The Design of Online Learning Algorithms

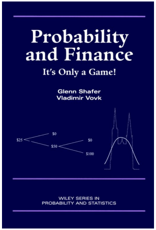

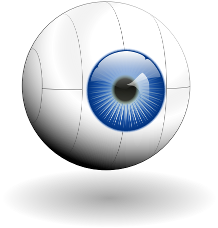

Wouter M. Koolen

CWI
Centrum Wiskunde & Informatica

Online Learning Workshop
Paris, Friday 20th October, 2017

Conclusion

A simple factor $(1 + \eta r_t)$ stretches surprisingly far.

Outline

Coin Betting

$\mathcal{K}_0 = 1$.
For $t = 1, 2, \dots$

- Skeptic picks $M_t \in \mathbb{R}$
- Reality picks $r_t \in [-1, 1]$
- $\mathcal{K}_t = \mathcal{K}_{t-1} + M_t r_t$

Coin Betting

$\mathcal{K}_0 = 1$.
For $t = 1, 2, \ldots$

- Skeptic picks $M_t \in \mathbb{R}$
- Reality picks $r_t \in [-1, 1]$
- $\mathcal{K}_t = \mathcal{K}_{t-1} + M_t r_t$

Pick an event E.
Skeptic wins if

1. $\mathcal{K}_t \geq 0$
2. $r_1 r_2 \cdots \in E$ or $\mathcal{K}_t \to \infty$.

Say: Skeptic can **force** E.

Forcing The Law of Large Numbers

Fix $\eta \in [0, 1/2]$. Suppose Skeptic plays $M_t = \mathcal{K}_{t-1}\eta$.
Then

$$\mathcal{K}_T \;=\; \mathcal{K}_{T-1} + \mathcal{K}_{T-1}\eta r_T \;=\; \mathcal{K}_{T-1}(1 + \eta r_T) \;=\; \prod_{t=1}^{T}(1 + \eta r_t)$$

Forcing The Law of Large Numbers

Fix $\eta \in [0, 1/2]$. Suppose Skeptic plays $M_t = \mathcal{K}_{t-1}\eta$.
Then

$$\mathcal{K}_T \;=\; \mathcal{K}_{T-1} + \mathcal{K}_{T-1}\eta r_T \;=\; \mathcal{K}_{T-1}(1 + \eta r_T) \;=\; \prod_{t=1}^{T}(1 + \eta r_t)$$

Now say $C \geq \mathcal{K}_T$. Then, using $\ln(1 + x) \geq x - x^2$,

$$\ln C \;\geq\; \sum_{t=1}^{T} \ln(1 + \eta r_t) \;\geq\; \sum_{t=1}^{T} \eta r_t - T\eta^2$$

Forcing The Law of Large Numbers

Fix $\eta \in [0, 1/2]$. Suppose Skeptic plays $M_t = \mathcal{K}_{t-1}\eta$.
Then

$$\mathcal{K}_T \;=\; \mathcal{K}_{T-1} + \mathcal{K}_{T-1}\eta r_T \;=\; \mathcal{K}_{T-1}(1 + \eta r_T) \;=\; \prod_{t=1}^{T}(1 + \eta r_t)$$

Now say $C \geq \mathcal{K}_T$. Then, using $\ln(1+x) \geq x - x^2$,

$$\ln C \;\geq\; \sum_{t=1}^{T} \ln(1 + \eta r_t) \;\geq\; \sum_{t=1}^{T} \eta r_t - T\eta^2$$

Hence

$$\frac{\ln C}{T\eta} \;\geq\; \frac{1}{T}\sum_{t=1}^{T} r_t - \eta \qquad \text{and so} \qquad \eta \geq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} r_t$$

Forcing The Law of Large Numbers

Finally, let Skeptic allocate a fraction γ_i of his initial $\mathcal{K}_0 = 1$ to η_i. Then

$$\mathcal{K}_T \;=\; \sum_{i=1}^{\infty} \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t)$$

Forcing The Law of Large Numbers

Finally, let Skeptic allocate a fraction γ_i of his initial $\mathcal{K}_0 = 1$ to η_i. Then

$$\mathcal{K}_T \;=\; \sum_{i=1}^{\infty} \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t)$$

