Reviewer 1 - (1). The reviewer's comments show a misunderstanding concerning what is achieved by our protocol and what is achievable by differential privacy (DP). In our model, a data holder wants to score an input x_i against a model M held by another party (model holder) such that, at the end of the protocol, no information about the input is leaked to the model holder (beyond the result of the classification) and no information about the model should leak to the data holder. Information about the model and the data also *should not be available to any other party involved in the computation*. Differential privacy is not useful in this scenario. In the central DP set-up, the data collector accesses the entire data set. Upon receiving a question, the data collector computes an answer based on the data set, adds noise to the answer and sends it to the party asking the question. In this case, unlike with the SMC approach, there is loss of accuracy because of the noise, and more importantly, all information, including the question, the dataset and the answer, is leaked to the data collector. In the local DP set-up, data owners add noise to their data entries and send them to a third party (the model holder / data collector). The data collector uses the noisy entries $(X = x_1, x_2, ...x_n)$ to answer a question $A(X)$, x_i represents the local data plus noise. The overall system is said to be differentially private if the view of the data collector does not change much if the data set is modified in just one entry. While in this scenario, privacy for the local inputs is possible, *the question A cannot depend solely on a single entry* by the very definition of differential privacy. DP cannot be used to single out individual "bad" entries. DP deals with global characteristics of the data set X. Our solution, on the other hand receives a single entry x_i and outputs the classification of x_i such that the only information leaked to the model holder is the class label and no information is leaked to the data holder at all, while having no loss in accuracy and without any trusted entity receiving information about the model or the input data. It is misleading to directly compare DP with SMC. They are different tools for achieving different notions of privacy in different situations. In many ways, they complement each other.
(2). It is an intrinsic characteristic of SMC that the function to be computed privately has to be represented as a circuit of addition and multiplication gates. Therefore, one has to come up with specific circuits and optimizations for each ML technique and algorithm. We have provided a fairly general solution to an important problem: text classification. Developing SMC protocols for all SOTA ML models is far beyond the scope of a single paper. Our work is the very first SMC based method for text classification. While we recognize that deep learning is SOTA for many NLP tasks, LR models on word n-grams for hate speech detection have been observed to be at par with CNN and LSTM model architectures (cfr. [33] Gröndahl et al., AISec 2018). Furthermore, we fully agree with reviewer 3 that many of the building blocks presented in the paper can be re-used or inspire the development of similar SMC building blocks for other kinds of ML models, stimulating future research as also pointed out by reviewer 2. We remark here that we have proved (in the appendix) that our building blocks can be securely composed.
(3). We disagree with the reviewer that the results in terms of accuracy and time are not good. The accuracy with SMC is the same as the accuracy in the clear (i.e. for the "non-private method"), in other words there is no loss of accuracy. The classification time is two orders of magnitude better than that of the only existing solution for privacy-preserving text classification and it comes with rigorous proofs of security. As pointed out, fast DP based solutions are of no help here, since the result of the classification depends solely on a single entry of the data set and the input data and the model should remain private.

Reviewer 2 - We thank reviewer 2 for his/her insightful comments. Indeed, to the best of our knowledge, we have designed the first efficient provably secure protocols for doing private text classification of individual entries. Making sure that all the computations happen over a ring, rather than a field, helped us to reduce the round complexity of our solution. We will cite works of SMC applied to other areas of ML, including clustering.

Reviewer 3 - We thank reviewer 3 for his/her careful review of our work. The reviewer's understanding of our paper is correct. The reviewer is correct in pointing out that the running time of the proposed protocols (seconds) is still higher than in the clear (a few milliseconds). SMC implementations are still substantially slower than solutions in the clear. Moreover, from an implementation perspective, we have paid a price in making our implementation modular and in Java. We probably could have decreased the running time by implementing everything in assembly and C. An association of Yao Garbled Circuits and Secret Sharing SMC combined with optimal ML protocols could give us an improvement over what we present here (but is beyond the scope of the paper). Improving those running times will demand substantial work from the ML and Cryptography communities, preferably in association. As it is, the biggest bottleneck in our solution is the private feature extraction. Coming up with improvements for it would have a big impact in our protocols. The relation between DP and our work is an interesting one. Locally-private DP offers the possibility of obtaining statistics and ML models trained over a data set distributed over many owners in a private way. It comes at a cost: there is no way to classify individual entries (see answer to Reviewer 1) and there is a substantial loss of accuracy, particularly when the number of entries is not large. However, it is fast. SMC gives us the possibility of scoring single entries against a ML model privately without any loss in accuracy. However, it is slow. We plan on investigating a mix of these two techniques, where SMC is used to decrease the noise/loss in the locally-private ML model. We thank the reviewer for pointing out this discussion topic and we plan on adding it to the paper. We believe that SMC and DP are complementary solutions; bringing the promise of privacy-preserving ML to practice will require an association of these and other paradigms. Exposing the ML community to works like ours is a necessary step towards that direction.