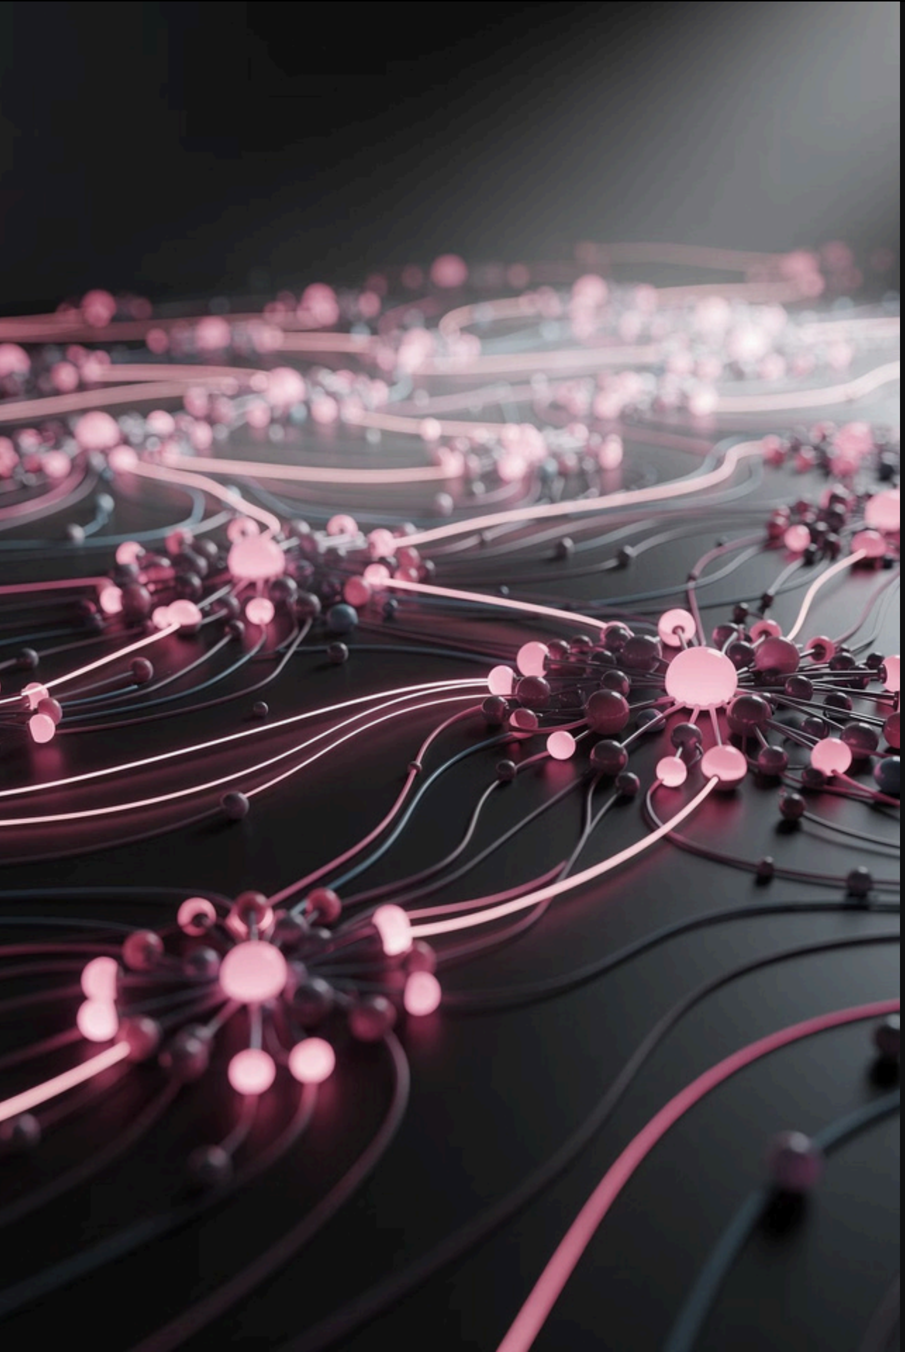

Taxonomy of AI Technology

The distinction between Artificial Intelligence, Machine Learning, and Deep Learning Learning—the foundation of every software engineer.

