

Unsupervised Learning: Deep Auto-encoder

How Much Information is the Machine Given during Learning?

- ▶ **“Pure” Reinforcement Learning (cherry)**
 - ▶ The machine predicts a scalar reward given once in a while.
 - ▶ **A few bits for some samples**
- ▶ **Supervised Learning (icing)**
 - ▶ The machine predicts a category or a few numbers for each input
 - ▶ Predicting human-supplied data
 - ▶ **10→10,000 bits per sample**
- ▶ **Self-Supervised Learning (cake génoise)**
 - ▶ The machine predicts any part of its input for any observed part.
 - ▶ Predicts future frames in videos
 - ▶ **Millions of bits per sample**

Yann LeCun Cake Analogy




<https://medium.com/syncedreview/yann-lecun-cake-analogy-2-0-a361da560dae>

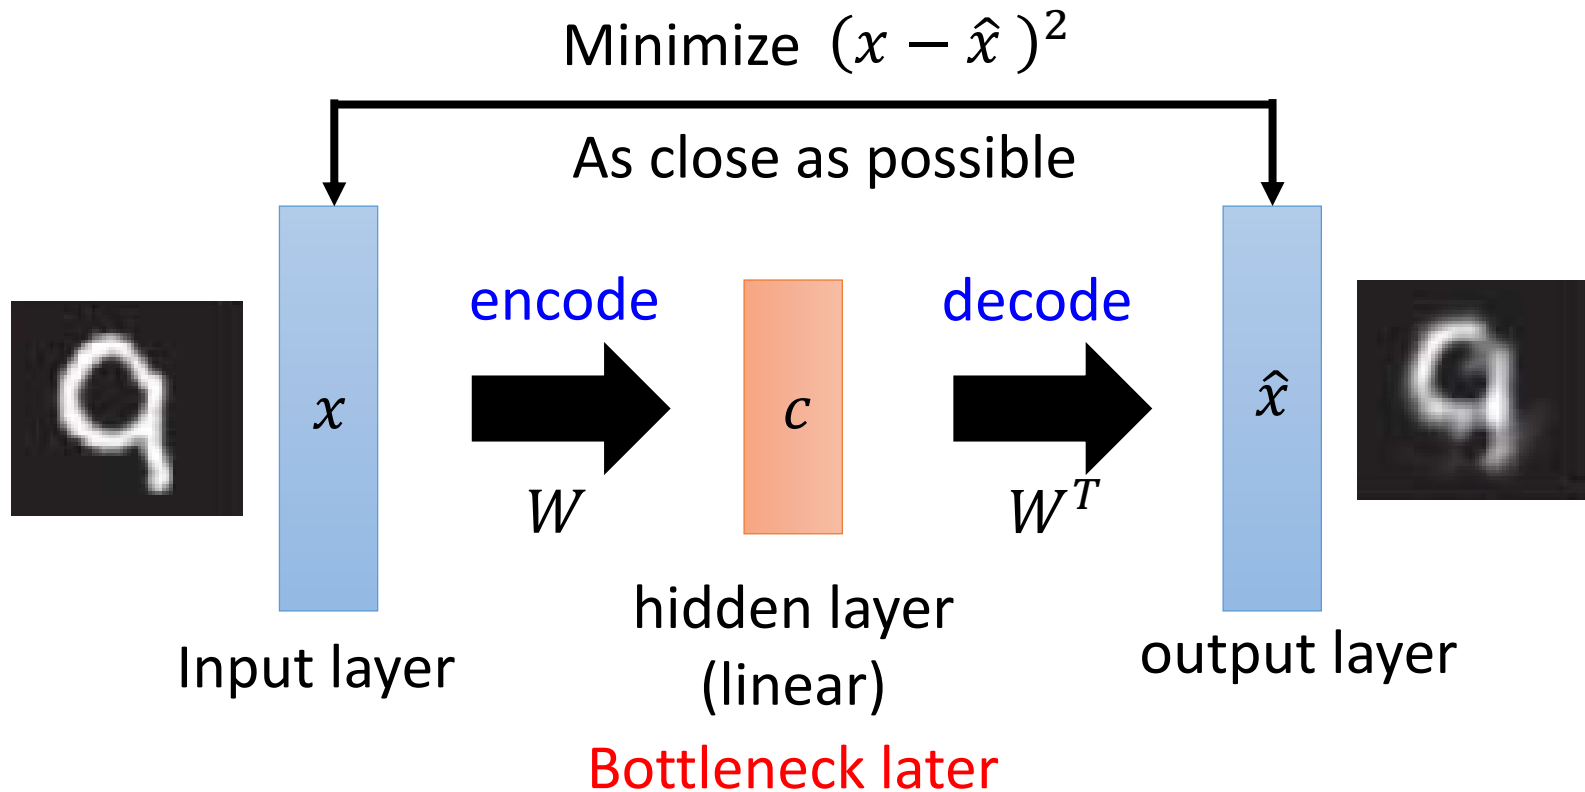
“We expect unsupervised learning to become far more important in the longer term. Human and animal learning is largely unsupervised: we discover the structure of the world by observing it, not by being told the name of every object.” – LeCun, Bengio, Hinton, Nature 2015

“If intelligence was a cake, unsupervised learning would be the cake, supervised learning would be the icing on the cake, and reinforcement learning would be the cherry on the cake. We know how to make the icing and the cherry, but we don't know how to make the cake.” - Yann LeCun, March 14, 2016 (Facebook)

Auto-encoder



Recap: PCA




Output of the hidden layer is the code

Deep Auto-encoder

Symmetric is not necessary.

