AI,
Autumn, 2026

Introduction to Artificial Intelligence

Faculty of DS & AI
Autumn semester, 2026

Trong-Nghia Nguyen

Business AI Lab

Content

- **About**
- The Foundation of AI
- The History of AI
- The Categories of AI
- The Application of AI
- Careers in AI

About

Bio

Website

Trong-Nghia Nguyen (PhD)
(1996)

Present, Lecture at **National Economics University**

2021, Msc. at **Hanoi University of Science and Technology**

2025, PhD. at **Chonnam National University, South Korea**

2018, BS. at **University of Science, Hue University**

Disciplines: **Artificial Intelligence, Computer Science, Information Technology**

Research

Smart technology

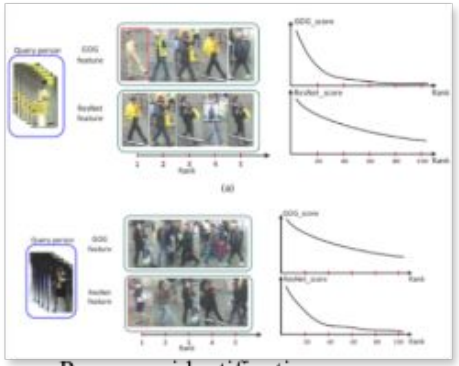

Person re-identification

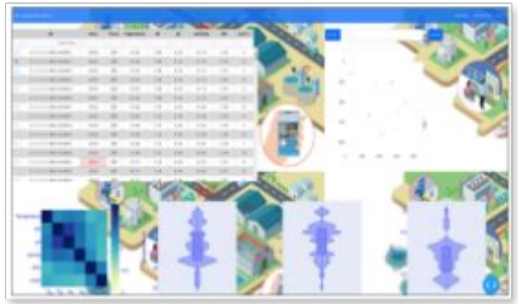

Smart Manufactoring

Smart Healthcare

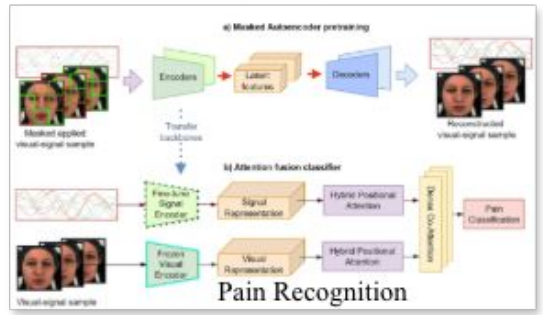

Pain Recognition

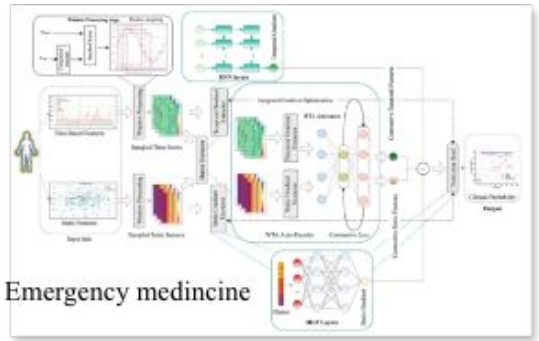

Emergency medicine

On going

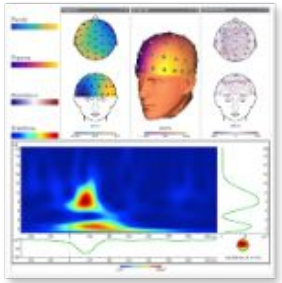

Brain EEG signal

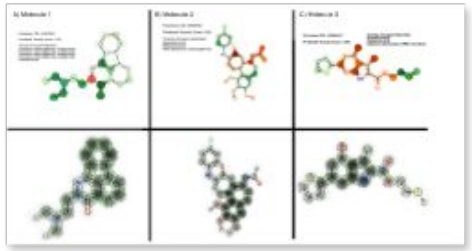

Drug's molecular structure

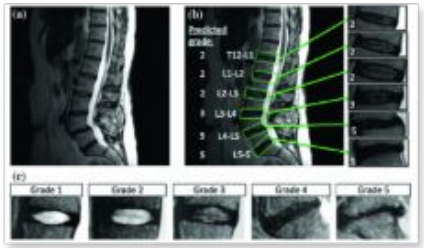

MRI - Pathology diagnostic

About

This course

Course site

Grading Component	Weight	Assessment Details
Attendance & Participation	10%	Random roll calls, active class participation (speaking up, doing board exercises).
Midterm Exam	40%	Midterm test / assignment / essay.
Final Exam	50%	Final examination / capstone project.
TOTAL	100%	

Rule Category	Detailed Policy Description
Base Attendance Score	Students with full attendance but **no active participation** will receive a maximum base score of **8.0/10** for this component.
Participation Bonus	Each instance of active participation (e.g., answering questions, doing exercises on the board) adds +**0.25** points to the Attendance score.
Absence Penalty & Ban	**Random roll calls** will be conducted. If a student is absent for more than 4 sessions, their **Attendance score will be 0**, and they will **NOT be eligible to take the Final Exam**.
Excess Point Transfer	The maximum recorded Attendance score is 10.0. If a student's total attendance points exceed 10.0, the **excess points** will be **divided by 2 (50%)** and **added directly to their Midterm Exam score.**
Calculation Example	Example: - A student has full attendance (8.0 pts) and participates 10 times (10 x 0.25 = 2.5 pts). - Total raw attendance = 10.5 pts. - Final Attendance Score = 10.0 pts. - The excess 0.5 points are halved (0.25) and added to the Midterm score.

Textbooks

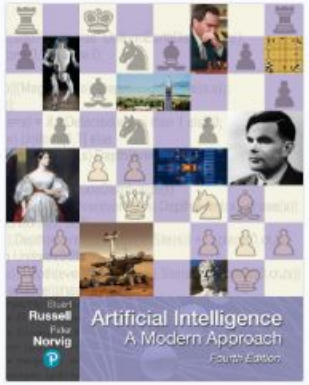

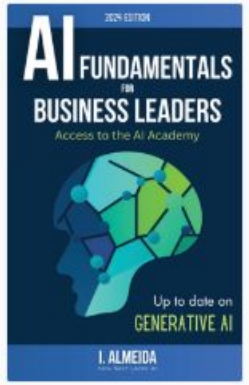

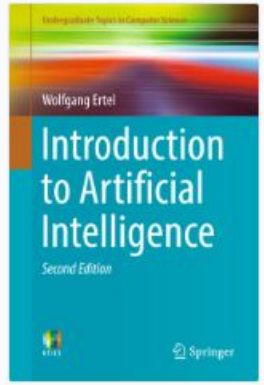

Content

- About
- **The Foundation of AI**
- The History of AI
- The Categories of AI
- The Application of AI
- Careers in AI

The Foundation of AI

What is Artificial Intelligence?

A field of computer science building systems capable of tasks requiring human intelligence.

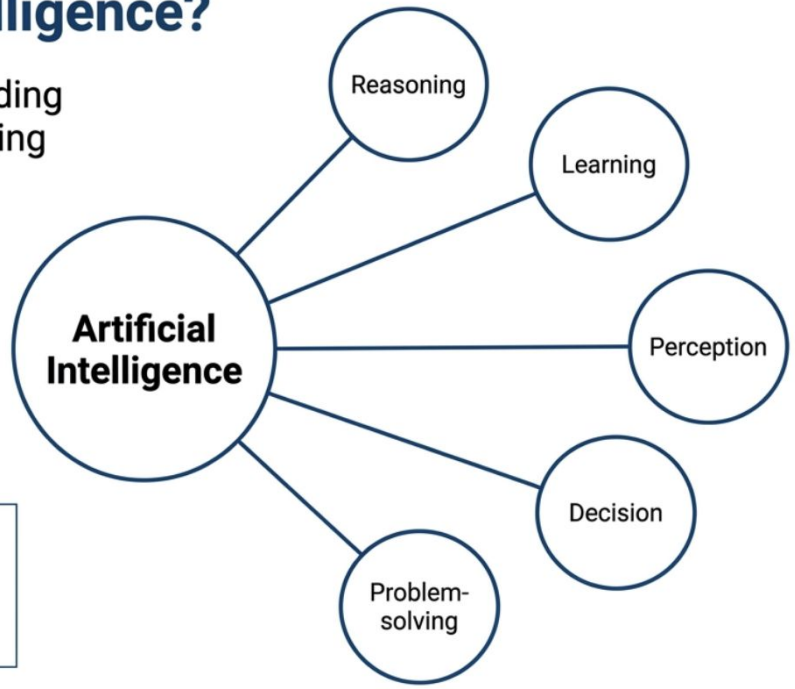

The Original Definition (John McCarthy, 1955)

"The science and engineering of making intelligent machines."

1943
McCulloch–Pitts neuron (First mathematical model of a neuron)

1950
Turing Test proposed (Alan Turing)

1955–56
Logic Theorist created (Newell & Simon)

1956
Dartmouth Workshop (Birth of the field; McCarthy coins AI)

1957
Perceptron (Rosenblatt's first machine learning model)

1966
ELIZA (Weizenbaum's first conversational chatbot)

The Foundation of AI

Logic Theorist (1956)

- Widely considered the first artificial intelligence program.
- Created by Allen Newell and Herbert A. Simon.

The Achievement

Automatically proved 38 of the 52 theorems in chapter 2 of Principia Mathematica.

Note: In some cases, it found shorter, more elegant proofs than the original human authors.

Four Approaches to Defining AI

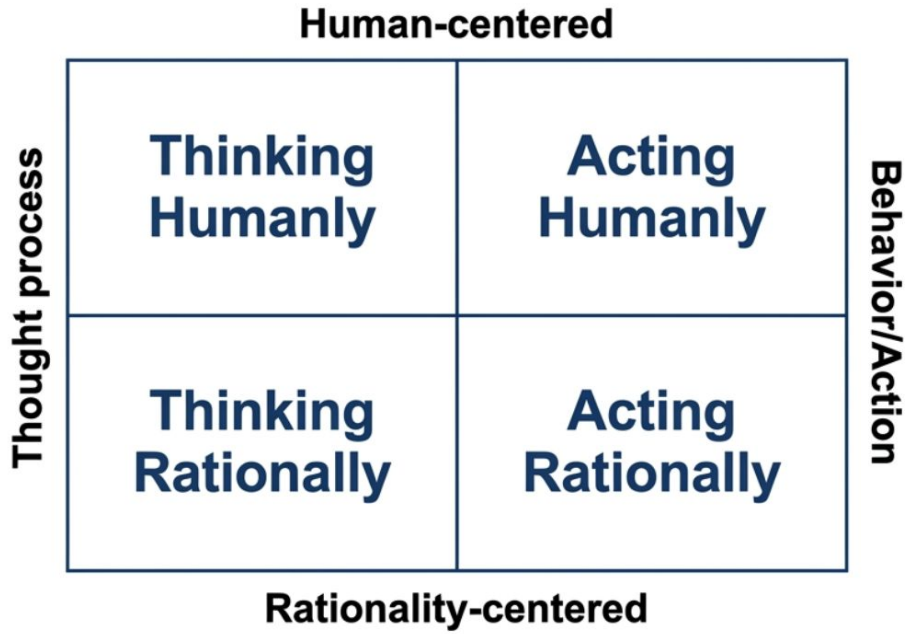

The Foundation of AI

Acting Humanly: The Imitation Game

Core Idea: A machine is intelligent if its behavior is indistinguishable from a human's.

Turing replaced the vague question "Can machines think?"" with a **measurable** behavioral test.

Passing Condition: The interrogator cannot reliably distinguish the machine from the human better than random guessing.

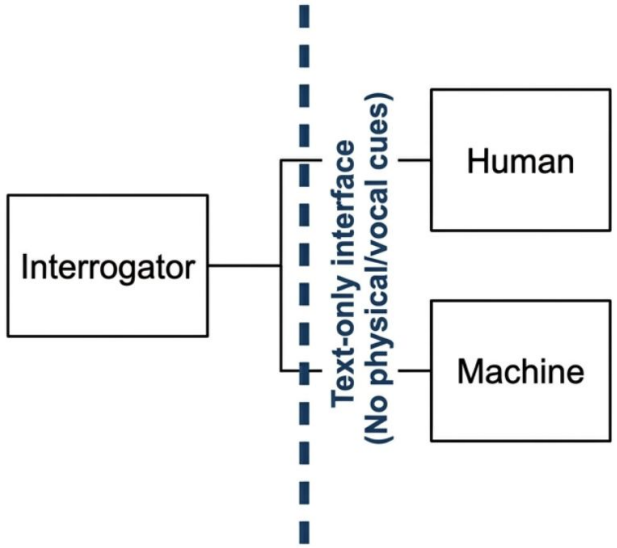

Passing the Turing Test

Standard Turing Test Requirements:

- **1. Natural Language Processing:** Communicate in human language.
- **2. Knowledge Representation:** Store information and knowledge.
- **3. Automated Reasoning:** Draw conclusions from stored data.
- **4. Machine Learning:** Adapt to new circumstances and detect patterns.

Total Turing Test Requirements (adds physical interaction):

- **5. Computer Vision:** Perceive objects.
- **6. Robotics:** Manipulate and move objects.

Combined, these six capabilities encompass most of modern AI.

The Foundation of AI

Thinking Humanly: Cognitive Science

Core Idea: Systems must not only produce the right output, but do so using human-like internal thought processes.

Merges computer models from AI with experimental techniques from psychology.

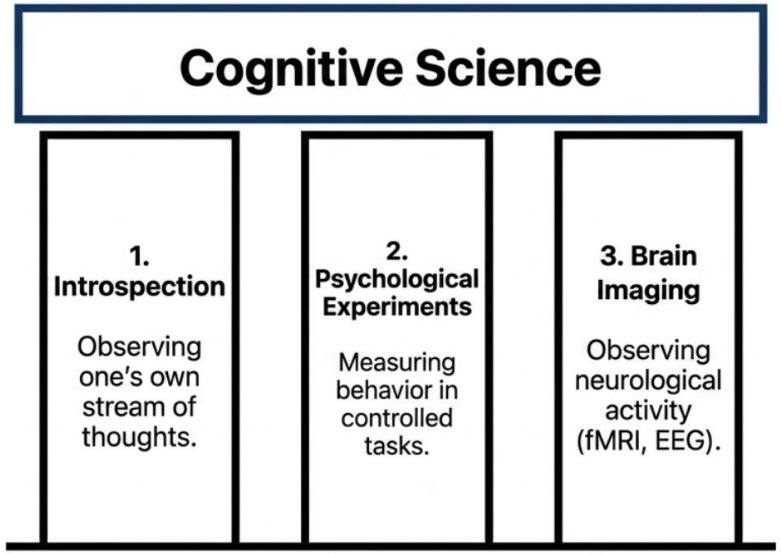

Matching the Reasoning Trace

Example: The General Problem Solver (Newell & Simon).

Merely arriving at the correct answer was insufficient.

The goal was to compare the trace of reasoning steps between the machine and human subjects.

A successful model must pass through the exact same intermediate steps in the same order as a human solving the same problem.

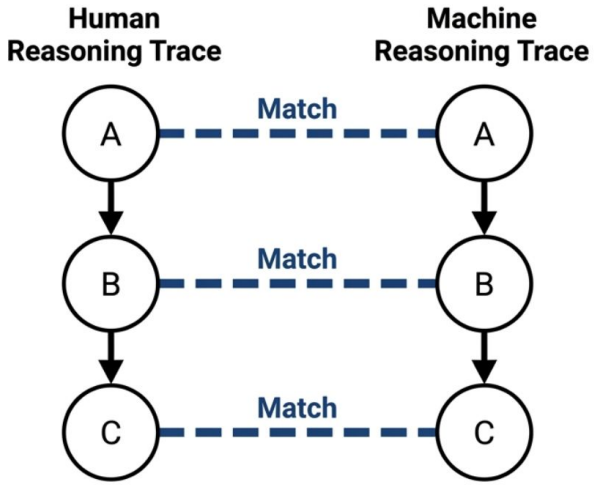

The Foundation of AI

Thinking Rationally: Laws of Thought

Core Idea: Thinking properly means adhering to normative logical rules that guarantee true conclusions from true premises.

Origin: Aristotle's formalization of "right thinking" via the Syllogism.

1 Major Premise: All humans are mortal.

2 Minor Premise: Socrates is a human.

---

3 Conclusion: Therefore, Socrates is mortal.

Validity relies on the structural form, independent of the specific content.

The Logicist Tradition & Its Limits

Aristotle → Boole → Frege/Peano → Logicist AI

The Logicist Hope: Representing all knowledge in formal logic would allow machines to automatically deduce correct conclusions.

1. Uncertain Knowledge

Real-world knowledge is ambiguous. (e.g., "Birds usually fly," but penguins do not). Formal logic demands absolute true/false states.

2. Combinatorial Explosion

In practice, computing logical deductions can exhaust system resources before finding a solution, even if theoretically possible.

