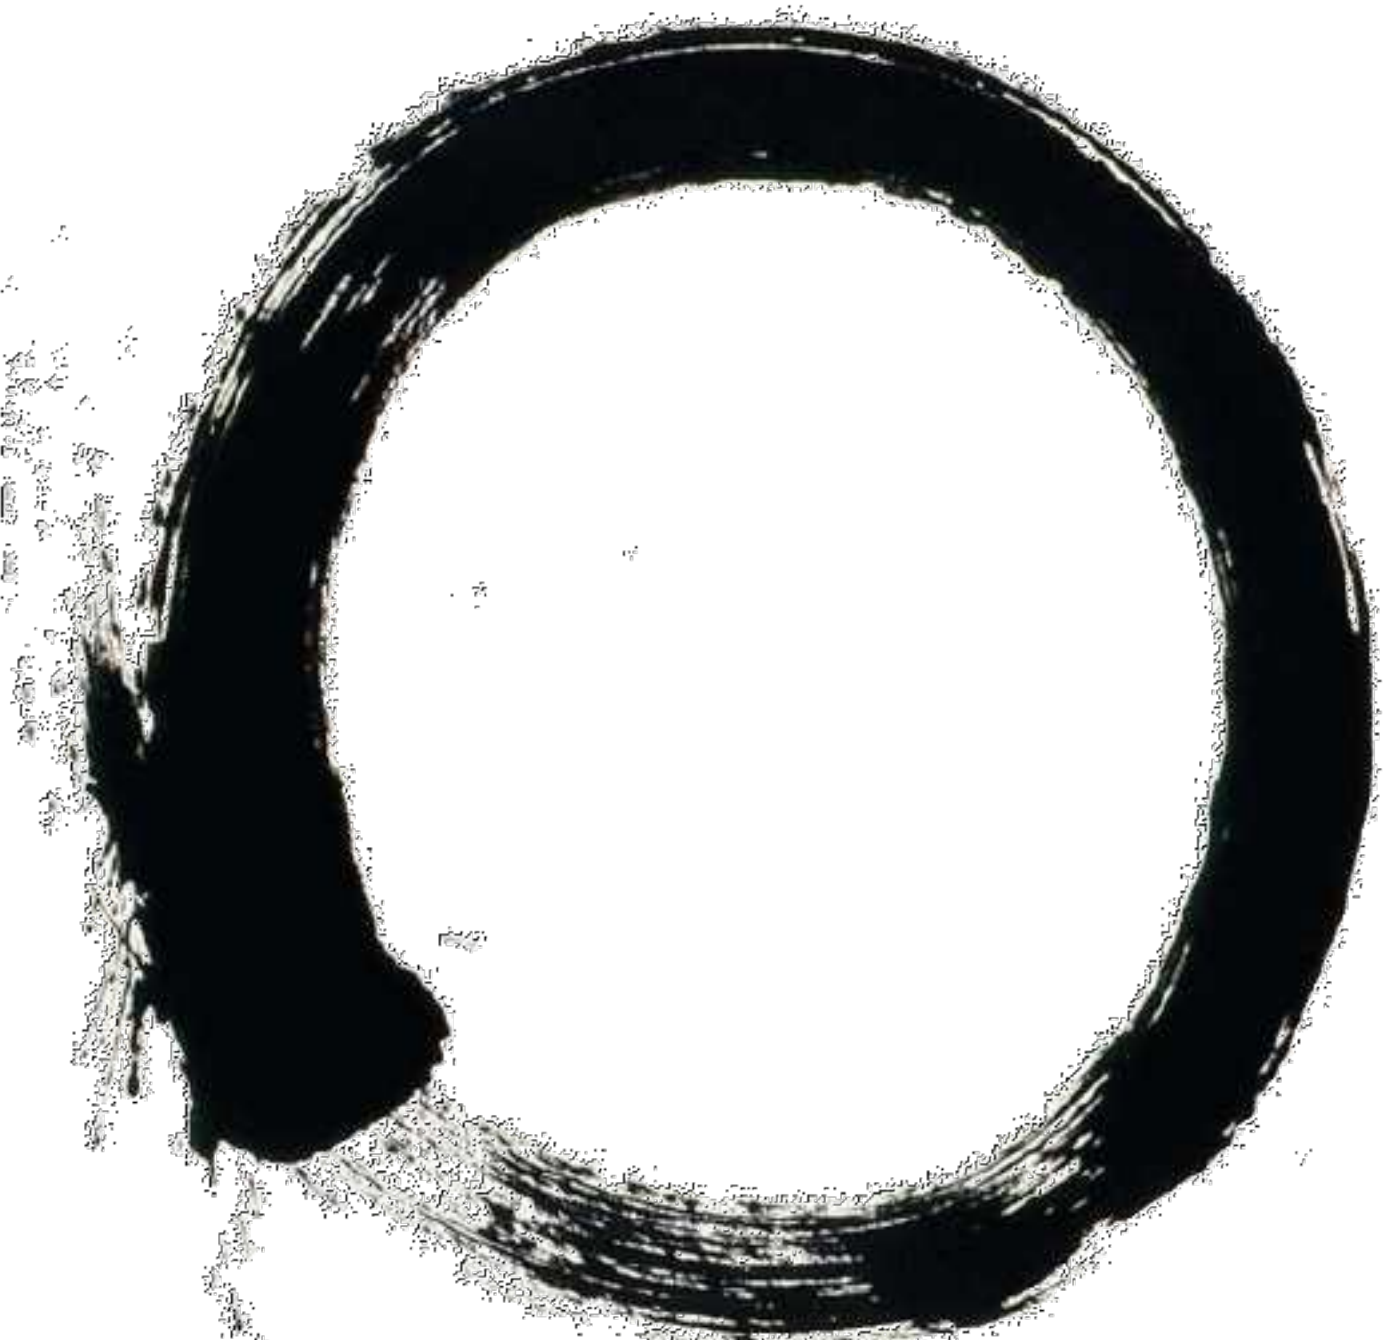


Jan W. Veldsink



平常之道

THE ART OF THE ARTIFICIAL

BALANCING DATA, TECHNOLOGY, AND HUMANITY

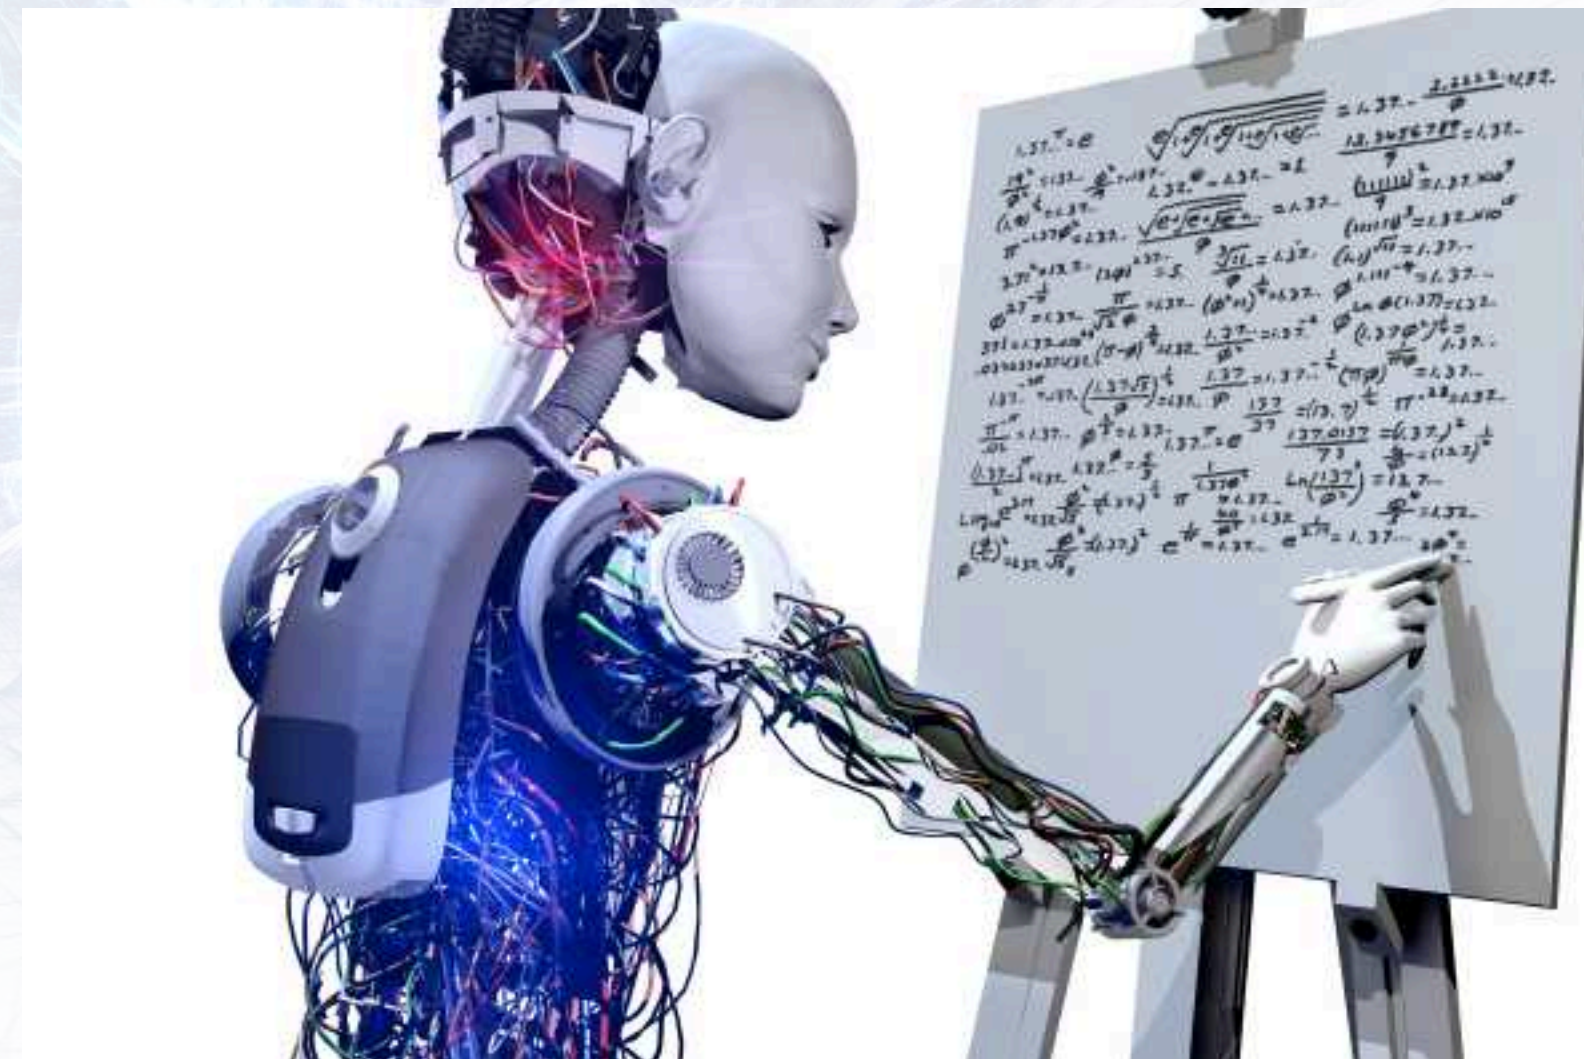
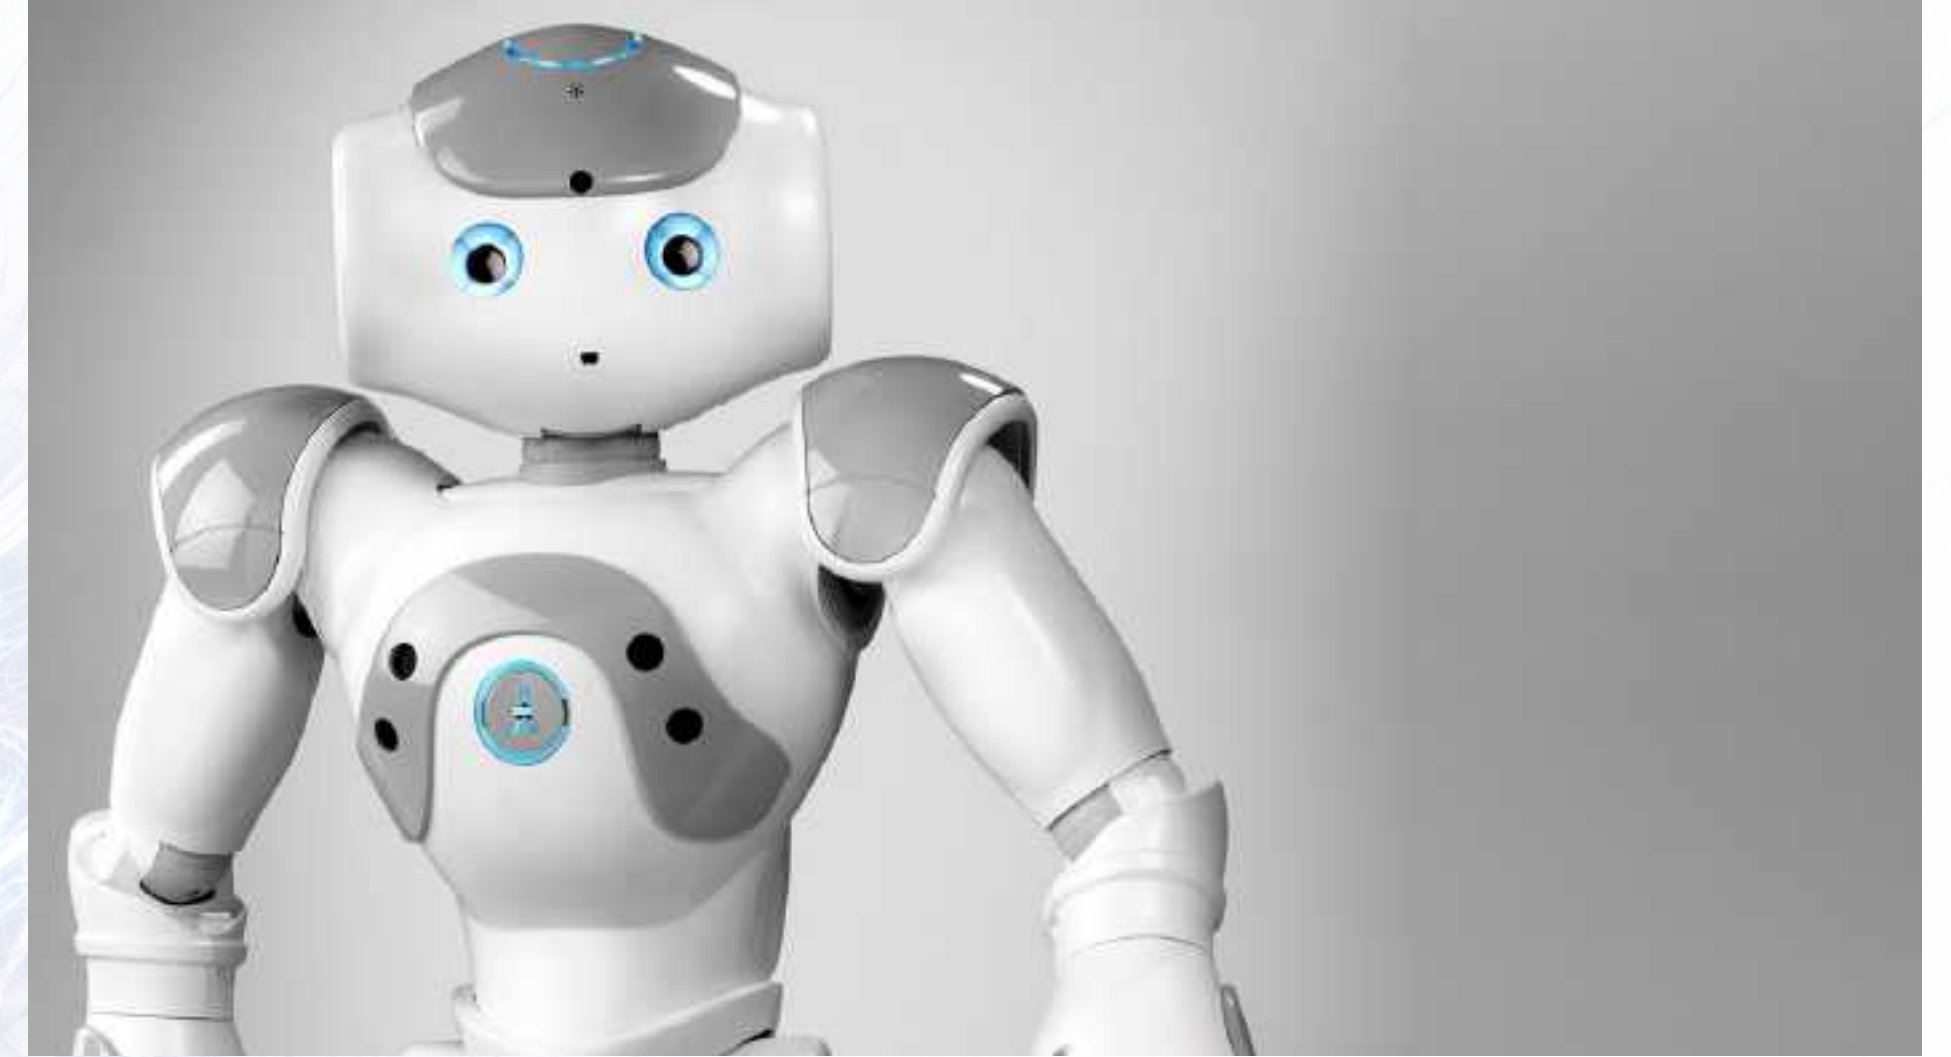
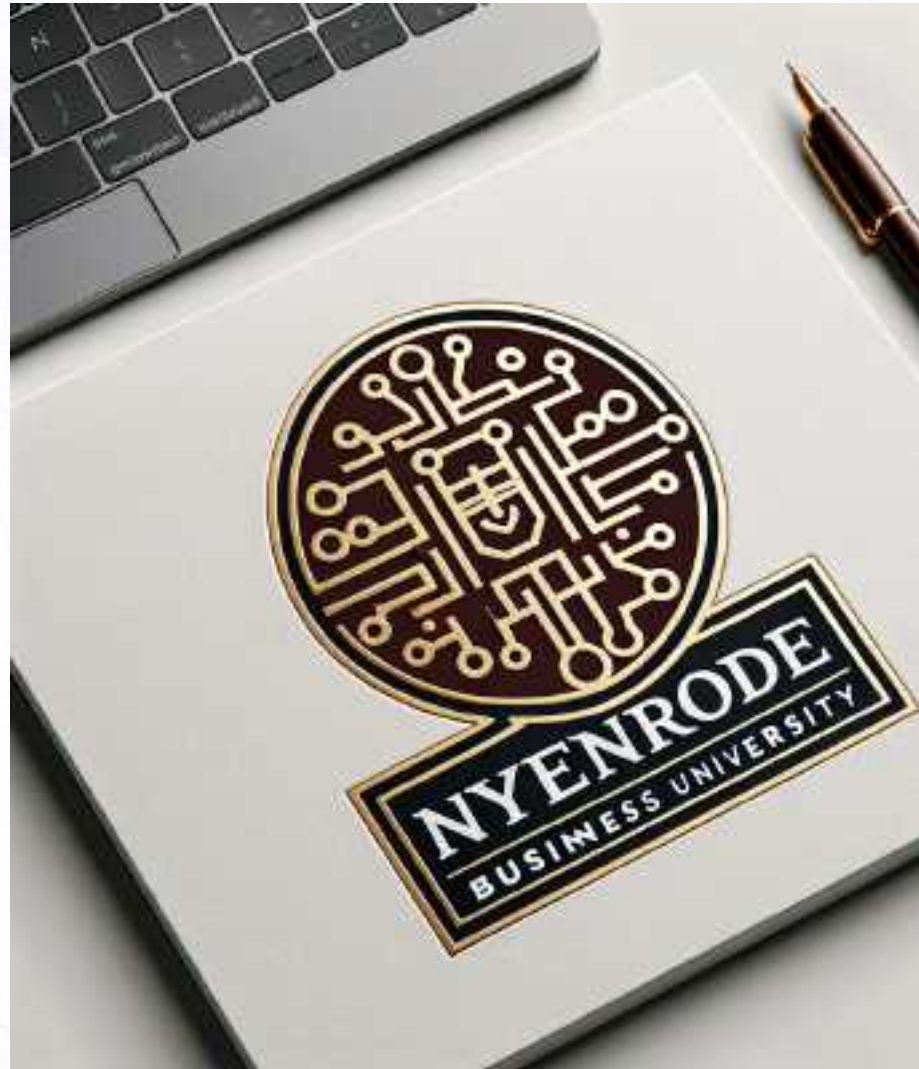
JAN W VELDSINK MSC

(MSC IN ARTIFICIAL INTELLIGENCE SINCE 2000)



Jan W. Veldsink MSC



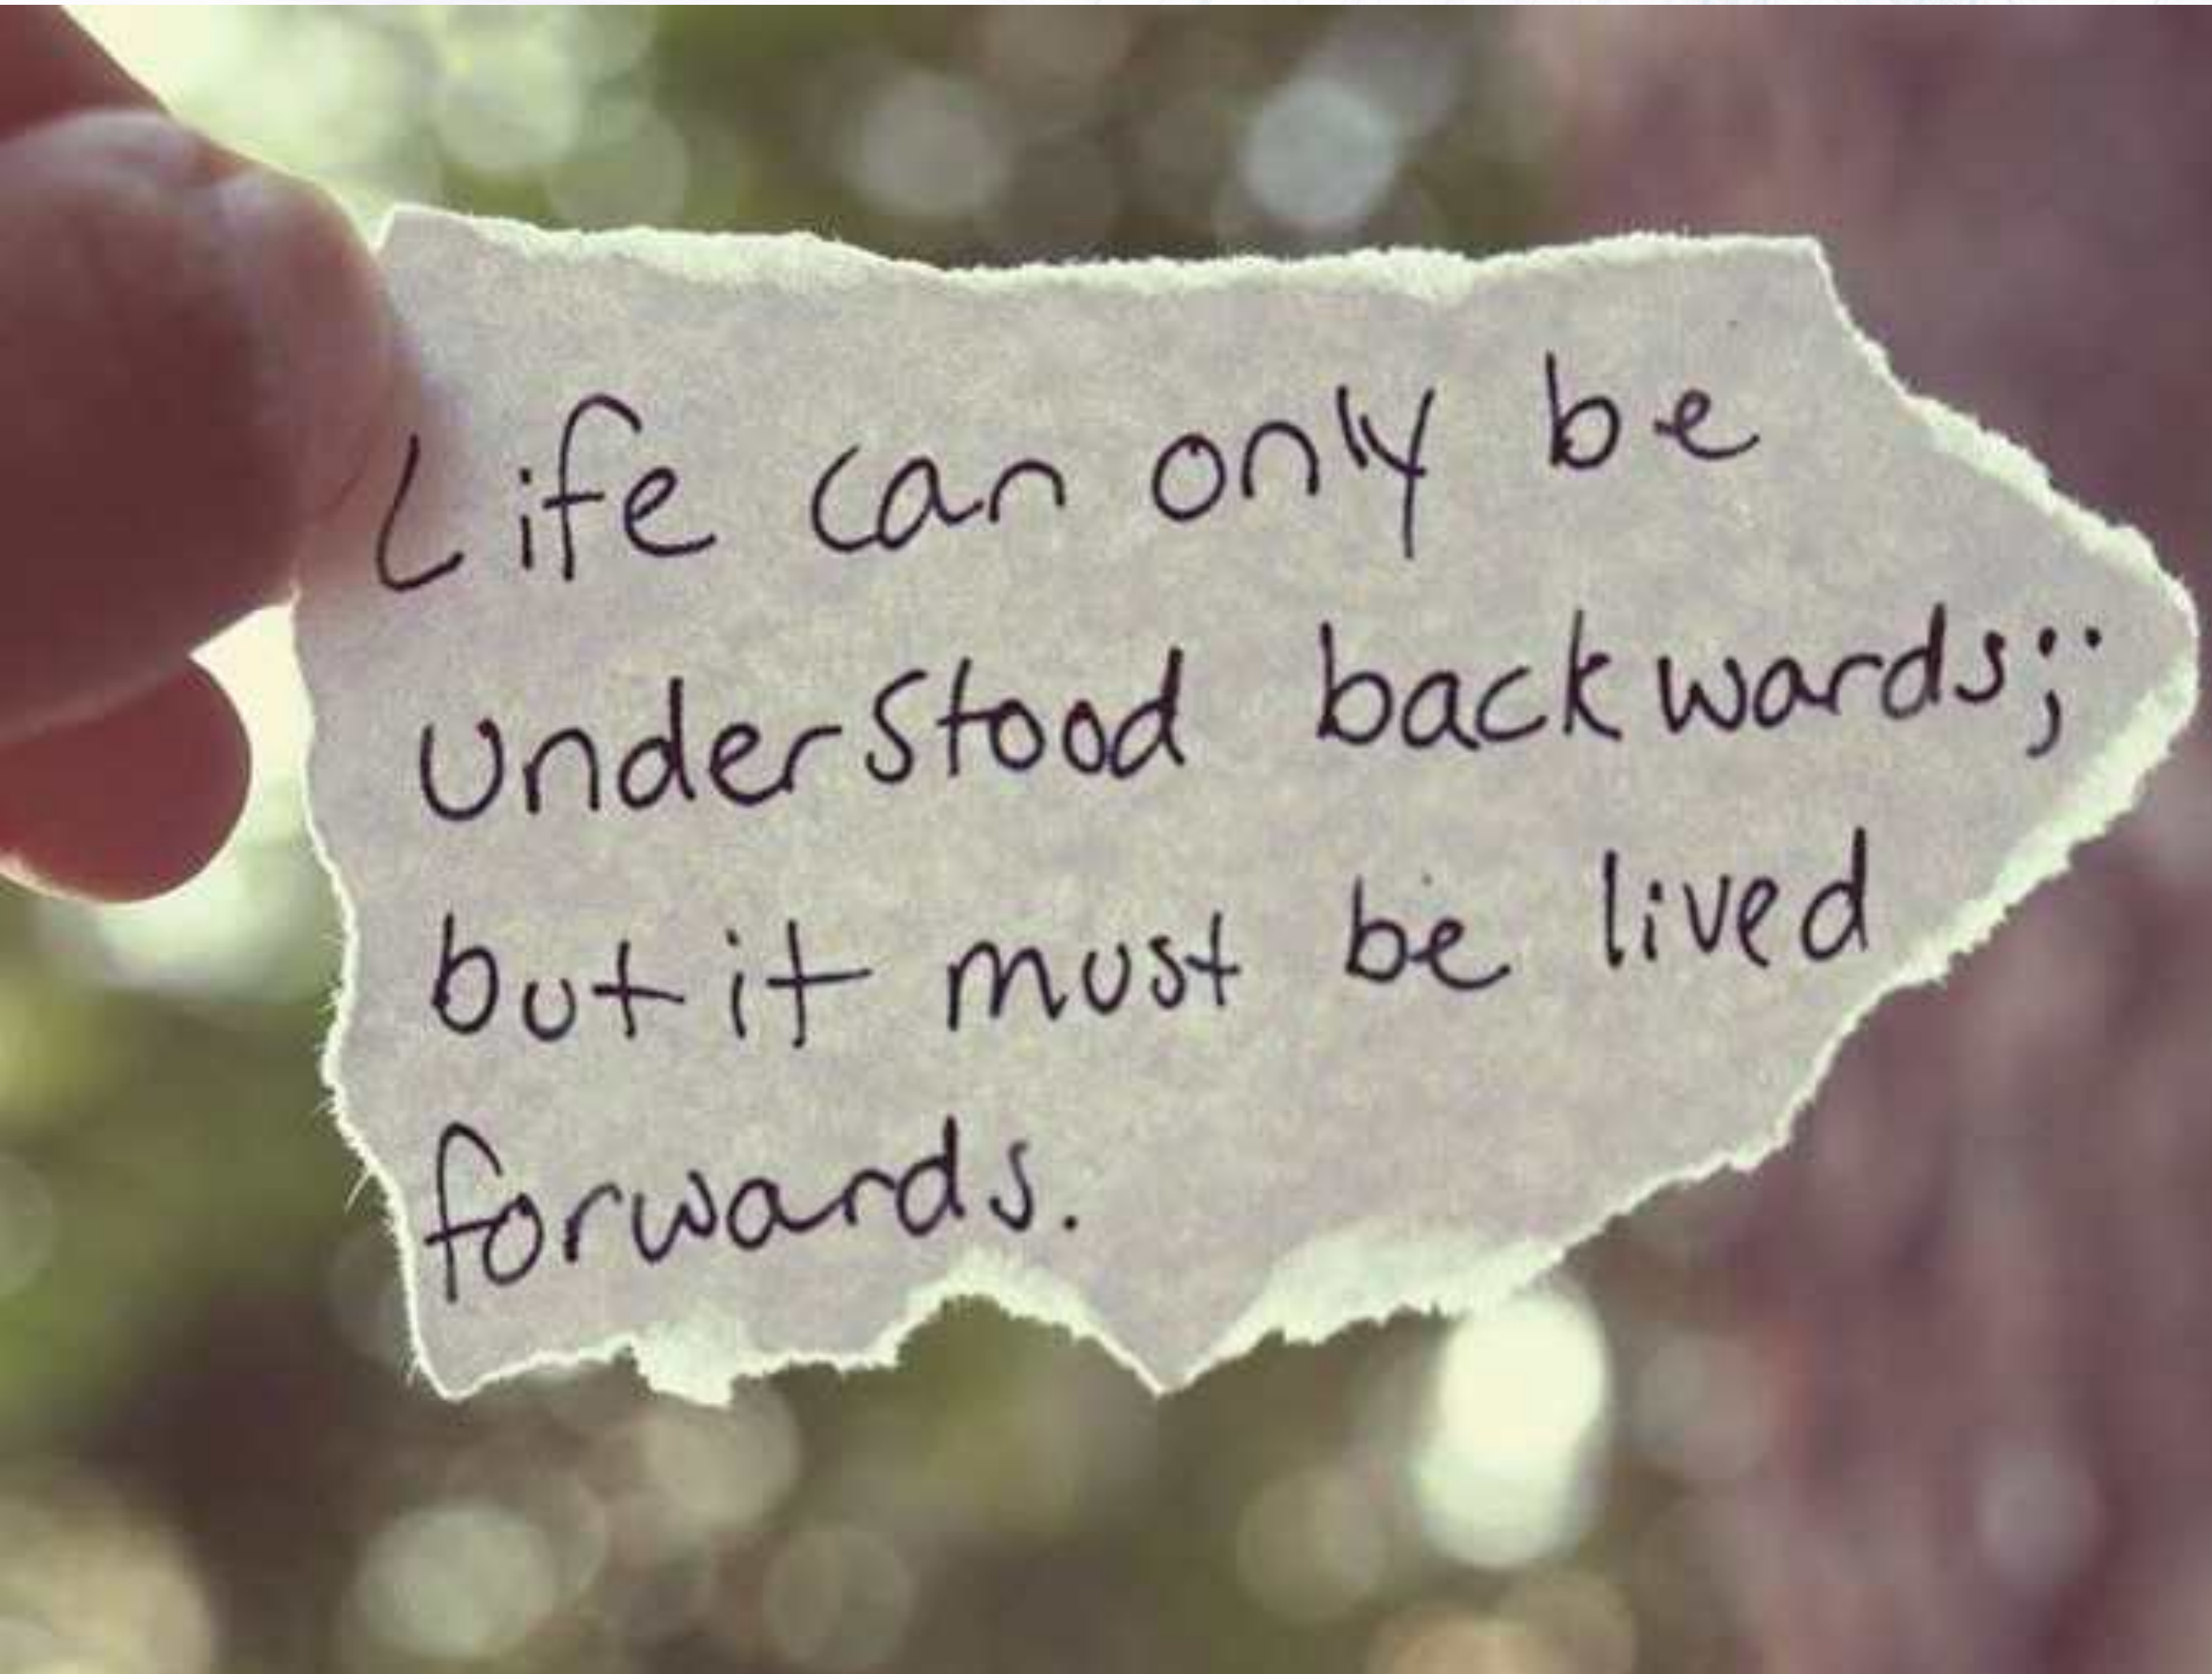


What is your question for today?

What is your goal for today?

- Take a piece of paper:
 - Write down your goals for today
 - Artificial Intelligence
 - ... (any subject in this context)
 - Think on what you can bring to the table: Context or Content
- Take 1 minute

At the story of my life.



Life can only be
understood backwards;
but it must be lived
forwards.



TECHNOLOGY CHANGES EVERYTHING

1. **Values** Technology can have built-in values.
2. **Behavior** Technology can take over what we previously had to do with laws/force
3. **Brain** Technology can take over what we previously had to do with our brains
4. **Will-power** Technology can take over what we previously had to do with our will-power



Progression of AI capabilities

Generative AI

AI systems that create new content or solutions autonomously.



Prescriptive AI

AI systems that recommend actions based on predictive insights.



Predictive AI

AI systems that use historical data to forecast future outcomes.



Descriptive AI

AI systems that analyze and summarize data to provide insights.



AI Adventure



AI Foundational Layer: Machine Learning (ML)

- Machine Learning is the core technology that powers most AI systems today. It enables computers to learn from data patterns and make predictions without being explicitly programmed.
- Supervised Learning: AI learns from labeled data (like image recognition).
- Unsupervised Learning: AI detects patterns in unlabeled data (like customer segmentation).
- Reinforcement Learning: AI learns through trial and error, receiving feedback for actions (like AlphaGo in board games).
- ML is the basis for all current AI applications and will underpin future AI developments.

History of AI



Neural Networks

1950-1970



Machine Learning

1980-2010

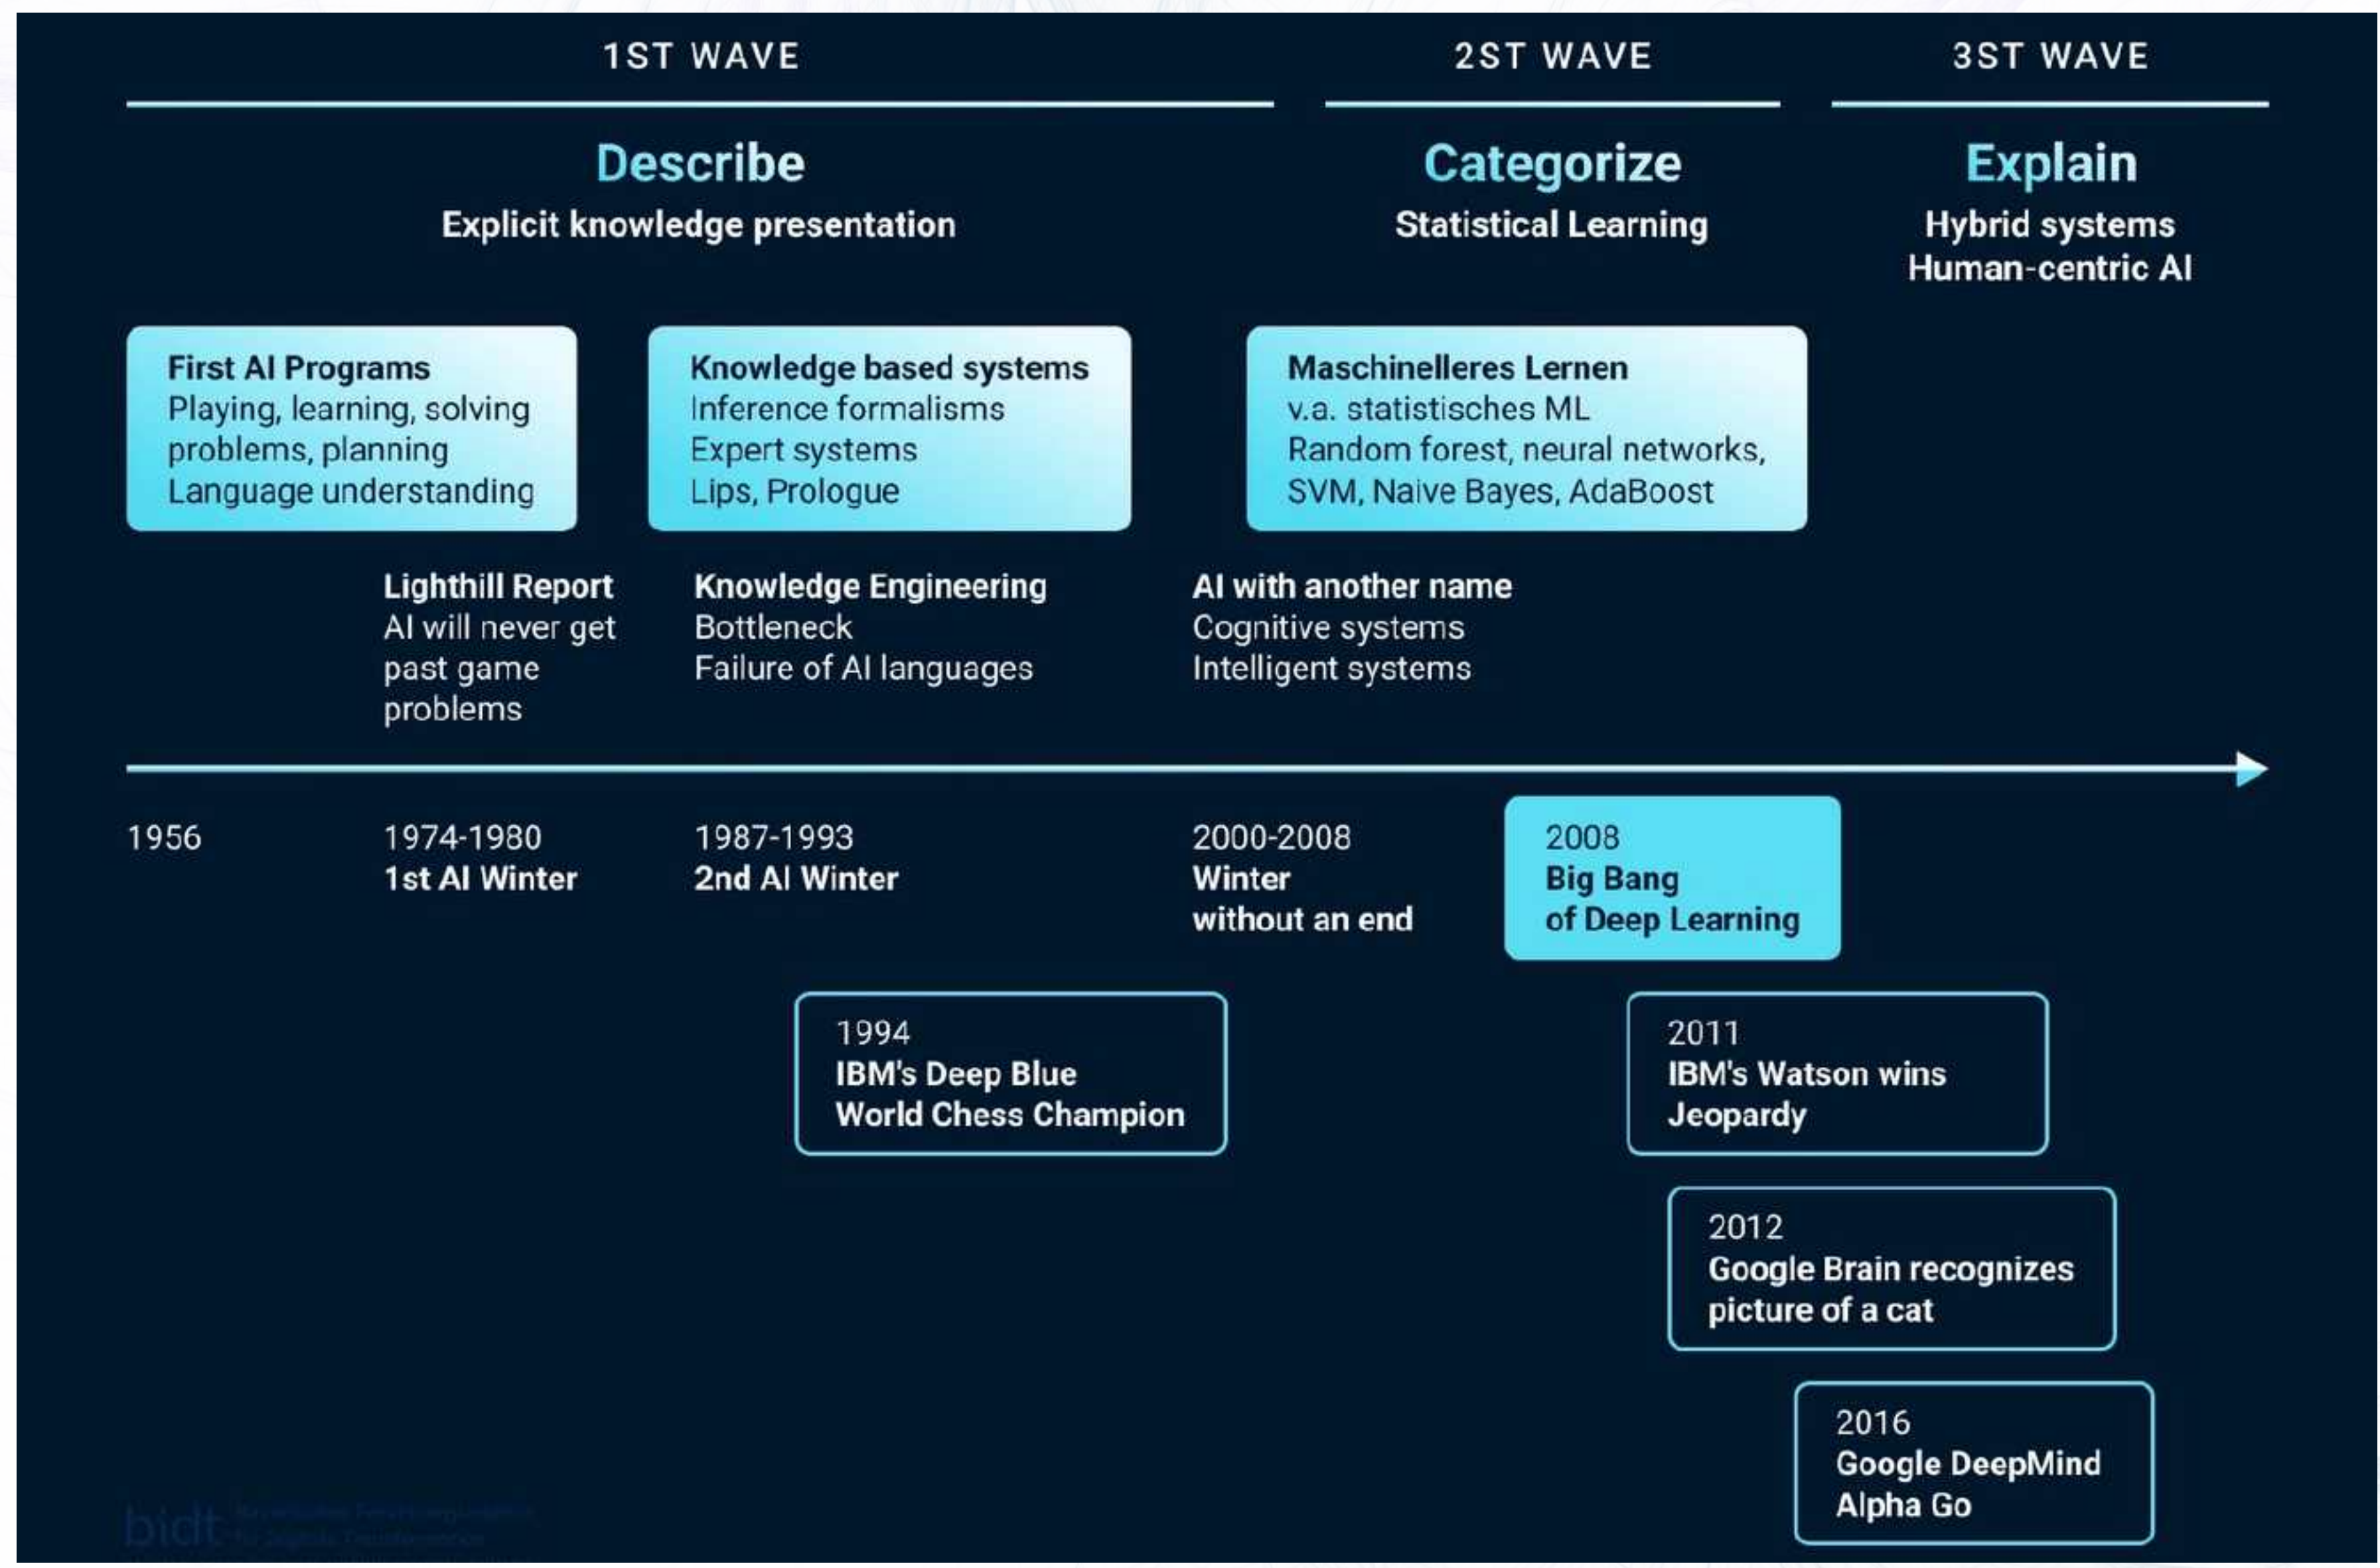


Deep Learning

Today



History of AI



AI versus AI

Generative uses of AI

Classifying uses of AI

What is genAI?



Generating new data by learning from large data sets and identifying patterns within – capabilities include **text, visual, sound, etc.**

Can **predict likelihood** of a **right response**

Open-ended and creative

Examples:

ChatGPT by OpenAI – for text

StyleGAN¹ by NVIDIA – for visuals

NSynth by Google – for sound



VS

What is “traditional” AI?



Solving specific tasks by making predictions based on previously analyzed sets of data and **predefined rules**

Can **recognize patterns** in data

Goal-oriented and specific

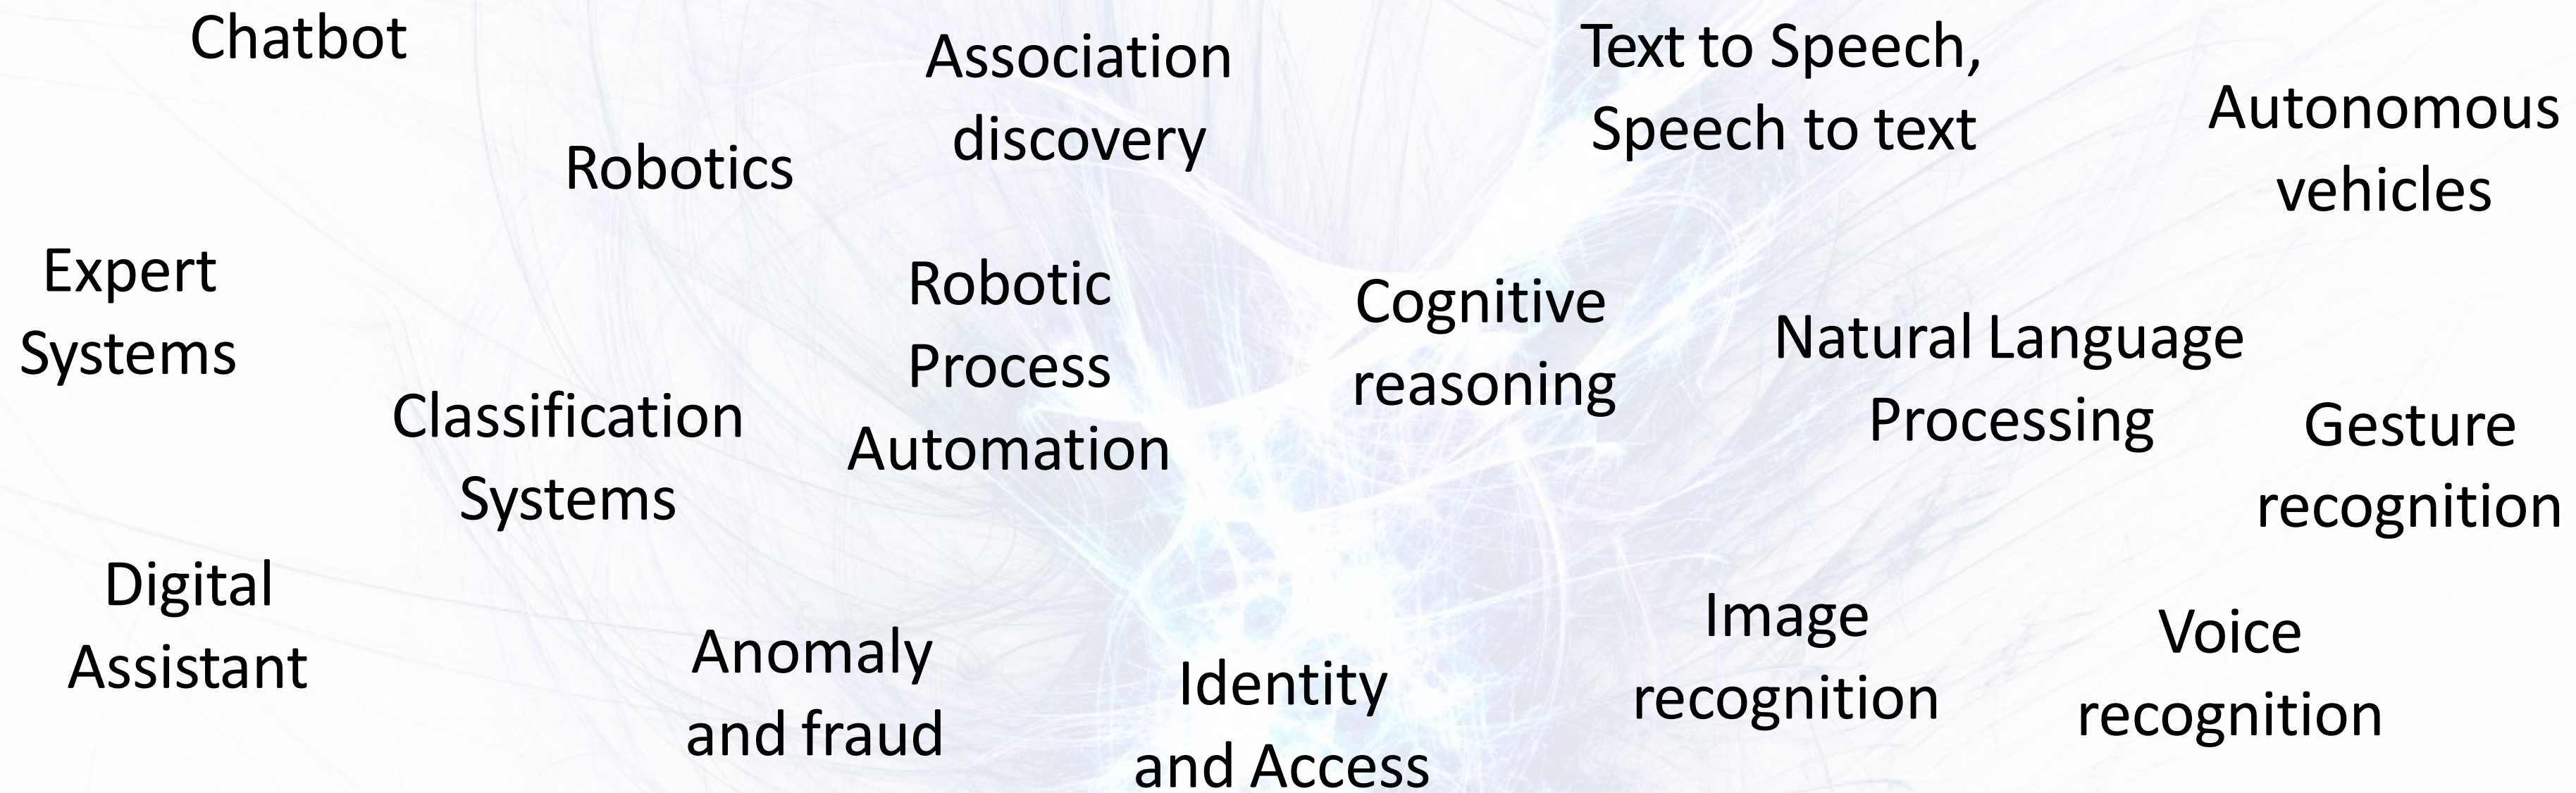
Examples:

Voice assistants

Text recognition / OCR

Recommender systems

AI models / applications based upon knowledge captured in DATA



Each application area has:

- Specific algorithms focused on answering a specified question
- Different types of data (numbers, text, sound, image etc)
- Optimized tools from different suppliers (BigML, Google, AWS, MS Azure, OpenAI or OpenSource etc)
- Requisite skills tailored to the application at hand

