AI: A Signal Processing Perspective

Brian Sadler

AI Landscape

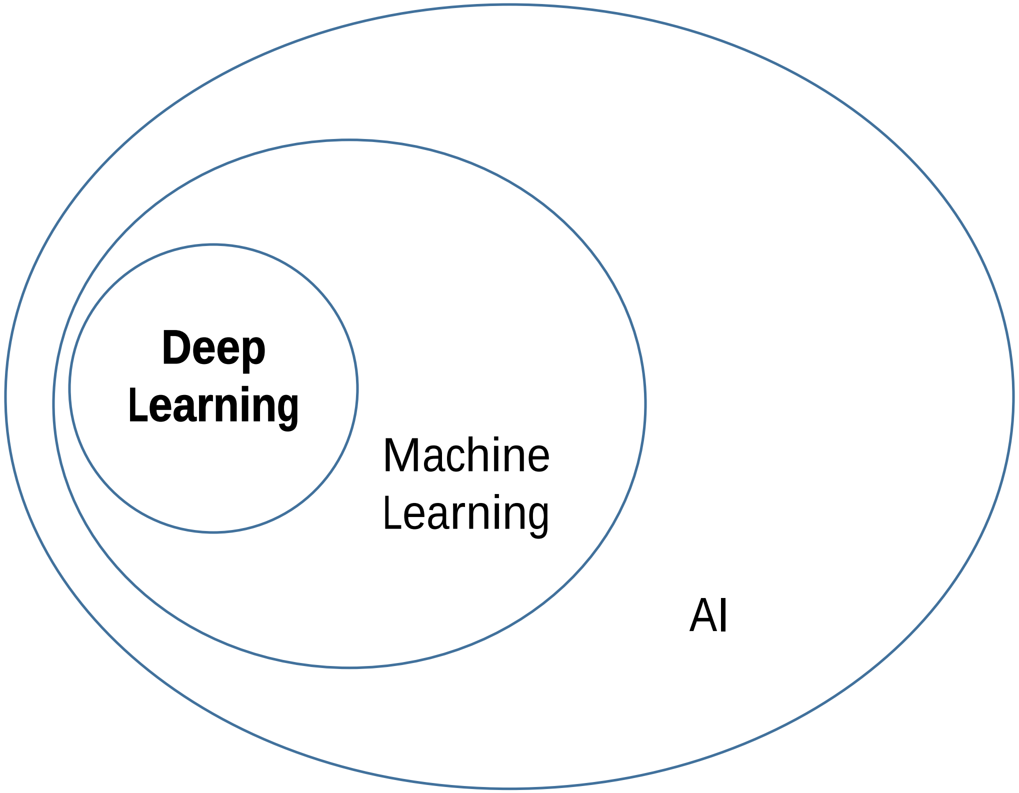

SP Landscape

Natural (e.g., speech, vision)

Man-Made Signals (e.g., communications, virtual reality)

Stochastic Models
Noise, Interference, Corruption

Information Science: Abstractions, Sensing, Transfer, ...

Estimation

Detection

Signals

Classification

Features
Representations

Manipulation

Dynamical Systems
Adaptive Filters
Kalman Filter
Sequence Estimation

Filtering
Separation
Morphing
Decomposition

Generation
Design

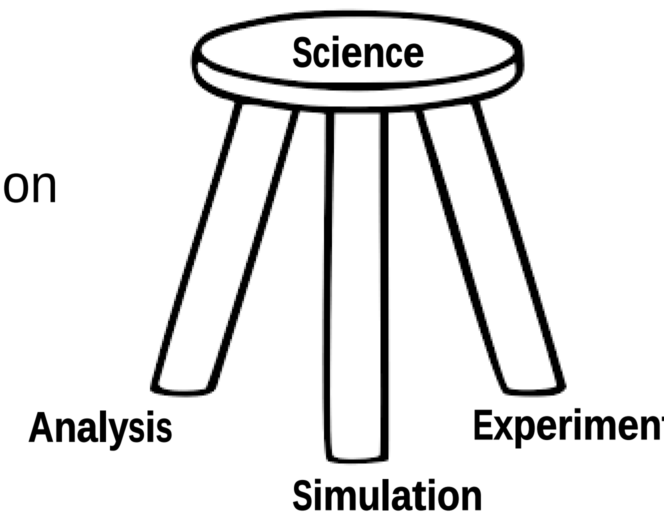

AI & SP

Estimation

Detection Signals Classification

Manipulation

Linear vs Non-Linear

New Function Classes

Representation & Manipulation

Data Driven

Performance Analysis & Statistical Confidence

Probability & Novelty

AI: Deep Learners

- Advances in nonlinear function approximation (learning)

$$\tilde{f}^* = \underset{f \in \mathscr{F}}{\operatorname{argmin}} \, \mathbb{E}_{\mathbf{x},y}[\ell(f(\mathbf{x}), y)]$$

Annotated Data
Input x, output y

Function Family
(DNN architecture)

Loss Function
(e.g., MSE)

- Functions:

  Probabilities, features, representations, manipulations

  Nonlinear and data-driven

AI: Deep Learners

- Advances in nonlinear function approximation (learning)
  - From high-dimensionality to Sparse Representation
  - Supervised estimation (mostly)
  - Continuous wrt model parameters: enables stochastic gradient descent (SGD) optimization
  - Leverage Moore's Law (specialized circuits, memory, computational architectures, multi-core machines)

Deep Learning (2016)

