



Artificial Intelligence

Dr. Felix Gonda

Associate Prof. Computer Science
uojdeeplearning@gmail.com

Convolutional Neural Network (CNN)



Backpropagation

$$\underbrace{\frac{\partial L}{\partial w} = \frac{\partial L}{\partial \hat{y}} \cdot \frac{\partial \hat{y}}{\partial z} \cdot \frac{\partial z}{\partial w}}_{\text{Chain Rule}} \rightarrow \frac{\partial L}{\partial w}$$

Weight Update

Precision Agriculture
Crop Health Mapping

Healthy
Moderate
Stressed

Infrastructure Networks
Stronger Communities
South Sudan

Deep Learning

Computer Vision

TensorFlow & Keras

South Sudan
Builds its
Future

SKILLED
YOUTH
STRONGER
SOUTH SUDAN

- ✓ Better Crops
- ✓ Healthy Communities
- ✓ Sustainable Future



PROMPT: A lone cowboy rides his horse across an open plain at beautiful sunset, soft light, warm colors

Source: <https://deepmind.google/technologies/veo/>



Prompt: A stylish woman walks down a Tokyo street filled with warm glowing neon and animated city signage. She wears a black leather jacket, a long red dress, and black boots, and carries a black purse. She wears sunglasses and red lipstick. She walks confidently and casually. The street is damp and reflective, creating a mirror effect of the colorful lights. Many pedestrians walk about.

Source: <https://openai.com/index/sora/>



Prompt:

Write a welcome message to students taking my Deep Learning course at University of Juba.

Response:

Dear Students,
I'm excited to welcome you to the Deep Learning course! This course will provide you with a solid foundation in the theory and practice of deep learning, a powerful tool for solving complex problems in various fields.


Gemini

<https://gemini.google.com/>



ChatGPT

<https://chatgpt.com/>




Meta AI

<https://www.meta.ai/>

Part I: The AI Paradigm Shift

State-of-the-Art Frontiers, Historical Foundations, and the Need for Representation Learning

The Frontier of Modern AI




Emergent Multimodal Capabilities

Modern frontier models seamlessly integrate vision, speech, code, and formal logic reasoning in real-time.



Scientific Breakthroughs



Deep Learning has solved 50-year challenges in biology (e.g., AlphaFold 3 for bio-molecular structures).



Generative Intelligence


High-fidelity diffusion models & autoregressive transformers create complex software and synthetic media.

Traditional Software (Heuristics)



THE CORE SHIFT


Deep Learning (Representation Learning)



Historical Evolution of Deep Learning

- Popularized by Rumelhart, Hinton, and Williams in 1986, **backpropagation** revolutionized artificial intelligence by applying the calculus chain rule to efficiently calculate gradients across multi-layered networks, solving the long-standing credit assignment problem and enabling the training of deep neural architectures.

- Winning the 2012 ImageNet competition by a massive margin, AlexNet demonstrated the immense power of deep Convolutional Neural Networks trained on GPUs, single-handedly reviving neural network research and catalyzing the modern era of deep learning.




- Proposed by Alan Turing in 1950, the **Turing Test** is a foundational benchmark evaluating whether an artificial intelligence can exhibit conversational behavior indistinguishable from that of a human.
- Introduced by Frank Rosenblatt in 1958, the **Perceptron** served as the first algorithmic and hardware realization of an artificial neuron, establishing the fundamental building block for modern deep neural networks.

- Introduced by Yann LeCun et al. in 1998, **LeNet-5** established the foundational architecture of Convolutional Neural Networks by combining local receptive fields, weight sharing, and spatial subsampling to automate visual feature extraction, setting the blueprint for modern computer vision.

The Hierarchy of Modern Intelligence: AI, ML & DL Explained

Artificial Intelligence (AI) is the broad field of simulating human intelligence. Within AI sits Machine Learning (ML), which focuses on algorithms that learn from data. Deep Learning (DL) is a further specialized subset of ML that uses multi-layered neural networks to automatically learn complex patterns from raw data.




The Deciding Factor: How They Process Data

Traditional ML: Manual Feature Engineering



A human expert must manually select and create relevant data features for the algorithm to process.

Deep Learning: Automatic Feature Extraction



The neural network learns the most important features directly from raw data during training.




Practical Trade-Offs

Criterion	Traditional Machine Learning	Deep Learning
Data Needs	Effective with small to medium datasets.	Requires very large datasets to perform well.
Hardware	Can run on standard CPUs.	Needs high-performance GPUs or TPUs.
Interpretability	High ("White Box" - decisions are traceable).	Low ("Black Box" - reasoning is opaque).

