Introduction to Generative AI & Fundamentals of ML, DL, LLM

Dr Eniafe Ayetiran

July 2025

Outline

What is Generative AI?

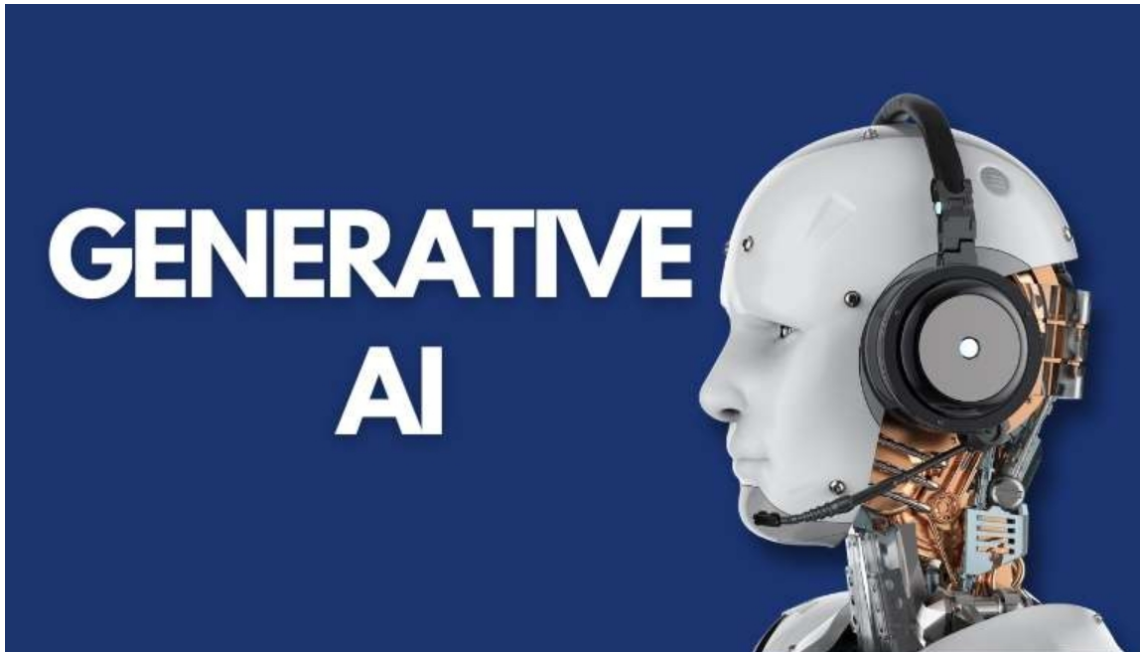

Figure: Generative AI

- **Definition**: AI systems that create new content (text, images, music, etc.) by learning patterns from data.

What is Generative AI?

