Proceedings of Machine Learning Research vol 75:1–3, 2018 31st Annual Conference on Learning Theory

Efficient Convex Optimization with Membership Oracles

Yin Tat Lee YINTAT@UW.EDU
University of Washington

Aaron Sidford SIDFORD@STANFORD.EDU
Stanford University

Santosh S. Vempala VEMPALA@GATECH.EDU
Georgia Tech

Editors: Sebastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

We consider the problem of minimizing a convex function over a convex set given access only to an evaluation oracle for the function and a membership oracle for the set. We give a simple algorithm which solves this problem with $\tilde{O}(n^2)$ oracle calls and $\tilde{O}(n^3)$ additional arithmetic operations. Using this result, we obtain more efficient reductions among the five basic oracles for convex sets and functions defined by Grötschel et al. (1988).[1]

Keywords: Optimization, Membership, Separation, Violation, Validity, Subgradient

1. Summary

Minimizing a convex function over a convex set is a fundamental problem with many applications. The problem stands at the forefront of polynomial-time tractability and its study has lead to the development of numerous general algorithmic techniques. In recent years, improvements to important special cases (e.g., maxflow) have been closely related to ideas and improvements for the general problem (Christiano et al., 2011; Sherman, 2013; Madry, 2013; Kelner et al., 2014; Lee et al., 2013; Lee and Sidford, 2014, 2015; Lee et al., 2015; Sherman, 2017).

Here we consider the very general setting where the objective function and feasible region are both presented only as oracles that can be queried, specifically an *evaluation oracle* for the function and a *membership oracle* for the set. We study the problem of minimizing a convex function over a convex set provided only these oracles as well as bounds $0 < r < R$ and a point $x_0 \in K$ s.t. $B(x_0, r) \subseteq K \subseteq B(x_0, R)$ where $B(x_0, r)$ is the ball of radius r centered at $x_0 \in \mathbb{R}^n$.

It is well-known that with a stronger *separation* oracle for the set (and subgradient oracle for the function), this problem can be solved with $\tilde{O}(n)$ oracle queries using any of Vaidya (1996); Bertsimas and Vempala (2004); Lee et al. (2015) or with $\tilde{O}(n^2)$ queries by the classic ellipsoid algorithm Grötschel et al. (1988). Moreover, it is known that the problem can be solved with only evaluation and membership oracles through reductions shown by Grötschel, Lovasz and Schrijver in their classic book Grötschel et al. (1988). However, the reduction in Grötschel et al. (1988) appears to take at least n^{10} calls to the membership oracle. This has been improved using the random walk method and simulated annealing to $n^{4.5}$ Kalai and Vempala (2006); Lovász and Vempala (2006)

1. Extended abstract. Full version appears as [arXiv:1706.07357]

and Abernethy and Hazan (2016) provides further improvements of up to a factor of $\sqrt{n}$ for more structured convex sets.

Our main result in this paper is an algorithm that minimizes a convex function over a convex set using only $\tilde{O}(n^2)$ membership and evaluation queries. Interestingly, we obtain this result by first showing that we can implement a separation oracle for a convex set and a subgradient oracle for a function using only $\tilde{O}(n)$ membership queries and then using the known reduction from optimization to separation. We state the result informally below.

Theorem 1 *Let K be a convex set specified by a membership oracle, a point $x_0 \in \mathbb{R}^n$, and numbers $0 < r < R$ such that $B(x_0, r) \subseteq K \subseteq B(x_0, R)$. For any convex function f given by an evaluation oracle and $\epsilon > 0$, there is a randomized algorithm that computes a point $z \in B(K, \epsilon)$ such that*

$$f(z) \leq \min_{x \in K} f(x) + \epsilon \left(\max_{x \in K} f(x) - \min_{x \in K} f(x) \right)$$

with constant probability using $O(n^2 \log^{O(1)}(\frac{nR}{\epsilon r}))$ calls to the membership oracle and evaluation oracle and $O(n^3 \log^{O(1)}(\frac{nR}{\epsilon r}))$ total arithmetic operations.

See the full version for the formal statement, proofs, comparison to prior work, and further consequences of this result.

Acknowledgments

The authors thank Sebastien Bubeck, Ben Cousins, Sham M. Kakade and Ravi Kannan for helpful discussions, and Yan Kit Chim for making the illustrations. Finally, we thank Xiaodi Wu for pointing out some typos in a previous version of the paper.

References

Jacob D. Abernethy and Elad Hazan. Faster convex optimization: Simulated annealing with an efficient universal barrier. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2520–2528, 2016. URL http://jmlr.org/proceedings/papers/v48/abernethy16.html.

Dimitris Bertsimas and Santosh Vempala. Solving convex programs by random walks. *Journal of the ACM (JACM)*, 51(4):540–556, 2004.

Paul Christiano, Jonathan A Kelner, Aleksander Madry, Daniel A Spielman, and Shang-Hua Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 273–282. ACM, 2011.

Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Algorithms and Combinatorics, 1988.

A. T. Kalai and S. Vempala. Simulated annealing for convex optimization. *Math. Oper. Res.*, 31(2): 253–266, 2006.

Jonathan A Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 217–226. SIAM, 2014.

Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in o(sqrt(rank)) iterations and faster algorithms for maximum flow. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 424–433. IEEE, 2014.

Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 230–249. IEEE, 2015.

Yin Tat Lee, Satish Rao, and Nikhil Srivastava. A new approach to computing maximum flows using electrical flows. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 755–764. ACM, 2013.

Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1049–1065. IEEE, 2015.

L. Lovász and S. Vempala. Fast algorithms for logconcave functions: sampling, rounding, integration and optimization. In *FOCS*, pages 57–68, 2006.

Aleksander Madry. Navigating central path with electrical flows: From flows to matchings, and back. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 253–262. IEEE, 2013.

Jonah Sherman. Nearly maximum flows in nearly linear time. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 263–269. IEEE, 2013.

Jonah Sherman. Area-convexity, l_∞ regularization, and undirected multicommodity flow. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 452–460, 2017. doi: 10.1145/3055399.3055501. URL http://doi.acm.org/10.1145/3055399.3055501.

P. M. Vaidya. A new algorithm for minimizing convex functions over convex sets. *Math. Prog.*, 73: 291–341, 1996.