BIG DATA and AI for business

Auto-encoder

Decisions, Operations & Information Technologies
Robert H. Smith School of Business
Fall, 2020

Unspervised Learning

- "We expect unsupervised learning to become far more important in the longer term. Human and animal learning is largely unsupervised: we discover the structure of the world by observing it, not by being told the name of every object."
– LeCun, Bengio, Hinton, Nature 2015

- Yann LeCun, March 14, 2016

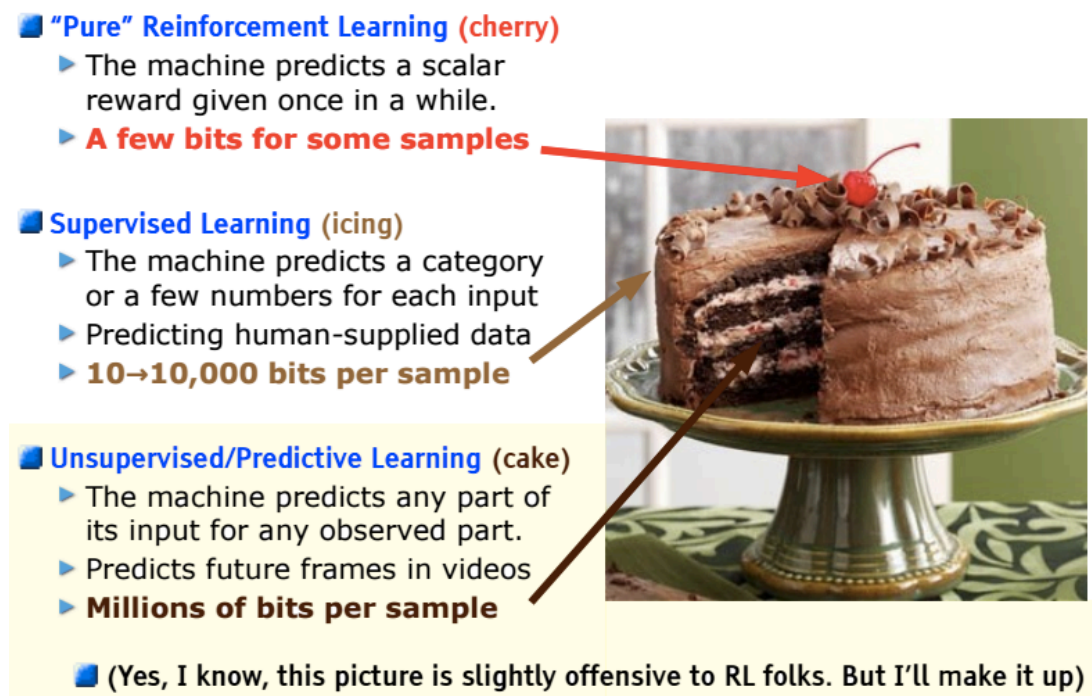

Auto-encoder

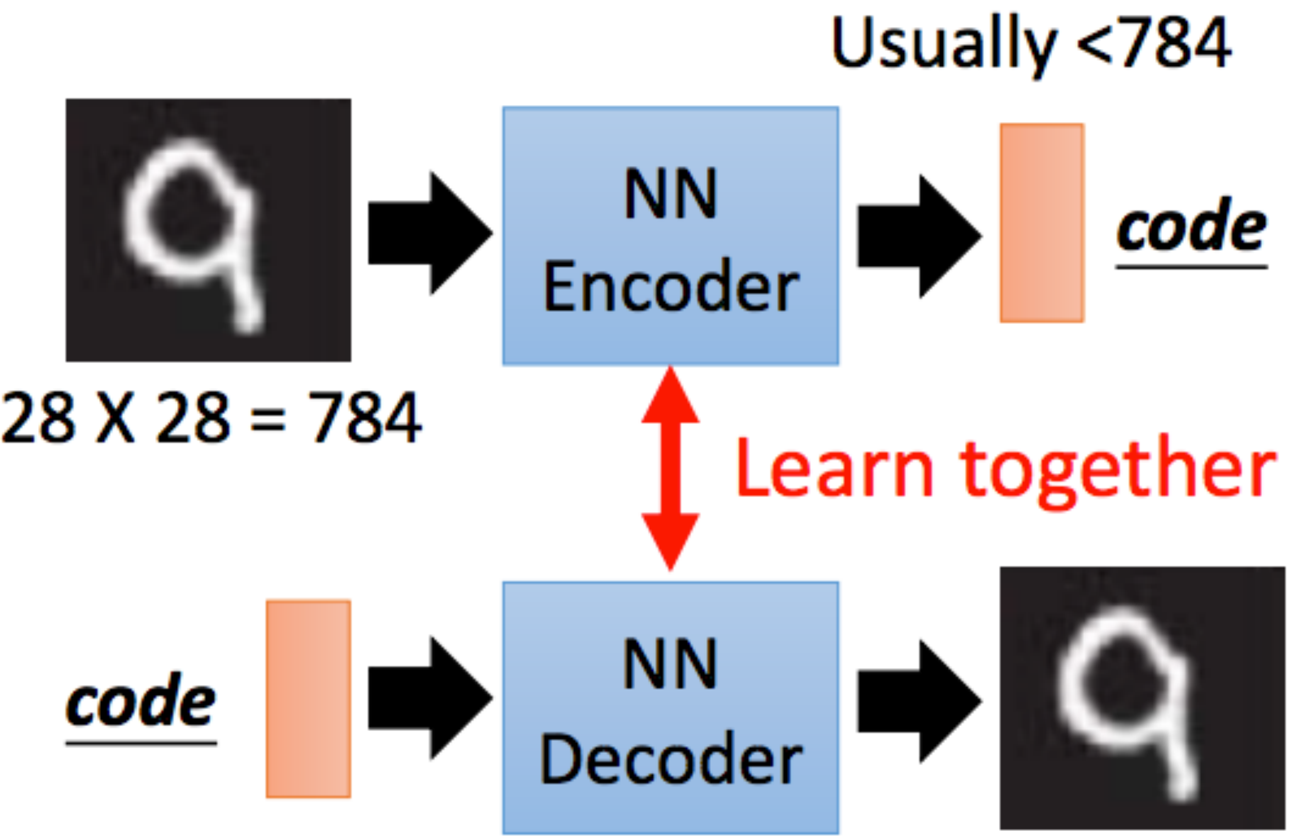

Compact representation of the input

Can reconstruct the original object

Recap: PCA

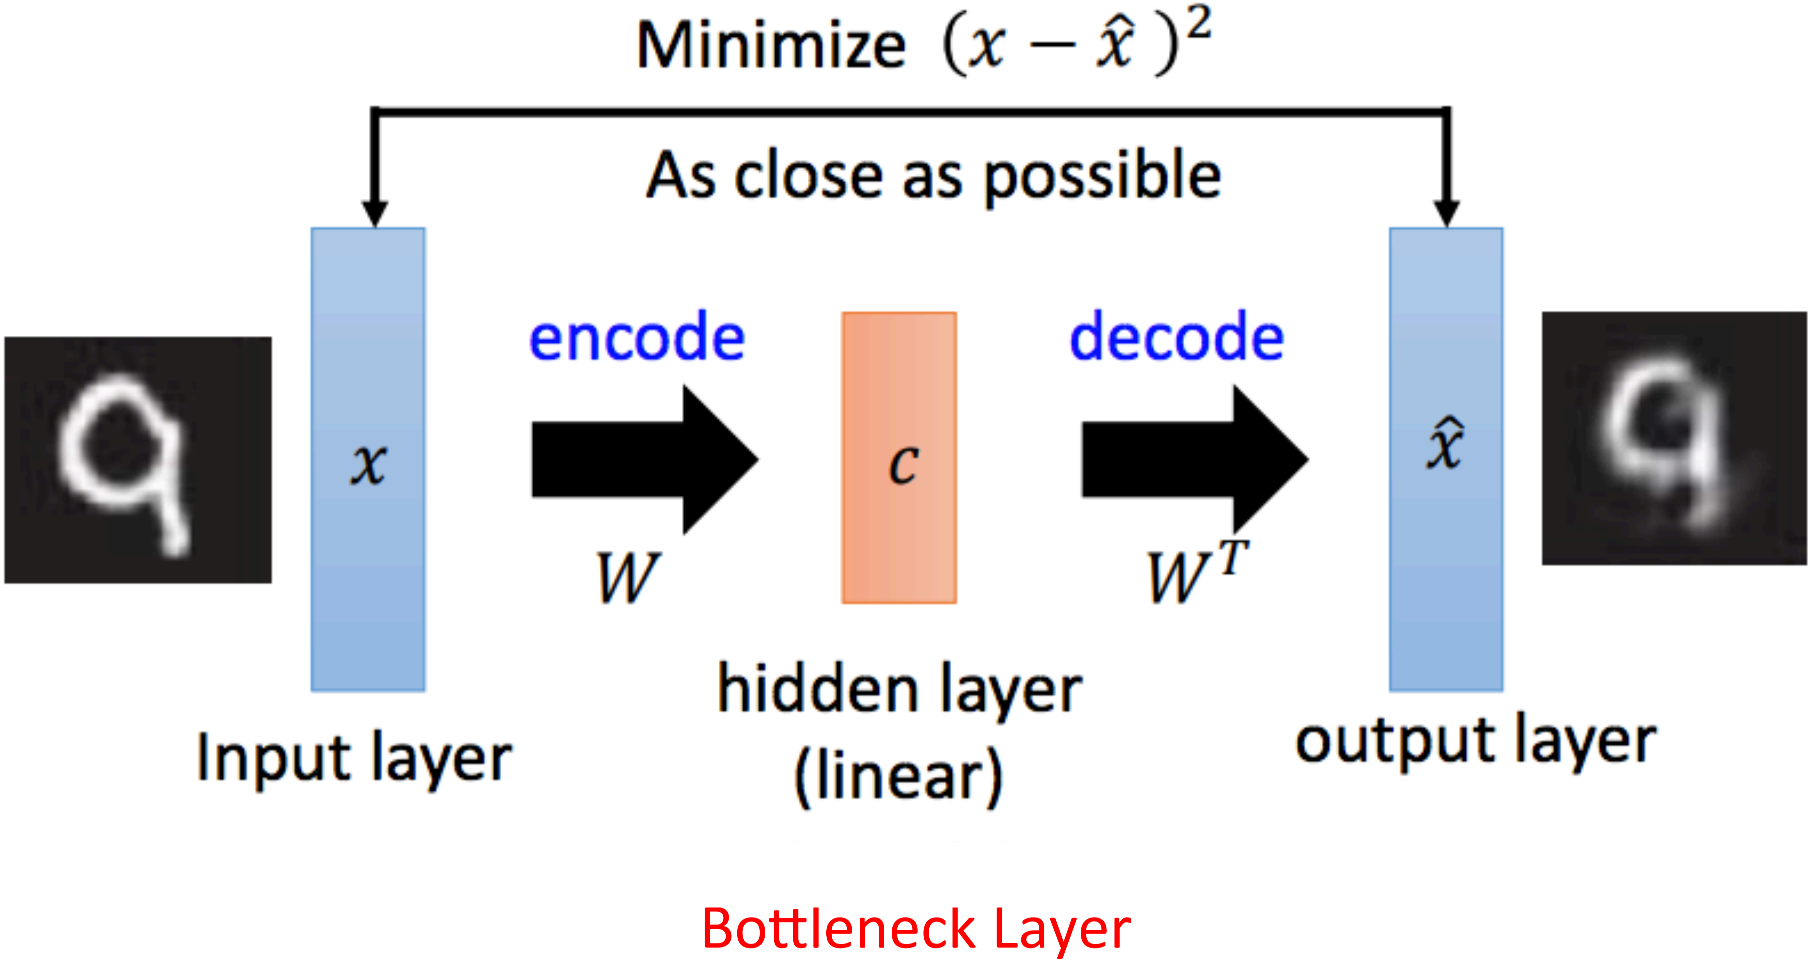

Output of the hidden layer is the code

Deep Auto-encoder

Symmetry is not necessary

- Of course, the auto-encoder can be deep

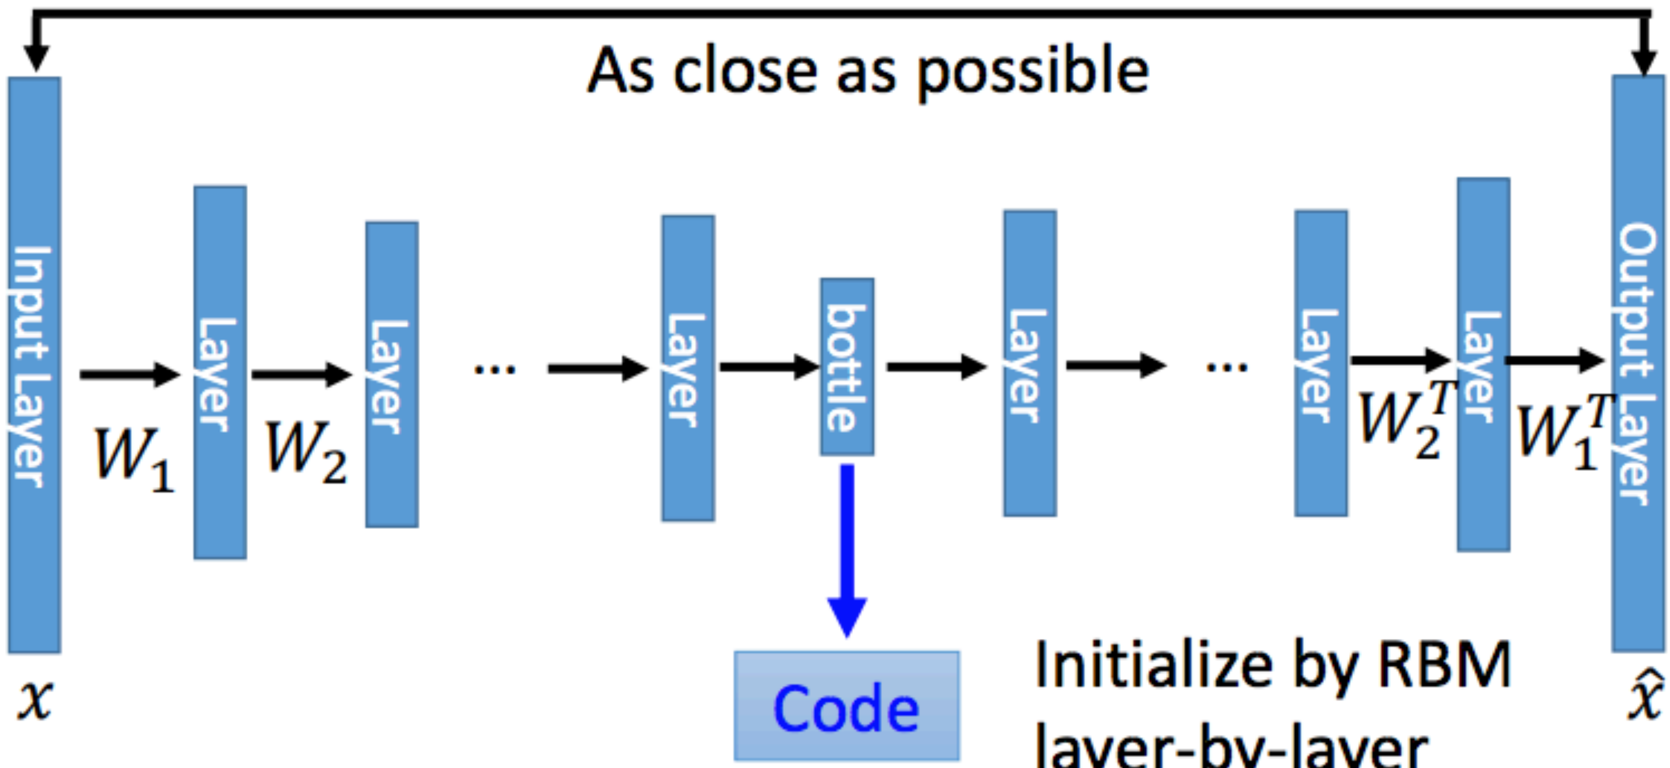

Hinton, Geoffrey E., and Ruslan R. Salakhutdinov. "Reducing the dimensionality of data with neural networks." Science 313.5786 (2006): 504-507

