Day 2

- 8:00 – 8:30 Reflection on day 1
- 8:30 – 9:00 Lecture: AI Meets Design intro
- 9:00 – 9:50 Exercise: Explore and Identify ML capacities
- 10:20 – 11:10 Lecture: AI needs Data and Data needs AI
- 11:20 – 12:10 Exercise: Play, Observe and understand Data

- 2:00 – 5:00 Exercise: AI Meets Design Activity
- 5:00 – 6:00 Review of the day

Welcome

Welcome to the AI meets Design toolkit: a set of tools to help design with and for machine intelligence. It is an invitation to designers and innovators everywhere to get involved by leveraging the opportunities and navigating the challenges of AI to create human-centered applications and meaningful user experiences.

What does this toolkit do?

This toolkit aims to build a bridge between the disciplines of design and the disciplines of machine learning. It will help you to:

- spot opportunities to leverage AI for user, business, and social value within your context
- adapt and apply the design thinking practice for AI concepts
- align user needs and guard human values within algorithmic systems
- communicate and collaborate with data scientists and ML engineers

Who is it for?

Designers and innovators with a role in shaping digital products and services. Anyone looking to transform this raw material into user, business, and social value. Engineers can also refer to it to promote adoption across disciplines.

Why a toolkit?

Based on interviews with >50 designers from across the world, we learned that while interest in AI is high, its high-tech character and a lack of practical tools keeps adoption low and design innovation limited. Building on the research findings, the goal is of this toolkit is to integrate AI with the design process, highlight its user-centric considerations, and make it accessible to all beyond the engineering field.

Why would I care?

AI is an important and inevitable development of our time. While offering huge potential, its implementation does come with a complex set of challenges. Our future is not determined and we can all play a role in designing AI to help, not harm, humanity at large.

How to use it?

As you see fit, really. It integrates with the steps of the design thinking process. You can go through the process end-to-end to develop new ideas. If you already have AI concepts to develop, or even on the market, you simply take what's useful to you. You can use it by yourself, within your team and across departments, with a client, or get us to facilitate it for you.

What's in this toolkit?

This toolkit contains a set of tools including exercises, worksheets, and card decks to assist designers at the different stages of the design thinking process.

Introduction

Welcome
What's in this toolkit
A crash course in AI+ML

Ideation

User-centered problem solving
Tech-driven opportunity spotting
Data-driven opportunity spotting
AI prompt card deck for ideation *

Start here if:
You want to explore opportunities for AI within your context

Concept development + idea selection

Impact matrix for idea selection
Value proposition design *
Assessing feasibility
Framing your task
Plotting your model *

You have a bunch of ideas and want to select and develop them in more detail

Prototyping + testing

User research & feedback
Prototyping & testing

You have a handful of ideas and want to validate which to move forward with

Design + implementation

Defining success and failure *
Mapping user needs to models
Metrics to evaluate by *
UX and design challenges of AI *
Capturing design tensions *
Consequence wheel *

You are ready to start building with your engineers and bring your idea to life

A crash course in AI + ML

A one-pager to get you up to speed on some core concepts including the difference between AI and ML, and the various types of machine learning.

artificial intelligence (AI)
= the science of getting machines to learn, think, act, and perform tasks in ways traditionally attributed to human intelligence

types of AI

narrow AI
= equals or exceeds human intelligence or efficiency at a very specific thing

general AI
= match human intelligence across domains + tasks

super AI
= exceeding human intelligence

not here (yet)

types of ML

machine learning (ML)
= the ability for machines to learn and infer from large sets of examples and experience instead of explicitly programming the rules

deep learning
= artificial neural networks inspired by the human brain capable of learning from data that is unstructured

types of learning

reinforcement learning
= collect data on the go and learn from trial and error to achieve an objective (below left)

supervised
= examples and data are labelled (below bottom)

unsupervised
= find patterns in large, non-labelled data sets (below right)

A crash course in AI + ML

A one-pager to get you up to speed on some core concepts including the difference between AI and ML, and the various types of machine learning.

