STOR566: Introduction to Deep Learning

Lecture 12: Generative Models

Yao Li
UNC Chapel Hill

Oct 04, 2022

Unsupervised Learning

- Working with datasets without a **response** variable
- Some Applications:
    - Clustering
    - Data Compression
    - Exploratory Data Analysis
    - Generating New Examples
    - ...
- Example: PCA, K-means, Autoencoders, GAN, etc

Autoencoder: Basic Architecture

- Autoencoder: A special type of DNN where the target (response) of each input is the input itself.

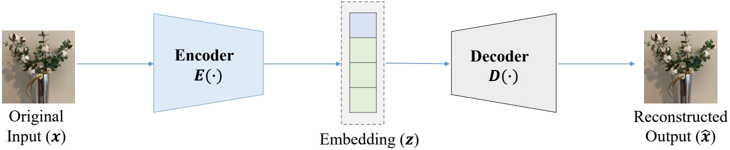

Autoencoder: Basic Architecture

- Autoencoder: A special type of DNN where the target (response) of each input is the input itself.

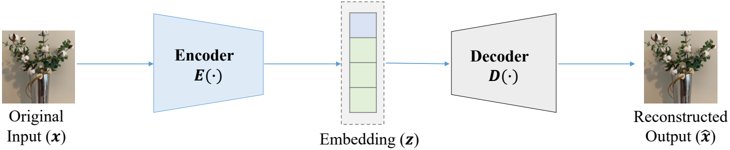

- Objective:

$$\|\boldsymbol{x} - \boldsymbol{D}(\boldsymbol{E}(\boldsymbol{x}))\|^2$$

Encoder: $\boldsymbol{E} : \mathbb{R}^n \to \mathbb{R}^d$

Decoder: $\boldsymbol{D} : \mathbb{R}^d \to \mathbb{R}^n$

Transposed Convolution

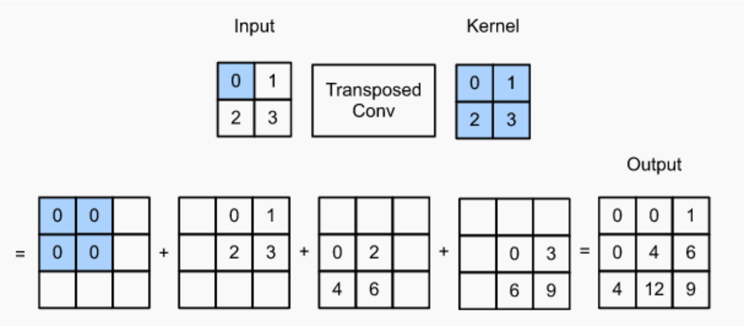

(Figure from Dive into Deep Learning)

- Multiple input and output channels: works the same as the regular convolution
- Number of weights: $k_1 \times k_2 \times d_{in} \times d_{out} + d_{out}$

Transposed Convolution

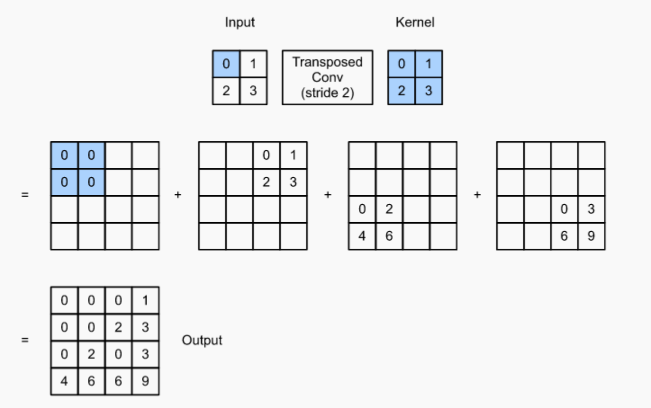

(Figure from Dive into Deep Learning)

- Strides are specified for the output feature map
- Padding: remove rows and columns from the output

Overfitting

- Overfitting is a problem
- Solutions:
  - Bottleneck layer: a low-dimensional representation of the data ($d < n$)
  - Denoise autoencoder
  - Sparse autoencoder
  - ...

Regularization

- Objective:

$$L(\boldsymbol{x}, \hat{\boldsymbol{x}}) + \text{regularizer},$$

Regularization

- Objective:

$$L(\boldsymbol{x}, \hat{\boldsymbol{x}}) + \text{regularizer},$$

$L(\cdot,\cdot)$: captures the distance between the input ($\boldsymbol{x}$) and the output ($\hat{\boldsymbol{x}}$).

  - Example: $\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|^2$

Regularization

- Objective:

$$L(\boldsymbol{x}, \hat{\boldsymbol{x}}) + \text{regularizer},$$

$L(\cdot, \cdot)$: captures the distance between the input ($\boldsymbol{x}$) and the output ($\hat{\boldsymbol{x}}$).