Deep Auto-encoder

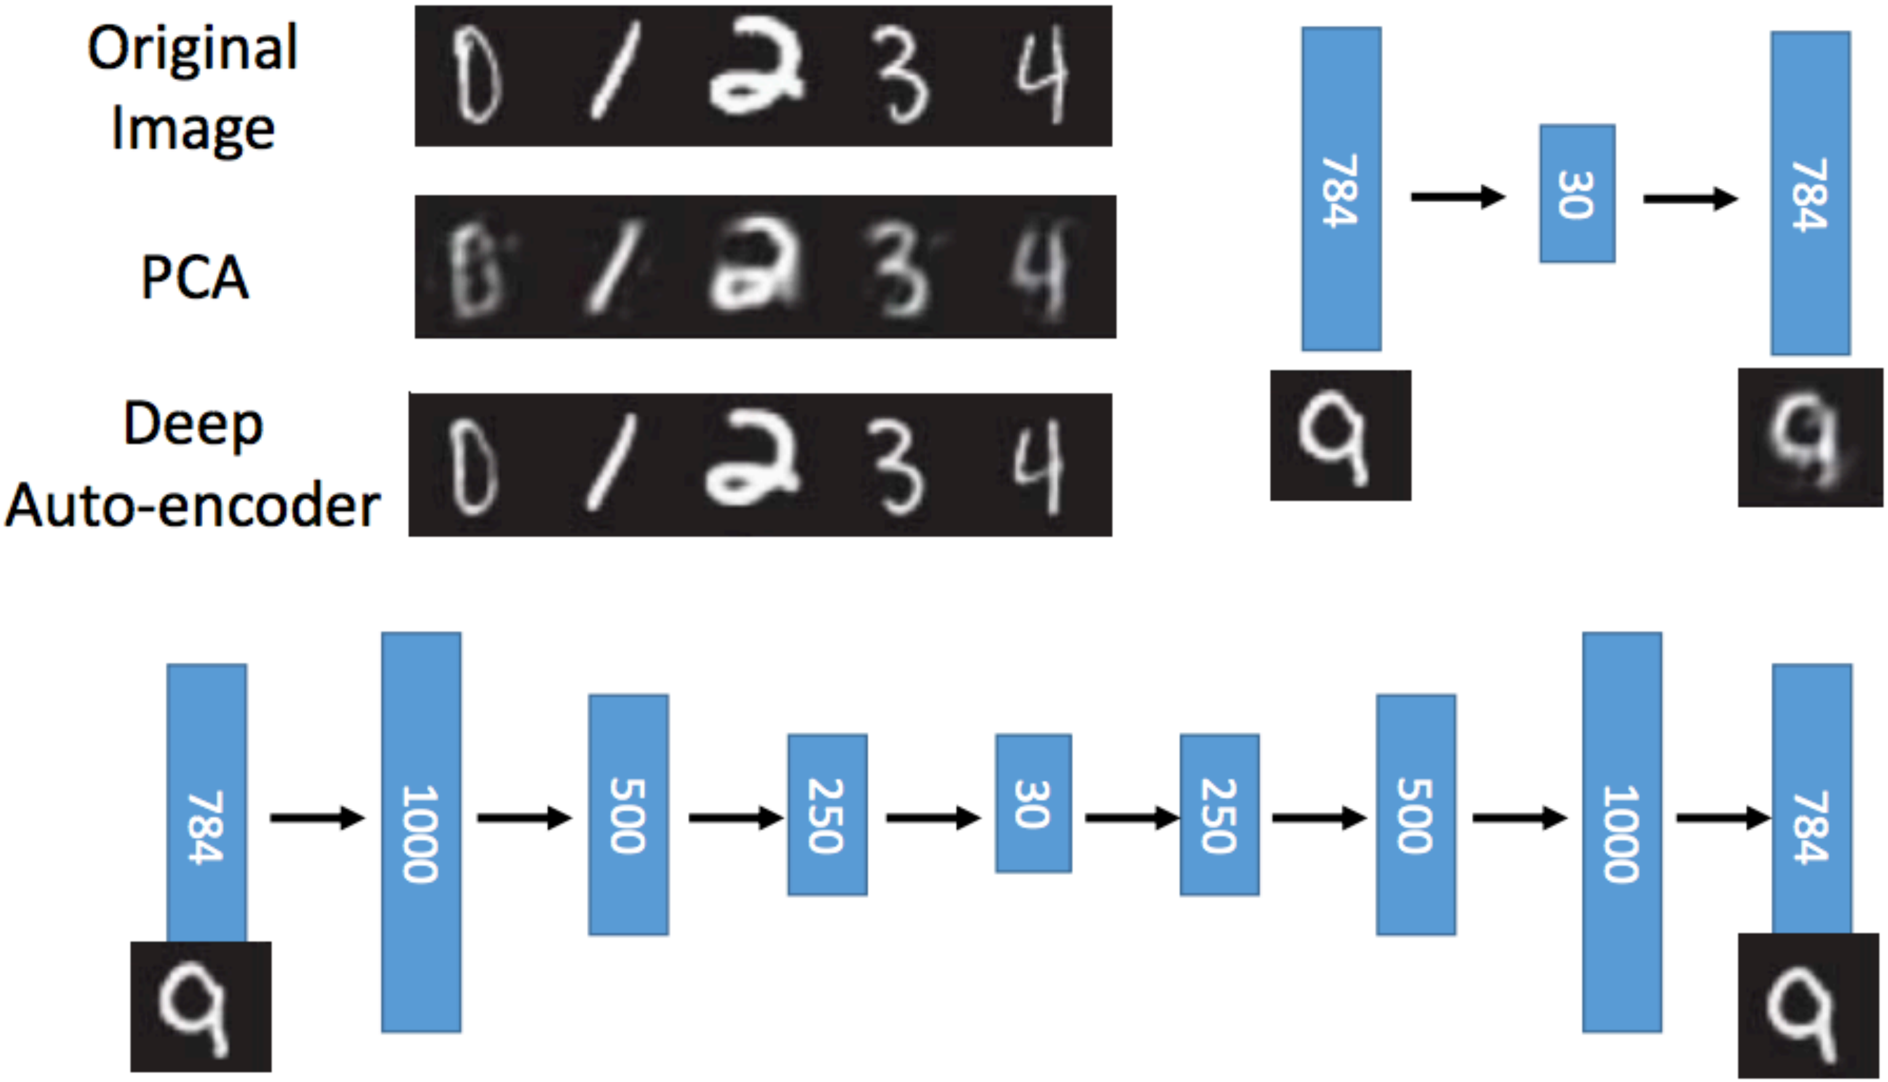

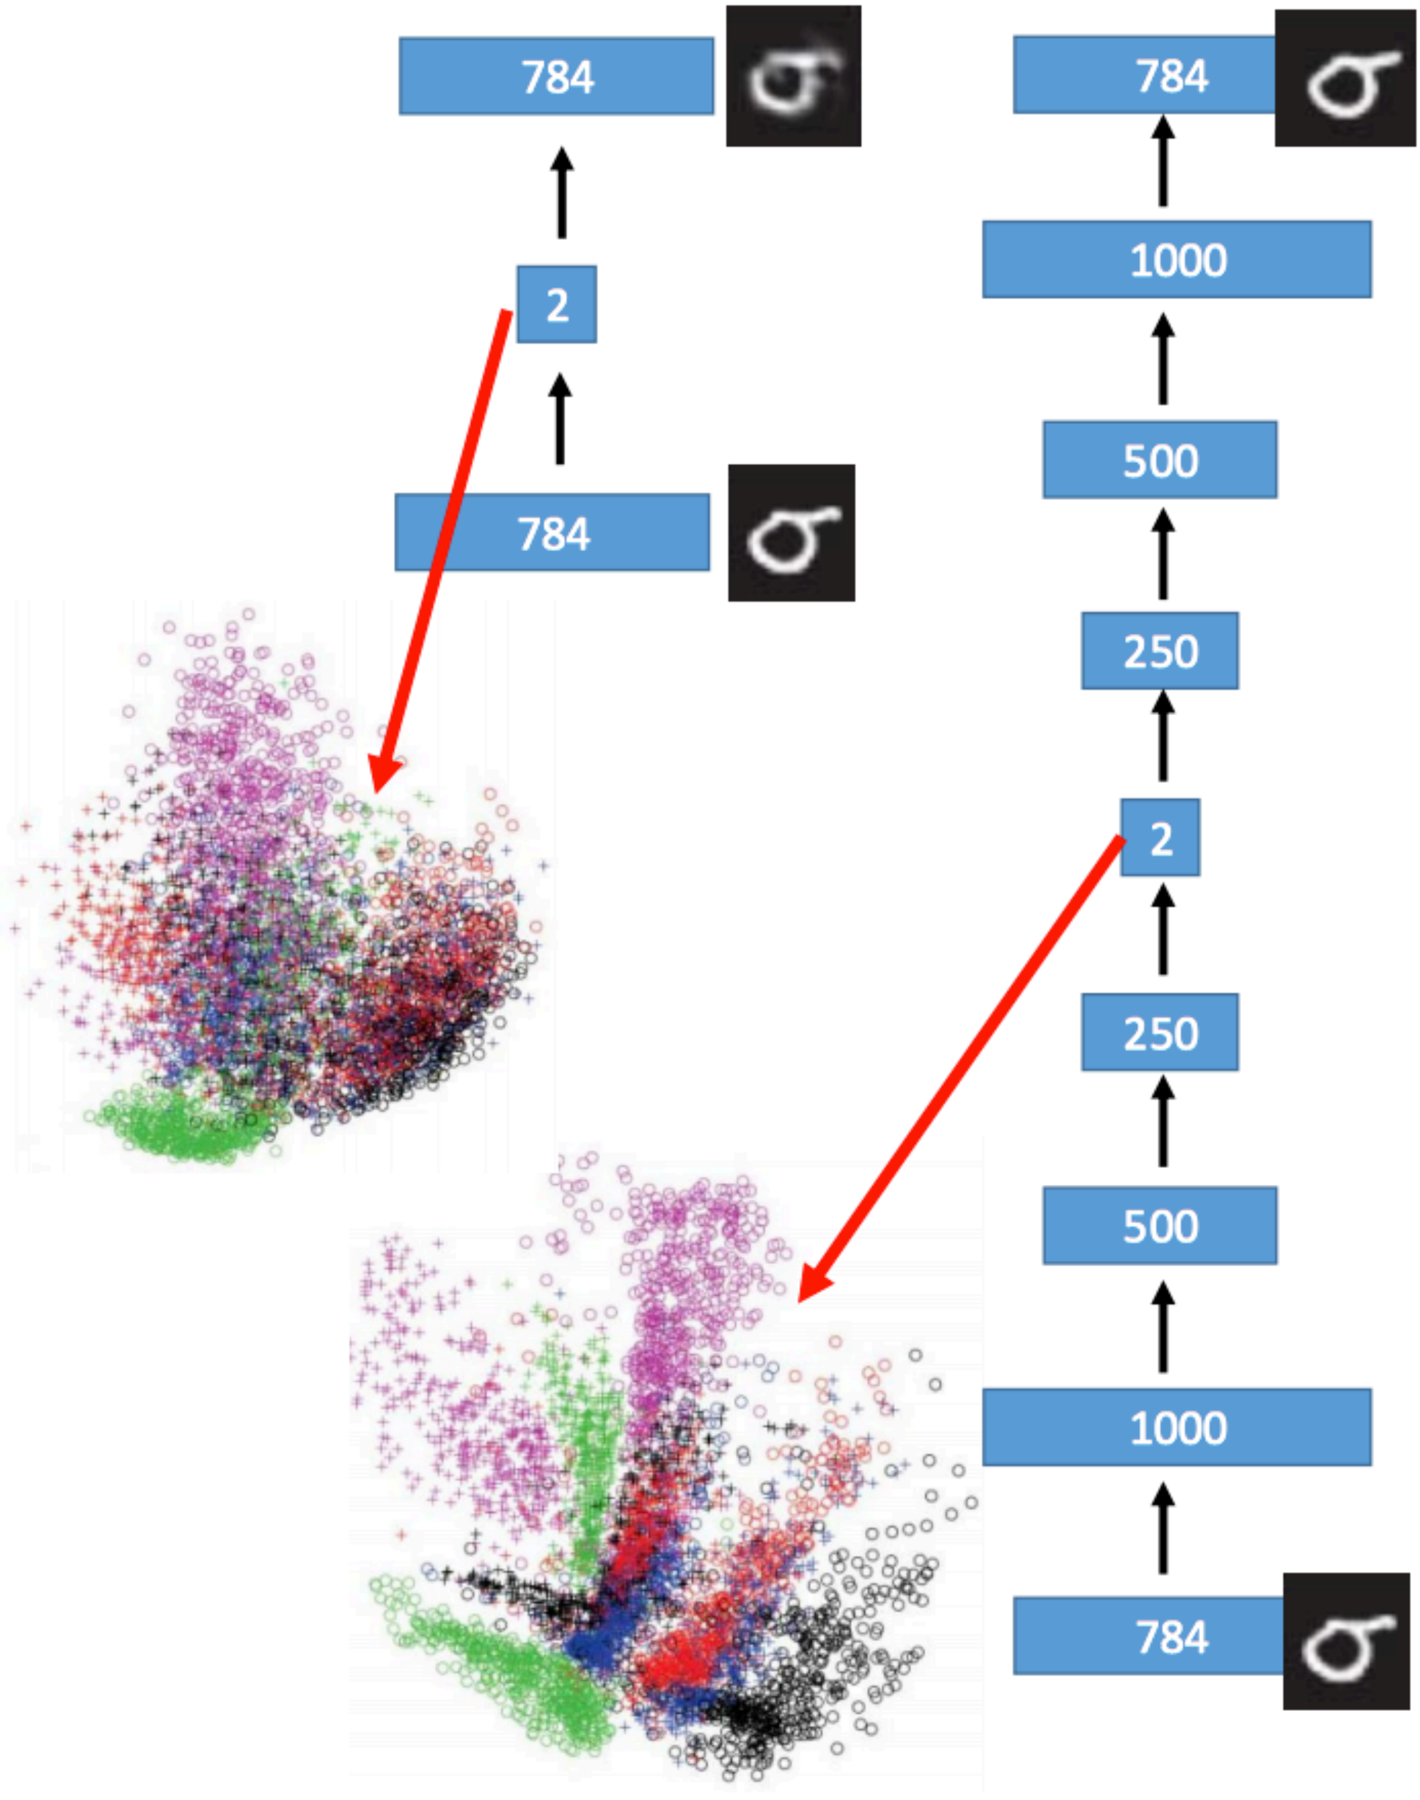

784
2
784
784
1000
500
250
2
250
500
1000
784

Deep Auto-encoder

- De-noising auto-encoder

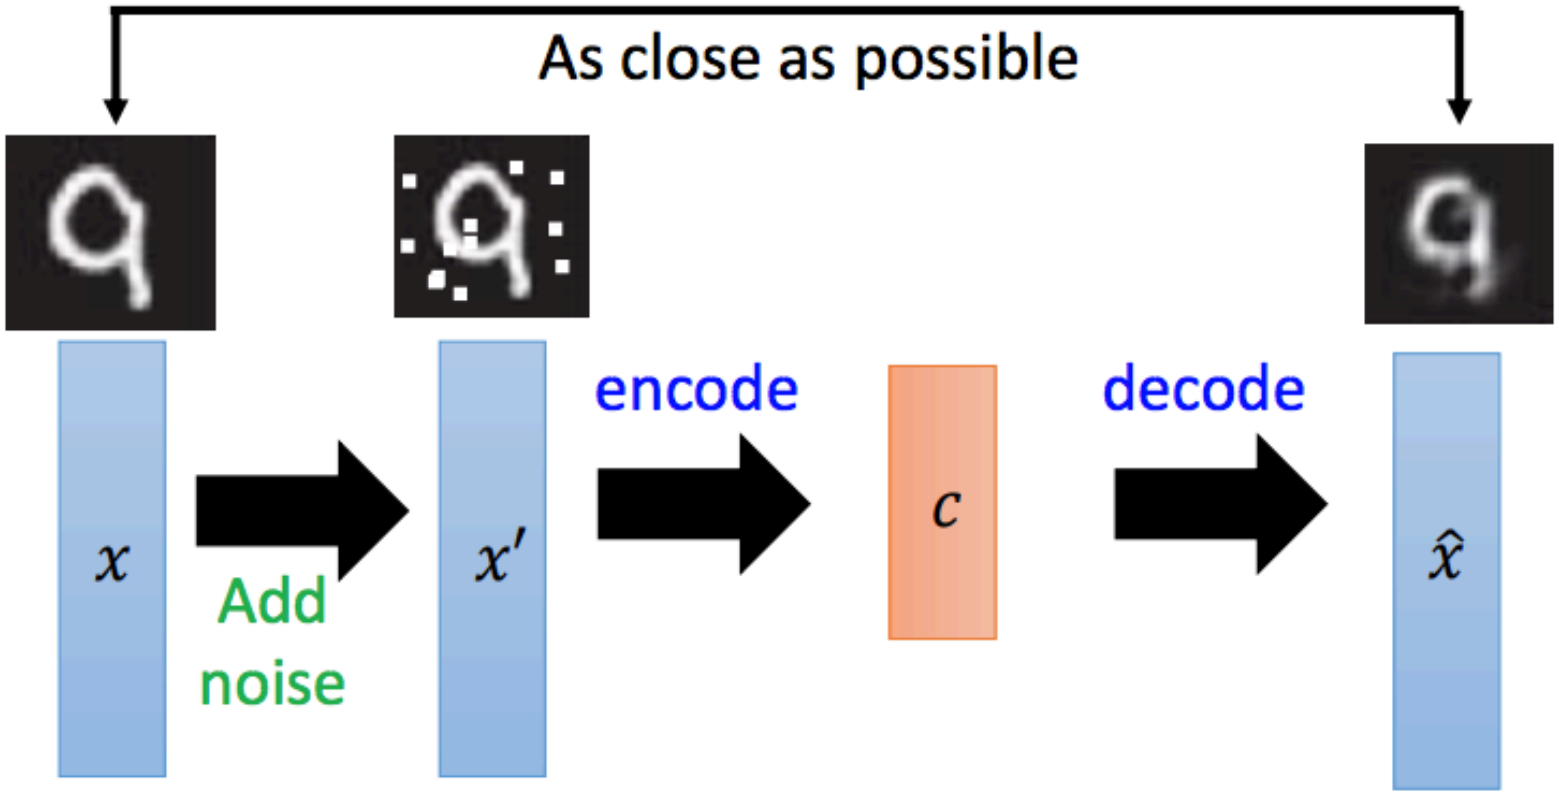

Vincent, Pascal, et al. "Extracting and composing robust features with denoising autoencoders." ICML, 2008.

