Minimax Optimal Quantile and Semi-Adversarial Regret via Root-Logarithmic Regularizers

Jeffrey Negrea*, Blair Bilodeau*, Nicolò Campolongo, Francesco Orabona, Daniel M. Roy

December 2021

Advances in Neural Information Processing Systems 35

*Equal contribution

Sequential Prediction with Expert Advice

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ
- Excess cumulative loss of the player against an expert distribution in hindsight

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ
- Excess cumulative loss of the player against an expert distribution in hindsight

$$\textbf{Regret:} \quad R_T(q) = \sum_{t=1}^{T} [\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t) - \mathbb{E}_{\theta \sim q}\, \ell_t(\theta)]$$

Familiar Setting

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ
- Excess cumulative loss of the player against an expert distribution in hindsight

$$\textbf{Regret:} \quad R_T(q) = \sum_{t=1}^{T} [\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t) - \mathbb{E}_{\theta \sim q}\, \ell_t(\theta)]$$

Familiar Setting

- The finite expert setting is when $|\Theta| = N < \infty$

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ
- Excess cumulative loss of the player against an expert distribution in hindsight

$$\textbf{Regret:} \quad R_T(q) = \sum_{t=1}^{T} \left[\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t) - \mathbb{E}_{\theta \sim q} \, \ell_t(\theta)\right]$$

Familiar Setting

- The finite expert setting is when $|\Theta| = N < \infty$
- Compete against q_T that is a point-mass on the best expert in hindsight

Sequential Prediction with Expert Advice

Let Θ be a generic space of *experts* (e.g., parameters or models).

For rounds $t = 1, \ldots, T$:

- Player selects a *distribution* over experts π_t
- Observe losses $\ell_t(\theta) \in [0, 1]$ for all $\theta \in \Theta$, generated by the environment
- Incur loss $\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t)$

Sequential prediction relies on a measure of the player's performance that is...

- Relative to the class of experts Θ
- Excess cumulative loss of the player against an expert distribution in hindsight

$$\textbf{Regret:} \quad R_T(q) = \sum_{t=1}^{T} \left[\mathbb{E}_{\theta_t \sim \pi_t} \ell_t(\theta_t) - \mathbb{E}_{\theta \sim q}\, \ell_t(\theta)\right]$$

Familiar Setting

- The finite expert setting is when $|\Theta| = N < \infty$
- Compete against q_T that is a point-mass on the best expert in hindsight

Primary Contributions

Primary Contributions

Generic FTRL Analysis

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Semi-Adversarial Regret Guarantees

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Semi-Adversarial Regret Guarantees

- Loss distribution is constrained in a structured way (introduced by [BNR20]).

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Semi-Adversarial Regret Guarantees

- Loss distribution is constrained in a structured way (introduced by [BNR20]).
- We achieve improved regret bounds for short time horizons.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Semi-Adversarial Regret Guarantees

- Loss distribution is constrained in a structured way (introduced by [BNR20]).
- We achieve improved regret bounds for short time horizons.
- We use root-logarithmic FTRL similarly to our quantile regret algorithm.

Primary Contributions

Generic FTRL Analysis

- We provide a novel local-norm analysis of follow-the-regularized-leader on general (potentially uncountable) expert spaces.
- Our result prescribes how to choose the regularizer for desired regret bounds, and provides an explicit expression for the FTRL output (i.e., π).

Quantile Regret Guarantees

- Compete against only the top quantile of experts (introduced by [CFH09]).
- We provide the first FTRL algorithm that achieves "root-KL" quantile regret.
- We provide novel lower bounds that show the root-KL guarantees are tight.
- A key application of this novel bound for uncountable expert spaces is we can compete against the terminal Bayesian posterior.

Semi-Adversarial Regret Guarantees

- Loss distribution is constrained in a structured way (introduced by [BNR20]).
- We achieve improved regret bounds for short time horizons.
- We use root-logarithmic FTRL similarly to our quantile regret algorithm.

FTRL with Linearly Decomposable Regularizers

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

Hedge optimizes a regularized objective...with any prior ν we like!

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} \mathrm{KL}\left(\pi \,\|\, \nu\right) \right\}.$$

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

Hedge optimizes a regularized objective...with any prior ν we like!

