Between Stochastic and Adversarial Online Convex Optimization

joint work with Cristóbal Guzmán (Twente),
Sarah Sachs and Tim van Erven (Amsterdam)

Hédi Hadiji, July 2022

Outline

- **(Adversarial) Online Convex Optimization:**
  - Setting, algorithm, rates, etc.

- Between Stochastic and Adversarial OCO
  - Optimism and Intermediate Rates
  - Two Applications
  - Future Directions

Online Convex Optimization

Protocol: Online Convex Optimization

1: **given**: (bounded) decision set $\mathcal{W} \subset \mathbb{R}^d$
2: **for** $t = 1, \ldots, T$
3: Player chooses $w_t \in \mathcal{W}$
4: Nature outputs convex loss $\ell_t : \mathcal{W} \to \mathbb{R}$

[Zinkevich '03]

Goal: minimize **regret**

$$R_T(u) = \sum_{t=1}^{T} \ell_t(w_t) - \sum_{t=1}^{T} \ell_t(u)$$

No assumptions on how the losses are generated

Nature can be an **adversary** who knows the algorithm!

Example: Online Spam Filtering

Training a linear model to filter spams

Email $x_t \in \mathbb{R}^d$, outcome $y_t = 1 - 2 \cdot \mathbf{1}\{x_t \text{ is spam}\} \in \{-1, +1\}$

Get w_t from an OCO algorithm

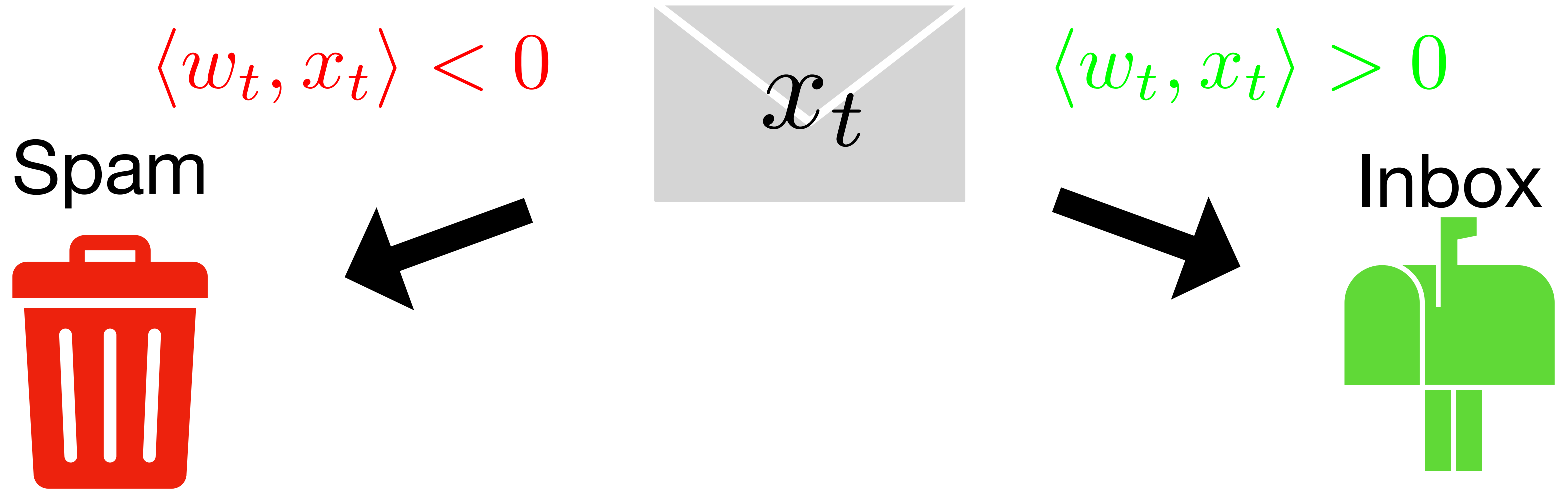

Proxy: Minimize loss: e.g., $\ell_t : w \mapsto (y_t - \langle w, x_t \rangle)^2$

