

---

---

---

---

---

---

---

---



---

---

---

---

---

---

---

---

### DISCLOSURES

- I have no financial conflicts of interest to report
- Research support from NIH and NSF

**UF** Herbert Wertheim  
College of Engineering  
UNIVERSITY of FLORIDA

---

---

---

---

---

---

---

---

## OBJECTIVES

1. Learn about limitations of conventional models
2. What do deep learning models look like?
3. Deep learning in action

---

---

---

---

---

---

---

---

## AI, MACHINE LEARNING, DEEP LEARNING

The diagram consists of three overlapping circles. The largest, outermost circle is yellow and labeled 'Artificial Intelligence (AI)'. Inside it is a smaller orange circle labeled 'Machine Learning'. Inside the orange circle is the smallest, innermost circle, which is red and labeled 'Deep Learning'. This illustrates that Deep Learning is a subset of Machine Learning, which is a subset of Artificial Intelligence.

---

---

---

---

---

---

---

---

## MACHINE LEARNING VS. STATISTICS

- Machine learning deals with high-dimensional data
  - Data can be unstructured
    - E.g. video, images, natural language data

The diagram shows two overlapping green circles. The left circle is labeled 'Statistics' and the right circle is labeled 'Machine Learning'. The overlapping area in the center represents the intersection of the two fields.

---

---

---

---

---

---

---

---

# MACHINE LEARNING TASKS

- Supervised Learning**
  - **Input:** data, labels
  - **Output:** label prediction for new data points
- Unsupervised Learning**
  - **Input:** data
  - **Output:** discover the latent structure of data points
- Reinforcement Learning**
  - **Input:** data, occasional feedback
  - **Output:** learned policy as series of actions

---

---

---

---

---

---

---

---

# CENTRAL PROBLEM

- To transform data to learn useful representations of the input data (encode), i.e. **representation learning**

---

---

---

---

---

---

---

---

# CONVENTIONAL MACHINE LEARNING MODELS

- There are already many different machine learning (ML) and statistical learning models
  - Logistic regression
  - Support Vector Machines (SVM)
  - Random forests

Decision Tree                      Support Vector Machine

---

---

---

---

---

---

---

---

# LIMITATIONS OF CONVENTIONAL ML MODELS

- Shallow models that are unable to model highly complex data
  - E.g. video or image analysis
- Features need to be provided to the model, but what features for each problem?
  - Feature engineering

---

---

---

---

---

---

---

---

---

---

# NEURAL NETWORKS HISTORY

GOES BACK TO 1940, WITH SEVERAL DARK AI WINTERS

---

---

---

---

---

---

---

---

---

---

# ELECTRONIC BRAIN - 1943

$$\Phi(x) = \begin{cases} 1 & \text{if } x \geq \theta \\ -1 & \text{otherwise} \end{cases}$$

- Warren McCulloch and Walter Pitts published the first concept of a simplified brain cell (1943).
- A simple logic gate with binary outputs
  - Multiple signals arrive at the dendrites.
  - Signals then integrated into the cell body.
  - If the accumulated signal exceeds a certain threshold, an output signal is generated that will be passed on by the axon.

---

---

---

---

---

---

---

---

---

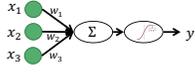
---



$$y = \sum_{j=1}^n w_j x_j + w_0 = \mathbf{w}^T \mathbf{x}$$

$$\mathbf{w} = [w_0, w_1, \dots, w_n]^T$$

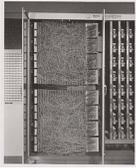
$$\mathbf{x} = [1, x_1, \dots, x_n]^T$$



**PERCEPTRON - 1957**

- The basic processing element (Rosenblatt, 1957)
- Associated with each input  $x_i$  is a connection weight  $w_i$
- Rosenblatt proposed an algorithm to automatically learn the optimal weights

The Mark I Perceptron machine was the first implementation of the perceptron algorithm. The machine was connected to a camera that used 20x20 cadmium sulfide photocells to produce a 400-pixel image. The main visible feature is a patchboard that allowed experimentation with different combinations of input features. To the right of that are arrays of potentiometers that implemented the additive weights.  
[Wiki - Perceptron](#)



