DS323: AI in Design (AIID)

Autumn 2023

Week 01 Lecture 01
AI + Design

Wan Fang

Southern University of Science and Technology

Welcome to DS323

AI in Design

Agenda

- Why should we care about AI?
  - Hear what people from different sectors talk about AI
  - Who are the stakeholders?
  - What is Artificial Intelligence?
- Course logistics

AI and Healthcare Industry

AI and Automotive Industry

AI and Bussiness

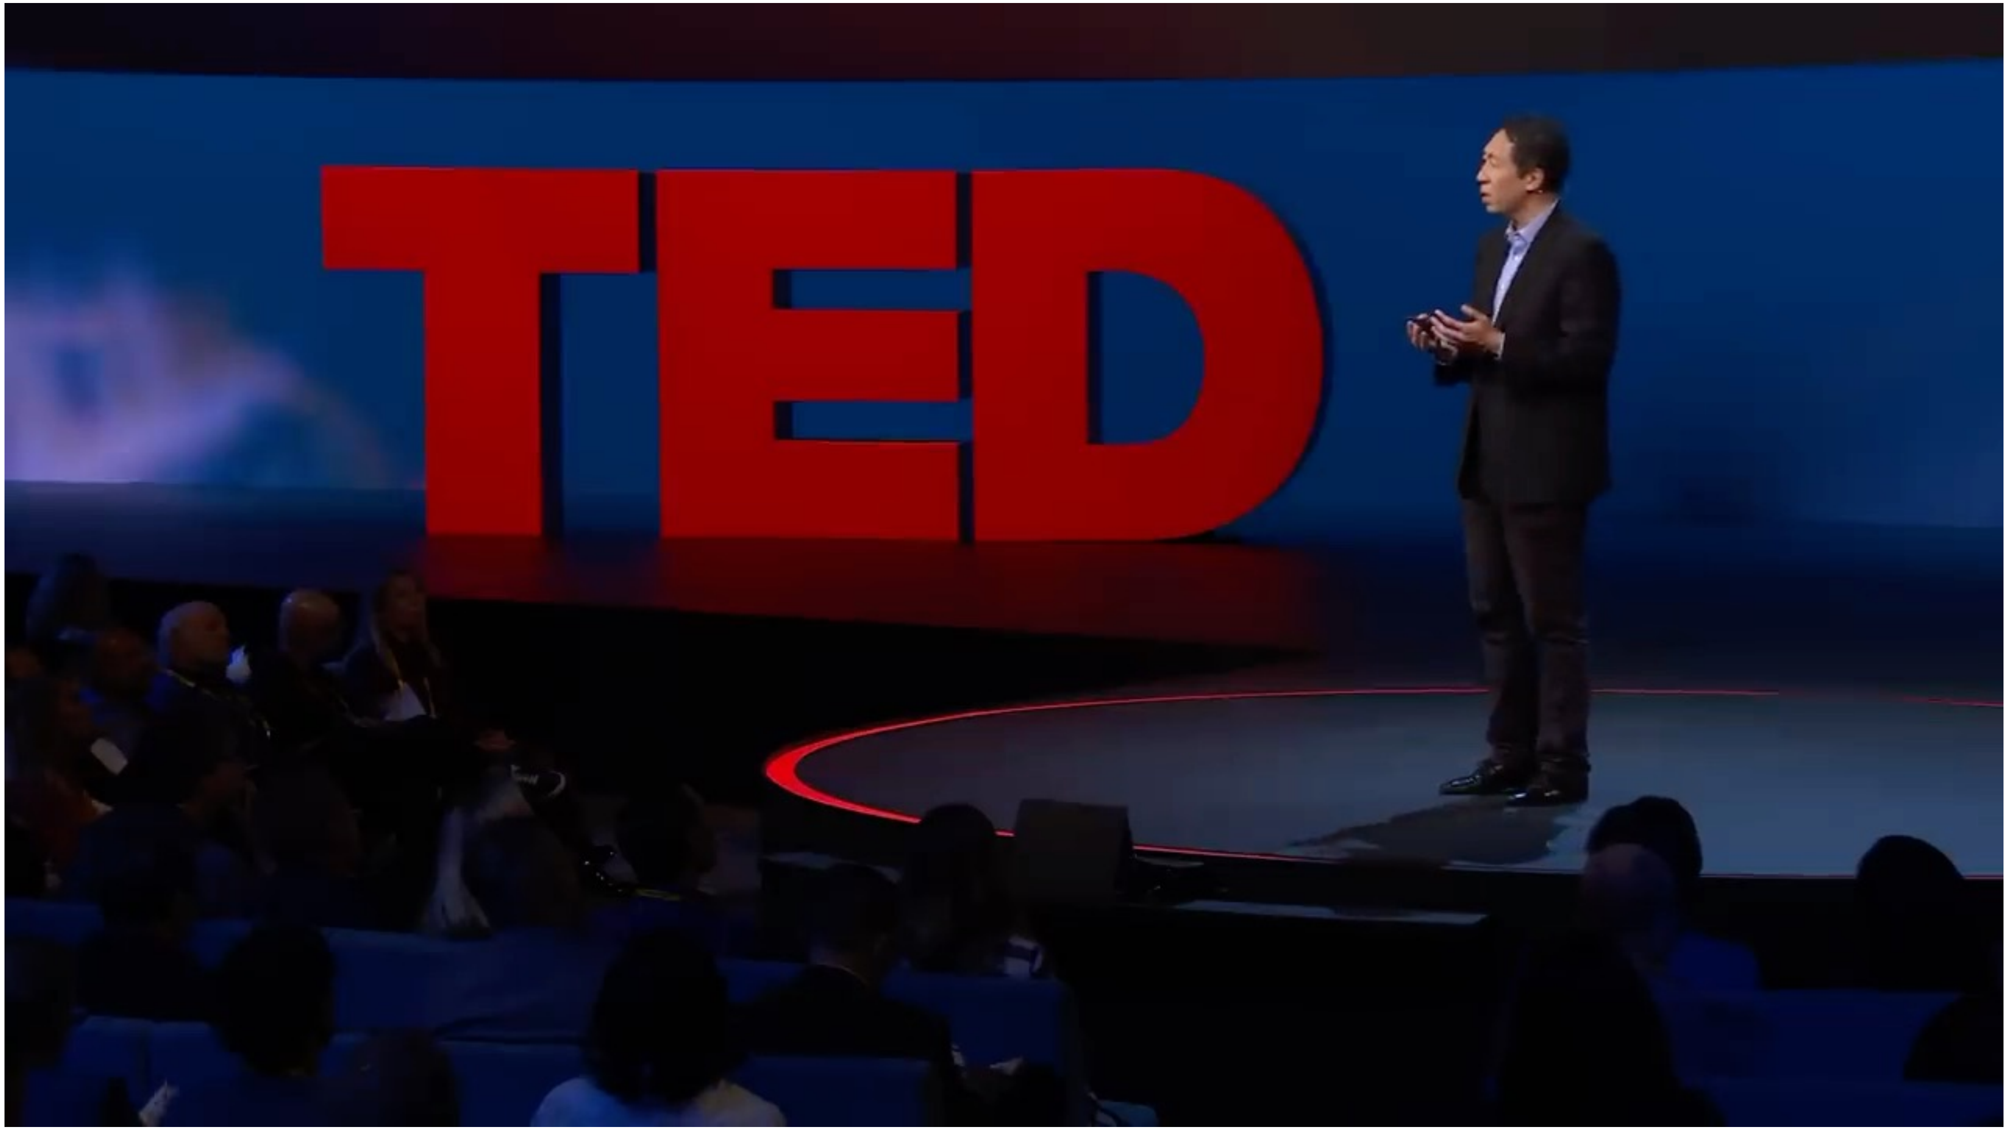

How will AI change the world

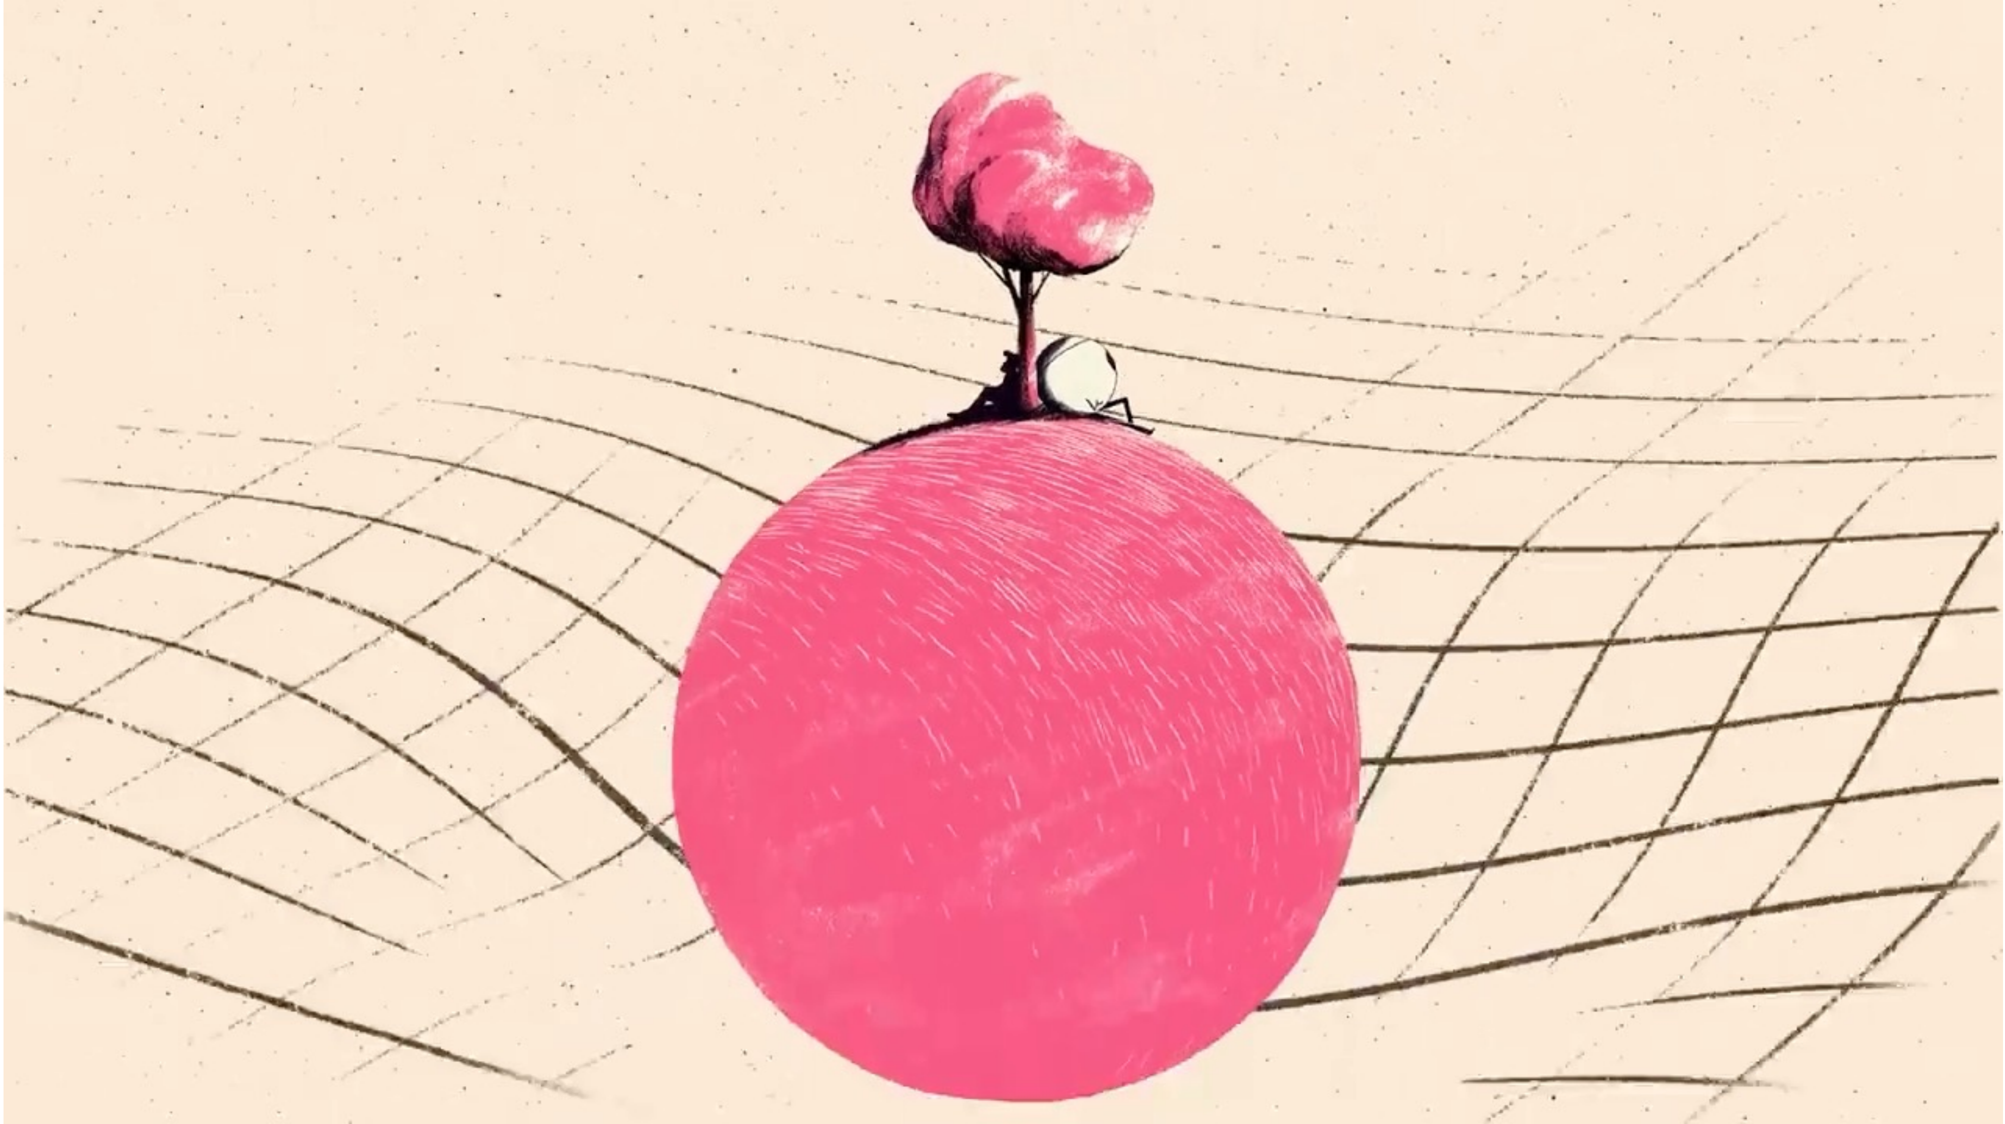

Who are the stakeholders?

AI Research Trends

Lifecycle and Key Dimensions of an AI System

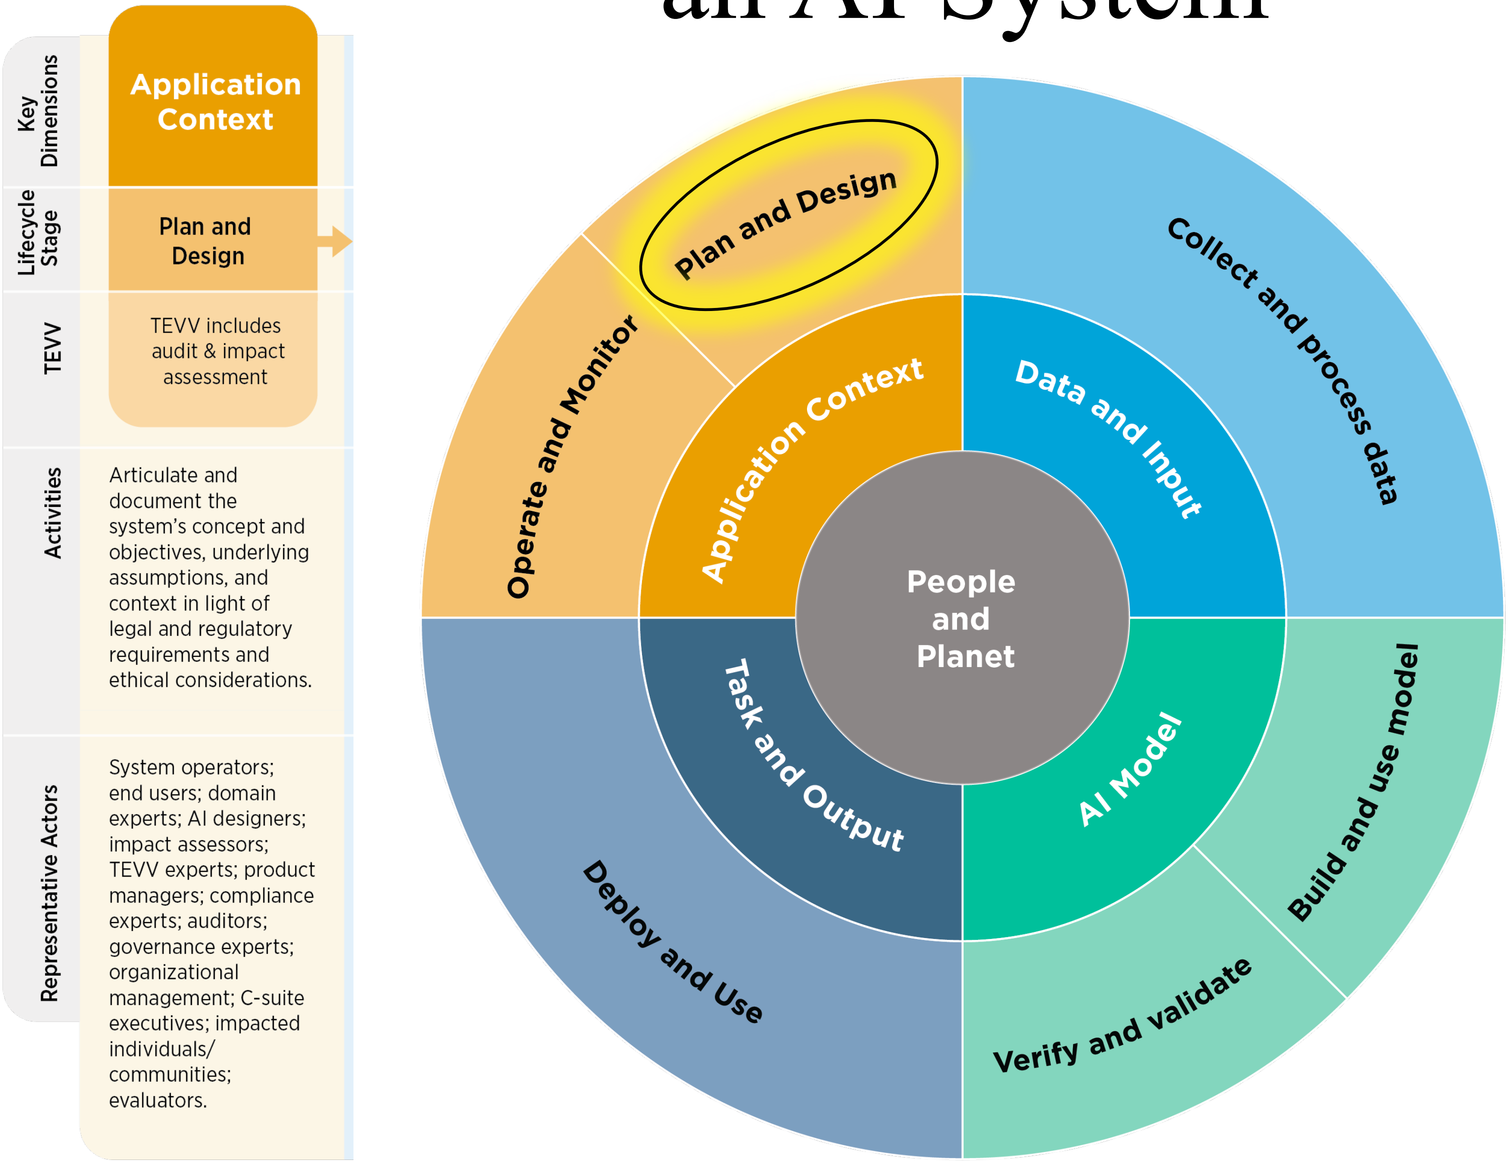

Artificial Intelligence Risk Management Framework by National Institute of Standards and Technology, USA

AI Design Manager in Microsoft

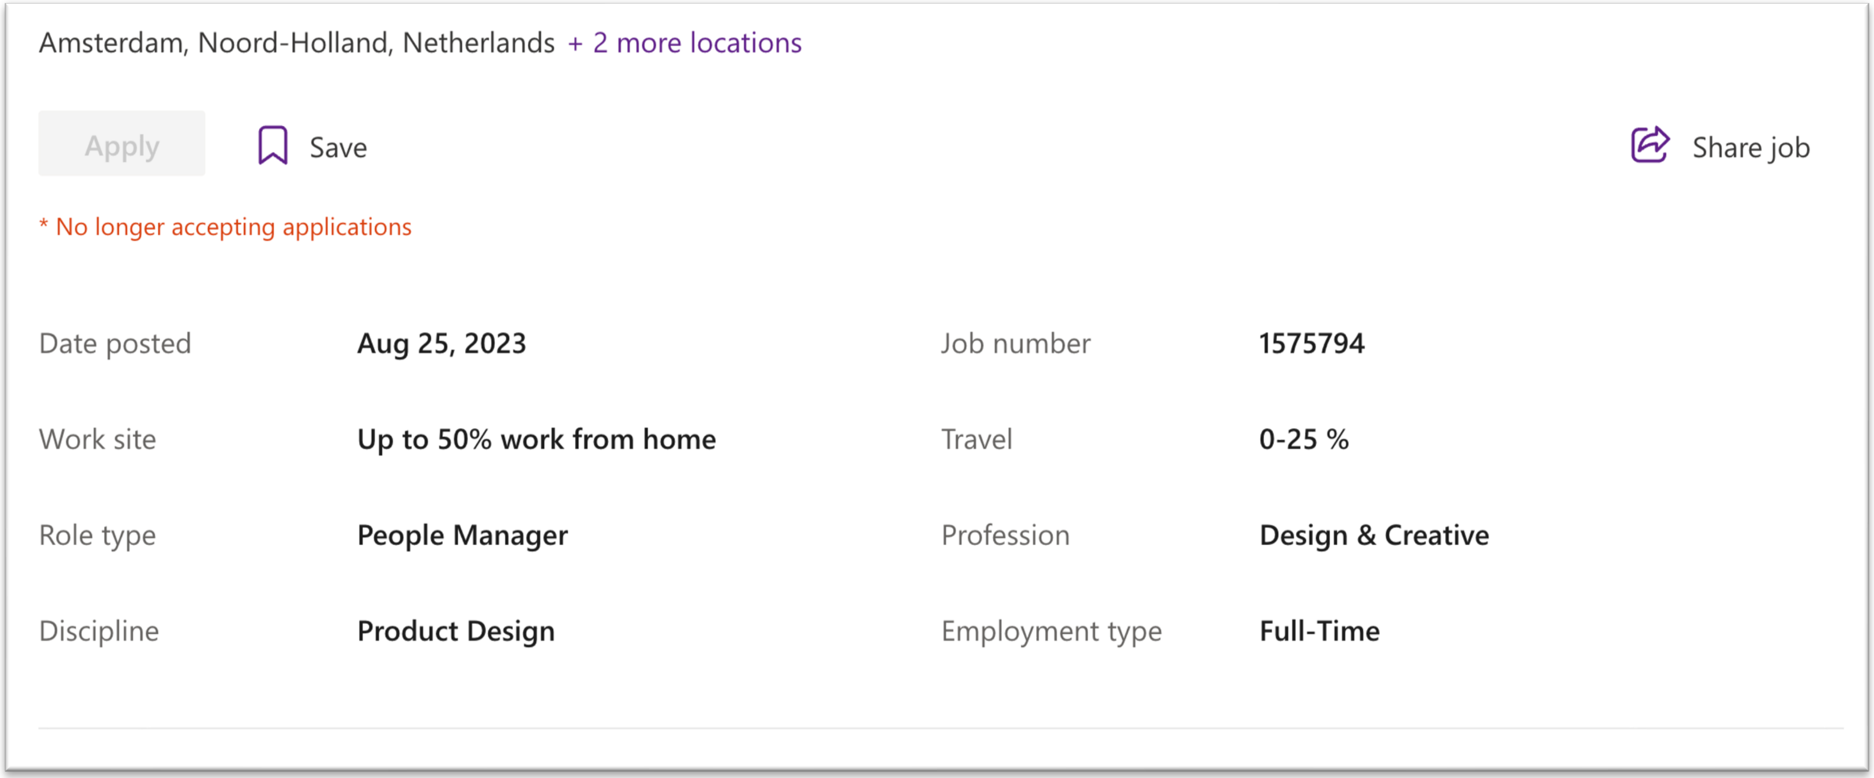
Amsterdam, Noord-Holland, Netherlands + 2 more locations

Apply Save Share job

* No longer accepting applications

Date posted	Aug 25, 2023	Job number	1575794
Work site	Up to 50% work from home	Travel	0-25 %
Role type	People Manager	Profession	Design & Creative
Discipline	Product Design	Employment type	Full-Time

Related Readings:

- https://jobs.careers.microsoft.com/global/en/job/1575794/AI-Design-Manager
- What Is the Role of an AI Designer? By Director of Product Design at Facebook
- How Is AI-Centered Product Design Different?

