Boosting

Boosting algorithms typically combine “weak” or mildly accurate predictors into a strong predictor that has a much better accuracy. This is achieved by iteratively calling the weak learning oracle with modified distributions over the data, And then, combine those learned classifiers from each round, and take the majority vote as the final predictor.

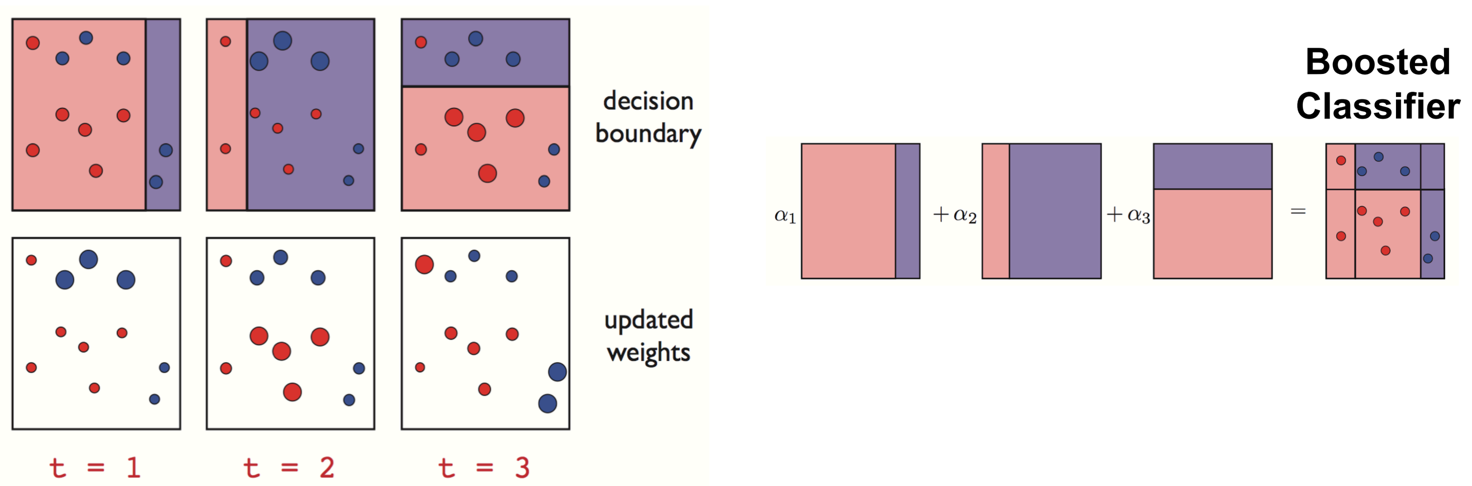

Classical Boosting holds in the Statistical and Realizable learning setting. Can we achieve in other learning settings, e.g. **Online** and/or **Agnostic**?

Agnostic learning removes the realizability assumption that the data is generated by a hypothesis in a given class, and only aims at competing with the best predictor in a given class, rather than perfectly fitting the data. In many practical problems the realizability assumption does not hold, and it is more realistic to assume the agnostic setting.

