University of Amsterdam
Informatics Institute

AMLAB
Amsterdam
Machine Learning Lab

Machine Learning 1

Lecture 9.1 - Unsupervised Learning
Latent Variable Models

Erik Bekkers

(Bishop 9.0)

Slide credits: Patrick Forré and Rianne van den Berg

Image credit: Kirillm | Getty Images

Unsupervised vs. Supervised learning

- Supervised
  - Data $D = \{\boldsymbol{X}, \boldsymbol{T}\}$
  - Goals $f(\boldsymbol{x}) \approx t, p(\boldsymbol{t}|\boldsymbol{x})$
  - Classification (discrete) or regression (continuous)

- Unsupervised
  - Data $D = \{\boldsymbol{X}\}$
  - Goals $p(\boldsymbol{x}), p(\boldsymbol{z}|\boldsymbol{x})$ or $p(\boldsymbol{x}|\boldsymbol{z})$
  - Density estimation, **clustering** (discrete) or **dimensionality reduction** (continuous)

z latent variable that influences how we observe x

Latent variable models

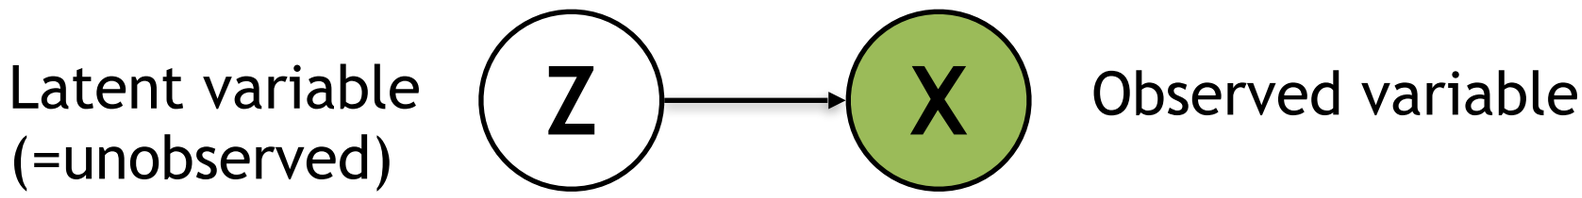

- Model complex distributions with more tractable representation by z
  - continuous:

$$p(\boldsymbol{x}) = \int p(\boldsymbol{x}, \boldsymbol{z}) d\boldsymbol{z} = \int p(\boldsymbol{x}|\boldsymbol{z}) p(\boldsymbol{z}) d\boldsymbol{z}$$

  - discrete:

$$p(\boldsymbol{x}) = \sum_{\boldsymbol{z}} p(\boldsymbol{x}, \boldsymbol{z}) = \sum_{\boldsymbol{z}} p(\boldsymbol{x}|\boldsymbol{z}) p(\boldsymbol{z})$$

