

Dimensionality Reduction

M. Falzari

January 23, 2022

How?

Linear

- LDA (Linear discriminant analysis)
- PCA (Principal component analysis)

How?

Linear

- LDA (Linear discriminant analysis)
- PCA (Principal component analysis)


Non linear

- LLE (Locally linear embedding)
- Isomap
- Autoencoder

Why autoencoders?

- Autoencoder \neq other methods
- Potentially detect repetitive structures
- [wang2016auto] [lin2020deep] [lonvcarevic2021robot]


Autoencoder



Bayesian and Variational inference


- find a posterior $p(Z|X)$ such that X = observation and Z latent variables.
- (many times) intractable since $p(X)$ is unknown
- Variational inference try to find a surrogate posterior given a family of distributions
- Usually KL(Kullback-Leibler) divergence is used to define how "close" the surrogate is to the desired posterior.

Variational Autoencoder




- [mescheder2017adversarial]
- fully implicit latent distribution
- problematic because KL_{div} is intractable
- use a discriminator in an adversarial manner to approximate the prior


MNIST Example AE



MNIST Example VAE



MNIST Example AVB



- [higgins2017darla] and others, propose to perform a separate phase for training the AE "offline"
 - Random agent
 - Fixed strategy agent
 - Human Data
- Online

Reference

./ref.bib