Overview

Do you enjoy designing innovative bleeding edge solutions and love to learn and tackle new challenges including building and managing the right team, creating goals for them to execute against, and delivering a high quality product? Do you like to direct the design of holistic experiences from storyboards, wireframes, motion studies to visual comps and lead the design and development of interactive prototypes?

As an **AI Designer Manager** for our AI Acceleration Studio, this is your opportunity to use your experience and expertise to design and develop concepts that optimise the emotional and functional experience of a product, software, device, or service. Your ability to bring the latest AI technologies to life will allow you to partner with some of our most strategic partners as you get to solve exciting business problems, produce compelling product stories, and collaborate with Microsoft product teams.

With your team, you will be working side-by-side with customers and their engineering teams to build visual and interactive designs that leverage generative AI and other leading edge AI capabilities. As part of our Industry Solutions Engineering (ISE) organisation, you will thrive as you work with a variety of customers across different industries and get to broaden your knowledge and skills as you interact with the latest technologies.

ISE is a global engineering organisation that works closely with our customers' engineers to jointly develop code for cloud-based solutions that can accelerate their organization. We work in collaboration with Microsoft product teams, partners, and open-source communities to empower our customers to do more with the cloud. We pride ourselves in making contributions to open source and making our platforms easier to use.

In addition, ISE is part of the Microsoft Industry Solutions organization and is a global organization of over 16,000 strategic sellers, industry experts, elite engineers, and world-class architects, consultants, and delivery experts who work together to bring Microsoft's mission of empowerment – and cutting-edge technology - to life for the world's most influential customers. We are on the front lines of innovation, working side-by-side with customers to drive value across the entirety of their digital transformation journey.

Our team prides itself on embracing a growth mindset, inspiring excellence, and encouraging everyone to share their unique viewpoints and be their authentic selves. Join us and help create life-changing innovations that impact billions around the world!

Watch this video to learn more about who we are and what we do: https://aka.ms/csevideo.

Qualifications

- Experience working in product or service design with an ability to use design thinking to solve problems
- Experience of managing a team of designers
- Experience in delivering products or solutions to customers following an end-to-end design process
- Familiarity with AI technologies
- Knowledge of responsible AI principles and artificial intelligence technologies
- Proficiency with prototyping technologies, and adept at iterative design processes
- Demonstrated ability of working with internal and external stakeholders at all levels in an enterprise organization
- Bachelor's or Master's Degree in Industrial Design, Product Design, Human Computer Interaction, or a related field OR equivalent experience

There is a lot of customer interaction in this role and travel can be up to 50% of the time.

#ISEngineering
#INDSOL

Emerging Tech
x
Design

If we want technology to represent all of us, it needs to be created by all of us.

https://dschool.stanford.edu/emerging-tech

DS323: AI in Design – Similar courses in China and abroad

SOUTHERN UNIVERSITY OF SCIENCE AND TECHNOLOGY 南方科技大学

SCHOOL OF DESIGN
创新创意设计学院

HNU DESIGN 湖南大学 设计艺术学院

DT03021 智能设计方法

MIT xPRO

DESIGNING AND BUILDING AI PRODUCTS AND SERVICES

d.

HASSO PLATTNER
Institute of Design at Stanford

DESIGNING MACHINE LEARNING

Dive into this popular class which takes a multidisciplinary, human-centered approach to designing systems of machine learning and AI.

Carnegie Mellon University

Human-Computer Interaction Institute

About | Academics | Careers | Giving | News & Events | People | Research

Home > Academics > Courses > Design of AI Products and Services

Design of AI Products and Services

Course Information
Course Number

HCI Undergraduate: 05-317
HCI Graduate: 05-617

Current status of representation

- Most of the AI systems are made by PhD in machine learning, who often lacks expertise in giving form to new products or services within human contexts and environments.

Current status of representation

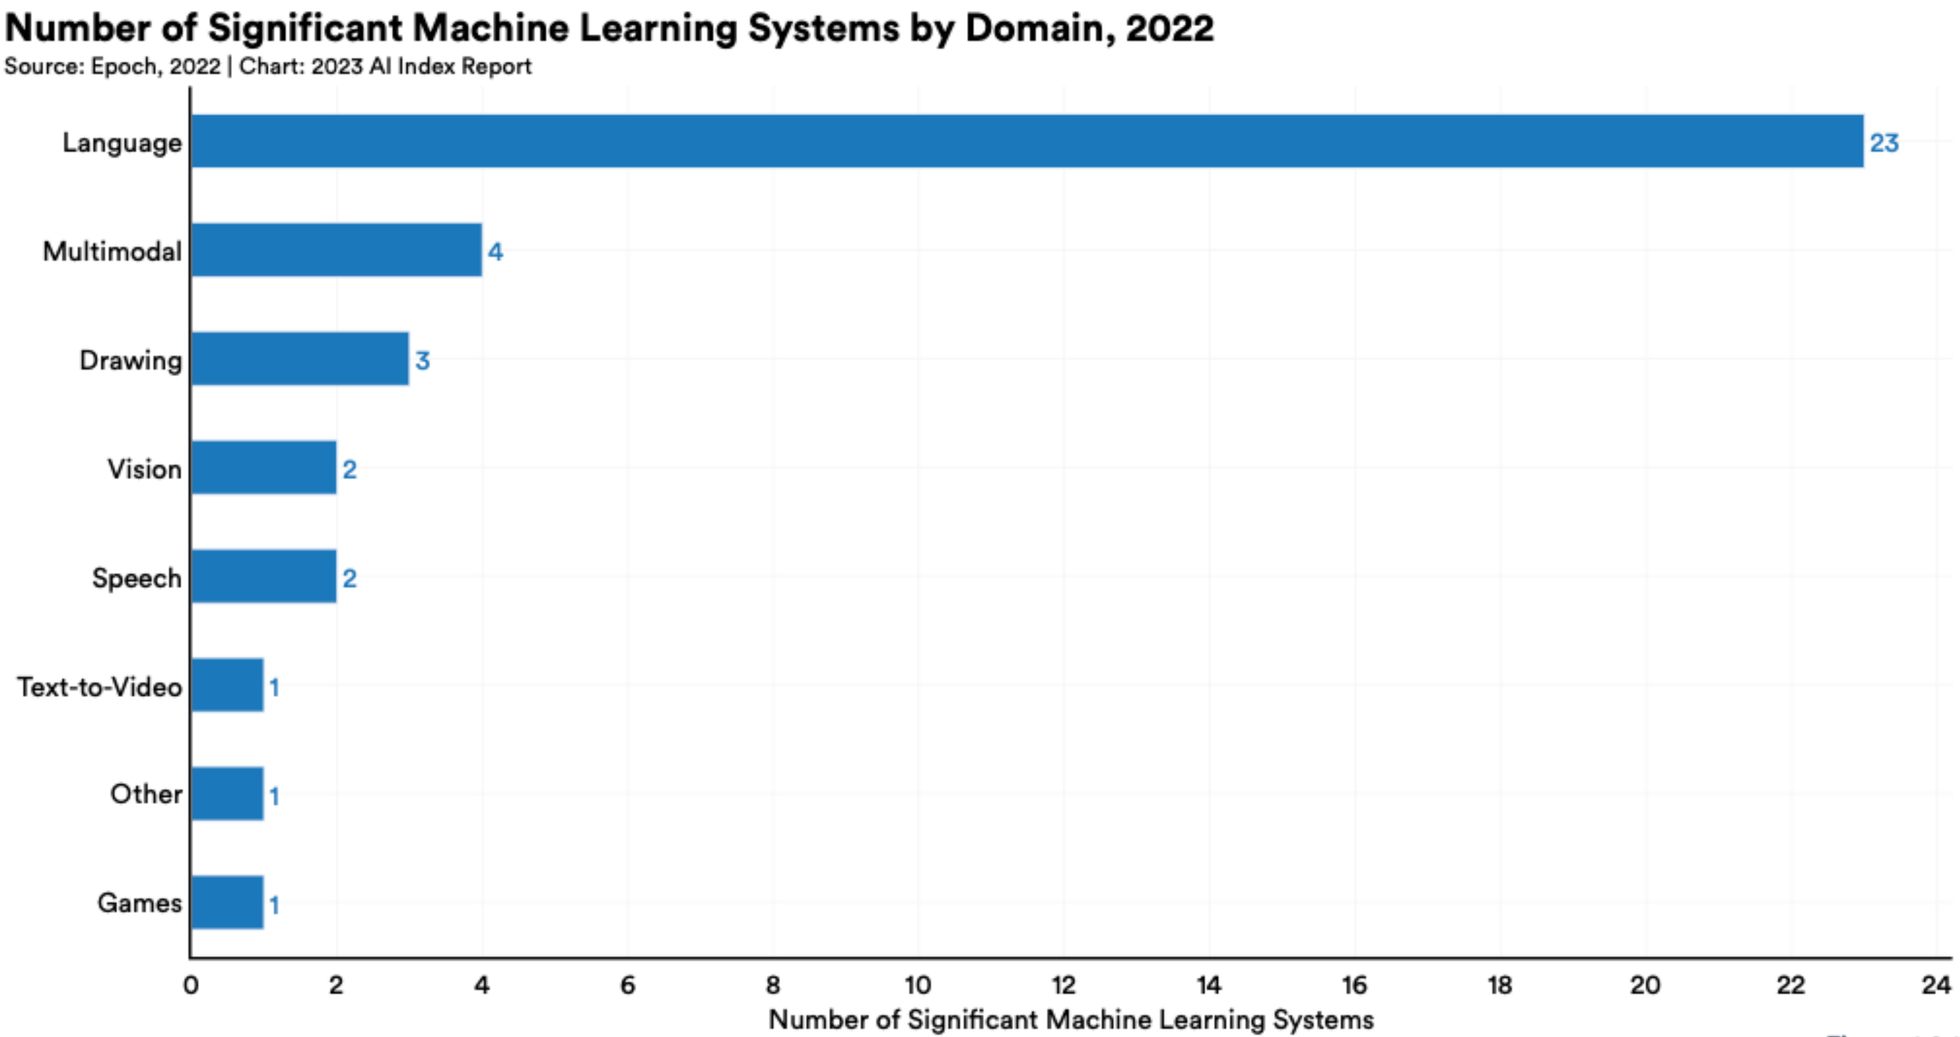

Figure 1.2.1[6]

6 There were 38 total significant AI machine learning systems released in 2022, according to Epoch; however, one of the systems, BaGuaLu, did not have a domain classification and is therefore omitted from Figure 1.2.1.