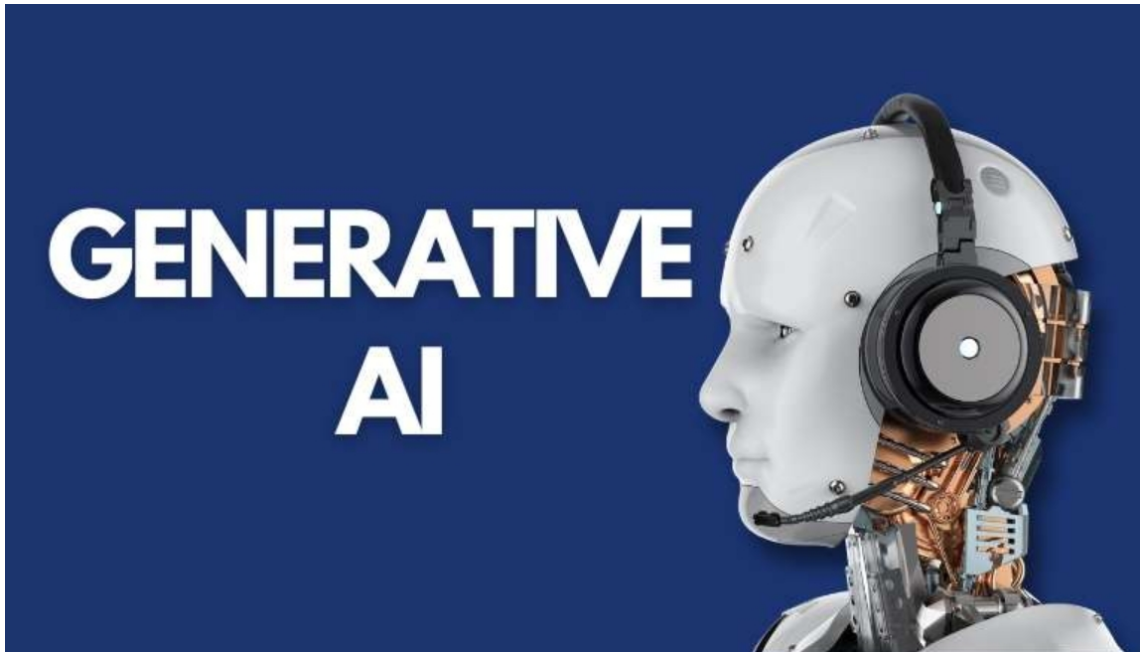

Figure: Generative AI

- **Definition**: AI systems that create new content (text, images, music, etc.) by learning patterns from data.
- **Key Idea**: Generates outputs that resemble training data but are novel.

What is Generative AI?

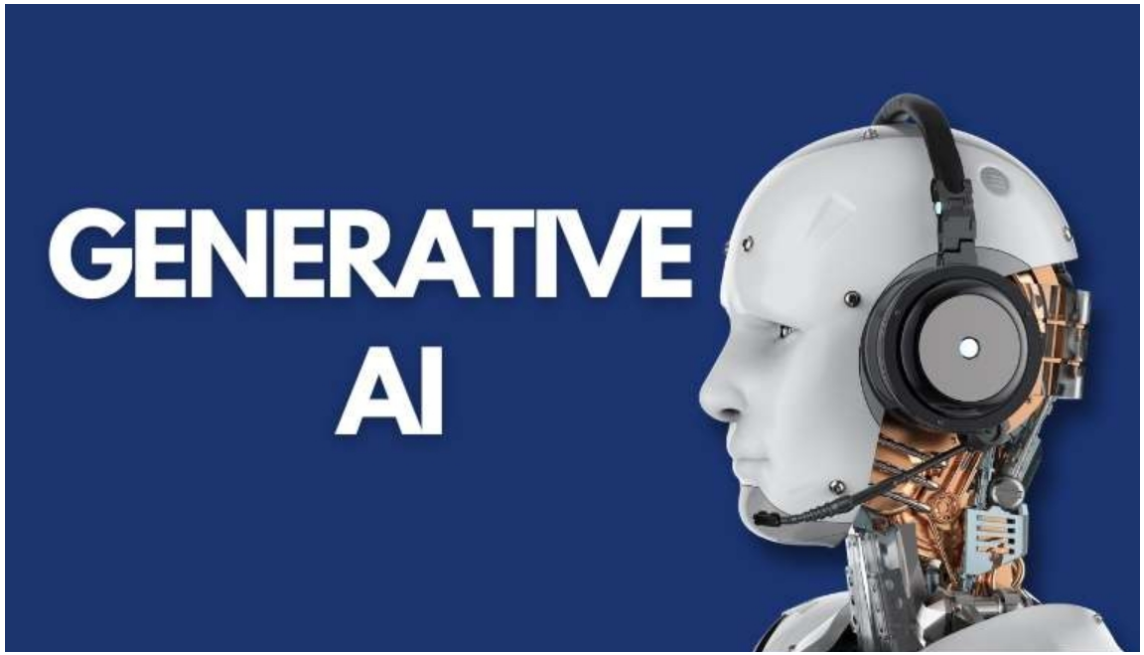

Figure: Generative AI

- **Definition**: AI systems that create new content (text, images, music, etc.) by learning patterns from data.
- **Key Idea**: Generates outputs that resemble training data but are novel.
- **Examples**:
  - Text generation (e.g., ChatGPT).
  - Image synthesis (e.g., DALL · E).
  - Music composition (e.g., Jukebox).

Why Generative AI Matters

- **Creativity**: Enables machines to produce human-like content.

Why Generative AI Matters

- **Creativity**: Enables machines to produce human-like content.
- **Applications**:
  - Content creation (writing, art).
  - Data augmentation (synthetic data).
  - Personalization (e.g., recommendation systems).

Why Generative AI Matters

- **Creativity**: Enables machines to produce human-like content.
- **Applications**:
  - Content creation (writing, art).
  - Data augmentation (synthetic data).
  - Personalization (e.g., recommendation systems).
- **Impact**: Transforming industries like entertainment, healthcare, and education.