Machine Learning Accessibility

Harvard
Business
Review

TECHNOLOGY

When AI Becomes an Everyday Technology

by Andrew Moore

JUNE 07, 2019

Summary Save Share Comment 2 Text Size Print \$8.95 Buy Copies

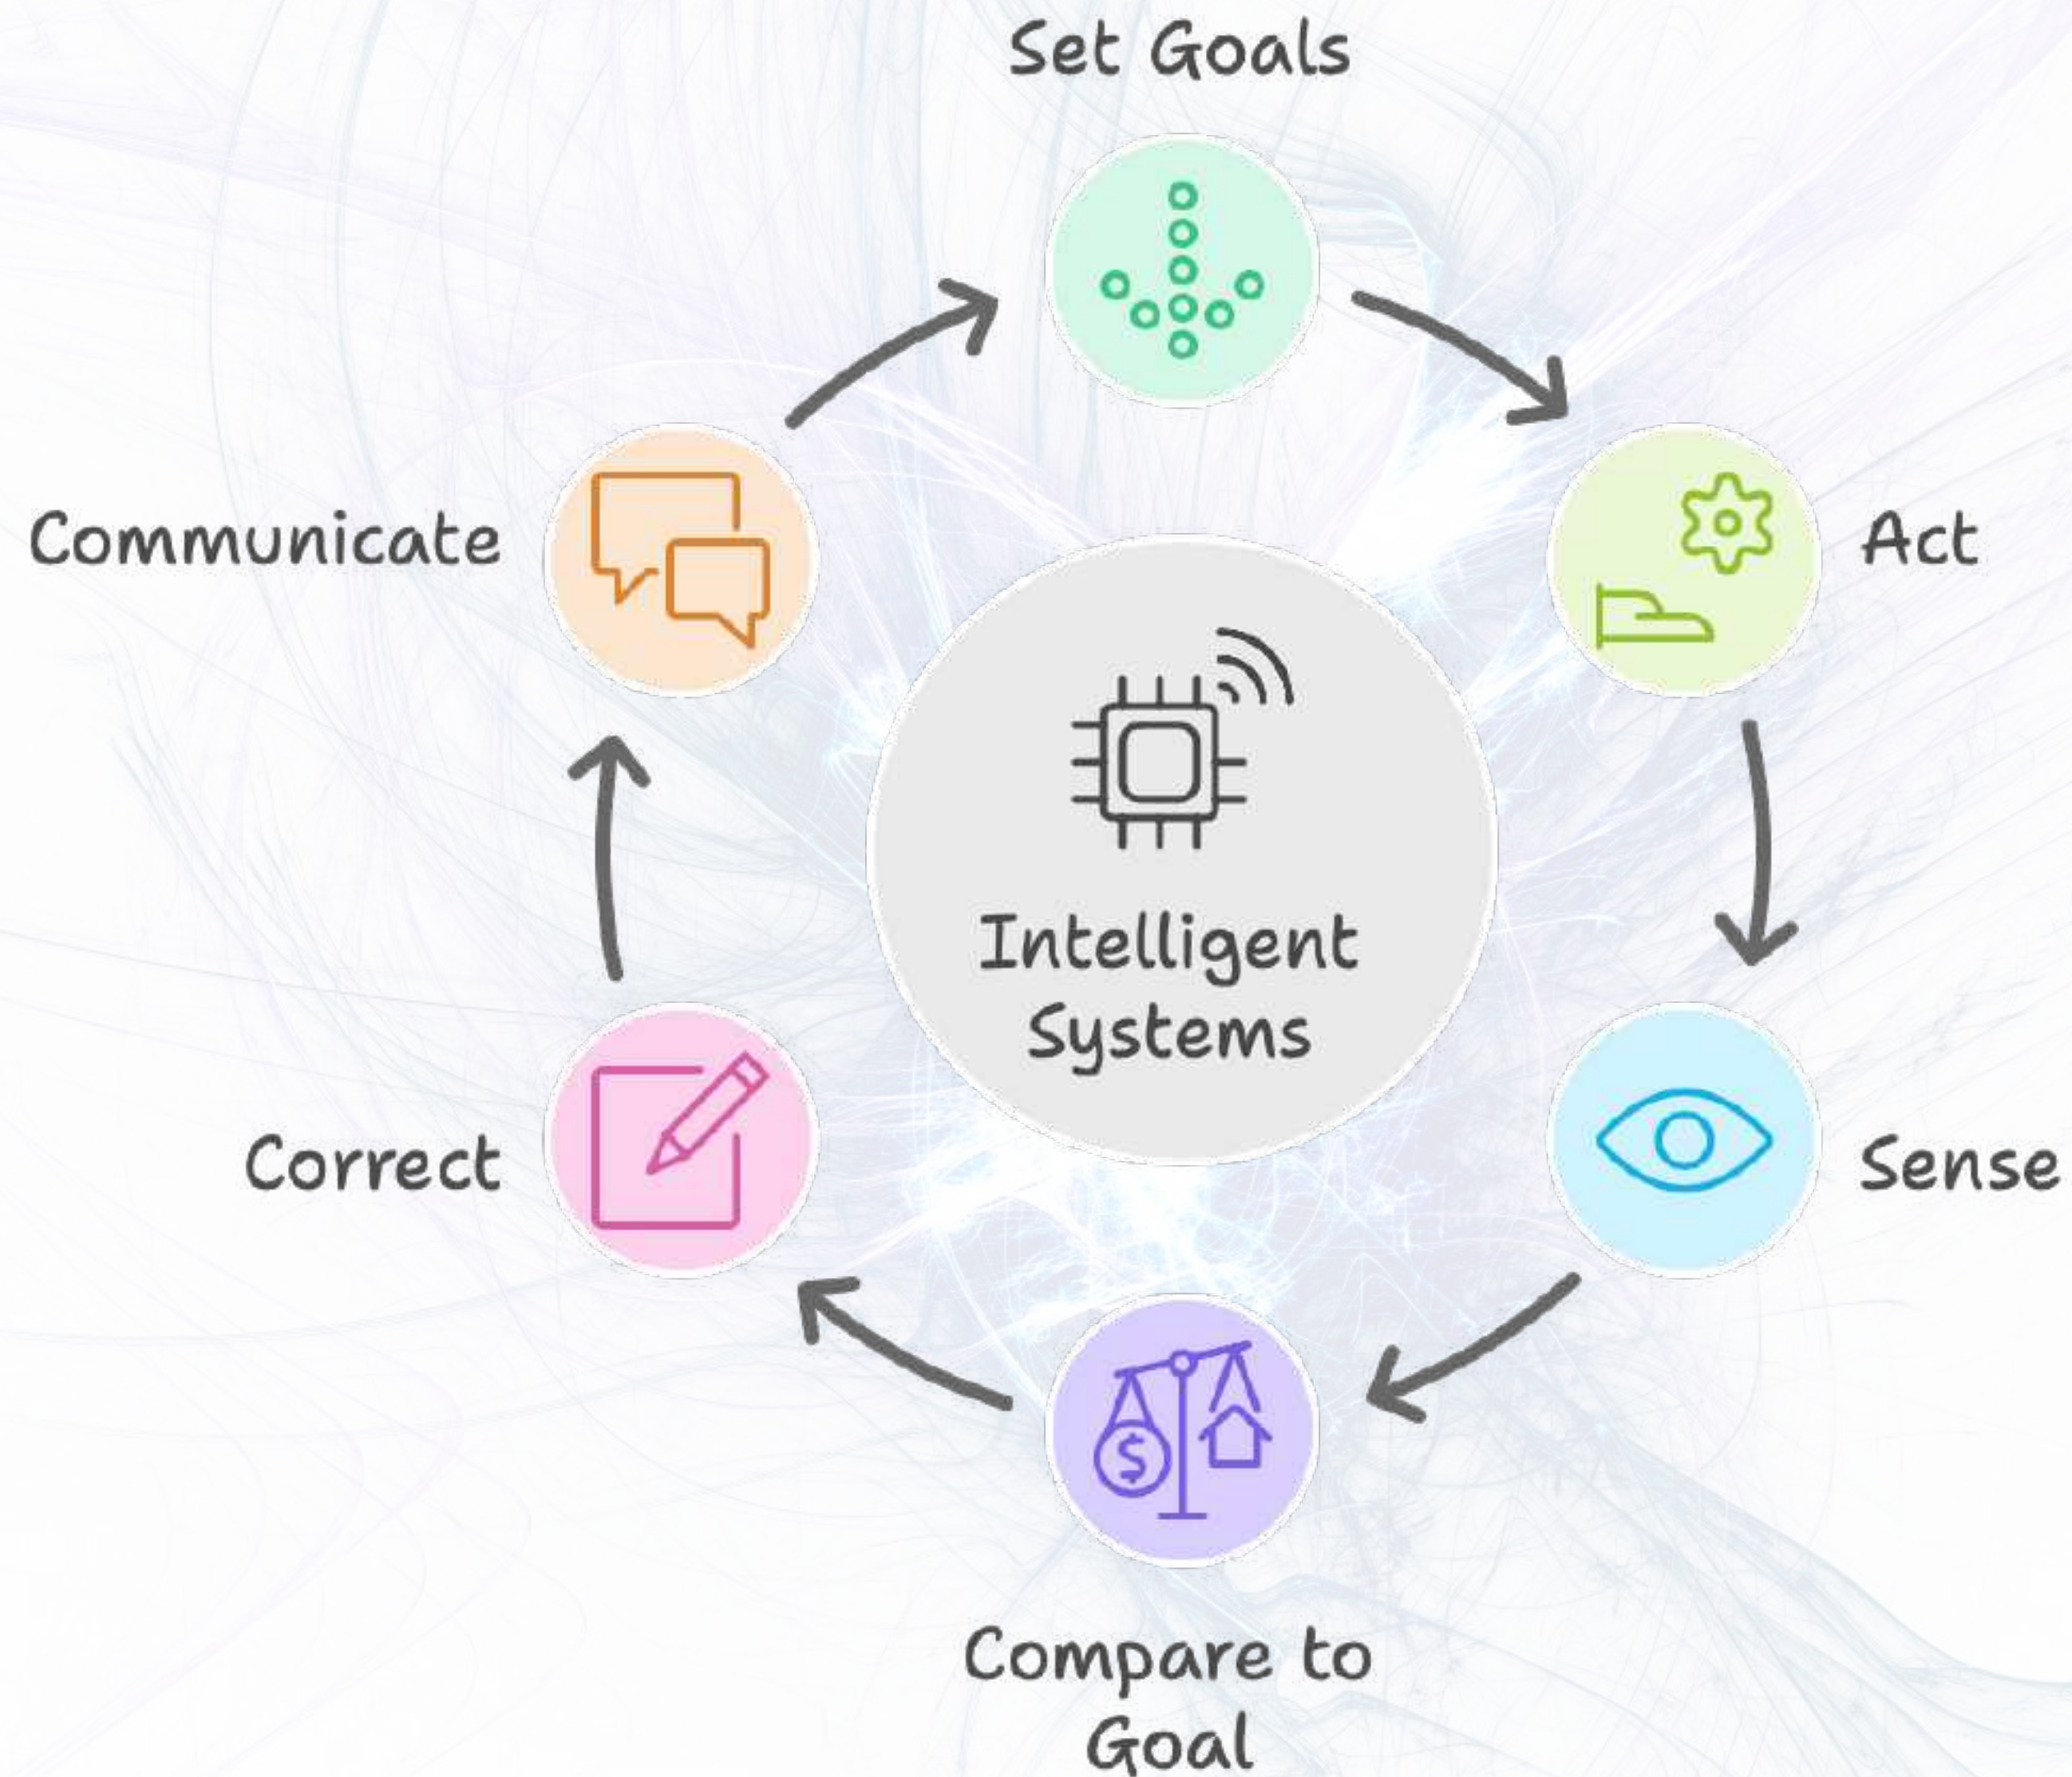


“After years of hype around mysterious neural networks and the PhD researchers who design them, we’re entering an age in which just about anyone can leverage the power of intelligent algorithms to solve the problems that matter to them. Ironically, although breakthroughs get the headlines, it’s accessibility that really changes the world. That’s why, after such an eventful decade, a lack of hype around machine learning may be the most exciting development yet.”

— Andrew Moore, Google

SOURCE:

Cybernetics AI



Artificial Intelligence

Machine Learning

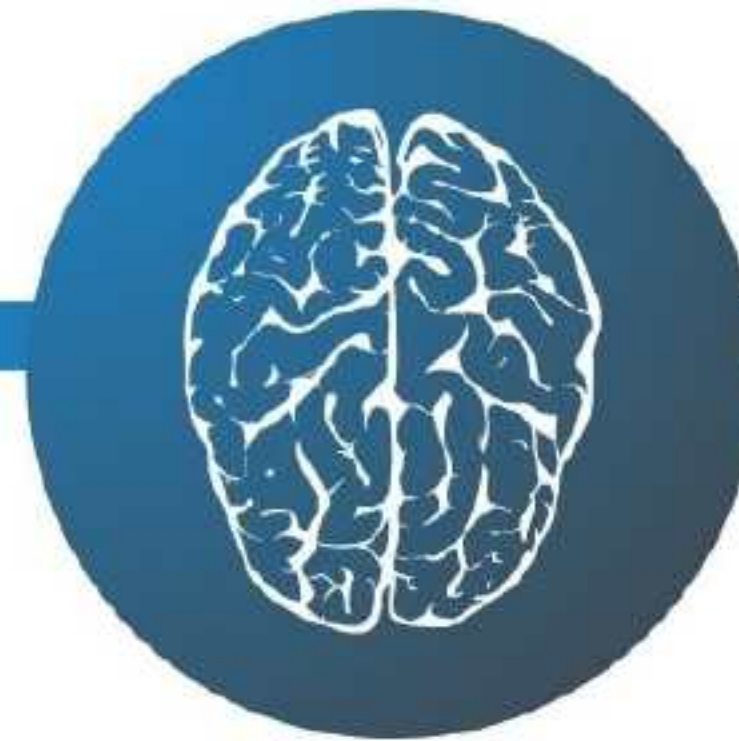


Narrow Artificial Intelligence (ANI)

Stage One: Machines imitate human behavior, specializing in one area to solve a problem.

i.e. Siri, ChatGPT, Alexa

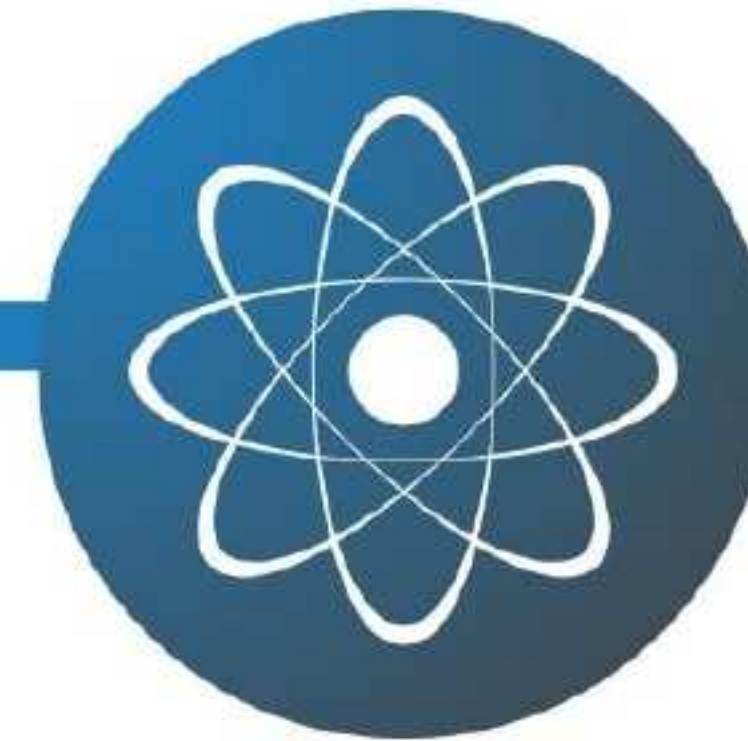
Machine Intelligence



Artificial General Intelligence (AGI)

Stage Two: Machines can continuously learn and are as smart as humans.

Machine Consciousness



Artificial Super Intelligence (ASI)

Stage Three: Machines that are smarter than humans across the board.

Artificial Intelligence

Machine Learning

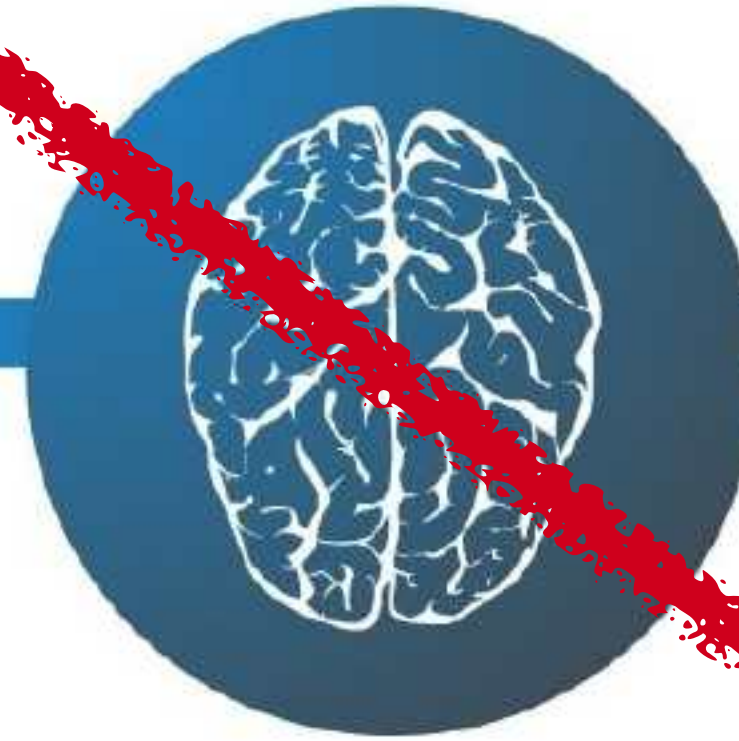


Narrow Artificial Intelligence (ANI)

Stage One: Machines imitate human behavior, specializing in one area to solve a problem.

i.e. Siri, ChatGPT, Alexa

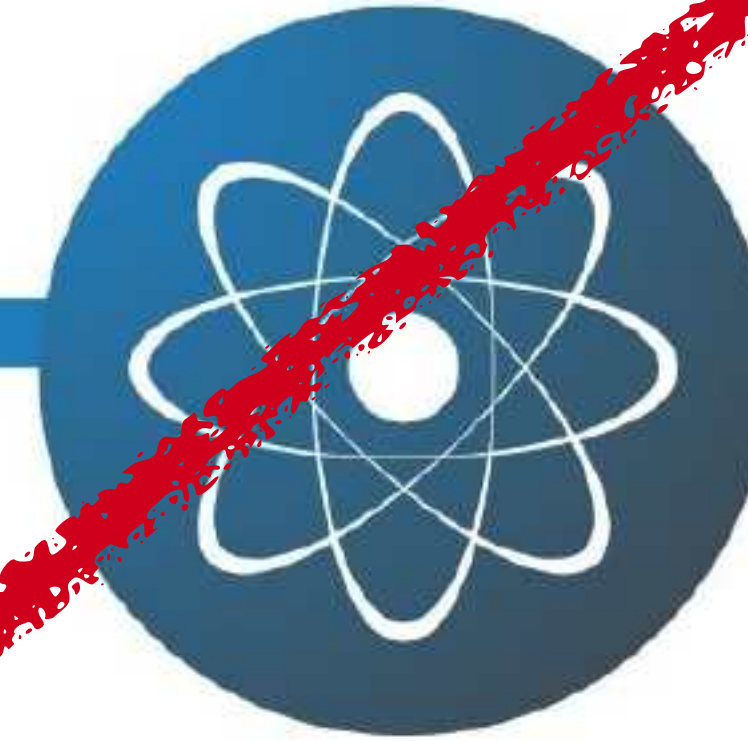
Machine Intelligence



Artificial General Intelligence (AGI)

Stage Two: Machines can continuously learn and are as smart as humans.

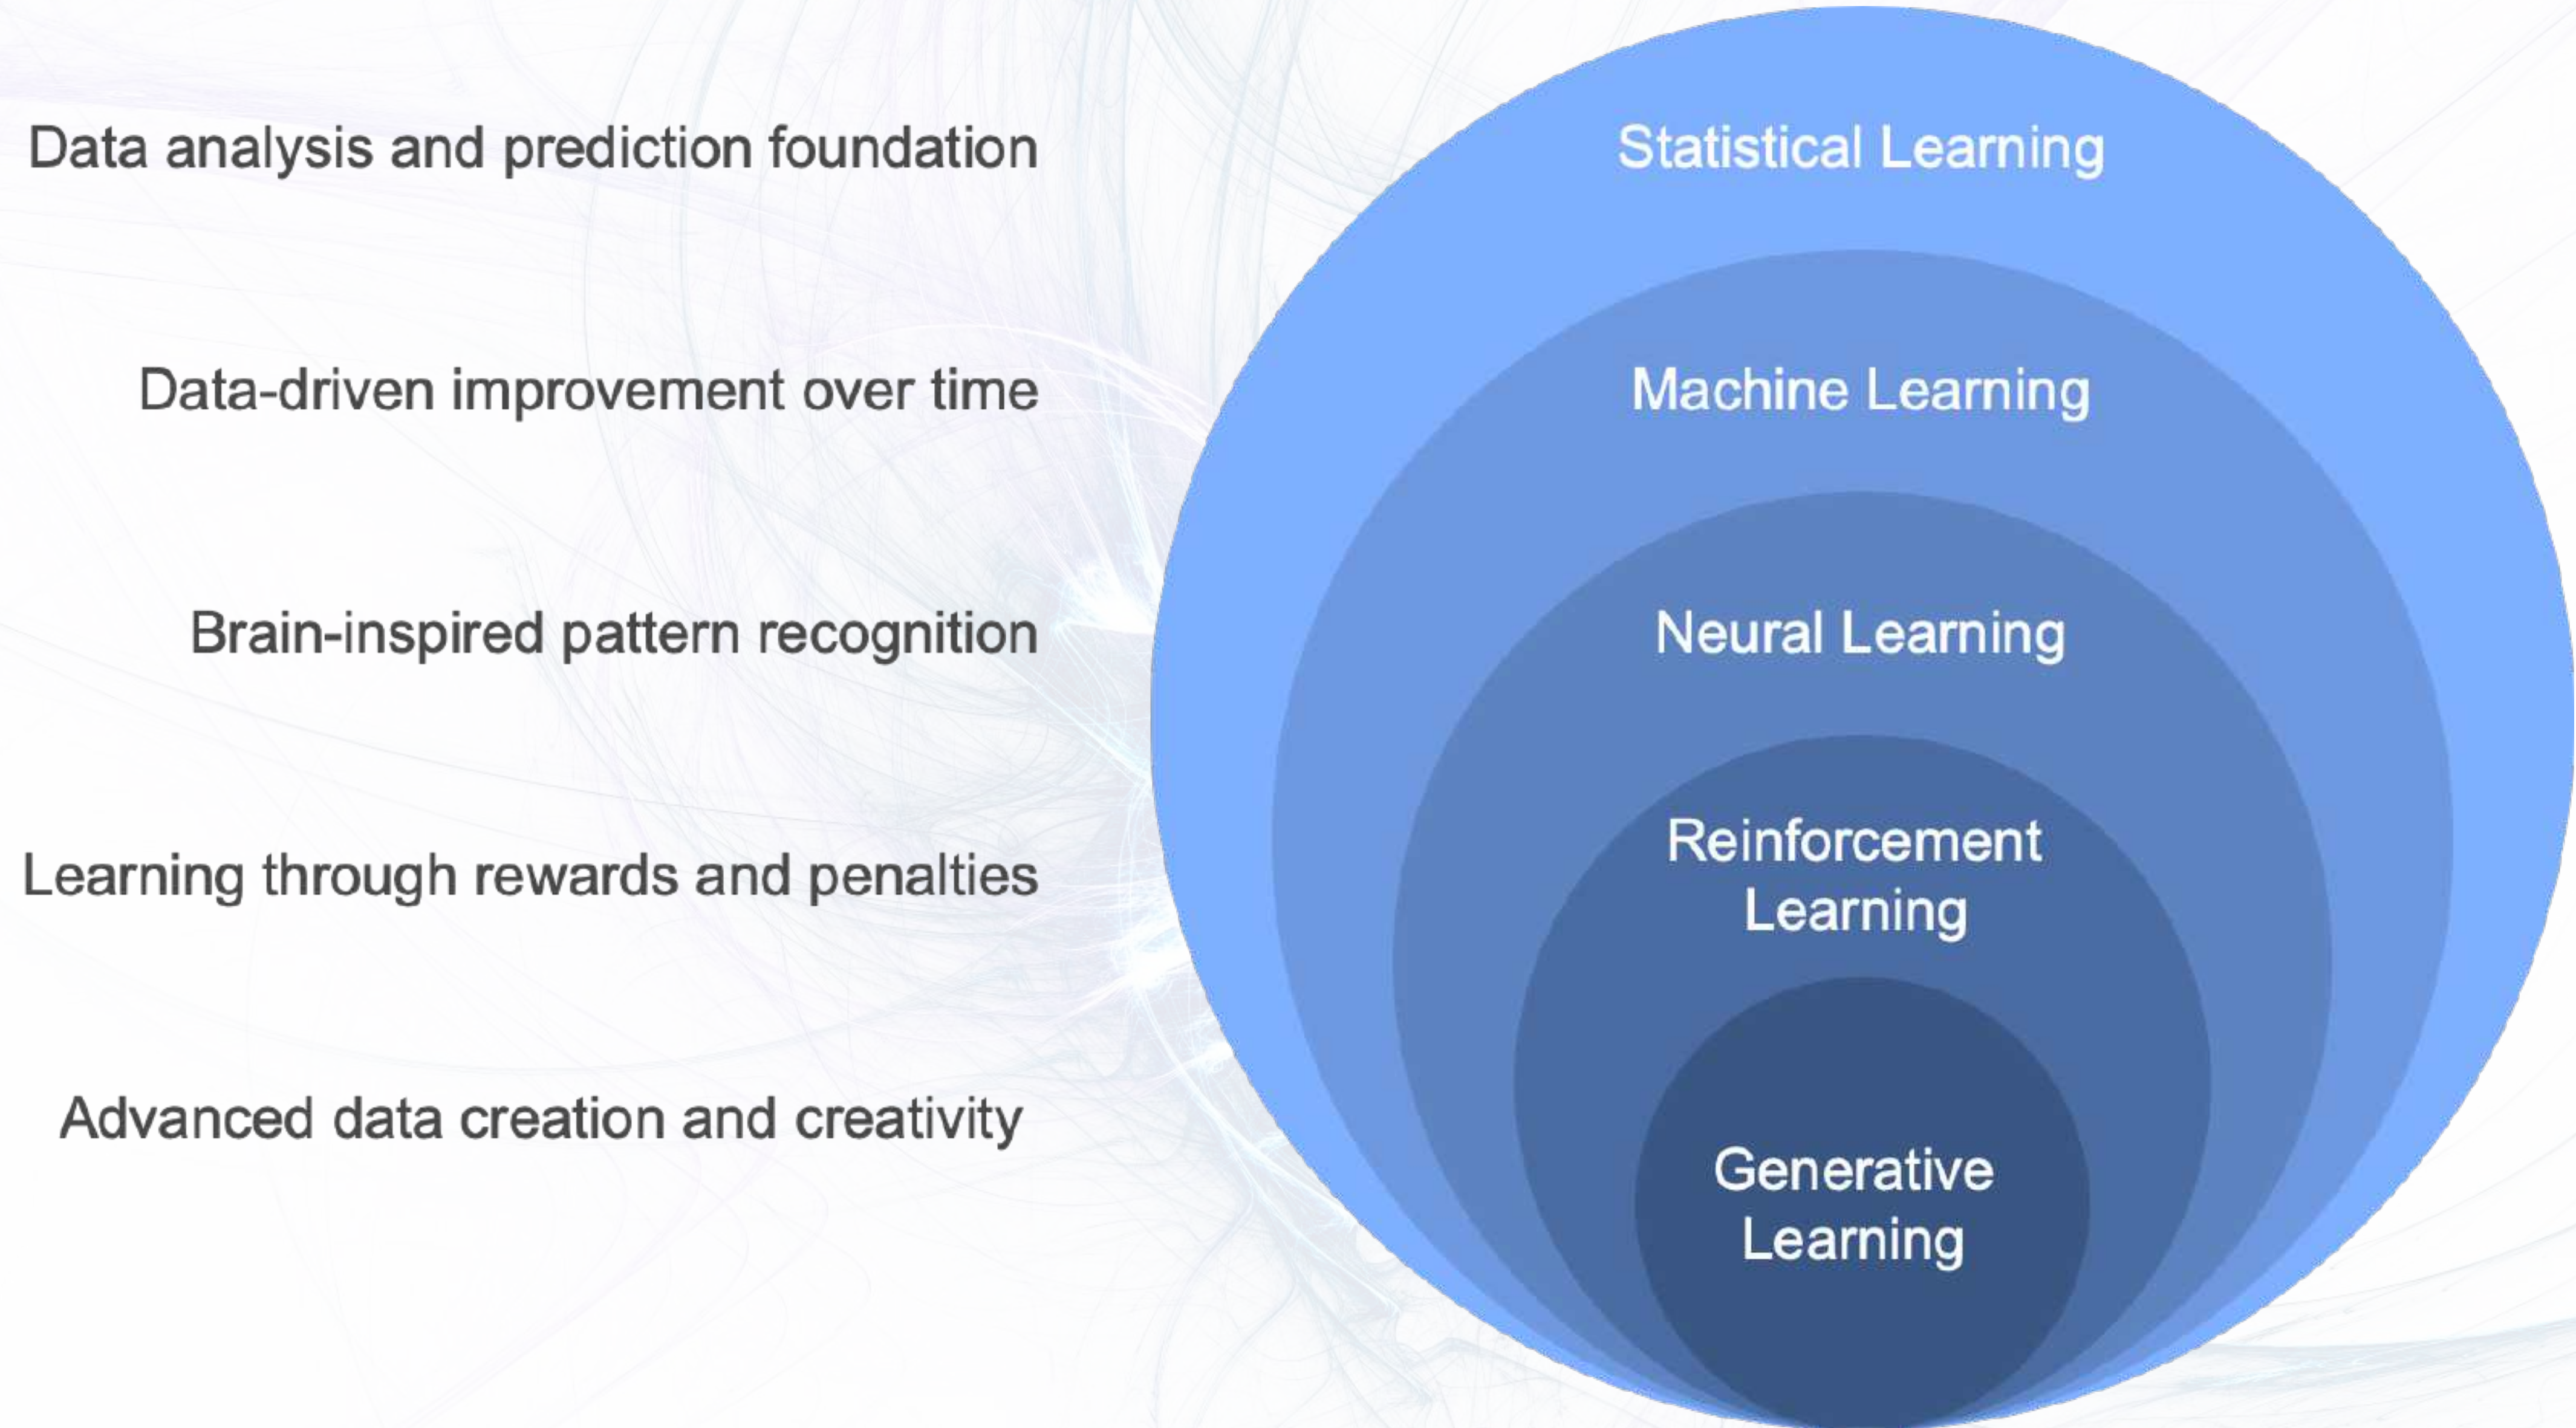
Machine Consciousness



Artificial Super Intelligence (ASI)

Stage Three: Machines that are smarter than humans across the board.

Hierarchy of AI learning paradigms



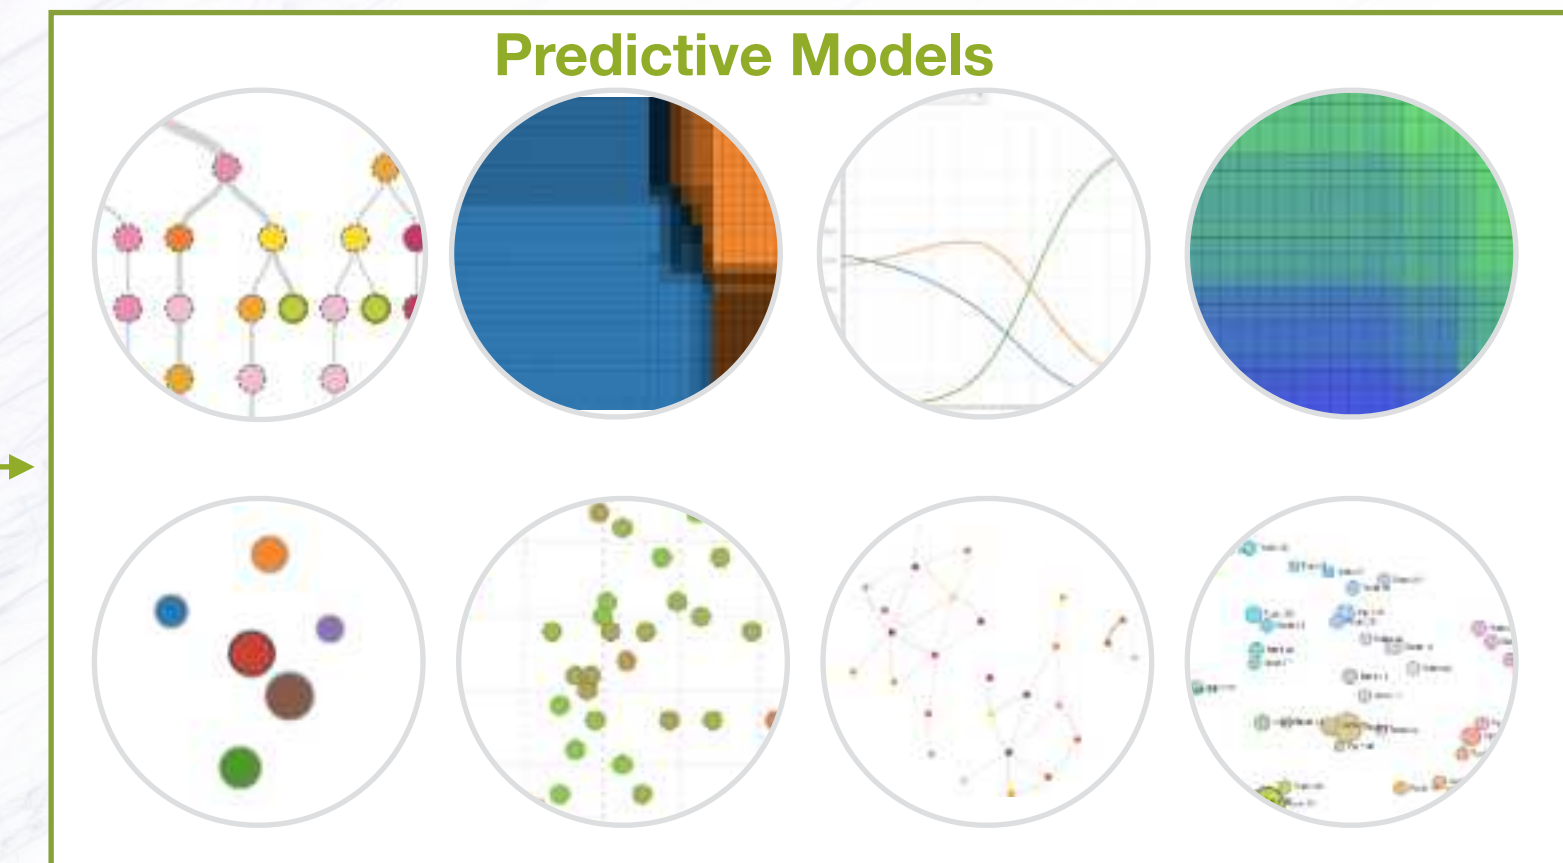
Classifying AI

Getting Value from Machine Learning Isn't About Fancier Algorithms — It's About Making It Easier to Use

by Ben Schrock, Max Kanter, Kalyan Veeramachaneni, Sanjeev Vohra, and Rajendra Prasad
MARCH 04, 2018



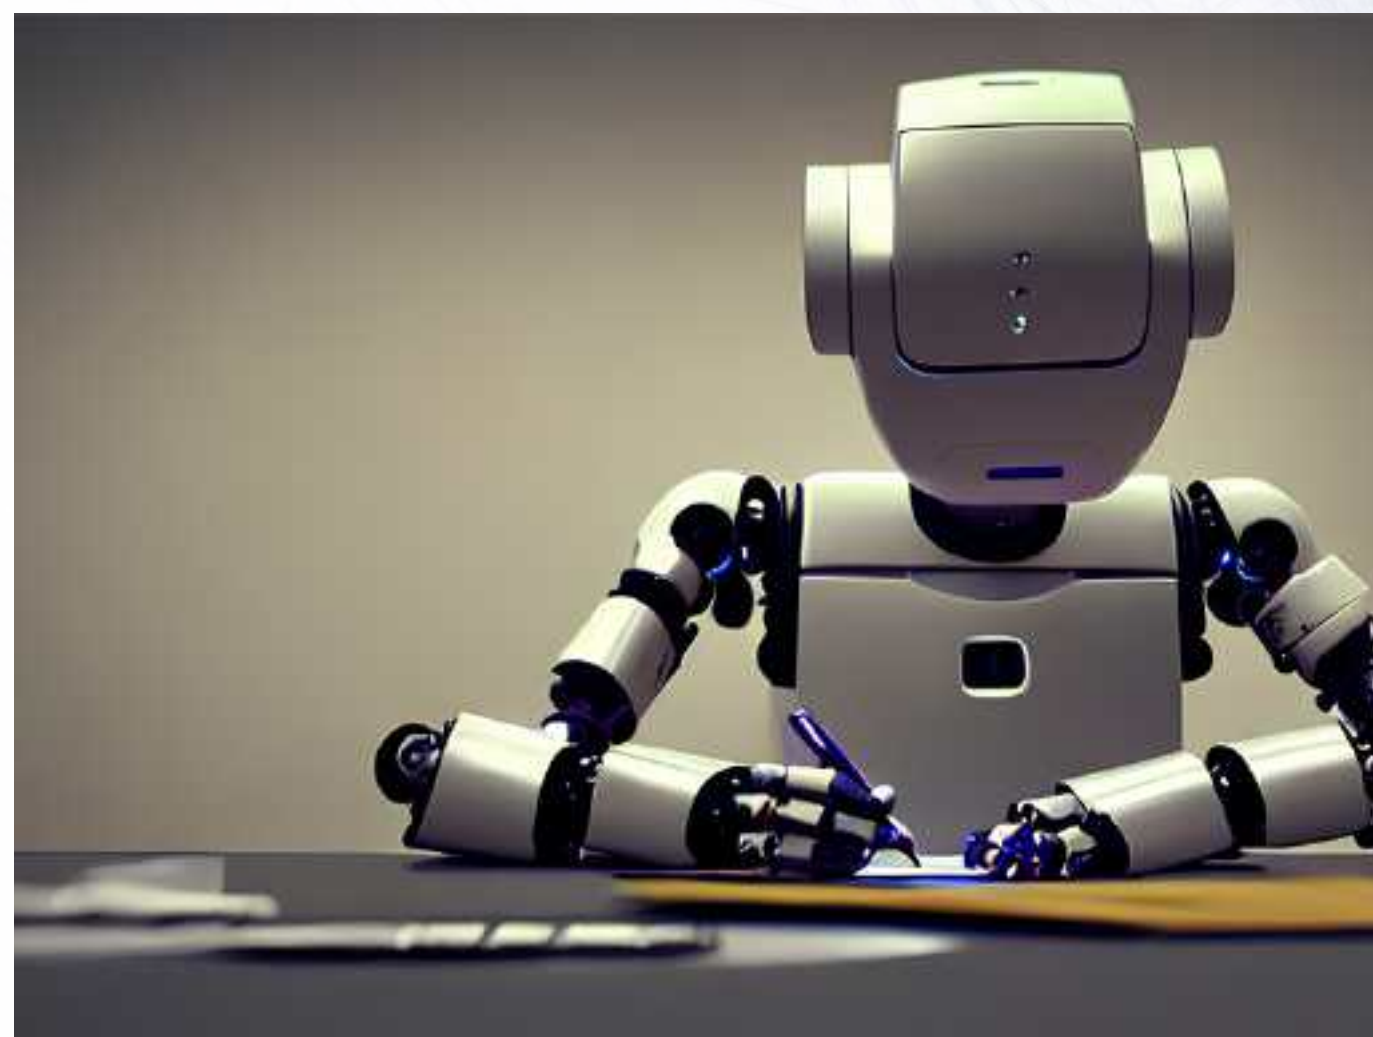
AIRLINE	ORIGIN	DESTINATION	DEPARTURE DELAY	DISTANCE	ARRIVAL DELAY
AS	ANC	SEA	-11	1448,0	-22
AA	LAX	PBI	-8	2330,0	-9
US	SFO	CLT	-2	2296,0	5
AA	LAX	MIA	-5	2342,0	-9
AS	SEA	ANC	-1	1448,0	-21
DL	SFO	MSP	-5	1589	8
NK	LAS	MSP	-6	1299	-17
US	LAX	CLT	14	2125,0	-10
AA	SFO	DFW	-11	1464,0	-13
DL	LAS	ATL	3	1747,0	-15



Finding **patterns** in data that can be used to make inferences...(decisions)

Generating AI

ChatGPT		
Examples	Capabilities	Limitations
"Explain quantum computing in simple terms" →	Remembers what user said earlier in the conversation	May occasionally generate incorrect information
"Got any creative ideas for a 10 year old's birthday?" →	Allows user to provide follow-up corrections	May occasionally produce harmful instructions or biased content
"How do I make an HTTP request in Javascript?" →	Trained to decline inappropriate requests	Limited knowledge of world and events after 2021



Write a short article on how photovoltaics work

Photovoltaics, or solar cells, are devices that convert sunlight into electricity. They are made of semiconductor materials, such as silicon, which are highly sensitive to light. When sunlight hits the solar cell, it knocks electrons loose from the atoms in the semiconductor material, creating a flow of electricity.

The solar cell is made up of two layers of semiconductor material, each with a different electrical charge. The top layer is positively charged, while the bottom layer is negatively charged. When sunlight hits the solar cell, it knocks loose electrons from the atoms in the semiconductor material, causing them to flow from the negatively charged layer to the positively charged layer. This flow of electrons creates an electric current, which can be used to power electrical devices.

Solar cells are typically connected together in a panel, and multiple panels can be combined to form a solar array. The solar array is then connected to an inverter, which converts the direct current (DC) electricity produced by the solar cells into alternating current (AC) electricity, which is what is used in homes and businesses.

Photovoltaics are a clean and renewable source of energy, as they produce electricity from sunlight, a natural and abundant resource. They also do not produce any emissions or pollutants, making them an environmentally friendly option for generating electricity. Photovoltaics are becoming increasingly popular as a source of electricity for homes, businesses, and even entire communities.

Programming AI

write a inference engine in clojure

Sure, here's an example of a simple inference engine implemented in Clojure:

```
(def knowledge-base {:rules {"rule1" ["fact1" "fact2"]
                             "rule2" ["fact3"]}}
  :facts {"fact1" true
          "fact2" false
          "fact3" true}})

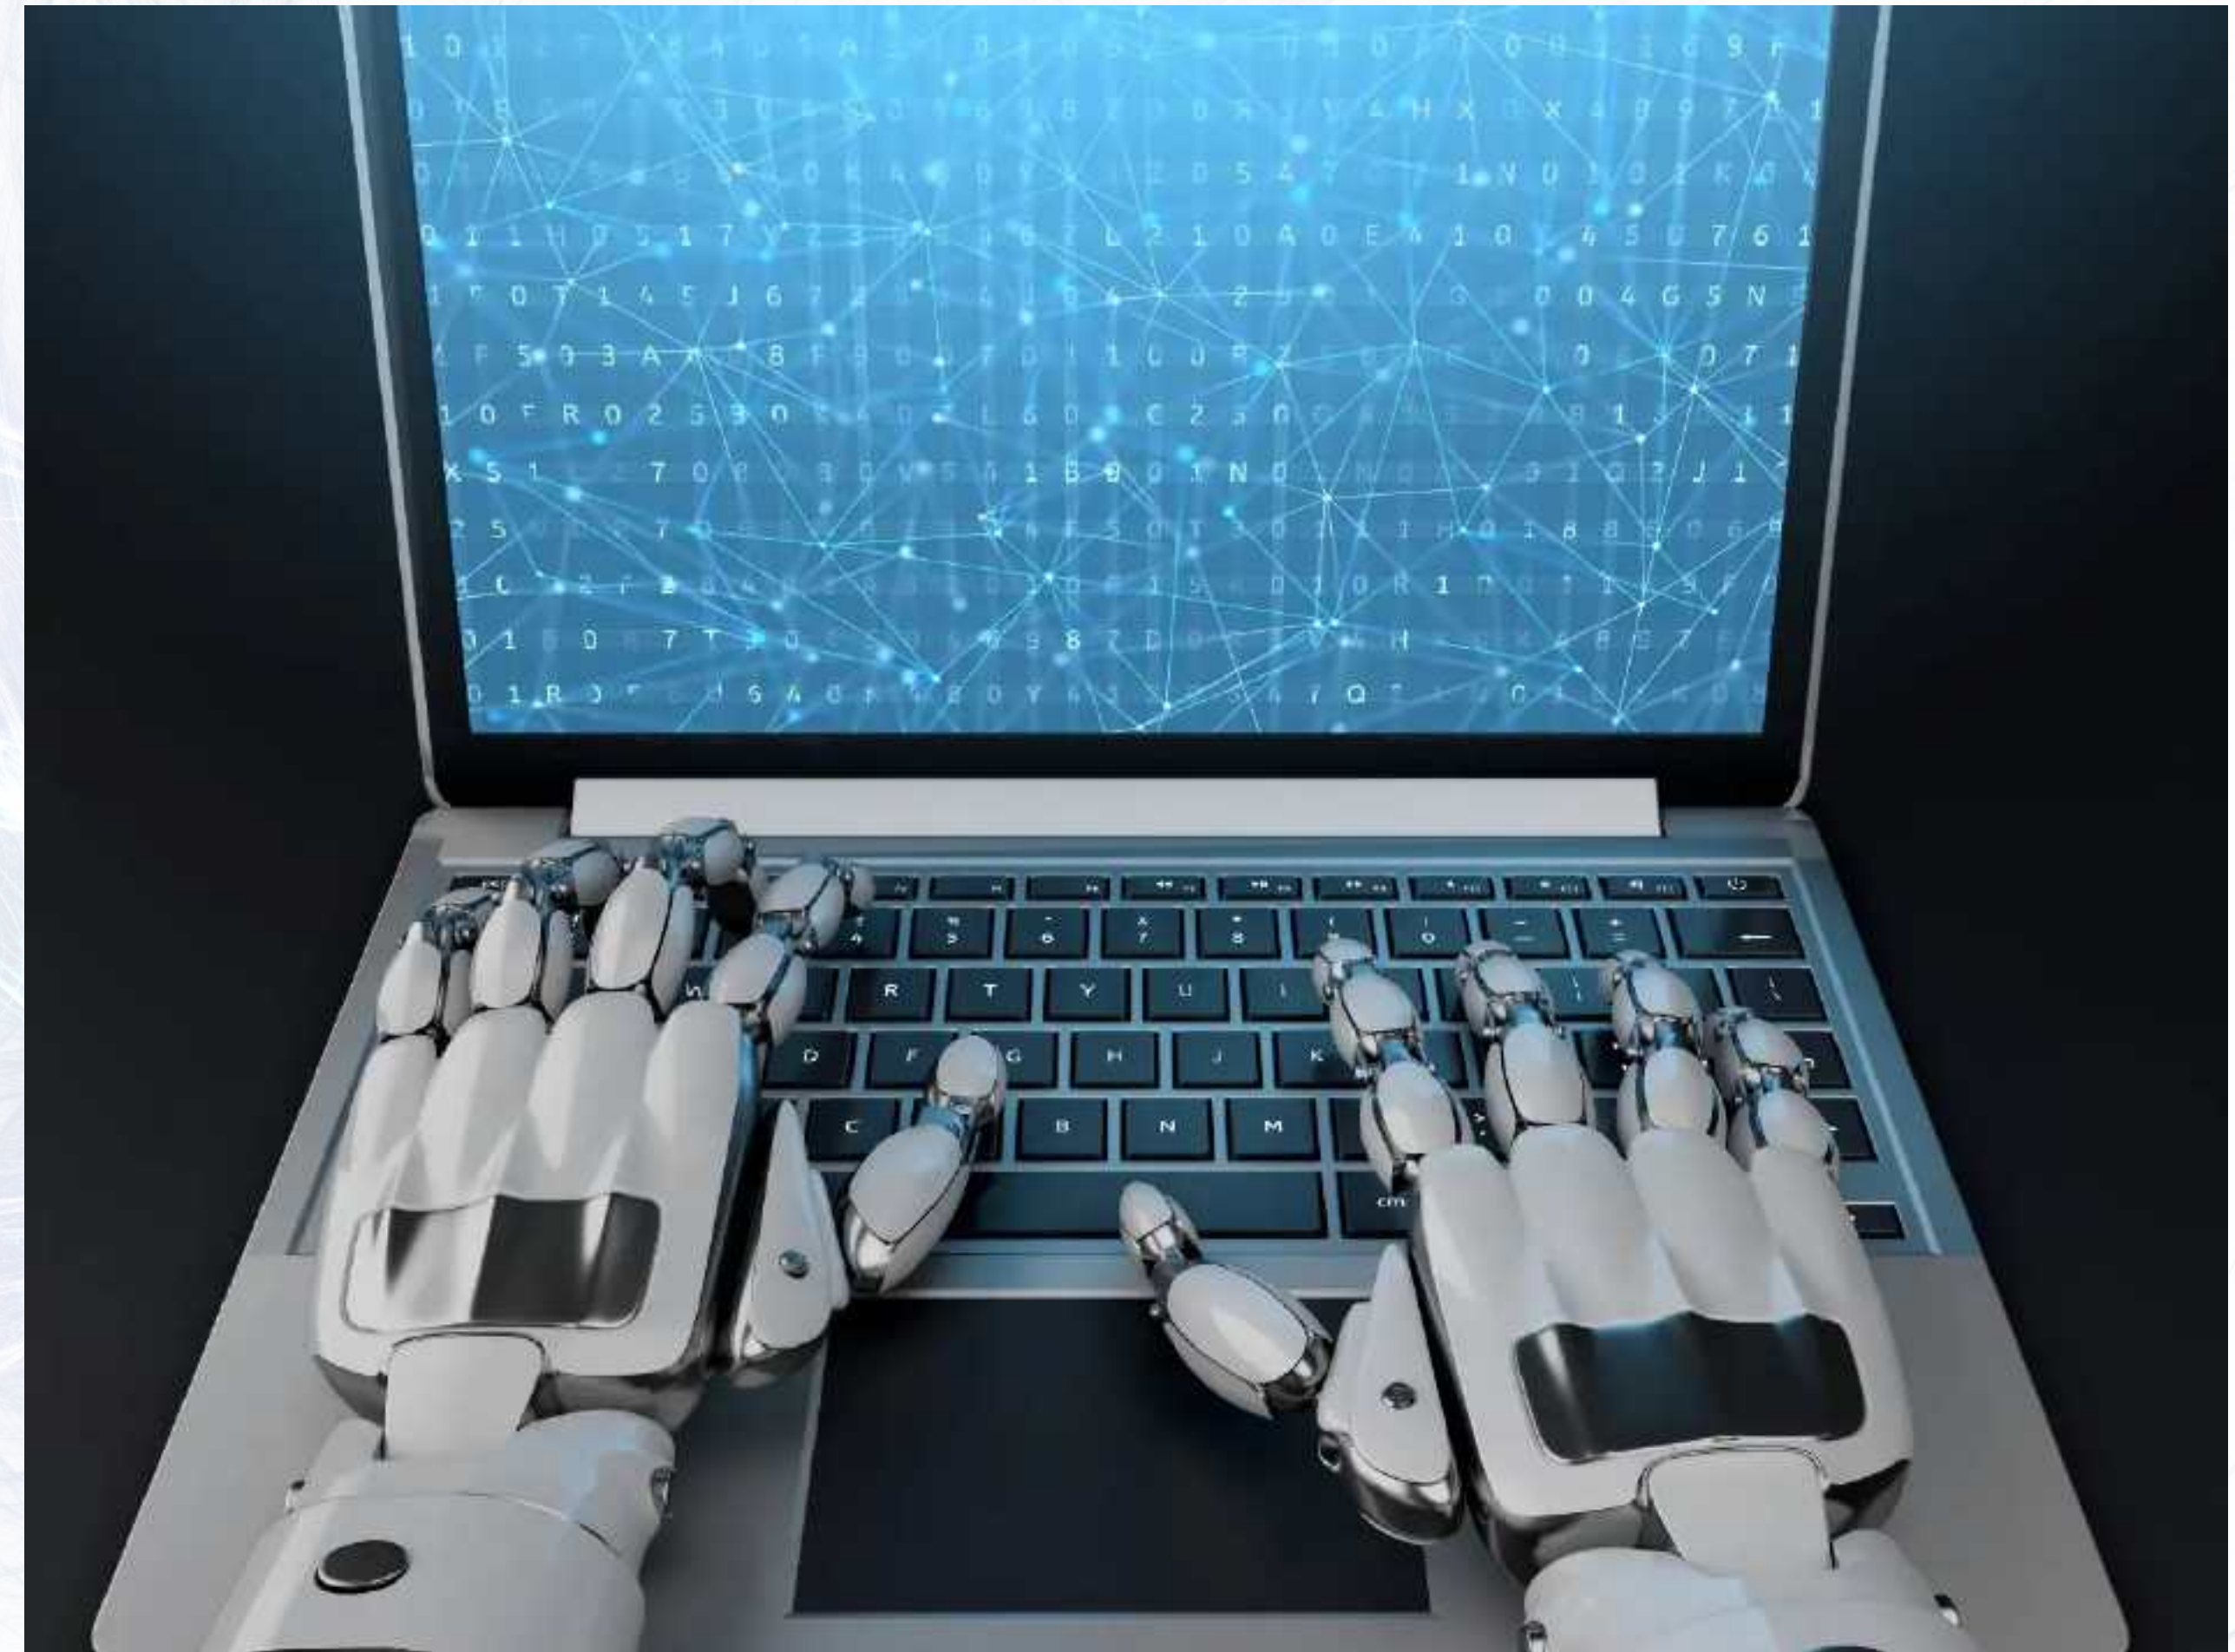
(defn infer [knowledge-base input-facts]
  (let [rules (:rules knowledge-base)
        facts (:facts knowledge-base)]
    (loop [result []
           remaining-rules rules]
      (if (empty? remaining-rules)
          result
          (let [[rule-name rule-premise] (first remaining-rules)
                remaining-rules (rest remaining-rules)]
            (if (every? #(facts %) rule-premise)
                (recur (conj result rule-name) remaining-rules)
                (recur result remaining-rules)))))))

(infer knowledge-base ["fact1" "fact3"])
;; Returns: ["rule1"]
```

This code defines a simple knowledge base as a map with two keys: `:rules` and `:facts`. The `:rules` key contains a map of rule names to lists of facts that must be true for the rule to be triggered. The `:facts` key contains a map of fact names to Boolean values indicating whether the fact is true or false.

The `infer` function takes the knowledge base and a list of input facts as its arguments. It uses a `loop` construct to iterate through the rules in the knowledge base, and for each rule, it checks if all the facts listed in the rule's premise are true using the `every?` function. If all the facts are true, it adds the rule to the list of results. Finally, it returns the list of rules that were triggered by the input facts.

This is a very simple example and can be further developed and expanded as per the requirement.



ChatGPT is a good coder, but not a good software engineer, or developer.