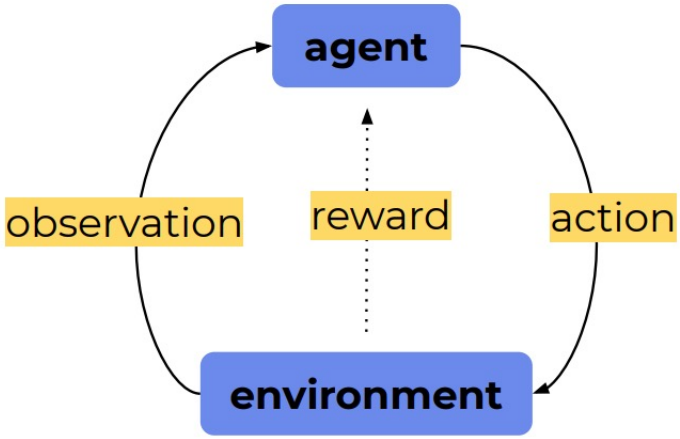

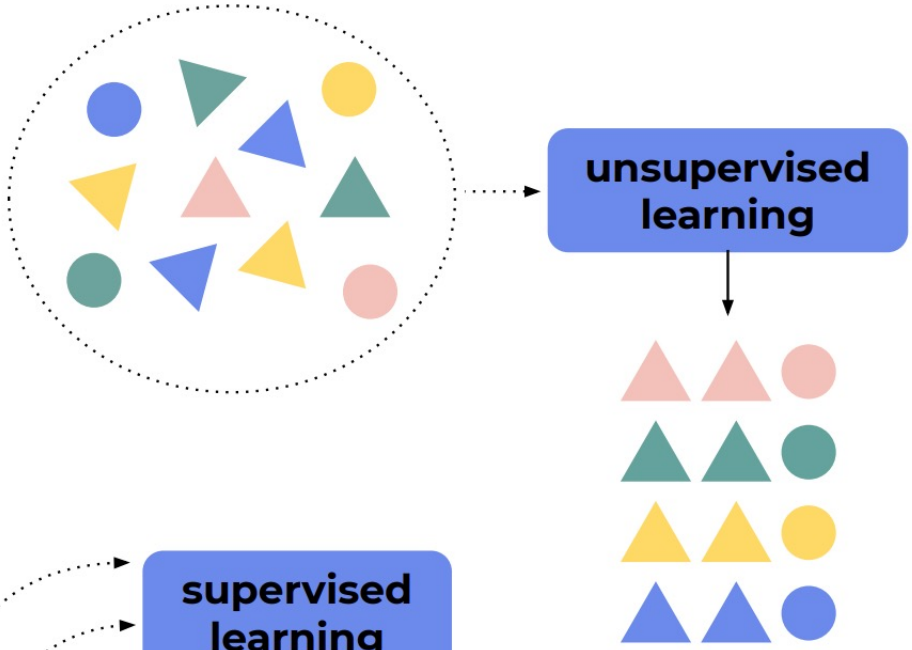

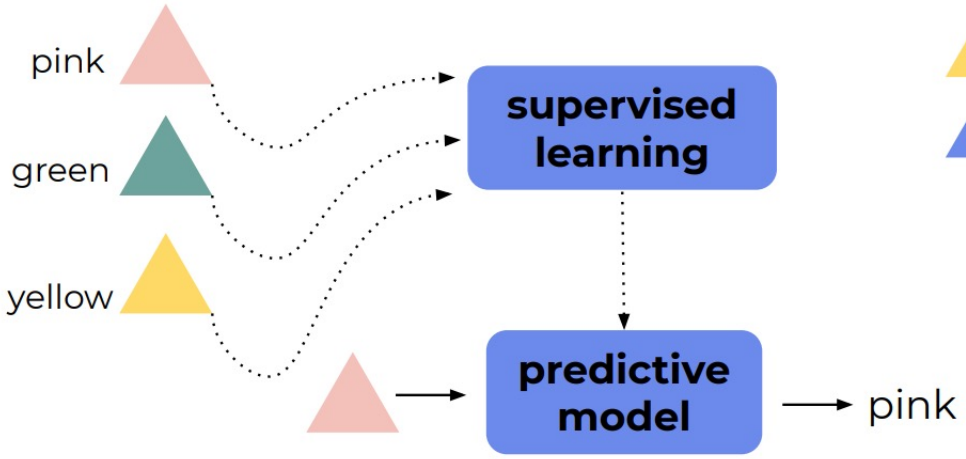

The ML process

To get acquainted with terms and understand how a model arrives at a prediction, it can be helpful to draw an analogy with a process we're familiar with: baking a cake.

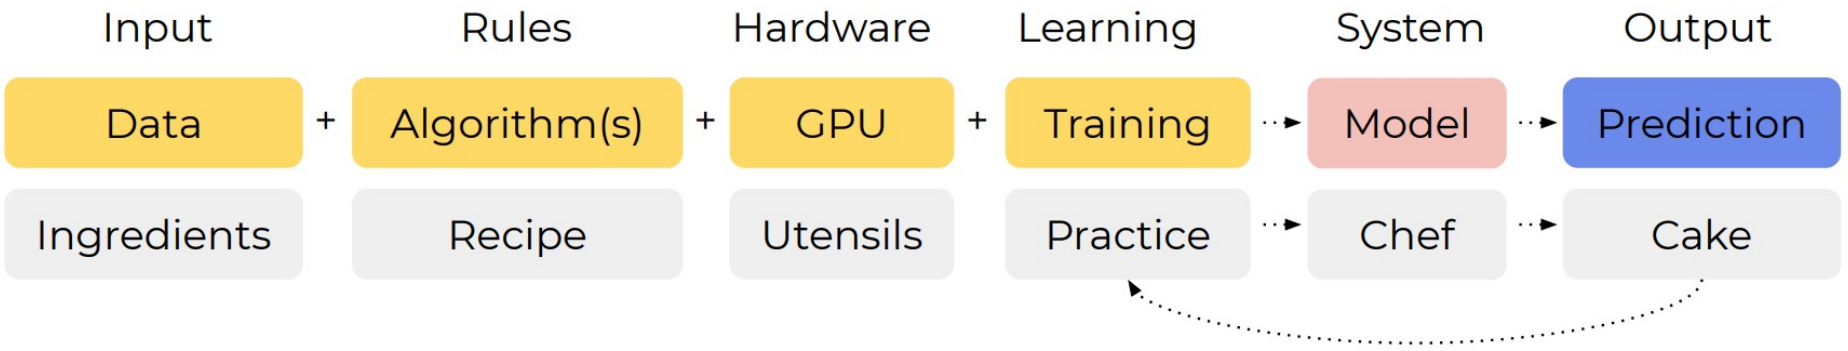

Data is the raw material you feed to the algorithm as input to produce a ML model.

