Early stopping for kernel boosting algorithms

Yuting Wei*, Fanny Yang*, Martin Wainwright
Department of Statistics and EECS, University of California, Berkeley

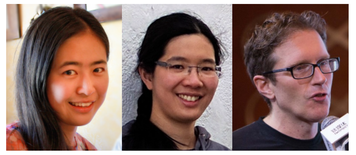

Problem Setting

- Given **arbitrary regular loss function** $\phi : \mathbb{R} \times \mathbb{R} \to [0, \infty)$, n fixed covariates x_i and corresponding random $Y_i \sim \mathbb{P}_{x_i}$
- Object of interest: minimizer of population loss function over some **function class** $\mathcal{F}$

$$\mathcal{L}(f) := \mathbb{E}_{Y_1^n}\Big[\frac{1}{n}\sum_{i=1}^n \phi(Y_i, f(x_i))\Big] \quad \text{and} \quad f^* := \arg\min_{f\in\mathcal{F}} \mathcal{L}(f)$$

- In practice: minimizer of empirical loss function based on observed $\{x_i, Y_i\}_{i=1}^n$

$$\mathcal{L}_n(f) := \frac{1}{n}\sum_{i=1}^n \phi(Y_i, f(x_i)) \quad \text{and} \quad \widehat{f} := \arg\min_{f\in\mathcal{F}} \mathcal{L}_n(f)$$

- If function class large $\mathcal{F} \to$ risk of **overfitting** to noise!
- Standard way to prevent overfitting: additive penalty function

Boosting algorithms

- Based on a sequence of additive updates (weak learners) to improve the fit of a function, see e.g. [1]
- Can be viewed as gradient descent on some $\mathcal{F}$ with updates

$$f^{t+1} = f^t - \alpha^t g^t \quad \text{with} \quad g^t \propto \arg\max_{\|d\|_{\mathcal{F}}\le 1} \langle \nabla \mathcal{L}_n(f^t),\, d(x_1^n)\rangle \qquad (1)$$

- We consider a **Reproducing Kernel Hilbert Space** (RKHS) $\mathcal{F} = \mathcal{H}$, with functions $f(\cdot) = \sum_{i=1}^n \omega_i \mathbb{K}(\cdot, x_i)$ with $\mathbb{K}$ kernel functions, e.g. Gaussian $\mathbb{K}(x,z) = \mathrm{e}^{\frac{(x-z)^2}{2\sigma^2}}$, Sobolev $\mathbb{K}(x,z) = 1 + \min\{x, z\}$

Early Stopping prevents overfitting

Running until convergence to $\widehat{f}$ again risks overfitting!
Illustration of boosting in functional space:

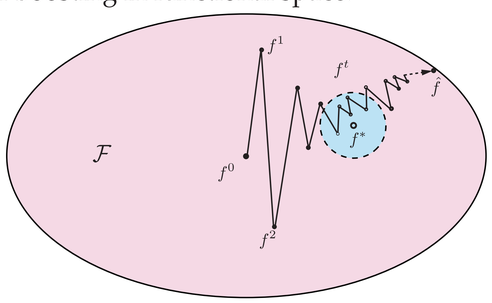

$\to$ better to stop at some **earlier time** $T(\mathcal{F}, n)$ when f^t closest to $f^\star$.
Early stopped boosting estimators are **consistent** (see e.g. [2,3])!

Early Stopped vs. penalty regularized

Distance to population optimum $f^\star$ over iterates:

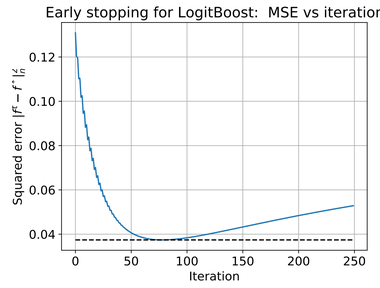

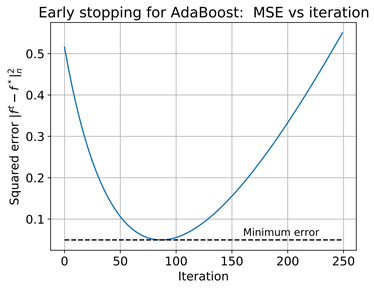

For **least-squares loss**, early stopped boosting (*algorithmic regular.*) and penalized estimators behave similarly statistically [4], i.e.

$$\|f_{\text{pen}} - f^*\|_n^2 \sim \|f^T - f^*\|_n^2$$

Main contributions

- Is there a common principle behind statistical behavior of algorithmic and penalized regularization?

  YES! We prove statistical rates for early stopping using ***the same key quantities as in penalized regularization***
- Can we extend to other loss functions?