Current status of representation

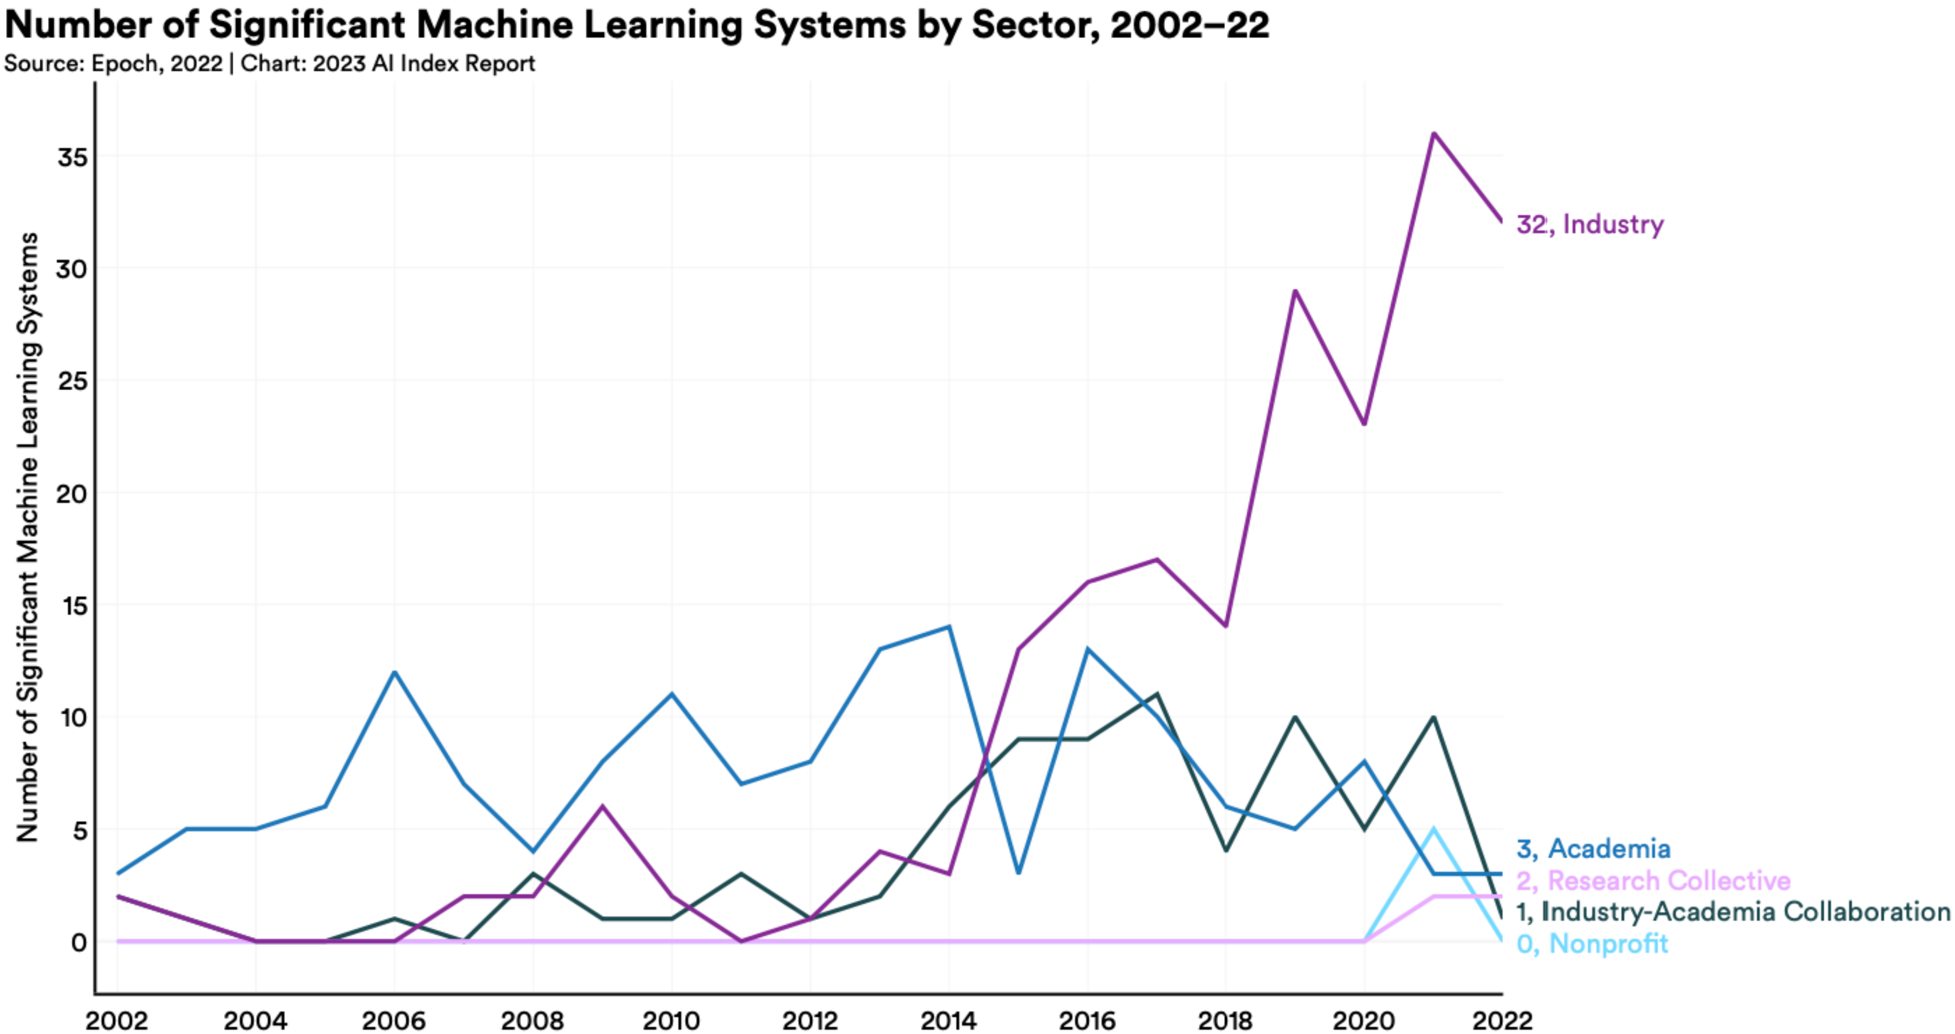

Figure 1.2.2

Current status of representation

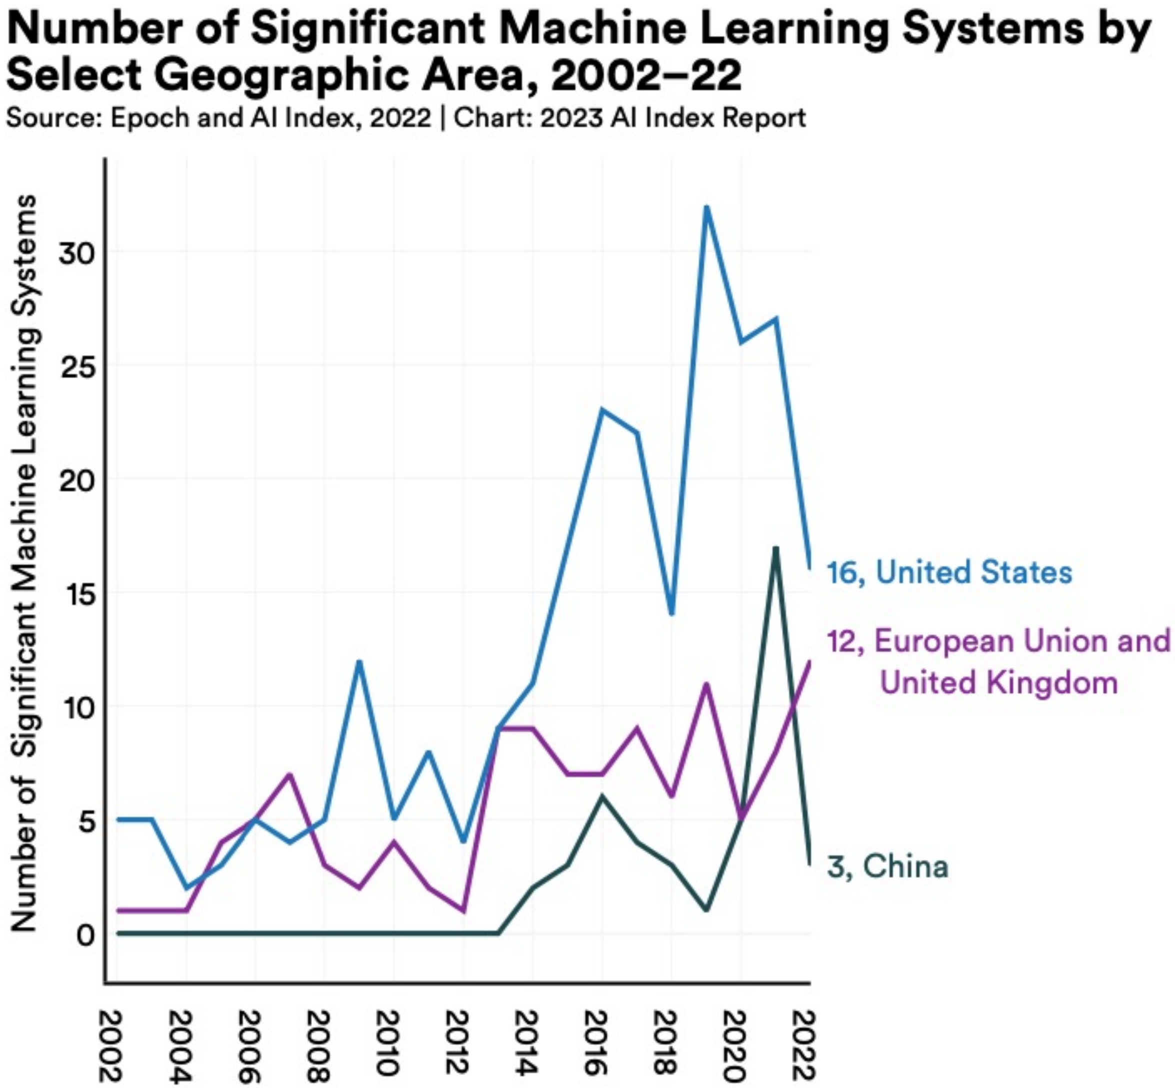

Figure 1.2.4

Current status of representation

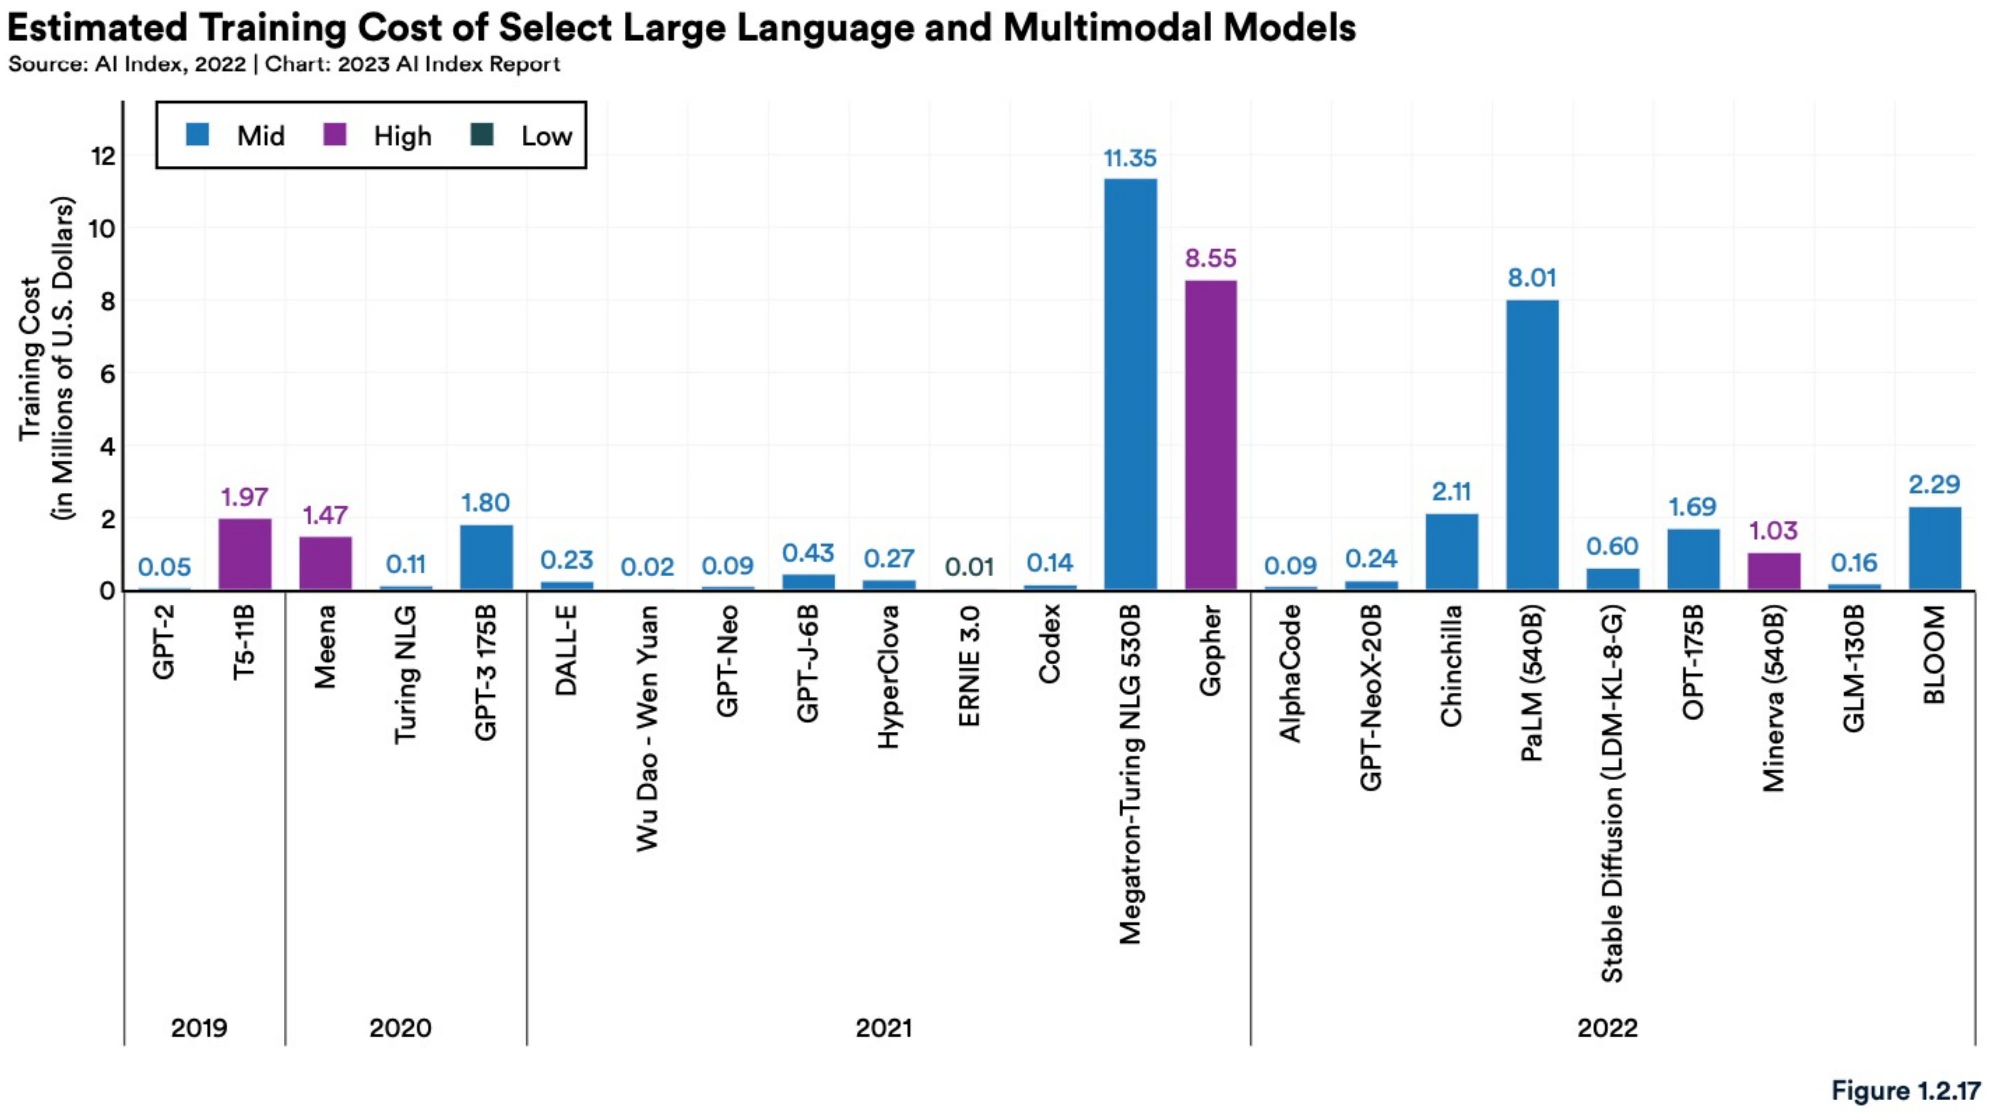

Current status of representation

- Frontier AI technology such as generative AI systems are introduced to the general public without ***enough*** risk control: ***Lack of trustworthiness in AI***.

What is your experience with generative AI?

Generative AI–related risks that organizations consider relevant and are working to mitigate,
% of respondents[1]

[1]Asked only of respondents whose organizations have adopted AI in at least 1 function. For both risks considered relevant and risks mitigated, n = 913.
Source: McKinsey Global Survey on AI, 1,684 participants at all levels of the organization, April 11–21, 2023

What is Artificial Intelligence?

AI in Classroom

Think about an AI product or service

- How is AI used in the product or service (what does it actually do)?
- What does the AI take as input and provide as output?
- How does the use of AI produce value for the user?
- How does the use of AI produce value for the product maker or service company?
- What kind of AI is used … what machine learning method, etc.? It's OK if you do not know. Many products and services do not make this clear.
- What kind of errors does the system make?
- How well does it need to work to be valuable?

What do we mean by Artificial Intelligence?

AI is a machine's ability to perform the cognitive functions
we associate with human minds,
such as perceiving, reasoning, learning, interacting with an environment, problem solving,
and even exercising creativity.

The Natural vs. The Artificials

The Natural

Let's begin by asking a simple question:

what is the natural, and what about the artificial?

The Artificial

The Natural vs. The Artificials

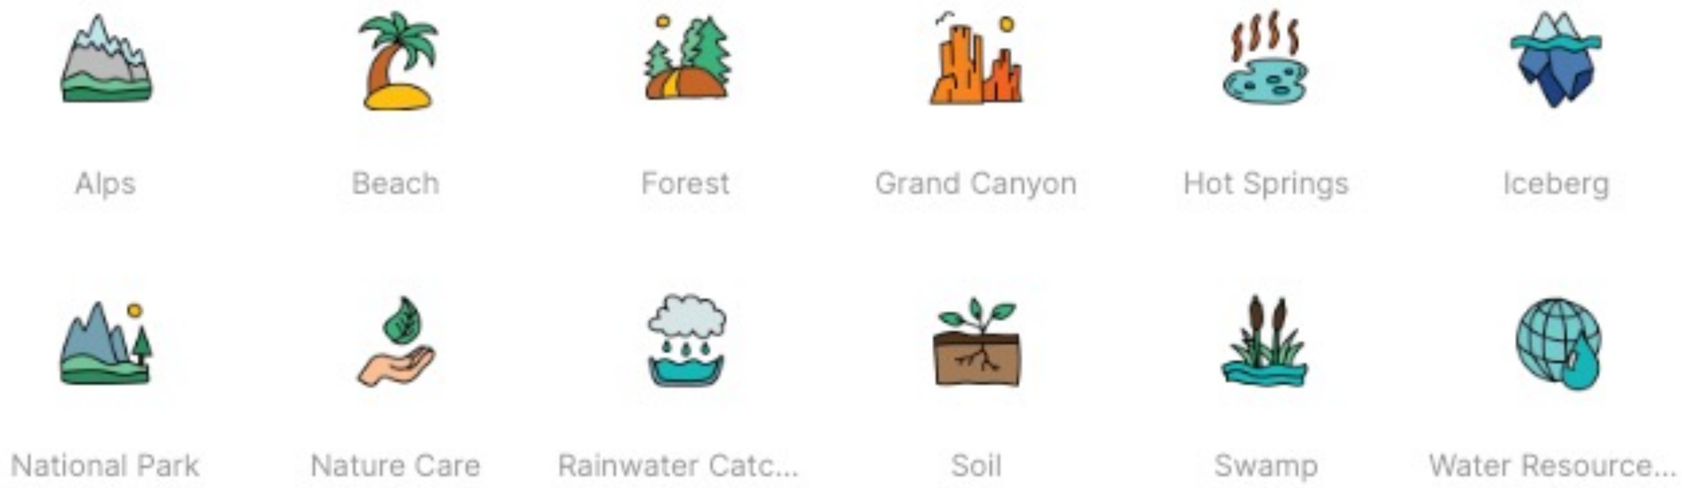

The Natural

We can generally describe ***the natural*** as anything that already existed on earth.

Or anything that is not made by the human, including the human

The Natural vs. The Artificials

We can generally describe ***the artificial*** as everything other than the natural.

