CS-523 - Final 2021
Most Repeated Errors

July 16, 2021

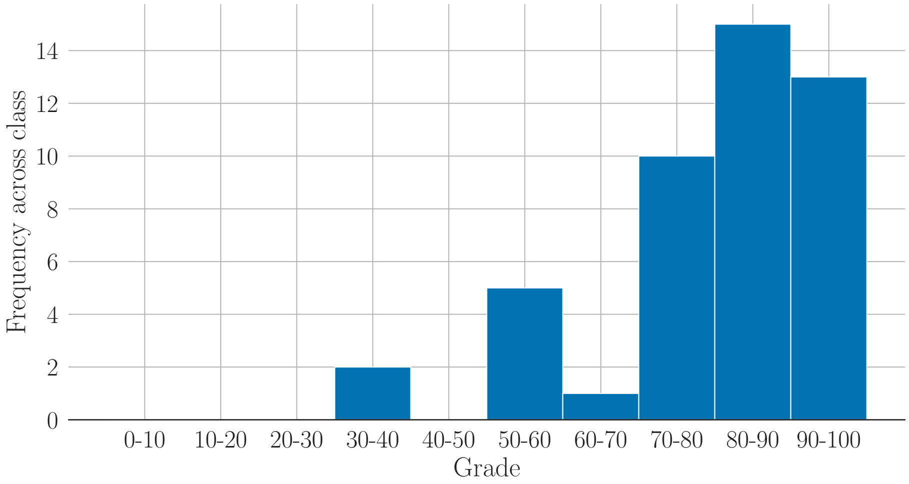

Figure 1: Distribution of grades across class

Question: Worker union at XenaWarrior

Error 1: Provide a censorship-resistant instead of answering the question Many answer to Part 2 do not answer whether the system proposed in Part 1 provides censorship resistance. Instead, they proposed a new system that could be used to evade the censor. These answers were not given full points as they did not answer the question.

Error 2: Not explain how the system should be used to achieve anonymity or censorship resistance Many students forgot to answer the part of the question about how unionizers should use the system. These answer were not given full points.

Error 3: Claims about properties or systems without explanations Many students made claims about properties (e.g., ”sender anonymity guarantees relationship anonymity”); or systems (e.g., ”Crowds does not provide anonymity”) without a justification in the context of the adversary model. When one makes such a claim, one needs to justify it (e.g., ”Crowds does not provide relationship anonymity because NSA is a global adversary and can trace messages through the network”). Unjustified claims were not counted as valid answers.

Question: Pear4Science

Error 1: in Part 1, mention the risk of membership or property inference attacks on the shared gradients without further explanation. The first part of this question asked you to describe what a potential adversary might learn and how. Many answered this question by mentioning the risk of common attacks, such as membership or property inference, without any description of how this might work or why this would constitute a privacy violation. This was not enough to get full points on this part of the question.

Given the simplicity of the model (see also below), it is unclear how complex attacks based on machine learning techniques would work.

Error 2: in Part 1, miss that the gradient $g = x$ directly reveals an individual user's coffee consumption. The algorithm described in the question leads to an obvious privacy violation where each gradient directly reveals a user's coffee consumption on this day as $\frac{\partial}{\partial a} L = x$.

Error 3: in Part 2, argue that the privacy guarantee of differentially private noise addition is dependent on group size. The protection that the differentially private noise addition provides against attacks on the published statistics does not depend on the size of the demographic group a user is in. It is determined by the value of the privacy parameter ε.

Error 4: in Part 2, re-state the general privacy-utility trade-off instead of giving a concrete disadvantage of the differentially private data aggregation. It was not enough to state that there is a trade-off between the accuracy of the published statistics and the privacy of individual users to get full points on the second part of this question. The existence of the trade-off was already part of the question. We asked for a potential disadvantage for the utility of the data in this specific scenario. For instance, the effect that the noise addition disproportionally affects minority subpopulation, i.e., that the ”poor get poorer”.

Question: Online Tracking

Error 1: Not specifying whether FloC IDs are saved in cookies. Part 1 of the question can be answered either using cookies (by saving the ID value in cookies and then performing redirects) or skipping them (reading the ID and sending requests with the ID in the URL). This also influences Part 2. Some of you mentioned that you will use cookie syncing to share the ID, but did not state how cookies were used in your method. For example, are the FloC IDs saved in the cookies. Cookie syncing does not only involve sharing values via redirects, the value has to be stored in a cookie. Note that some of you used cookies to store user-specific information and shared this along with the ID; this was accepted as correct.

Question: Private Set Intersection

Error 1: in Part 1, not specify the circuit using the operations we have given (binary OR, binary AND, unary NOT). As stated in the question, the response had to build the circuit using the basic operations we have provided. It was not enough to describe the high-level boolean function to be computed as there are multiple ways to build the circuit representing this function; such a response got partial grade.

Error 2: in Part 2, not backing up the security argument. The security argument had to be backed up by, e.g., the security of Garbled Circuits. Claiming only that the client cannot learn more than needed without any explanation about why this was the case in the adversarial model of the question was not enough to get the full grade.

Error 3: in Part 2, assume the client could get access to the intermediate values of the circuit, or to the inputs of the server. Such knowledge goes against the security model of Garbled Circuits. Without an argument why the client could have gotten access to these values, such response did not get a full grade.