Bayesian networks (or directed acyclic graphical models) have been a cornerstone of the modeling of causal inference. They have also led to the study of rich new combinatorial notions such as d-separation and Markov equivalence. Although these structures have been studied extensively since at least the 80s, many natural combinatorial questions about them—including basic ones regarding counting and uniform sampling—are still not fully understood. This talk will be a broad survey of some such problems, and recent progress on them.

While a small part of this talk will be based on joint work with collaborators, most of it will be a survey of work by others, and will be based on discussions with—and work of—Vidya Sagar Sharma, a PhD student at TIFR.