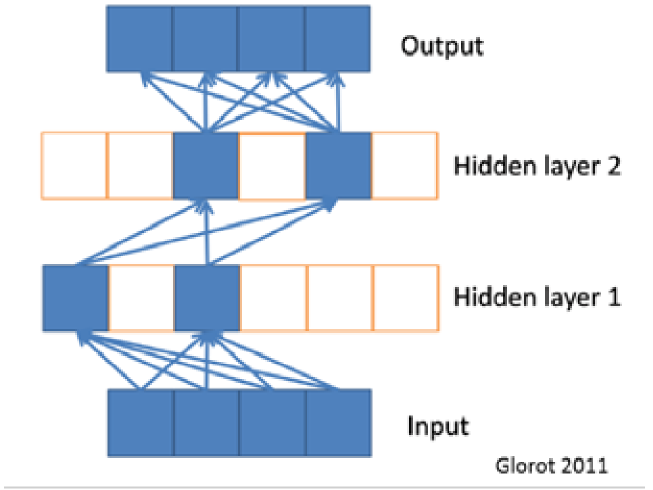

AI: Deep Learners

- Issues
  - Model interpretability difficult
  - Massive models generally require massive training
  - Many hyper-parameters, set by experiment
  - Performance analysis & bounds not easily derived
  - Modeling for robustness & outliers not straightforward
  - Adaptive updating not straightforward

- Normalization
- Dropout
- Stride
- Hand-designed training augmentation
- Parameter choices: guess & check
- SGD: Gradient clipping, momentum, …

- Good news: code sharing community

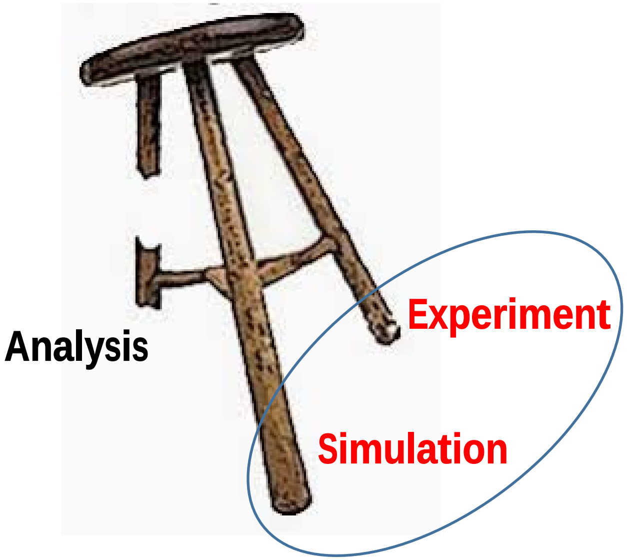

Qualitative DNN History

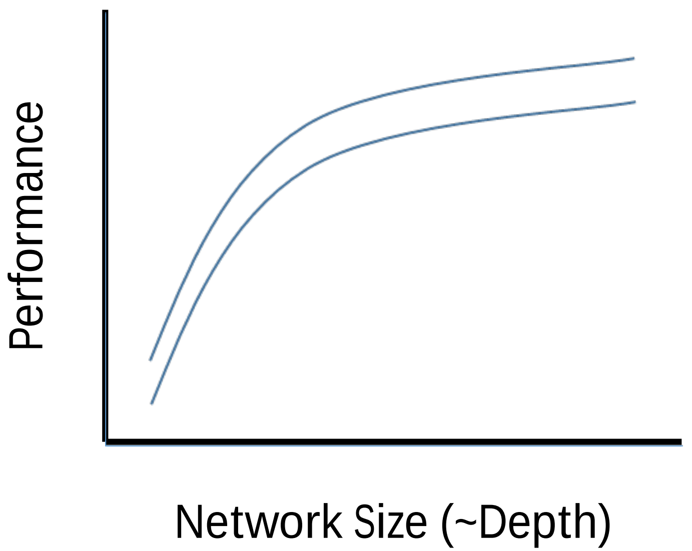

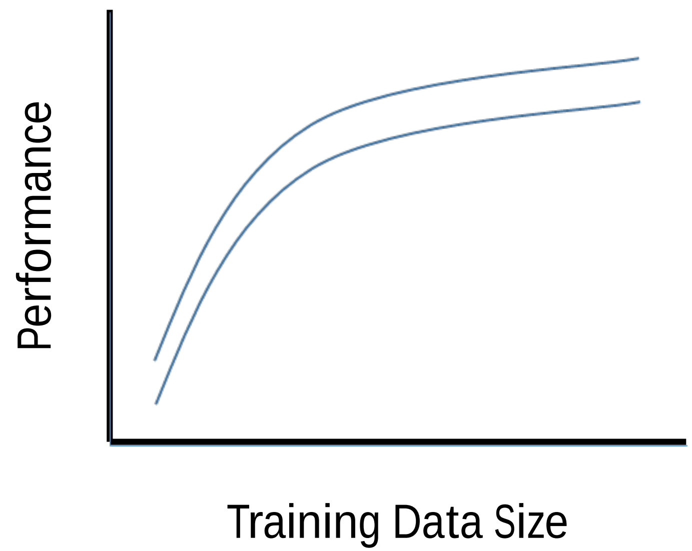

CPU to GPU for training speedup

Historical ebb & flow with unsupervised pre-training, & fully supervised training

ImageNet
2009, 14Mil images

DNN Classifiers

ImageNet Error Rates

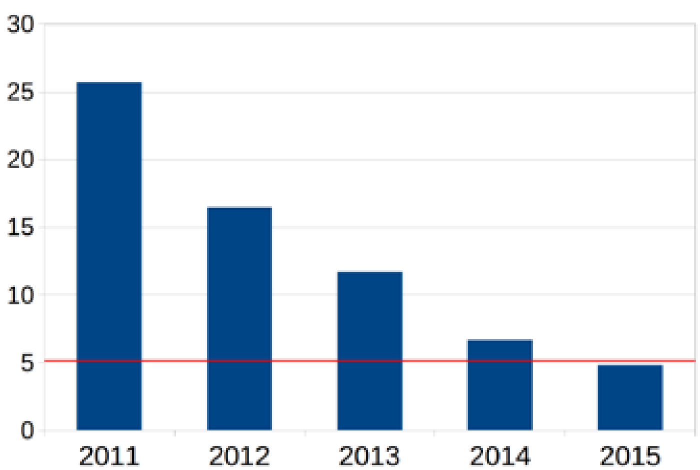

At or exceeding “human performance”