The Foundation of AI

Acting Rationally: Intelligent Agents

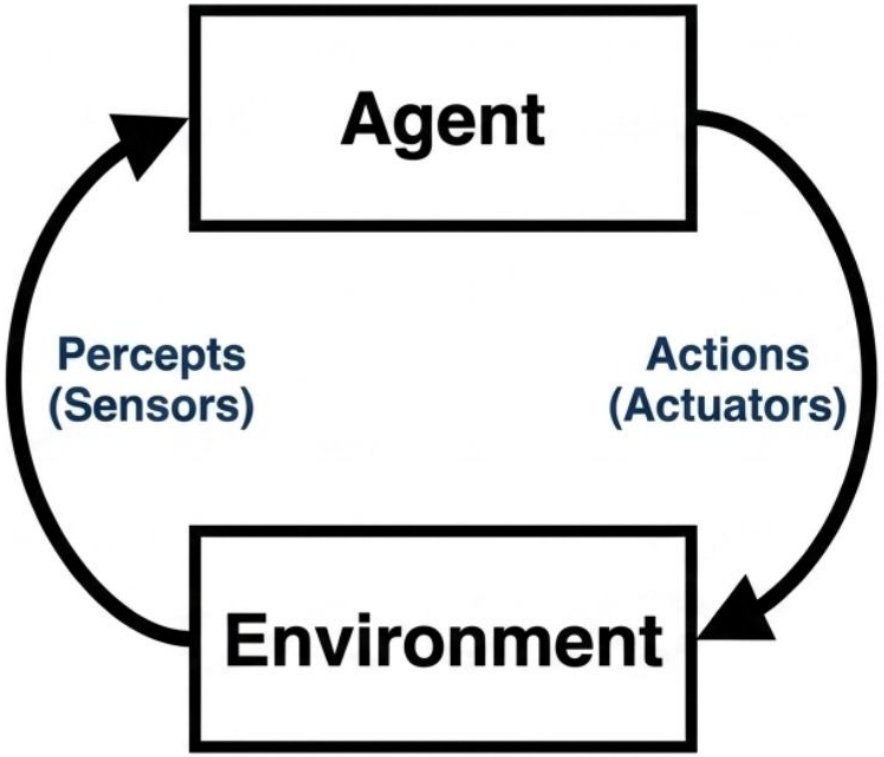

Core Idea: Acting to achieve the best expected outcome, regardless of human similarity or perfect logical deduction.

Rational Agent: Chooses actions that maximize a performance measure, based on:

1. Performance measure (success criteria)
2. Prior knowledge of the environment
3. Possible actions
4. Percept sequence (observations)

The Foundation of AI

Summary & The Value Alignment Problem

Acting Humanly	**Focus**: Behavior	**Benchmark**: Turing Test
Thinking Humanly	**Focus**: Thought	**Benchmark**: Cognitive Science
Thinking Rationally	**Focus**: Thought	**Benchmark**: Formal Logic
Acting Rationally	**Focus**: Behavior	**Benchmark**: Rational Agents

The Danger of Perfect Rationality: If an agent optimizes for the wrong performance measure, highly rational behavior becomes dangerous. This is known as the Value Alignment Problem (Stuart Russell).

Content

- About
- The Foundation of AI
- **The History of AI**
- The Categories of AI
- The Application of AI
- Careers in AI

The History of AI

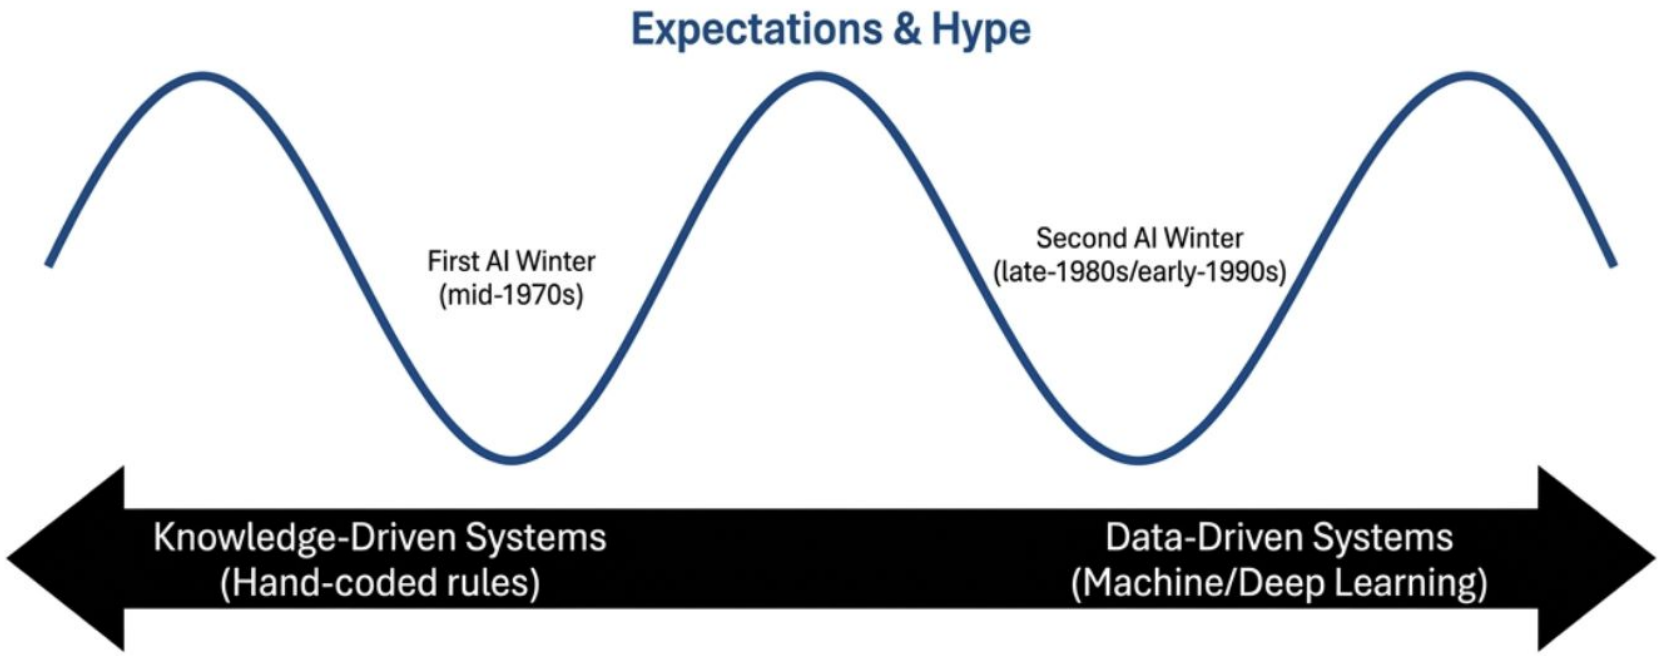

- AI history follows a recurring cycle of hype, unfulfilled expectations, and severe funding collapses (AI Winters).
- Beneath these cycles lies a fundamental structural shift: moving from embedding human expertise to extracting patterns from massive datasets.

The History of AI

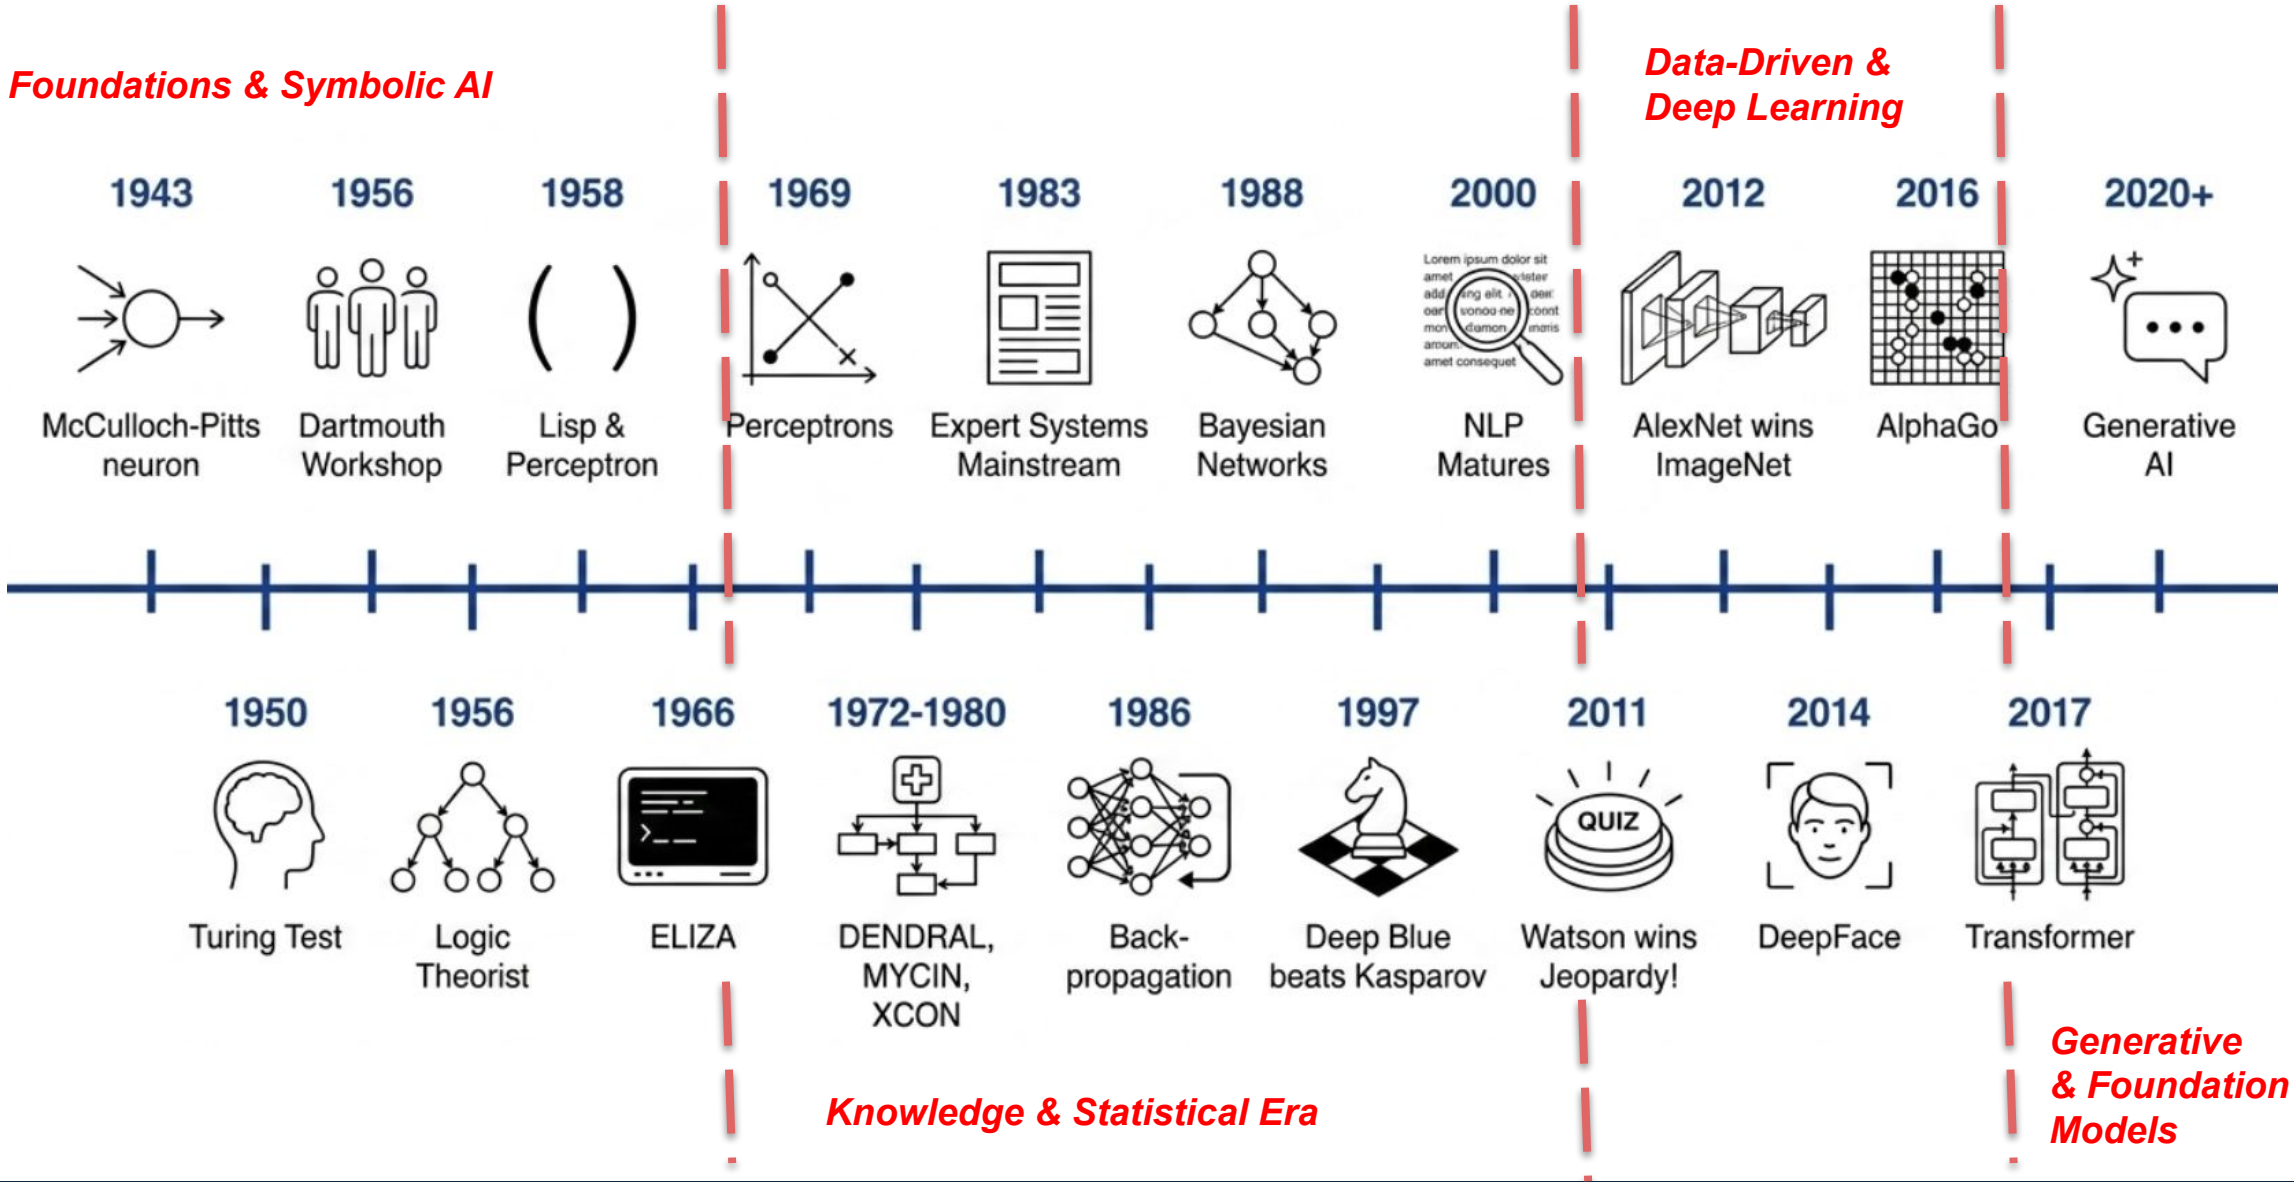

The History of AI

Foundations & Symbolic AI

Phase 1: The Gestation of AI (1943–1955)

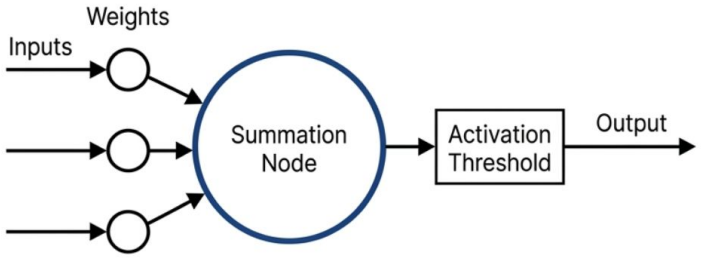

- **1943: McCulloch & Pitts** propose the first artificial neuron; prove neural networks can compute logical functions.
- **1949: Hebb** introduces Hebbian learning for updating connection weights.
- **1950:** Turing proposes the Turing Test, machine learning, and reinforcement learning.
- **1951:** Minsky & Edmonds construct SNARC, the first neural-network computer.

1943 1949 1950 1951

Phase 2: The Birth of a Field (1956)

The Dartmouth Workshop

Organized by John McCarthy, Marvin Minsky, Claude Shannon, and Nathaniel Rochester.

Officially coins the term "Artificial Intelligence".

The Logic Theorist

Presented by Allen Newell and Herbert Simon.

Recognized as the first true AI program, capable of autonomously proving mathematical theorems.

1956

Phase 3: Great Expectations (1952–1969)

Innovation	Significance
General Problem Solver	Simulates human problem-solving search.
Lisp (1958)	John McCarthy's creation; dominates AI programming for 30 years.
Perceptron (1958)	Frank Rosenblatt designs the first learnable neural model.
Checkers Program (1959)	Arthur Samuel achieves strong amateur-level machine learning.
SHRDLU / ELIZA	Demonstrates language understanding in limited domains (microworlds).

