Score-matching & Diffusion Generative Models

Deep Learning II - uvald2c.github.io

Efstratios Gavves - University of Amsterdam

Overview

Introduction to score-matching
Noise conditional score networks
Score-based generation via SDEs
Diffusion models

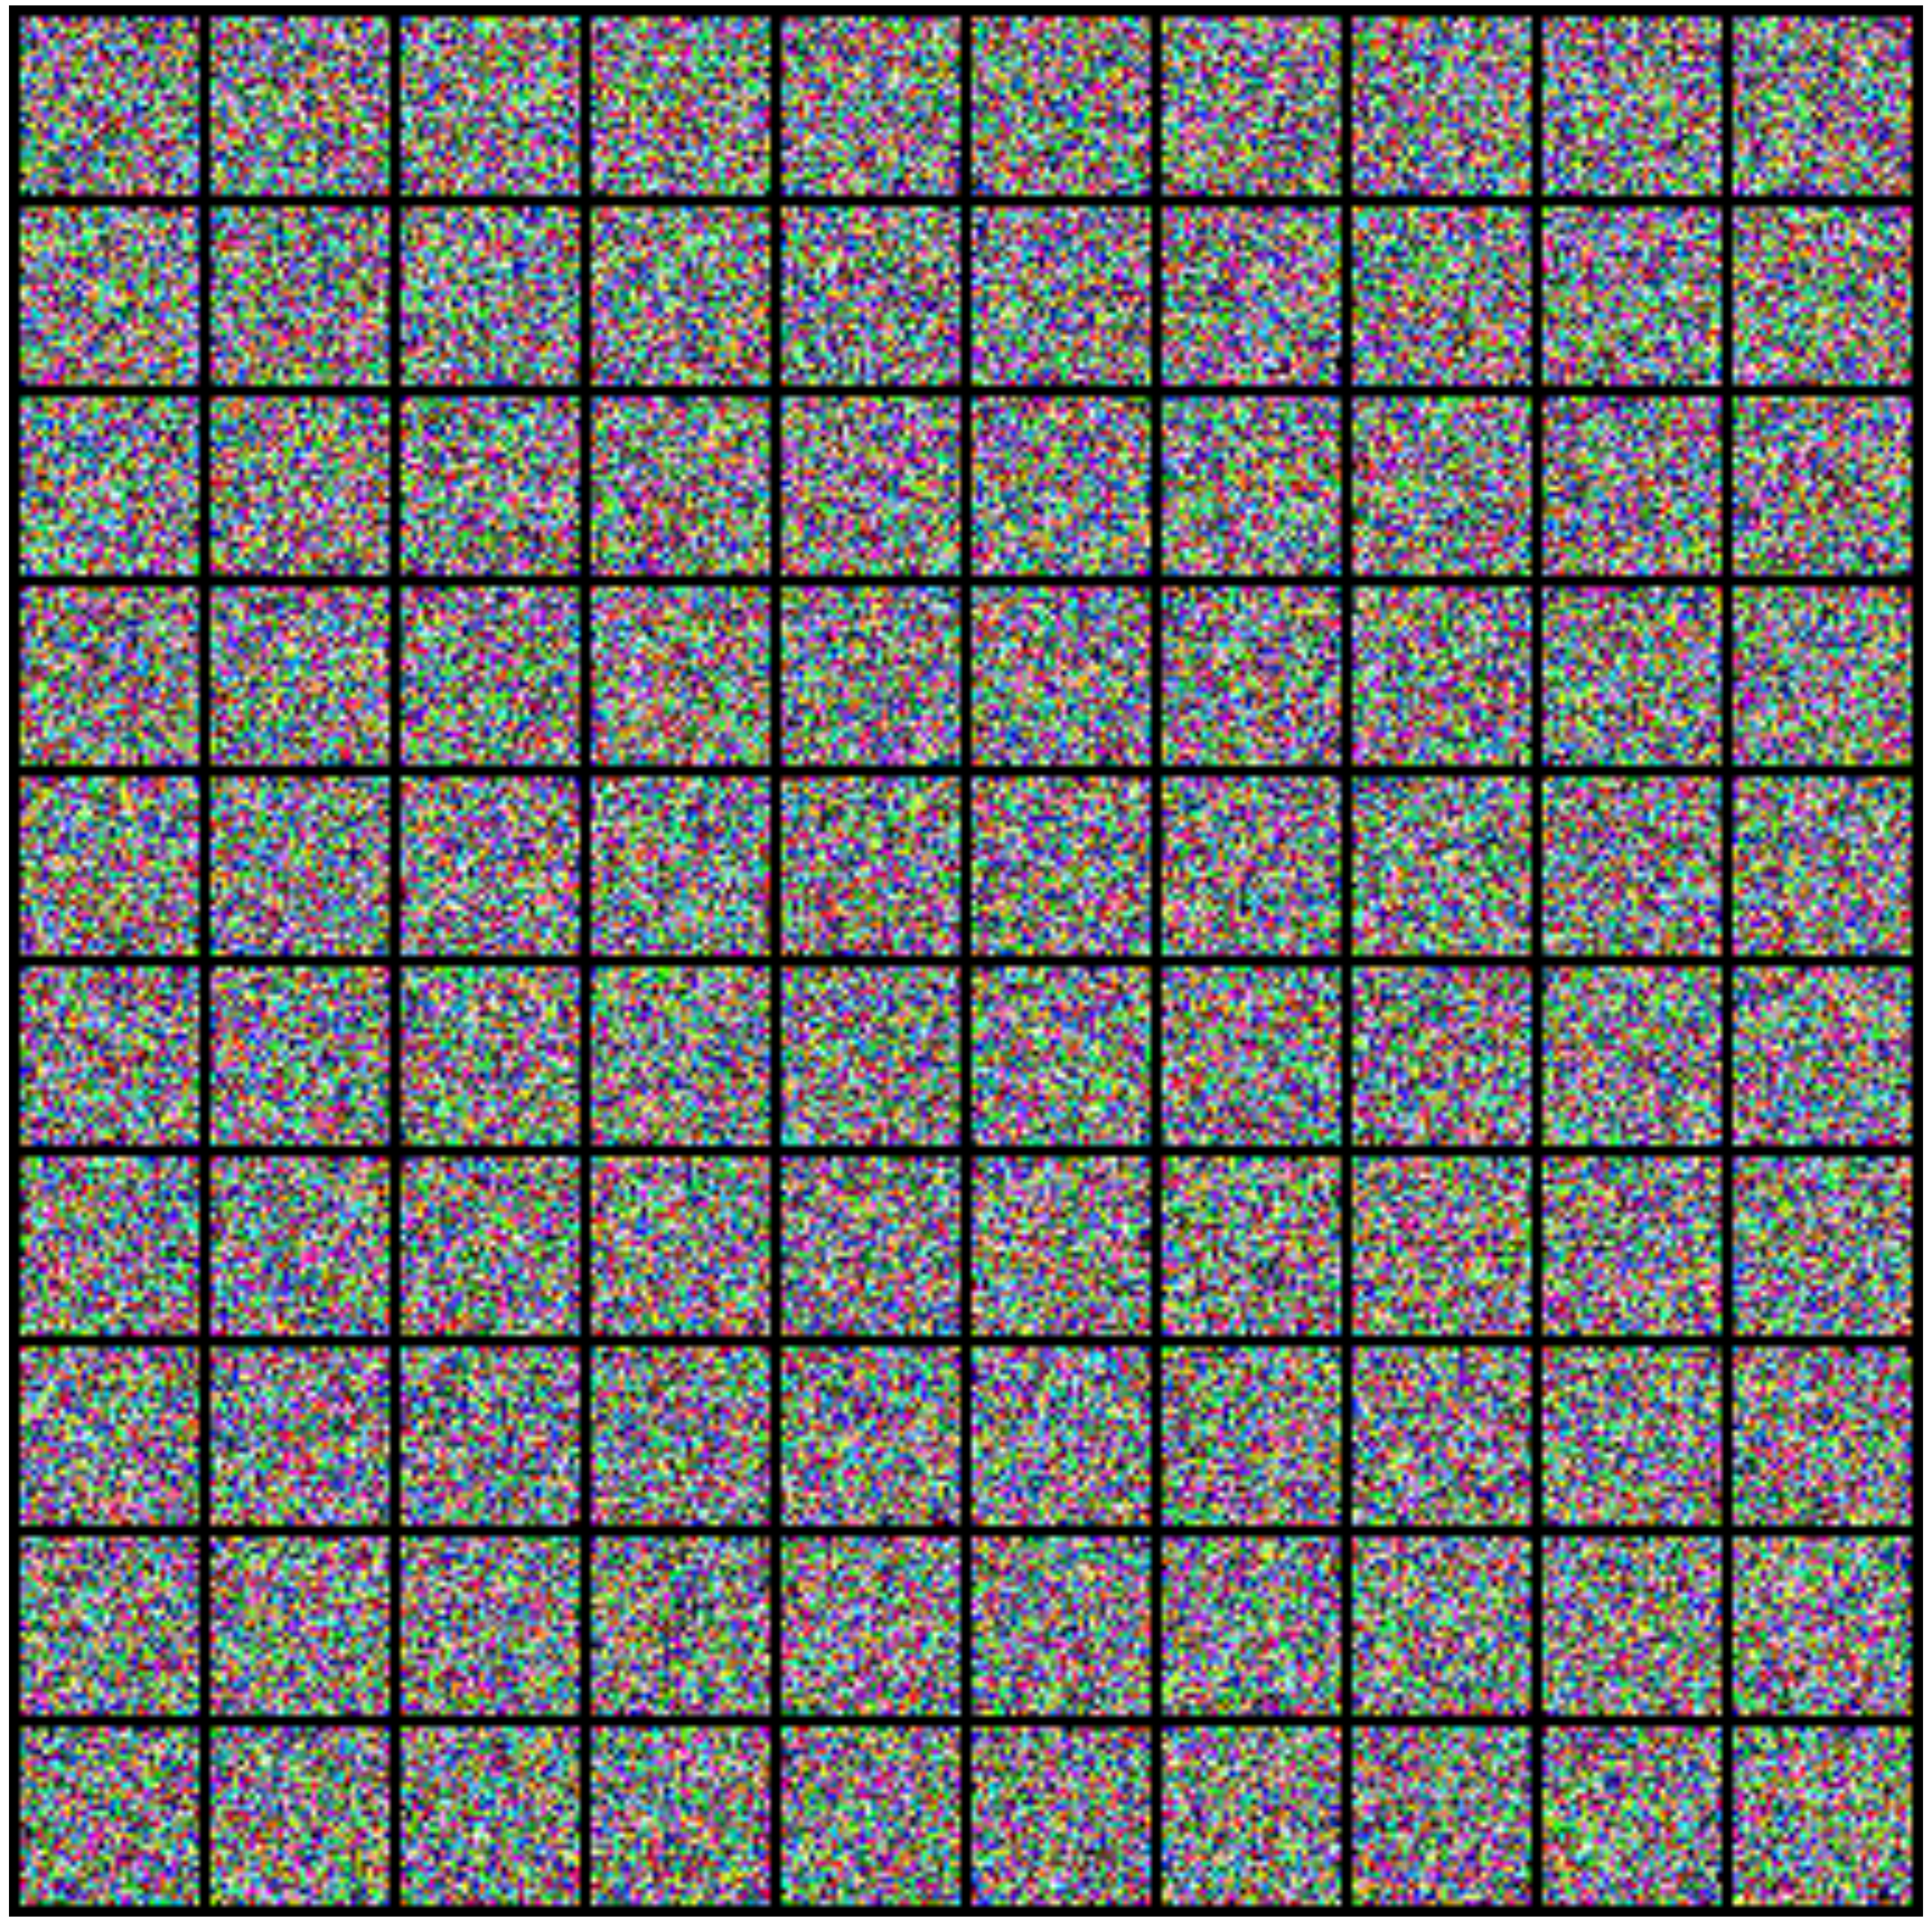

https://yang-song.github.io/blog/2021/score/

Tractability v. Flexibility

- In generative modelling there are two opposing forces: tractability and flexibility
- Tractable models are usually analytically computable, thus easy to evaluate and fit
- But they are usually not flexible enough to learn the true data structure
- Flexible models can fit arbitrary structures in data
- But they are usually expensive to evaluate, fit, or sample from