- Of course, the auto-encoder can be deep



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

Deep Auto-encoder



Original Image




PCA



Deep Auto-encoder






More: Contractive auto-encoder

Auto-encoder

Ref: Rifai, Salah, et al. "Contractive auto-encoders: Explicit invariance during feature extraction." *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. 2011.


- 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-words

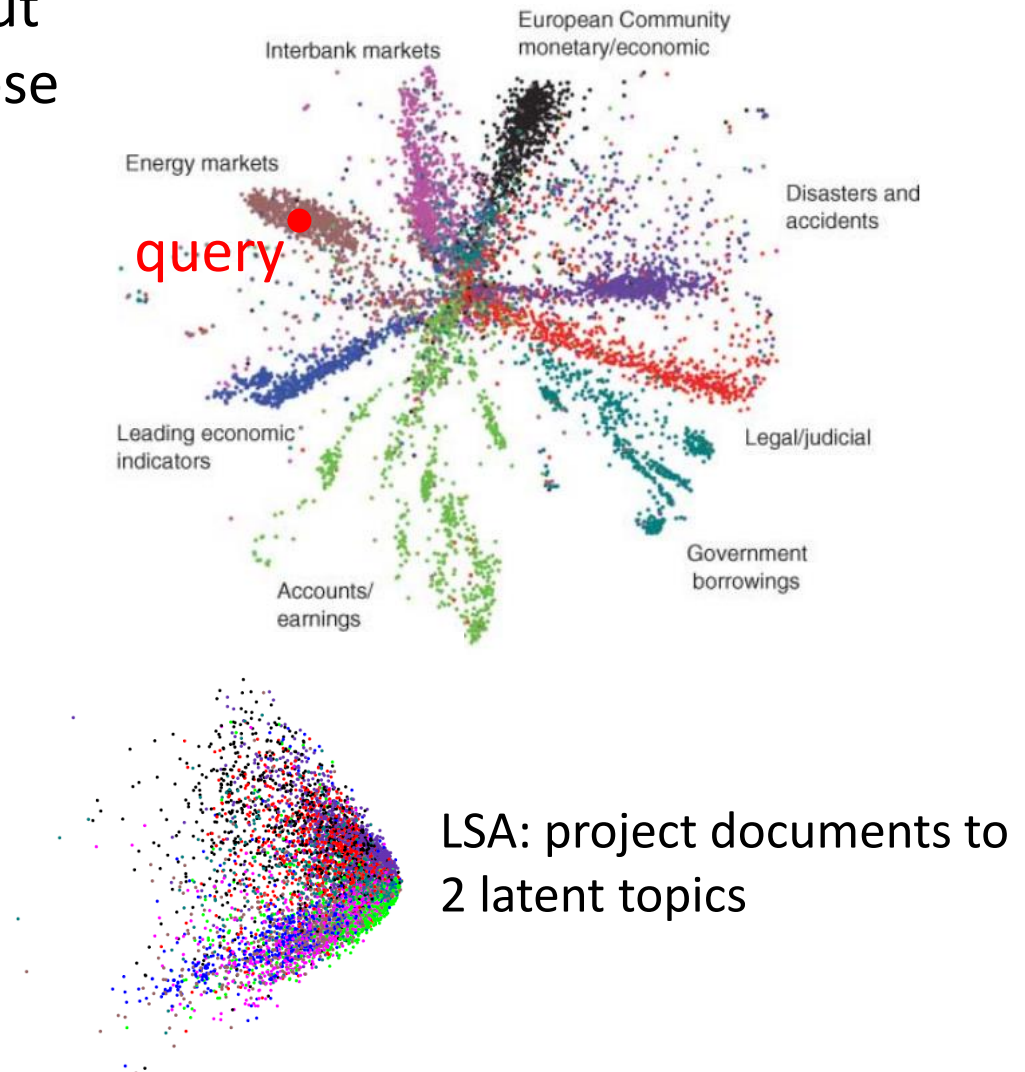

word string:
"This is an apple"

this	●	1
is	●	1
a	●	0
an	●	1
apple	●	1
pen	●	0
⋮	●	

Semantics are not considered.