Why Taxonomy Matters

Industry problem

The term "AI" has become a universal buzzword—the same word describes a simple email filtering script and a multi-billion dollar language model. This blurs the lines between deterministic systems and adaptive algorithms.

Consequences for the engineer

- Using a neural network where linear regression is sufficient → over-engineering Trying if-else rules for image recognition → failure
- **design**
- Incorrect hardware selection (CPU vs. GPU) and underestimated
- data. Enormous technical debt for future developers.

The AI taxonomy is a hierarchical system for classifying methods, from the broadest concept (AI) to the most specialized subset (Deep Learning).

Artificial Intelligence - The Umbrella Concept

AI is a superordinate field of computer science encompassing all systems capable of simulating human cognitive abilities: problem-solving, pattern recognition, planning and language understanding.

AI Examples

Deep Blue (chess, 1997), GPS navigation, virtual assistants in smartphones

Analogy

AI is like "motor vehicle"—an umbrella term for scooters, trains, and airplanes. It says nothing about the engine, only the purpose: getting from A to B.

Three Levels of AI

1

Superintelligence (ASI)

A hypothetical machine surpassing human intellect in all areas—creativity, wisdom, self-improvement—does not exist.

2

General AI (AGI)

A theoretical system that learns and adapts like a human in any domain. The subject of long-term research. Does not exist.

3

AI Narrow (ANI)

The only thing that exists today. Designed for a single task—spam filters, Netflix recommendations, facial recognition. Even GPT-4 is ANI.

Key takeaway: Everything the IT industry deals with today—including the most powerful LLMs—is still Narrow AI (ANI). Powerful, but single-purpose calculators.

Three Approaches to Building AI

Symbolic (GOFAI)

Hand-coded if-then-else rules. Advantages: full transparency, determinism, no need for data. Disadvantages: extremely rigid—any unknown situation causes a failure.

Connectionist

Artificial neural networks that learn from data. Advantages: flexibility, unstructured data. Disadvantages: "black box" nature, requires massive data sets and GPU power.

Hybrid (Neuro-symbolic)

Combines DL (perception, patterns) with logical modeling (deduction). It maintains adaptivity while dramatically improving the explainability of algorithms.

Machine Learning—A Subset of AI

Machine Learning specialized field AI: algorithms automatically analyze data, identify hidden statistical patterns and make decisions — *with out being explicitly programmed for a specific task.*

Example

A credit risk algorithm at a bank that has independently learned to predict loan repayments based on the histories of hundreds of thousands of customers.

Analogy

A child learning to ride a bike — not by reading a physics textbook, but through thousands of trials and errors, creating their own model of balance.

Reversing the Programming Paradigm

Traditional programming

Data + Rules → Answer

The programmer writes by hand: `if amount > 100: discount = 10%` The result is deterministic and predictable.

Machine learning

Data + Results → Rules The algorithm will conclude that "customers aged 25–30, buying on Tuesdays

in the evening from the mobile app, have the highest probability of renewing the subscription."

Thanks to this breakthrough, ML solves problems with complexity far beyond the capabilities of human multidimensional perception.

Three Paradigms of Machine Learning

1

Supervised Learning

Labeled data. The model minimizes the error between the prediction and the true answer.

- **Classification: Spam Detection** (spam / not spam)
- **Regression: Predicting the Price of an Apartment**

2

Unsupervised Learning

Data without labels. The algorithm itself discovers hidden structures and patterns.

- **Clustering: Customer Segmentation**
- **Anomaly Detection: Card Fraud** credit

3

Reinforcement Learning

The agent learns by interacting with the environment – rewards and punishments shape the optimal strategy.