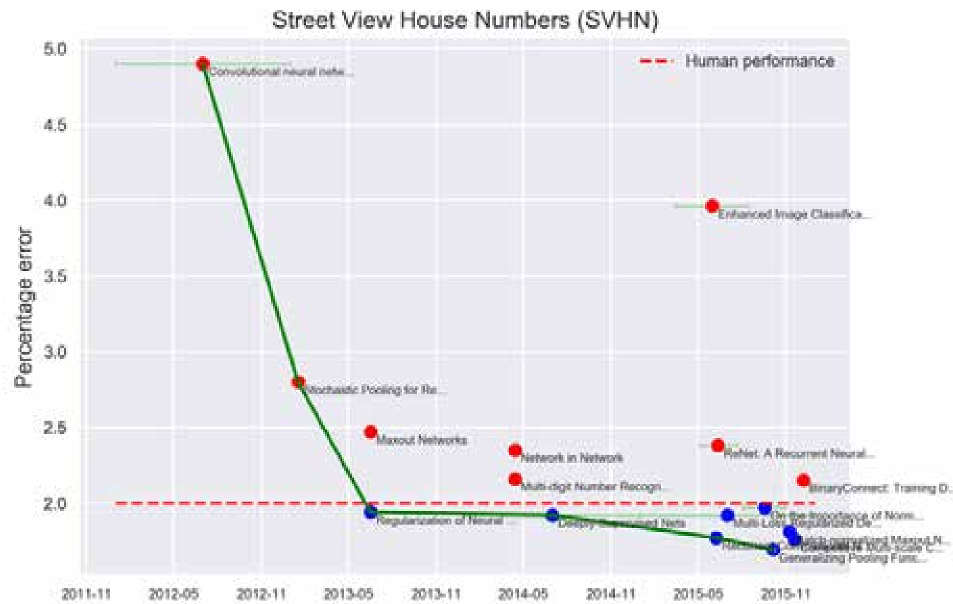

Google

CNN: Convolutional Neural Network

- Modular, intuitive, feed-forward, efficient
- Dominant approach for image processing, manipulation, object detection, object classification

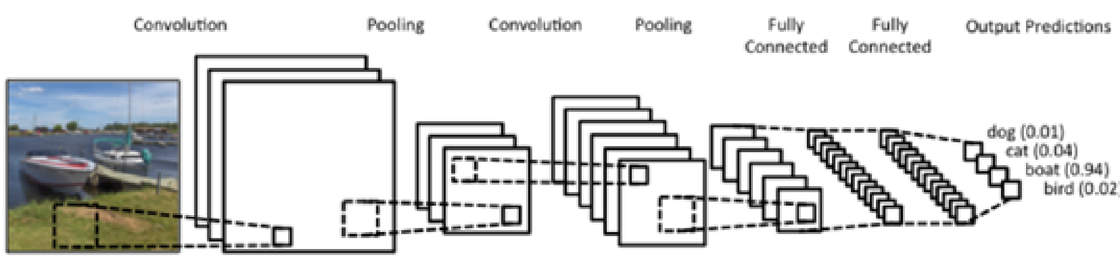

Graphic: Random Grab from Internet

CNN Layer

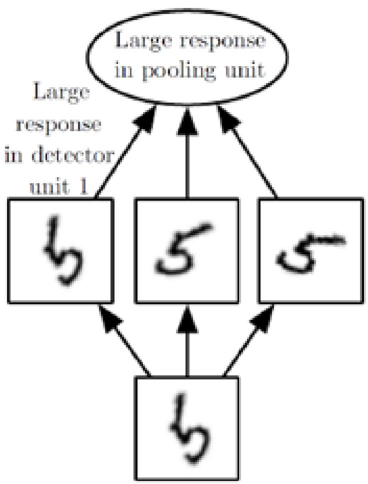

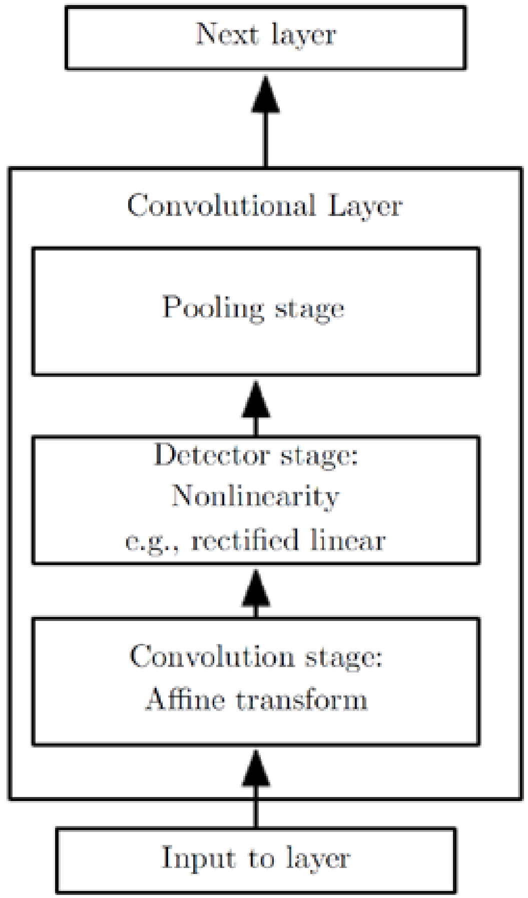

Stride (down sampling)

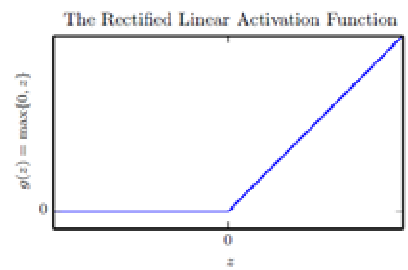

Rectified Linear Unit (ReLU)
(Half-Wave Rectifier)

Deep Learning (2016)