Regret: learn to filter as well as the 'best' linear model in hindsight

Plenty of Other Examples

see surveys by [Hazan ’16, Orabona ’19]

- Prediction with Expert Advice [Freund and Schapire ’97]
- Portfolio Selection [Cover ’91]
- Online Routing
- Batch Optimization

- … and many more

Worst-Case Regret in OCO

Theorem: Minimax Rates for OCO

$$\min_{\text{Alg}} \max_{\ell_{1:T}} R_T \asymp DG\sqrt{T}$$

D is the diameter of $\mathcal{W}$

$G \geqslant \|g_s\|$

$g_s := \nabla \ell_s(w_s)$

Linearized regret $$\sum_{t=1}^{T} \ell_t(w_t) - \ell_t(u) \leqslant \sum_{t=1}^{T} \langle \nabla \ell_t(w_t),\, w_t - u \rangle$$

Upper bound: **FTRL**

$$\eta_t = \frac{D}{G\sqrt{T}}$$

$$w_t \in \operatorname*{arg\,min}_{w \in \mathcal{W}} \left\{ \sum_{s=1}^{t-1} \langle g_s, w \rangle + \frac{\|w\|^2}{\eta_t} \right\}$$

Lower bound: linear losses, pure noise (nothing to learn)

(Online Gradient Descent (aka SGD) also reaches the minimax rates)

Easy Data

Losses are often far from worst-case

- Small gradients [Zinkevich’03, Duchi’10], Small comparator [Orabona, Cutkosky]

- Both [Mhammedi, Koolen ’20, Mayo, Hadiji, van Erven ’22]

- Predictable gradients [Rakhlin, Sridharan’13]

- Many more… (curved losses, extra information available, etc.)

<u>This work:</u>

losses = (smooth) stochastic + slowly-varying adversarial

Stochastic Data

Stochastic Optimization, Online-to-Batch

Losses from i.i.d data $\mathbb{E}\big[\ell_t(x)\big] = F(x)$ for all t

Learner gets w_t from an OCO algorithm $\overline{w}_T = \frac{1}{T}\sum_{t=1}^{T} w_t$

$$\mathbb{E}\big[F(\overline{w}_T)\big] - \min_{u \in \mathcal{W}} F(u) \leqslant \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} F(w_t) - F(w^\star)\right]$$

$$= \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(w_t) - \ell_t(w^\star)\right] = \frac{\mathbb{E}\big[R_T(w^\star)\big]}{T} \leqslant \frac{DG}{\sqrt{T}}$$

Conversely, given $(\ell_s)_{s \leqslant t-1}$, learner can simulate an optimization algorithm

Stochastic Data

Faster Rates with Smoothness

[Allen-Zhu, Orecchia '17]

Smoothness: $w \mapsto \nabla F(w)$ is L-Lipschitz

If F is smooth: optim error $\mathcal{O}\left(\frac{DG}{\sqrt{T}}\right)$ improves to $\mathcal{O}\left(\frac{D\sigma}{\sqrt{T}} + \frac{LD^2}{T^2}\right)$

where $\sigma^2 \geqslant \mathbb{E}\left[\|g_t - \nabla F(w_t)\|^2\right]$

➡ Regret should scale as $\mathbb{E}[R_T] \leqslant \mathcal{O}\left(D\sigma\sqrt{T} + LD^2\right)$

i.e., replace **magnitude of the largest gradient** with **variance**

Outline

- (Adversarial) Online Convex Optimization:
    - Setting, algorithm, rates, etc.