Auto-encoder - Text Retrieval

Vector Space Model

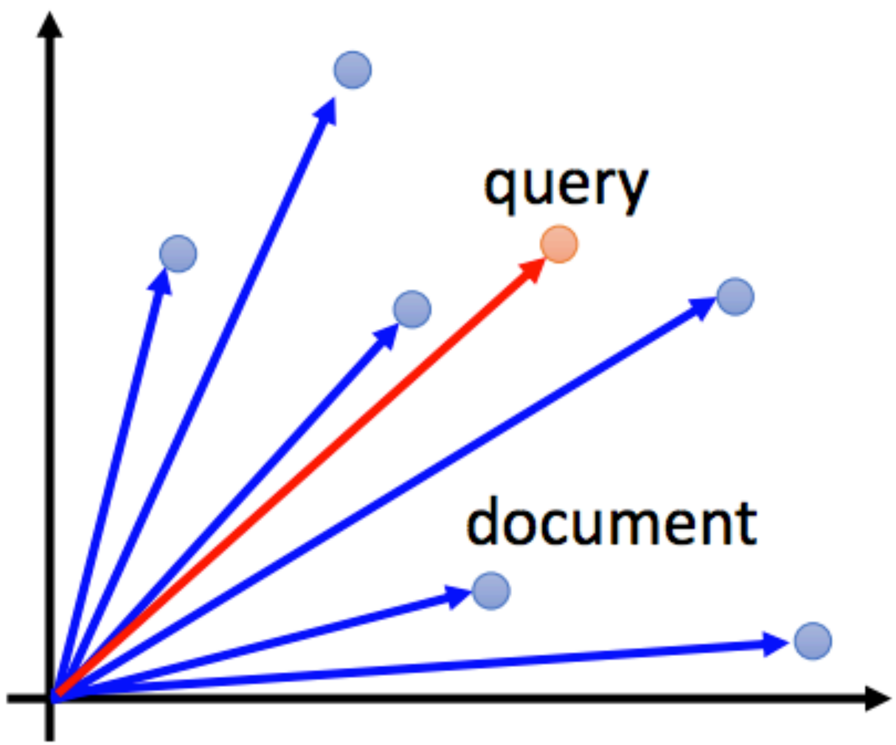

Bag-of-word

word string:
"This is an apple"

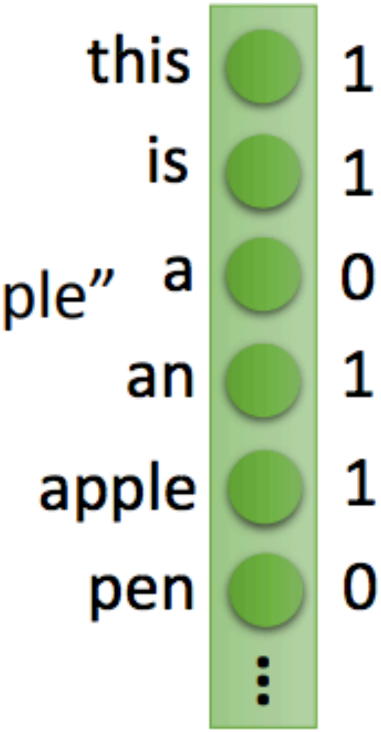

Semantics are not considered.

Auto-encoder - Text Retrieval

The documents talking about the same thing will have close code.

