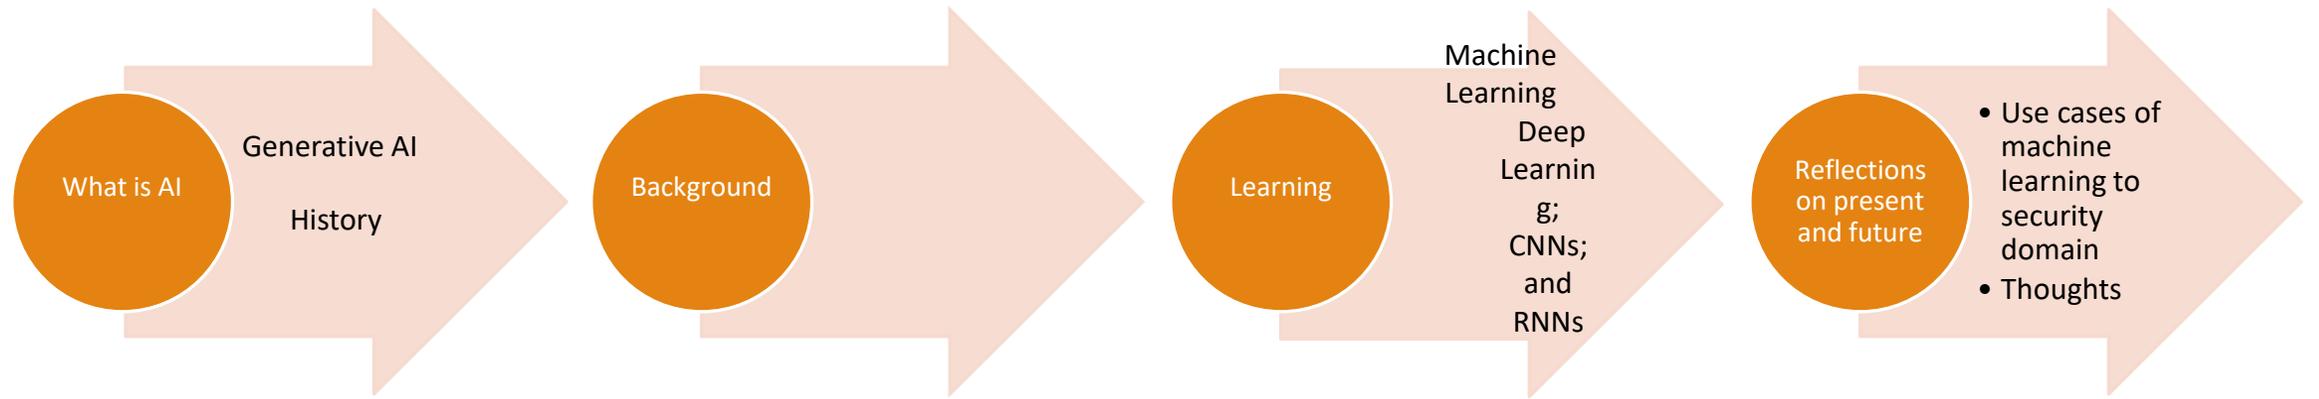


Foundations of AI and its Future Prospects

VIR V PHOHA, PH.D.
PROFESSOR OF EE & CS
SYRACUSE UNIVERSITY
NEW YORK, USA

Overview



What is AI

Science and technology to simulate human intelligence by machines and computers

- Examples: Autonomous vehicles, crop analysis; drug discovery; medical diagnosis; Generative AI tools

Machine Learning; Deep Learning; Rule based

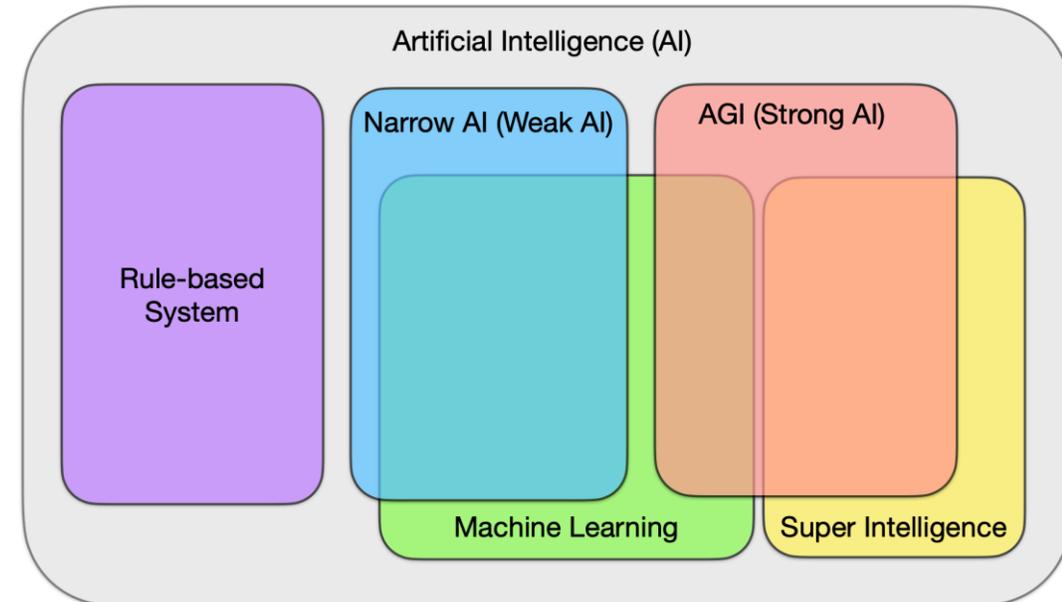
Based on the human decision-making processes that can learn and generalize

Narrow AI (ANI-Artificial Narrow Intelligence)

- Trained and focused on specific tasks– Apple’s Siri, Amazon’s Alexa, IB Watsonx™, Self-Driving Vehicles

Strong AI (AGI-Artificial General Intelligence or Artificial Super Intelligence)

- Aims for intelligence equivalent to humans—Self aware; consciousness ability to solve problems; plan for the future
- ASI- Surpass the human abilities and intelligence– at present mostly theoretical
- Examples: Superhuman and rogue computer assistant in 2001: A Space Odyssey



Generative AI



Deep learning that takes raw data (Wikipedia; works of Rembrandt, etc) and can generate statistically probable outputs

Encode a simplified description of training data and generate new work that may be similar but not identical to the original data

History



1950

Alan Turing –
*Computing Machinery
and Intelligence.*

Can Machines Think?
Turing Test.



1956

John McCarthy
“artificial
intelligence”

Newell, Shaw,
Simon create the
first running AI
software program
Logic Theorist



1967

Mark Rosenblatt –
Mark 1
Perceptron, NN
that learns

Minsky and Papert
book Perceptrons
- argument
against NN



1980

BP becomes
widely used



1997

IBMs Deep Blue
beats world chess
champion Garry
Kasparov



2004

John McCarthy
writes a paper
“What is Artificial
Intelligence” and
defines AI



2011

IBM Watson beats
champions Ken
Jennings and
Brad Rutter at
Jeopardy



2015

Baidu’s Minwa
supercomputer
uses convolutional
neural networks
to identify and
categorize images
better than an
average human



2016

DeepMind’s
AlphaGo program
using deep neural
network beats Lee
Sedol, the world
Go player
(possible moves
14.5 trillion after
four moves)



