




---

# Into the Black Box: On Artificial Neural Networks

---


Alexander G. Ororbia II  
Introduction to Machine Learning  
CSCI-635  
11/15/2023

We still have a lot to learn....




# Artificial Neural Networks (ANNs): Neurobiological Motivations


- ▶ Human **brain** = a good candidate inference engine and learning algorithm
  - ▶ Evidence of layered architectures in neuroscientific research, i.e., cortical structures
- ▶ Early success of specialized yet deep architectures
  - ▶ Convolutional Networks, NeoCognitron




<http://cs.brown.edu/~tld/projects/cortex/>



**An organic take on a neural system (Rummelhart, Hinton, & McClelland, 1986).**



**Threshold Output Function**




$$net_i = \sum_j \omega_{ij} o_j(t)$$

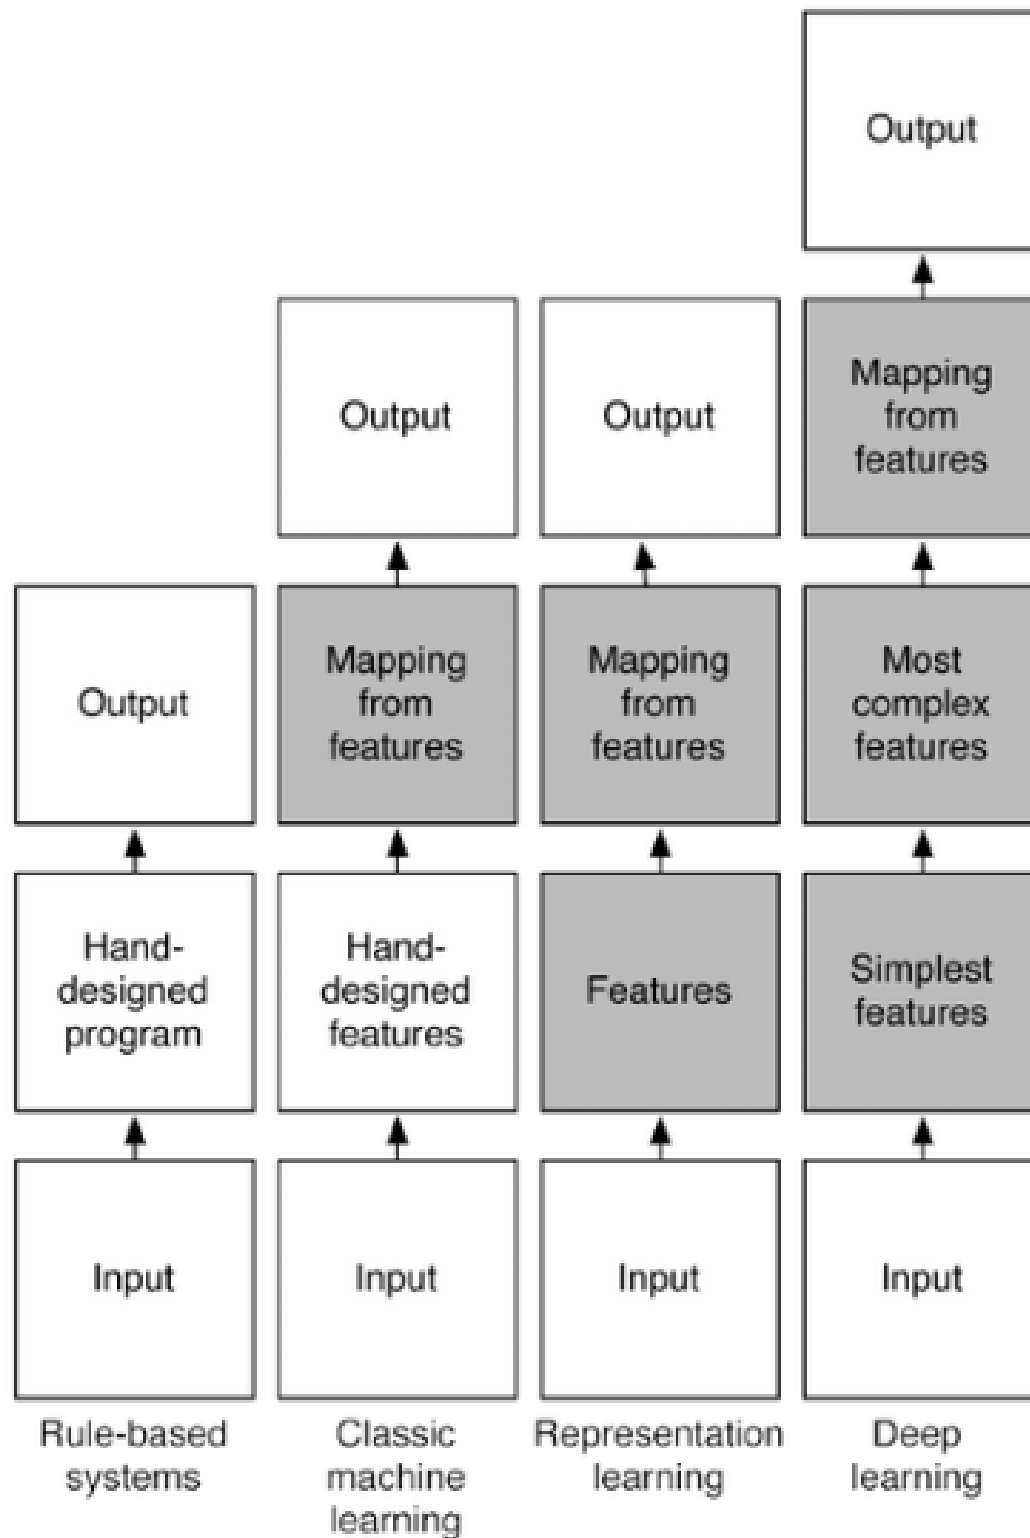
**Sigmoid Activation Function**


**FIGURE 1.** The basic components of a parallel distributed processing system.

# Why? Feature Abstraction

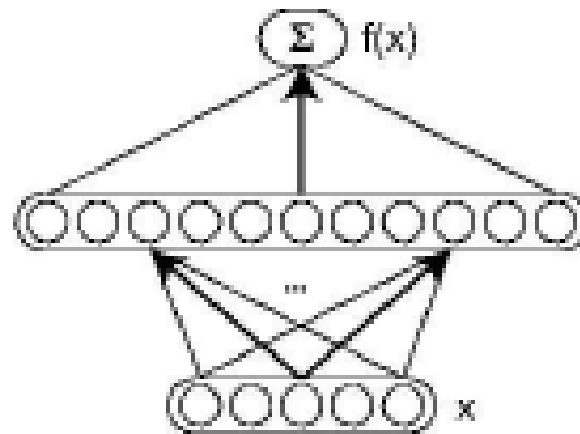
- ▶ Raw features, e.g., pixel values of image, viewed as “low-level” data representation
  - ▶ Can be complex & high-dimensional
  - ▶ Observed variables: “nature” sensory samples, observed/recorded
- ▶ **Abstract representations** = layers of feature detectors
  - ▶ Latent /unobserved variables that describe observed variables