CNN: Sparse Activation

- **Invariance:** combine training & prior information
- Translation, rotation, scaling, elastic deformations
- CNN learns the features, and learns to pool them, to achieve invariance
- Sparse use of layered (distributed) feature representation is efficient and avoids the curse of dimensionality

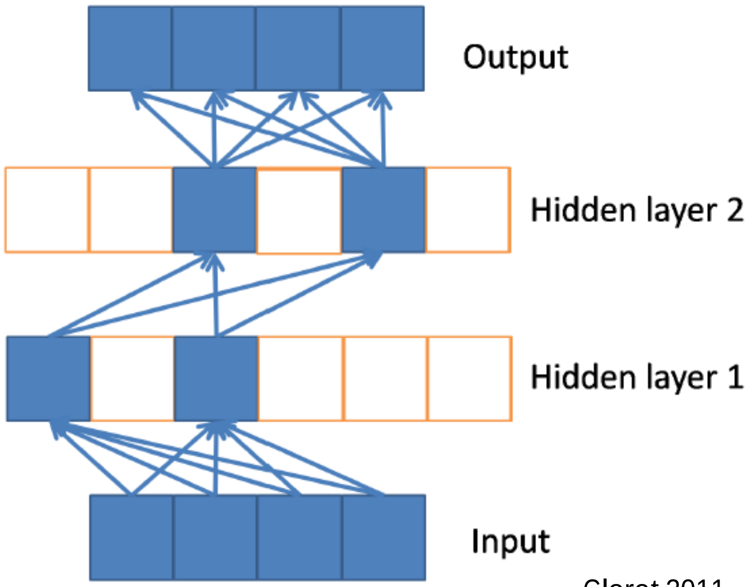

Glorot 2011

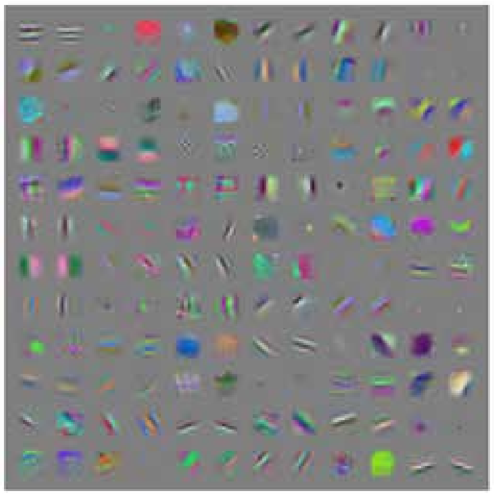

Learned Filters: 1st Layer

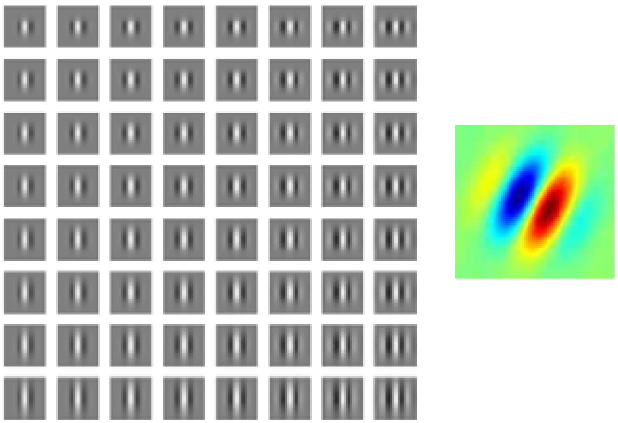

Gabor Filters

Deep Learning (2016)

CNN: Classification

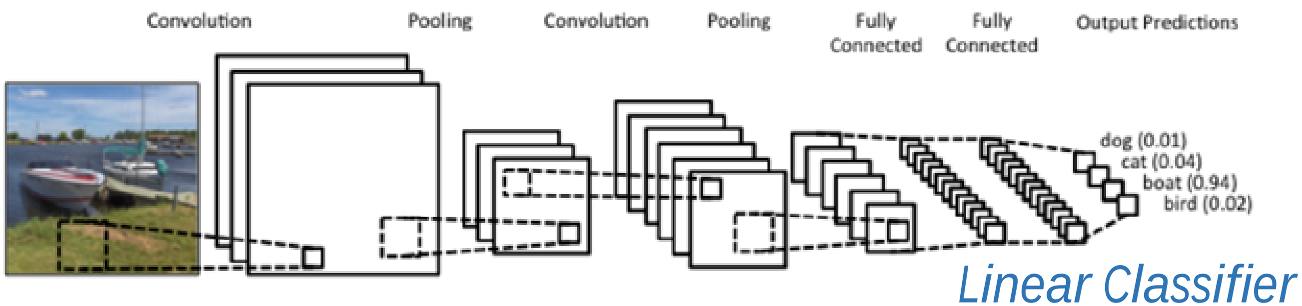

Although images lie on a highly nonlinear manifold, CNN maps images to **representations** that are **linearly** separable

After decades of feature selection & classification ...

Tractable bio-inspired function class, computational feasibility, sufficient data, persistent experimentation