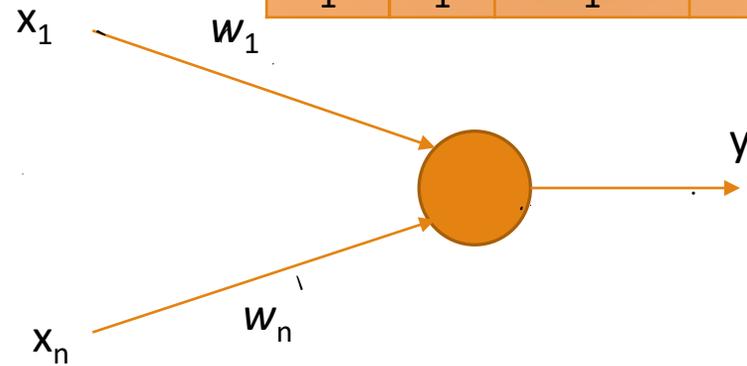
2023

Rise in Large
Language Models,
such as ChatGpt

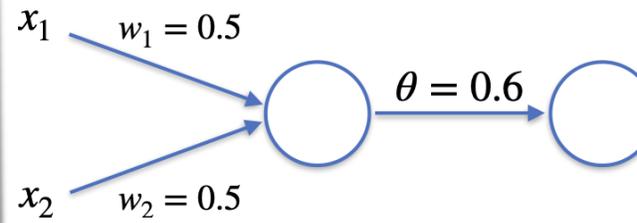
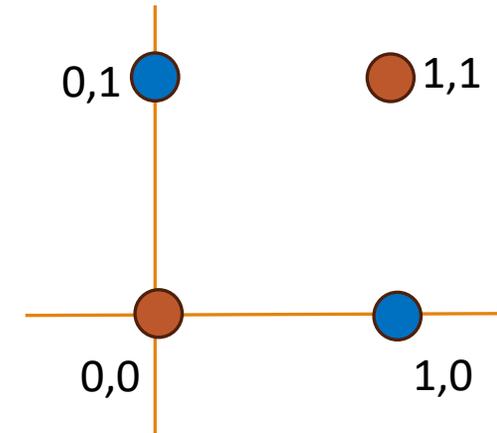
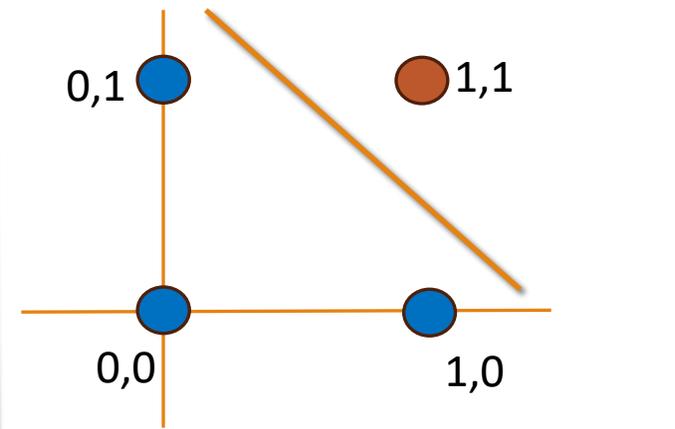
Background – Neuron --Perceptrons

Implementation of AND XOR?

Input		And	XOR
x_1	x_2		
0	0	0	0
0	1	0	1
1	0	0	1
1	1	1	0



$$y = f\left(\sum w_i x_i\right) = \begin{cases} 0, & \text{if } f < \theta \\ 1, & \text{if } f \geq \theta \end{cases}$$



$$y = f(x_1 * w_1 + x_2 * w_2)$$

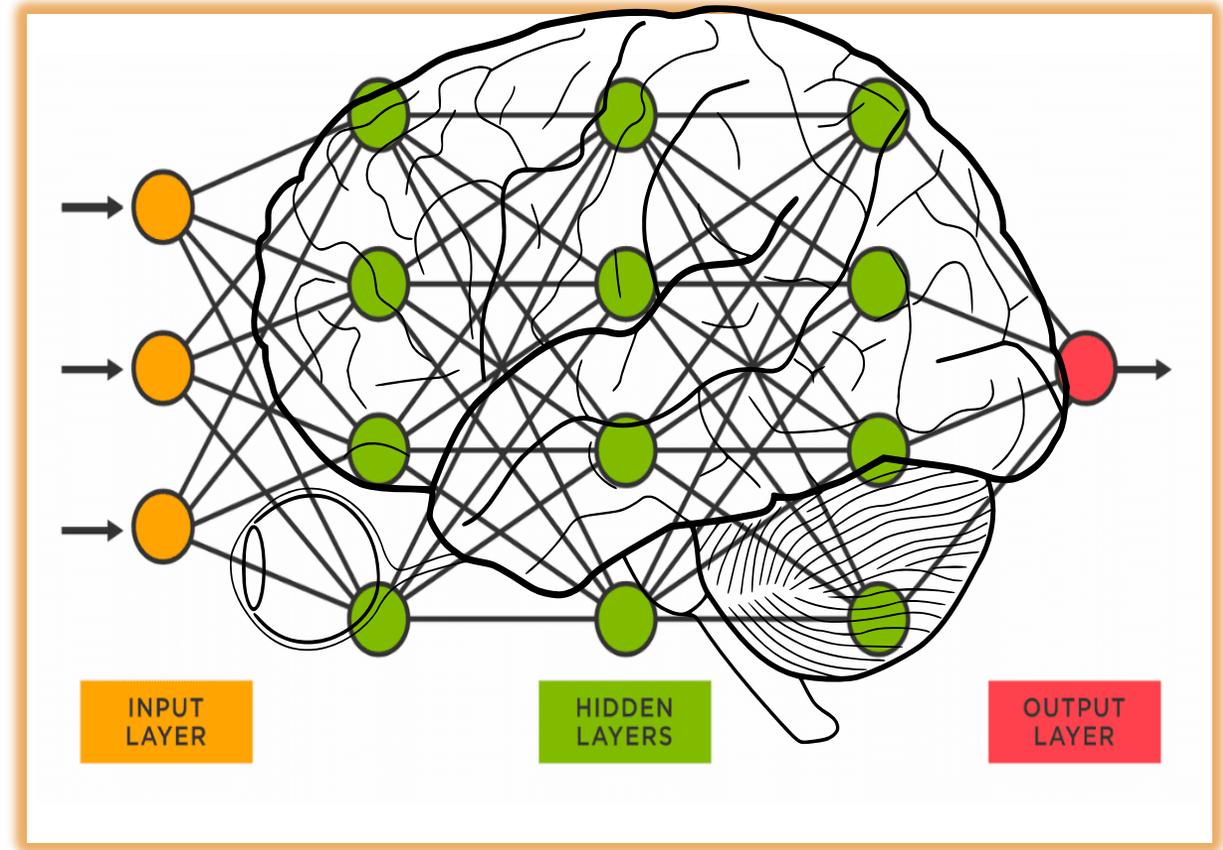
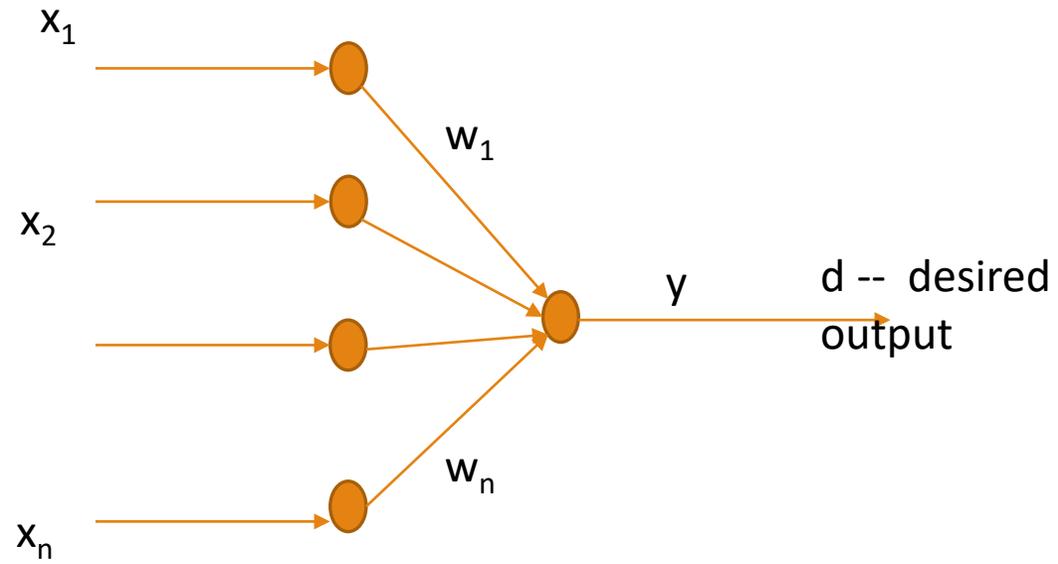
$$y = f(0 \times 0.5 + 0 \times 0.5) = f(0) = 0$$

