Minimax Rates for Continuum-Armed Bandits

A Function Space Perspective

Shashank Singh

Carnegie Mellon University, Pittsburgh, USA

↑Link to paper↑

Carnegie Mellon University

Introduction

- **Continuum-Armed Bandits: Optimize a smooth function f in as few queries as possible.**
- Also known as *blackbox*, 0^{th}*-order*, or *gradient-free optimization*.
  - Applications:
    - Hyperparameter tuning
    - Optimizing functions with unknown/expensive gradients
    - Nonconvex optimization
- **We provide the first tight analyses of continuum-armed bandits across several settings.**

Hölder and Besov Spaces

Definition 1 (Hölder Space)**.** *Let $s \in [0,\infty)$. Then, $f \in \mathcal{L}^\infty$ lies in the Hölder ball $\mathcal{C}^s$ iff*

$$\|f\|_{\mathcal{C}^s} := \sup_{\beta\in\mathbb{N}:\|\beta\|_1=\lfloor s\rfloor} \sup_{x\neq y} \frac{\left|f^\beta(x)-f^\beta(y)\right|}{\|x-y\|^{s-\lfloor s\rfloor}} < \infty.$$

For example, the case $s=1$ corresponds to Lipschitz continuity.

Definition 2 (Besov Space)**.** *Let $\beta_{j,k}$ denote coefficients of a function f in a wavelet basis. Let $\sigma \geq 0$ and $p,q \in [1,\infty]$. Then, $f \in \mathcal{L}^2$ lies in the Besov ball $\mathcal{B}^\sigma_{p,q}(L)$ iff*

$$\|f\|_{\mathcal{B}^\sigma_{p,q}} := \left\| \left\{ 2^{j(\sigma+D(1/2-1/p))} \left\| \{\beta_\lambda\}_{\lambda\in\Lambda_j} \right\|_{l^p} \right\}_{j\in\mathbb{N}} \right\|_{l^q} < \infty.$$

The case $\mathcal{B}^s_{\infty,\infty}$ is equivalent to the Hölder space $\mathcal{C}^s$.
The case $\mathcal{B}^\sigma_{2,2}$ is equivalent to the Sobolev space $\mathcal{H}^\sigma$.
The parameter q does not affect convergence rates, so we omit it in sequel.