Learns the features, and learns to pool them, simultaneously

Adversarial Examples

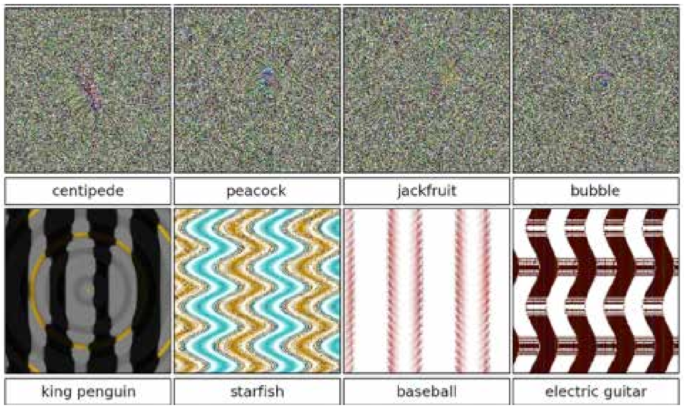

Nguyen et al., 2015

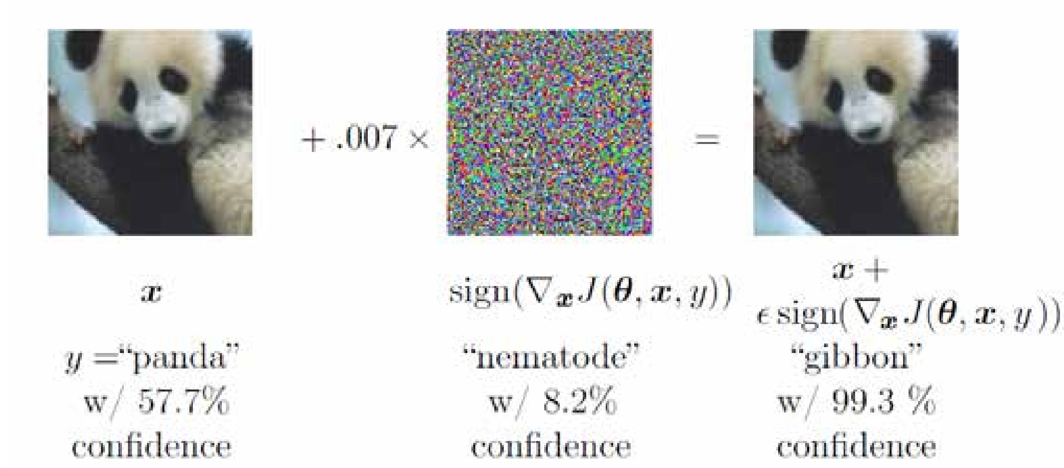

- Training leads to concentration around a low-Dim manifold
- f(x) may behave correctly near the manifold but not off it
- Unreliable estimation may occur with input "far" from the distribution of the training data

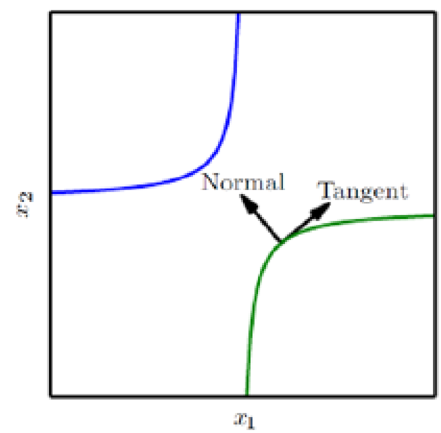

Deep Learning (2016)

Example 1: Probability and Novelty Detection

- **Autonomous navigation**: train DNN to estimate probability of collision given camera input & action (direction, velocity)
- Control: Infer safe navigation from visual content & structure
- Robustness issue: Detect novel environment (untrained upon) and act accordingly

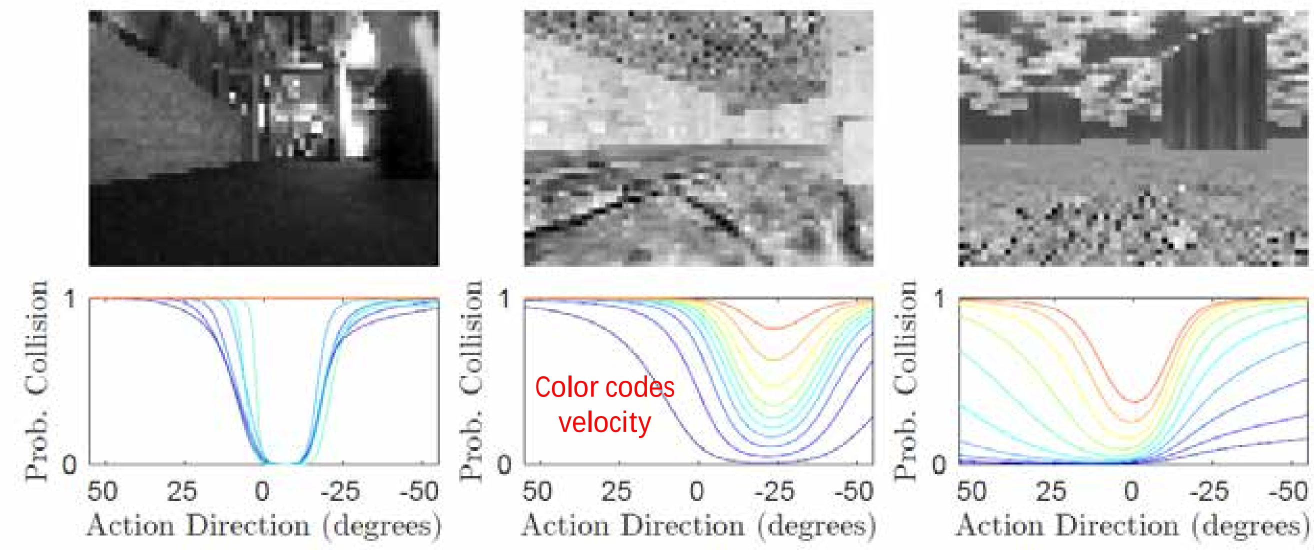

(a) Real hallway. (b) Simulated hallway. (c) Simulated forest.