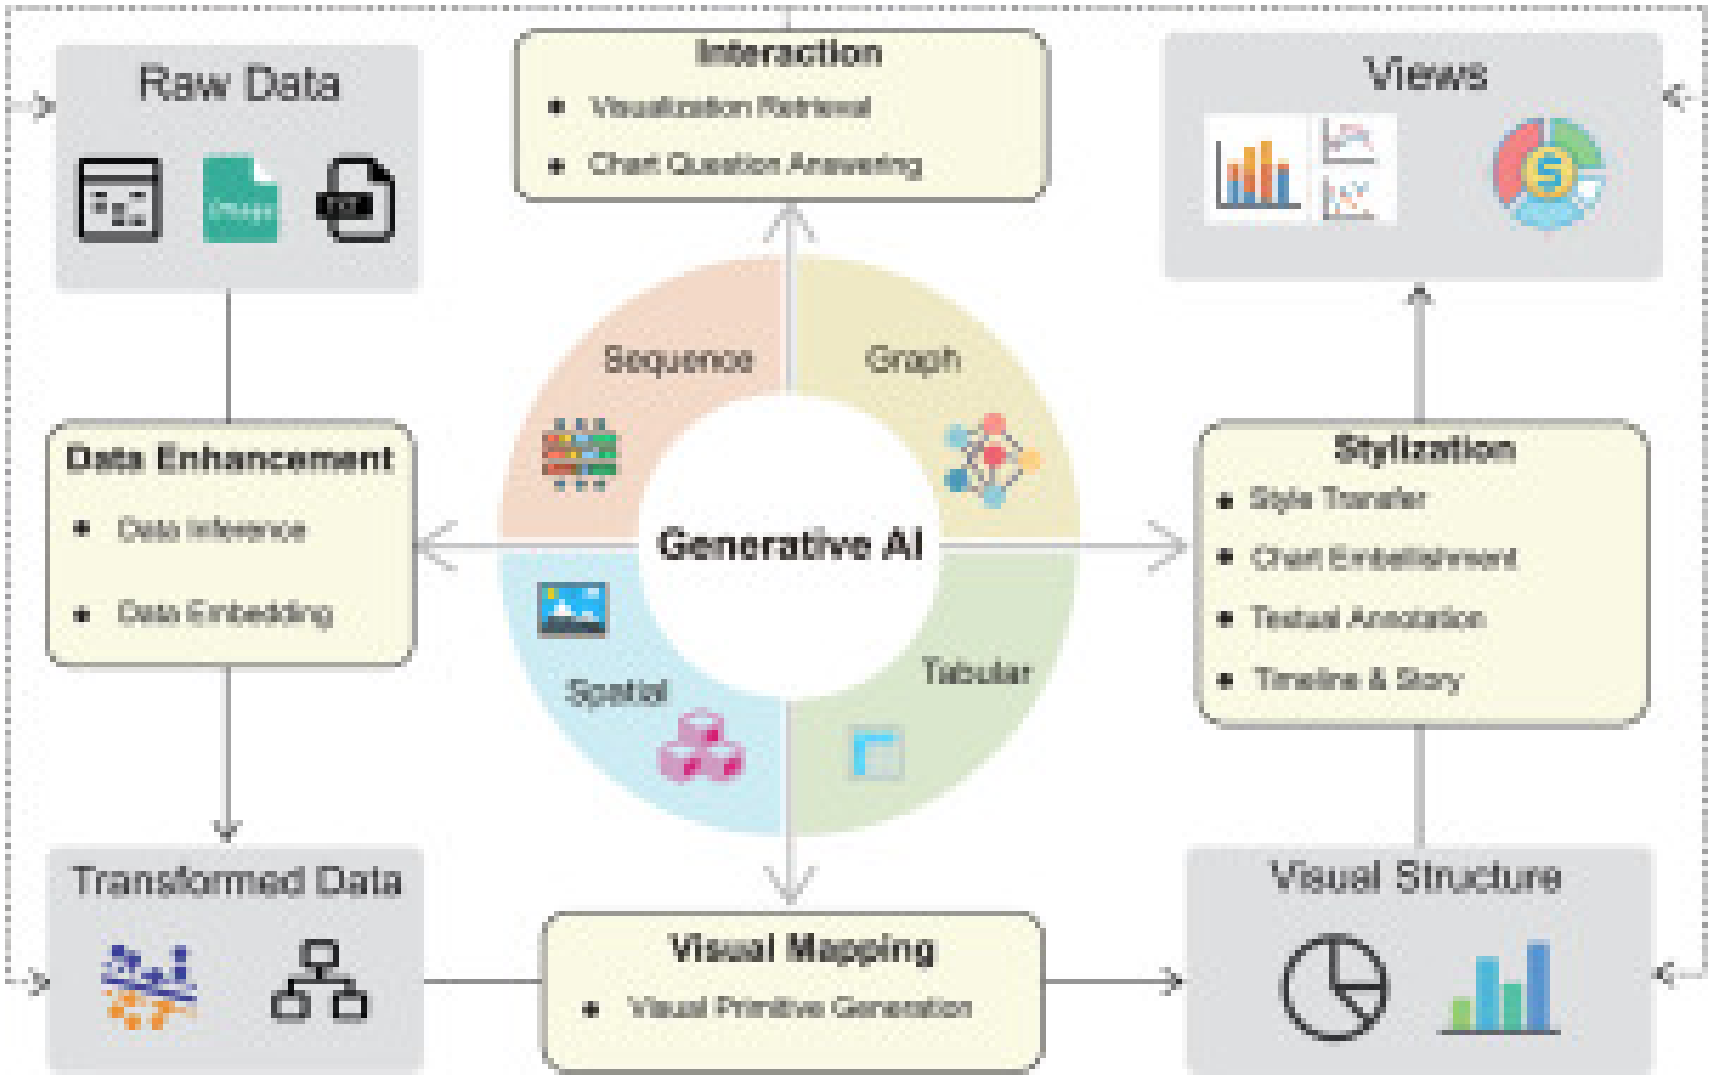

Figure: Generative AI Visualization

What is Machine Learning?

- **Definition**: Algorithms that learn patterns from data to make predictions or decisions.

What is Machine Learning?

- **Definition**: Algorithms that learn patterns from data to make predictions or decisions.
- **Types of ML**:
  - Supervised: Learn from labeled data (e.g., predicting house prices).
  - Unsupervised: Find patterns in unlabeled data (e.g., clustering).
  - Reinforcement: Learn through rewards (e.g., game AI).

What is Machine Learning?

- **Definition**: Algorithms that learn patterns from data to make predictions or decisions.
- **Types of ML**:
  - Supervised: Learn from labeled data (e.g., predicting house prices).
  - Unsupervised: Find patterns in unlabeled data (e.g., clustering).
  - Reinforcement: Learn through rewards (e.g., game AI).
- **Connection to Generative AI**: Generative models often use unsupervised or semi-supervised learning.

Key ML Concepts

- **Data**: Training, validation, test sets.

Key ML Concepts

- **Data**: Training, validation, test sets.
- **Features**: Input variables (e.g., pixel values, word counts).

Key ML Concepts

- **Data**: Training, validation, test sets.
- **Features**: Input variables (e.g., pixel values, word counts).
- **Model**: Mathematical function mapping inputs to outputs.

Key ML Concepts

- **Data**: Training, validation, test sets.
- **Features**: Input variables (e.g., pixel values, word counts).
- **Model**: Mathematical function mapping inputs to outputs.
- **Loss Function**: Measures error (e.g., mean squared error).

Key ML Concepts

- **Data**: Training, validation, test sets.
- **Features**: Input variables (e.g., pixel values, word counts).
- **Model**: Mathematical function mapping inputs to outputs.
- **Loss Function**: Measures error (e.g., mean squared error).
- **Optimization**: Minimize loss (e.g., gradient descent).

