

Representation Learning

Neil Gong


Representation/Feature/Embedding learning

- Key of deep learning
- Via unsupervised or self-supervised learning

Image

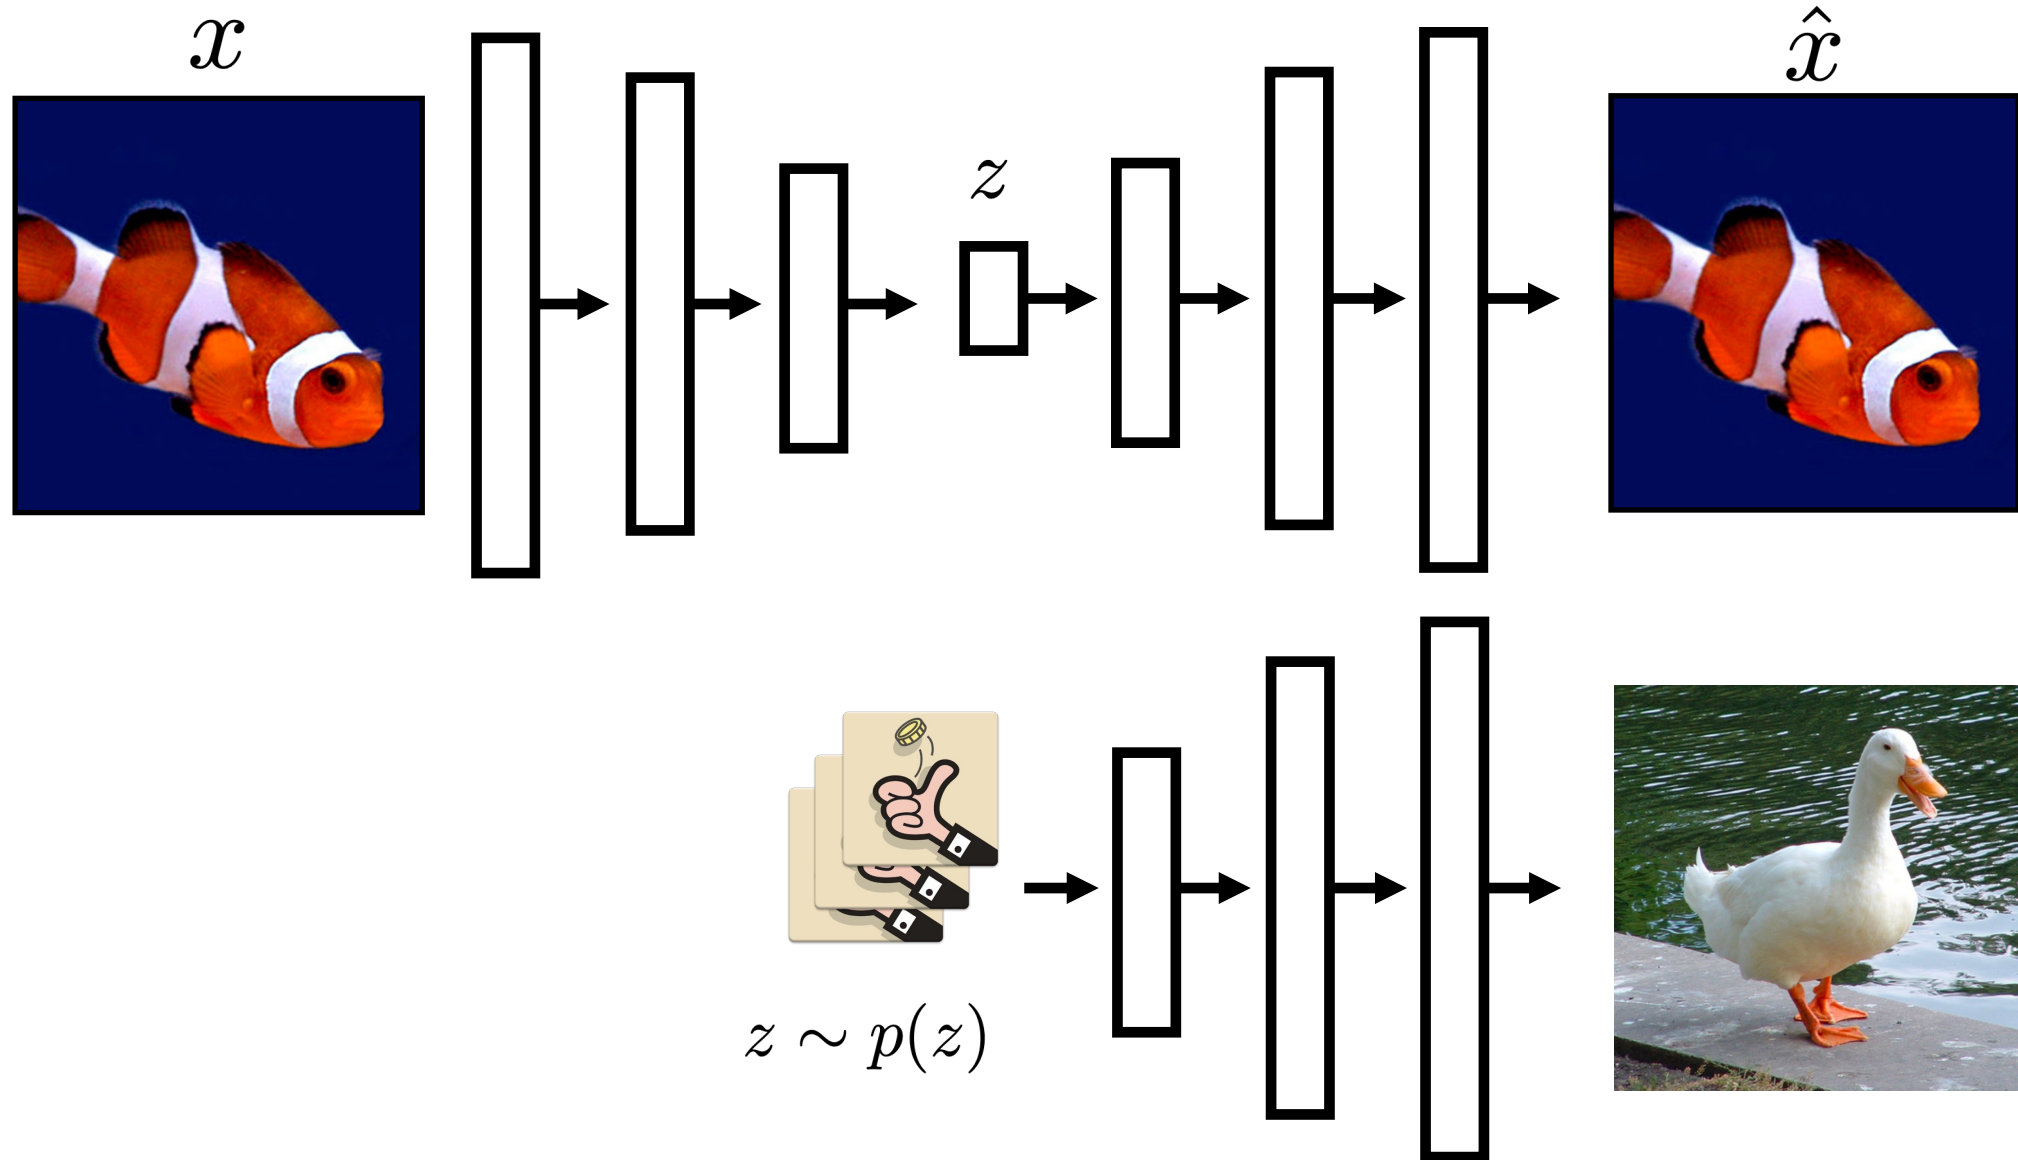
- Autoencoder
- Variational autoencoder
- Contrastive learning