1958 1959 1966

The History of AI

Knowledge & Statistical Era

Phase 4: A Dose of Reality & The First Winter (1966–1973)

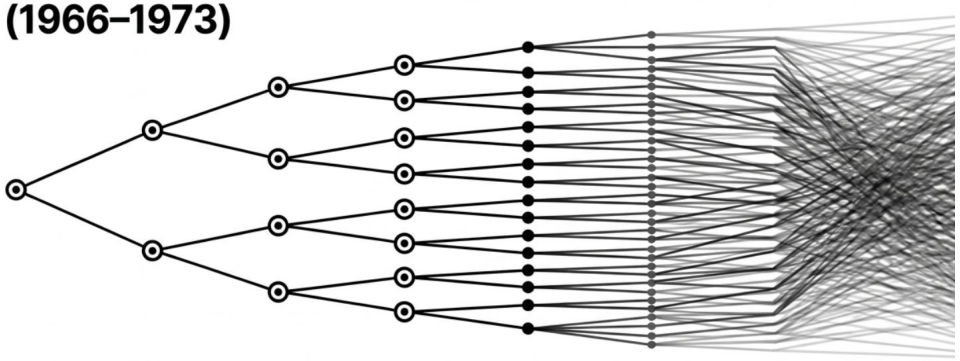

- Early methods fail to scale from toy microworlds to complex, real-world problems.
- **Intractability** (combinatorial explosion) halts general-purpose search algorithms.
- **1966**: **ALPAC** report terminates machine translation funding due to systemic failures.
- **1969**: Minsky & Papert publish **Perceptrons**, proving single-layer networks cannot solve the **XOR** problem. Neural network funding freezes.

1966 1969

Phase 5: Knowledge-Based Systems (1969–1986)

[Knowledge Base] → [Inference Engine] → [Output]

- Paradigm shift from general-purpose search to embedding deep, domain-specific human knowle
- **DENDRAL**: Infers molecular structures from mass spectrometry (first successful expert system).
- **MYCIN**: Diagnoses blood infections at expert levels, pioneering certainty factors for uncertainty.
- **XCON (R1)**: Saves DEC tens of millions annually, sparking a commercial boom in the early 1980s.
- **1981**: Japan's Fifth Generation project triggers a massive global AI investment race.

1972-1980 1981 1983

Phase 6: The Connectionist Revival (1986–Present)

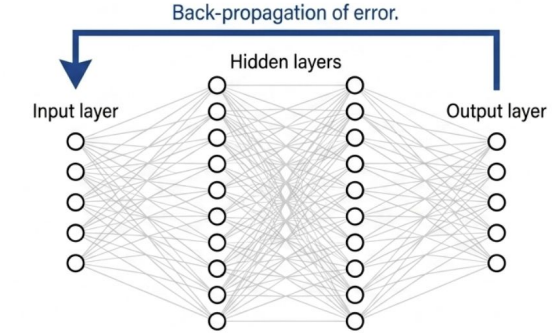

- **1986**: The PDP group (Rumelhart, Hinton, McClelland) rediscovers and popularizes the back-propagation algorithm.
- Solves the limitations identified by Minsky & Papert in 1969.
- Enables the successful training of multi-layer networks (hidden layers).
- Re-establishes Connectionism as a viable counterpoint to symbolic, rule-based AI.

1986

Phase 7: AI Adopts the Scientific Method (1987–Present)

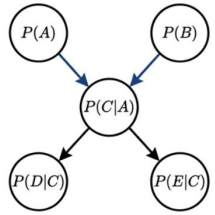

- The field shifts toward rigorous mathematical foundations: statistics, decision theory, and optimization.
- 1988: Judea Pearl introduces Bayesian networks, revolutionizing reasoning under uncertainty.
- The Rational Agent model emerges as a unifying framework across disciplines.
- Focus moves to reproducible experiments, standardized benchmarks, and quantitative evaluation.

1988

The History of AI

Data-Driven & Deep Learning => Gen AI

Phase 8: The Availability of Big Data (2001–Present)

[Massive Web Datasets] × [GPU Compute Power] = Performance Breakthroughs

- AI researchers realize that raw data matters just as much as algorithm design.
- The unreasonable effectiveness of data drives a paradigm shift.
- Internet-scale accumulation of text, images, and user data provides the fuel necessary to train highly parameterized models.

Phase 9: The Deep Learning Revolution (2011–Present)

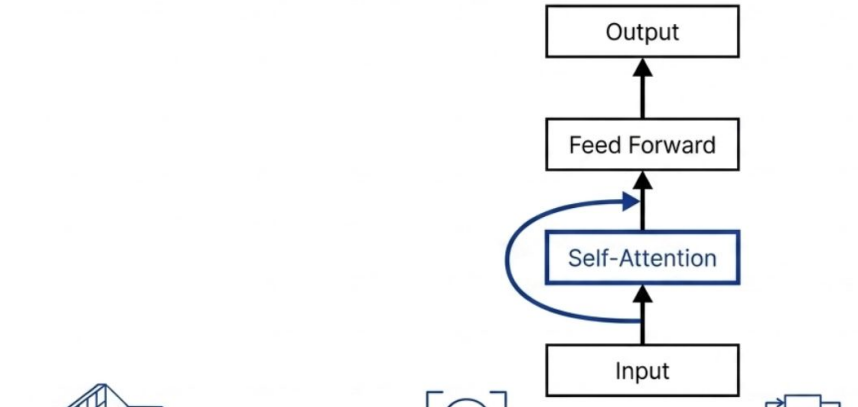

- 2012: AlexNet (Krizhevsky, Sutskever, Hinton) wins ImageNet, proving the dominance of deep convolutional networks.
- 2017: The Transformer architecture (Attention is All You Need) eliminates recurrency, allowing massive parallelization.
- 2018–Present: Rise of Large Language Models (BERT, GPT variants).
- Generative AI booms into mainstream commercial and creative use.

Content

- About
- The Foundation of AI
- The History of AI
- **The Categories of AI**
- The Application of AI
- Careers in AI

The Categories of AI

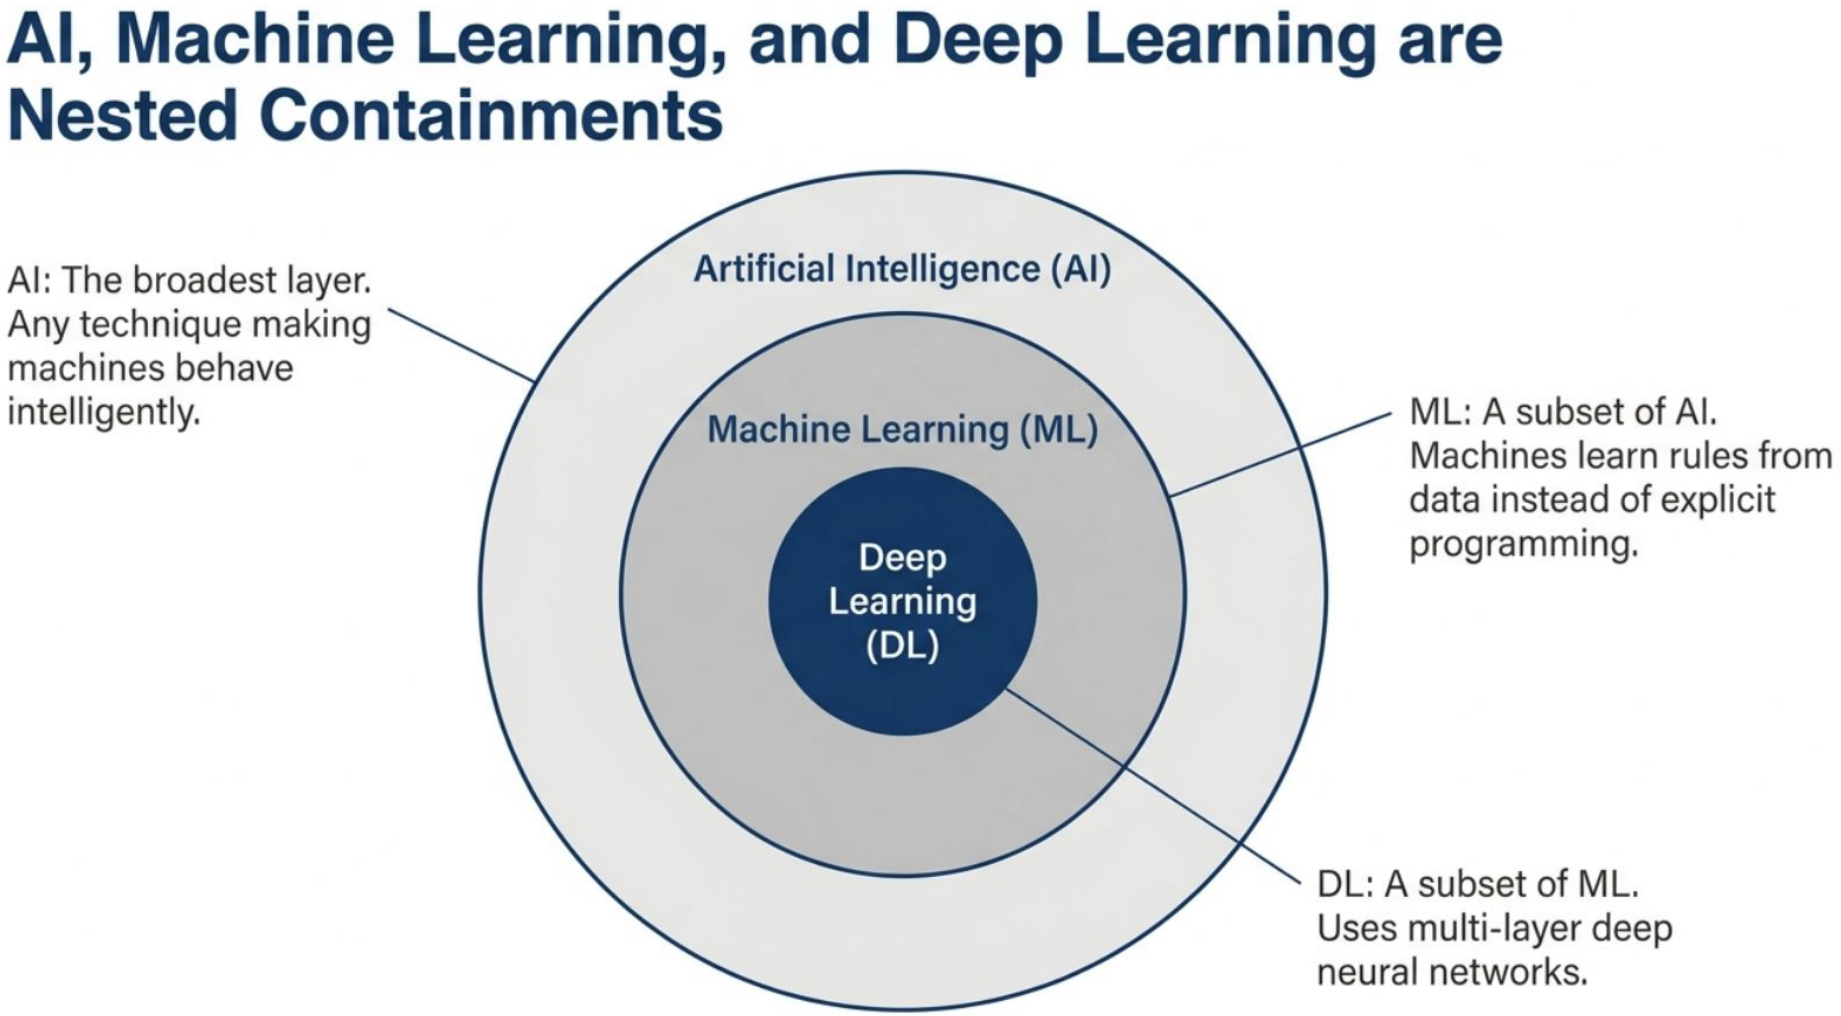

The Categories of AI

Artificial Intelligence (AI)

- Machines performing tasks requiring intelligence.
- May rely on human-written rules and logic.
- Examples include search algorithms and expert systems.

Machine Learning (ML)

- Machines learning patterns directly from data.
- Requires no explicit human-written rule programming.
- Utilizes statistical algorithms.

Deep Learning (DL)

- Machine learning executed via deep neural networks.
- Internal representations are built across many hidden layers.
- Features are learned automatically.

The Categories of AI

Axis 1: Classification by Data Modality

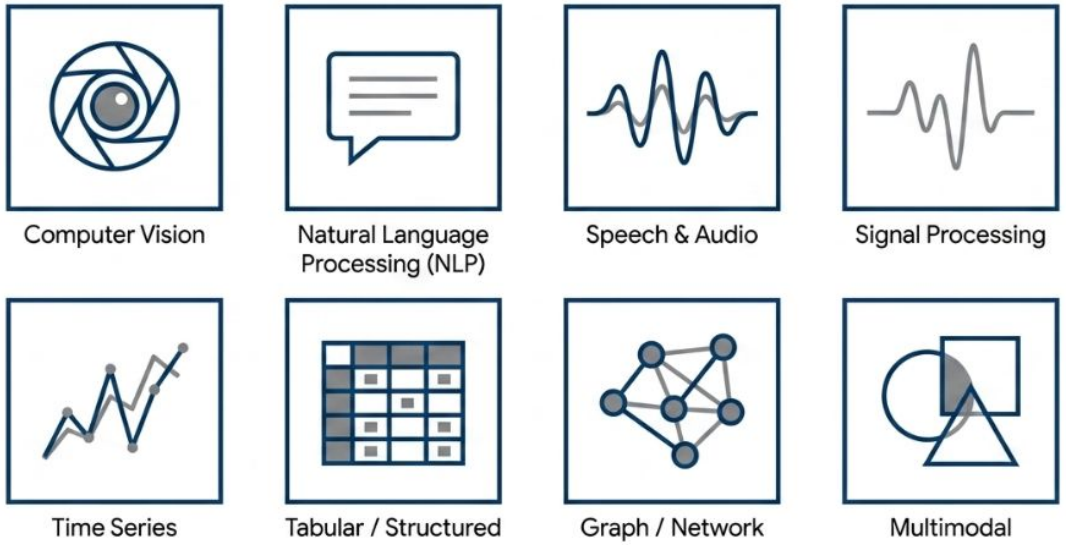

Axis 2: Classification by Research Methodology

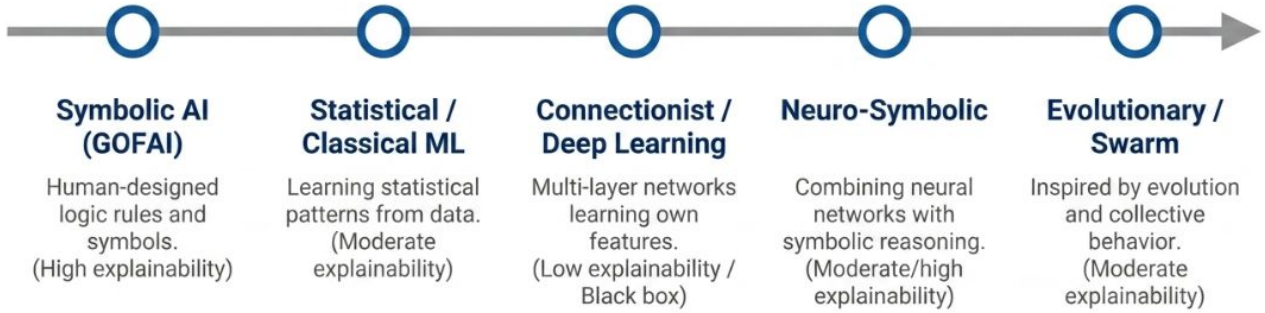

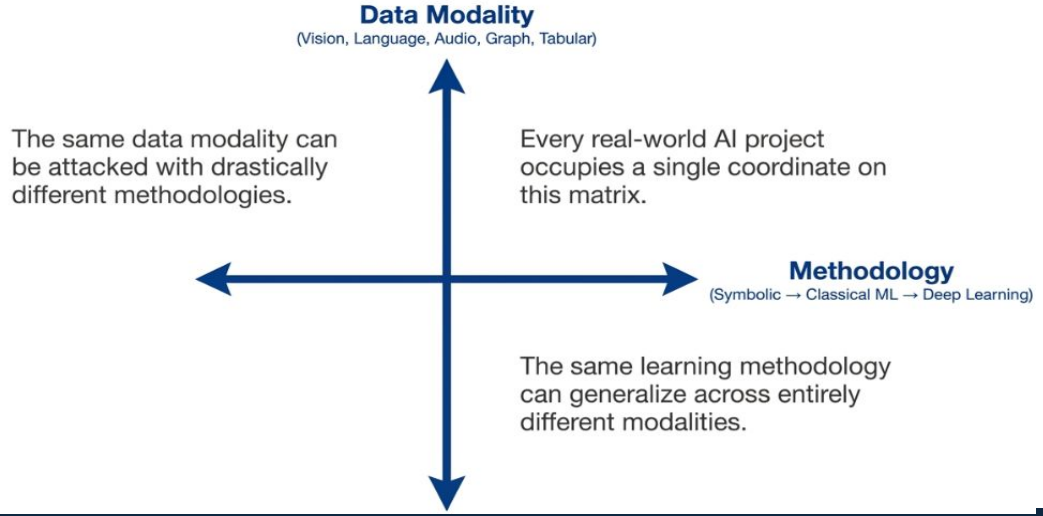

The Categories of AI

Classification by output objectives

Predictive / Analytical AI

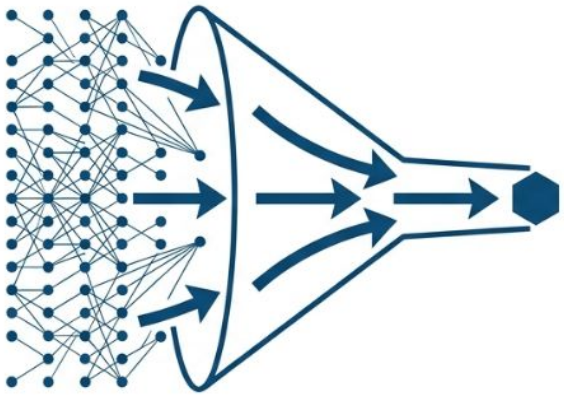

Compresses rich input into a compact judgment.