---

---

---

---

---

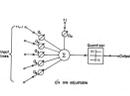
---

---

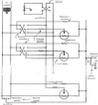
---

---

---







Adaline

$$E(w) = \frac{1}{2} \sum_t (y_t - \Phi(z_t))^2$$

**ADALINE - 1960**

- Proposed by Widrow & Hoff
- Illustrates the key concept of defining and minimizing a cost function
- The groundwork for understanding more advanced techniques
- Key difference from perceptron
  - Compared to perceptron, we can use a continuous function to compute the error
  - Differentiable

---

---

---

---

---

---

---

---

---

---



Winter is coming

AI WINTER  
(1974 - 1980)

Picture & quote from HBO series, Game of Thrones

---

---

---

---

---

---

---

---

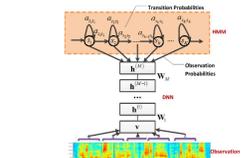
---

---

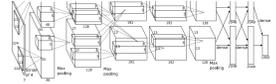


## FIRST IMPRESSIVE RESULTS

STARTING IN  
2010-2012



Dahl, George E., Dong Yu, Li Deng, and Alex Acero. "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition." *IEEE Transactions on audio, speech, and language processing* 20, no. 1 (2012): 30-42.



Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." In *Advances in neural information processing systems*, pp. 1097-1105. 2012.

---

---

---

---

---

---

---

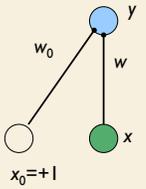
---

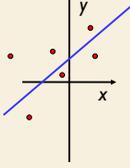
---

---

## REGRESSION

$$y = \sum_{j=1}^d w_j x_j + w_0 = \mathbf{w}^T \mathbf{x}$$





---

---

---

---

---

---

---

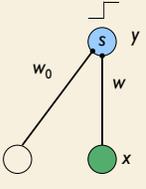
---

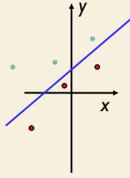
---

---

## CLASSIFICATION

$$\Phi(x) = \begin{cases} 1 & \text{if } x \geq \theta \\ -1 & \text{otherwise} \end{cases}$$





---

---

---

---

---

---

---

---

---

---

## CLASSIFICATION + POSTERIOR PROBABILITY

$y = \text{sigmoid}(\mathbf{w}^T \mathbf{x}) = \frac{1}{1 + \exp[-\mathbf{w}^T \mathbf{x}]}$

---

---

---

---

---

---

---

---

## PERCEPTRON ONLINE LEARNING

- We do not write the error function over the whole sample at once
  - But on individual instances at each step

- Start from random weights
- At each iteration, adjust parameters a little bit to minimize error based on current input

$$\Delta w_j = \eta (y_i - \hat{y}_i) x_i^j$$

Labels: Learning Rate ( $\eta$ ), Input Feature Value ( $x_i^j$ ), Actual value ( $y_i$ ), Predicted value ( $\hat{y}_i$ ), Weight update ( $\Delta w_j$ )

---

---

---

---

---

---

---

---

## FULLY CONNECTED NETWORKS (FCN)

- Composed of multiple hidden layers
  - Every neuron in layer  $i$  is connected to every other neuron in layer  $i + 1$

Input    Hidden Layers    Output

- Even simple NN can be helpful for classification:
  - Normal beat
  - Congestive heart failure beat
  - Ventricular tachyarrhythmia beat
  - Atrial fibrillation beat

Rumelhart, Hinton, Williams (1986)  
Güler, İnan, "ECG beat classifier designed by combined neural network model," Pattern recognition 38.2 (2005): 199-206.