$$f_c(c|i_t, a_t)$$

Richter, Roy, RSS 2017

Example 1: Autoencoder & Novelty Detector

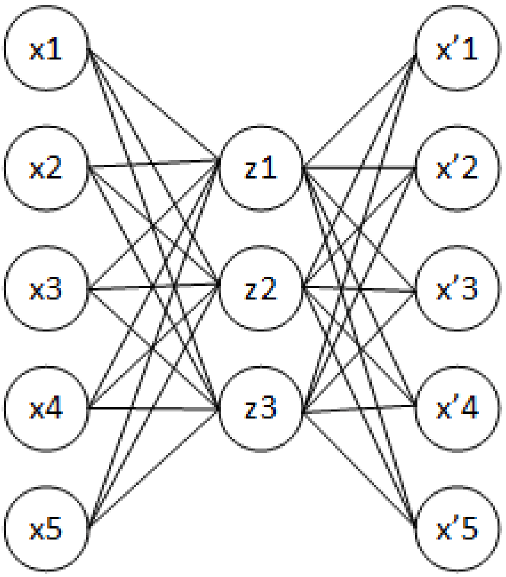

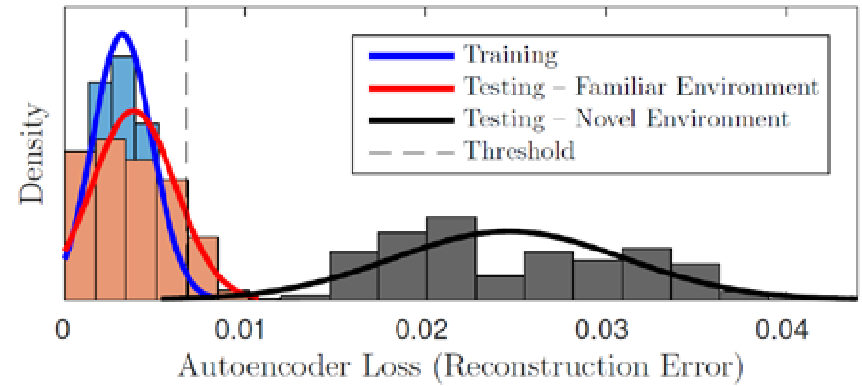

$$L_n(i, \hat{i}) = \frac{1}{K} \sum_{k=1}^{K} (i^k - \hat{i}^k)^2$$

- DNN: Compress & reconstruct
- Unsupervised: minimize reconstruction error
- Learns a compressed representation of the signal class

- Novelty Detector: Does trained autoencoder faithfully reproduce a new input?
- Reconstruction error grows with input "novelty"

Richter, Roy, RSS 2017

Example 1: Robust Control

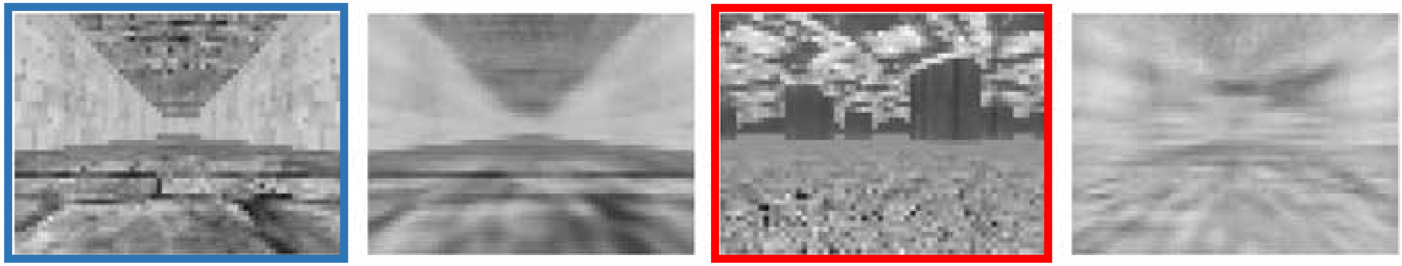

(a) Input. (b) Output. (c) Input. (d) Output.

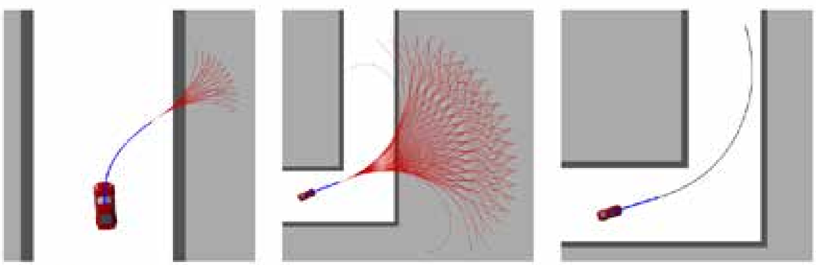

(a) "Collision". (b) "Collision". (c) "Non-Collision".

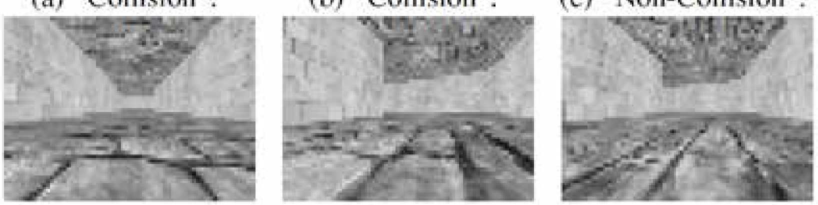

(d) Image for 1(a). (e) Image for 1(b). (f) Image for 1(c).

$$f_i(i_t, b_t, a_t) = \begin{cases} \boxed{f_{\text{net}}(i_t, a_t)} & \text{if } f_{\text{novel}}(i_t) = 0 \\ \boxed{P_{\text{pr.}}(b_t, a_t)} & \text{if } f_{\text{novel}}(i_t) = 1 \end{cases}$$