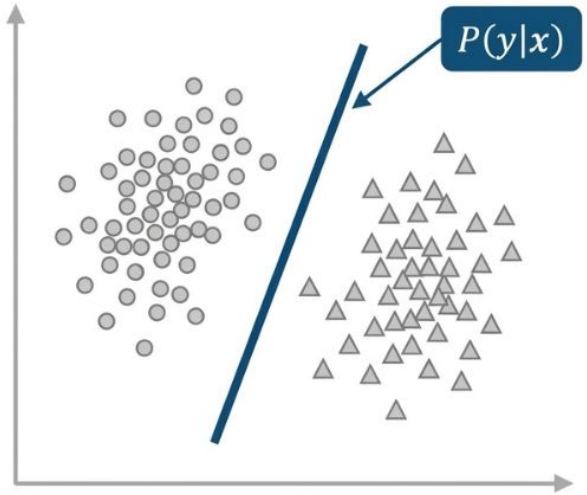

Learns **conditional** probability.

Calculates probability of label y given input x.

Focuses on separating classes, not understanding the whole.

Generative AI

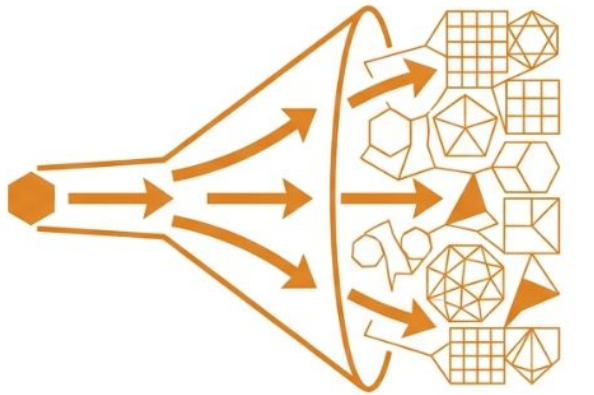

Expands a compact input into a rich, novel output.

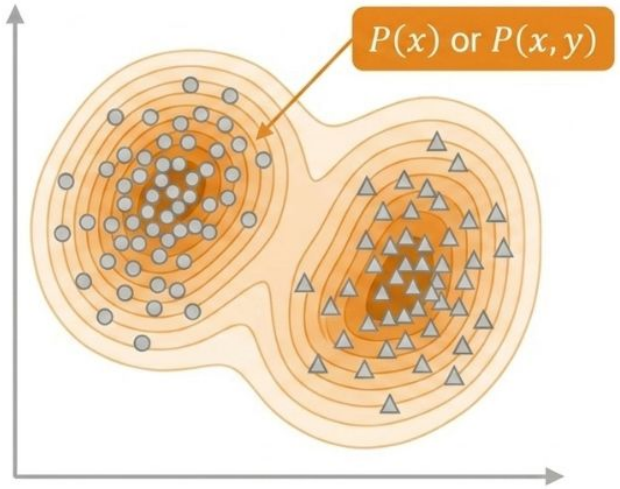

Learns the underlying distribution.

Maps how the data is structured as a whole.

Allows for sampling to produce entirely new instances.

The Categories of AI

Predictive / Analytical AI

Cancer Detection

Stock Forecasting

Fraud Detection

Generative AI

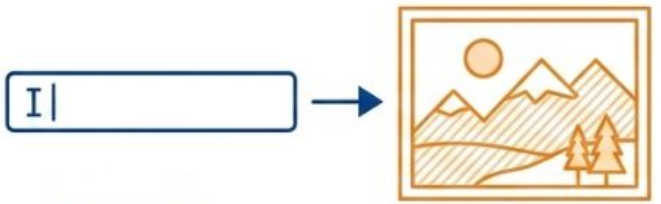

Image Generation

Document Writing

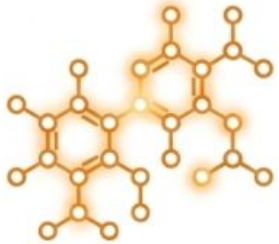

Molecule Design

The Categories of AI

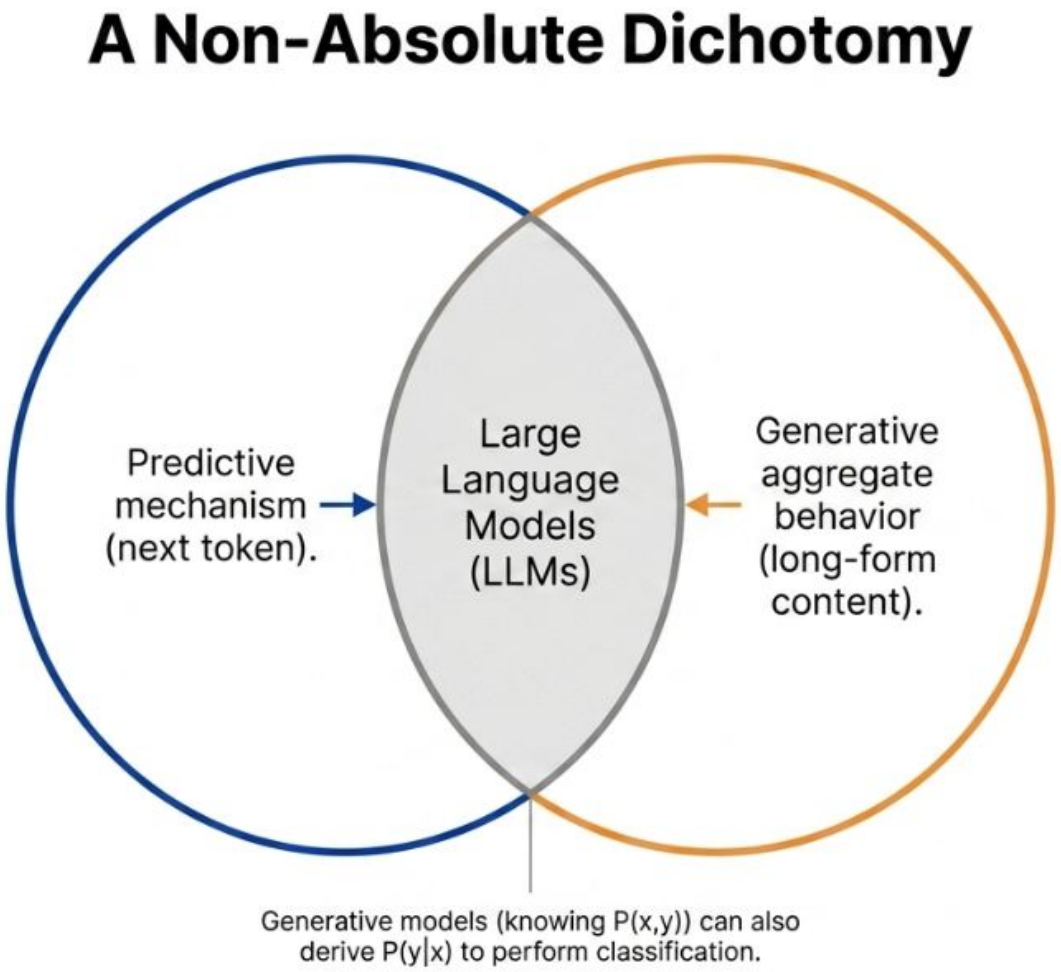

Content

- About
- The Foundation of AI
- The History of AI
- The Categories of AI
- **The Application of AI**
- Careers in AI

The Application of AI

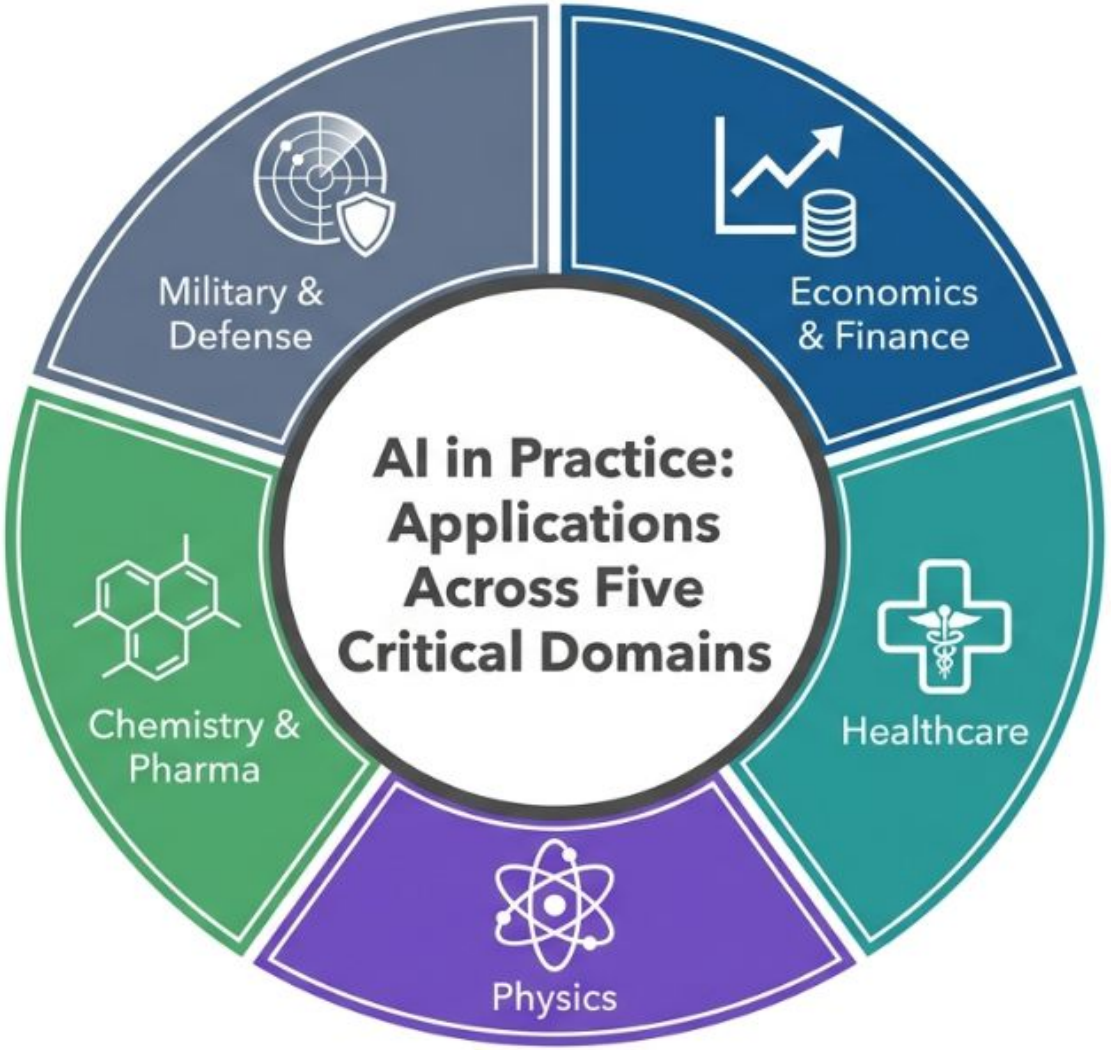

The Application of AI

Economics & Finance

Operates in milliseconds for networks like Visa/Mastercard.

Heavily regulated; applicants possess legal right to explanation.

The Application of AI

Healthcare: Modality Scaling

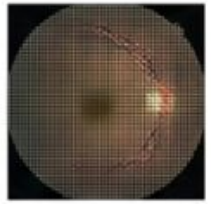

Image Classifier

Severity Grade (1-5)

Population-level screening where specialists are scarce.

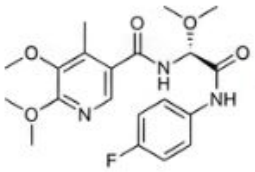

Graph Neural Network

Hazard/Toxicity Probability

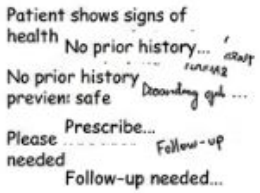

Generative LM

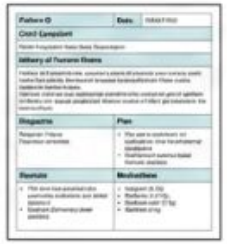

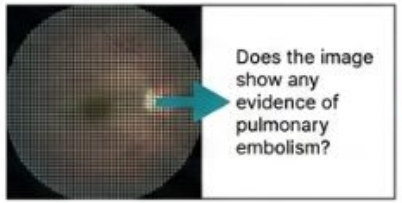

Multimodal System

The Application of AI

Healthcare: Modality Scaling

Medical Imaging Diagnosis: The Explainability Layer

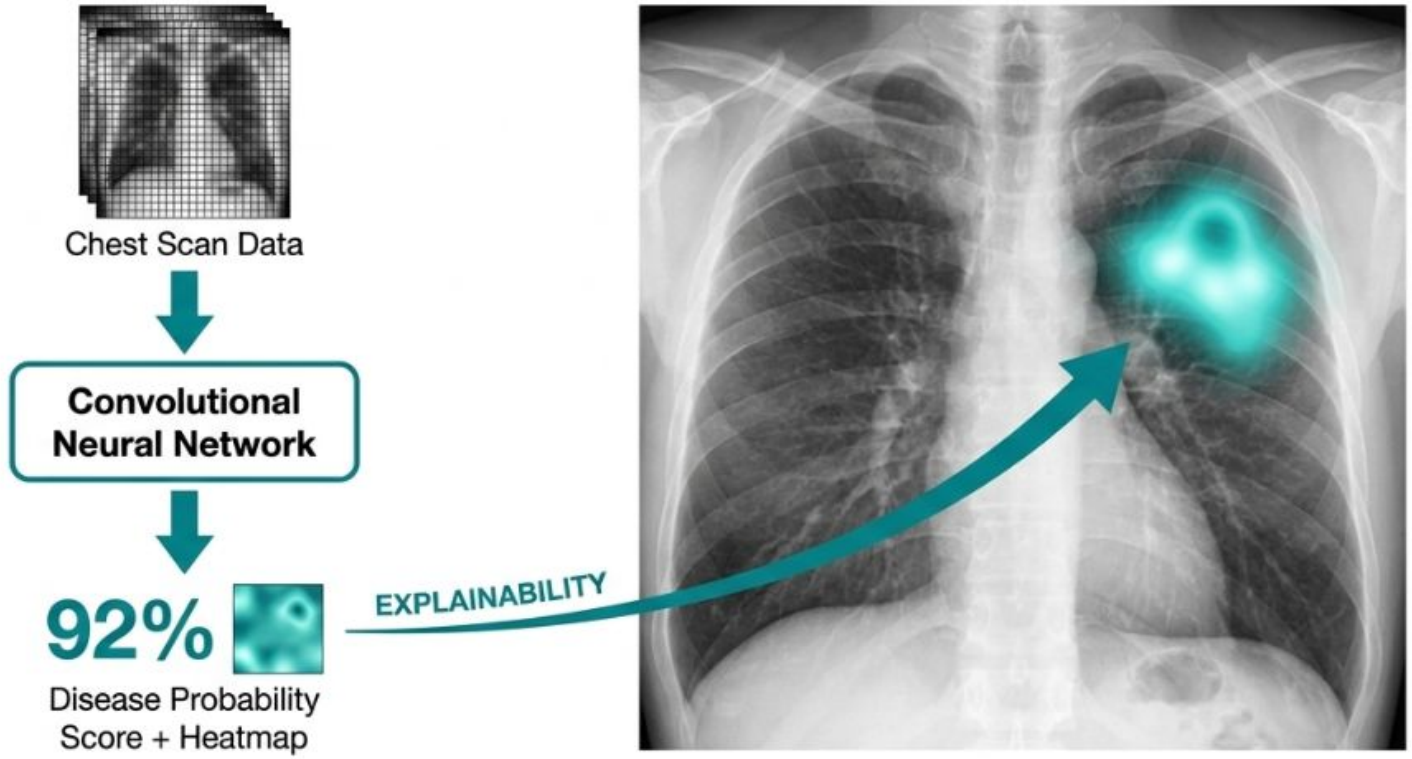

A raw probability score is insufficient for a clinician. The heatmap translates a black-box decision into a usable clinical insight by highlighting the exact anatomical region that drove the prediction.