Auto-encoder – Text Retrieval

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



Auto-encoder – Similar Image Search

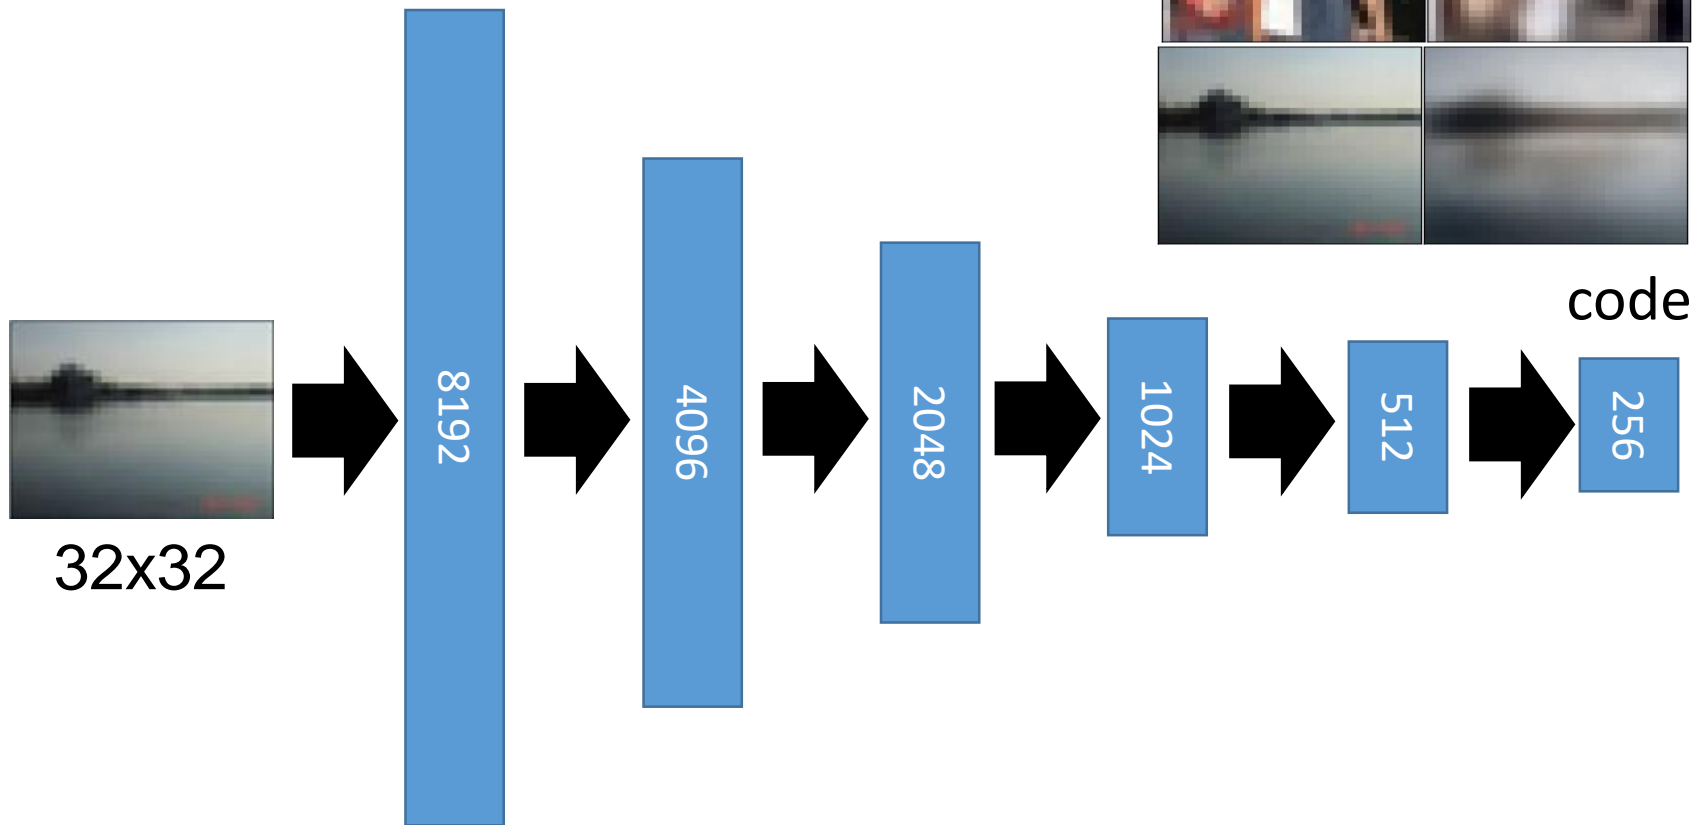
Retrieved using Euclidean distance in pixel intensity space



(Images from Hinton's slides on Coursera)


Reference: Krizhevsky, Alex, and Geoffrey E. Hinton. "Using very deep autoencoders for content-based image retrieval." *ESANN*. 2011.

Auto-encoder – Similar Image Search



(crawl millions of images from the Internet)

