Week 5 Exercises (ECE 598 DA)

Exercise (Private Mean Algorithm): Prove formally that the Private Mean Algorithm (in the lecture notes) can be instantiated with Gaussian noise, instead of Laplace, and would satisfy (ε, δ)-DP and derive the $(1-\beta)$ high-probability error bound.

Exercise (DP Mechanism for Variance): Show that if $x_i \in [0,1]$ for all $i \in [n]$, the global sensitivity of $S_2 = \sum_{i=1}^{n} x_i^2$ is 1 and derive an $(\varepsilon, 0)$-DP mechanism for variance by appropriately noising S_2 and $S_1 = \sum_{i=1}^{n} x_i$. Why might the resulting DP variance estimator be biased? How does increasing n affect this bias?

Exercise (Gradient Clipping): Why is gradient clipping necessary in DP-SGD? Consider removing the clipping step and adding Gaussian noise to the full (unclipped) batch gradient. Would the mechanism still be differentially private? Discuss what could go wrong if we do not clip, especially when some individual's gradient is extremely large.