



AI Concept Flashcards

A rapid AI knowledge-revision system — scan one slide, refresh an entire domain in 30–60 seconds. Each card covers ~15 related concepts, grouped into visual clusters with micro-definitions. Built for engineers, students, and product leaders who need to think clearly about AI.


20 FLASHCARD MODULES

300+ AI CONCEPTS

RAPID REVISION

AI Foundations

The hierarchy every AI practitioner must internalize — from broad intelligence to specific model parameters.



Dataset

Collection of examples used to train

Feature

Input variable the model analyzes

Label

Correct answer in supervised data

Algorithm

Step-by-step learning procedure

Training

Optimizing model on labeled data

Inference

Using a trained model to predict

Parameters

Learned values inside the model

Hyperparameters

Settings chosen before training begins

Supervised Learning

Train on labeled input-output pairs

Unsupervised Learning

Find hidden patterns without labels

Reinforcement Learning

Learn via reward and penalty signals

Model

Mathematical function mapping input → output

Flashcard 02

Machine Learning

The ML toolbox — algorithms, objectives, pitfalls, and evaluation methods every practitioner must recognize at a glance.



Classification

Predict discrete class labels



Clustering

Group similar data without labels



Random Forest

Ensemble of many decision trees



Regression

Predict continuous numeric output



Decision Trees

Rules split on feature thresholds



SVM

Maximize margin between classes

K-Means

Assign points to K cluster centers

Gradient Descent

Minimize loss by stepping downhill

Loss Function

Measures model prediction error

Overfitting

Memorizes training, fails on new data

Underfitting

Too simple to capture patterns

Bias

Systematic error from wrong assumptions

Variance

Sensitivity to small data fluctuations

Cross-Validation


Rotate test sets to evaluate robustness

Accuracy

Fraction of correct predictions made

Deep Learning

From a single artificial neuron to the transformer architectures powering today's AI — understand the building blocks and the progression.




Neuron Single computation unit with weights	Layer Group of neurons at same depth	Activation Function Adds non-linearity to neuron output	Forward Propagation Input flows forward to produce output
Backpropagation Errors flow back to update weights		Weight Learned strength of a connection	
Epoch One full pass through training data	Batch Subset of data per training step	Learning Rate Step size during gradient descent	CNN Learns spatial patterns in images
RNN Processes sequential data with memory	LSTM RNN with long-term memory gates	Transformer Attention-based modern architecture	Bias (node) Constant offset added to neuron sum

Flashcard 04

Large Language Models (LLMs)

The complete LLM lifecycle — from raw data to a deployed, prompt-responsive intelligence.



Data

Pre-training

Fine-tuning

Prompt

Inference

LLM

Massive model trained on text corpora

Transformer

Architecture powering all modern LLMs

Token

Small unit an LLM processes

Context Window

Max tokens model can process at once

Parameter

Learned numeric value inside model

Attention

Mechanism to weigh token importance

Pre-training

Learn language from massive corpus

Fine-tuning

Adapt pretrained model to a task

Instruction Tuning

Train model to follow instructions

RLHF

Human feedback shapes model behavior

Temperature

Controls randomness of model output

Hallucination

Model confidently states false information

Reasoning Model


LLM that thinks step-by-step

Self-Attention

Tokens attend to each other in context

Transformers

The architectural revolution behind GPT, BERT, and every modern LLM — understand each component and its role in producing output.



Encoder

Reads and encodes input sequence

Decoder

Generates output tokens one by one

Query (Q)

What this token is looking for

Key (K)

What each token offers to others

Value (V)

Content retrieved when key matches

Positional Encoding

Injects word order into embeddings

Multi-head Attention

Parallel attention across many subspaces

Layer Normalization

Stabilizes training across layers

Feed-forward Network

Applies learned transformations per token

Context


All tokens the model currently sees

Next-token Prediction

Core training objective of language models

Prompt Engineering

The art and science of instructing AI — master these techniques to unlock reliable, precise, and controllable model behavior.



System Prompt

Instructions defining model persona and rules



Few-shot

Provide 2–5 examples in the prompt



Tool Calling

Model invokes external functions or APIs



Zero-shot

No examples — model uses pure knowledge



Chain-of-Thought

Prompt model to reason step by step



Guardrails

Rules that constrain model outputs safely

Role Prompting

Assign model a persona or expert role

Prompt Chaining

Pass one output as the next prompt input

Structured Output

Request JSON or specific format response

Delimiters

Symbols that separate prompt sections clearly

Context Engineering

Curate what context the model receives

Grounding

Anchor model responses in factual sources

Retrieval-Augmented Generation (RAG)

RAG extends LLMs with external knowledge — grounding every response in real, retrievable information rather than memorized weights.




