Artificial Intelligence Between Myth and Reality:

Understanding Models, Limits, and Learning

Gaetano Salina

Istituto Nazionale di Fisica Nucleare

Sezione di Roma Tor Vergata

salina@roma2.infn.it

42

It is the answer to the ultimate question of life, the universe and everything !

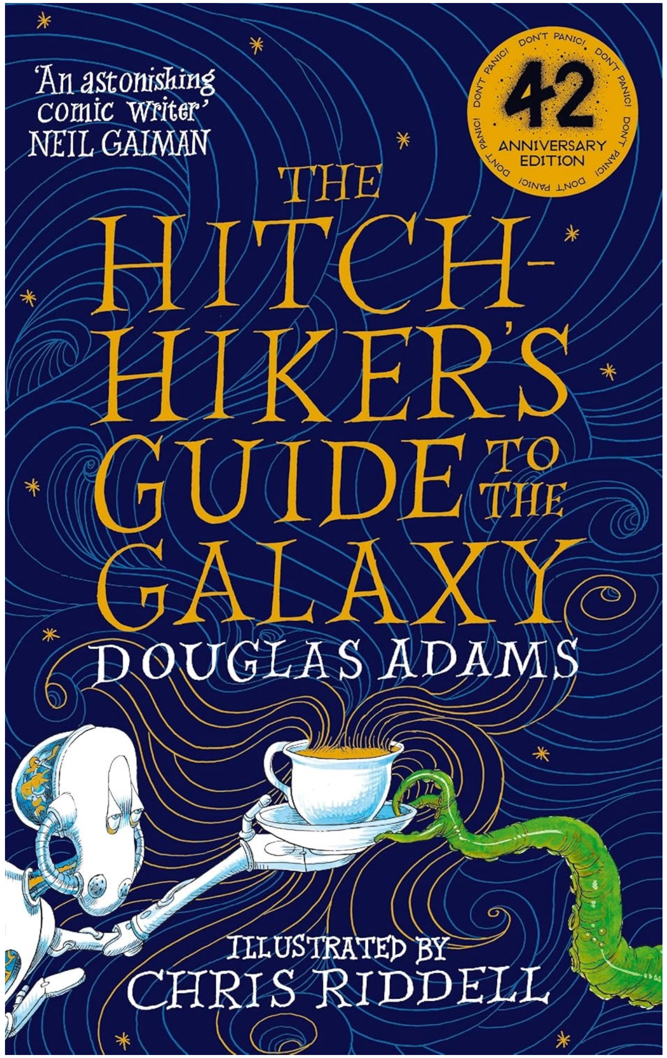

The Hitchhiker's Guide to the Galaxy

D. Adams 1982

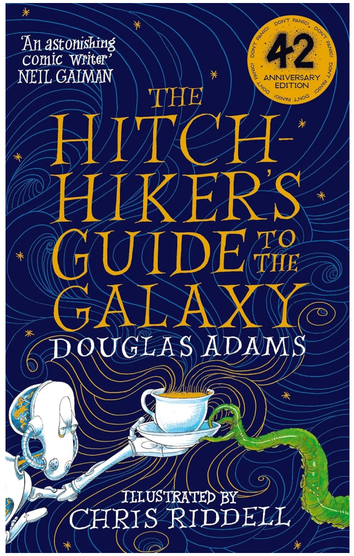

A group of scientists—actually the pandimensional projections of a super-intelligent species—build Deep Thought to solve the Ultimate Question of Life, the Universe, and Everything, wait seven and a half million years… and get back “42.”

“Forty-two!” shouted Loonquawl, his voice cracking across the void. “That’s all? After seven and a half million years of work?”

“I’ve checked it thoroughly,” said the computer, cold and unflinching. “That is the answer. The problem… is that you’ve never known the question.”

All You Need Is … the Scientific Method !

- In the Middle Ages, truth was defined by the scholastics as *argumentum ad auctoritatem.*
- That is, a thesis is accepted only by virtue of the authority of the person presenting it.
- The scientific method transcends authority (l'ipse dixit); it is defined by a rigorous process, not by the opinions of scientists.

Without a clear question, even a correct answer is meaningless

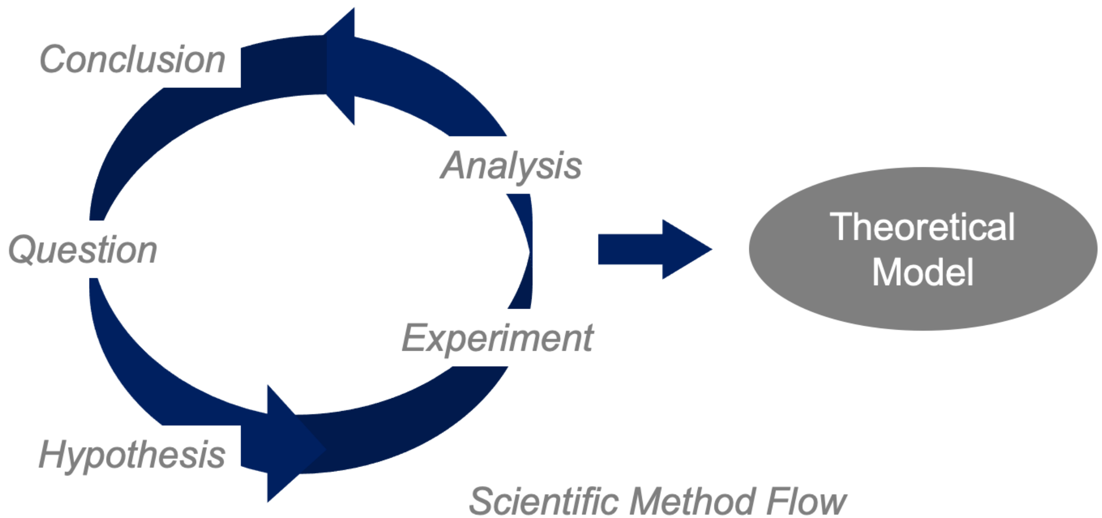

Scientific Method Flow

All You Need Is … the Scientific Method !

Correva l'anno 1632 …

Ma sapete, signor Simplicio, quel che accade? Sì come a voler che i calcoli tornino sopra i zuccheri, le sete e le lane, bisogna che il computista faccia le sue tare di casse, invoglie ed altre bagaglie, così, quando il filosofo geometra vuol riconoscere in concreto gli effetti dimostrati in astratto, ***bisogna che difalchi gli impedimenti*** *della materia; che se ciò saprà fare io vi assicuro che le cose si riscontreranno non meno aggiustatamente che i computi aritmetica. Gli errori dunque non consistono né nell'astratto né nel concreto, né nella geometria o nella fisica, ma nel* ***calcolatore****, che non sa fare i conti giusti.*

G. Galilei, *Dialogo sopra i due massimi sistemi, tolemaico e copernicano,* Einaudi, Torino, p. 252

To apply abstract truths to the real world, **one must remove unnecessary obstacles**, make the right approximations.
Done properly, reality matches calculation; errors arise only in the **one who miscalculates**.

Constructing a model of a Natural Phenomenon

The scientific method transforms observation, data, into knowledge

- To integrate apparently disparate empirical findings
- To predict previously unobserved aspects of the system's behavior
- The possibility of constructing an artifact that implements (subsets of) the system's relevant behavioral characteristics

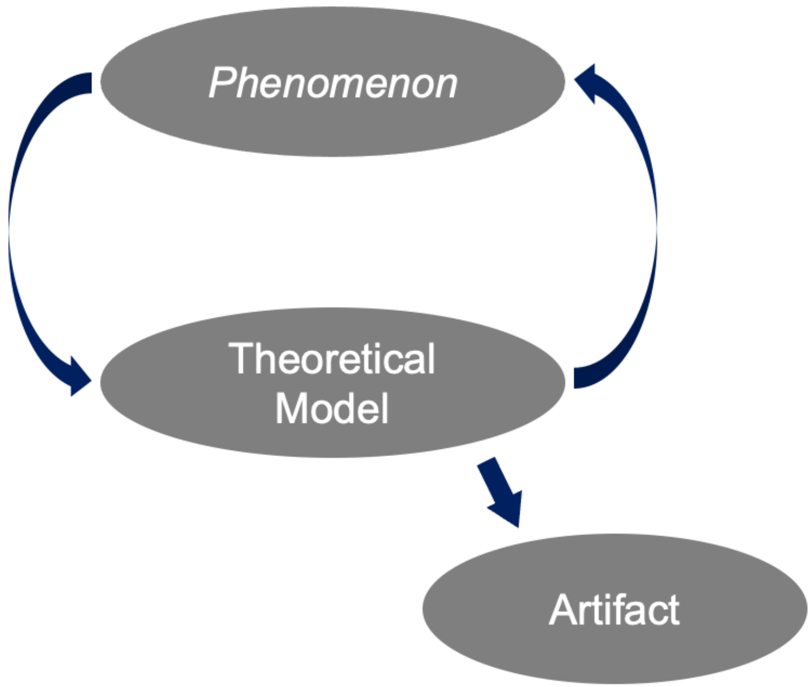

Constructing a model of a Natural Phenomenon

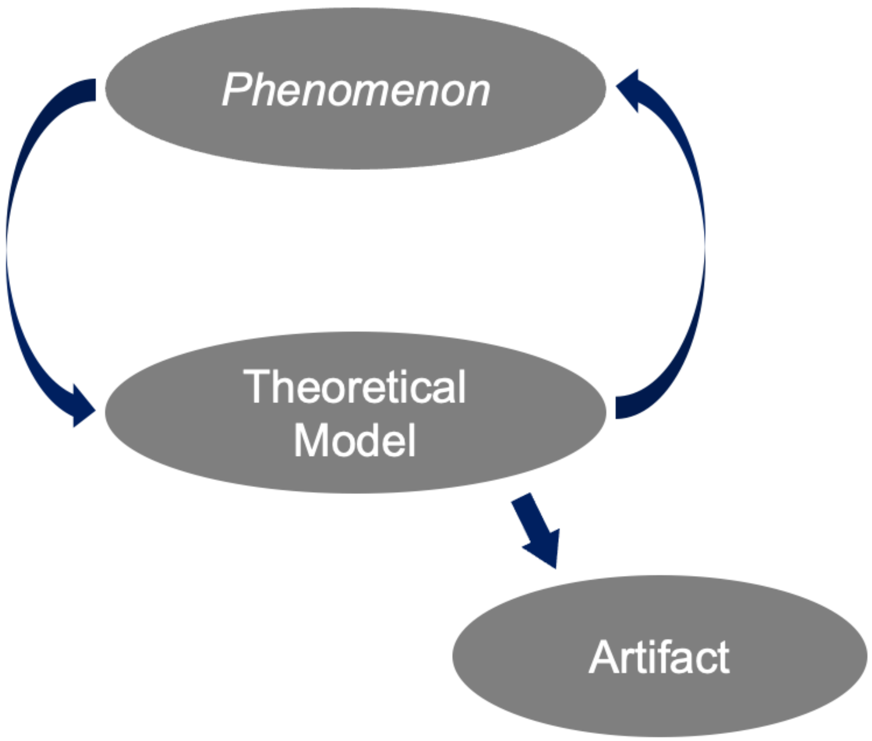

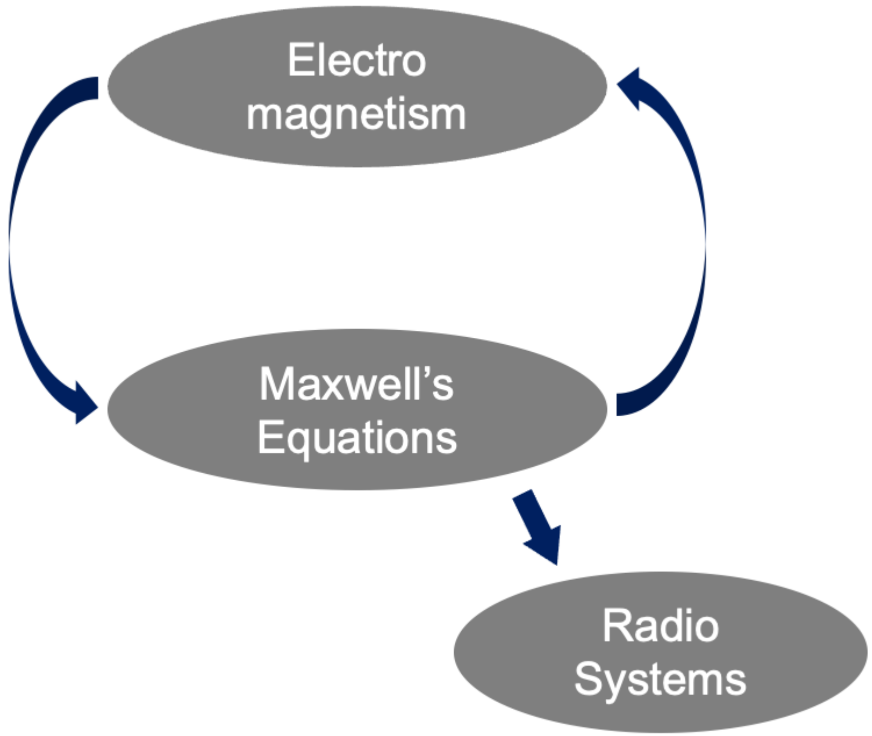

A new paradigm emerges with the data surge in the digital era, enabled by approaches that extract knowledge from data faster and better than conventional scientific methods, aiming directly at generating revenue."

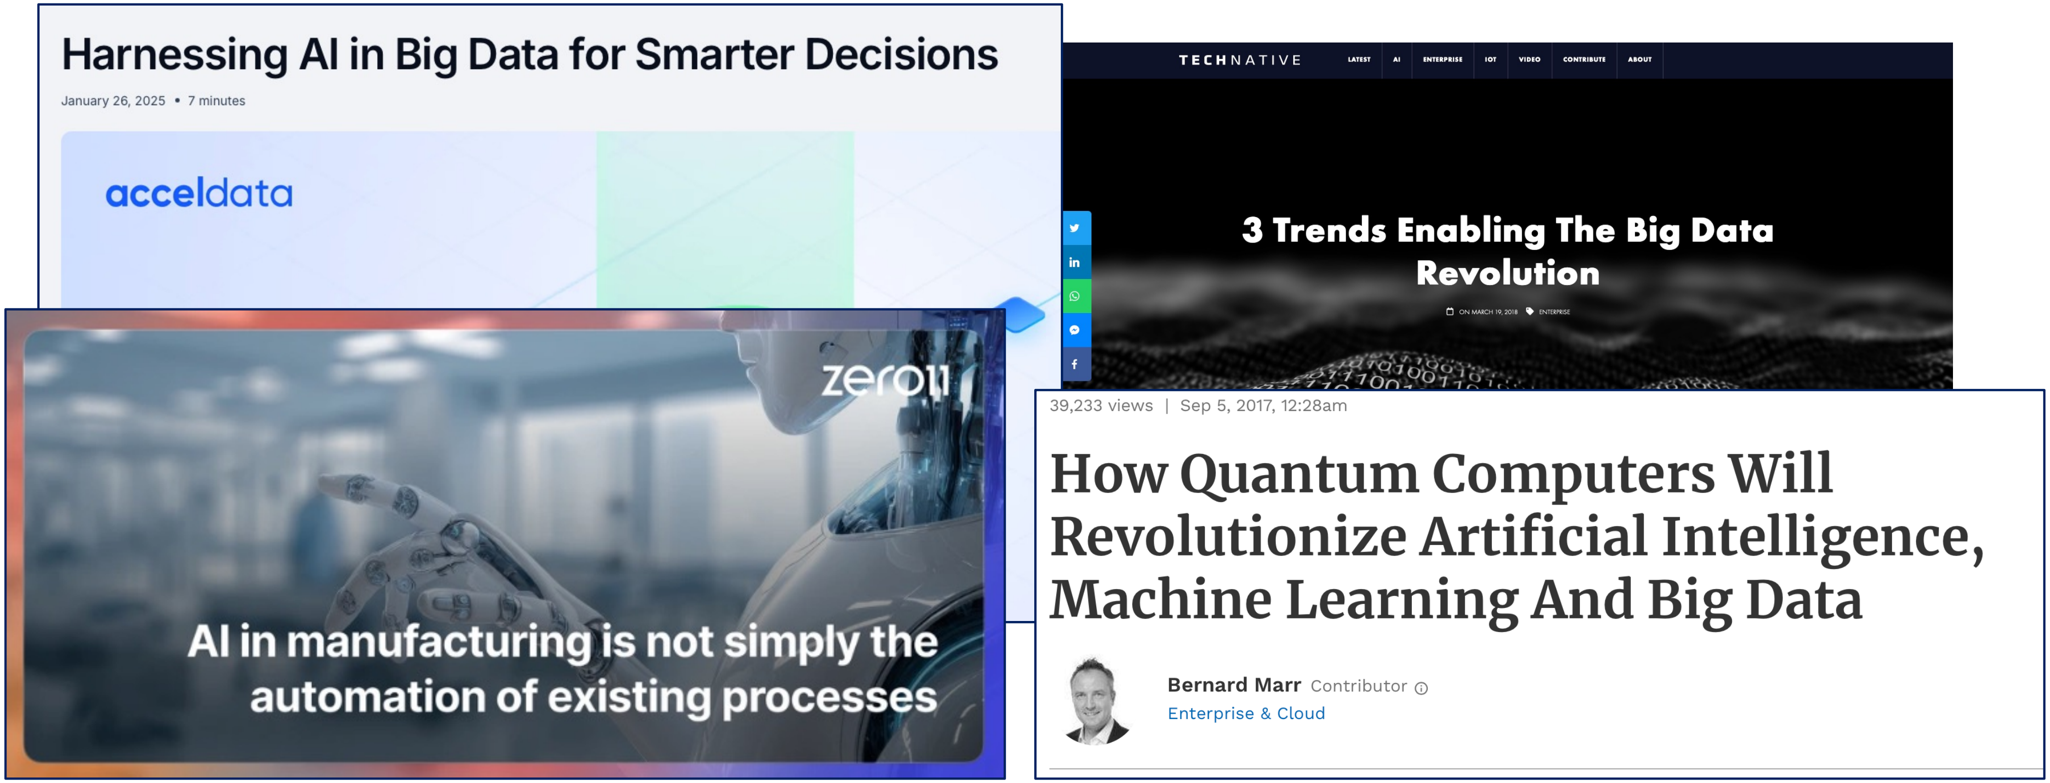

All You Need Is … Big Data & AI Approach

Correva l'anno 2008 …

There is can analyze the data without hypotheses about what it might show. We can throw the numbers now a better way. Petabytes allow us to say: "Correlation is enough." *We can stop looking for models.* ***We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot.***

C. Anderson, *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete,* https://www.wired.com/2008/06/pb-theory/

With enough data, the numbers speak for themselves

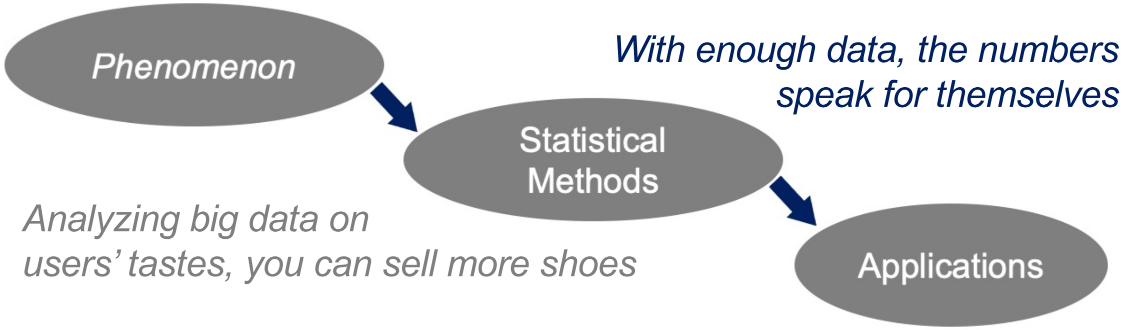

Analyzing big data on users' tastes, you can sell more shoes

CHRIS ANDERSON SCIENCE 06.23.08 12:00 PM

THE END OF THEORY: THE DATA DELUGE MAKES THE SCIENTIFIC METHOD OBSOLETE

Illustration: Marian Bantjes

"**All models are wrong**, but some are useful."

This article has been reproduced in a new format and may be missing content or contain faulty links. Contact wiredlabs@wired.com to report an issue.

So proclaimed statistician George Box 30 years ago, and he was right. But what choice did we have? Only models, from cosmological equations to theories of human behavior, seemed to be able to consistently, if imperfectly, explain the world around us. Until now. Today companies like Google, which have grown up in an era of massively abundant data, don't have to settle for wrong models. Indeed, they don't have to settle for models at all.

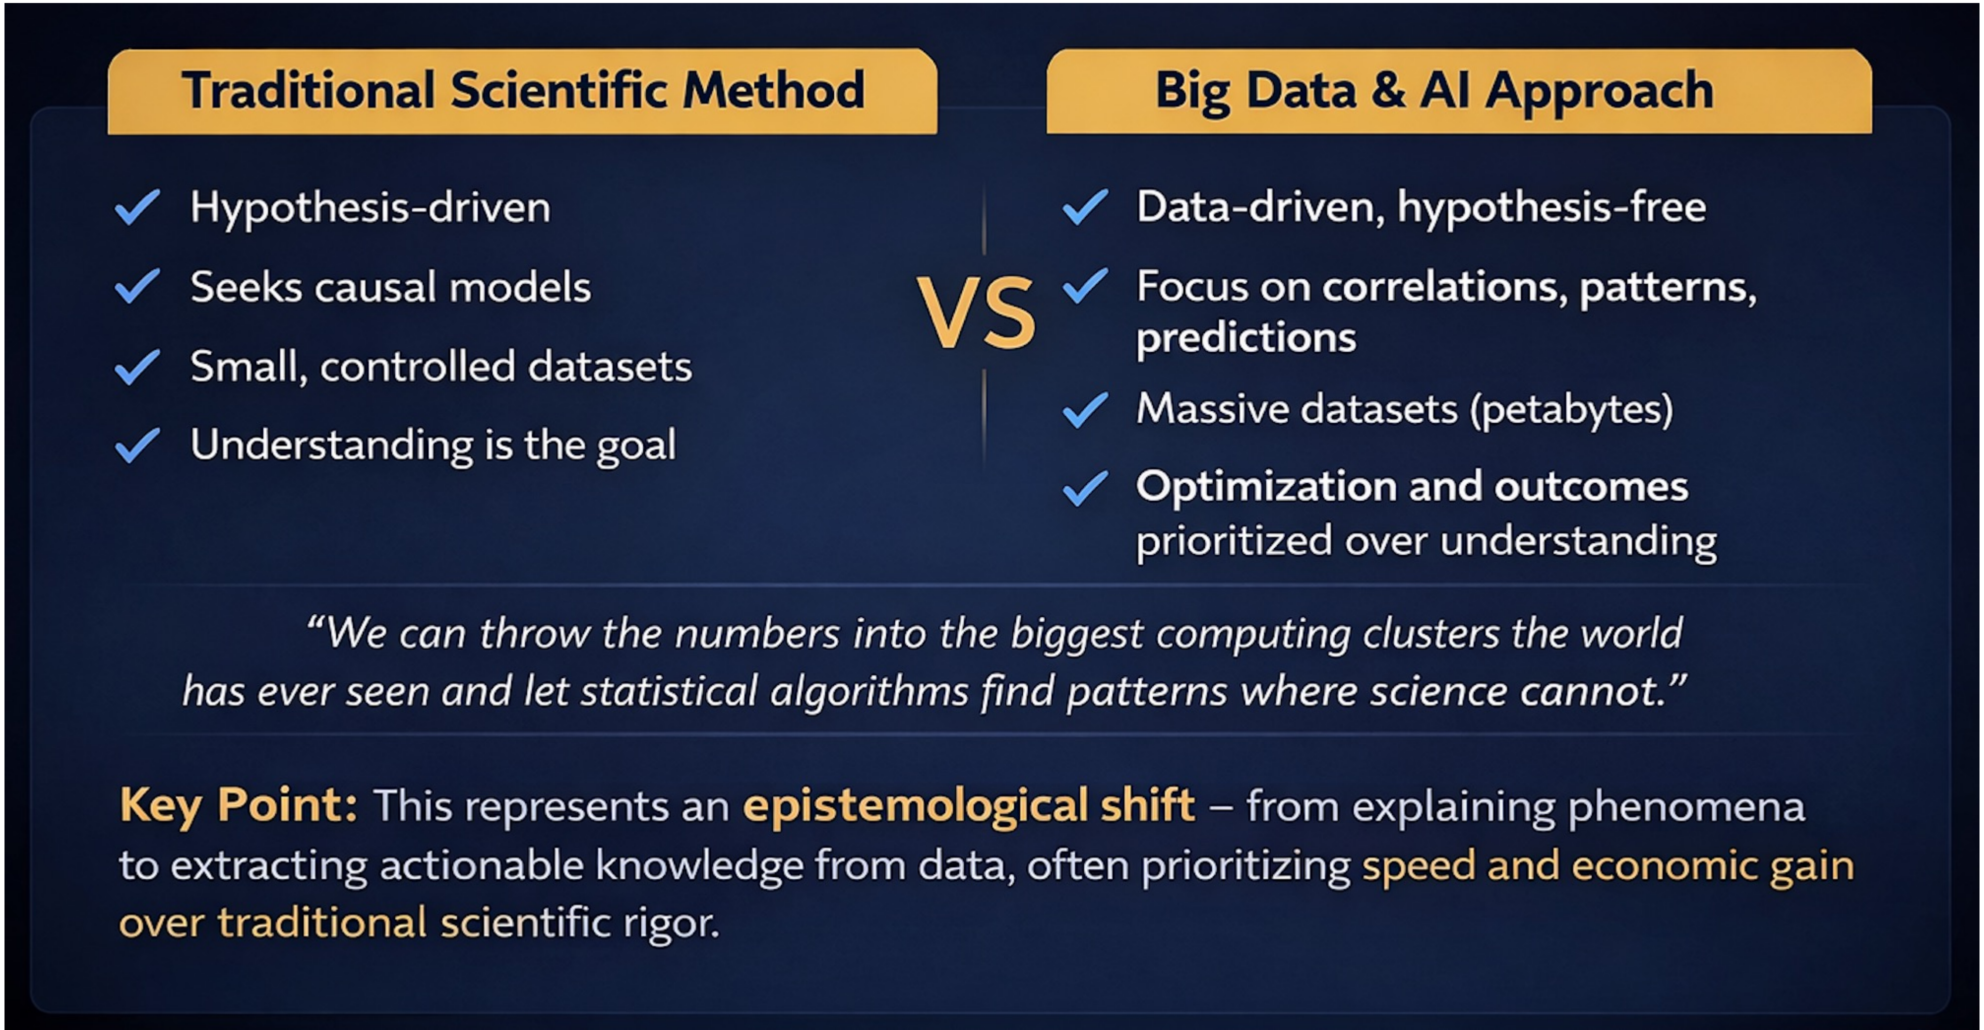

AI Generated Slide

My Methodological Bias

- Technologies currently framed as *revolutionary* (AI, Big Data) largely originate outside academia and are promoted through market-driven narratives in which the role of academia is marginal and ancillary, rather than grounded in mature epistemological reflection.
- The data-driven approach, when treated as an autonomous methodology of knowledge, is epistemically insufficient. Moreover, empirical evidence shows that purely data-driven methods rarely produce groundbreaking scientific discoveries.
- When integrated within the traditional scientific method, data-driven approaches can yield highly effective tools.
- The primacy of market logic, especially under monopolistic conditions, tends to undermine methodological and cultural pluralism, favoring technological stasis over genuine scientific progress.