$$y = f(0 \times 0.5 + 1 \times 0.5) = f(0.5) = 0$$

$$y = f(1 \times 0.5 + 0 \times 0.5) = f(0.5) = 0$$

$$y = f(1 \times 0.5 + 1 \times 0.5) = f(1) = 1$$

Perceptron



Machine Learning

Machine Learning

Machine learning covers algorithms that learn from and make predictions on data

Characteristics

- Generalization -- programming explicit algorithm is infeasible.
- Approximation is acceptable -- Pattern recognition, Face recognition, Handwriting recognition

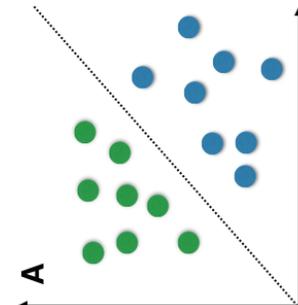
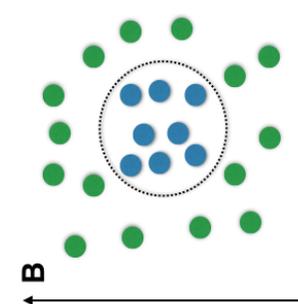
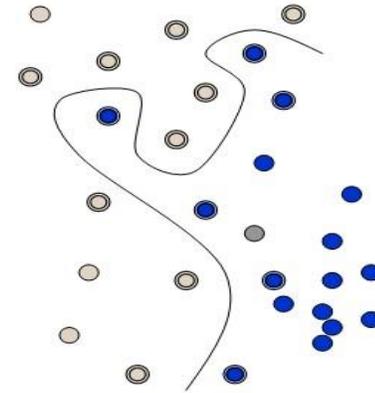
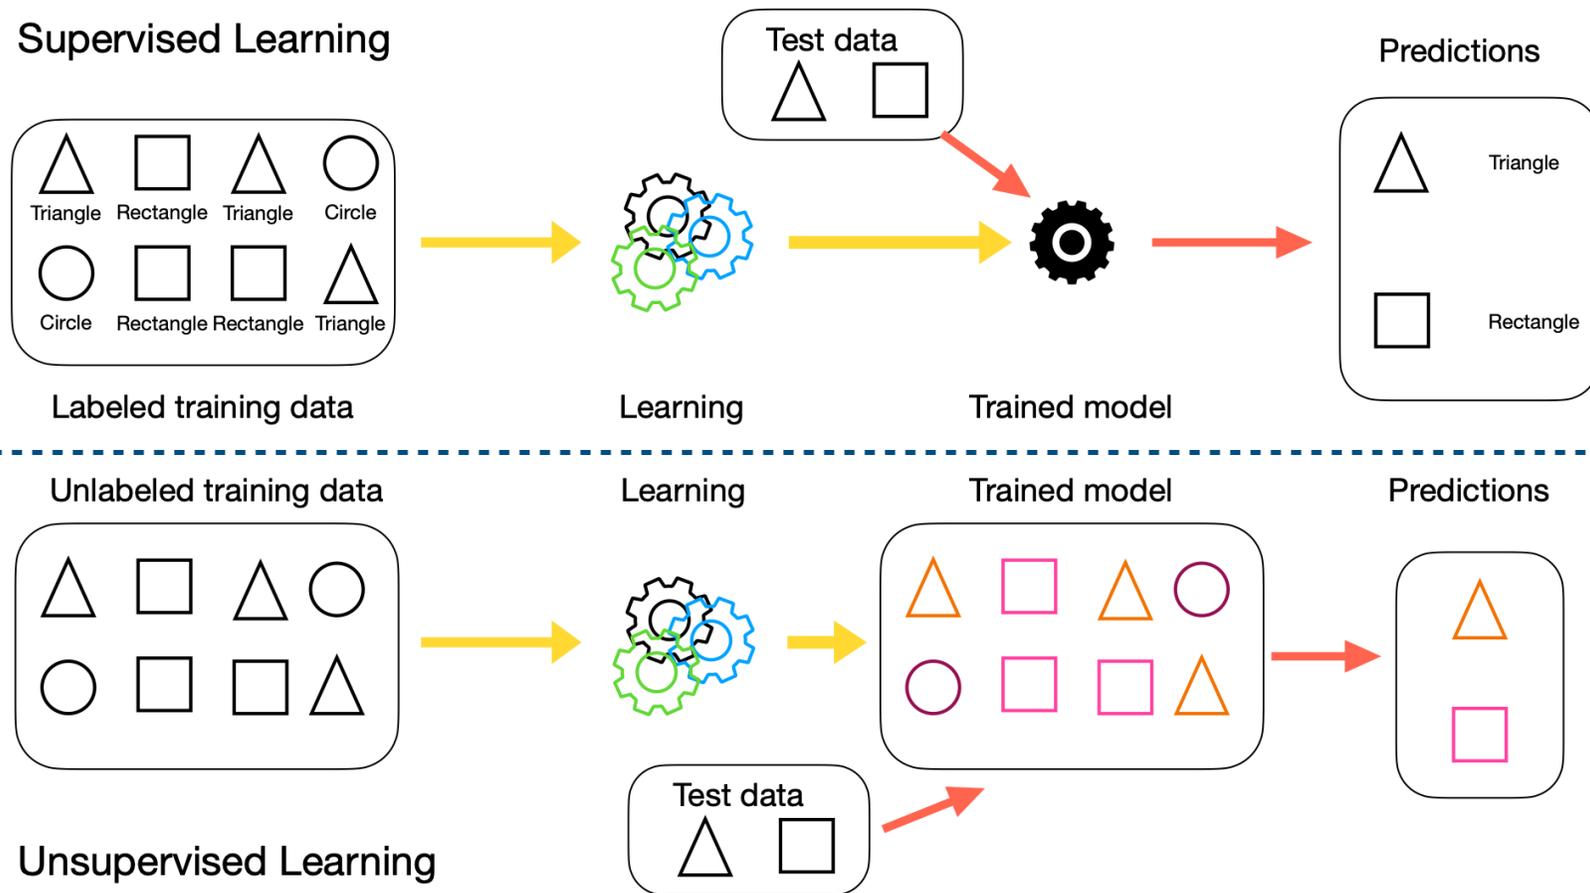


Illustration of a learning system

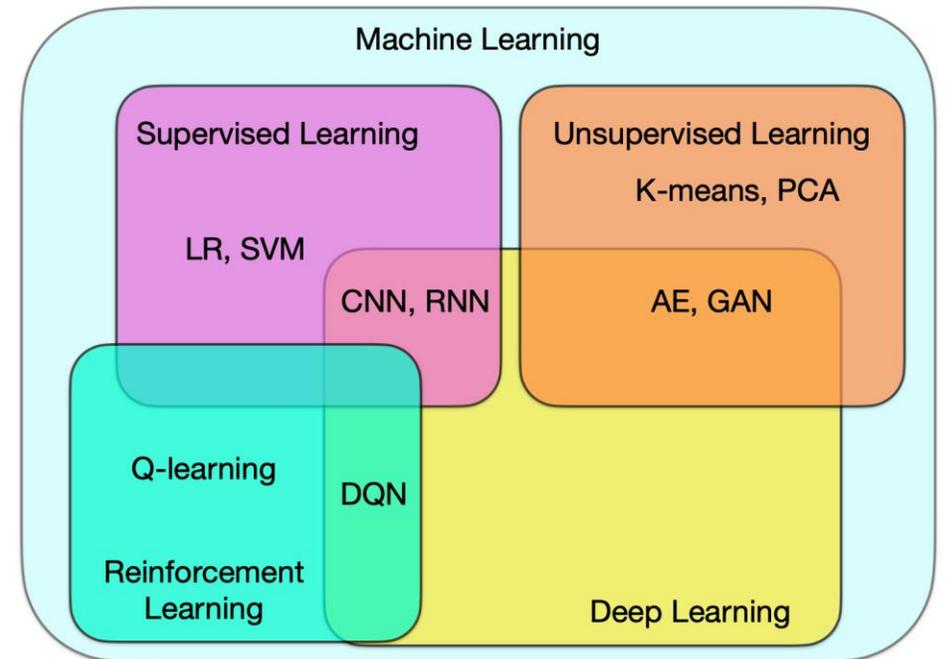


Deep Learning vs. Machine Learning

Machine learning focuses on the development of algorithms that enable computers to learn from and make decisions based on data.