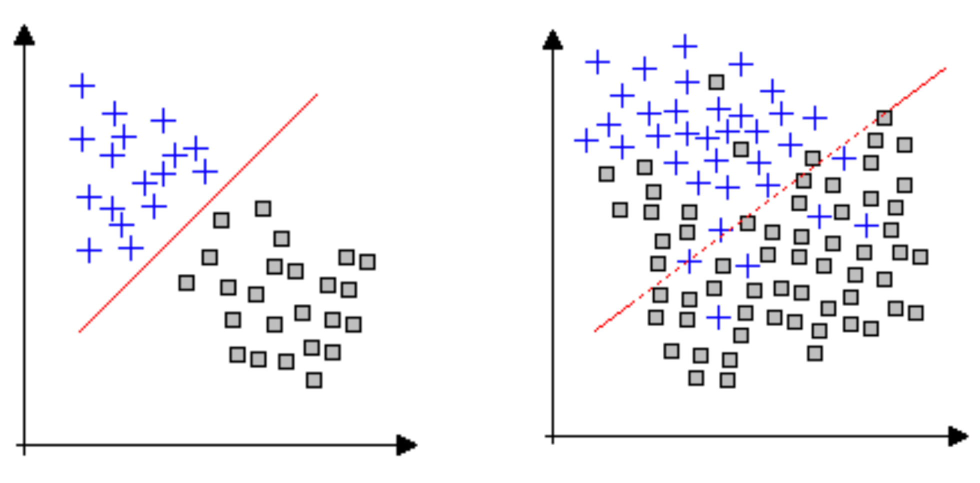

Online Convex Optimization (OCO)

OCO is an optimization process that is sequential. Given a sequence of convex loss functions; and a set of allowed actions K, at each round the optimizer picks an action to minimize the unknown loss function. Then the player observes the loss function The goal is to minimize the **regret:** the difference between the sum of losses incurred by the player, and the sum of losses with the best single action in hindsight.

- Convex loss functions: $\ell_1, \ell_2, \ldots, \ell_T$
- Allowed action: $x \in \mathcal{K}$
- At each round t,
  1. Player chooses $x_t \in \mathcal{K}$ to optimize (unknown) ℓ_t
  2. Player observes loss ℓ_t

Minimize Regret: $R(T) = \sum_{t=1}^{T} \ell_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} \ell_t(x)$

Online Agnostic Boosting via OCO

Boosting in different learning settings:

	Realizable	***Agnostic***
Statistical	Freund and Schapire (1997) *(AdaBoost)*	Kalai and Kanade (2009) ; Feldman (2009)
Online	Chen et al (2012) ; Beygelzimer et al (2015)	**This work**

Online Agnostic Boosting

Setting:
Hypotheses class $\mathcal{H} \subseteq \{\pm 1\}^{\mathcal{X}}$
Data $(x_1, y_1), \ldots, (x_t, y_t) \in \mathcal{X} \times \{\pm 1\}$

Weak *online agnostic* learner:

$$\forall \tilde{\mathbf{y}} \in \{\pm 1\}^T, \quad \sum_{t=1}^{T} \mathcal{W}(x_t)\, \tilde{y}_t \geq \gamma \cdot \max_{h^* \in \mathcal{H}} \sum_{t=1}^{T} h^*(x_t)\, \tilde{y}_t - o(T)$$

OCO (Online Convex Optimiz.)

Strong *online agnostic* learner:

$$\sum_{t=1}^{T} S(x_t) y_t \geq \max_{h^* \in \mathcal{H}} \sum_{t=1}^{T} h^*(x_t) y_t - \epsilon$$

Combination of $\mathcal{W}_1, \ldots, \mathcal{W}_N$

Intuition: A zero-sum game between **Weak learner** and **OCO**, w.r.t carefully designed value function. Iterative game playing results in a **boosted** learner.