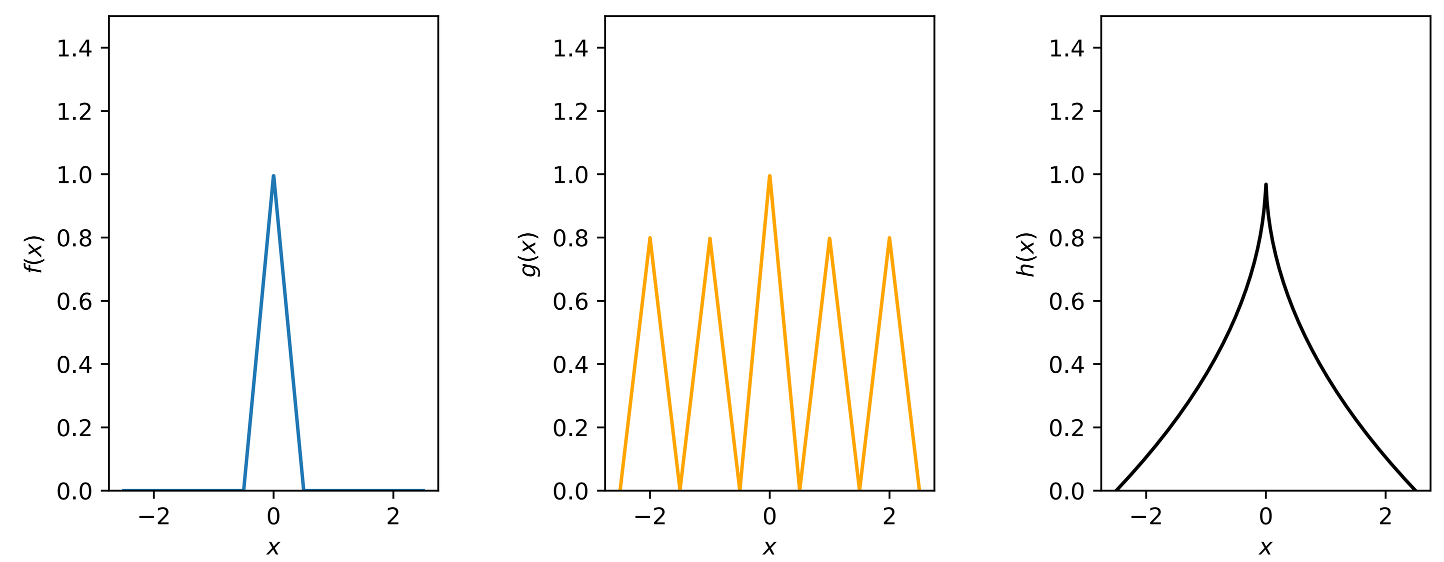

$\|f\|_{\mathcal{B}^\sigma_{p,q}} \ll \|f\|_{\mathcal{C}^s}$ $\quad$ $\|f\|_{\mathcal{B}^\sigma_{p,q}} \approx \|f\|_{\mathcal{C}^s}$ $\quad$ $\|f\|_{\mathcal{B}^1_{2,2}} < \infty, \|f\|_{\mathcal{C}^1} = \infty$

Theorem 1 (Besov Embedding [3])**.** *Let $\sigma \in (d/p,\infty)$, $p,q \in [1,\infty]$. Then, $\mathcal{B}^\sigma_{p,q} \subseteq \mathcal{C}^{\sigma-d/p}$.*

Theorem 2 (Besov Regression [2]; Informal)**.** *Regression over $\mathcal{B}^\sigma_{p,q}$ is easier than over $\mathcal{C}^{\sigma-d/p}$ (using spatially adaptive methods, e.g., thresholded wavelets [2] or adaptive splines [6]).*

Intuition: If $\|f\|_{\mathcal{B}^\sigma_{p,q}} \ll \|f\|_{\mathcal{C}^s}$, focus on estimating where f is non-smooth.

Simple and Cumulative Regrets

Simple Regret	Cumulative Regret
$R_S(\widehat{X},f) := \sup_{x\in\mathcal{X}} f(x) - f(\widehat{X}_T)$	$R_C(\widehat{X},f) := \sum_{n=1}^T \sup_{x\in\mathcal{X}} f(x) - f(\widehat{X}_n)$
$M_S(\mathcal{F}) := \inf_{\widehat{x},Z} \sup_{f\in\mathcal{F}} R_S(\widehat{x}_T,f)$	$M_C(\mathcal{F}) := \inf_{\widehat{x},Z} \sup_{f\in\mathcal{F}} R_C(\widehat{X}_T,f).$

Minimax Rates for Noiseless Case

Theorem 3 (Minimax Rates, Noiseless Case)**.** *Suppose we can query f exactly. For $\sigma > d/p$,*

$$M_S\left(\mathcal{B}^\sigma_{p,q}\right) \asymp M_S\left(\mathcal{C}^{\sigma-d/p}\right) \asymp T^{1/p-\sigma/d},$$

and

$$M_C\left(\mathcal{B}^\sigma_{p,q}\right) \asymp M_C\left(\mathcal{C}^{\sigma-d/p}\right) \asymp T^{1+1/p-\sigma/d}.$$

Proof Ideas:
Upper Bound: Follows from Besov Embedding and existing results for Hölder spaces [5].
Lower Bound: Use wavelet basis to construct a large family Θ of bump functions with disjoint supports. Show that, for any algorithm A, there exists $f \in \Theta$ such that, with high probability, A fails to find the support of f within n samples.

Minimax Rates for Noisy Case

Theorem 4 (Minimax Rates, Noisy Case)**.** *Suppose we can query f subject to IID additive sub-Gaussian noise with variance proxy η^2. For $\sigma > d/p$,*

$$M_S\left(\mathcal{B}^\sigma_{p,q}\right) \asymp M_S\left(\mathcal{C}^{\sigma-d/p}\right) \asymp \max\left\{ \left(\frac{\eta^2 \log\frac{T}{\eta^2}}{T}\right)^{\frac{\sigma-d/p}{2(\sigma-d/p)+d}}, T^{1/p-\sigma/d} \right\},$$

and

$$M_C\left(\mathcal{B}^\sigma_{p,q}\right) \asymp M_C\left(\mathcal{C}^{\sigma-d/p}\right) \asymp \max\left\{ T^{\frac{d+\sigma-d/p}{2(\sigma-d/p)+d}} \left(\eta^2 \log\frac{T}{\eta^2}\right)^{\frac{\sigma-d/p}{2(\sigma-d/p)+d}}, T^{1+1/p-\sigma/d} \right\}.$$

Proof Ideas:
Upper Bound: Follows from Besov Embedding and existing results for Hölder spaces [1, 4].
Lower Bound: Consider same family Θ from Theorem 3. Bound the maximum possible information gain at any possible step of the algorithm. Apply Fano's Lemma to lower bound error probability of identifying the true f.