Deep learning is a specialized subfield of machine learning that uses neural networks with many layers (hence “deep”) to model complex patterns in data.

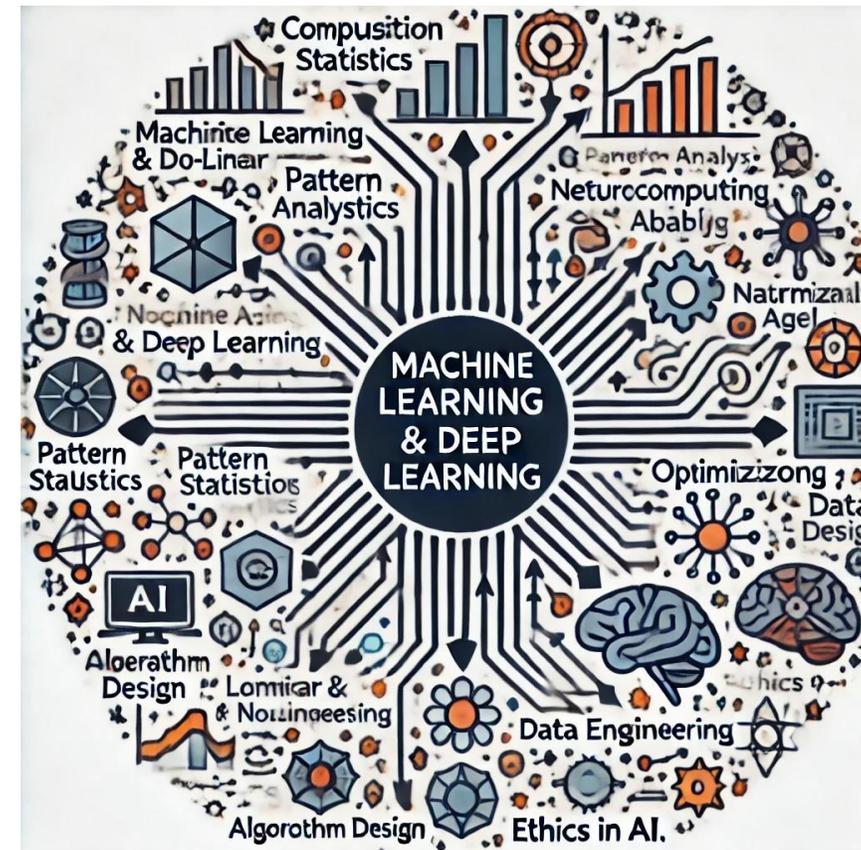
	Machine Learning	Deep Learning
Contemporary Usage	Traditional machine learning algorithms and simpler neural networks	Neural networks with deep architectures
Features	Feature engineering with manually extracted and selected features	Automatically extracted by the neural network
Applications	Well-suited to smaller datasets and problems where interpretability is important	Excels in handling large, complex datasets, especially in CV, NLP, and LLM



Component fields of Machine Learning

Areas include –

- Computational statistics
- Pattern analysis
- Neurocomputing
- Linear and non-linear algebra
- Optimization
- Algorithm design, etc.



How deep learning works

Multiple layers of interconnected nodes

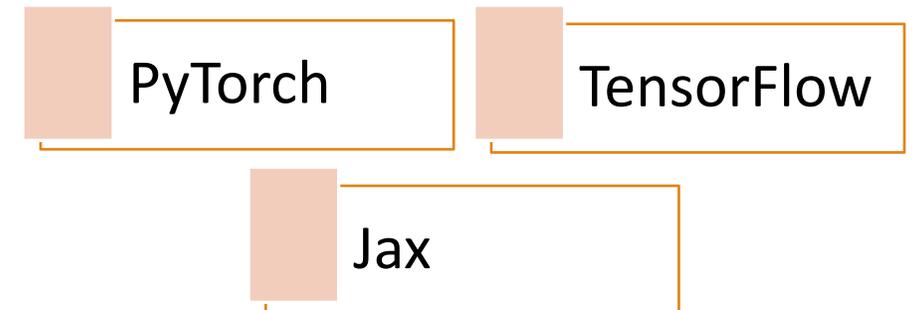
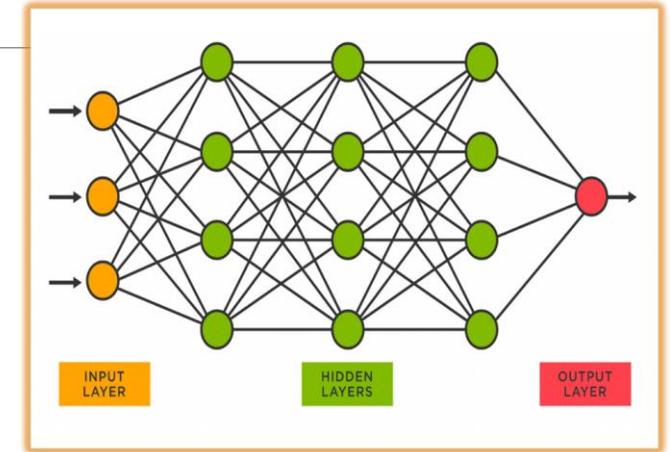
Backpropagation- Error is backpropagated to adjust weights to learn the patterns

Massive set of neurons – require massive resources

High performance graphical processing units – calculations in multiple cores with large memory

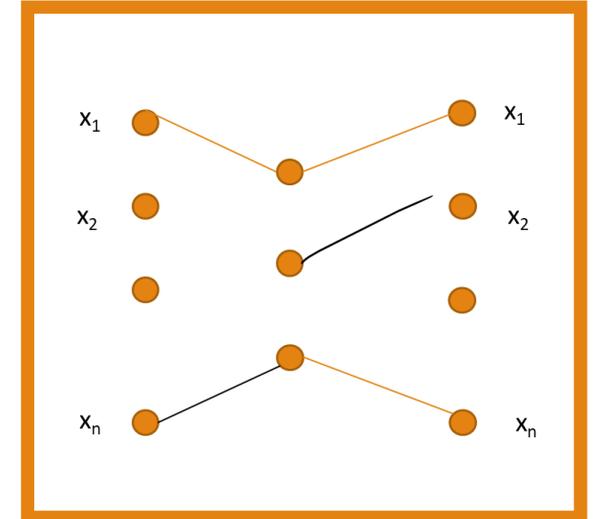
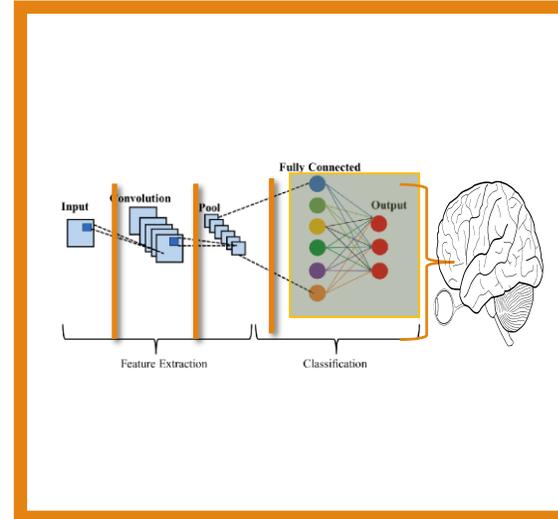
Distributed cloud computing is good