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta\sim\pi} L_t(\theta) + \eta_t^{-1} \mathrm{KL}\left(\pi \| \nu\right) \right\}.$$

Obvious (but not obviously useful): we replace KL with something more generic.

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta\sim\pi} L_t(\theta) + \eta_t^{-1} \int \left(f \circ \frac{\mathrm{d}\pi}{\mathrm{d}\nu} \right) \mathrm{d}\nu \right\}.$$

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

Hedge optimizes a regularized objective...with any prior ν we like!

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} \mathrm{KL}\left(\pi \| \nu \right) \right\}.$$

Obvious (but not obviously useful): we replace KL with something more generic.

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu \right) \right\}.$$

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

Hedge optimizes a regularized objective...with any prior ν we like!

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} \mathrm{KL}\left(\pi \| \nu\right) \right\}.$$

Obvious (but not obviously useful): we replace KL with something more generic.

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

KL corresponds to $f(x) = x \log x$ and χ^2 corresponds to $f(x) = x^2 - 1$ [Alq21].

FTRL with Linearly Decomposable Regularizers

Familiar: Hedge Algorithm for Uncountable Experts (Bayesian Inference)

Hedge optimizes a regularized objective...with any prior ν we like!

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} \mathrm{KL}\left(\pi \| \nu\right) \right\}.$$

Obvious (but not obviously useful): we replace KL with something more generic.

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

KL corresponds to $f(x) = x \log x$ and χ^2 corresponds to $f(x) = x^2 - 1$ [Alq21].

In this generality we...

- Provide a form for π_t with a generic f.
- Provide a novel local-norm analysis of the regret of playing π_t.
- Use these results to prove new guarantees for quantile regret, semi-adversarial regret, and other applications.

Novel Quantile Regret Guarantees

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \;\forall \epsilon, \text{ or } \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \;\forall q.$$

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \; \forall \epsilon, \text{ or } \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \; \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \ \forall \epsilon, \quad \text{or} \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \ \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Theorem [NEW] (abNormal Root-KL Bound using FTRL)

Using $f(x) = \int_1^x \sqrt{2 \log(1+s)} \mathrm{d}s$ and any ν, for all $q \ll \nu$

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \mathrm{KL}(q \| \nu))} + \sqrt{8T}.$$

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)}\ \forall \epsilon, \text{ or } \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))}\ \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Theorem [NEW] (abNormal Root-KL Bound using FTRL)

Using $f(x) = \int_1^x \sqrt{2 \log(1+s)}\mathrm{d}s$ and any ν, for all $q \ll \nu$

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \mathrm{KL}(q \| \nu))} + \sqrt{8T}\,.$$

Taking ν to be uniform over all experts and q to be uniform over the top $\lfloor \varepsilon N \rfloor$ recovers the $\sqrt{T \log(1/\varepsilon)}$ bound for finite experts,

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \; \forall \epsilon, \text{ or } \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \; \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Theorem [NEW] (abNormal Root-KL Bound using FTRL)

Using $f(x) = \int_1^x \sqrt{2 \log(1+s)} \mathrm{d}s$ and any ν, for all $q \ll \nu$

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \mathrm{KL}(q \| \nu))} + \sqrt{8T} .$$

Taking ν to be uniform over all experts and q to be uniform over the top $\lfloor \varepsilon N \rfloor$ recovers the $\sqrt{T \log(1/\varepsilon)}$ bound for finite experts,

. . . and we've significantly generalized to uncountable experts.

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \; \forall \epsilon, \quad \text{or} \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \; \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Theorem [NEW] (abNormal Root-KL Bound using FTRL)

Using $f(x) = \int_1^x \sqrt{2 \log(1+s)} \mathrm{d}s$ and any ν, for all $q \ll \nu$

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \mathrm{KL}(q \| \nu))} + \sqrt{8T}\,.$$

Taking ν to be uniform over all experts and q to be uniform over the top $\lfloor \varepsilon N \rfloor$ recovers the $\sqrt{T \log(1/\varepsilon)}$ bound for finite experts,

...and we've significantly generalized to uncountable experts.