Examples

$$p(x) = \sum_z p(x, z) = \sum_z p(x \mid z)\, p(z)$$

Clustering

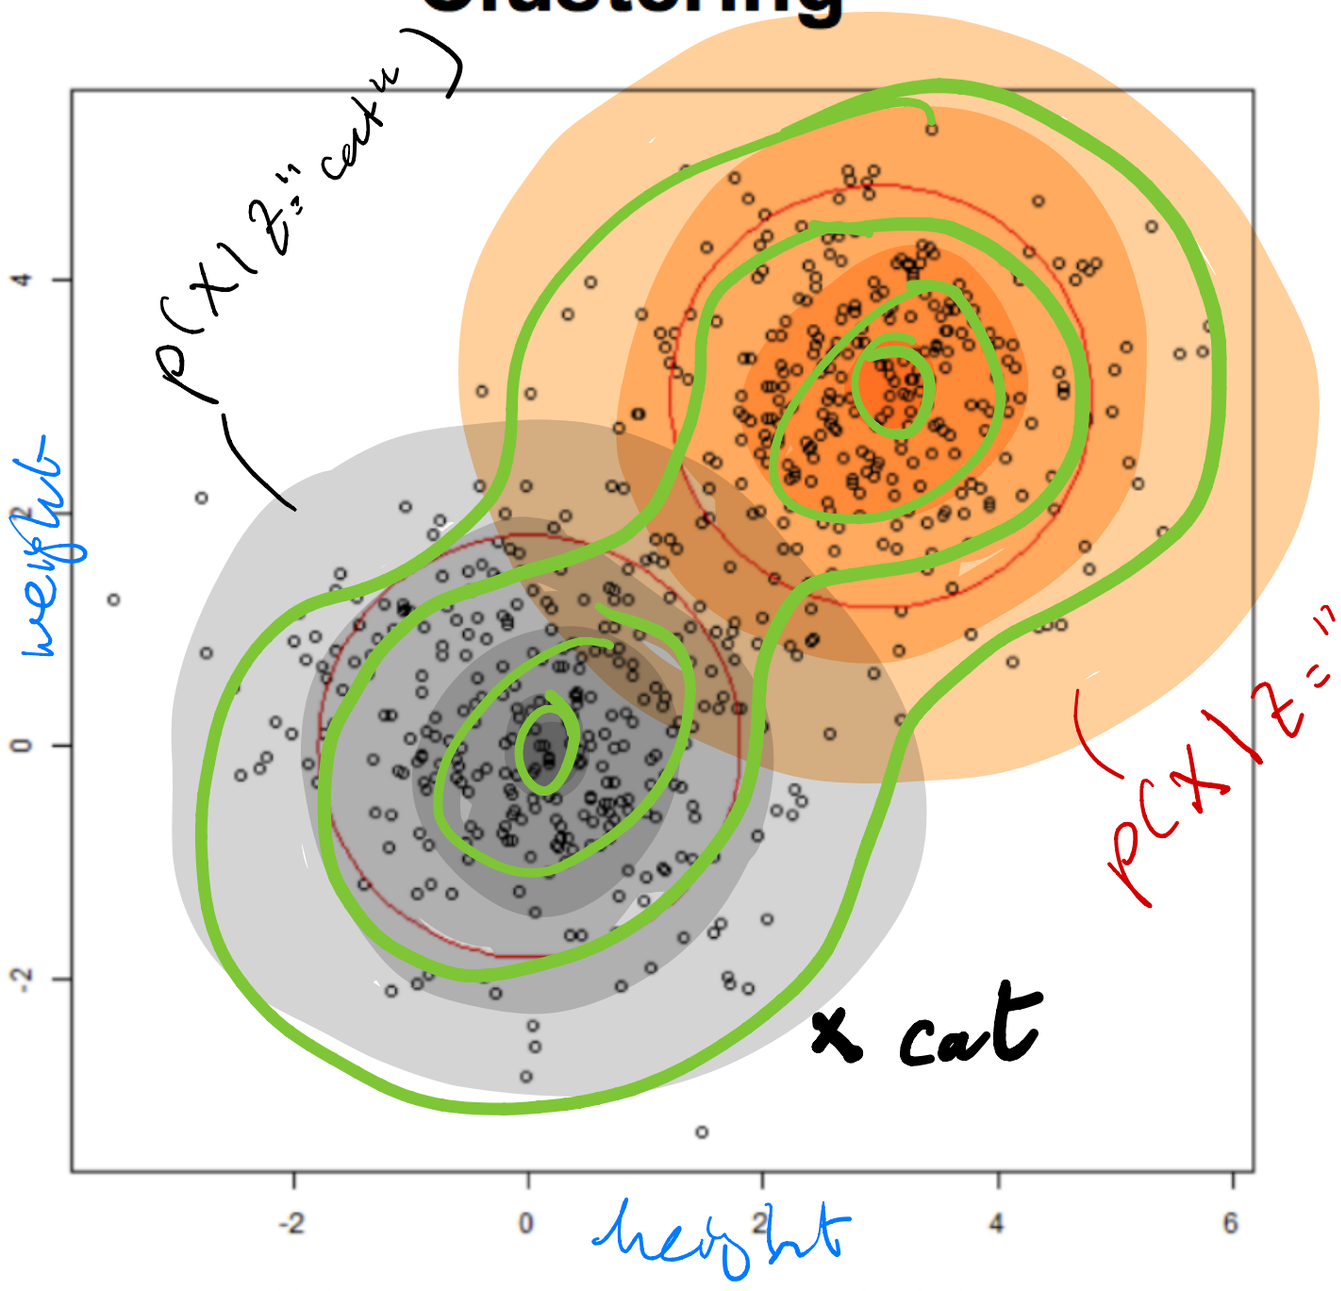