An algorithm is a set of rules or step-by-step instructions to solve a problem.

The model requires GPU, and sometimes other resources, to run on.

Training process taking time and tweaking to learn, create and improve its model.

A model is a mathematical representation based on the algorithm(s) and data that is able to predict or produce an output and continues to learn over time.

Disclaimer: Please note this is a highly simplified representation of the real process which is a lot more complex and consists of plenty subtasks.

The ML process

To get acquainted with terms and understand how a model arrives at a prediction, it can be helpful to draw an analogy with a process we're familiar with: baking a cake.

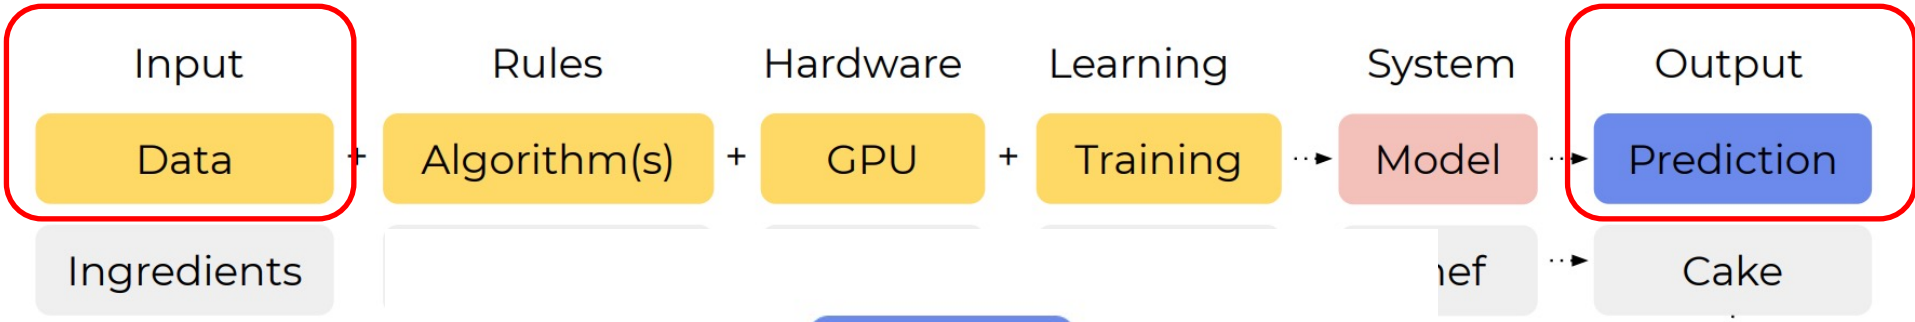

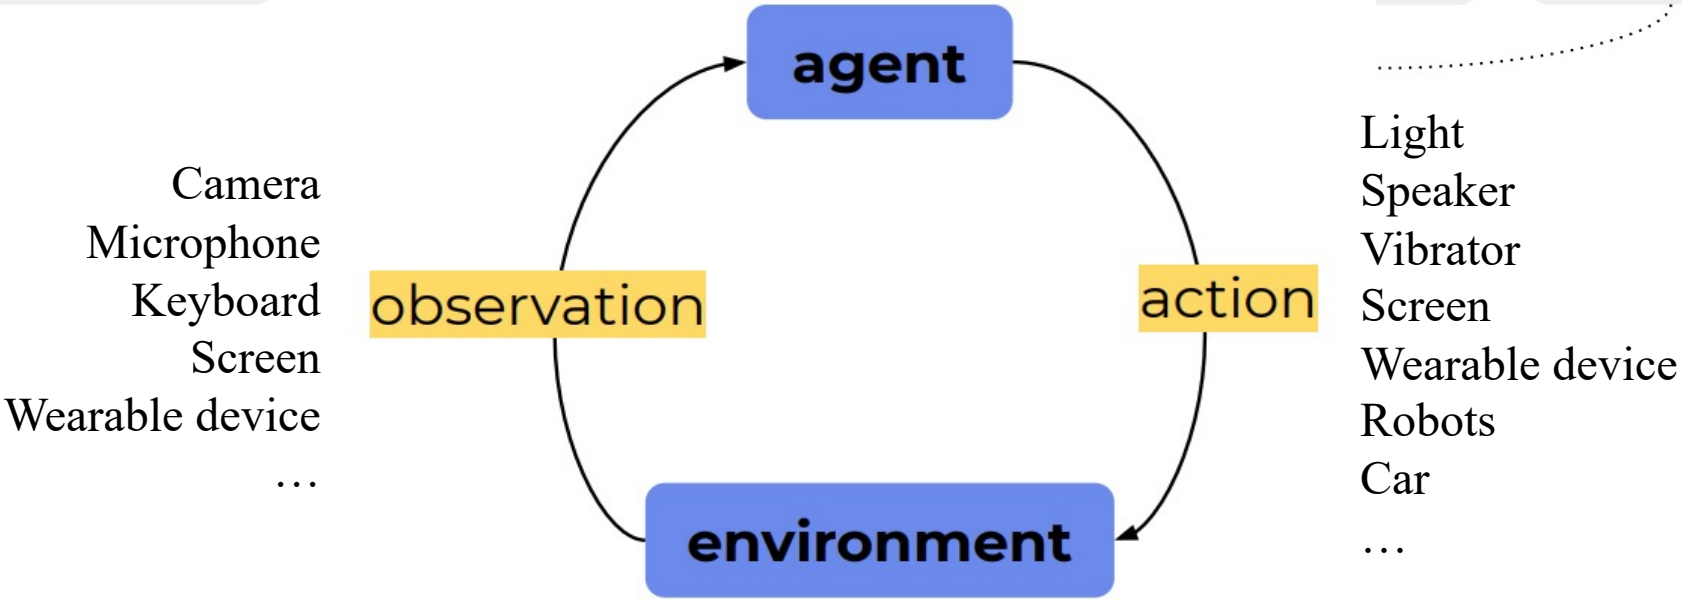

Camera
Microphone
Keyboard
Screen
Wearable device
…

Light
Speaker
Vibrator
Screen
Wearable device
Robots
Car
…

Space, Time, Human, Objects…
Virtual world
Learn from virtual world and apply in real world

Common ML tasks

- **Regression** predicts numerical values for a future instance. Used for example to estimate future housing or stock prices.
- **Clustering** seeks out (hidden) patterns in data and groups instances accordingly. Used for example to segment customers or reviews.
- **Classification** predicts to which category an individual instance belongs (in discrete values). Used for example to filter out SPAM emails or diagnose illness.

Dimension reduction helps narrow down relevant data points from big data sets. Used for feature selection and extraction.

Testing and matching compares different data sets to each other.

Association rules discovers interesting relations between variables in large data sets.

Multivariate querying aims to query and find similar instances.

Density estimation finds the likelihood and frequency of instances.

GANs generate increasingly realistic multimedia material.

Ideation

You want to explore the potential of AI but are not sure how to get started. You wonder how you might leverage it to solve problems, uncover opportunities, and create value for your users, community, and organization.

What can AI actually do? Which of its capabilities are relevant to my context? How do I spot these opportunities?

In this chapter you will find:

User-centered problem solving
to explore how AI could help solve user needs in a unique way

Tech-driven opportunity spotting
to spot opportunities for AI capabilities to create value

Data-driven opportunity spotting
to understand how to leverage private and public data

AI prompt card deck for ideation *
to prompt creative idea generation based on AI capabilities

