

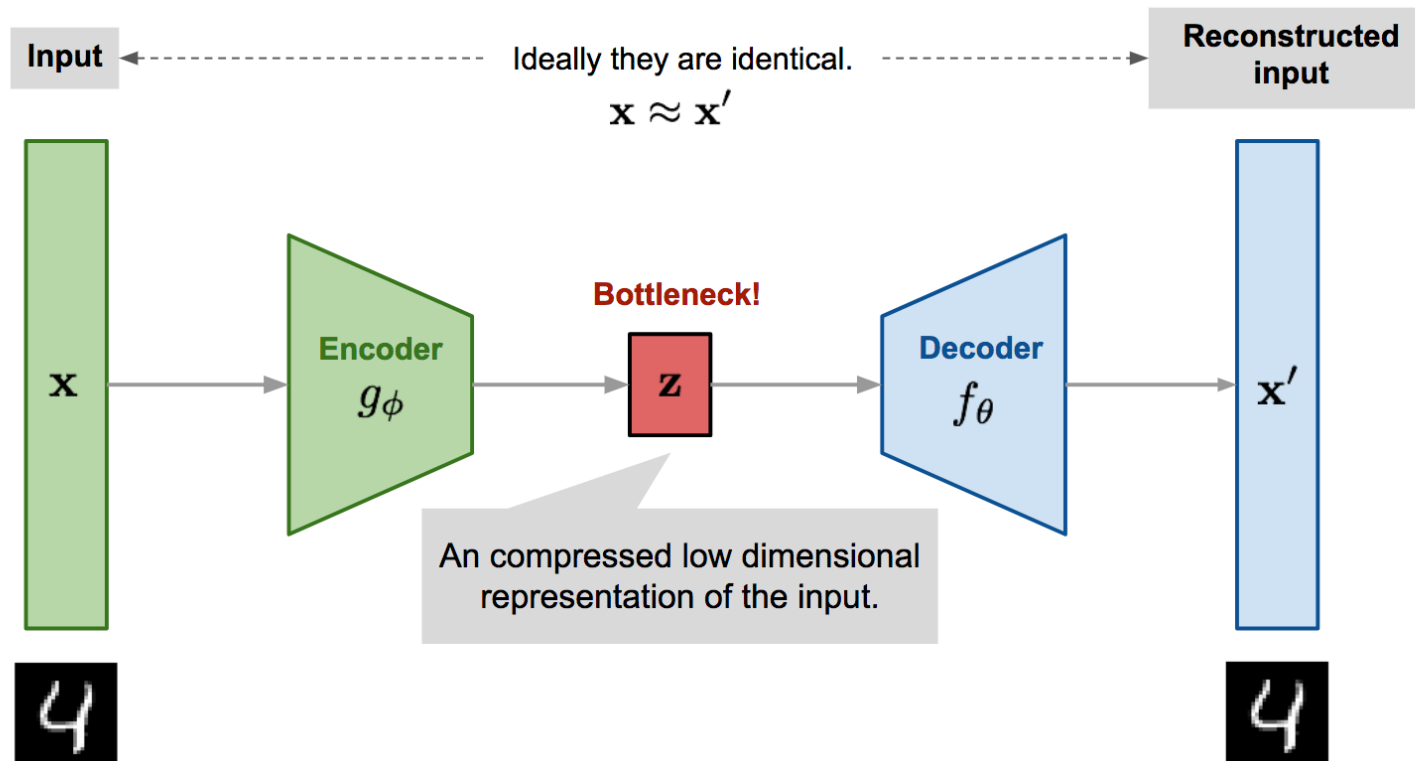
What's in a good representation?

John Tan Chong Min

Aim

- Find a way to encode a suitable representation to perform decision making
- Such representation can also be how we store memories for use in the future

Autoencoders: Representation via Reconstruction



Do you need to predict everything?



