Thanks to all reviewers for their extensive feedback. The goal of this paper was two-fold, we proposed a new and important problem setting of active learning for combinatorial pool-based FDR control - a problem of tantamount importance in the sciences that has received little to no attention and gave a state of the art algorithm and analysis for pool-based active classification that achieves optimal bounds in several important settings. Achieving this required several technical advancements that show structural connections between bandits and active learning that has so far gone unnoticed in the machine learning community. We are excited to share these contributions and ideas with the NeurIPS community. We now dive into these ideas more carefully and address specific comments by the reviewers.

Reviewer #1: Thank you for the encouraging review. We will add a simplified statement of the theorem, which while not being optimal will illustrate the two key terms - the sample complexity of FDR verification combined with demonstrating the highest TPR set. The proof relies heavily on the sampling scheme and the choice of estimators. We will expand on this proof in the main body.

Reviewer #3: Thank you for your comments. We first address your concerns about significance. Utilizing a bandit-based analysis for active classification is a new strategy that promises to lead to new breakthroughs in active learning algorithms. In a recent work ([10]), the authors explicitly state that an pure exploration elimination strategy doesn't seem to address problems in active learning - a hurdle we overcome (partially due to not needing to sample each arm once) while giving improved sample complexity bounds in this setting. Vice versa, empirical process techniques could lead to new methods of dealing with union bounds in combinatorial bandit based algorithms.

This also relates to your concerns in points 3 and 4. Bounds in statistical learning theory based on VC (local) dimensions are well-understood and standard in the passive case to replace finite class union bounds with bounds over infinite classes. They have received less attention in the bandit and active learning literature. The expression given in line 178 and the result of Theorem 1 are novel contributiona and are motivated more thoroughly in the appendix (see A.1) and in the form of the confidence bound given there. We will move a summary of this discussion into the main body.

In the introduction we will clarify the plot adding a sentence like "... a particularly informative feature (Buried NPSA) are shown in each round for two different protein topologies (notated $\beta\alpha\beta\beta$, and $\alpha\alpha\alpha$)". We motivated algorithm 1 in lines 159-175, and algorithm 2 in lines 233-252 along with Figure 2. We will provide additional motivation emphasizing the algorithm simultaneously samples to guarantee FDR-control while removing sets with sub-optimal TPR.

Reviewer #4: Thank you for thorough review and comments. We understand that Theorem 2 can be hard to unpack, but we hoped the discussion in lines 267-282 along with Figure 2 would clarify it - a discussion we are happy to expand. We will provide additional settings where our algorithm will perform significantly better than uniform sampling and [4]. This includes cases when there are a large number of sets that can be FDR-controlled quickly, but then are eliminated based on TPR - a setting in which our confidence intervals are much tighter than those given in [4].

Regarding multiple testing, we explain in the related works section (lines 145-157) that this paper (though it borrows terminology from that domain) is not at all related to multiple testing. We are familiar with the works you referenced - and made a deliberate choice to not discuss them further, though we agree that we should cite them. The multiple testing setting has a different goal and less structure from ours - the goal there is to find which individuals $i \in [n]$ have means μ_i above a threshold. In our setting we are seeking to return a hypothesis in a fixed class that has high TPR under an FDR constraint.

We appreciate your concerns about the writing and will clarify the algorithm further. In the discussion before Algorithm 1, we were assuming that the reader was reading the discussion in 164- 175 simultaneously while viewing the algorithm display - we will signpost this better. Sampling in the symmetric difference is a strategy that has appeared before in [4,10], however, building estimators that do not require each item to be sampled once, work in an agnostic noise setting, an analysis based on gaps, and connections to VC dimension are all improvements on these works.

The estimator $\widehat{\mu}_{\pi',k} - \widehat{\mu}_{\pi,k} = \frac{n}{t}\sum_{s=1}^{t} R_{I_t,s}(\mathbf{1}(I_s \in \pi' \setminus \pi) - \mathbf{1}(I_s \in \pi \setminus \pi'))$ depends on *all* t samples up to the t-th rounds, each of which is uniformly and independently drawn at each step. Thus each summand is an unbiased estimate of $\mu_{\pi'} - \mu_{\pi}$. However, for π, π' active in round k, a summand is only non-zero if $I_s \in \pi\Delta\pi' \subset T_k$ hence we only need to observe $R_{I_t,s}$ if $I_t \in T_k$. This is equivalent to the rejection sampling procedure given: take a sample from an appropriate binomial distribution (not a geometric as pointed out) and then drawing that many indices from T_k and observing rewards from their associated distributions. We will say this more formally in the paper.

Regarding point 3 in the stylistic comments. Thank you for the math typo, indeed a minus sign is missing. Section 1.1 motivates active classification which we address in Section 3 with the goal of motivating sampling in the symmetric difference. This is an idea that is critical in Algorithm 2 for determining the set with the highest TPR. Addressing other style comments:1. We will change this to say something like"$Y_{I_t,t}$ is an independent R.V. For any i, $Y_{i,t}$ for all t are iid". 2. By "finite class", we meant "smaller class". 4. Thank you for the typos. 5. We mentioned Benjamini-Hochberg in the related works, and agree that citing the original work for this terminology is warranted.