Control falls back to more conservative prior when in unknown environment

Richter, Roy, RSS 2017

Representation

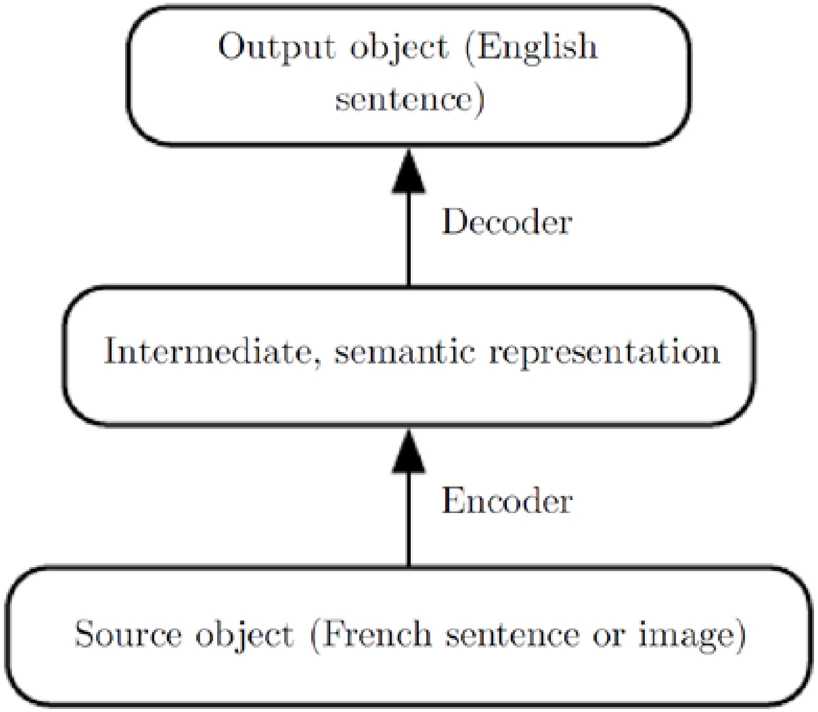

context variable

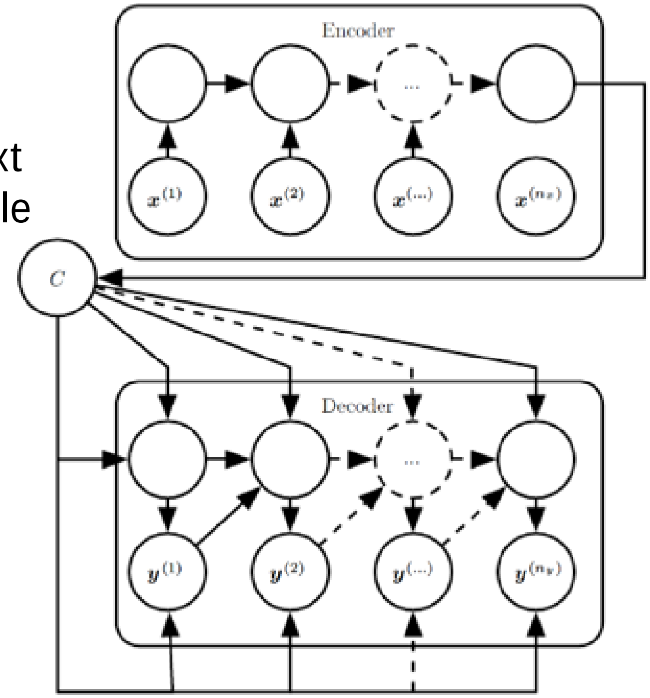

Sequence to sequence mapping

Learn encoder-representation-decoder simultaneously

Deep Learning (2016)

Representation & Manipulation

- Representations as features: classification
- Representations as sufficient measures: language translation
- Manipulating Representations

Example 2: Image Manipulation by Deep Feature Interpolation (DFI)