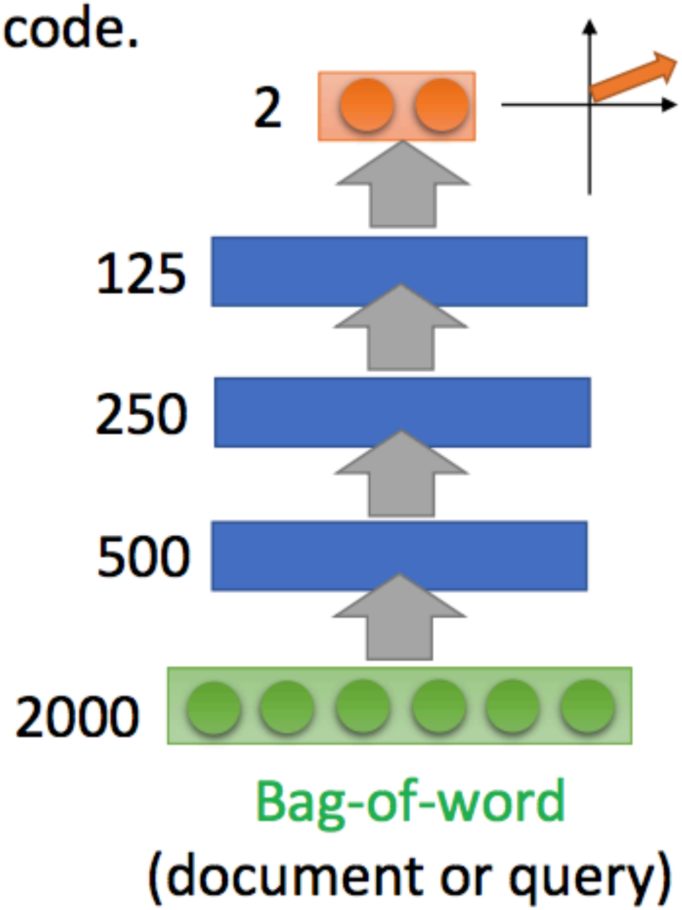

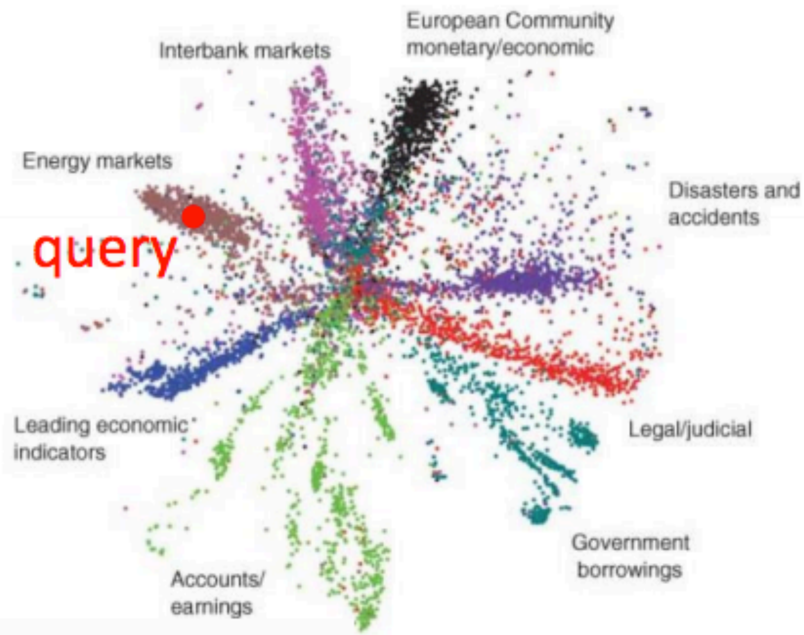

LSA: project documents to 2 latent topics

Auto-encoder – Similar Image Search

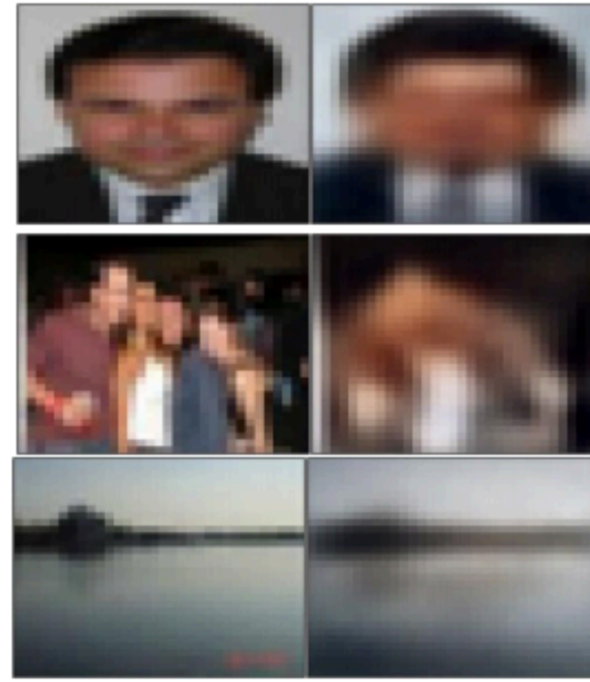

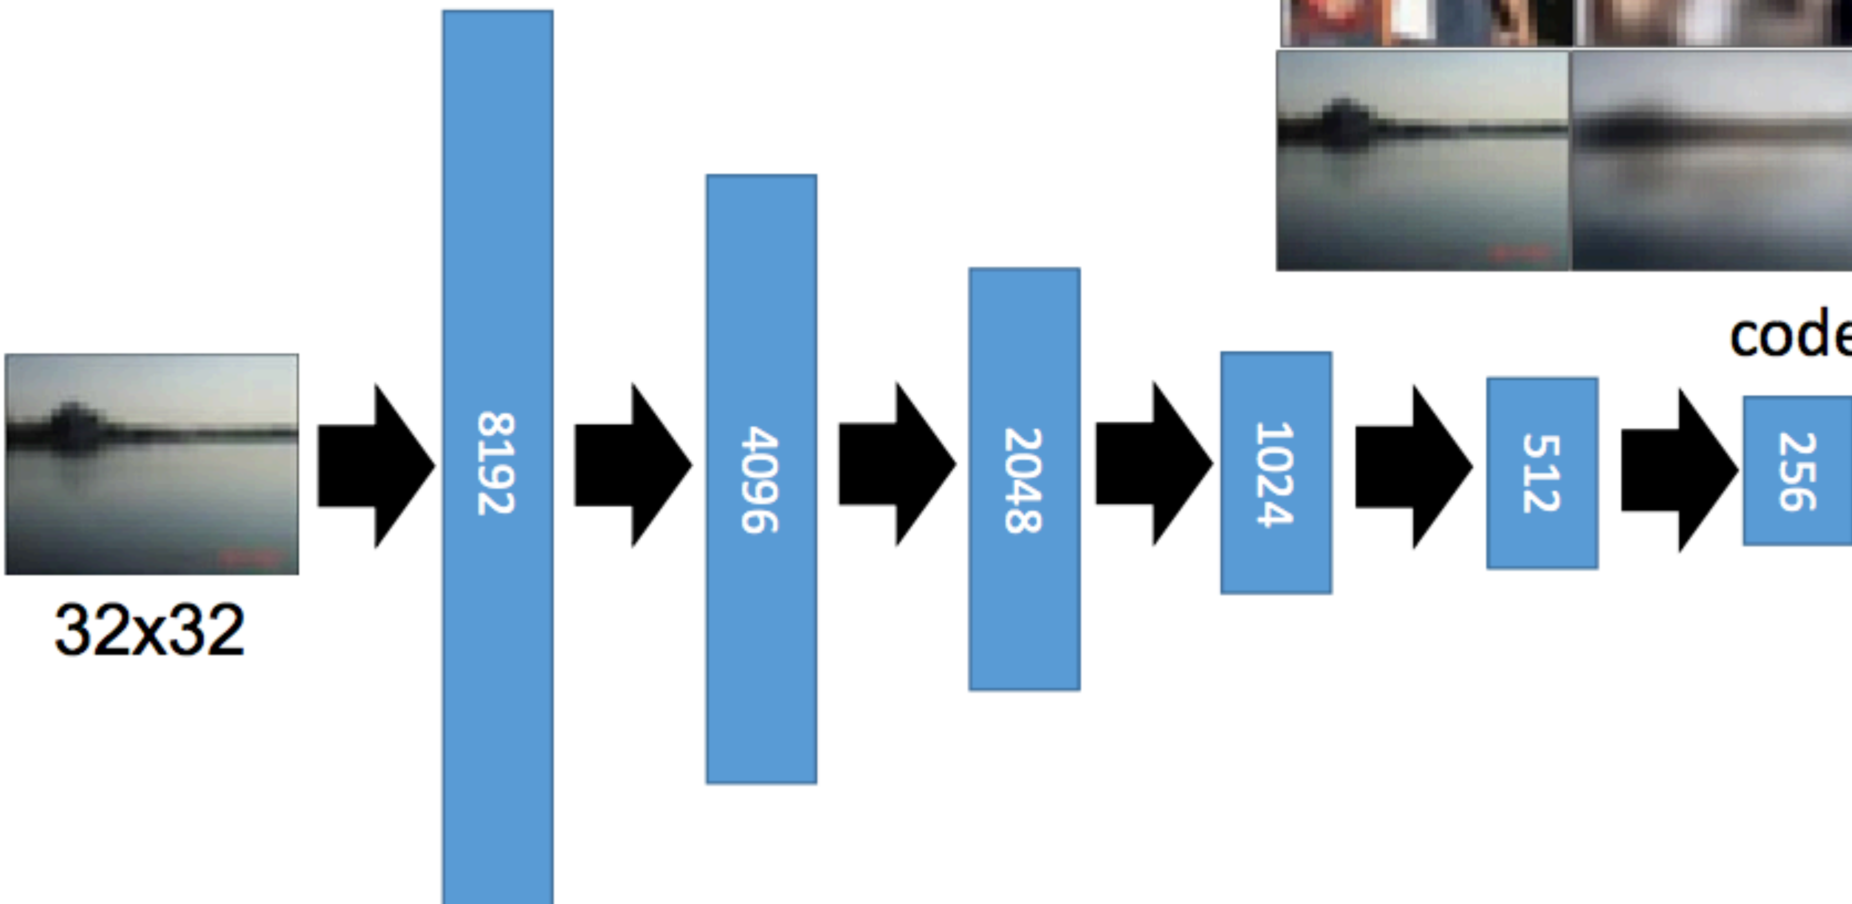

Retrieved using Euclidean distance in pixel intensity space

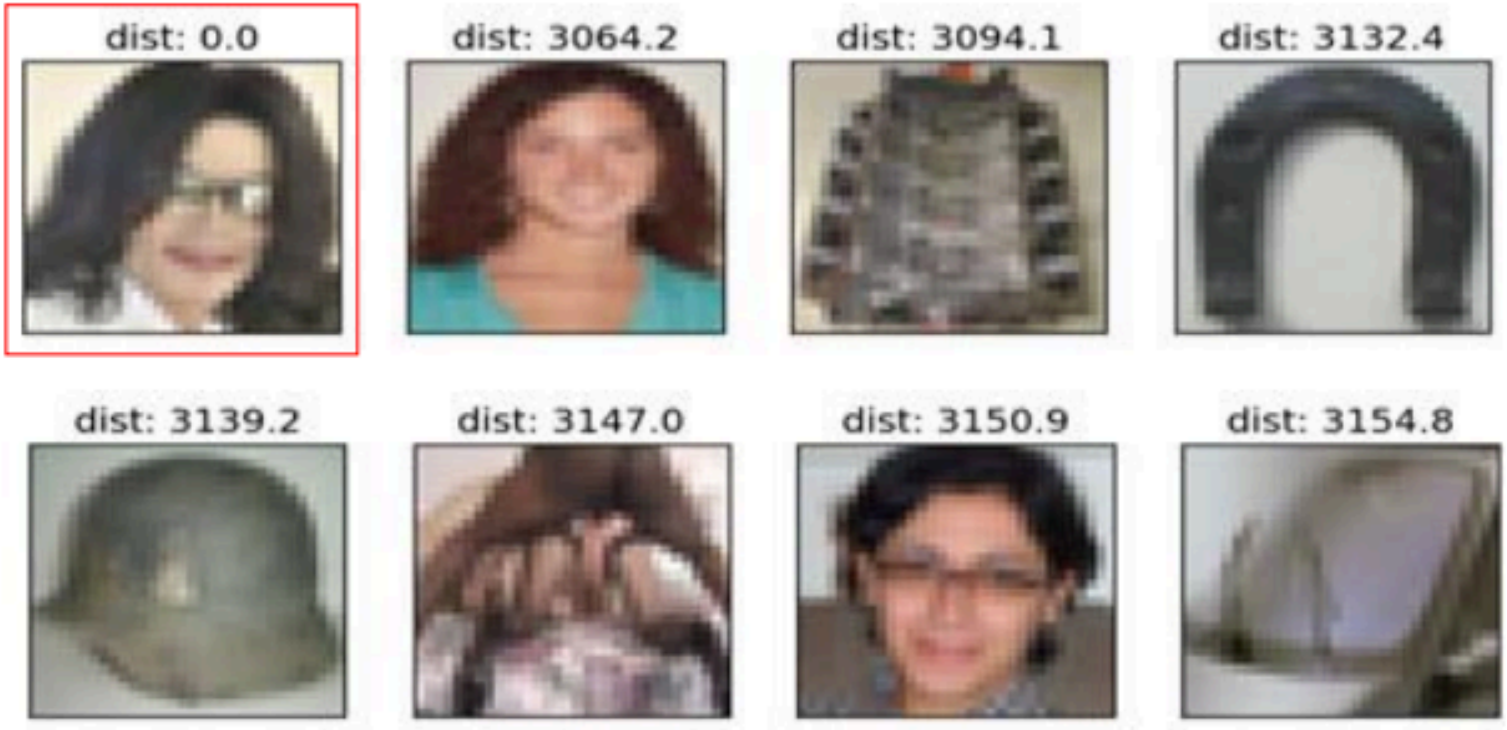

retrieved using 256 codes

Auto-encoder - CNN

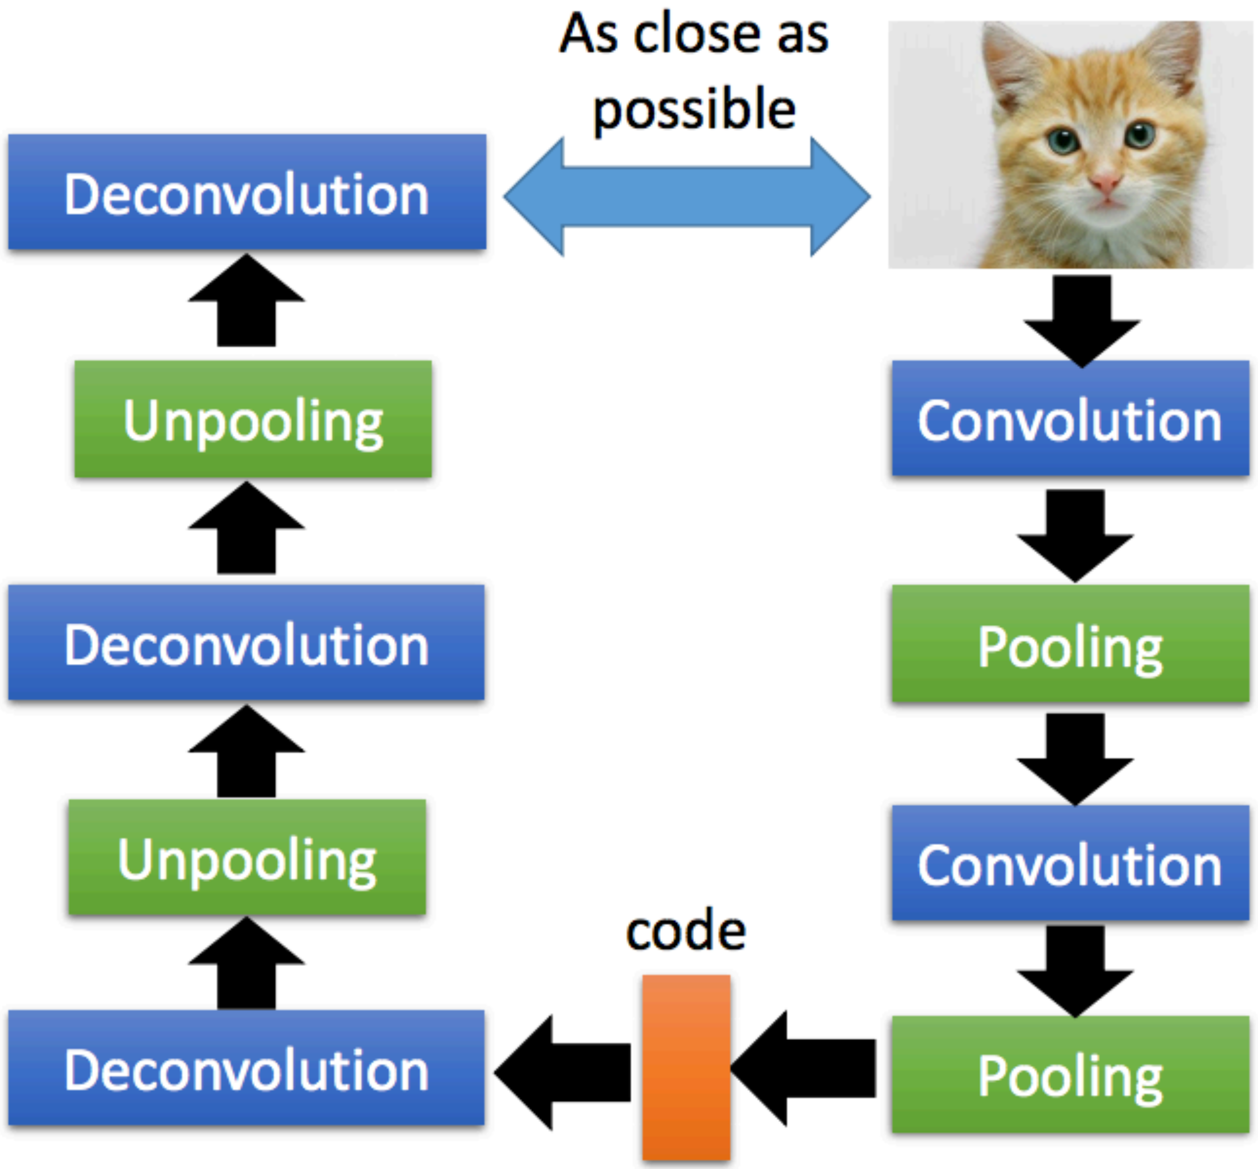

CNN -Unpooling

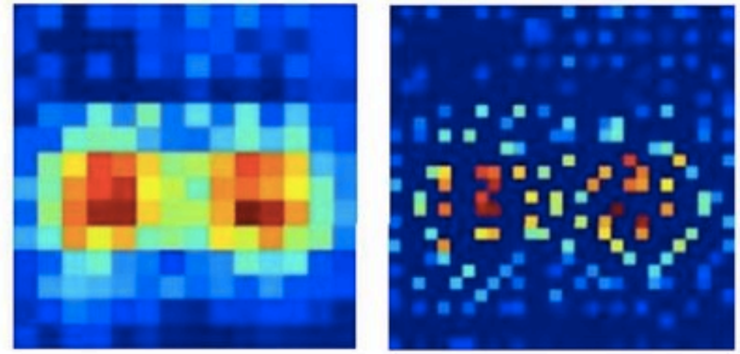

14 x 14 28 x 28

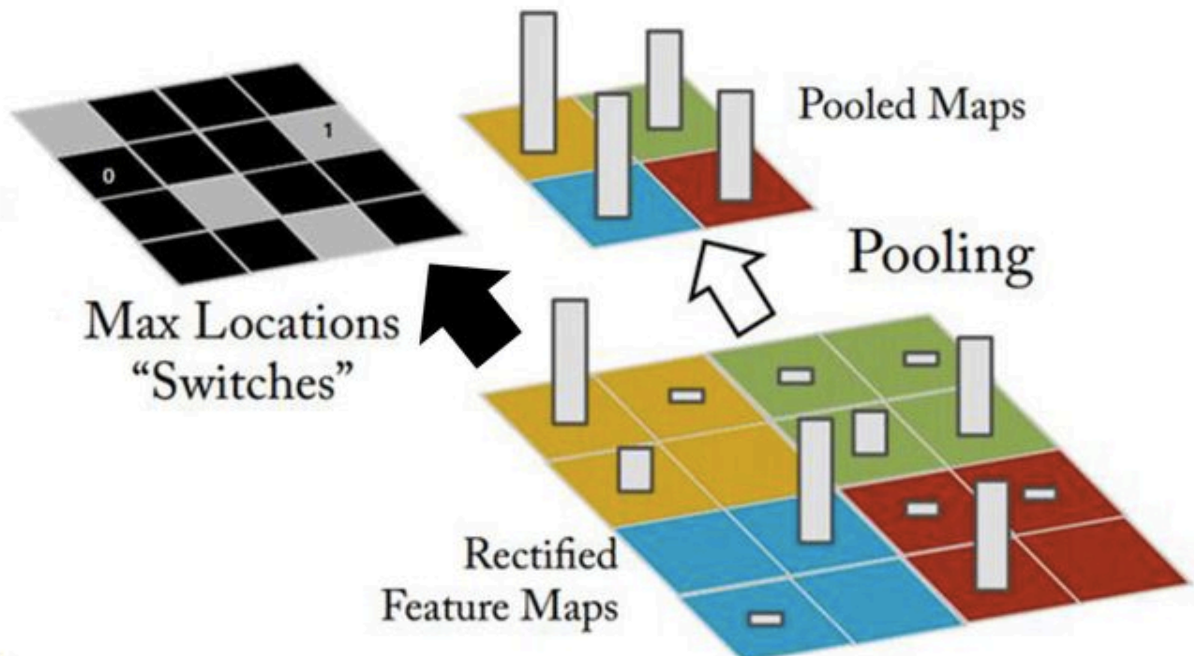

Alternative: simply repeat the values

Source of image :
https://leonardoaraujosantos.gitbooks.io/artificial-inteligence/content/image_segmentation.html

Actually, deconvolution is convolution.