  - Example: $\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|^2$

Regularizer example:

  - L_1 penalty: $\sum_j |h_j^l|$
  - h_j^l: hidden output of j-th neuron in l-th layer

Sparse Autoencoder

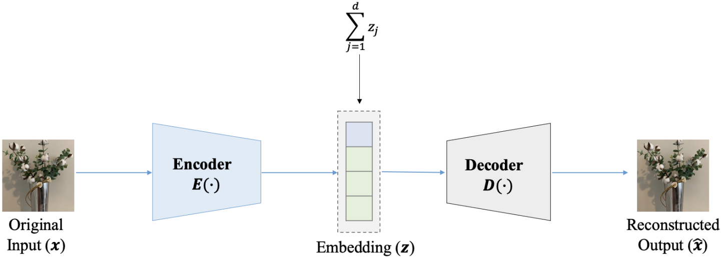

- Objective:

$$\|\boldsymbol{x} - \boldsymbol{D}(\boldsymbol{E}(\boldsymbol{x}))\|^2 + \lambda \sum_j |z_j|$$

- Iterate over layers.

Sparse Autoencoder

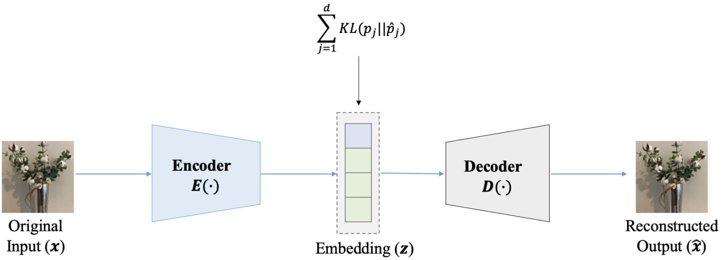

- Another regularizer:

$$\|x - D(E(x))\|^2 + \lambda \sum_j KL(p_j||\hat{p}_j)$$

- Convert value of z to $[0, 1]$. (e.g., sigmoid activation)
- p_j: probability of activation for neuron j in the bottleneck layer
- $\hat{p}_j = \frac{1}{B} \sum_{i=1}^{B} z_{ij}$

Denoising Autoencoder

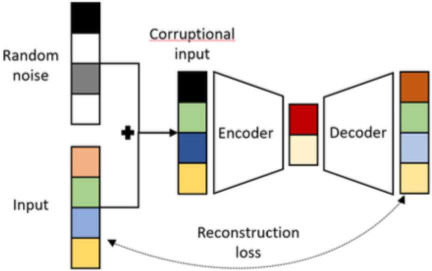

Figure from Bank, Dor, Noam Koenigstein, and Raja Giryes. "Autoencoders." (2020).

- Another regularizer:

$$\|\boldsymbol{x} - \boldsymbol{D}(\boldsymbol{E}(\boldsymbol{x} + \delta))\|^2$$

- δ: Random noise

Variational Autoencoder (VAE)

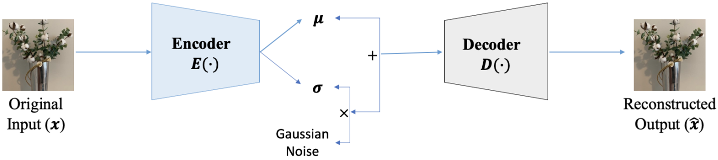

- Probabilistic model: will let us generate data from the model
- Encoder outputs $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$
- Draw $\tilde{\boldsymbol{z}} \sim N(\boldsymbol{\mu}, \boldsymbol{\sigma})$
- Decoder decodes this **latent** variable $\tilde{\boldsymbol{z}}$ to get the output

Variational Autoencoder (VAE)

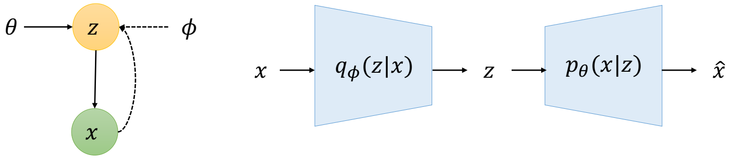

- Maximum likelihood approach: $\Pi_i p(\boldsymbol{x}_i)$
- Variational lower bound as objective:
  - End-to-End reconstruction loss (e.g., square loss)
  - Regularizer: $KL\left(q_\Phi(\boldsymbol{z}|\boldsymbol{x})||p(\boldsymbol{z})\right)$
- Objective:

$$L(\boldsymbol{x}, \hat{\boldsymbol{x}}) + KL\left(q_\Phi(\boldsymbol{z}|\boldsymbol{x})||p(\boldsymbol{z})\right)$$

Re-parameterization Trick

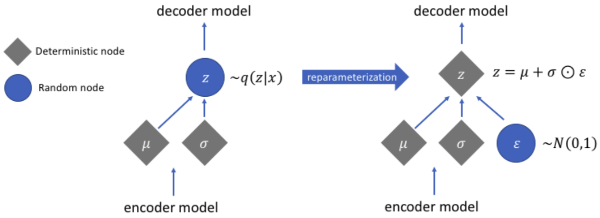

Figure from Jeremy Jordon Blog

- Cannot back-propagate error through random samples
- Reparameterization trick: replace $\tilde{\boldsymbol{z}} \sim N(\boldsymbol{\mu}, \boldsymbol{\sigma})$ with $\epsilon \sim N(0, I)$, $\boldsymbol{z} = \epsilon \boldsymbol{\sigma} + \boldsymbol{\mu}$

Variational Lower Bound

- Variational lower bound:

$$\log p(x) \geq E_{q(\mathbf{z}|\mathbf{x})} \left(\log p(x|z) \right) + KL \left(q(z|x) || p(z) \right)$$

- How to derive the variational lower bound from the likelihood?

Conclusions

- Autoencoder
- Regularization
- Variational Autoencoder

Questions?