---

---

---

---

---

---

---

---

# GRADIENT DESCENT

- A simple, yet powerful optimization algorithm
- Think of it as **climbing down a hill** until a local or global minimum is reached.
- In each iteration, we take a step away from the gradient where the step size is determined by the value of the **learning rate** as well as the slope of the gradient.

$$\Delta w = -\eta \nabla E(w) \text{ Update}$$

$$\frac{\partial E}{\partial w_j} = \frac{\partial}{\partial w_j} \frac{1}{2} \sum_i (y_i - \Phi(z_i))^2$$

Gradient

$$= - \sum_i (y_i - \Phi(z_i)) x_i^j$$




---

---

---

---

---

---

---

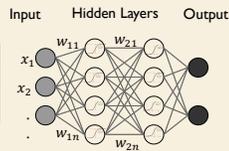
---

# FULLY CONNECTED NETWORKS (FCN)

- Backpropagation algorithm is key

$$h_i = \sigma \left( \sum_{j=1}^d x_j w_{ij} + b_{ij} \right)$$

1. Sample a batch of data
2. Forward through network, compute loss
3. Backpropagate to calculate the gradients
4. Update the weights using the gradients



Rumelhart, Hinton, Williams (1986)

---

---

---

---

---

---

---

---

# DEEP LEARNING

- Learning successive layers of increasingly meaningful representations
  - Mostly learned via models known as neural networks

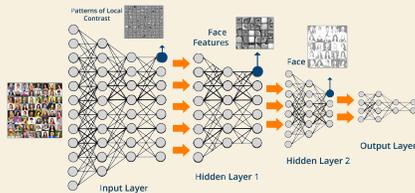


Image from <https://www.edureka.co/blog/what-is-deep-learning>

---

---

---

---

---

---

---

---



# CNNs HAVE BEEN AROUND FOR A LONG TIME

- Goes back to 1989, first successful application *LeNet*
- In recent years, it scaled up due to hardware improvements, algorithmic advances, and data availability

Yann LeCun

LeNet (1989)

---

---

---

---

---

---

---

---

---

---

# IMPRESSIVE PROGRESS

- Error rates for ImageNet challenge have fallen from 28.5% to below 2.5% since 2010

Year	AlexNet	Deep Learning	Human
2010	~26.2%	-	~10.0%
2011	~21.4%	-	~10.0%
2012	~16.7%	~16.7%	~10.0%
2013	~11.7%	~11.7%	~10.0%
2014	~6.7%	~6.7%	~10.0%
2015	~2.8%	~2.8%	~10.0%

- Many biomedical imaging applications
  - Image segmentation & registration
  - Automatic labeling (e.g. lesion detection)
  - Computer-aided diagnosis

Lee, Joon-Seok, Sanghoon Jun, Young-Pil Cho, Hyungsik Lee, Guk-Bae Kim, Joon Beom Seo, and Namik Park, "Deep Learning in Medical Imaging: General Challenges", *Korean Journal of Radiology*, 18, no. 4 (2017), 570-584.

---

---

---

---

---

---

---

---

---

---

# RECURRENT NEURAL NETWORKS (RNN)

- CNNs useful when data has spatial structure
- Recurrent neural networks (RNNs) useful for sequentially ordered data
  - Time series, natural language
- Popular RNN variants
  - Long short-term memory (LSTM)
  - Gated recurrent unit (GRU)

---

---

---

---

---

---

---

---

---

---

# RECURRENT NEURAL NETWORKS (RNN)

- Physiological signals can be processed by LSTMs

---

---

---

---

---

---

---

---

# AUTOENCODERS

- Typically used for unsupervised representation learning
- To encode input into a lower dimensional space
  - Reconstruction and decoding