Software requirements– typically JAX, PyTorch, TensorFlow

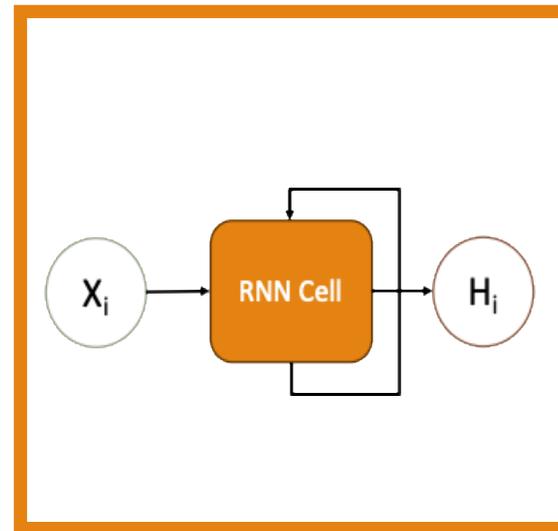


Types of deep learning -- 1

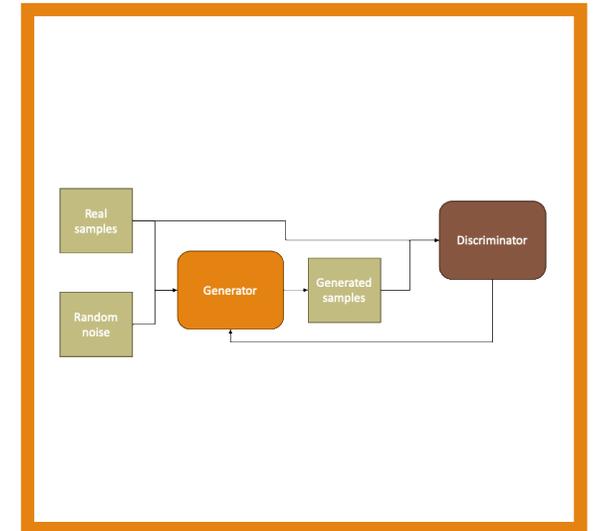
Convolutional neural networks (CNNs)



Recurrent Neural Networks (RNNs) – Uses a BPTT



Autoencoders and variational autoencoders



Generative Adversarial Networks

Convolutional Neural Networks

Used primarily in image processing and vision

At least three main types of layers: a convolutional layer, pooling layer and fully connected (FC) layer.

“convolution”—working and reworking the original input—detailed patterns can be discovered.

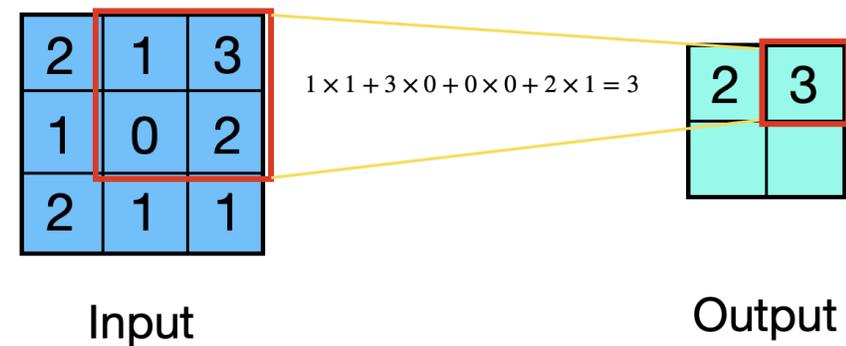
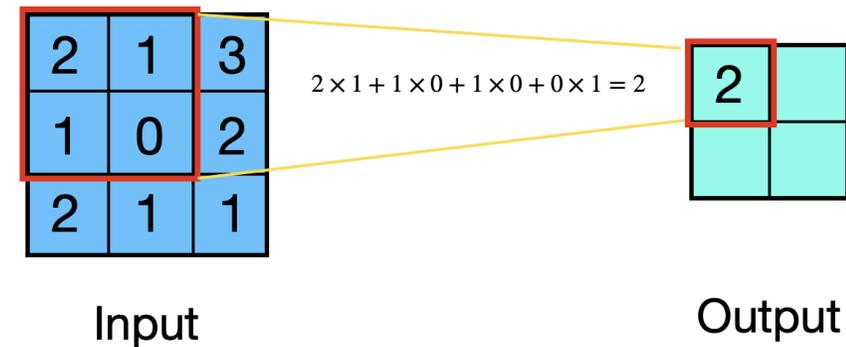
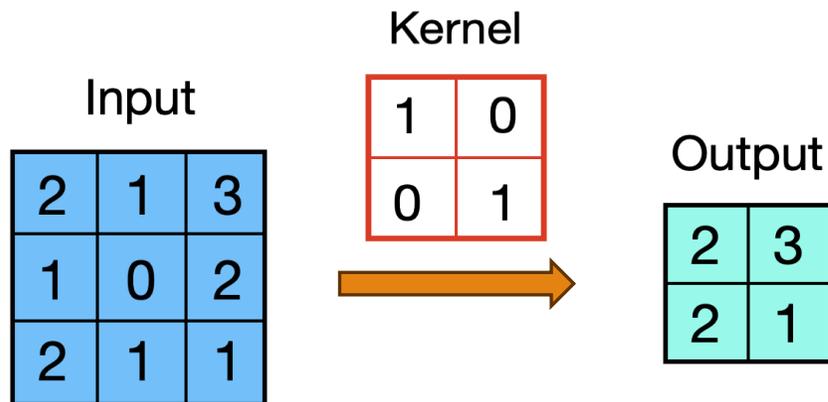
Earlier layers of a CNN detect simple features like colors and edges. As data moves through the layers, the network recognizes larger shapes and eventually identifies the object.

Convolutional Neural Networks

A simple example shows how CNN works:

Suppose we have a 3*3 matrix as input.