Transformers: Representation via Prediction

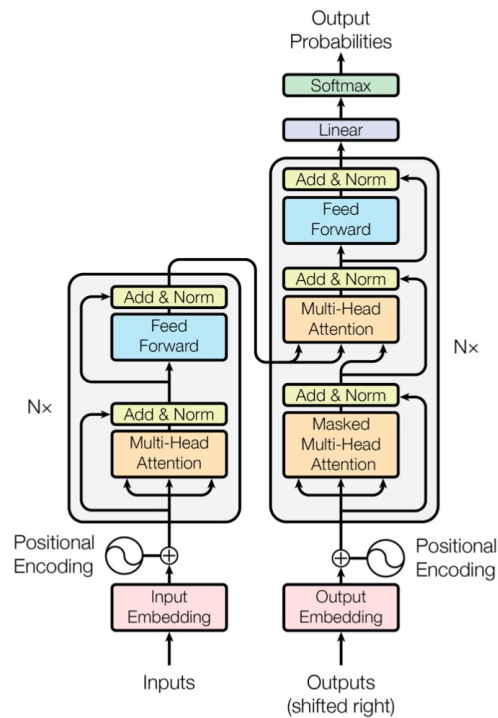



Figure 1: The Transformer - model architecture.

Scaled Dot-Product Attention



Multi-Head Attention




Figure 2: (left) Scaled Dot-Product Attention. (right) Multi-Head Attention consists of several attention layers running in parallel.


$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Taken from: Attention is all you need. Vaswani et al. (2017)


Large Self-Supervised Learning

- Self-supervised learning helps to learn better manifolds across large data
- Can work zero-shot on a new sample