Autoencoder

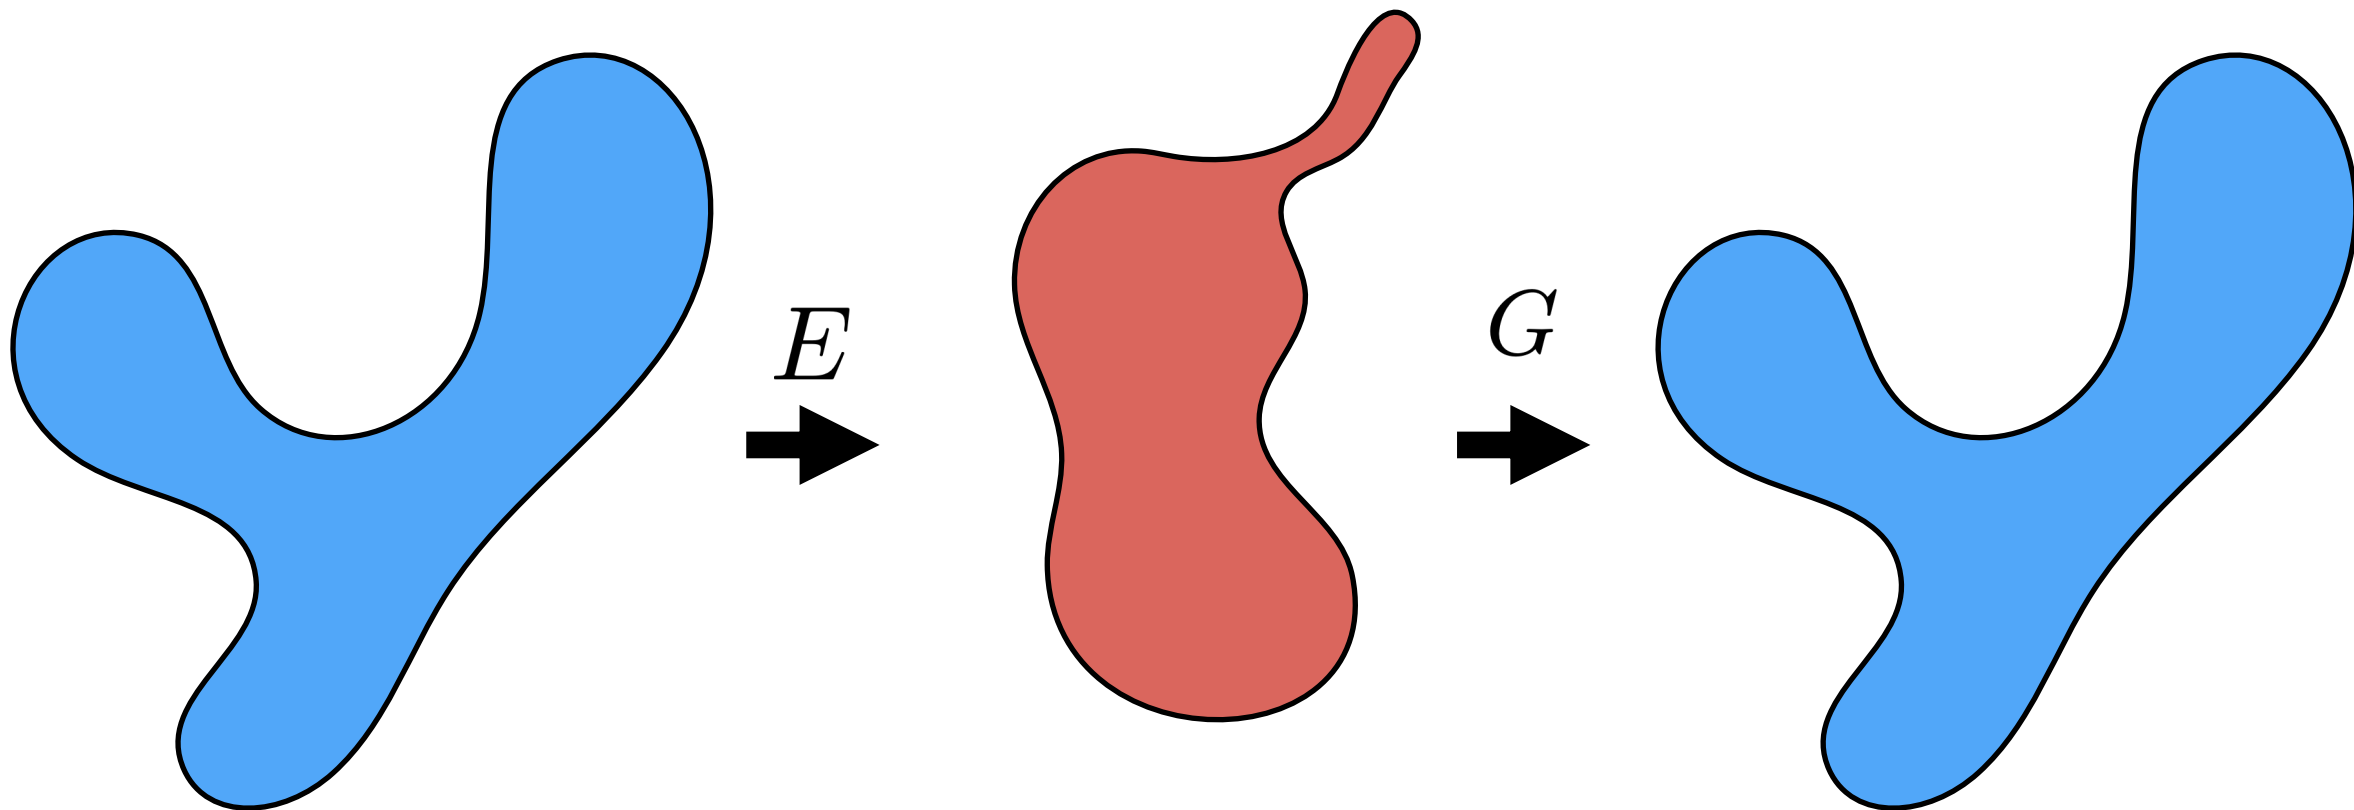


Slides adapted from Sara Beery, Kaiming He, Vincent Sitzmann, Mina Konaković Luković