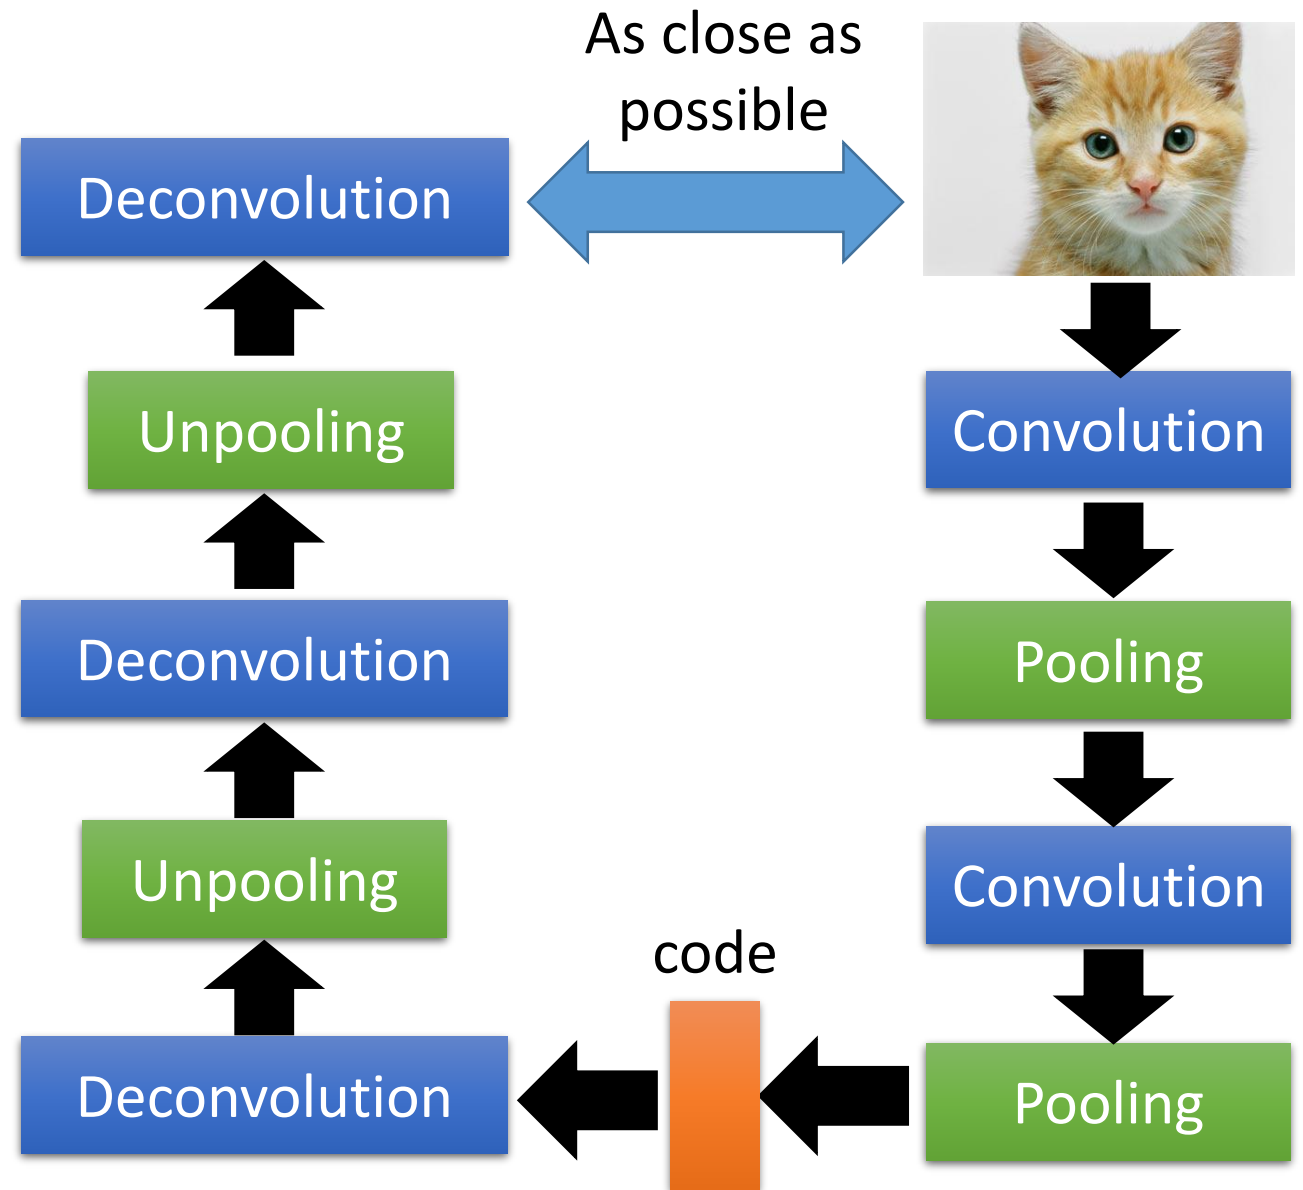
Retrieved using Euclidean distance in pixel intensity space




retrieved using 256 codes




Auto-encoder for CNN




CNN -Unpooling



14 x 14



28 x 28





Alternative: simply
repeat the values

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


Actually, deconvolution is convolution.

