


1 Dear reviewers, first we want to say that your feedback was very helpful. We see a few genuine concerns and some
2 misunderstandings, so we appreciate the opportunity to address them. Specific issues raised by the reviewers (R1, R2,
3 R3, R4) have been organized into paragraphs below. Thank you! — The Authors.

4 **The gradient-free baseline should be (e.g.) BARON, not COBYLA (R2).** BARON is not applicable: a gradient-free
5 baseline must support either bilevel natively or must allow black-box objective and constraints (i.e., callbacks). This
6 also rules out CPLEX and Gurobi. IPOPT, BONMIN and Couenne are all gradient-based. YALMIP (MATLAB)
7 supports bilevel via reformulation as a single-level model, but imposes restrictions on how parameters appear in the
8 inner problem. Pyomo’s `bilevel` support is restrictive and deprecated. We would like to clarify that our aim was to
9 compare *local* NLP methods like COBYLA and SQP, not global methods like BARON, ARGONAUT [Boukouvola
10 et al., 2016], or DOMINO [Beykal et al., 2020]. Of the many gradient-free solvers provided by the well-known
11 NLOpt package, the documentation explains that only COBYLA supports arbitrary non-linear equality and inequality
12 constraints. (Likewise, the only such gradient-based solver in NLOpt happens to be SLSQP.) Finally, note that COBYLA
13 was used as a state-of-the-art local NLP baseline in the ARGONAUT and and DOMINO papers.

14 **Compare to OptNet/CvxpyLayers (R2, R3).** We agree that such a comparison is important for the gradient
15 computation component of our work. In fact, method SQP_{impl} (implicit differentiation) in our experiments was
16 intended to be this comparison: it is our efficient implementation of OptNet specialized to LPs, and performed
17 well. However, as `cvxpylayers` [Agrawal et al. 2019] can implement OptNet efficiently, we agree that a compar-
18 ison to that implementation is informative. We used `cvxpylayers` to implement method SQP_{cvx} , shown at right.
19 As in our previous experiments, we use the same vectorized LP generation,
20 solve in batch mode, enable efficient re-solves, and use double precision.
21 We confirmed that timing differences were due to the speed of the internal
22 `scs` solve (fast conic solver written in C) and not any overhead. The
23 only difference is that the default `max_iters` for `cvxpylayers` can
24 terminate long before reaching the requested precision (10^{-8} , the SciPy
25 default), but we were careful to ensure that this did not limit its maximum
26 achievable success rate. We found `scs` can be very slow when held to the
27 same precision standard as the rest of our experiments (dotted line). We
28 in no way “relied upon” [Agrawal et al. 2019] in developing our work.



29 **Compare to [Tan et al. 2019] (R1).** A comparison with their work was done but not highlighted due to two reasons:
30 their method is not applicable to objective error loss when learning cost with constraints, as it returns arbitrarily bad
31 solutions like w_1 in Figure 1 of the paper; their forward solver implementation is much slower and occasionally fails.

32 **Compare to specialized inverse methods (R2, R3).** To the best of our knowledge, there are no other inverse
33 optimization (IO) methods in the literature that can be directly applied to the problem instances we investigated. Most
34 IO methods make strong assumptions or address one aspect of the problem. Our approach is qualitatively more flexible.

35 **Scalability concerns (R2).** We are actually very excited about the scalability opportunities of our approach! First,
36 consider that the forward solver can be any specialized algorithm that returns primal and dual variables for the problem
37 at hand, not just a general-purpose LP solver; this means the forward solve takes a much smaller fraction of compute.
38 Second, with fast forward solvers the speed advantage of our closed form gradient becomes important, because solving
39 a large system (OptNet/CvxpyLayers-style) scales much worse than simply computing an outer product (our approach);
40 even on these modest instances, our closed form gradients are >100x faster than `CvxpyLayer.backward`. Other
41 well-known IO approaches, such as single-level models [Keshavarz et al., 2011], cannot exploit fast solvers for the
42 forward problem (i.e., the inner problem of our formulation) and must rely on general-purpose machinery like CPLEX.

43 **Isn’t Theorem 1 a special case of Appendix B in [Agrawal et al., 2019]? (R3).** No. Although LPs are a special case
44 of convex programs, our theorem applies only to objective error loss (not considered by [Amos & Kolter, 2017] nor
45 [Agrawal et al., 2019]) and cannot be derived by merely specializing the steps of their existing proofs.

46 **Prove SQP convergence to second-order stationary points (R2).** We agree that such a proof would be informative,
47 but characterizing the theoretical properties of a particular gradient-based NLP method (such as SQP) in this setting
48 was not our goal. Different parametrizations may or may not satisfy the technical conditions required (e.g., to guarantee
49 that B matrix remains a good Hessian estimate), but such analysis was not the focus of our work.

50 **The paper just applies SQP to an NLP, which is obvious (R2).** The IO formulation itself is our contribution and,
51 despite seeming “natural” (we agree!), has not been proposed before in either the `cvxopt`/OptNet body of work nor
52 in the IO-specific literature. So, the paper is not just about “applying SQP to an NLP.” Our framework provides the
53 most general approach for inverse LPs to date. Further, due to the importance of LPs and the existence of many fast
54 algorithms for specialized LPs, we believe that our fast gradient formula has long-term value.