- Image manipulation by linear interpolation in feature space

K. Bala, et al., 2017

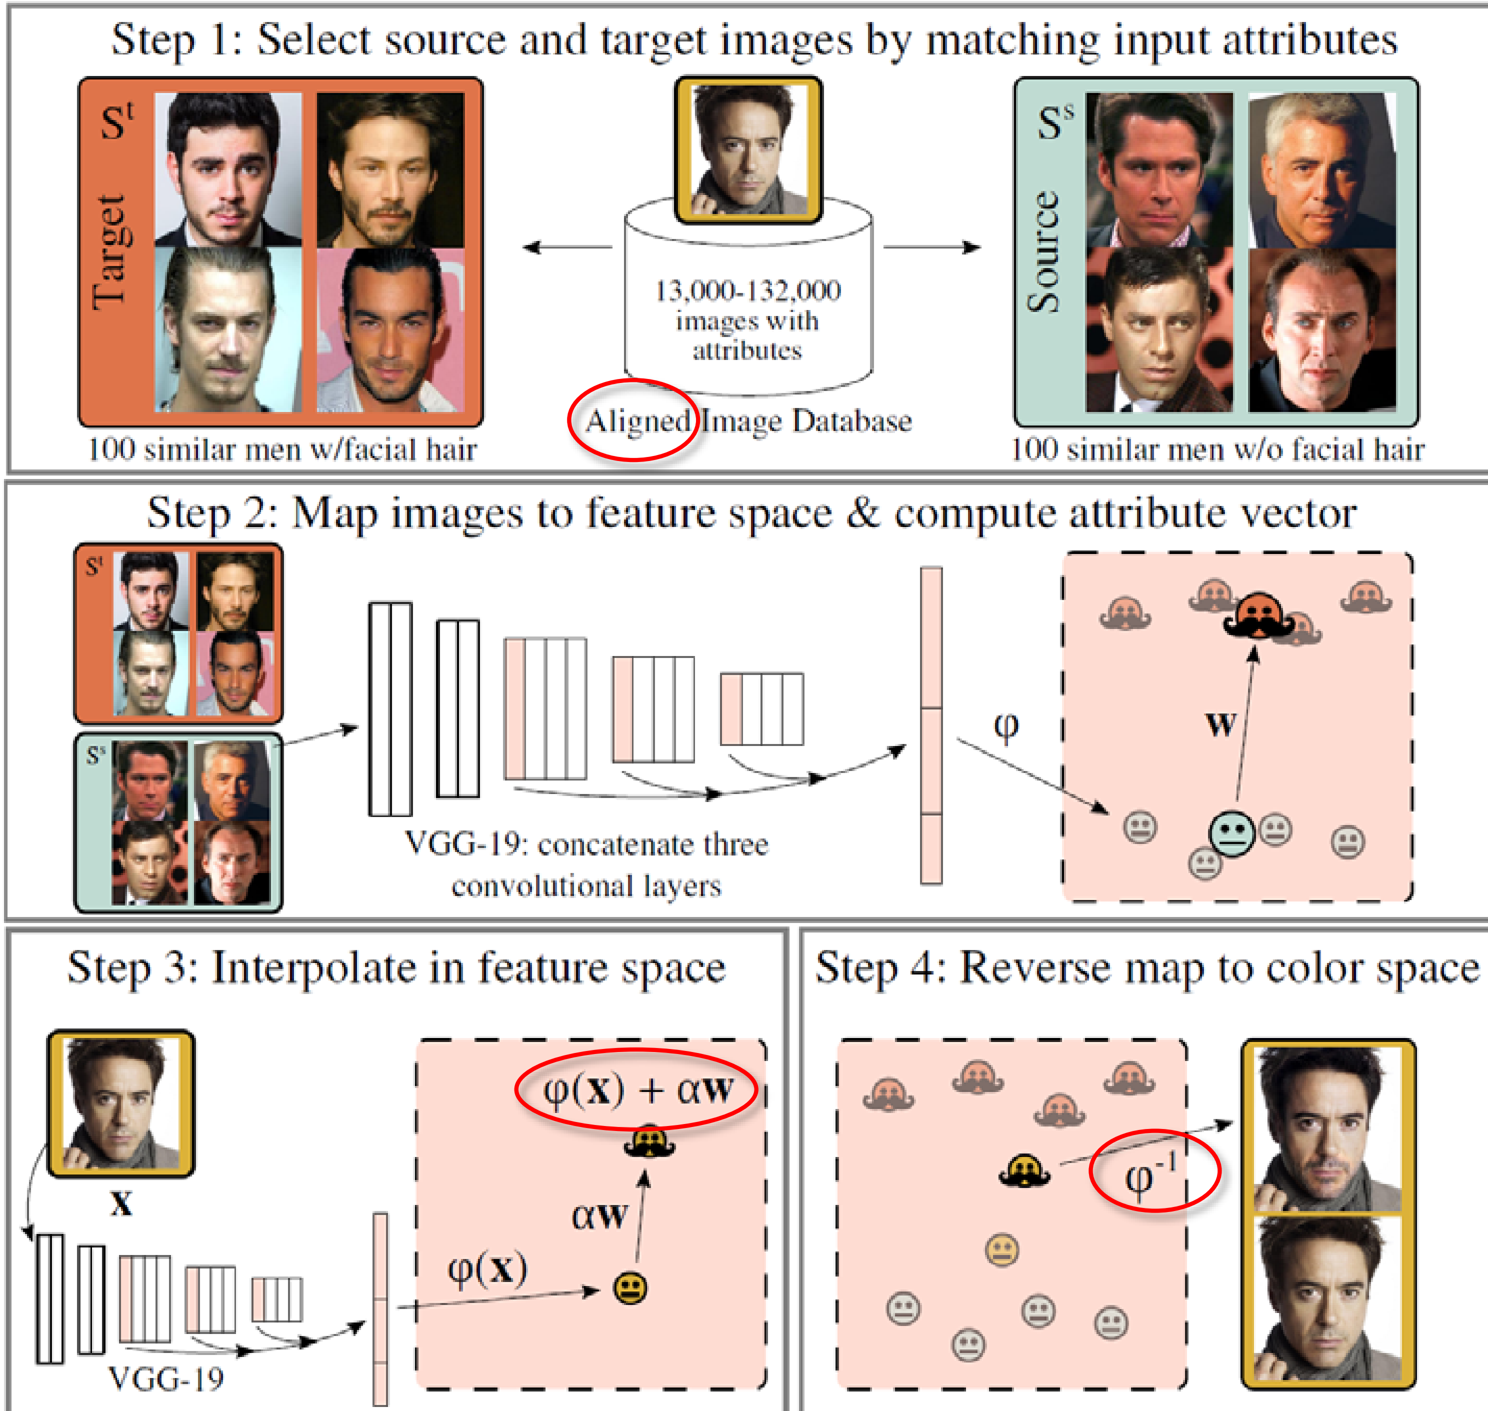

K. Bala, et al., 2017

Example 2: Image Reconstruction

$$\mathbf{z} = \arg\min_{\mathbf{z}} \frac{1}{2} \|(\phi(\mathbf{x}) + \alpha \mathbf{w}) - \phi(\mathbf{z})\|_2^2 + \lambda_{V^\beta} R_{V^\beta}(\mathbf{z})$$

- Reconstruction via Optimization
- Find the image that best maps to the new representation
- Total variation regularization for smoothness

Example 2: DFI Outputs

Representation Interpolation

Attribute Matching

K. Bala, et al., 2017

AI & SP

Estimation

Detection **Signals** Classification

Manipulation

Linear vs Non-Linear

New Function Classes

Representation & Manipulation

Data Driven

Performance Analysis & Statistical Confidence

Probability & Novelty

AI: Discussion

- Our response to AI
  - Human imagination of AI dramatically outperforms our ability to implement AI
  - Attracted & repelled simultaneously
  - Want smart machines, and fear what this means
  - Hills & valleys of AI progress perceived with very sharp gradients

- AI has a time-varying definition
  - Understanding & implementation breeds acceptance & normalcy

- Can we do better than biology?
  - Sometimes (e.g., wheel)

Intelligent Systems

Autonomy / Swarms

Knowledge Bases & HPC

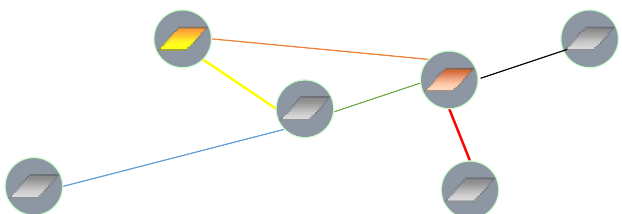

Network

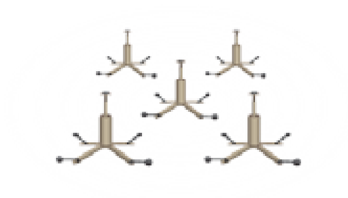

Sensors

Experts

END