We apply a CNN layer with a 2*2 kernel, padding is 0, and stride length is 1



Recurrent Neural Networks

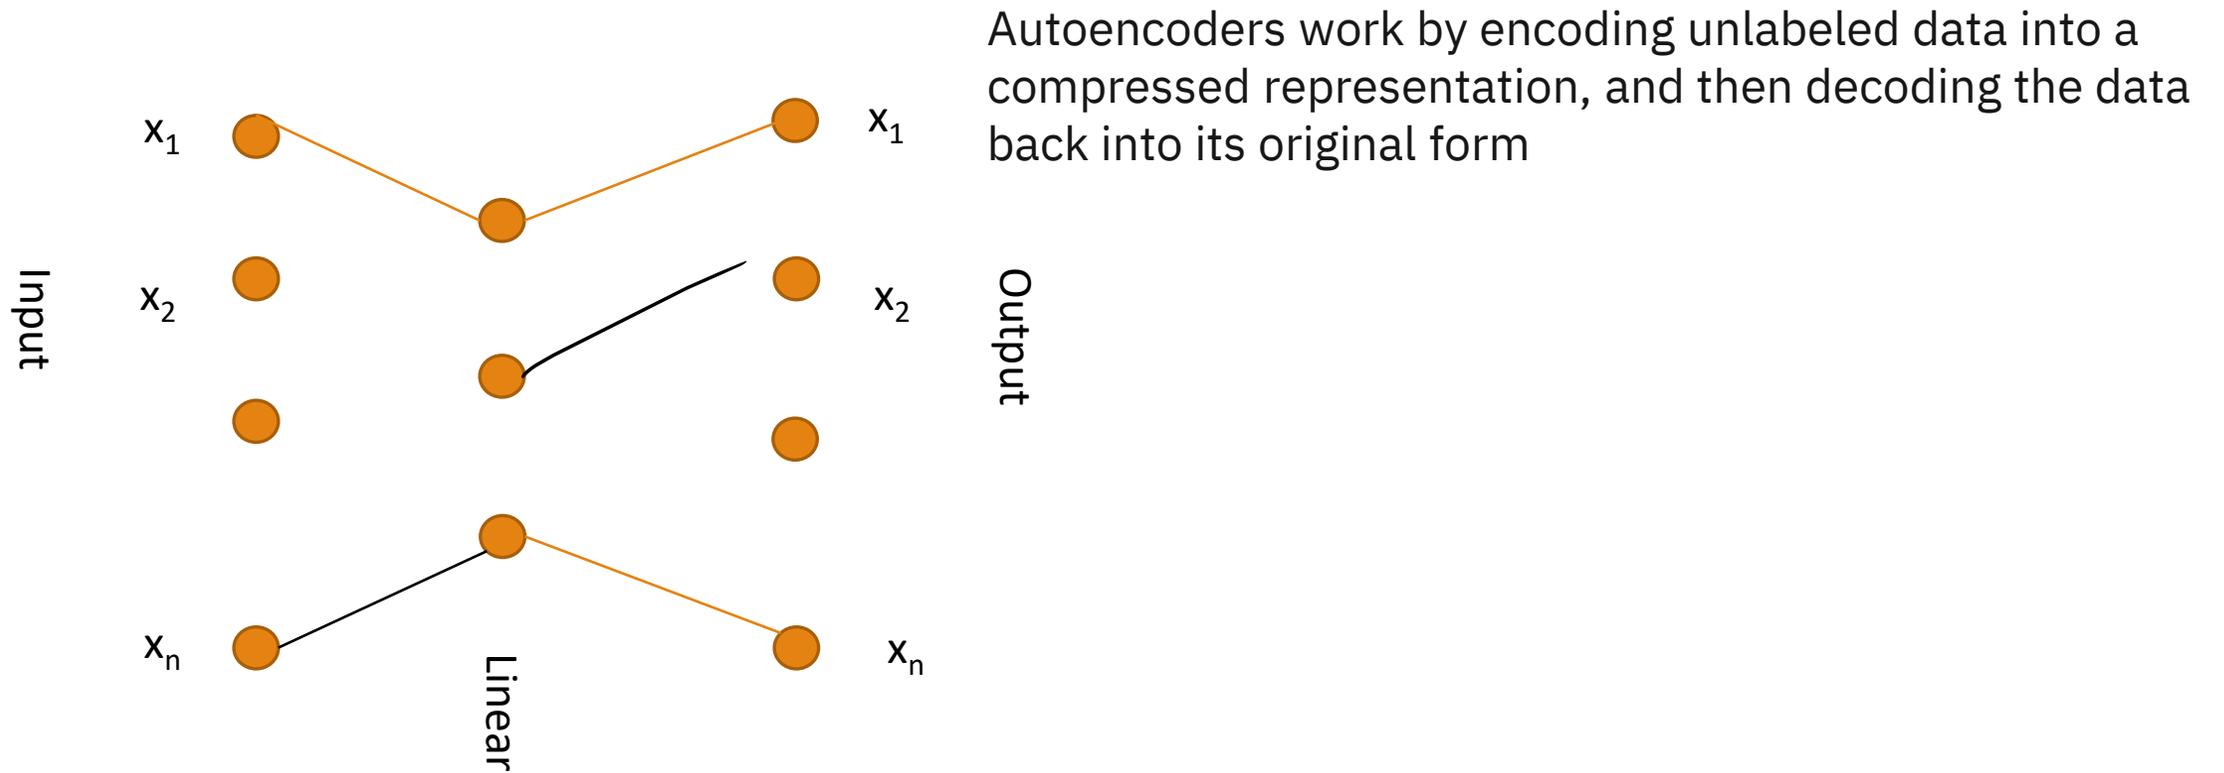
used in natural language and [speech recognition](#) applications as they use sequential or time-series data

“memory” as they take information from prior inputs to influence the current input and output

Uses Back Propagation through time

BPTT differs from the traditional approach in that BPTT sums errors at each time step,

Autoencoder



Generative Adversarial Networks

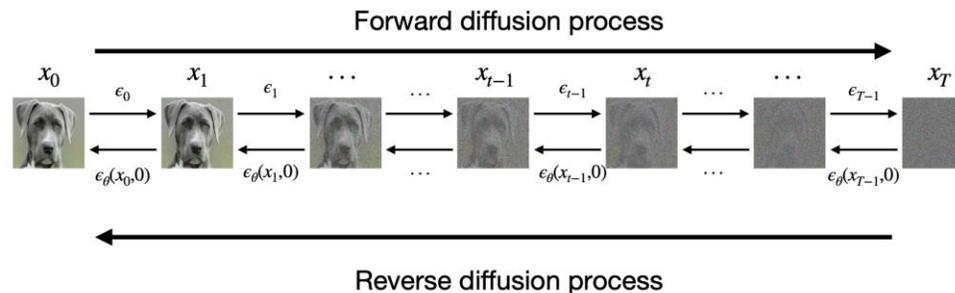
GANs

- The **generator** creates *something*: images, video or audio and then producing an output with a twist. For example, a horse can be transformed into a zebra with some degree of accuracy. The result depends on the input and how well-trained the layers are in the generative model for this use case.
- The **discriminator** is the adversary, where the *generative* result (fake image) is compared against the *real* images in the dataset. The discriminator tries to distinguish between the real and fake images, video or audio.

Types of deep learning – Diffusion models

1. Forward Diffusion Process

The model starts with a clean image (original image) and gradually adds noise to it over a series of time steps.

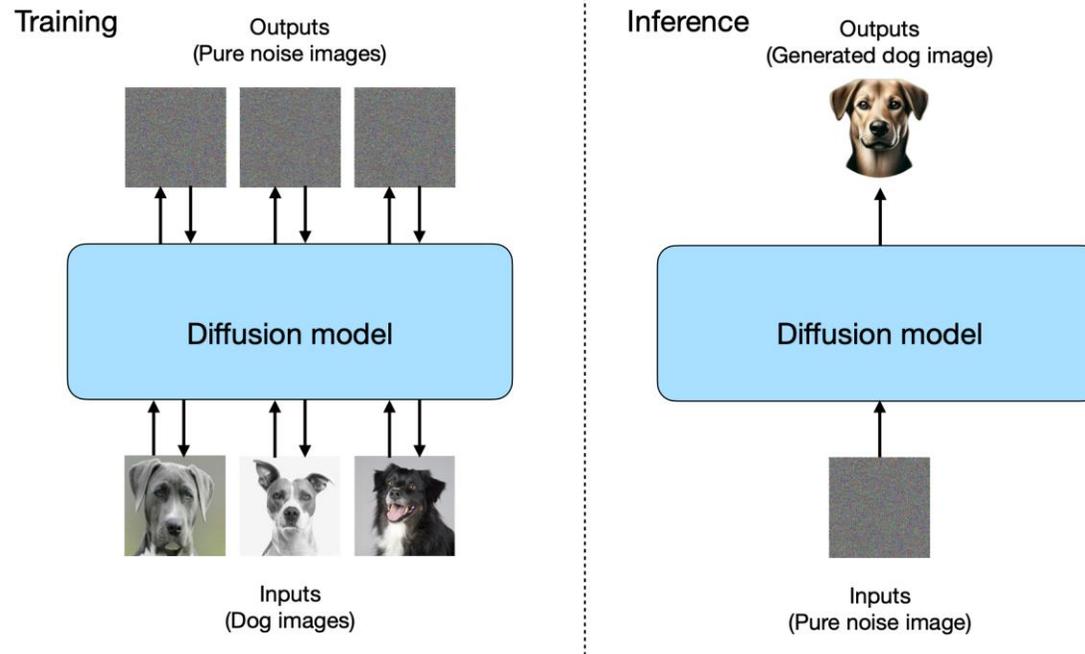


2. Reverse Diffusion Process

The goal of the model is to learn the reverse process: given a noisy image at any time step, the model learns to predict the slightly less noisy version of the image from the previous time step.