From Narrative to Methodology

Setting aside the dominant narrative, this seminar focuses on methodological questions: what kind of knowledge do data-driven systems actually produce, under which assumptions, and within which epistemological limits.

- The goal is not to reject data-driven approaches, but to reposition them within the broader framework of the scientific method.
- Artificial Intelligence is treated here as a set of tools, not as an autonomous epistemology.
- The central question shifts from *What can AI do?* to *What can we legitimately know through AI-based methods?*

Today’s Talk Overview

- ***Knowledge Discovery Process***
  A modern name for well-known method
- ***AI & Machine Learning: what them are & how them work***
  A new name for well-defined technologies
- ***Computability Theory and Computational models***
  Not everything can be computed
- ***Some notes on Computational Paradigms***
  Analog Computing and Neural Networks
- ***Myth & Reality***
  The Rise and Fall of IBM Watson’s AI Medical System
- ***Three suggested books as Conclusion***
  To improve the understanding of some topics covered

Tomorrow's Talk Overview

- *Biological Neural Networks as a Functional Paradigma*
  Brain inspired approach
- *McCulloch-Pitts Formal Neuron*
  Propositional Calculus
- Multilayer Perceptron (MLP)
  Stacked Artificial Neuron Layers
- *Feed Forward and Attractors Neural Network*
  Layers, Dynamics, and Memory
- *Hopfield Model*
  Dynamic Neural Architecture
- *All You Need is Attention… or Hopfield*
  Physics before AI Models

Knowledge Discovery Process

A modern name for well-known method

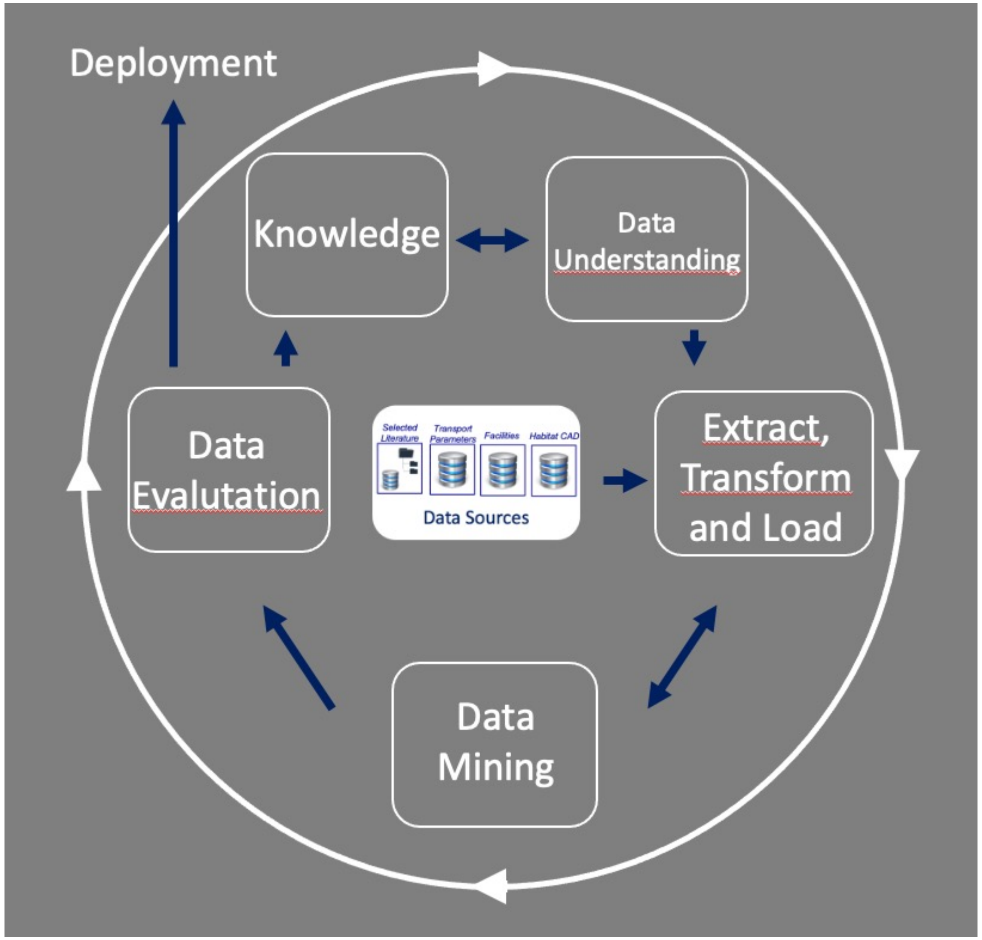

Knowledge Discovery Process: Definitions

- The KDP aims to identify hidden relationships, patterns, and trends in data, with the goal of generating predictions, recommendations, and informed decisions. It is a systematic approach to uncover knowledge that is not immediately apparent.
- The KDP is intended to provide a knowledge base (KB) to support decision-making, transforming raw data into actionable insights while maintaining rigor and reproducibility.

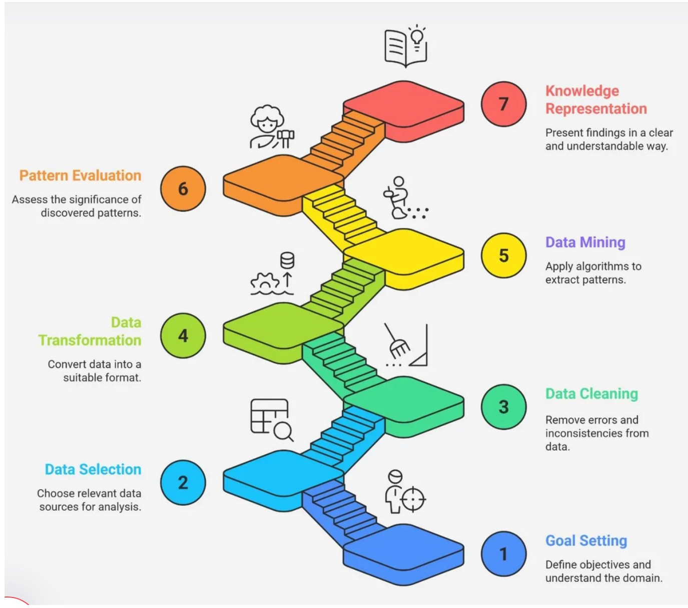

- In essence, the KDP mirrors the scientific method applied to data: it formulates hypotheses based on theoretical models, tests them against evidence, and iteratively refines understanding to build reliable knowledge.

Knowledge Discovery Process: Data Cemeteries?

- The ability to store data has far exceeded the human capacity to analyze it.
- Data Archives ⇄ Data Cemeteries
- Data contain knowledge of great scientific and economic interest.
- Research in KDP aims to define tools to transform data into information.

Knowledge Discovery Process: Five distinct steps

- Selection
- Pre-processing
- Transformation
- Data Mining
- Evaluation and Interpretation

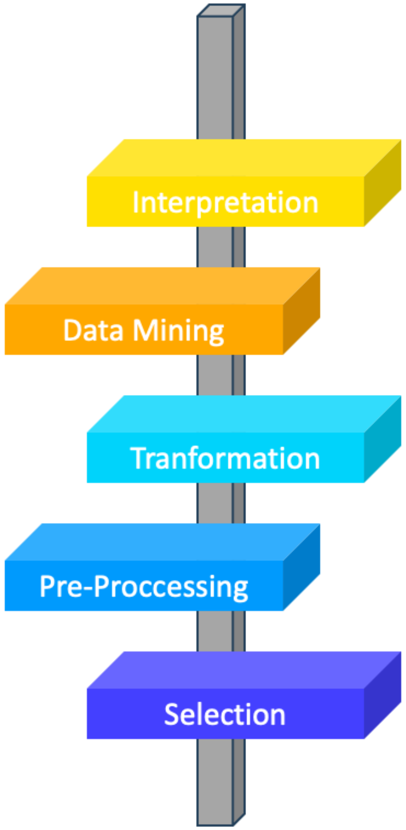

ETL (Extract, Transform, and Load) process

Knowledge Discovery Process: ETL process

- Selection: Acquisition of prior knowledge, identification of objectives, and development of the application domain → selection of appropriate data samples.
- Pre-processing: Handling missing values, identifying (and correcting) errors, removing duplicates, matching, merging, and resolving conflicts in data collected from multiple sources.
- Transformation: The cleaned dataset is transformed into a format suitable for analysis.

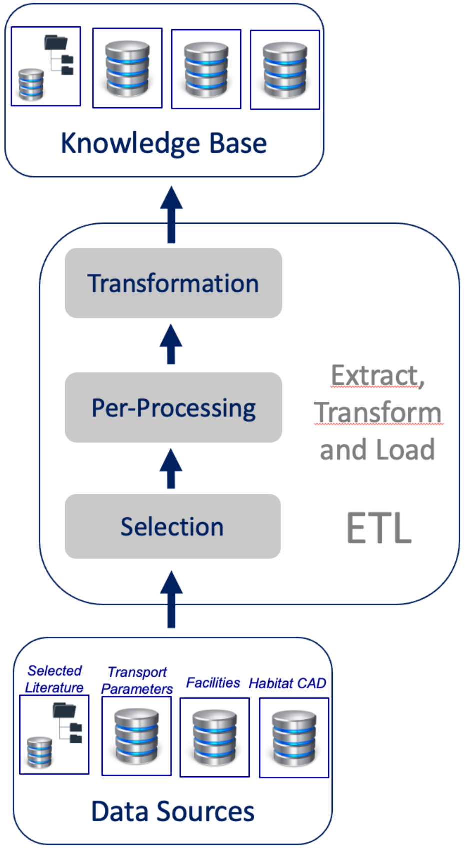

The ETL process enables an integrated Knowledge Base (KB) by with unique, consistent, and high-quality data and metadata.

Knowledge Discovery Process: ETL Process

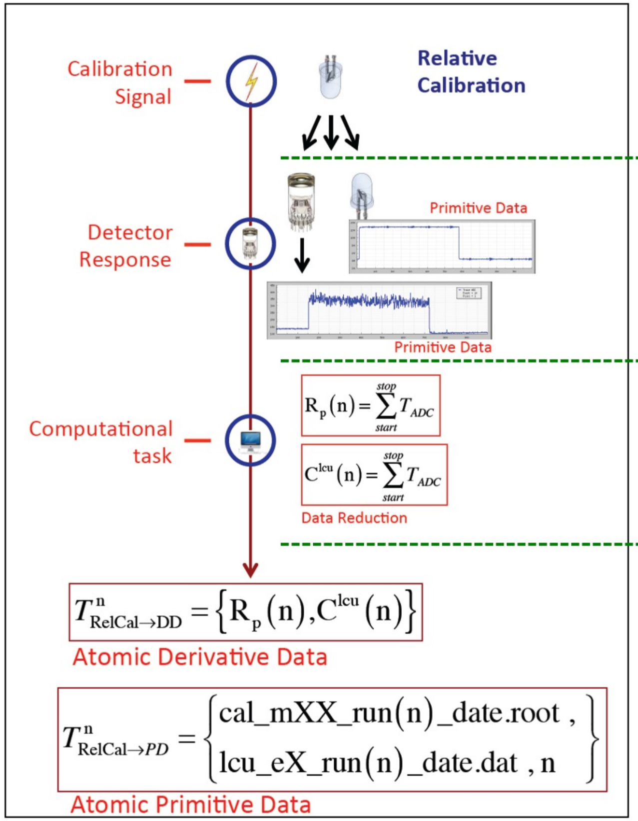

Physical Systems

Knowledge → Primitive Data

Data Acquisition:
Selection

Knowledge → Derivative Data

Pre-processing

Transformation

Data

Knowledge Base

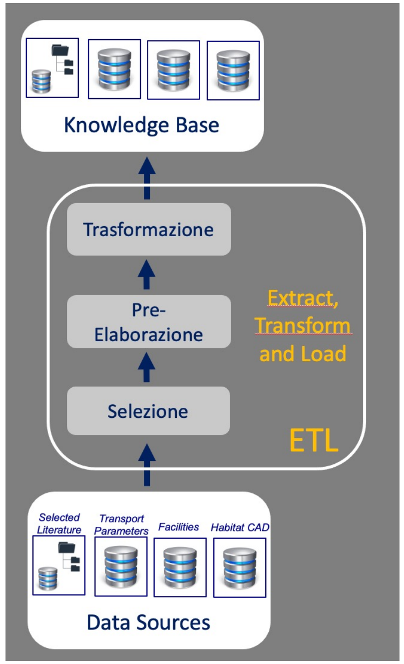

Knowledge Discovery Process: Data Mining

- Heuristic approach: data mining techniques (classification, regression, clustering) applied to transformed data under explicit researcher-defined hypotheses; model selection and hyperparameter tuning are driven by KDP objectives.
- Theoretical approach: transformed data are used to identify and estimate parameters of a theoretical dynamical model; the KDP coincides with the formal understanding of the system's physical dynamics.
- Statistical approach: transformed data are processed using data-driven machine learning models (e.g., neural networks), where the hypothesis space is implicit and learned from training examples rather than explicitly specified.

Knowledge Discovery Process: Data Mining in Physics

Theory
We know the relevant observables for describing a system and their interaction laws. Computational models allow us to solve system with a large number of degree of freedom and complex.

Heuristic
No theory is available and we are not sure that the observables are relevant. We use trial errors heuristic approach to define the observables interactions and define some numerical algorithms.

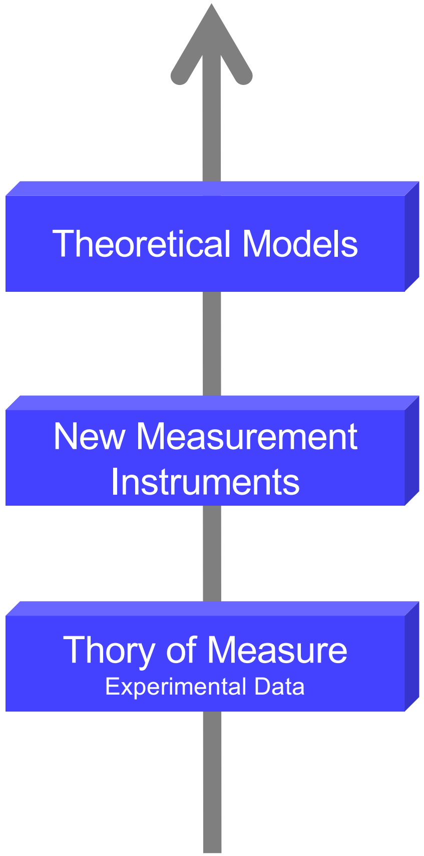

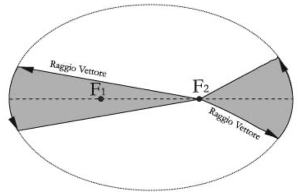

Kepler used Brahe's vast and precise data to formulate his laws of planetary motion. These laws laid the foundation for modern celestial mechanics.

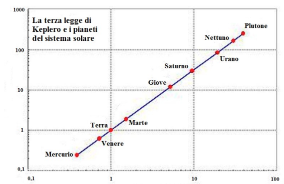

Brahe was a towering figure who ran a castle-like observatory with top self-made instruments and a NASA-like budget, making naked-eye measurements with high precision and amassing an enormous number of observations.

Knowledge Discovery Process: Evaluation and Interpretation

Evaluation and interpretation: the knowledge, patterns, and models obtained are assessed for their validity and practical relevance in a given application. This process also involves visualizing the data and the extracted patterns.

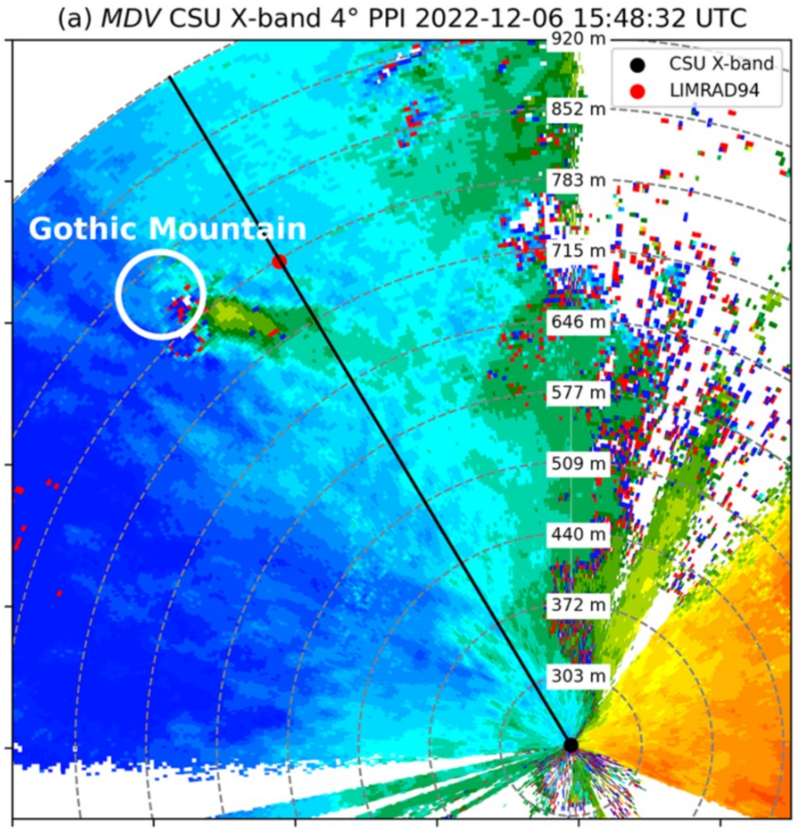

Knowledge Discovery Process: Data-Driven Predictions

- This latter view has recently gained some attention in response to the availability of unprecedented amounts of data and increasingly sophisticated algorithmic analytic techniques.

- The purpose of this note is to assess critically the role of big data in reshaping the key aspects of forecasting and in particular the claim that bigger data leads to better predictions.

FORECASTING IN LIGHT OF BIG DATA

HYKEL HOSNI AND ANGELO VULPIANI

Abstract. Predicting the future state of a system has always been a natural motivation for science and practical applications. Such a topic, beyond its obvious technical and societal relevance, is also interesting from a conceptual point of view. This owes to the fact that forecasting lends itself to two equally radical, yet opposite methodologies. A reductionist one, based on the first principles, and the naïve-inductivist one, based only on data. This latter view has recently gained some attention in response to the availability of unprecedented amounts of data and increasingly sophisticated algorithmic analytic techniques. The purpose of this note is to assess critically the role of *big data* in reshaping the key aspects of forecasting and in particular the claim that *bigger* data leads to *better* predictions. Drawing on the representative example of weather forecasts we argue that this is not generally the case. We conclude by suggesting that a clever and context-dependent compromise between modelling and quantitative analysis stands out as the best forecasting strategy, as anticipated nearly a century ago by Richardson and von Neumann.

Nothing is more practical than a good theory (L. Boltzmann)

Knowledge Discovery Process: Data-Driven Predictions

- We may ask whether, based on our knowledge of a system's past states, it is possible to make meaningful predictions about its future.
- Our answer is largely negative, as significant difficulties arise even in highly abstract and simplified scenarios.

- As we will see, the key challenge lies in determining the *appropriate level* of abstraction for the system.

Knowledge Discovery Process: Data-Driven Predictions

Data-Driven Predictions: An Inductive Approach
Predicting the future behavior of a system solely from knowledge of its past. Two key assumptions are necessary to make such predictions meaningful:

- Analogy: Similar conditions are expected to yield similar outcomes.
- Determinism: Systems that exhibit a particular behavior are assumed to continue doing so.

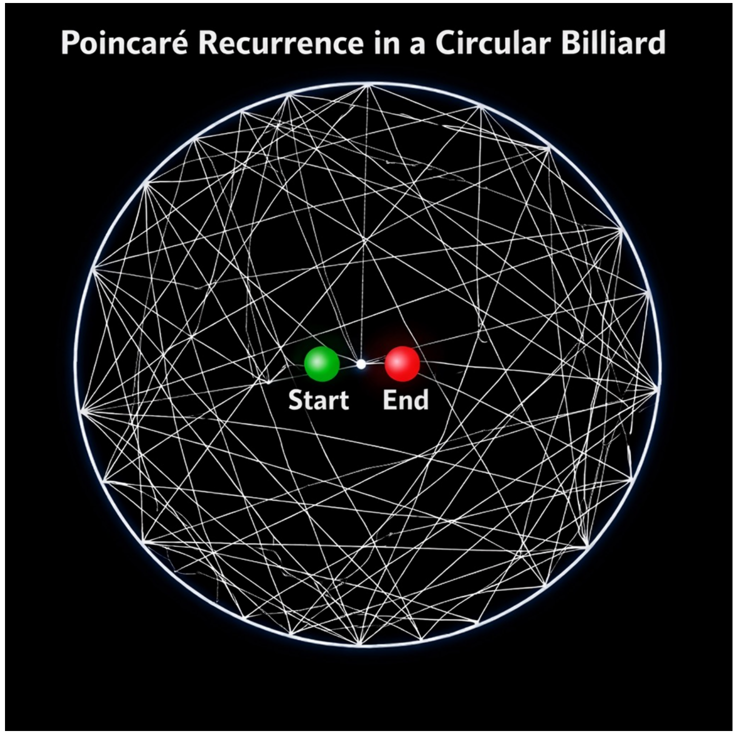

Knowledge Discovery Process: ***Data-Driven Predictions***

Consider a system defined by the series $\{x_1, x_2, \ldots, x_M\}$, where each x_i is a *D-dimensional* vector representing the system state at time $i\Delta t$.

Our goal is to predict x_{M+1}.

By analogy, we search the past for a state x_k sufficiently close to x_M (with $k < M$) such that $|x_k - x_M| < \varepsilon$), where ε denotes the desired precision.

Once such a vector is identified, we predict the future state at time $M + n > M$ by assuming $x_{M+n} = x_{k+n}$.

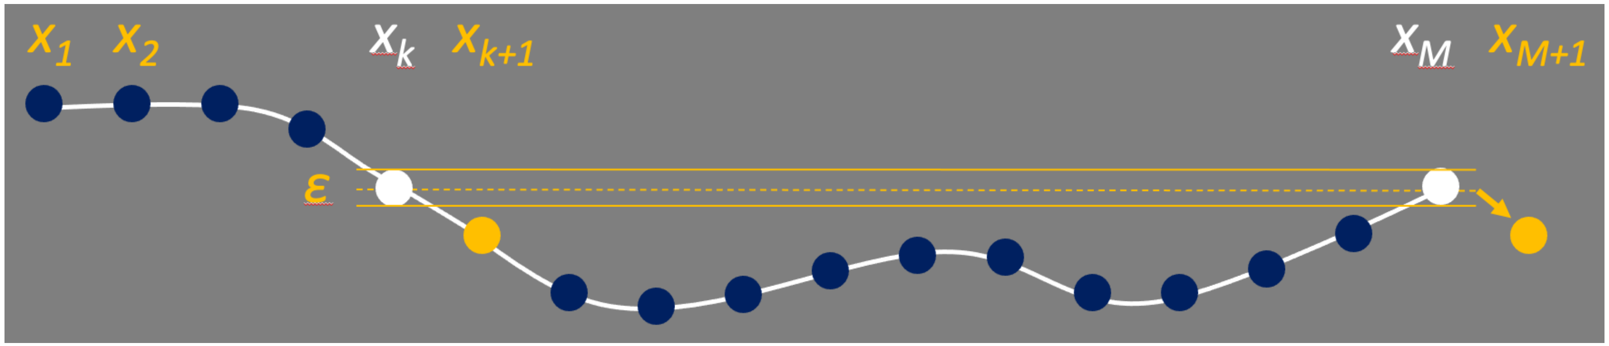

Knowledge Discovery Process: Data-Driven Predictions

Poincaré Recurrence Theorem

A deterministic system with a bounded phase space returns arbitrarily close to its initial state after a time $<T_R>$.

Phase Space:

An abstract space used to describe the time evolution of a physical system. Its dimension depends on the number of variables that determine the state of the system.

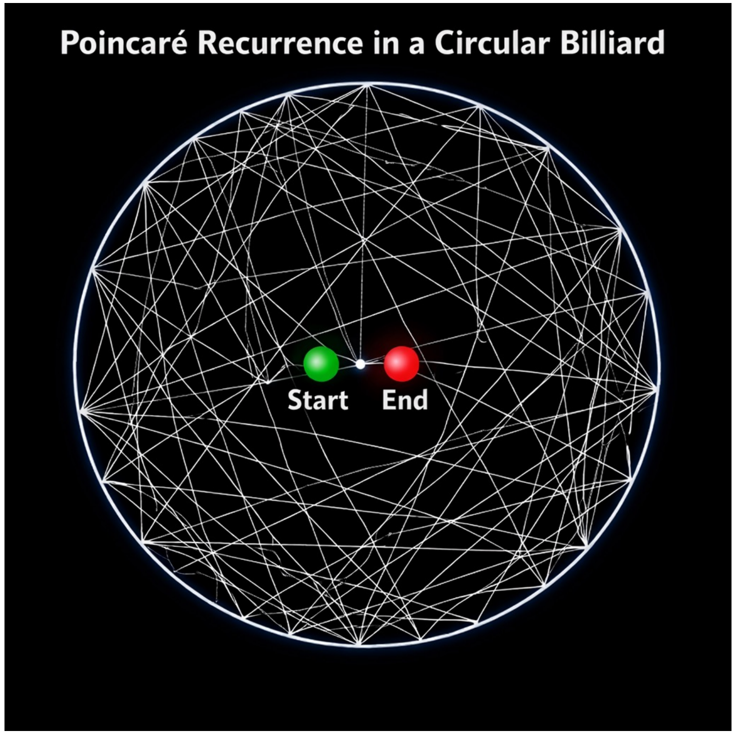

Knowledge Discovery Process: Data-Driven Predictions

Kac's Lemma

The average return time to a region A is inversely proportional to *P(A)*

$$\langle T_R \rangle \propto \frac{1}{P(A)}$$

Here, *P(A)* is the probability that the system is in A,

$$P(A) \propto \varepsilon^D \mapsto \langle T_R \rangle = O\left(\varepsilon^{-D}\right)$$

having

$$D \approx 10 \quad \varepsilon \approx 0.01 \quad \Rightarrow \langle T_R \rangle = O\left(10^{20}\right)$$

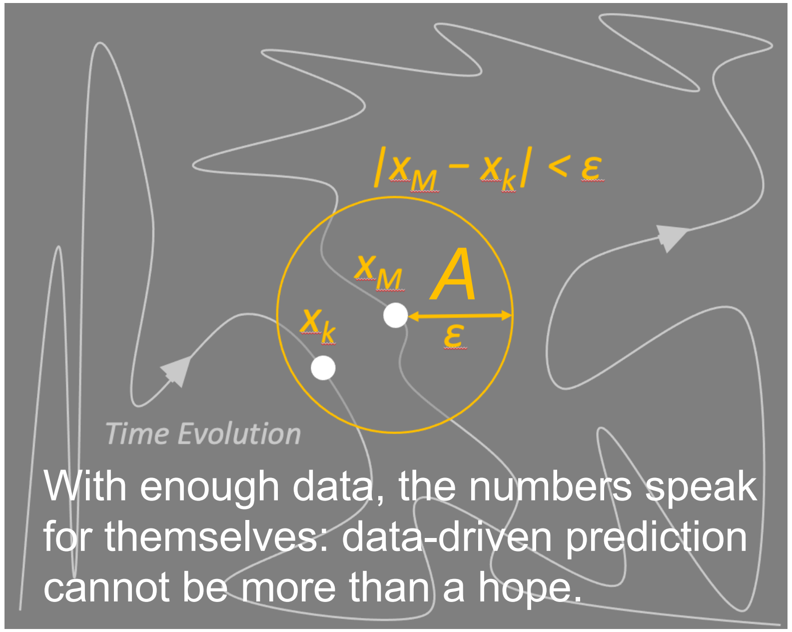

Knowledge Discovery Process: Data-Driven Predictions

- In socio-economic systems, the gap between available data and our understanding is typically substantial.
- Leveraging the abundance of data, advanced data-driven KDP methods have the potential to enhance our knowledge of such systems.
- However, any resulting insights must be validated against other analytical approaches to ensure their plausibility.