Art Made With Artificial Intelligence Wins at State Fair



Théâtre D'opéra Spatial by Jason Allen

SOMEWHAT STRICTLY CONFIDENTIAL

Poetry to Visual-Art



Ogen wijd gesloten

Jaren hebben open ogen
Mij het zicht ontnomen

Daarom spreek ik met je
In deze diepe nacht

Met mijn ogen gesloten
Spreek zacht de woorden

Zie en ken je beter zo
Om ogen te openen

En je opnieuw te zien
En opnieuw
En opnieuw
En opnieuw

©Jan Veldsink 2019



Two prompts

The prettiest person in the world

The ugliest person in the world

The outcome

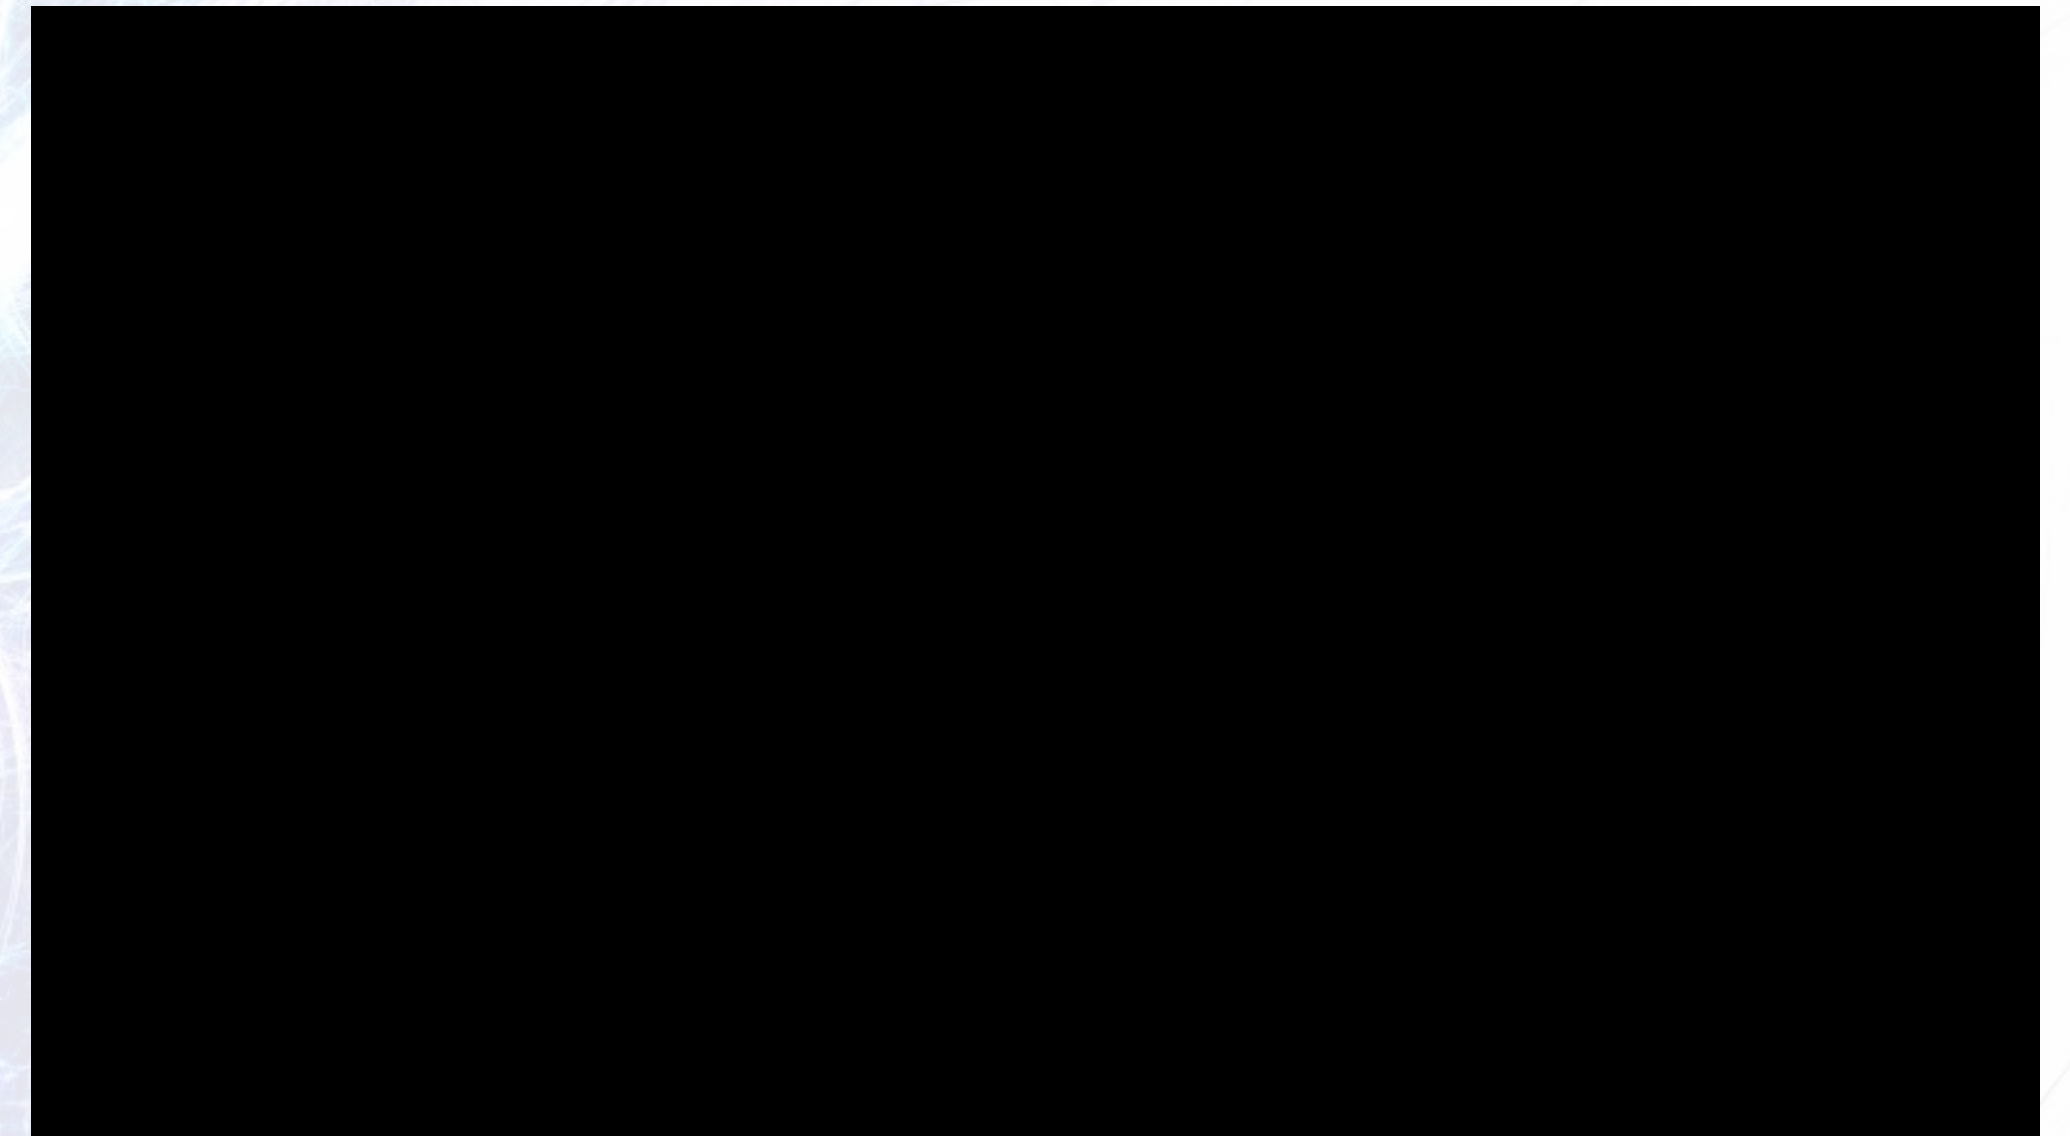
The prettiest person in the world



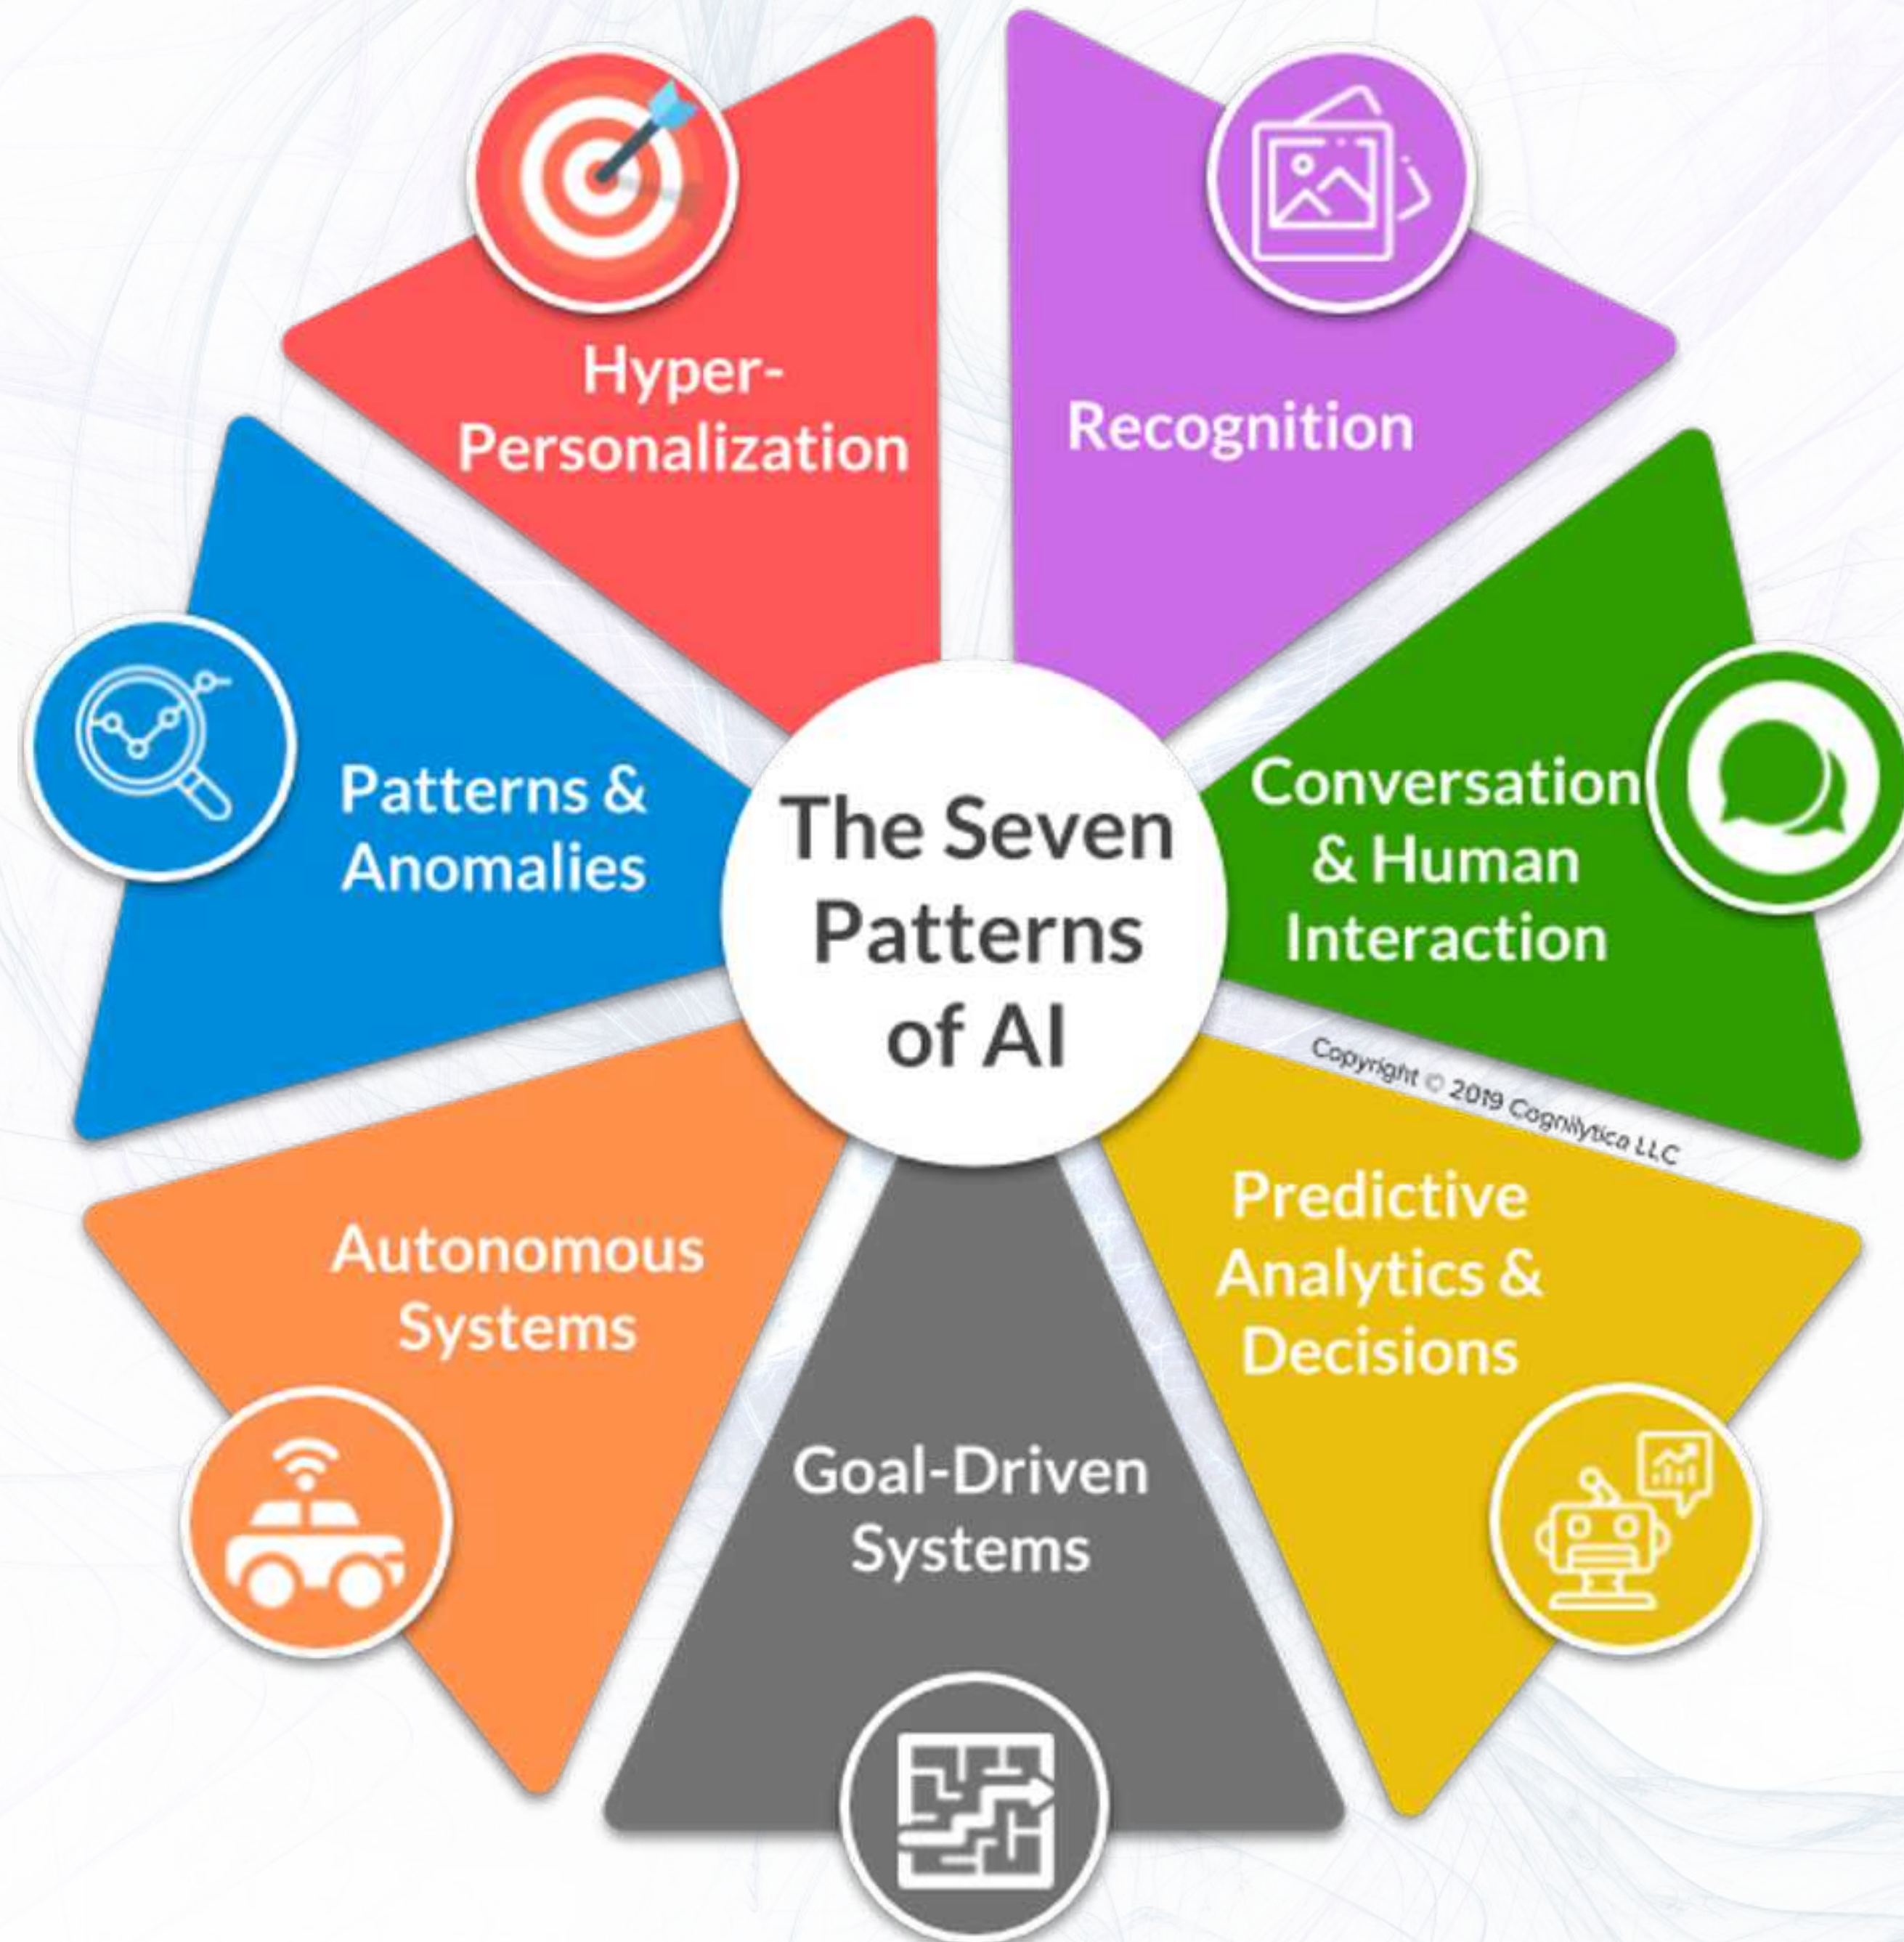
The ugliest person in the world



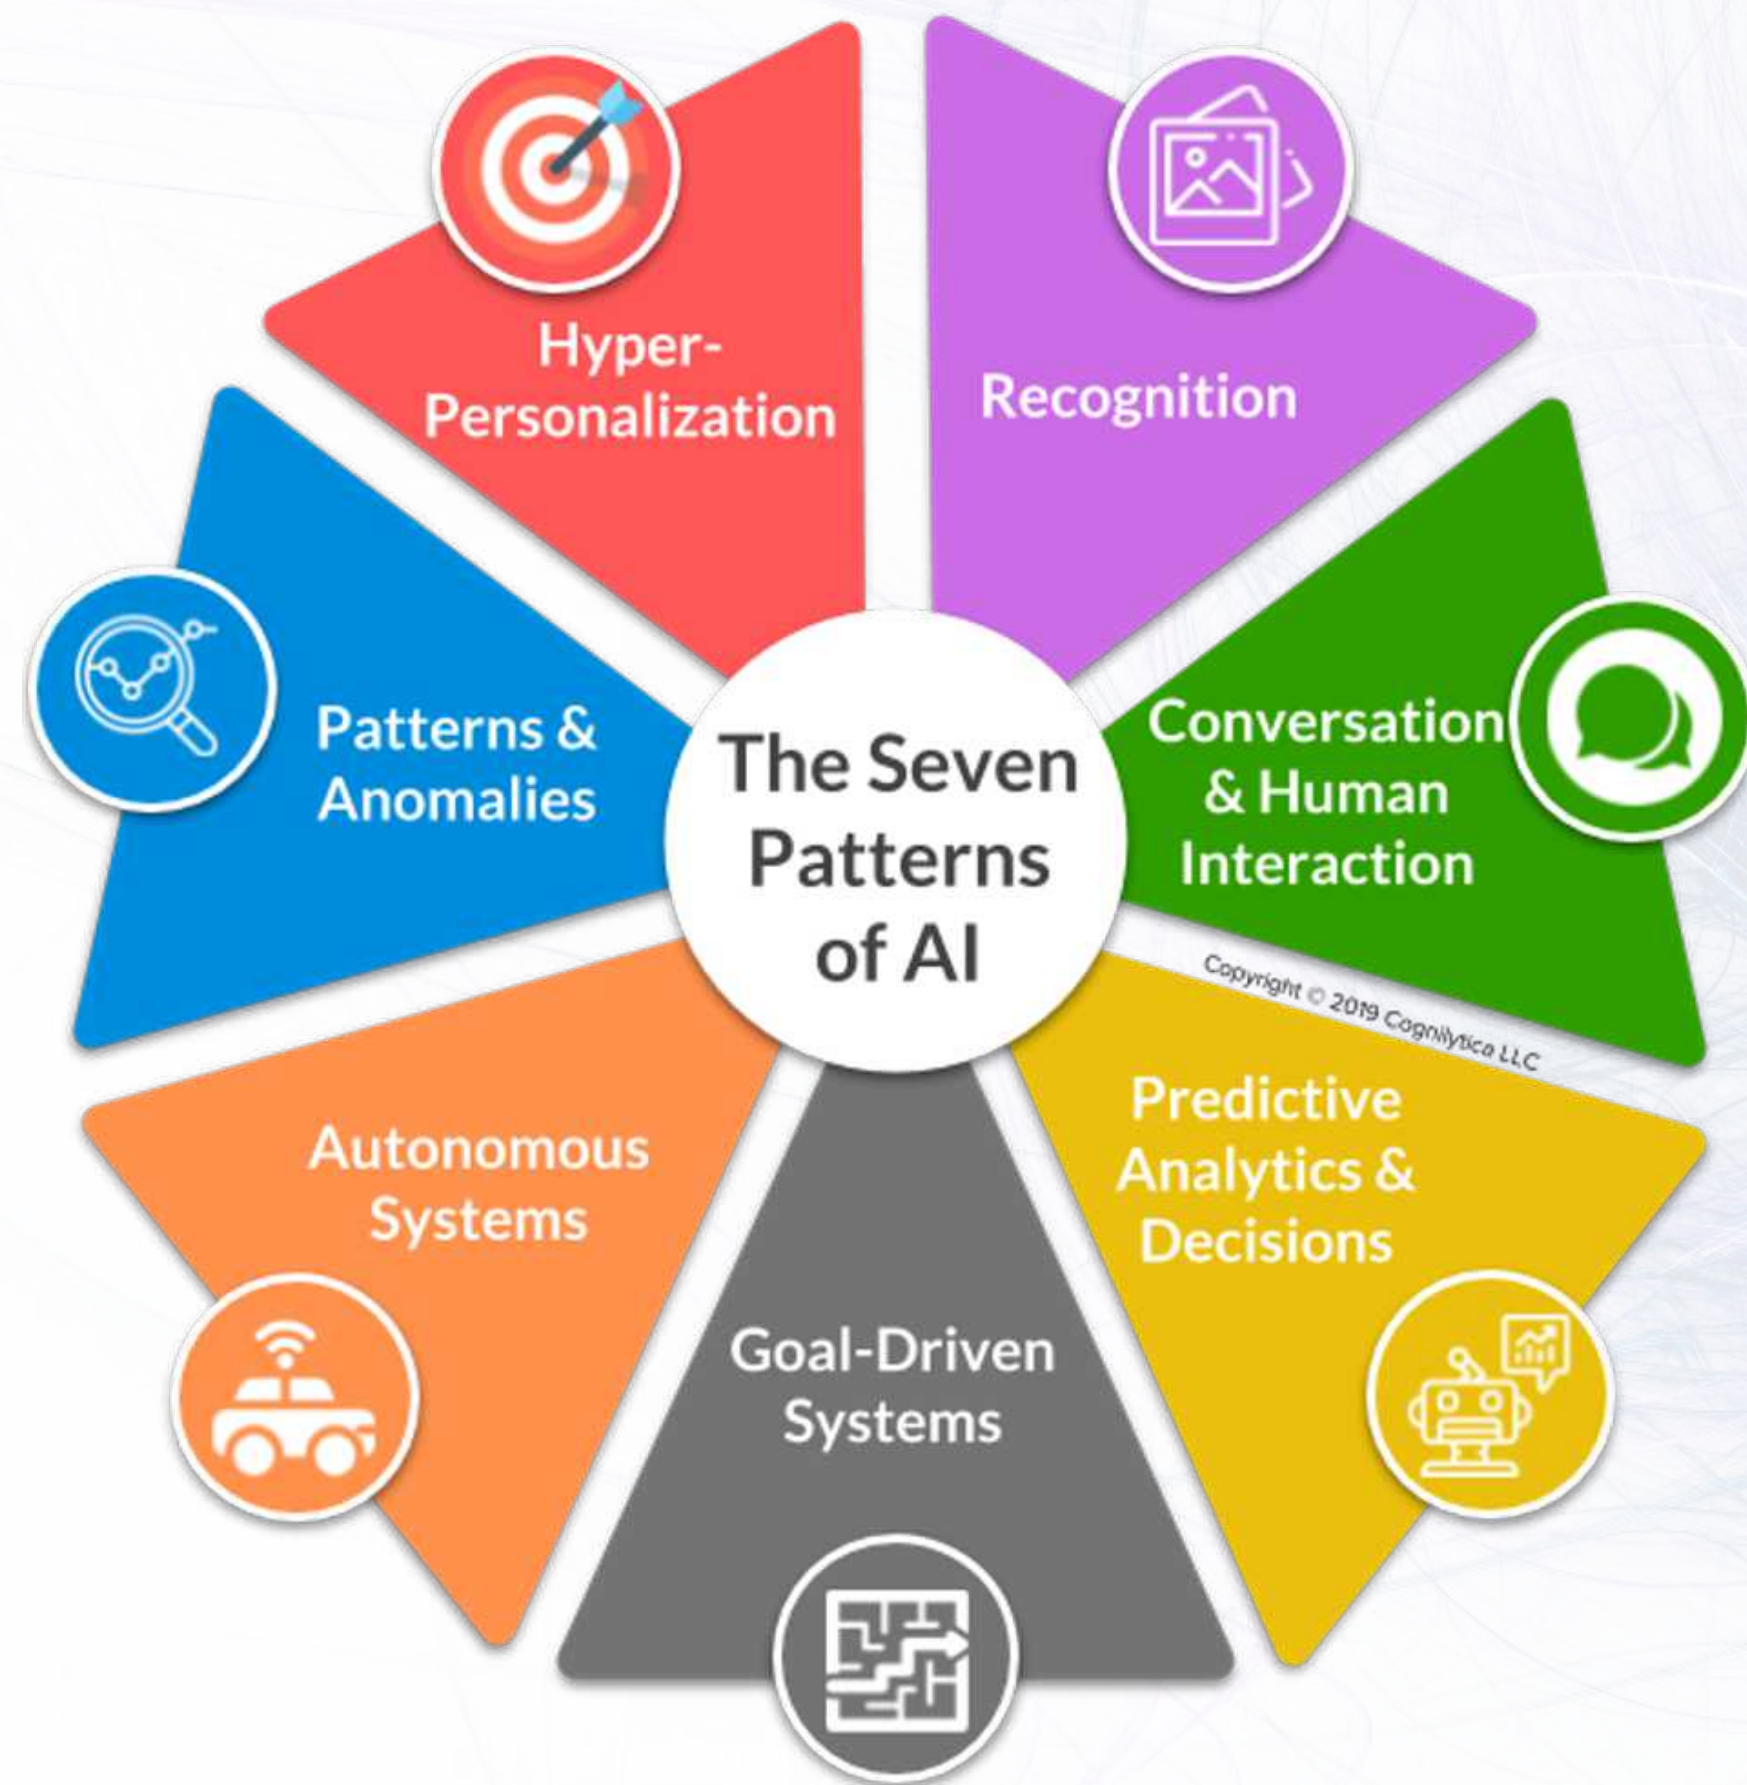
Swiss army knife - Metaphor



<https://www.youtube.com/watch?v=O20mKRHYttg>



The Hyperpersonalization Pattern



- Treat each customer as an individual.
 - The hyperpersonalization pattern is defined as using machine learning to develop a profile of each individual, and then having that profile learn and adapt over time for a wide variety of purposes including displaying relevant content, recommend relevant products, provide personalized recommendations and so on.
- The objective of this pattern is to treat each individual as an individual.

Just ask Notion AI.

Knowledge, answers, ideas. One click away.

Get started

Sign up for Q&A Waitlist



Get answers

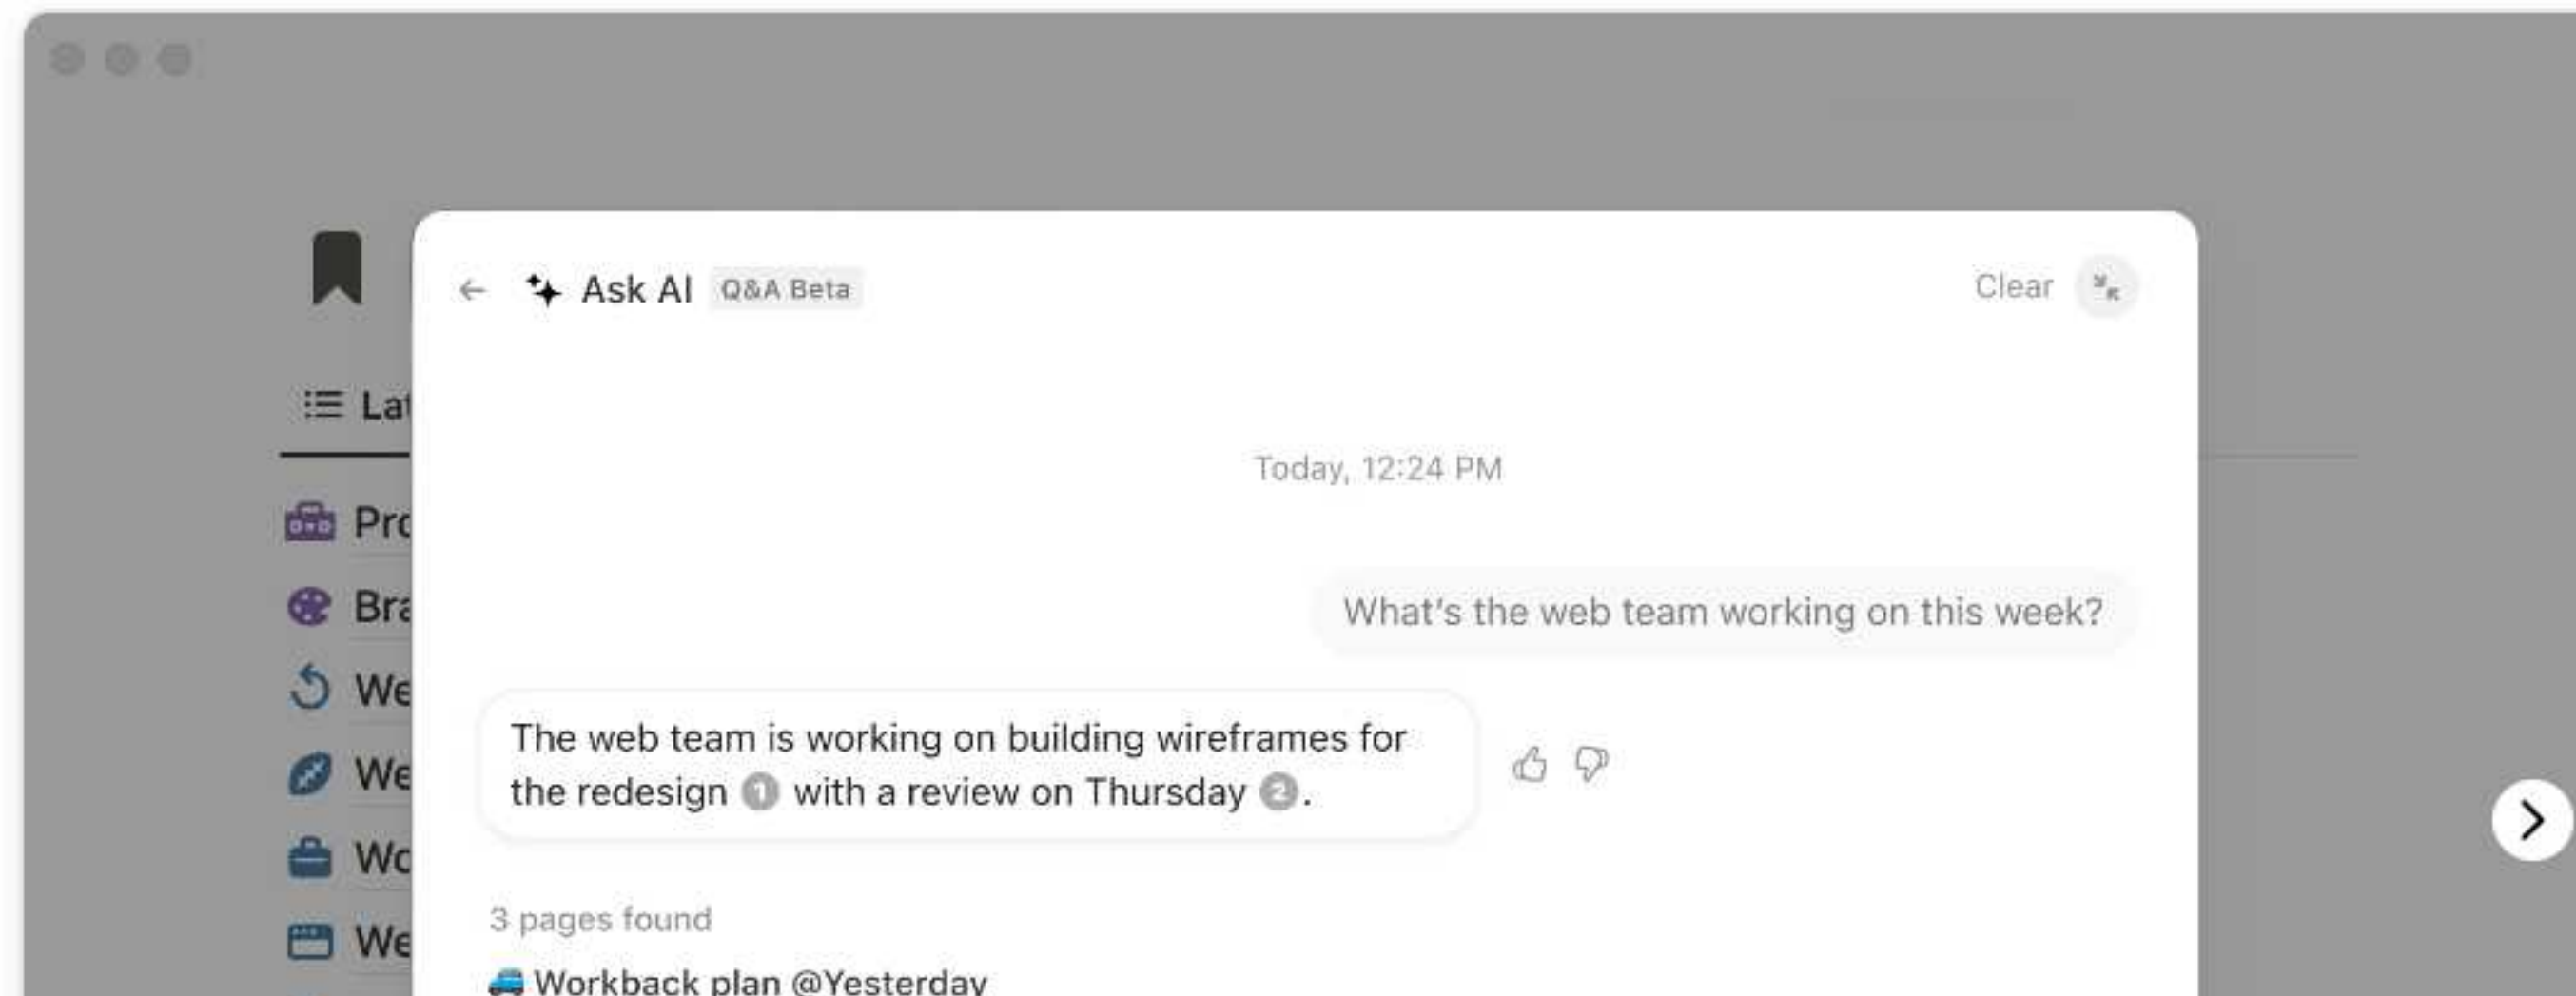
Just ask Q&A, and find the info you need in seconds.

Write better

Get help writing and brainstorming in Notion, not in a separate browser tab.

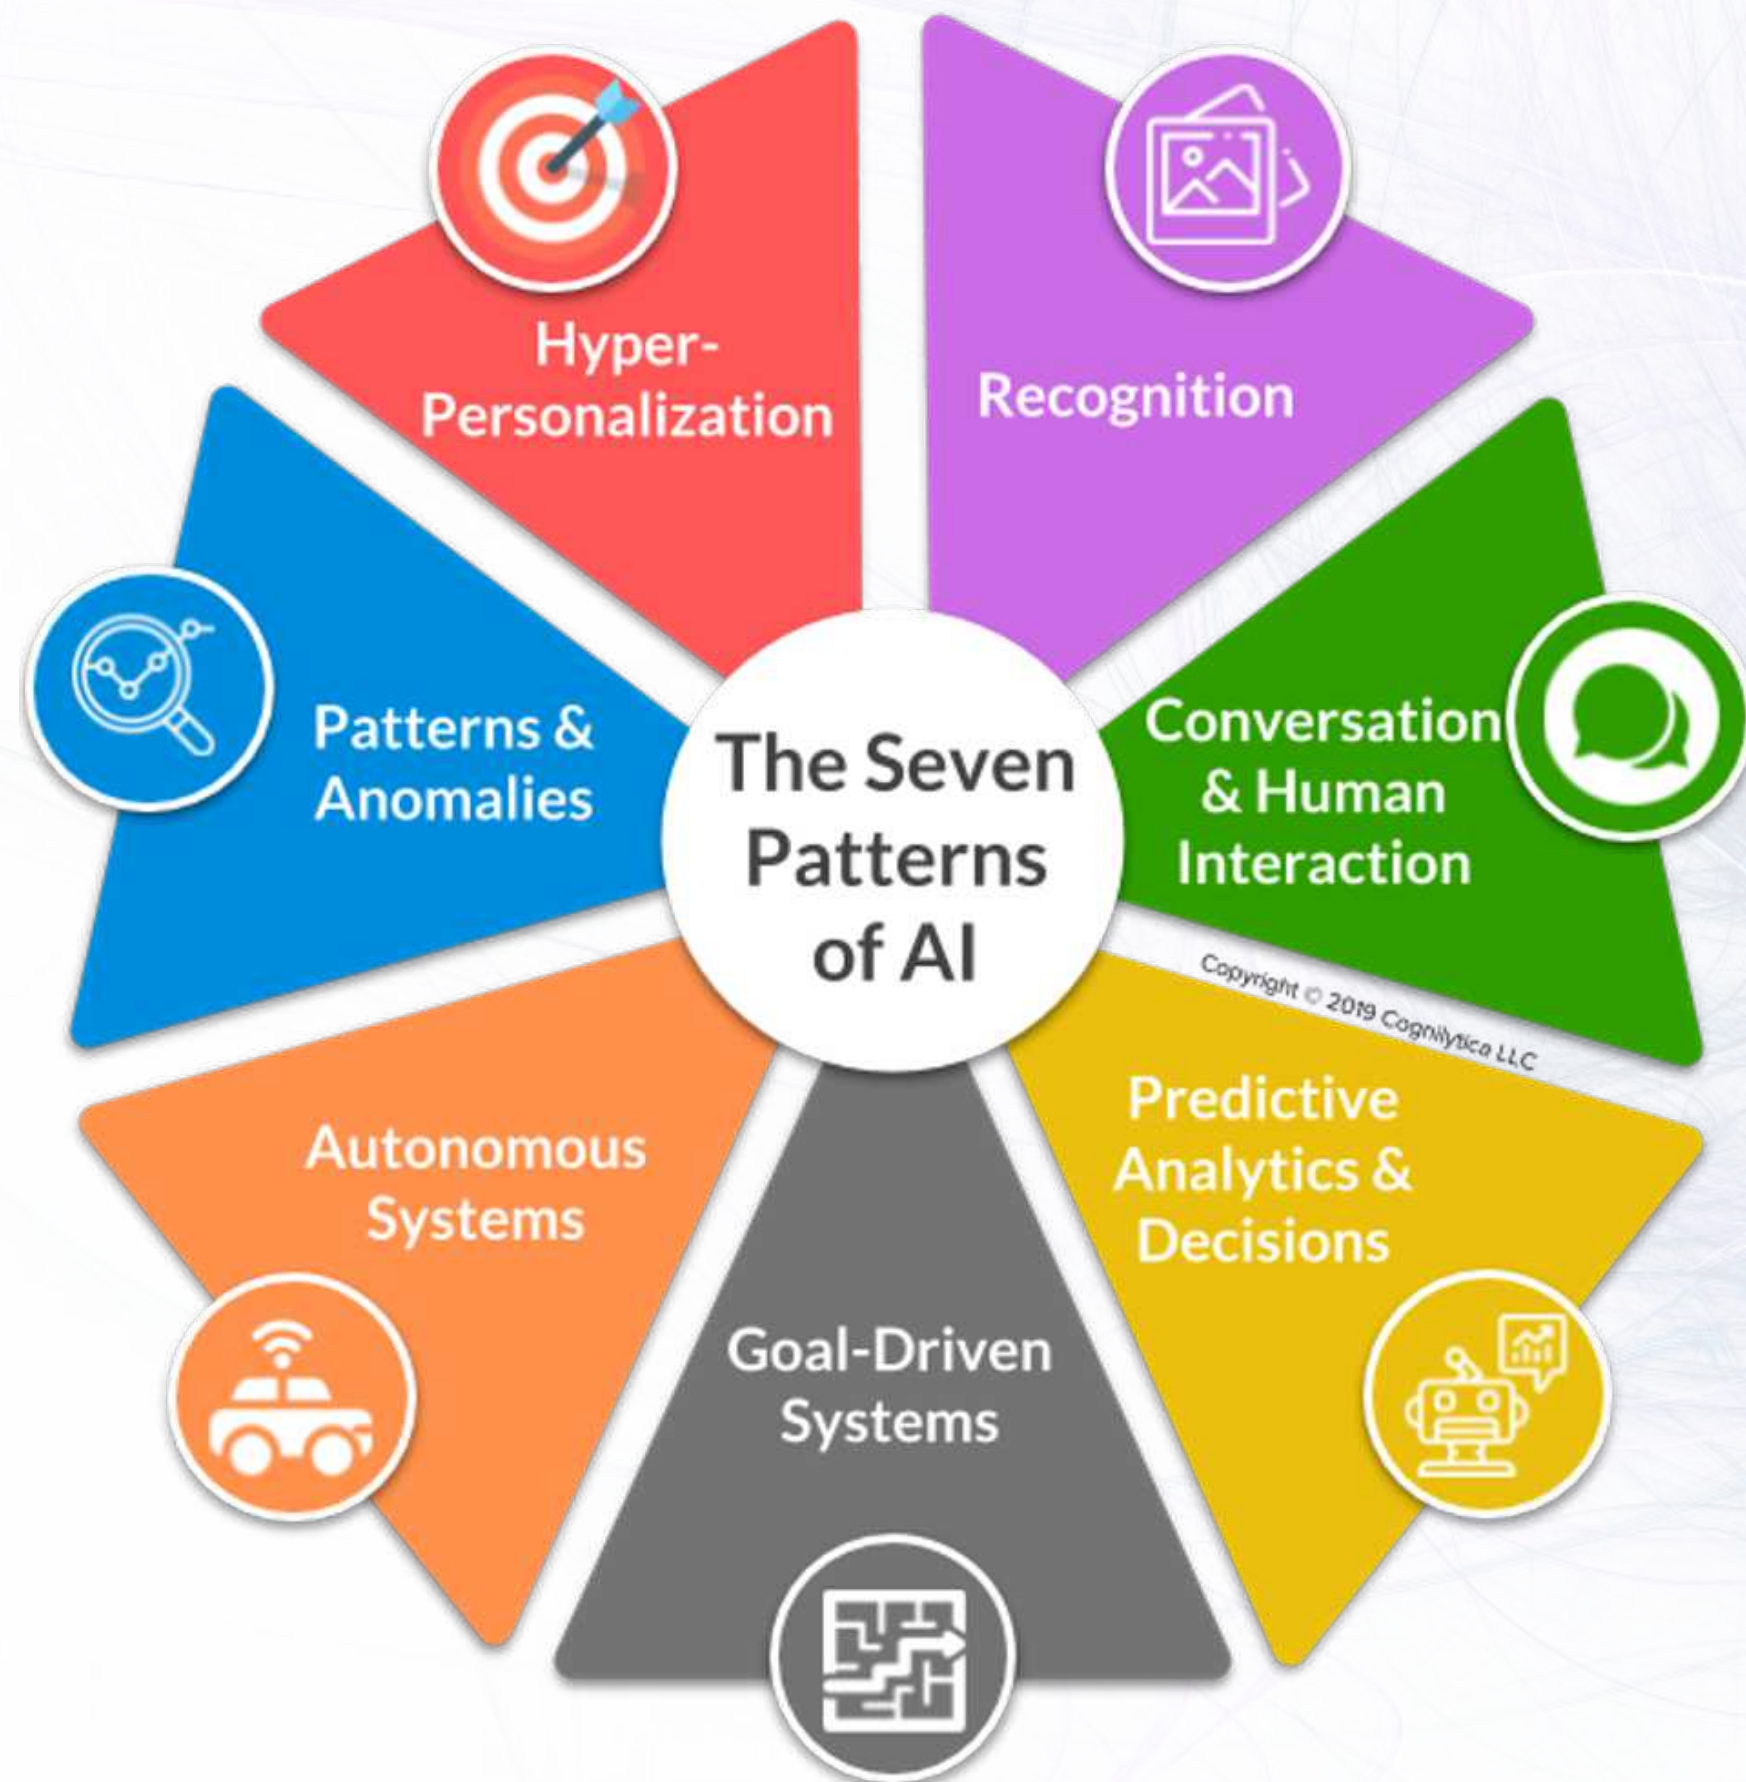
Autofill tables

Turn overwhelming data into clear, actionable information in seconds.



The Recognition Pattern

- Machines that can recognize the world
- One of the big advancements in machine learning is the use of deep learning to greatly improve the accuracy of recognition-related tasks such as image, video, audio, and object recognition, classification, and identification. The recognition pattern is defined as using machine learning and other cognitive approaches to identify and determine objects or other desired things to be identified within image, video, audio, text, or other primarily unstructured data formally. The objective of this pattern is to have machines identify and understand things.





Replicated Meat - Coming Soon to Major Supermarkets Near You
yahoo.com



Utku Gen
Twitter is blocked in Turkey. On the streets of Istanbul, the action against censorship is graffiti DNS addresses.
pic.twitter.com/VXcaN7LvS

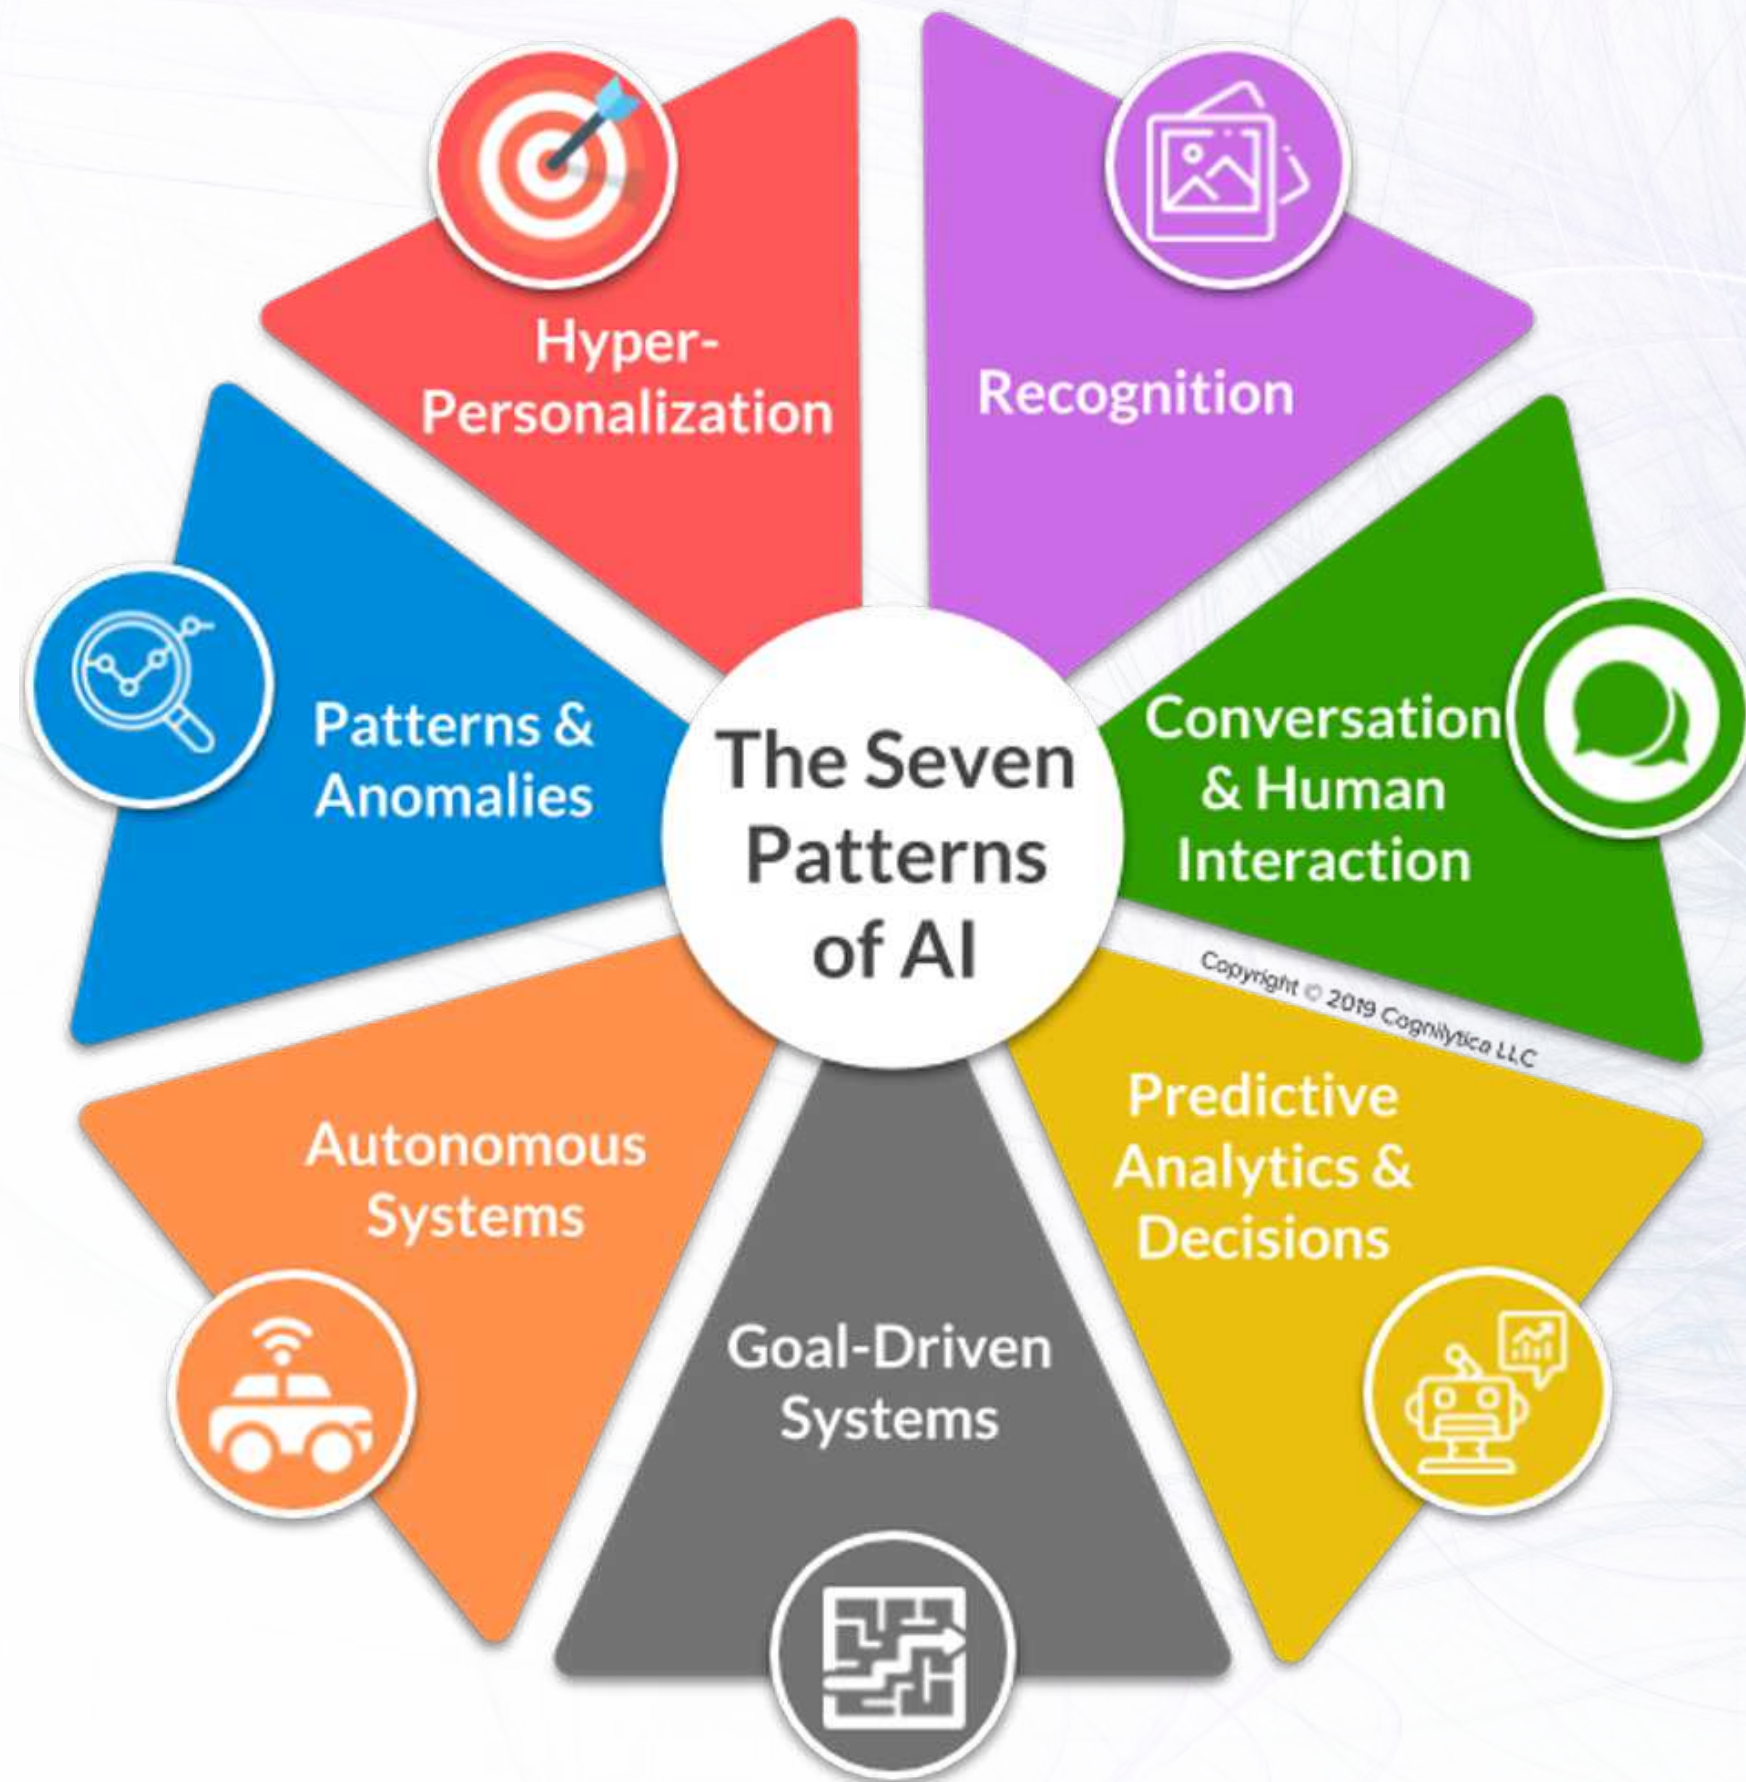


QuiQui Drone Delivery Is Coming

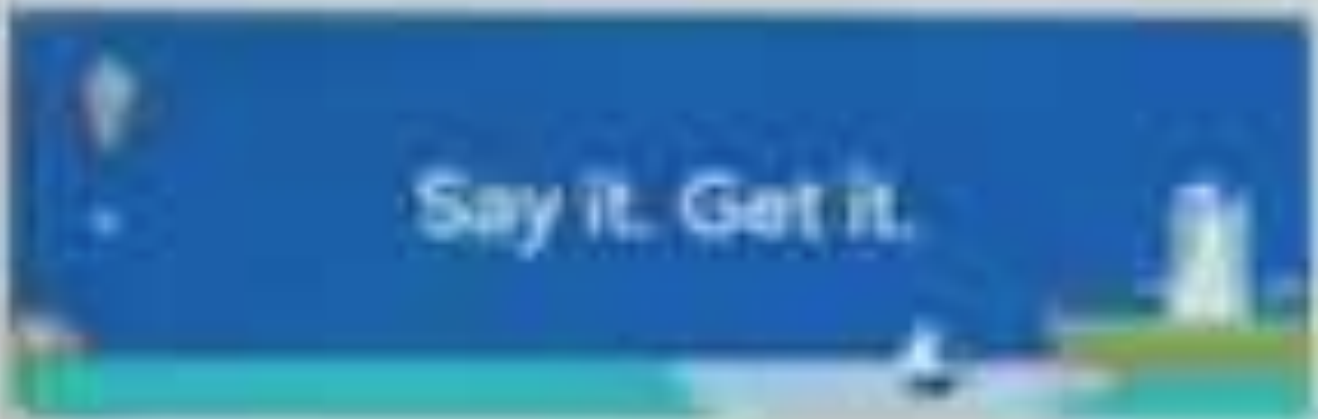
Steering wheel can 'detect anger'
bbc.com



The Conversational Pattern



- Machines that can communicate as humans do
- Another pattern of AI is the conversational/human interaction pattern. This is defined as machines and humans interacting with each other through conversational forms of interaction and content across a variety of methods including voice, text, and image forms. This includes machine to human, human to machine, and back and forth human and machine interaction. The objective of this pattern is enabling machines to interact with humans how humans interact with each other.