AI –> Machine Learning–> Deep Learning

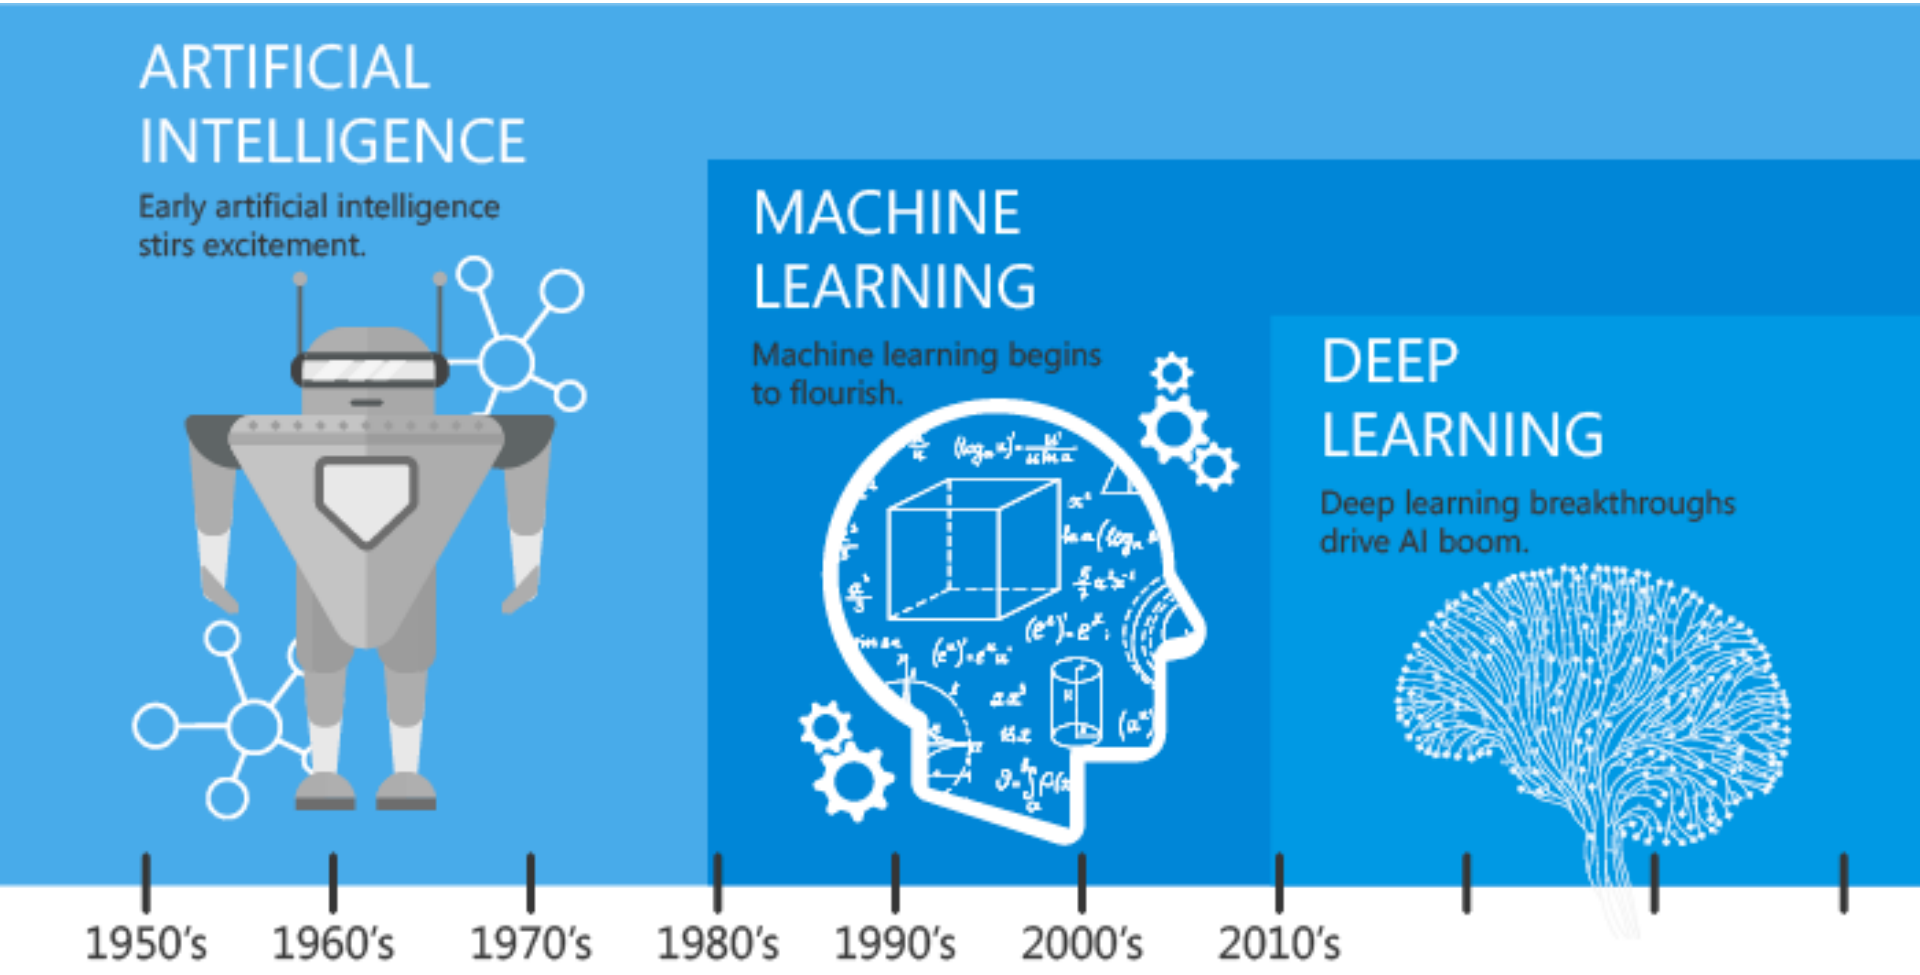

Image: Linked In | Machine Learning vs Deep learning

What is Machine Learning

- Machine Learning algorithms enable the computers to learn from data, and even improve themselves, without being explicitly programmed.