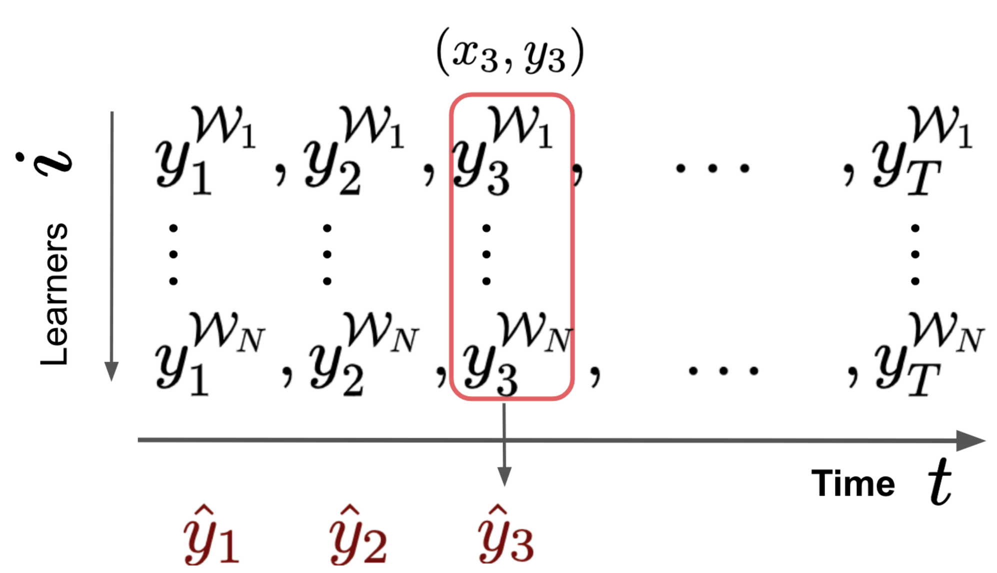

Our approach generalizes to other learning setting:

	Realizable	***Agnostic***
Statistical	Freund and Schapire (1997) *(AdaBoost)* $O(\frac{1}{\gamma^2}\log(\frac{1}{\epsilon}))$ $O(\frac{1}{\epsilon^2\gamma^2})$	Kalai and Kanade (2009) ; Feldman (2009) $O(\frac{1}{\epsilon^2\gamma^2})$ $O(\frac{1}{\epsilon^2\gamma^2})$
Online	Chen et al (2012) ; Beygelzimer et al (2015) $O(\frac{1}{\epsilon\gamma^2})$ $O(\frac{1}{\epsilon^2\gamma^2})$	**This work:** $T = O(\frac{1}{\epsilon^2\gamma^2})$

Note: our complexity bounds can be suboptimal compared to previous work’s bounds. However, it demonstrates the tight connection between boosting and regret minimization and unifies the existing theory.

Online Boosting with Bandit Feedback

- Setting: Online ***regression*** learning with partial information
- Motivation: Reinforcement learning, sub-modular optimization, etc.

- **Challenge** 1: Exact gradients are infeasible/unknown
- **Challenge** 2: Projection onto constraint set is expensive

- **Goal**: **Projection-free** OCO with **Stochastic Gradients**.

Can we just plugin stochastic gradients into Frank-Wolfe?

Projection-free (offline opt.) - Exact gradients

How to solve $\min_{\mathbf{x} \in \mathcal{K}} \ell(\mathbf{x})$?

Frank-Wolfe algorithm — **Doesn't work!**

for $t = 1 \ldots T$
$\tilde{\mathbf{v}}_t \leftarrow \arg\min_{\mathbf{x} \in \mathcal{K}} \left\{ \mathbf{x}^\top \tilde{\nabla} \ell(\mathbf{x}_t) \right\}$
$\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t + \frac{1}{t}(\mathbf{v}_t - \mathbf{x}_t)$

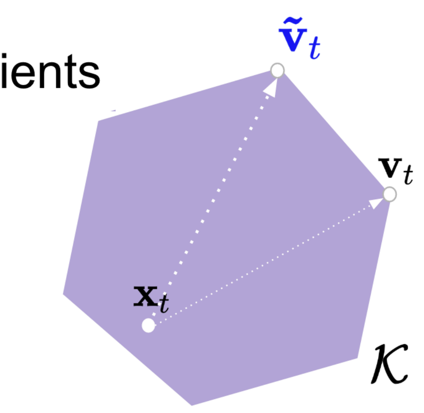

$\mathbb{E}[\tilde{\nabla}_t] = \nabla \ell(\mathbf{x}_t)$
But:
$\mathbb{E}[\tilde{\mathbf{v}}_t] \neq \mathbf{v}_t$

Our algorithm: projection-free OCO w. stochastic gradients

Oracle access: Online linear optimizers (OLO) $\mathcal{A}_1, \ldots, \mathcal{A}_N$.