Now suppose $C \geq \mathcal{K}_T$. Then for each i:

$$\ln C \;\geq\; \ln \gamma_i + \sum_{t=1}^{T} \ln(1 + \eta_i r_t)$$

Forcing The Law of Large Numbers

Finally, let Skeptic allocate a fraction γ_i of his initial $\mathcal{K}_0 = 1$ to η_i. Then

$$\mathcal{K}_T \;=\; \sum_{i=1}^{\infty} \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t)$$

Now suppose $C \geq \mathcal{K}_T$. Then for each i:

$$\ln C \;\geq\; \ln \gamma_i + \sum_{t=1}^{T} \ln(1 + \eta_i r_t)$$

So for each i

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} r_t \;\leq\; \eta_i$$

and hence

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} r_t \;\leq\; 0$$

What else

Skeptic can force many laws of probability. For example the LIL

$$\limsup_{T \to \infty} \frac{\sum_{t=1}^{T} r_t}{\sqrt{2 T \ln \ln T}} \leq 1$$

What else

Skeptic can force many laws of probability. For example the LIL

$$\limsup_{T \to \infty} \frac{\sum_{t=1}^{T} r_t}{\sqrt{2T \ln \ln T}} \leq 1$$

Small deviations?

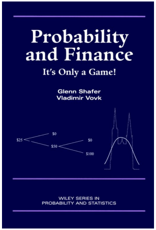

Outline

Experts

Let's play the experts game.

- Learner picks $\boldsymbol{w}_t \in \triangle_K$
- Reality picks $\boldsymbol{\ell}_t \in [0,1]^K$
- Learner incurs $\langle \boldsymbol{w}_t, \boldsymbol{\ell}_t \rangle$

Experts

Let's play the experts game.

- Learner picks $\boldsymbol{w}_t \in \triangle_K$
- Reality picks $\boldsymbol{\ell}_t \in [0,1]^K$
- Learner incurs $\langle \boldsymbol{w}_t, \boldsymbol{\ell}_t \rangle$

Goal: make sure regret compared to any expert k is **sublinear**.

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} r_t^k \ \leq\ 0 \qquad \text{where} \qquad r_t^k \ =\ \langle \boldsymbol{w}_t, \boldsymbol{\ell}_t \rangle - \ell_t^k$$

Experts

Let's play the experts game.

- Learner picks $\boldsymbol{w}_t \in \triangle_K$
- Reality picks $\boldsymbol{\ell}_t \in [0,1]^K$
- Learner incurs $\langle \boldsymbol{w}_t, \boldsymbol{\ell}_t \rangle$

Goal: make sure regret compared to any expert k is **sublinear**.

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} r_t^k \;\leq\; 0 \qquad \text{where} \qquad r_t^k \;=\; \langle \boldsymbol{w}_t, \boldsymbol{\ell}_t \rangle - \ell_t^k$$

Key idea: **Defensive Forecasting**

1. Fix a strategy for Skeptic that forces this goal.
2. Play $\boldsymbol{w}_t$ so that Skeptic does not get rich

Stategy

Stragegy for Skeptic:

Split capital $\mathcal{K}_0 = 1$ over experts k with weights π_k and η_i with γ_i.

$$\mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

Stategy

Stragegy for Skeptic:
Split capital $\mathcal{K}_0 = 1$ over experts k with weights π_k and η_i with γ_i.

$$\mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

How to play? Make sure $\mathcal{K}_T$ does not grow big.

$$\begin{aligned}
\mathcal{K}_{T+1} - \mathcal{K}_T \;&=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k) \eta_i r_{T+1}^k \\
&=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k) \eta_i \left(\langle \boldsymbol{w}_{T+1}, \boldsymbol{\ell}_{T+1} \rangle - \ell_{T+1}^k \right) \\
&=\; 0
\end{aligned}$$

Stategy

Stragegy for Skeptic:

Split capital $\mathcal{K}_0 = 1$ over experts k with weights π_k and η_i with γ_i.

$$\mathcal{K}_T = \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

How to play? Make sure $\mathcal{K}_T$ does not grow big.