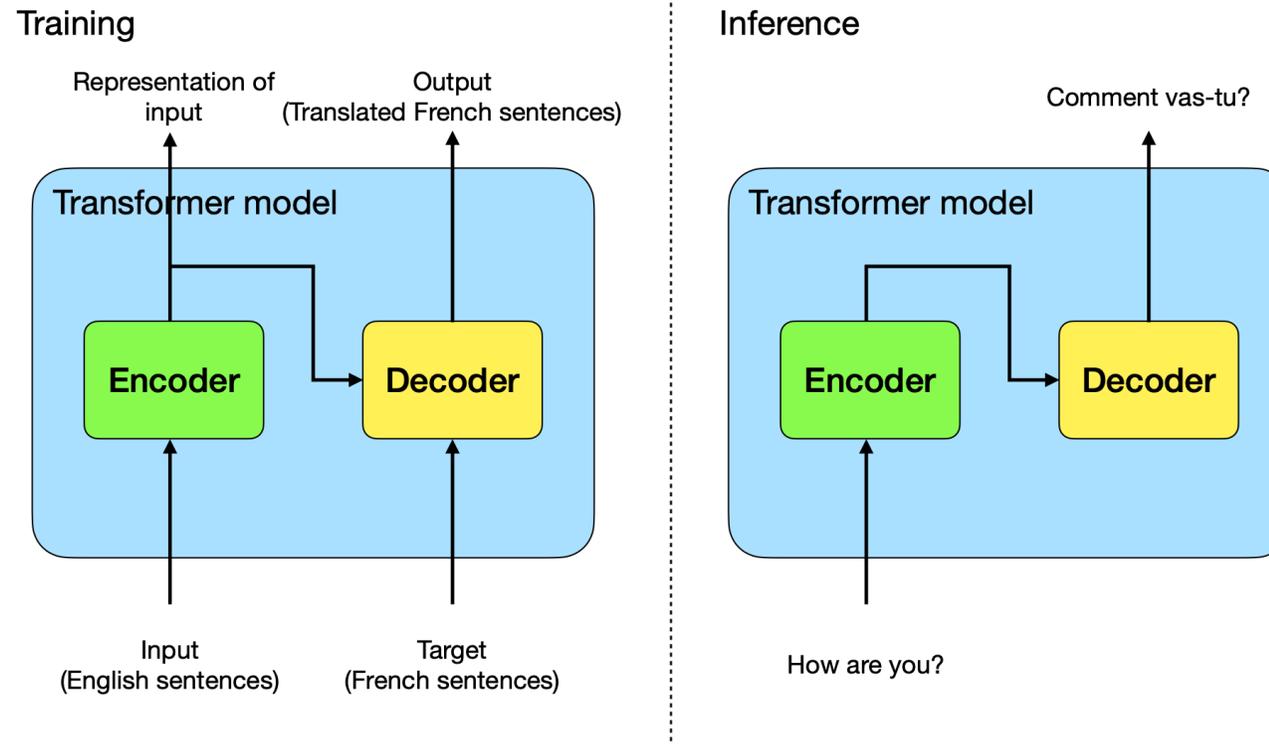
3. Image Generation

During inference, starting with pure noise, the trained model gradually removes the noise step by step, eventually reconstructing an image that resembles the learned data distribution.



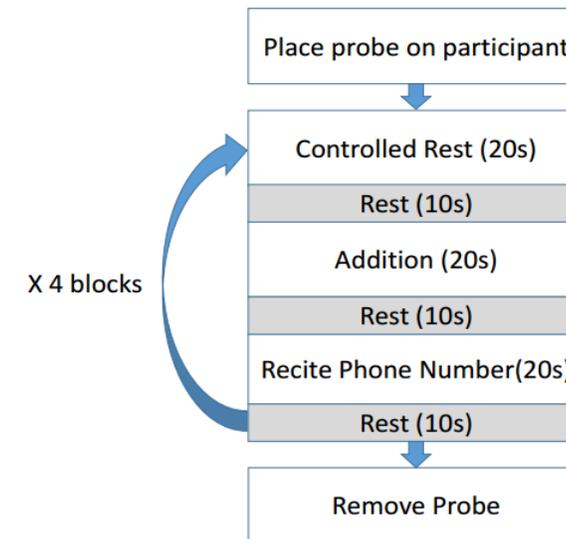
Types of deep learning – Transformer models

1. **Encoder:** Maps the input sequence to a continuous representation.
2. **Decoder:** Generates the output sequence from the continuous representation provided by the encoder.



Demo Video

Brain Activity



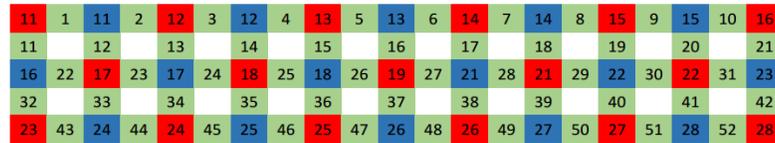
EXPERIMENT TASKS

- ❖ THREE TASKS CHOSEN BASED UPON A REVIEW OF PSYCHOLOGICAL AND NEUROSCIENCE LITERATURE TO PRODUCE CONSISTENT PATTERNS OF BRAIN ACTIVATION FOR THE LATER IDENTIFICATION OF SUBJECTS.

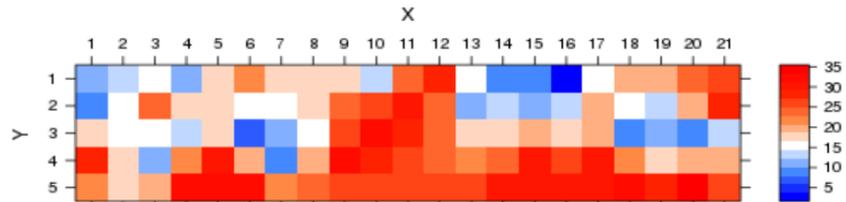
EXPERIMENT PROTOCOL

- ❖ 50 PARTICIPANTS (37 MALE, 13 FEMALE) FOR 4 IDENTICAL SESSIONS
- ❖ NEWLY PLACED FNIRS CAP ON THE SUBJECT AT THE BEGINNING OF EACH SESSION,
- ❖ PROBE CENTERED ON EACH PARTICIPANT'S FOREHEAD.

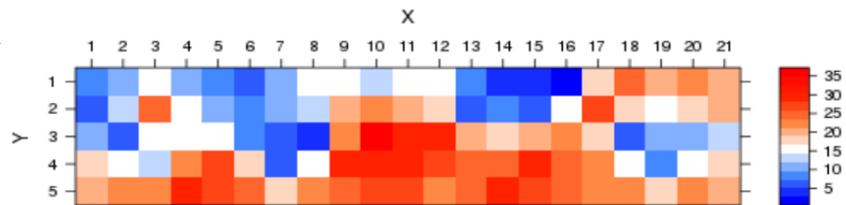
A. Serwadda, V. **Phoha**, S. Poudel, L. Hirshfield, D. Bandara, S. Bratt, *et al.*, "fNIRS: A New Modality for Brain Activity-Based Biometric Authentication," in *The IEEE Seventh International Conference on Biometrics: Theory, Applications and Systems (BTAS 2015)*, Washington, DC, 2015.



Legend: Channels (green), Detectors (blue), Sources (red).
 (a) Locations of the 52 channels relative to each other and relative to the two eyes (approximately).