- **Example: AlphaGo (DeepMind) —** millions of parties with yourself
- **Example: Robot control** humanoid

Classic ML Algorithms and Tools

Key algorithms

→ **Linear/Logistic Regression - Value Prediction and Two-Class Classification**

→ **Decision Trees and Random Forest - Sets of Logical Trees for Hard Decisions**

→ **SVM - finding the optimal hyperplane separating classes**

→ **K-Means - Unsupervised Clustering around Centroids**

Standard tool: scikit-learn

A Python library with an elegant, unified API. It hides mathematical complexity in simple calls:

- .fit() — start training .predict() —
- get prediction

Often, just a dozen or so lines of code are enough to train and evaluate a model.

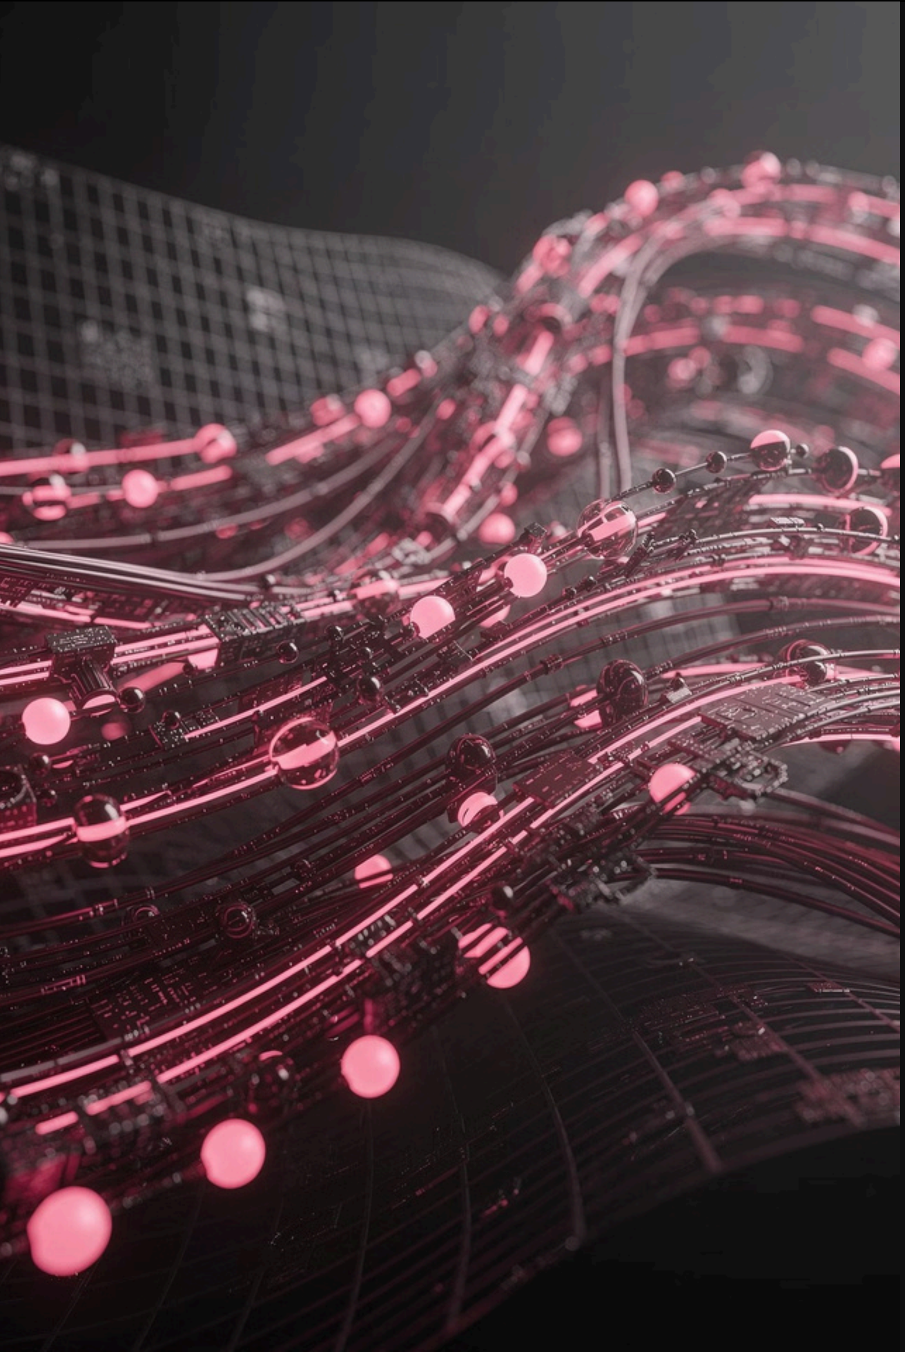

Deep Learning – A Subset of ML

Deep Learning is a specialized subset of MLo based on giant neural networks with many hidden layers. Since its breakthrough in 2012, AlexNet has revolutionized image recognition, machine translation, and enabled the emergence of LLM.

