Shifted Interpolation for Differential Privacy*

Jinho Bok
UPenn
jinhobok@upenn.edu

Weijie Su
UPenn
suw@upenn.edu

Jason M. Altschuler
UPenn
alts@upenn.edu

February 29, 2024

Noisy gradient descent and its variants are the predominant algorithms for differentially private machine learning. It is a fundamental question to quantify their privacy leakage, yet tight characterizations remain open even in the foundational setting of convex losses. This paper improves over previous analyses by establishing (and refining) the "privacy amplification by iteration" phenomenon in the unifying framework of f-differential privacy—which tightly captures all aspects of the privacy loss and immediately implies tighter privacy accounting in other notions of differential privacy, e.g., (ε, δ)-DP and Rényi DP. Our key technical insight is the construction of *shifted interpolated processes* that unravel the popular shifted-divergences argument, enabling generalizations beyond divergence-based relaxations of DP. Notably, this leads to the first *exact* privacy analysis in the foundational setting of strongly convex optimization. Our techniques extend to many settings: convex/strongly convex, constrained/unconstrained, full/cyclic/stochastic batches, and all combinations thereof. As an immediate corollary, we recover the f-DP characterization of the exponential mechanism for strongly convex optimization in Gopi et al. (2022), and moreover extend this result to more general settings.

*The full version of the paper is available here: https://arxiv.org/abs/2403.00278.