Foundations for Fair Algorithmic Decision Making

Krishna P. Gummadi*

*Max Planck Institute for Software Systems (MPI-SWS), Allemagne
gummadi@mpi-sws.org,
https://people.mpi-sws.org/ gummadi/

Summary

Algorithmic (data-driven learning-based) decision making is increasingly being used to assist or replace human decision making in a variety of domains ranging from banking (rating user credit) and recruiting (ranking applicants) to judiciary (profiling criminals) and journalism (recommending news-stories). Recently concerns have been raised about the potential for discrimination and unfairness in such algorithmic decisions. Against this background, in this talk, I will discuss the following foundational questions about algorithmic unfairness :

1. How do algorithms learn to make unfair decisions ?
2. How can we quantify (measure) unfairness in algorithmic decision making ?
3. How can we control (mitigate) algorithmic unfairness ? i.e., how can we re-design learning mechanisms to avoid unfair decision making ?