Or anything that is made by the human

The Artificial

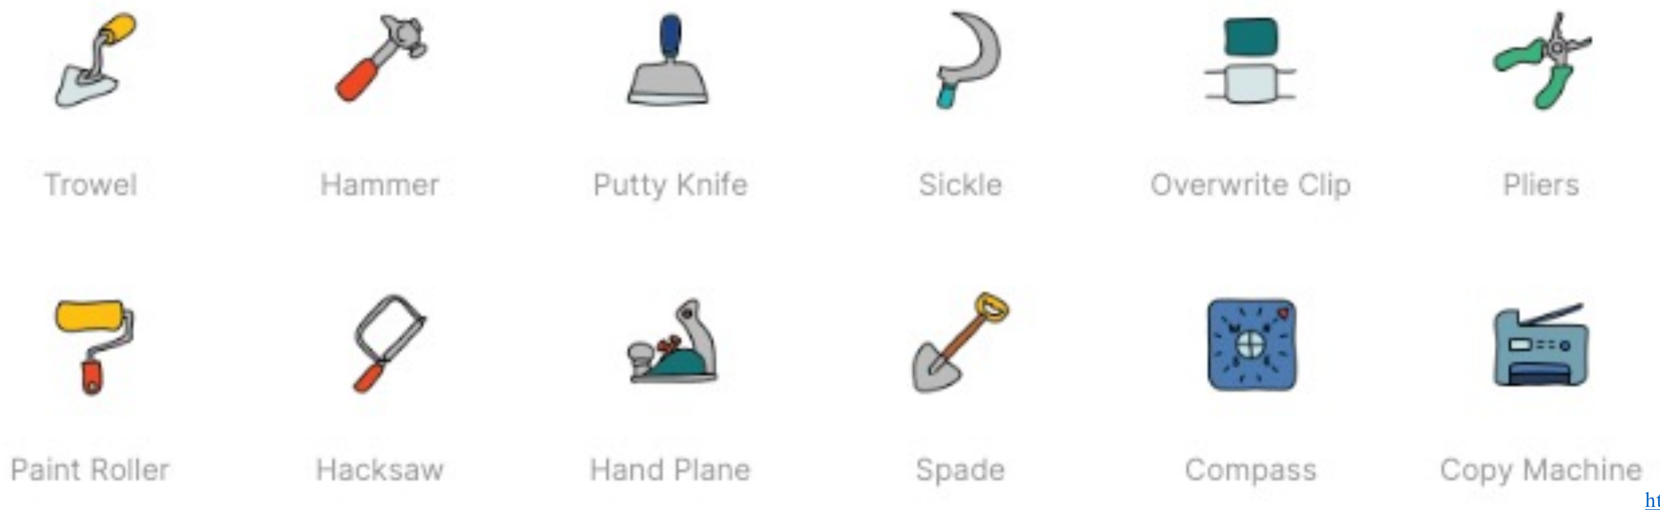

https://icons8.com/

The Natural vs. The Artificials

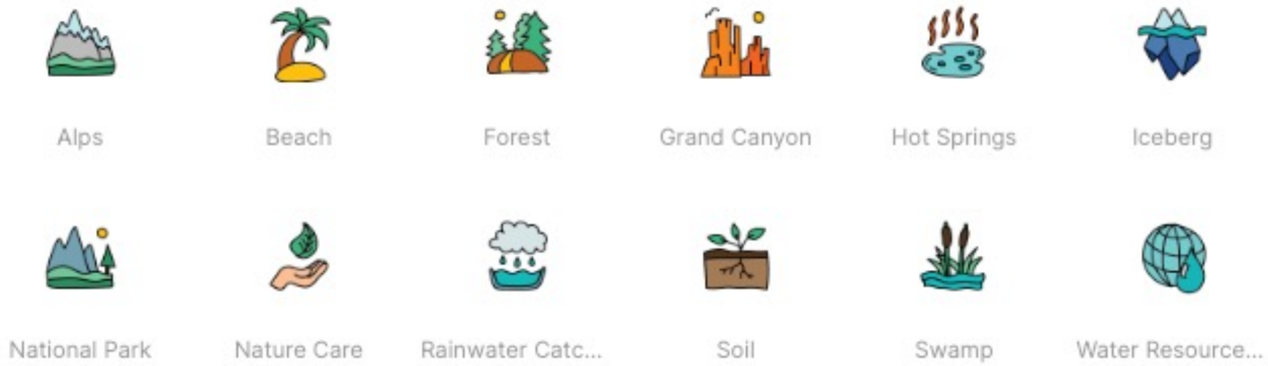

Natural

Human

Artificial

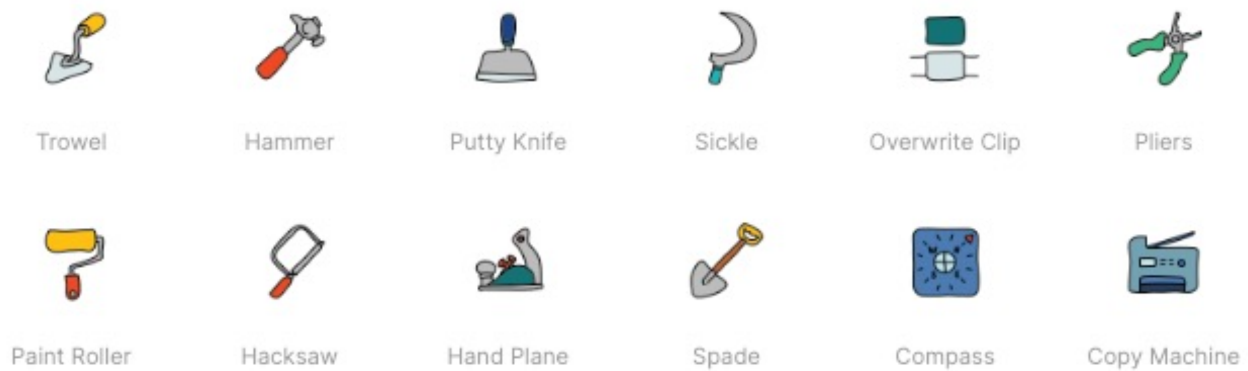

https://icons8.com/

The Natural vs. The Artificials

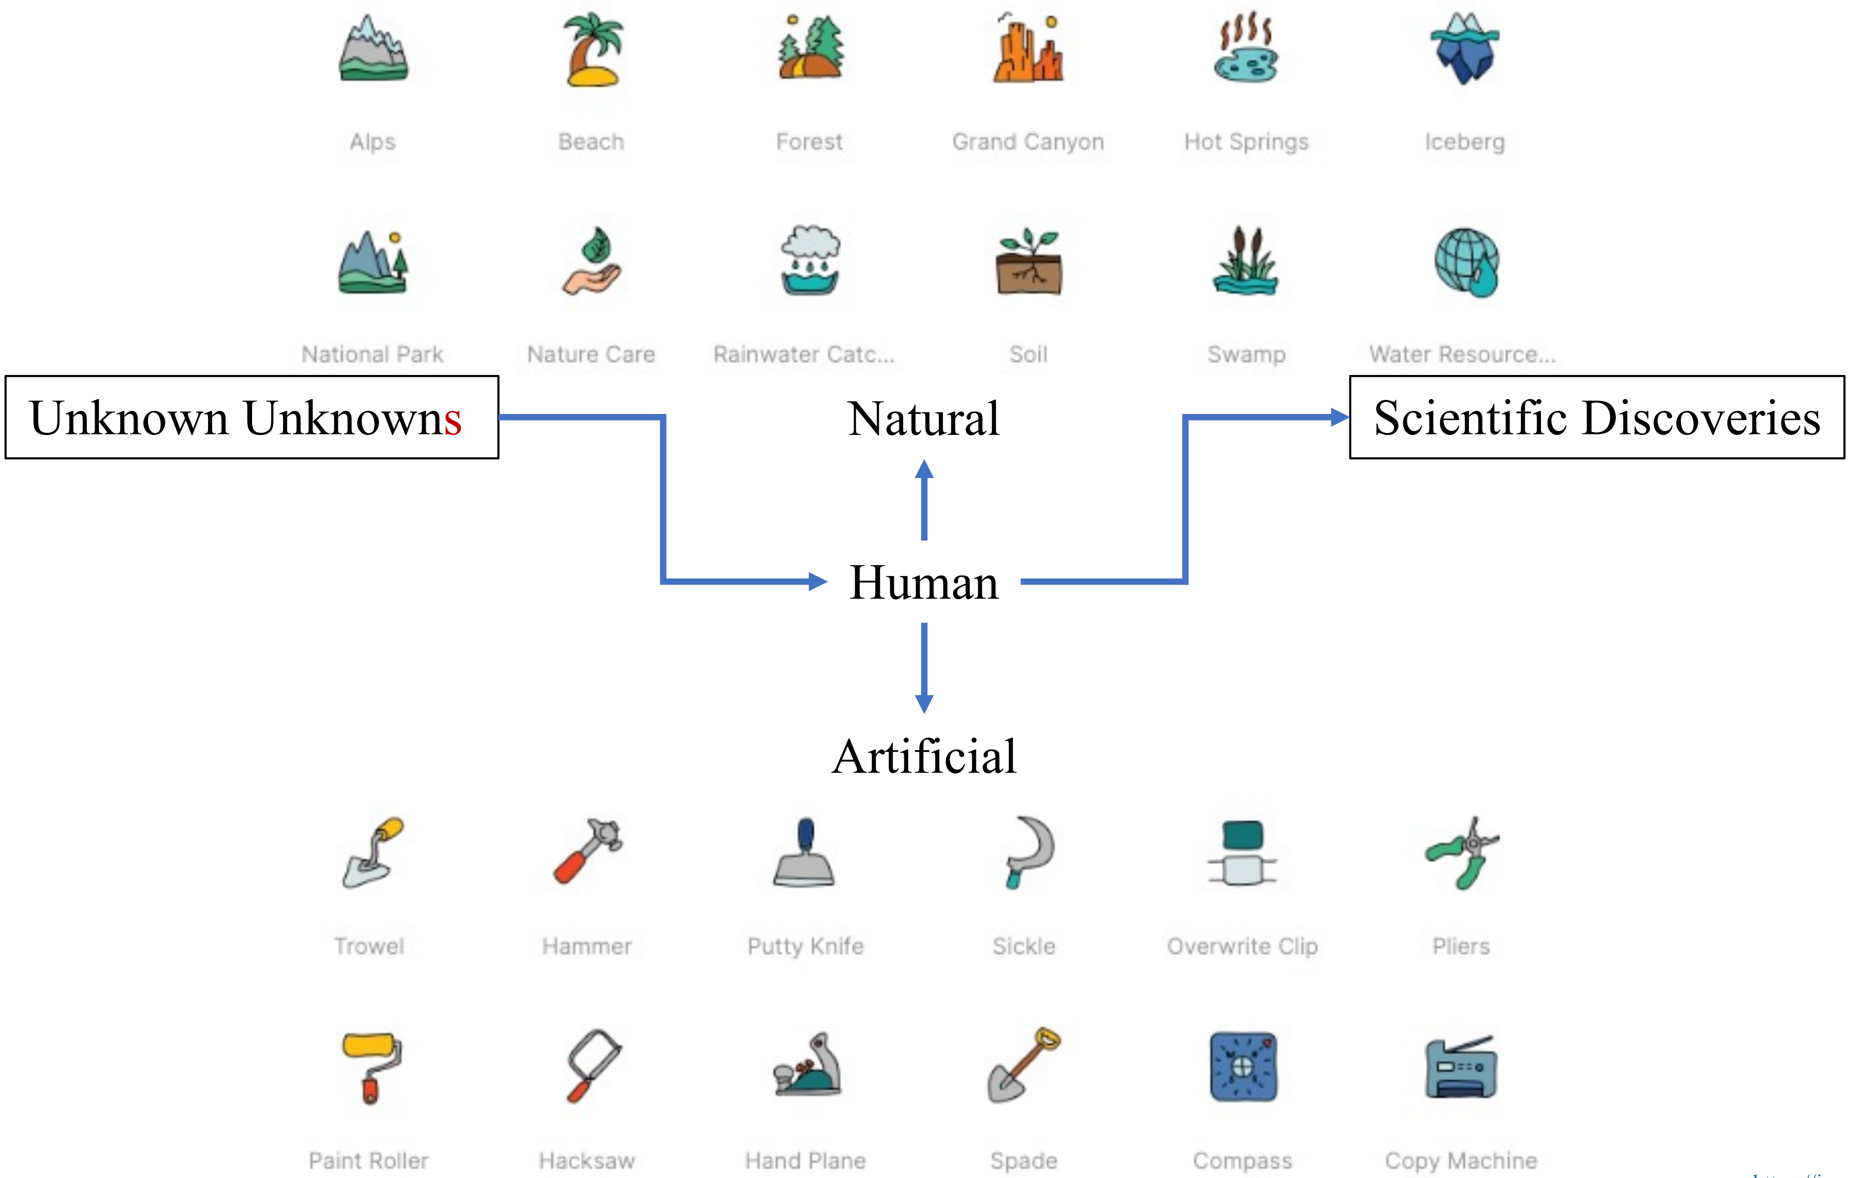

https://icons8.com/

The Natural vs. The Artificials

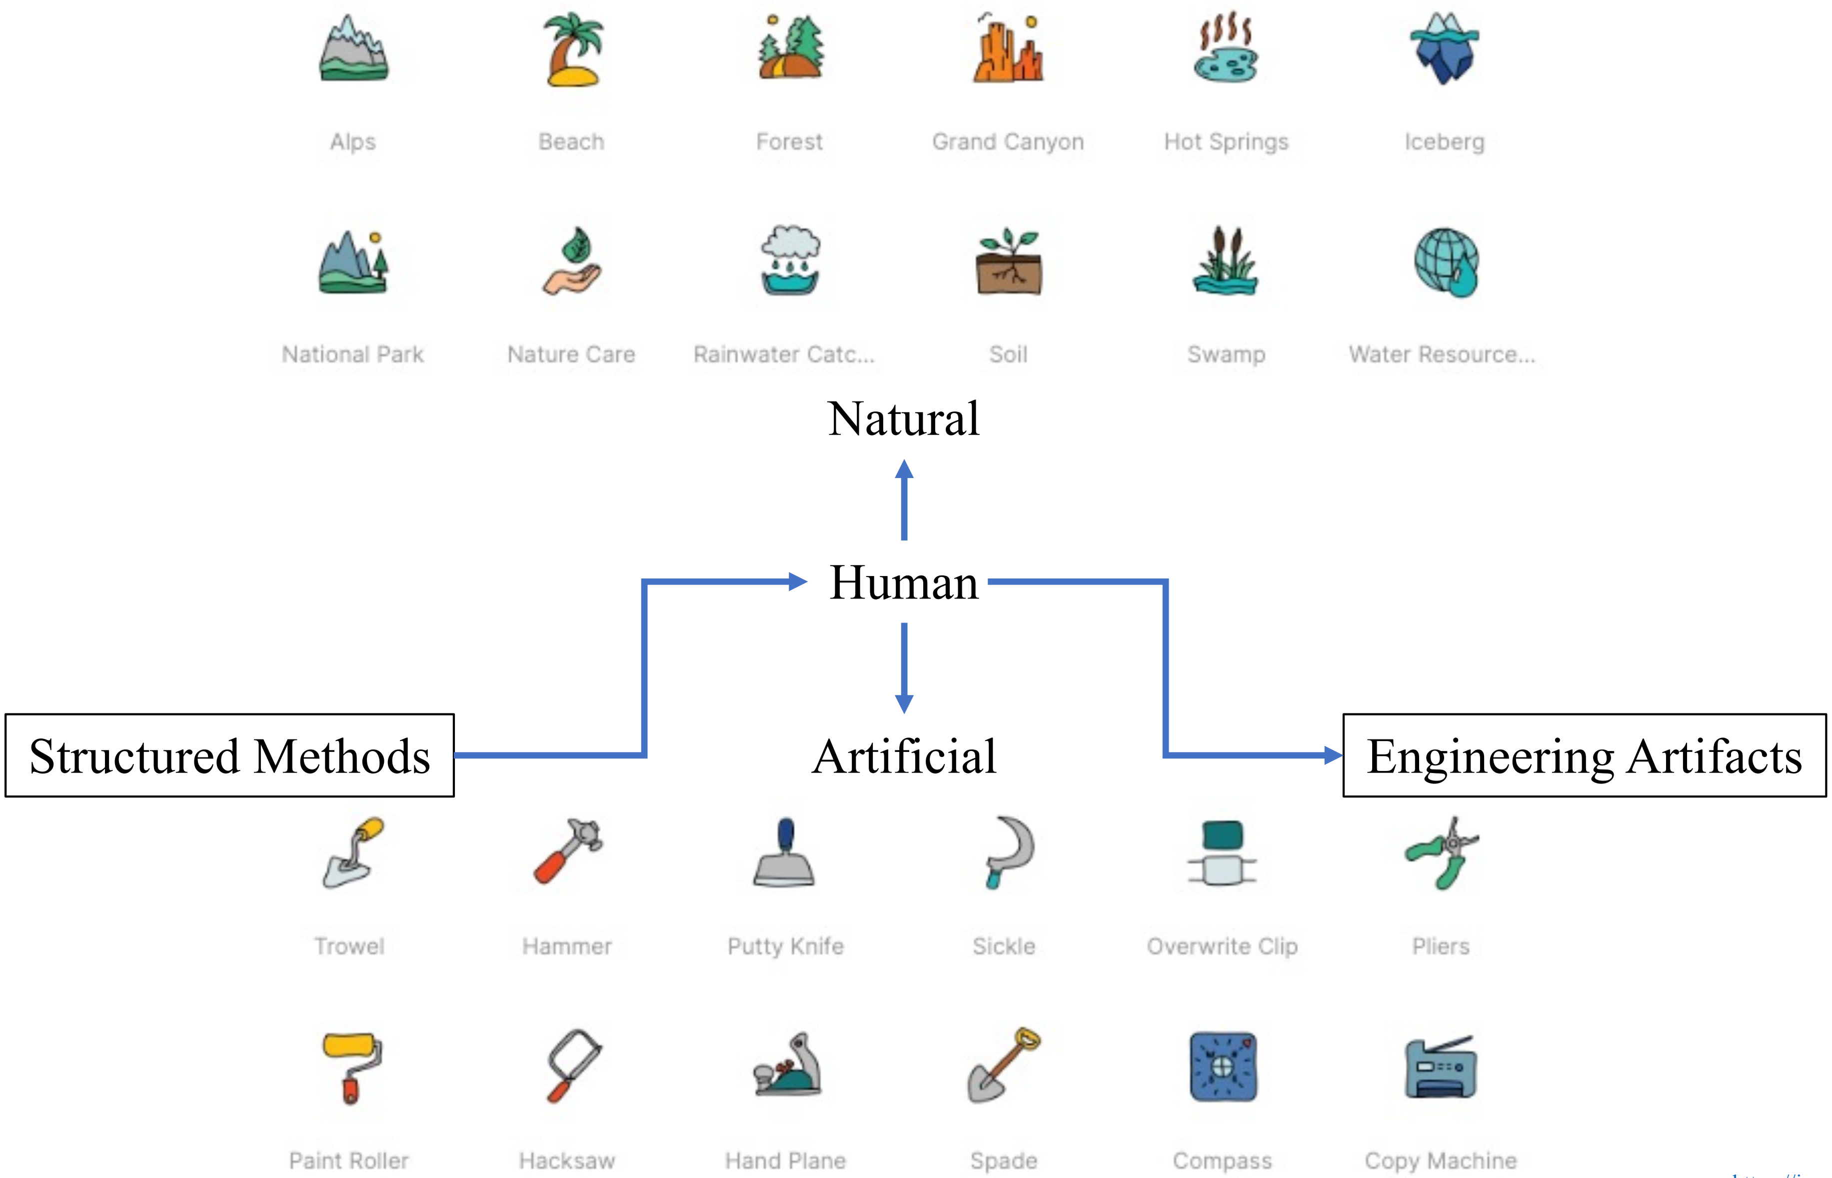

https://icons8.com/

The Natural vs. The Artificials

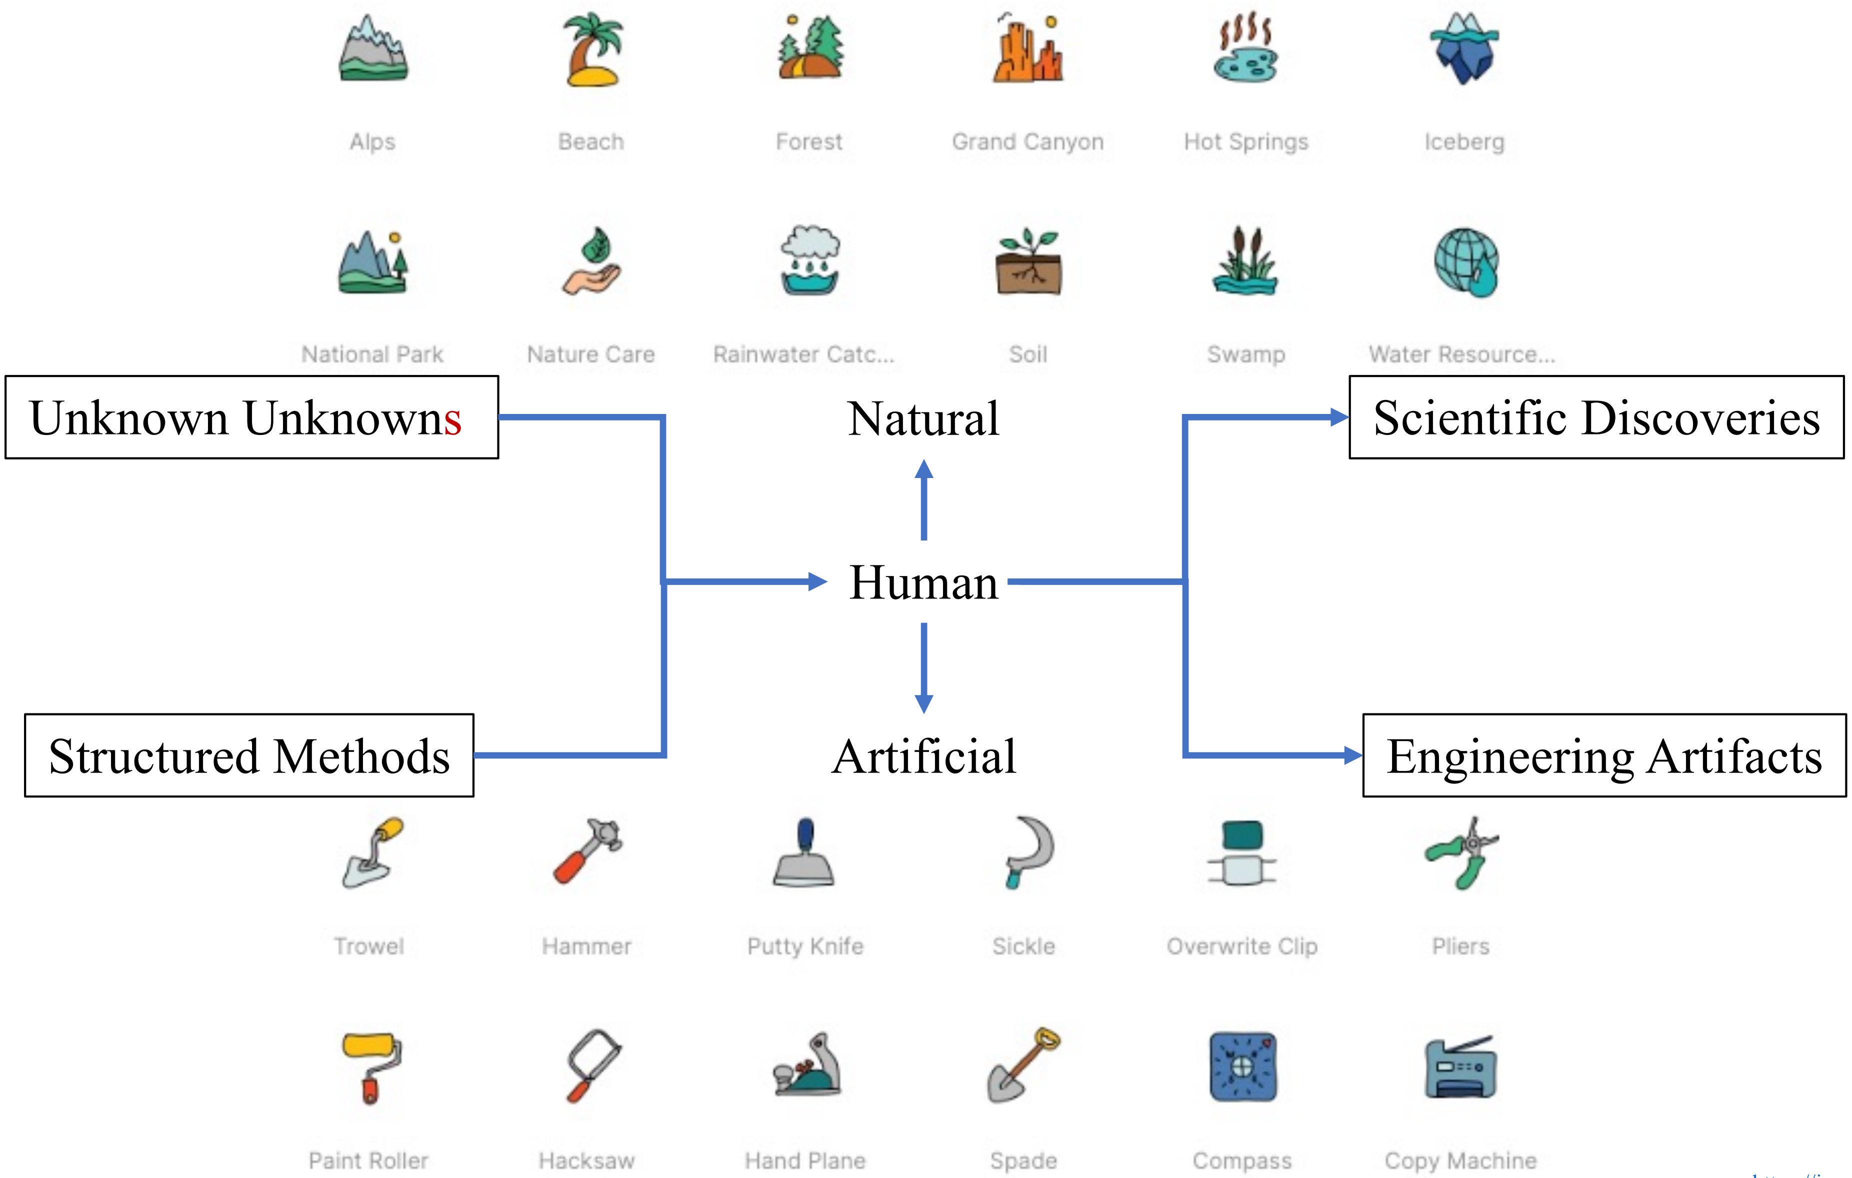

https://icons8.com/

The Natural vs. The Artificials

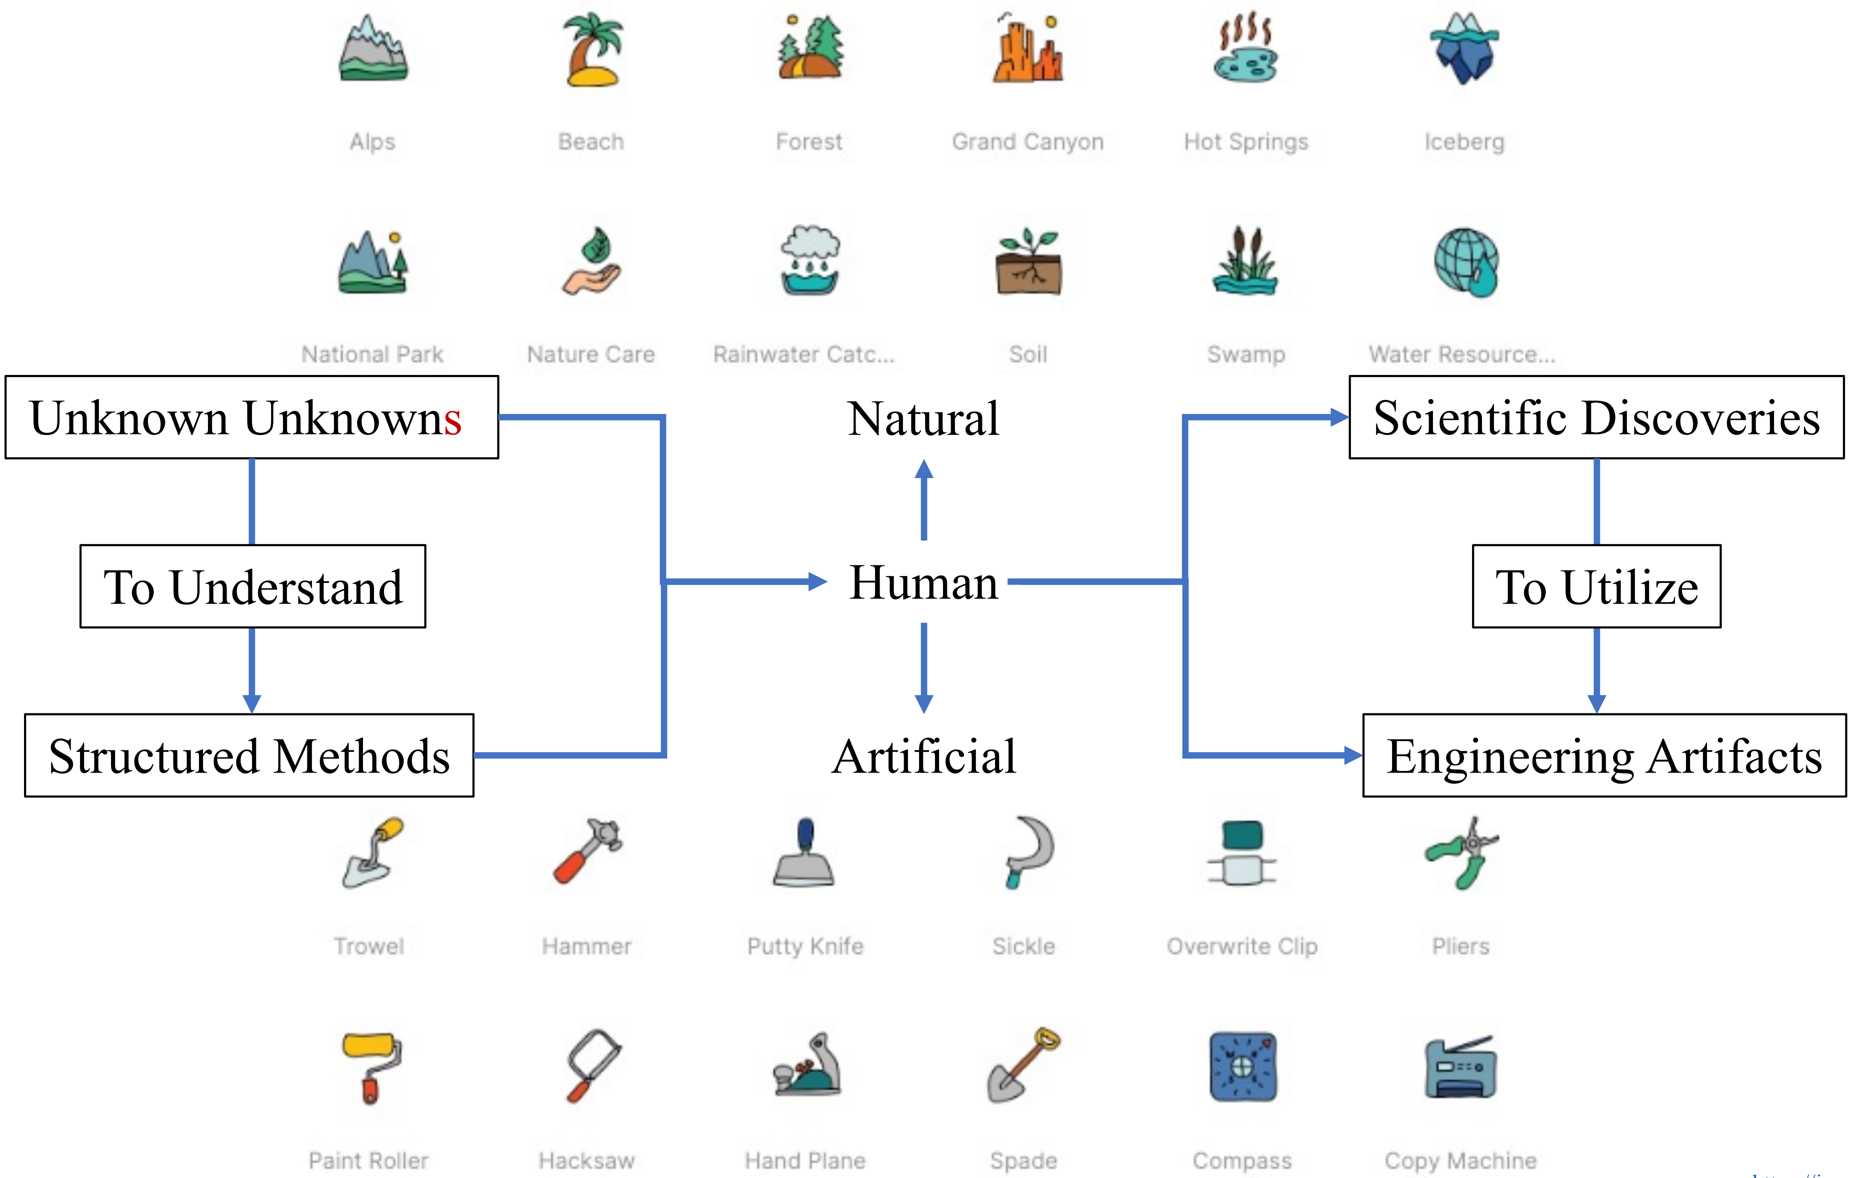

https://icons8.com/

The Natural vs. The Artificials

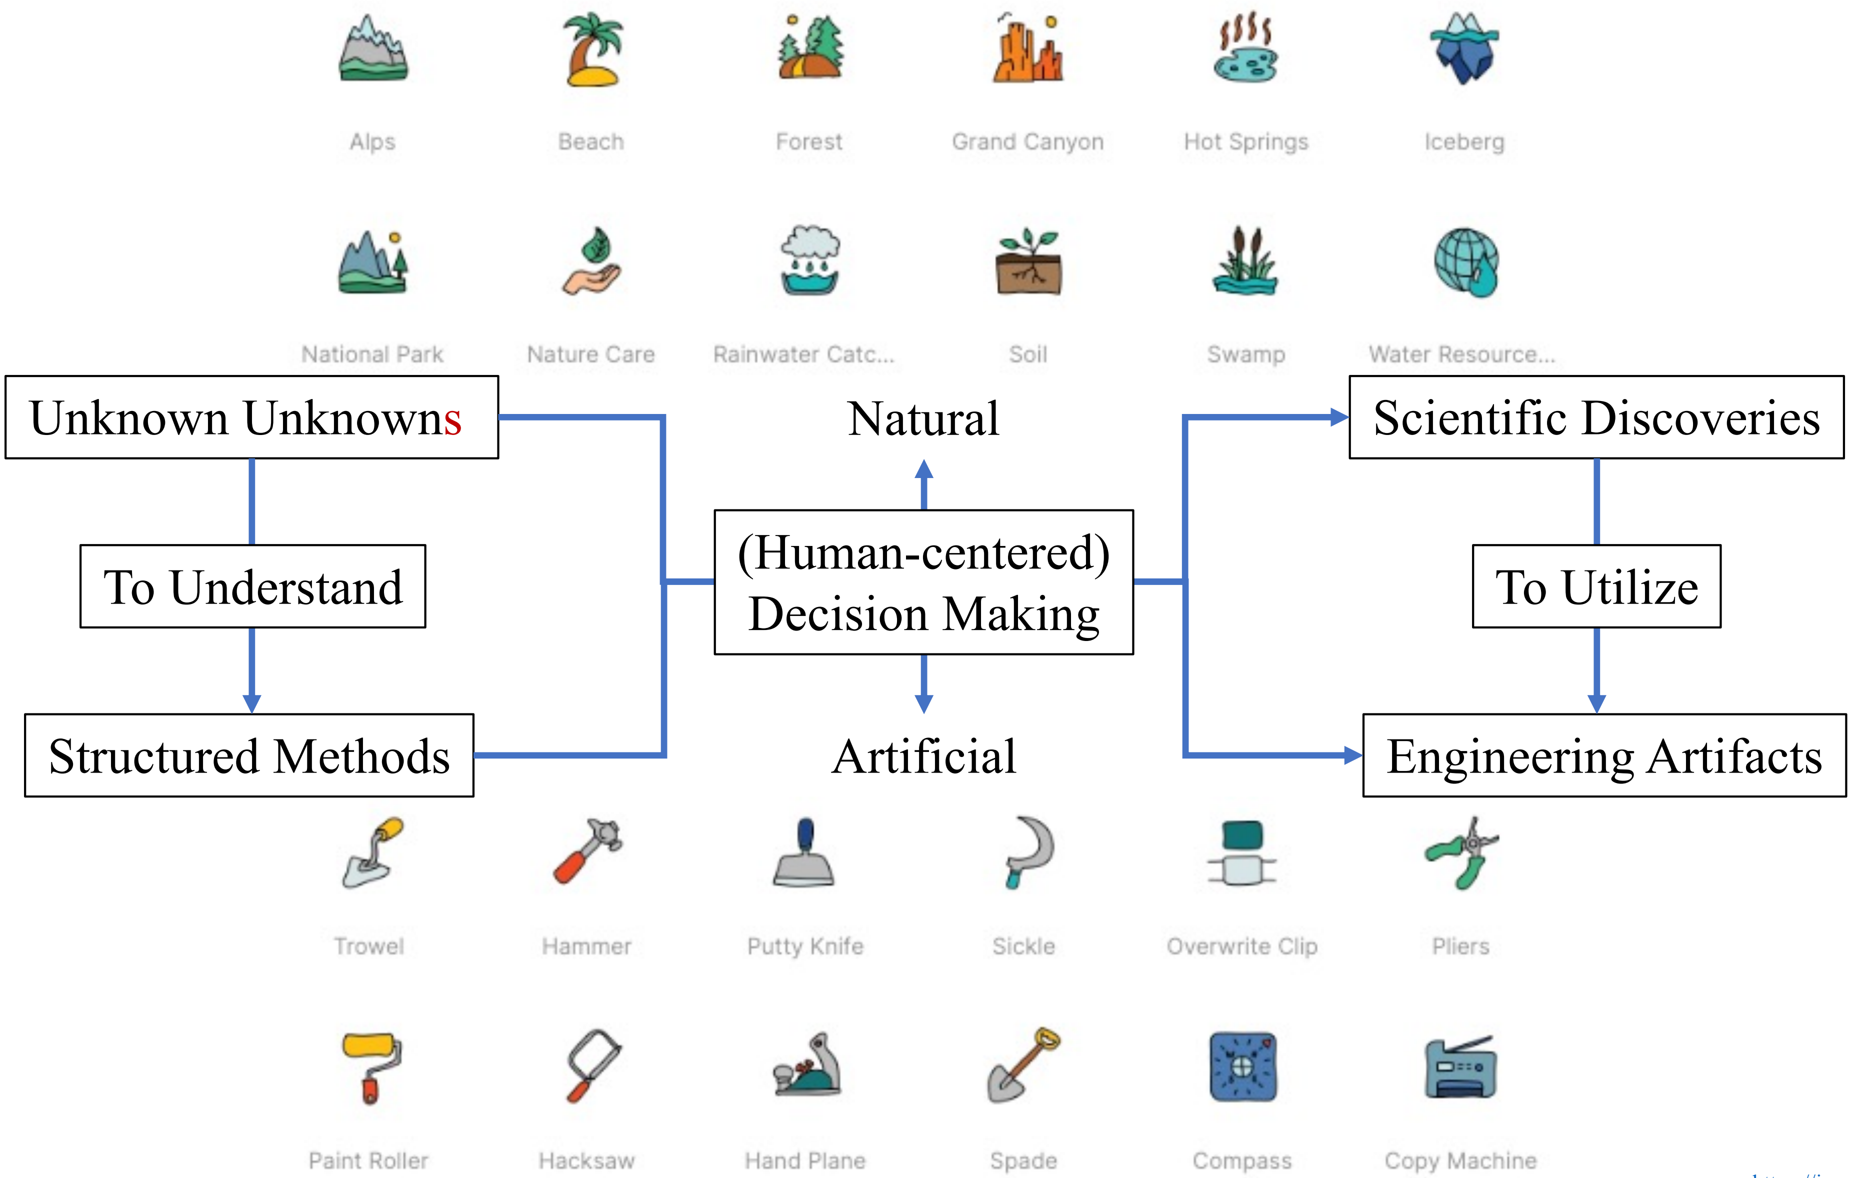

https://icons8.com/

Decision Making as an Optimization Problem

- Optimization is an inherently mathematical subject.
  - It is about maximizing or minimizing some mathematical function to arrive at the best possible solution to a problem, and involves creating design options that are shaped by certain outcomes as they are being created.
- Optimization problems arise in all kinds of fields, from aerospace engineering to architectural design.
  - However, regardless of domain, every optimization problem has three features:
    - An objective function.
    - Constraints.
    - Data.

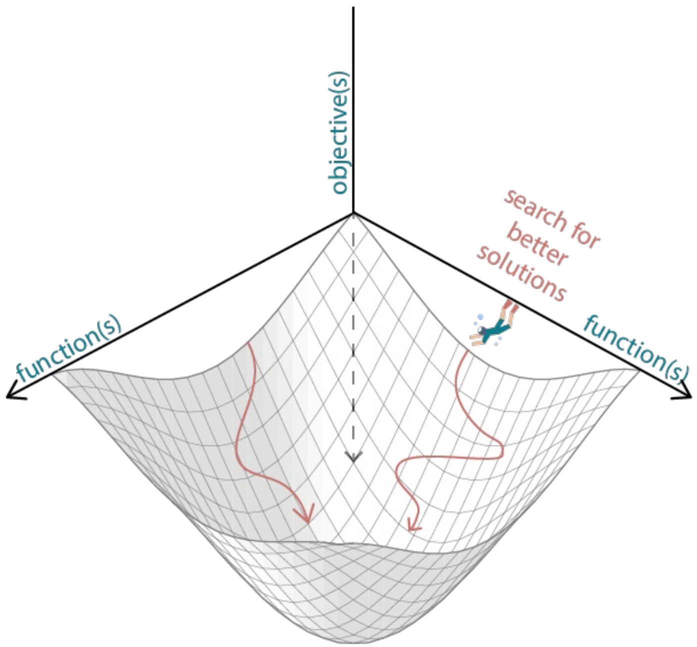

Decision Making as an Optimization Problem

- Without any loss of generality an optimization problem can be defined by:

If to max f_m,
then use $-f_m$ instead.

$$
\begin{aligned}
\min \quad & f_m(x) & x_i^L m &= 1, \dots, M \\
\text{s.t.} \quad & g_j(x) \leq 0 & x_i^L j &= 1, \dots, J \\
& h_k(x) = 0 & x_i^L k &= 1, \dots, K \\
& x_i^L \leq x_i \leq x_i^U & x_i^L i &= 1, \dots, N \\
& x \in \Omega & x_i^L &
\end{aligned}
$$

s.t. means subject to

- where
  - x_i represents the i-th variable to be optimized, x_i^L and x_i^U its lower and upper bound,
    - The variables that defines describes the problems, more if the problem is complex, less if simple
  - f_m the m-th objective function,
    - Your goal, or the decision you want to make
  - g_j the $j-$th inequality constraint and
    - A type of constraints that you need to satisfy within a range
  - h_k the k-th equality constraint.
    - Another type of constraints that you MUST/WANT TO satisfy for sure

Objective Function

- An objective function is the output that you want to maximize or minimize. It is what you will measure designs against to decide which option is best.