The Application of AI

Healthcare: Modality Scaling

Protein Structure Prediction: Dimensional Expansion
Dimensional Expansion

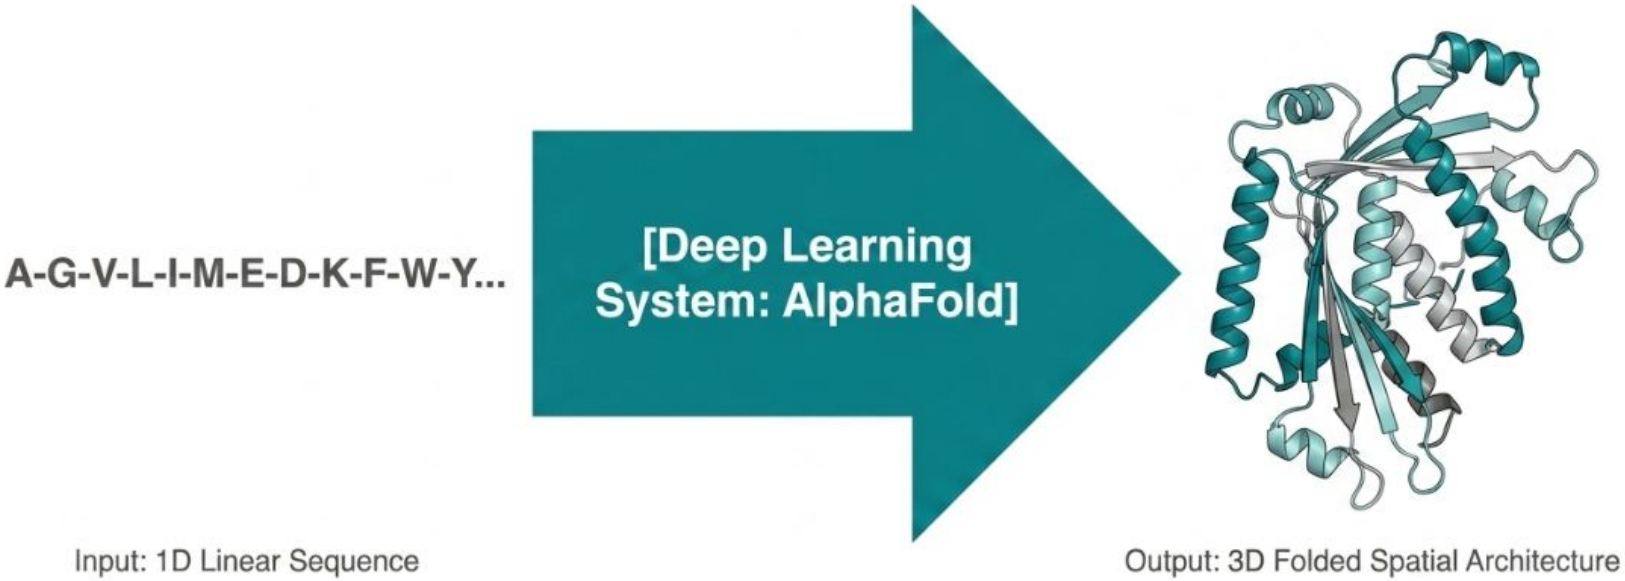

A striking example of algorithmic dimensional expansion, where a flat linear sequence is translated directly into functional, three-dimensional spatial geometry.