  YES! Our new proof technique allows to extend to a ***broad class of loss functions (e.g. AdaBoost, LogitBoost ...)***

Key quantities and main theorem

- Key quantity I: **Localized Gaussian complexity**
  $\mathcal{G}_n\big(\mathcal{E}(\delta,1)\big) := \mathbb{E}\Big[\sup_{g\in\mathcal{E}(\delta,1)} \frac{1}{n}\sum_{i=1}^n w_i g(x_i)\Big], \;\; w_i \overset{i.i.d}{\sim} \mathcal{N}(0,1)$
  where: $\mathcal{E}(\delta,1) := \Big\{f - g \;\mid\; f, g \in \mathcal{H}, \|f-g\|_{\mathcal{H}} \le 1, \|f-g\|_n \le \delta\Big\}$
- Key quantity II: **Critical radius** δ_n smallest δ s.t. $\frac{\mathcal{G}_n(\mathcal{E}(\delta,1))}{\delta} \le \frac{\delta}{\sigma}$

Theorem 1. *Given some regular loss function ϕ and the iterates $\{f^t\}_{t=0}^\infty$ in (1), for all iterations $T = 0, 1, \ldots \lfloor 1/(8\delta_n^2)\rfloor$, the averaged estimate $\bar{f}^T$ satisfies w.h.p.*

$$\|\bar{f}^T - f^*\|_n^2 \le O\Big(\frac{1}{T} + \delta_n^2\Big).$$

Optimality

Early stopped boosted estimators are usually statistically optimal:

Corollary 1. *For the class of regular kernels and any function f^* with $\|f^*\|_{\mathcal{H}} \le 1$, running $T := \lfloor \frac{1}{8\delta_n^2} \rfloor$ iterations*

$$\mathbb{E}\|\bar{f}^T - f^*\|_n^2 \asymp \inf_{\widehat{f}} \sup_{\|f^*\|_{\mathcal{H}}\le 1} \mathbb{E}_{Y_1^n}\|\widehat{f} - f^*\|_n^2.$$

Examples for specific kernel spaces:

- γ**-exponential decay**: the kernel eigenvalues $\mu_j \le c_1 \exp(-c_2 j^\gamma)$, when stopped after $T \asymp \frac{n}{\log^{1/\gamma} n}$ steps:

$$\|\bar{f}^T - f^*\|_n^2 \lesssim \frac{\log^{1/\gamma} n}{n}$$

- β**-polynomial decay**: the kernel eigenvalues $\mu_j \le c_1 j^{-2\beta}$, when stopped after $T \asymp n^{2\beta/(2\beta+1)}$ steps:

$$\|\bar{f}^T - f^*\|_n^2 \lesssim n^{\frac{-2\beta}{2\beta+1}}$$

Numerical results

$\mathcal{H}$: first order Sobolev space, stop iterates after $T = (7n)^\kappa$

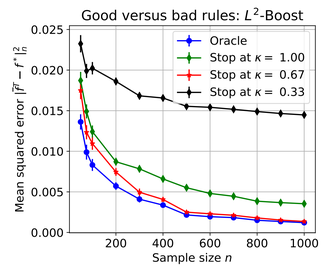

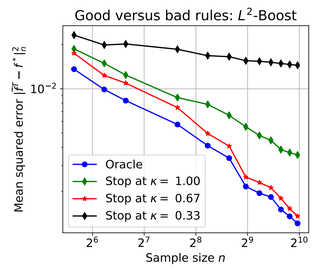

Log-scale: Error scales worse when stopped too early or too late

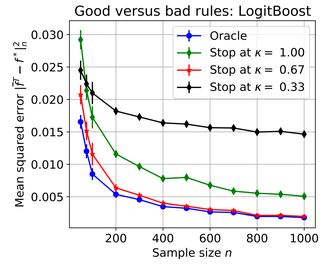

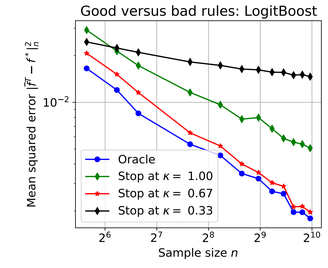

References

[1] R. E. Schapire, "The boosting approach to machine learning: An overview," in *Nonlinear estimation and classification*. Springer, 2003, pp. 149–171.
[2] T. Zhang and B. Yu, "Boosting with early stopping: Convergence and consistency," *Annals of Statistics*, vol. 33, no. 4, pp. 1538–1579, 2005.
[3] P. L. Bartlett and M. Traskin, "Adaboost is consistent," *Journal of Machine Learning Research*, vol. 8, no. Oct, pp. 2347–2368, 2007.
[4] G. Raskutti, M. Wainwright, and B. Yu, "Early stopping and non-parametric regression: An optimal data-dependent stopping rule," *Journal of Machine Learning Research*.