  - The objective function can be thought of as the goal of your generative design process.
  - In finance, the objective function is usually to maximize portfolio value; in aerospace engineering the objective is often to minimize weight.

- In generative design workflows, we are not limited to one objective ('single objective optimization') - we can also have multiple objectives or goals that we are trying to optimise our design against ('multi-objective optimization').

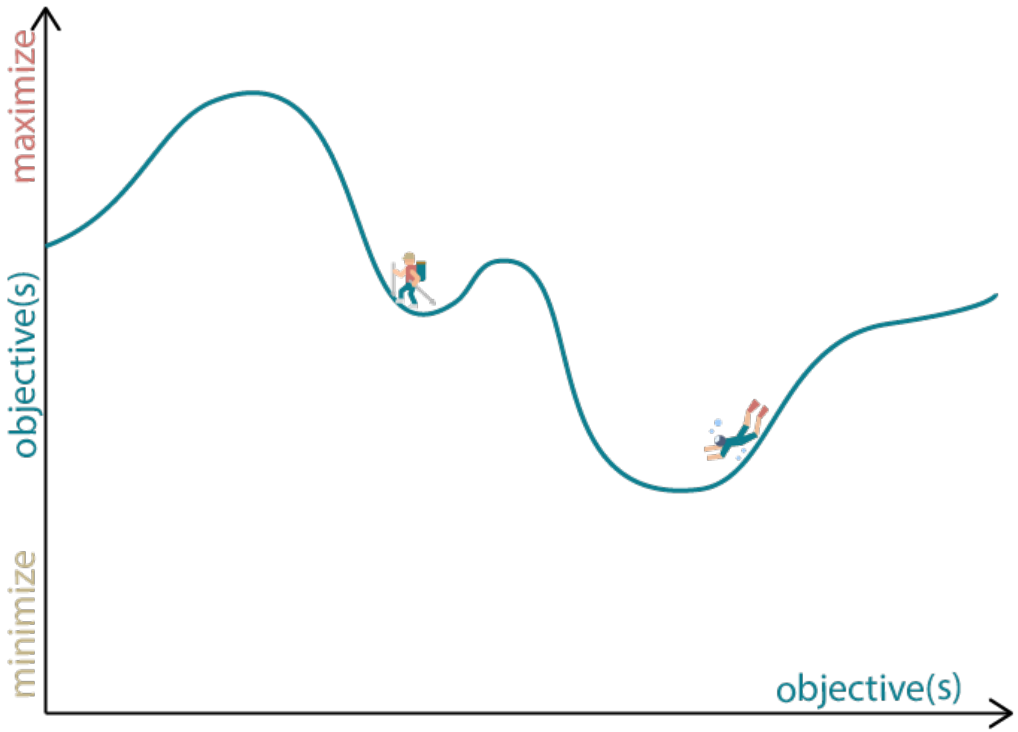

Constraints

- A constraint is a condition that the solution of an optimization problem must satisfy. In the table example we saw earlier, the constraints could be:
  - 'the table must have four legs'
  - 'the table must be at least 50cm wide'
  - 'the table may be no more than 1m tall', or
  - 'the table cannot be blue'.

- Constraints give a model its realism
  - they ensure that a solution only includes realistic values or values that the user knows are critical to the design brief.
  - If a model is unconstrained, it's likely to return absurd results that aren't useful
  - for example, here it could be a circular table that is three metres high with a single leg that balances on a point.

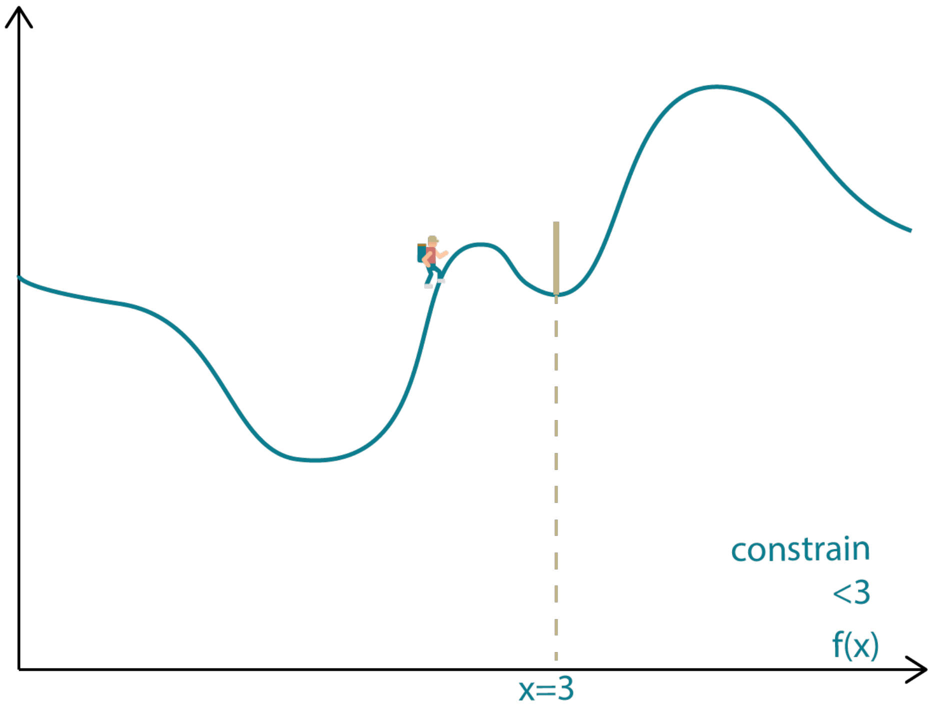

- Constraints are vital because they ensure that a generative design algorithm outputs something realistic and reasonable.

Data

- Data is the fundamental 'stuff' that you feed into your optimization model - the inputs and outputs that the objective function and constraints use to produce design solutions.
- In design, data could be the density or price of construction materials, how many hours of sunlight a room can expect to receive in a day, or any goals that are important to your design exploration that you can define mathematically. In finance, data could be the assets you can buy or sell, and their prices; or, in the aerospace industry, data could be the unit weights and costs of carrying a certain kind of fuel.

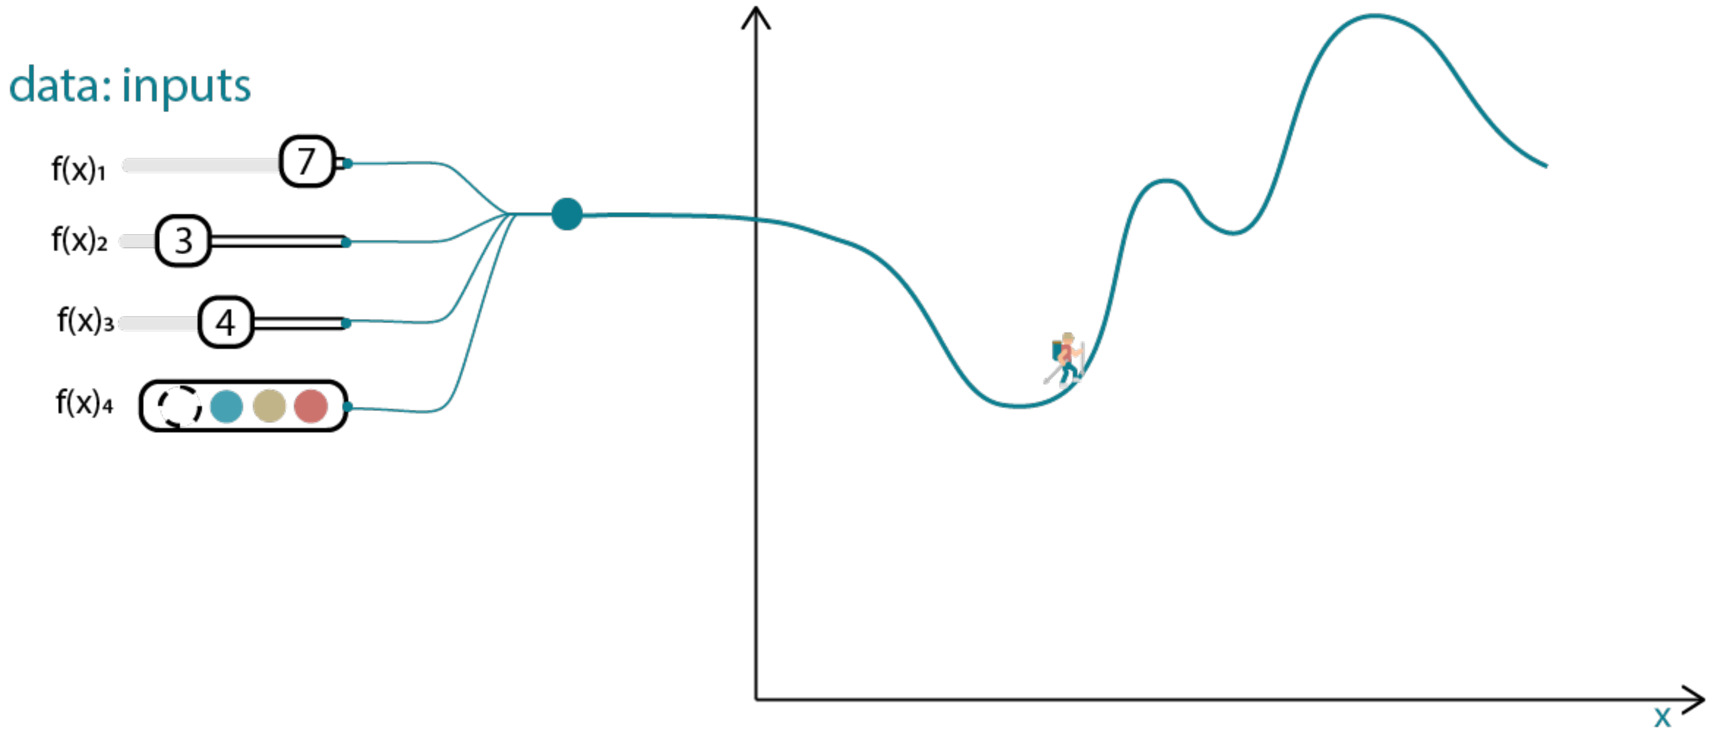

Data

- The optimization model in the generative design toolkit takes this data and uses it to maximize or minimize values as specified by the designer.
- Real-world optimization problems are invariably solved algorithmically and there are often many algorithms that can solve the same problem. The most common algorithm used in Generative Design for architectural and engineering workflows is called the 'genetic algorithm'. We will cover this later on.