The Application of AI

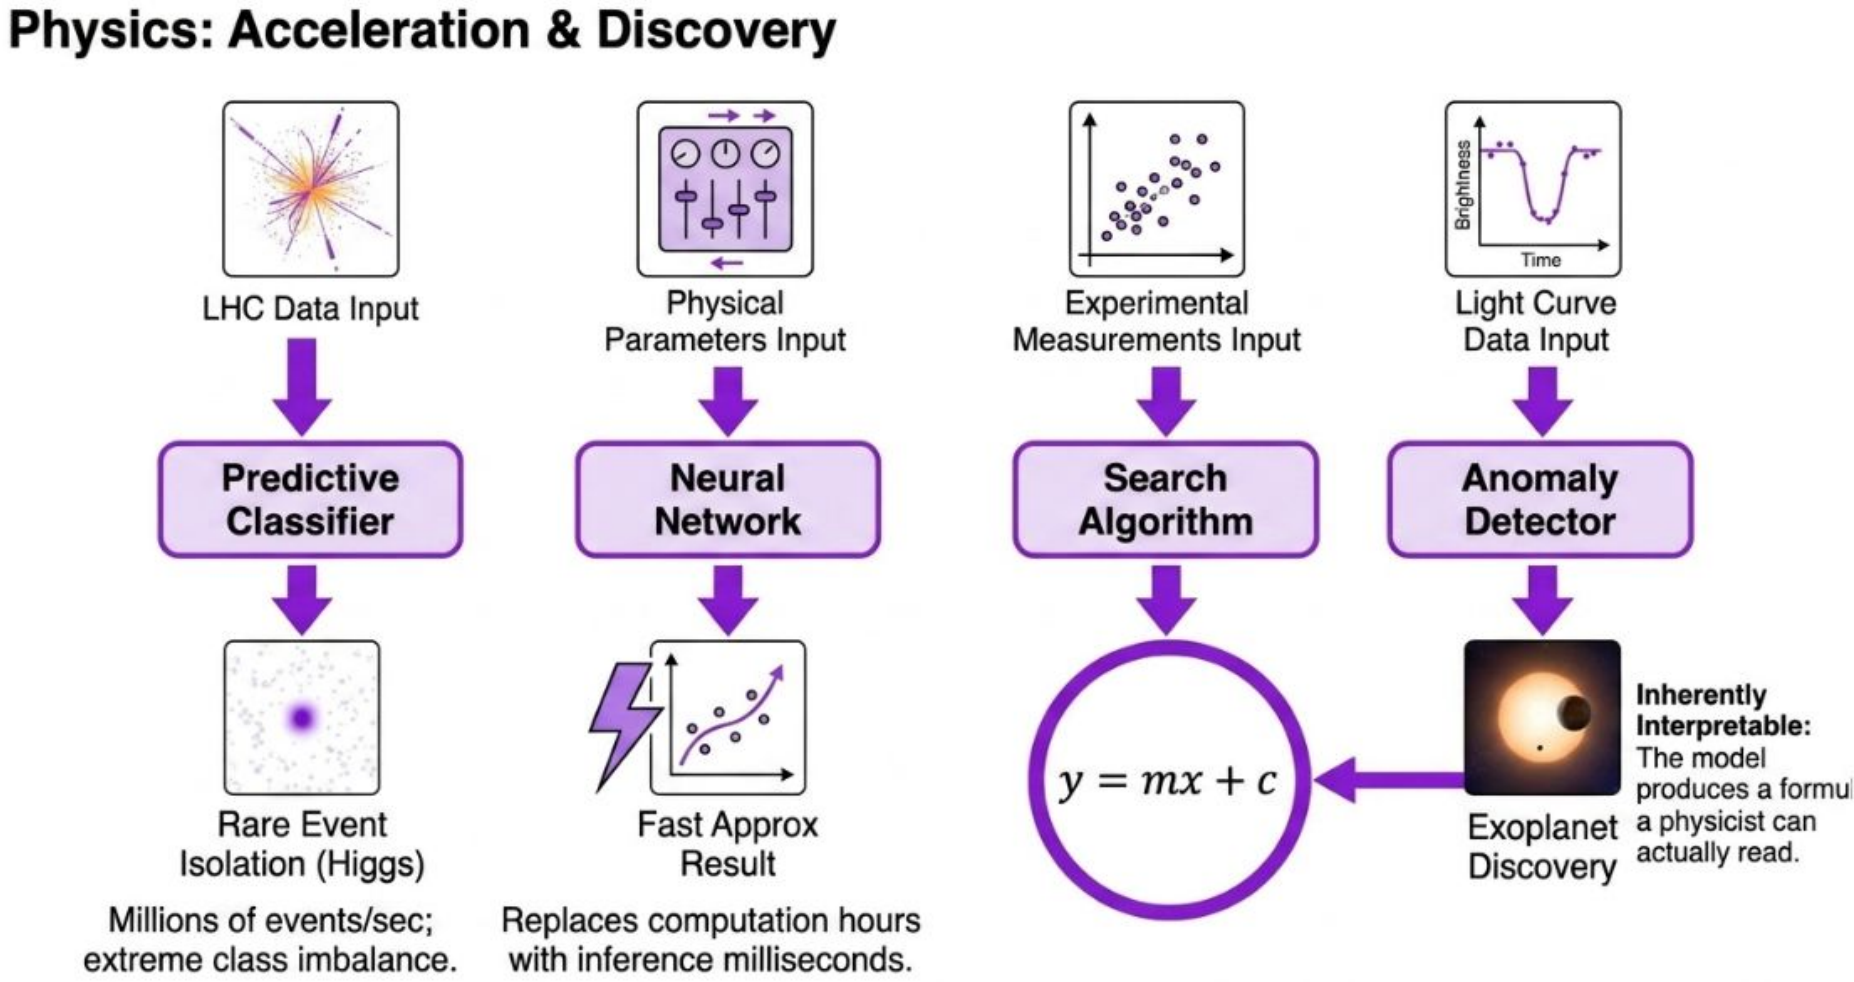

The Application of AI

Physics: Acceleration & Discovery

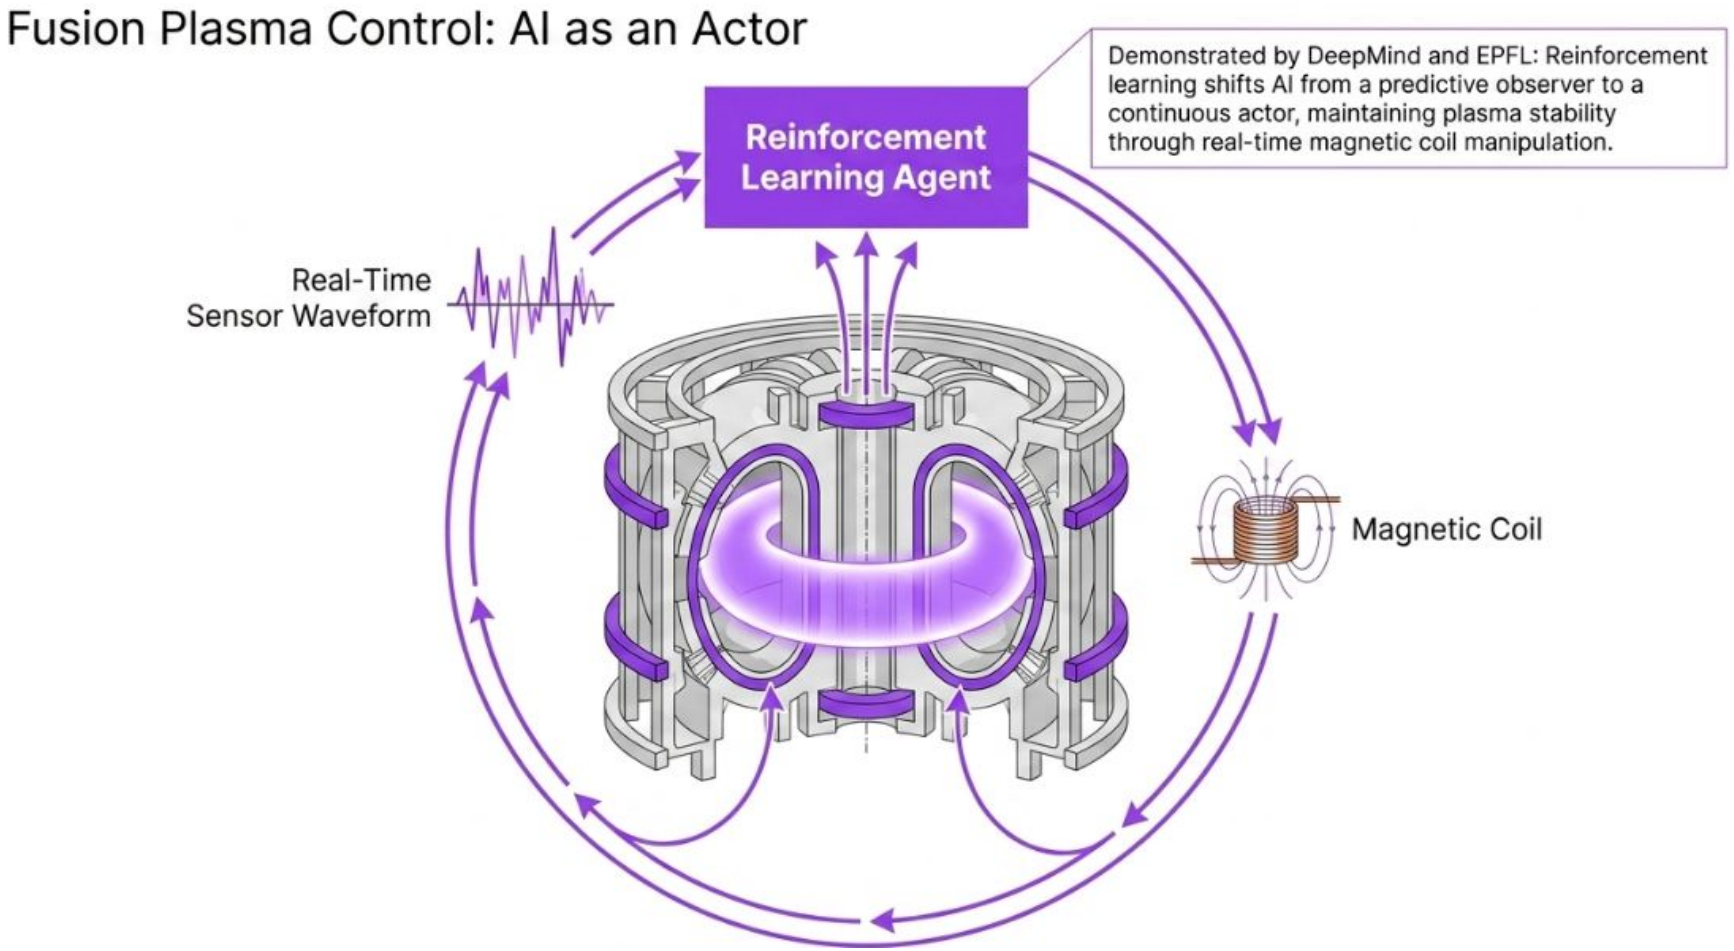

The Application of AI

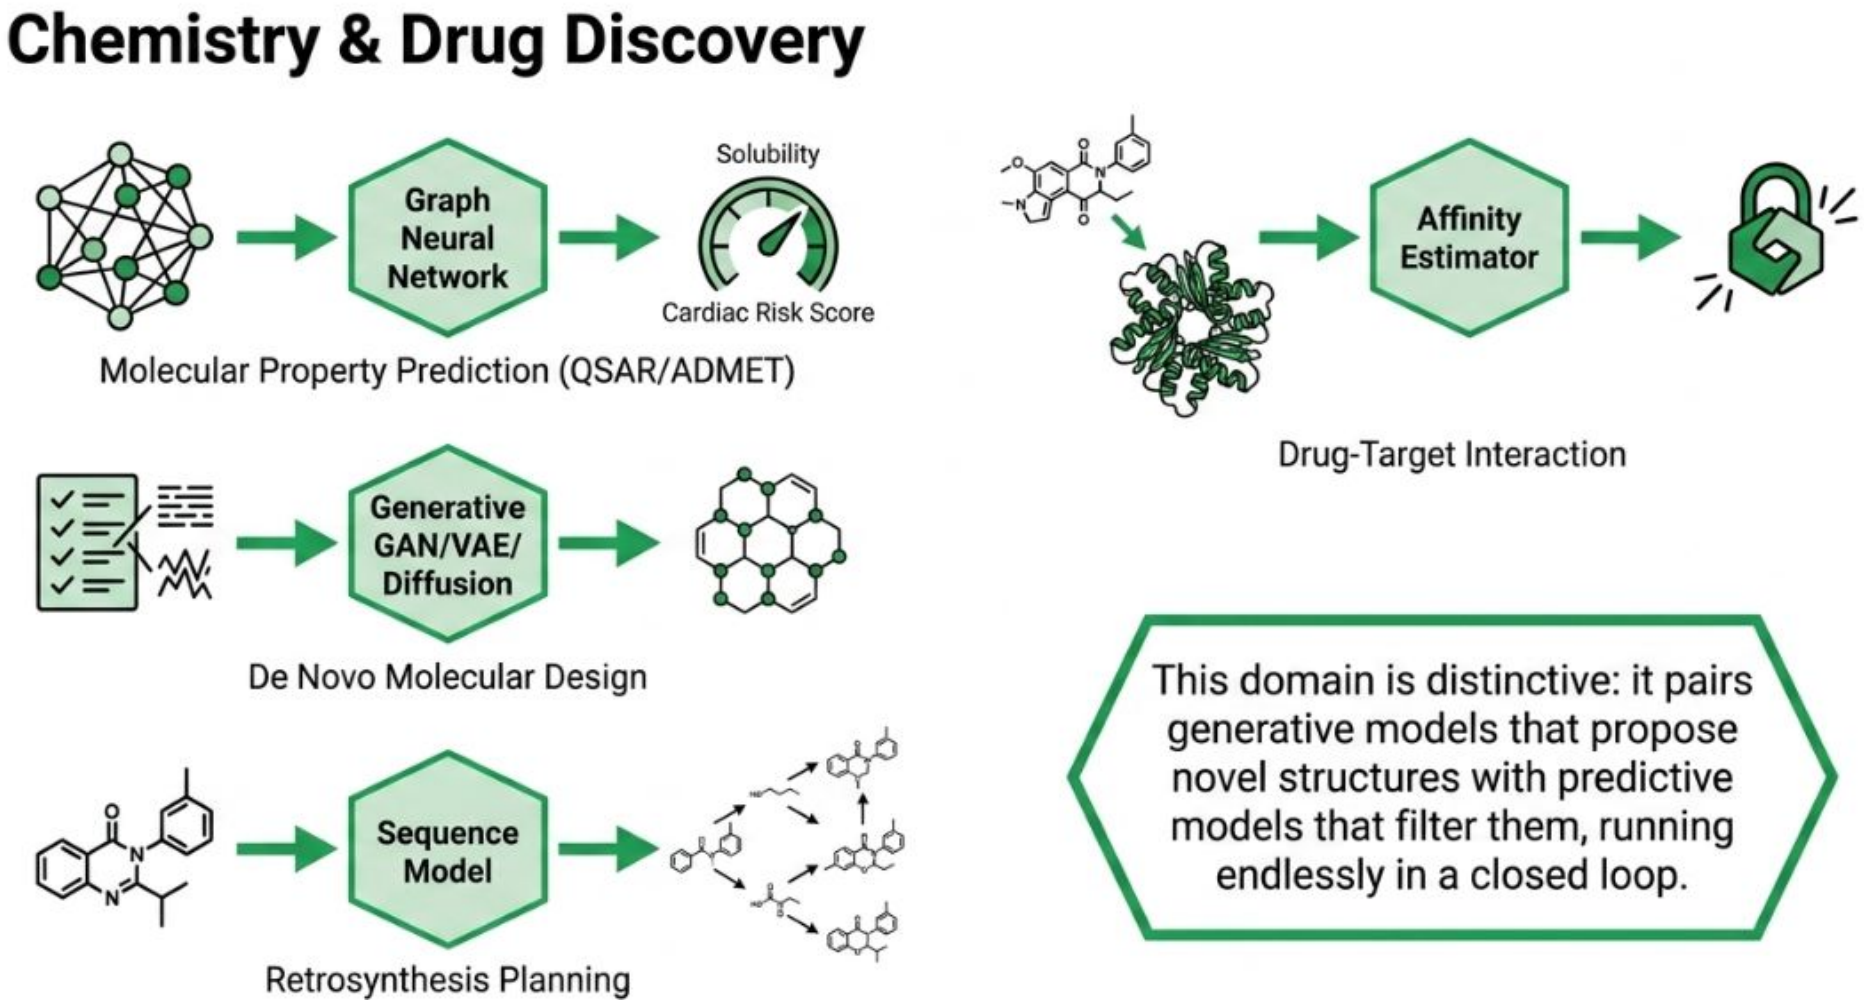

The Application of AI

Chemistry & Drug Discovery

Accelerating Discovery: Loops and Funnels

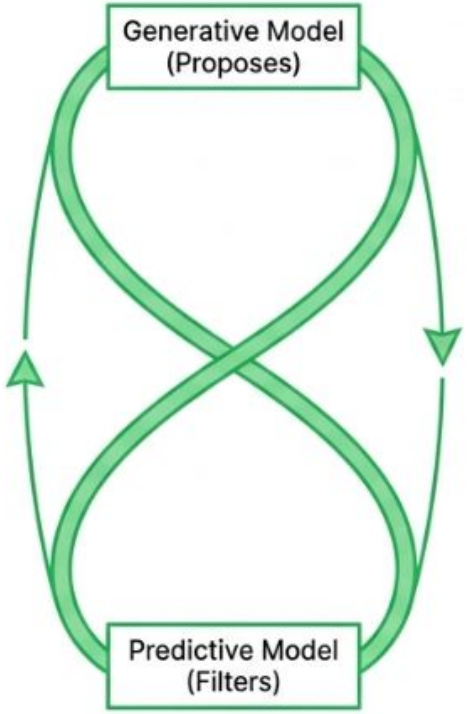

Drug discovery uniquely pairs generative and predictive paradigms in a continuous, automated closed loop.

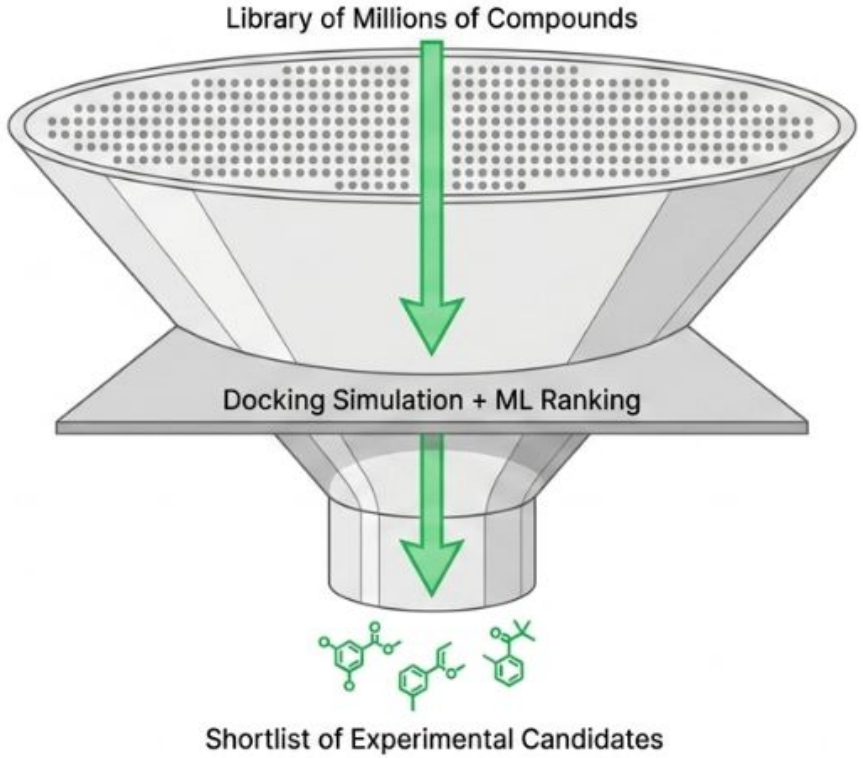

Virtual screening reduces an otherwise intractable search space to a manageable shortlist of experimental candidates.

The Application of AI

Chemistry & Drug Discovery

The Explainability Imperative in Pharma

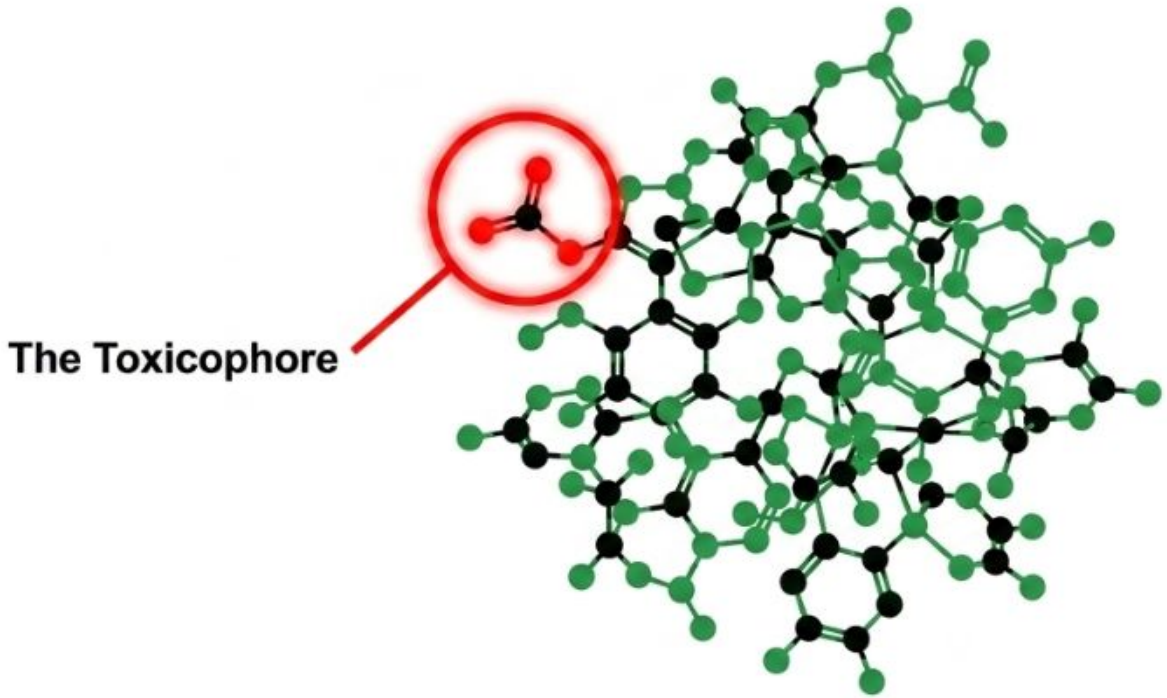

A toxic prediction (e.g., hERG inhibition) is a dead end. Identifying the specific responsible substructure—the toxicophore—converts an opaque probability score into actionable insight that guides molecular redesign.

The Application of AI

Military & Defense

Autonomous Systems

Electronic Warfare

Cyber Defense

Logistics & Wargaming

The Application of AI

Military & Defense

ISR and the Precondition of Oversight

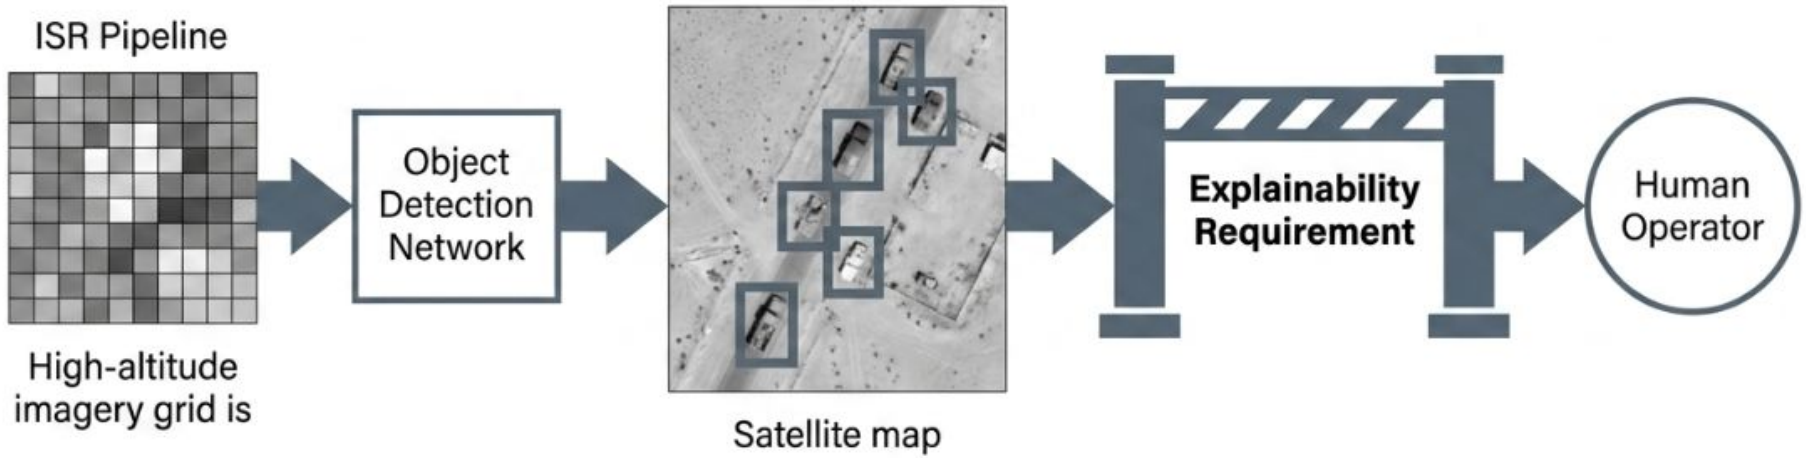

Lethal autonomous systems raise unresolved questions about accountability. Keeping a human in the decision loop depends entirely on that operator's ability to understand why the system recommended a target. Explainability is not an academic nicety—it is the precondition for meaningful oversight.

Content

- About
- The Foundation of AI
- The History of AI
- The Categories of AI
- The Application of AI
- **Careers in AI**