RAG Retrieve knowledge before generating response	Embedding Meaning represented as numeric vectors	Vector List of numbers encoding semantic meaning	Vector Database Stores and searches by vector similarity
Chunking Split documents into searchable segments		Retrieval Fetch most relevant chunks for prompt	
Semantic Search Find meaning-similar content via vectors	Keyword Search Find exact term matches in documents	Hybrid Search Combine semantic and keyword retrieval	Re-ranking Score and reorder retrieved results
Metadata Attributes that filter retrieval results	Grounding Responses anchored to retrieved facts	Knowledge Base Curated document store for retrieval	Graph RAG Retrieval enhanced with knowledge graph

Flashcard 08

AI Agents

Agents are AI systems that don't just answer — they plan, act, use tools, observe results, and iterate autonomously toward a goal.



AI Agent

System that perceives, plans, and acts autonomously



Memory

Stores context across agent steps



Human-in-the-Loop

Human approves critical agent decisions



Tool Calling

Agent invokes external APIs or functions



Multi-Agent System

Multiple specialized agents cooperating



Orchestration

Coordinate multiple agents toward goal

Planning

Break goal into executable subtasks

Reasoning

Logical inference guiding next action

Reflection

Agent evaluates its own past actions

Action

Executable step the agent performs

Observation


Feedback received after taking action

Autonomous Workflow

Multi-step process running without human

AI Engineering

Building production AI systems requires more than model selection — it demands infrastructure, optimization, and operational excellence.



API

Interface to access model capabilities

SDK

Toolkit to build AI-powered apps

Model Serving

Deploy model to handle live requests

Deployment

Move model from dev to production

MLOps

DevOps practices applied to ML systems

LLMOps

MLOps specialized for large language models

Model Registry

Versioned store for trained models

Quantization

Reduce model size with lower precision

Distillation

Train small model to mimic large one

Fine-tuning

Adapt pretrained model to specific task

Latency

Time from request to response delivery

Throughput

Requests processed per unit of time

GPU


Parallel processor optimized for AI training

TPU

Google's custom AI accelerator chip

AI Architecture

Enterprise AI is a layered stack — each layer serves a distinct purpose, from raw data ingestion to the user-facing interface.



Data Pipeline

Automates data movement and transformation

Cloud

Scalable remote compute infrastructure

Kubernetes

Orchestrates containers at scale

Model Gateway

Routes and manages model API traffic

Monitoring

Track model performance in production

Governance

Policies ensuring responsible AI use

Flashcard 11

Generative AI

GenAI doesn't retrieve — it creates. Text, images, video, audio, and code, all generated from learned patterns in massive datasets.



Text Generation

LLMs produce language from prompts



Image Generation

Diffusion models create visual content



Video Generation

AI creates temporal visual sequences



Audio Generation

AI synthesizes speech and music

LLM

Large-scale language generation model

Diffusion Model

Generates by denoising random noise

GAN

Generator vs. discriminator adversarial training

Foundation Model

Massive pretrained base for many tasks

Multimodal AI

Understands and generates multiple modalities

Synthetic Data

AI-generated training data at scale

Fine-tuning

Specialize a foundation model for use case

AI Copilot

AI assistant embedded in workflow tools


RAG

Ground GenAI in real retrieved knowledge

Flashcard 12

AI Governance & Responsible AI

Trust is built through structure. Responsible AI demands accountability at every layer — from training data to production deployment and beyond.



Bias

Unfair skew embedded in training data

Fairness

Equal outcomes across demographic groups

Explainability

Understand why model made a decision

Interpretability

Inspect internal model decision logic

Privacy

Protect personal data in AI systems

Security

Defend AI from adversarial attacks

Transparency

Open disclosure of model capabilities and limits

Human Oversight

Humans retain control over AI decisions

Guardrails

Technical limits on model behavior outputs

Model Risk

Potential harm from model failures

Auditability


Ability to trace and review AI actions

AI Regulation

Laws governing AI development and use

AI Strategy

Strategy transforms AI from experiment to enterprise capability — by aligning use cases, investment, and operating models with business value.



AI Vision

Where AI takes the organization long-term



Use Case Prioritization

Rank opportunities by value and feasibility



Build vs. Buy

Make or procure AI capabilities strategically



AI Maturity

Organization's current AI capability level



AI Portfolio

Balanced set of AI initiatives and investments



AI Roadmap

Sequenced plan from pilot to scale

ROI

Measured return on AI investment

Automation

AI replaces repetitive manual tasks

Augmentation

AI enhances human decision-making

AI Operating Model

How AI capabilities are organized and governed

Flashcard 14

Enterprise AI

Enterprise AI is not a single tool — it is a capability system built on platforms, talent, governance, and operating discipline.



AI Factory

Industrialized AI development at enterprise scale



Model Governance

Policies for model risk and compliance



Change Management

Guide organization through AI adoption



AI Center of Excellence

Central hub for AI standards and expertise



AI Talent

Engineers, scientists, and translators needed



AI Literacy

Baseline AI understanding across the workforce

AI Platform

Shared infrastructure for AI development

Data Governance

Rules for data quality and access

AI Adoption

Embedding AI in everyday workflows

Citizen AI

Non-technical users leverage AI tools

AI Risk

Threats from AI failures or misuse


AI Value Capture

Realizing business outcomes from AI investment

Flashcard 15

AI Value Creation

AI creates value across a spectrum — from operational efficiency at the base to transformational business reinvention at the peak.