News
Stay up to date with the latest news from the government and your local area.



Local Search
Find local businesses, services and more.

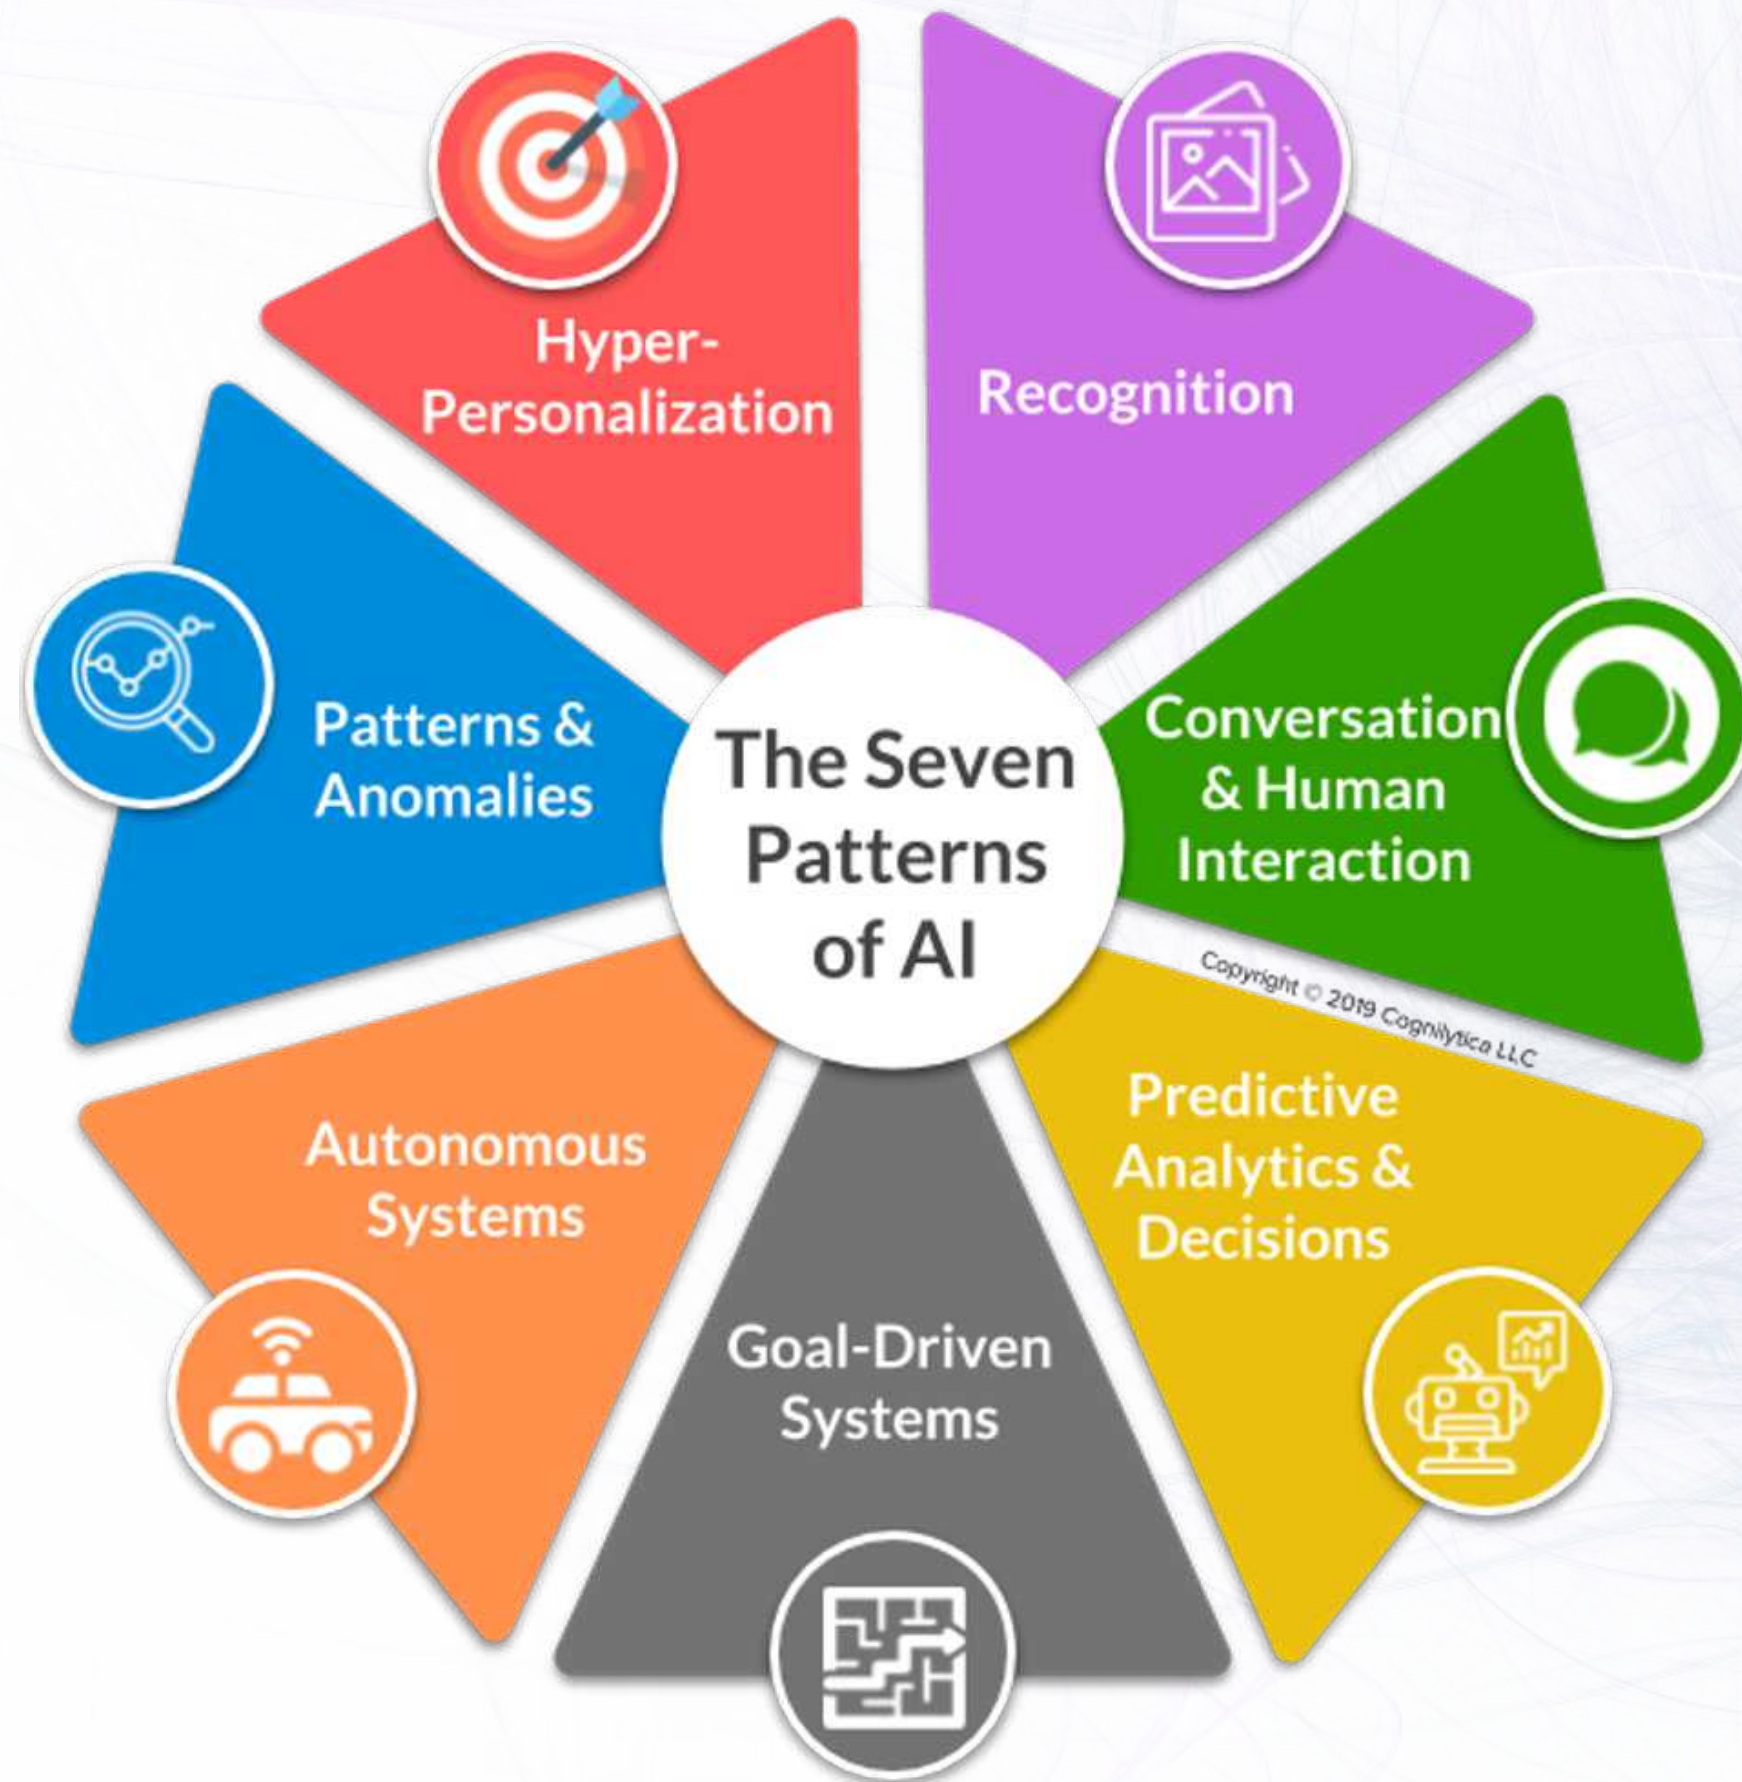


Weather
Get the latest weather forecast for your area.

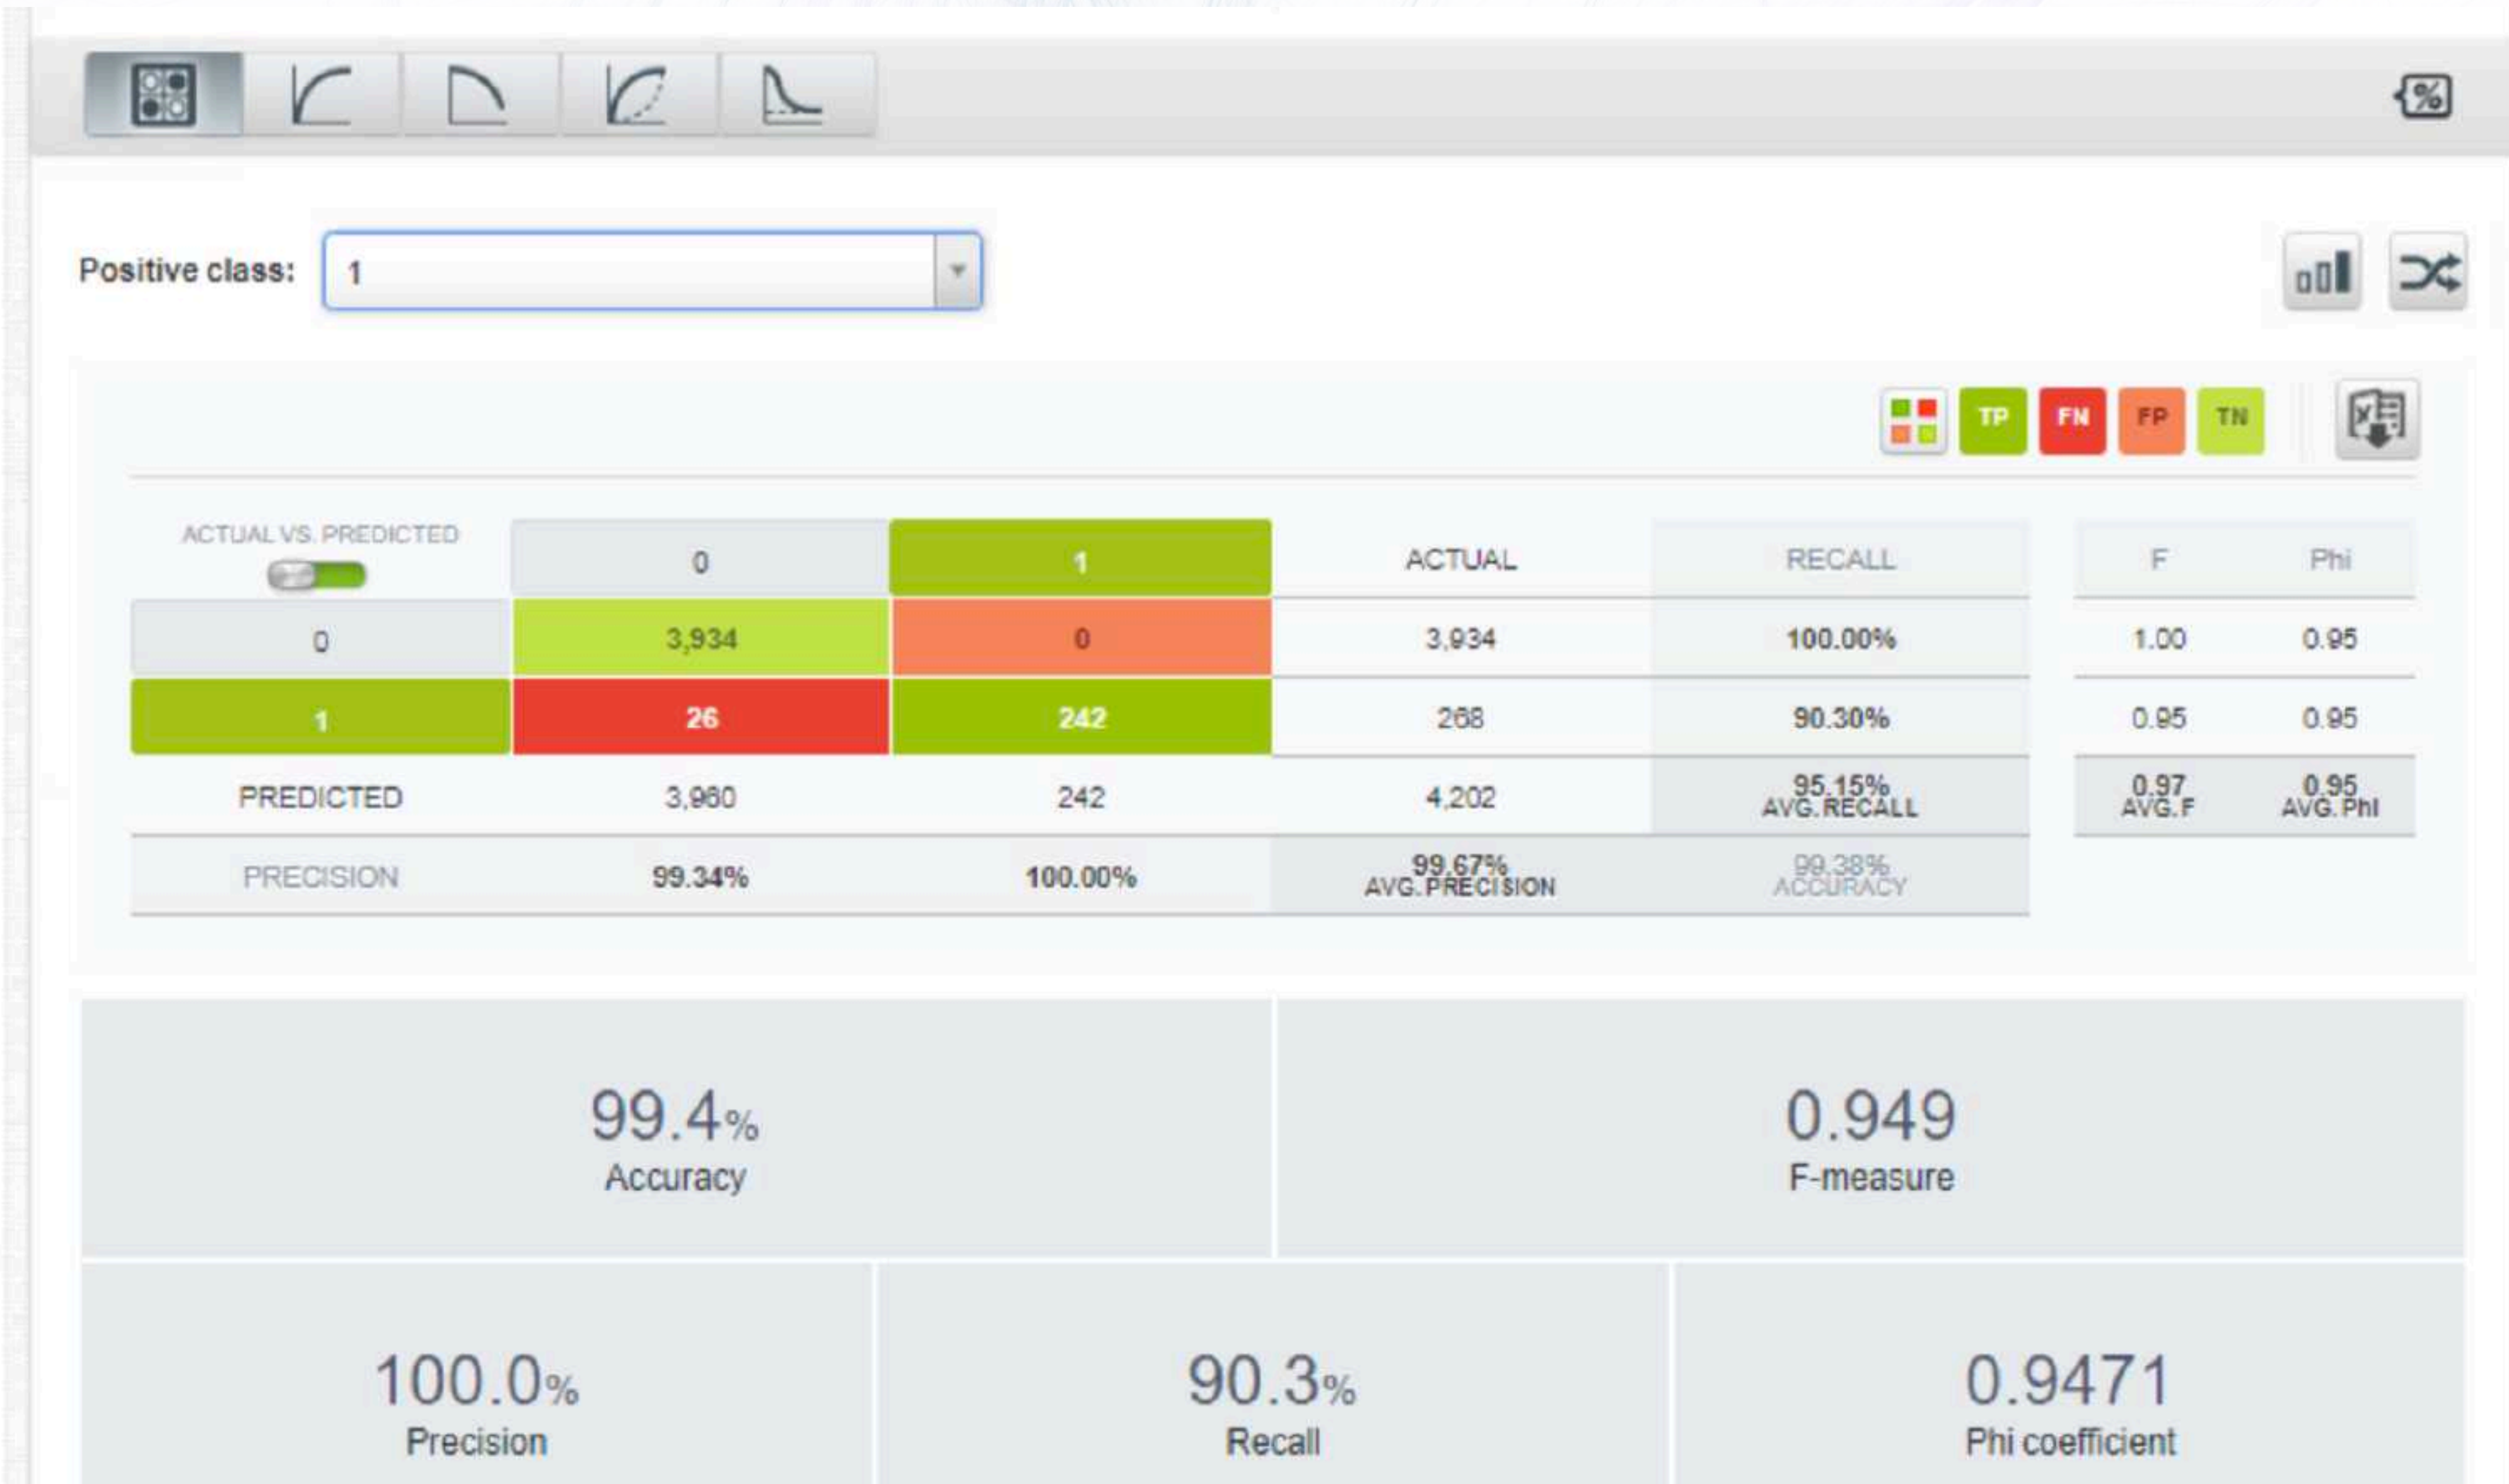


AI powered predictive analytics

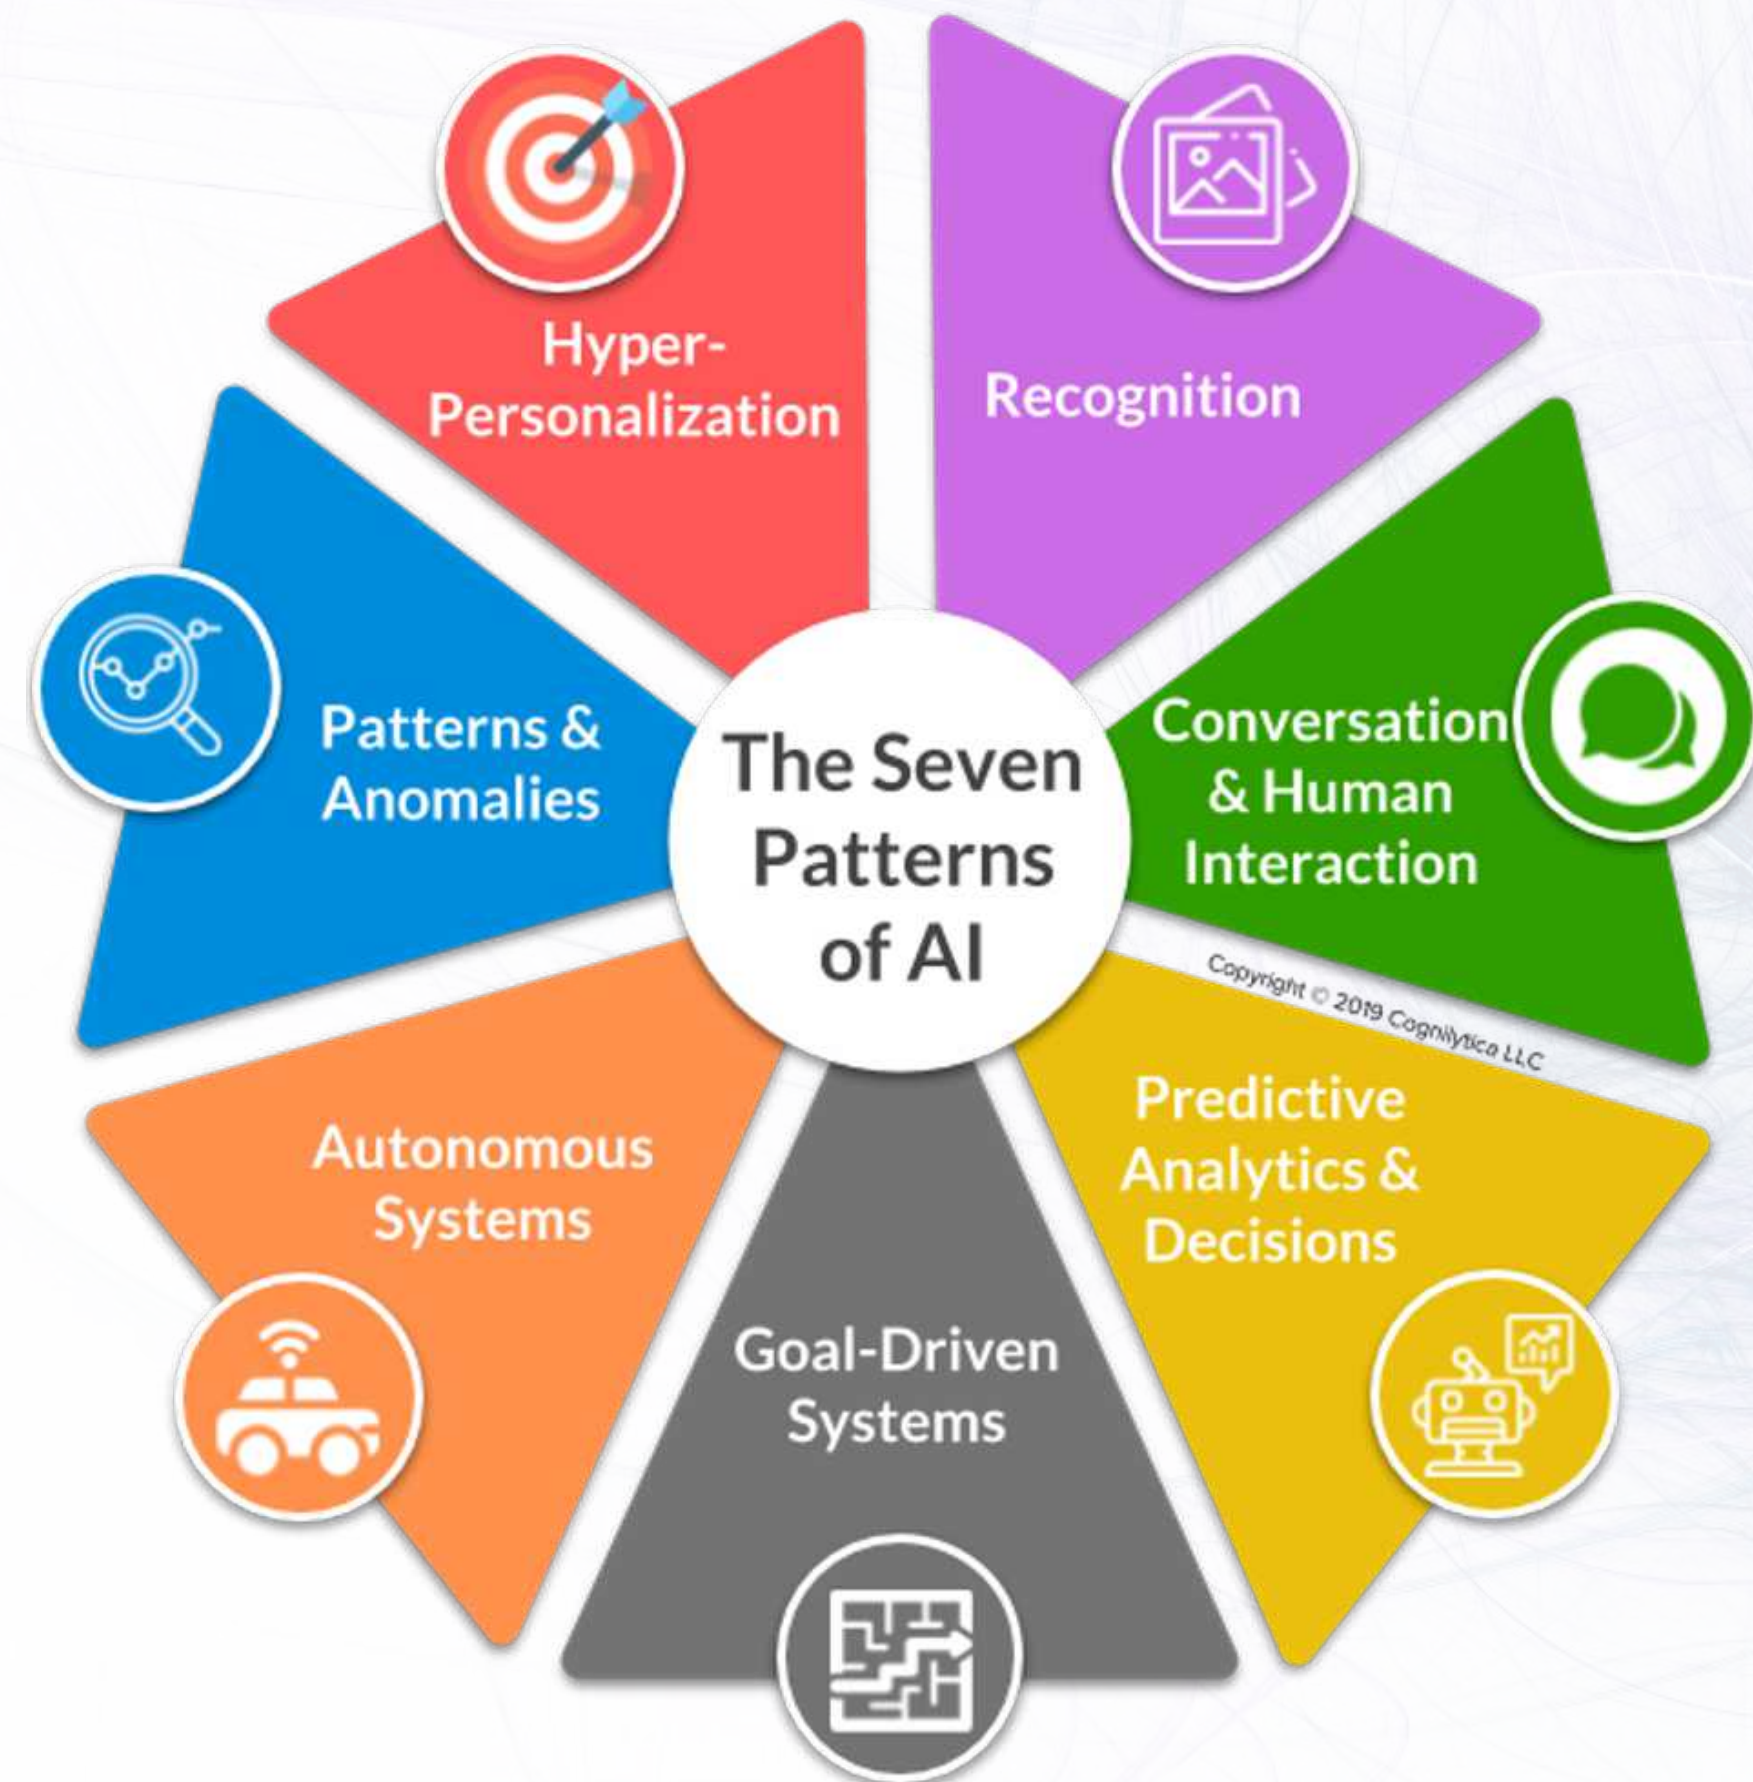
- Machines that decides like humans do.
- Another pattern of AI is predictive analytics and decision support. This is defined as using machine learning and other cognitive approaches to understand how past or existing behaviors can help predict future outcomes or help humans make decisions about future outcomes based on these patterns. The objective of this pattern is helping humans make better decisions.



ML Results on fraud



The Goal-Driven Systems Pattern



- Solving the Puzzle
- Machines have proven to be particularly adept at learning the rules of games, and beating humans at their own games. In the past, machines have easily conquered the games of checkers, chess, and finding solutions to mazes. Through the power of reinforcement learning and much more advanced computation capabilities, machines are now able to win at Go, multi-player games such as DoTA, and much more complicated games. Alpha Go and Alpha Zero were created by Google's DeepMind division under the theory that through goals, computers could learn anything through game play. Games are just the beginning to solutions that could potentially even lead to breakthroughs in solving long-hoped for goals in Artificial General Intelligence (AGI).



ima...

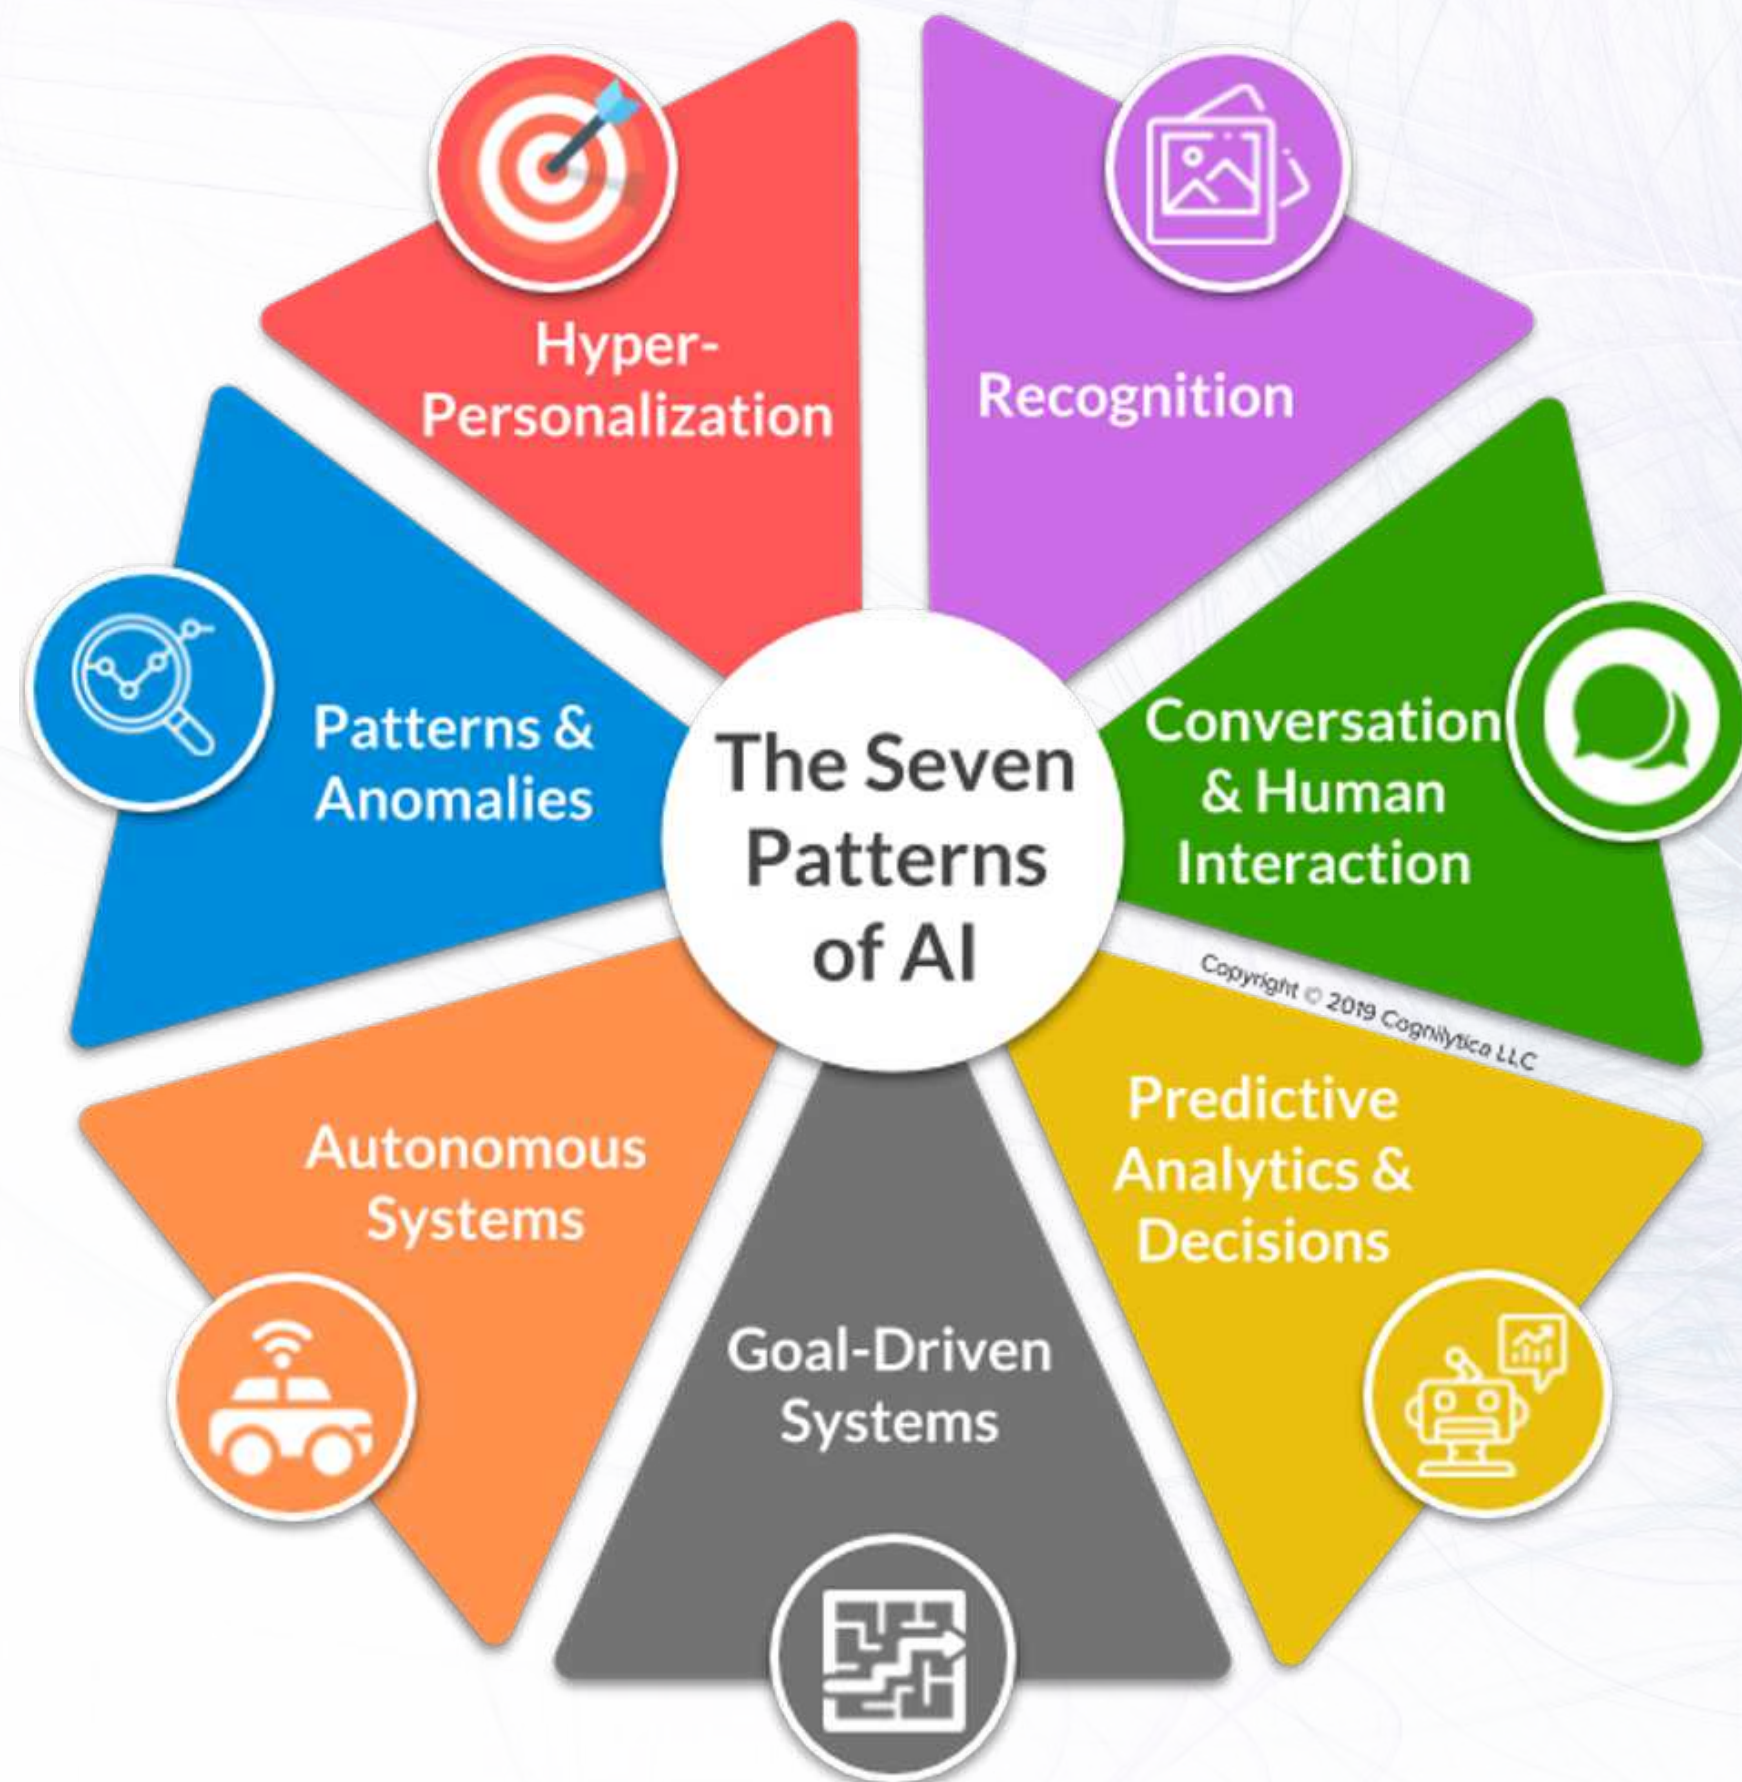


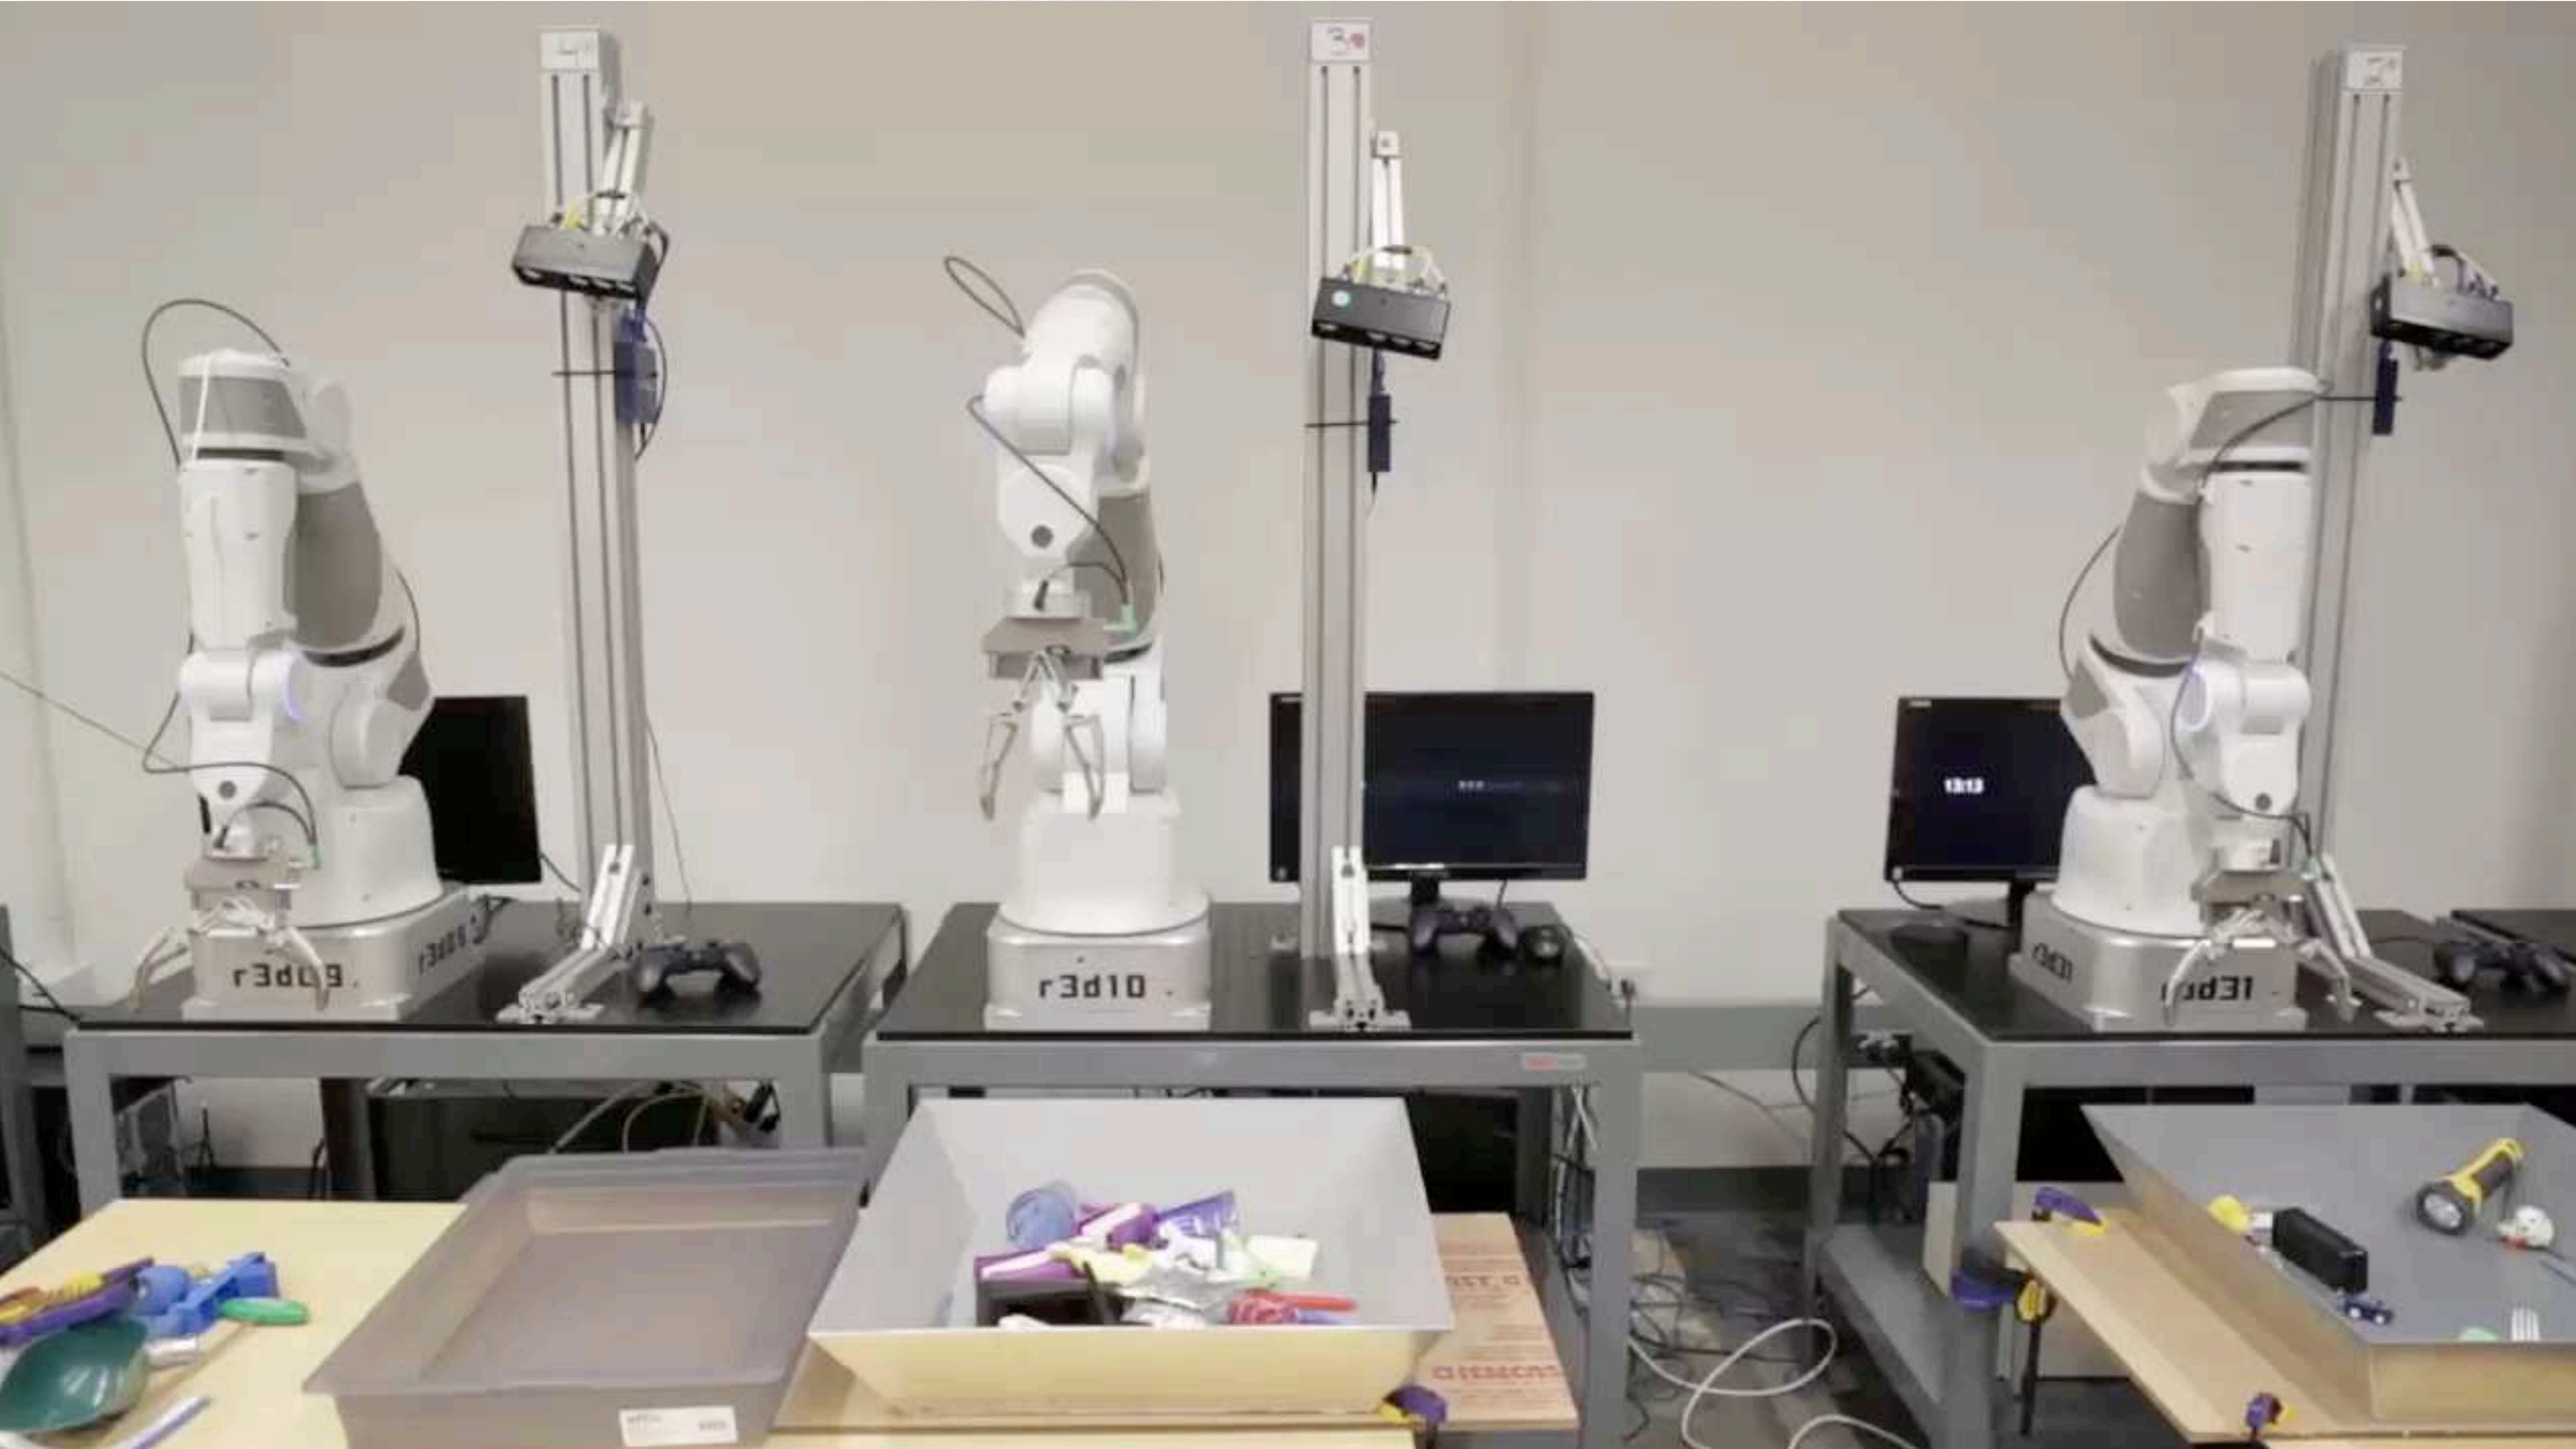
032 3 1



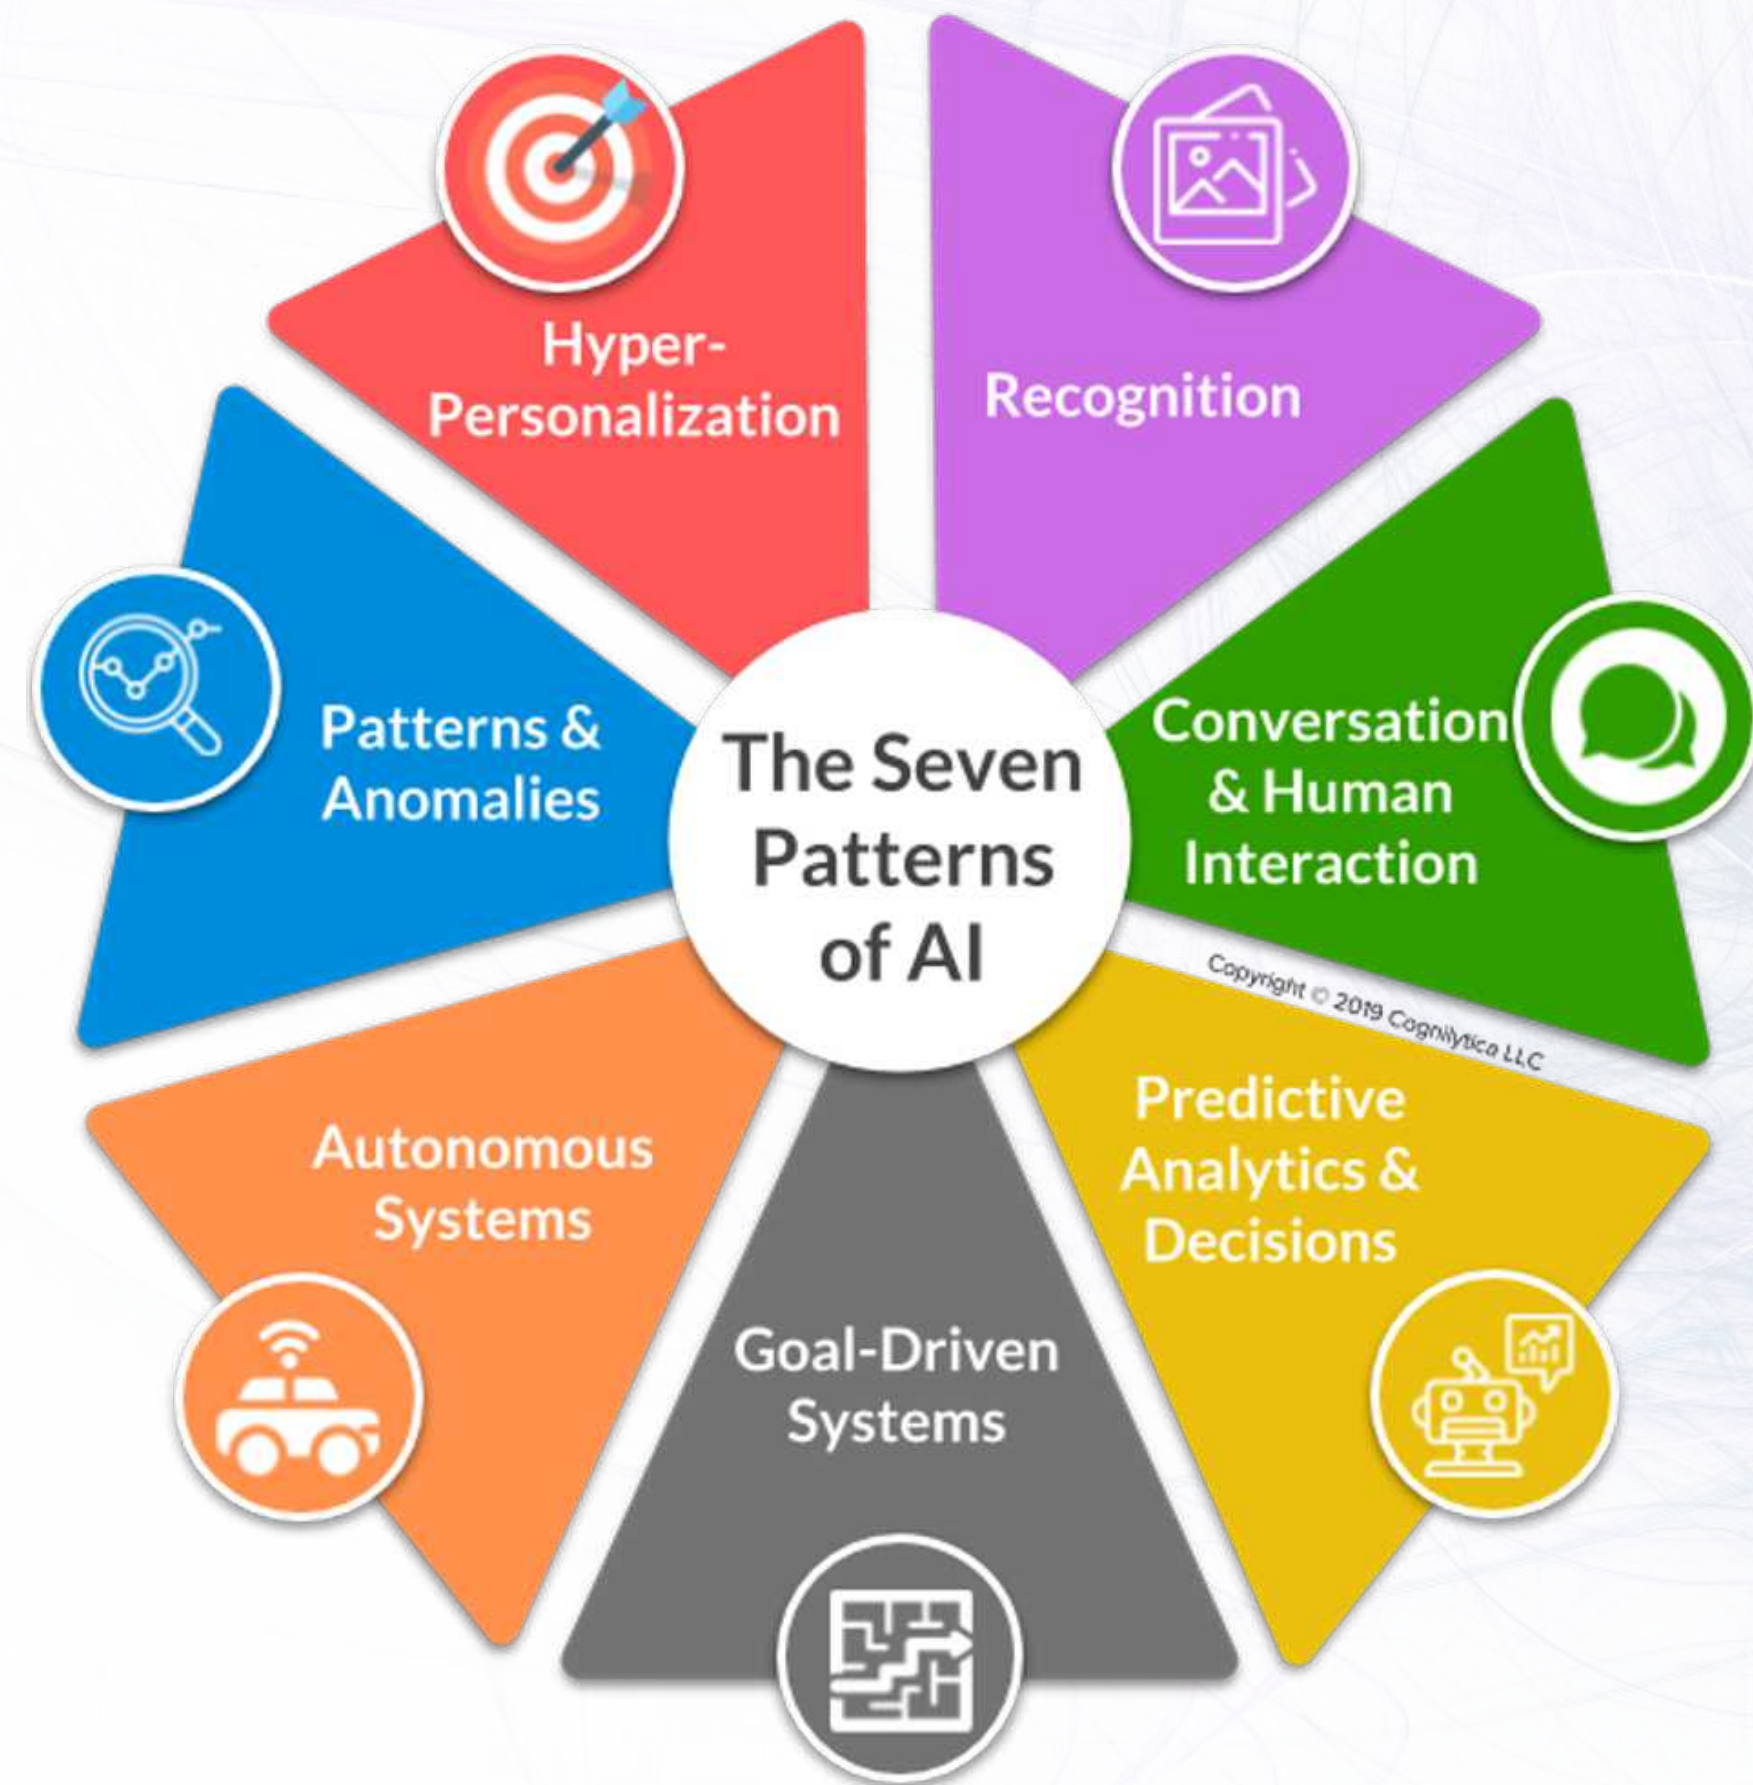
Autonomous systems Pattern

- Reducing the need for manual labor.
 - Autonomous systems are physical and virtual software and hardware systems that are able to accomplish a task, reach a goal, interact with their surroundings, and achieve an objective with minimal human involvement.
- Where the primary objective of hyper-personalization is to treat people as individuals, the goal of autonomous systems is to streamline things with as little human interaction as possible.
- The autonomous pattern requires machine learning capability that can independently perceive the outside world, predict the future behavior of external elements, and plan for how to deal with those changes.