- **Between Stochastic and Adversarial OCO**
    - Optimism and Intermediate Rates
    - Two Applications
    - Future Directions

Stochastic Extended Adversary

Protocol: OCO with a Stochastic Extended Adversary

1: **given**: (bounded) decision set $\mathcal{W} \subset \mathbb{R}^d$
2: **for** $t = 1, \ldots, T$
3: Player chooses $w_t \in \mathcal{W}$
4: Adversary picks distribution $\mathcal{D}_t$ over convex losses, outputs $\ell_t \sim \mathcal{D}_t$

Goal: minimize
expected regret

$$\mathbb{E}[R_T(u)] = \mathbb{E}\left[\sum_{t=1}^{T} \ell_t(w_t) - \ell_t(u)\right]$$

Generalizes

- Stochastic $\qquad \mathcal{D}_1 = \cdots = \mathcal{D}_T$
- Fully Adversarial $\qquad \mathcal{D}_t = \delta_{\ell_t}$

Algorithm: Optimistic FTRL

Exploiting Regularity/Predictability of the Data

Optimistic FTRL

[Rakhlin, Sridharan '13]

$$w_t \in \arg\min_{w \in \mathcal{W}} \left\{ \sum_{s=1}^{t-1} \langle g_s, w \rangle + \langle M_t, w \rangle + \frac{\|w\|^2}{\eta_t} \right\}$$

Better bounds if $M_t \approx g_t$ (but M_t is chosen before g_t)

Smoothness of expected losses: $g_t \approx g_{t-1}$ in expectation

Adaptive tuning of the learning rate, à la AdaHedge [de Rooij et al. '11]

O-FTRL and the SEA

$$\mathbb{E}_t[\,\cdot\,] := \mathbb{E}_{\ell_t \sim \mathcal{D}_t}[\,\cdot\,]$$

Assumptions: $F_t : w \mapsto \mathbb{E}_t[\ell_t(w)]$ is L-smooth

$$\mathbb{E}_t\left[\|\nabla \ell_t(w) - \nabla F_t(w)\|^2\right] \leqslant \sigma_t^2$$

Theorem: Optimistic AdaFTRL Regret Bound

$$\mathbb{E}[R_T(u)] \leqslant \mathcal{O}\left(D\sqrt{\sum_{t=1}^{T}\sigma_t^2 + \sum_{t=1}^{T-1}\sup_{w\in\mathcal{W}}\|\nabla F_t(w) - \nabla F_{t+1}(w)\|^2 + C}\right)$$

- Fully adaptive (only input is D)
- $C = C(G, L)$
- Key technical ingredient: keeping negative stability terms in analysis (à la [Nemirovski ’05])

Analysis
(up to constants and additive terms)

$$R_T(u) \leqslant \frac{D^2}{\eta_T} + \sum_{t=1}^{T} \eta_t \|g_t - M_t\|^2 - \frac{\|w_t - w_{t+1}\|^2}{2\eta_t}$$

O-FTRL Analysis

$$\leqslant D\sqrt{\sum_{t=1}^{T} \|g_t - g_{t-1}\|^2} - \sum_{t=1}^{T} \frac{\|w_t - w_{t+1}\|^2}{2\eta_t}$$

Adaptive Learning Rate

$M_t = g_{t-1}$

$$\mathbb{E}\big[R_T(u)\big] \leqslant D\sqrt{\sum_{t=1}^{T} \mathbb{E}\big[\|g_t - g_{t-1}\|^2\big]} - \mathbb{E}\left[\sum_{t=1}^{T} \frac{\|w_t - w_{t+1}\|^2}{2\eta_t}\right]$$

Usually thrown away

Analysis II

$$\mathbb{E}\big[R_T(u)\big] \leqslant D\sqrt{\sum_{t=1}^{T}\mathbb{E}\big[\|g_t - g_{t-1}\|^2\big]} - \mathbb{E}\left[\sum_{t=1}^{T}\frac{\|w_t - w_{t+1}\|^2}{2\eta_t}\right]$$

Bounded in expectation by

$$\|g_t - g_{t-1}\|^2 \leqslant 4\Big(\, \|\nabla \ell_t(w_t) - \nabla F_t(w_t)\|^2 \qquad \sigma_t^2 \ ✓$$

$$+ \|\nabla F_t(w_t) - \nabla F_t(w_{t-1})\|^2 \qquad L^2\|w_t - w_{t-1}\|^2 \ ✗$$

$$+ \|\nabla F_t(w_{t-1}) - \nabla F_{t-1}(w_{t-1})\|^2 \qquad \text{Itself} \ ✓$$

$$+ \|\nabla F_{t-1}(w_{t-1}) - \nabla \ell_{t-1}(w_{t-1})\|^2 \,\Big) \qquad \sigma_{t-1}^2 \ ✓$$