Careers in AI

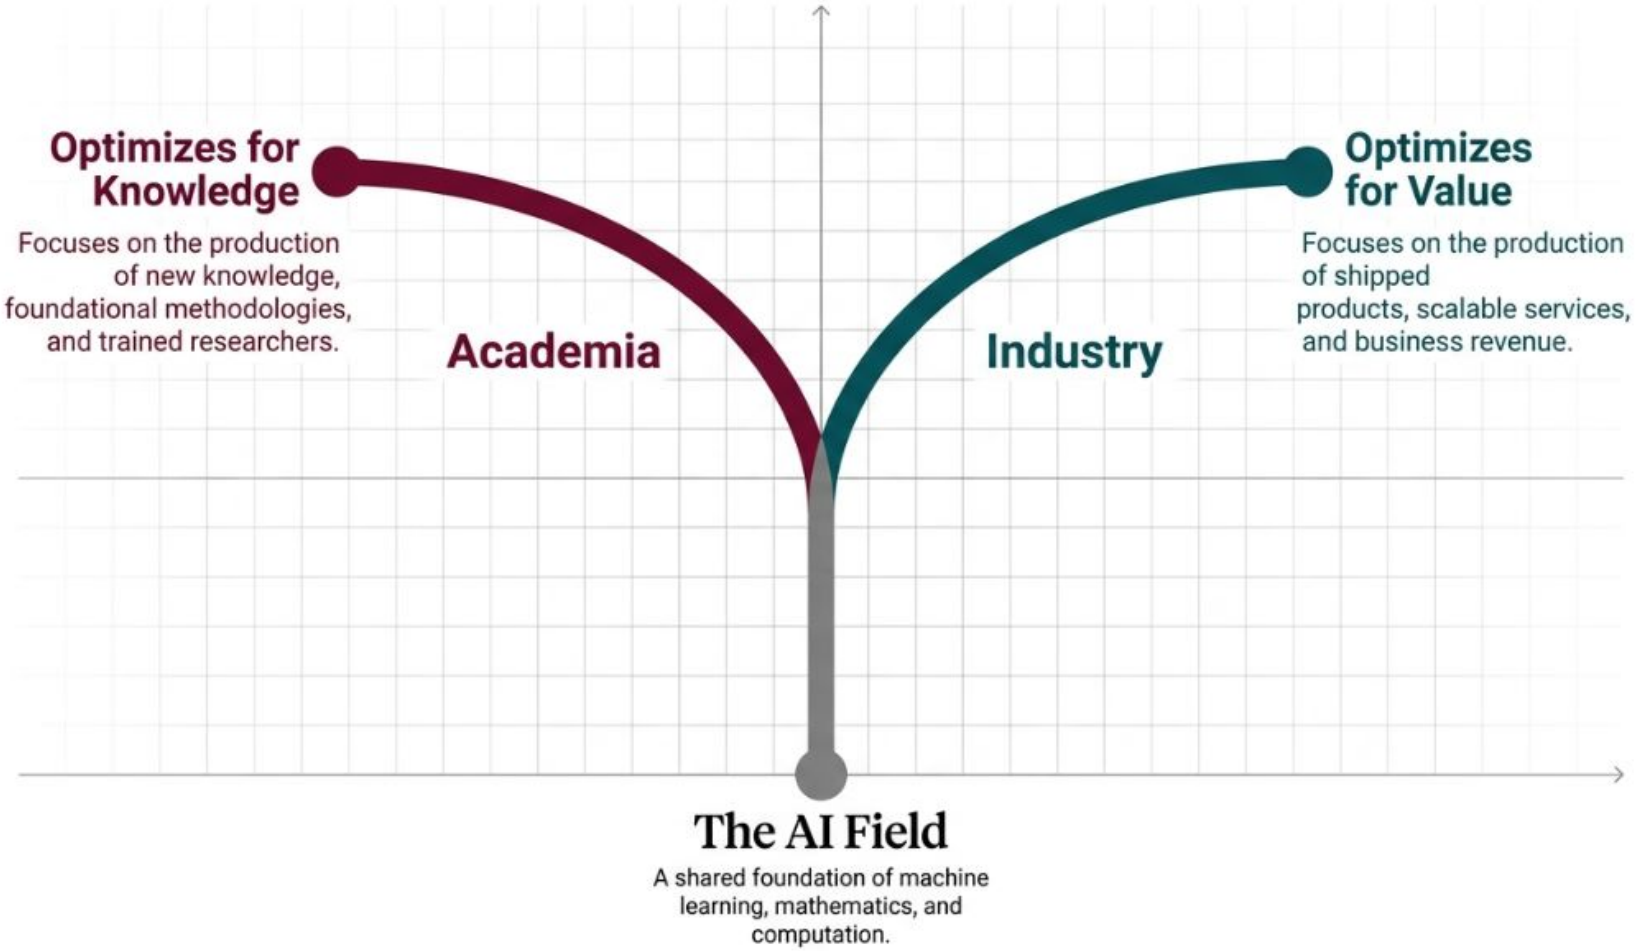

Careers in AI

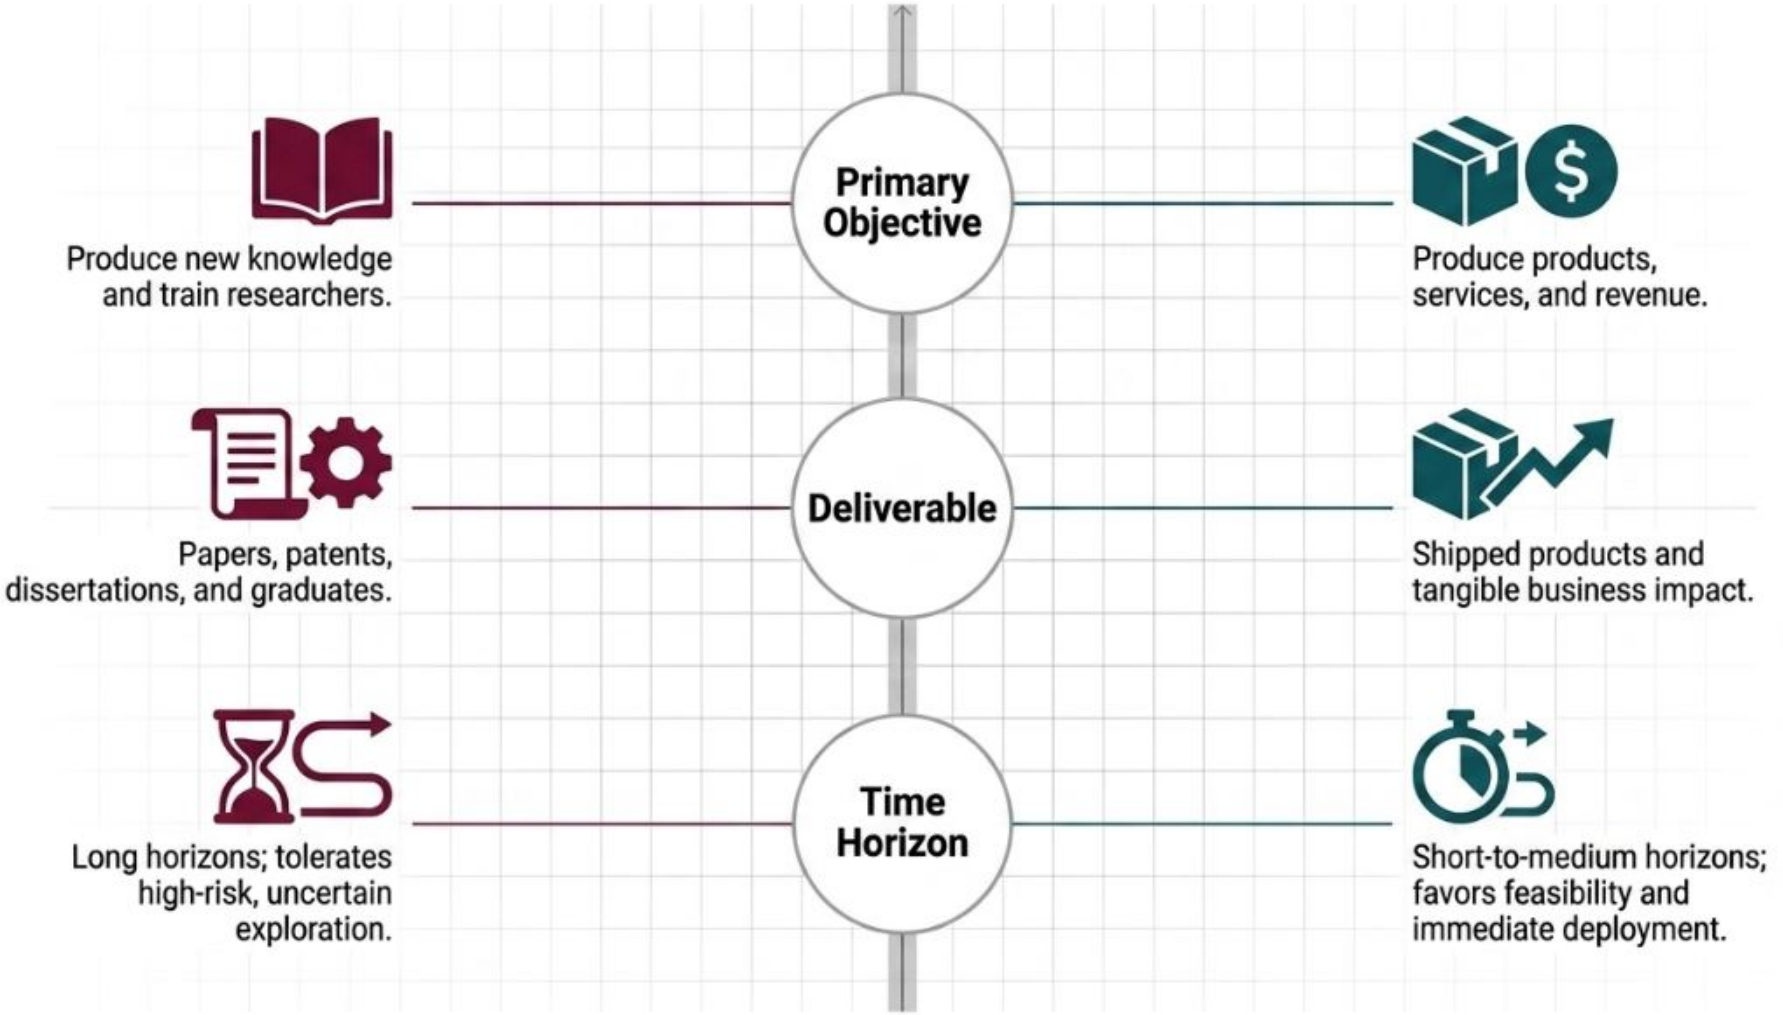

Careers in AI

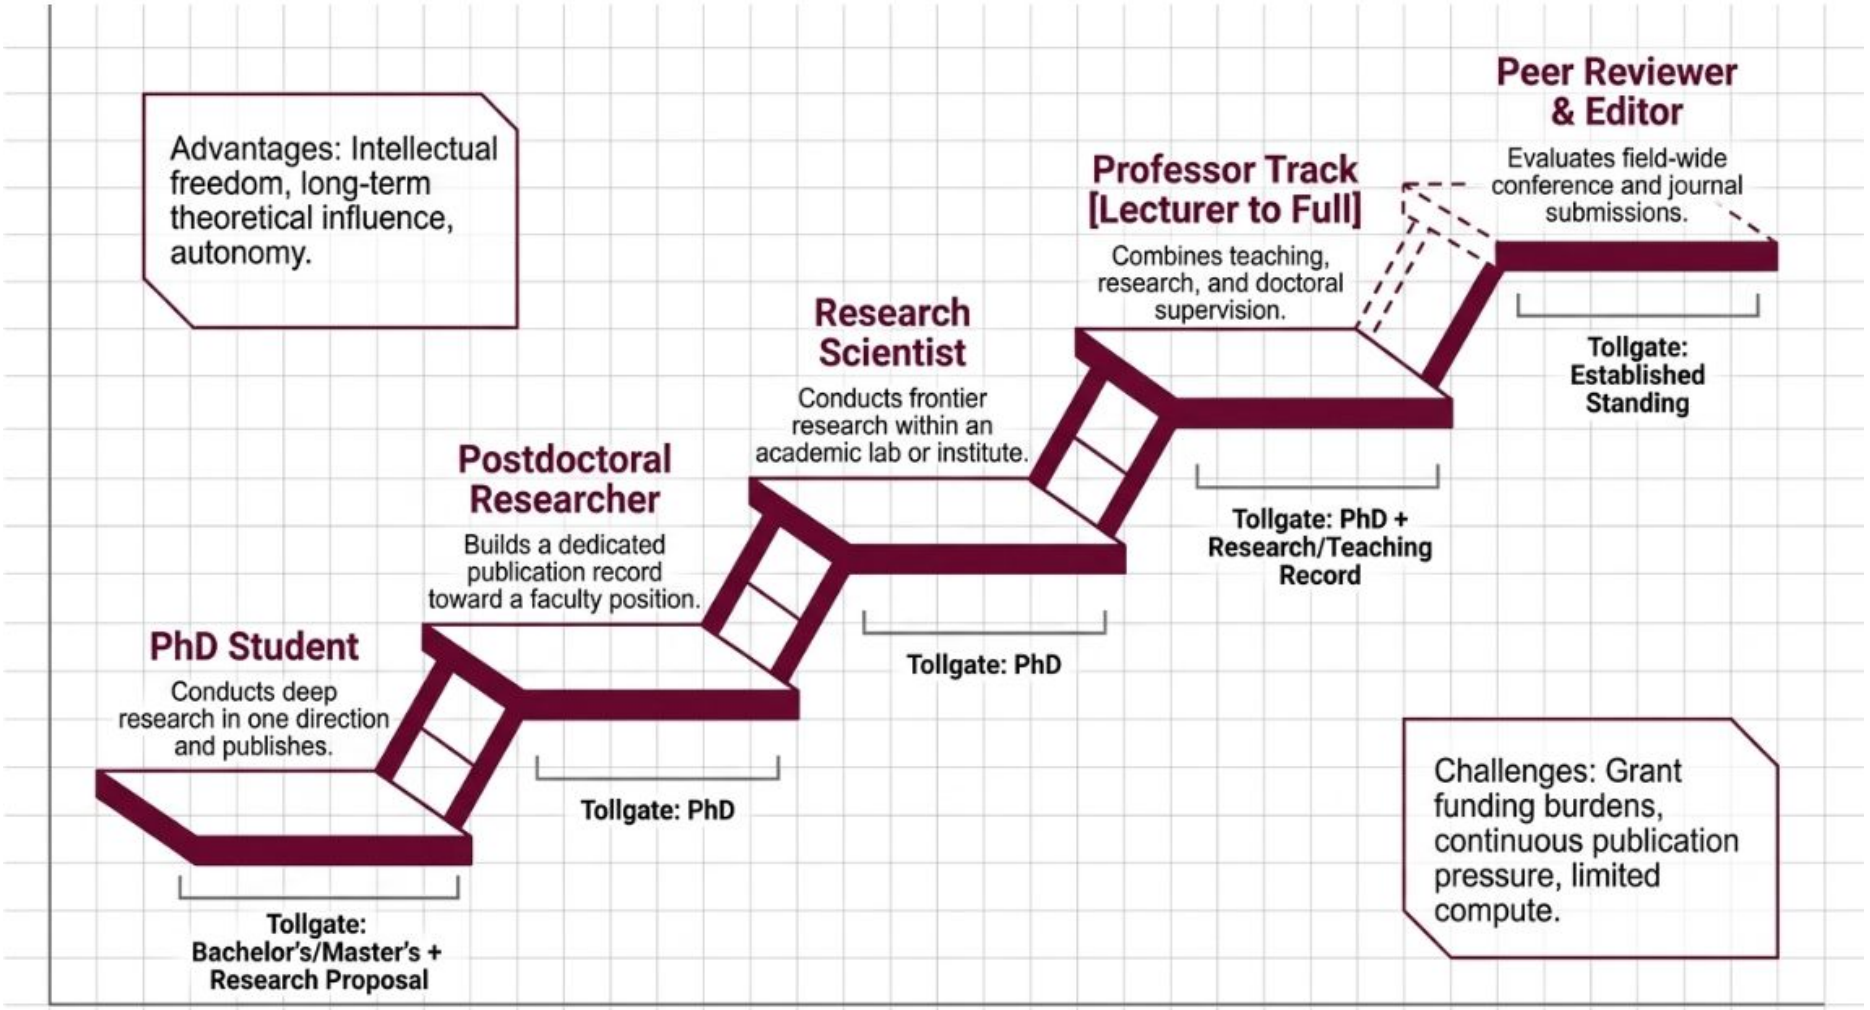

Careers in AI