AI & Machine Learning: AI & Turing Test

- C interacts with A and B.
- If C cannot distinguish the computer from the human, the test is passed.
- In this case, A is an intelligent artificial agent!

An artificial agent is considered intelligent if it behaves like a human being.

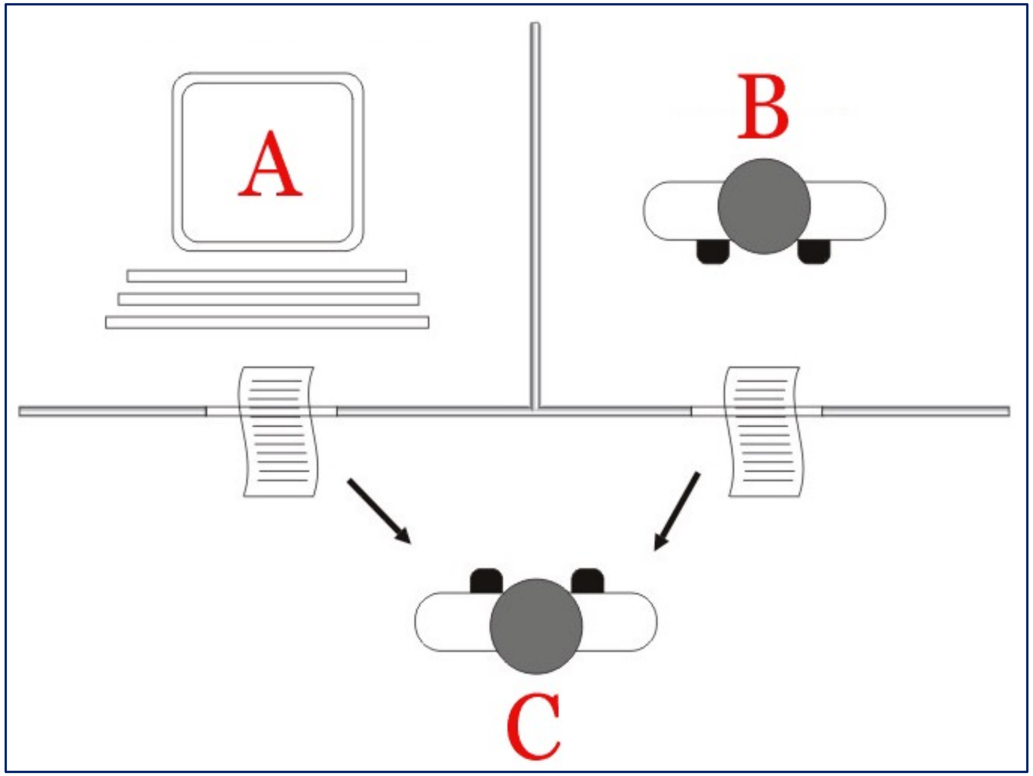

AI & Machine Learning: AI & EU Regulation

- AI is the ability of a machine to demonstrate human-like capabilities such as reasoning, learning, planning, and creativity.

- An AI system solves problems by autonomously interacting with the environment and analyzing its own history.

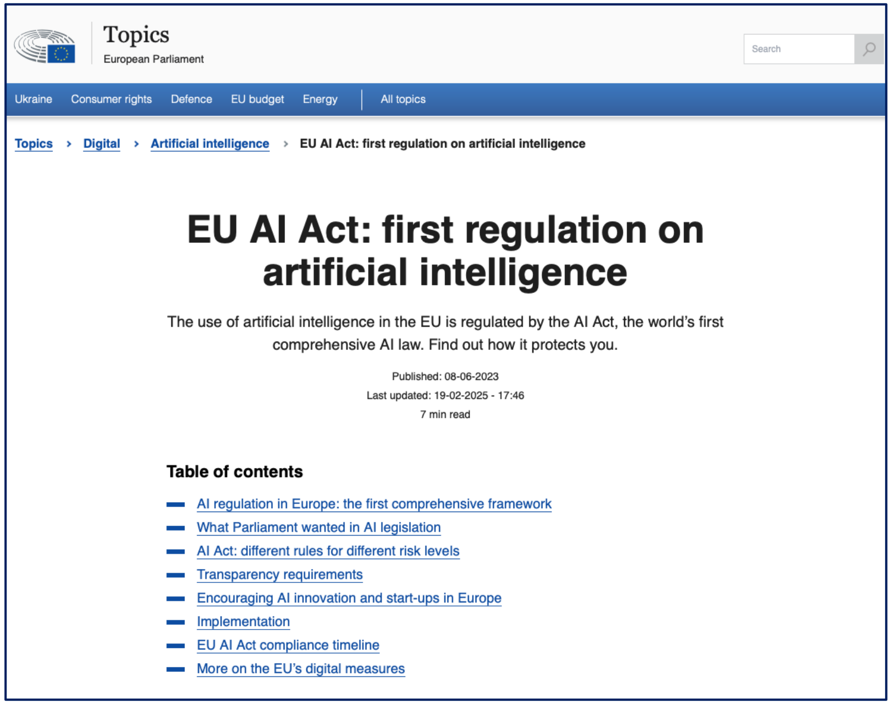
Topics
European Parliament
Search
Ukraine | Consumer rights | Defence | EU budget | Energy | All topics
Topics › Digital › Artificial intelligence › EU AI Act: first regulation on artificial intelligence

EU AI Act: first regulation on artificial intelligence

The use of artificial intelligence in the EU is regulated by the AI Act, the world's first comprehensive AI law. Find out how it protects you.

Published: 08-06-2023
Last updated: 19-02-2025 - 17:46
7 min read

Table of contents

An artificial agent is considered intelligent if it behaves like a human being.

AI & Machine Learning: AI & Proposed Taxonomy

- Thinking like humans
  Human Cognitive Processes, focuses on understanding and modeling human thought
- Acting like humans
  Human-like Behavioral Test, focuses on replicating human behavior
- Thinking rationally
  Formal Logical Reasoning, focuses on correct and systematic logical reasoning
- Acting rationally
  Rational Goal-directed Action, focuses on taking the right goal-directed actions

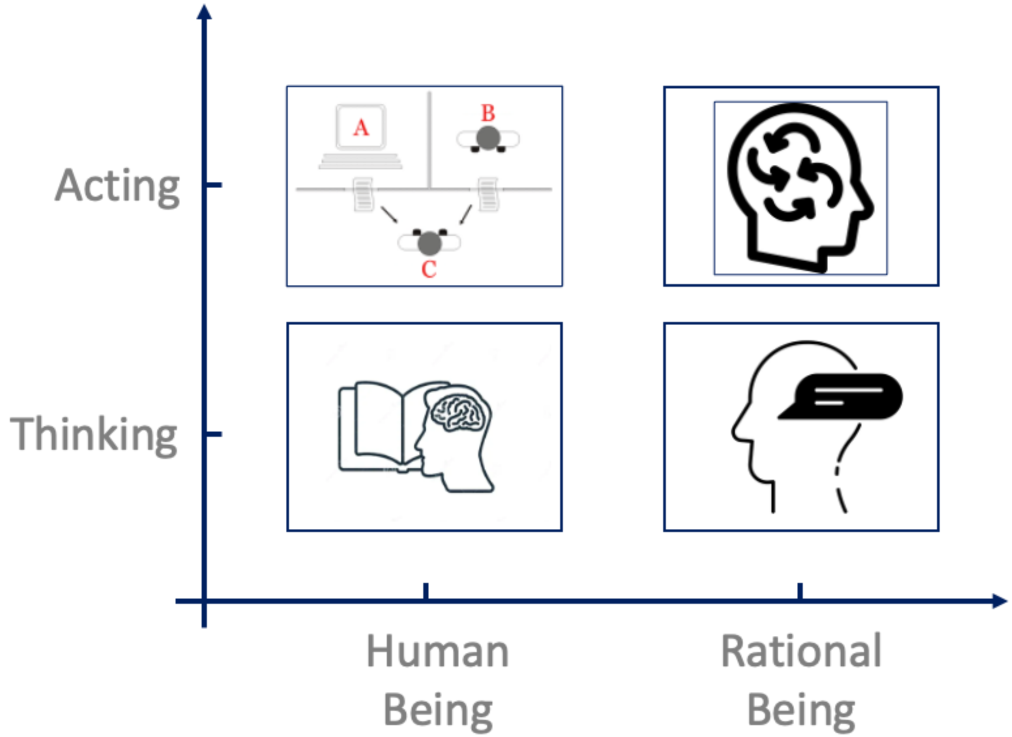

A Modern Approach, S. Russell and P. Norvig, Prentice Hall, 2020

AI & Machine Learning: AI & Operational Definition

An artificial agent is intelligent if it has:

- the ability to receive information and take actions
- a function $\mathcal{F}$ mapping sequences of perceptions to sequences of actions

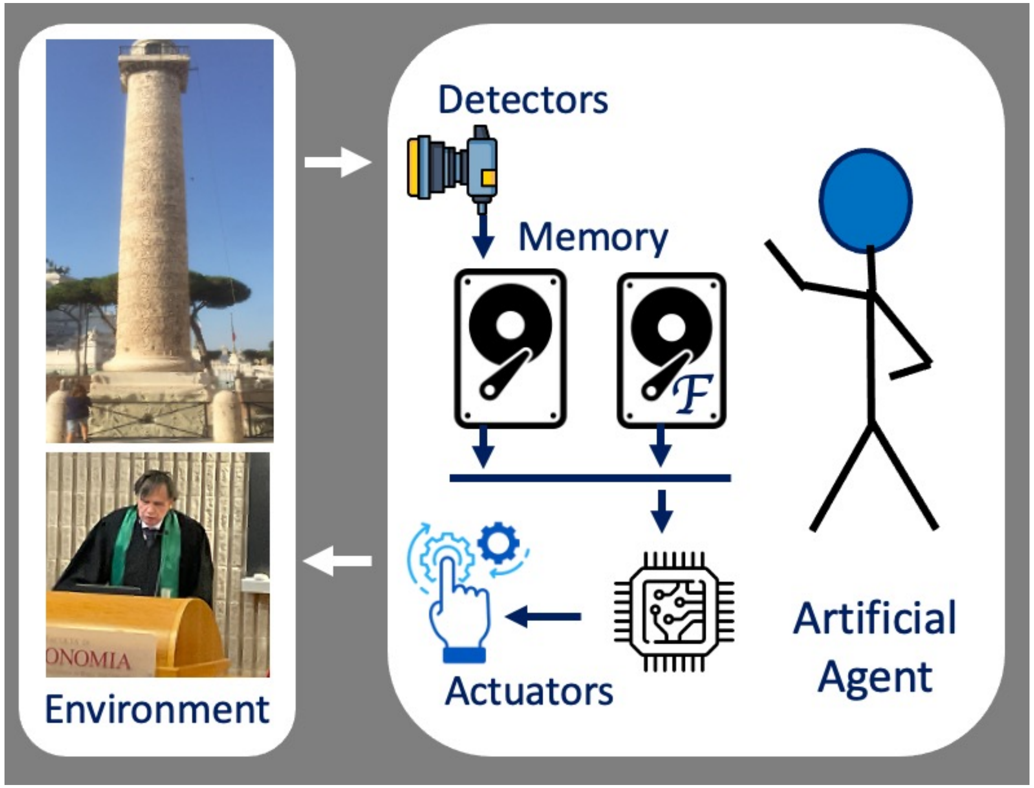

A Modern Approach, S. Russell and P. Norvig, Prentice Hall, 2020

AI & Machine Learning: AI & Operational Definition

An artificial agent is intelligent if it has:

- the ability to receive information and take actions
- a function $\mathcal{F}$ mapping sequences of perceptions to sequences of actions

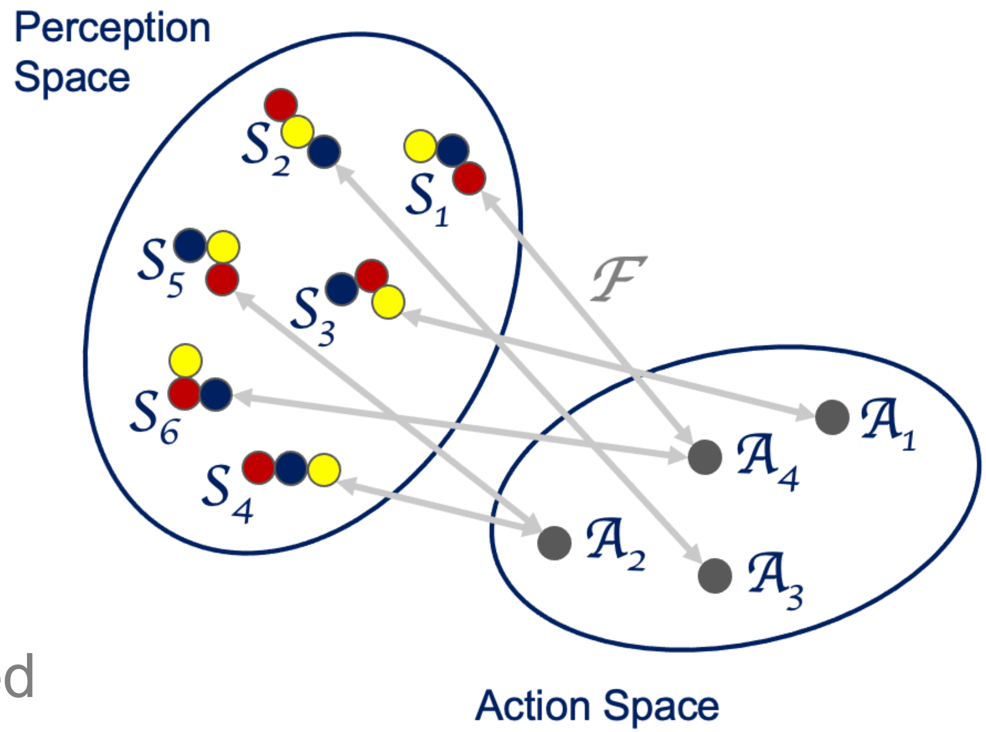

Learning and Operation conceptualized as Computational Tasks

A Modern Approach, S. Russell and P. Norvig, Prentice Hall, 2020

AI & Machine Learning: Intelligence as Computation ?

Q: Yes, but what is intelligence?

The Gentle Vultures

Isaac Asimov 1958

"But if they have discovered atomic energy, where do they conduct their tests, their explosions?"

"On their own planet, sir."

Naron straightened to his full six meters in height and thundered, *"On their own planet?"*

"Yes, sir."

Slowly, Naron took up his pen and drew a line through the latest entry in the little book. It was an unprecedented act, but Naron was very, very wise, and could see the inevitable more clearly than anyone else in the galaxies.
"Idiots!" he muttered.

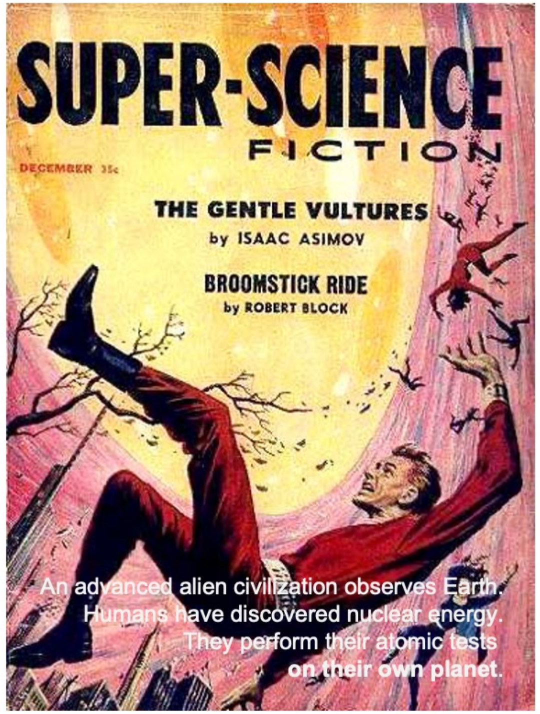

AI & Machine Learning: *Intelligence as Computation ?*

Q: Yes, but what is intelligence?

A: Intelligence is the computational aspect of an agent's ability to achieve goals.

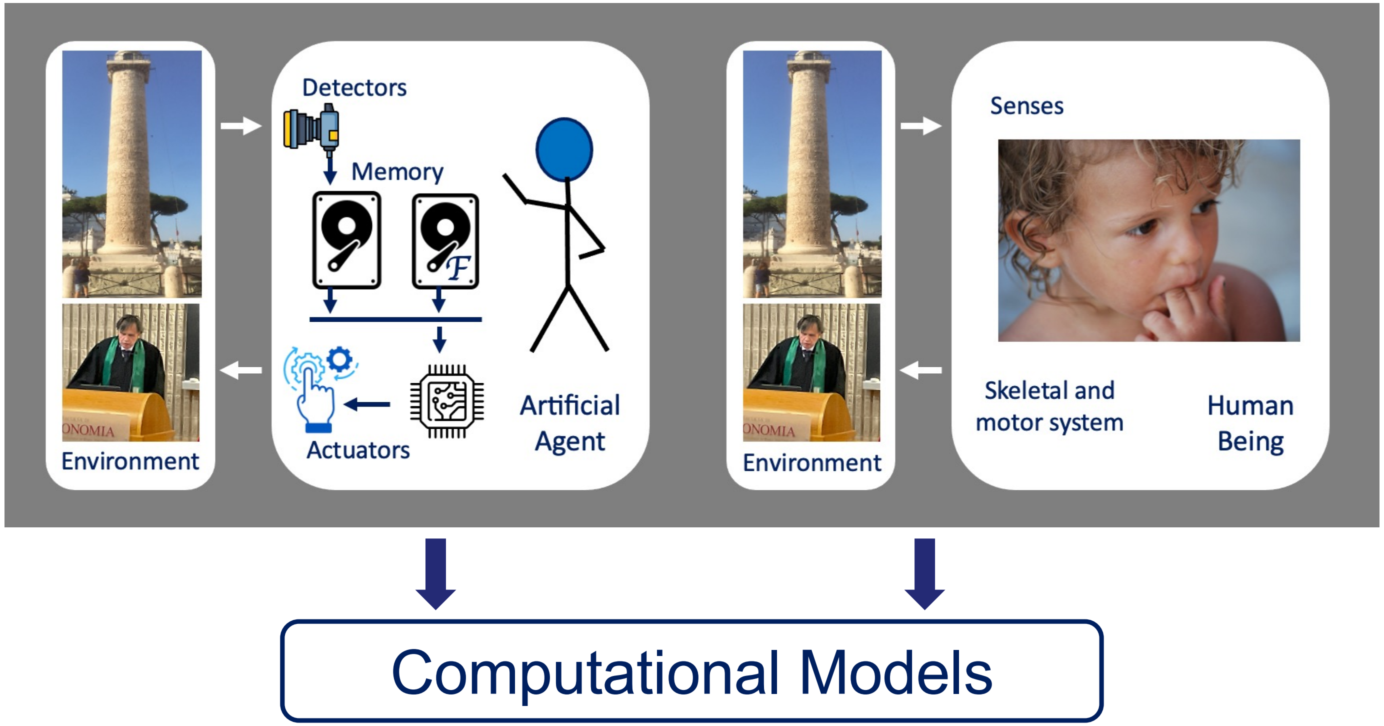

AI & Machine Learning: What is Machine Learning?

Machine learning is a branch of AI focused on building applications that learn from data and improve accuracy over time without explicit programming

Algorithms are trained to find patterns in large datasets to make decisions and predictions

Growth Factors

- Early studies in the 1950s
- Exponential growth in the last 10 years
- Better algorithms
- Greater computing power
- More data

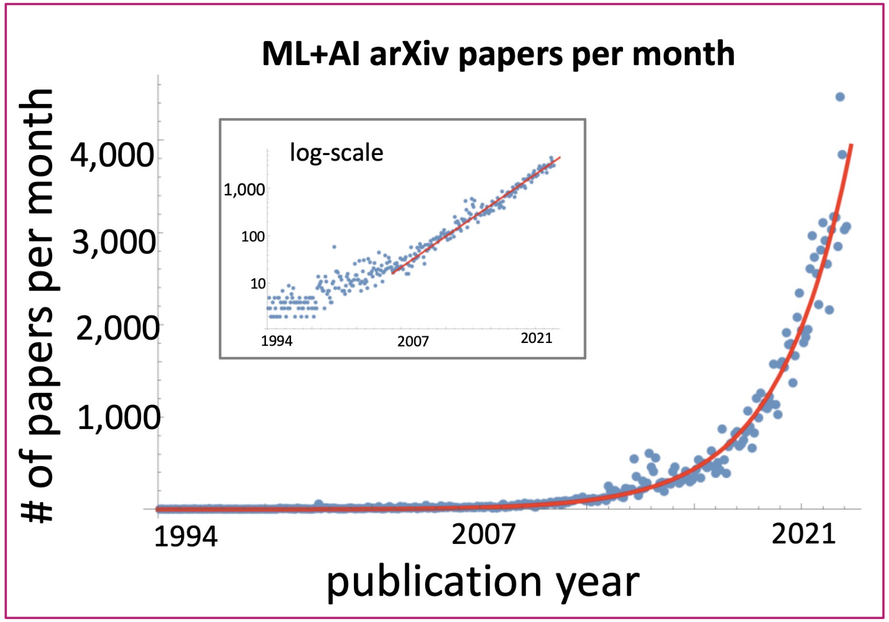

The publications in ML is growing exponentially.

AI & Machine Learning: What is Machine Learning?

- Traditional computing: write a program to get an answer
- Machine Learning: train a model to extract rules from examples

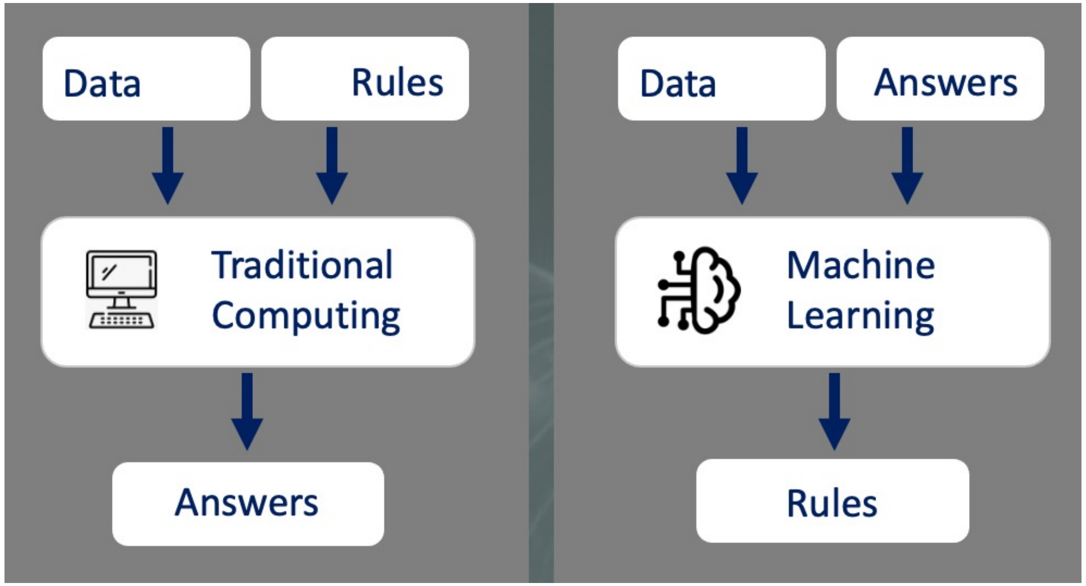

Machine Learning vs Traditional Computing

AI & Machine Learning: What is Machine Learning?

- *Historical view*

Machine learning refers to computational methods that emulate human learning from experience (learning from examples), rather than being explicitly programmed

- *Modern abstraction*

Learning framework

  - T – Task to be performed
  - E – Experience (data or examples)
  - P – Performance measure

A program learns if its performance P on task T improves with experience E.

Clustering

Remark: Many traditional statistical methods (e.g., regression or clustering) can be seen as machine learning when used for data-driven model estimation

AI & Machine Learning: What is Machine Learning?

Linear Regression
- Task (T): Predict a numerical value
- Experience (E): Past data points
- Performance (P): Prediction accuracy

Ellipse Recognition (Geometric Shapes)
- Task (T): Identify ellipses in images
- Experience (E): Labeled images of shapes
- Performance (P): Correct detection rate

Sorting Potatoes by Color (Clustering)
- Task (T): Classify potatoes by color
- Experience (E): Sample labeled potatoes
- Performance (P): Classification accuracy

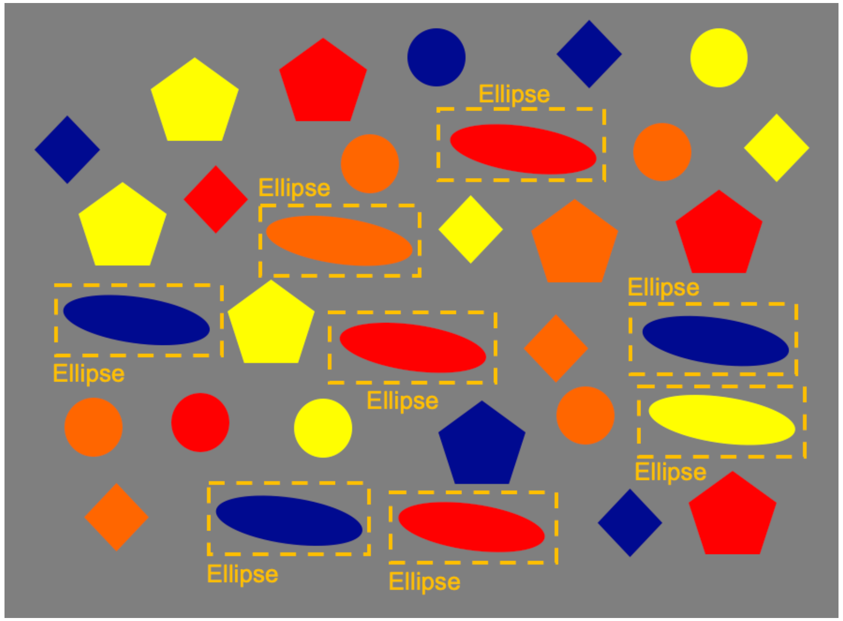

Examples of Learning Algorithms

AI & Machine Learning: What is Machine Learning?