$$\hat{x} = \sigma(W'z + b')$$

$$z = \sigma(Wx + b)$$


---

---

---

---

---

---

---

---

# AUTOENCODERS

**Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records**

---

---

---

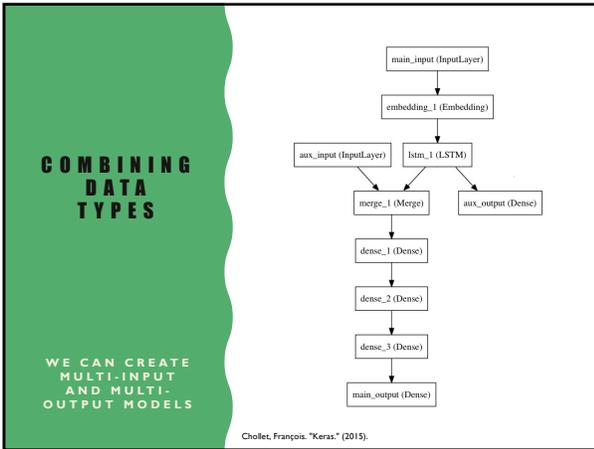
---

---

---

---

---




---

---

---

---

---

---

---

---

**REINFORCEMENT LEARNING**

- Which actions to take to maximize cumulative reward
- Tradeoff between exploration & exploitation

$\alpha = \pi(s)$

Environment

Interpreter

Reward

State

Agent

Action

Success 1000

Shier, David, Julian Schrittwieser, Karan Simonyan, Ioannis Antonoglou, AJa Huang, Arthur Guez, Thomas Hubert et al. "Mastering the game of go without human knowledge." *Nature* 550, no. 7676 (2017): 354.

Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Ruus, Joel Veness, Marc G. Bellemare, Alex Graves et al. "Human-level control through deep reinforcement learning." *Nature* 518, no. 7560 (2015): 529-533.

Image via Wiki - Reinforcement Learning

---

---

---

---

---

---

---

---

**SUCCESS STORIES**

- Near-human-level
  - Image classification, speech recognition, handwriting transcription
- Improved
  - Machine translation, text-to-speech, autonomous driving
- Superhuman
  - Go playing

---

---

---

---

---

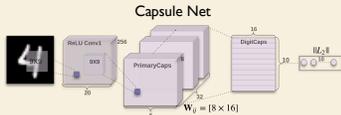
---

---

---

## RECENT PROGRESS - 2017

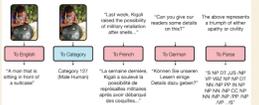
- Capsule net
- One model for all
- Alpha-zero



Sabour, Sara, Nicholas Frosst, and Geoffrey E. Hinton. "Dynamic routing between capsules." In *Advances in Neural Information Processing Systems*, pp. 3859-3869. 2017.



Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert et al. "Mastering the game of go without human knowledge." *Nature* 550, no. 7678 (2017): 354.



Kaiser, Lukasz, Aidan N. Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. "One model to learn them all." *arXiv preprint arXiv:1708.05137* (2017).

---

---

---

---

---

---

---

---

---

---

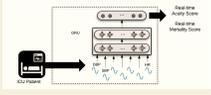
## OUR USE CASES

- Time series: LSTM
  - SOFA Score Prediction
- Video feed: CNN
  - Functional status detection
- Text: LSTM
  - Mental health, clinical notes

AUC @ 24 hours (early detection)



■ Deep Learning  
■ SOFA Baseline




Shickel, Benjamin, Tyler Loffus, Tezcan Ozrazgat Bastanti, Azra Bihorac, and Parisa Rashidi. "1619: Increasing Sofa Score Granularity With Deep Learning." *Critical Care Medicine* 46, no. 1 (2018): 794.

Shickel, Benjamin, Marlin Heesacker, Sherry Benton, and Parisa Rashidi. "Hashtag Healthcare: From Tweets to Mental Health Journals Using Deep Transfer Learning." *arXiv preprint arXiv:1708.01372* (2017).

---

---

---

---

---

---

---

---

---

---

## AI FUTURE

- Human-level general intelligence: still far far away
- Many opportunities for AI in medicine
- Best to moderate our expectations, there is a short-term hype, but there is also a long-term vision



---

---

---

---

---

---

---

---

---

---

**K** Keras

arXiv.org  
open access  
to the world's  
science

CNTK

TensorFlow

ANAconda  
Powered by Scientific Python

jupyter

mxnet

### WHERE TO GO FROM HERE?

- Our recent survey paper on use of deep learning with EHR data [link](#)
- Look for the latest developments on ArXiv
- Many open software tools
  - Keras
  - TensorFlow
  - CNTK
  - Theano
  - MXNet

---

---

---

---

---

---

---

---