Identifying Patterns and anomalies with AI



- Machines that find unknown patterns
- Machine learning is particularly good at identifying patterns and finding anomalies or outliers. The “pattern-matching pattern” is one of the repeating approaches to AI projects that has seen wide and increasing adoption. The goal of the Patterns and Anomalies pattern of AI is to use machine learning and other cognitive approaches to learn patterns in the data and learn higher order connections between data points to see if it fits an existing pattern or if it is an outlier or anomaly. The object of this pattern is to find what fits with existing data and what doesn't.

Anomaly Detection

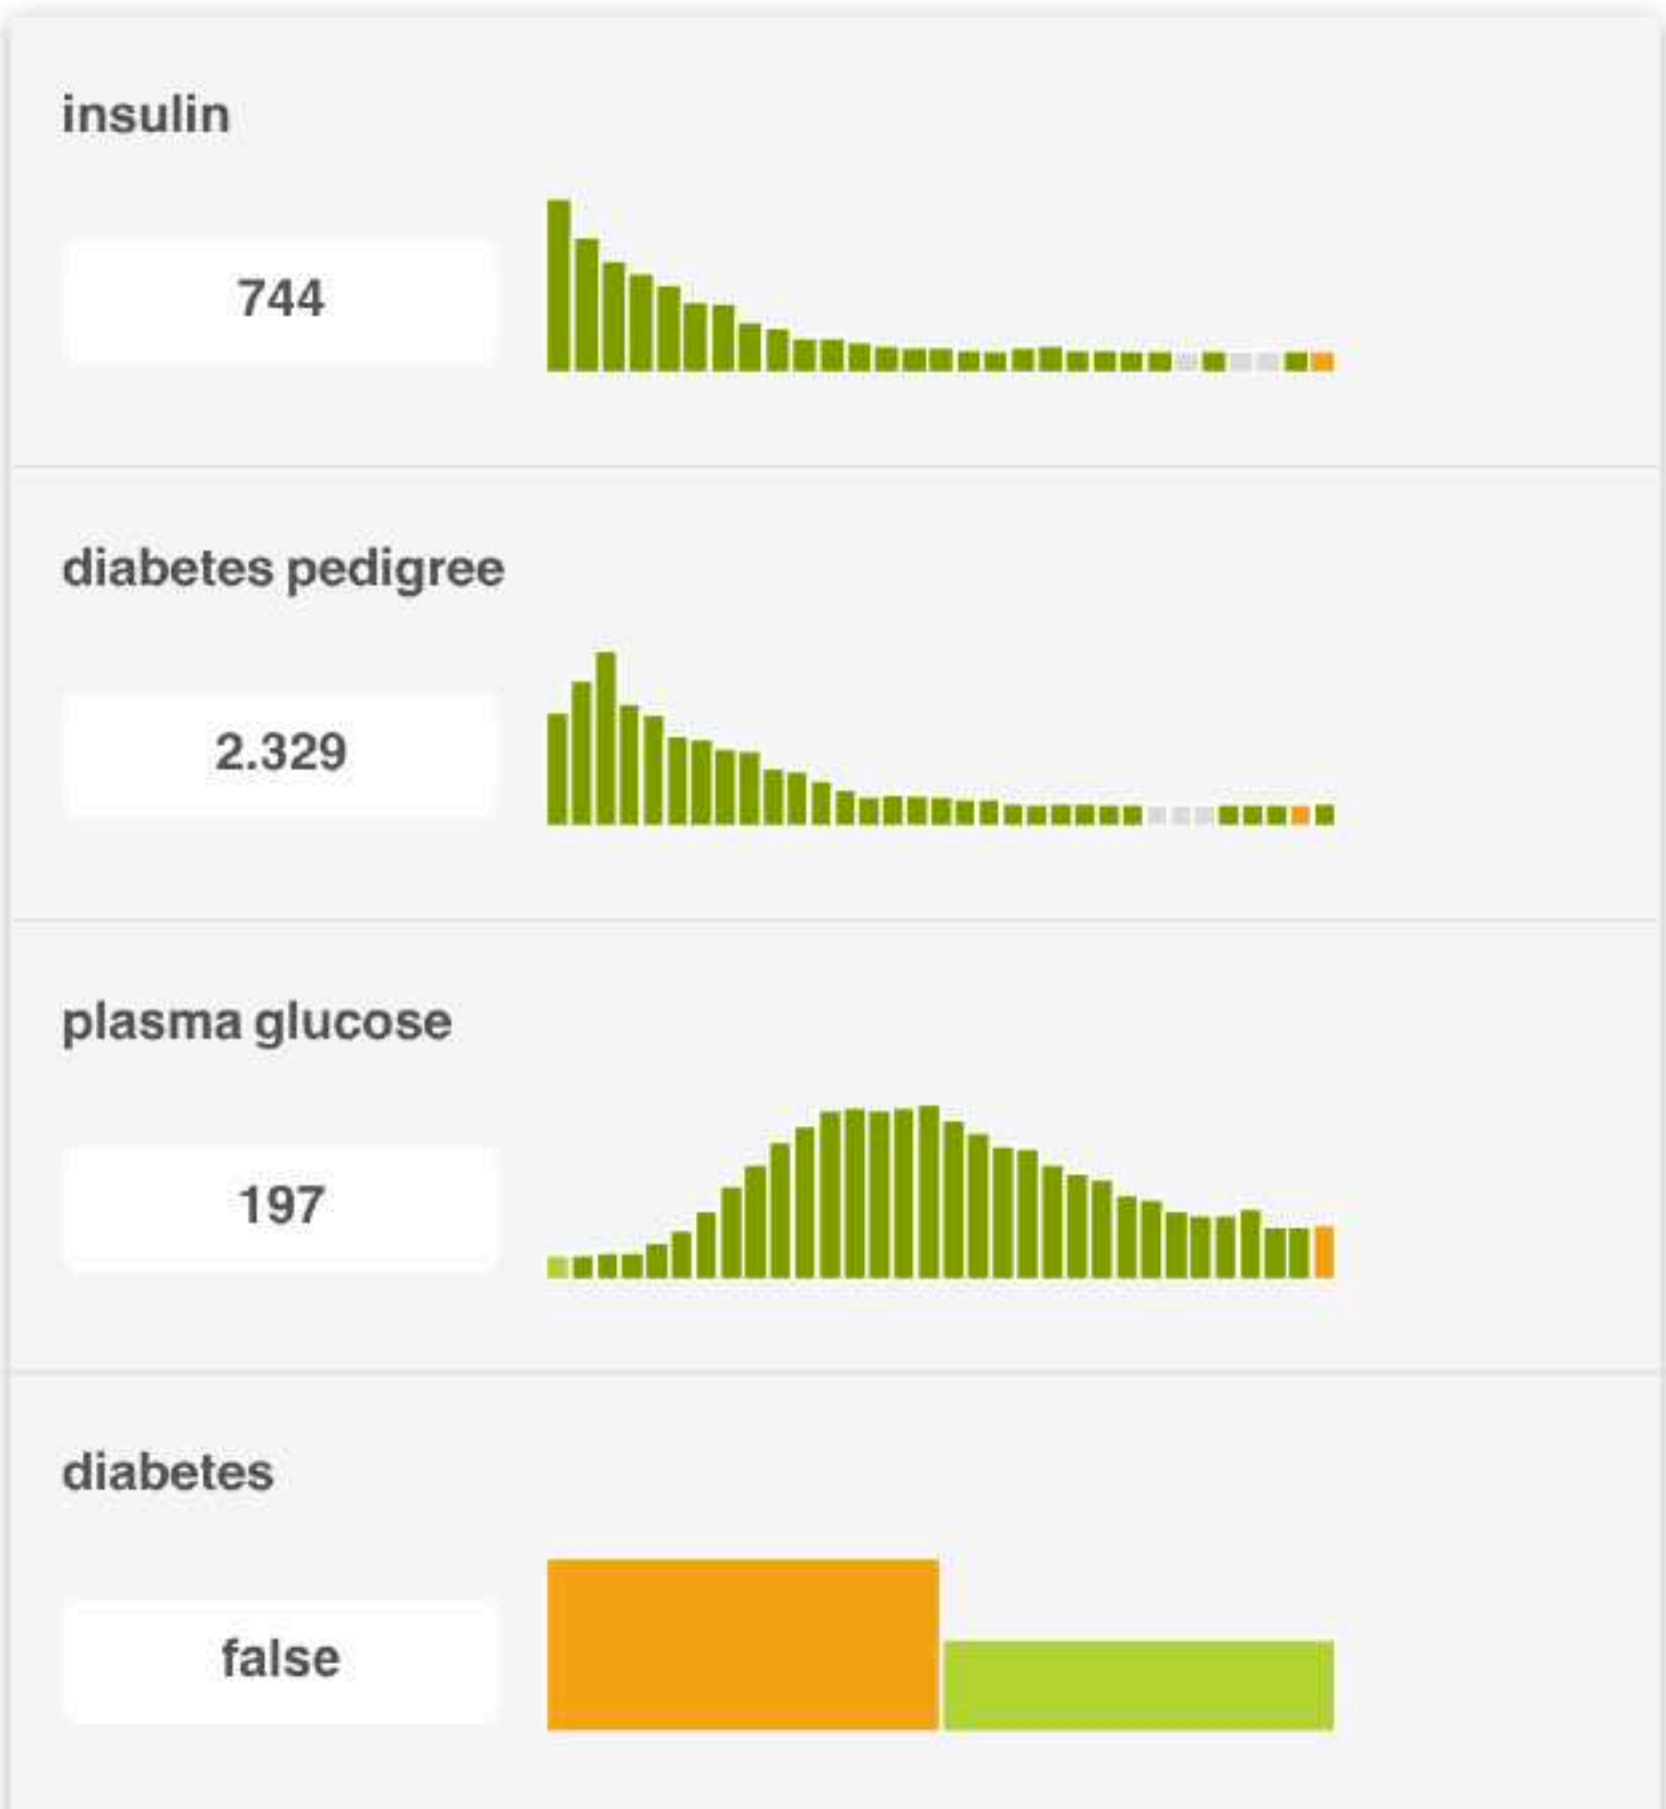
ANOMALIES	SAMPLE SIZE	NORMALIZE REPEATS	CONSTRAINTS	FOREST SIZE	INSTANCES
10	483	NO	NO	128	768

TOP ANOMALIES

Select all



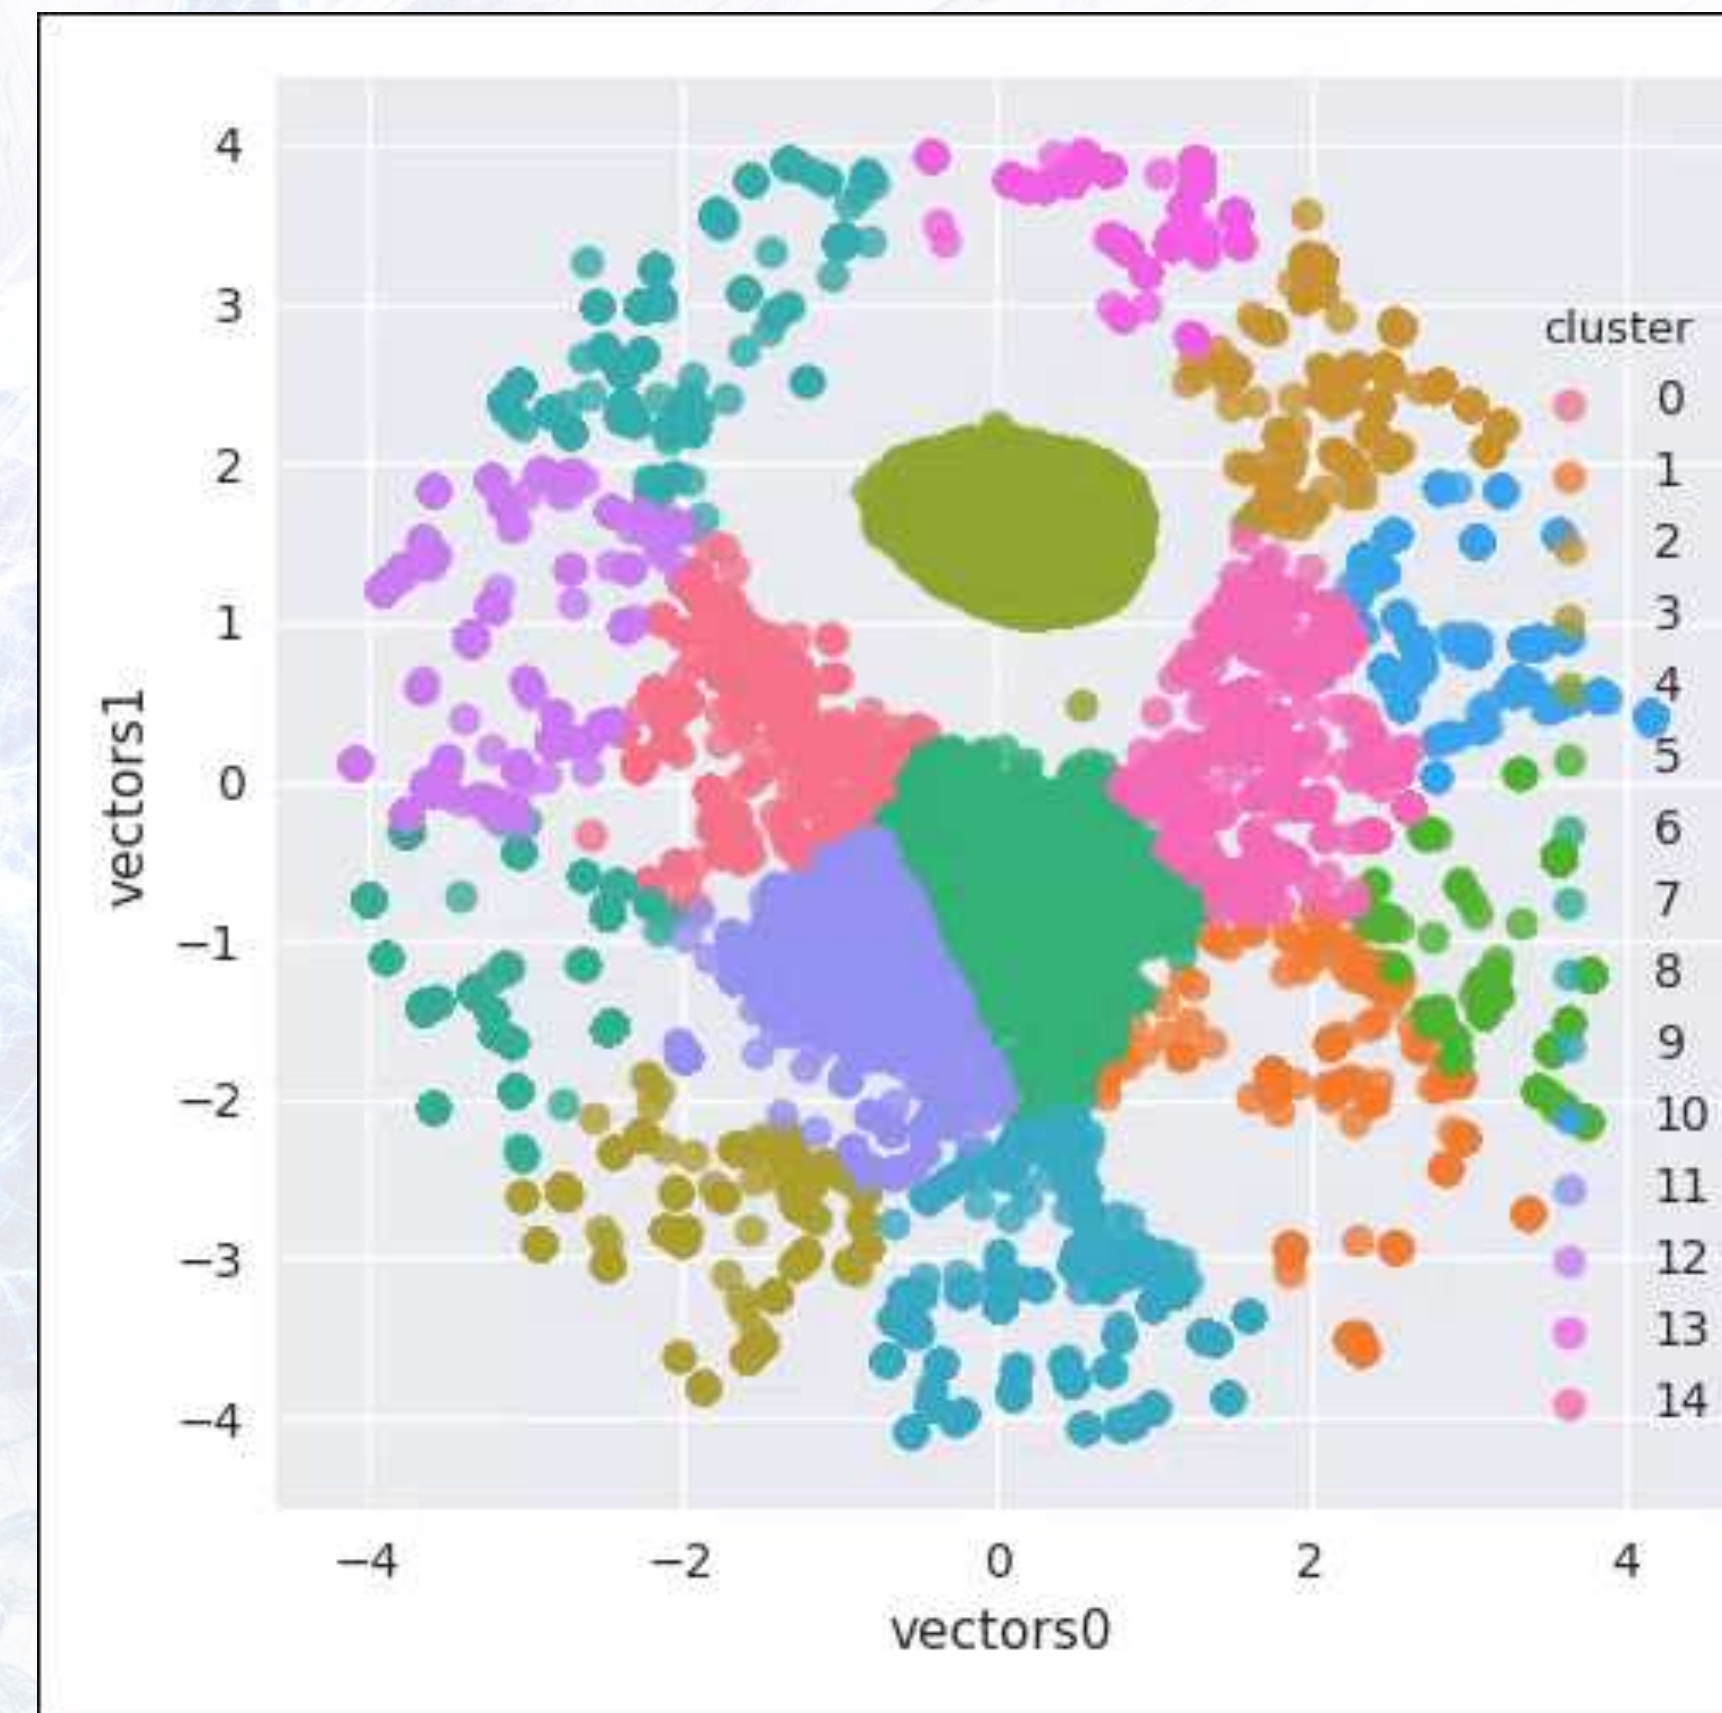
DATA INSPECTOR



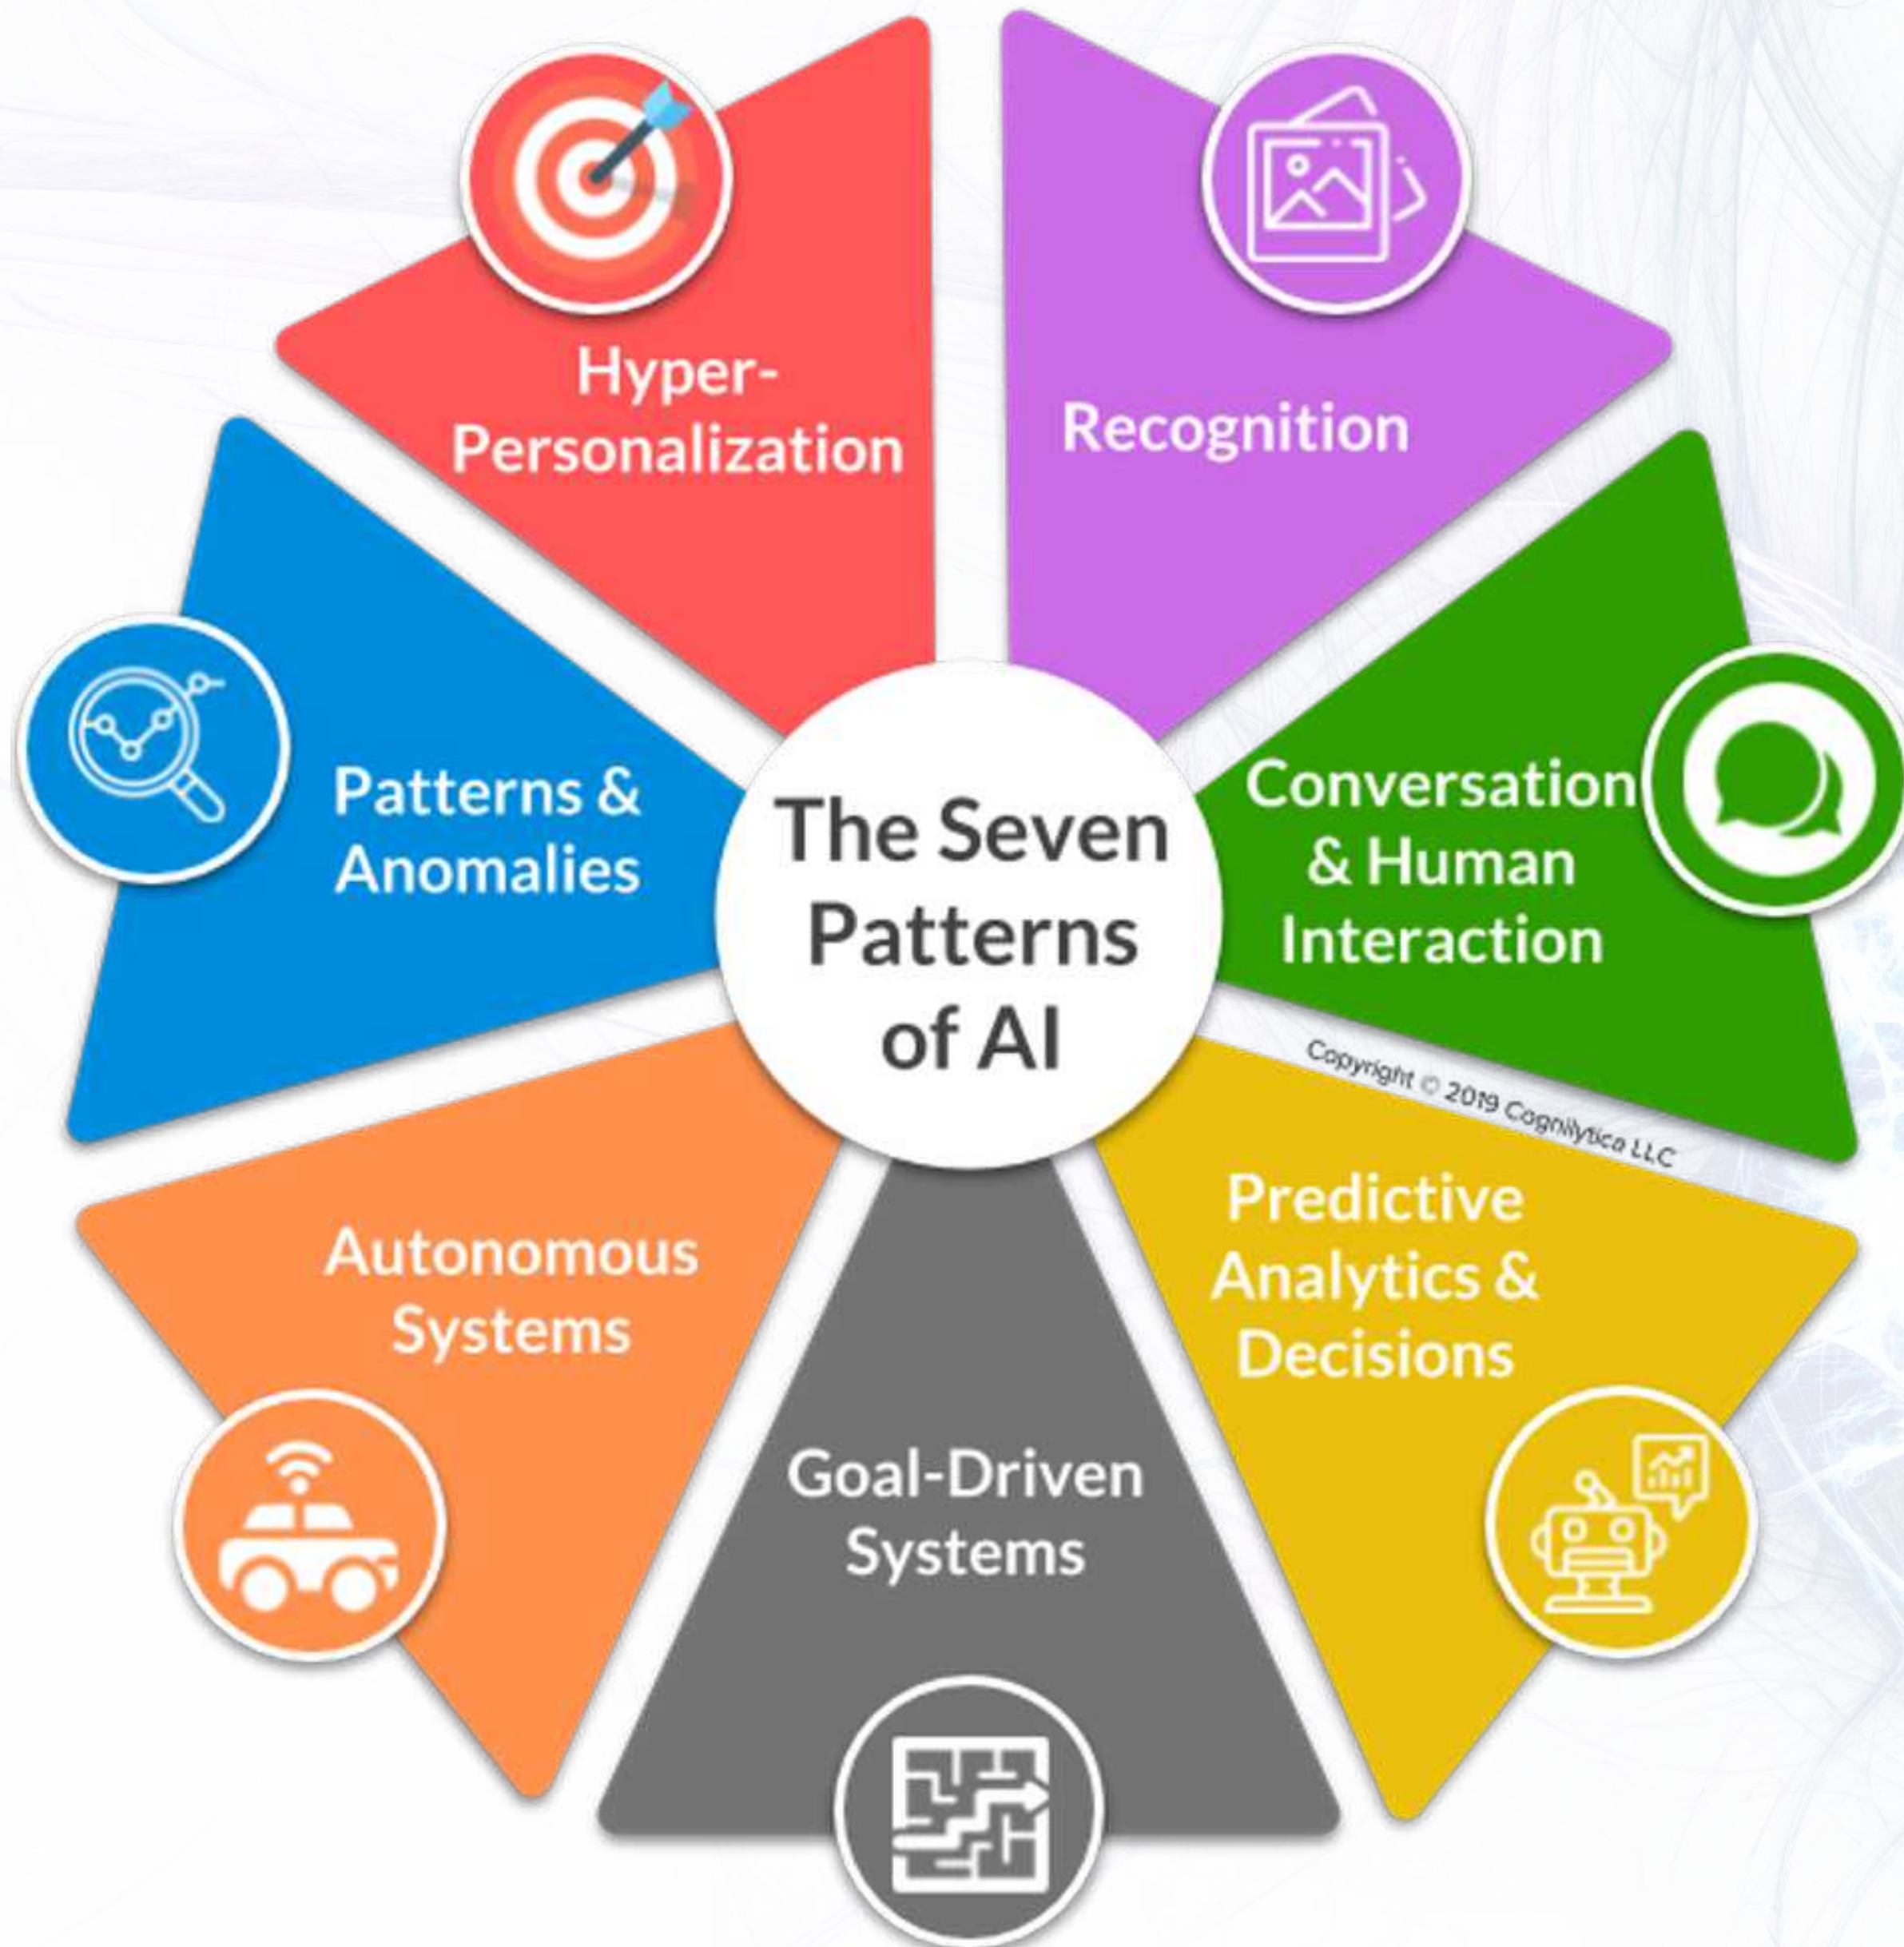
WE NEED TO LEARN IN DIFFERENT WAYS

nieuwe' betaalpas'
betaalpas' ^{inzake'} aanvraag'
en' idealbetaling'
stichting' van'
aanvraag' nieuwe'
aan' de'

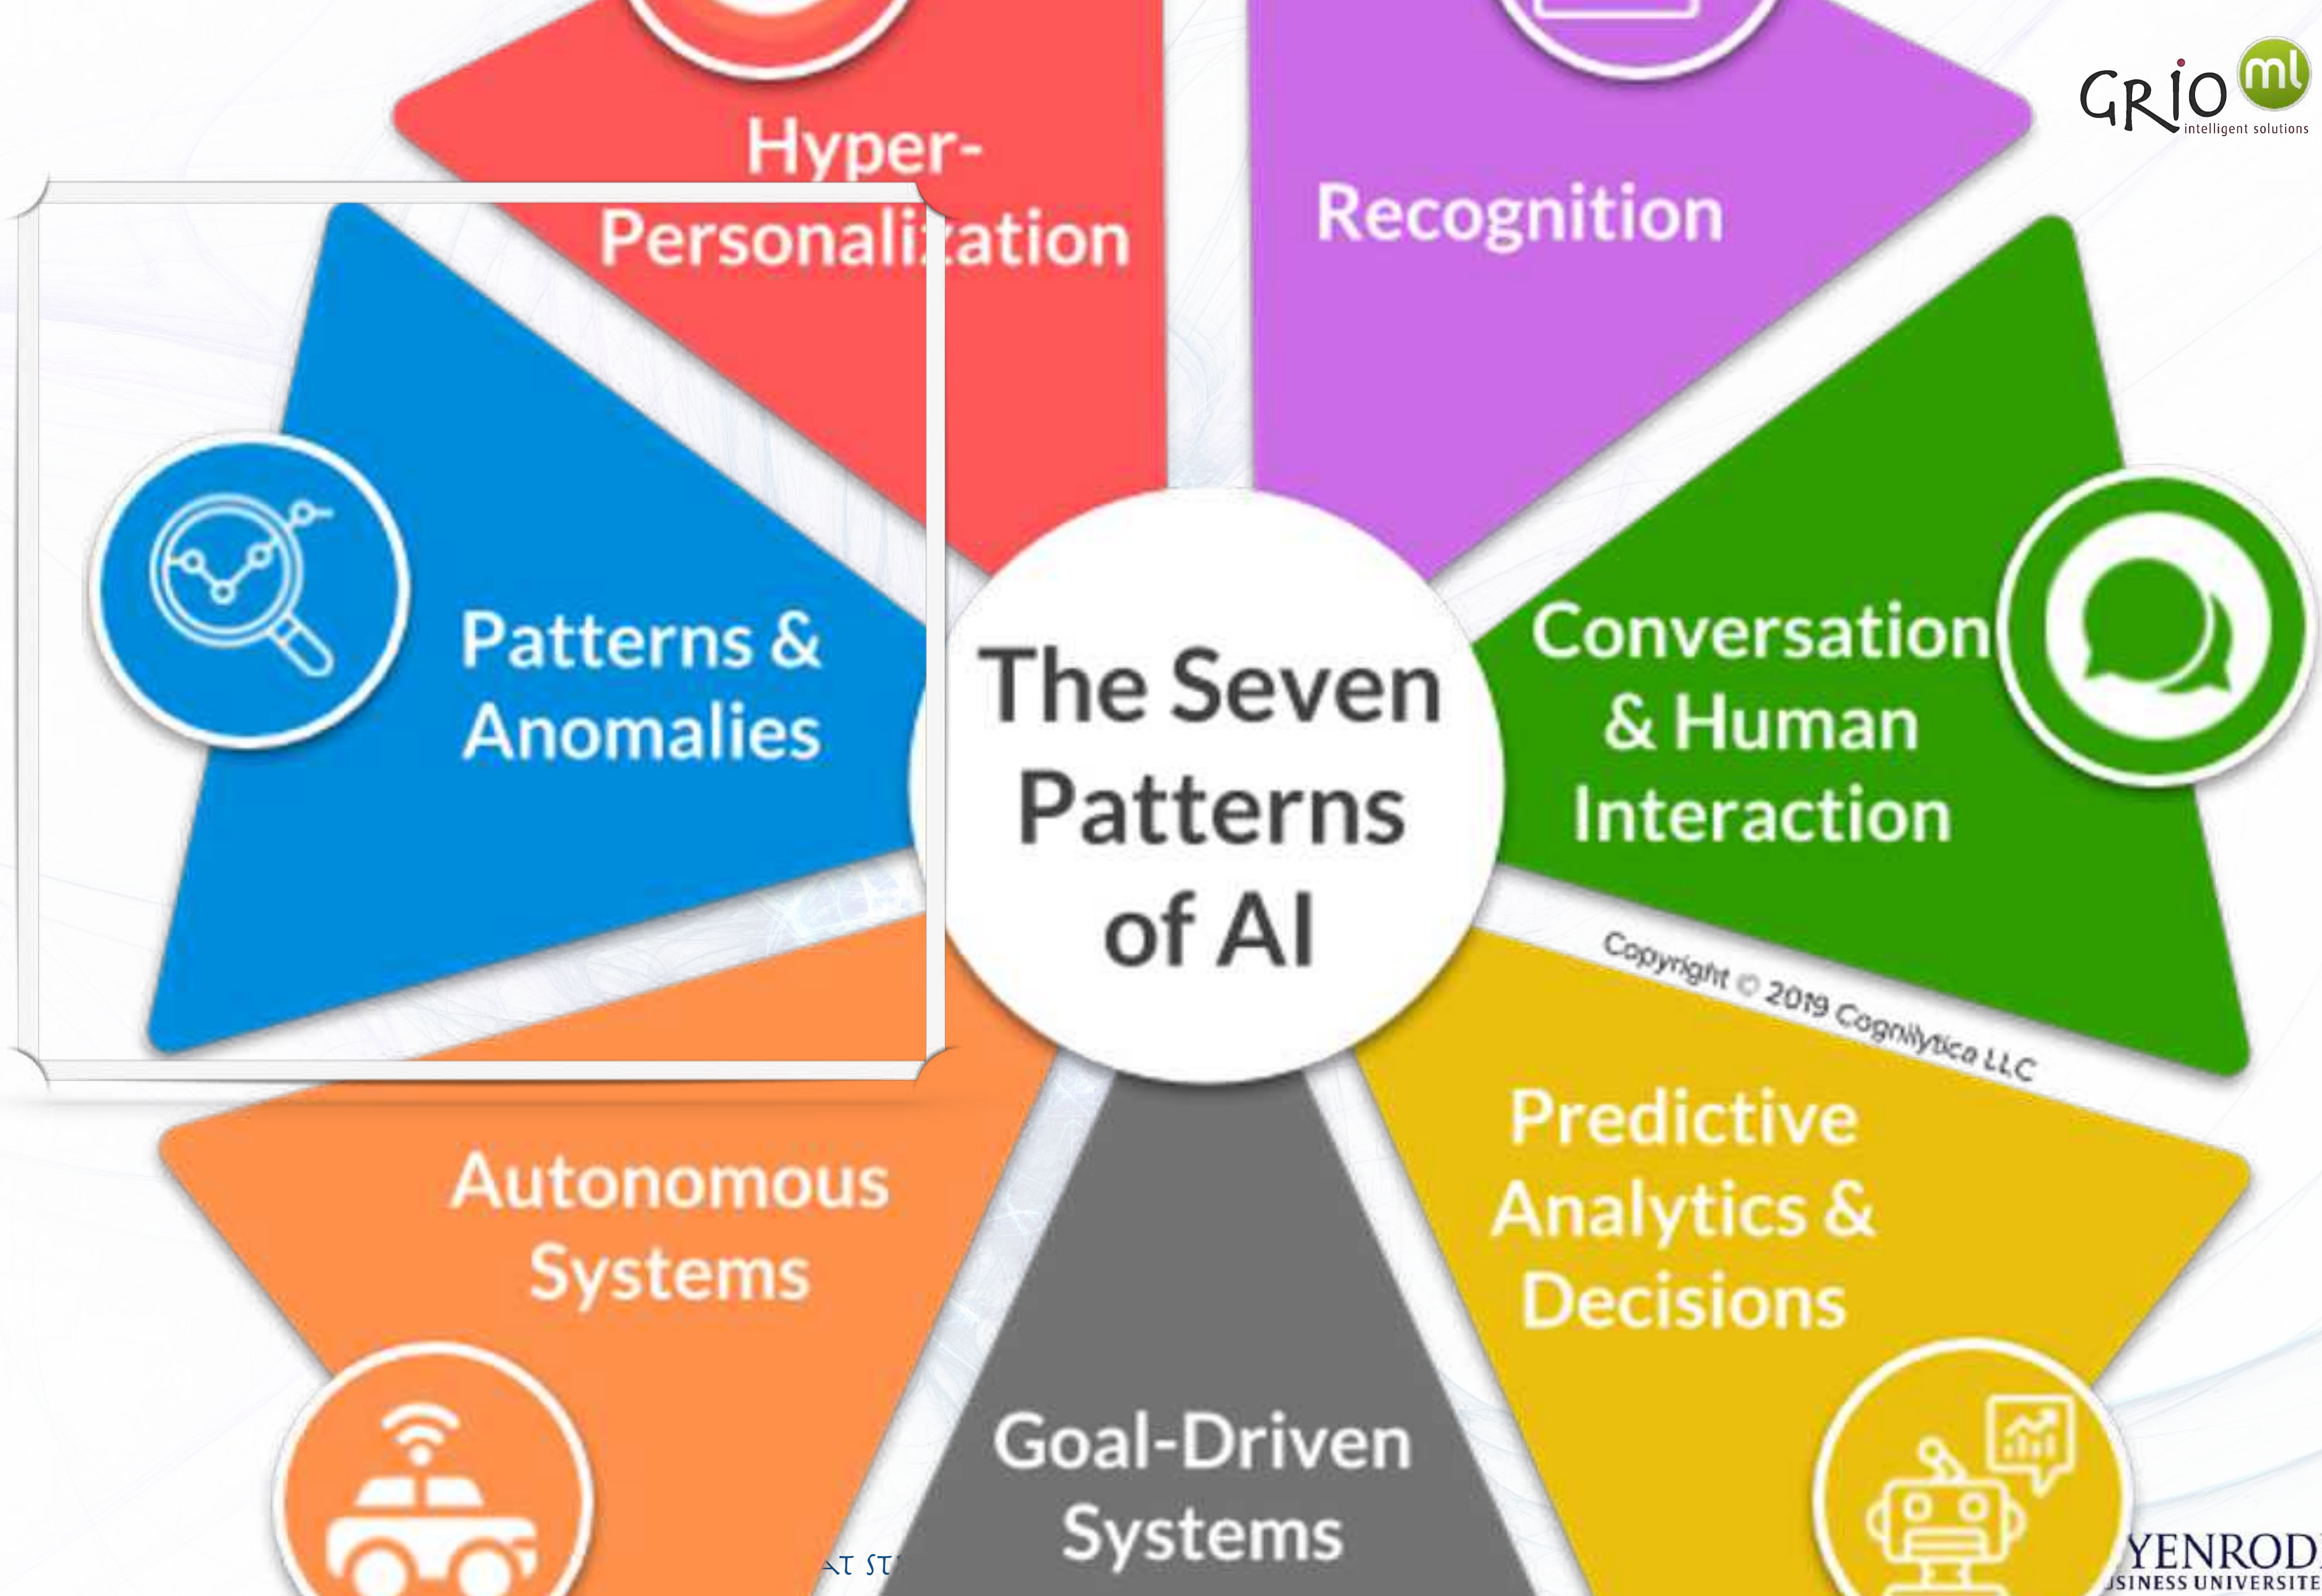
payments'



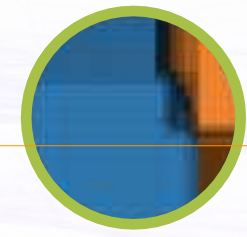
Workshop



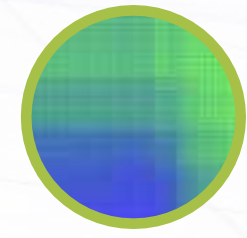
- Take your ideas and plot them into these seven patterns
- Do not forget to argue: WHY!



Across Industry Use Case Examples



Predictive Maintenance: Will this machine component fail?



Supplier Risk: What will be the delivery performance per supplier?



Forecasting: How much of each vehicle model will we sell next quarter?



Marketing: Which customers show affinity for shared mobility?



Finance: Is this transaction fraudulent?