Zoo of Machine Learning Algorithms

- Linear Regression: Predicts values
- Logistic Regression: Predicts probabilities
- Decision Trees: Splits data into branches
- Random Forests: Ensemble of trees
- k-Nearest Neighbors (k-NN): Based on closest examples
- Support Vector Machines (SVM): Finds best boundary
- Naive Bayes: Probability-based
- k-Means Clustering: Groups data into clusters
- PCA: Reduces dimensionality
- Neural Networks: Learns complex patterns

Traditional data analysis algorithms

AI & Machine Learning: What is Machine Learning?

AI & Machine Learning: What is Machine Learning?

Supervised Learning

- Each data point has an associated label (Labeled Data)

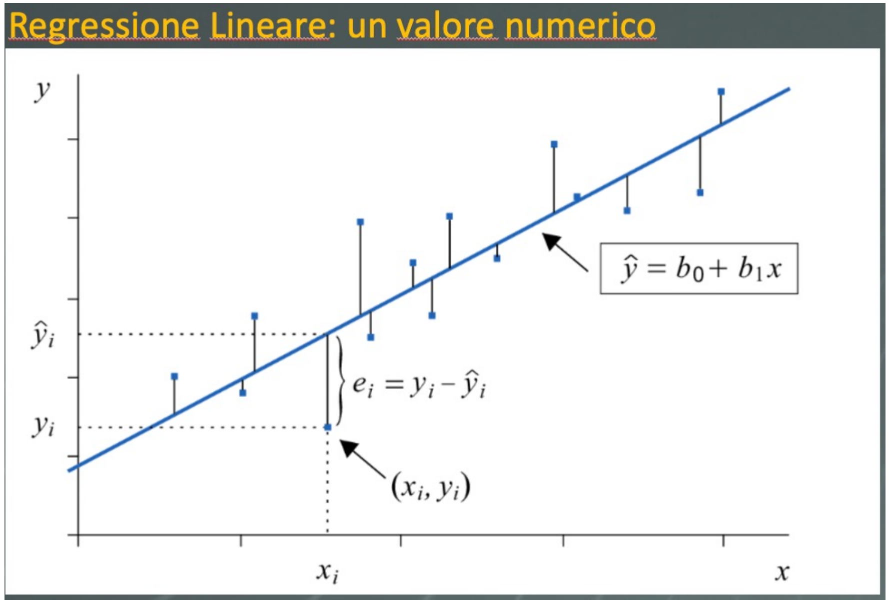

- Goal: minimize classification error or maximize accuracy

AI & Machine Learning: *What is Machine Learning?*

Unsupervised Learning (Self-Supervised)

- No information about the correct answer (Unlabeled Data)

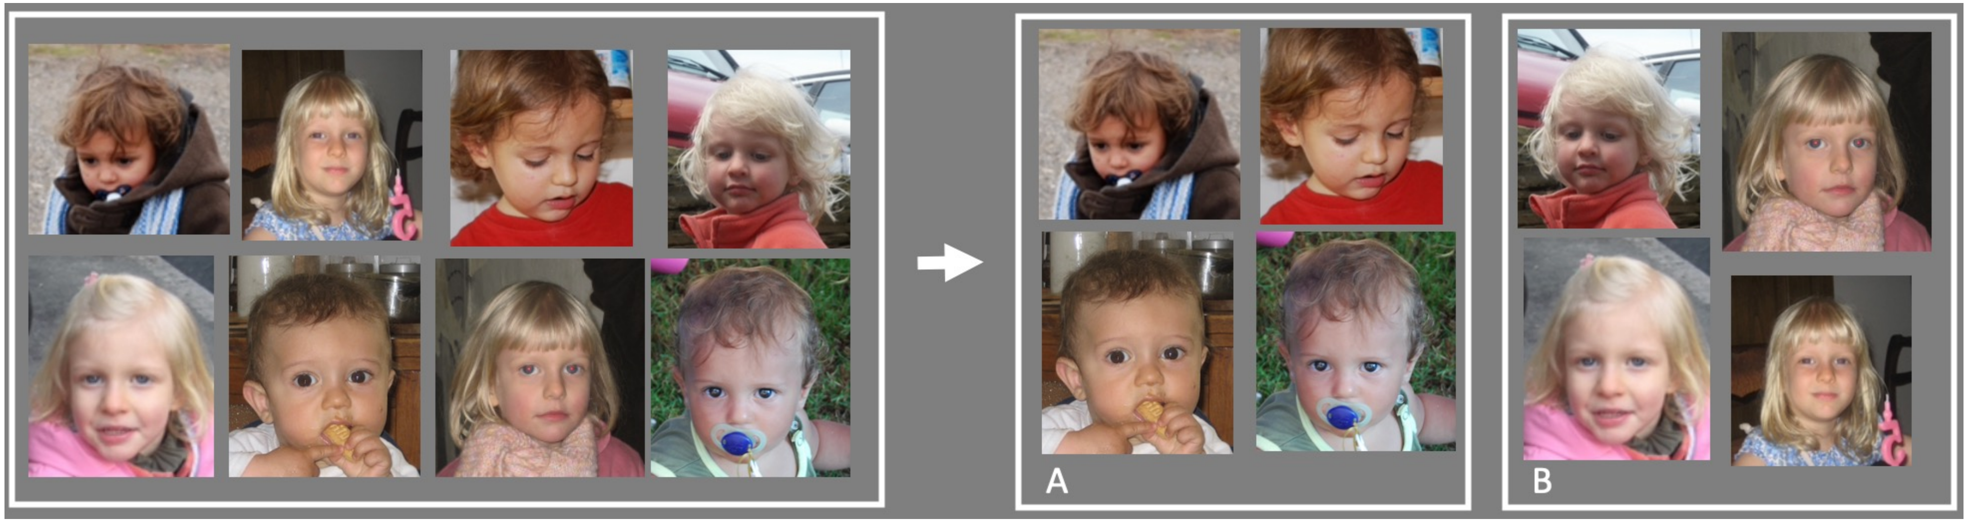

- Goal: create groups/clusters from input data

AI & Machine Learning: What is Machine Learning?

Hybrid Learning

- Uses both labeled & unlabeled data

- Goal: leverage both to improve learning performance

AI & Machine Learning: *What is Machine Learning?*

Reinforcement Learning

- Inspired by behavioral psychology
- No fixed set of examples
- Interaction with the environment
- Continuous feedback: action ↔ state
- Action ↔ Reward or Punishment

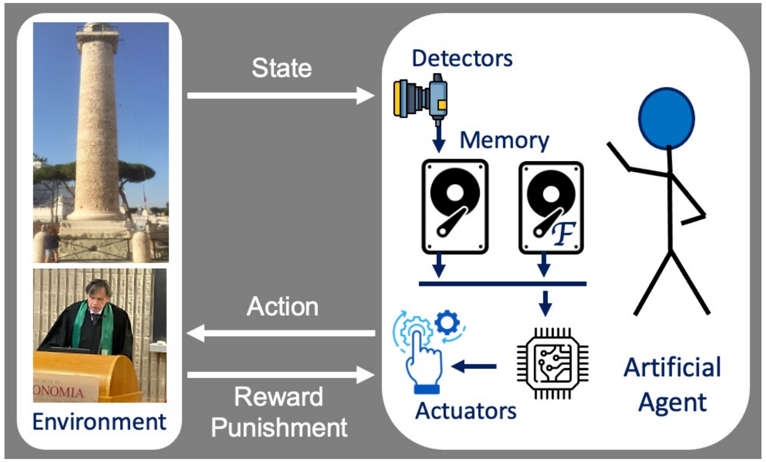

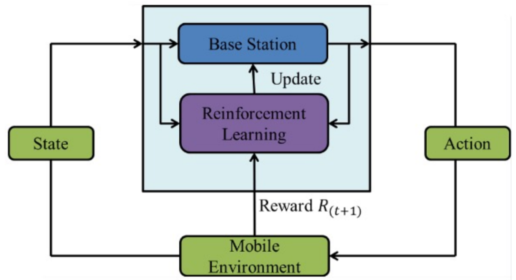

AI & Machine Learning: What is Machine Learning?

- Selection
- Pre-processing
- Transformation
- Data Mining
- Evaluation and Interpretation

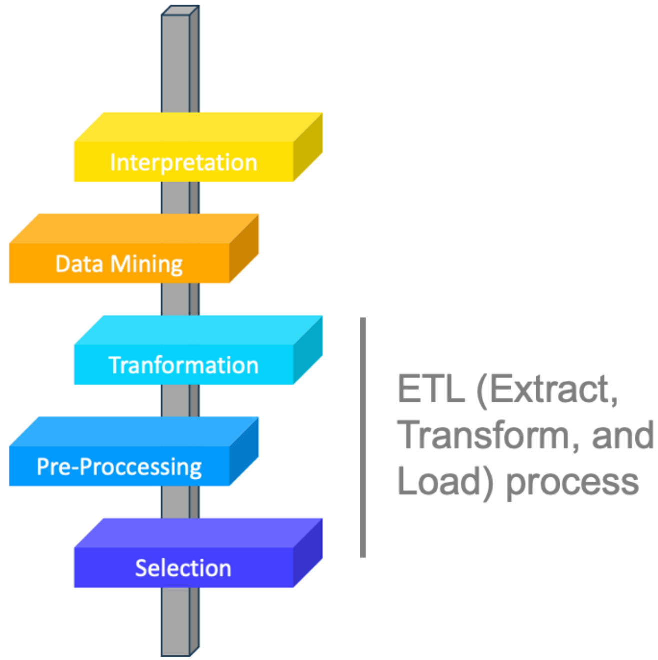

Machine Learning as Knowledge Discovery Process

Recipe for Machine Learning

How machine learning work !

https://www.ibm.com/cloud/learn/machine-learning#toc-what-is-ma-qhM6PX35

There are **five basic steps** for building a machine learning application (or **model**).

These are typically performed by data scientists working closely with the business professionals for whom the model is being developed.

Step 1: Select and prepare a training data set

Training data is a data set representative of the data the machine learning model will ingest to solve the problem it's designed to solve.

- labeled data: 'tagged' to call out features and classifications
- unlabeled data, and the model will need to extract those features and assign classifications on its own.

The training data needs to be properly prepared—randomized, de-duped, and checked for imbalances or biases that could impact the training.

It should also be divided into two subsets: the **training subset**, which will be used to train the application, and the **evaluation subset**, used to test and refine it.

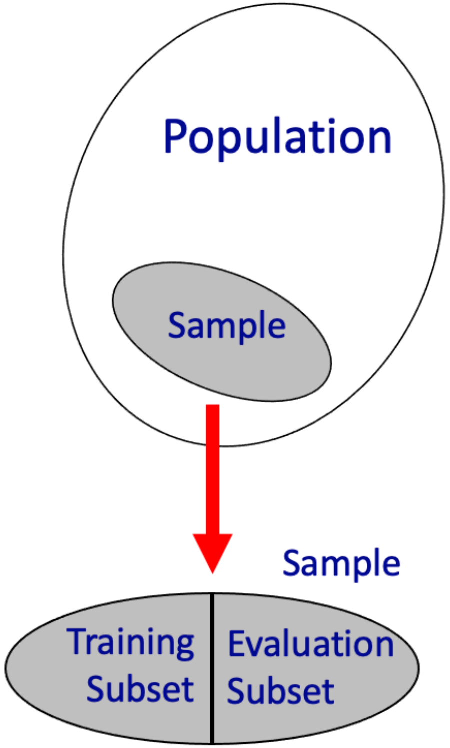

Recipe for Machine Learning

How machine learning work !

There are **five basic steps ...**

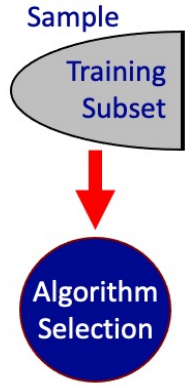

Step 2: Choose an algorithm

Algorithm is a set of statistical processing steps. The algorithm depends on the type and amount data set and on the problem to be solved.

Labeled Data

- **Regression algorithms:** Linear regression predicts the value of a dependent variable based on the value of an independent variable. Logistic regression is used when the dependent variable is binary in nature.
- **Decision trees:** Decision trees use classified data to take decisions based on a set of decision rules.
- **Instance-based algorithms:** K-Nearest Neighbor. It uses classification to estimate how likely a data point is to be a member of one group or another based on its proximity to other points.

Unlabeled Data

- **Clustering algorithms:** Identifying groups of similar records and labeling the records according to the group. This is done without prior knowledge about the groups and their characteristics.
- **Association algorithms:** Association algorithms find patterns and relationships in data and identify frequent 'if-then' relationships called *association rules*. These are similar to the rules used in data mining.
- **Neural networks:** Them were vaguely inspired by the inner workings of the human brain. A neural network is an algorithm that defines a layered network of calculations featuring an input layer, at least one hidden layer, where calculations are performed make different conclusions about input; and an output layer. where each conclusion is assigned a probability.

Recipe for Machine Learning

How machine learning work !

There are **five basic steps ...**

Step 3: Training the algorithm (create the model)

Training the algorithm is an iterative process comparing the output with the results it should have produced, adjusting weights and biases within the algorithm that might yield a more accurate result, and running the variables again until the algorithm returns the correct.

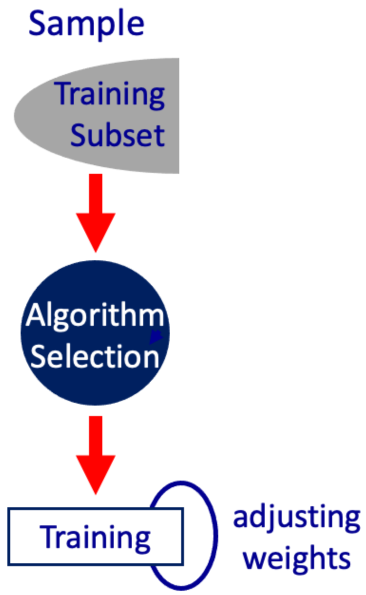

Recipe for Machine Learning

How machine learning work !

There are **five basic steps ...**

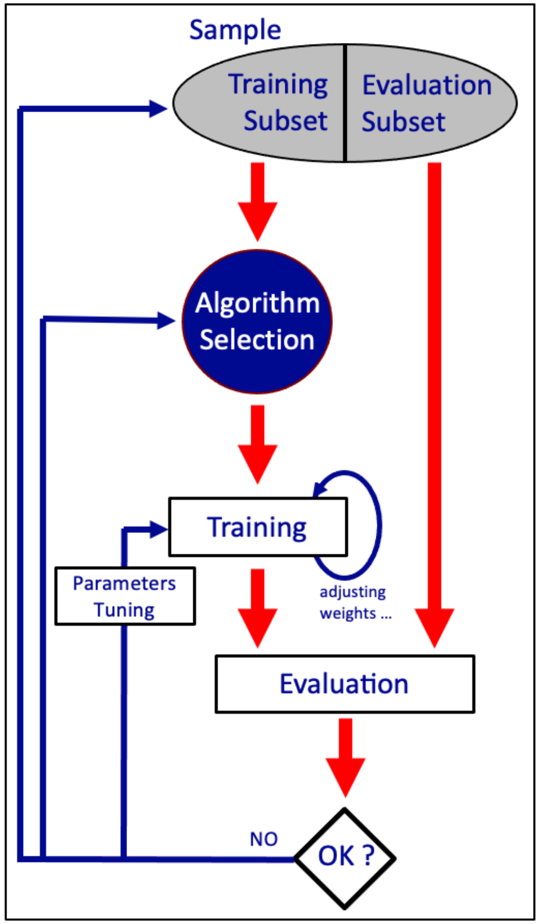

Step 4: Evaluation of the model

Once the model is defined, its performance are evaluate using the *Evaluation Subset*. The hope and goal is that model learns a relationship that generalizes to new examples beyond the *Training Subset*.

Recipe for Machine Learning

How machine learning work !

There are **five basic steps ...**

Step 5: Using and improving the model

The final step is to use the model with new data and, in the best case, for it to improve in accuracy and effectiveness over time.

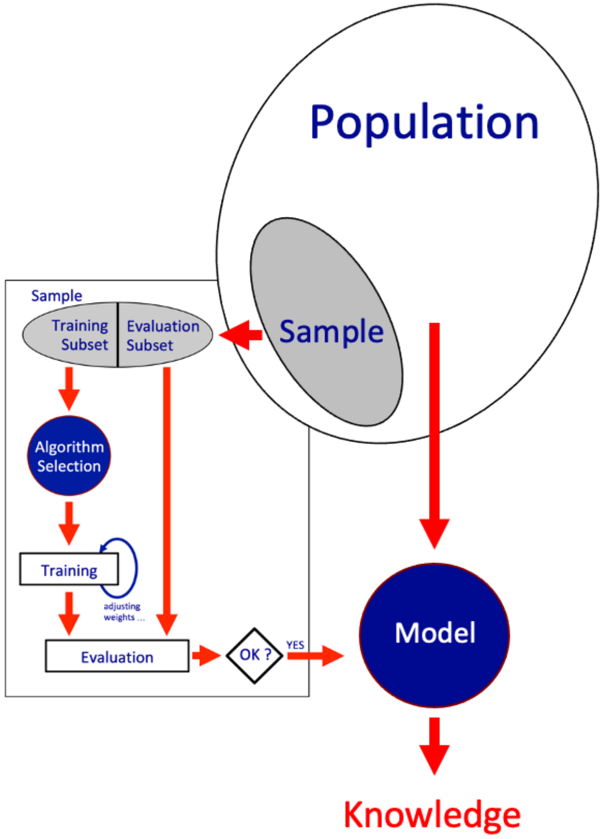

Machine Learning Styles

- Supervised machine learning trains itself on a labeled data set. For example, a computer vision model designed to identify purebred German Shepherd dogs might be trained on a data set of various labeled dog images.
- Unsupervised machine learning uses unlabeled data and uses algorithms to extract meaningful features needed to label, sort, and classify the data in real-time, without human intervention. An unsupervised learning algorithm can analyze huge volumes of emails and uncover the features and patterns that indicate spam.
- Semi-supervised learning offers a happy medium between supervised and unsupervised learning. During training, it uses a smaller labeled data set to guide classification and feature extraction from a larger, unlabeled data set.

Machine Learning as Knowledge Discovery Process

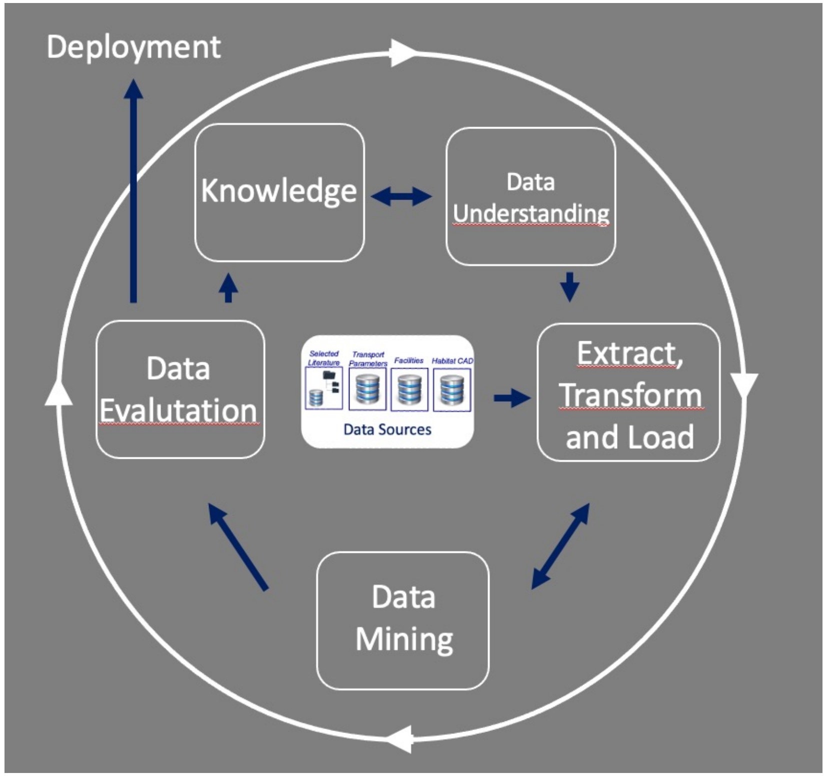

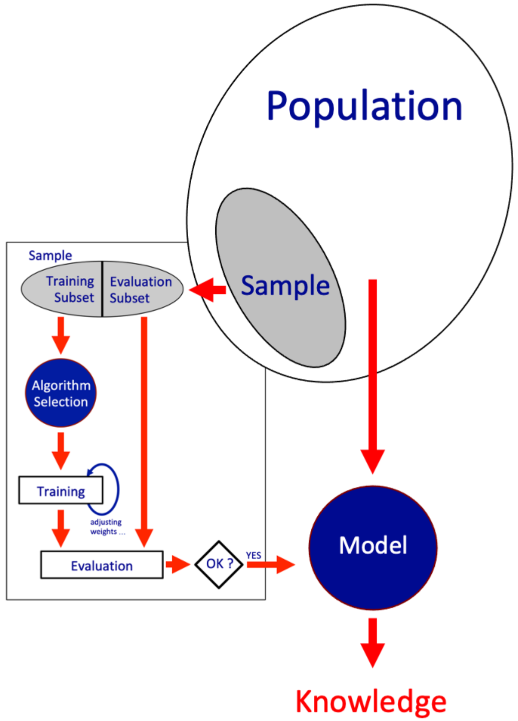

Machine Learning as Knowledge Discovery Process

- Data is being generated fast and needs to be processed quickly
- Online Data Analytics
- Late decisions lead to missed opportunities

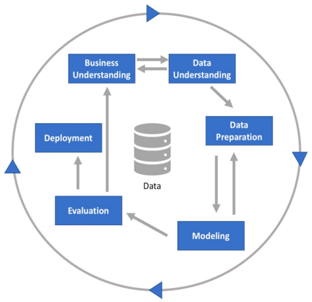

E-Promotions Systems:

- Based on your location, purchase, and likes
- Send promotions when near the store

Healthcare monitoring:

- Sensors track your activities and body metrics
- Immediate response to abnormal readings

Whether an AI/ML approach works depends on the level of error an application can tolerate. In more critical settings, a deeper understanding of the method’s underlying mechanisms is required, potentially demanding substantial methodological improvements.

Machine Learning as Knowledge Discovery Process

We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot.

Chris Anderson

My business has increased by 10%

E-Promotions:
Based on your location, purchase, and likes. Send promotions when near the store

Only 10% of the abnormal situations are detected

Healthcare monitoring:
Sensors track your activities and body metrics. Immediate response to abnormal readings

When is a KDP successful?

Machine Learning as Knowledge Discovery Process

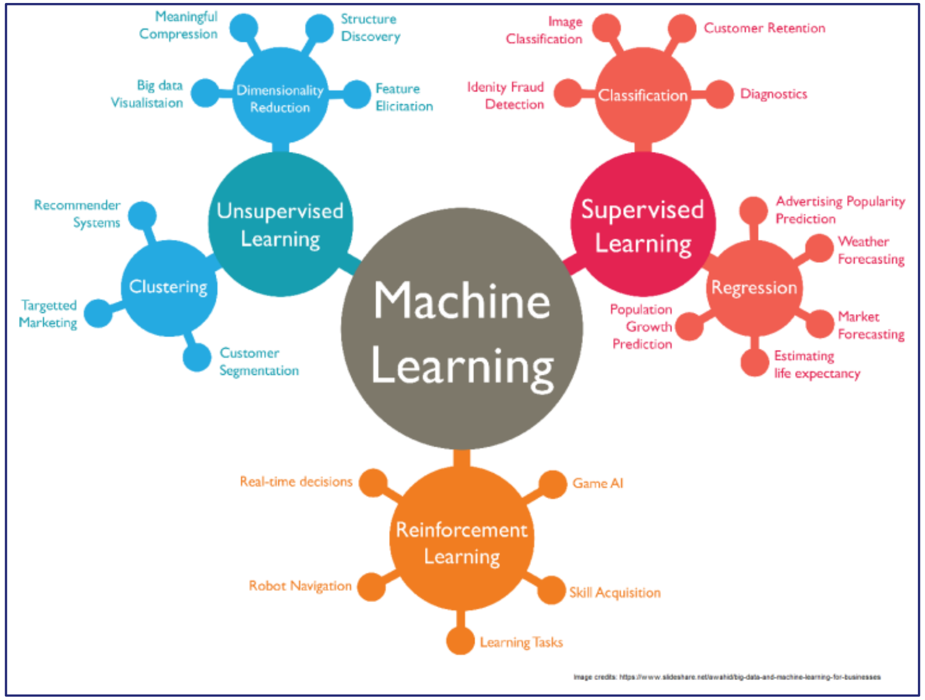

Machine Learning Zoo

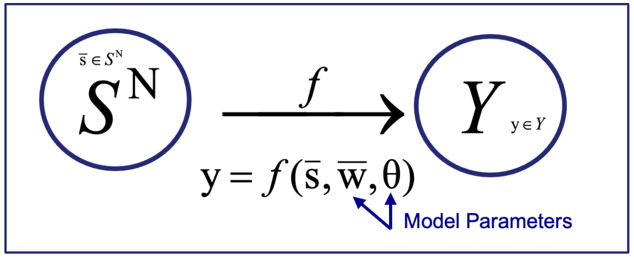

Machine Learning Method

Old wine in a new bottle !

Computability Theory and Computational models

Not everything can be computed

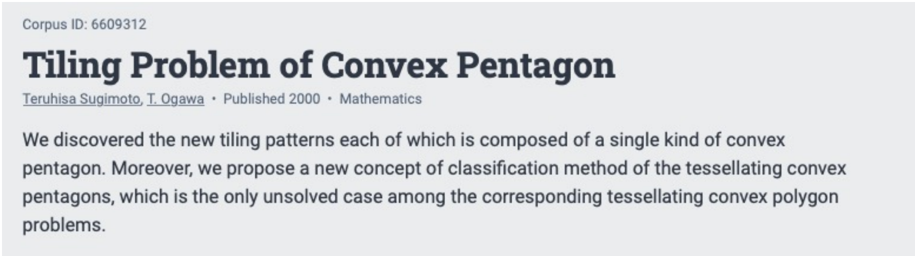
Corpus ID: 6609312

Tiling Problem of Convex Pentagon

Teruhisa Sugimoto, T. Ogawa • Published 2000 • Mathematics

We discovered the new tiling patterns each of which is composed of a single kind of convex pentagon. Moreover, we propose a new concept of classification method of the tessellating convex pentagons, which is the only unsolved case among the corresponding tessellating convex polygon problems.