Autoencoder \rightarrow Variational Autoencoder



Classical Autoencoder





Variational Autoencoder



Data Augmentation

Augmented views


Image



Crop and resize



Horizontal flip



Rotation

Contrastive Learning— SimCLR [ICML'20]

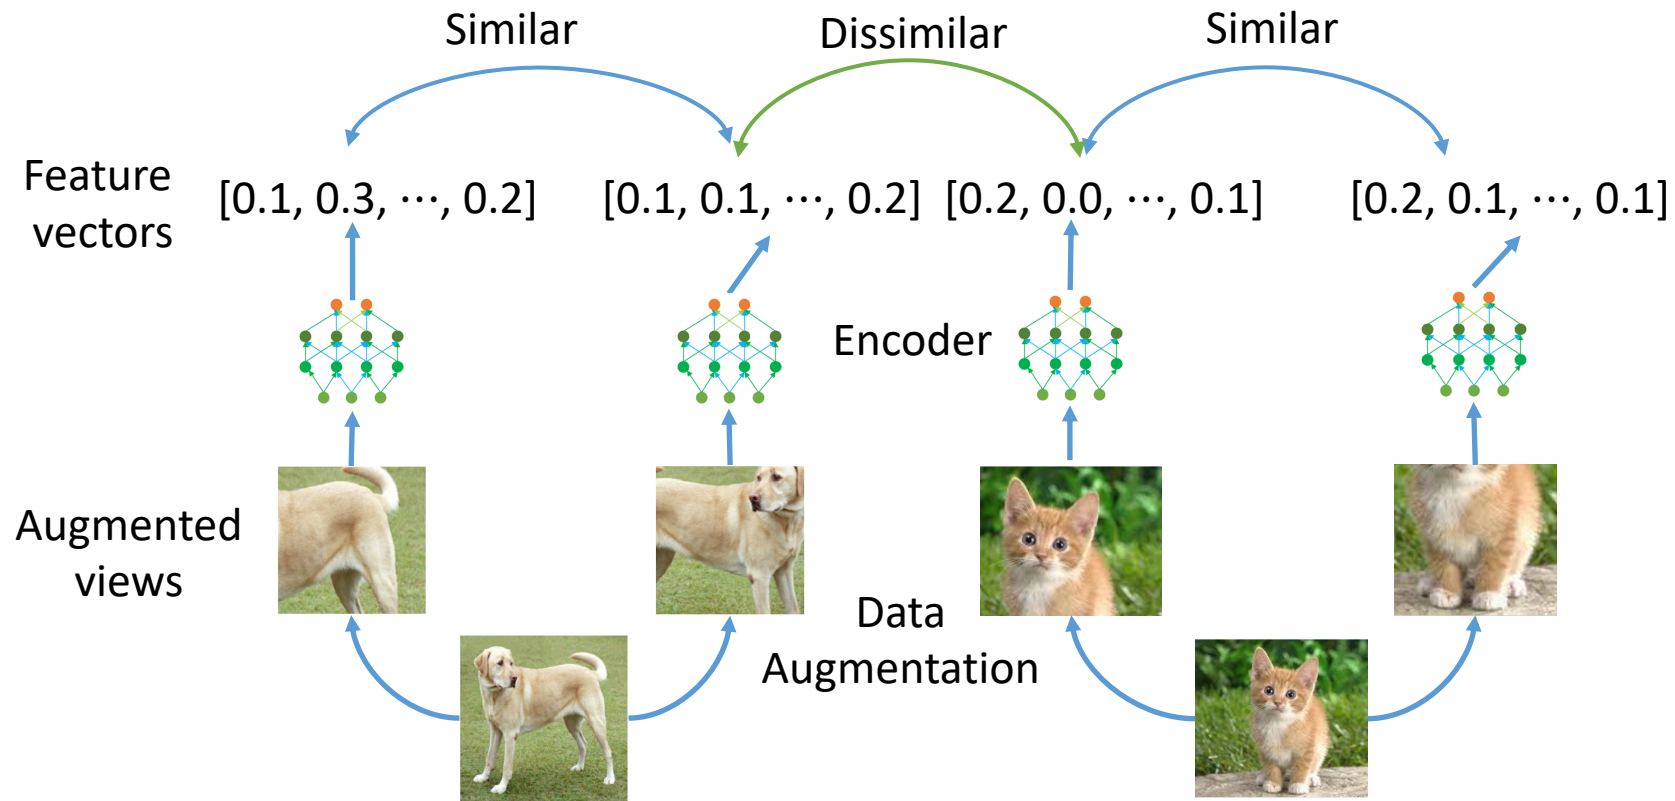
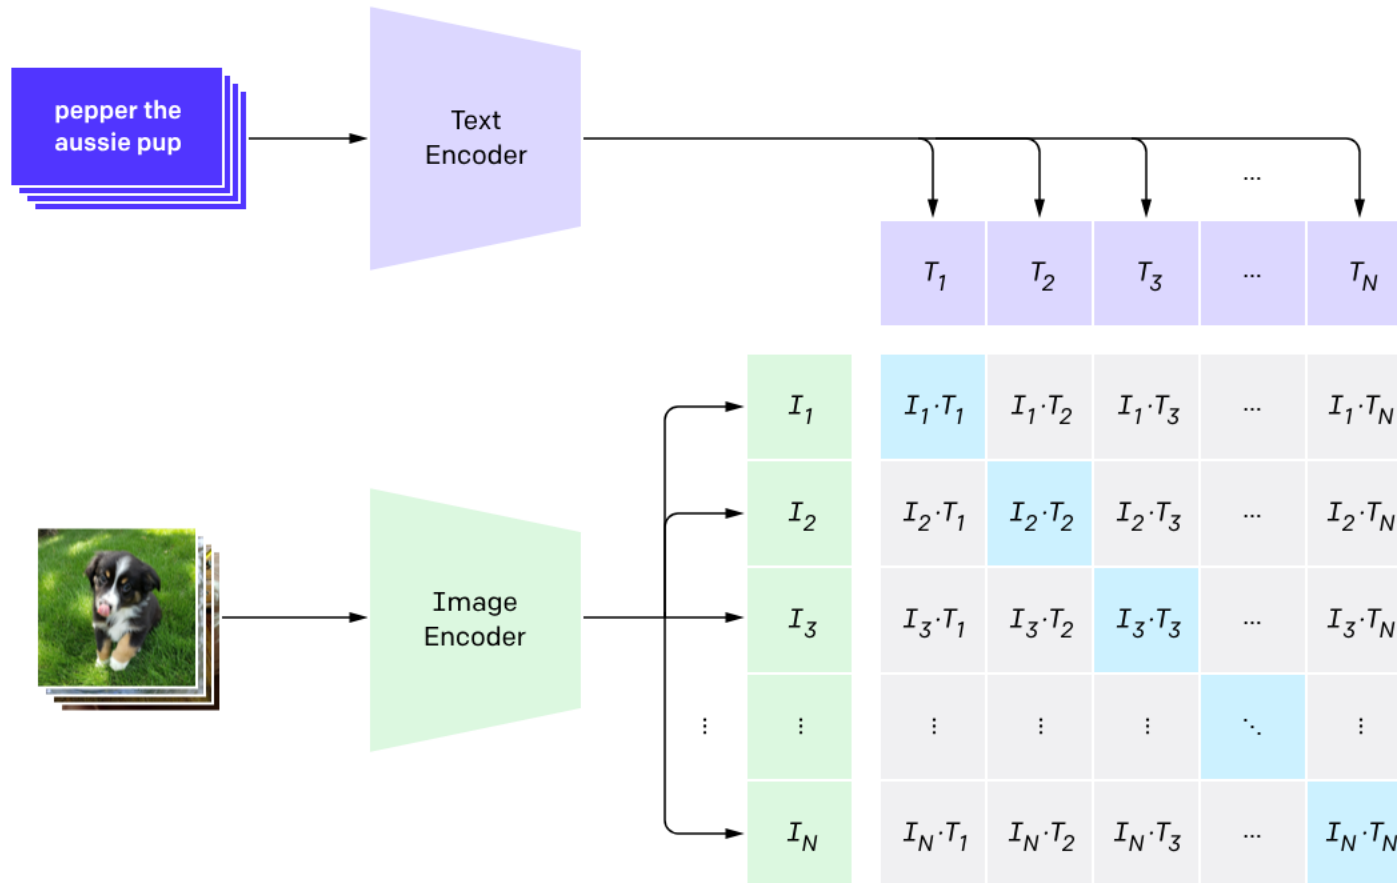


Image-text pairs (CLIP)



Text

- Next-word prediction