

1 We thank the reviewers for their effort and their very detailed reviews.

2 Reviewer 1

3 **Q:** Comparison to probabilistic programming languages.

4 **A:** You are right, PPLs also allow flexible modeling and problem solving. We will add a review/comparison to PPLs  
5 and discuss the differences to our approach in the related work section.

6 **Q:** Why is GENO so much better than other solvers.

7 **A:** Other general purpose approaches/solvers (CVXPY) need to transform any problem instance into some standard  
8 form like an LP, QP, or SDP in standard form. The transformation increases in the problem size in terms of optimization  
9 variables and/or constraints. The transformed problems are typically addressed by Newton-type solvers that are very  
10 general but do not scale with the problem dimension.

11 **Q:** Running times, what is reported? Table 2 / comparison to CVXPY?

12 **A:** We will make this more clear. The reported timings include only the running times of the individual solvers.  
13 Generating the solvers by GENO takes only a few milliseconds and is not included into the timings. Anyway, the solver  
14 has to be generated only once. CVXPY was way too slow to be included with the other experiments in Figure 2. Hence,  
15 we ran a few much smaller problems for CVXPY and present the results in Table 2.

16 **Q:** Distinction between well-engineered and recent state-of-the-art solvers.

17 **A:** We consider solvers like LIBLINEAR and LIBSVM that are maintained for more than ten years well-engineered.  
18 Such solvers often outperform recently published solvers that implement new algorithmic ideas or adapt to some  
19 problem structure. We refer to the latter as recent state-of-the-art solvers.

20 **Q:** GENO does not transform a problem but a whole problem class?

21 **A:** In a problem class, parameters like a data matrix remain abstract while they are concrete (numerical values) in any  
22 problem instance. Traditional modeling languages take a problem instance and transform it into a problem instance of  
23 either a standard LP, QP, or SDP. Essentially, GENO takes a problem class that is specified by an objective function and  
24 constraints and transforms it into another problem class with smooth objective function but without constraints.

25 **Q:** Workflow example.

26 **A:** Initially, we had a workflow example but removed it due to space limitations. We will add it again to the appendix.

27 Reviewer 2

28 **Q:** Special structure in the problem.

29 **A:** The solver can exploit some special structure in the data. One can specify that a matrix is symmetric and/or sparse.

30 **Q:** Source code availability and license (also asked by Reviewer 3).

31 **A:** There is a link on the bottom of the anonymized GENO webpage that points to the github repository that contains  
32 the anonymized source code along with installation instructions. It seems, the reviewer has missed this link. We plan to  
33 make the original code available via github under the GPLv3 license.

34 Reviewer 3

35 **Q:** Limited practical impact. Highly efficient solvers have existed for years.

36 **A:** The purpose of GENO is to provide highly efficient solvers for new problems. Actually, our work on GENO is  
37 motivated from our repeated experience that there was no efficient readily available solver for some model that we were  
38 considering. Hence, we wanted GENO to be flexible but at the same time the generated code to be highly efficient. Only  
39 to prove our point that the generated code is indeed efficient we compared it to solvers that have existed for years. We  
40 also compared GENO to recently published solvers (NeurIPS 2018 and ICML 2019) for special problems (non-linear  
41 least-squares, compressed sensing). GENO outperforms these specialized solvers.

42 **Q:** Exploit convex duality for efficiency for efficiency.

43 **A:** Convex duality can only be exploited for convex problems. For this, the problem needs to be transformed into a  
44 standard form. The solvers Gurobi and Mosek exploit convex duality. However, they are a few orders of magnitude  
45 slower than our approach. It is usually believed that in order to be very efficient one needs a specialized solver. We  
46 challenged this claim by showing that one approach can be as efficient as specialized, well-established solvers.

47 **Q:** Deep learning problems out of scope.

48 **A:** As detailed in the introduction, there are a number of well-established, efficient deep learning frameworks. Such a  
49 framework was missing for classical ML. We want to fill this gap with GENO.

50 **Q:** GENO seems to be more convenient than, e.g., CVXPY (problem dimensions need not be specified).

51 **A:** This is not only a matter of convenience. A key aspect of GENO is that it generates a new solver for every problem  
52 class, while existing approaches like CVXPY transform a problem instance (where the dimensions and values are  
53 specified) into some standard form. This fundamentally different approach leads to an increase in efficiency of a few  
54 orders of magnitude.