We thank all the reviewers for their insightful comments and suggestions.

Reviewer 1. Thank you very much for your review and helpful suggestions. Detailed response is included below.

- "The results in main contribution section are not explained clearly, in particular about the relationship between n, δ, ϵ, and what is approaching zero/infinity.": Thank you for this comment. We'll add to each bullet point in the main contributions in the introduction a formal statement as in theorems 1-4 so that the relationships between the problem parameters are formally stated already at the introduction.
- "For given values of n, ϵ, and δ, what is the best algorithm to use according to your best knowledge...?": Generally speaking, ABALEH has best sample complexity whenever $n > 10^5$ and $n > 1/\delta$. When these conditions do not hold, the naive approach should be taken. SABA (which make assumptions on the input) and ABA are used mostly for didactic purposes to present and analyze the construction of ABALEH. We will discuss this in the body of the paper.
- Regarding all other comments: supplementary material compilation, line 206 and Hoeffding's bound. Thank you very much. We'll fix compilation and add the Hoeffding bound in the main body of the paper.

Reviewer 2. Thank you very much for your review and helpful suggestions. Detailed response is included below.

- "It would be interesting to see if using this algorithm is a subroutine improves the performance...": Agreed. That's an excellent idea and we would add such an analysis for gap elimination and other algorithms using MEDIAN ELIMINATION as a subroutine. Indeed, the complexity of such algorithms largely depends on MEDIAN ELIMINATION, thus as the results in appendix H. our algorithms will make a substantial improvements in these settings as well.
- "It would be useful if the authors summarized succintly the central insights that led to these results": This is a great idea, and we will add such a summary as a technical overview in the introduction.

Reviewer 3. Thank you very much for your review and helpful suggestions. The comments seem to stem from parts of the paper that were overlooked. Empirical evaluation was indeed performed and can be found in Appendix H and we included a discussion about instance-based analysis and why it is not applicable for the PAC setting studied here. We elaborate further in the comments below and hope you will consider revising your score based on this response.

- "though the analysis is deep and technical the result may not be immediately applicable. The range of δ and n proposed are not practical for many settings. If the authors really feel that this algorithm improves over naive elimination (which suffers a bad union bound over n), then they should demonstrate this empirically - in general experiments would have helped the paper.": Perhaps it has been overlooked, but **Appendix H is dedicated to empirical evaluation**. It shows dramatic benefit of ABALEH, even for reasonable values of n and δ.
- Regarding concerns when comparing to instance dependent results: **Please see paragraph starting on line 88 titled: "'From Instance-based to worst case analysis"**. In particular, one of the challenges with comparing PAC bounds to instance specific bounds, is that instance specific algorithms assume that n is **constant** and δ goes to zero, but do not have a simple closed form expression which is based only on n that determines the rate at which δ must go to zero to make the analysis work, and what happens at given (finite) values of n and δ. As a concrete example for how this is a problem, if $\delta \ll 1/n$, then we need to worry about the $\log 1/\delta$ more than about the $\log n$ in naive elimination. OTOH, if $\delta \ll 2^n$, we need to worry about the $\log(1/\delta)$ factor more than about the n factor in naive elimination - this is the n factor that instance optimal algorithms try to save in the first place (the biggest difference between instance optimal algorithms and worst case algorithms like ours is when there is a unique best arm, one arm which is almost ϵ-close, and $n-2$ arms which are always zero). Moreover, the gain in instance-based algorithms is bounded not just by n (which is assumed to be constant) but also by $1/\epsilon^2$, since their gain comes from the difference between one ϵ-far arm (which makes it difficult for the worst case algorithm) and the other arms that can be always zero, and the smaller epsilon is the largest this difference is. To summarize, both research directions on instance-based and worst case (i.e. PAC) learning algorithms are valid, but are useful for completely different parameter domains.

Reviewer 4. Thank you for your time and efforts.

- "...an empirical evaluation... could be included in the main text. Empirically, it would be useful to also show the performance of a track-and-stop algorithm (Ref [14]) and a upper confidence bound algorithm (Ref [17]).": Thank you for this comment. We will revise the manuscript to include the empirical evaluation of appendix H in the main body of the paper. Regarding the track-and-stop algorithm, we ran code provided us by authors of these papers. Unfortunately, the track-and-stop algorithms can only be run for small values of n (roughly $n = 100$). We discuss the inherent implementation and running time bottlenecks of track-and-stop in appendix G.
- "line 3 of algorithm 2: shouldn't it say "$\lfloor |A_i|(\delta + \phi(n)) \rfloor$"": The line as written is technically correct but writing it as you suggest is clearer and we will change it – thank you.
- "third argument of NaiveElimination: should it be "δ" in line 181": True. Thank you.