

1 We thank the reviewers for their detailed comments and suggestions. We will address the minor presentation comments  
2 in the paper, and will address the remaining comments below.

3 **Reviewer 1:** For the first reviewer, we would like to emphasize that while the case of  $p = 2$  is central in practice, the  
4 case of  $p = 1$  is also very important, and in fact is often more desirable than  $p = 2$  for its robustness properties. For  
5 instance, in the rank regression estimator also considered in this paper, it is necessary that regression is carried out with  
6 respect to absolute deviation ( $p = 1$ ), and not  $p = 2$ . For  $p < 2$ , our results demonstrate an even larger improvement  
7 over the prior work, which had runtime that grew super-linearly in  $\sum_i \text{nnz}(A_i)$ . Moreover, in several interesting cases  
8 where the vector  $b$  itself is a Kronecker product (such as all-pairs regression), our  $p < 2$  algorithms do not even require  
9 the  $\text{nnz}(b)$  runtime, thus matching our results for  $p = 2$ . Additionally, regarding the remark made by reviewers 1 and 2  
10 about a comparison of our work with SGD, we note that getting provable relative-error regression guarantees with first  
11 order methods such as SGD would indeed require a condition number assumption on the matrices  $A_1, \dots, A_q$  (note  
12 that the condition number of  $\otimes_{i=1}^q A_i$  is the product of the condition numbers of  $A_i$ ). Our algorithms, interestingly, do  
13 not depend at all on the condition number of the  $A_i$ 's, and make no assumptions on these matrices.

14 **Reviewer 2** In response to the reviewer's worries about practicality and applicability of our results, we would like  
15 to reiterate that Kronecker products do arise naturally in many applications. Significantly, Kronecker product regres-  
16 sion comes up very frequently in tensor related-problems, which are relevant whenever the data has more than two  
17 associated dimensions. For instance, tensor regression is used centrally in [1] as well as in [2]. Kronecker products  
18 naturally arise when solving matrix equations, which show up in many different settings (see e.g. [4]). For instance,  
19 solving  $AX - XB = C$ , which is the Sylvester equation, can be written as  $(A \otimes B)\text{vec}(x) = \text{vec}(c)$ . Additionally,  
20 Kronecker products of more than two matrices arise when looking at partial differential equations such as a Poisson  
21 equation. Another important motivation to study Kronecker product regression is that a common way of solving low  
22 rank approximation (LRA) is via alternating minimization. Namely, if you want a LRA to a matrix  $A \in \mathbb{R}^{n \times n}$ , you  
23 can first fix some  $U \in \mathbb{R}^{n \times k}$  and solve for the  $V \in \mathbb{R}^{k \times n}$  that minimizes  $\|A - UV\|$  via regression. Once you have  
24 a  $V$ , then you solve for the best  $U$  given the  $V$ , and so on. The same procedure is used for third-order tensor low rank  
25 approximation, where you have  $A \approx U \otimes V \otimes W$  and use regression to solve for one of the factors at a time with the  
26 other two fixed, which is now a Kronecker product regression problem.

27 With regards to dependence on  $d$  in Theorem 3.1, we are indeed missing a factor of  $d$  in the definition of  $m$  – we  
28 will fix this typo. With respect to the reviewer's comment on the run-time of this theorem, we first note that the main  
29 computation taking place is the leverage score computation from Proposition A.3. For a constant number of input  
30 matrices  $q$  (as is generally the case in applications), the term  $d^{O(1)}$  in Proposition A.3 to approximate leverage scores  
31 is  $O(d^3)$ . The remaining computation is to compute the pseudo-inverse of a  $d/\epsilon^2 \times d$  matrix, which requires  $O(d^3/\epsilon^2)$   
32 time, so the additive term in Theorem 3.1 can be replaced with  $O(d^3/\epsilon^2)$ . This implies that even for  $q = 2$  matrices,  
33 if  $n > d^2$ , the size of the Kronecker product will be  $\Omega(d^4) \times d$ , which is larger than the runtime of our algorithm.

34 Additionally, the regime where  $n \gg d$ , known as the over-constrained regime, is motivated by many common situ-  
35 ations in practice, and as a result has been the focus of a large amount of algorithmic research over the past decade  
36 (as cited in our submission). The special case of Kronecker products is no different, and also frequently arises in the  
37  $n \gg d$  setting, such as [3] which considers very rectangular matrices (see their experiments in Table 5.1). This setting  
38 addresses another comment of the second reviewer regarding the relative sizes of  $\text{nnz}(b)$ , and  $\sum_i \text{nnz}(A_i)$ . Even in the  
39 above example with  $n$  only slightly larger than  $d$ , we would have  $\text{nnz}(b) = \Omega(d^4)$  but our runtime would be  $O(d^3)$ .  
40 The above shows that our algorithms become much more practical than both traditional linear algebraic algorithms  
41 and SGD (which can have a bad dependence on the condition number) even in a mildly over-constrained setting.

42 **Reviewer 3** With regards to the comment on the impact of our algorithms in applications, we first point to the above  
43 paragraphs which discuss settings in practice where our algorithms have been proven to perform substantially better  
44 than prior techniques. In addition, we remark that our experiments section demonstrates strong improvements on the  
45 run-time of our algorithms when compared to the prior work of Diao et al. This suggests that our algorithms are even  
46 faster than the provable guarantee of  $O(d^3)$ , which may be pessimistic because it is a worst-case theoretical bound.

47 [1] O. A. Malik and S. Becker. Low-rank tucker decomposition of large tensors using tensorsketch. In *Advances in*  
48 *Neural Information Processing Systems*, pages 10096–10106, 2018.

49 [2] O. A. Malik and S. Becker. Fast randomized matrix and tensor interpolative decomposition using counts sketch.  
50 *arXiv preprint arXiv:1901.10559*, 2019.

51 [3] G. Pisinger and A. Zimmermann. Linear least squares problems with data over incomplete grids. *BIT Numerical*  
52 *Mathematics*, 47(4):809–824, 2007.

53 [4] V. Simoncini. Computational methods for linear matrix equations. *SIAM Review*, 58(3):377–441, 2016.