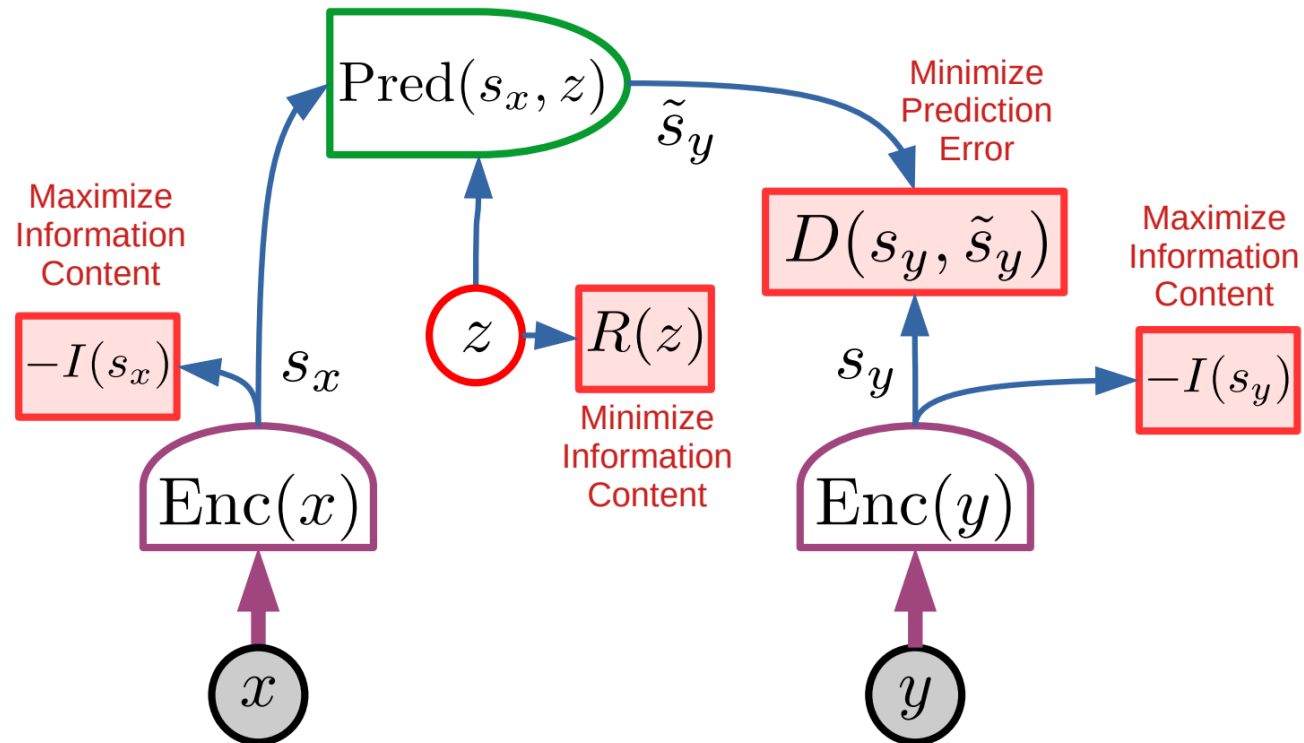
Manifold of having only 2 classes





Manifold of having 4 classes




JEPA - Only use whatever is necessary to predict



Natural Fixed Biases

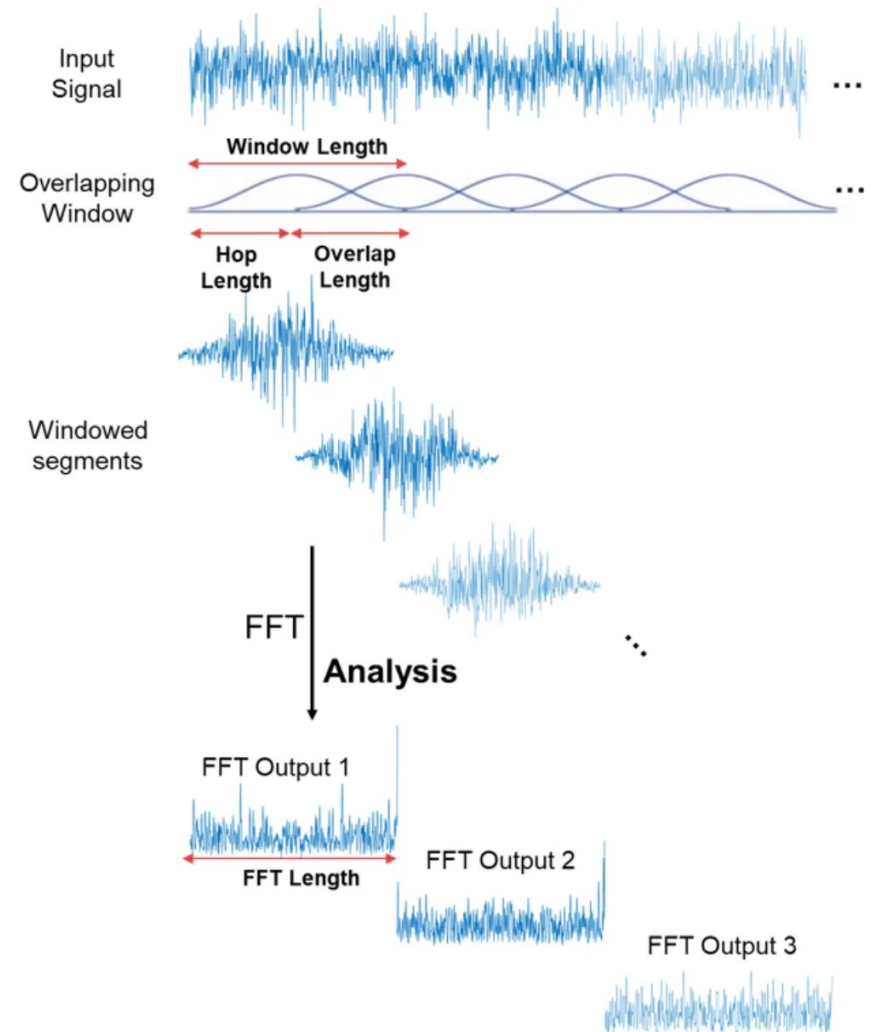