CNN
- Deconvolution

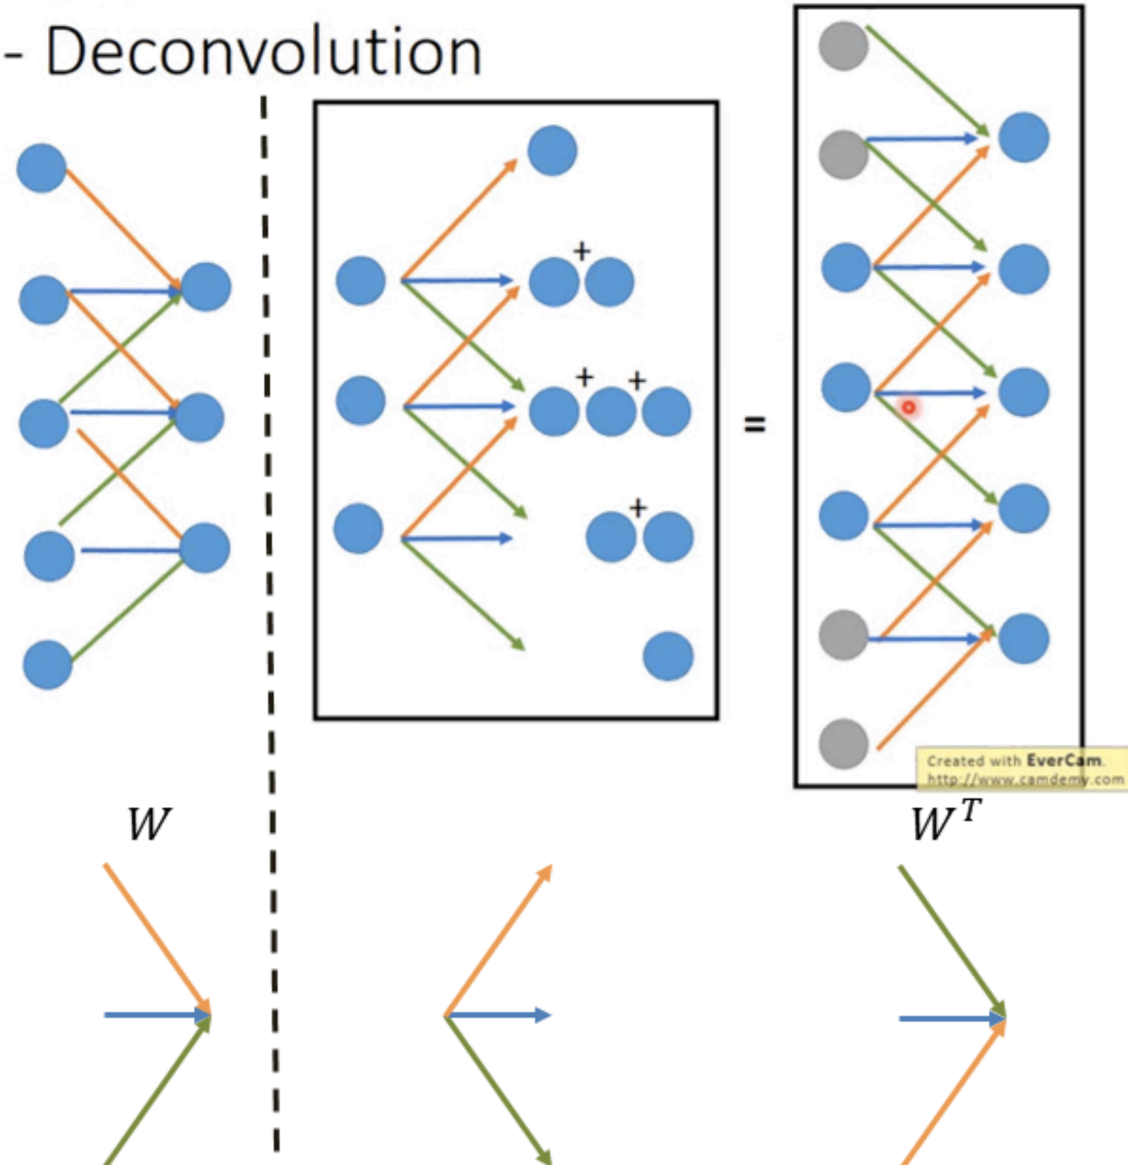

$conv(W)$ | $deconv(W) = conv(W^T)$

In Tensorflow $= conv\_transpose(W)$

Auto-encoder – Pre-training

- Greedy layer-wise pre-training

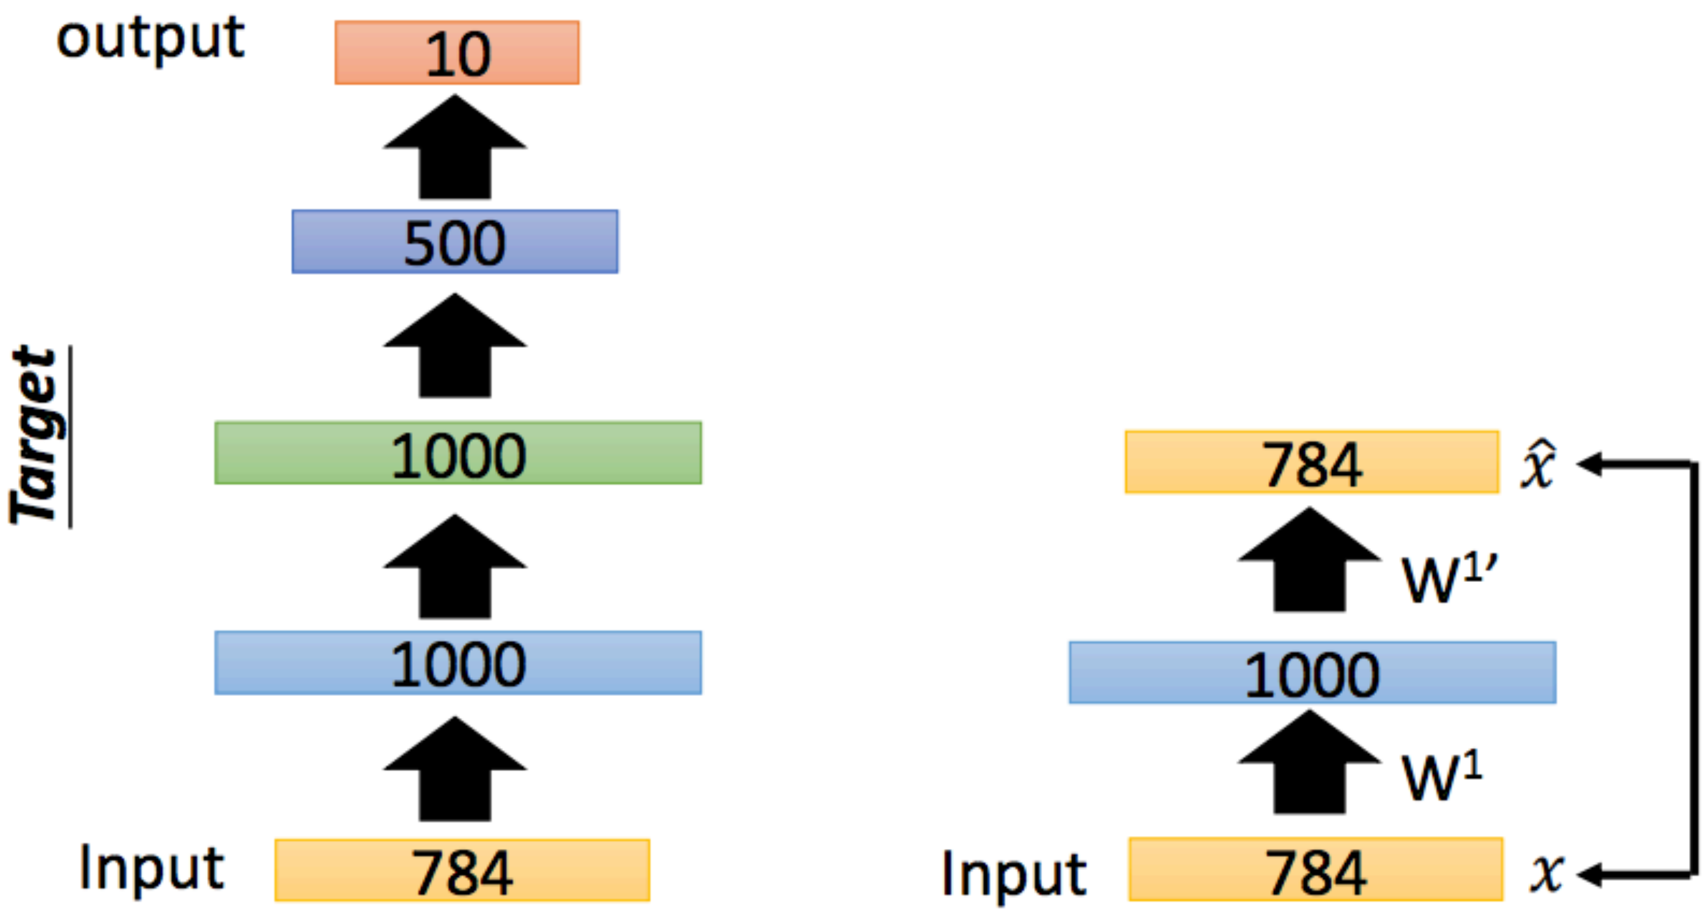

Auto-encoder – Pre-training

- Greedy layer-wise pre-training

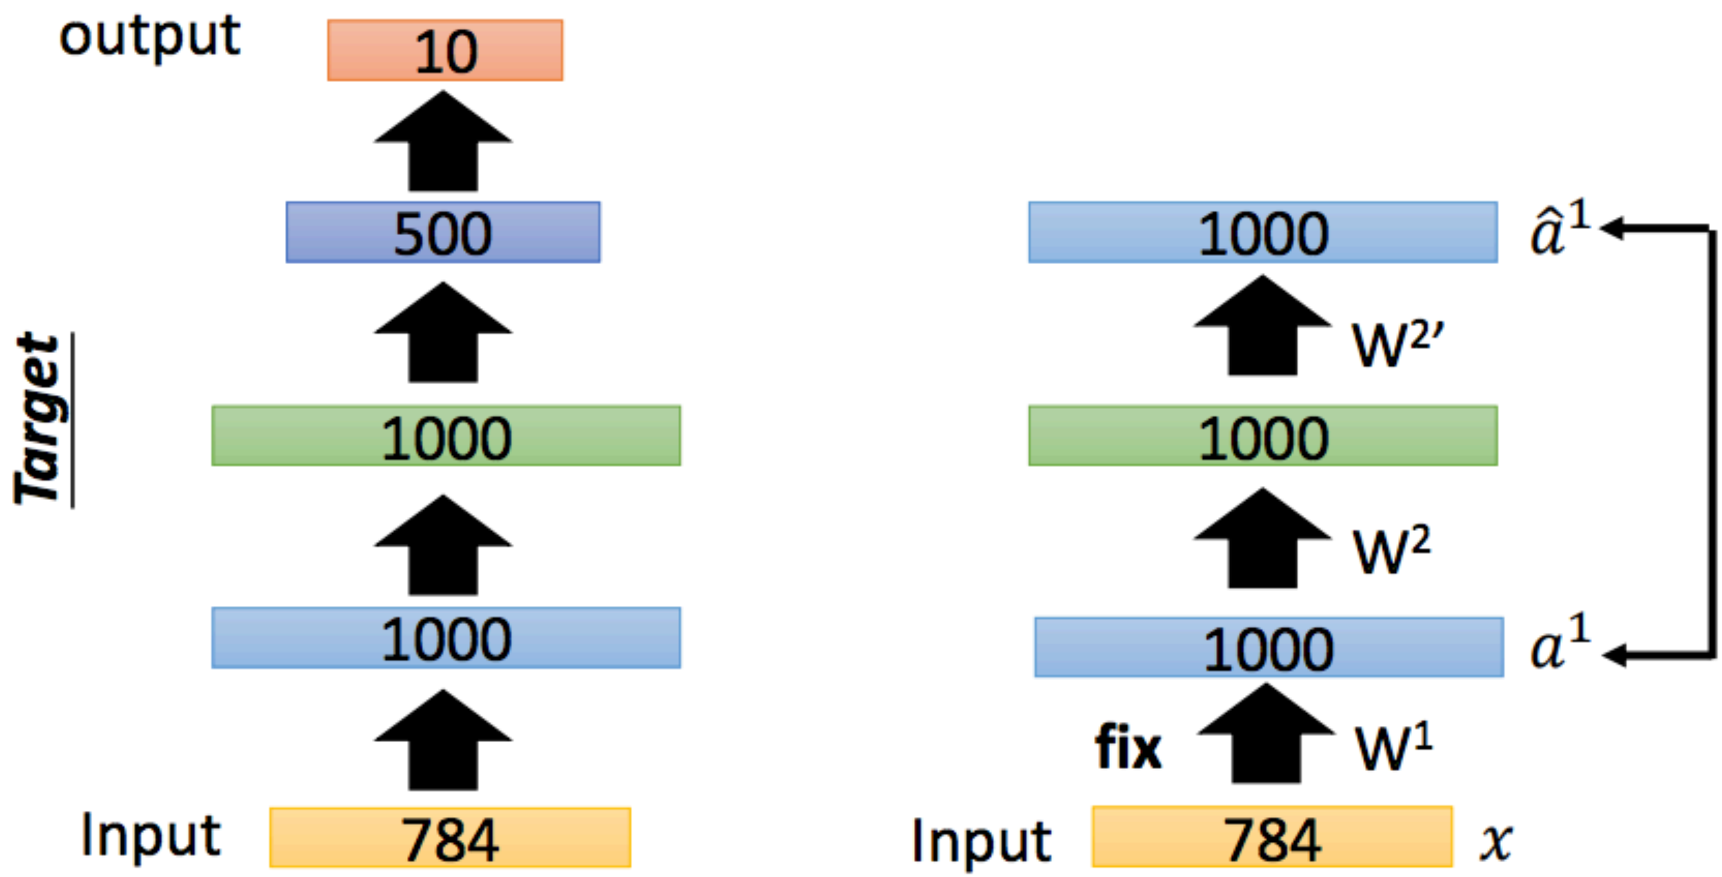

Auto-encoder – Pre-training

- Greedy layer-wise pre-training

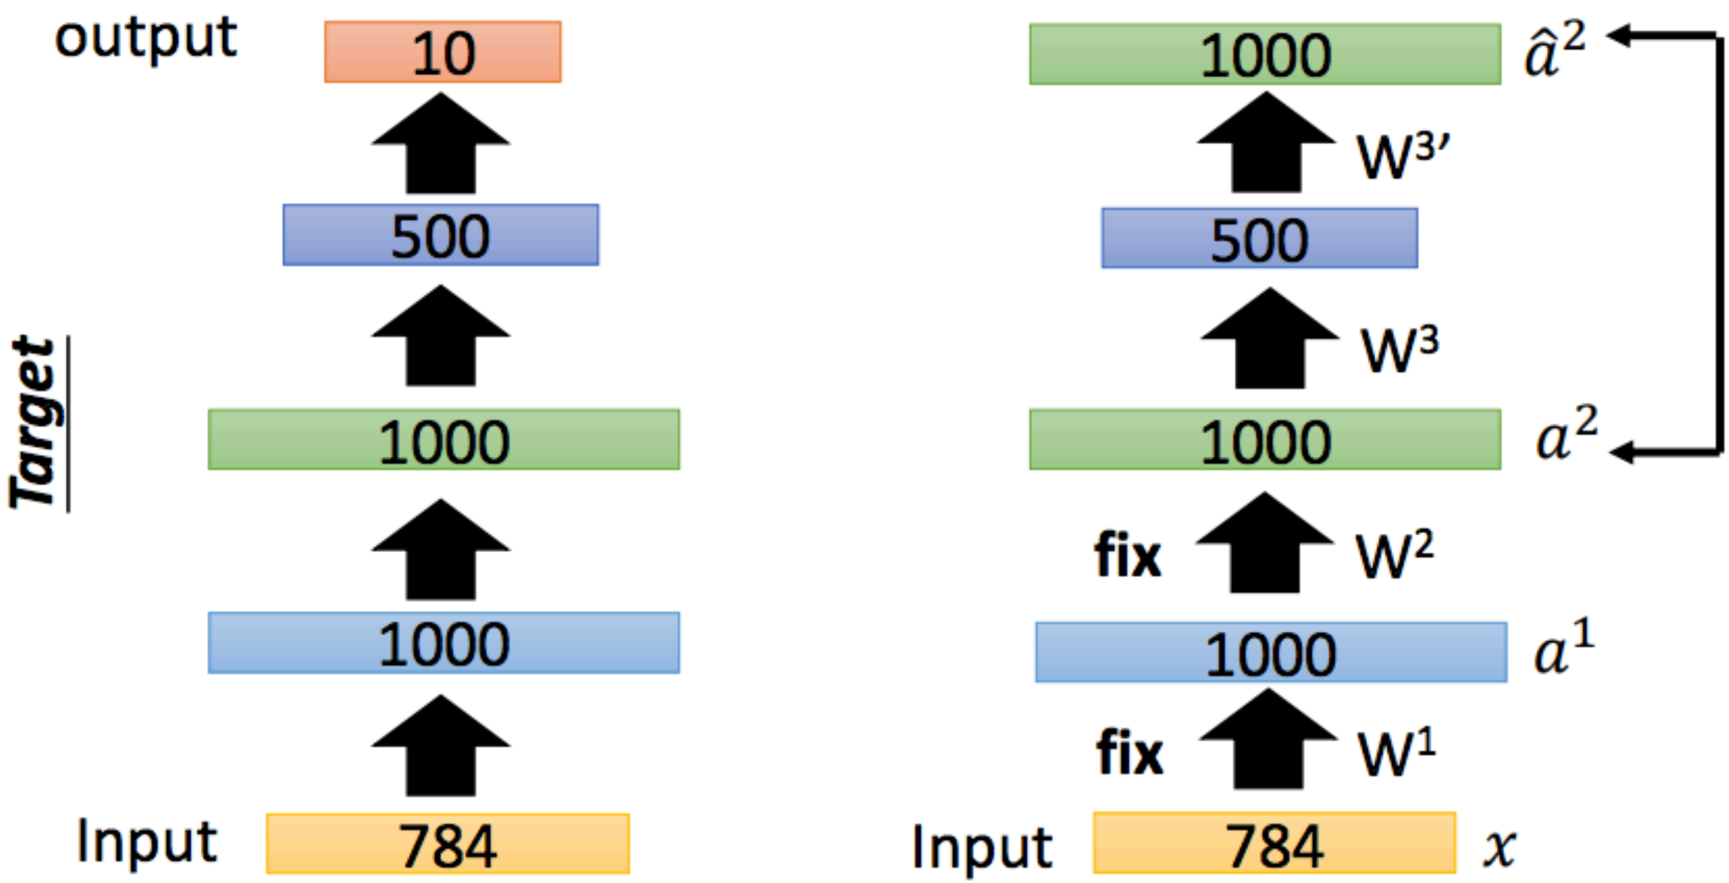

Auto-encoder – Pre-training

Find-tune by backpropagation

- Greedy layer-wise pre-training

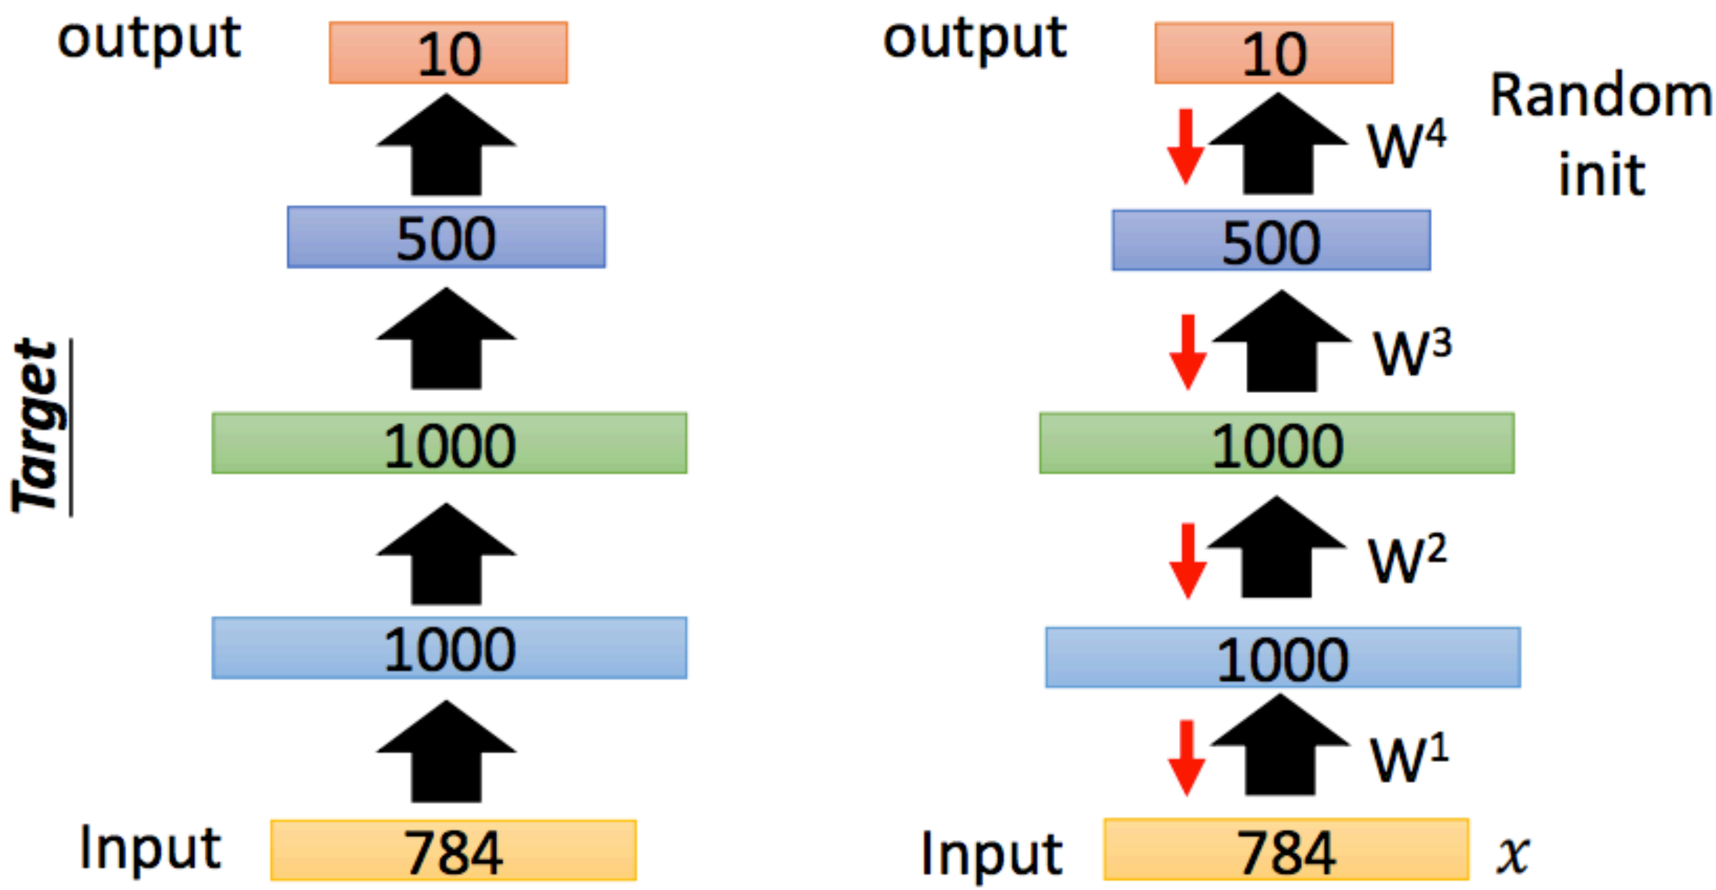

More about Auto-encoder

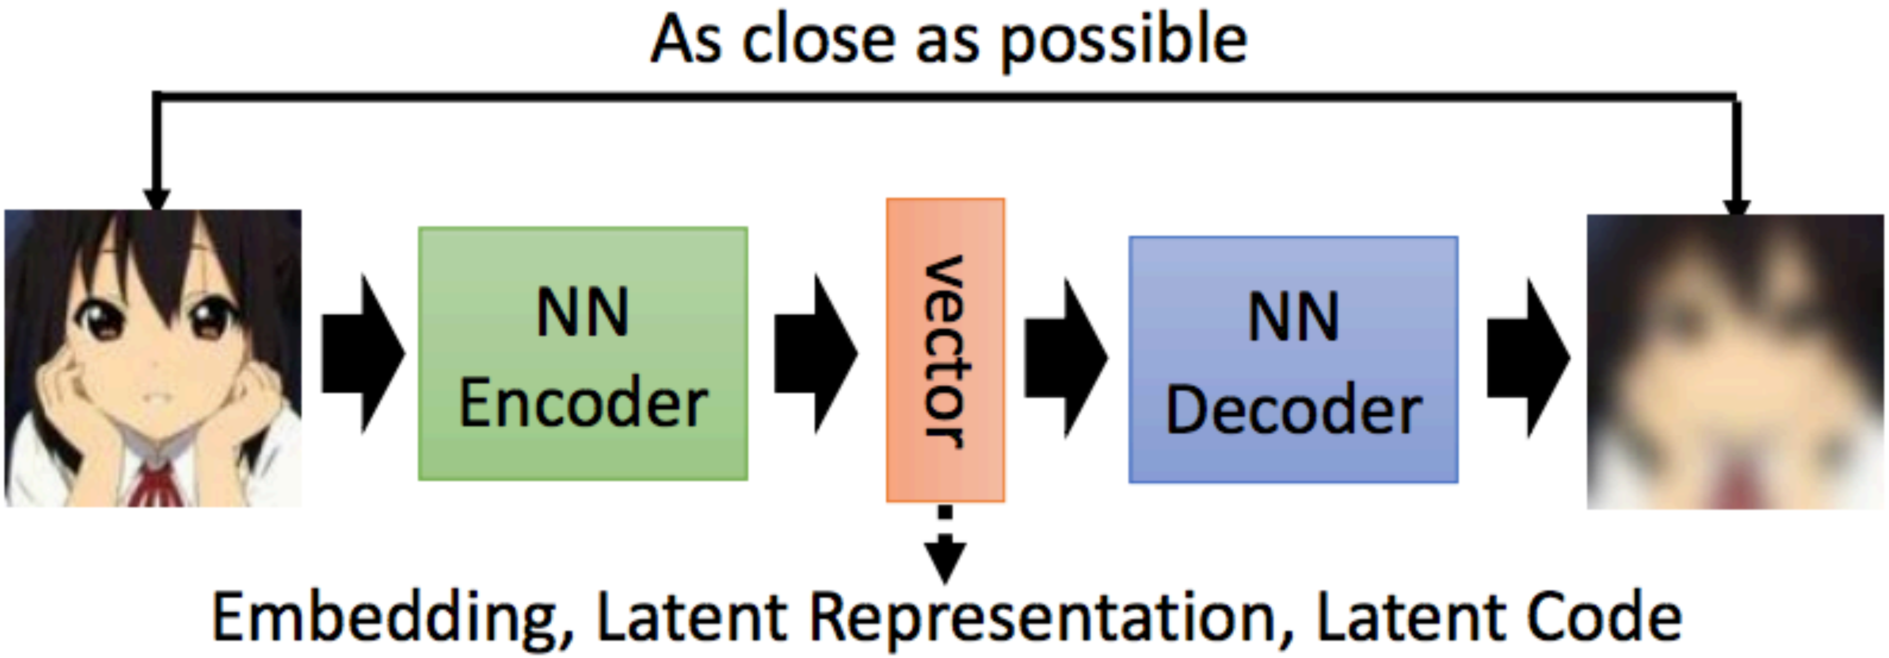

- More than minimizing reconstruction error
- More interpretable embedding

1. Beyond Reconstruction

- How to evaluate an encoder?
  - Loss of the classification task is L_D

Train ϕ to minimize L_D

$$L_D^* = \min_{\phi} L_D$$

Small L_D^* ➡ The embeddings are representative.