$$\begin{aligned}
\mathcal{K}_{T+1} - \mathcal{K}_T &= \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k) \eta_i r_{T+1}^k \\
&= \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k) \eta_i \left(\langle \boldsymbol{w}_{T+1}, \boldsymbol{\ell}_{T+1} \rangle - \ell_{T+1}^k \right) \\
&= 0
\end{aligned}$$

when we pick

$$w_{T+1}^k = \frac{\sum_i \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k) \eta_i}{\sum_{j,i} \pi_j \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^j) \eta_i} \qquad \text{(iProd)}$$

Outline

Squint

$$w_{T+1}^{k} = \frac{\sum_i \pi_k \gamma_i \prod_{t=1}^{T}(1+\eta_i r_t^k)\eta_i}{\sum_{j,i} \pi_j \gamma_i \prod_{t=1}^{T}(1+\eta_i r_t^j)\eta_i} \qquad \text{(iProd)}$$

Needs work:

- Rates (how sublinear are the regrets?)
- Computation

Rates

By design,

$$1 = \mathcal{K}_T = \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

Rates

By design,

$$1 \;=\; \mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

So for each i and k,

$$-\ln \pi_k - \ln \gamma_i \;\geq\; \sum_{t=1}^{T} \ln(1 + \eta_i r_t^k) \;\geq\; \eta_i \sum_{t=1}^{T} r_t^k - \eta_i^2 \sum_{t=1}^{T} (r_t^k)^2$$

Rates

By design,

$$1 \;=\; \mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

So for each i and k,

$$-\ln \pi_k - \ln \gamma_i \;\geq\; \sum_{t=1}^{T} \ln(1 + \eta_i r_t^k) \;\geq\; \eta_i \sum_{t=1}^{T} r_t^k - \eta_i^2 \sum_{t=1}^{T} (r_t^k)^2$$

That is, abbreviating $v_t^k = (r_t^k)^2$,

$$\sum_{t=1}^{T} r_t^k \;\leq\; \min_i \left(\eta_i \sum_{t=1}^{T} v_t^k + \frac{-\ln \pi_k - \ln \gamma_i}{\eta_i} \right)$$

Rates

By design,

$$1 \;=\; \mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

So for each i and k,

$$-\ln \pi_k - \ln \gamma_i \;\geq\; \sum_{t=1}^{T} \ln(1 + \eta_i r_t^k) \;\geq\; \eta_i \sum_{t=1}^{T} r_t^k - \eta_i^2 \sum_{t=1}^{T} (r_t^k)^2$$

That is, abbreviating $v_t^k = (r_t^k)^2$,

$$\sum_{t=1}^{T} r_t^k \;\leq\; \min_i \left(\eta_i \sum_{t=1}^{T} v_t^k + \frac{-\ln \pi_k - \ln \gamma_i}{\eta_i} \right)$$

Theorem (Koolen and van Erven [2015])

$$R_T^k \;\leq\; O\left(\sqrt{V_T^k \left(-\ln \pi_k + \ln \ln V_T^k\right)} \right)$$

Computation

$$\mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

We are using $\ln(1 + r) \geq r - r^2$ in the **analysis**. Put it in **algorithm**?

Computation

$$\mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

We are using $\ln(1 + r) \geq r - r^2$ in the **analysis**. Put it in **algorithm**?

Indeed, we can start from **supermartingale**

$$\mathcal{K}_T \;\geq\; \Phi_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} e^{\eta_i r_t^k - \eta_i^2 v_t^k} \;=\; \sum_{k,i} \pi_k \gamma_i e^{\eta_i R_T^k - \eta_i^2 V_T^k}$$

Computation

$$\mathcal{K}_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} (1 + \eta_i r_t^k)$$

We are using $\ln(1+r) \geq r - r^2$ in the **analysis**. Put it in **algorithm**?
Indeed, we can start from **supermartingale**

$$\mathcal{K}_T \;\geq\; \Phi_T \;=\; \sum_{k,i} \pi_k \gamma_i \prod_{t=1}^{T} e^{\eta_i r_t^k - \eta_i^2 v_t^k} \;=\; \sum_{k,i} \pi_k \gamma_i e^{\eta_i R_T^k - \eta_i^2 V_T^k}$$

One choice of weights to keep this small:

$$w_{T+1}^k \;=\; \frac{\sum_i \pi_k \gamma_i e^{\eta_i R_T^k - \eta_i^2 V_T^k} \eta_i}{\sum_{j,i} \pi_j \gamma_i e^{\eta_i R_T^j - \eta_i^2 V_T^j} \eta_i} \qquad \text{(Squint)}$$