> The term "deep" refers only to the physical depth of the network—the number of mathematical layers. It does not mean "depth of understanding."

3

Success factors

Big Data, GPU power and algorithmic breakthroughs (backpropagation, ReLU)

2012

The year of the breakthrough

AlexNet Wins ImageNet Competition, Inaugurating the Era of Deep Learning

100+

Hidden Layers

Modern DL architectures can have hundreds of mathematical layers

Key Deep Learning Architectures

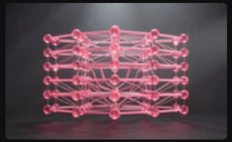

Dense / Fully Connected

Each neuron is connected to every other neuron in the next layer. The foundation of larger systems, good for tabular data.

CNN - Convolutional Networks

They move filters across an image to detect edges and shapes. They dominate computer vision, object detection, and autonomous vehicles.

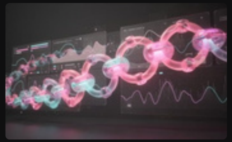

RNN / LSTM / GRU

They have a "memory" of previous sequence elements. LSTM is revolutionizing NLP, time series analysis, and speech generation.

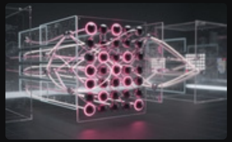

Transformer

The self-attention mechanism analyzes all words simultaneously. The basis of GPT, Claude, and Gemini. Dominates NLP thanks to parallelization of computations.

Hierarchy AI ⊃ ML ⊃ DL

Computer Vision (CV) and Natural Language Processing (NLP) are application domains, not algorithms—they cut across the entire hierarchy. LLM (GPT, Claude, Gemini) reside at the heart of DL, based on the Transformer architecture.

When is DL and when is Classic ML?

Criterion	Classic ML	Deep Learning
Input data	Structured, tabular	Unstructured (image, sound, text)
Amount of data	Thousands–hundreds of thousands of records	Millions+ (Big Data)
Equipment	CPU (standard processor)	GPU / TPU (accelerators)
Interpretability	High (decision trees)	Low - "black box"
Training time	Seconds–minutes	Weeks–months
Implementation cost	Low-medium	Very high

Overfitting is a particular threat of DL with small datasets – the model memorizes training data instead of generalizing patterns.

Applications in the IT Industry

E-commerce

Product recommendations → ML (K-Means, decision trees) Product image categorization → DL (CNN)

Finance and Banking

Fraud Detection → Unsupervised ML (Anomaly Detection) Credit Scoring → Supervised ML (Random Forest, Logistic Regression)

Medicine

MRI scan analysis, tumor detection → DL (CNN) Drug dosing according to protocols → AI Symbolic (expert systems)

Industrial Automation

Robotic arm control → DL (Reinforcement Learning) NPC navigation in games → Symbolic AI (A* algorithm)

Cybersecurity

Spam Filtering → Supervised ML (Naive Bayes) Job Offer Generation → DL / GenAI (LLM, Transformer)

Common Conceptual Mistakes and Pitfalls

"AI is a physical robot with a personality"

A myth ingrained in pop culture, AI is an intangible script residing on servers—a mathematical model that trains itself on data. It doesn't need a body or robots to be artificial intelligence.

"ML eliminates the point of writing algorithmic rules"

Pathological over-engineering. In closed business environments, where decision rules can be written down on paper, classic if/else programming is always the optimal economic and engineering response.

