We thank the reviewers for their feedback and suggestions. We appreciate this opportunity to clarify important aspects of our work and to describe proposed improvements to our manuscript that address the reviewers' concerns.

1. Intellectual fit at NeurIPS: We are very excited for the opportunity to present our work introducing the new application area of *adaptive electrode selection algorithms* to the NeurIPS community. NeurIPS has already been instrumental in advancing algorithms for applications in neuroscience and "5. Neuroscience and Cognitive Science" (subcategories: Brain Imaging and Brain–Computer Interfaces), and "2. Applications" are main categories in the call-for-papers. NeurIPS should now promote the area of adaptive electrode selection algorithms that comprise a growth area in neurotechnology. High-quality datasets are essential for developing new applications (e.g., ImageNet, MNIST, etc.). In this work, we collected and will publicly release a first-in-class, curated dataset. Therefore, along with introducing a new application area, our work provides a benchmark data set to support further development. Finally, we present CBS, which we show is a highly-effective, fast, reliable algorithm that successfully monitors almost all the available neurons. Thus, we propose our work has excellent intellectual fit with NeurIPS.

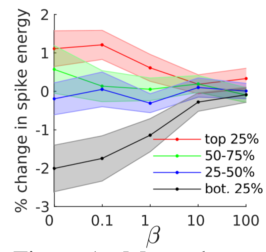

Figure 1: Mean change in recorded spike energy relative to unmodified CBS. Neurons are grouped via `TaskScore` quantiles.

2. Incorporating behaviorally-informative neurons: Multiple reviewers raise a very important point: electrode selection can also target behaviorally-informative neurons. We have now extended our work to prioritize neurons with greater task relevance, quantified by a metric denoted as `TaskScore`. `TaskScore` reweights contributions of neurons to the CBS objective: $S_w^\theta = \sum_{i=1}^{N} \widetilde{\alpha}_i \Sigma_i^\theta$ and $S_b^\theta = \sum_{i=1}^{N} \widetilde{\alpha}_i (\mu_i^\theta - \mu^\theta)(\mu_i^\theta - \mu^\theta)^T$, where $\widetilde{\alpha}_i = \alpha_i / \sum_j \alpha_j$ and $\alpha_i = \texttt{TaskScore}(i) + \beta$. $\beta \in [0, \infty]$ is a smoothing parameter that controls how much task-relevance is weighted. If $\beta = 0$, task information dominates selection and the weights are the task scores. When $\beta \to \infty$, we recover the original task-blind score and each neuron is weighted equally. We have implemented this extension. Figure 1 presents the results. When $\beta = 0$, the top 25% most task-relevant neurons were more heavily monitored, resulting in a 1.1% higher recorded spike energy, whereas task-irrelevant neurons (bottom 25%) were less densely monitored, incurring a 2% decrease in spike energy. Interestingly, while statistically significant, the effect of the bias towards task-relevant neurons was relatively small in our data, suggesting that task-blind CBS sufficiently monitors almost all available neurons. As intuited by the reviewers, our results suggest that task-relevant weighting will have larger impact when a smaller fraction of the available neurons can be recorded. We will include this extension as a supplementary section.

3. Significance and relation to earlier work: We agree with Reviewers #2 and #3 that Fisher's Linear Discriminant (FLD) has a long history and the connection to FLD should be explicit. We will amend the text appropriately. Nevertheless, as recognized by the reviewers, new applications, solvers, and approximations based on FLD can provide meaningful and important extensions to earlier work. Reviewer #2 points to the relevant work of Lei et al. 2012, who present a greedy approach to FLD-based feature selection. We will add this citation and discussion. We note that Lei et al. improve efficiency via a simplified criterion, while we use the full criterion and instead exploit 1) PCA to reduce waveform dimensionality and 2) banded structure for inverting S_w^θ. The differences between application areas should also be noted. In the Lei et al. study, the application (face recognition) leads to use of a simplified criterion, whereas this is not required for our application.

Reviewer #1 also brings up work on ECoG electrode selection (Saboo et al. 2019) that selects task-relevant electrodes without constraints, like in region-of-interest (ROI) estimation. We will add this citation. We note that we address a physical bottleneck that arises in Neuropixels and other high density arrays. Unlike in Saboo et al., we seek the jointly optimal "view" of the neuronal population under physical constraints. This differs from sub-selecting all highly relevant electrodes in the absence of constraints (i.e., wiring or bandwidth).

4. Success metric & use of training data: Reviewer #2 questions the focus on the # of neurons discovered and recommends matching sub-sampled waveform data to training templates as an alternative to metrics computed after re-running spike-sorting. We should first clarify that we focused on maximizing the # of discovered neurons because this is a fundamental goal which supports all subsequent population-level analyses. We further clarify that in one of our metrics, we find matches in spike timing between sub-sampled and densely recorded neurons. We also see the value of the suggested metric since it directly evaluates sub-sampled waveforms with their training templates. We thank the reviewer for their insight and will add this as part of the validation.

5. Bottleneck clarification: Reviewer #1 was concerned about the data bottleneck of spike-sorting and recommended fast spike-sorting methods. We appreciate the opportunity to clarify that the bottleneck we reference in that statement is not a computational bottleneck in spike-sorting, but a data-acquisition bottleneck in modern devices. We will clarify this point in the revision.

6. Clarifications: We appreciate all the reviewers' recommended points to clarify (i.e., notation, probe layout, median subtraction, alternative ascent strategies, etc.) We will update and improve the manuscript to address these points.