- Traditional Programming

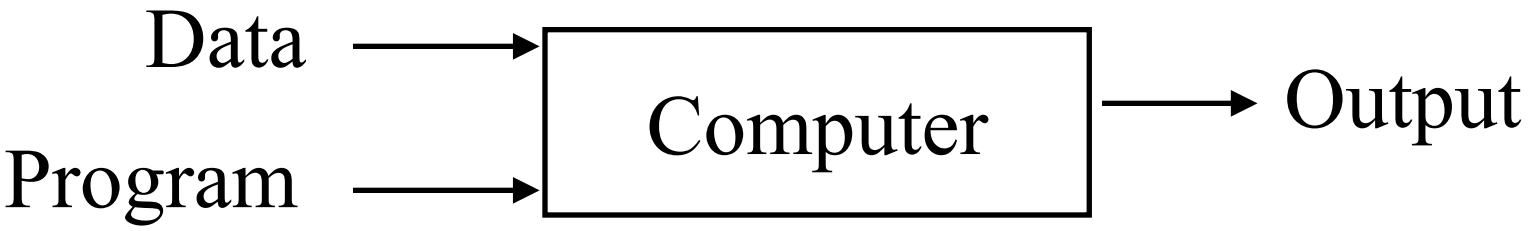

- Machine Learning

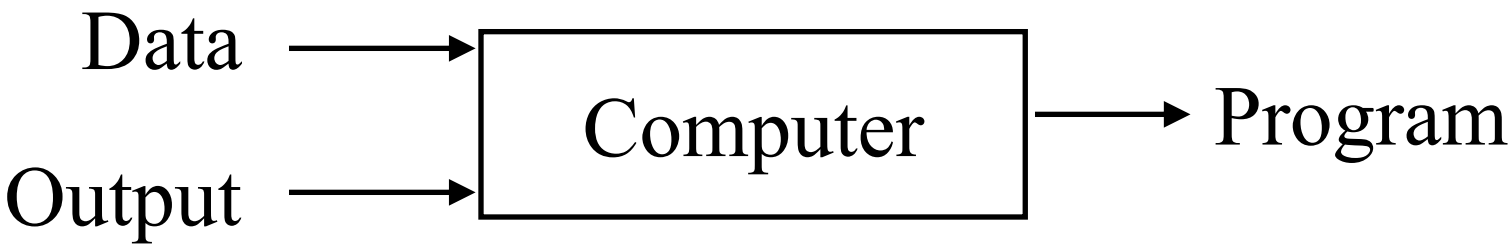

Types of Machine Learning?