Theorem [NEW] (Quantile Lower Bound)

For all N there is a loss distribution s.t. for all $\varepsilon < 1/4$ and T and any algorithm

$$\mathbb{E} R_T(q_\varepsilon) \geq \sqrt{(T/2)(\log(1/\varepsilon) - 2\log 2)} - 2\log N - 2\,.$$

Novel Quantile Regret Guarantees

Existing Upper Bounds

For N experts, if q_ε is a point-mass on the (unknown) $\lfloor \varepsilon N \rfloor$ best expert, then non-FTRL algorithms [CFH09; CV10; OP16] achieve

$$R_T(q_\varepsilon) \lesssim \sqrt{T \log(1/\varepsilon)} \ \forall \epsilon, \ \text{or} \quad R_T(q) \lesssim \sqrt{T(1 + \mathrm{KL}(q \| \nu))} \ \forall q.$$

Hedge can only achieve this with knowledge of ε or $\mathrm{KL}(q \| \nu)$ in advance.

Theorem [NEW] (abNormal Root-KL Bound using FTRL)

Using $f(x) = \int_1^x \sqrt{2 \log(1 + s)} \mathrm{d}s$ and any ν, for all $q \ll \nu$

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \mathrm{KL}(q \| \nu))} + \sqrt{8T}\,.$$

Taking ν to be uniform over all experts and q to be uniform over the top $\lfloor \varepsilon N \rfloor$ recovers the $\sqrt{T \log(1/\varepsilon)}$ bound for finite experts,

...and we've significantly generalized to uncountable experts.

Theorem [NEW] (Quantile Lower Bound)

For all N there is a loss distribution s.t. for all $\varepsilon < 1/4$ and T and any algorithm

$$\mathbb{E} R_T(q_\varepsilon) \geq \sqrt{(T/2)(\log(1/\varepsilon) - 2\log 2)} - 2 \log N - 2\,.$$

Examples and Applications

Examples and Applications

Recovering Predictive Variance Bound for Hedge

$$f(x) = x \log x \quad \Longrightarrow \quad R_T(q) \lesssim \mathrm{KL}\left(q \| \nu\right) \sqrt{\textstyle\sum_{t=1}^{T} \text{Predictive Loss Variance}} \quad \forall q$$

Novel Prior Variance Bound for χ^2

$$f(x) = x^2 - 1 \quad \Longrightarrow \quad R_T(q) \lesssim \chi^2(q) \sqrt{\textstyle\sum_{t=1}^{T} \text{Prior Loss Variance}} \quad \forall q$$

Examples and Applications

Recovering Predictive Variance Bound for Hedge

$$f(x) = x \log x \quad \Longrightarrow \quad R_T(q) \lesssim \mathrm{KL}\left(q \| \nu\right) \sqrt{\sum_{t=1}^{T} \text{Predictive Loss Variance}} \quad \forall q$$

Novel Prior Variance Bound for χ^2

$$f(x) = x^2 - 1 \quad \Longrightarrow \quad R_T(q) \lesssim \chi^2(q) \sqrt{\sum_{t=1}^{T} \text{Prior Loss Variance}} \quad \forall q$$

Novel Regret against the Terminal Posterior

$$f(x) \approx x\sqrt{\log(1+x)} \quad \Longrightarrow \quad R_T(\hat{\pi}_T) \lesssim \sqrt{T \, \mathrm{KL}\left(\hat{\pi}_T \| \nu\right)}$$

Examples and Applications

Recovering Predictive Variance Bound for Hedge

$$f(x) = x \log x \quad \Longrightarrow \quad R_T(q) \lesssim \mathrm{KL}\,(q \| \nu) \sqrt{\textstyle\sum_{t=1}^T \text{Predictive Loss Variance}} \quad \forall q$$

Novel Prior Variance Bound for χ^2

$$f(x) = x^2 - 1 \quad \Longrightarrow \quad R_T(q) \lesssim \chi^2(q) \sqrt{\textstyle\sum_{t=1}^T \text{Prior Loss Variance}} \quad \forall q$$

Novel Regret against the Terminal Posterior

$$f(x) \approx x\sqrt{\log(1+x)} \quad \Longrightarrow \quad R_T(\hat{\pi}_T) \lesssim \sqrt{T\,\mathrm{KL}\,(\hat{\pi}_T \| \nu)}$$

Novel Model Selection Regret for Infinite Classes

Known countable union of model classes $\Theta = \cup_{m \geq 1} \Theta_m$, prior $\nu(\theta) \propto [m^2 \, |\Theta_m|]^{-1}$:

$$f(x) \approx x\sqrt{\log(1+x)} \quad \Longrightarrow \quad R_T(\theta) \lesssim \sqrt{T(\log|\Theta_m| + \log m)} \quad \forall m, \theta \in \Theta_m$$