Overview of generative models

Likelihood-based generative models

Implicit generative models

Likelihood-based generative models

- Typically make strong assumptions to ensure tractability of likelihood
- Specifically of the normalising constant $Z(x)$ in $p(x) = \dfrac{\tilde{p}(x)}{Z(x)}$
- For instance, VAEs assume a tractable variational approximation
- Autoregressive models require causal convolutions
- Normalizing Flows require invertibility in the network architecture

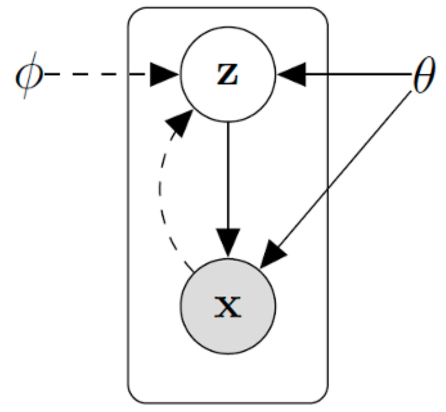

Bayesian networks
(e.g., VAEs)

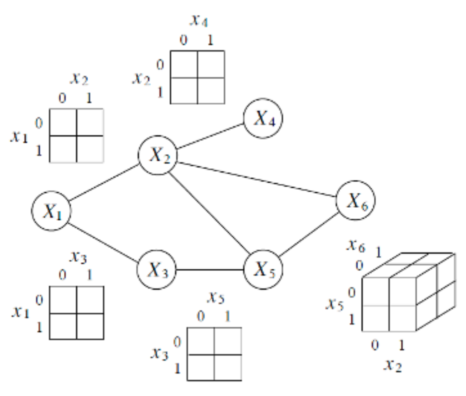

MRF

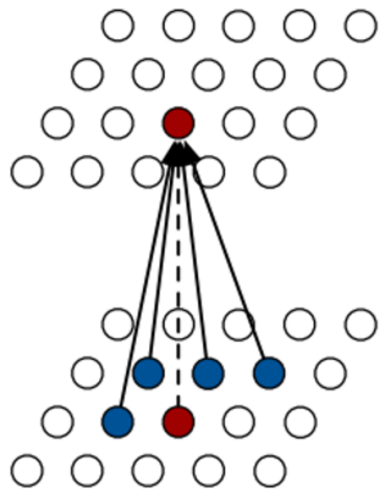

Autoregressive
models

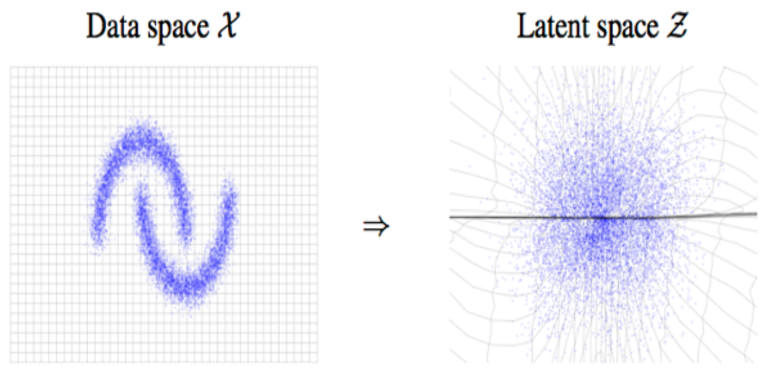

Flow models

Implicit generative models

- Adversarial training for implicit generative models is very unstable
- Adversarial training leads often to mode collapse and reduced sampling variance
- Implicit generative models cannot compute likelihood of a sample, they just sample