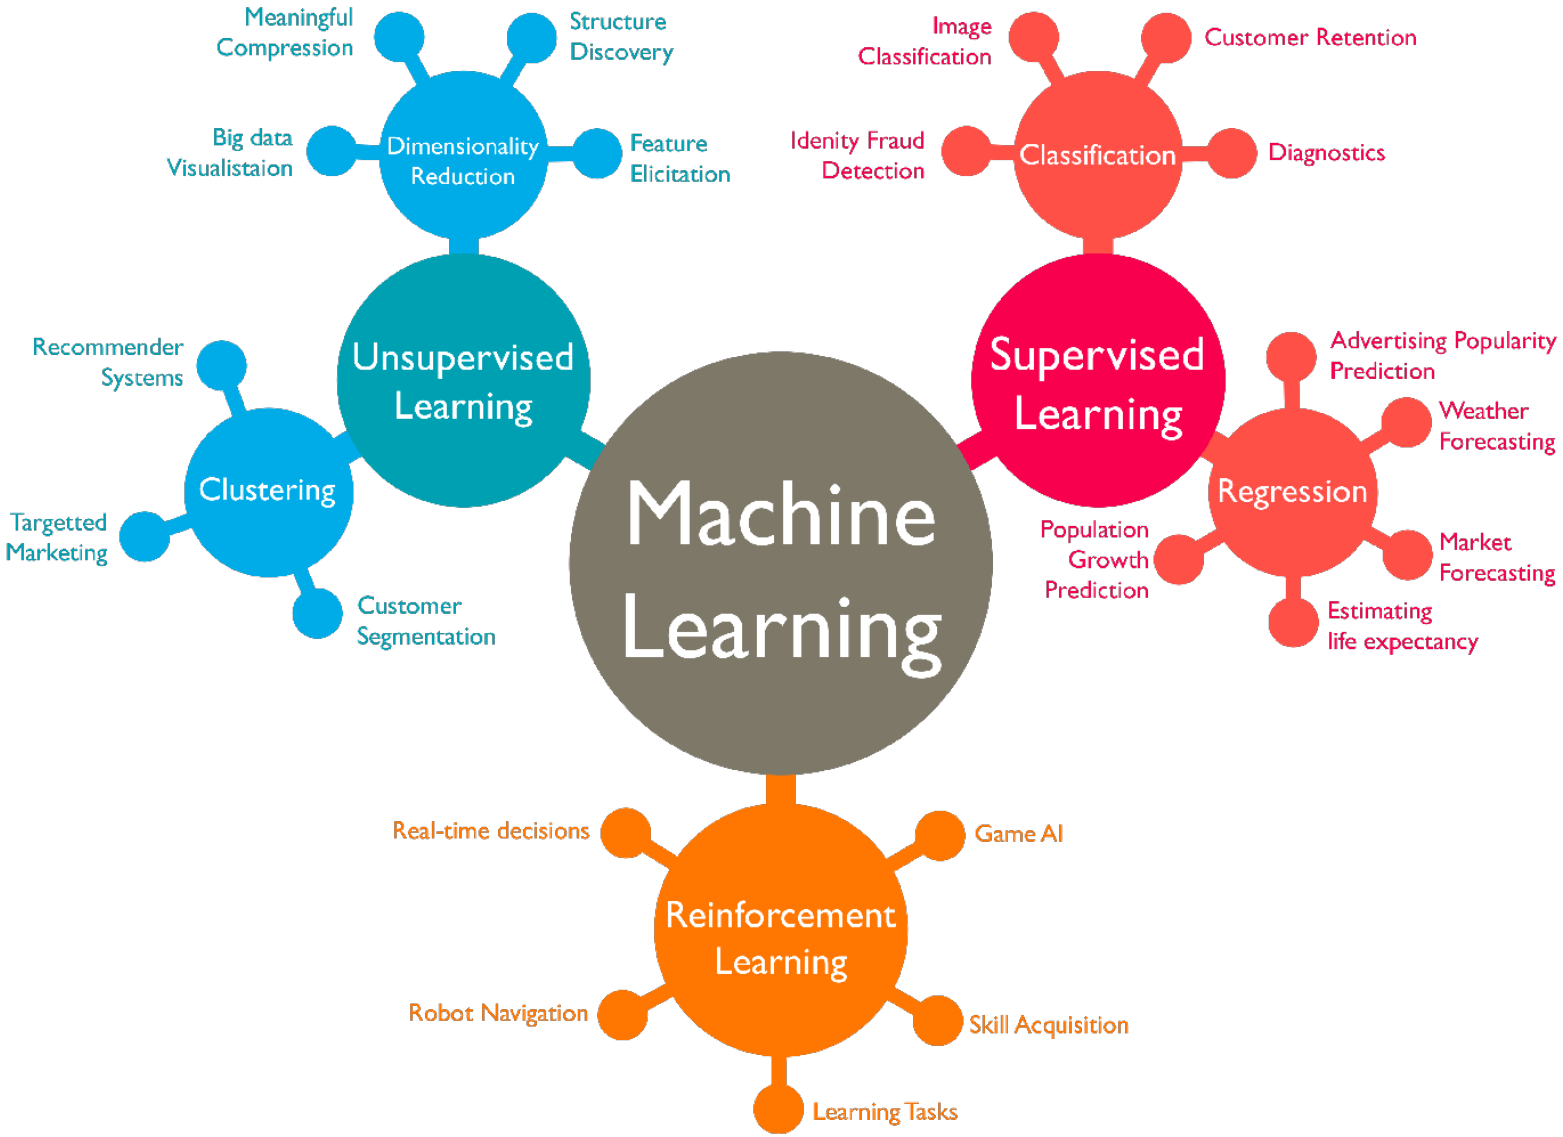

Types of Machine Learning?

	Supervised	Unsupervised
Continuous	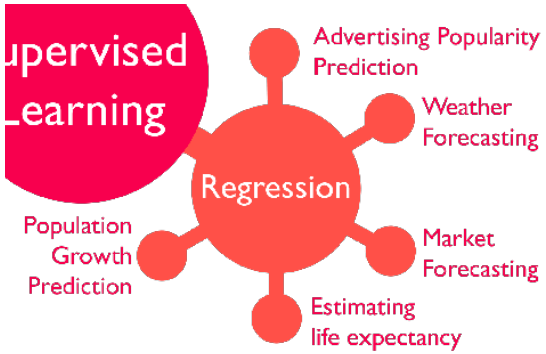	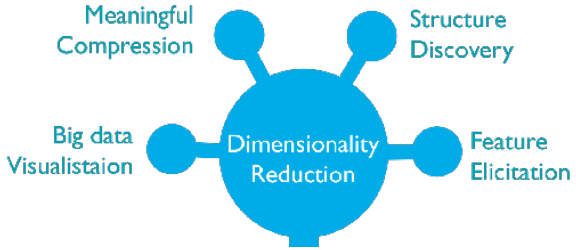
Discrete	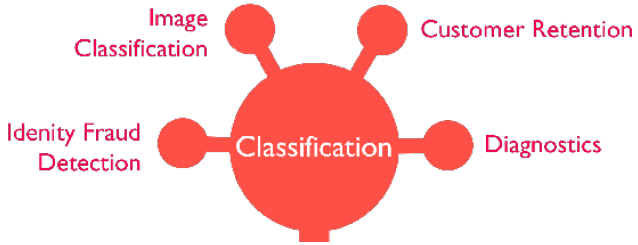	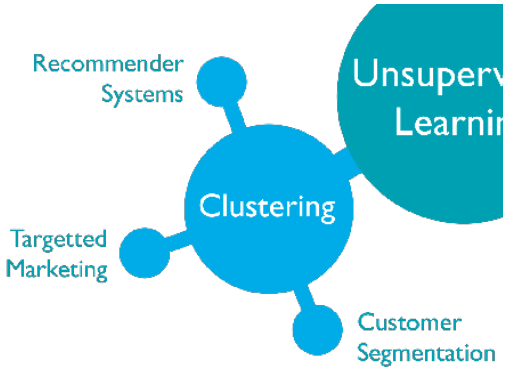

Exercise
Explore and Identify ML capacities

Wan Fang
Southern University of Science and Technology

Exercise

- Visit Google's collection of experiments with AI or other resources
- https://experiments.withgoogle.com/collection/ai

- Select 1 ~ 2 applications and share your understanding of the AI models or capacities behind.

Lecture
ML needs Data, Data needs ML

Wan Fang
Southern University of Science and Technology

What is Data?

- Data can be defined as a representation of facts, concepts, or instructions in a formalized manner,

- Suitable for communication, interpretation, or processing by human or electronic machines.

What is Data?