Audio – Freq Modelling



Audio signal converted into frequencies

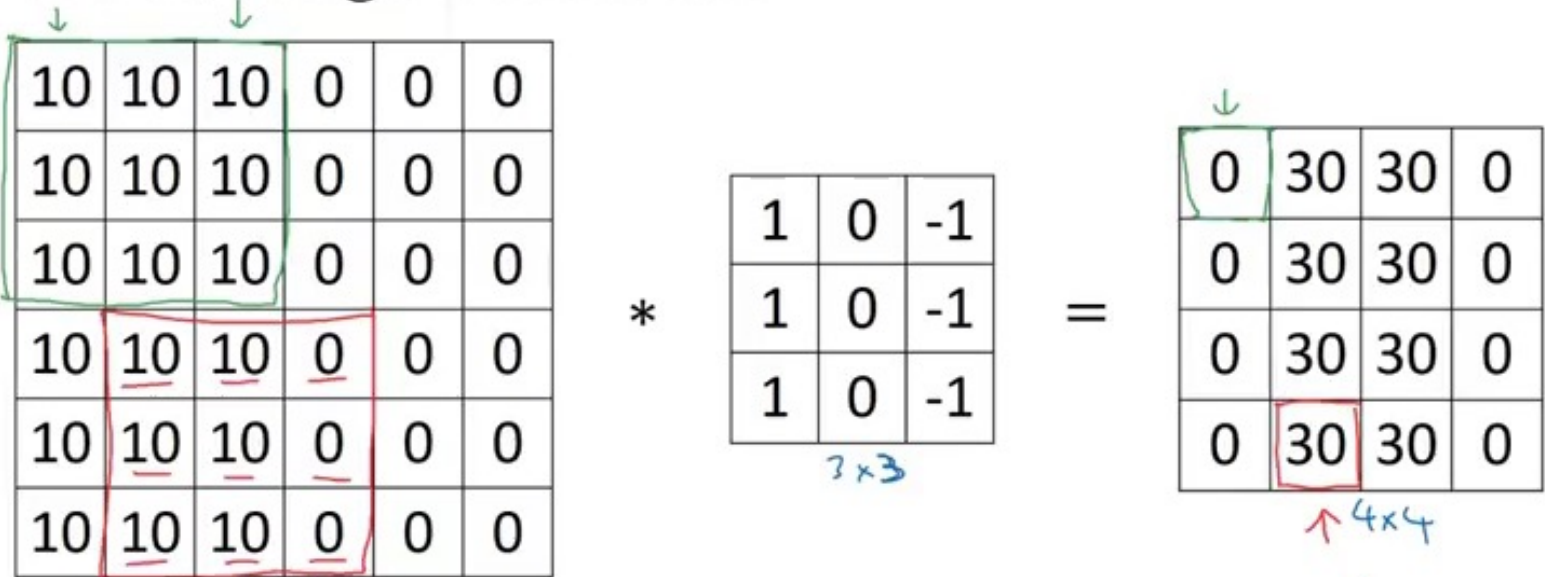
Images taken from: <https://medium.com/analytics-vidhya/understanding-the-mel-spectrogram-fca2afa2ce53>