CNN

- Deconvolution

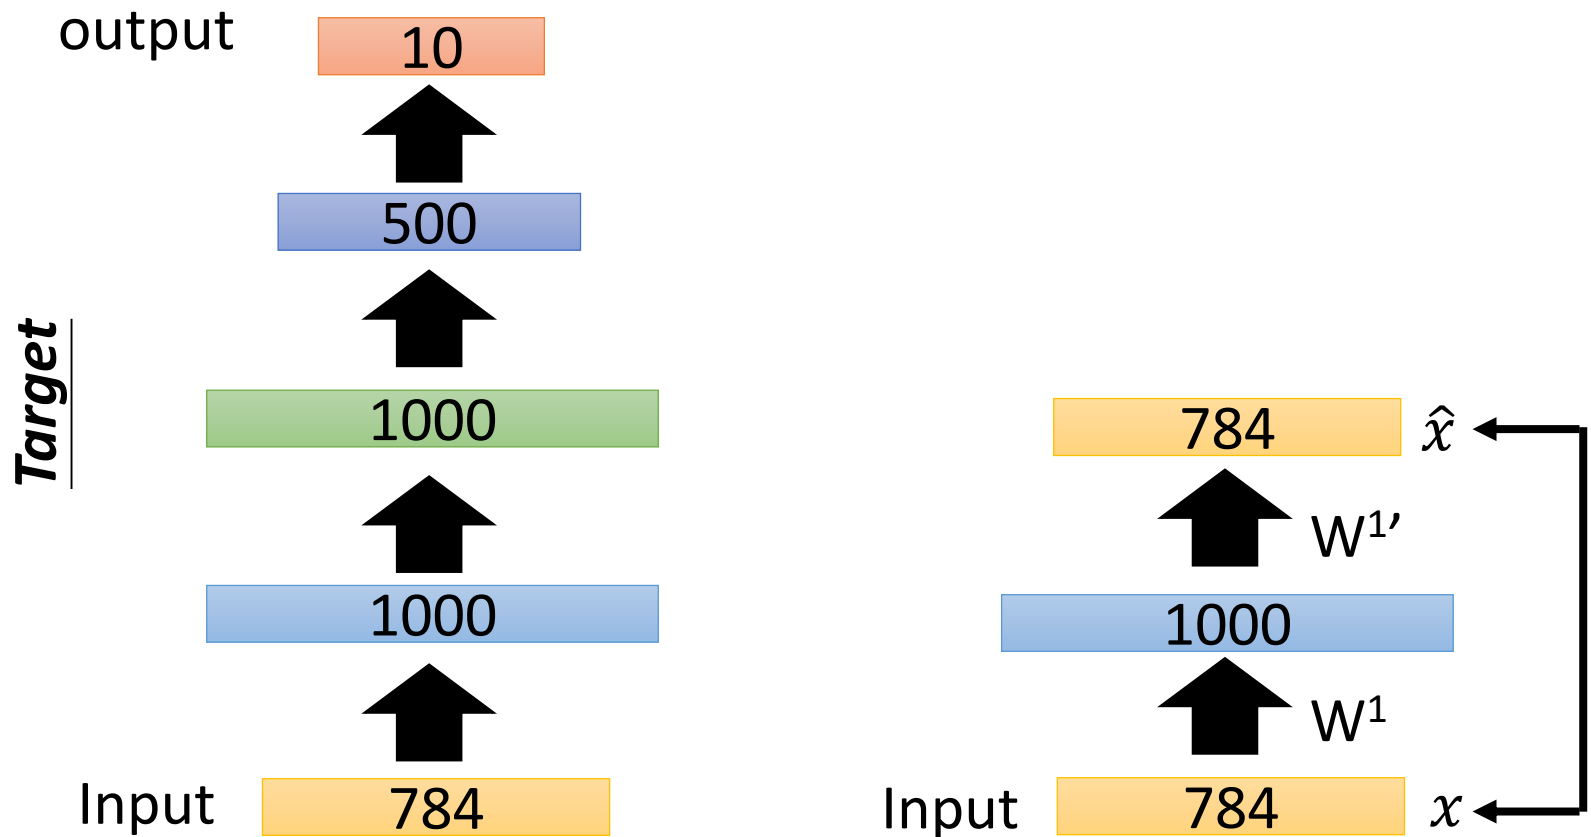


=



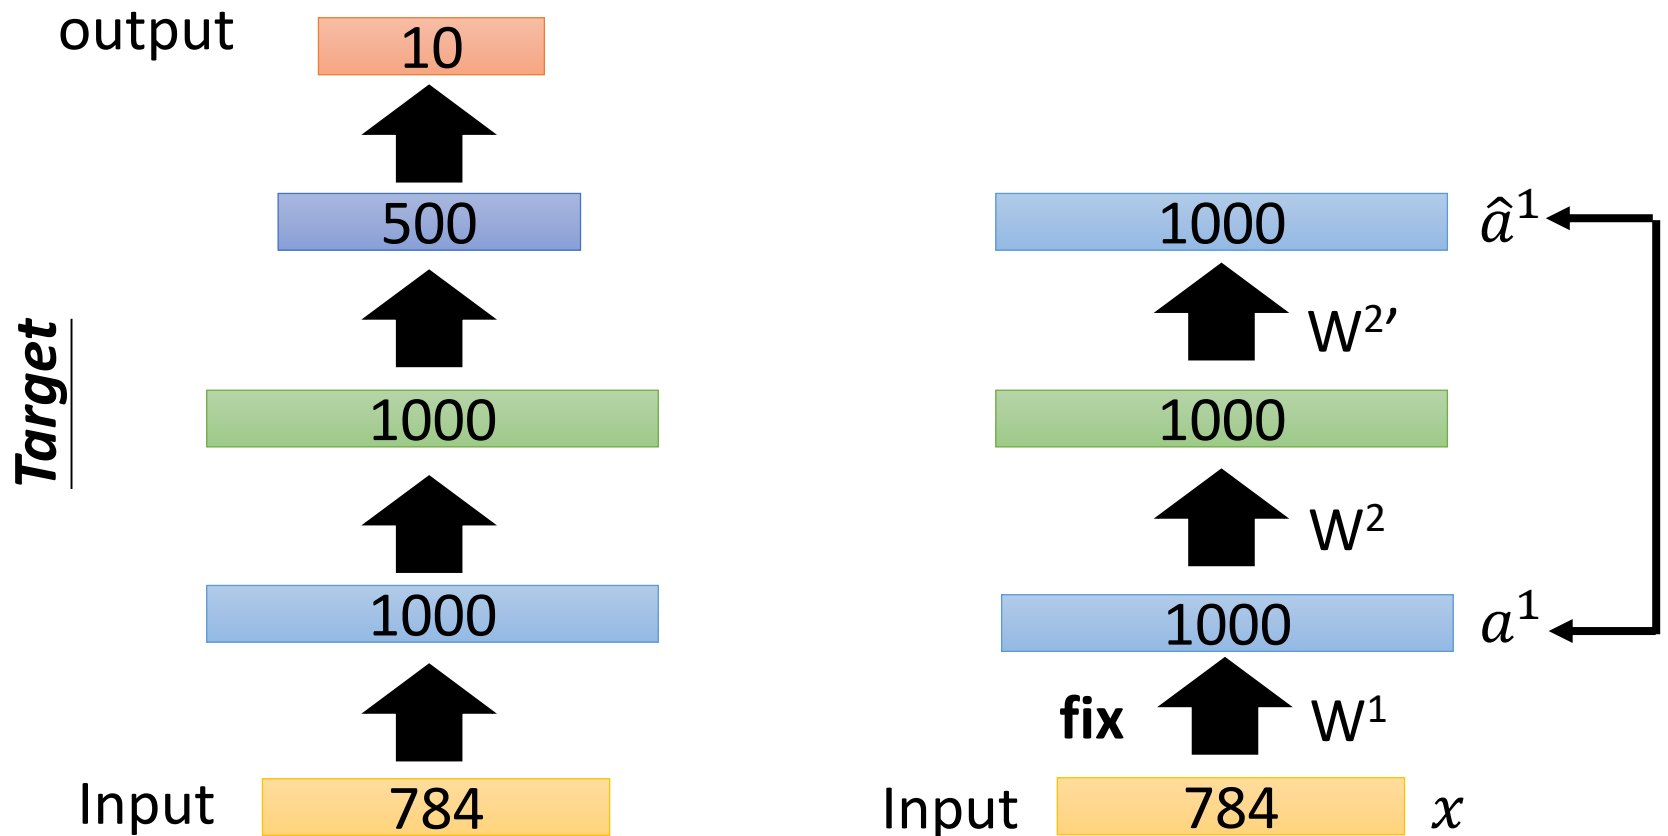
Auto-encoder – Pre-training DNN