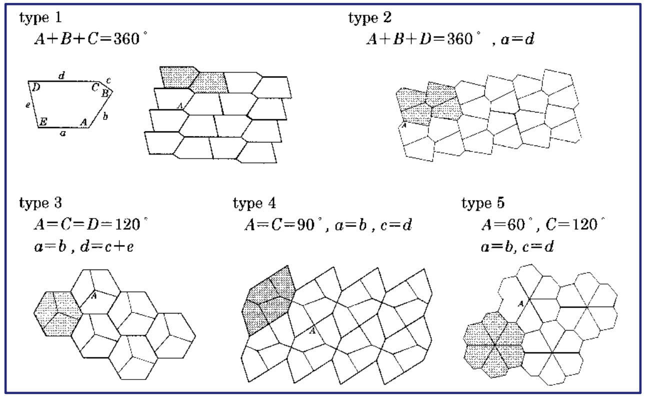

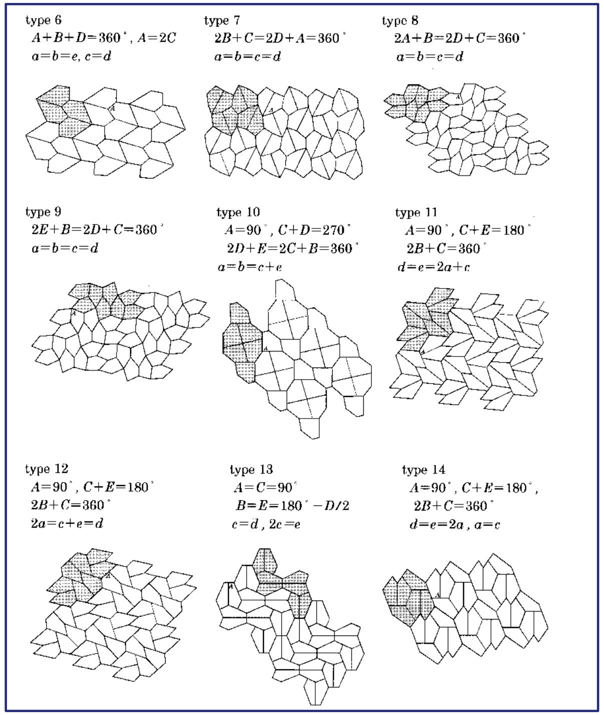

Computability Theory and Computational models

Contrary to popular belief, computer science is not about programming computers. It is about solving problems and determining which problems are possible to solve; the machine or the language used is largely irrelevant.

- For each problem, several algorithms can be used. Which one is the "best"?
- While mathematical elegance may guide algorithm choice, in practice the required resources are more important.
- Time and space are the key resources; impractical requirements make a problem effectively unsolvable.
- Since small problems are generally easy, the focus is on how resource usage scales with the problem size, *N*.

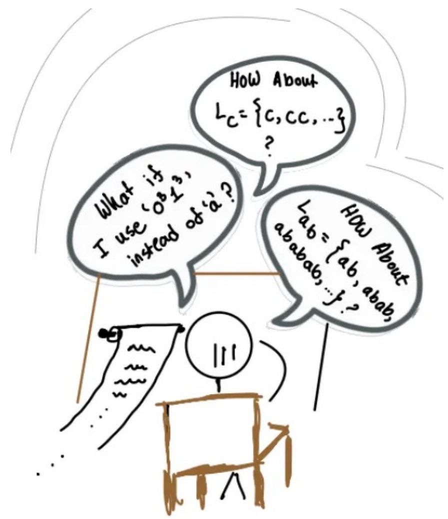

Computability Theory and Computational models

Decision Problem

- For a class of mathematical or logical questions, it is the task of finding an algorithm or procedure that gives a definite “yes” or “no” answer.
- The method performs a finite sequence of steps according to predetermined rules.
- The term is often used for determining whether a given well-formed formula is provable in a formal system.

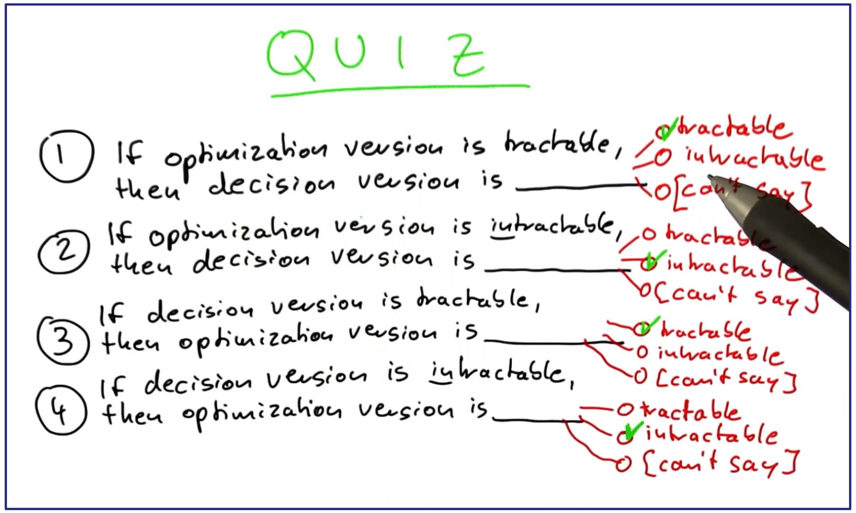

Computability Theory and Computational models

Classifying Problems

- Decision (computational) problems can be classified into **decidable and undecidable**.
- This result was obtained by Turing and others in the 1930', before the invention of computers.
- After the invention of computers, it became clear that it would be useful to **classify decidable problems**, to distinguish **harder problems** from **easier problems**
- This led to the development of computational complexity theory in the 1960's and 1970's

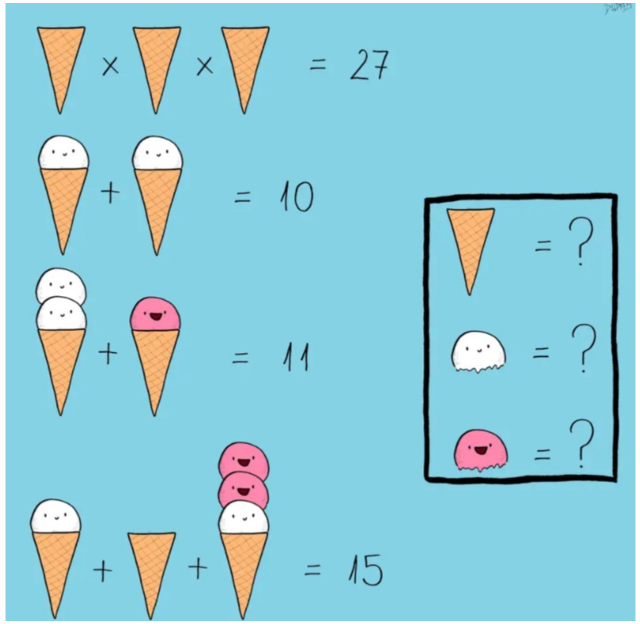

Hard Math Problem

Computability Theory and Computational models

Complexity Measures

Every decidable problem has a set of algorithms that solve it. What property of this set of algorithms could we measure to classify the problem ?

- The difficulty of constructing such an algorithm ?
- The length of the shortest possible algorithm ?
- The efficiency of the most efficient possible algorithm ?

Using a Turing Machine as the computational model, the relevant metrics are computational time and computational space.

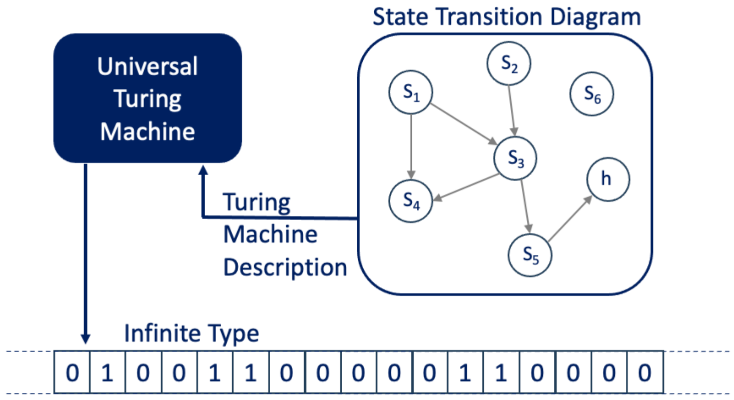

Computability Theory and Computational models

Complexity Measures:

- The most critical computational resource is often time, so the most useful complexity measure is often time complexity.
  - Definition: The time complexity of a Turing Machine T is the function Time_T such that $\mathrm{Time}_T(x)$ is the number of steps taken by the computation $T(x)$.

- Another computational resource is amount of memory used, that is space. The complexity measure is space complexity.
  - Definition: The space complexity of a Turing Machine T is the function Space_T such that $\mathrm{Space}_T(x)$ is the number of distinct cells visited during the computation $T(x)$.

Computability Theory and Computational models

See H. Wang, *Popular Lectures on Mathematical Logic.*
Dover, New York,1981.

We can distinguish two different types of questions:

- The first regards what can be done at all
- The second what can be done with limited resources

One of the simplest unsolvable physical problems is the tiling problem: given a set of shapes, can they tile the plane?

- Easy if the set is restricted (e.g., one polygon).
- In general, determining if an arbitrary set can tile the plane is uncomputable.

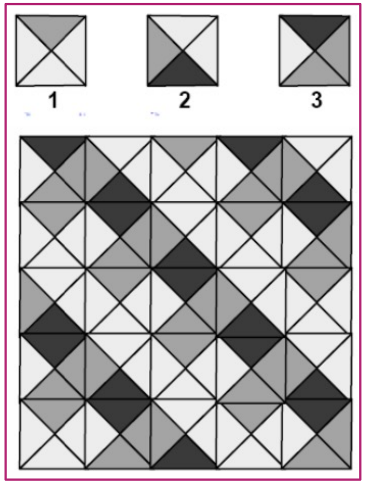

Computability Theory and Computational models

Some problems are harder than others

- Intractable problems are those for which the fastest known algorithms require non-polynomial (exponential) time, typically classified as NP problems
- A problem is tractable if there exists a polynomial-time algorithm that solves it, called a P problem
- The boundary between these classes is not sharp: for many problems, it is unknown whether a polynomial-time solution exists (whether P = NP or not), which is a central question in computational complexity theory

The Towers of Hanoi problem

Intractable Problem

Only one plate can be moved at a time, and larger plates cannot go on smaller ones

Computability Theory and Computational models

Some problems are harder than others:

The Max & Min problem

$\{x_i\} \rightarrow x_i \in A; x_i > 0; \quad i = 1,2,3, \ldots, N$

$\{x_i^B\}_k \rightarrow x_{i,k}^B \in B; \quad i = 1,2,3 \ \ldots, n_1; \quad k = 1,2,3, \ldots, 2^N$

$\{x_i^C\}_k \rightarrow x_{i,k}^C \in C; \quad i = 1,2,3 \ \ldots, n_2; \quad k = 1,2,3, \ldots, 2^N$

$A = B \cup C; \quad B \cap C = \emptyset;$

$B \neq \emptyset; \quad C \neq \emptyset;$

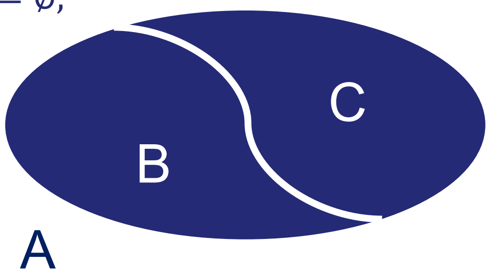

Partition the set A into subsets B and C so as to maximize or minimize Δ_k

$$\Delta_k = \left| \left(\sum_i^{n_1} x_{i,k}^B \right) - \left(\sum_i^{n_2} x_{i,k}^C \right) \right|$$

$$\Delta_{max} = \max_k \Delta_k$$

$$\Delta_{min} = \min_k \Delta_k$$

Since there are 2^Npartitions, the algorithm runs in $\mathrm{T(x)} = O(2^N)$ time $\longleftrightarrow$ Intractable Problem

Computability Theory and Computational models

Some problems are harder than others:

Is the Max problem really intractable ?

To answer we use a **magic guess-and-check** procedure:

- Consider the **N-1** partitions in which one subset (say, **B**) contains only a single element.
- It is easy to see that the partition where this singleton element is $x_1^B = \min_i\{x_i\}$ maximizes

$$\Delta_k = \left|\left(\sum_i^{n_1} x_{i,k}^B\right) - \left(\sum_i^{n_2} x_{i,k}^C\right)\right|$$

Check

Suppose we have an arbitrary partition $\boldsymbol{a}$.

Let C be the subset with the smaller sum.

$$\left(\sum_i^{n_1} x_{i,a}^B\right) > \left(\sum_i^{n_2} x_{i,a}^C\right)$$

If C contains more than one element, take any element $x_i \in C$ and move it to B.

The new difference becomes:

$$\left(\sum_i^{n_1} x_{i,b}^B\right) - \left(\sum_i^{n_2} x_{i,b}^C\right) =$$

$$\left(\sum_i^{n_1} x_{i,a}^B + x_i\right) - \left(\left(\sum_i^{n_2} x_{i,a}^C\right) - x_i\right) = \left(\sum_i^{n_1} x_{i,a}^B\right) - \left(\sum_i^{n_2} x_{i,a}^C\right) + 2x_i$$

Since $x_i > 0$, the difference **always increases**.

Repeating this process, the lighter subset is reduced to a **single element** in the optimal partition. QED.

Computability Theory and Computational models

Why the "Magic" Works: The MAX partition problem is easy not because of a clever algorithm, but because the structure of optimal solutions can be characterized. This collapses an exponential search space to a linear one.

Key message: *The apparent magic comes from building a conceptual model of the problem. This step is often dismissed as unnecessary in data-driven approaches, but without it no genuine understanding, and often no efficient solution, is possible.*

Takeaway

- Algorithms explore solution spaces
- Models explain why most of that space is irrelevant
- Computational power cannot replace conceptual understanding

Computability Theory and Computational models

Some problems are harder than others:
Levinthal paradox

Protein Folding

Atomistic Approach

If a protein has 100 amino acids, and each amino acid has n (n~2) conformational states, then the protein has n^{100} conformational states, of which only one is the native state.

Therefore: Proteins cannot sample all possible conformations while folding.

Therefore: a folding pathway must exist.

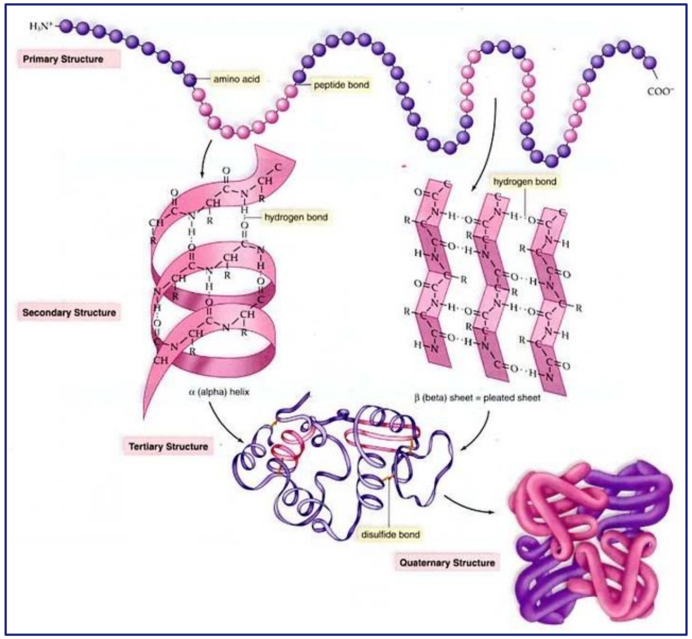

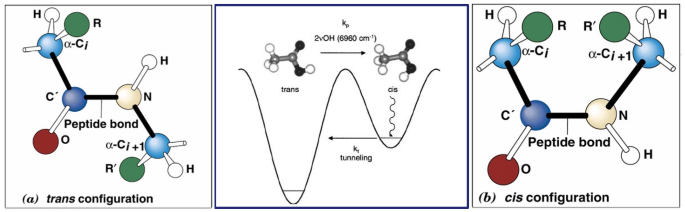

Intractable Problems

Computability Theory and Computational models

Some problems are harder than others:

Problems that can be solved in *nondeterministic* polynomial time

- If we could try all possible solutions simultaneously, we could identify the correct one in polynomial time.
- Alternately: If we had a magic procedure that always makes the correct choice at each decision point, we could design an algorithm that solves the problem in polynomial time.

Spin Systems

... the partition function depends on 2^N configurations, each of which is specified by N numbers $\{ s_i \}$, the energy levels are in general highly degenerate.

For example, the 10 × 10 square lattice with open boundary conditions $2^{100} \sim 10^{30}$ possible configurations, but only about 20000 ~ 10^4 different energy levels.

Even with the largest computer, 10^{30} is too many calculations, but managing a polynomial with 10^4 terms is practicable. ...

MonteCarlo Methods …

Some notes on Computational Paradigms

Analog Computing and Neural Networks

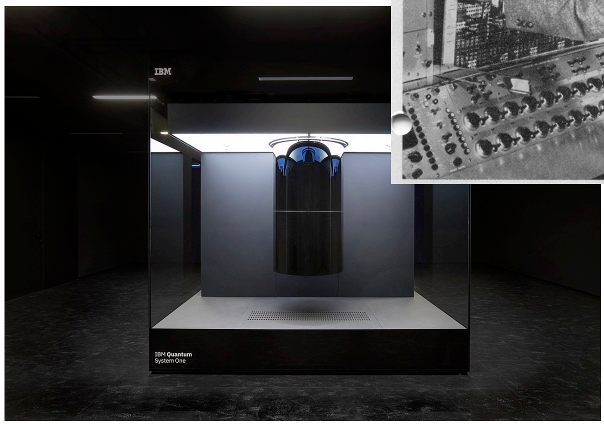

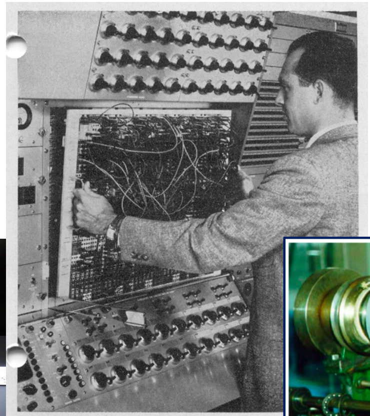

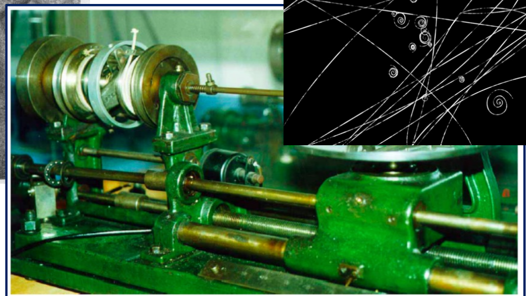

Fig. Vista delle ruote e dischi di integrazione della macchina. I dispositivi che sembrano dei motori sulla sinistra sono gli "amplificatori di torsione di Nieman", dispositivi a guida d'argano che funzionavano come amplificatori meccanici per poter procedere nelle successive operazioni di integrazione senza perdita di accuratezza.

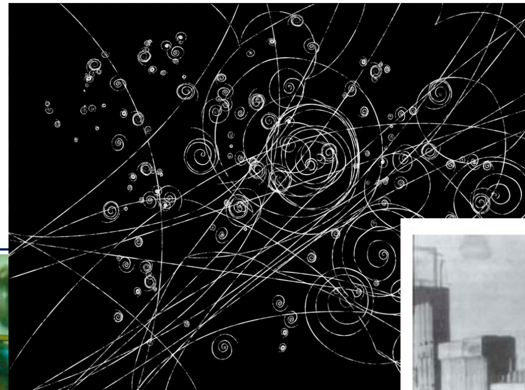

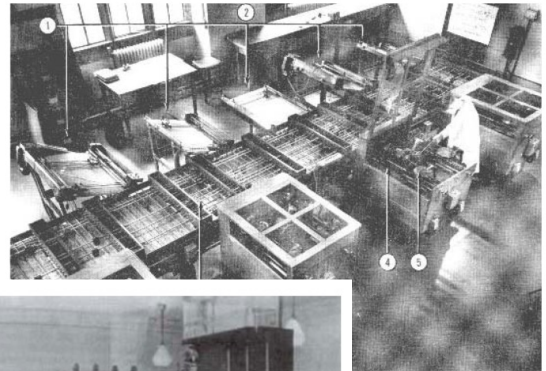

Some notes on Computational Paradigms

Computational Models

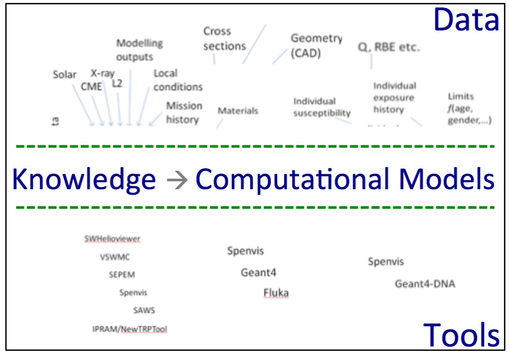

Computational models are mathematical models that are simulated using computation to study complex systems.

…

The parameters of the mathematical model are adjusted using computer simulation to study different possible outcomes.

www.nature.com/subjects/computational-models

This may be a restrictive definition

Computational Model means:

- Heuristic approach: data mining techniques (classification, regression, clustering) applied to transformed data under explicit researcher-defined hypotheses; model selection and hyperparameter tuning are driven by KDP objectives.
- Theoretical approach: transformed data are used to identify and estimate parameters of a theoretical dynamical model; the KDP coincides with the formal understanding of the system's physical dynamics.
- Statistical approach: transformed data are processed using data-driven machine learning models (e.g., neural networks), where the hypothesis space is implicit and learned from training examples rather than explicitly specified.

Some notes on Computational Paradigms

Data do not speak for themselves.
They acquire meaning only within a well-defined conceptual model.

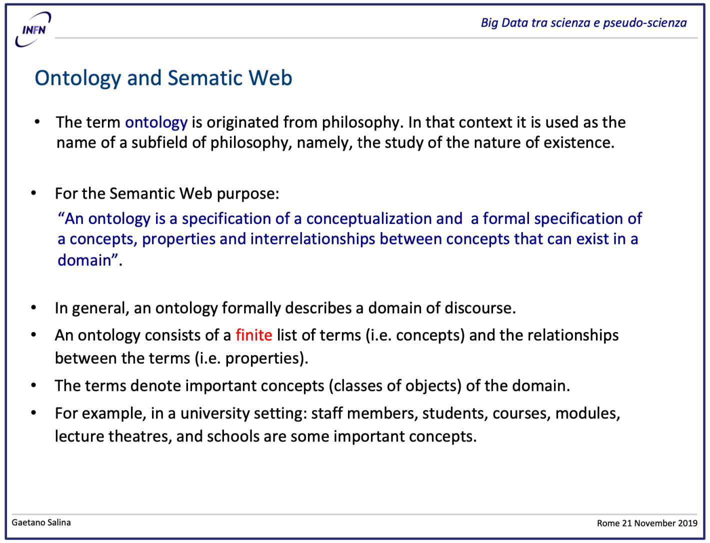
Big Data tra scienza e pseudo-scienza

Ontology and Sematic Web

- The term ontology is originated from philosophy. In that context it is used as the name of a subfield of philosophy, namely, the study of the nature of existence.
- For the Semantic Web purpose:
  "An ontology is a specification of a conceptualization and a formal specification of a concepts, properties and interrelationships between concepts that can exist in a domain".
- In general, an ontology formally describes a domain of discourse.
- An ontology consists of a finite list of terms (i.e. concepts) and the relationships between the terms (i.e. properties).
- The terms denote important concepts (classes of objects) of the domain.
- For example, in a university setting: staff members, students, courses, modules, lecture theatres, and schools are some important concepts.

Gaetano Salina Rome 21 November 2019

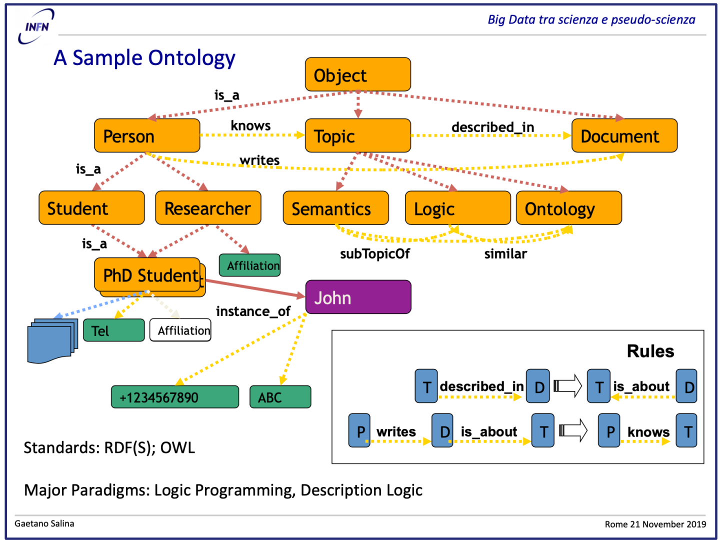

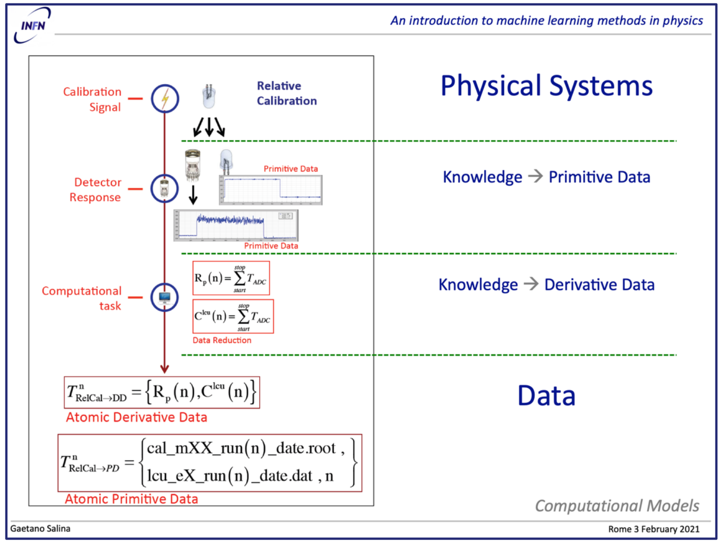

Some notes on Computational Paradigms

Do Computational Models mean Algorithms ?

Algorithm:

An algorithm is a finite sequence of well-defined, computer implementable instructions, typically to solve a class of problems or to perform a computation.
Algorithms are always unambiguous and are used as specifications for performing calculations, data processing, automated reasoning, and other tasks.

Computational Task:

Compare two words, S_1 e S_2 and determine if $S_1 = S_2$ or $S_1 \mathrel{!}= S_2$.