Automation

AI executes repetitive tasks autonomously

Augmentation

AI amplifies human decision-making speed

Productivity

More output per unit of human effort

Revenue Growth

AI drives new sales or pricing opportunities

Cost Reduction

Lower operational expense through AI

Personalization

Tailor experiences to each individual user

Decision Intelligence

AI-enhanced, data-driven business decisions

Hyperautomation

End-to-end automation across entire processes


AI Copilot

Embedded AI assistant in user workflow

AI vs. Automation vs. GenAI vs. Agentic AI

Understanding the evolution from rule-following machines to autonomous AI systems is critical for choosing the right technology for the right problem.

	Approach	Learns	Creates	Acts Autonomously
Traditional Automation	Rules-based, deterministic	No learning	No generation	No
RPA	Scripted actions	No reasoning	No generation	No
Machine Learning	Learn from data	Statistical patterns	No generation	No
Predictive AI	Forecast outcomes	Data-driven	No generation	No
Generative AI	Creates content	Probabilistic	Yes	No
Copilot	AI-assisted human	Process-driven	Yes (suggestions)	No (human-in-loop)
AI Agent	Autonomous goal-pursuit	Task-oriented	Yes (actions/content)	Yes (multi-step)
Agentic AI	Multi-agent, self-directing	Adaptive strategies	Yes (complex workflows)	Yes (complex)



Flashcard 17

AI Data Concepts

Every AI system is only as good as its data — understanding data types, pipelines, and quality concepts is foundational to all AI work.



Structured Data

Organized in rows and columns (tables)



Data Pipeline

Automated process moving and transforming data



Data Warehouse

Processed, structured data for analytics



Unstructured Data

Text, images, video with no fixed schema



Data Lake

Raw data stored at massive scale



Knowledge Graph

Entities and relationships in graph form

Training Data

Labeled examples used to fit model

Synthetic Data

AI-generated data for training augmentation

Data Quality

Accuracy, completeness, and consistency of data

Metadata

Data describing attributes of other data

Embeddings

Dense vector representations of data meaning

Data Lineage

Trace data origin and transformation history

Ground Truth

Verified correct labels for evaluation

Data Governance

Policies managing data quality and access

Flashcard 18

AI Security

AI systems face unique attack surfaces — from adversarial inputs at the model layer to data leakage in retrieval pipelines. Security must be designed in, not bolted on.



Prompt Injection

Malicious input hijacks model behavior



Adversarial Attack

Inputs designed to fool model deliberately



Red Teaming

Adversarial testing to find model weaknesses



Data Leakage

Sensitive data exposed in model output



Jailbreaking

Bypass model safety guardrails via prompts



Zero Trust

Never trust — always verify every access request

Model Theft

Replicating model via API queries

Toxic Output

Harmful or offensive model-generated content

Privacy

Protect user data within AI workflows

Access Control

Restrict who can query the model

Sandboxing

Isolate agent actions from production systems

Guardrails

Rules that block unsafe model outputs

Encryption

Protect data at rest and in transit

Monitoring

Detect anomalous model behavior in prod

Flashcard 19

Emerging AI

The frontier is moving fast. These are the concepts reshaping what AI can do — and what it will become over the next 3–5 years.



Reasoning Models

LLMs that think step-by-step before answering



Multimodal AI

Models understanding text, images, audio, video



World Models

AI that simulates how the world works



Small Language Models

Efficient models for on-device deployment



Physical AI

AI governing robots in the physical world



AI Coding

AI writes, reviews, and debugs code autonomously

Agentic AI

Autonomous systems
executing multi-step goals

Synthetic Data

AI-generated training data at
any scale

AI Companions

Persistent, personalized AI
relationships

Autonomous Systems

Self-directing machines in
complex environments

Multi-Agent Systems

Teams of AI agents
collaborating on tasks

AI Scientists

AI systems that conduct
research autonomously

AI-Native Companies

Organizations built entirely
around AI workflows

AI + Robotics

Embodied intelligence in the
physical world

Flashcard Complete

You've covered the full AI stack.

From foundational ML concepts to emerging agentic systems, you now have a mental model of the entire AI landscape. Use these cards to review, teach, or build.

- 1 AI Foundations
- 2 Machine Learning
- 3 Deep Learning
- 4 Large Language Models
- 5 Transformers
- 6 Prompt Engineering
- 7 Retrieval-Augmented Generation
- 8 AI Agents
- 9 AI Engineering
- 10 AI Architecture
- 11 Generative AI
- 12 AI Governance & Responsible AI
- 13 AI Strategy
- 14 Enterprise AI
- 15 AI Value Creation
- 16 AI vs. Automation vs. GenAI vs. Agentic AI
- 17 AI Data Concepts
- 18 AI Security
- 19 Emerging AI

Keep building. Keep learning. The frontier is yours.