Discrete latent variables

animal type

Dimensionality reduction

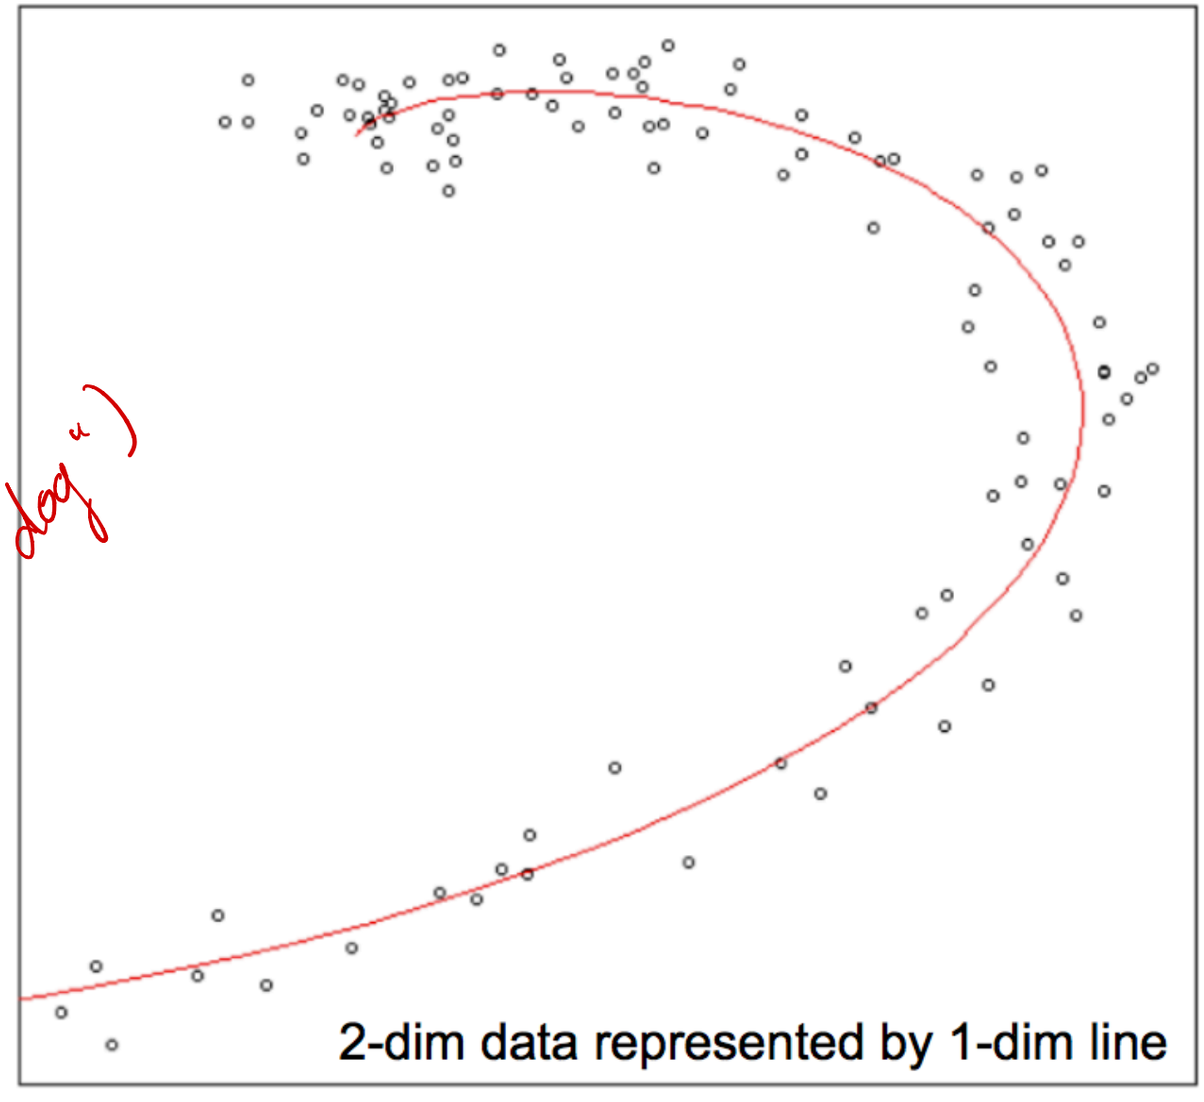

Continuous latent variables