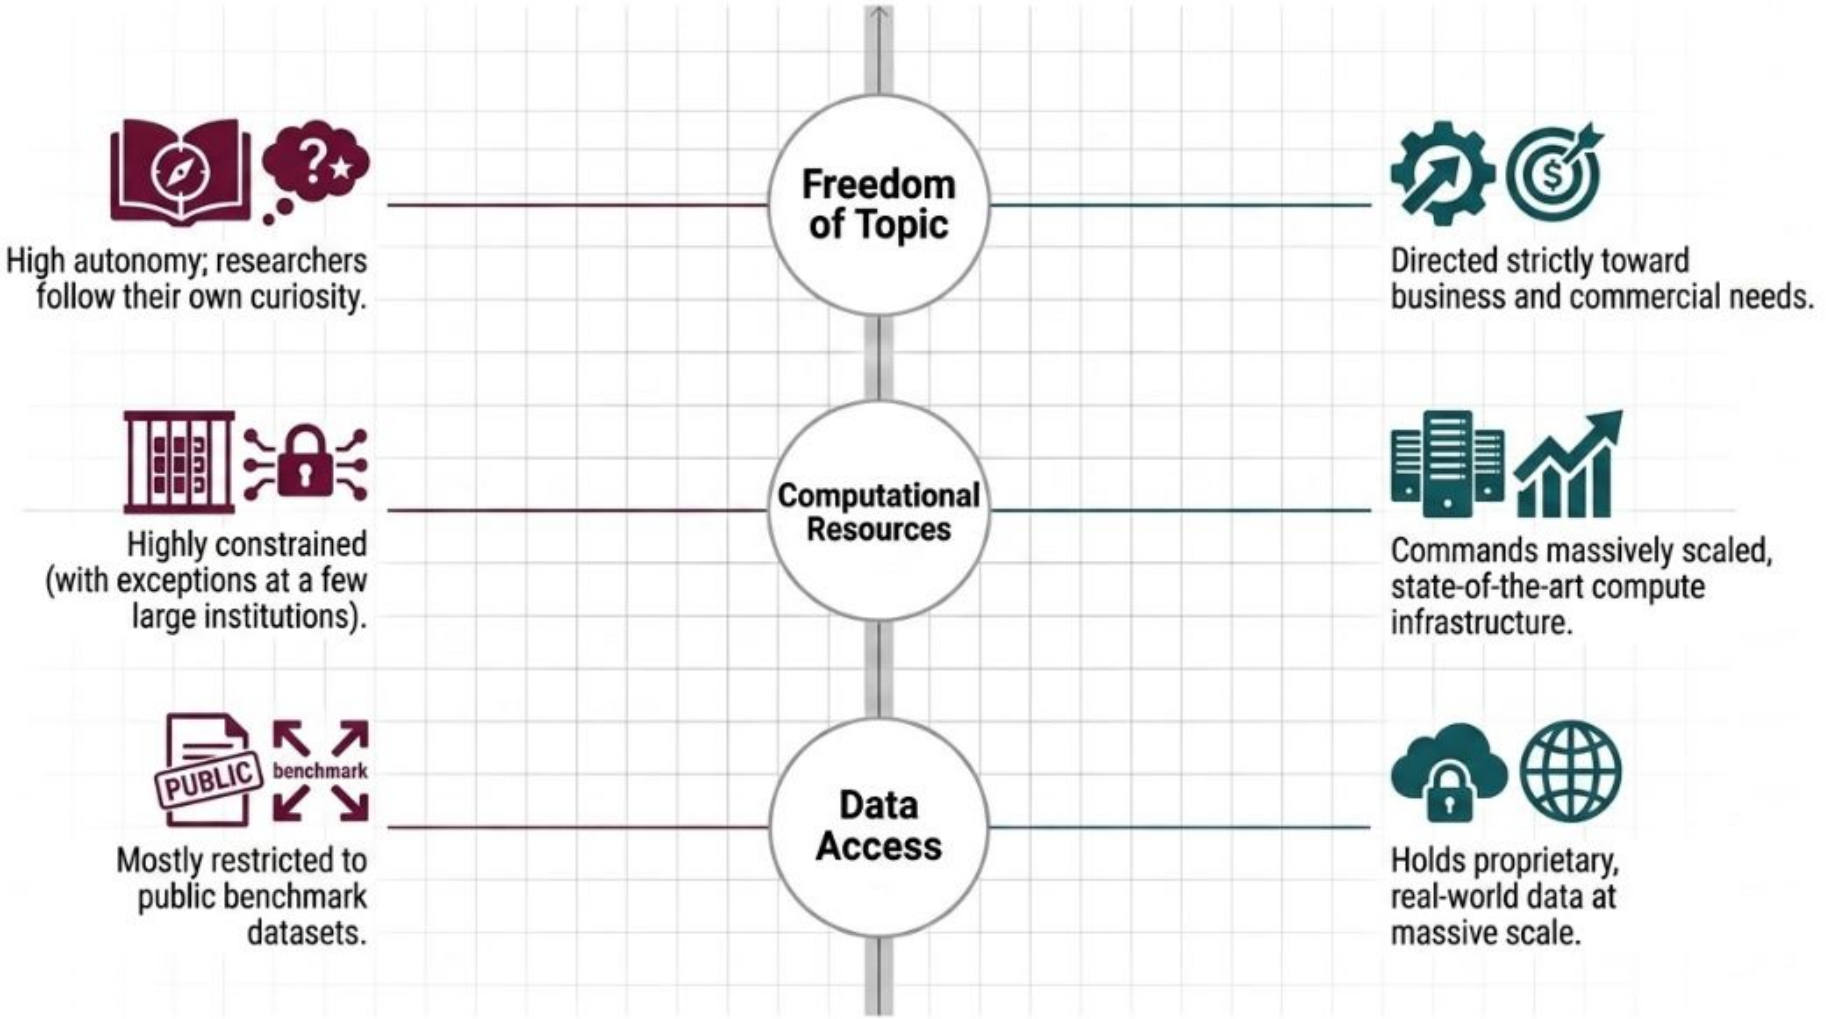

Careers in AI

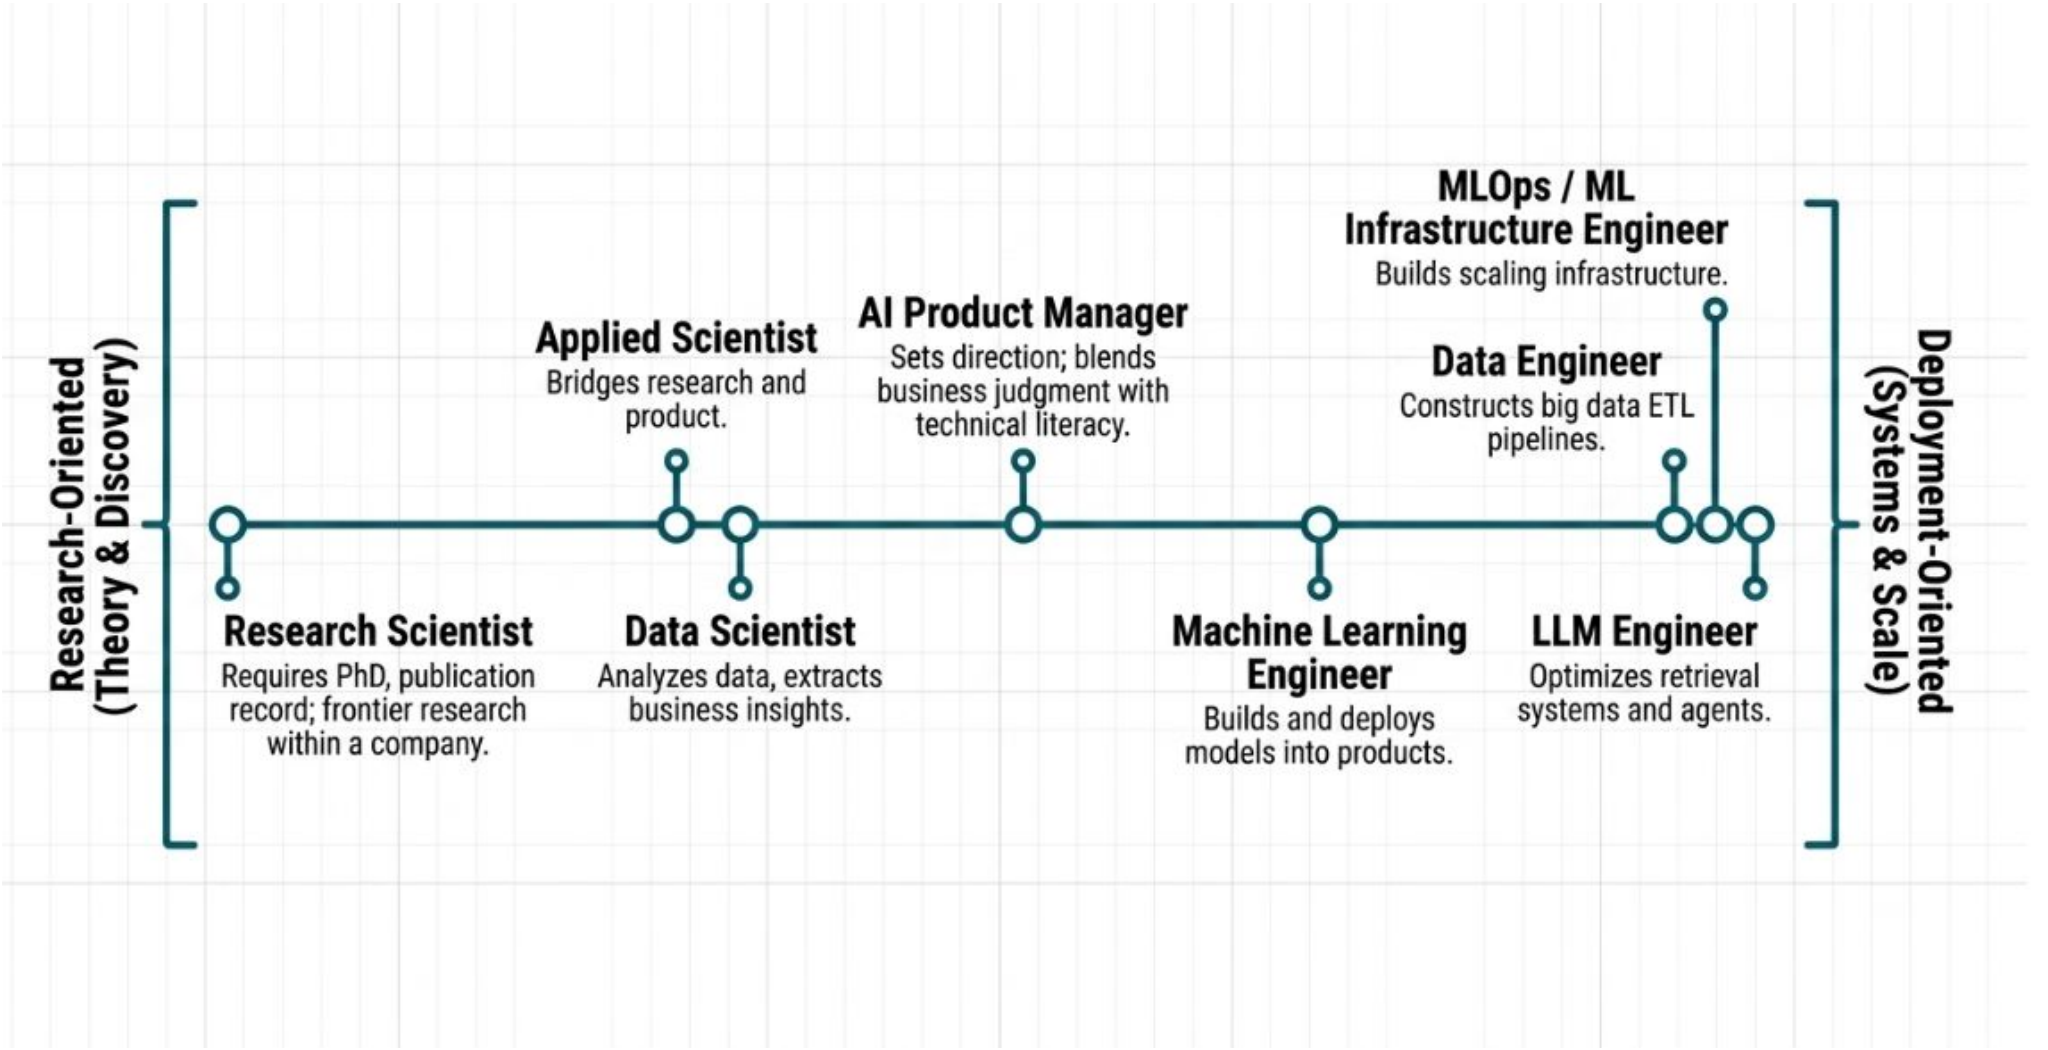

Careers in AI

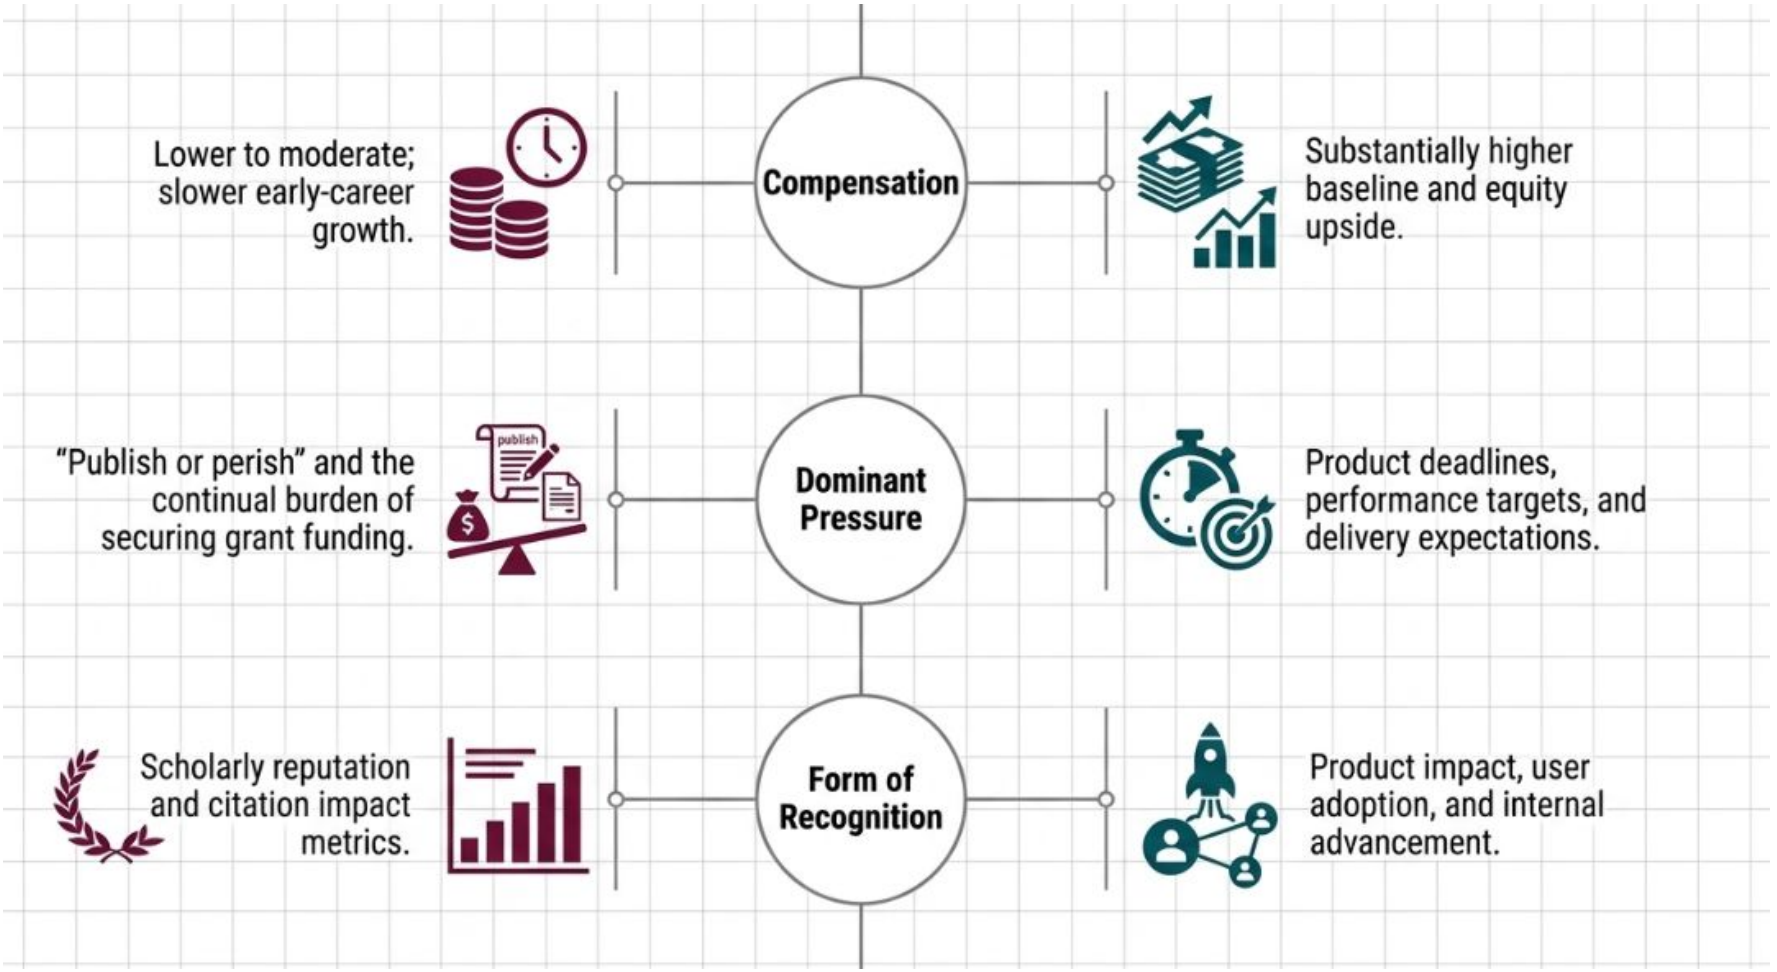

Careers in AI

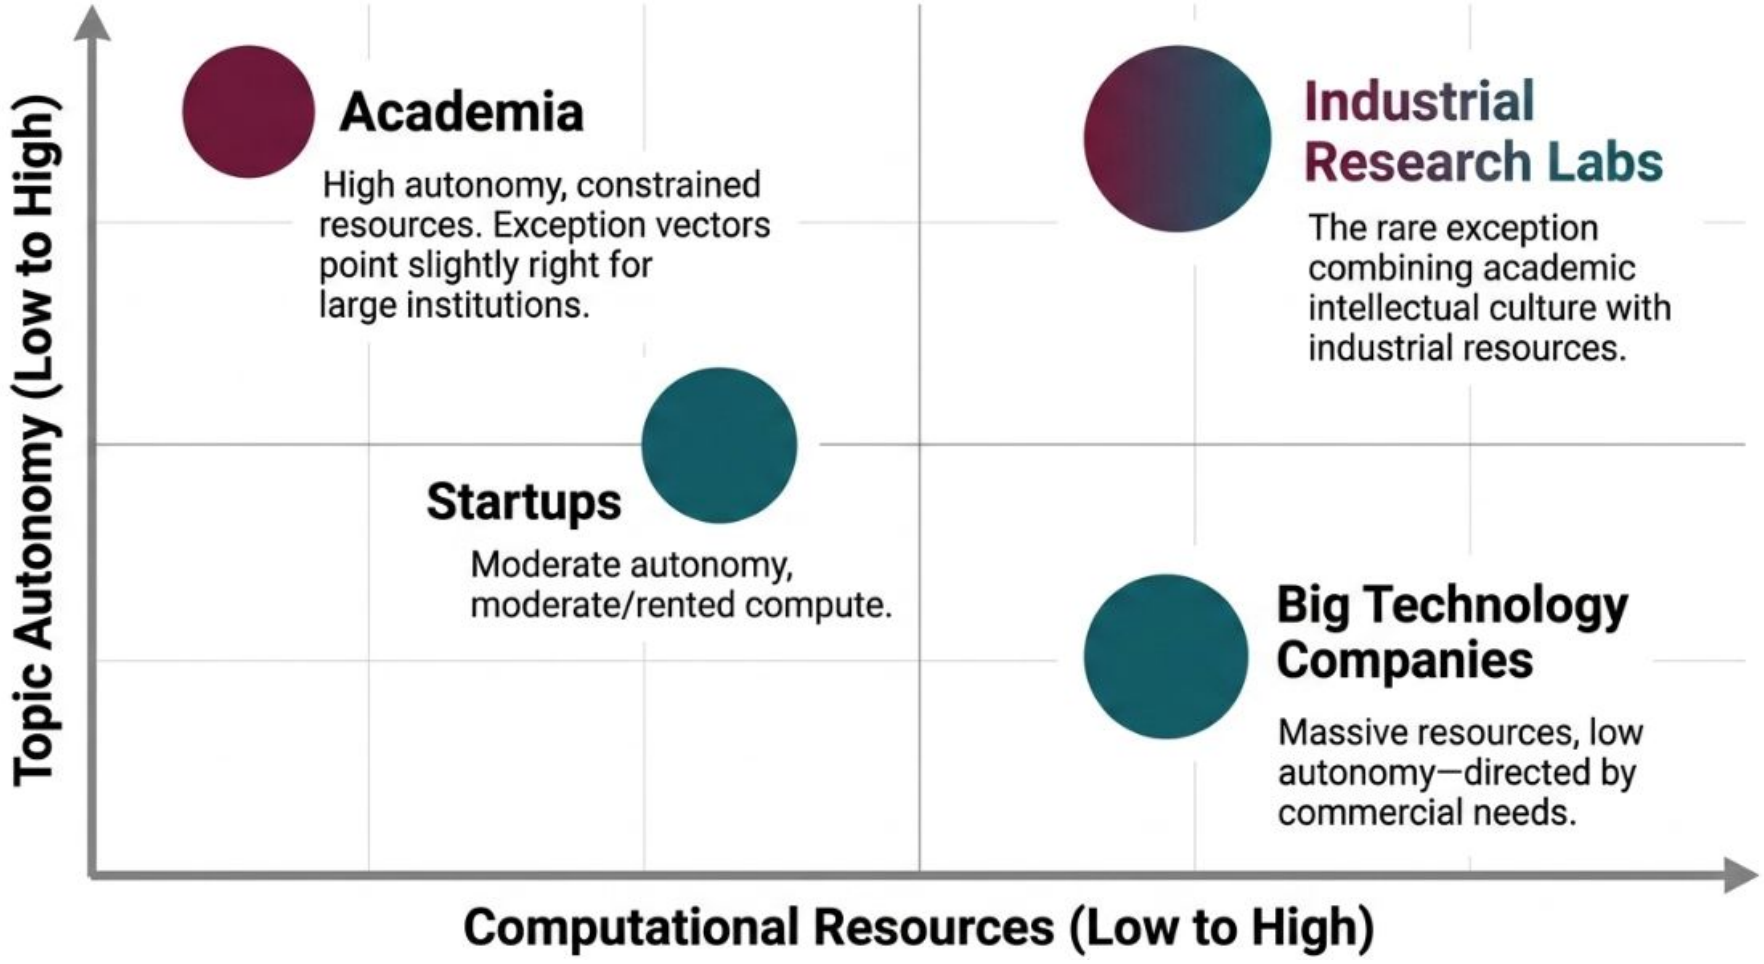

Thank you!