- Greedy Layer-wise Pre-training *again*



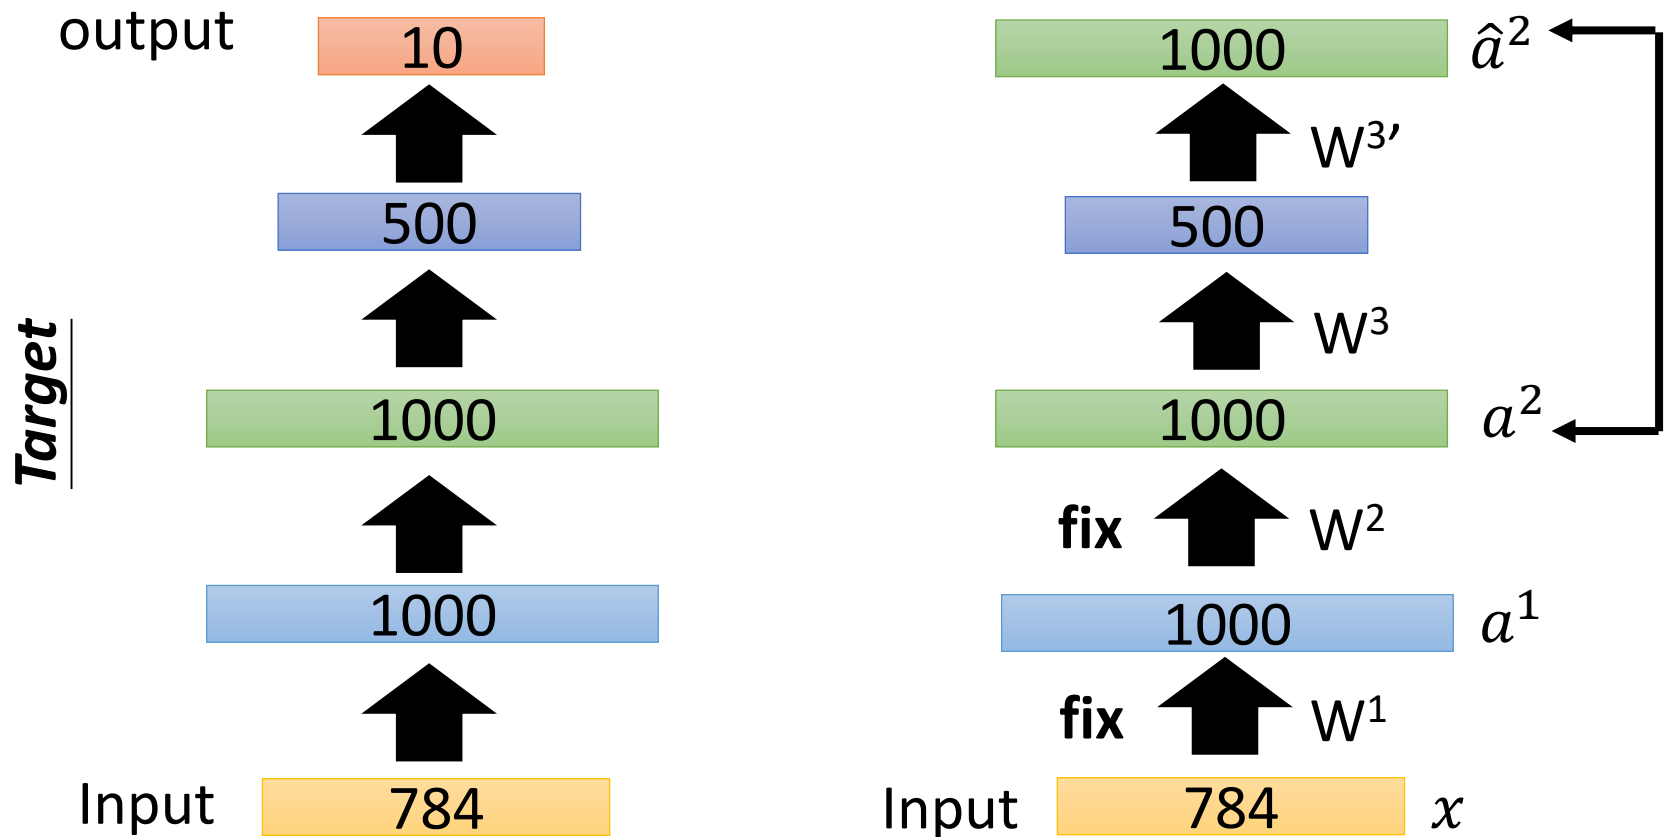
Auto-encoder – Pre-training DNN

