We thank reviewers for their constructive comments on how to improve our paper.

Common Questions:

Regarding the lower bound: The main contributions of this paper are the two algorithms for the coarse ranking problem. Our lower bounds are used to *complement* the upper bounds: our upper bounds cannot be improved by more than a logarithmic factor even in the special case (recall that when $\mathbf{m} = (0, 1, n)$, the coarse ranking problem degenerates to the best arm identification problem). The lower bound for a general parameter vector $\mathbf{m}$ is an interesting open question for future study.

Regarding the clarity: In the next version of this paper, we will add a table of notations in the introduction and a conclusion. We will also give more intuition and technical overviews for the proofs.

Below are our responses to individual reviewers.

Reviewer #1: We thank the reviewer for the comments on our paper. We have already addressed the comments on the lower bounds and the presentation in "Common Questions".

Reviewer #2: We thank the reviewer for the comments on our paper. Regarding the comparison between our batched algorithm and the fully adaptive algorithm (LUCBRanking) in [30] in the fixed confidence setting, we would like to mention two items. First, as briefly mentioned in Line 83 of the submission, it seems difficult to adapt UCB-type algorithms to the batched setting, since the arm pulls in UCB-type algorithms are inherently sequential. Second, the complexity measure (instance complexity) in [30] is different from ours, and thus the sample complexities of the two algorithms are not directly comparable. We will add these discussions to the next version of this paper.

In Algorithm 1, yes, the reviewer is right that R should be part of the input.

In Line 12, we believe the current writing is correct. First recall that $I_r = I \setminus \left(\bigcup_{j=1}^{k} C_j^{(r)} \right)$ and $I_r = \hat{C}_1^{(r)} \cup \ldots \cup \hat{C}_k^{(r)}$. At the r-th round we partition E_r into k subsets $\left\{ E_r \cap \hat{C}_j^{(r)} \right\}_{j=1}^{k}$. According to Line 12 in Algorithm 1, $C_j^{(r+1)}$ is a superset of $C_j^{(r)}$, and $C_j^{(r+1)} \setminus C_j^{(r)} = E_r \cap \hat{C}_j^{(r)}$. Thus in the general case, $C_j^{(r+1)}$ is *not* equal to $C_j^{(r)}$.

Regarding the code of algorithms, the implementation of our algorithms has already been included in the supplementary materials of our submission (see the file `code-2904.zip`).

Reviewer #3: We thank the reviewer for the detailed feedback on our paper. We will address all the comments on the presentation in the next version.

In "Common Questions", we have already discussed the functionality of our lower bound results and some writing improvements that we will conduct in the next version.

On the comment about correctness, we thank the reviewer for pointing out this typo. The definition for E_r^* should be: "*Let E_r^* be the set of $(n_r - n_{r+1})$ arms in I_r with the largest gaps $\Delta_i^{\langle \mathbf{m} \rangle}(I)$*".

Regarding LUCBRank in the experimental studies, for the first question, yes, the reviewer is right. In each round of adaptivity, LUCBRank makes at most $2k$ pulls. In our experiments we set $k = 5$, and thus at time T the number of rounds is at least $T/(2k) = T/10$. For the second question, It is not the averaged number of rounds; it is the lower bound on the number of rounds.

Regarding the $\log(2n)$ factor, yes, the reviewer is right; we don't need this term. Thank you for pointing this out.

Reviewer #4: Regarding Line 51, we will add citations in the next version of the paper. For the best arm identification problem, both the fixed budget variant and the fixed confidence variant were considered in references [15, 28]. The fixed budget variant of the top-m problem was considered in reference [6].