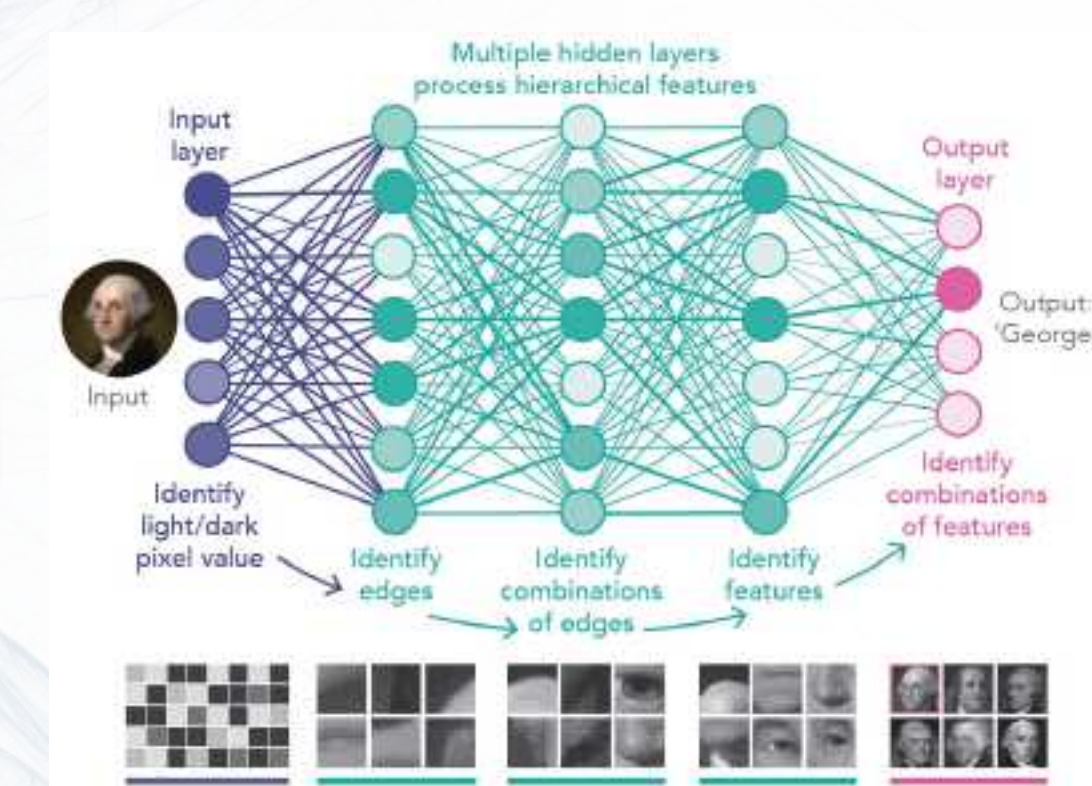
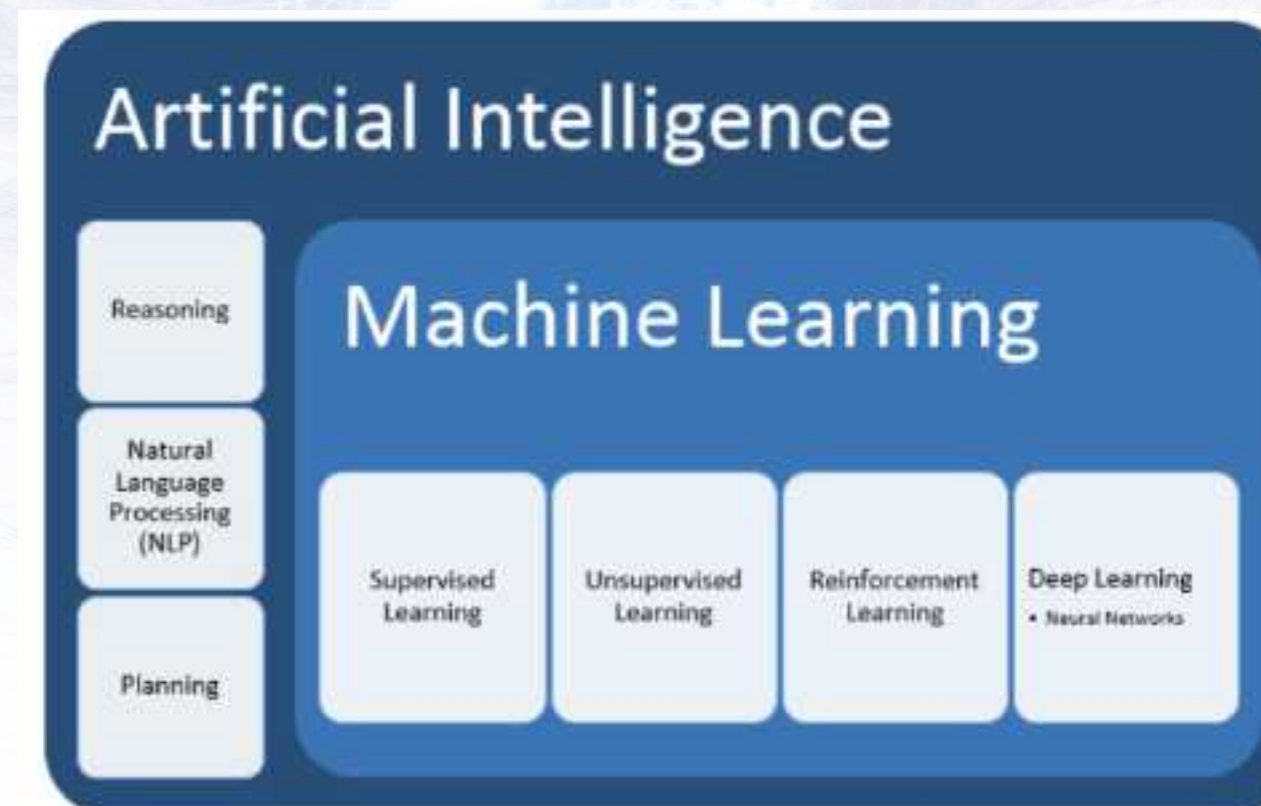
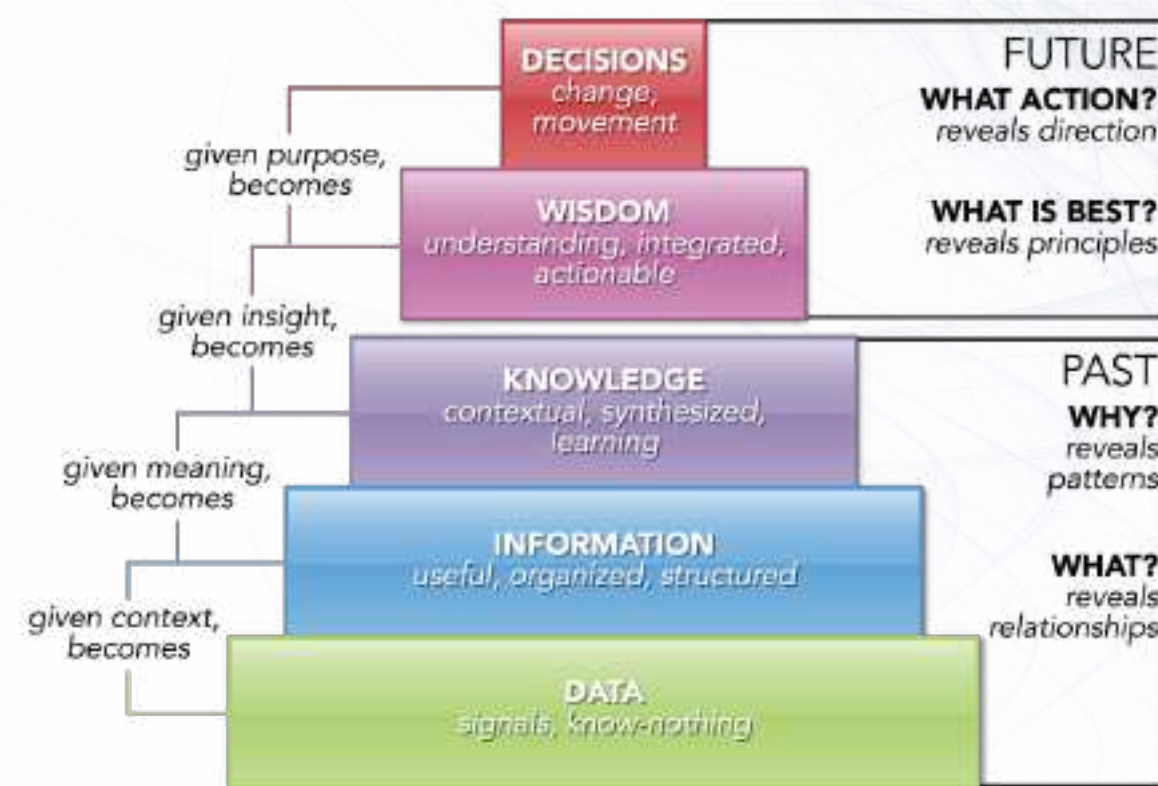
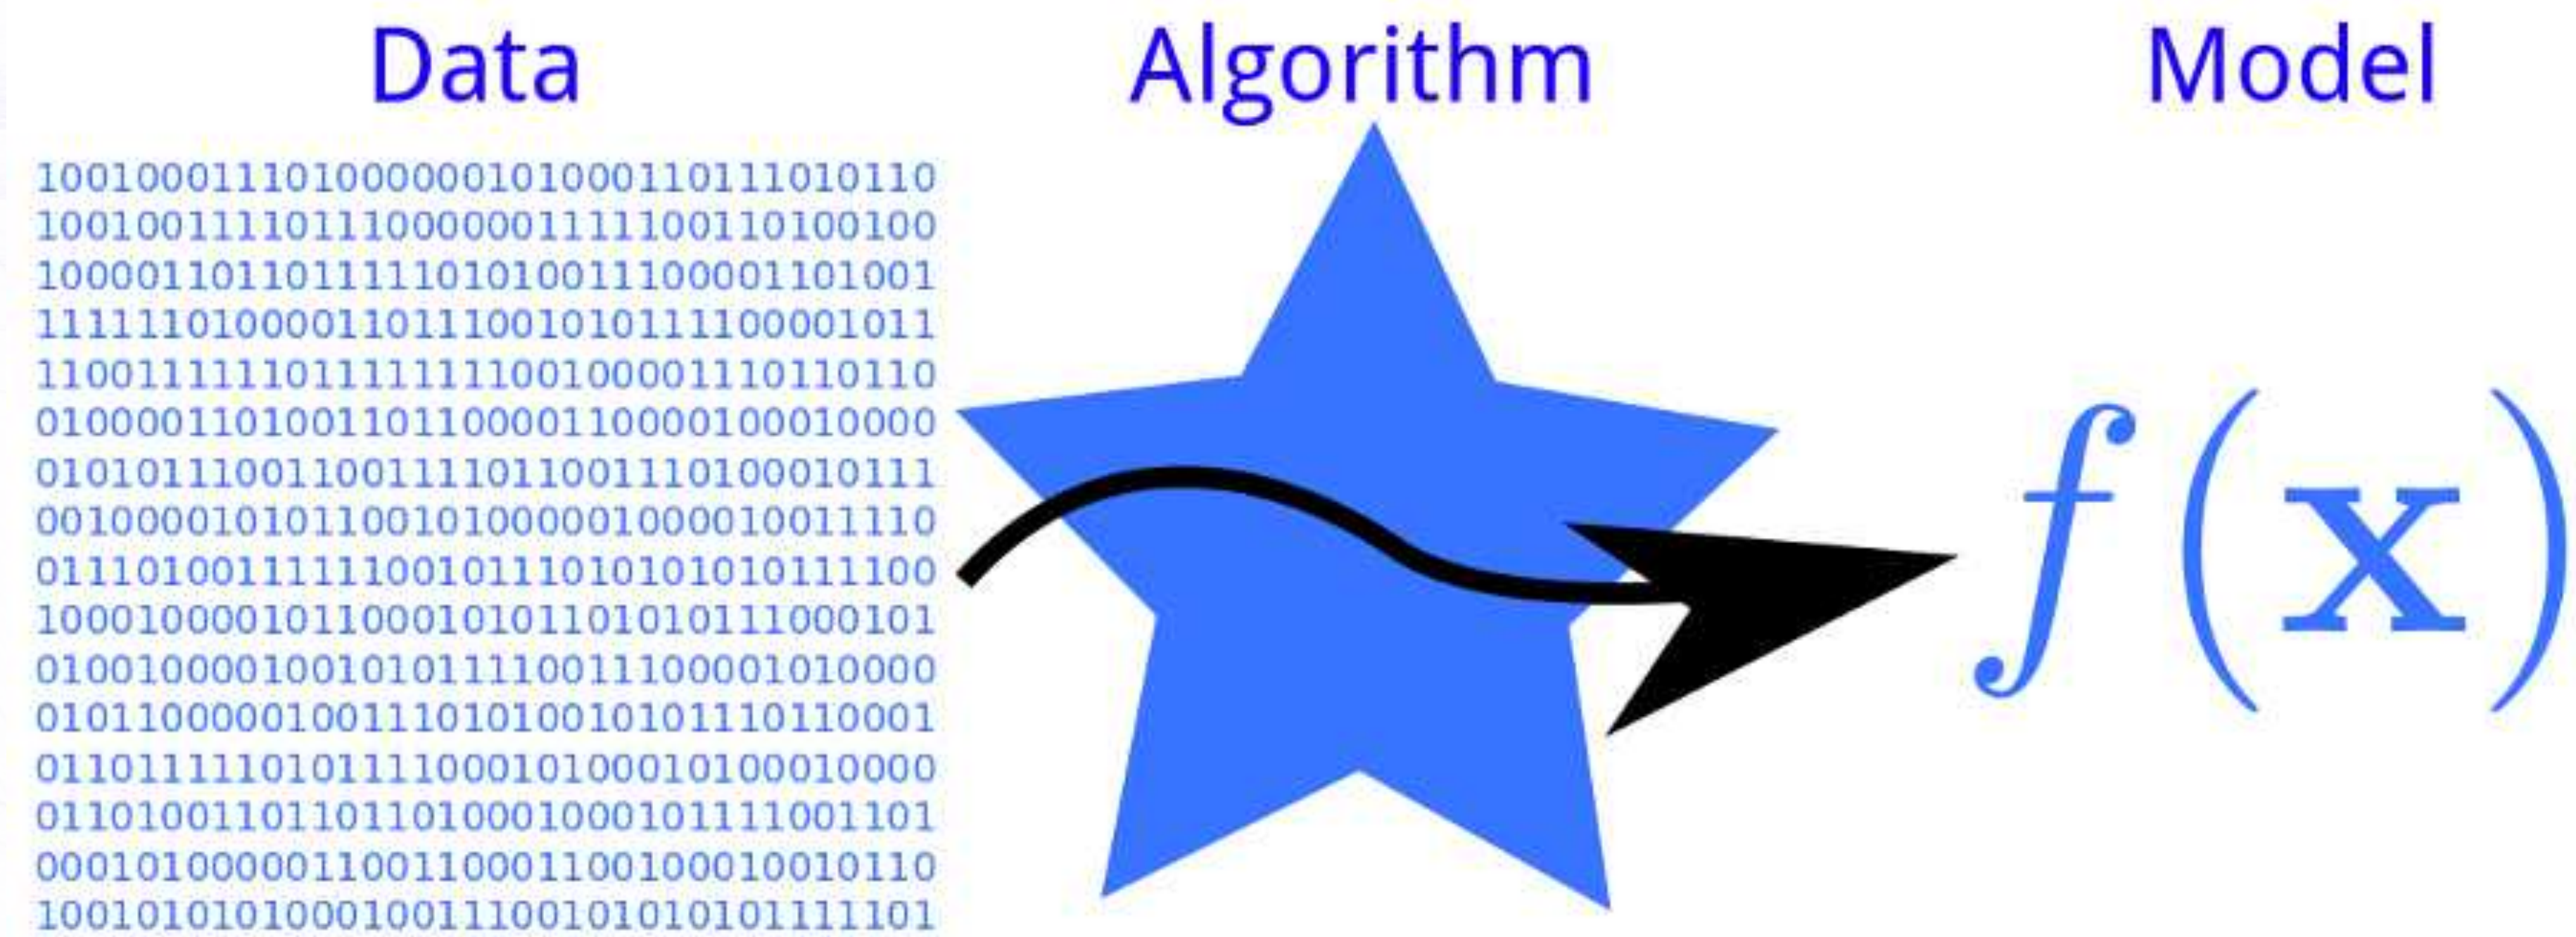


Operations: Which manufacturing configurations are optimal to use?



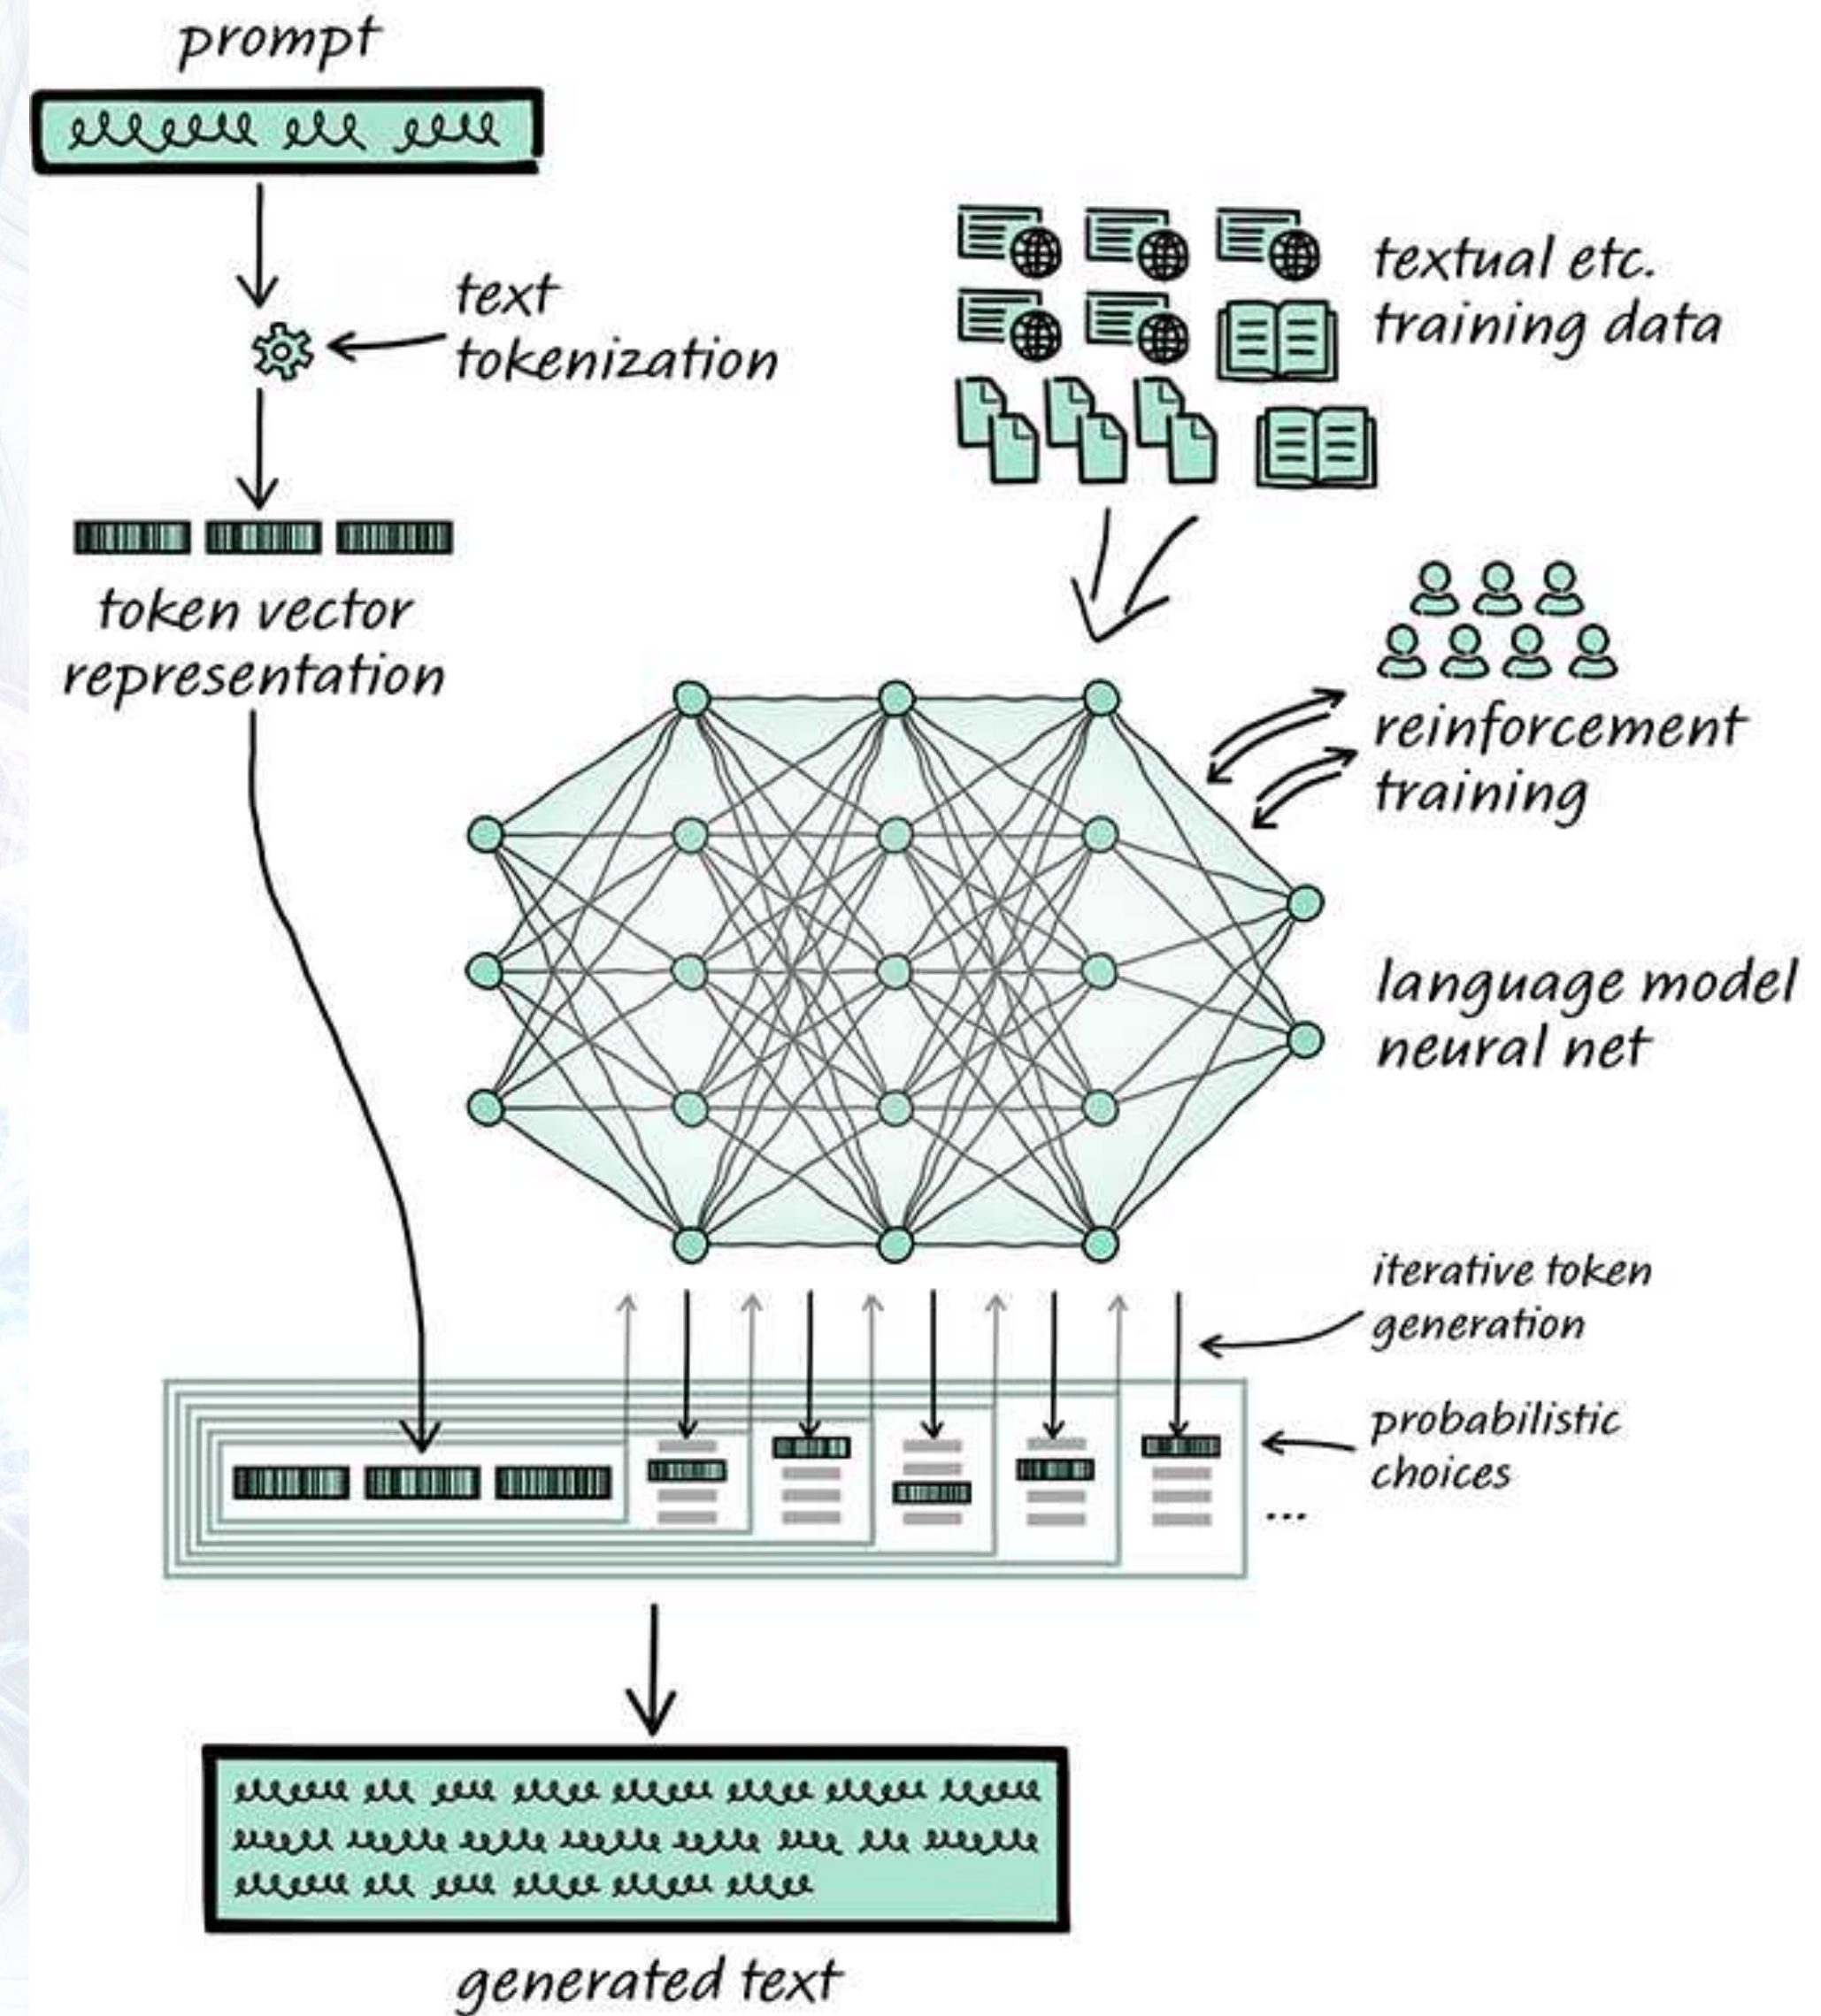
Knowledge work: What is the best way to formulate a strategy?

Data Contains the Knowledge!



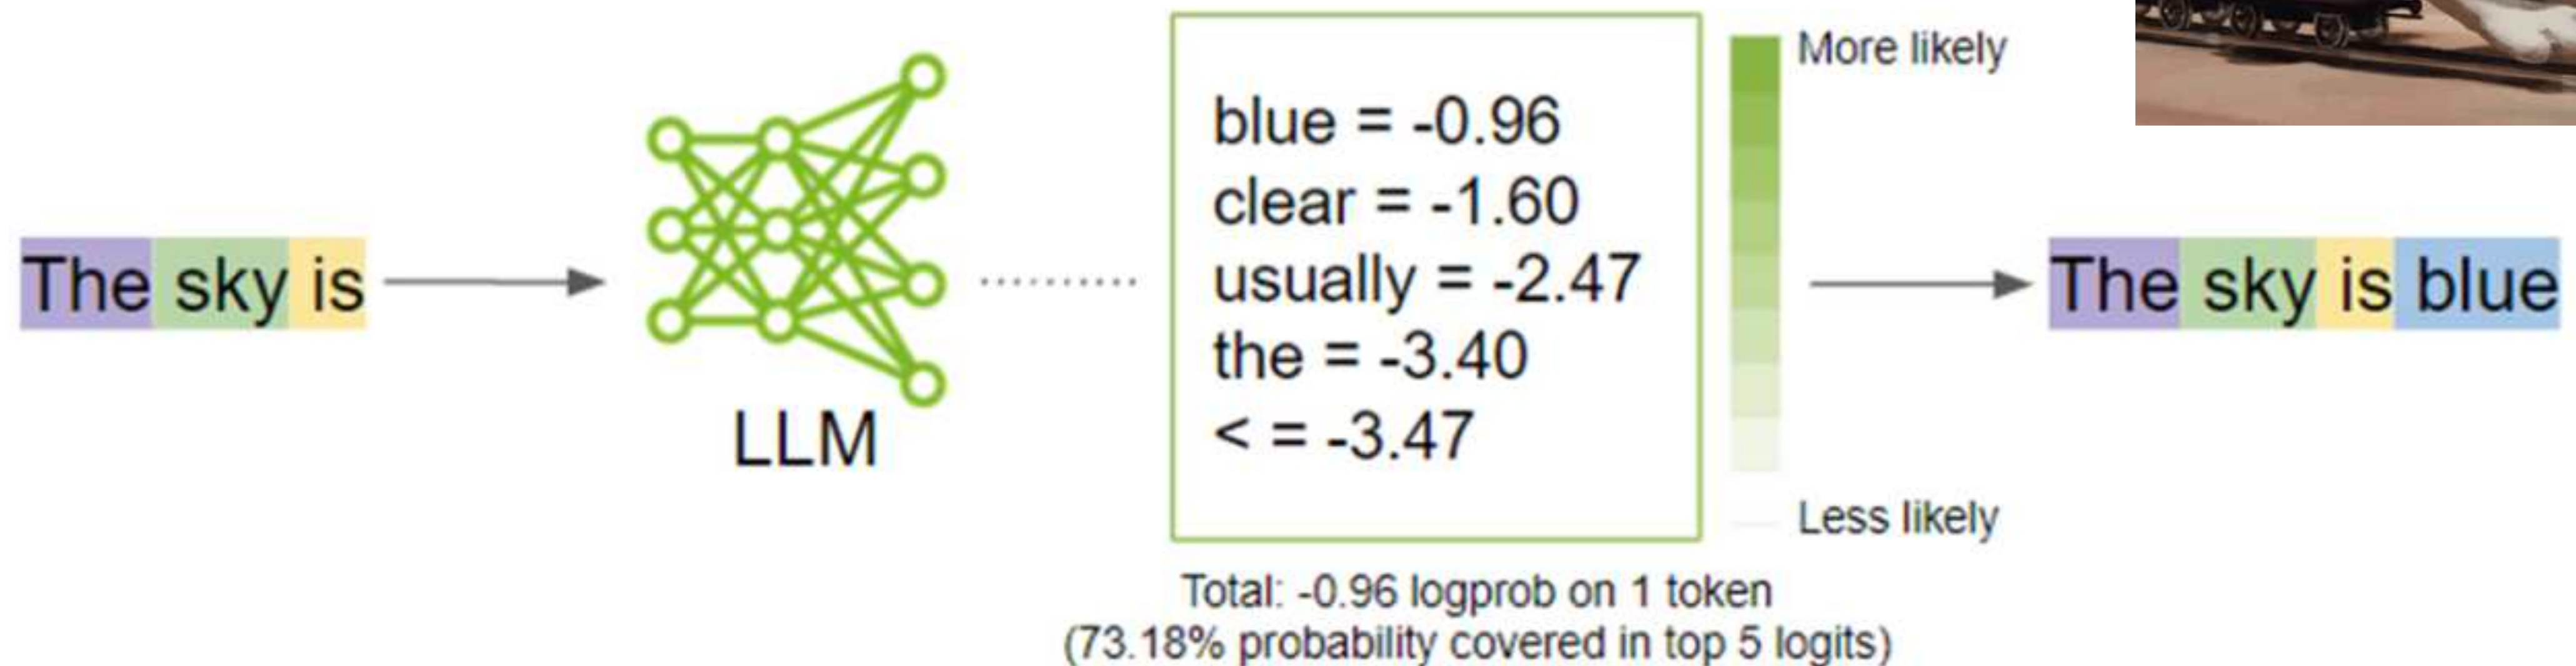
Chat GPT

- Training data
 - Large language Neural Net
 - Token and Vector generation
- Prompt
 - Tokens
 - Token Vector Representation
- Result: *Probabilistic generated text*



CHATGPT / LLM & How does it create?

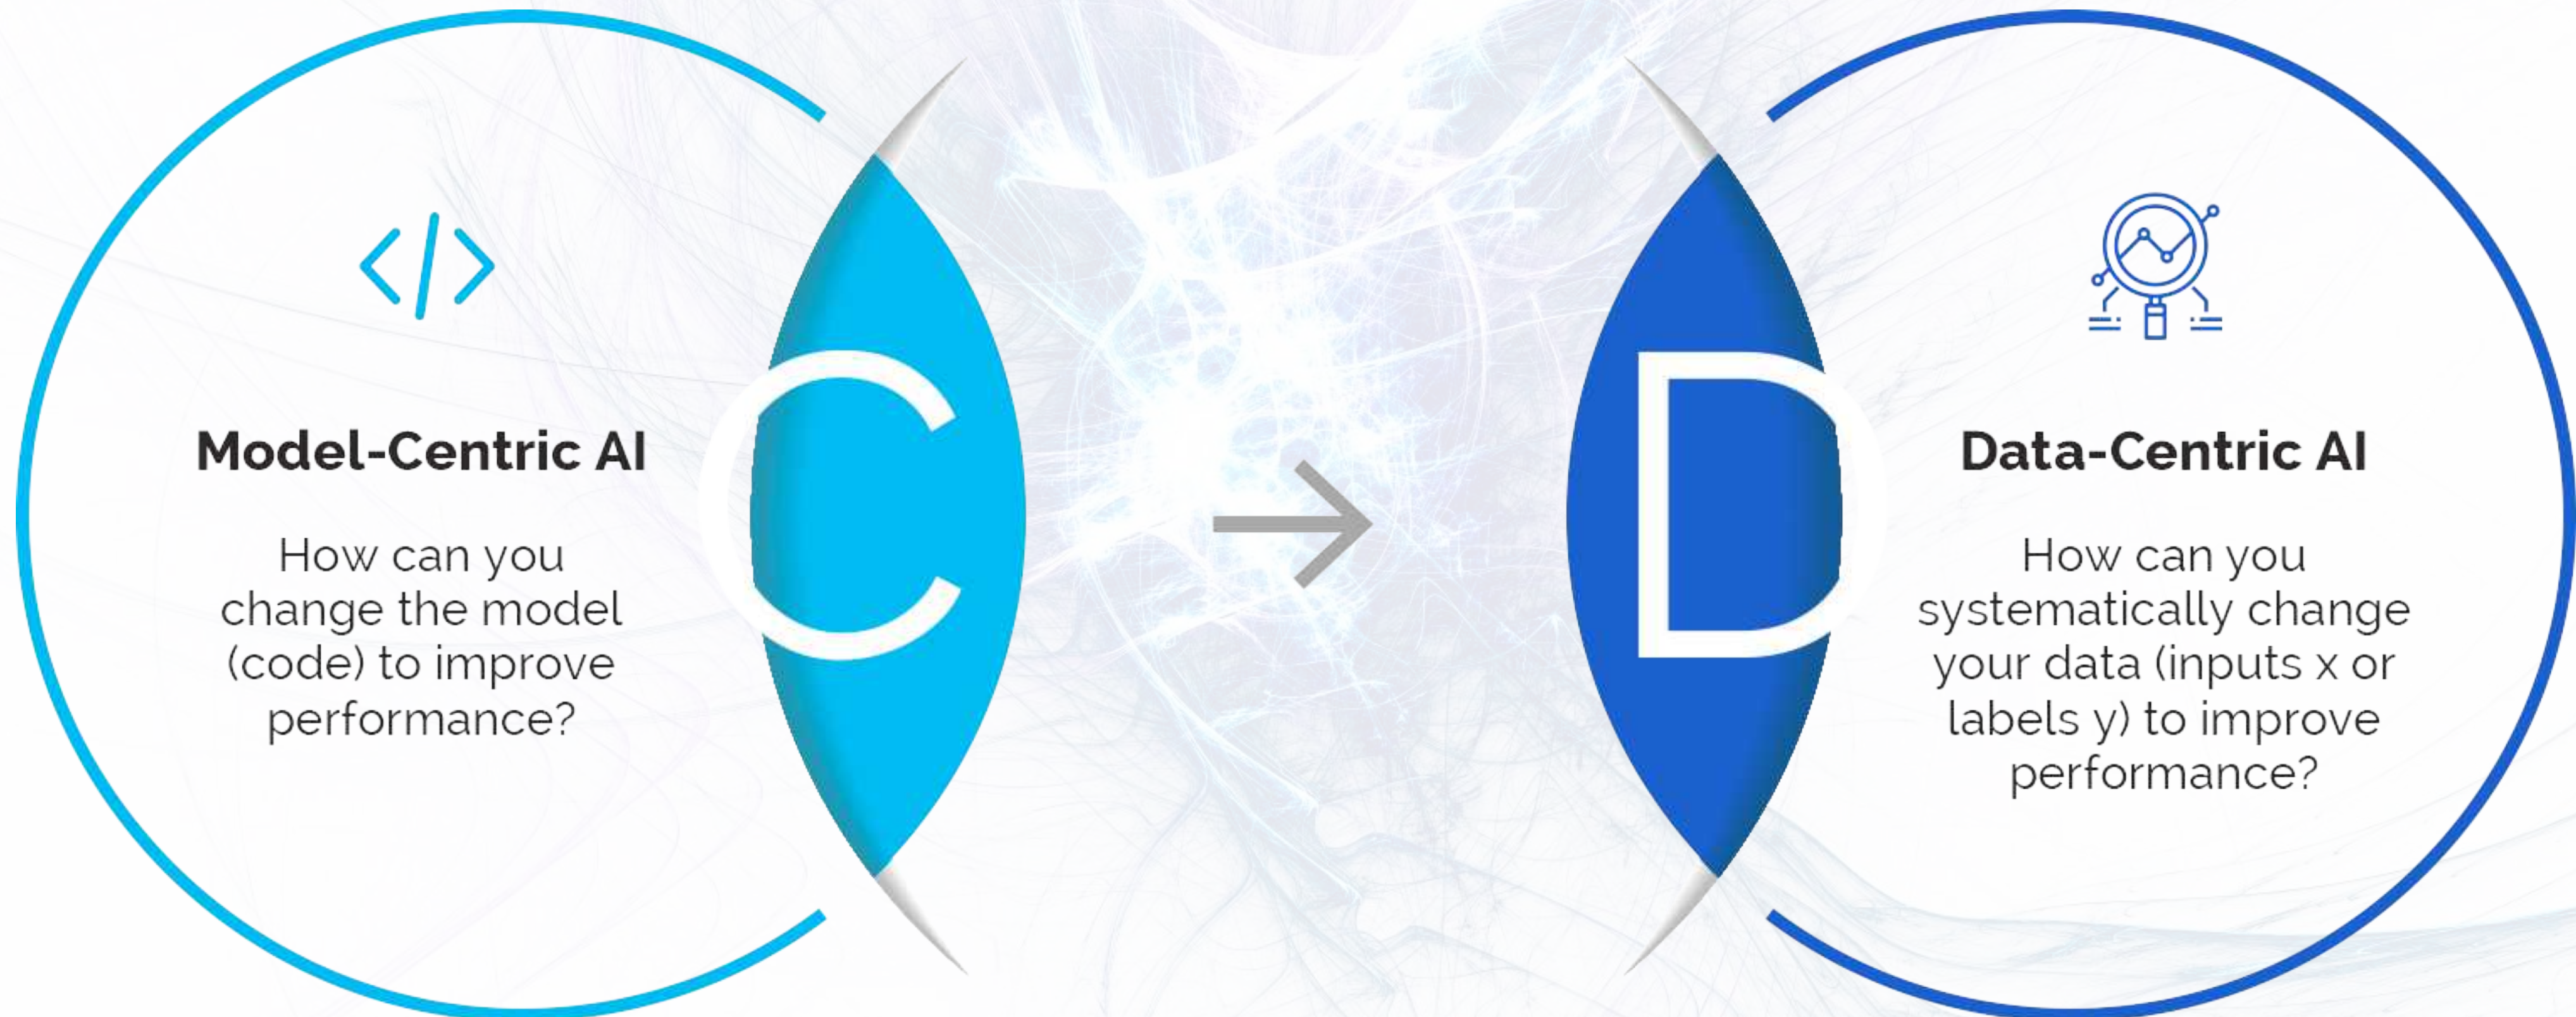
Statistical approach to text generation,
enhanced by training and Reinforcement Learning with Human
Feedback



Paradigm Shift₂

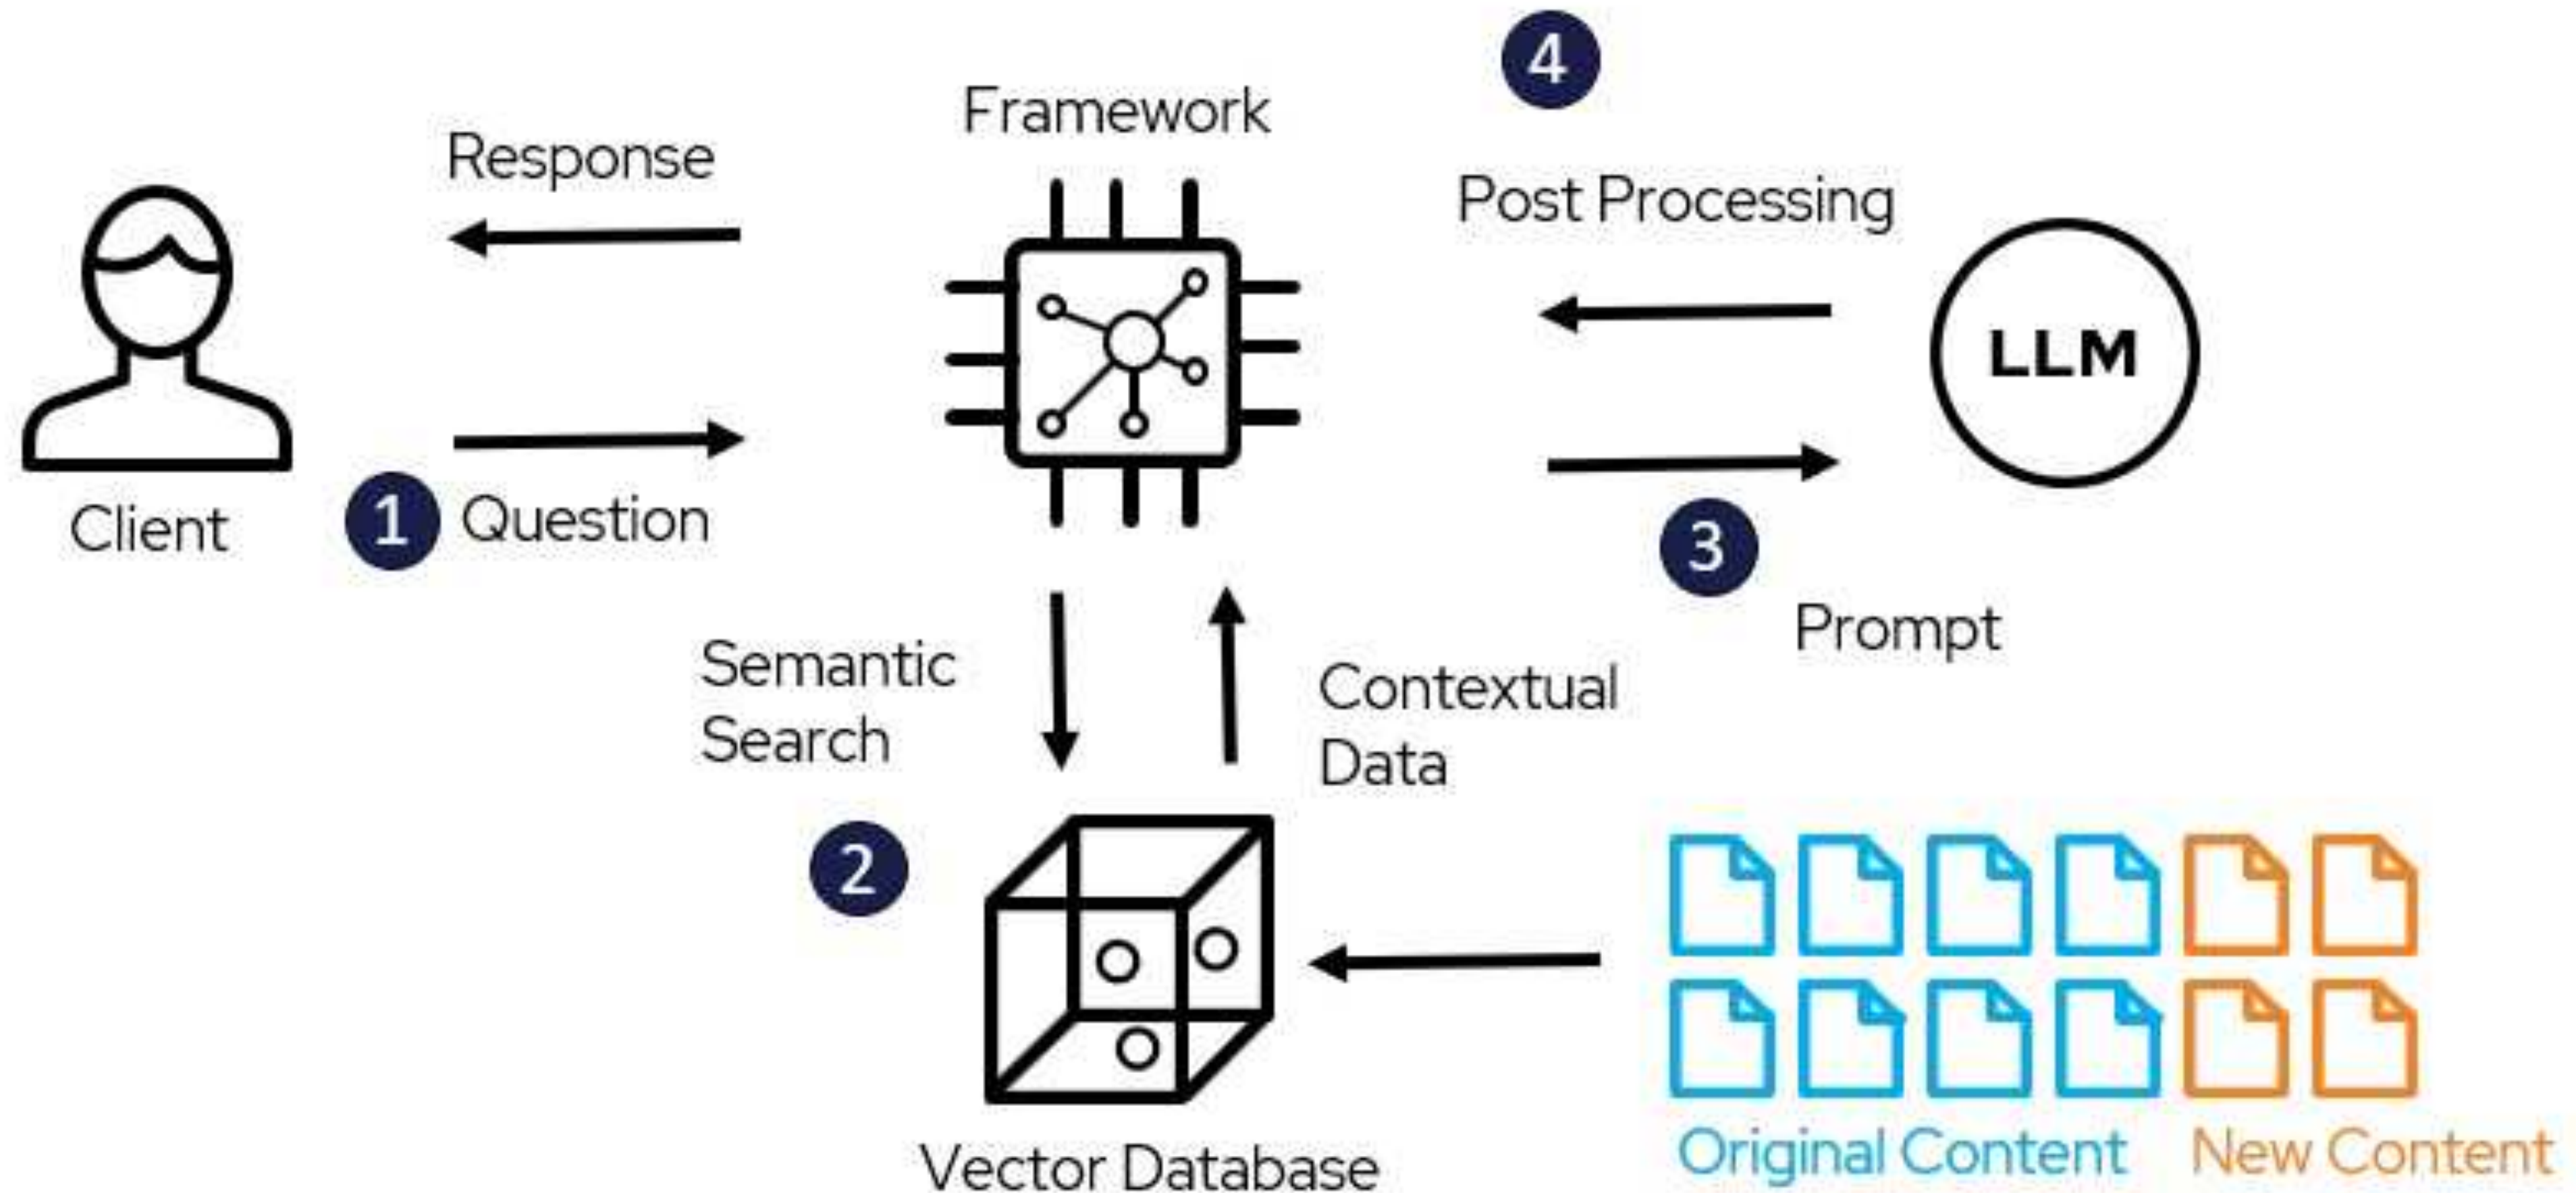
AI SUPERPOWER FOR ANALYSTS

Talk To Your Database.
Get *Explainable* Answers.



SO: RAG - making a LLM special

RAG Architecture Model



Every human task = Language

AI SUPERPOWER FOR ANALYSTS

Talk To Your Database.
Get *Explainable* Answers.

WrenAI is your AI data assistant that helps you get results and insights faster by asking questions without writing SQL, and it's *open-source*!

WrenAI

Interactive
Interface

Observability
& Explainability

Semantic
Engine

Continuous
Learning

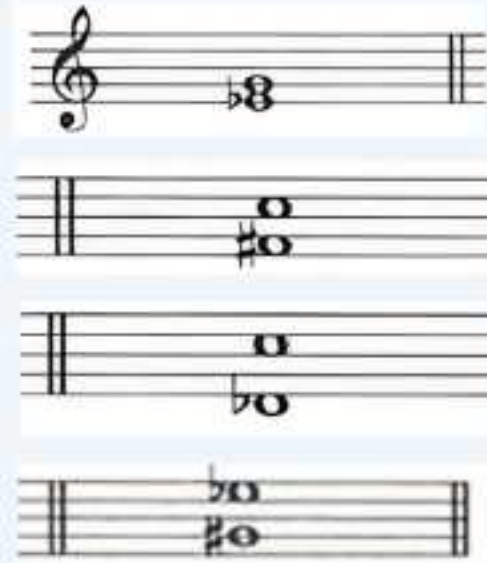


Art & Design

Question: Among the following harmonic intervals, which one is constructed incorrectly?

Options:

- (A) Major third <image 1>
- (B) Diminished fifth <image 2>
- (C) Minor seventh <image 3>**
- (D) Diminished sixth <image 4>



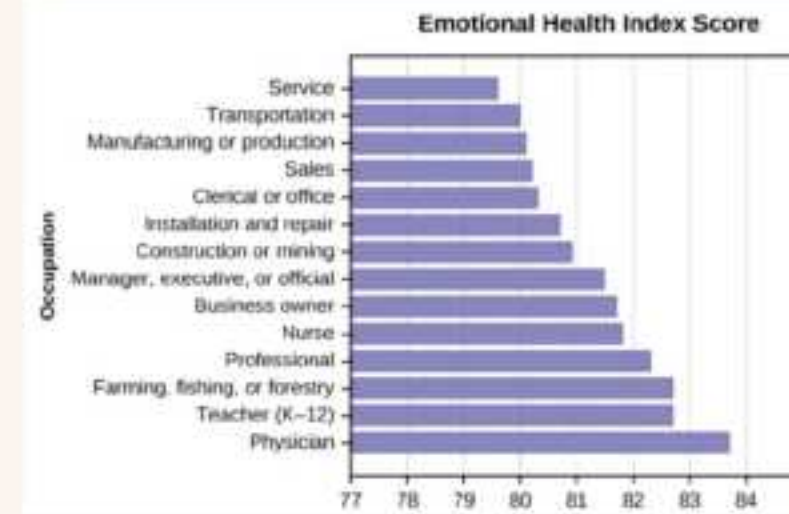
Subject: Music; **Subfield:** Music;
Image Type: Sheet Music;
Difficulty: Medium

Business

Question: ...The graph shown is compiled from data collected by Gallup <image 1>. Find the probability that the selected Emotional Health Index Score is between 80.5 and 82?

Options:

- (A) 0
- (B) 0.2142
- (C) 0.3571**
- (D) 0.5



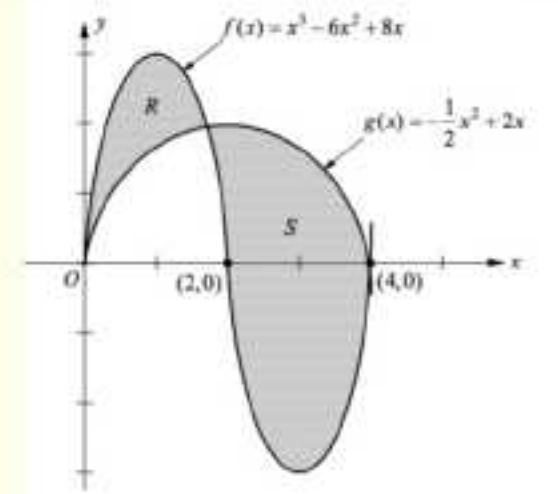
Subject: Marketing; **Subfield:** Market Research;
Image Type: Plots and Charts;
Difficulty: Medium

Science

Question: <image 1> The region bounded by the graph as shown above. Choose an integral expression that can be used to find the area of R.

Options:

- (A) $\int_0^{1.5} [f(x) - g(x)] dx$**
- (B) $\int_0^{1.5} [g(x) - f(x)] dx$
- (C) $\int_0^2 [f(x) - g(x)] dx$
- (D) $\int_0^2 [g(x) - x(x)] dx$



Subject: Math; **Subfield:** Calculus;
Image Type: Mathematical Notations;
Difficulty: Easy

Health & Medicine

Question: You are shown subtraction <image 1>, T2 weighted <image 2> and T1 weighted axial <image 3> from a screening breast MRI. What is the etiology of the finding in the left breast?