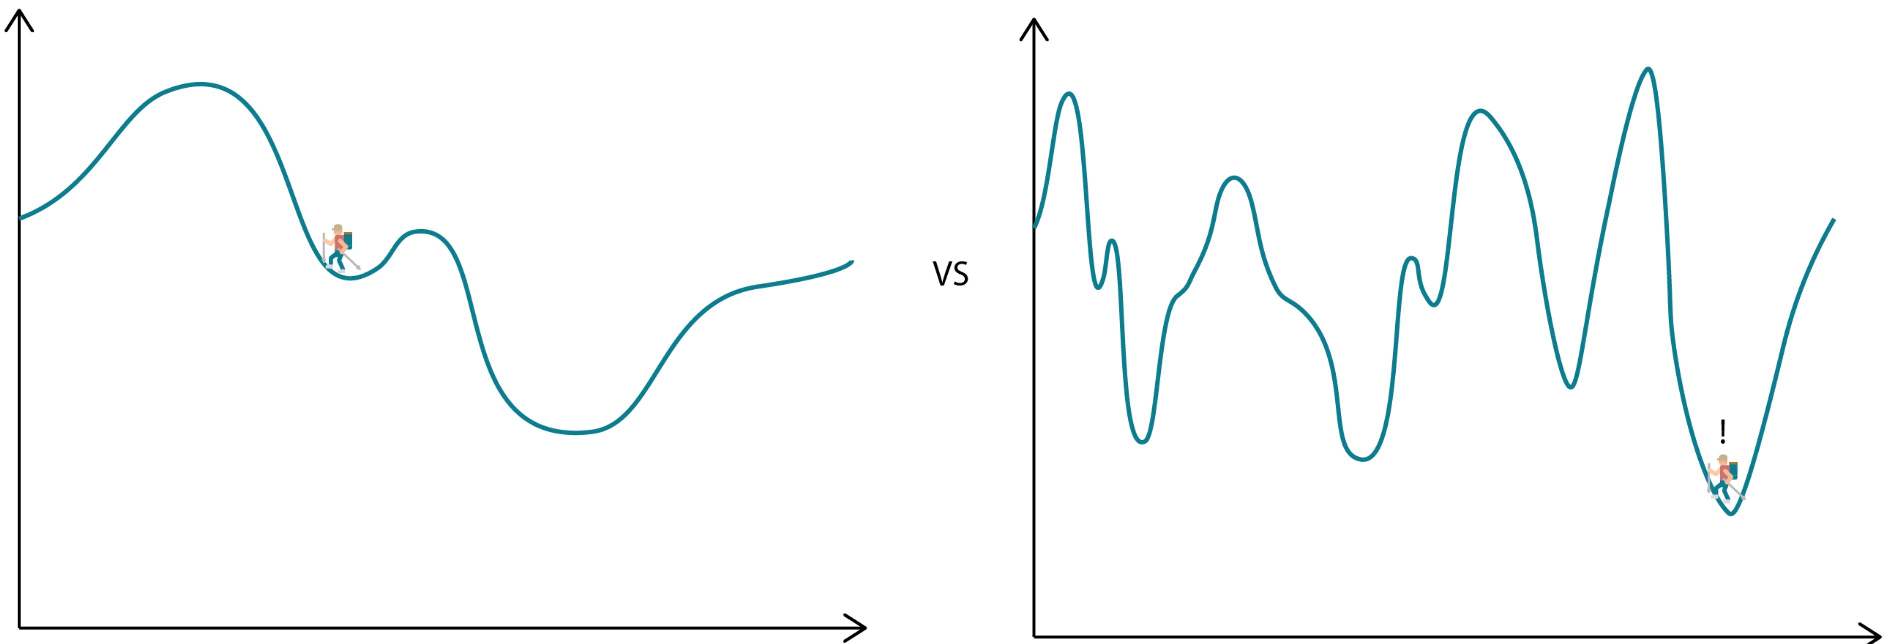

easy to navigate graph

hard to navigate graph

Defining Goals

- When it comes to generative design processes, it is vital that you know your design parameters inside and out. Every good generative design project starts with a clear and precise understanding of the design problem and a clear description of the goals.

- Algorithms are great at churning through thousands of design options very quickly, but they don’t perform nearly as well if they're given vague or imprecise instructions. You must be able to define your problem in a mathematical way (i.e., with some sort of number that can be used to rank outcomes).

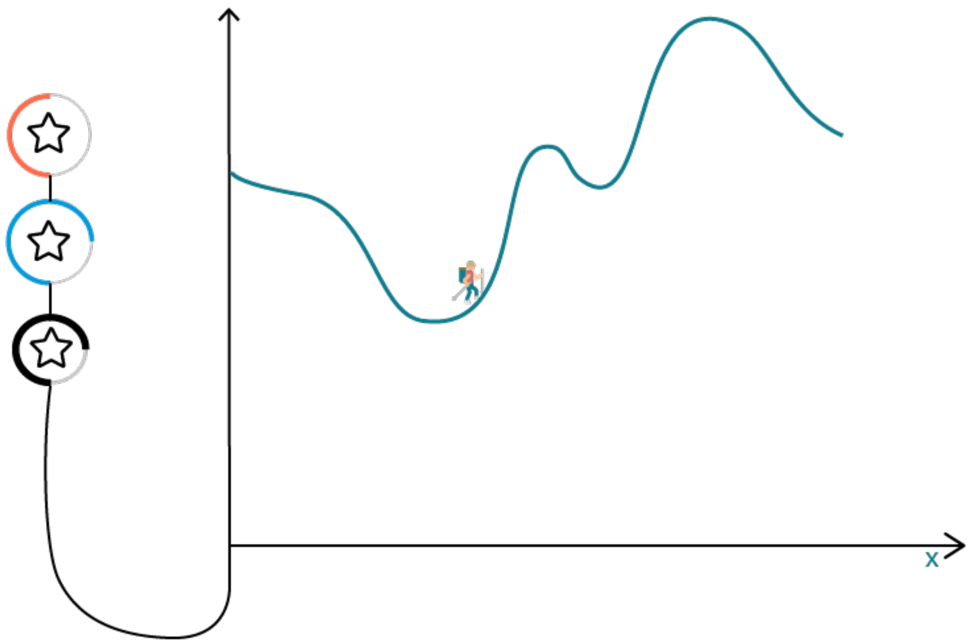

Defining Goals

- Some good questions to ask when formulating design goals are:
  - What do you want to achieve?
  - Which features must your ideal design have?
  - Which features cannot appear in your ideal design?
  - Do you simply want to see a lot of design options?
  - Do you want to optimize your design for some specific characteristic?
  - Do you want to optimize your design for multiple competing characteristics?
  - What would you like to maximize? Why?
  - What would you like to minimize? Why?
  - Can your maximization or minimization question be quantified mathematically? If so, how precisely?

- Being able to confidently answer at least some of the questions above is a good first step to figure out precisely which objectives your computational procedure should have.

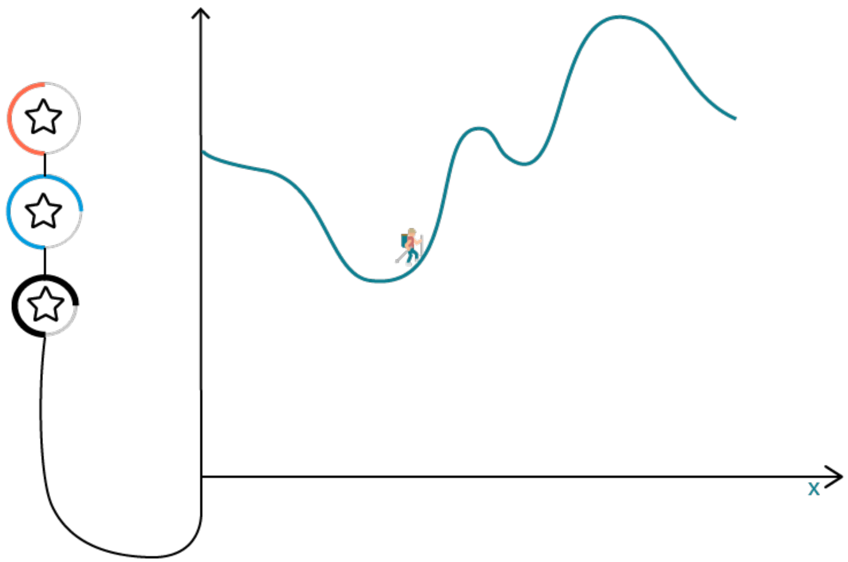

Decision Making as an Optimization Problem

- Be aware of the complete optimization problem helps you
  - to identify the challenging facets of your optimization problem and, thus,
  - to select a suitable algorithm

Refer to the text if interested

- Variable Types
- Number of Variables
- Number of Objectives
- Constraints
- Multi-modality
- Differentiability
- Evaluation Time
- Uncertainty

Variable Types.

- The variables span the search space Ω of your optimization problem. Thus, the type of variables is an essential aspect of the problem to be paid attention to. Different variables types, such as continuous, discrete/integer, binary, or permutation, define the characteristics of the search space. In some cases, the variable types might be even mixed, which increases the complexity further.

Number of Variables.

- Not only the type but also the number of variables (N) is essential. For either a very small or large number, different algorithms are known to work more efficiently. You can imagine that solving a problem with only ten variables is fundamentally different from solving one with a couple of thousand. For large-scale optimization problems, even the second-order derivate becomes computationally very expensive, and efficiently handling the memory plays a more important role.

Number of Objectives.

- Some optimization problems have more than one conflicting objective ($M>1$) to be optimized. Before researchers have investigated multi-objective optimization, single-objective problems were the main focus. Single-objective optimization is only a particular case where $M=1$. In multi-objective optimization, the solution's domination relation generalizes the comparison of two scalars in single-objective optimization. Moreover, having more than one dimension in the objective space, the optimum (most of the time) consists of a set of non-dominated solutions. Because a set of solutions should be obtained, population-based algorithms have mainly been used as solvers.

Constraints.

- Optimization problems have two types of constraints, inequality (g) and equality (h) constraints. From an end-user perspective, constraints have a priority over objective values. No matter how good the solution's objectives are, it is considered infeasible if it turns out to violate just a single constraint. Constraints can have a big impact on the complexity of the problem. For instance, if only a few islands in the search space are feasible or a large number of constraints ($|J|+|K|$) need to be satisfied. For genetic algorithms satisfying equality constraints can be rather challenging. Thus, this needs to be addressed differently, for instance, by mapping the search space to a utility space where the equality constraints are always satisfied or injecting the knowledge of the equality constraint through customization.

Multi-modality.

- Most aspects discussed so far are most likely known or to be relatively easy to define. However, the nature of the fitness landscape is less obvious bet yet essential to be aware of. In the case of multi-modal fitness landscapes, optimization becomes inevitably more difficult due to the existence of a few or even many local optima. For the solution found, one must always ask if the method has explored enough regions in the search space to maximize the probability of obtaining the global optimum. A multi-modal search space quickly shows the limitation of local search, which can easily get stuck.

Differentiability.

- A function being differentiable implies the first or even second-order derivative can be calculated. Differentiable functions allow gradient-based optimization methods to be used, which can be a great advantage over gradient-free methods. The gradient provides a good indication of what direction shall be used for the search. Most gradient-based algorithms are point-by-point based and can be highly efficient for rather unimodal fitness landscapes. However, in practice, often functions are non-differentiable, or a more complicated function requires a global instead of a local search. The research field addressing problems without knowing their mathematical optimization is also known as black-box optimization.

Evaluation Time.

- Many optimization problems in practice consist of complicated and lengthy mathematical equations or domain-specific software to be evaluated. The usage of third-party software often results in a computationally expensive and time-consuming function for evaluating objectives or constraints. For those types of problems, the algorithm's overhead for determining the next solutions to be evaluated is often neglectable. A commercial software performing an evaluation often comes with various more practical issues such as distributed computing, several instances to be used in parallel and software license, and the software's possible failure for specific design variable combinations.

Uncertainty.

- Often it is assumed that the objective and constraint functions are of a deterministic manner. However, if one or multiple target functions are nondeterministic, this introduces noise or also referred to as uncertainty. One technique to address the underlying randomness is to repeat the evaluation for different random seeds and average the resulting values. Moreover, the standard deviation derived from multiple evaluations can be utilized to determine the performance and the reliability of a specific solution. In general, optimization problems with underlying uncertainty are investigated by the research field called stochastic optimization.