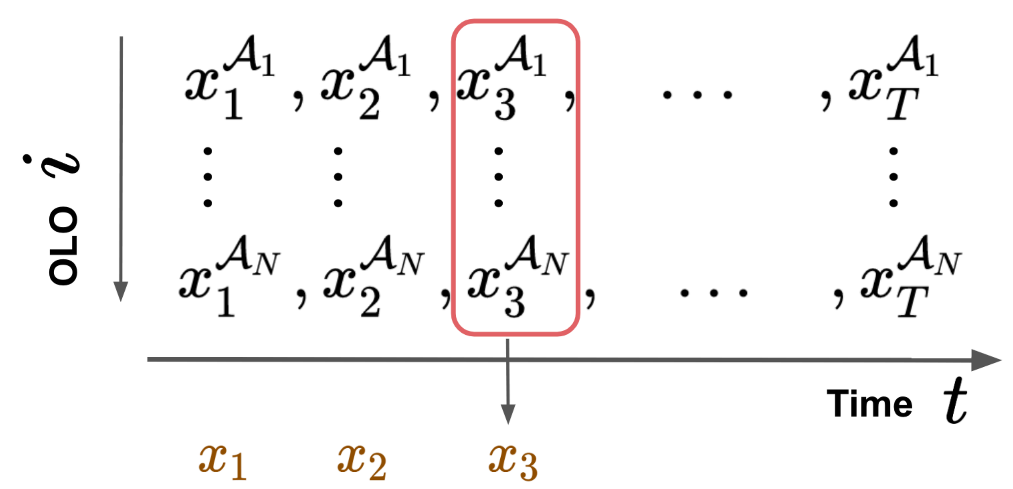

Algorithm	Regret	Per-round Cost	Feedback	Guarantee
Online-FW (Hazan and Kale, 2012)	$O(T^{3/4})$	$O(1)$	Full	deterministic
Meta-FW (Chen et al., 2018)	$O(\sqrt{T})$	$O(T^{3/2})$	Stochastic	in expectation
MORGFW (Xie et al., 2019)	$\tilde{O}(\sqrt{T})$	$O(T)$	Stochastic	w.h.p.
Our results	$\tilde{O}(\sqrt{T})$	$O(\sqrt{T})$	Stochastic	w.h.p.

Online Boosting for Regression

OLO → ***Weak*** *learner*
OCO → ***Strong*** *learner*
Stochastic *grads* → ***Bandit*** *feedback*

References

Lin Chen, Christopher Harshaw, Hamed Hassani, and Amin Karbasi. Projection-free online optimization with stochastic gradient: From convexity to submodularity. In *International Conference on Machine Learning*, pages 814–823, 2018.

Jiahao Xie, Zebang Shen, Chao Zhang, Hui Qian, and Boyu Wang. Stochastic recursive gradient-based methods for projection-free online learning. *arXiv preprint arXiv:1910.09396*, 2019.

Elad Hazan and Satyen Kale. Projection-free online learning. In *29th International Conference on Machine Learning, ICML 2012*, pages 521–528, 2012.

Alina Beygelzimer, Satyen Kale, and Haipeng Luo. Optimal and adaptive algorithms for online boosting. In *International Conference on Machine Learning*, pages 2323–2331, 2015.

Shang-Tse Chen, Hsuan-Tien Lin, and Chi-Jen Lu. An online boosting algorithm with theoretical justifications, 2012.

Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.

Vitaly Feldman. Distribution-specific agnostic boosting. *arXiv preprint arXiv:0909.2927*, 2009.

Varun Kanade and Adam Kalai. Potential-based agnostic boosting. In *Advances in neural information processing systems*, 2009.