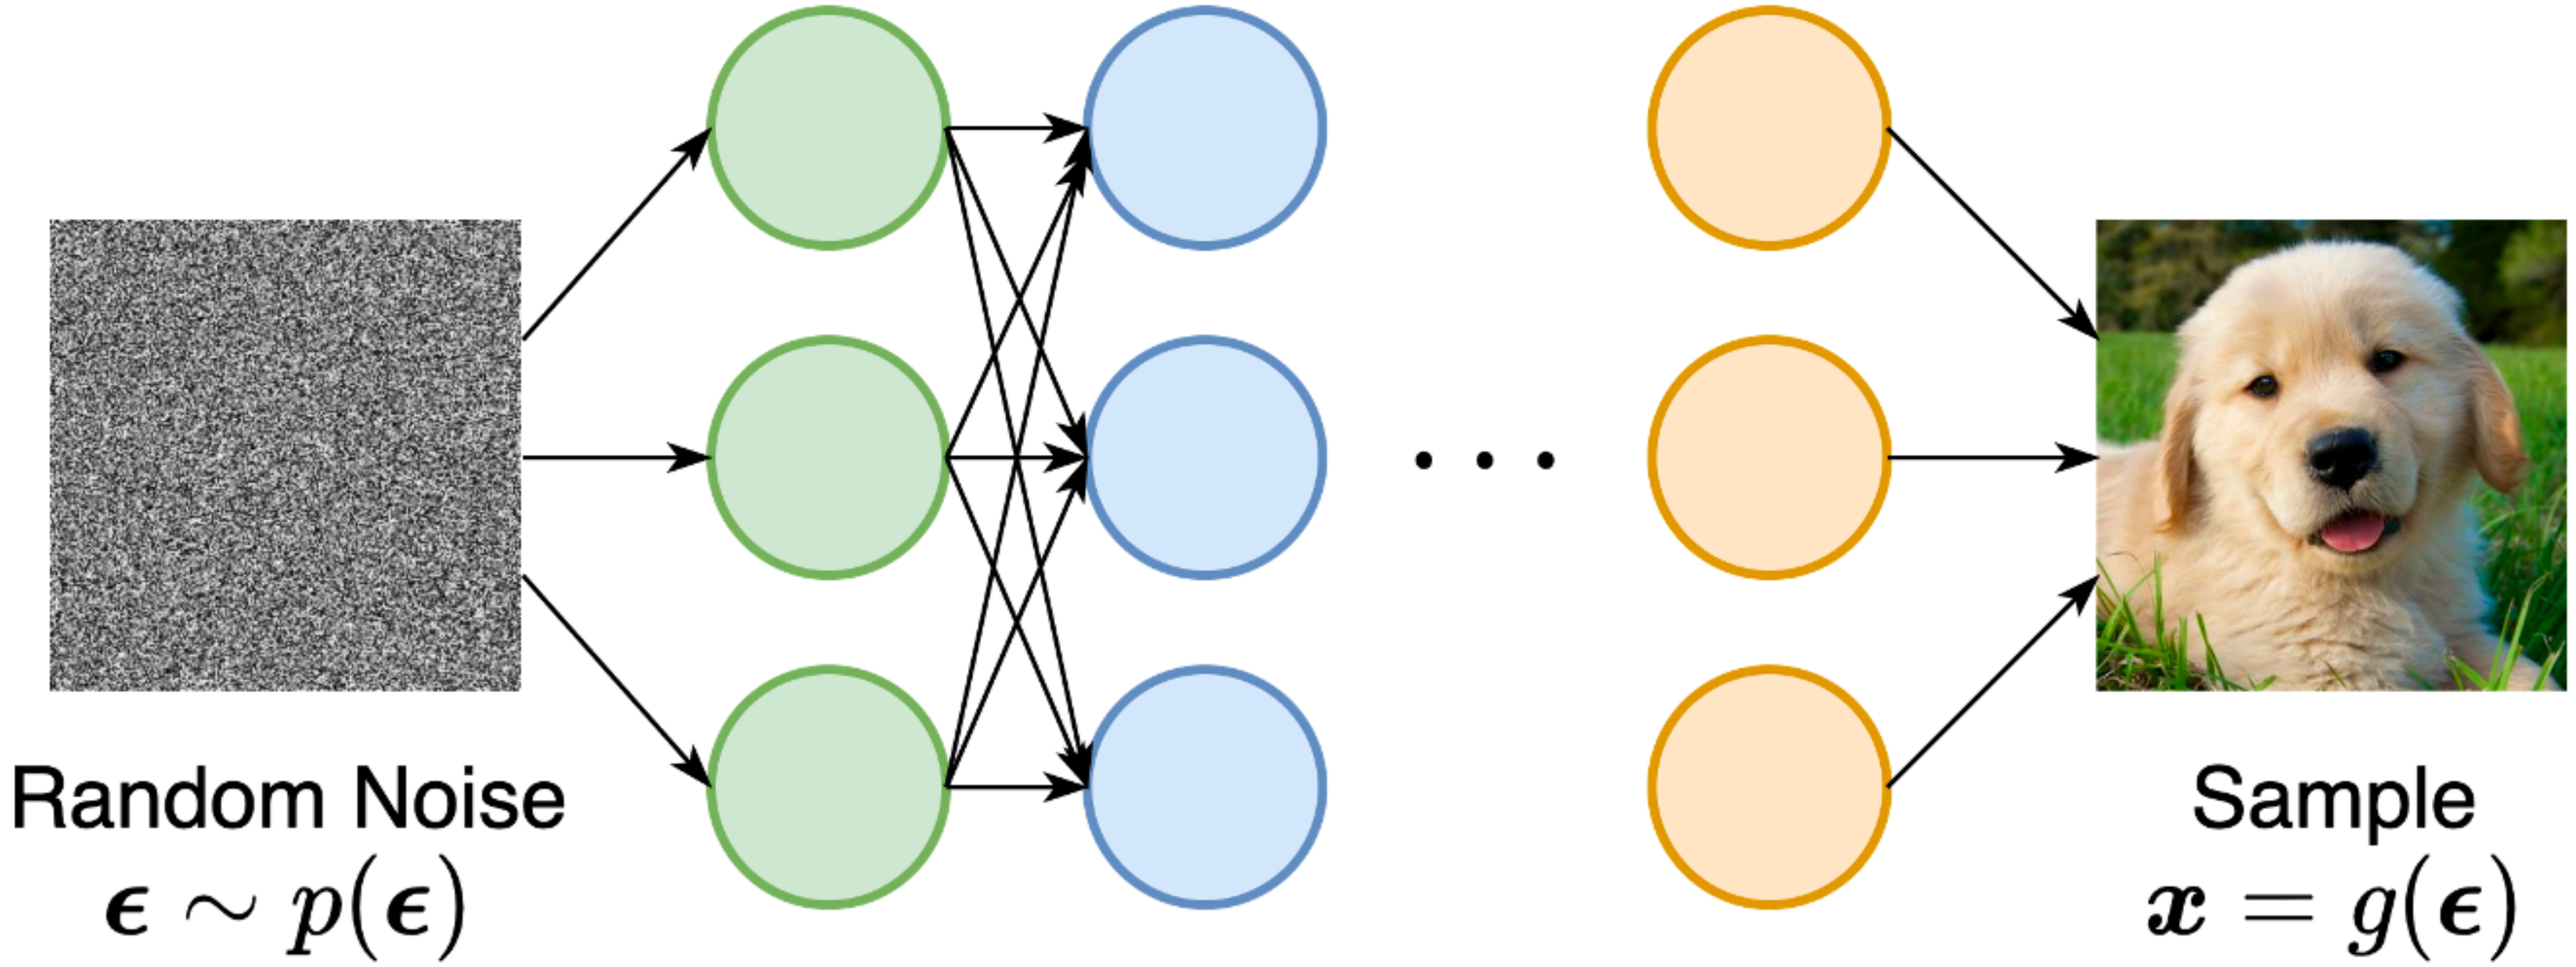

Overview of generative models

Likelihood-based generative models

Implicit generative models

Score-based generative models

Tractability v. Flexibility

- In generative modelling there are two opposing forces: tractability and flexibility
- Tractable models are usually analytically computable, thus easy to evaluate and fit
- But they are usually not flexible enough to learn the true data structure
- Flexible models can fit arbitrary structures in data
- But they are usually expensive to evaluate, fit, or sample from
- <u>Diffusion/score-matching models are both tractable and flexible</u>