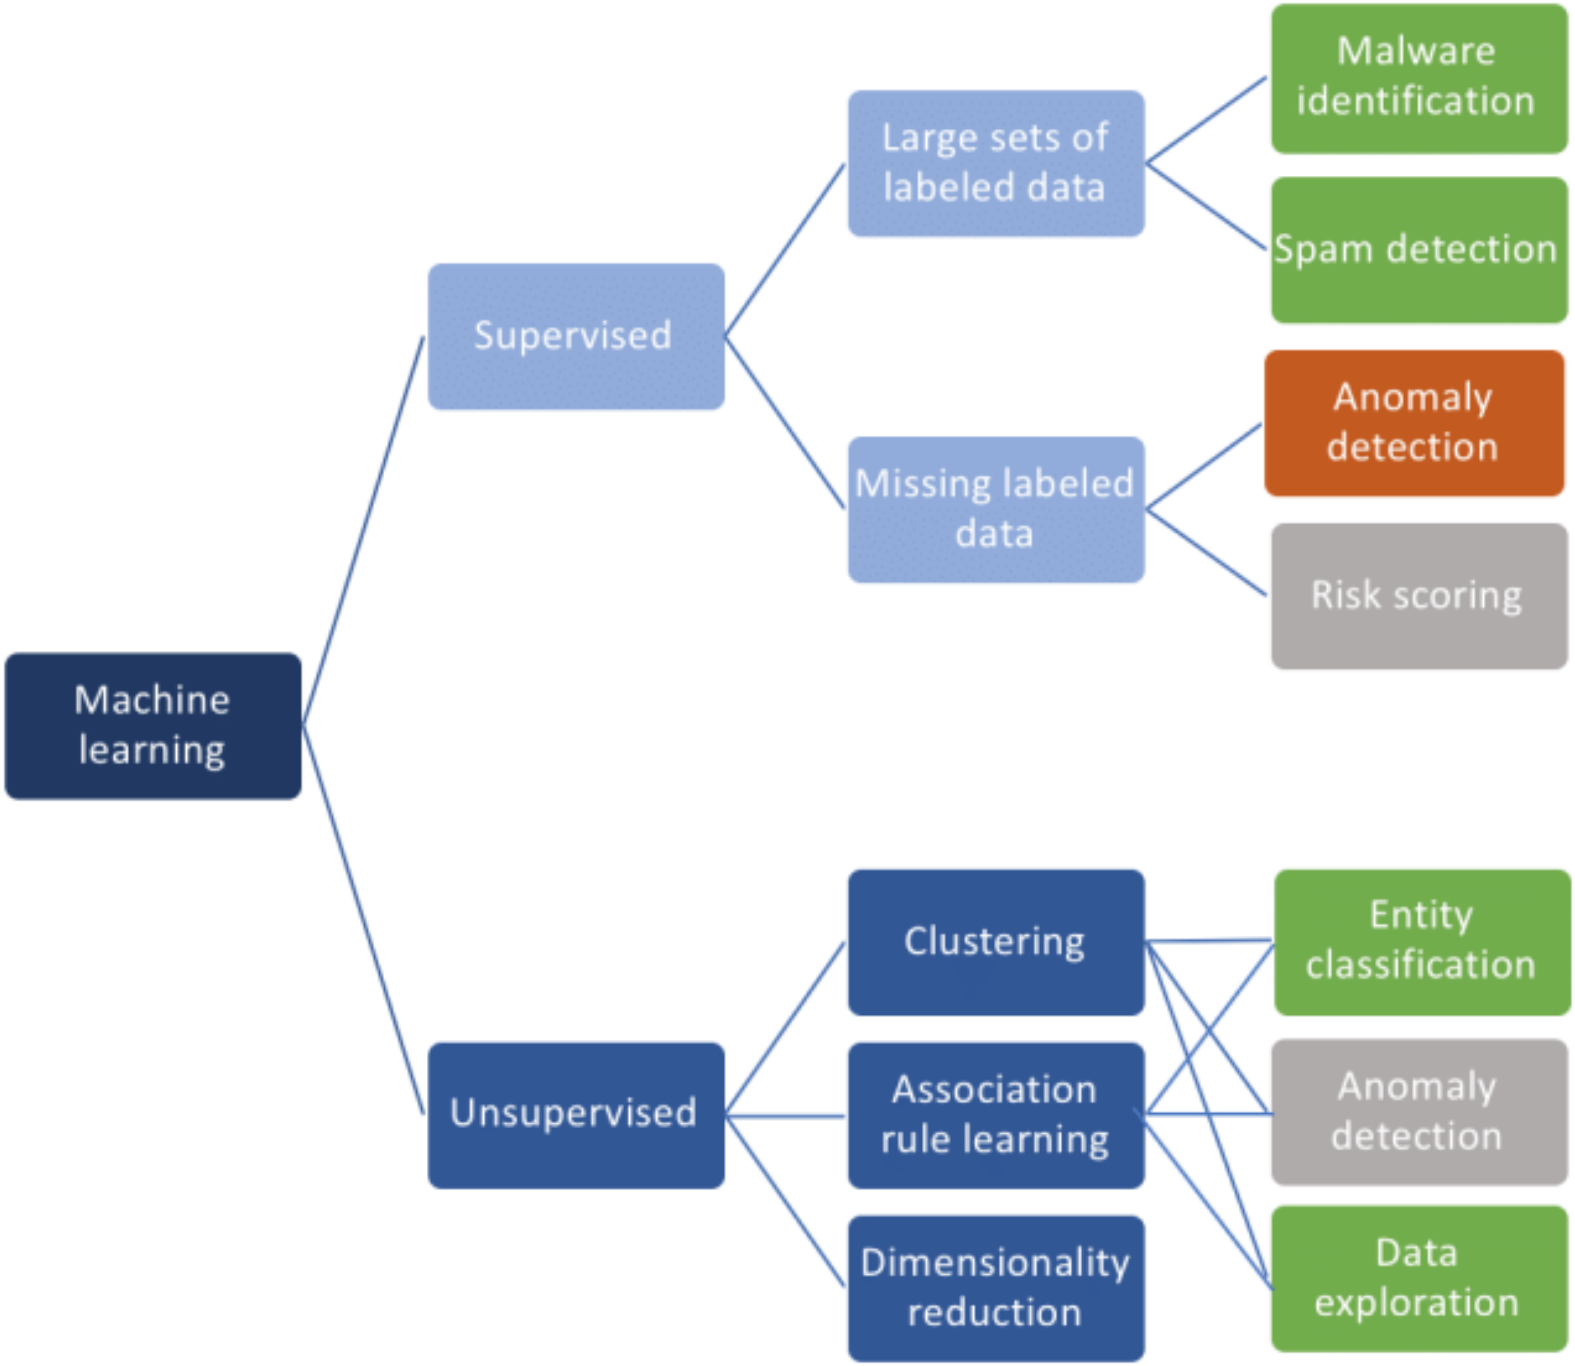
Machine learning
Supervised
Large sets of labeled data
Malware identification
Spam detection
Missing labeled data
Anomaly detection
Risk scoring
Unsupervised
Clustering
Entity classification
Association rule learning
Anomaly detection
Dimensionality reduction
Data exploration

