1 \\ 1 & 1000 \end{pmatrix}$$

Application: two-player zero-sum games

$$\langle \boldsymbol{p}^\star,\ A\boldsymbol{q}^\star \rangle = \min_{\boldsymbol{p}} \max_{\boldsymbol{q}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle = \max_{\boldsymbol{q}} \min_{\boldsymbol{p}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle$$

$$\frac{\sigma_{\max}}{t}$$

$$A = \begin{pmatrix} 0.1 & -1 \\ 1 & 1000 \end{pmatrix}$$

Application: two-player zero-sum games

$$\langle \boldsymbol{p}^\star,\ A\boldsymbol{q}^\star \rangle = \min_{\boldsymbol{p}} \max_{\boldsymbol{q}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle = \max_{\boldsymbol{q}} \min_{\boldsymbol{p}} \langle \boldsymbol{p},\ A\boldsymbol{q} \rangle$$

$$A = \begin{pmatrix} 0.1 & -1 \\ 1 & 1000 \end{pmatrix}$$

$$\frac{\sigma_{\max}}{t}$$

↓ Mulstiscale

$$\frac{\sigma_{\text{saddle}-\text{point}}}{t}$$

Conclusion

	Worst case	Margin condition
MUSCADA	$\sigma_{k^*}\sqrt{t \ln K}$	$O(1)$

Conclusion

	Worst case	Margin condition
MUSCADA	$\sigma_{k^*}\sqrt{t \ln K}$	$O(1)$

Thanks!