(b) Color map showing how the discriminative power of the mean of the first ten points varied across the 52 channels (or brain regions). The red regions cover a large area relative to the blue regions, meaning that there is a good number of channels for which this feature was highly discriminative.



(c) Color map showing how the discriminative power of the standard deviation of the last ten points varied across the 52 channels (or brain regions). The blue regions cover a larger area relative to their coverage area in Figure 3(b), meaning that this feature was not as discriminative as the feature represented in Figure 3(b).

Classifier	Mean EER		% Change in EER
	All Channels	Best Channels	
SVM	0.043	0.036	17.1
Naïve Bayes	0.063	0.046	28.3

Conclusion

- ❖ While there is still a need to evaluate fNIRS for a wider range of mental tasks, these results suggest that fNIRS holds promise as an AA modality.
- ❖ Major part of our ongoing research is to carry out analysis on a wider variety of tasks and to more rigorously evaluate the dependence of authentication performance on specific brain regions.

A. Serwadda, V. Phoha, S. Poudel, L. Hirshfield, D. Bandara, S. Bratt, *et al.*, "fNIRS: A New Modality for Brain Activity-Based Biometric Authentication," in *The IEEE Seventh International Conference on Biometrics: Theory, Applications and Systems (BTAS 2015)*, Washington, DC, 2015.

LLMs Current and Future



Near future -Trends

- 1.Small language-based Models
- 2.Development of one-shot and few shot learning
- 3.Complex non-hierarchical Reasoning
- 3.Economic Revolution;
- 4.Lot of Funding and Developer Interest

- Modeling and translating languages
- One-shot and few shot learning
- Explainable AI
- Summarizing Generating and Classifying text
- Image analysis and content recommendation
- Analyzing sentiment in text and images



Future - Trends

- 1.Contextual and Personalized Content Generation
- 2.Advanced autonomy
- 2.Advanced Conversational Features
- 3.Domain specific and complex Solutions



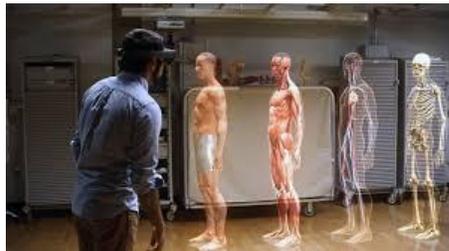
Near Future (3 to 5 years)

Philosophical questions

IoT includes wearables; Cloud

Beyond 5 years
Deep learning directed applications

Augmented Reality By definition
includes Network centric and IoT



Current Needs



Faster Solutions to hard problems

Diseases Hunger Energy solved

Less need of physical bodies— people will be minds –augmented by machines

Can we achieve

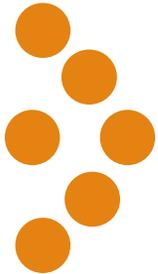
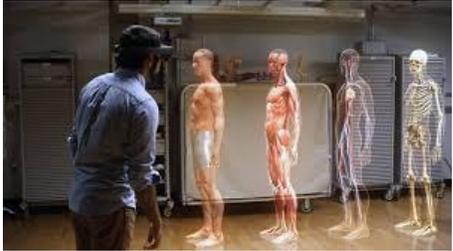
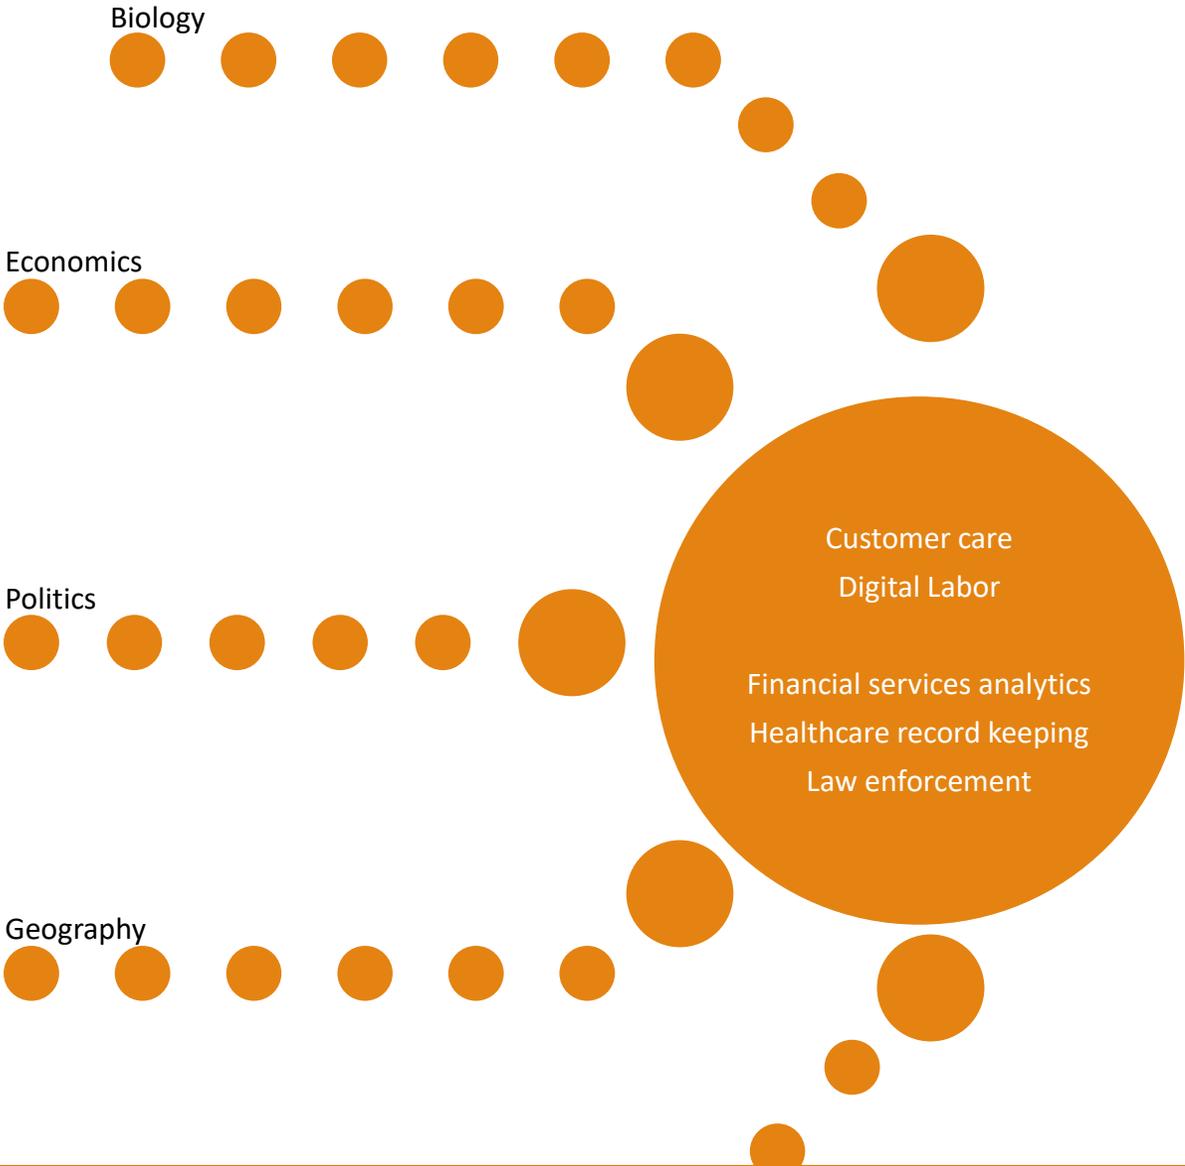
Empathy Love Creativity

Ride the train and not jump in front of it

Machines become
more intelligent
than humans

Singularity -- a hypothetical future where technology growth is out of control and irreversible.

Convergence



Applied sciences



Thank you

?