ML Example: Linear Regression

- **Problem**: Predict house prices based on size.

ML Example: Linear Regression

- **Problem**: Predict house prices based on size.
- **Model**: $y = wx + b$ (price y, size x).

ML Example: Linear Regression

- **Problem**: Predict house prices based on size.
- **Model**: $y = wx + b$ (price y, size x).
- **Learning**: Adjust w and b to minimize error.

ML Example: Linear Regression

- **Problem**: Predict house prices based on size.
- **Model**: $y = wx + b$ (price y, size x).
- **Learning**: Adjust w and b to minimize error.
- **Visualization**:

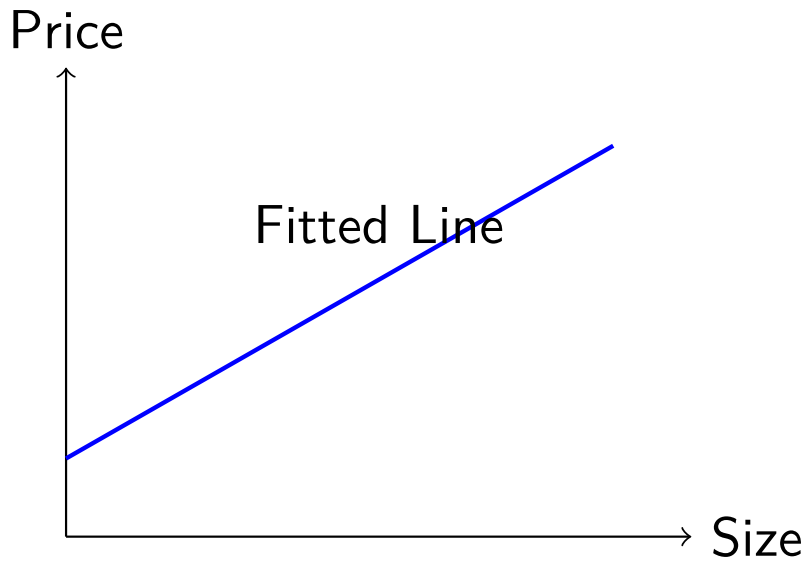

What is Deep Learning?

- **Definition**: ML using neural networks with multiple layers.

What is Deep Learning?

- **Definition**: ML using neural networks with multiple layers.
- **Neural Networks**: Mimic human brain's structure to learn complex patterns.

What is Deep Learning?