Examples and Applications

Recovering Predictive Variance Bound for Hedge

$$f(x) = x \log x \quad \Longrightarrow \quad R_T(q) \lesssim \mathrm{KL}\left(q \,\|\, \nu\right) \sqrt{\textstyle\sum_{t=1}^{T} \text{Predictive Loss Variance}} \quad \forall q$$

Novel Prior Variance Bound for χ^2

$$f(x) = x^2 - 1 \quad \Longrightarrow \quad R_T(q) \lesssim \chi^2(q) \sqrt{\textstyle\sum_{t=1}^{T} \text{Prior Loss Variance}} \quad \forall q$$

Novel Regret against the Terminal Posterior

$$f(x) \approx x\sqrt{\log(1+x)} \quad \Longrightarrow \quad R_T(\hat{\pi}_T) \lesssim \sqrt{T\,\mathrm{KL}\left(\hat{\pi}_T \,\|\, \nu\right)}$$

Novel Model Selection Regret for Infinite Classes

Known countable union of model classes $\Theta = \cup_{m \geq 1} \Theta_m$, prior $\nu(\theta) \propto [m^2 \, |\Theta_m|]^{-1}$:

$$f(x) \approx x\sqrt{\log(1+x)} \quad \Longrightarrow \quad R_T(\theta) \lesssim \sqrt{T(\log|\Theta_m| + \log m)} \quad \forall m, \theta \in \Theta_m$$

Semi-Adversarial Paradigm

Real data $\neq$ stochastic

Semi-Adversarial Paradigm

Real data $\neq$ stochastic$\leftarrow$Optimistic

Semi-Adversarial Paradigm

Real data $\neq$ stochastic$\leftarrow$Optimistic

Real data $\neq$ adversarial

Semi-Adversarial Paradigm

Real data $\neq$ stochastic$\leftarrow$Optimistic

Real data $\neq$ adversarial$\leftarrow$Pessimistic

Semi-Adversarial Paradigm

Real data $\not\equiv$ stochastic$\leftarrow$Optimistic

Real data $\not\equiv$ adversarial$\leftarrow$Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

Semi-Adversarial Paradigm

Real data $\not\equiv$ stochastic$\leftarrow$Optimistic

Real data $\not\equiv$ adversarial$\leftarrow$Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

Intuitively, fix a "neighbourhood" of distributions $\mathcal{D} \subseteq \mathcal{M}([0,1]^N)$.

Semi-Adversarial Paradigm

Real data $\not\equiv$ stochastic $\leftarrow$ Optimistic

Real data $\not\equiv$ adversarial $\leftarrow$ Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

- Intuitively, fix a "neighbourhood" of distributions $\mathcal{D} \subseteq \mathcal{M}([0,1]^N)$.
- Each data point drawn from an arbitrary distribution in "neighbourhood".

Semi-Adversarial Paradigm

Real data $\neq$ stochastic$\leftarrow$Optimistic   Real data $\neq$ adversarial$\leftarrow$Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

Intuitively, fix a "neighbourhood" of distributions $\mathcal{D} \subseteq \mathcal{M}([0,1]^N)$.

Each data point drawn from an arbitrary distribution in "neighbourhood".

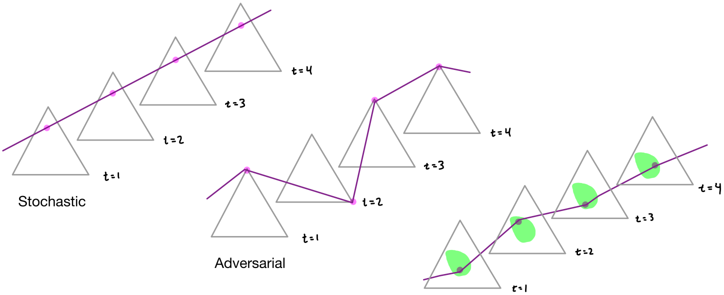

Semi-Adversarial Paradigm

Real data $\not\equiv$ stochastic←Optimistic   Real data $\not\equiv$ adversarial←Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

Intuitively, fix a "neighbourhood" of distributions $\mathcal{D} \subseteq \mathcal{M}([0,1]^N)$.