- Greedy Layer-wise Pre-training *again*



Auto-encoder – Pre-training DNN

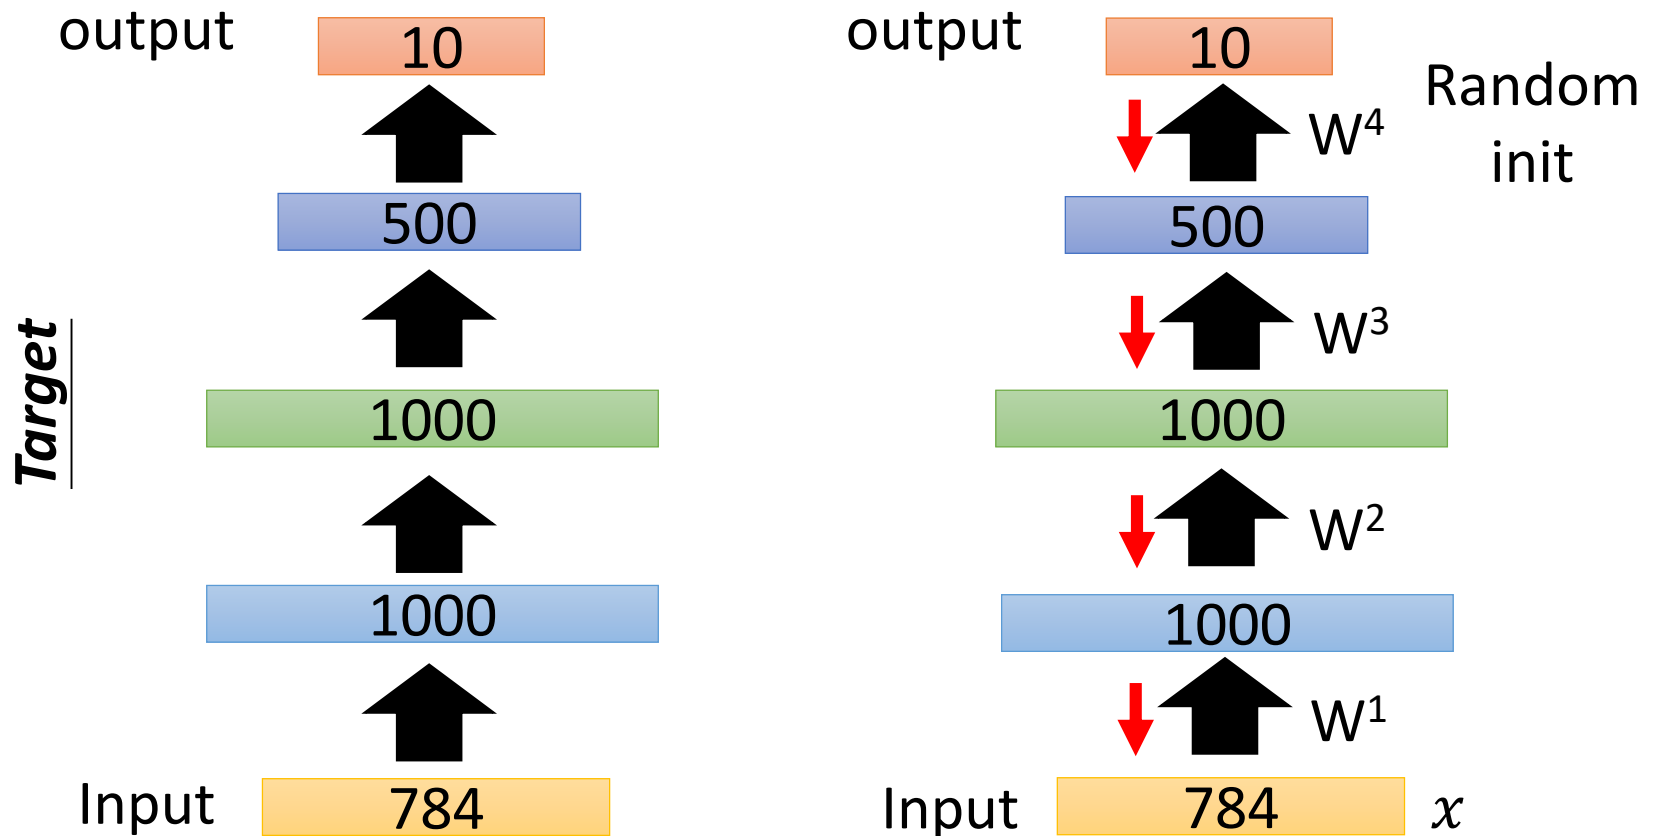
- Greedy Layer-wise Pre-training *again*