$Read\ S_1;\ Read\ S_2;$

$Evaluate\ L_1(S_1);\ Evaluate\ L_2(S_2);$

$if\ (L_1 \mathrel{!}= L_2)\ \ than$

$\quad print("S_1\ \mathrel{!}= S_2");$

$\quad Exit;$

$else$

$\quad for\ i\ =\ 1\ to\ L_1$

$\quad\quad if\ (S_1(i) \mathrel{!}= S_2(i))\ \ than$

$\quad\quad\quad print("S_1\ \mathrel{!}= S_2");$

$\quad\quad\quad Exit;$

$\quad\quad endif$

$\quad endfor$

$\quad print("S_1 = S_2");$

$endif$

Some notes on Computational Paradigms

Non-algorithmic computation !

Where is the yellow ellipse ?

What algorithm did your brain run?

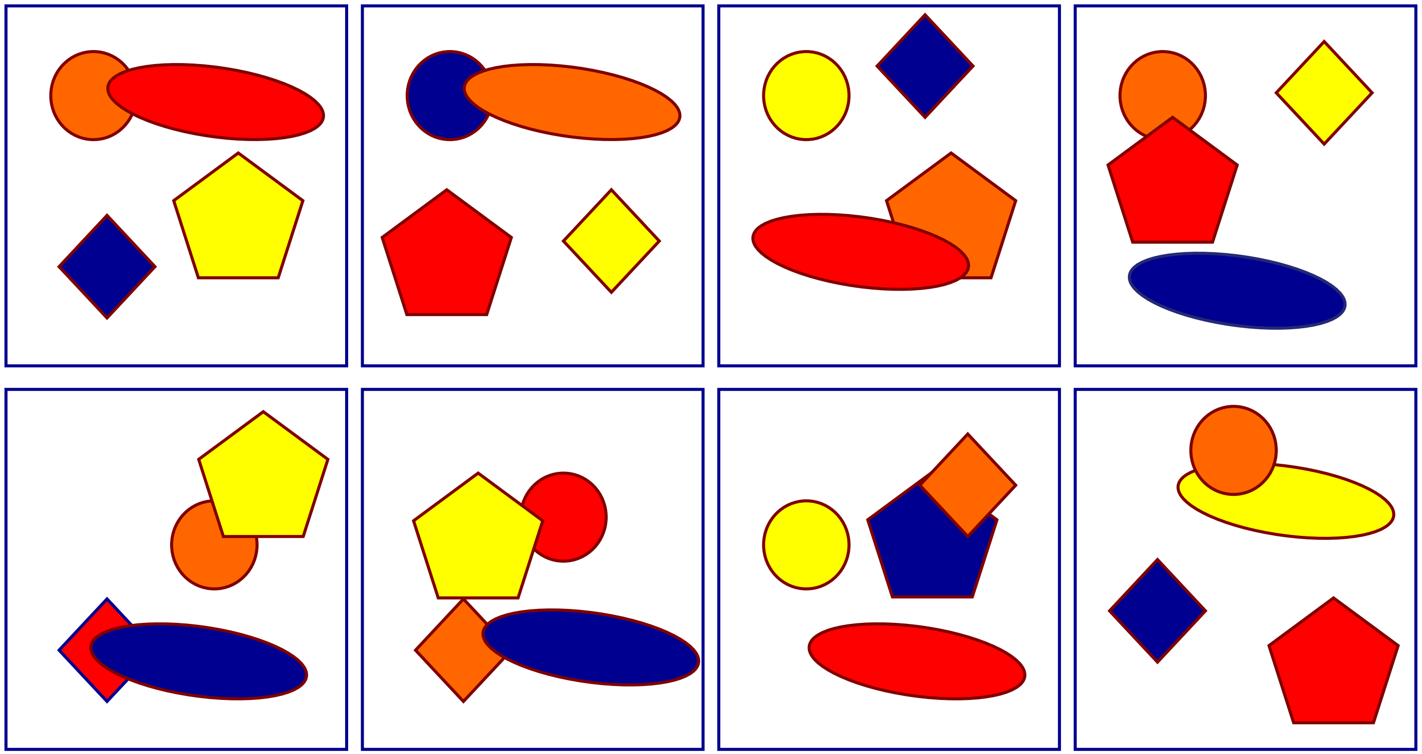

Some notes on Computational Paradigms

Are Computational Models just Algorithms ?

Computational Task:
Compare two words, S_1 e S_2 and determine if S_1 = S_2 or S_1 != S_2.

S_1 = Esprimeremo questo algoritmo in un linguaggio un po' rigido, questa descrizione informale è il punto di partenza anche per il disegno di nuovi algoritmi.

S_2 = Esprimeremo questo algoritmo in un limguaggio un po' rigido, cuesta descrizione informale è il punto di partenza anche per il disegno di guovi algoritmi.

A

S_1 = COLORE

S_2 = FELICE

B

Algorithm

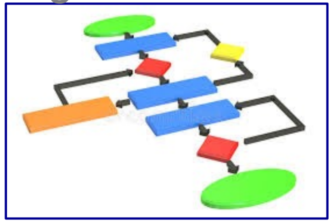

$$A) \Rightarrow (S_1 \neq S_2)$$

$$B) \Rightarrow (S_1 \neq S_2)$$

Brain

Non-algorithmic computation !

$$A) \Rightarrow (S_1 \neq S_2)$$

$$B) \Rightarrow (S_1 = S_2)$$

Except for typos

Some notes on Computational Paradigms

A look at the past to guess the future

- bubble chamber
  - Vessel, filled (e.g.) with liquid hydrogen at a temperature above the normal boiling point but held under a pressure of about 10 atmospheres by a large piston to prevent boiling.
  - When particles have passed, and possibly interacted in the chamber, the piston is moved to reduce the pressure, allowing bubbles to develop along particle tracks.
  - After about 3 milliseconds have elapsed for bubbles to grow, tracks are photographed using flash photography. Several cameras provide stereo views of the tracks.
  - The piston is then moved back to recompress the liquid and collapse the bubbles before boiling can occur.
- Invented by Glaser in 1952 (when he was drinking beer)

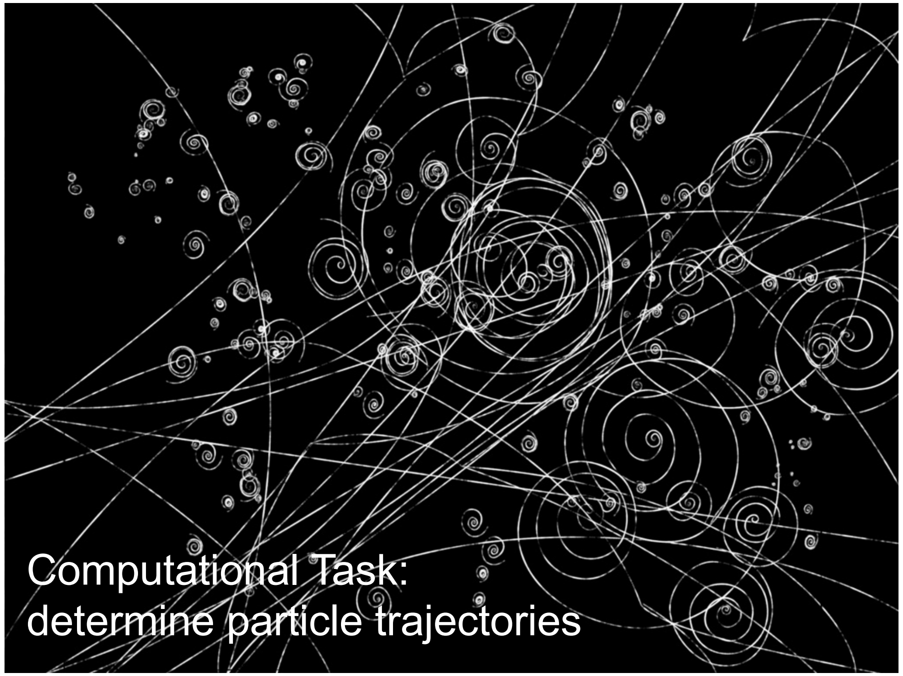

Bubble Chamber as Analog Computer

Some notes on Computational Paradigms

A look at the past to guess the future

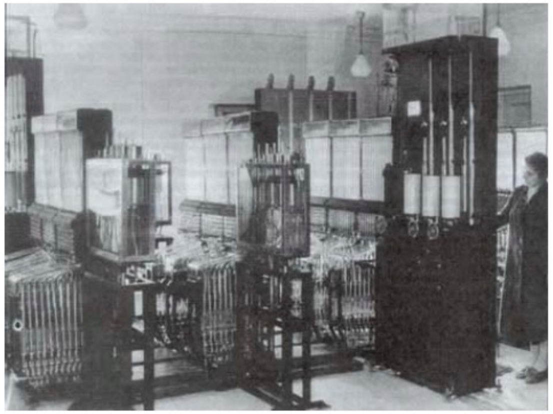

The Water Integrator was an early analog computer, built in the Soviet Union in 1936 by Vladimir Lukyanov. Using interconnected pipes and pumps, water levels represented numbers and flows represented operations, allowing the machine to solve inhomogeneous differential equations.

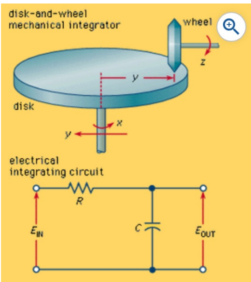

The Differential Analyser is a mechanical analogue computer designed to solve differential equations by integration, using wheel-and-disc mechanisms to perform the integration. It was one of the first advanced computing devices to be used operationally.

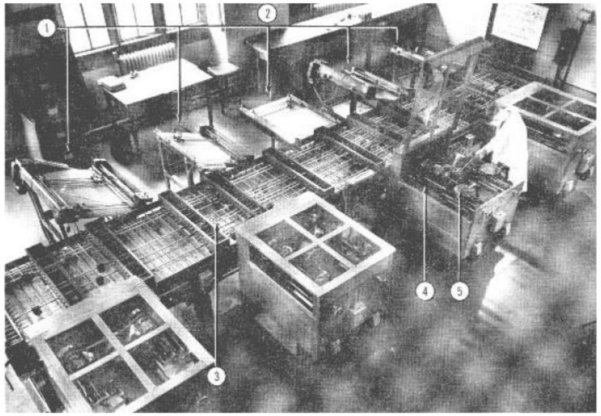

Some notes on Computational Paradigms

A look at the past to guess the future

The analog computer

The analog computer → any device that 'computes' by means of an analog between real, physical, CONTINUOUS, quantities and some other set of variables.

An analog/analogue computer → is a form of computer that uses electronic or mechanical phenomena to model the problem being solved by using one kind of physical quantity to represent another.

Real quantities → the distance between points on a scale, the angular displacement, the velocity, or the acceleration of a rotating shaft, a quantity of some liquid, the electrical current in a conductor.

Shannon introduced the General Propose Analog Computer (GPAC) as a mathematical model of an analog device, the Differential Analyzer in 1941. The Differential Analyzer was used from the 30s to the early 60s to solve numerical problems, especially differential equations, for example, in ballistics problems. These devices were first built with mechanical components and later evolved to electronic versions. A GPAC may be seen as a circuit built of interconnected black boxes, whose behavior is given by Fig. 1, where inputs are functions of an independent variable called the time. Shannon's model was extended by others to neural networks and extended Analog computers.

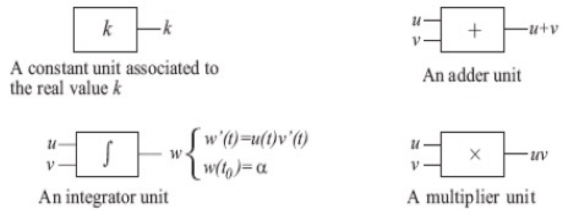

Analog Computing

Some notes on Computational Paradigms

A look at the past to guess the future

The Computational Tasks are defined by the connection topology of the *elementary* computing elements (Shannon)

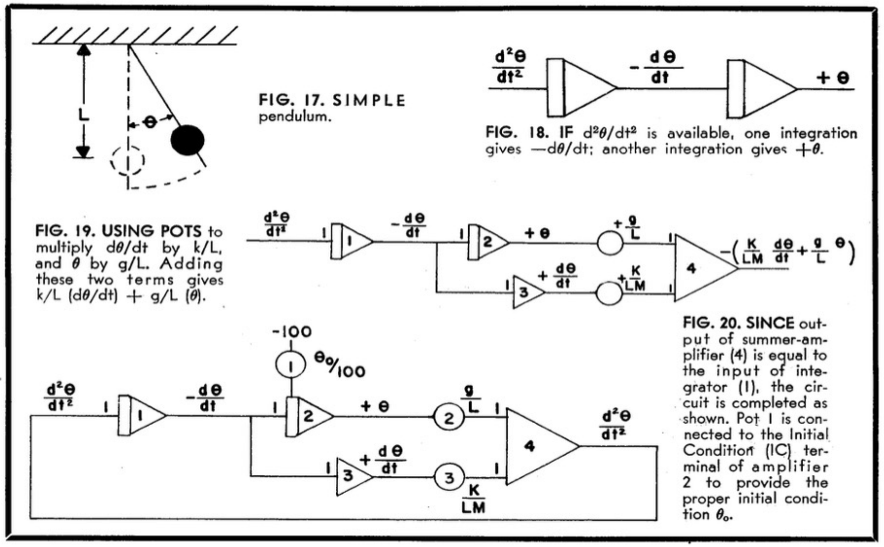

FIG. 17. SIMPLE pendulum.

FIG. 18. IF $d^2\theta/dt^2$ is available, one integration gives $-d\theta/dt$; another integration gives $+\theta$.

FIG. 19. USING POTS to multiply $d\theta/dt$ by k/L, and θ by g/L. Adding these two terms gives $k/L\,(d\theta/dt) + g/L\,(\theta)$.

FIG. 20. SINCE output of summer-amplifier (4) is equal to the input of integrator (1), the circuit is completed as shown. Pot 1 is connected to the Initial Condition (IC) terminal of amplifier 2 to provide the proper initial condition θ_0.

It all began in 1955 with the very first analog computer ever built by Telefunken:

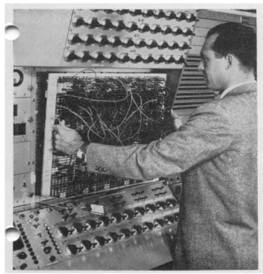

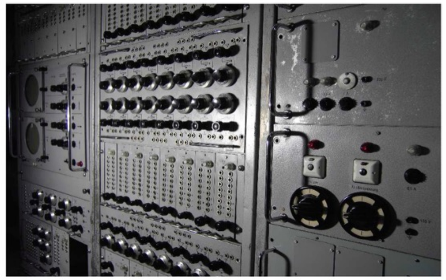

Analog Computers
1950-1970

Some notes on Computational Paradigms

A look at the past to guess the future

Analog Computers
1950-1970

The Computational Tasks are defined by the connection topology of the *elementary* computing elements

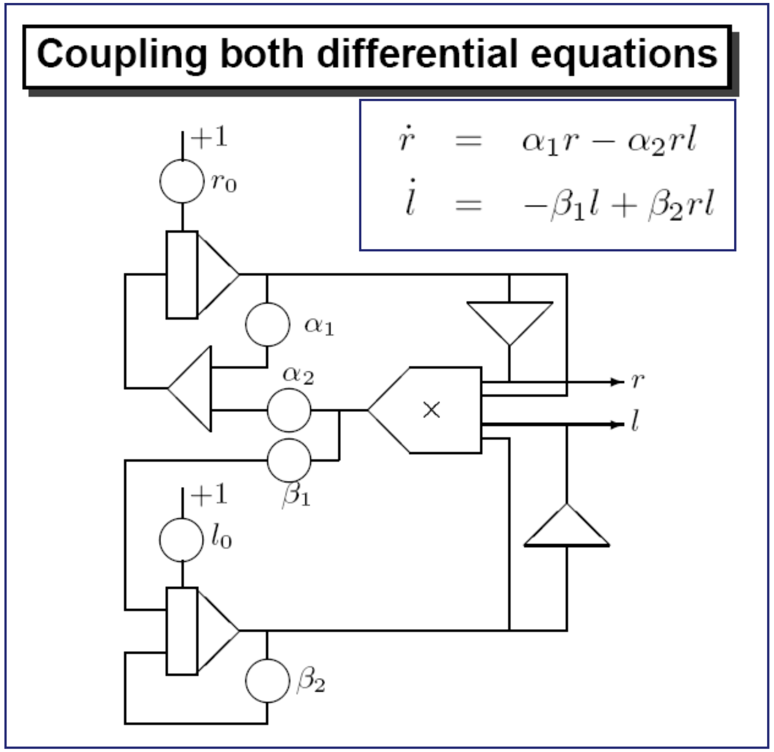

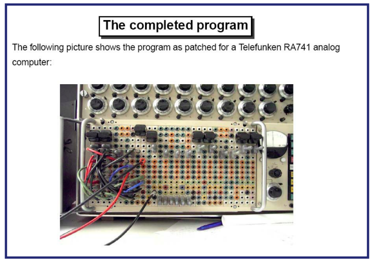

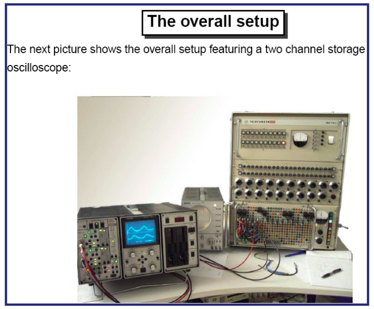

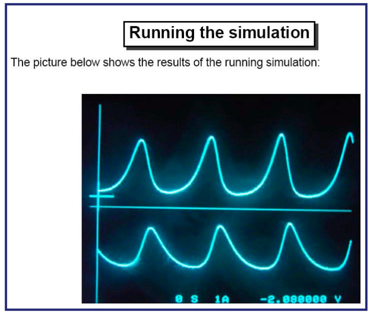

The Computational Tasks are defined by the connection topology of the *elementary* computing elements

Some notes on Computational Paradigms

A look at the past to guess the future

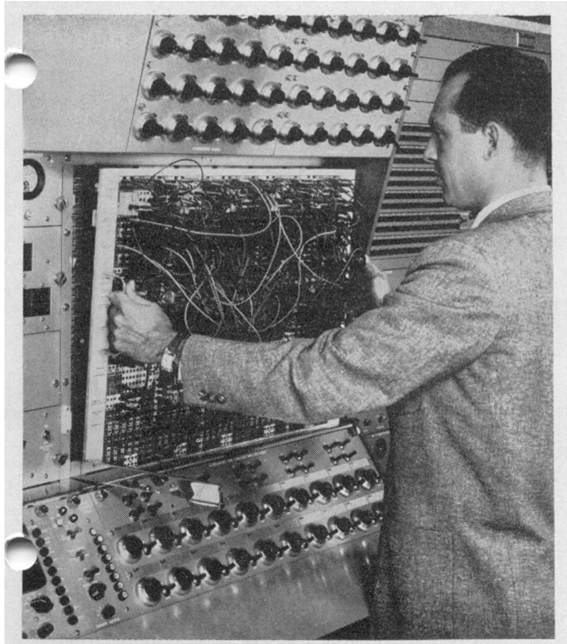

Analog Computers
1950-1970

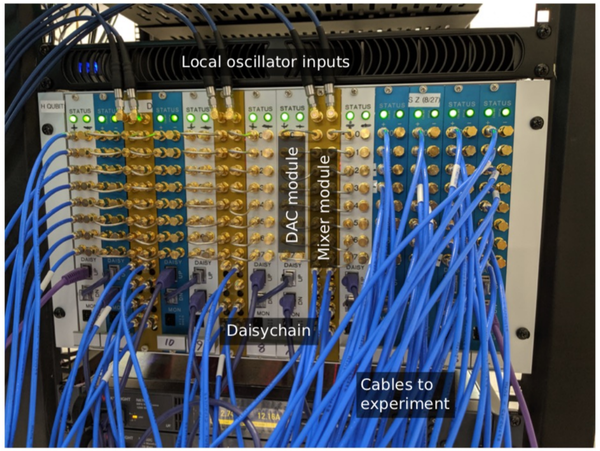

FIG. S4. **Electronics chassis.** Each chassis supports 14 DAC and/or mixer modules. Local oscillators are connected at the top of each mixer module. A set of daisychain cables connects from each ADC module to the next. Control signals exit the chassis through coaxial cables.

~~Quantum~~ (Analog ?)
Computers 2020

Some notes on Computational Paradigms

A look at the past to guess the future

Real-time, biologically plausible classification processes

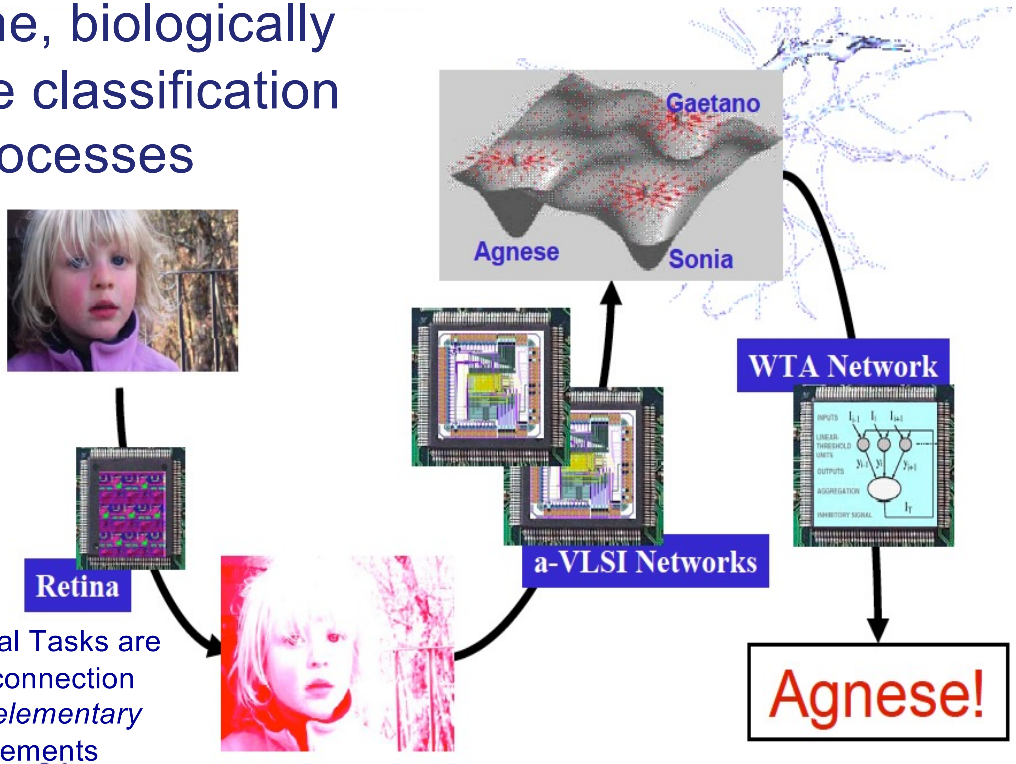

The Computational Tasks are defined by the connection topology of the *elementary* computing elements

Neuromorphic Hardware:

It encompasses any electronic device which **mimics** the natural biological structures of our nervous system.

D.J. Amit and G. Salina 1990-2006

- The goal is to impart cognitive abilities to machines by implementing neurons in silicon.
- Due to its superior energy efficiency and inherent parallelism, this approach is being considered as an alternative to conventional computational architectures

Computational Paradigms

Electronic implementation of an analog attractor neural network with stochastic learning

September 21, 1994

Davide Badoni, Stefano Bertazzoni, Stefano Buglioni, Gaetano Salina
INFN, Sezione di Roma II, Dipartimento di Fisica
Università di Roma "Tor Vergata", V.le della Ricerca Scientifica 1, 00133 Roma

Daniel J. Amit
INFN, Sezione di Roma, Dipartimento di Fisica
Università di Roma, La Sapienza, P.le Aldo Moro, Roma and
Racah Institute of Physics, Hebrew University, Jerusalem

Stefano Fusi
INFN, Sezione dell'Istituto Superiore di Sanità, Viale Regina Elena 299, Roma

Abstract

We describe and discuss an electronic implementation of an attractor neural network with plastic synapses. The network undergoes double dynamics, for the neurons as well as the synapses. Both dynamical processes are unsupervised. The synaptic dynamics is autonomous, in that it is driven exclusively and perpetually by neural activities. The latter follow the network activity via the developing synapses and the influence of external stimuli. Such a network self-organizes and is a device which converts the gross statistical characteristics of the stimulus input stream into a set of attractors (reverberations). To maintain for long time the acquired memory the analog synaptic efficacies are discretised by a stochastic refresh mechanism. The discretised synaptic memory has indefinitely long life time in the absence of activity in the network. It is modified only by the arrival of new stimuli. The stochastic refresh mechanism produces transitions at low probability which ensures that transient stimuli do not create significant modifications and that the system have large palimpsestic memory. A change in the attractor structure represents a major, macroscopic change in the statistics of the input stream, which may deform attractors, may create new ones and may eliminate others.

The electronic implementation is completely analog, stochastic and asynchronous. The circuitry of the first prototype is discussed in detail as well as the tests performed on it. In carrying out the implementation we have been guided by biological considerations and by electronic constraints. Both are discussed and new insights and lessons for the learning process are proposed.

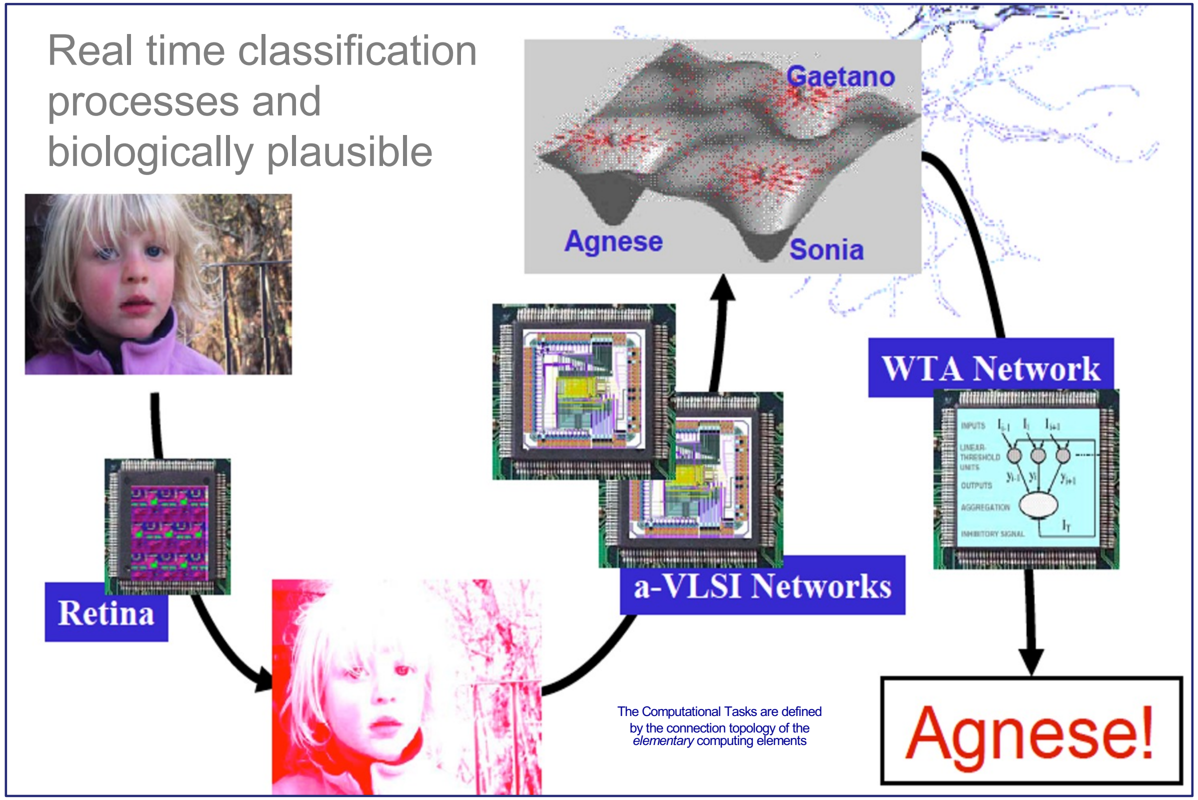

The electronic implementation is completely analogue, stochastic and asynchronous. The circuitry of the first prototype is discussed in detail as well as the tests performed on it. In carrying out the implementation we have been guided by biological considerations and by electronic constraints. Both are discussed and new insights and lessons for the learning process are proposed.