Large L_D^* ➡ Not representative

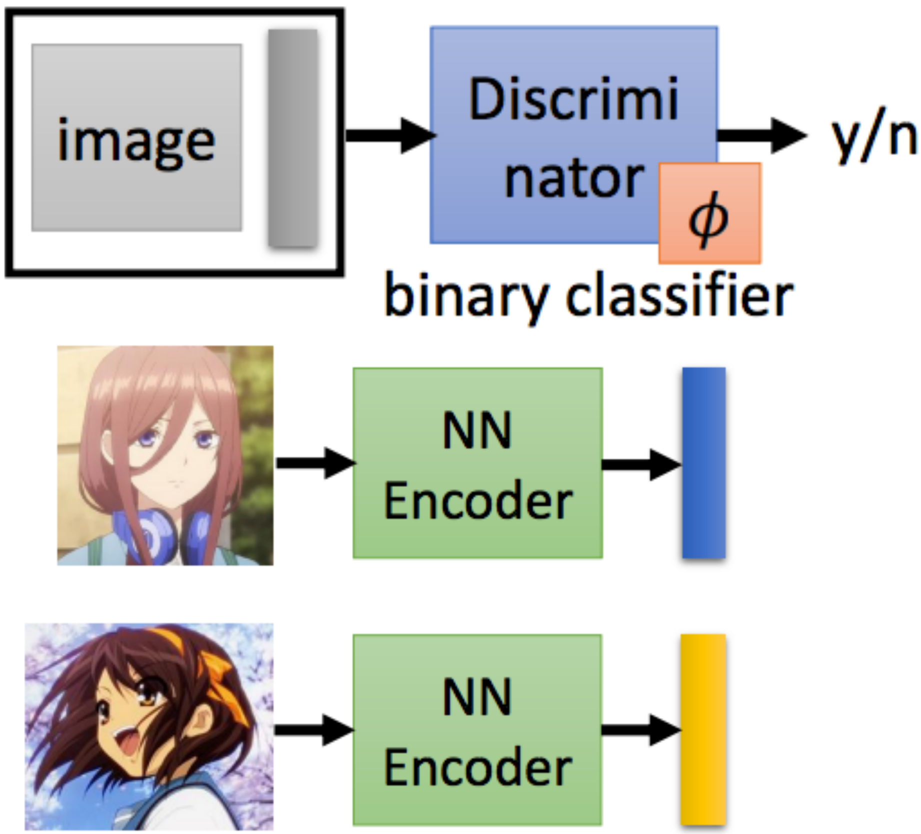

1. Beyond Reconstruction

- How to evaluate an encoder?
  - Loss of the classification task is L_D

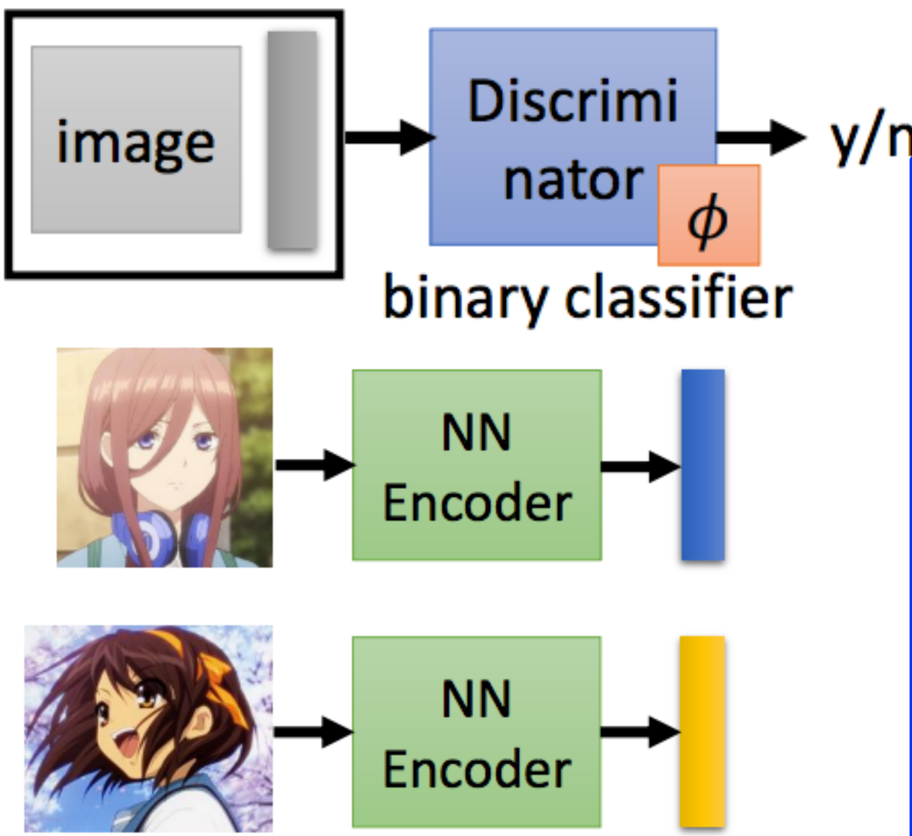

Train ϕ to minimize L_D

$$L_D^* = \min_{\phi} L_D$$

Small L_D^* ➡ The embeddings are representative.

Large L_D^* ➡ Not representative

Train θ to minimize L_D^*

$$\theta^* = arg \min_{\theta} L_D^*$$

$$= arg \min_{\theta} \min_{\phi} L_D$$

Train the encoder θ and discriminator ϕ to minimize L_D

Deep InfoMax (DIM)

(c.f. training encoder and decoder to minimize reconstruction error)