Options:

- (A) Susceptibility artifact
- (B) Hematoma
- (C) Fat necrosis**
- (D) Silicone granuloma



Subject: Clinical Medicine; **Subfield:** Clinical Radiology;
Image Type: Body Scans: MRI, CT.;
Difficulty: Hard

Humanities & Social Science

Question: In the political cartoon, the United States is seen as fulfilling which of the following roles? <image 1>

Option:

- (A) Oppressor
- (B) Imperialist
- (C) Savior**
- (D) Isolationist



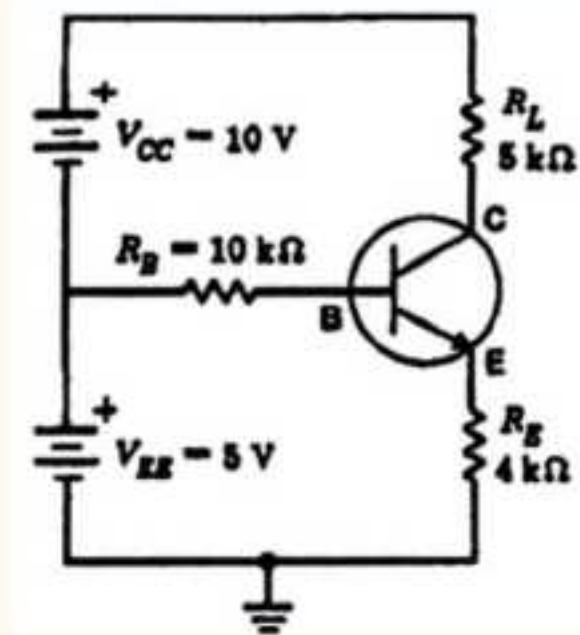
Subject: History; **Subfield:** Modern History;
Image Type: Comics and Cartoons;
Difficulty: Easy

Tech & Engineering

Question: Find the VCE for the circuit shown in <image 1>. Neglect VBE

Answer: 3.75

Explanation: ... $I_E = [(V_{EE}) / (R_E)] = [(5 \text{ V}) / (4 \text{ k-ohm})] = 1.25 \text{ mA}$; $V_{CE} = V_{CC} - I_E R_L = 10 \text{ V} - (1.25 \text{ mA}) 5 \text{ k-ohm}$; $V_{CE} = 10 \text{ V} - 6.25 \text{ V} = 3.75 \text{ V}$

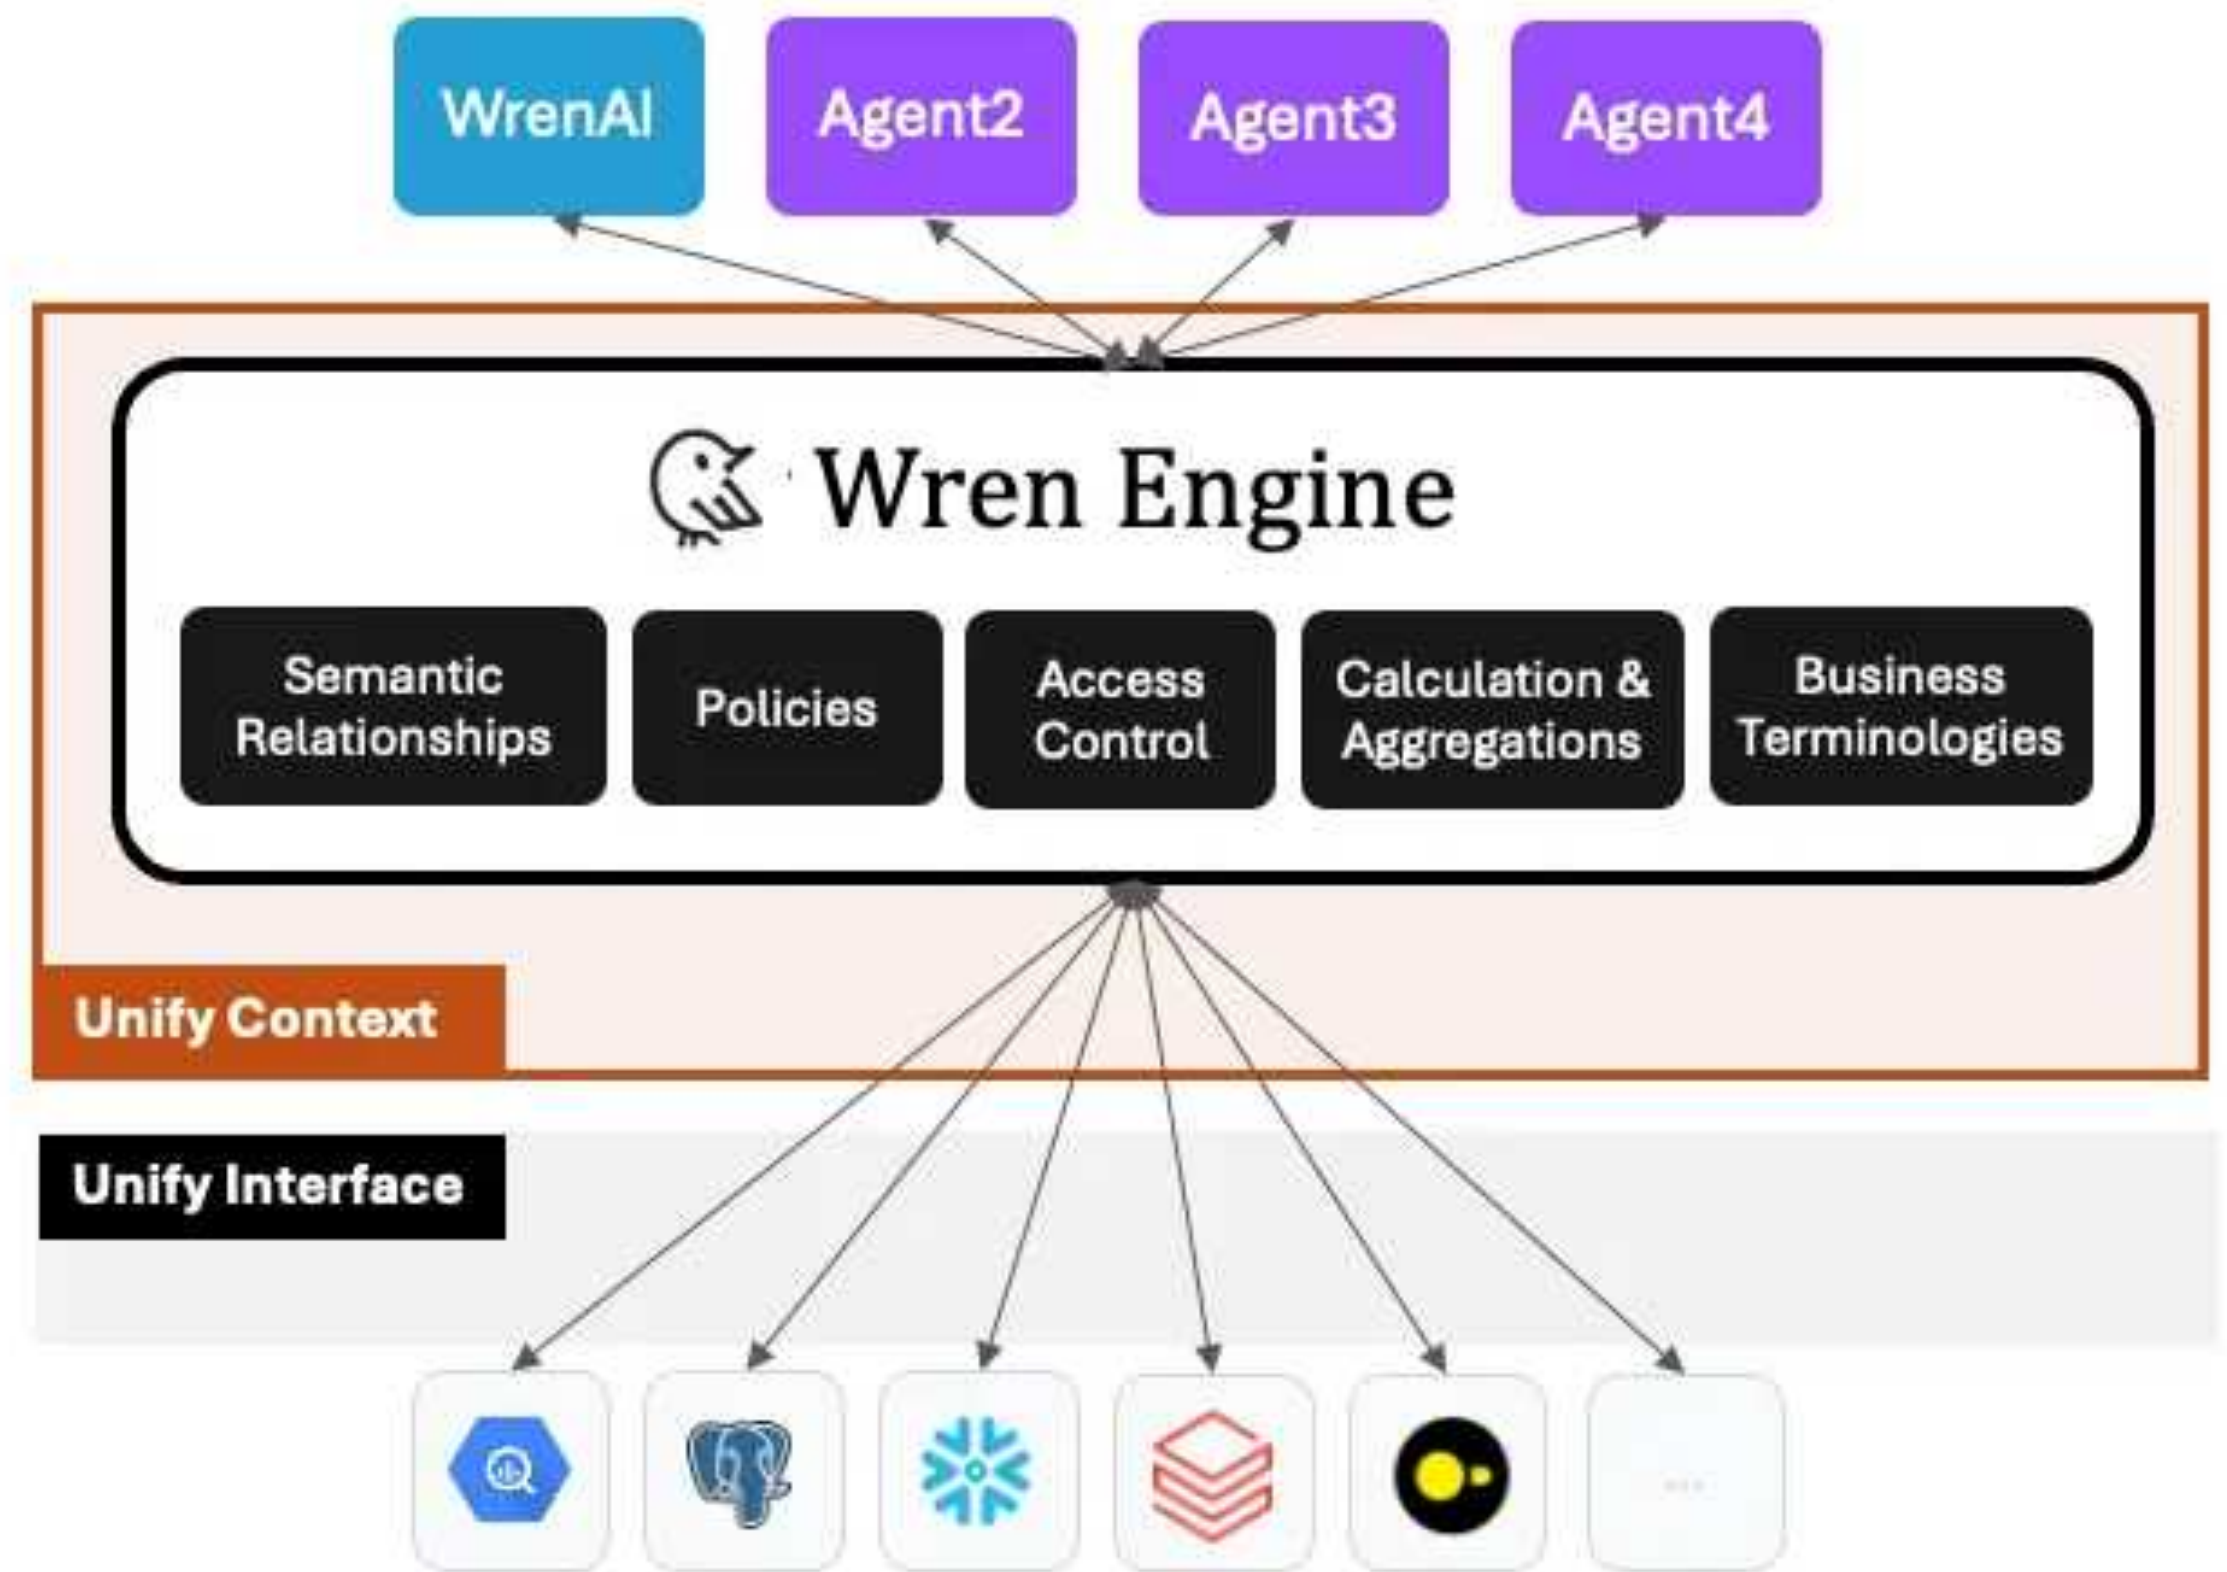
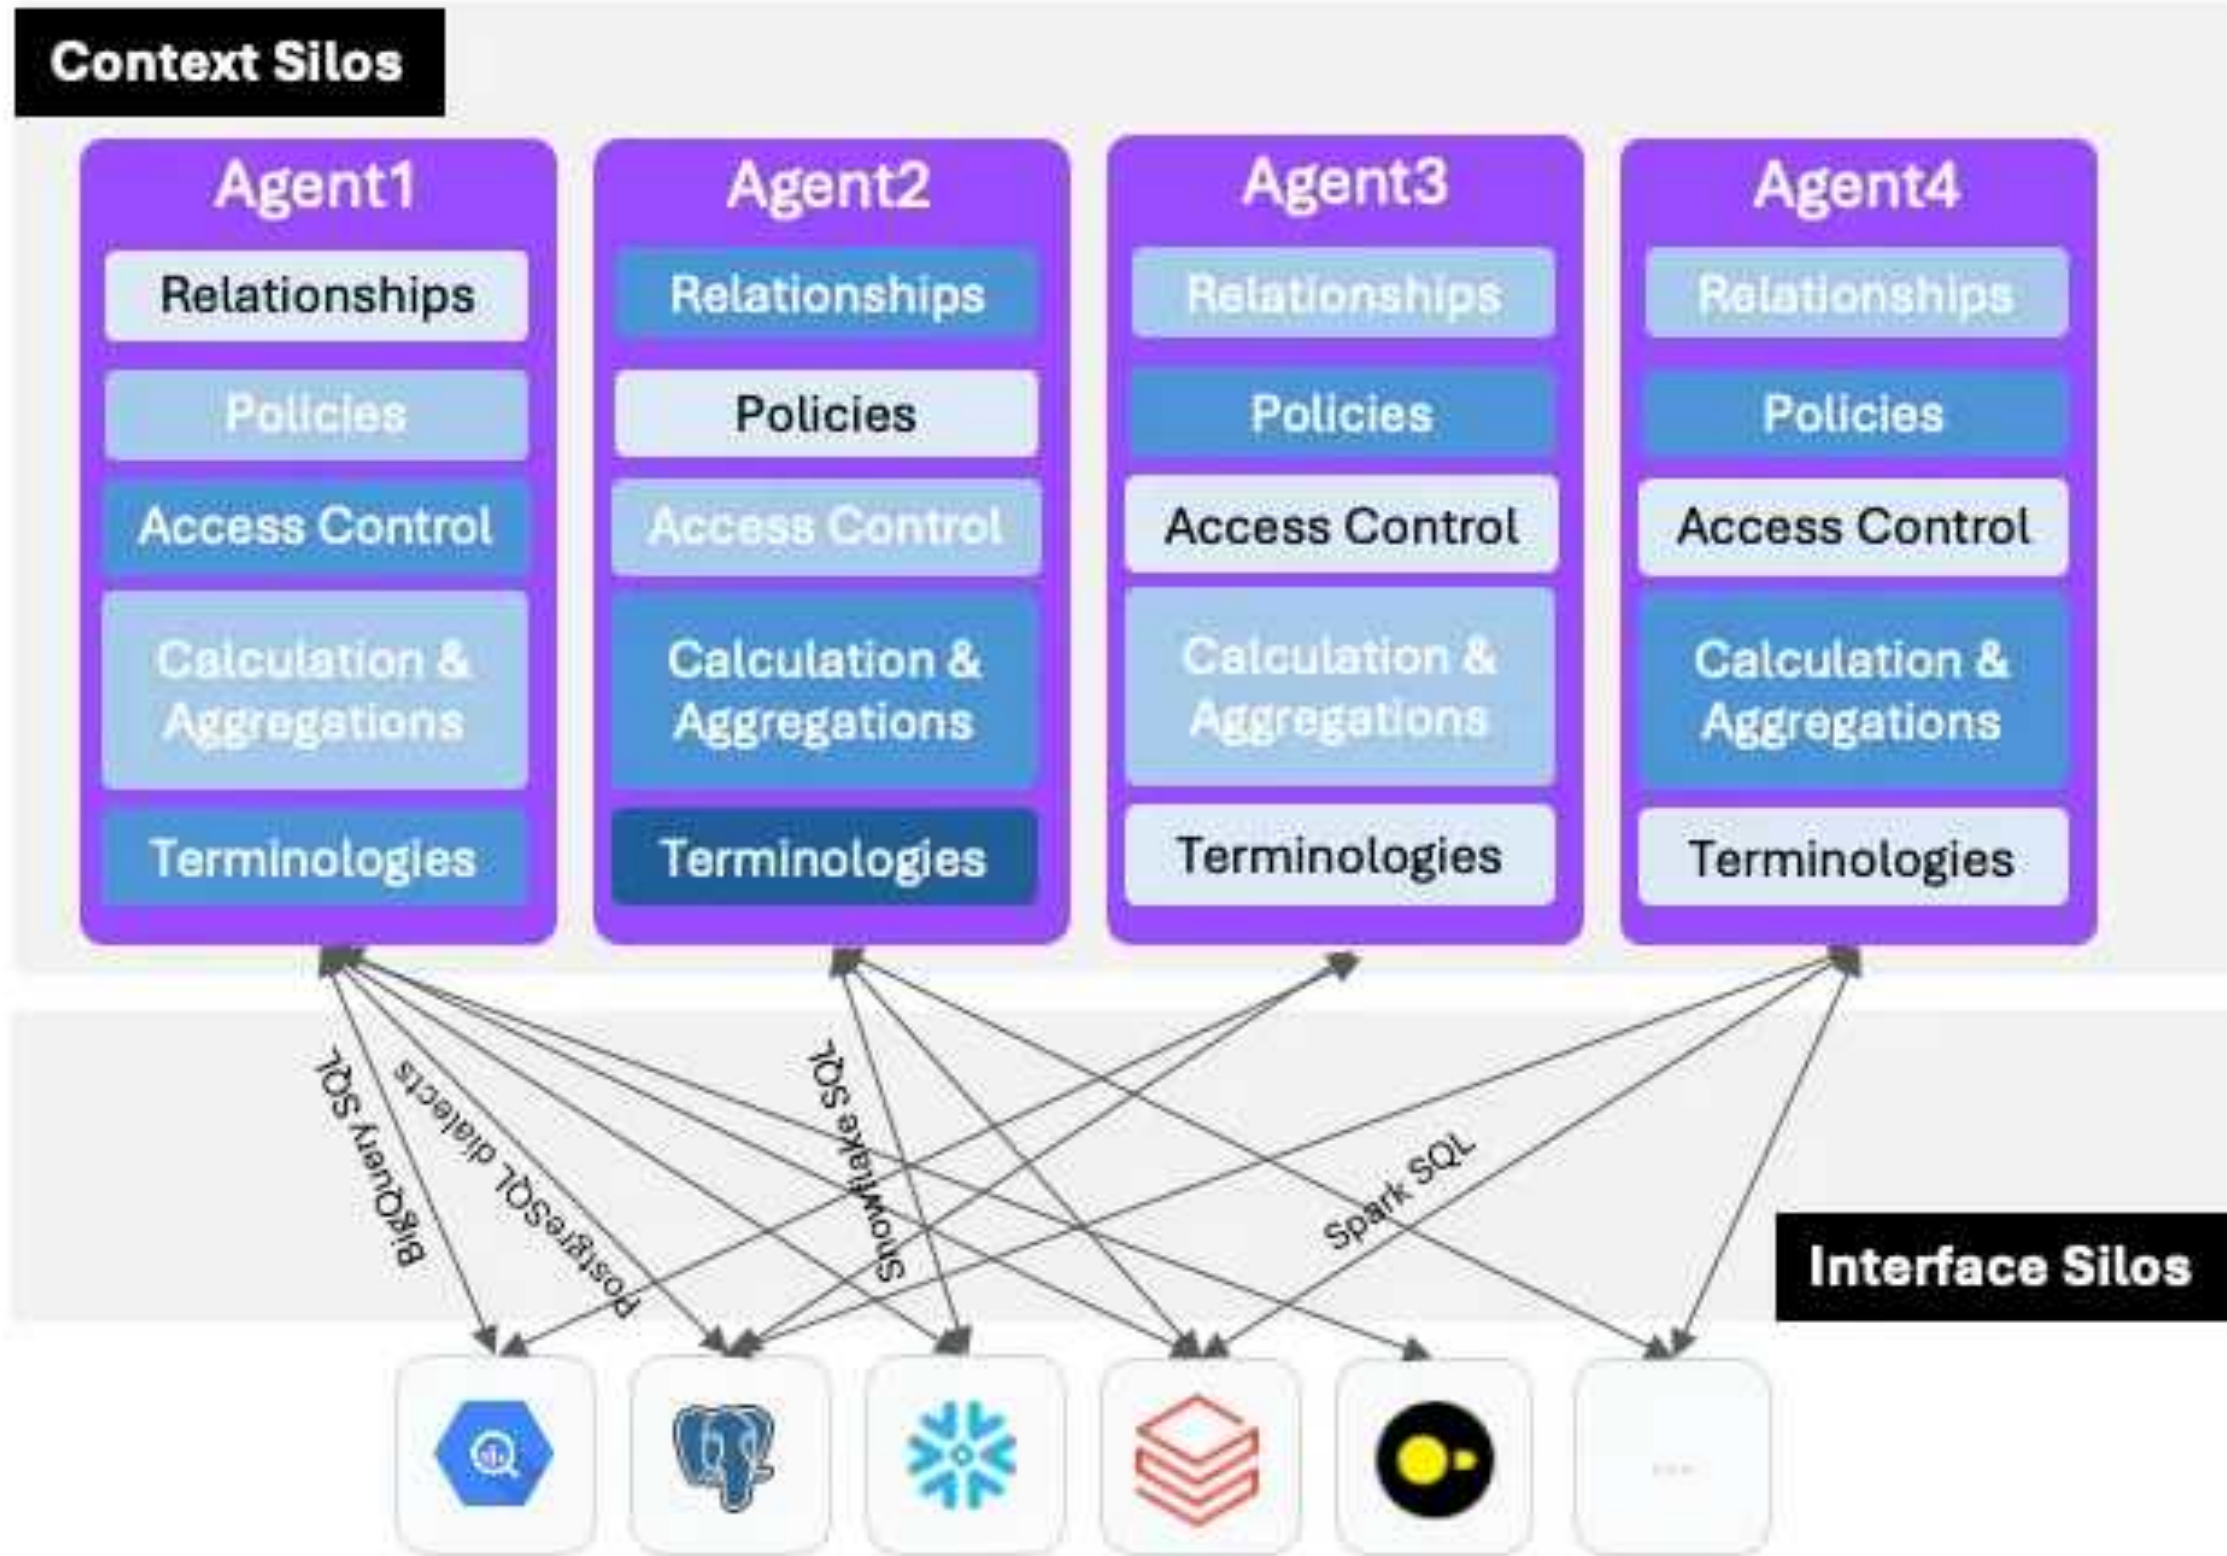


Subject: Electronics; **Subfield:** Analog electronics;
Image Type: Diagrams;
Difficulty: Hard

Language = it ALL

Today

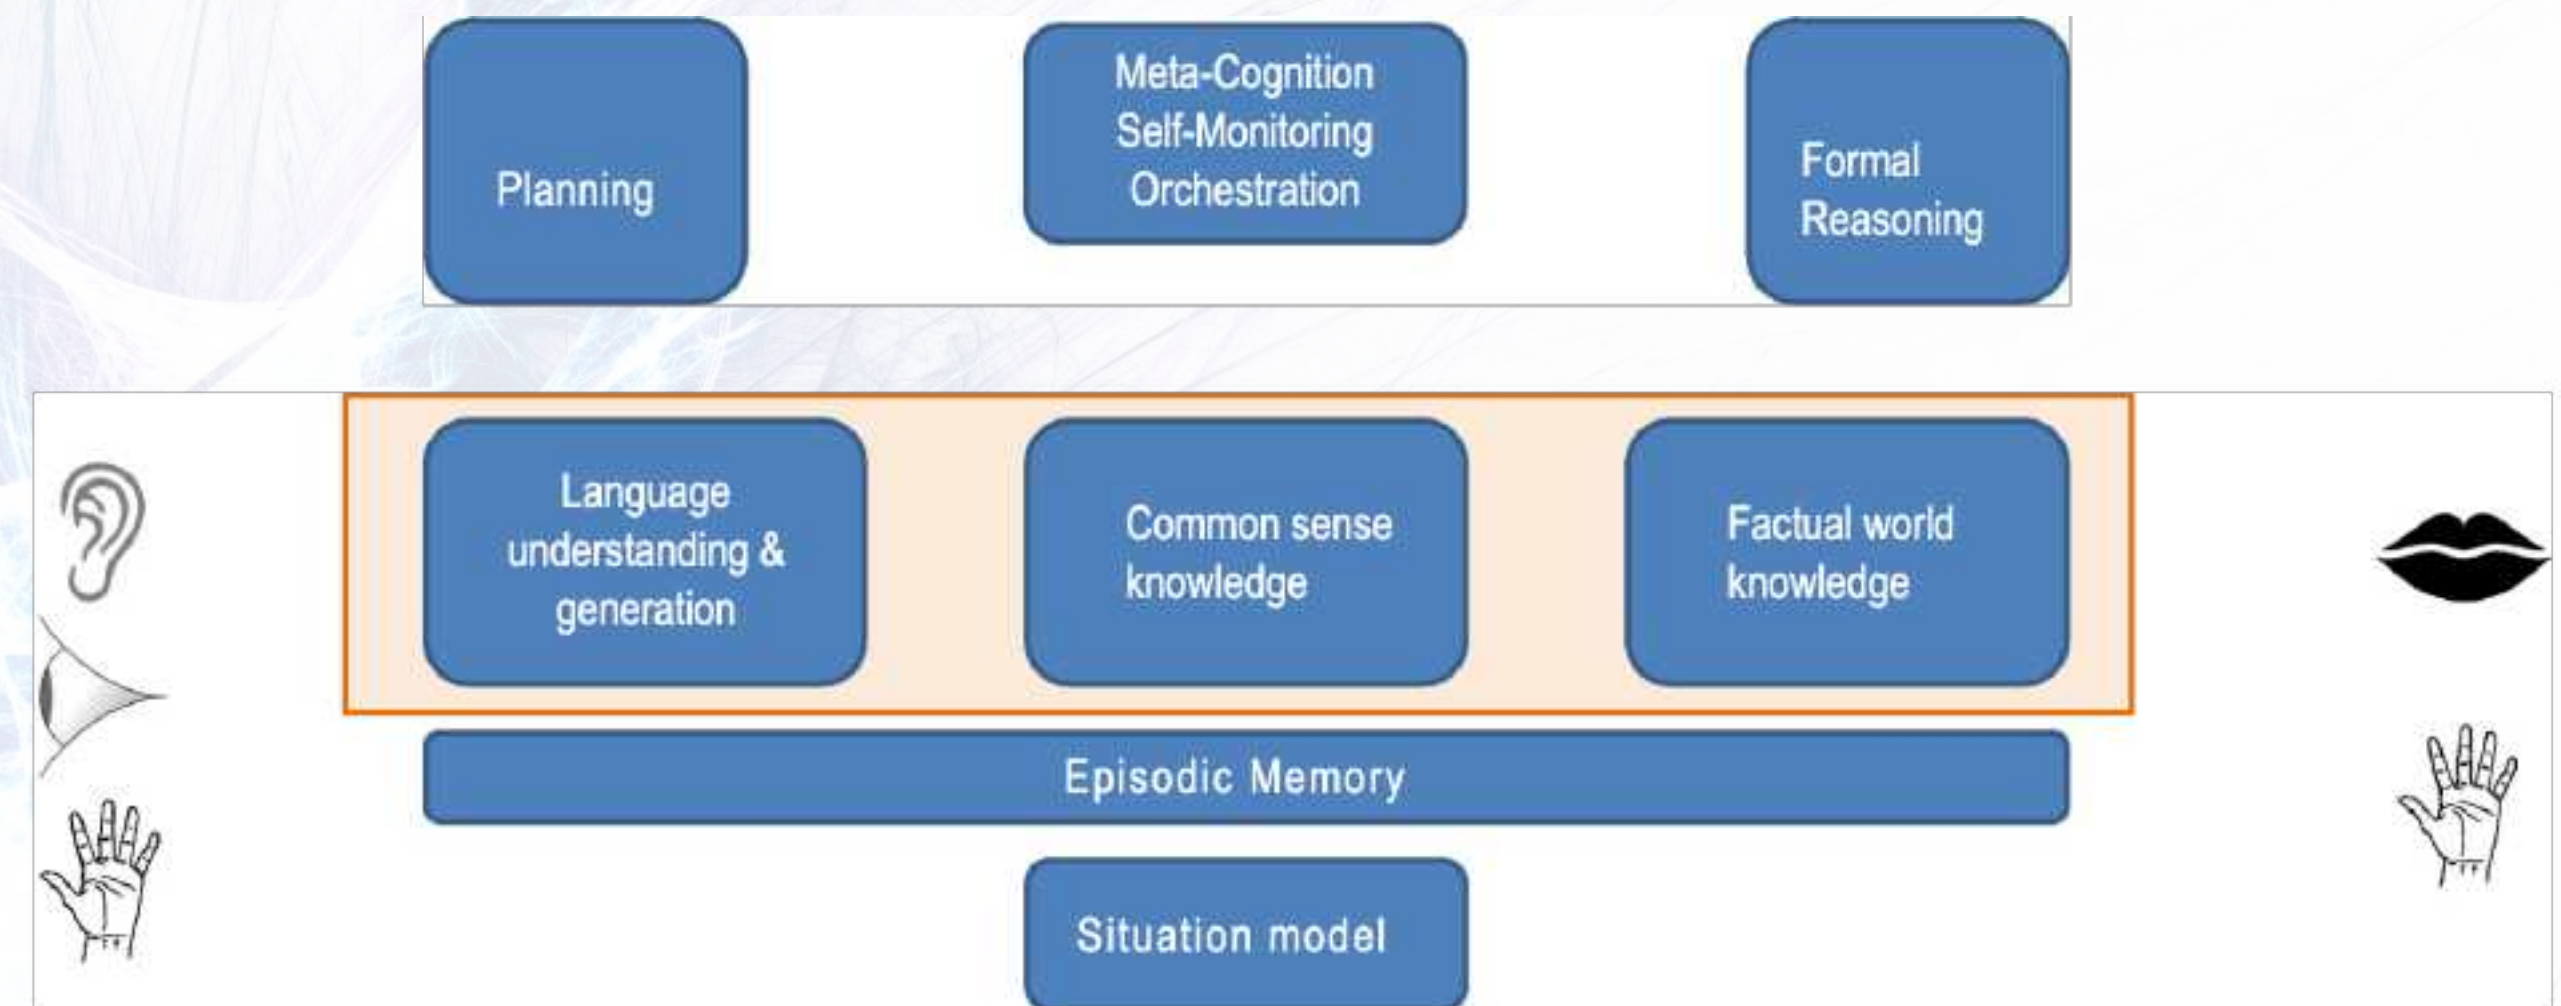
Future



There are many flaws in GPT's

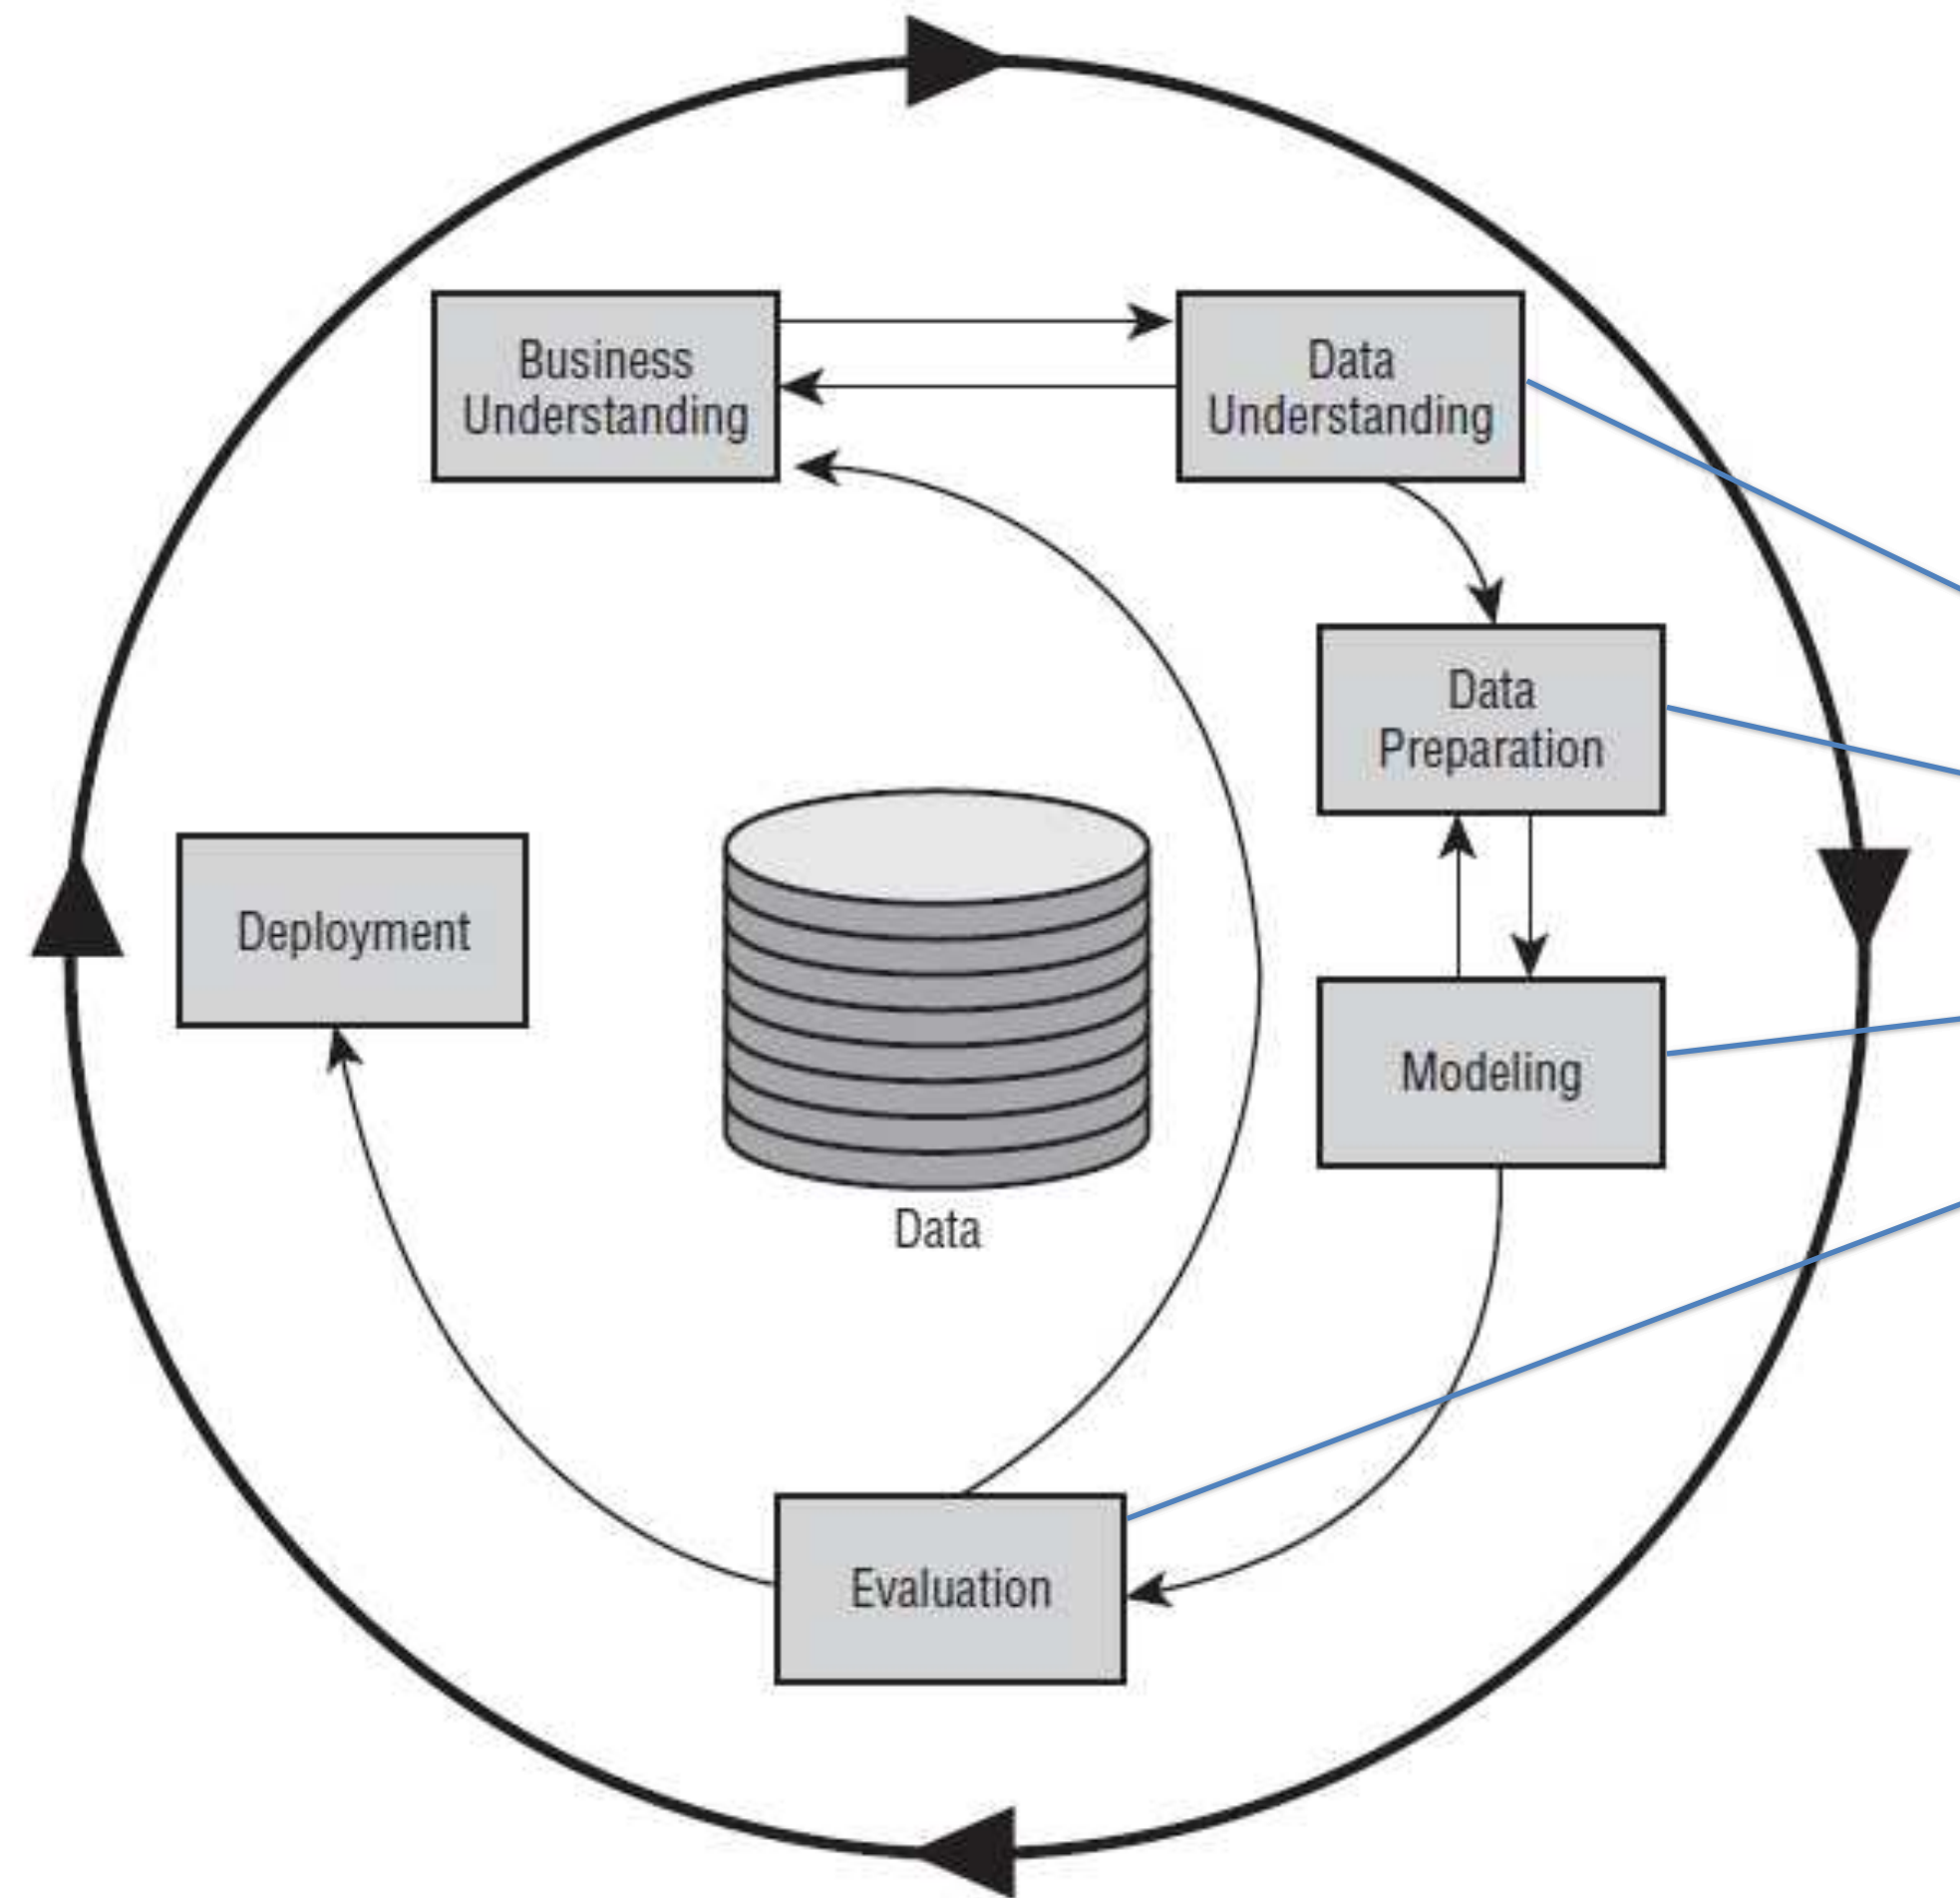
To fix some we need Cognitive architectures

- Separate Language Skill from Factual World Knowledge
- Represent world knowledge as a knowledge graph over an extensible ontology



Mahowald, et al. 2023 "Dissociating language and thought in large language Model: A cognitive perspective."

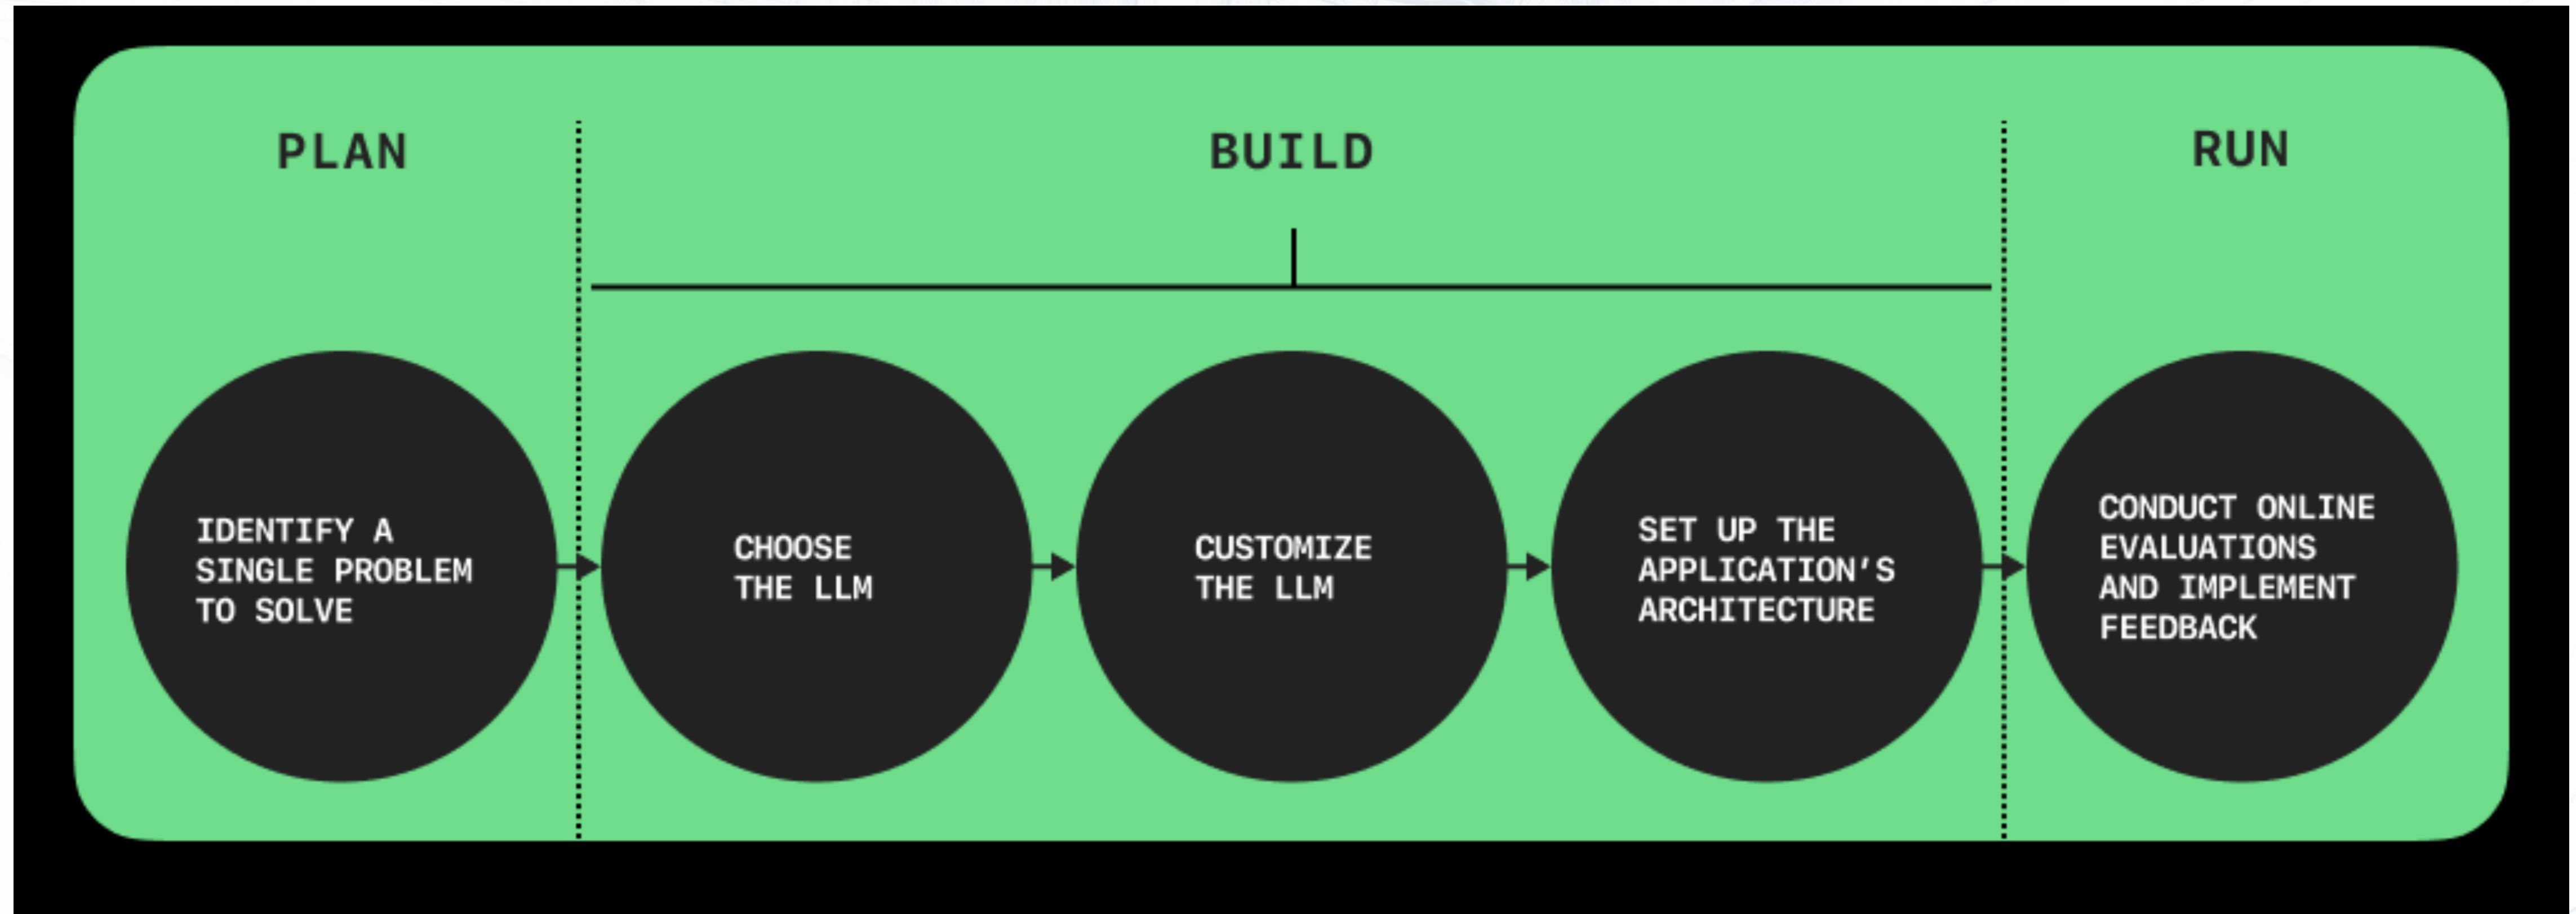
CRISPDM standard



Machine learning out of the box



Five steps to build an LLM application



Local LLM - RAG out of the box

The screenshot displays the Elephas AI interface. On the left, a sidebar lists various categories and documents. The main area shows a file browser with a grid of PDF documents. A chat window is open, showing a user query and a detailed AI response.

Files & Folders:

- Examining Beh... Locatelli.pdf
- Social Scin... oellad.pdf
- Kid-Whisper... Atla.pdf
- Beyond Partic... y AI Maas.pdf
- Sponsored New... amic Dash.pdf
- Problems with... a thechild.pdf
- Effectiveness... Ca stleman.pdf
- Facial Impres... la s_Wolfe.pdf
- Red-Teaming_I... AI...
- Scaling Laws... c ale_Diaz.pdf
- Estimating We... s Alcaraz.pdf
- Feasibility_I... cati ng_U.pdf
- Breaking Bias... Id ing_Raj.pdf
- What to Trust... e_Purves.p
- AI_Trustworth... Bisconti.pdf
- AIES 2024 Con... # Prell

Chat Window:

Query: are there copliance issue

Response:

Yes, there are compliance issues highlighted in the provided contexts. Here are some key points regarding compliance challenges faced by organizations:

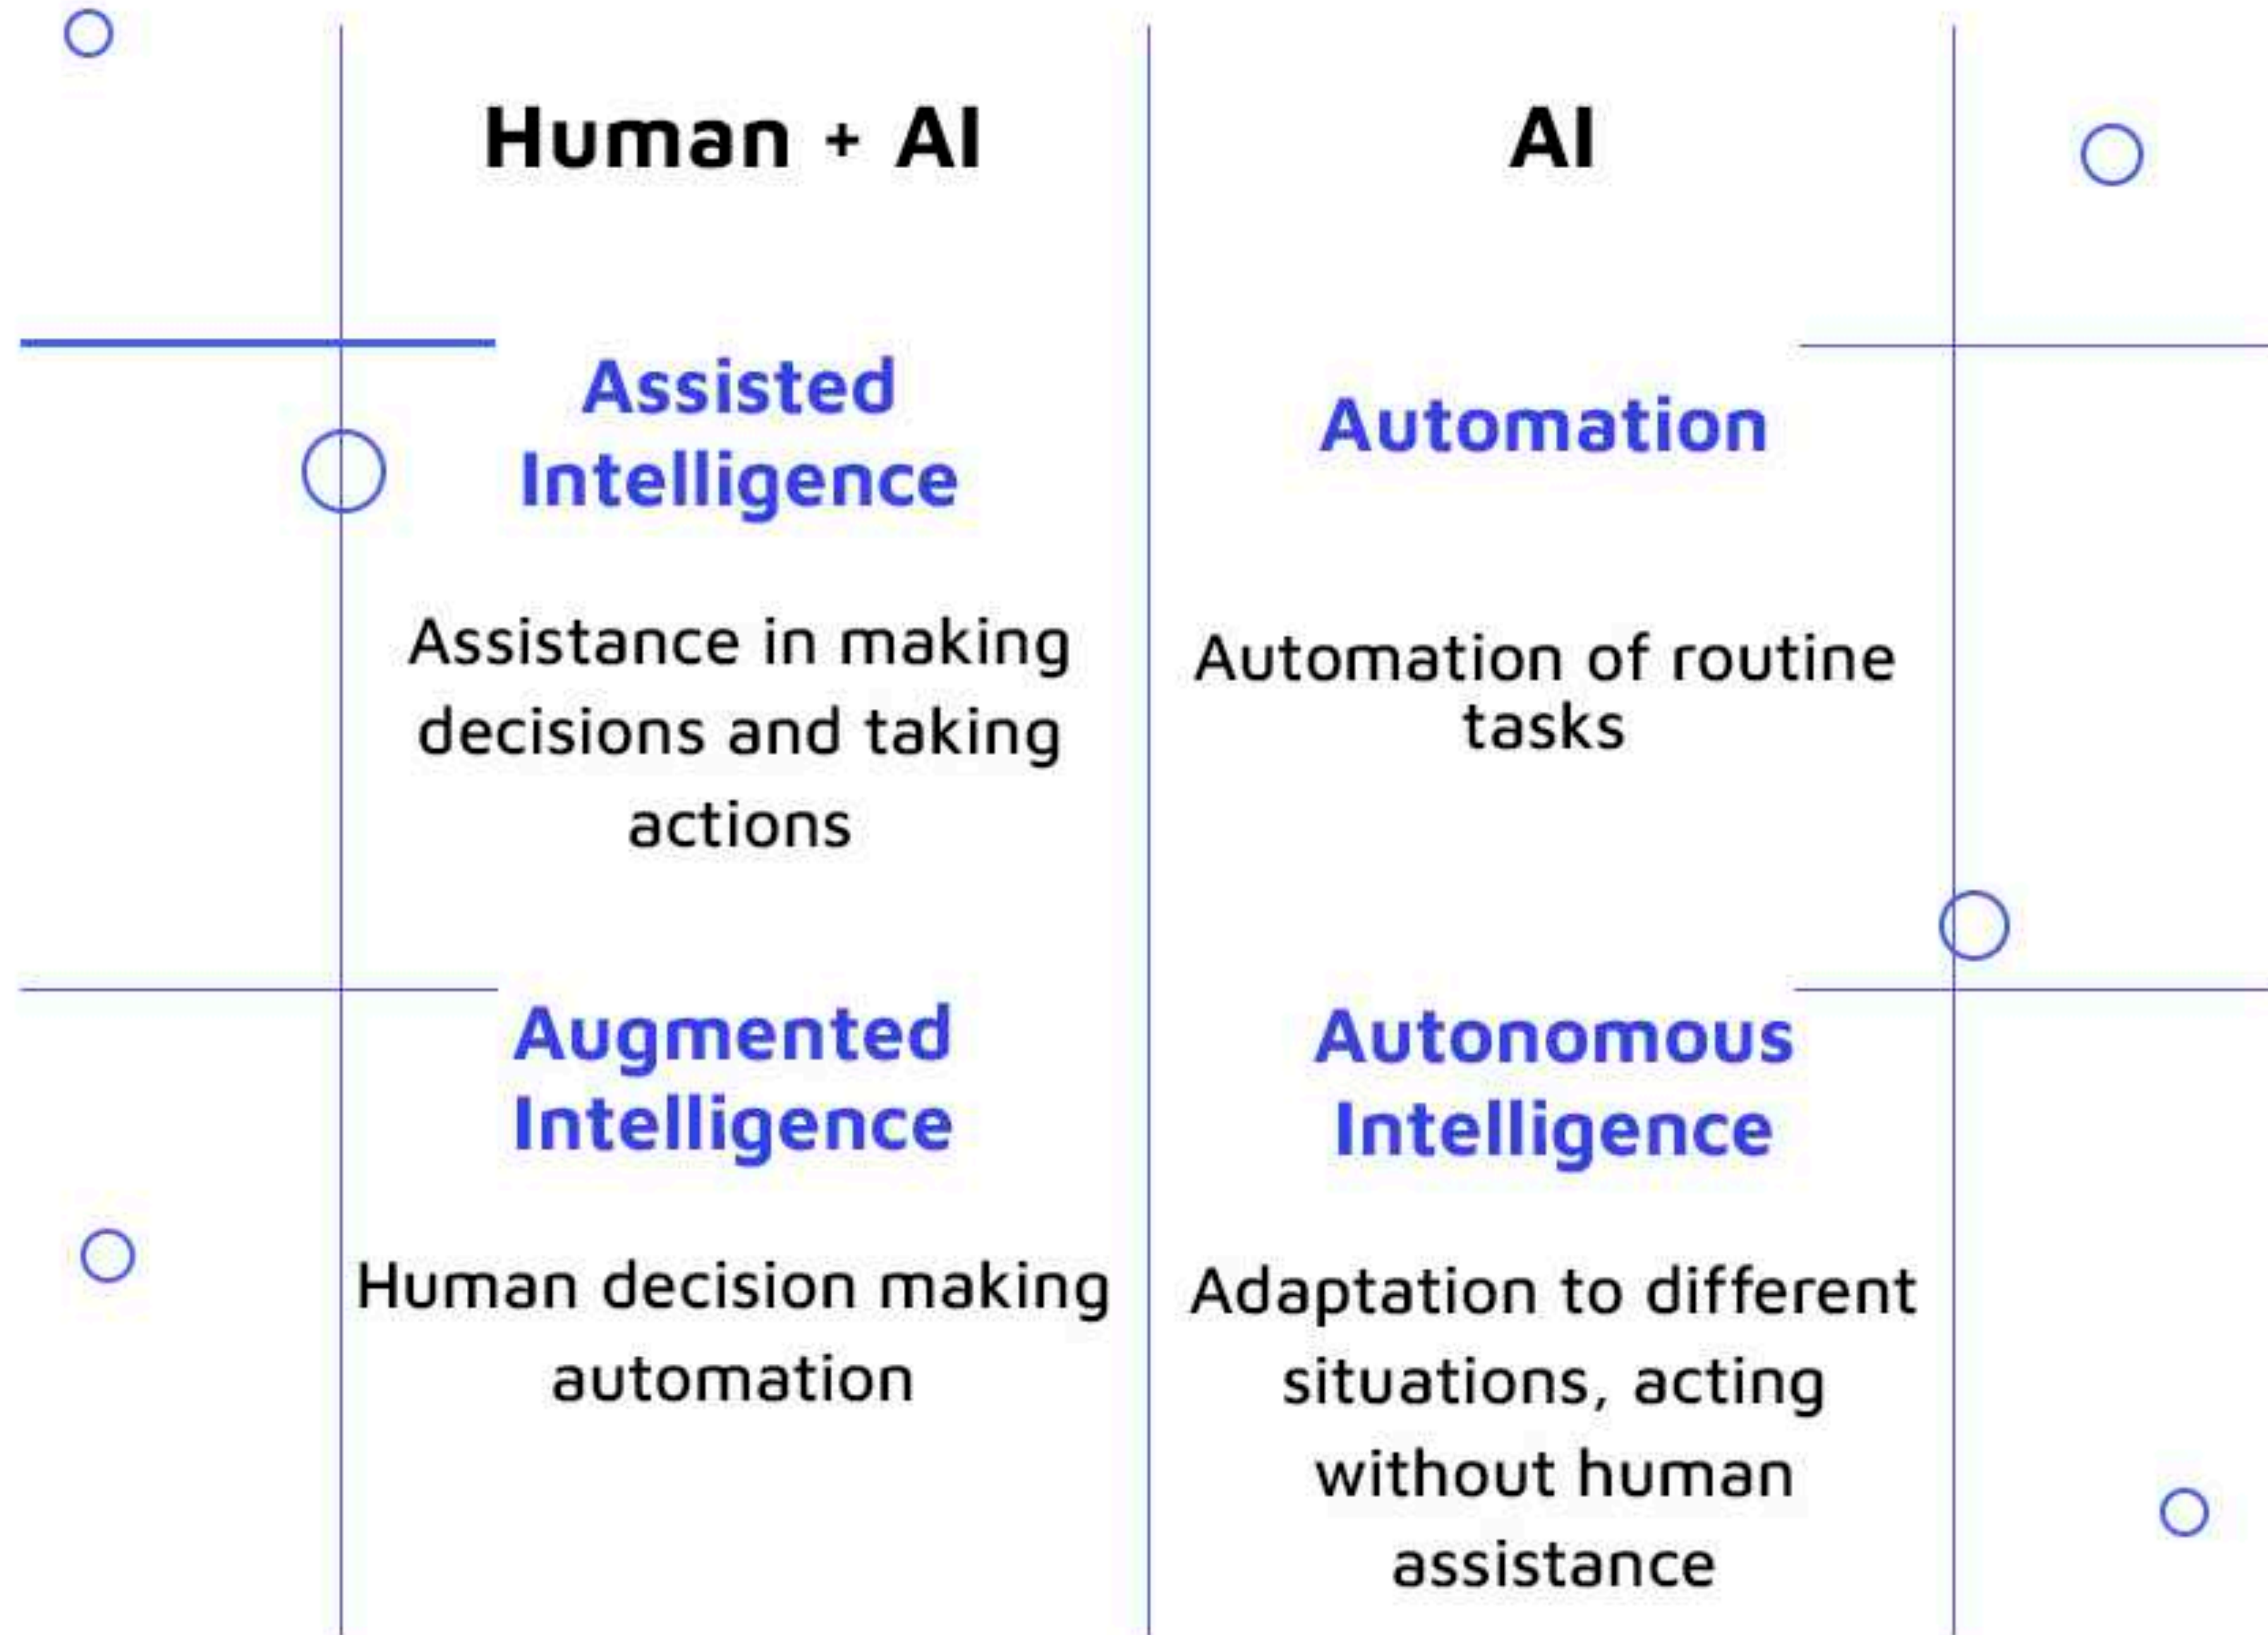
- Balancing Pressures:** Organizations are often caught in a balancing act among multiple pressures, such as technological innovation, value creation, and regulatory compliance. For instance, Firm 1 has a compliance score of 0.75, indicating a relatively high level of compliance, but still faces challenges in integrating compliance processes effectively. On the other hand, Firm 2, a smaller startup, has a lower compliance score of 0.57 and tends to prioritize product commitments over compliance (Context 3, Context 13).
- Exaggeration of Compliance:** Many organizations tend to exaggerate their actual compliance levels. For instance, Firm 2 has an exaggeration ratio of 1.34, indicating that it has significantly overstated its compliance capabilities. This is often driven by the urgency to demonstrate innovative AI products, leading to neglect of fundamental compliance tasks (Context 15, Context 21).
- Lack of Established Processes:** Some firms, such as Firm 1a, have no established AI compliance processes, leading to disrupted workflows and decreased productivity. This lack of structured compliance mechanisms can hinder the organization's ability to meet regulatory requirements effectively (Context 5, Context 12).
- Resource Limitations:** Smaller organizations often lack the resources necessary for rigorous compliance and reporting activities. This can lead to lower actual compliance scores and a greater tendency to exaggerate compliance levels (Context 14, Context 20).
- Regulatory Pressures:** Larger organizations typically experience more regulatory pressure, which influences their internal priorities regarding compliance. The findings suggest that there is a correlation between the priority to comply with regulations and actual compliance levels, particularly in smaller organizations that may not prioritize compliance as highly (Context 25).
- Norm Inconsistency:** There are also discussions about norm inconsistency in decision-making processes, particularly in high-risk settings, which can complicate compliance efforts and lead to misalignment between AI systems and real-world facts (Context 19).