Auto-encoder – Pre-training DNN

- Greedy Layer-wise Pre-training *again*

Find-tune by
backpropagation



Learning More

- Restricted Boltzmann Machine


- Neural networks [5.1] : Restricted Boltzmann machine – definition
 - https://www.youtube.com/watch?v=p4Vh_zMw-HQ&index=36&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH
- Neural networks [5.2] : Restricted Boltzmann machine – inference
 - https://www.youtube.com/watch?v=lekCh_i32iE&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH&index=37
- Neural networks [5.3] : Restricted Boltzmann machine - free energy
 - https://www.youtube.com/watch?v=e0Ts_7Y6hZU&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH&index=38

Learning More



- Deep Belief Network

- Neural networks [7.7] : Deep learning - deep belief network
 - <https://www.youtube.com/watch?v=vkb6AWYXZ5I&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH&index=57>
- Neural networks [7.8] : Deep learning - variational bound
 - <https://www.youtube.com/watch?v=pStDscJh2Wo&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH&index=58>
- Neural networks [7.9] : Deep learning - DBN pre-training
 - <https://www.youtube.com/watch?v=35MUIYCColk&list=PL6Xpj9I5qXYEcOhn7TqghAJ6NAPrNmUBH&index=59>


Next




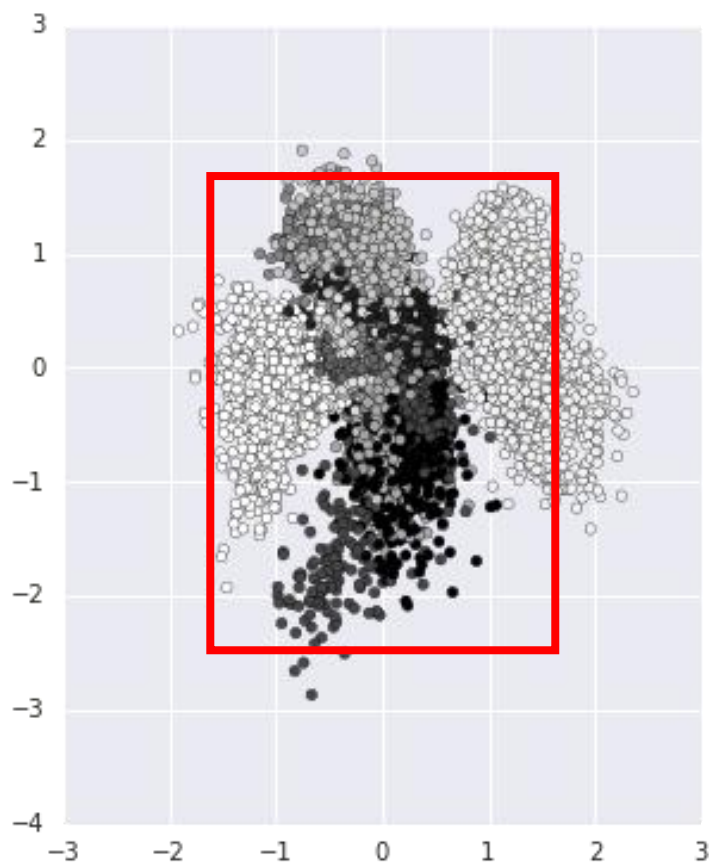
- Can we use decoder to generate something?



Next



- Can we use decoder to generate something?



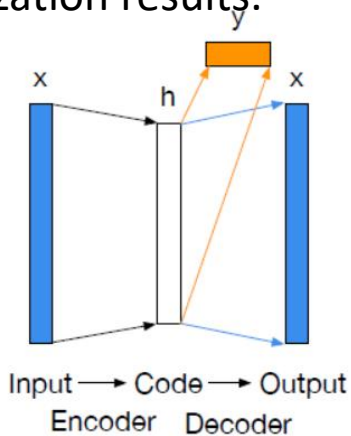
Appendix

Supervised AutoEncoder

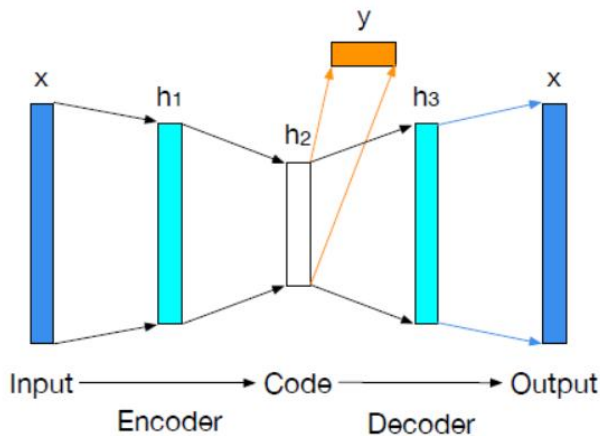
Goal: Improving generalization performance (i.e., reducing overfitting) in supervised learning.

Idea: Auto-encoder design. Jointly predicts target and input (reconstruction) from the code.


Reason: By adding the reconstruction MSE as regularization term, uniform stability can be proved (under mild conditions) assuming linear AE, which in general implies better generalization results.




(a) (Linear) Supervised Autoencoder




(b) Deep Supervised Autoencoder



Reference: (a) Sigmoid activation



(b) ReLU activation



(c) Kernel Representation

Lei Le, Andrew Patterson, Martha White, "Supervised autoencoders: Improving generalization performance with unsupervised regularizers." NeurIPS 2018: 107-117