

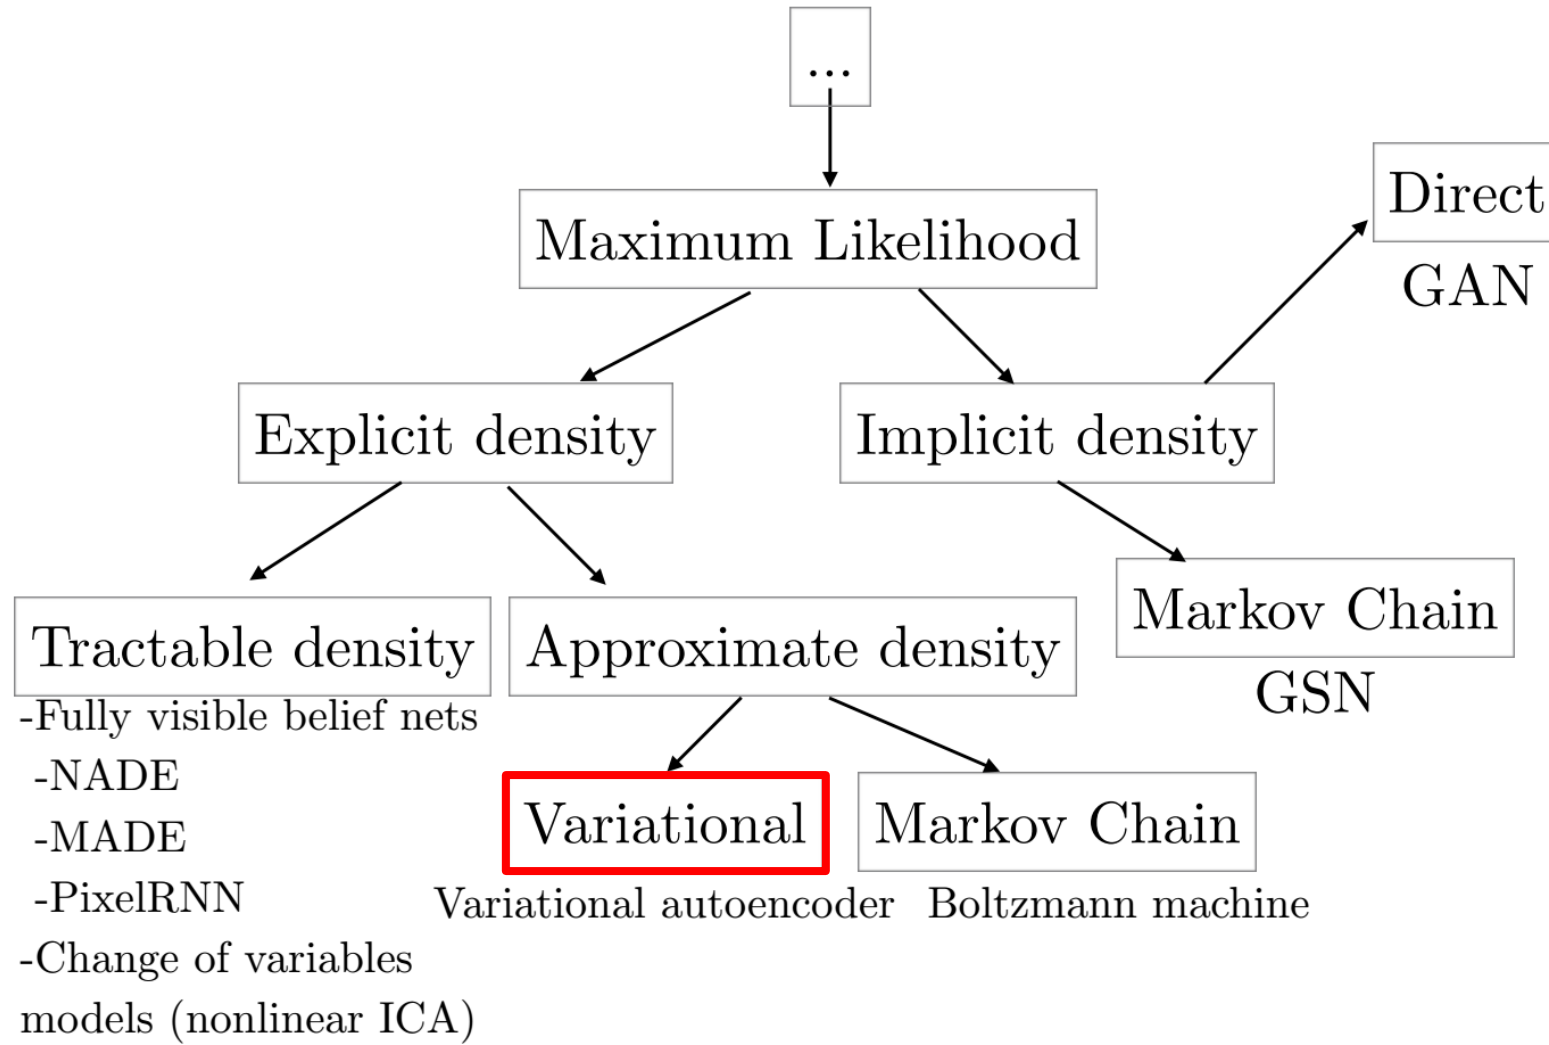
Lecture 9: Deep Variational Inference

Efstratios Gavves

Lecture overview



- Latent variable models
- Autoencoders
- Variational inference
- Variational autoencoders
- Reparameterization trick

A map of generative models




Motivation behind latent variables

- Latent variables are high-level features
 - in combination they generate the data $p(\mathbf{x}|\mathbf{z})$
- We want to model the inverse process
 - → Figure out what these latents are





Why bother?

- By making convenient choices for latents
 - We can model much more complex x
- Without latents we can have exploding number of parameters
 - → intractability
 - Remember (regular) Boltzmann machines



Distributed representations

- Distribute the ‘representation’ of our data over multiple neurons
- And each neuron models a distribution of concepts
- The latent layer learns to encode combinations of patterns for efficiency



Hinton, CSC321: 2011 Introduction to Neural Networks and Machine Learning