Computation, ctd

$$w^k_{T+1} = \frac{\sum_i \pi_k \gamma_i e^{\eta_i R^k_T - \eta_i^2 V^k_T} \eta_i}{\sum_{j,i} \pi_j \gamma_i e^{\eta_i R^j_T - \eta_i^2 V^j_T} \eta_i} \qquad \text{(Squint)}$$

Computation, ctd

$$w_{T+1}^k = \frac{\sum_i \pi_k \gamma_i e^{\eta_i R_T^k - \eta_i^2 V_T^k} \eta_i}{\sum_{j,i} \pi_j \gamma_i e^{\eta_i R_T^j - \eta_i^2 V_T^j} \eta_i} \qquad \text{(Squint)}$$

Maybe a continuous prior on η could help? How to make

$$\int_0^{1/2} \gamma(\eta) e^{\eta R_T^k - \eta^2 V_T^k} \eta \, \mathrm{d}\eta$$

fast?

Computation, ctd

$$w_{T+1}^{k} = \frac{\sum_i \pi_k \gamma_i e^{\eta_i R_T^k - \eta_i^2 V_T^k} \eta_i}{\sum_{j,i} \pi_j \gamma_i e^{\eta_i R_T^j - \eta_i^2 V_T^j} \eta_i} \qquad \text{(Squint)}$$

Maybe a continuous prior on η could help? How to make

$$\int_0^{1/2} \gamma(\eta) e^{\eta R_T^k - \eta^2 V_T^k} \eta \, \mathrm{d}\eta$$

fast?

- Conjugate $\gamma(\eta) \propto e^{-a\eta - b\eta^2}$. $\Rightarrow$ Truncated Guassian mean.
- Improper $\gamma(\eta) \propto \frac{1}{\eta}$. $\Rightarrow$ Gaussian CDF.

Computation, ctd

$$w^k_{T+1} \;=\; \frac{\sum_i \pi_k \gamma_i e^{\eta_i R^k_T - \eta_i^2 V^k_T} \eta_i}{\sum_{j,i} \pi_j \gamma_i e^{\eta_i R^j_T - \eta_i^2 V^j_T} \eta_i} \qquad \text{(Squint)}$$

Maybe a continuous prior on η could help? How to make

$$\int_0^{1/2} \gamma(\eta) e^{\eta R^k_T - \eta^2 V^k_T} \eta \, \mathrm{d}\eta$$

fast?

- Conjugate $\gamma(\eta) \propto e^{-a\eta - b\eta^2}$. $\Rightarrow$ Truncated Guassian mean.
- Improper $\gamma(\eta) \propto \frac{1}{\eta}$. $\Rightarrow$ Gaussian CDF.

Theorem (Koolen and van Erven [2015])

$$R^k_T \;\le\; O\left(\sqrt{V^k_T \left(-\ln \pi_k + \ln \ln T\right)}\right)$$

Squint Conclusion

Computation:

- $O(1)$ time per round (like Hedge, Adapt-ML-Prod, ...)
- Library: `https://bitbucket.org/wmkoolen/squint`

Regret:

- Adaptive $\sqrt{V_T^k(\ln K + \ln\ln T)}$ bound.
- Implies L^* bound, T bound.
- Constant regret in stochastic gap case.

Outline

Online Convex Optimisation

Let's play the OCO game.

For $t = 1, 2, \ldots$

- Learner plays $\boldsymbol{w}_t \in \mathcal{U}$ (convex, bounded).
- Reality picks $f_t : \mathcal{U} \to \mathbb{R}$ (convex, bounded gradient)
- Learner incurs $f(\boldsymbol{w}_t)$ and observes $\nabla f_t(\boldsymbol{w}_t)$

Online Convex Optimisation

Let's play the OCO game.
For $t = 1, 2, \ldots$

- Learner plays $\boldsymbol{w}_t \in \mathcal{U}$ (convex, bounded).
- Reality picks $f_t : \mathcal{U} \to \mathbb{R}$ (convex, bounded gradient)
- Learner incurs $f(\boldsymbol{w}_t)$ and observes $\nabla f_t(\boldsymbol{w}_t)$

