We thank the reviewers for their valuable feedback, which will improve the paper. We will address all typos, errors, and clarity recommendations the reviewers suggest, turning now to the main concerns of each reviewer.

Response to reviewer 1:

Regarding the reviewer's comments about applications, we chose to limit the number of applications to three because of space limitations. However, we note that more applications are possible; for example we may consider any setting where a smooth sensitivity algorithm has been applied where our algorithms will offer similar advantages (e.g. statistics on graphs, Ullman and Sealfon 2019). We also note that all of our algorithms have relatively simple noise distributions with efficient implementation and so, as requested by the reviewer, we will add a more detailed explanation that explains the efficiency of our algorithms for each application we consider.

The reviewer recommends clarifying the comparison of our algorithms to the smooth sensitivity framework, which we sprinkled in different sections throughout the paper (probably haphazardly). We will add a unified discussion that explains the differences between the two frameworks and the advantages of the inverse sensitivity approach. Briefly, the smooth sensitivity algorithms may not be instance-optimal (as Asi & Duchi 2020 show) and require noise distributions with heavy tails (e.g. Cauchy, which has unbounded variance), in contrast to our mechanisms. This yields worse concentration—and hence performance—as, for example, our results for PCA demonstrate. Smooth sensitivity algorithms also require computing the smooth sensitivities, which may be challenging (for example, in linear regression).

As requested, we will add a discussion about related work on lower bounds for private mechanisms.

Response to reviewer 2:

For the reviewer's main comment on the contributions of this paper with regard to Asi & Duchi 2020, we believe that the approximation frameworks we develop in this paper significantly improve the applicability and scope of instance-optimal mechanisms. Indeed, the mechanisms of Asi & Duchi do not have efficient implementations for general functions and so currently apply only to limited settings. Our approximations, which rely only on the local sensitivities that are usually used in standard private mechanisms, open the door to several new applications of inverse sensitivity algorithms that enjoy instance-optimal behavior. Our applications provide a great example of this: the algorithms of Asi & Duchi 2020 do not have efficient implementation for these, in contrast to our algorithms, which have simple and efficient implementations. Moreover, the work of Asi & Duchi, while an important and interesting step to measure (instance) optimality in private learning and estimation, is essentially limited to 1-dimensional functionals, as it does not establish instance-optimal bounds for vector-valued functions or provide a family of mechanisms for such functions. Such general (vector-valued) functions are the main focus of this submission.

We thank the reviewer for bringing our attention to the Reimherr & Awan's K-norm mechanism (2019), which certainly is relevant, and we will add it to related work, as we were unaware of it. We briefly mention a few differences: their work provides asymptotic utility analyses, without finite sample guarantees on the performance of their algorithms, and they propose an approximate implementation of their mechanisms using an MCMC procedure without providing privacy guarantees for the implementation, which (in our view) is an important component of any putative privacy-preserving algorithm. We will discuss this work more carefully in the final version.

We will adopt all of the the reviewer's suggestions to improve the clarity of the paper in the final version.

Response to reviewer 3:

The reviewer discusses limitations of our first notion of instance-optimality against unbiased mechanisms by presenting an example where unbiasedness may not be desirable. Such weaknesses motivate the second notion of instance-optimality we consider, based on local minimax complexity. This is similar to classical statistics (e.g. Le Cam's local asymptotic normality), where local-minimax lower bounds alleviate the weaknesses of lower bounds for unbiased estimators. Many standard mechanisms (including smooth sensitivity algorithms and the mechanisms in this paper), satisfy our definition of unbiasedness, motivating a desire for instance-lower bounds for this large (and important) family of mechanisms.

We also agree with the reviewer's comment on the local-minimax definition as our main focus in this paper is primarily on ε-differential privacy. Extending this appropriately to (ε, δ)-DP is an interesting future question, which will likely require modifying the definition (especially the radius) and relying on lower bounds via fingerprinting codes (e.g. the work of Steinke and Ullman).

We will also add a discussion about related work for linear regression as the reviewer requests.

Response to reviewer 4:

We thank the reviewer for the comments on clarity and will address these in the final version of the paper.