


An Introduction to Probabilistic Graphical Models (Jordan, 2003)

Chapter 2: Conditional Independence and Factorization

Satvik Saha (ss7228)

October 14, 2025

This chapter introduces a framework for moving back and forth between a (potentially very complicated) joint probability distribution and its graphical representation. One advantage of the latter is that assertions of *conditional independence* are much easier to read off (almost a triviality in *undirected* graphical models): the task of carefully factorizing a joint distribution is translated into examining *graph separations* between groups of nodes. This can be done via a *Bayes ball algorithm*: two groups of nodes are conditionally independent given a third group if ‘information’ cannot travel along the graph between the former groups, according to a simple set of rules. In an undirected graphical model, this reduces to checking connectedness between the former groups after removing the latter!



For example, in the above (Gaussian mixture) model, we may conclude that each $z_i \perp z_j \mid \mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\lambda}, \pi$ for $i \neq j$, simply by observing that the paths between z_i and z_j are blocked. Furthermore, $\pi \perp \{\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\lambda}\} \mid \mathbf{z}$.

This is often useful when writing full conditionals for Gibbs sampling: the full conditional of a node depends solely on its *Markov blanket*, consisting of parents, children, and co-parents! This tells us that the updates for

- π involve only $\mathbf{z}$,
- z_i involve only $x_i, \boldsymbol{\mu}, \boldsymbol{\lambda}, \pi$,
- μ_k involve only $\mathbf{x}, \boldsymbol{\lambda}, \mathbf{z}$,
- λ_k involve only $\mathbf{x}, \boldsymbol{\mu}, \mathbf{z}$.

References

Michael Jordan. *An Introduction to Probabilistic Graphical Models*. 2003.