Use overlapping window to model waveform over time

Vision – Pixel Proximity


Vertical edge detection





6x6

3x3

4x4

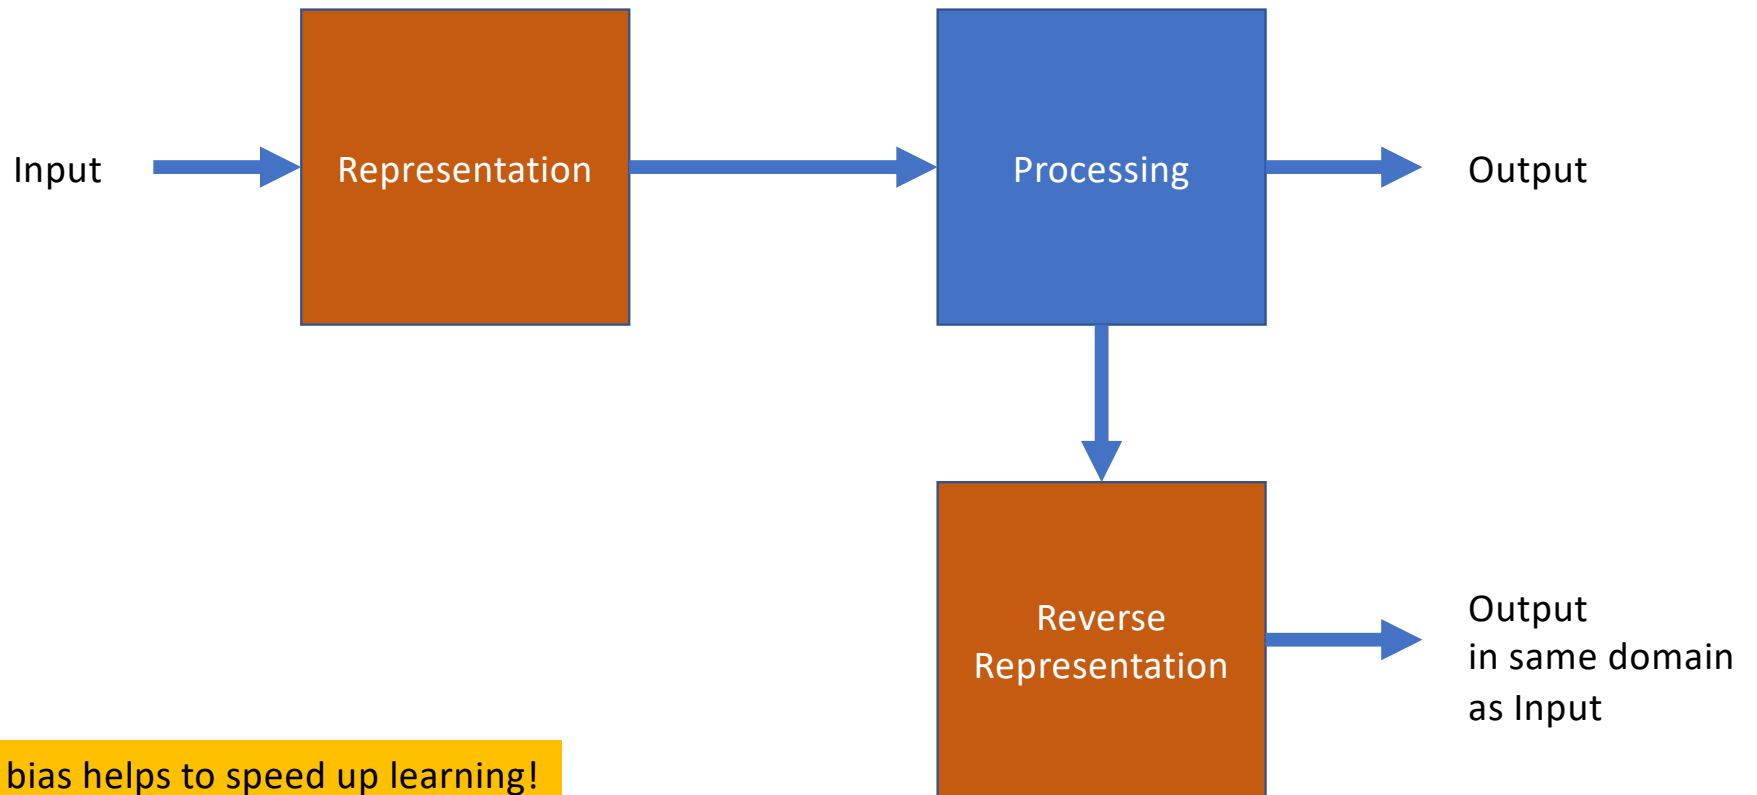


*



Andrew Ng

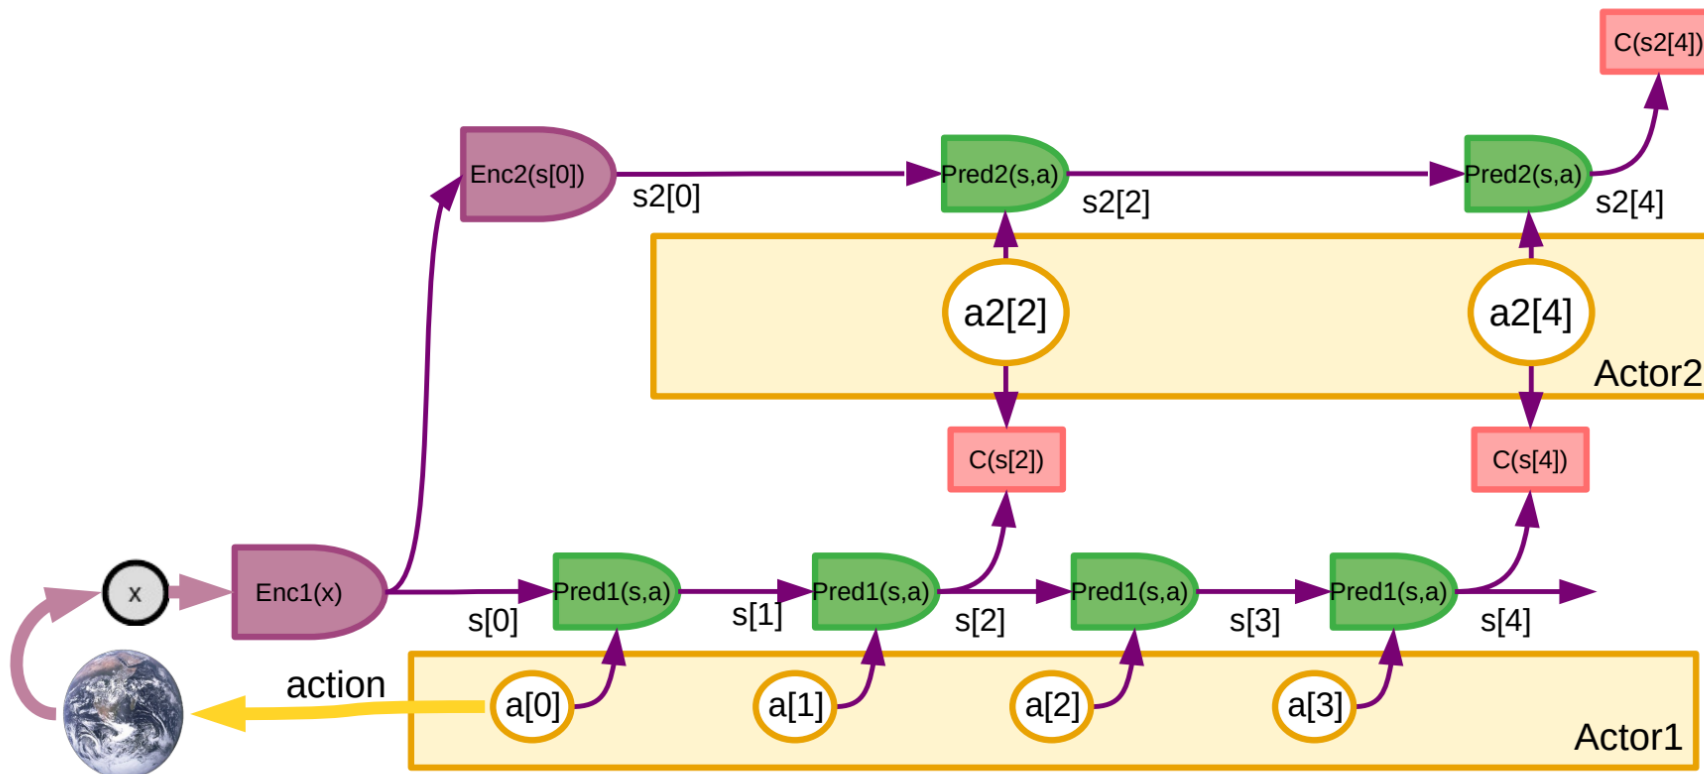