Design as Decision Making

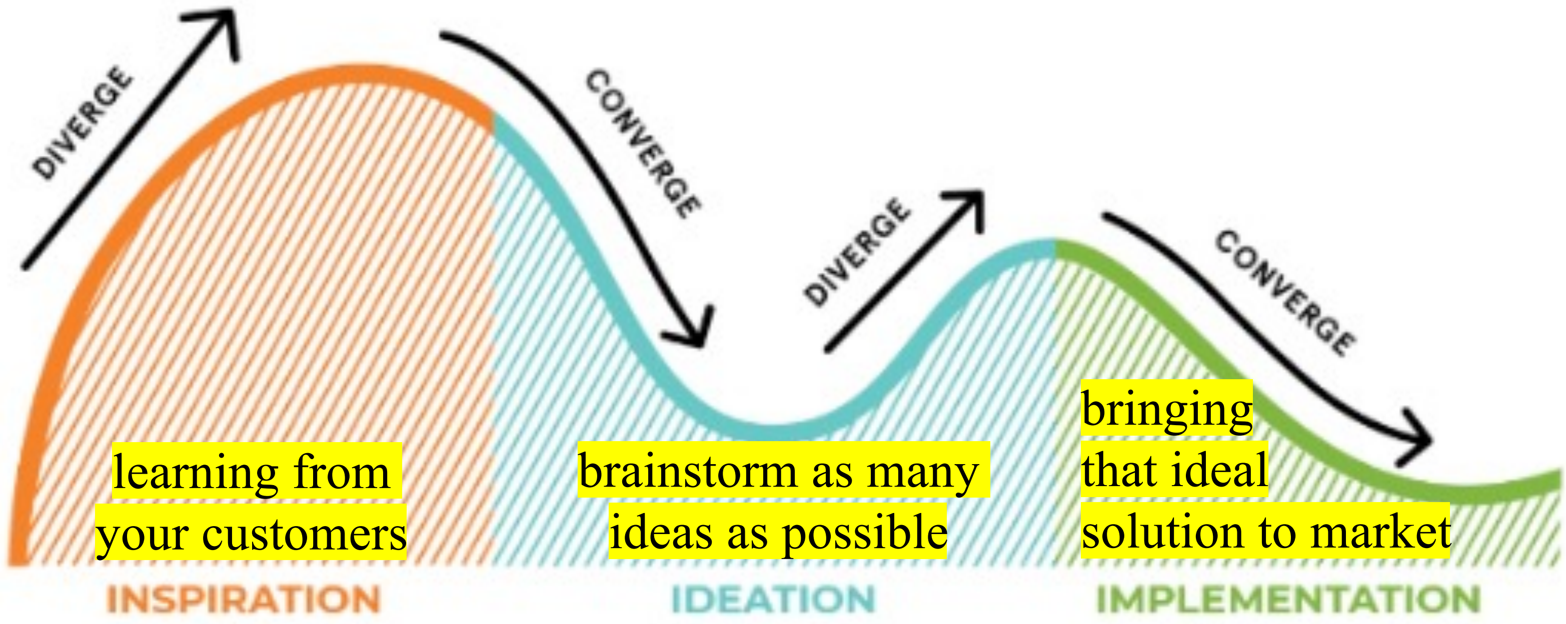

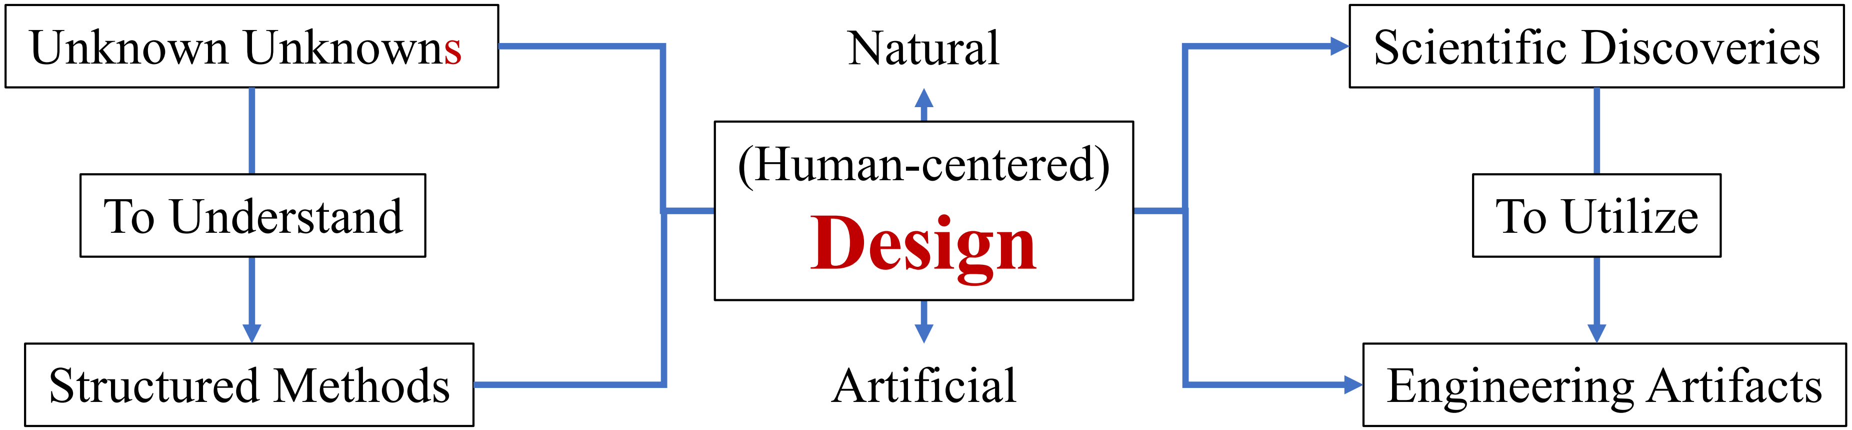

Principals of Data-driven Technology

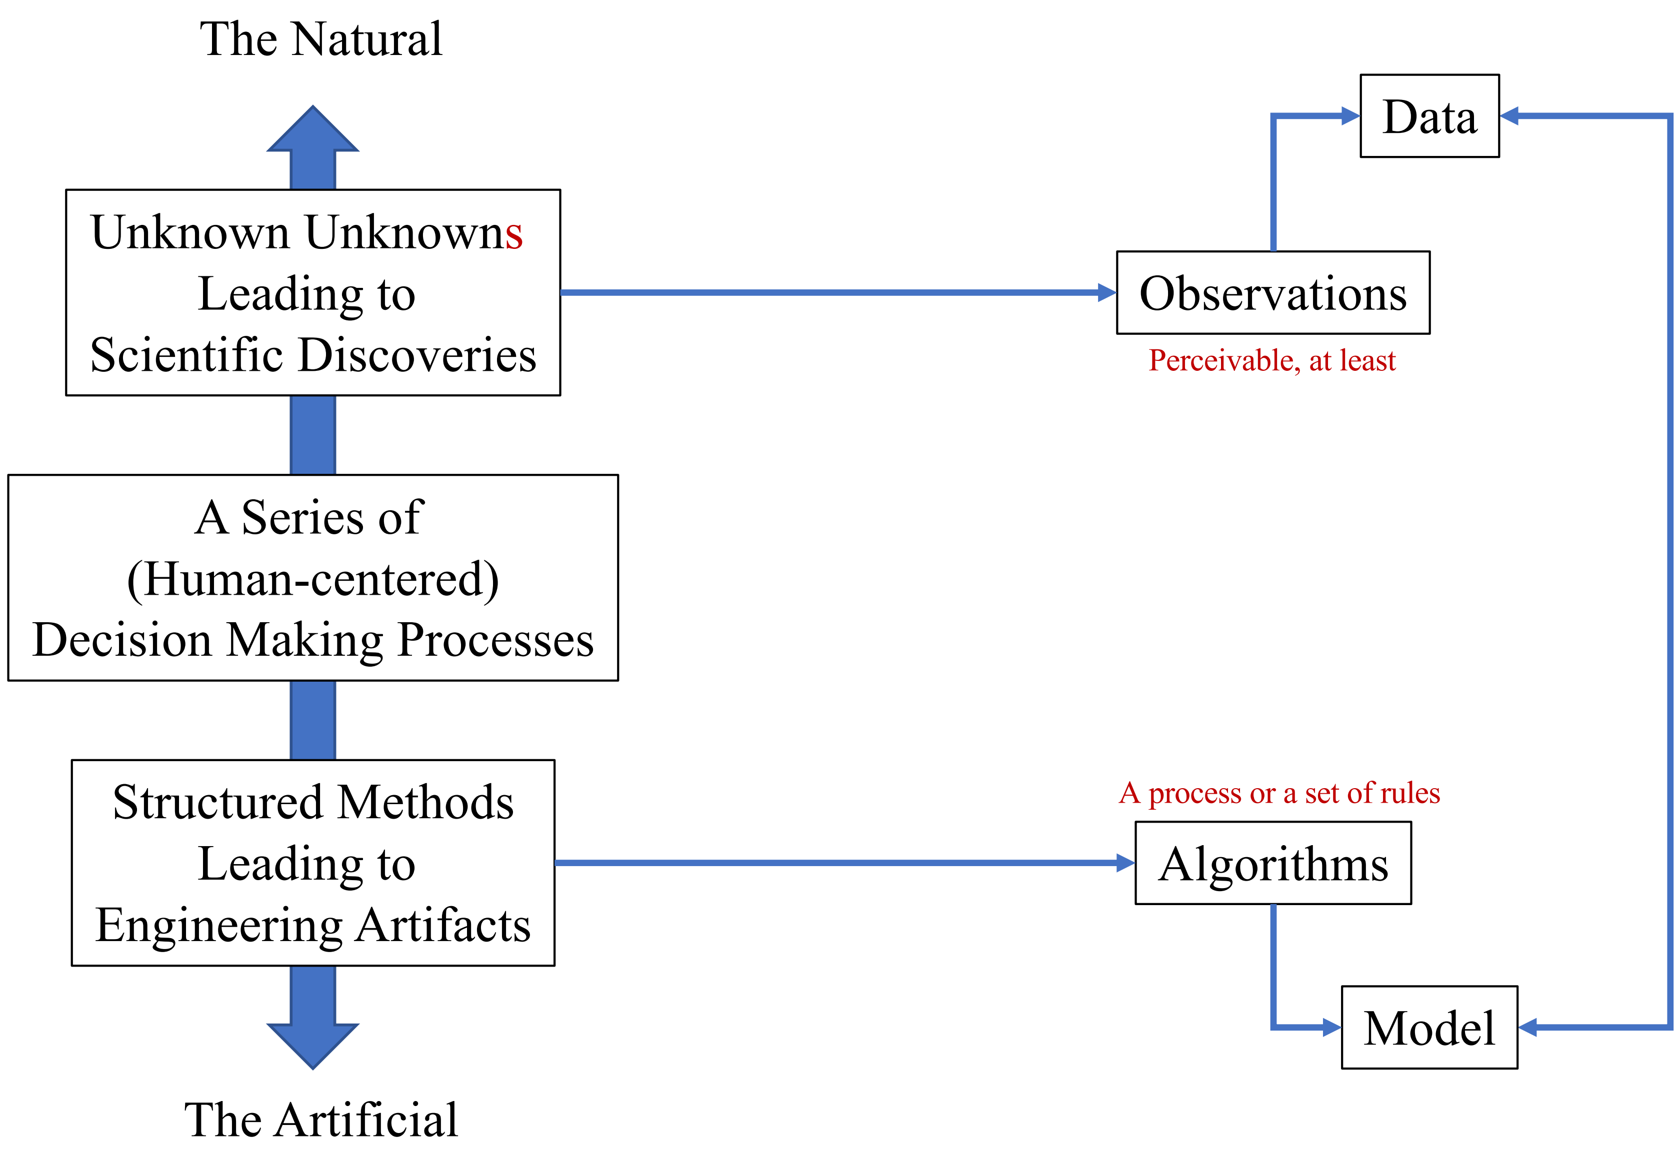

Principals of Data-driven Technology

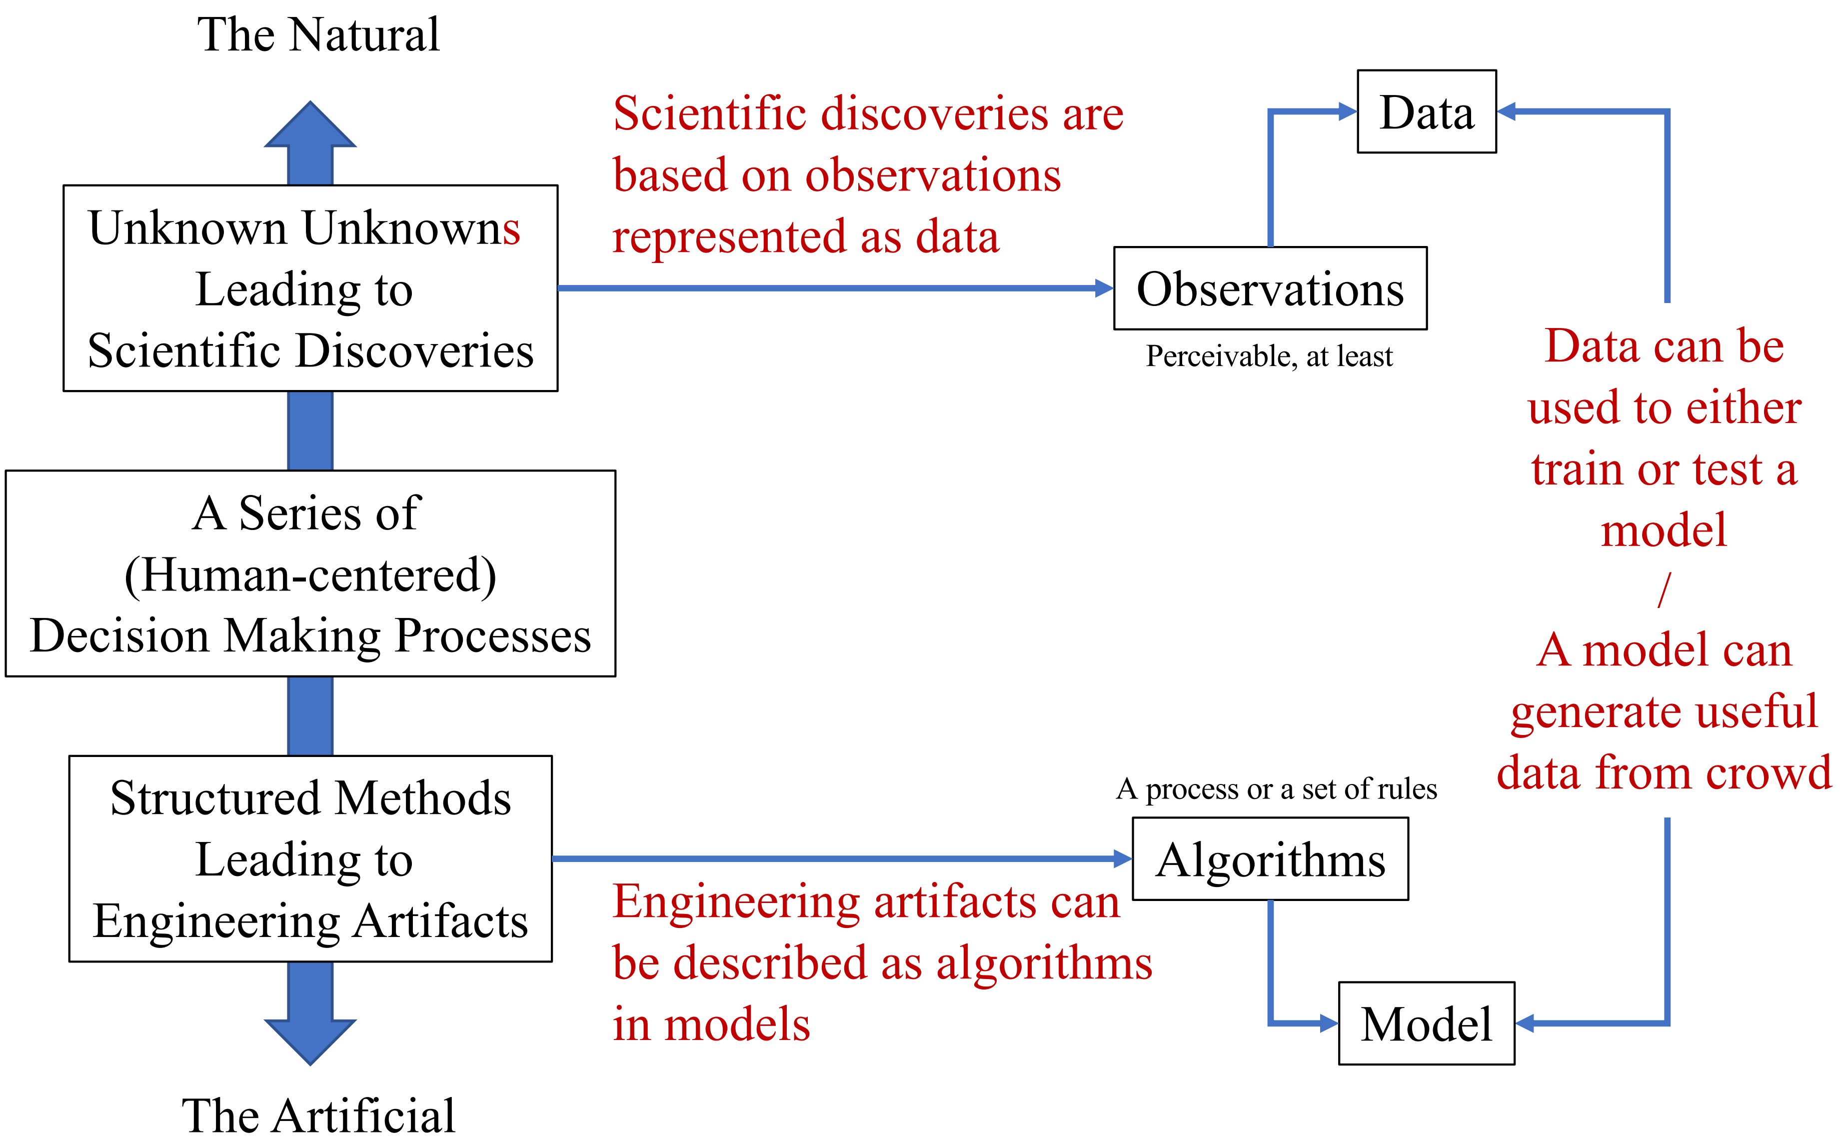

Principals of Data-driven Technology

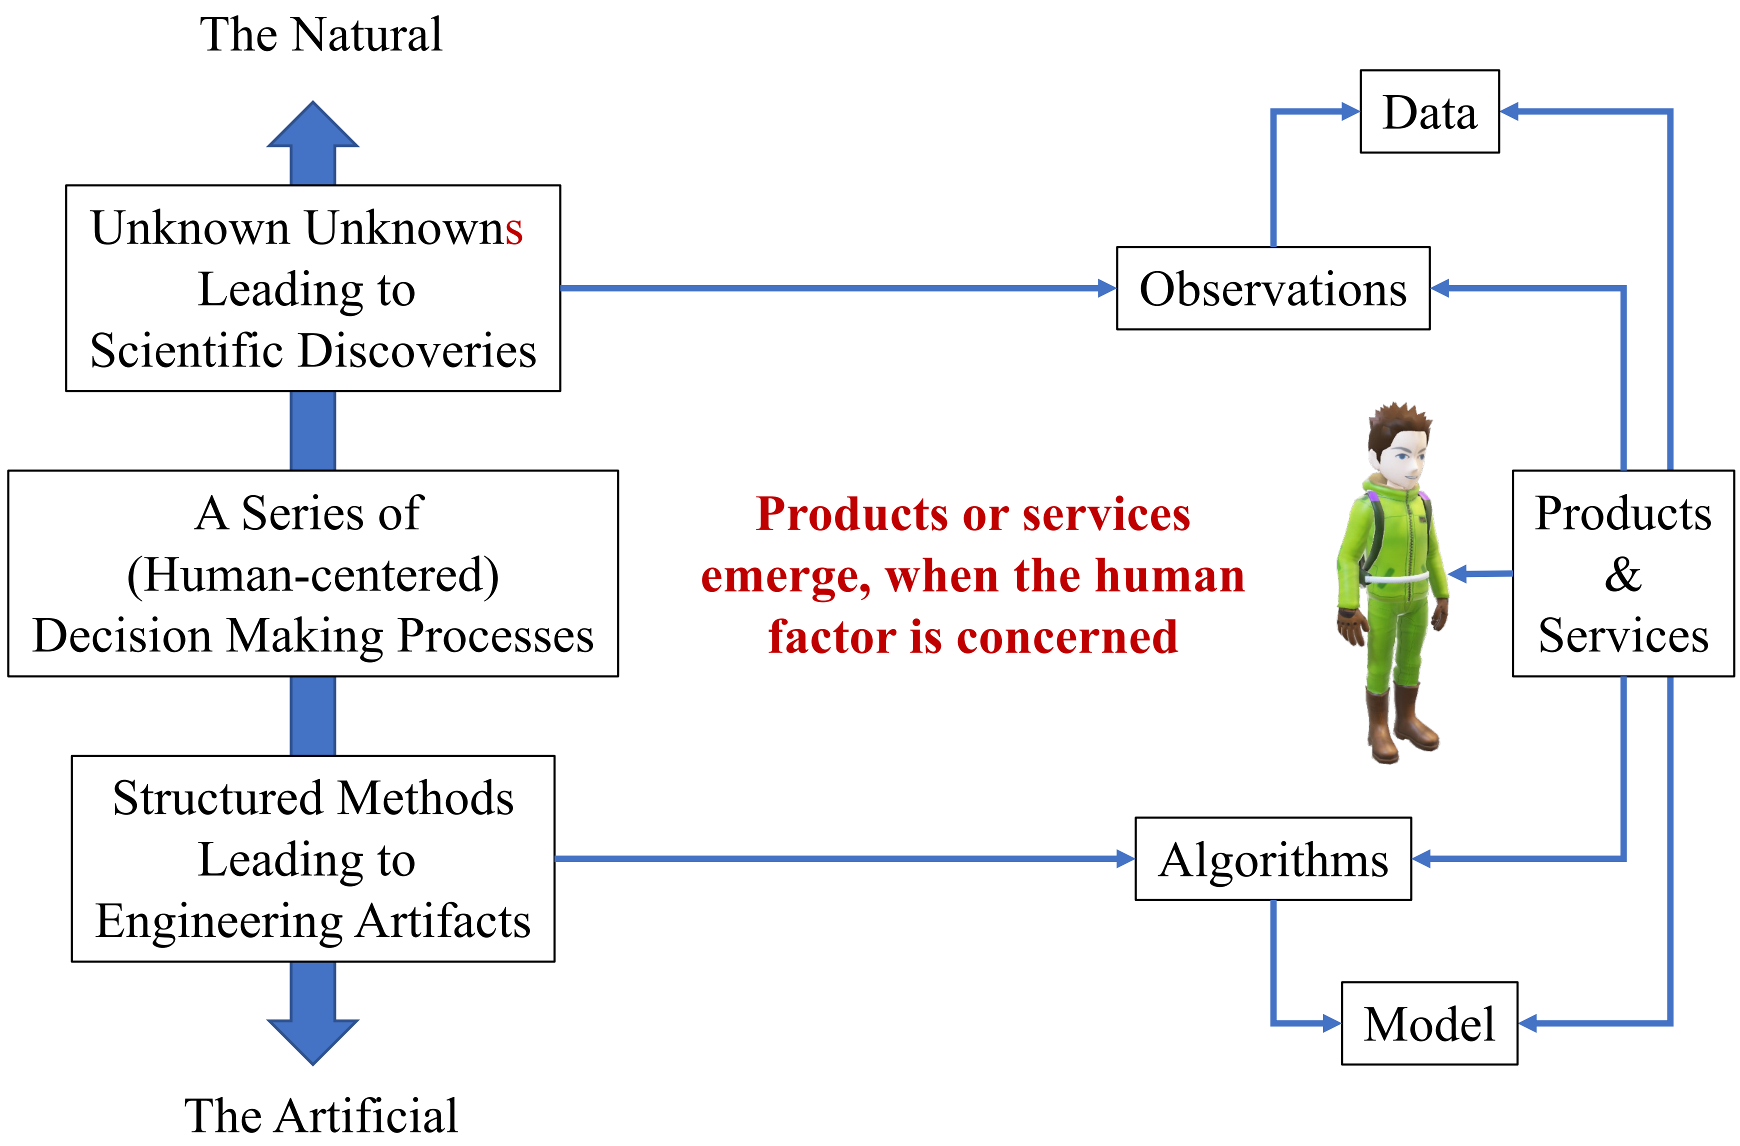

Principals of Data-driven Technology

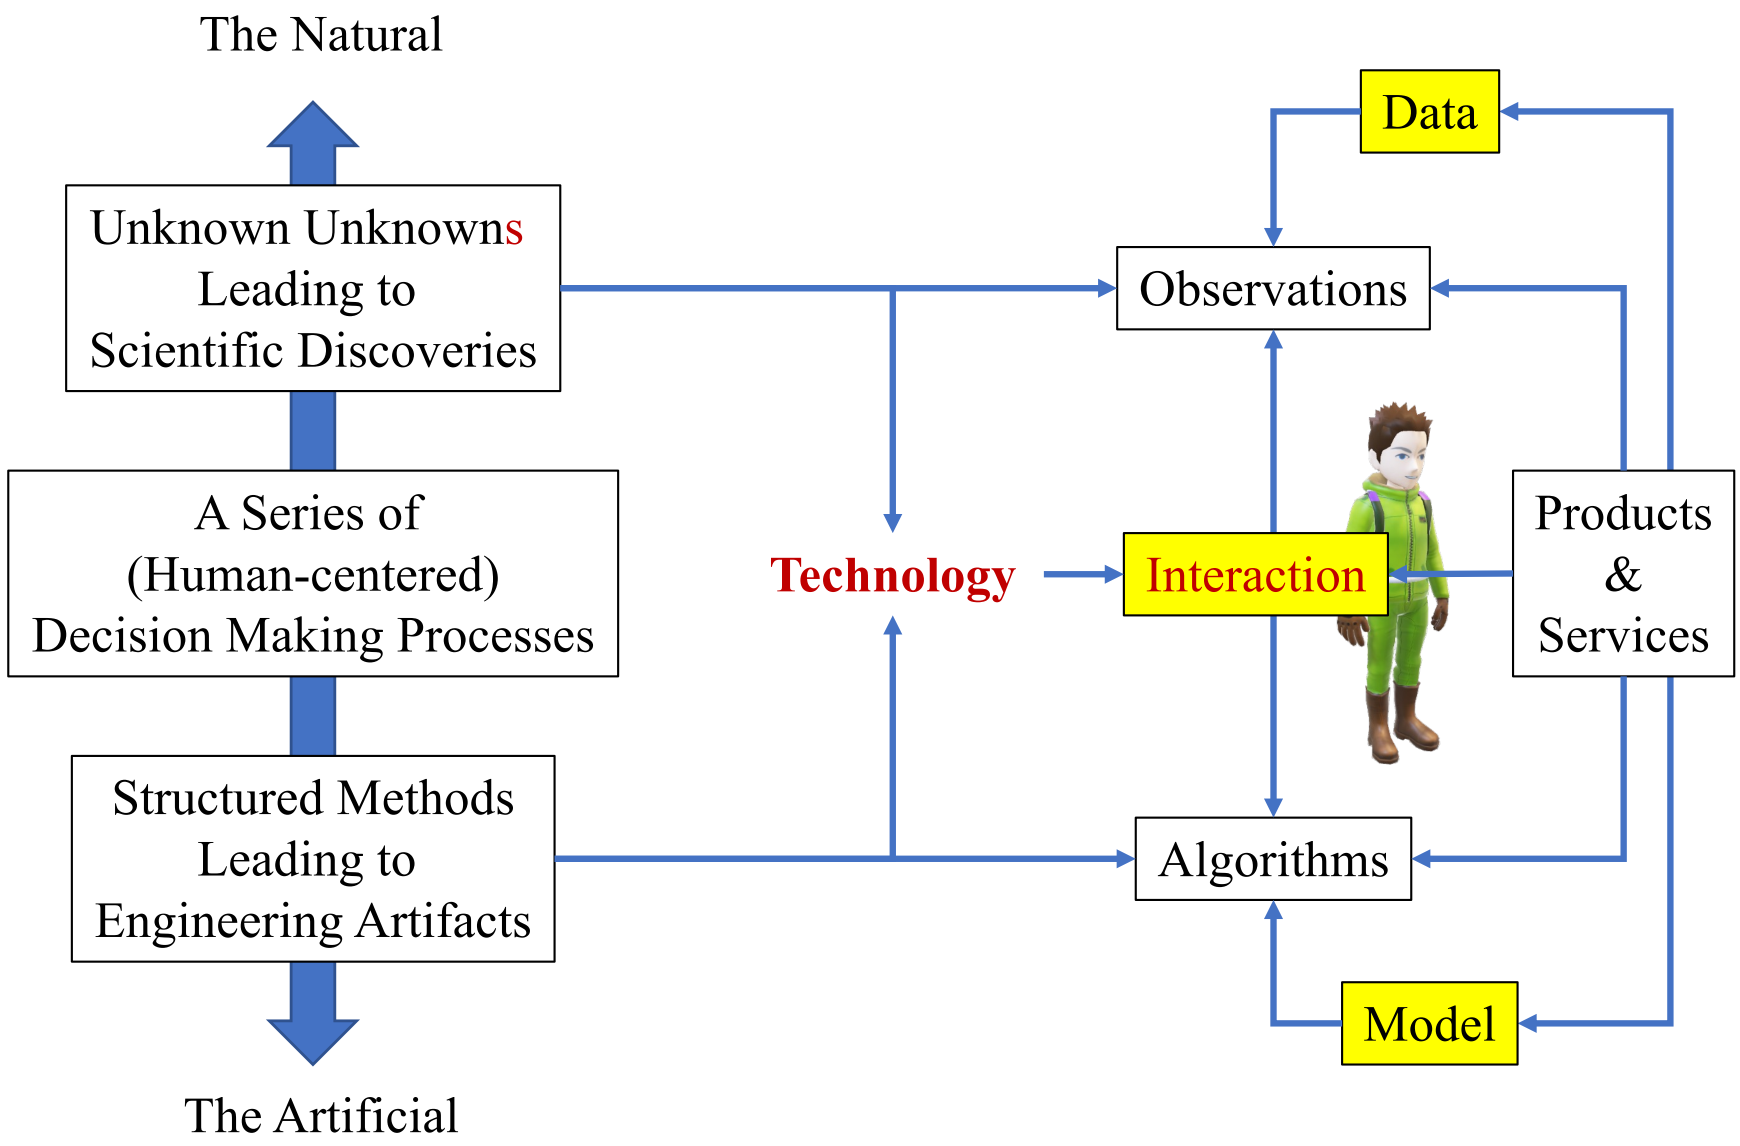

The definition of AI

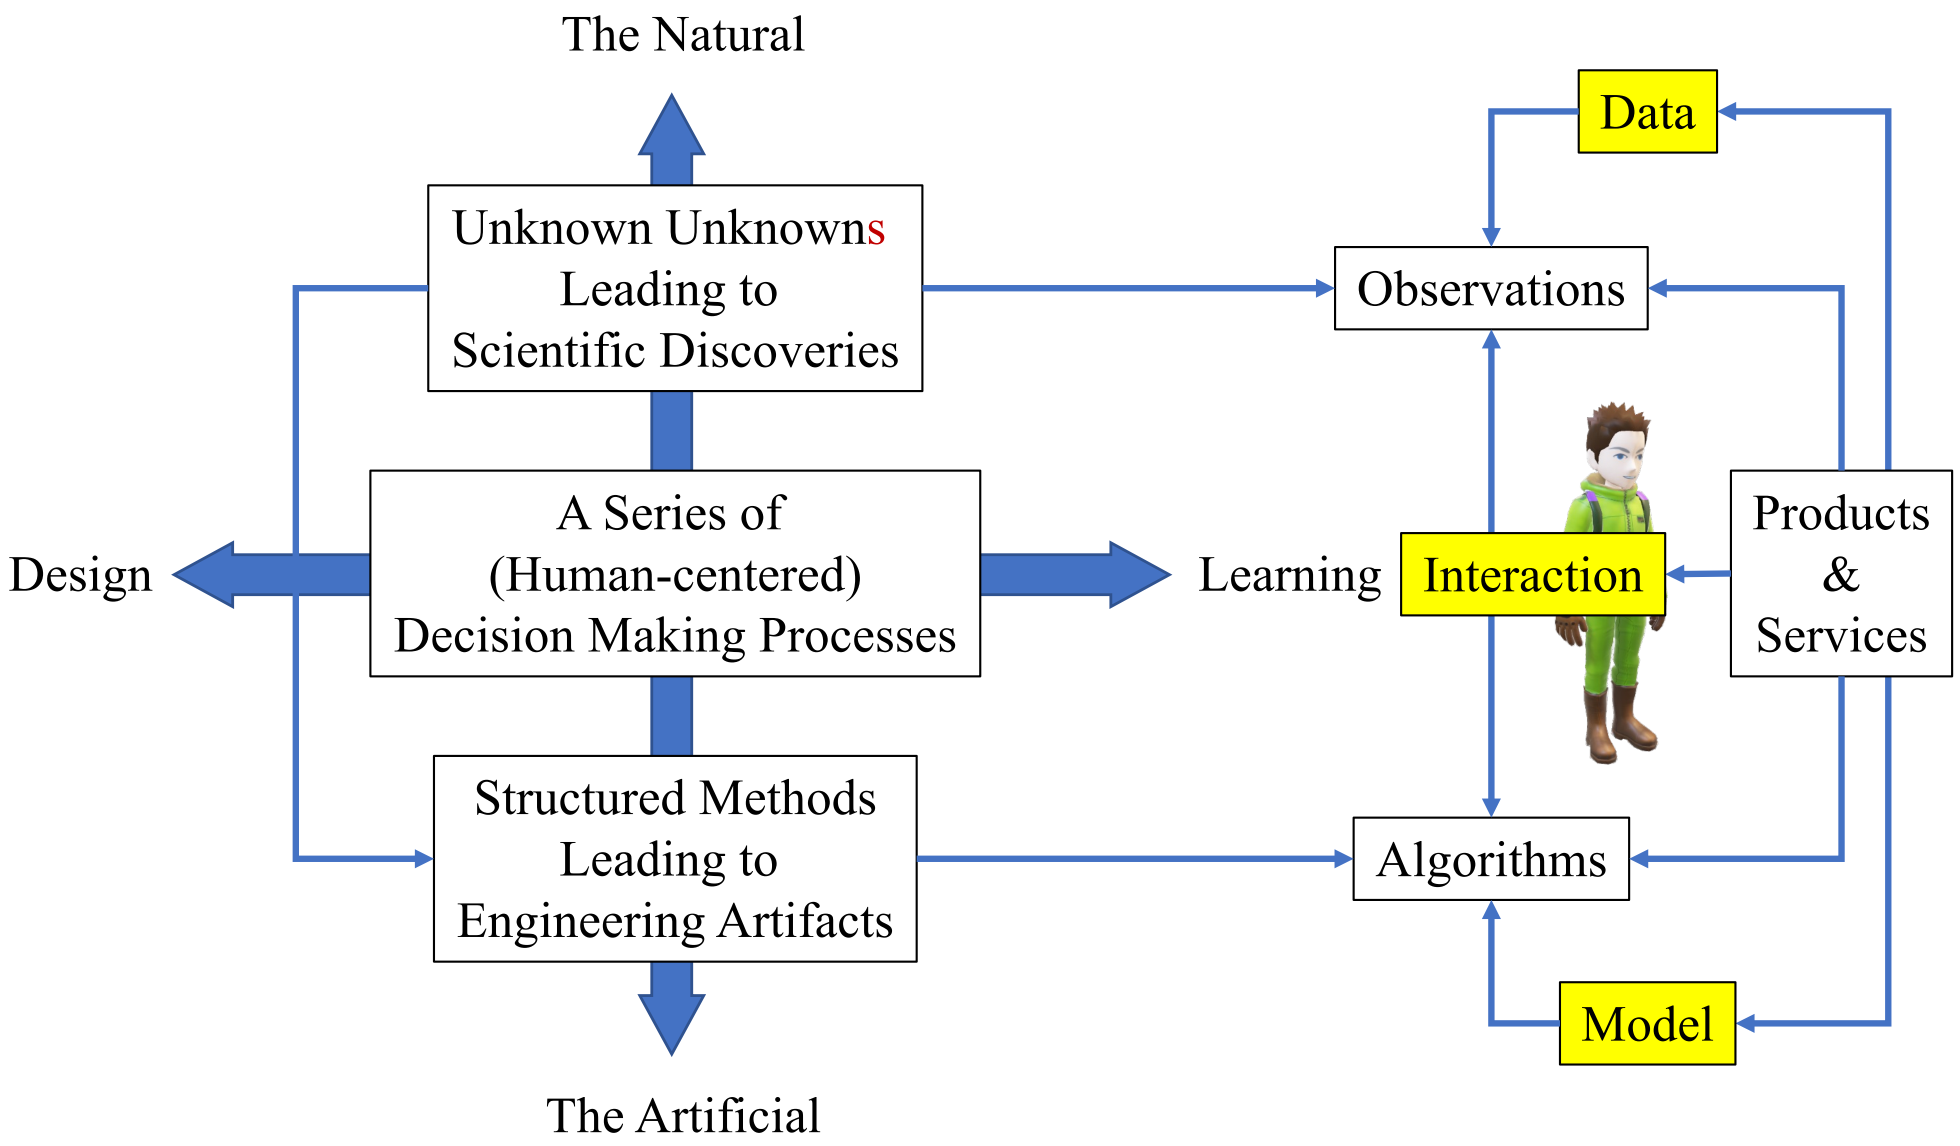

What do we mean by design?

The typical design process

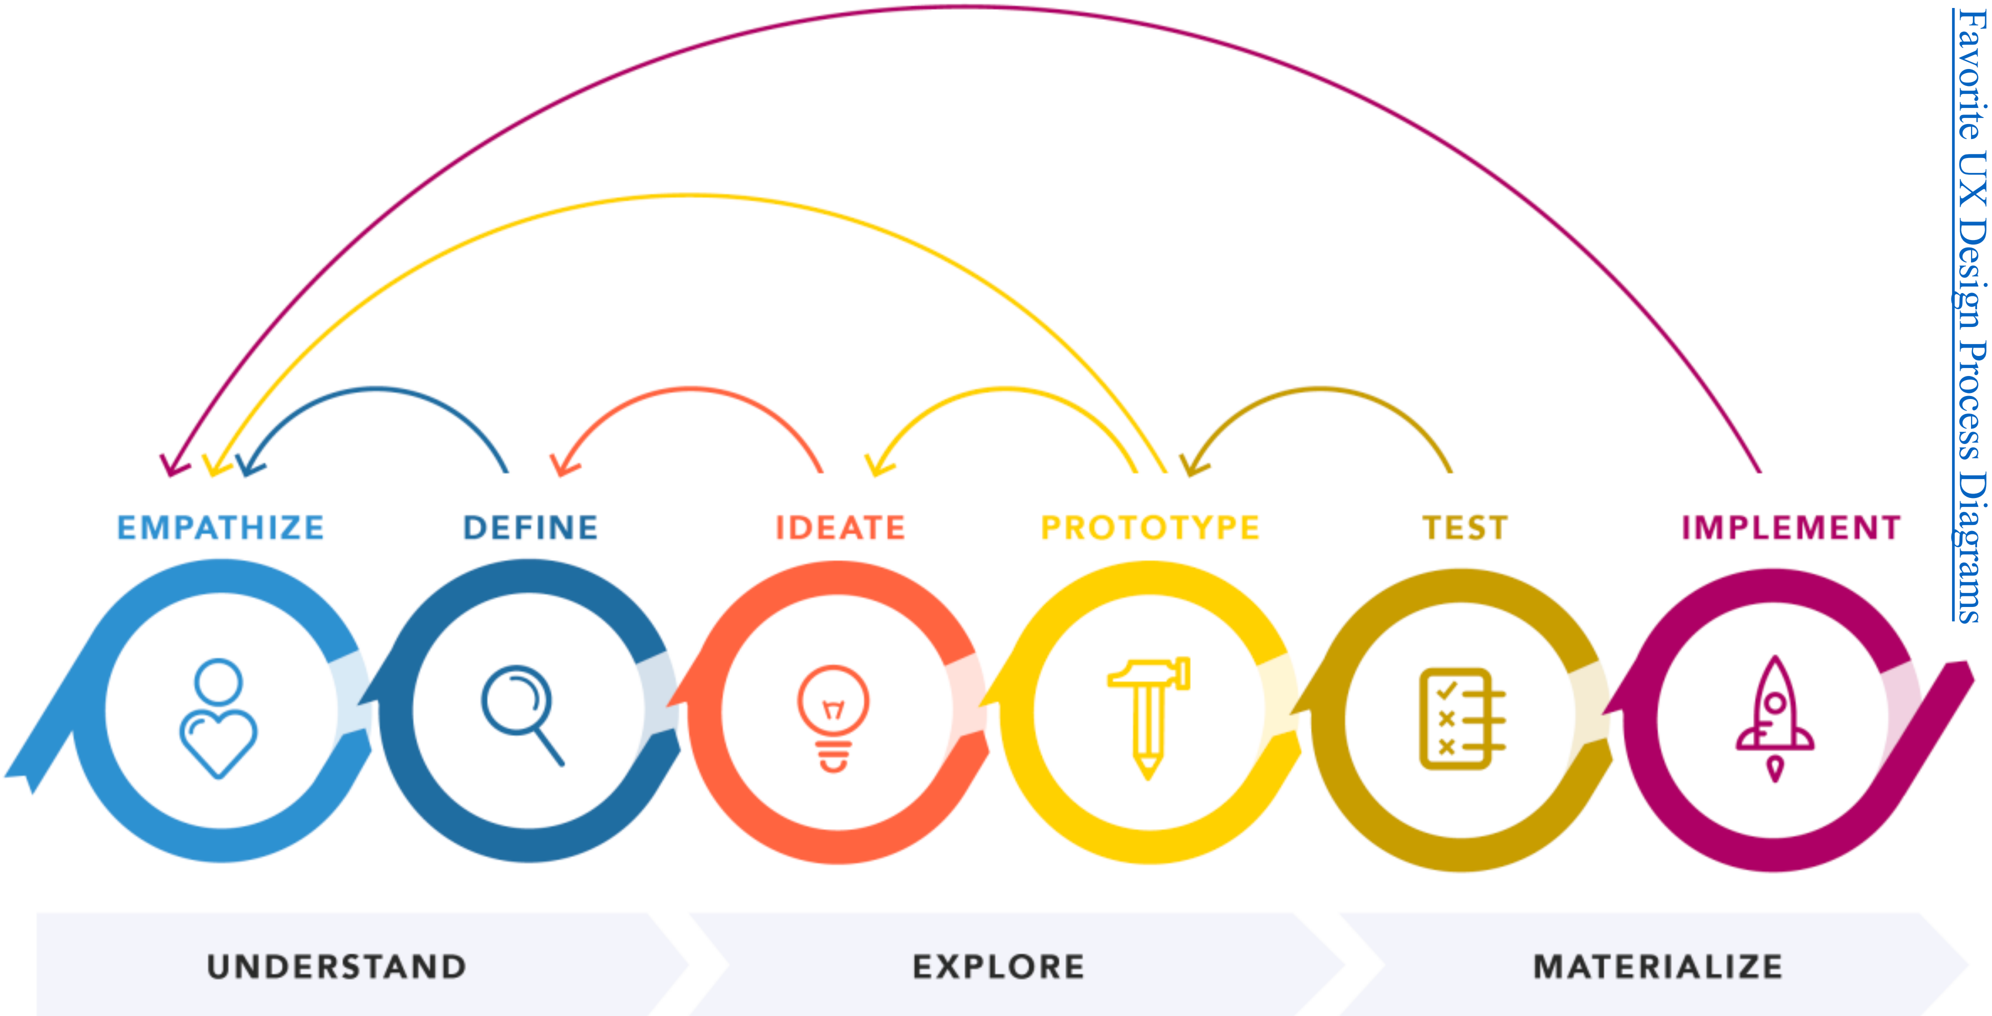

Favorite UX Design Process Diagrams

The task of "designing AI products"

- If you were a designer and work with Tencent AI lab to develop a new product/service
  - Will you embark on a user-centered design process and select a group of target users?

  - There is no guarantee that the user research will reveal a need for AI.
  - Even if designers do find a user need for AI, there is no guarantee that there will already be (or ever be) data to train the system.