Visualization of Rates

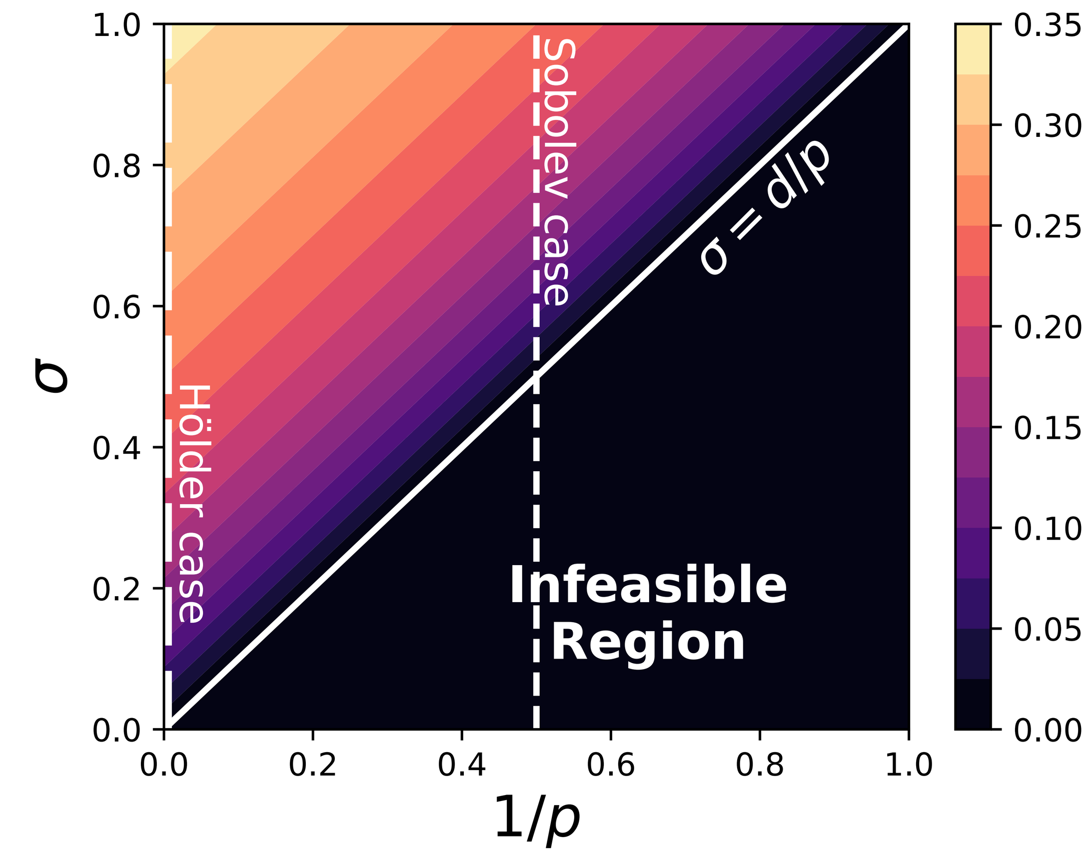

Phase diagram of minimax simple regret rates in the noisy case, as a function of σ and $1/p$, in the case $d=1$. Dashed lines indicate Hölder and Sobolev special cases. Contours (level sets of $\sigma - d/p$) show the relationship between Besov spaces and the embedding Hölder space.

1 Conclusions

1. First matching upper and lower bounds across all Besov spaces.
2. Spatially adaptive methods *do not* improve convergence rates for continuum-armed bandits.
3. When analyzing algorithms for continuum-armed bandits, it suffices to consider Hölder spaces, in which other spaces can simply be embedded.
4. In the Hölder case, our lower bounds improve existing lower bounds by polylog factors.
5. In the Sobolev case, our results improve existing results by polynomial factors and disprove a conjecture by Scarlett et al. [7].

References

[1] Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*. Springer, 2007.

[2] David L Donoho and Iain M Johnstone. Minimax estimation via wavelet shrinkage. *The Annals of Statistics*, 26(3):879–921, 1998.

[3] Dorothee D Haroske and Cornelia Schneider. Besov spaces with positive smoothness on rn, embeddings and growth envelopes. *Journal of Approximation Theory*, 161(2):723–747, 2009.

[4] Yusha Liu, Yining Wang, and Aarti Singh. Smooth bandit optimization: Generalization to h\” older space. *arXiv preprint arXiv:2012.06076*, 2020.

[5] Cedric Malherbe and Nicolas Vayatis. Global optimization of lipschitz functions. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

[6] David Ruppert and Raymond J Carroll. Theory & methods: Spatially-adaptive penalties for spline fitting. *Australian & New Zealand Journal of Statistics*, 42(2):205–223, 2000.

[7] Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. Lower bounds on regret for noisy gaussian process bandit optimization. In *Conference on Learning Theory*, 2017.