Information Pipeline – Bias for Representation



The next level


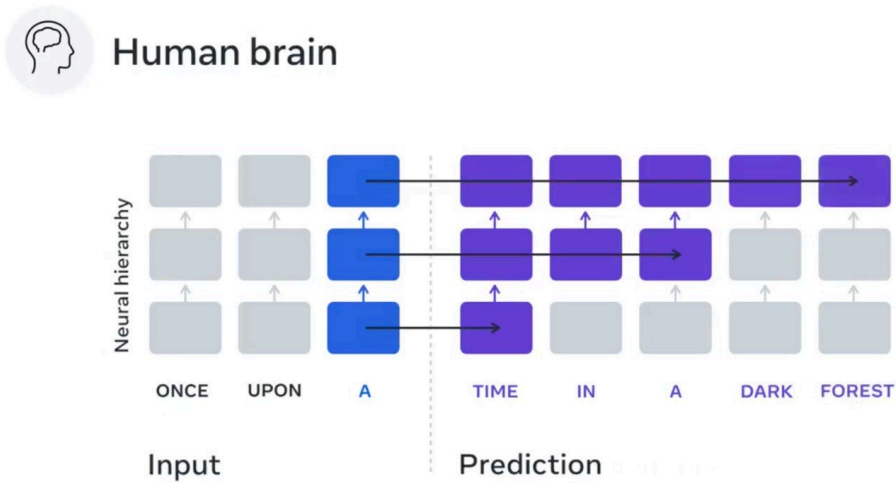
Hierarchical Prediction

Hierarchical JEPA



A Path towards Autonomous Machine Intelligence. Yann LeCun. 2022.