Potential "seeds" to drive envisionment

Dataset

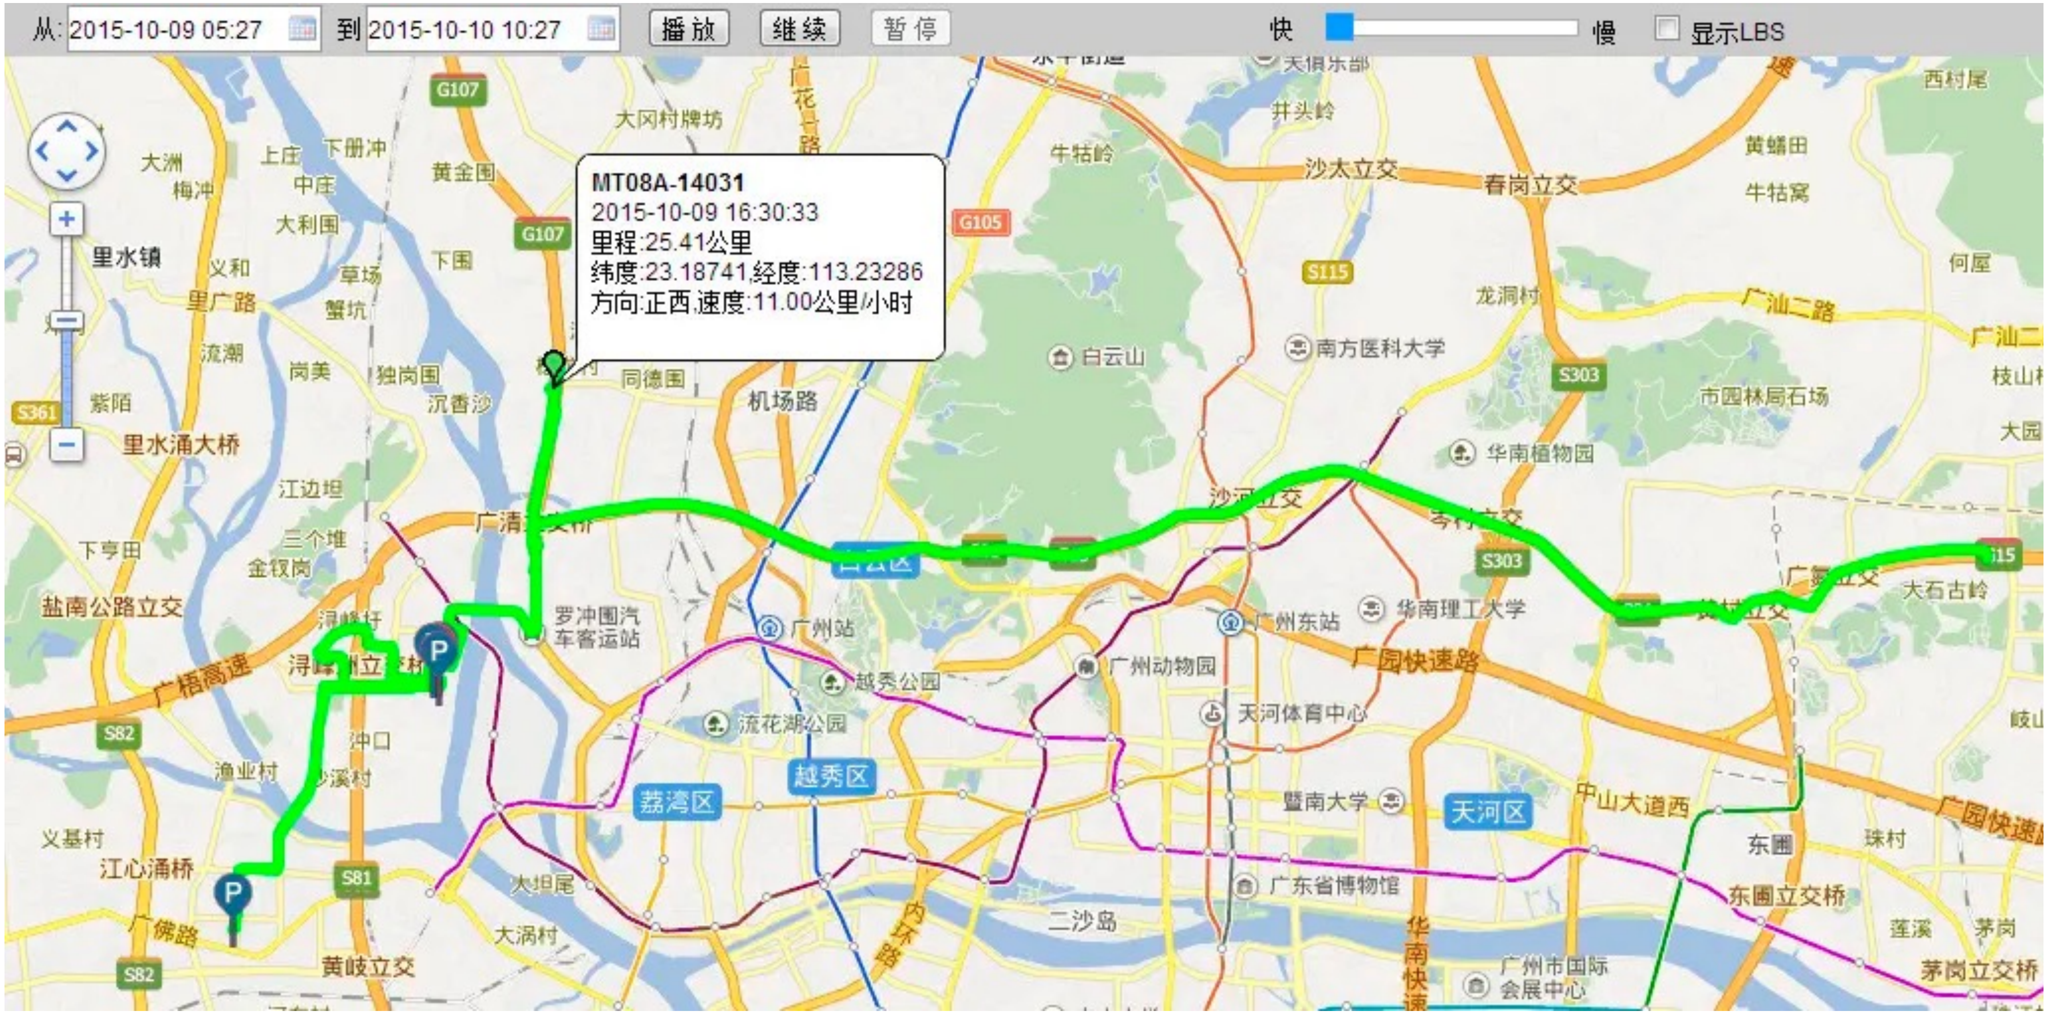

Potential “seeds” to drive envisionment

A technical capability (a patent)

Design through matchmaking

- AI meets Design Toolkit

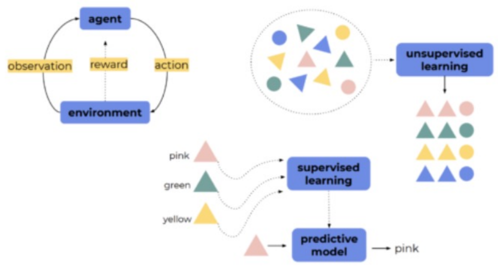

A 2-page crash course in AI/ML to help you get up to speed on different kinds of artificial intelligence and machine learning

Prompts to start spotting opportunities based on user needs, AI capabilities, and data availability and a card deck with 30 prompts of common AI applications for ideation sessions

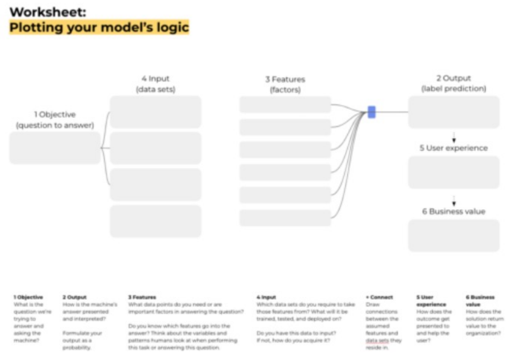

Exercises to align with your machine learning engineers and data scientists on the model, confusion matrix, and evaluation metrics

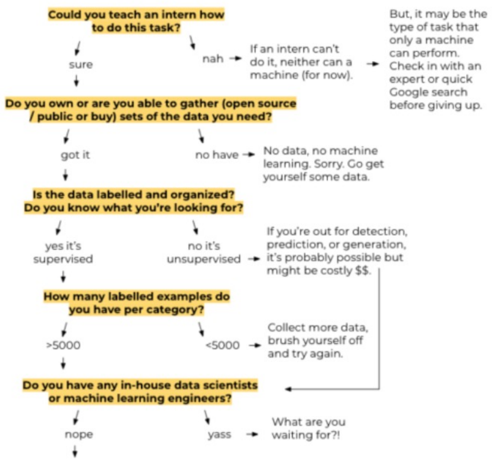

Assessing feasibility, viability & desirability

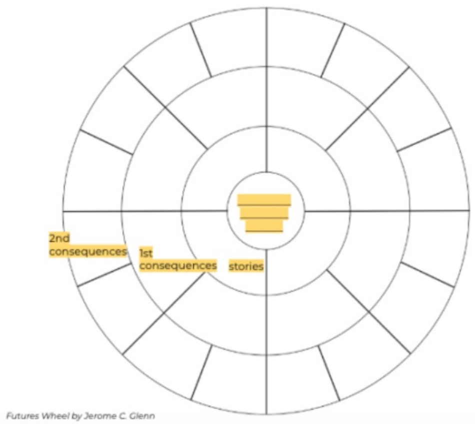

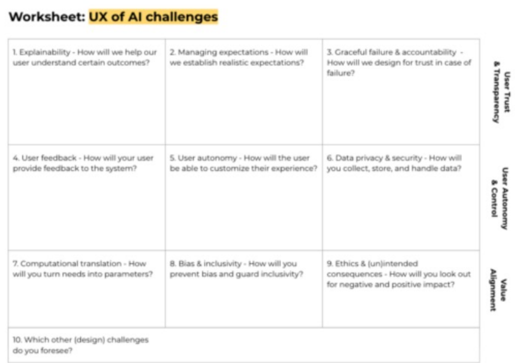

An overview of 9 UX and design challenges of AI as a material

Thank you~

Wan Fang
Southern University of Science and Technology