Each data point drawn from an arbitrary distribution in "neighbourhood".

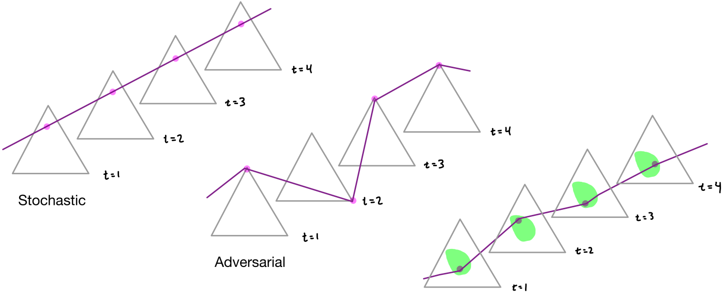

$\mathcal{I}_0 = \{$experts that are optimal for some $\mu \in \mathcal{D}\}$ $N_0 = |\mathcal{I}_0|$

$\Delta_0 = \inf_{\mu \in \mathcal{D}} \{\mu$-expected difference in loss of best expert and best expert not in $\mathcal{I}_0\}$

Semi-Adversarial Paradigm

Real data $\not\equiv$ stochastic←Optimistic   Real data $\not\equiv$ adversarial←Pessimistic

[BNR20] introduces a spectrum between i.i.d. and adversarial.

Intuitively, fix a "neighbourhood" of distributions $\mathcal{D} \subseteq \mathcal{M}([0,1]^N)$.

Each data point drawn from an arbitrary distribution in "neighbourhood".

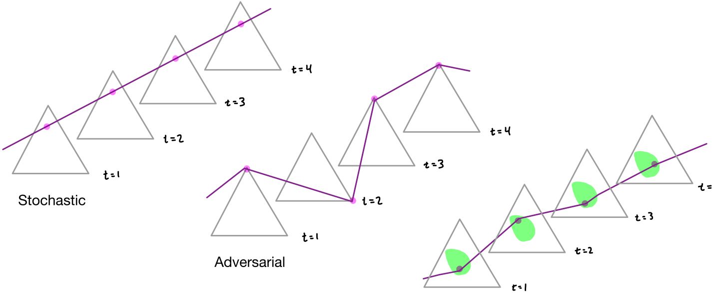

$\mathcal{I}_0 = \{$experts that are optimal for some $\mu \in \mathcal{D}\}$ $N_0 = |\mathcal{I}_0|$

$\Delta_0 = \inf_{\mu \in \mathcal{D}} \{\mu$-expected difference in loss of best expert and best expert not in $\mathcal{I}_0\}$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E} R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$

- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$

- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$
- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$
- META-CARE algorithm achieves
  $\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}_{[N_0=1]}(\log N)/\Delta_0 + \mathbb{I}_{[N_0>1]}(\log N)^{3/2}/\Delta_0$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$
- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$
- META-CARE algorithm achieves $\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}_{[N_0=1]}(\log N)/\Delta_0 + \mathbb{I}_{[N_0>1]}(\log N)^{3/2}/\Delta_0$

Theorem [NEW] (Semi-Adversarial Regret Bound for FTRL-CARL)

Using $f(x) = \int_1^x -\sqrt{2\log(1/s)}\mathrm{d}s$,

$$\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + (\log N)/\Delta_0.$$

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$
- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$
- META-CARE algorithm achieves
  $\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}_{[N_0=1]}(\log N)/\Delta_0 + \mathbb{I}_{[N_0>1]}(\log N)^{3/2}/\Delta_0$

Theorem [NEW] (Semi-Adversarial Regret Bound for FTRL-CARL)

Using $f(x) = \int_1^x -\sqrt{2\log(1/s)}\mathrm{d}s$,

$$\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + (\log N)/\Delta_0.$$

Improved dependence on (N, Δ_0).

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$
- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$
- META-CARE algorithm achieves
  $\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}_{[N_0=1]}(\log N)/\Delta_0 + \mathbb{I}_{[N_0>1]}(\log N)^{3/2}/\Delta_0$

Theorem [NEW] (Semi-Adversarial Regret Bound for FTRL-CARL)

Using $f(x) = \int_1^x -\sqrt{2\log(1/s)}\mathrm{d}s$,

$$\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + (\log N)/\Delta_0.$$

Improved dependence on (N, Δ_0).