Some notes on Computational Paradigms

A look at the past to guess the future

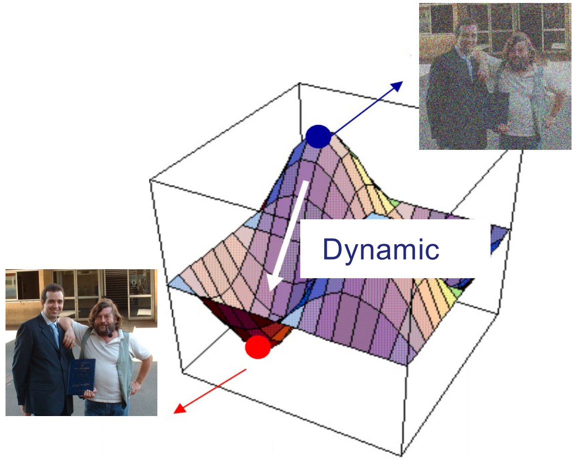

Memory addressable by content.

The recall of the information is triggered by a partial knowledge of the information

Dynamic System

Stocastic Neurons

$$H(t) = -\frac{1}{2}\sum_{ij} w_{ij} s_i(t) s_j(t)$$

$$\boxed{w_{ij} = \frac{1}{N}\sum_{\mu=1}^{p} \xi_i^{\mu}\xi_j^{\mu}}$$

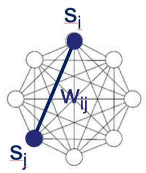

- Attractors (memorized patterns) are local minima of the Energy surface
- The Energy of the System decreases by virtue of its dynamics

Hopfield Model & Associative Memory

Some notes on Computational Paradigms

A look at the past to guess the future

Amorphous Computiong

"How does one **engineer** pre-specified, coherent behavior from the **cooperation** of immense numbers of unreliable parts that are interconnected in unknown, irregular, and time-varying ways?"

H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsky, T. Knight, R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss, *Amorphous Computing*, Communications of the ACM **43**, 74 (2000)

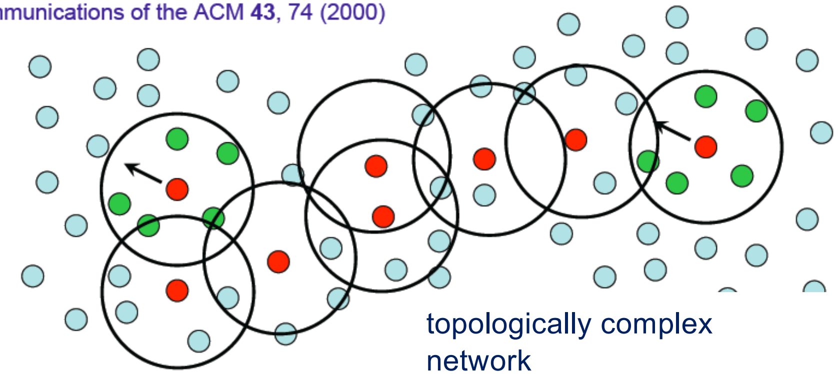

topologically complex network

Smart City as Amorphous System

Some notes on Computational Paradigms

A look at the past to guess the future

Computational paradigms and topologically Complex Systems

- Systems with a high number of degrees of freedom: **elementary processors** are connected by a **topologically complex network** that mediates their *interaction.*
- **Collective and naturally parallel** behaviour, little influenced by the detailed nature of the temporal evolution of the single processor and the connection network.
- Computational paradigms not defined a priori in algorithmic terms. System *code* is in its topological structure.

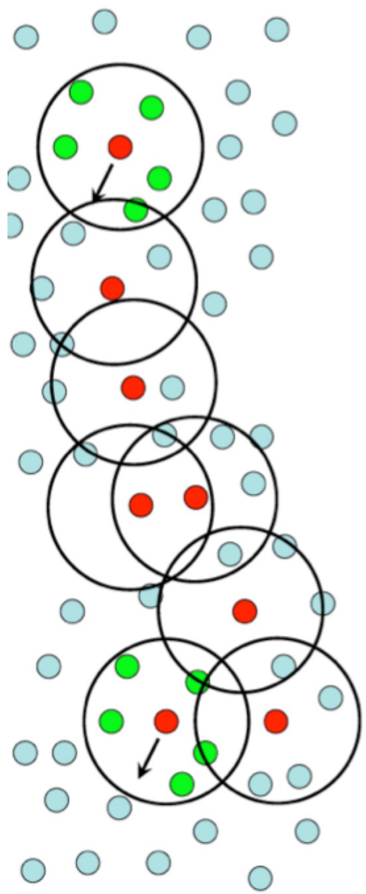

Computing Devices as Complex Systems?

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

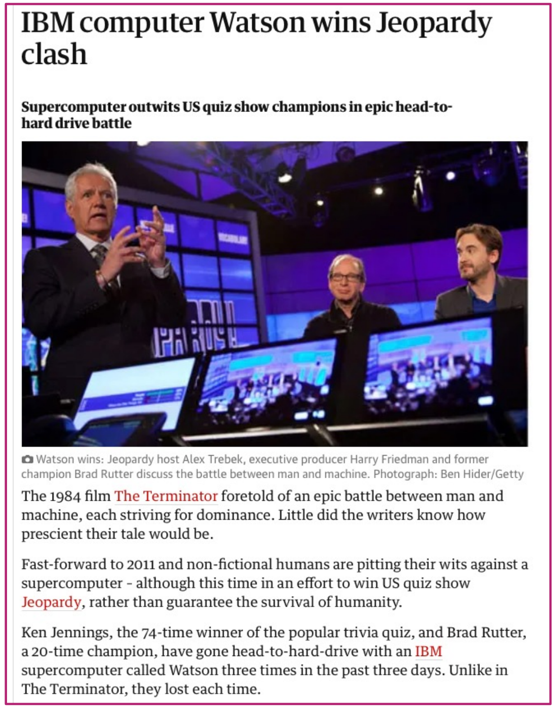

IBM computer Watson wins Jeopardy clash

Supercomputer outwits US quiz show champions in epic head-to-hard drive battle

Watson wins: Jeopardy host Alex Trebek, executive producer Harry Friedman and former champion Brad Rutter discuss the battle between man and machine. Photograph: Ben Hider/Getty

The 1984 film The Terminator foretold of an epic battle between man and machine, each striving for dominance. Little did the writers know how prescient their tale would be.

Fast-forward to 2011 and non-fictional humans are pitting their wits against a supercomputer - although this time in an effort to win US quiz show Jeopardy, rather than guarantee the survival of humanity.

Ken Jennings, the 74-time winner of the popular trivia quiz, and Brad Rutter, a 20-time champion, have gone head-to-hard-drive with an IBM supercomputer called Watson three times in the past three days. Unlike in The Terminator, they lost each time.

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

Precision Medicine:
Medical approach tailoring care to individual variability, beyond one-size-fits-all medicine.

Key components:
- Genomic and molecular data, clinical data (medical history, imaging, tests), lifestyle/environmental factors
- Integration across multiple data sources (data-driven approch)

Core objective:
- Improve risk assessment, treatment selection, and outcomes
- Reduce ineffective treatments and adverse effects

Conceptual foundation:
- Rooted in biomedical knowledge and clinical reasoning
- Made scientifically actionable by large-scale data analysis.

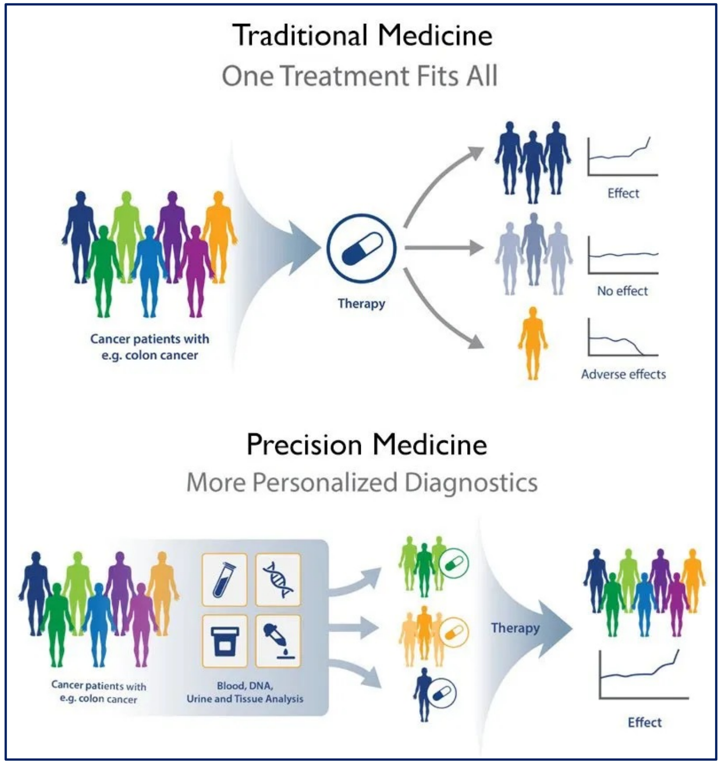

The Myth and the Reality

The Rise and Fall of IBM Watson’s AI Medical System

Precision Medicine and the Big Data Narrative

A data-driven turning point:

- Precision Medicine emerges with Big Data and AI, using genomics & clinical data
- Emphasis on data integration, pattern recognition, prediction
- Focus on integrating diverse datasets to identify patterns and predict outcomes

The hype component

- Data analysis assumed to fill understanding clinical gaps
- Expectation of quick results can cause over-optimism

Critical issue:

- Risk of relying on correlations over causation
- Modeling, clinical reasoning & biology may be downplayed

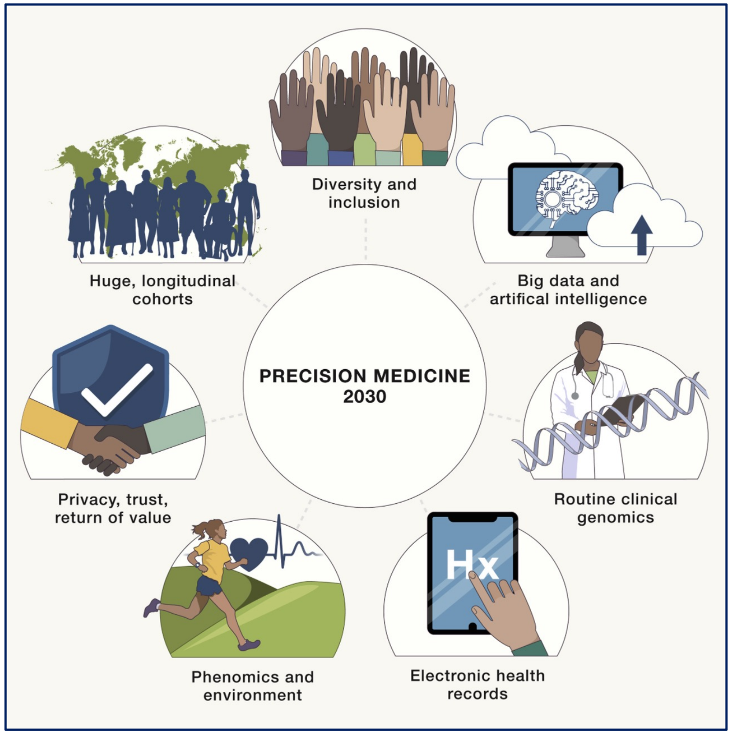

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

Why Watson Was Believed Capable in Oncology

AI credibility from Jeopardy success
Watson's performance on the quiz show showcased rapid, large-scale information processing.

Massive investments in data and healthcare tech
From 2017 IBM spent over 4 billion USD on companies and clinical datasets and electronic health Records (HER) to train Watson

Belief in AI + EHR integration
Access data was expected to give it deep real-world insight

Promise of rapid, personalized cancer guidance
Aimed to support faster, data-informed cancer treatment, with progress expected in a few years.

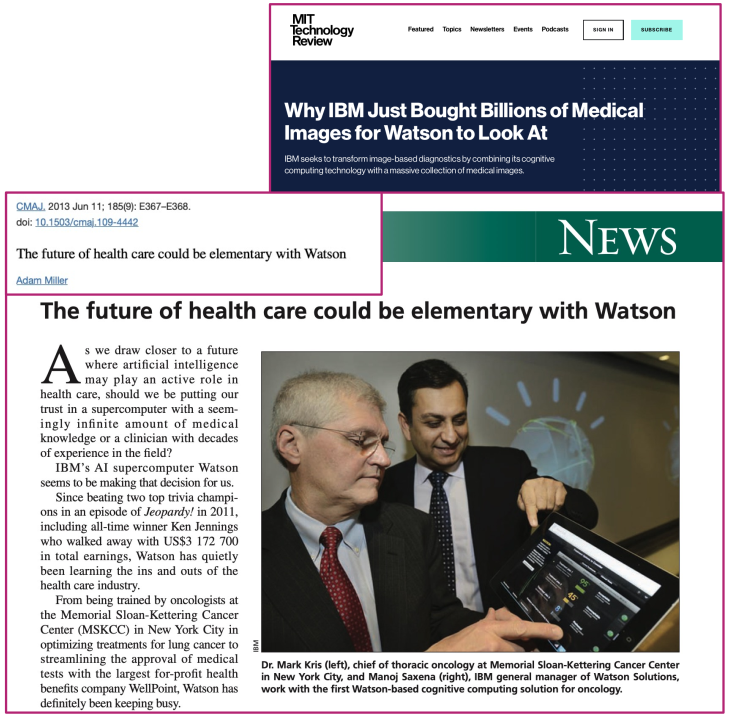
MIT Technology Review
Featured Topics Newsletters Events Podcasts SIGN IN SUBSCRIBE

Why IBM Just Bought Billions of Medical Images for Watson to Look At

IBM seeks to transform image-based diagnostics by combining its cognitive computing technology with a massive collection of medical images.

CMAJ. 2013 Jun 11; 185(9): E367–E368.
doi: 10.1503/cmaj.109-4442

The future of health care could be elementary with Watson

Adam Miller

NEWS

The future of health care could be elementary with Watson

As we draw closer to a future where artificial intelligence may play an active role in health care, should we be putting our trust in a supercomputer with a seemingly infinite amount of medical knowledge or a clinician with decades of experience in the field?

IBM's AI supercomputer Watson seems to be making that decision for us.

Since beating two top trivia champions in an episode of *Jeopardy!* in 2011, including all-time winner Ken Jennings who walked away with US$3 172 700 in total earnings, Watson has quietly been learning the ins and outs of the health care industry.

From being trained by oncologists at the Memorial Sloan-Kettering Cancer Center (MSKCC) in New York City in optimizing treatments for lung cancer to streamlining the approval of medical tests with the largest for-profit health benefits company WellPoint, Watson has definitely been keeping busy.

IBM

Dr. Mark Kris (left), chief of thoracic oncology at Memorial Sloan-Kettering Cancer Center in New York City, and Manoj Saxena (right), IBM general manager of Watson Solutions, work with the first Watson-based cognitive computing solution for oncology.

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

Scientific Method vs Market Logic

- Financial perspective: Business plans justify investment, show market potential, and promise ROI.
- Scientific/technical perspective: Often insufficient, no detailed architecture, algorithms, validation; promises may be unrealistic.
- Why scientists matter less early on: Scientific method requires verification, testing, and time; business requires convincing narratives, rapid timelines, and perceived scalability.
- Result: Early-stage AI projects often have big promises driven by finance, while technical feasibility lags behind.

Bending scientific rigor and technological feasibility to market logic leads straight to intellectual and technological suicide.

Vision and Critical View

Watson is presented as able to analyse large volumes of clinical and scientific data, emphasizing business potential over detailed technical or clinical validation

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

Reasons for Decline:

- Overhyped expectations: Public and media expectations exceeded Watson's real capabilities in clinical settings.
- Overselling economic benefits: Sales narratives emphasized rapid transformation and ROI, creating unrealistic expectations.
- Limited stakeholder engagement: Clinicians were insufficiently involved, leading to resistance and poor adoption.

CASE STUDY

Challenges in Commercial Deployment of AI: Insights from The Rise and Fall of IBM Watson's AI Medical System

Inspection Copy

Educator Material

By Quy Huy, Timo Vuori, Tero Ojanpera, Lisa Simone Duke

Published 14 Feb 2023

Reference 6753

Topic Strategy

Industry Hospital & Health Care

Region North America

Length 16 page(s)

Language English

Summary

When IBM set about commercializing its artificial intelligence-driven Watson AI in the healthcare market, its early successes were widely publicized. Senior managers and the media claimed that its diagnostic features would soon surpass those of the sharpest doctors. The case describes the large gap between what was promised and what happened in practice, offering insider insights on why IBM's projects failed. As the corporate commitment to AI escalated in response to successful lab results, cognitive dissonance arose between managers' expectations and what they could actually deliver. How could that have happened? Three reasons for Watson's downfall are explored: 1) The tendency for societal expectations to exceed the actual technical capabilities, leading to a gap in perception between AI in the lab and AI in the field. 2) Overselling of the economic benefits of AI by the salesforce; 3) Failure to secure the cooperation of key stakeholders, notably doctors who were asked to improve the performance of AI but were undermined by claims that AI could outperform them.

https://publishing.insead.edu/case/challenges-commercial-deployment-ai-insights-rise-and-fall-ibm-watsons-ai-medical-system

Key Issue: Underestimating real-world complexity. Executives relied on simplified lab tests and ideal scenarios, overestimating feasibility.

The Myth and the Reality

The Rise and Fall of IBM Watson’s AI Medical System

Lessons Learned:

- Align technical feasibility with expectations: Ensure AI capabilities match promises made. Hype can backfire
- Engage end-users early: Clinicians must be involved in design and testing to ensure adoption and trust.
- Balance business narrative and scientific rigor: Marketing and investment narratives cannot override technical validation and clinical evidence.
- Incremental deployment over “big bang” launch: Start with pilots, validate outcomes, then scale. Avoid rushing adoption for financial reasons.
- Transparent communication: Clearly communicate AI limitations and confidence levels; manage expectations

CASE STUDY

Challenges in Commercial Deployment of AI: Insights from The Rise and Fall of IBM Watson’s AI Medical System

Inspection Copy

Educator Material

By Quy Huy, Timo Vuori, Tero Ojanpera, Lisa Simone Duke

Published 14 Feb 2023

Reference 6753

Topic Strategy

Industry Hospital & Health Care

Region North America

Length 16 page(s)

Language English

Summary

When IBM set about commercializing its artificial intelligence-driven Watson AI in the healthcare market, its early successes were widely publicized. Senior managers and the media claimed that its diagnostic features would soon surpass those of the sharpest doctors. The case describes the large gap between what was promised and what happened in practice, offering insider insights on why IBM’s projects failed. As the corporate commitment to AI escalated in response to successful lab results, cognitive dissonance arose between managers’ expectations and what they could actually deliver. How could that have happened? Three reasons for Watson’s downfall are explored: 1) The tendency for societal expectations to exceed the actual technical capabilities, leading to a gap in perception between AI in the lab and AI in the field. 2) Overselling of the economic benefits of AI by the salesforce; 3) Failure to secure the cooperation of key stakeholders, notably doctors who were asked to improve the performance of AI but were undermined by claims that AI could outperform them.

https://publishing.insead.edu/case/challenges-commercial-deployment-ai-insights-rise-and-fall-ibm-watsons-ai-medical-system

The Myth and the Reality

The Rise and Fall of IBM Watson's AI Medical System

What Ever Happened to IBM's Watson?

- Ambitious general-purpose vision: Watson was promoted as capable of transforming healthcare, finance, law, and more.
- Reality check in healthcare: Integration with real-world clinical data and workflows proved extremely difficult; adoption was limited.
- Financial strain: Massive investments in data, partnerships, and infrastructure did not translate into expected revenue; profitability pressures emerged.
- Shift to special-purpose tools: IBM repositioned Watson as a suite of AI tools for specific tasks (automation, analytics, support), moving away from broad "moonshot" ambitions.
- Lesson learned: Success in lab or quiz shows (e.g., Jeopardy) does not translate directly to complex real-world applications.

What Ever Happened to IBM's Watson?

IBM's artificial intelligence was supposed to transform industries and generate riches for the company. Neither has panned out. Now, IBM has settled on a humbler vision for Watson.

https://www.nytimes.com/2021/07/16/technology/what-happened-ibm-watson.html

Three suggested books as Conclusion

To improve the understanding of some topics covered

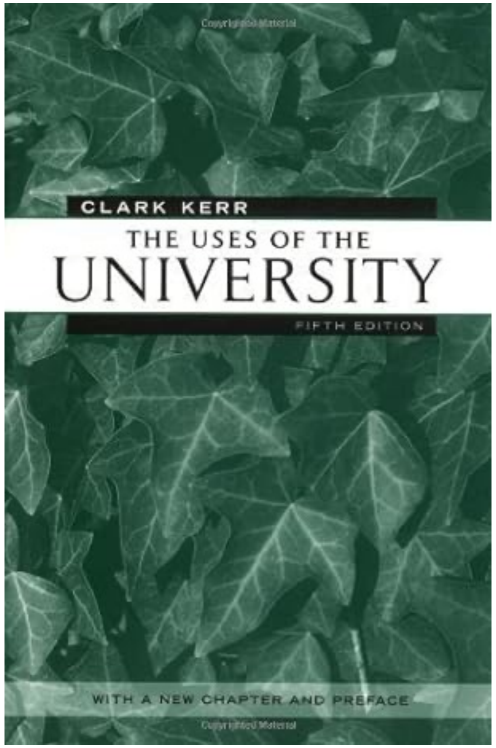

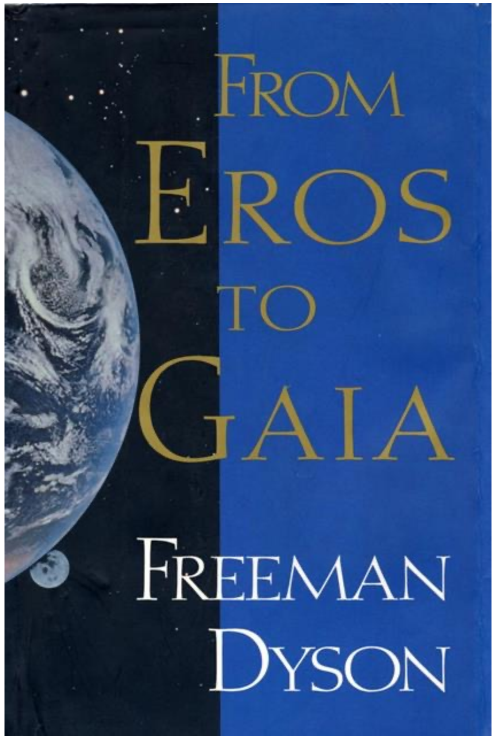

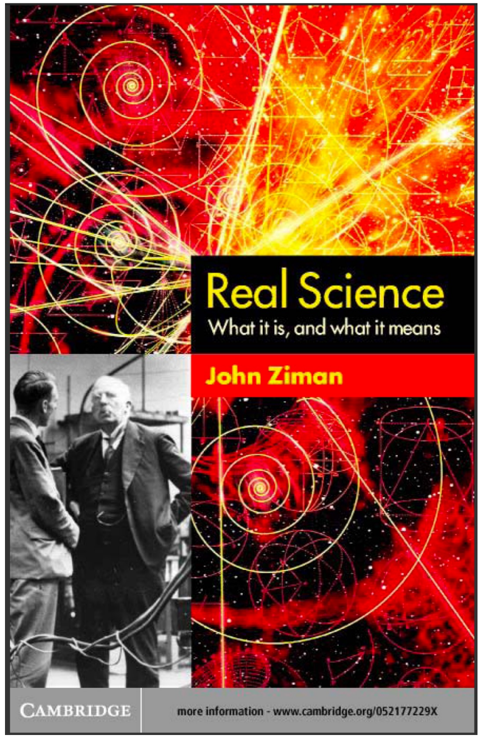

From the University to the Multiversity

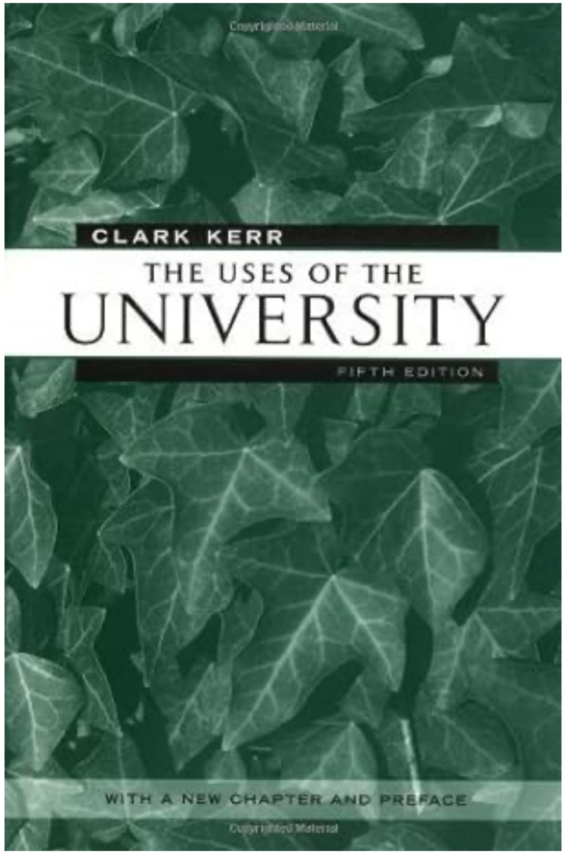

Clark Kerr, The uses of the University, Berkley 1963

The university was no longer cloistered, but now central to society, a prime instrument of national purpose. Its job was to produce *new knowledge.*

New knowledge is the most important factor in *economic and social growth*.

What the railroads did for the second half of the last century, and the automobile for the first half of this century, may be done for the second half of this century by the knowledge industry: that is, to serve as the focal point for national growth.

The university and segments of industry are becoming more and more alike.

As the university becomes tied into the world of work, the *professor* - at least in the natural and some of the social sciences - *takes on the characteristics of an entrepreneur...The two worlds are merging, physically and psychologically*...