Use negative terms to cancel

$$D\sqrt{L^2\sum_{t=1}^{T}\mathbb{E}\big[\|w_t - w_{t-1}\|^2\big]} - \frac{1}{\eta_1}\sum_{t=1}^{T}\|w_t - w_{t-1}\|^2$$

$$\leqslant \sup_{X\in\mathbb{R}}\left\{LDX - \frac{X^2}{2\eta_1}\right\} = \frac{L^2D^2\eta_1}{2}$$

Application: Random Order OCO

[Garber et al. ’20, Sherman et al. ’21]

Adversary selects a set of loss functions, but not the order

$$\mathcal{L} = \{\ell_s \mid s \in [T]\}$$

Practitioners shuffle their data before optimization

Fits in the SEA framework: $\pi(s)$ is the s-th observed loss

$$\mathcal{D}_t = \text{Unif}\left(\mathcal{L} \setminus \{\ell_{\pi(1)}, \ldots, \ell_{\pi(t-1)}\}\right), \quad F_t(w) = \frac{1}{T-t+1}\sum_{s=t}^{T} \ell_{\pi(s)}(w)$$

Almost i.i.d. with expected loss $F(w) = \frac{1}{T}\sum_{s=1}^{T} \ell_s(w)$ **but not quite**

Sampling with replacement vs. Sampling without replacement

Random Order OCO II

Analysis

Theorem: Random Order OCO

$$\mathbb{E}[R_T] \leqslant \mathcal{O}\left(D\sigma\sqrt{T\log\left(\frac{G}{\sigma}\wedge T\right)} + C\right)$$

$$\sigma^2 = \max_{w\in\mathcal{W}} \frac{1}{T}\sum_{t=1}^{T}\left\|\nabla \ell_t(w) - \frac{1}{T}\sum_{s=1}^{T}\nabla \ell_s(w)\right\|^2$$

i.i.d. rate up to the log factor

Proof: a few lines

$$\|\nabla F_t(x) - \nabla F_{t-1}(x)\|^2 \leqslant \frac{4G^2}{(T-t+2)^2}$$

$$\sigma_t^2 \leqslant \frac{T}{T-t+1}\sigma^2$$

$$\mathbb{E}[\sigma_t^2] \leqslant \frac{1}{T}\sum_{t=1}^{T}\max_{x\in\mathcal{W}}\left\|\nabla \ell_t(x) - \frac{1}{T}\sum_{s=1}^{T}\nabla \ell_s(x)\right\|^2 \leqslant (T\sigma^2)\wedge(2G^2)$$

Application: Adversarial Perturbations

- Adversarially perturbed stochastic data:

$$\mathbb{E}_t[\ell_t(w)] = F(w) + c_t(w)$$

- F is smooth and c_t are small adversarial perturbations

Theorem: Regret against Perturbed Losses

$$\mathbb{E}[R_T(u)] \leqslant \mathcal{O}\left(D\sigma\sqrt{T} + D\sqrt{\sum_{t=1}^{T} \sup_{w \in \mathcal{W}} \|\nabla c_t(w)\|^2} + C\right)$$

Corollary: Regret against ``True" Losses

$$\mathbb{E}\left[\sum_{t=1}^{T} F(w_t) - F(w)\right] \leqslant \mathcal{O}\left(D\sigma\sqrt{T} + D\sum_{t=1}^{T} \sup_{w \in \mathcal{W}} \|\nabla c(w)\| + C\right)$$

More in the paper and Future Work

- In the paper:
  - Strong convexity
  - Lower bounds
- Future Work
  - In practice, SGD performs well: extend analysis and identify its limits
  - Bounded regret for stochastic experts with perturbations (Ito ’21): unifying analysis + optimality with arbitrary decision set?
  - Unifying analyses that keep the negative terms (acceleration, games)
  - Dynamic regret, Minimax Optimization
  - Improve applications/find lower bounds

Thank you!