Improvement by an arbitrarily large multiplicative factor for "small" time horizons.

Novel Semi-Adversarial Regret Guarantees

Existing Results [BNR20]

- For every algorithm and $\mathcal{D}$, $\mathbb{E}R_T \gtrsim \sqrt{T \log N_0} + (\log N)/\Delta_0$
- Hedge algorithm can't do better than $\mathbb{E}R_T \gtrsim (\log N_0)\sqrt{T} + (\log N)/\Delta_0$
  - The standard tuning even gets $\mathbb{E}R_T \gtrsim \mathbb{I}_{[N_0>1]}\sqrt{T \log N} + (\log N)/\Delta_0$
- META-CARE algorithm achieves $\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}_{[N_0=1]}(\log N)/\Delta_0 + \mathbb{I}_{[N_0>1]}(\log N)^{3/2}/\Delta_0$

Theorem [NEW] (Semi-Adversarial Regret Bound for FTRL-CARL)

Using $f(x) = \int_1^x -\sqrt{2\log(1/s)}\mathrm{d}s$,

$$\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + (\log N)/\Delta_0.$$

Improved dependence on (N, Δ_0).

Improvement by an arbitrarily large multiplicative factor for "small" time horizons.

Technical Details for Generic Bounds

$$\pi_{t+1} = \operatorname*{arg\,min}_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} D_f \left(\pi \| \nu \right) \right\}.$$

Technical Details for Generic Bounds

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta\sim\pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

Lemma [NEW] (Generic FTRL Solution)

We provide an analytic expression for π_{t+1} whenever $f'' > 0$.

Technical Details for Generic Bounds

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta\sim\pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

Lemma [NEW] (Generic FTRL Solution)

We provide an analytic expression for π_{t+1} whenever $f'' > 0$.

Theorem [NEW] (Generic FTRL Local-Norm Bound)

If $f'' > 0$, for all $(m_t)_{t\in\mathbb{N}}$ there exist $\mu_{t+1} \in \mathrm{Conv}(\pi_t, \pi_{t+1})$ s.t. for all $q \ll \nu$

$$R_T(q) \le \underbrace{\eta_{T+1}^{-1} D_f\left(q \| \nu\right)}_{\text{rate of regret}} - \underbrace{\eta_1^{-1} \min_{\pi\ll\nu} D_f\left(\pi \| \nu\right)}_{\text{lower order}} + \sum_{t=1}^{T} \int \Bigg[\underbrace{\frac{\eta_t(\ell_t(\theta) - m_t)^2}{2f''(\mu_{t+1}(\theta)/\nu(\theta))}}_{\text{local-norm variance}} - \underbrace{\left(\eta_{t+1}^{-1} - \eta_t^{-1}\right) f\left(\frac{\pi_{t+1}(\theta)}{\nu(\theta)}\right)}_{\text{regularizer increments}} \Bigg] \nu(\mathrm{d}\theta)$$

Technical Details for Generic Bounds

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta\sim\pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

Lemma [NEW] (Generic FTRL Solution)

We provide an analytic expression for π_{t+1} whenever $f'' > 0$.

Theorem [NEW] (Generic FTRL Local-Norm Bound)

If $f'' > 0$, for all $(m_t)_{t\in\mathbb{N}}$ there exist $\mu_{t+1} \in \mathrm{Conv}(\pi_t, \pi_{t+1})$ s.t. for all $q \ll \nu$

$$R_T(q) \le \underbrace{\eta_{T+1}^{-1} D_f\left(q \| \nu\right)}_{\text{rate of regret}} - \underbrace{\eta_1^{-1} \min_{\pi\ll\nu} D_f\left(\pi \| \nu\right)}_{\text{lower order}}$$

$$+ \sum_{t=1}^{T} \int \Bigg[\underbrace{\frac{\eta_t(\ell_t(\theta) - m_t)^2}{2f''(\mu_{t+1}(\theta)/\nu(\theta))}}_{\text{local-norm variance}} - \underbrace{\left(\eta_{t+1}^{-1} - \eta_t^{-1}\right) f\left(\frac{\pi_{t+1}(\theta)}{\nu(\theta)}\right)}_{\text{regularizer increments}} \Bigg] \nu(\mathrm{d}\theta)$$