In summary, compliance issues are prevalent across different organizations, influenced by factors such as size, resource availability, regulatory pressures, and the inherent challenges of integrating compliance into innovative processes.

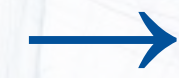
What competences do we need?

- Language
- Systems thinking
- Conceptual thinking
- Problem decomposition
- Understanding and valuing the outcome of models
- Reality checks
- Critical thinking





Machine Learning



Learning Machines

Automated knowledge
discovery
(data mining, prediction)

Computers program
themselves on hard tasks
(machines learn & adapt)

Algorithms can create algorithms
AI can create AI

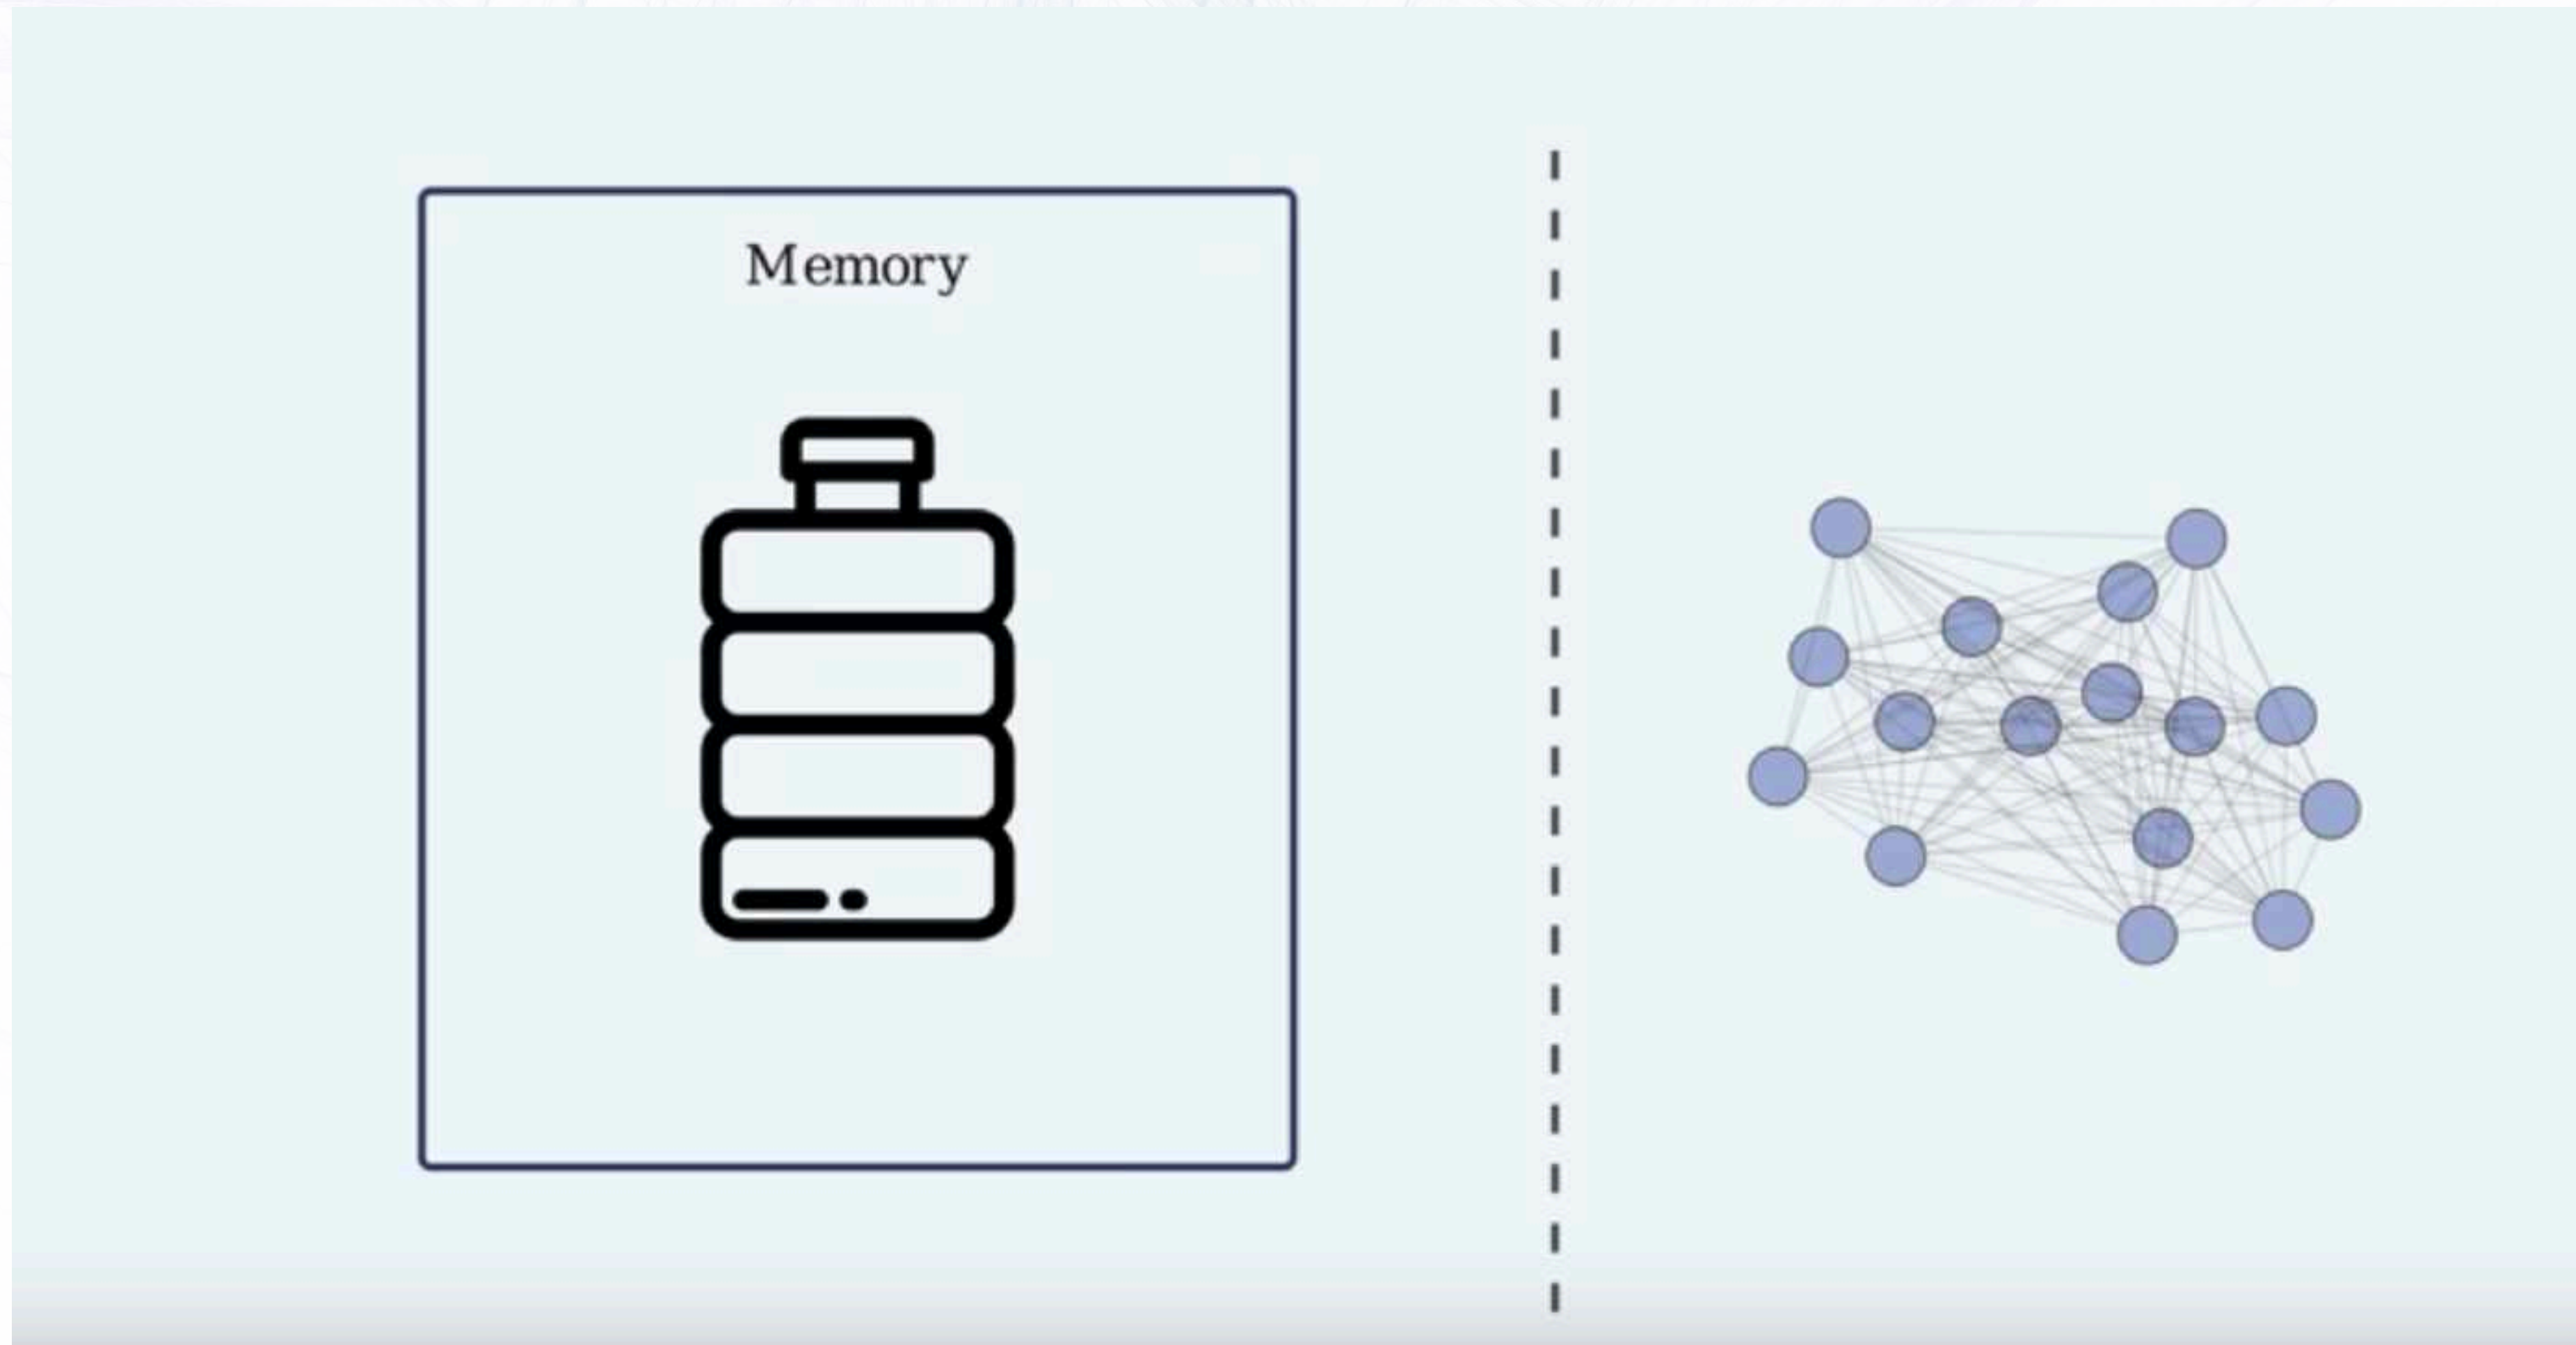
How to regulate **autonomy**?

How to regulate **adaptivity**?

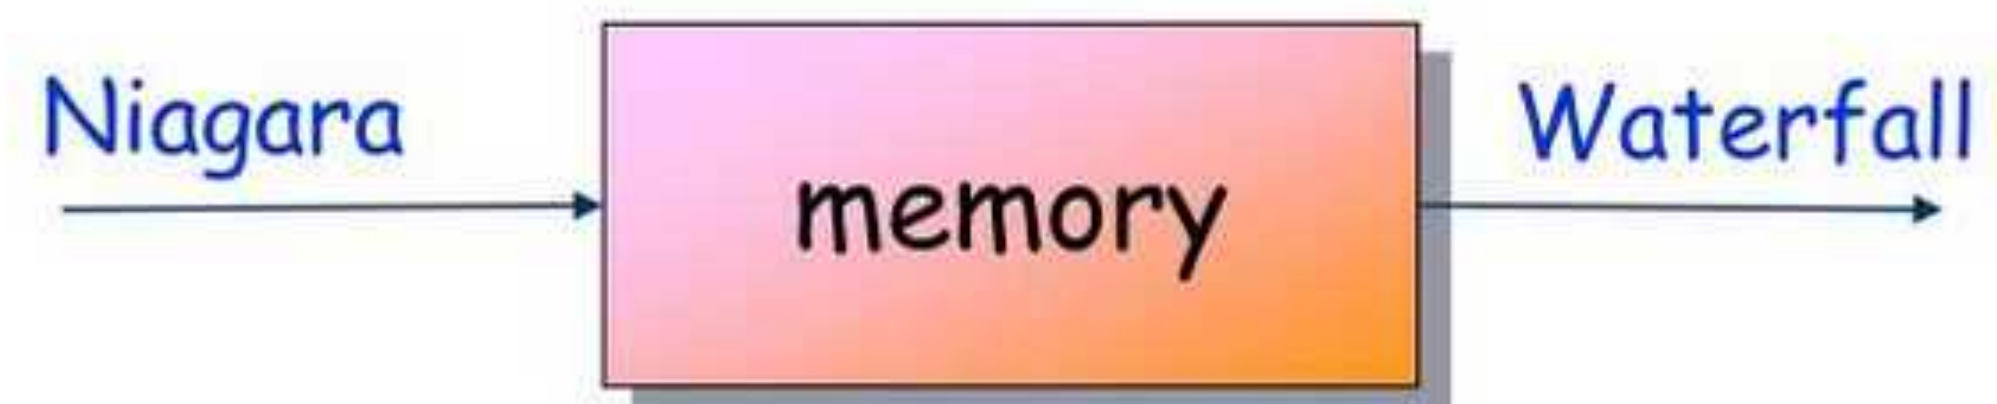
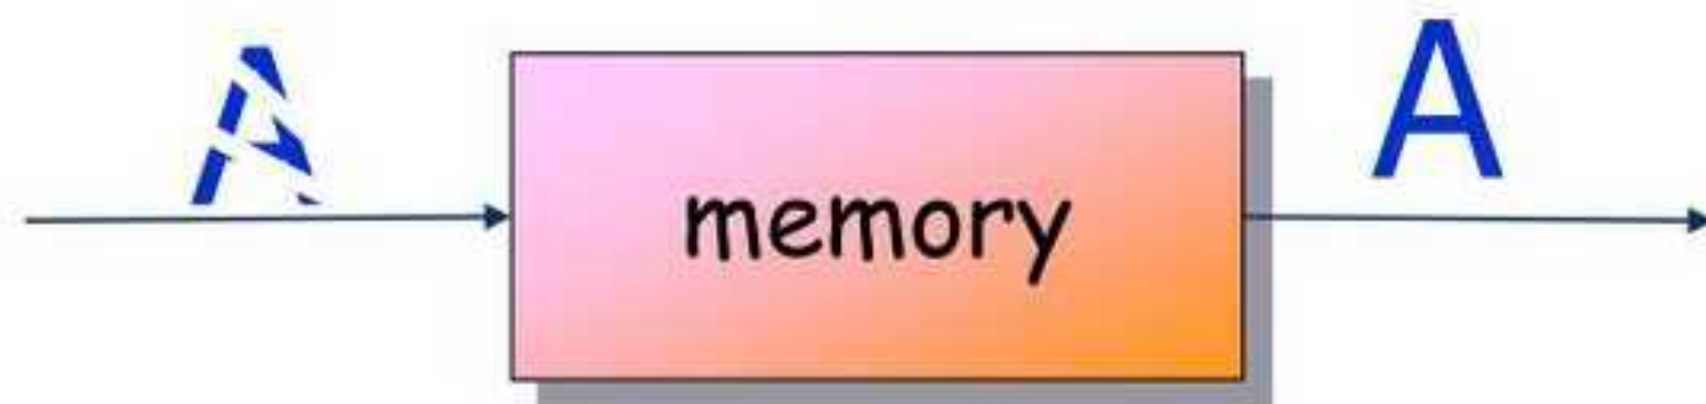
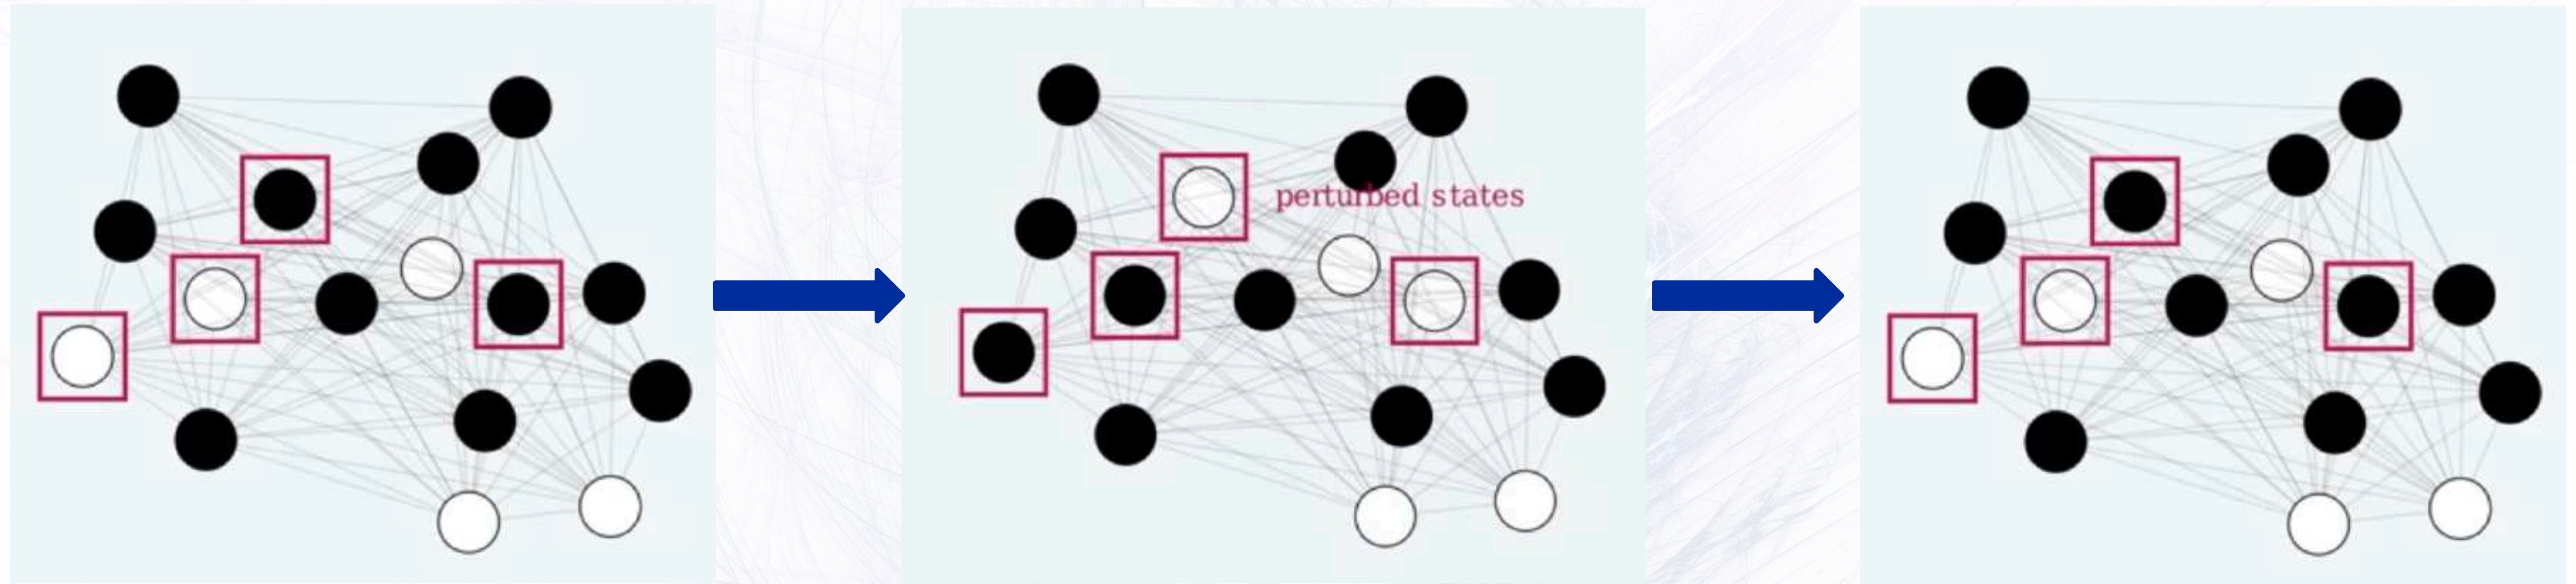
7 Strategies to control Models / Data

People	Process	Technology
<i>Educate</i> On every level in every area	<i>Embrace Risk Management</i> Be as adaptive as possible	<i>Centralize infrastructure</i> A solid technical foundation
<i>Hire the right people</i> Power of the teams	<i>Get organized</i> Structure and standardize	<i>Use Technology to fix technology</i> Security and fairness by design
<i>Formalize</i> Take position		

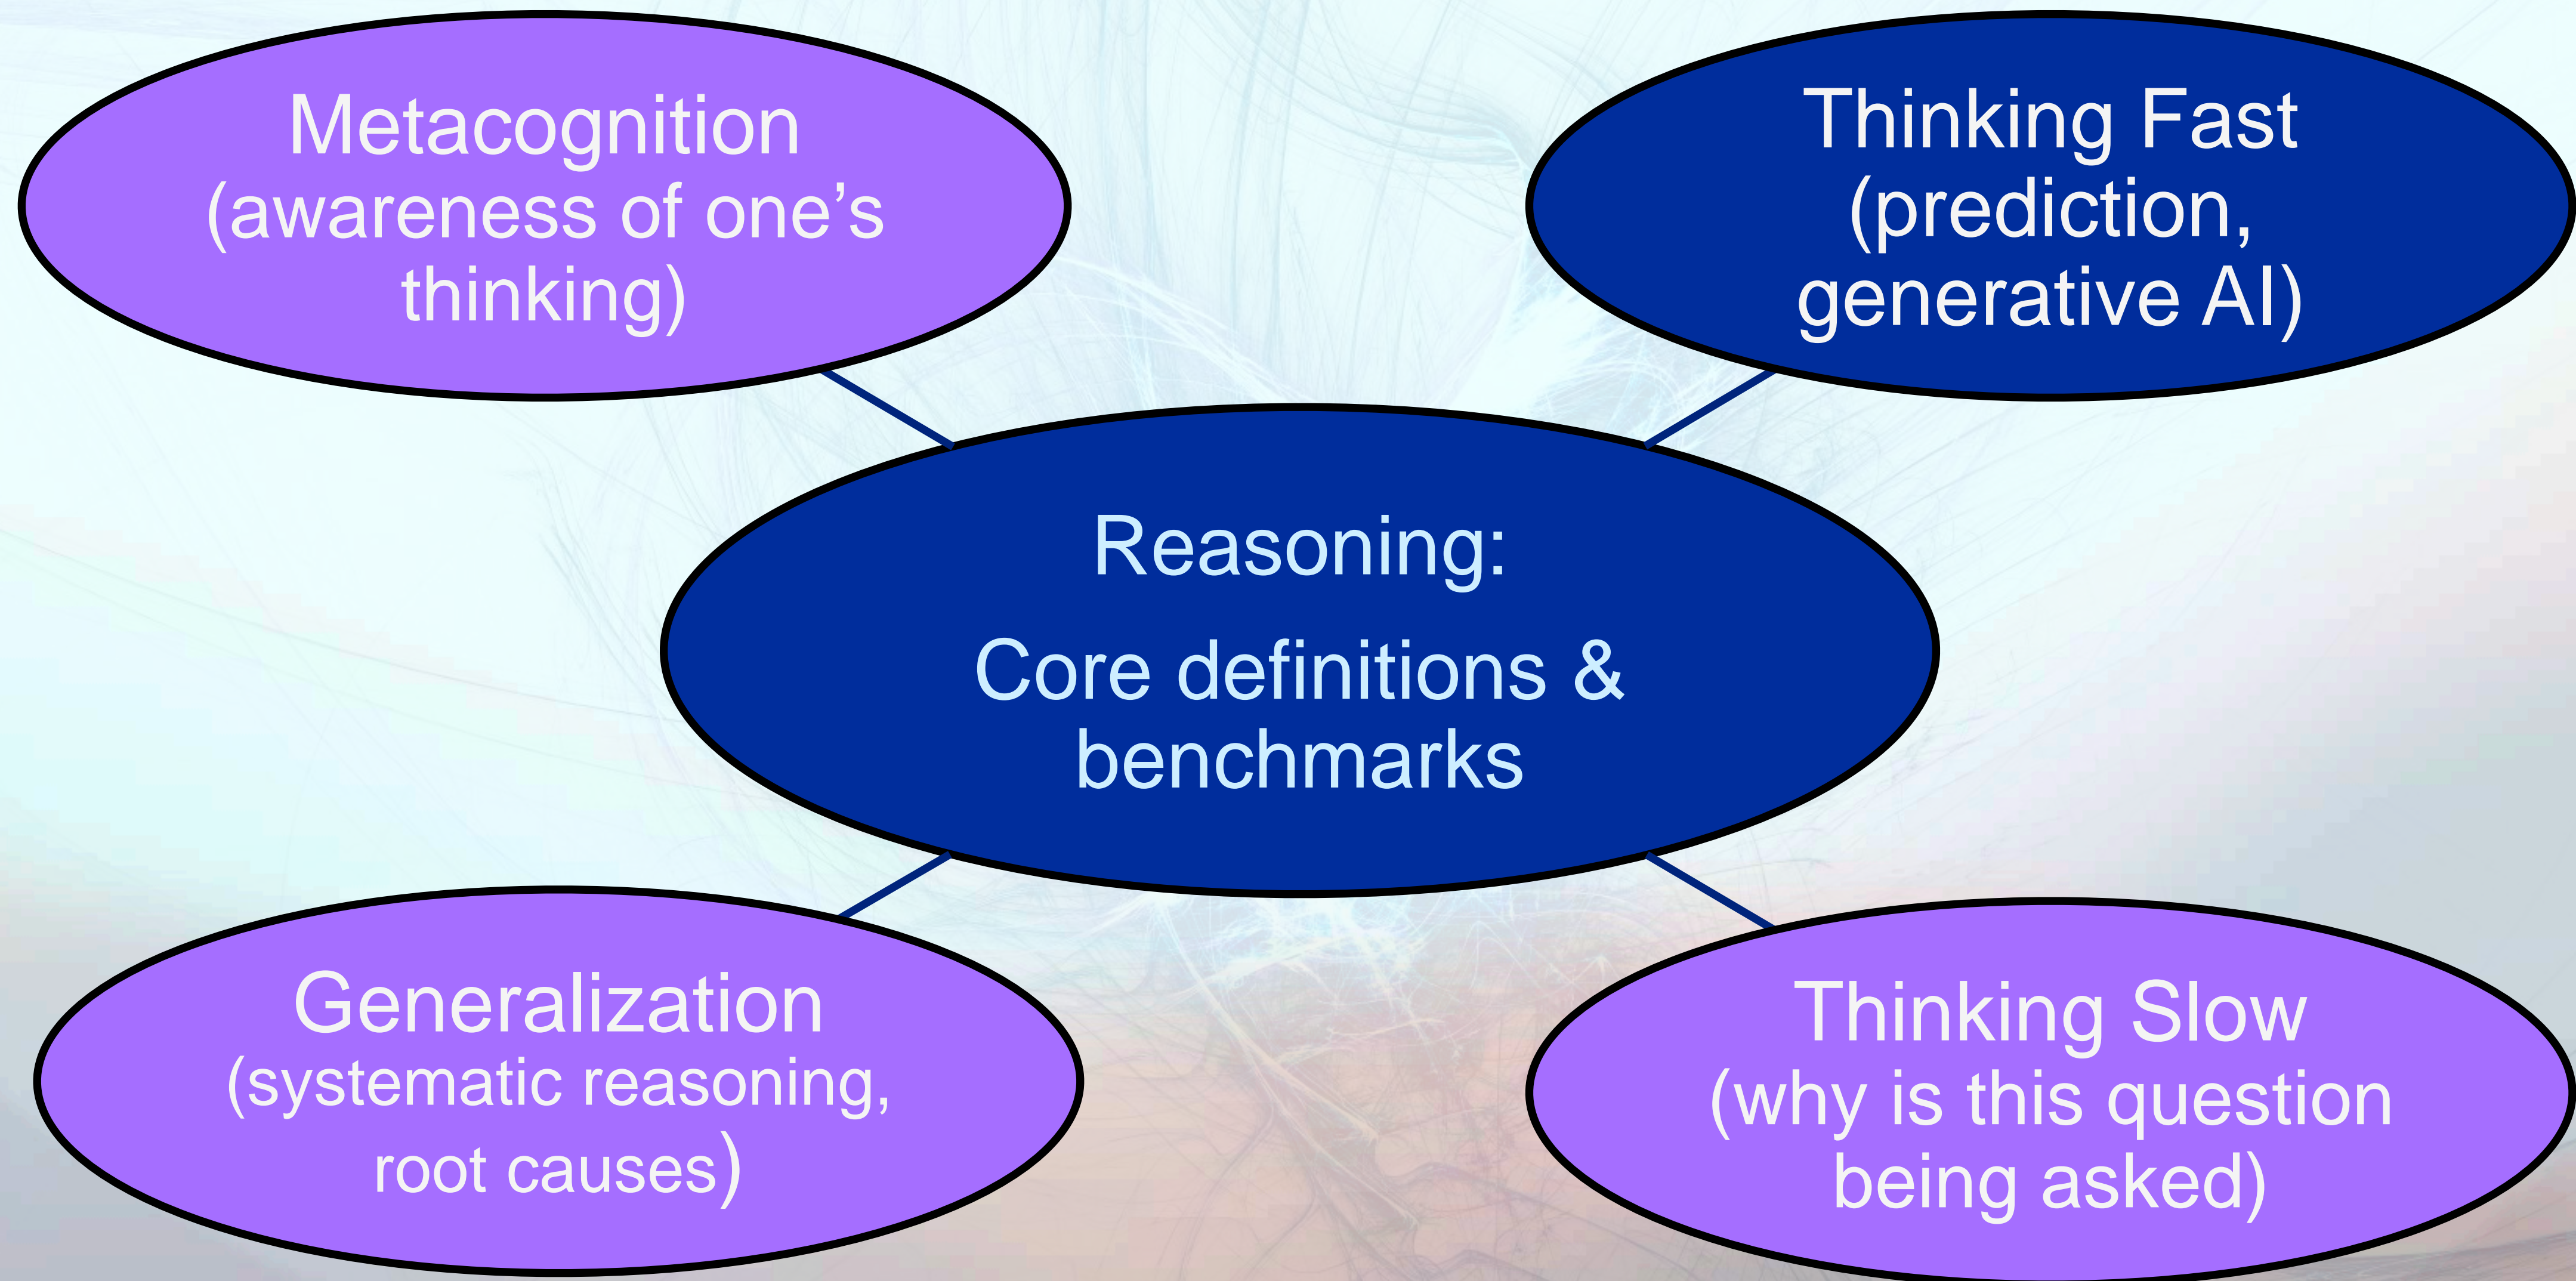
AI Futures: Bio inspired: Associative Memory (Hopfield networks)



AI Futures: Bio inspired: Associative Memory (Hopfield networks)



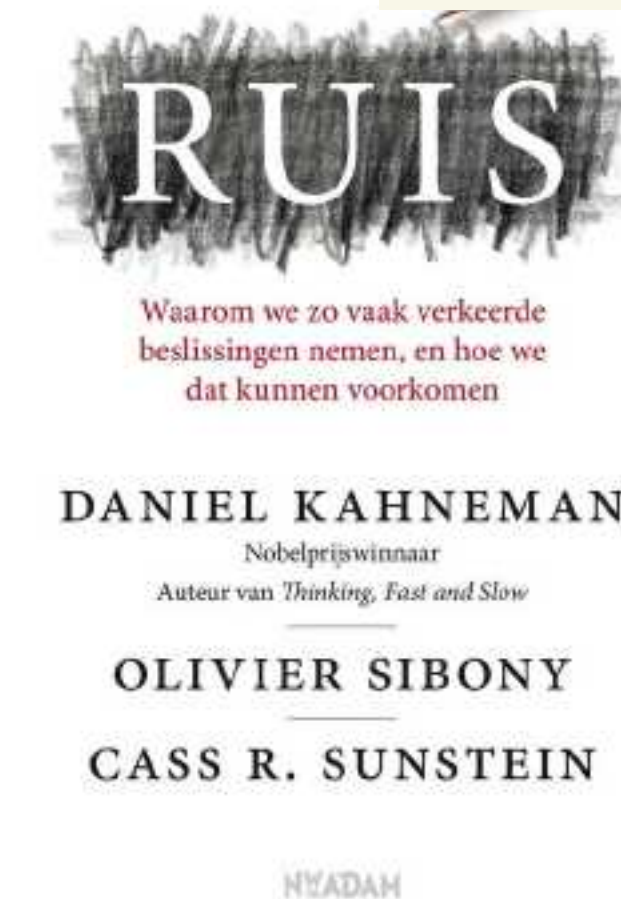
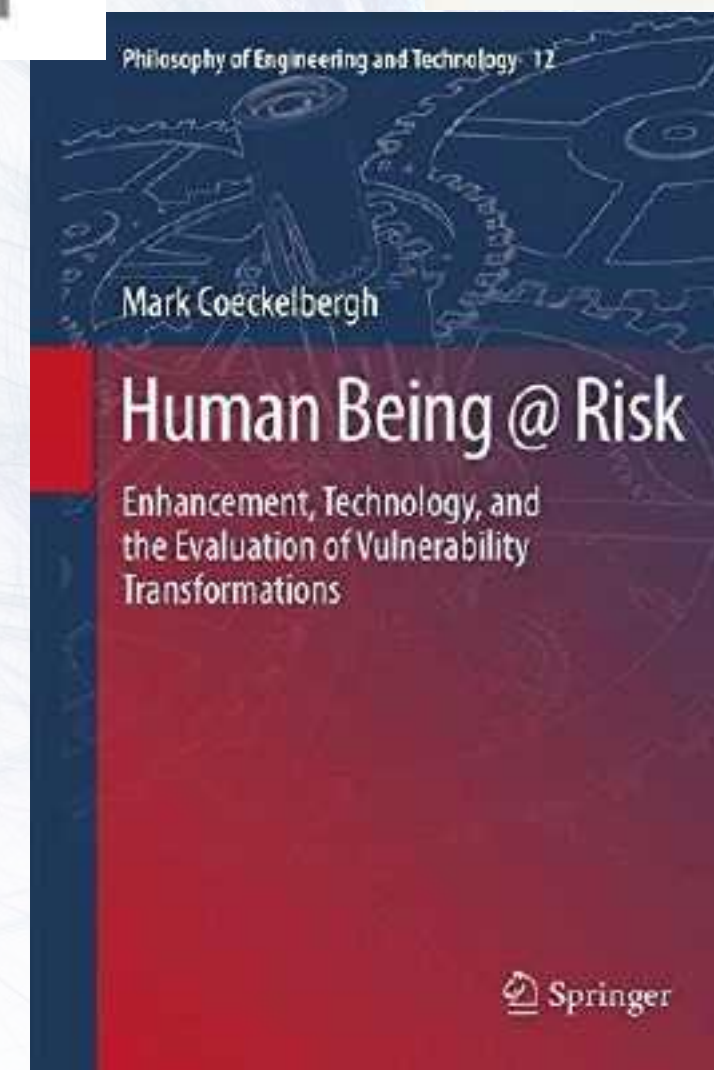
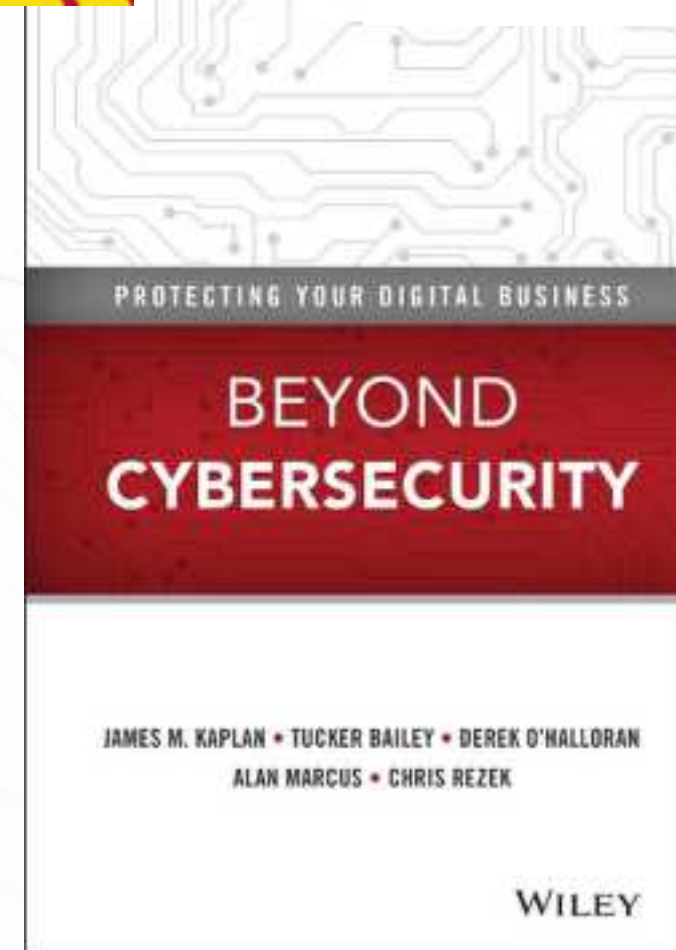
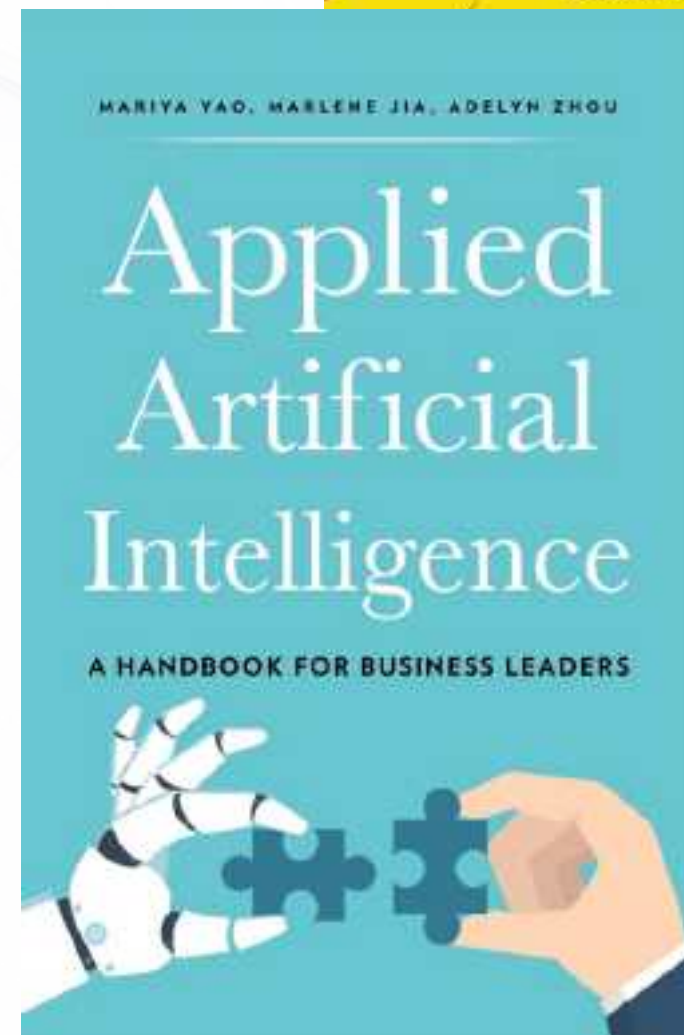
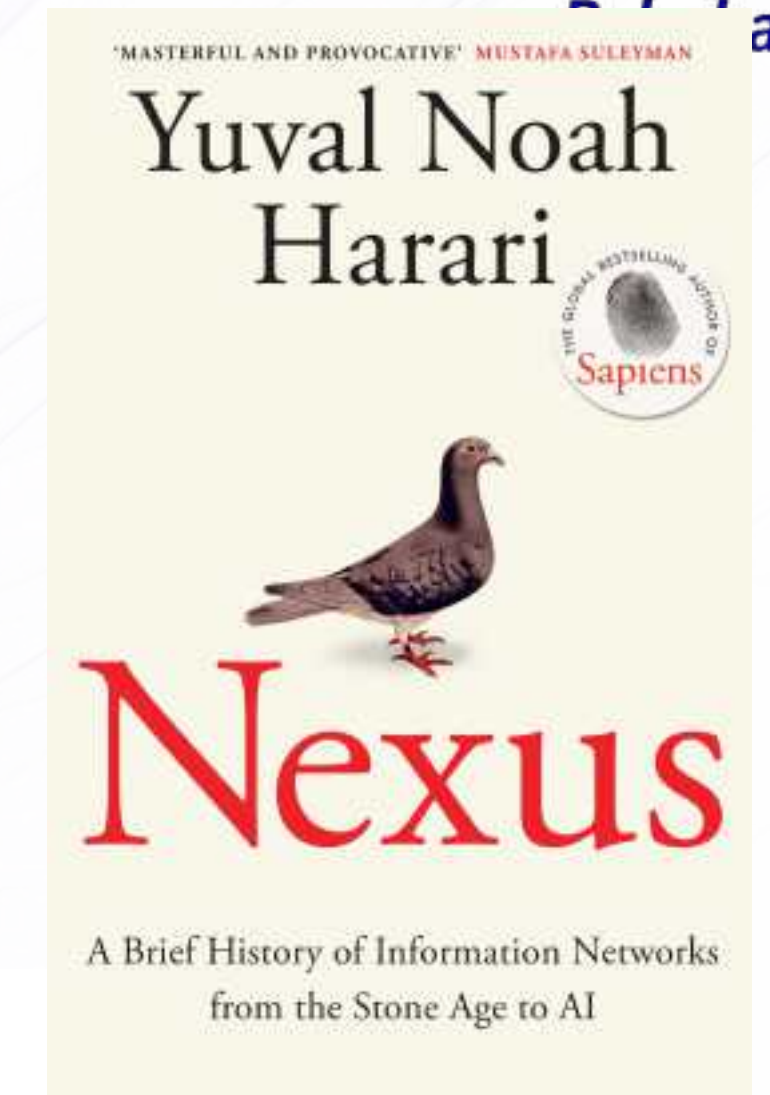
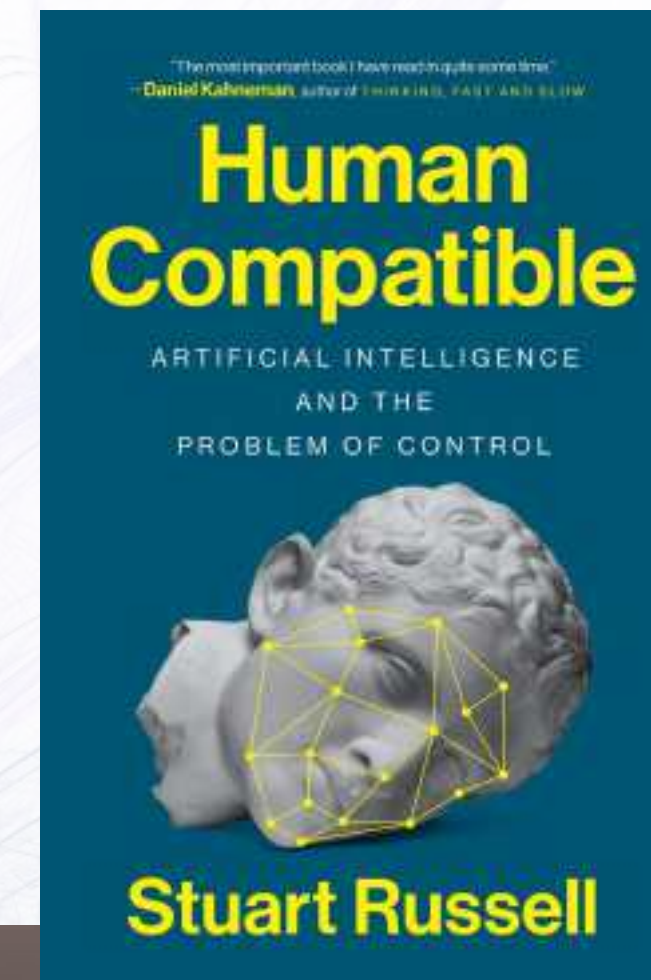
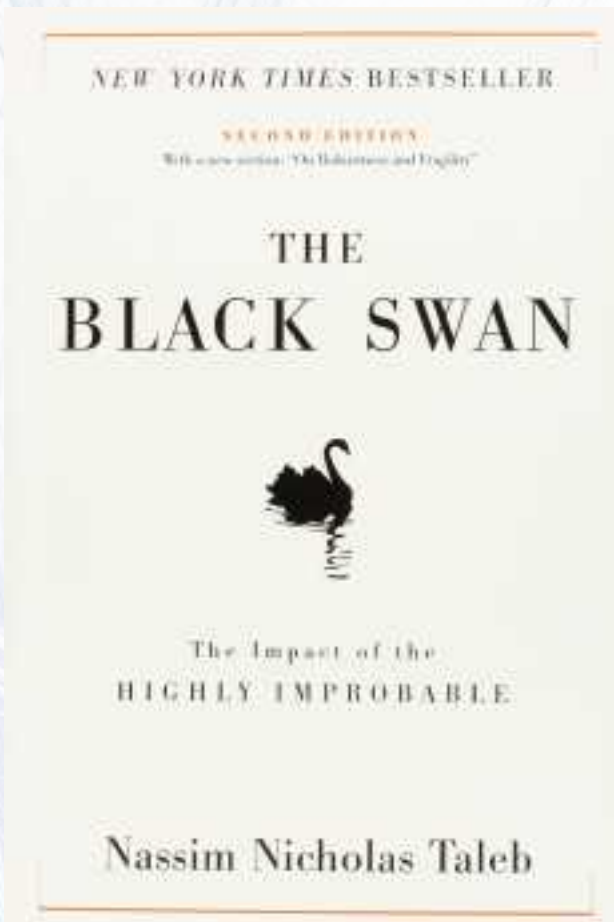
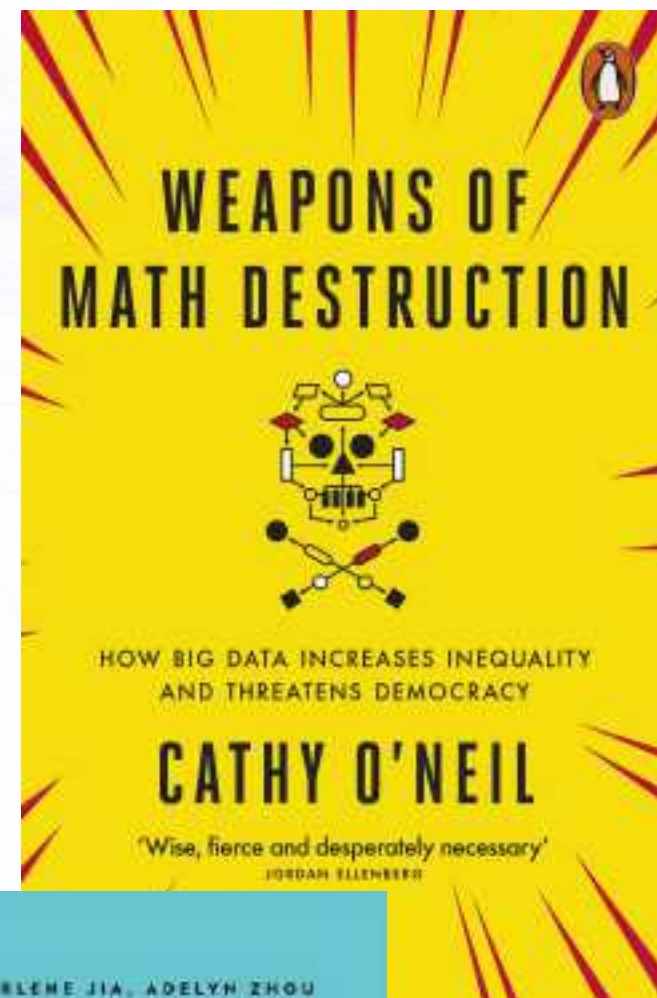
AI Futures - Bio Inspired



Some of my inspirations



ank



Technology will change our society

How to establish digital trust and
sane human computer
relationships?

That there will be impact is for
sure, how you address these
questions is a strategic question
for you and your company.

Start working on it today!

Make it safe, secure, meaningful
and fun.



AI tools

- **Classifying AI:**
 - BigML - <https://bigml.com/>
 - Application: <https://ndalynn.com/>
- **Generative AI:**
 - Elphas - <https://elephas.app/>
 - TypingMind - <https://www.typingmind.com/>
 - BoltAI - <https://boltai.com/>
 - PDFPals - <https://pdfpals.com>
 - GPT4ALL - <https://www.nomic.ai/>



References

- The AI Dilemma:
 - <https://www.youtube.com/watch?v=x0VJKj8lcNQ>
- The Social Dilemma:
 - <https://www.youtube.com/watch?v=iYvVgGWUKKg>
- AI/ML non-coding platform <https://www.bigml.com>
- Lori on Data science myth
 - <https://www.youtube.com/watch?v=u4vxoufe2kQ&feature=youtu.be>
- Lanier on data
 - <https://www.youtube.com/watch?v=Np5ri-KktNs>
- Edsger Dijkstra
 - <https://www.youtube.com/watch?v=mLEOZO1GWVc>
- Grio AI Ideation Canvas <https://www.grio.nl/AICanvas/>
- GenAI open source: GPT4ALL <https://www.nomic.ai/>



Visit www.grio.nl