The campus and society are undergoing a somewhat reluctant and cautious merger, already well advanced.

Free Speech Movement
(Berkeley, October 1964)

https://alchetron.com/Mario-Savio

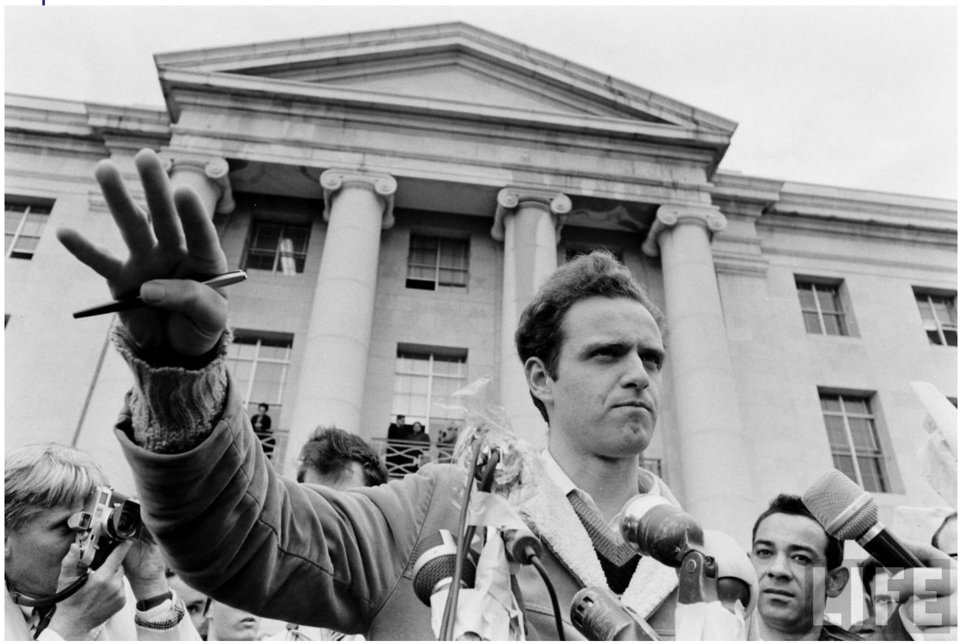

Mario Savio: Sproul Hall Steps, Dec 2, 1964

… and if President Kerr in fact is the manager, then I tell you something: the faculty are a bunch of employees!

And we're the raw material!

But we're a bunch of raw materials that don't mean to have any process upon us, don't mean to be made into any product, don't mean to end up being bought by some clients of the University, be they the government, be they industry, be they organized labor, be they anyone!

We're human beings!

The 1963 C. Kerr Godkin lectures

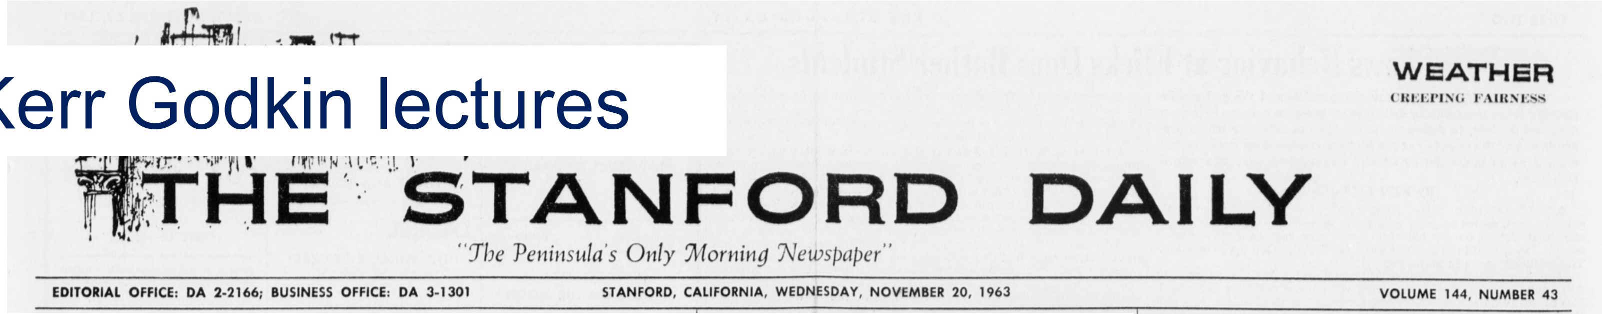
WEATHER
CREEPING FAIRNESS
THE STANFORD DAILY
"The Peninsula's Only Morning Newspaper"
EDITORIAL OFFICE: DA 2-2166; BUSINESS OFFICE: DA 3-1301
STANFORD, CALIFORNIA, WEDNESDAY, NOVEMBER 20, 1963
VOLUME 144, NUMBER 43

Lecture 3 – Key Themes

- Two key shifts in HE: Universal access and Progress through science
- The Multiversity as driver of economic growth and a comprehensive problem-solver for society
- The rise of the *knowledge industry* as central to business and government
- Universities increasingly resembling industrial organizations
- Concentration of resources in units producing *useful* knowledge, growing stratification in higher education (HE)
- Intense inter-university rivalry

Although Clark Kerr did not use the term *academic capitalism*, he clearly described many of its emerging features

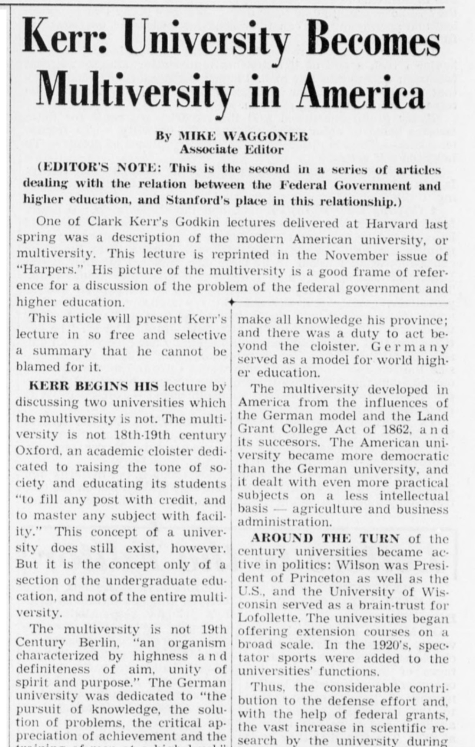
Kerr: University Becomes Multiversity in America

By MIKE WAGGONER
Associate Editor

(EDITOR'S NOTE: This is the second in a series of articles dealing with the relation between the Federal Government and higher education, and Stanford's place in this relationship.)

One of Clark Kerr's Godkin lectures delivered at Harvard last spring was a description of the modern American university, or multiversity. This lecture is reprinted in the November issue of "Harpers." His picture of the multiversity is a good frame of reference for a discussion of the problem of the federal government and higher education.

This article will present Kerr's lecture in so free and selective a summary that he cannot be blamed for it.

KERR BEGINS HIS lecture by discussing two universities which the multiversity is not. The multiversity is not 18th-19th century Oxford, an academic cloister dedicated to raising the tone of society and educating its students "to fill any post with credit, and to master any subject with facility." This concept of a university does still exist, however. But it is the concept only of a section of the undergraduate education, and not of the entire multiversity.

The multiversity is not 19th Century Berlin, "an organism characterized by highness and definiteness of aim, unity of spirit and purpose." The German university was dedicated to "the pursuit of knowledge, the solution of problems, the critical appreciation of achievement and the

make all knowledge his province; and there was a duty to act beyond the cloister. Germany served as a model for world higher education.

The multiversity developed in America from the influences of the German model and the Land Grant College Act of 1862, and its succesors. The American university became more democratic than the German university, and it dealt with even more practical subjects on a less intellectual basis — agriculture and business administration.

AROUND THE TURN of the century universities became active in politics: Wilson was President of Princeton as well as the U.S., and the University of Wisconsin served as a brain-trust for Lofollette. The universities began offering extension courses on a broad scale. In the 1920's, spectator sports were added to the universities' functions.

Thus, the considerable contribution to the defense effort and, with the help of federal grants, the vast increase in scientific research by the university during

Globalisation, Societies and Education, 2014
http://dx.doi.org/10.1080/14767724.2013.876313

Routledge

The European Round Table of Industrialists and the restructuring of European higher education

Ilkka Kauppinen[a,b]*

[a]*Department of Social Sciences and Philosophy, University of Jyväskylä, Jyväskylä, Finland;* [b]*Institute of Higher Education, University of Georgia, Athens, GA, USA*

(*Received 10 July 2012; accepted 6 December 2013*)

The restructuring of European higher education (EHE) since the 1980s is a widely studied subject. However, this paper argues that previous studies have paid insufficient attention to the role of transnational policy-making groups in this complex and multilevel process. This argument is supported by focusing on how the European Round Table of Industrialists (ERT) has participated in this restructuring since the mid-1980s. This paper's focus is especially in two ERT documents that were published in the 1980s. The main finding is that the current restructuring of EHE reflects interests of the ERT that represents the emerging transnational capitalist class (TCC) at European level.

Keywords: the European Round Table of Industrialists; European higher education; transnational capitalist class

Introduction

It is widely recognised that transnationalisation of higher education is no longer a marginal topic for universities, but rather one of the most central issues of strategic importance even if universities' activities are still strongly guided by national regulatory and funding frameworks (e.g., Beerkens and van der Wende 2007). Despite this continuing role of national governments in steering higher education, nation-state-centric approaches are becoming increasingly problematic, for instance, in European context, given the development of both the European Higher Education Area (EHEA) and European Research Area (ERA), and how they are linked to wider economic strategies at European level (e.g., Lisbon strategy and Europe 2020 growth strategy).

This paper is based on the assumption that previous studies on the restructuring of European higher education (EHE) have not taken sufficiently into consideration the role of transnationally oriented economic actors in this multilevel process. Instead, a lot of attention has been paid to the role of the European Union (EU), the Organization for Economic Co-operation and Development (OECD) and individual member states and how they have

*Corresponding author. Email: ilkka.j.kauppinen@jyu.fi

This paper is based on the assumption that previous studies on the restructuring of European higher education (EHE) have not taken sufficiently into consideration the role of transnationally oriented economic actors in this multilevel process.

Instead, a lot of attention has been paid to the role of the European Union (EU), the Organization for Economic Co-operation and Development (OECD) and individual member states.

This is a suitable starting point to explore how the European Round Table of Industrialists (ERT) has, in the context of globalising knowledge capitalism, strategically participated in the restructuring of EHE since the 1980s.

A University is a place whither students come from every quarter for every kind of knowledge; [ ] a place for the communication and circulation of thought, by means of personal intercourse. [ ] It is the place to which a thousand schools make contributions; in which the intellect may safely range and speculate. It is a place where inquiry is pushed forward, [ ] discoveries verified and perfected, and [ ] error exposed, by the collision of mind with mind, and knowledge with knowledge. Education is one of the great and incessant occupations of human society.

J. H. Newman in *The Idea of a University,* 1852

The university was no longer cloistered, but now central to society, a prime instrument of national purpose. Its job was to produce new knowledge.
New knowledge is the most important factor in economic and social growth.
What the railroads did for the second half of the last century, and the automobile for the first half of this century, may be done for the second half of this century by the knowledge industry: that is, to serve as the focal point for national growth.

C. Kerr in *The Uses of the University,* 1963

Big and Small Science competition:

The large and fashionable squeezes out the small and unfashionable

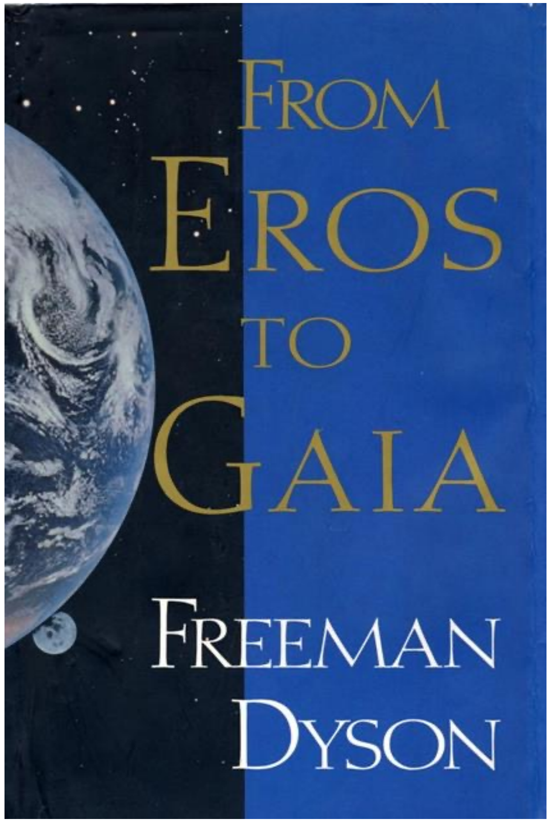

Freeman Dyson, From Eros to Gaia, 1992

Cap 3. Six Cautionary Tales for Scientists

To represent the third world. I choose the village of Ngon, a village in Central Africa. The main problem in Ngon is water. [ ]

To represent the second world. I choose the great Soviet astronomical observatory (six-meter telescope) at Zelenchukskaya in the Caucasus Mountains. [ ]

“My third cautionary tale concerns our own world, the so-called first world. [ ] I shall talk only about the problems of space-based astronomy. [ ]

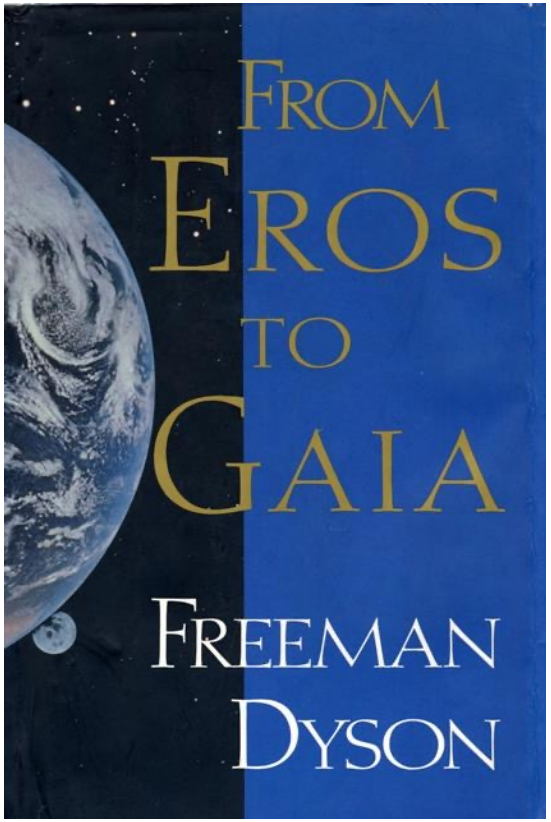

The game of status seeking, organized around committees, *is played in roughly the same fashion* in Africa and in America and in the Soviet Union.

The game has had profound consequences for science. In science, as in the quest for a village water supply, *big projects bring enhanced status; small projects do not.*

In the competition for status, big projects usually win, whether or not they are scientifically justified. As the committees of academic professionals compete for power and influence, big science becomes more and more preponderant over small science.

The large and fashionable squeezes out the small and unfashionable.

From CUDOS to PLACE:
The post-academic research

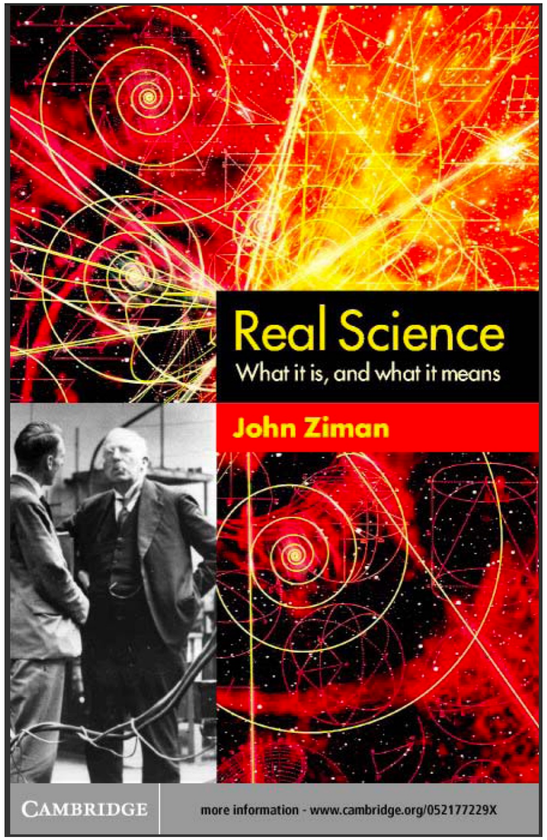

John Ziman, Real Science, 2000

The Mertonian norms CUDOS:

Robert Merton, The normative structure of science, 1942

Communalism, that scientific results should be shared as widely and quickly as possible;

Universalism, that science is independent of the personal or cultural status of the scientist;

Disinterestedness, that scientific results should be free from personal or corporate biases and dishonesty;

Originality, that scientific results should contribute something new;

Scepticism, that scientific results must be able to withstand systematic doubt.

John Ziman contrasts Merton's CUDOS with its technological pendant, the acronym PLACE:

John Ziman, Real Science, 2001

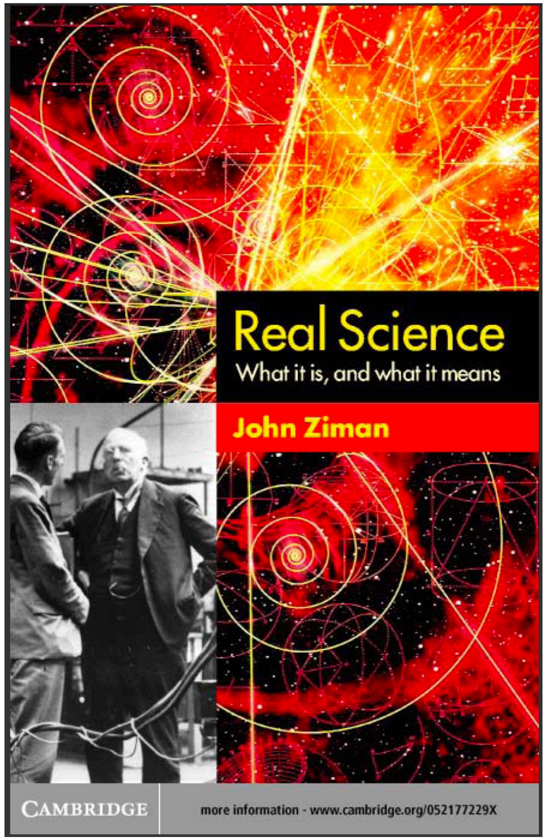

The results produced are *Proprietary*, and therefore not necessarily communal;

Researchers concentrate on *Local*, technical problems which may not contribute to general understanding;

Authority is vested in a managerial hierarchy, not in the individual researcher;

Work is Commissioned to solve specific problems, not as a contribution to knowledge as a whole;

The scientist is valued as an Expert rather than a source of creativity

Three Critical Conclusions

These three books describe different moments, but they converge on the same tension: science is more powerful, more visible, more useful and at the same time less free.

1. University → Multiversity

Market logic increasingly shapes academic life

→ *Scientific norms subordinated to economic priorities*

2. Big vs Small Science

Scale and visibility dominate competition

→ *Small, creative science is structurally marginalized*

3. Post-Academic Research

Accountability replaces autonomy

→ *Research agendas are set outside the scientific community*

“There are only engineers — ‘researcher’ is a relic term from academia.” — Elon Musk

→ **The ultimate challenge** to the traditional ideal of free academic research

AI Generated Slide

Scientific Method vs AI & Big Data

AI and Big Data offer major opportunities as analytical tools where conceptual frameworks are weak or rapid, application-driven results are required.

- *Dominant framework*

Market logics dominate post-academic research
Narratives and language framing (AI hype, Big Data) suggest market priorities

- *Architectural convergence*

Market-driven AI favors convergent solutions, limiting architectural diversity
Need for new architectures, new algorithms, and algorithmic optimization

- *Scientific autonomy*

Market competition imposes similar architectures and shaping both technology and public funding
Only autonomous academic research enables alternative and fundamental innovation.

AI innovation risks convergence rather than progress without scientific autonomy, and this can only be avoided by preserving rigorous scientific methodology.

The Power of AI Hype: Speculative Risks, Real Politics

https://innovitalia.esteri.it/notizia/artificial-intelligence-safety-summit-bletchley-park-12-nov-2023

At the global AI summit in Bletchley Park, representatives from 28 countries, debated alleged existential risks posed by advanced artificial intelligence.

Vision of AGI (OpenAI)

Artificial General Intelligence is portrayed as a technology capable of elevating humanity — expanding economic abundance, accelerating scientific discovery, and providing universal cognitive tools that amplify human creativity and intelligence

Critical Perspective

When examined through a rigorous scientific lens, the catastrophic scenarios invoked in discussions of AGI remain closer to science fiction literature than to empirically grounded research

Closing Statement

The real danger does not lie in imaginary superintelligent machines, but in the political and economic ideologies underpinning these a-scientific approaches, and in their influence over the allocation of public research funding

https://www.nytimes.com/2023/06/30/opinion/artificial-intelligence-danger.html

Il pericolo più grande dell'umanità è qualcosa che non esiste. Si chiama Agi: Intelligenza artificiale generale. Si tratta dell'ipotesi che l'intelligenza artificiale possa diventare super intelligente. Capace non solo di replicare quello che fanno gli umani, ma di capire, imparare e svolgere tutti i compiti che può svolgere un essere umano. E farlo meglio. Tanto meglio che forte della propria superiorità e dell'autocoscienza acquisita, potrebbe decidere di cancellarci dalla faccia della Terra qualora non ci ritenesse più necessari. L'Agi da un momento all'altro potrebbe svegliars, come Cthulhu nei romanzi di H.P. Lovecraft, non appena gli algoritmi che la animano avranno trovato la giusta formula per farlo. Ma l'Agi oggi è poco più di uno scenario possibile. Non si sa quando, né se un giorno ci sarà.

Eppure è per paura del suo avvento che qualche mese fa 350 tra imprenditori e accademici hanno firmato una lettera per chiedere una moratoria di sei mesi allo sviluppo di intelligenze artificiali (tra loro Elon Musk, capo di Tesla, e Steve Wozniak, co-fondatore di Apple e studiosi di primo piano di apprendimento delle macchine). Paura condivisa da Sam Altman, numero uno di OpenAi, la società che ha creato ChatGpt, che lo scorso mese ha fatto il giro del mondo per stringere le mani ai capi di stato e illustrare sia i rischi possibili, ma anche le opportunità concrete dell'intelligenza artificiale. Le due cose probabilmente si tengono.

https://openai.com/blog/planning-for-agi-and-beyond

La nostra missione è garantire che l'intelligenza artificiale generale, ovvero sistemi di intelligenza artificiale generalmente più intelligenti degli esseri umani, porti benefici a tutta l'umanità.

In caso di successo, questa tecnologia potrebbe aiutarci a migliorare l'umanità aumentando la ricchezza, potenziando l'economia globale e contribuendo alla scoperta di nuove conoscenze scientifiche.
L'AGI ha il potenziale di fornire a tutti incredibili nuove capacità; possiamo immaginare un mondo in cui tutti abbiamo accesso a strumenti utili per qualsiasi ogni compito cognitivo, incrementando il nostro ingegno umano e la nostra creatività.

D'altro canto, l'AGI comporterebbe anche seri rischi di abuso, incidenti drastici e perturbazioni sociali. Poiché i vantaggi dell'AGI sono così grandi riteniamo auspicabile che la società la società e gli sviluppatori dell'AGI capiscano come farlo nel modo giusto

The Gentle Vultures

Isaac Asimov 1958

"But if they have discovered atomic energy, where do they conduct their tests, their explosions?"

"On their own planet, sir."

Naron straightened to his full six meters in height and thundered, *"On their own planet?"*

"Yes, sir."

Slowly, Naron took up his pen and drew a line through the latest entry in the little book. It was an unprecedented act, but Naron was very, very wise, and could see the inevitable more clearly than anyone else in the galaxies.
"Idiots!" he muttered.

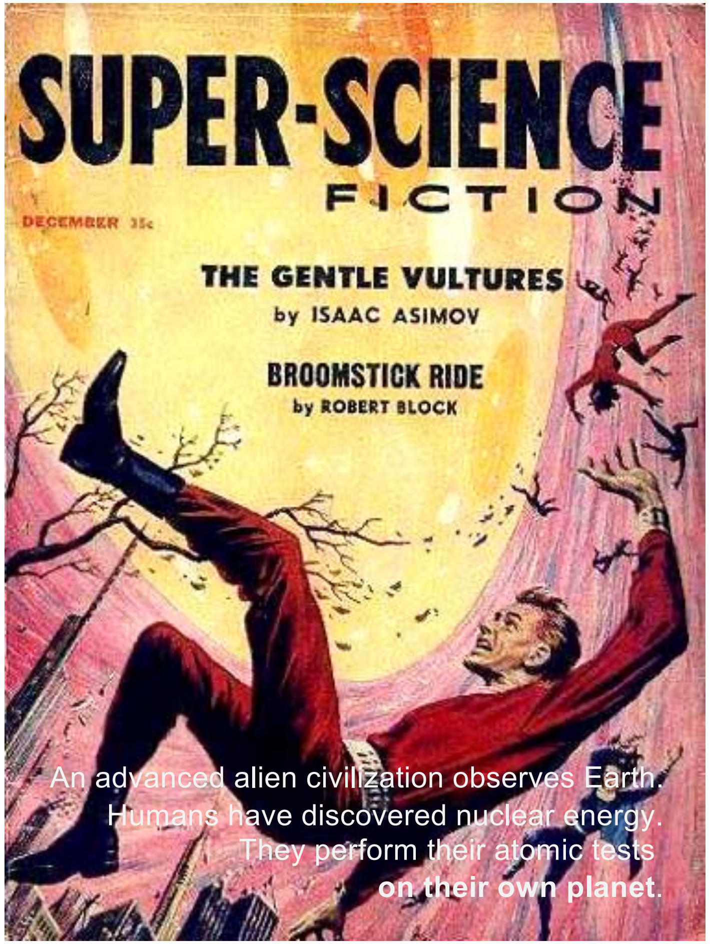

Suppose we have an arbitrary partition a.

Let C be the subset with the smaller sum.

$$\left(\sum_{i}^{n_1} x_{i,a}^{B}\right) > \left(\sum_{i}^{n_2} x_{i,a}^{C}\right)$$

If C contains more than one element, take any element $x_i \in C$ and move it to B.

The new difference becomes:

$$\left(\sum_{i}^{n_1} x_{i,b}^{B}\right) - \left(\sum_{i}^{n_2} x_{i,b}^{C}\right) =$$

$$\left(\sum_{i}^{n_1} x_{i,a}^{B} + x_i\right) - \left(\left(\sum_{i}^{n_2} x_{i,a}^{C}\right) - x_i\right) = \left(\sum_{i}^{n_1} x_{i,a}^{B}\right) - \left(\sum_{i}^{n_2} x_{i,a}^{C}\right) + 2x_i$$

Since $x_i > 0$, the difference always increases.

Repeating this process, the lighter subset is reduced to a single element in the optimal partition. QED.