Interpretation of Terms

- Choice of f removes need to place desired rate of regret in η_t

[OLD] Choose η_t to match rate of regret with local-norm for fixed f

[NEW] Choose f to match local-norm with regularizer increments: $f \approx \pm[f'']^{-1}$

Technical Details for Generic Bounds

$$\pi_{t+1} = \arg\min_{\pi \ll \nu} \left\{ \mathbb{E}_{\theta \sim \pi} L_t(\theta) + \eta_t^{-1} D_f\left(\pi \| \nu\right) \right\}.$$

Lemma [NEW] (Generic FTRL Solution)

We provide an analytic expression for π_{t+1} whenever $f'' > 0$.

Theorem [NEW] (Generic FTRL Local-Norm Bound)

If $f'' > 0$, for all $(m_t)_{t \in \mathbb{N}}$ there exist $\mu_{t+1} \in \mathrm{Conv}(\pi_t, \pi_{t+1})$ s.t. for all $q \ll \nu$

$$R_T(q) \leq \underbrace{\eta_{T+1}^{-1} D_f\left(q \| \nu\right)}_{\text{rate of regret}} - \underbrace{\eta_1^{-1} \min_{\pi \ll \nu} D_f\left(\pi \| \nu\right)}_{\text{lower order}}$$

$$+ \sum_{t=1}^{T} \int \Bigg[\underbrace{\frac{\eta_t(\ell_t(\theta) - m_t)^2}{2 f''(\mu_{t+1}(\theta)/\nu(\theta))}}_{\text{local-norm variance}} - \underbrace{\left(\eta_{t+1}^{-1} - \eta_t^{-1}\right) f\left(\frac{\pi_{t+1}(\theta)}{\nu(\theta)}\right)}_{\text{regularizer increments}} \Bigg] \nu(\mathrm{d}\theta)$$

Interpretation of Terms

- Choice of f removes need to place desired rate of regret in η_t

[OLD] Choose η_t to match rate of regret with local-norm for fixed f

[NEW] Choose f to match local-norm with regularizer increments: $f \approx \pm [f'']^{-1}$

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.
- Obtain the first root-KL quantile regret bounds with an FTRL algorithm.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.
- Obtain the first root-KL quantile regret bounds with an FTRL algorithm.
- Prove novel lower bounds for quantile regret that show root-KL is tight.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.
- Obtain the first root-KL quantile regret bounds with an FTRL algorithm.
- Prove novel lower bounds for quantile regret that show root-KL is tight.
- Design a novel regularizer with improved semi-adversarial regret bounds.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.
- Obtain the first root-KL quantile regret bounds with an FTRL algorithm.
- Prove novel lower bounds for quantile regret that show root-KL is tight.
- Design a novel regularizer with improved semi-adversarial regret bounds.
- Applications to variance regret, posterior prediction, and model selection.

Conclusion

- Provide a novel local-norm analysis of FTRL for uncountable experts.
- Include an explicit expression for the FTRL solution with generic regularizers.
- Obtain the first root-KL quantile regret bounds with an FTRL algorithm.
- Prove novel lower bounds for quantile regret that show root-KL is tight.
- Design a novel regularizer with improved semi-adversarial regret bounds.
- Applications to variance regret, posterior prediction, and model selection.

References

- [Alq21] P. Alquier. "Non-Exponentially Weighted Aggregation: Regret Bounds for Unbounded Loss Functions". In: *Proceedings of the 38th International Conference on Machine Learning*. 2021.
- [BNR20] B. Bilodeau, J. Negrea, and D. M. Roy. *Relaxing the I.I.D. Assumption: Adaptively Minimax Optimal Regret via Root-Entropic Regularization*. arXiv:2007.06552. 2020.
- [CFH09] K. Chaudhuri, Y. Freund, and D. J. Hsu. "A Parameter-Free Hedging Algorithm". In: *Advances in Neural Information Processing Systems 22*. 2009.
- [CV10] A. Chernov and V. G. Vovk. "Prediction with Advice of Unknown Number of Experts". In: *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*. 2010.
- [OP16] F. Orabona and D. Pál. "Coin Betting and Parameter-Free Online Learning". In: *Advances in Neural Information Processing Systems 29*. 2016.