- **Definition**: ML using neural networks with multiple layers.
- **Neural Networks**: Mimic human brain's structure to learn complex patterns.
- **Why Deep?**: More layers $\rightarrow$ better at capturing complex relationships (e.g., image recognition).

What is Deep Learning?

- **Definition**: ML using neural networks with multiple layers.
- **Neural Networks**: Mimic human brain's structure to learn complex patterns.
- **Why Deep?**: More layers $\rightarrow$ better at capturing complex relationships (e.g., image recognition).
- **Generative AI Connection**: Many generative models (e.g., GANs) use deep learning.

Neural Network Basics

- **Neuron**: Computes weighted sum of inputs, applies activation (e.g., ReLU).

Neural Network Basics

- **Neuron**: Computes weighted sum of inputs, applies activation (e.g., ReLU).
- **Layers**: Input, hidden, output layers.

Neural Network Basics

- **Neuron**: Computes weighted sum of inputs, applies activation (e.g., ReLU).
- **Layers**: Input, hidden, output layers.
- **Training**: Backpropagation adjusts weights to minimize loss.

Neural Network Basics

- **Neuron**: Computes weighted sum of inputs, applies activation (e.g., ReLU).
- **Layers**: Input, hidden, output layers.
- **Training**: Backpropagation adjusts weights to minimize loss.
- **Visualization**:

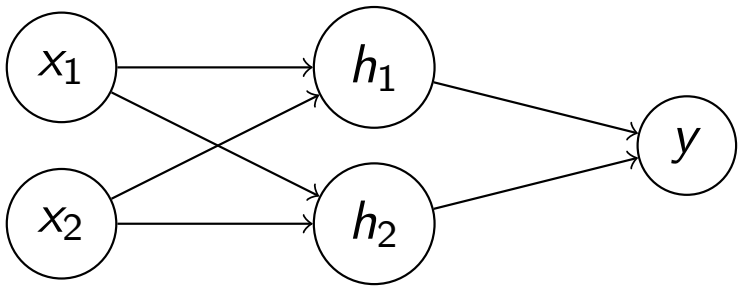

DL Example: Image Classification

- **Problem**: Classify images as cats or dogs.

DL Example: Image Classification

- **Problem**: Classify images as cats or dogs.
- **Model**: Convolutional Neural Network (CNN) with layers to detect edges, shapes.

DL Example: Image Classification

- **Problem**: Classify images as cats or dogs.
- **Model**: Convolutional Neural Network (CNN) with layers to detect edges, shapes.
- **Training**: Feed labeled images, adjust weights via backpropagation.

DL Example: Image Classification

- **Problem**: Classify images as cats or dogs.
- **Model**: Convolutional Neural Network (CNN) with layers to detect edges, shapes.
- **Training**: Feed labeled images, adjust weights via backpropagation.
- **Outcome**: High accuracy in distinguishing cats from dogs.

What are Large Language Models?

- **Definition**: Deep learning models trained on massive text data to generate or understand language.

What are Large Language Models?

- **Definition**: Deep learning models trained on massive text data to generate or understand language.
- **Architecture**: Transformer-based (e.g., GPT, BERT).

What are Large Language Models?

- **Definition**: Deep learning models trained on massive text data to generate or understand language.
- **Architecture**: Transformer-based (e.g., GPT, BERT).
- **Capabilities**: Text generation, translation, question answering.

What are Large Language Models?

- **Definition**: Deep learning models trained on massive text data to generate or understand language.
- **Architecture**: Transformer-based (e.g., GPT, BERT).
- **Capabilities**: Text generation, translation, question answering.
- **Generative AI Connection**: LLMs are a key type of generative AI for text.

Transformer Architecture

- **Key Components**:
  - Attention: Focuses on relevant words in a sentence.
  - Encoder-Decoder: Processes input and generates output.

Transformer Architecture

- **Key Components**:
  - Attention: Focuses on relevant words in a sentence.
  - Encoder-Decoder: Processes input and generates output.
- **Why Powerful?**: Captures long-range dependencies in text.

Transformer Architecture

- **Key Components**:
  - Attention: Focuses on relevant words in a sentence.
  - Encoder-Decoder: Processes input and generates output.
- **Why Powerful?**: Captures long-range dependencies in text.
- **Visualization**:

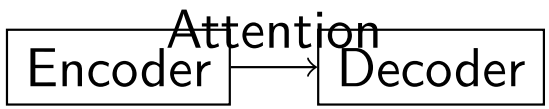

LLM Example: Text Generation

- **Problem**: Generate a story from a prompt.

