Sam Straus

The Reconstruction in Causal and and Bayesian Invference

For my course, Stochastic Processes on Graphs (graphical models), I read papers and organized material and gave a survey presentation on recovering directed acyclic graph (DAG) structure from graphical models. In my presentation, I covered the theroy of causality following the work of Judea Pearl, and concluded by presenting a paper establishing bounds on the number of samples needed to infer structure from Bayesian nets.

Causality inference is, in my opinion, one of the most commonly misrepresented terms in science, and even by using DAGs "in the obvious way" the theory we usually refer to still doesn't entirely allow us to make the types of inferences we want to allow. This is because causal inference, in some sense corresponds to applying an external change to a system, and guessing what will happen in the new, changed, system. However, the new system represents a fundamentally different probability distribution than the one you measured when you conducted your measurements. A good example is if in a baseball game, if the ball goes out of the park it causes the teams score to go up, but if I apply an outside change to the system and cause the ball to go out of the park by storming the field, running away with the ball, and throwing it over the fence, it does not cause the score to go up. This is because by applying the change, I'm changing the environment of what should happen, and the assumptions of the system no longer hold. I presented a formalization that allows us to talk about these types of inferences precisely, and presents a path towards inferring from data things that cause each other.

In addition, I presented work that covered learning bayesian networks more generally (not just causality). This work showed in many ways that constructing graphs by adding a penalty term via graph complexity was an effective way to solve the problem (convergence in limit), and estimated the number of samples needed to achieve accurate estimates.