1. Beyond Reconstruction

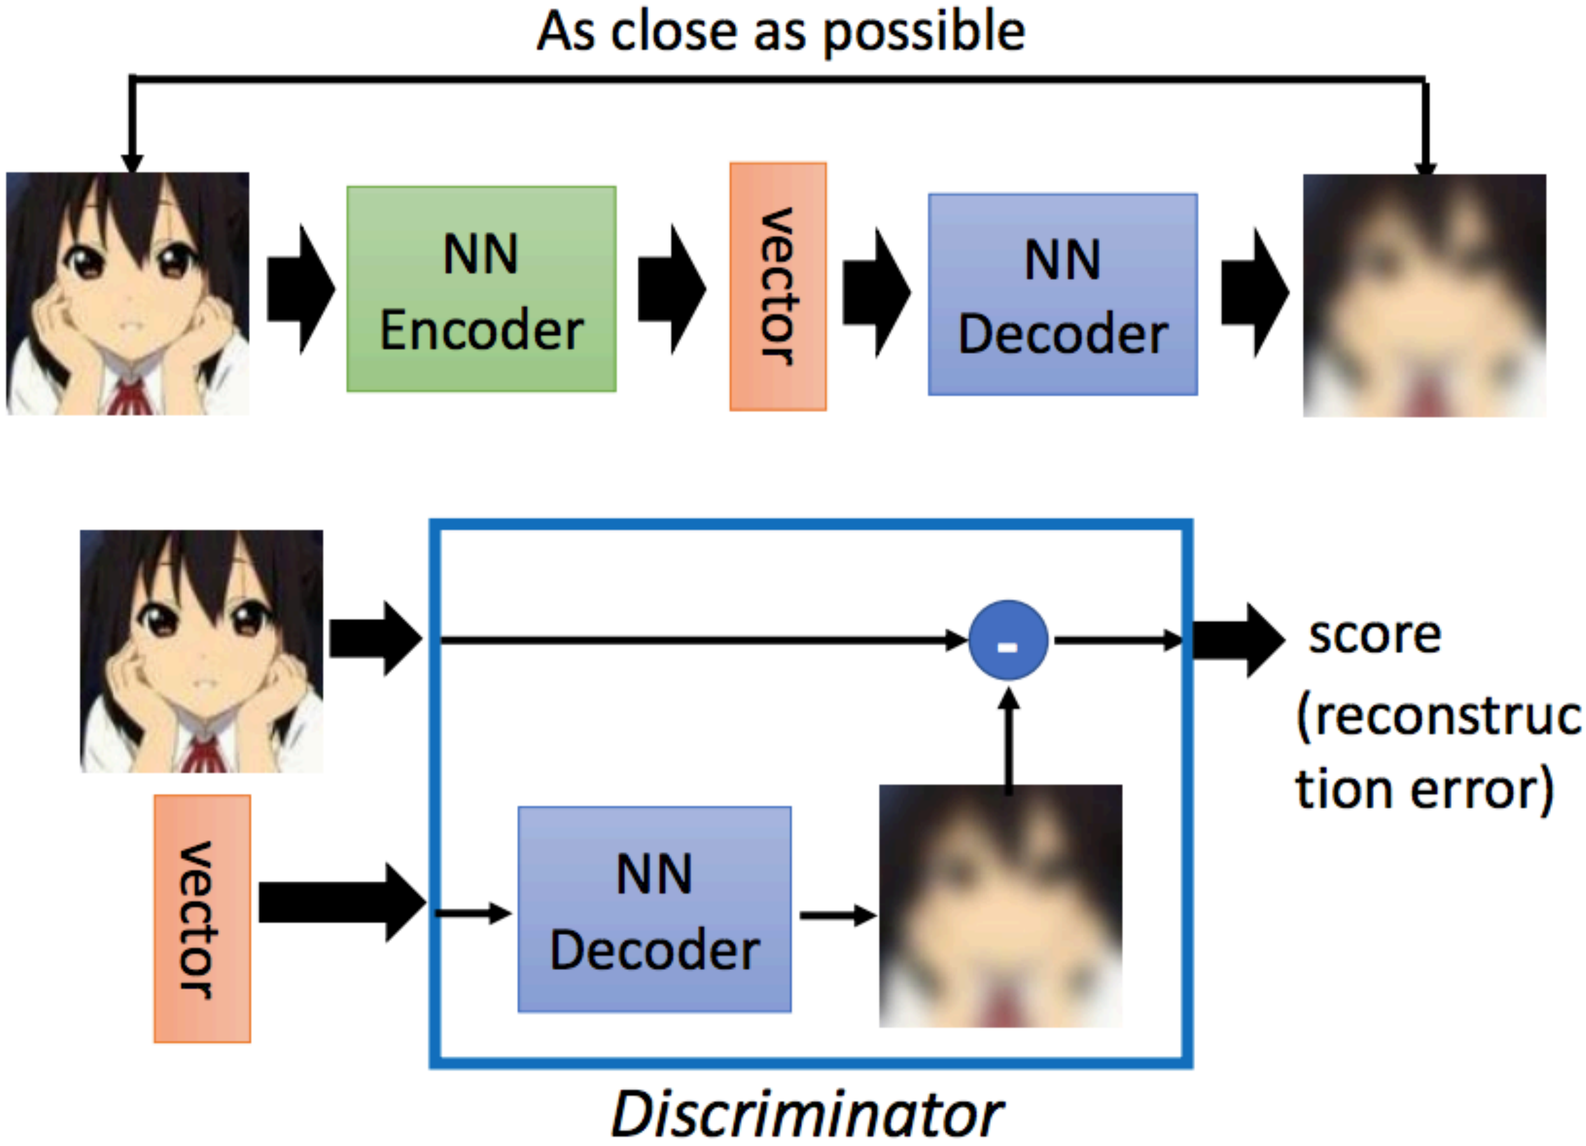

2. Sequential Data

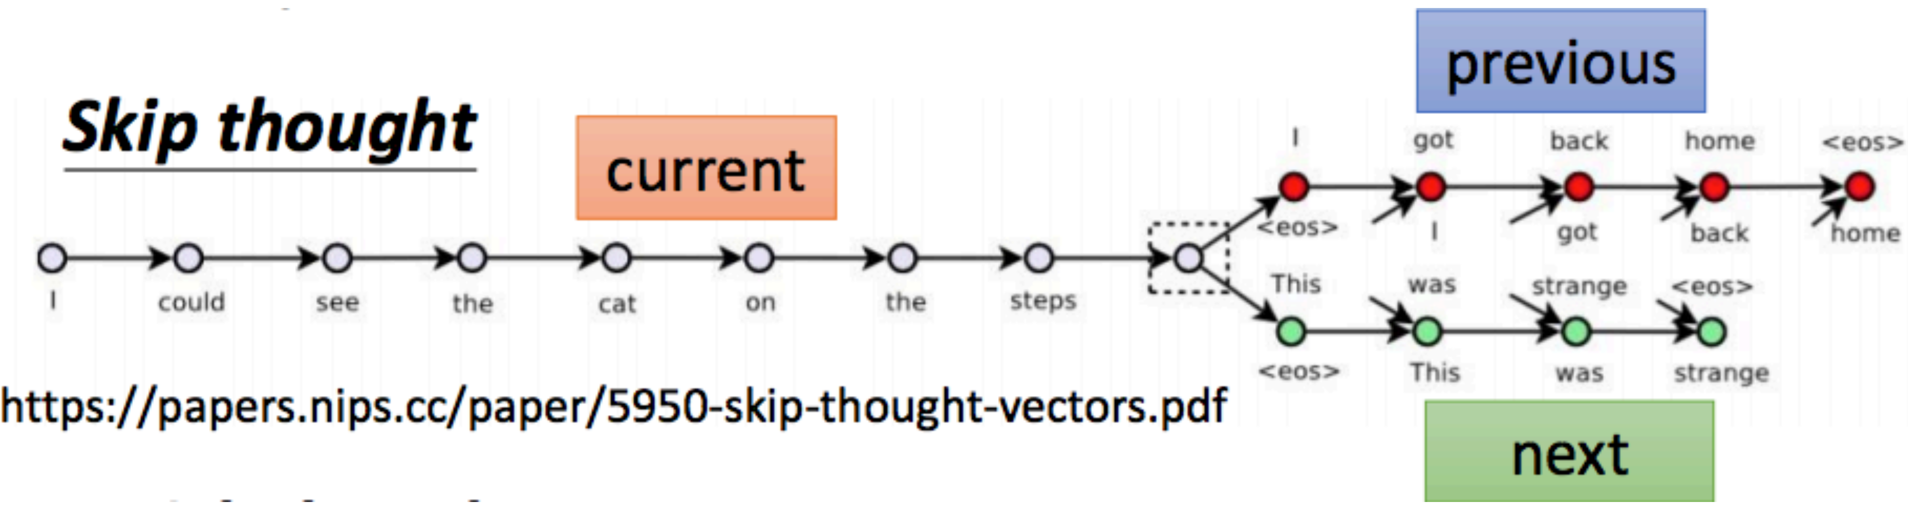

https://papers.nips.cc/paper/5950-skip-thought-vectors.pdf

2. Sequential Data

Quick thought

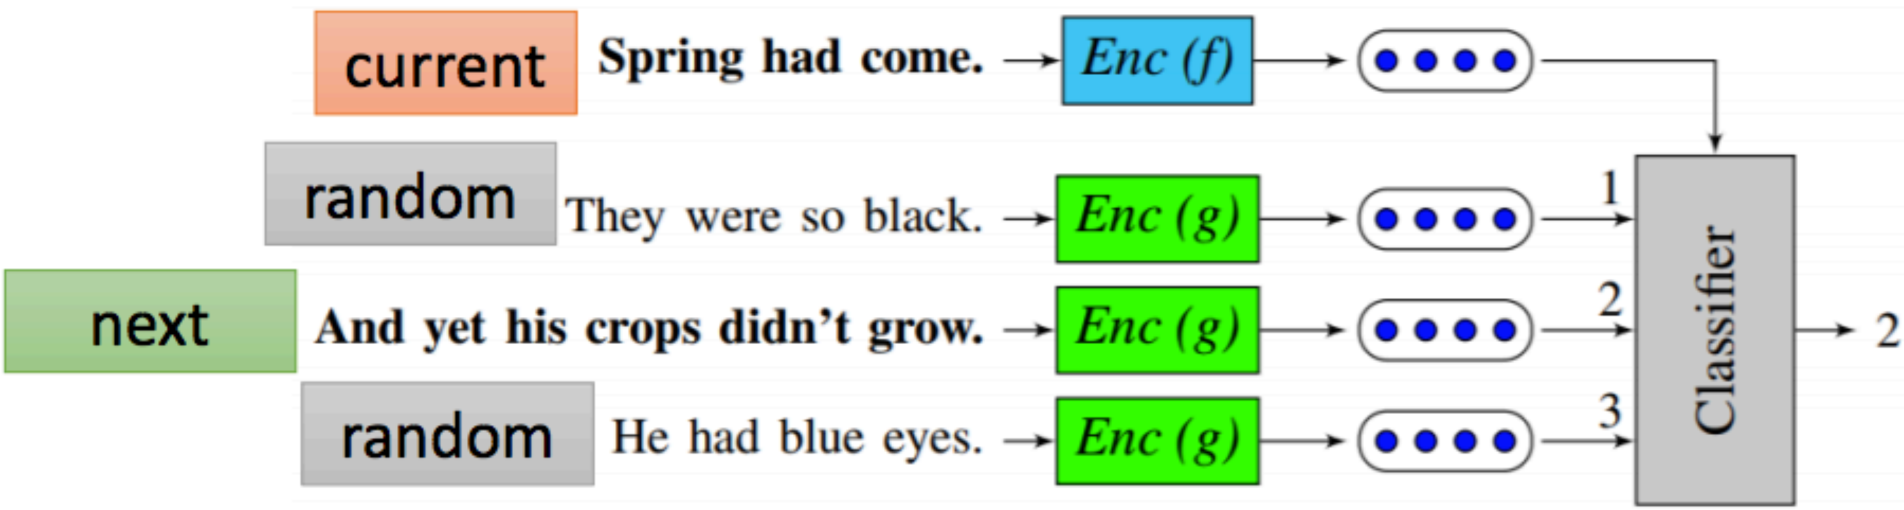

https://arxiv.org/pdf/1803.02893.pdf

3. Feature Disentangle

- An object contains multiple aspect information

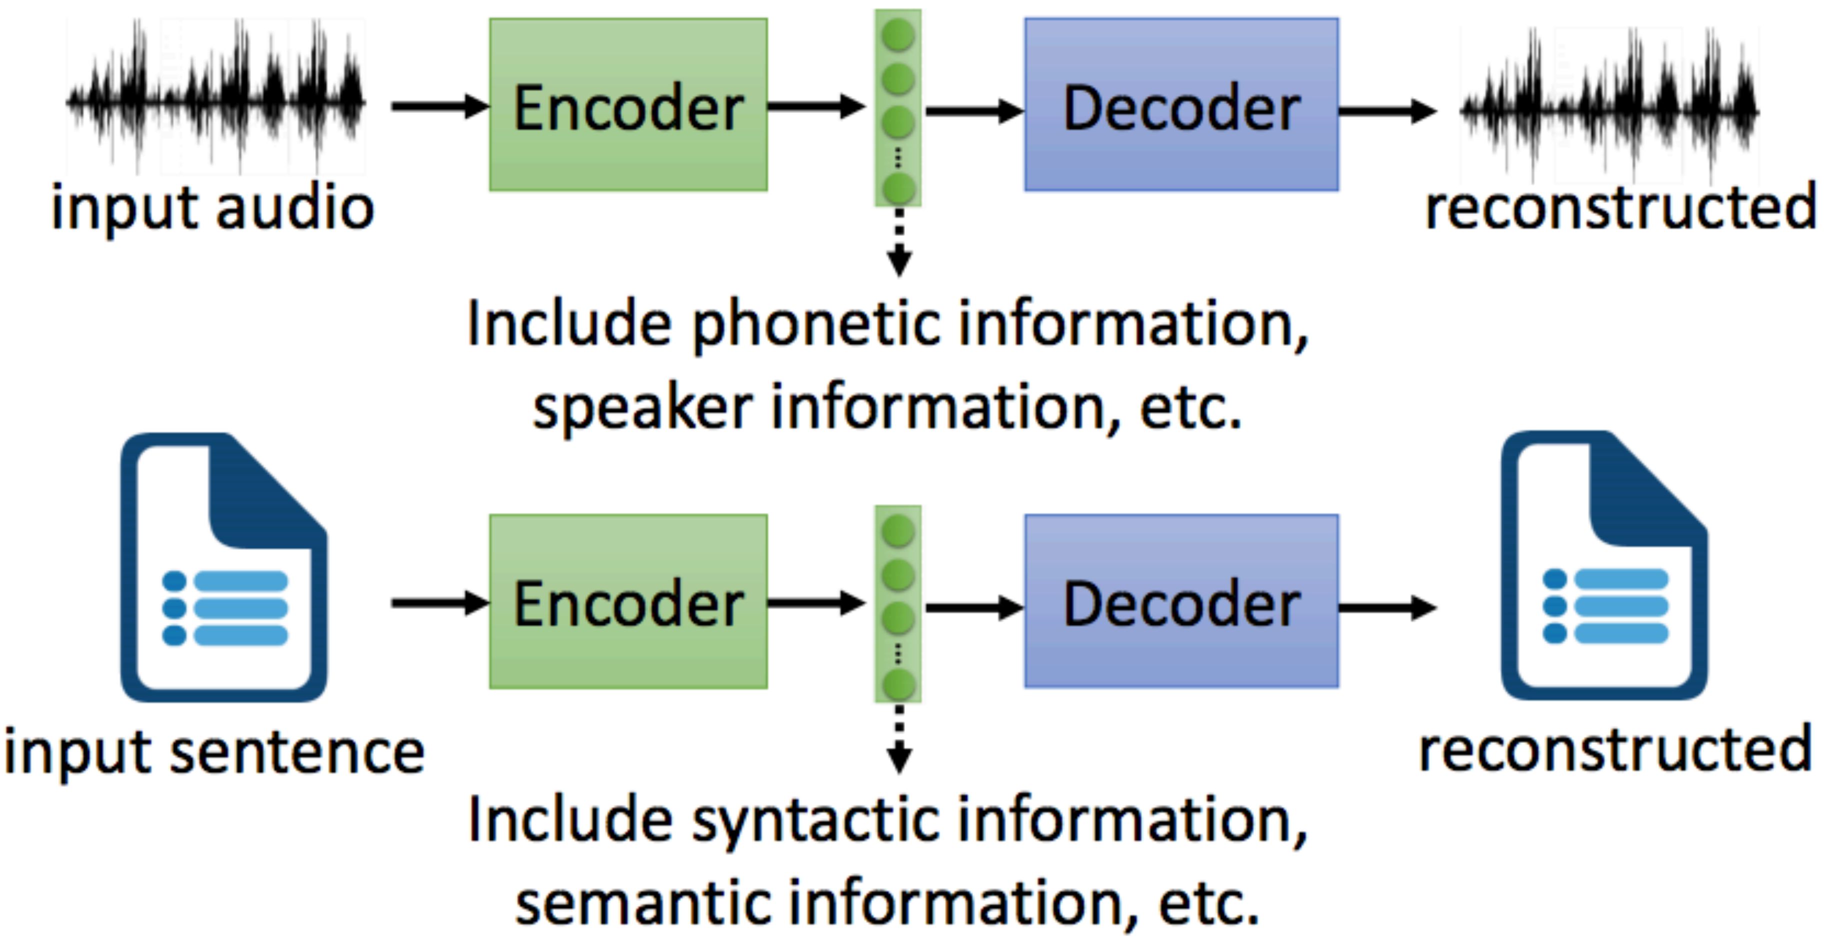

3. Feature Disentangle

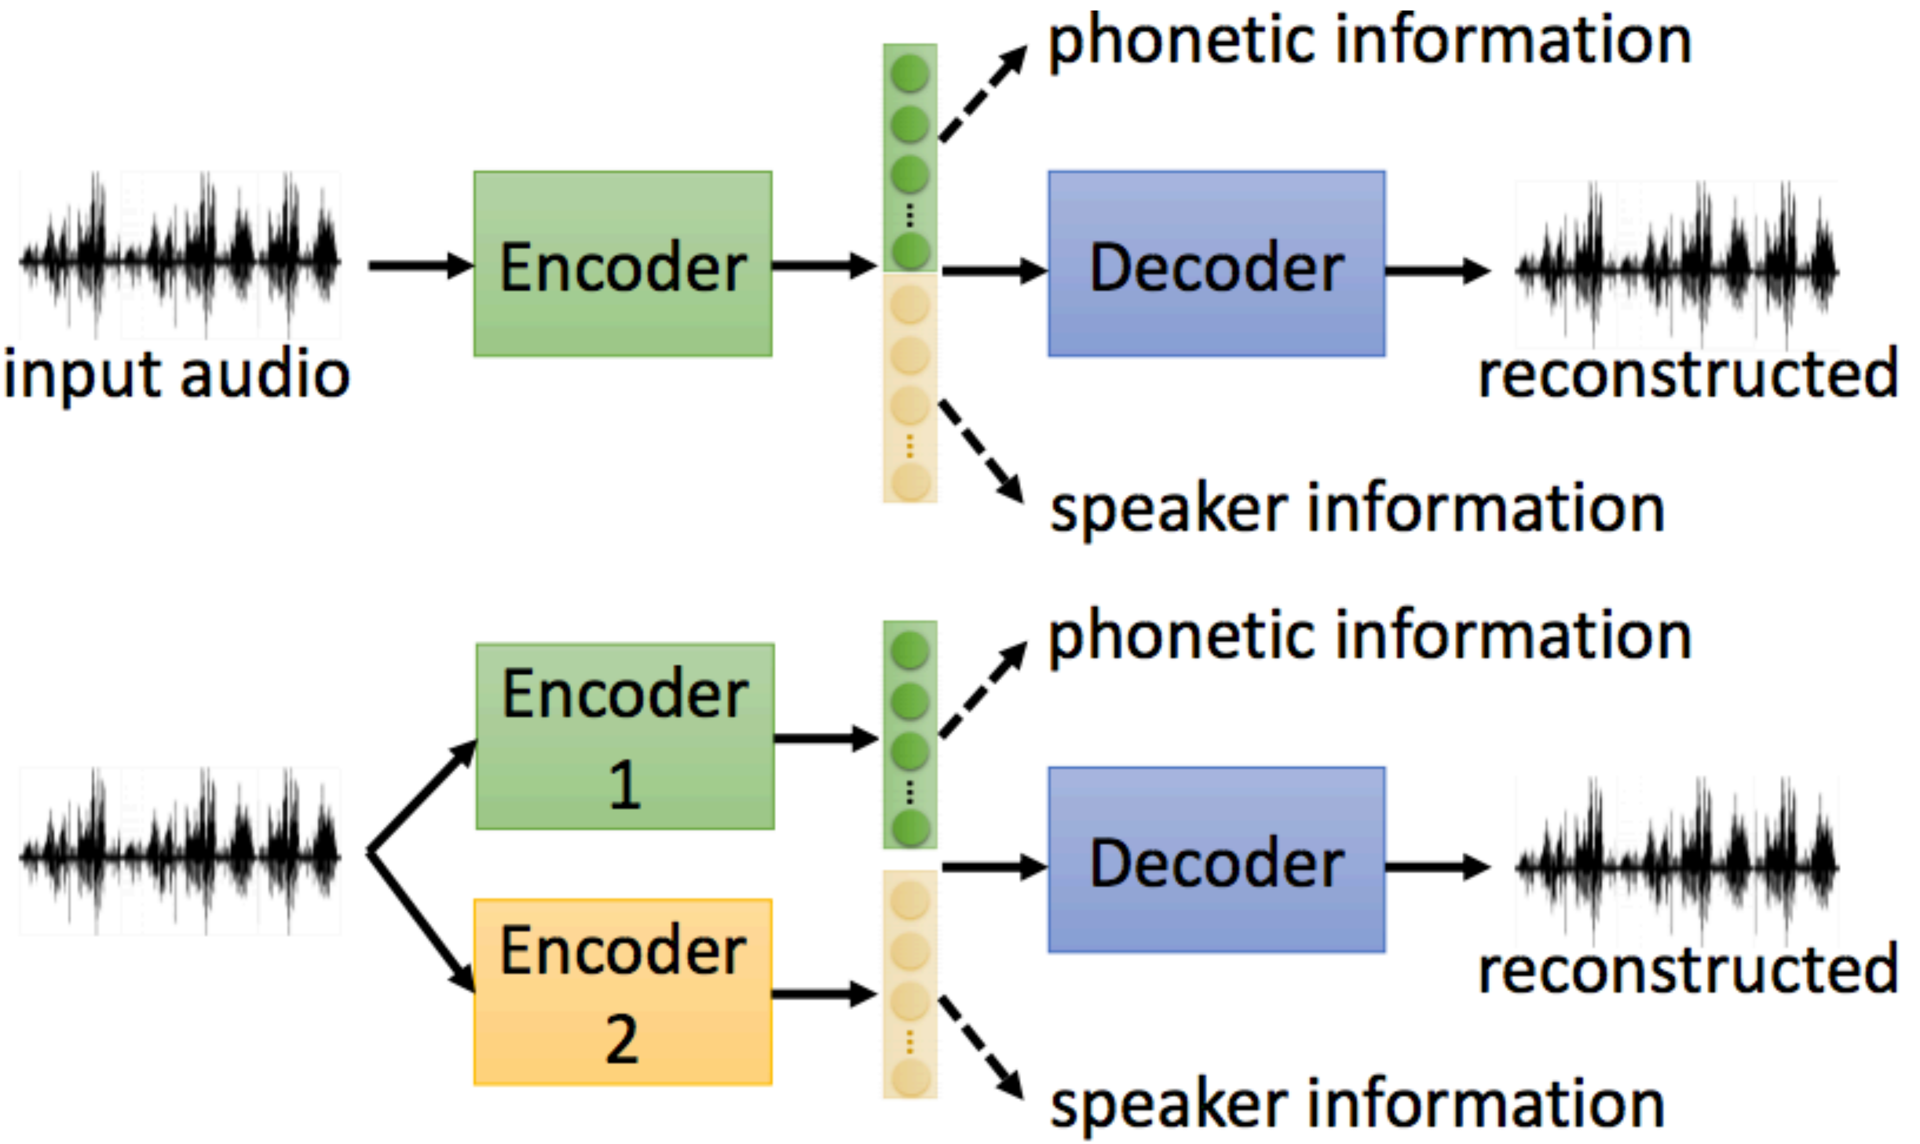