Goal: small regret w.r.t. all $\boldsymbol{u} \in \mathcal{U}$

$$R_T^{\boldsymbol{u}} = \sum_{t=1}^{T} \left(f_t(\boldsymbol{w}_t) - f_t(\boldsymbol{u}) \right)$$

Online Convex Optimisation

Let's play the OCO game.
For $t = 1, 2, \ldots$

- Learner plays $\boldsymbol{w}_t \in \mathcal{U}$ (convex, bounded).
- Reality picks $f_t : \mathcal{U} \to \mathbb{R}$ (convex, bounded gradient)
- Learner incurs $f(\boldsymbol{w}_t)$ and observes $\nabla f_t(\boldsymbol{w}_t)$

Goal: small regret w.r.t. all $\boldsymbol{u} \in \mathcal{U}$

$$R_T^{\boldsymbol{u}} \;=\; \sum_{t=1}^{T} \left(f_t(\boldsymbol{w}_t) - f_t(\boldsymbol{u}) \right)$$

Step 1: let's play a harder game with linearised loss

$$f_t(\boldsymbol{w}_t) - f_t(\boldsymbol{u}) \;\leq\; \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle =: r_t^{\boldsymbol{u}}$$

iProd/Squint for OCO

Goal: keep regret small for all $\boldsymbol{u} \in \mathcal{U} \Rightarrow$ prior π on $\boldsymbol{u}$.

$$\mathcal{K}_T = \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \,\mathrm{d}\boldsymbol{u} \,\mathrm{d}\eta$$

Kan we keep this small?

iProd/Squint for OCO

Goal: keep regret small for all $\boldsymbol{u} \in \mathcal{U} \Rightarrow$ prior π on $\boldsymbol{u}$.

$$\mathcal{K}_T \;=\; \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

Kan we keep this small?

$\mathcal{K}_{T+1} - \mathcal{K}_T$

$$= \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \left\langle \boldsymbol{w}_{T+1} - \boldsymbol{u}, \nabla f_{T+1}(\boldsymbol{w}_{T+1}) \right\rangle \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

iProd/Squint for OCO

Goal: keep regret small for all $\boldsymbol{u} \in \mathcal{U} \Rightarrow$ prior π on $\boldsymbol{u}$.

$$\mathcal{K}_T = \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

Kan we keep this small?

$\mathcal{K}_{T+1} - \mathcal{K}_T$

$$= \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \left\langle \boldsymbol{w}_{T+1} - \boldsymbol{u}, \nabla f_{T+1}(\boldsymbol{w}_{T+1}) \right\rangle \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

Need

$$\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \left(\boldsymbol{w}_{T+1} - \boldsymbol{u} \right) \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta = 0$$

iProd/Squint for OCO

Goal: keep regret small for all $\boldsymbol{u} \in \mathcal{U} \Rightarrow$ prior π on $\boldsymbol{u}$.

$$\mathcal{K}_T = \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

Kan we keep this small?

$\mathcal{K}_{T+1} - \mathcal{K}_T$

$$= \int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \left\langle \boldsymbol{w}_{T+1} - \boldsymbol{u}, \nabla f_{T+1}(\boldsymbol{w}_{T+1}) \right\rangle \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta$$

Need

$$\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \left(\boldsymbol{w}_{T+1} - \boldsymbol{u} \right) \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta = 0$$

which mandates

$$\boldsymbol{w}_{T+1} = \frac{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \boldsymbol{u} \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}$$

OCO iProd:

$$\boldsymbol{w}_{T+1} = \frac{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \boldsymbol{u} \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) \prod_{t=1}^{T} (1 + \eta r_t^{\boldsymbol{u}}) \eta \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}$$

Work needed:

- Picking priors
- Computation
- Rates

Prod bound to the rescue

We might also use

$$\boldsymbol{w}_{T+1} = \frac{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \boldsymbol{u} \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}$$

$$R_T^{\boldsymbol{u}} = \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle$$

$$V_T^{\boldsymbol{u}} = \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle^2$$

Prod bound to the rescue

We might also use

$$\boldsymbol{w}_{T+1} = \frac{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \boldsymbol{u} \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}$$

$$R_T^{\boldsymbol{u}} = \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle$$

$$V_T^{\boldsymbol{u}} = \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle^2$$

linear/quadratic in $\boldsymbol{u}$. $\Rightarrow$ suggests $\pi(\boldsymbol{u})$ multivariate Normal.