"DL understands the analyzed text or phenomenon"

LLMs are remarkably complex statistical calculators designed solely for precise pattern matching. The model calculates that the word "Morning" should statistically be followed by "good"—it doesn't understand the idea behind the words.

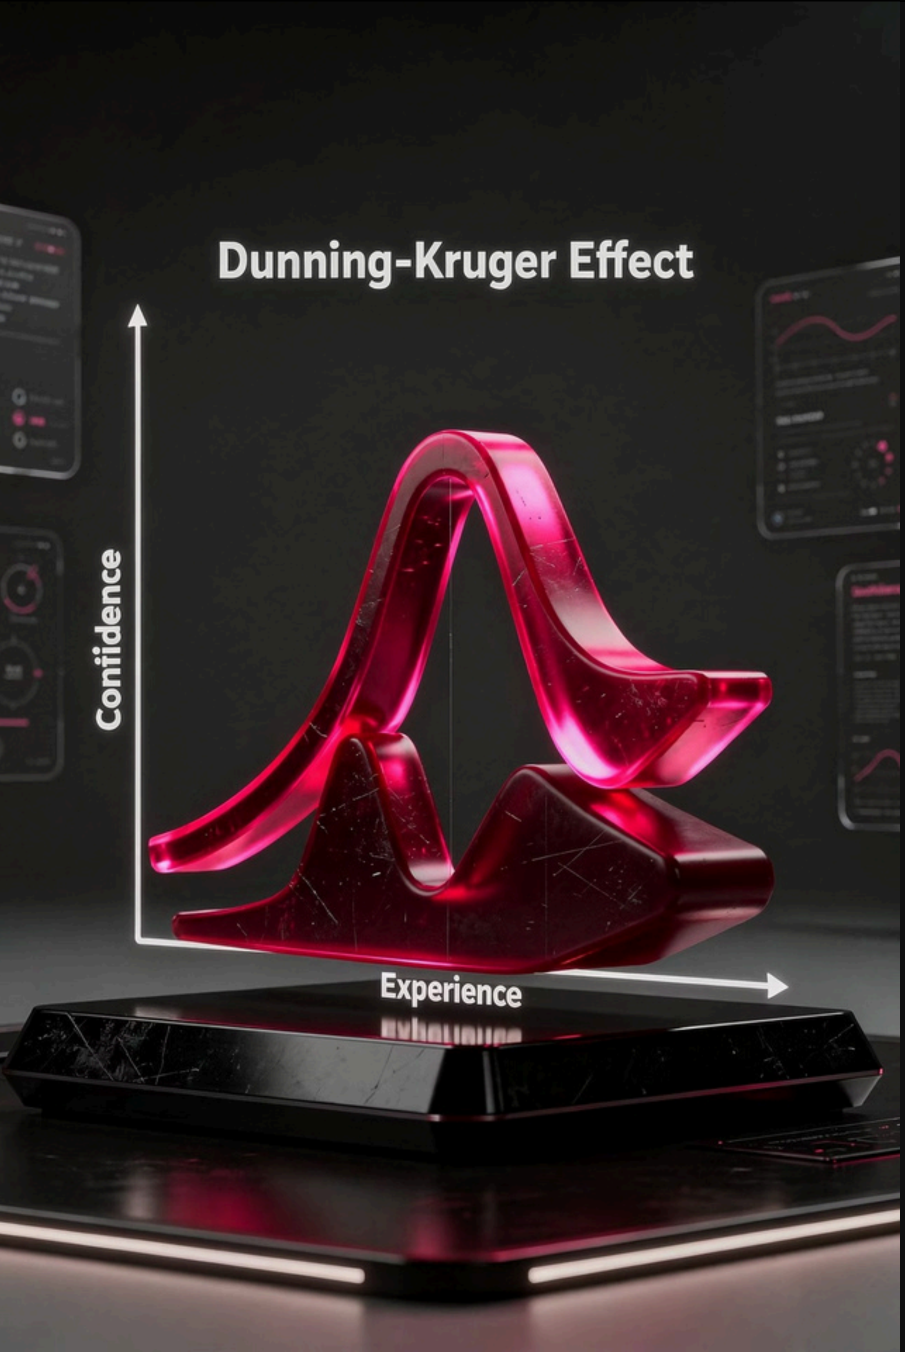

Push-Kruger Effect AI

The Dunning-Kruger effect—people with low domain knowledge dramatically overestimate its competences—gained a new iteration of the generative AI.