3. Feature Disentangle

- Voice conversion

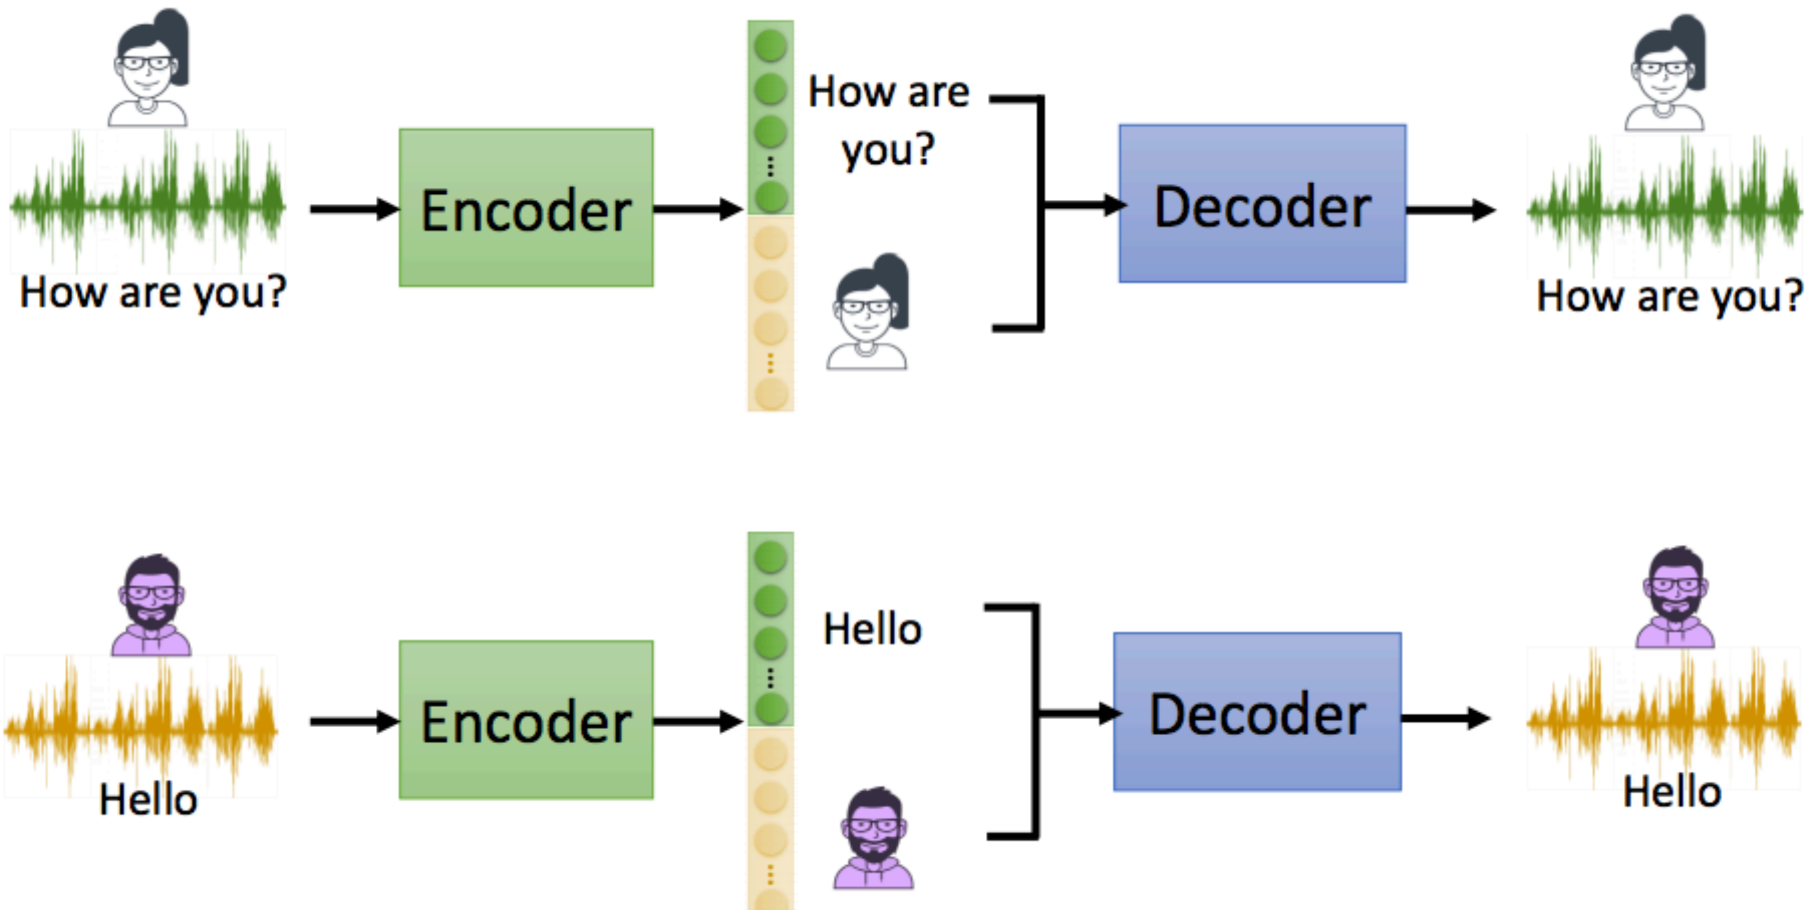

3. Feature Disentangle

- Voice conversion

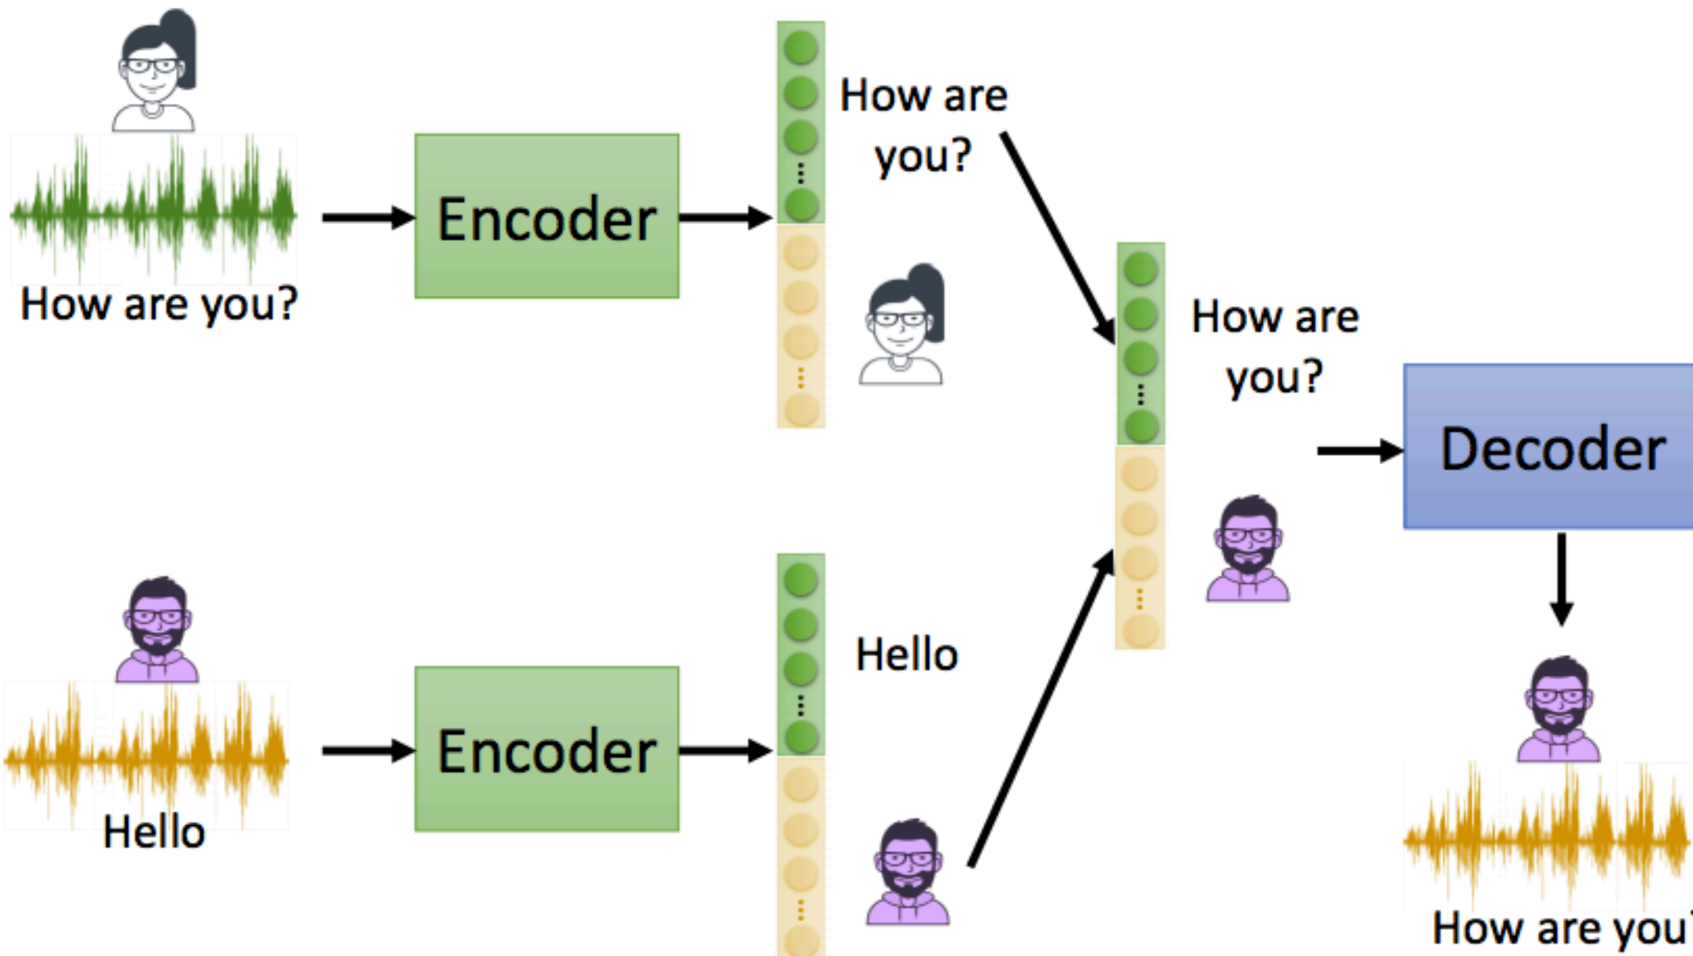

4. Discrete Representation

- Easier to interpret or clustering

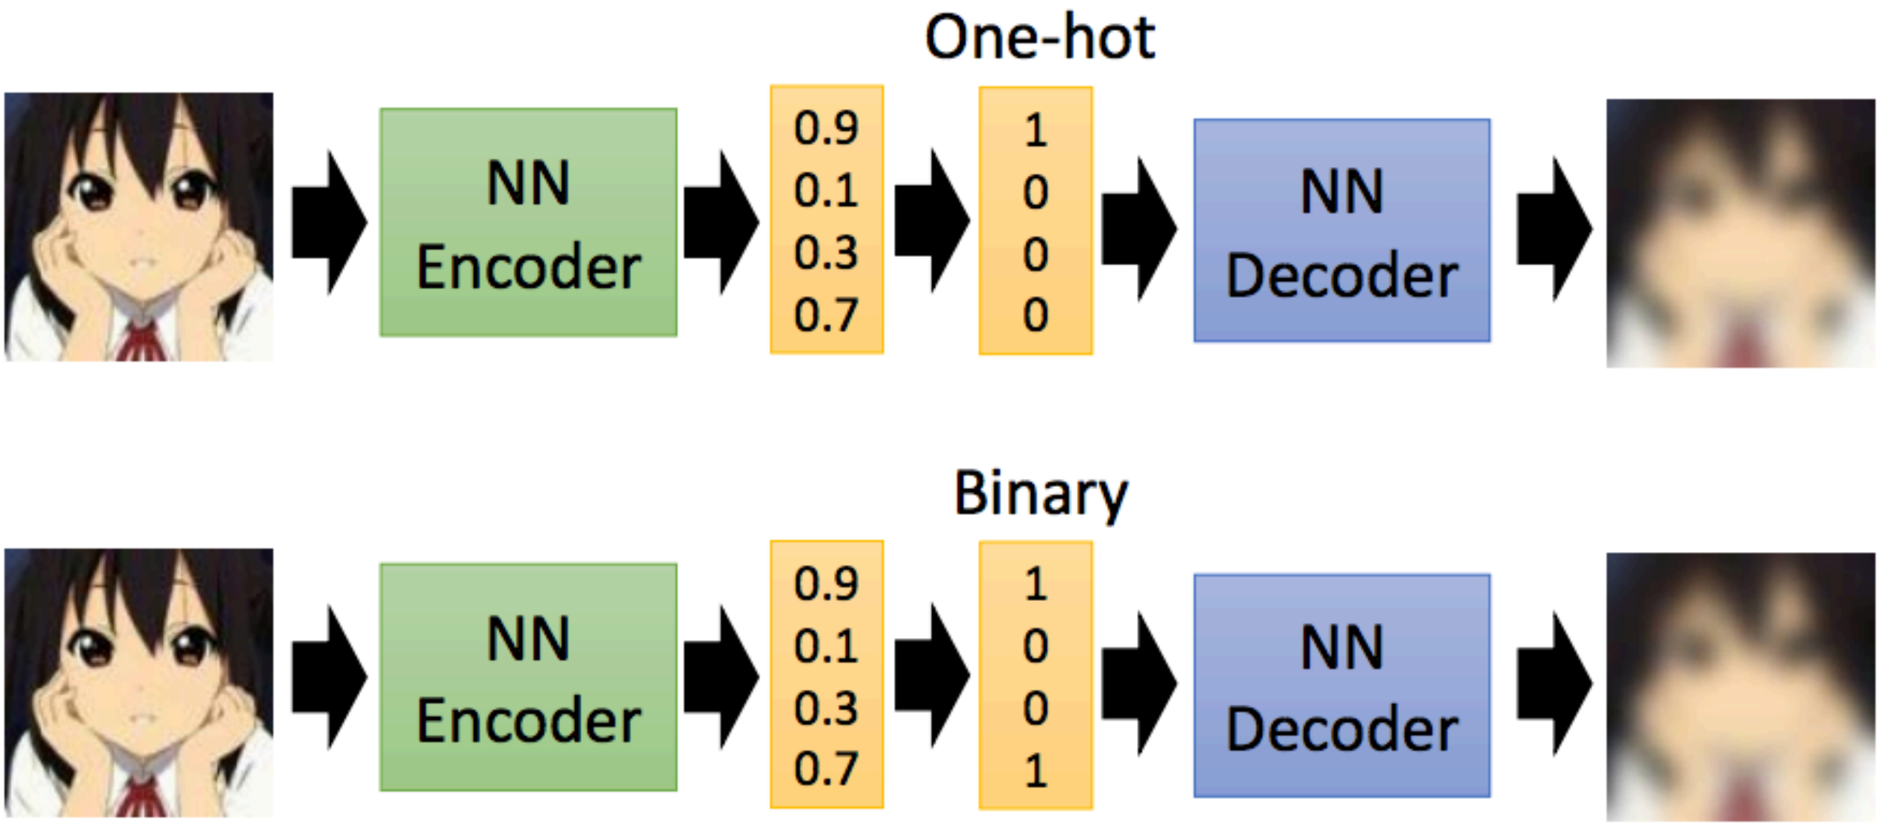

Concluding Remarks

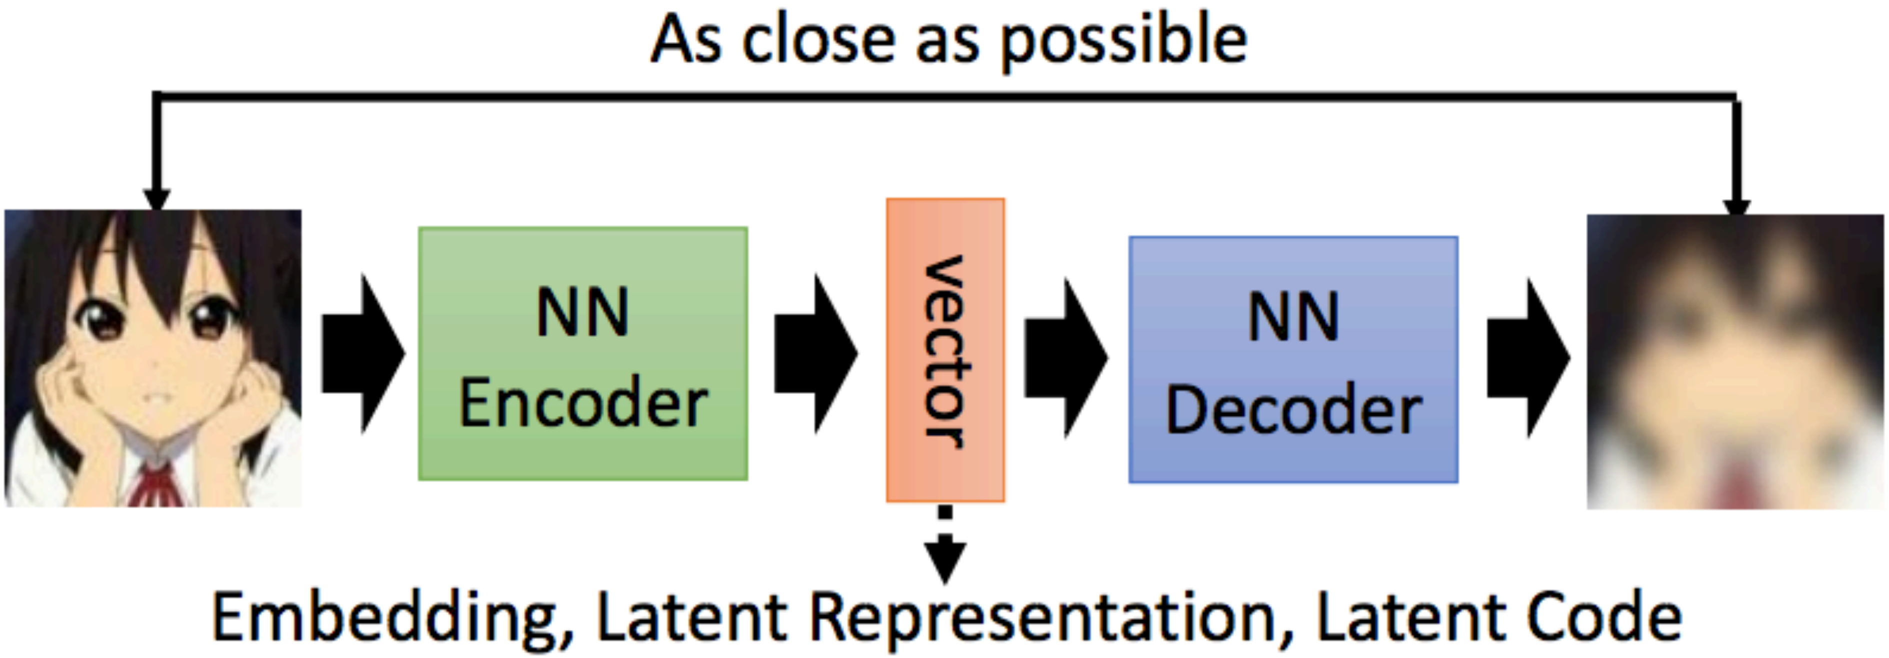

- More than minimizing reconstruction error
  - Using discriminator
  - Sequential data
- More interpretable embedding
  - Feature disentangle
  - Discrete representation