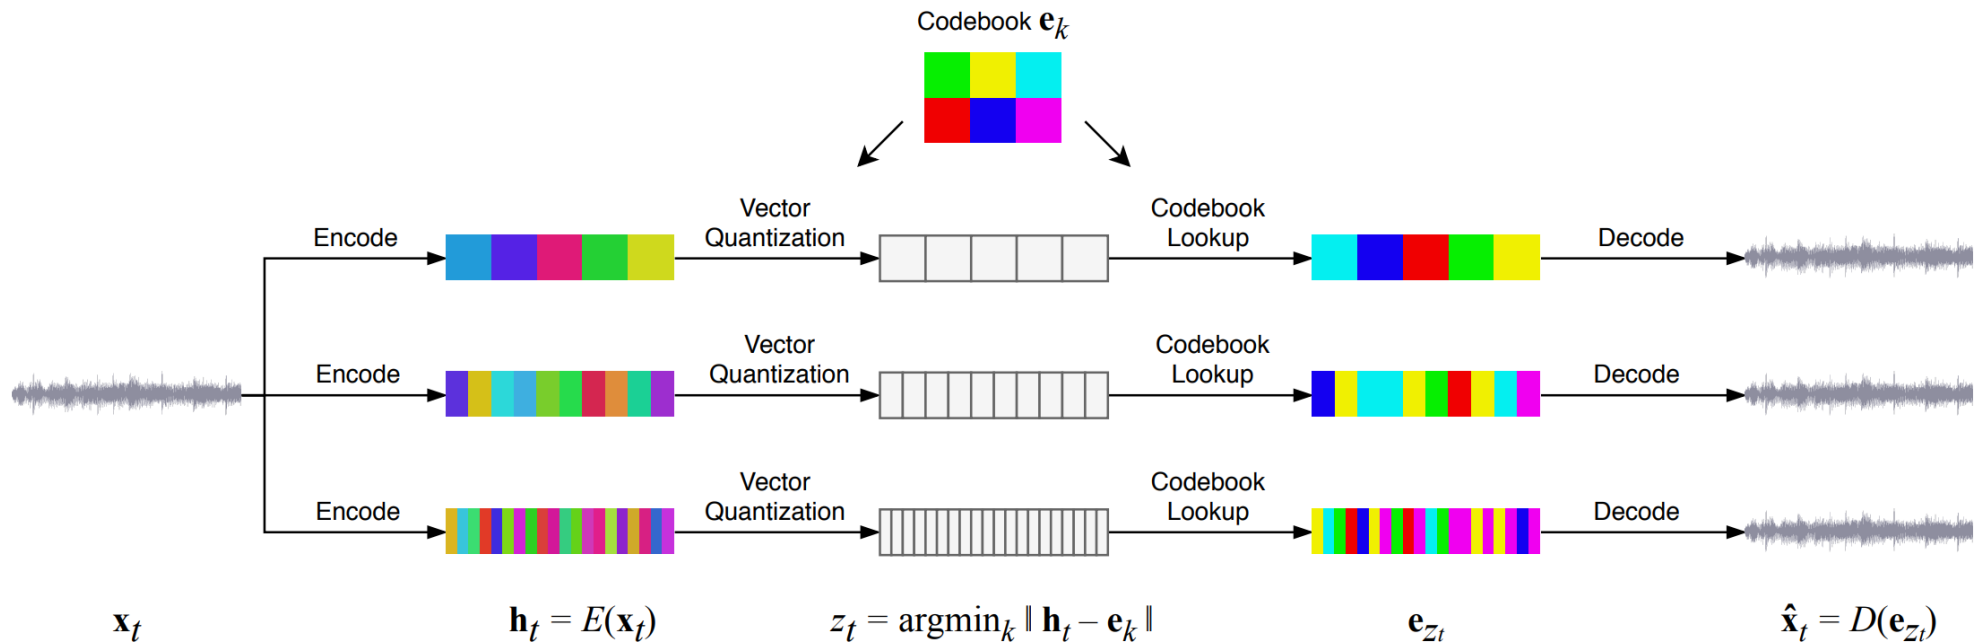
Hierarchical Prediction?



Evidence of a predictive coding hierarchy in the human brain listening to speech. Caucheteux. 2022. Nature Human Behaviour.


How to represent part-whole hierarchies in a neural network. Hinton. 2021.

Hierarchical Prediction – Jukebox (OpenAI)




Jukebox: A Generative Model for Music. Dhariwal et. al. 2020.


Hierarchical Prediction - Feature Pyramid Network




(a) Featurized image pyramid



(b) Single feature map



(c) Pyramidal feature hierarchy



(d) Feature Pyramid Network

Discussion