Proceedings of Machine Learning Research vol 178:1–2, 2022 35th Annual Conference on Learning Theory

Corruption-Robust Contextual Search through Density Updates

Renato Paes Leme RENATOPPL@GOOGLE.COM
Google Research NYC
Chara Podimata PODIMATA@G.HARVARD.EDU
Harvard
Jon Schneider JSCHNEI@GOOGLE.COM
Google Research NYC

Editors: Po-Ling Loh and Maxim Raginsky

Contextual Search is a fundamental primitive in online learning with binary feedback with applications to dynamic pricing (Kleinberg and Leighton, 2003) and personalized medicine (Bastani and Bayati, 2016). In each round, the learner chooses an action based on contextual information and observes only a single bit of feedback (e.g., "yes" or "no"). In the classic (realizable and noise-free) version, there exists a hidden vector $\theta^\star \in \mathbb{R}^d$ with $\|\theta^\star\| \leq 1$ that the learner wishes to learn over time. Each round $t \in [T]$ begins with the learner receiving a context $u_t \in \mathbb{R}^d$ with $\|u_t\| = 1$. The learner then chooses an action $y_t \in \mathbb{R}$, learns the sign $\sigma_t = \mathsf{sign}(\langle u_t, \theta^\star\rangle - y_t) \in \{+1, -1\}$ and incurs loss $\ell(y_t, \langle u_t, \theta^\star\rangle)$. Importantly, the learner does not get to observe the loss they incur, only the feedback. A sequence of recent papers (Amin et al., 2014; Cohen et al., 2016; Lobel et al., 2017; Leme and Schneider, 2018; Liu et al., 2021) obtained the optimal regret[1] bound for various loss functions, as highlighted on Table 1. The matching (up to $\log d$) upper and lower bounds in

Loss	$\ell(y_t, y_t^\star)$	**Lower Bound**	**Upper Bound**
ε-ball	$\mathbb{1}\{\lvert y_t^\star - y_t\rvert \geq \varepsilon\}$	$\Omega(d \log(1/\varepsilon))$	$O(d \log(1/\varepsilon))$ (Lobel et al., 2017)
absolute	$\lvert y_t^\star - y_t\rvert$	$\Omega(d)$	$O(d \log d)$ (Liu et al., 2021)
pricing	$y_t^\star - y_t \mathbb{1}\{y_t \leq y_t^\star\}$	$\Omega(d \log \log T)$	$O(d \log \log T + d \log d)$ (Liu et al., 2021)

Table 1: Optimal regret guarantees for realizable contextual search.

Table 1 indicate that the noiseless version of the problem is well understood. However, a lot of questions remain when the feedback is perturbed by some type of noise (as is often the case in practical settings). In the noisy model, the target value $y_t^\star = \langle u_t, \theta^\star\rangle$ is perturbed to $y_t^\star = \langle u_t, \theta^\star\rangle + z_t$. Most of the literature thus far has focused on stochastic noise models (Javanmard and Nazerzadeh, 2016; Cohen et al., 2016; Javanmard, 2017; Shah et al., 2019; Liu et al., 2021; Xu and Wang, 2021, 2022).

A recent trend in machine learning is the study of *adversarial noise models*, often also called *corrupted noise models*. In this model, most of the data follows a learnable pattern but an adversary can corrupt a small fraction of it. The goal is to design learning algorithms whose performance is a function of how much corruption was added to the data. For the ε-ball loss, we give a tight regret bound of $O(C + d \log(1/\varepsilon))$ improving over the $O(d^3 \log(1/\varepsilon)) \log^2(T) + C \log(T) \log(1/\varepsilon))$ bound of Krishnamurthy et al. (2021). For the symmetric loss, we give an efficient algorithm with regret $O(C + d \log T)$. Our techniques are a significant departure from prior approaches. Specifically, we keep track of carefully maintained density functions over the candidate vectors instead of a knowledge set consisting of the candidate vectors consistent with the feedback obtained.[2]

1. We use the terms "regret" and "total loss" interchangeably.
2. Extended abstract. Full version appears as http://arxiv.org/abs/2206.07528.

References

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 622–630, 2014. 1

Hamsa Bastani and Mohsen Bayati. Online decision-making with high-dimensional covariates. *Working paper, Stanford University*, 2016. 1

Maxime C Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 817–817. ACM, 2016. 1

Adel Javanmard. Perishability of data: dynamic pricing under varying-coefficient models. *The Journal of Machine Learning Research*, 18(1):1714–1744, 2017. 1

Adel Javanmard and Hamid Nazerzadeh. Dynamic pricing in high-dimensions. *Working paper, University of Southern California*, 2016. 1

Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 594–605. IEEE, 2003. 1

Akshay Krishnamurthy, Thodoris Lykouris, Chara Podimata, and Robert Schapire. Contextual search in the presence of irrational agents. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 910–918, 2021. 1

Renato Paes Leme and Jon Schneider. Contextual search via intrinsic volumes. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 268–282, 2018. 1

Allen Liu, Renato Paes Leme, and Jon Schneider. Optimal contextual pricing and extensions. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1059–1078, 2021. 1

Ilan Lobel, Renato Paes Leme, and Adrian Vladu. Multidimensional binary search for contextual decision-making. *Operations Research*, 2017. 1

Virag Shah, Ramesh Johari, and Jose Blanchet. Semi-parametric dynamic contextual pricing. *Advances in Neural Information Processing Systems*, 32, 2019. 1

Jianyu Xu and Yu-Xiang Wang. Logarithmic regret in feature-based dynamic pricing. *Advances in Neural Information Processing Systems*, 34, 2021. 1

Jianyu Xu and Yu-Xiang Wang. Towards agnostic feature-based dynamic pricing: Linear policies vs linear valuation with unknown noise. *arXiv preprint arXiv:2201.11341*, 2022. 1