Sorting Fruits

An example problem

Suppose we have a business that receive baskets of mixed fruits (apple, bananas, and oranges) and we want to sort them into different baskets. Apples should go into one basket, bananas into another basket, and oranges into a separate basket.




Sorting Fruits

Manual Solution

A manual solution involves hiring employees to sort the fruits by inspecting their labels or sort them from memory.

- ☒ Works well on small number of fruits.
- ☐ Throughput relies on speed of workforce (more workers required at large scale)
- ☐ Potential of human errors.




Sorting Fruits

Machine Learning Solution

An ML solution resolves the problem of missing labels.

- ✓ First, a human expert creates a blueprint that describes the features and attributes that characterize each fruit. (E.g: size, color, shape, etc).
- ✓ Then the ML algorithm is trained to learn the features of each fruit.
- ✓ The trained algorithm is then used to automatically classify the fruits.
- The ML model, however, cannot classify a fruit it has never seen before.



Feature Extraction

A human expert defines a blueprint containing attributes and features of each fruit.

Features:

- Size
- Shape
- Texture
- Color

Classification

The algorithm is trained to learn the features of each fruit so it can automatically classify fruits.


Sorting Fruits

Deep Learning Solution

- ✓ A deep learning algorithm can sort any fruit by learning features of the fruits from images.
- ✓ The algorithm consists of multiple layers of a neural network, which automatically extract features.
- ✓ Similar to how our brain works to classify objects through many layers of neurons.
- ✓ The DL algorithm will categorized the fruits based on their statistical similarities.
- ❑ But it requires lots of data and significant computational power.



Input: A vast amount of image data of many kinds of fruits.




Why Deep Learning?

Hand design features are time consuming to create, brittle, and not scalable in practice.

So, can we learn the underlying features directly from the data?

Low Level Features



These are basic elements like **edges, lines, and corners**. A deep learning algorithm can learn to identify these simple patterns in the early layers.

Mid-level Features



As the network progresses, it starts to recognize more complex patterns, such as facial components like **eyes, noses, and ears**. These mid-level features are built upon the low-level features.


High Level Features



In the deeper layers, the network learns to represent the overall **facial structure, including the spatial arrangement of facial components** and identity-specific characteristics. These high-level features are crucial for distinguishing between different individuals.

Scaling Laws

Performance vs Data



Unlike traditional models whose accuracy saturates as data grows, Deep Neural Networks obey empirical Power Laws: performance scales predictably with compute, parameters, and training data volume.

Factors

What drives DL Growth?

Availability of data

Innovations in cloud storage have made it possible for governments, research institutes, and businesses to unlock large datasets for training algorithms.

kaggle

Google

VisualData

UC Irvine
Machine Learning
Repository



OpenML

DATA
HUB

DATA

HealthData

The NLP Index


Advancements in Hardware

New breed of processors, such as the graphic processing unit (GPU), have made training of deep learning algorithms easier.



Advancements in Deep Learning


Continuous innovations in deep learning algorithms in image and video analysis, text analysis, audio processing enables rapid development of AI solutions.




Part II: Course Logistics & Roadmap

Syllabus Structure, Grading Breakdown, Environment Setups, and Office Hours


Optional Textbooks



<https://www.deeplearningbook.org/>



<https://mml-book.github.io/book/mml-book.pdf>



<https://d2l.ai/>

Course Website

<https://uojai.github.io/deeplearning>

References

- [Turing 1950] Turing, A. M. "Computing Machinery and Intelligence." *Mind*, 59(236), 433-460.
- [Rosenblatt 1958] Rosenblatt, F. "The Perceptron: A Probabilistic Model for Information Storage and Organization." *Psychological Review*.
- [Rumelhart et al. 1986] Rumelhart, D. E., Hinton, G. E., & Williams, R. J. "Learning representations by back-propagating errors." *Nature*.
- [LeCun et al. 1989] LeCun, Y., et al. "Backpropagation applied to handwritten zip code recognition." *Neural Computation*.
- [Vaswani et al. 2017] Vaswani, A., et al. "Attention Is All You Need." *Advances in Neural Information Processing Systems (NeurIPS)*
- [MIT/Stanford OpenCourseWare] Stanford CS231n: CNNs for Visual Recognition; MIT 6.S191: Introduction to Deep Learning.