1 **Cognitive offloading**

Novices uncritically delegate to the model the task of checking logical solutions and writing complex code, treating the model's output as an absolute fact.

2 **Ignoring hallucinations**

One laconic prompt, taking the result for granted—without validation. Models have built-in artificial overconfidence through manufacturer calibration methods.

3 **The Engineer's Duty**

A critical approach, testing and validation of AI results is an absolute must for every computer scientist working with generative models.

Exercise 1: Taxonomy of Business Problems

Categorize each problem into: Symbolic AI / Supervised ML / Unsupervised ML / Deep Learning

#	Business problem	Answer
1	Spam/non-spam email filtering based on thousands of flagged messages	ML Supervised
2	Vacuum cleaner with rules: if sensor=true then turn_left	AI Symbolic
3	Automatic face tagging from airport surveillance	Deep Learning (CNN)
4	Segmenting app users by click patterns	ML Unsupervised
5	Car price prediction from 100,000 listings	Supervised ML (Regression)
6	Generating photorealistic product graphics from text descriptions	Deep Learning (GenAI)
7	Diagnosis of skin diseases from photos of moles	Deep Learning (CNN)
8	TV series recommendations that look for correlations in the behavior of similar users	ML Unsupervised
9	A banking chatbot operating on a decision tree with a ready-made Q&A list	AI Symbolic
10	Real-time Polish to Japanese translator	Deep Learning (Transformer)

Exercise 2: Iris Classifier in scikit-learn

```
# 1. Import necessary modules from sklearn.datasets
import load_iris from sklearn.model_selection import
train_test_split
from sklearn.linear_model import LogisticRegression

# 2. Loading the collection
X, y = load_iris(return_X_y=True)

# 3. Set partition and initialization of the classical model
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3)
model = LogisticRegression(max_iter=200)

#4. Learning process (fit) and prediction
model.fit(X_train, y_train)

#5. Validation of results
print(f"Effectiveness: {model.score(X_test, y_test)}")
```

Is it ML or DL? Definitely classic ML. The model uses LogisticRegression—one of the oldest supervised learning mechanisms, without a single hidden neural network layer, working on perfectly structured tabular data.

Glossary of Key Terms

- **Fundamentals of taxonomy**
  - **AI (Artificial Intelligence) — a superior field that simulates human cognitive abilities ANI / AGI / ASI — narrow (existing) /**
  - general (theoretical) / superintelligence (hypothetical)
  - **ML (Machine Learning) - algorithms** learning from data without explicit rule programming
  - **DL (Deep Learning) - ML with** multi-layer neural networks, requires GPU and Big Data

- **Learning paradigms**
  - **Supervised learning - data with labels, model minimizes prediction error Unsupervised learning - none** labels, discovering hidden structures
  - **Reinforcement learning - agent,** rewards and penalties, maximizing cumulative reward
  - Feature Engineering - Handcrafted preparing input features for classic ML
  -

- **Key technical terms**
  - **Overfitting / Underfitting - overfitting (data memorization) / underfitting (too simple a model) Classification / Regression / Clustering**
  - — discrete categories / continuous value / unlabeled grouping CNN / LSTM / Transformer - DL
  - architectures for images / sequences / natural language Hyperparameter -
  - settable value by the programmer before training (e.g. number of layers)

Summary: AI Engineer Competency

You don't have to know all the matrix formulas by heart. Key Competency is the ability to precisely map a business problem to the appropriate technological paradigm.

The right choice tools

Simple rule → Symbolic AI. Tabular data → classic ML. Image/audio/text at scale → Deep Learning.

Precise communication

Engineers need to be sure that "supervised learning" means the same, well-defined optimization process for both parties.

Critical validation

Test, verify, and don't blindly trust AI models. Hallucinations and the Dunning-Kruger effect are real occupational hazards.