LLM Example: Text Generation

- **Problem**: Generate a story from a prompt.
- **Model**: GPT-3 predicts next words based on context.

LLM Example: Text Generation

- **Problem**: Generate a story from a prompt.
- **Model**: GPT-3 predicts next words based on context.
- **Output**: Coherent, human-like text.

LLM Example: Text Generation

- **Problem**: Generate a story from a prompt.
- **Model**: GPT-3 predicts next words based on context.
- **Output**: Coherent, human-like text.
- **Example Prompt**: "Once upon a time..." $\rightarrow$ Model continues the story.

Types of Generative AI Models

- **Generative Adversarial Networks (GANs)**: Generator vs. discriminator to create realistic images.

Types of Generative AI Models

- **Generative Adversarial Networks (GANs)**: Generator vs. discriminator to create realistic images.
- **Variational Autoencoders (VAEs)**: Learn latent representations for data generation.

Types of Generative AI Models

- **Generative Adversarial Networks (GANs)**: Generator vs. discriminator to create realistic images.
- **Variational Autoencoders (VAEs)**: Learn latent representations for data generation.
- **Autoregressive Models**: Predict next element (e.g., LLMs like GPT).

Types of Generative AI Models

- **Generative Adversarial Networks (GANs)**: Generator vs. discriminator to create realistic images.
- **Variational Autoencoders (VAEs)**: Learn latent representations for data generation.
- **Autoregressive Models**: Predict next element (e.g., LLMs like GPT).
- **Diffusion Models**: Gradually refine noise into data (e.g., Stable Diffusion).

GANs: How They Work

- **Generator**: Creates fake data (e.g., images).

GANs: How They Work

- **Generator**: Creates fake data (e.g., images).
- **Discriminator**: Judges if data is real or fake.

GANs: How They Work

- **Generator**: Creates fake data (e.g., images).
- **Discriminator**: Judges if data is real or fake.
- **Training**: Generator improves to fool discriminator.

GANs: How They Work

- **Generator**: Creates fake data (e.g., images).
- **Discriminator**: Judges if data is real or fake.
- **Training**: Generator improves to fool discriminator.
- **Visualization**:

Applications of Generative AI

- **Content Creation**: Writing articles, generating art.

Applications of Generative AI

- **Content Creation**: Writing articles, generating art.
- **Healthcare**: Synthetic medical images for training.

Applications of Generative AI

- **Content Creation**: Writing articles, generating art.
- **Healthcare**: Synthetic medical images for training.
- **Education**: Personalized learning materials.

Applications of Generative AI

- **Content Creation**: Writing articles, generating art.
- **Healthcare**: Synthetic medical images for training.
- **Education**: Personalized learning materials.
- **Entertainment**: AI-generated music, movies.

Challenges in Generative AI

- **Ethical Concerns**: Misinformation, deepfakes.

Challenges in Generative AI

- **Ethical Concerns**: Misinformation, deepfakes.
- **Bias**: Models reflect biases in training data.

Challenges in Generative AI

- **Ethical Concerns**: Misinformation, deepfakes.
- **Bias**: Models reflect biases in training data.
- **Computational Cost**: Training requires significant resources.

Challenges in Generative AI

- **Ethical Concerns**: Misinformation, deepfakes.
- **Bias**: Models reflect biases in training data.
- **Computational Cost**: Training requires significant resources.
- **Interpretability**: Hard to understand model decisions.

Key Takeaways

- Generative AI creates novel content using ML and DL techniques.

Key Takeaways

- Generative AI creates novel content using ML and DL techniques.
- ML provides the foundation, DL scales complexity, and LLMs excel in language tasks.

Key Takeaways

- Generative AI creates novel content using ML and DL techniques.
- ML provides the foundation, DL scales complexity, and LLMs excel in language tasks.
- Applications span industries, but challenges like ethics and bias remain.

Key Takeaways

- Generative AI creates novel content using ML and DL techniques.
- ML provides the foundation, DL scales complexity, and LLMs excel in language tasks.
- Applications span industries, but challenges like ethics and bias remain.
- **Future**: Generative AI will continue to transform technology and society.

References

- Goodfellow, I., et al. (2014). Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 27.

References

- Goodfellow, I., et al. (2014). Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 27.
- Vaswani, A., et al. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems*, 30.