Consistent Closures: A Novel Methodology for Analyzing Privacy Definitions

Daniel Kifer, Penn State University

In this talk we introduce a novel methodology for analyzing statistical privacy definitions. When sanitizing data for public release, it is important to understand what privacy protections are being guaranteed by a chosen privacy definition. Although many privacy definitions exist, it is not always clear what they guarantee and when they can be used. Often, privacy definitions are evaluated by using specific attacks such as linkage to external records. However, these approaches tend to be brittle as it is easy to design algorithms that leak private information yet evade attacks encoded in a fixed set of software libraries.

Instead of directly attacking the sanitized data generated by algorithms that satisfy a given privacy definition, we take a different approach. We propose a methodology, called the consistent closures methodology, that analyzes a privacy definition directly. The first step of our methodology rewrites the privacy definition in a normal form, called its consistent closure, which is easier to analyze. The second step then extracts the privacy guarantees -- statements relating an attacker's prior and posterior beliefs after observing sanitized data.

We apply our methodology to tuple-perturbation definitions such as randomized response, PRAM, and FRAPP to bring new insights into what these privacy definitions protect -- we show that their privacy protections are equivalent to protecting the parity of a dataset. We also apply our methodology to partition-based privacy definitions like k-anonymity to demonstrate that the consistent closure of a privacy definition can often quickly identify its weaknesses.