Prod bound to the rescue

We might also use

$$\boldsymbol{w}_{T+1} = \frac{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \boldsymbol{u} \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}{\int_0^{1/2} \gamma(\eta) \int_{\mathcal{U}} \pi(\boldsymbol{u}) e^{\eta R_T^{\boldsymbol{u}} - \eta^2 V_T^{\boldsymbol{u}}} \eta \, \mathrm{d}\boldsymbol{u} \, \mathrm{d}\eta}$$

$$\begin{aligned} R_T^{\boldsymbol{u}} &= \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle \\ V_T^{\boldsymbol{u}} &= \sum_{t=1}^{T} \langle \boldsymbol{w}_t - \boldsymbol{u}, \nabla f_t(\boldsymbol{w}_t) \rangle^2 \end{aligned}$$

linear/quadratic in $\boldsymbol{u}$. $\Rightarrow$ suggests $\pi(\boldsymbol{u})$ multivariate Normal.

But then $\boldsymbol{w}_t$ may end up outside $\mathcal{U}$. And $r_t^{\boldsymbol{u}}$ not bounded.

MetaGrad

Let's do it anyway. Turns out it works.

$$\boldsymbol{w} = \frac{\sum_i \gamma_i \eta_i \boldsymbol{w}_i}{\sum_i \gamma_i \eta_i}$$

$$\gamma_i \leftarrow \gamma_i e^{-\eta_i r_i - \eta_i^2 r_i^2}$$

where

$$r_i = (\boldsymbol{w}_i - \boldsymbol{w})^\top \nabla f_t$$

Tilted Exponential Weights

$$\Sigma_i \leftarrow (\Sigma_i^{-1} + 2\eta_i^2 \nabla f_t \nabla f_t^\top)^{-1}$$

$$\boldsymbol{w}_i \leftarrow \Pi_{\mathcal{U}} (\boldsymbol{w}_i - \eta_i \Sigma_i \nabla f_t (1 + 2\eta_i r_i))$$

$\approx$ Online Newton Step

MetaGrad

Let's do it anyway. Turns out it works.

$$\boldsymbol{w} = \frac{\sum_i \gamma_i \eta_i \boldsymbol{w}_i}{\sum_i \gamma_i \eta_i}$$

$$\gamma_i \leftarrow \gamma_i e^{-\eta_i r_i - \eta_i^2 r_i^2}$$

where

$$r_i = (\boldsymbol{w}_i - \boldsymbol{w})^\top \nabla f_t$$

Tilted Exponential Weights

$$\Sigma_i \leftarrow (\Sigma_i^{-1} + 2\eta_i^2 \nabla f_t \nabla f_t^\top)^{-1}$$

$$\boldsymbol{w}_i \leftarrow \Pi_{\mathcal{U}} (\boldsymbol{w}_i - \eta_i \Sigma_i \nabla f_t (1 + 2\eta_i r_i))$$

$\approx$ Online Newton Step

Theorem

The regret of MetaGrad is bounded by

$$R_T = O\left(\min\left\{\sqrt{T}, \sqrt{V_T^{u^*} d \ln T}\right\}\right)$$

Consequences

What's new with $\sqrt{V_T d \ln T}$?

Consequences

What's new with $\sqrt{V_T d \ln T}$?

Corollary

For α-**exp-concave** *or* α-**strongly convex** *losses, MetaGrad ensures*

$$R_T = O(d \ln T)$$

without knowing α.

Consequences

What's new with $\sqrt{V_T d \ln T}$?

Corollary

*For α-**exp-concave** or α-**strongly convex** losses, MetaGrad ensures*

$$R_T \;=\; O\left(d \ln T\right)$$

without knowing α.

Corollary (Koolen, Grünwald, and van Erven [2016])

*For any β-**Bernstein** $\mathbb{P}$, MetaGrad keeps the expected regret below*

$$\mathbb{E}\, R_T^* \;\leq\; O\left(\left(d \ln T\right)^{\frac{1}{2-\beta}}\, T^{\frac{1-\beta}{2-\beta}}\right)$$

without knowing β.

Conclusion

A simple factor $(1 + \eta r_t)$ stretches surprisingly far.