We would like to thank all the reviewers for their thoughtful comments, especially given the difficult times. Your individual comments are addressed below (itemized by reviewer number and a short description of the comment).

R2 (Empirical usefulness): We will clarify parts of the introduction to make sure the paper does not promise more than it delivers. The goal of our paper is to establish a formal connection between min-hash based pairwise alignment and a rank-one crowdsourcing model and show how such a model can be efficiently solved using adaptivity. Our main motivating application is the pairwise alignment of third-generation (PacBio and ONT) sequencing data. Since these technologies produce noisy reads, most practical assembly pipelines employ min-hash based schemes to perform pairwise alignment [5,6,28] (citation numbers from the original manuscript). In particular, it has been previously shown by Baharav et al. [5] that spectral methods can be used to improve pairwise alignment accuracy via extensive experimentation on PacBio datasets from the NCTC 3000 project. Hence, our focus was on deriving a framework to compute confidence intervals for this spectral estimator, which allows bandit algorithms to be used to speed up the estimation. In the revised version, we will provide results on additional sequencing datasets and results based on the thresholding bandits, thus providing additional experimental validation of our approach. Moreover, following the comments of R4, we will be making our code available to maximize the potential practical impact.

R1 (Paper Organization): Following R1's suggestion, we will revise the paper for clarity and move some of the technical content from Section 3 into the appendix. This will free up space, which we will use for more experimental results. In particular, we will provide results on other PacBio datasets and results for the thresholding bandit algorithm.

R1 (Unspecified constants): While the constants were unspecified in our theorem statements, they can be computed explicitly based on the proofs in the appendix. In the revised paper, to maintain clarity, we will keep the constants unspecified in the main text, but will include detailed re-statements of the main theorems in the appendix with explicit constants. We will also clarify in the main text that these constants are nonasymptotic and point the reader to the appendix for their specific values. We will also include a remark explaining that the derived constants are quite loose, and that, in practice, a standard trick for improving performance is to first run the method on a dataset with ground truths to better approximate tighter (empirical) constants and then use these on the real problem.

R1 (Consistency of terminology): To avoid the back and forth between workers, questions, reads, and k-mers, we will discuss the connection with crowdsourcing in the Introduction, and then keep the rest of the discussion in terms of read alignments and hash functions.

R1, R3 (Experiments): Since the main purpose of the paper was to introduce adaptivity in the context of spectral estimation of pairwise alignments, the natural baseline method to compare with is the non-adaptive one. This emphasizes the gains obtained from adaptivity. We will include comparisons on two more datasets from the NCTC 3000 project (one of them, NCTC 4174, is shown in figures (b) and (c) below). These experiments focus on the gains provided by adaptivity, as the empirical usefulness of spectral methods for pairwise alignment had been previously studied in [5].

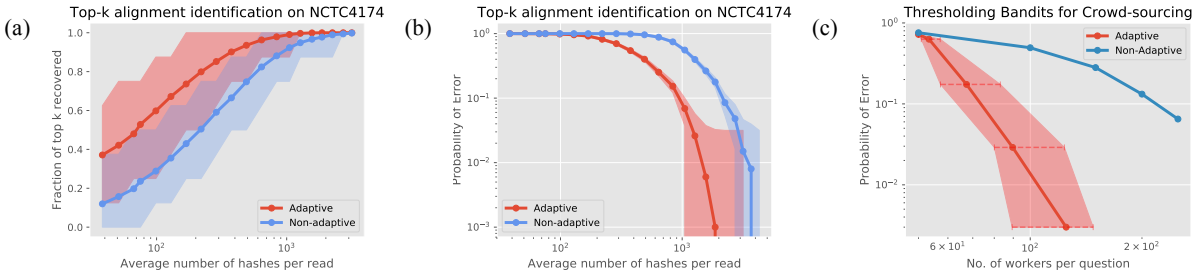

R1 (Related Works): We thank the reviewer for pointing us to Mash. We will add references and a short discussion regarding more recent alignment methods such as Mash and Mashmap in the revised version of the paper.

R4 (Thresholding Bandits): We thank the reviewer for the suggestion. We will present some of our discussion on threhsolding bandits in the main body of the paper, and we will include some thresholding bandits results in the Empirical Results section. In figure (c) above we show thresholding bandits results for the crowdsourcing problem (the set-up is identical to that of Fig. 2(c) in the paper but the task is to return a list that includes all products liked by at least 65% of the population and none liked by less than 50% of the population). As we see, the achievable points of the thresholding bandit algorithms are significantly better than the non-adaptive algorithm (the non-adaptive algorithm is the probability of error obtained at different fixed budgets and hence has no CIs around the number of workers).

R4 (Figure 2): We will change the ribbons to be 95% CI instead of 1 standard deviation in our plots.

R4 (Software): We thank the reviewer for the enthusiasm regarding software for our method. We plan on making our codebase publicly available once the review period is over.