    - ▶ Capture key aspects of underlying stochastic process
    - ▶ Many concepts can be represented as (strict) hierarchies (taxonomy of species) → goal of model is to “learn” a plausible, structured unknown hierarchy
  - ▶ Idea: extracting “structure” from “unstructured”/messy data
- ▶ **Automatic feature engineering**



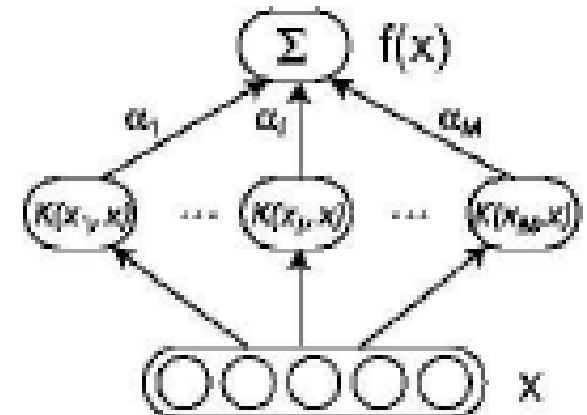




(a) Linear model architecture




(b) Single layer neural network architecture



(c) Kernel SVM architecture

**Most of machine learning can be viewed as a type of ANN...if you squint hard enough...**







## Deep cat detector!




Face detector



Human body detector



Cat detector





## Background

# A Recipe for Machine Learning

1. Given training data:

$$\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$$

2. Choose each of these:

– Decision function

$$\hat{\mathbf{y}} = f_{\boldsymbol{\theta}}(\mathbf{x}_i)$$

– Loss function

$$\ell(\hat{\mathbf{y}}, \mathbf{y}_i) \in \mathbb{R}$$

3. Define goal:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^N \ell(f_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{y}_i)$$

4. Train with SGD:

(take small steps opposite the gradient)

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta_t \nabla \ell(f_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{y}_i)$$

# QUESTIONS?