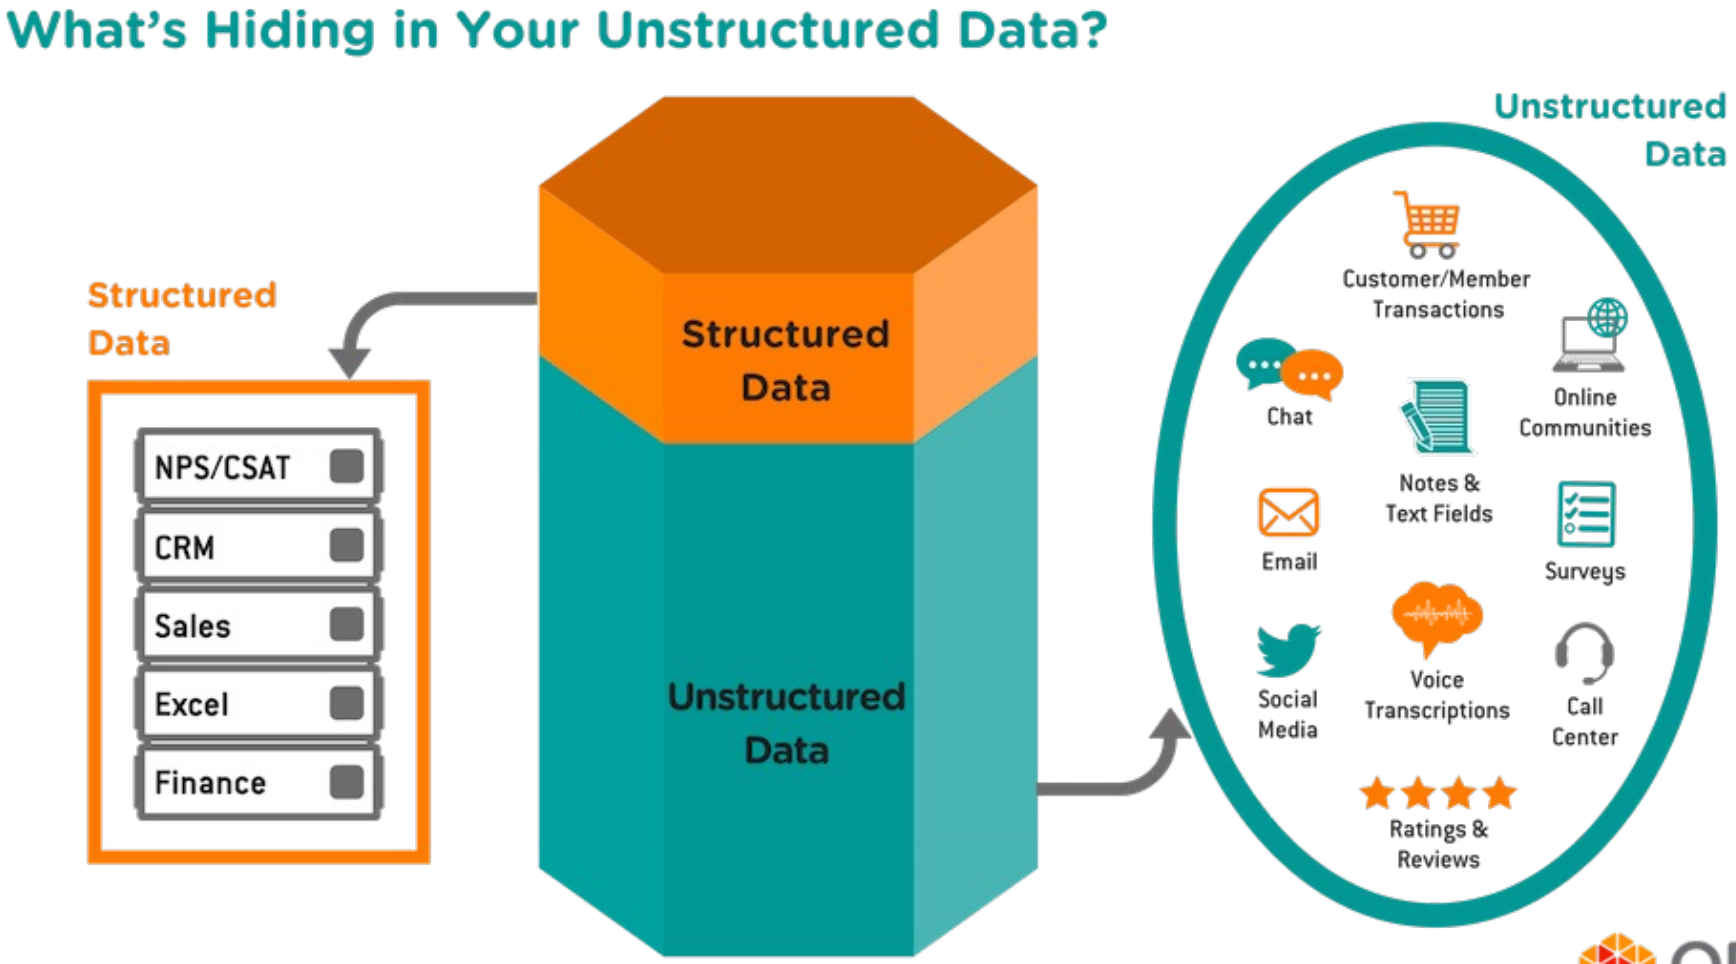

Source: Graphic adapted from January 2018 CXPA Presentation "The Why Behind the What," Jim Kitterman

What is Data?

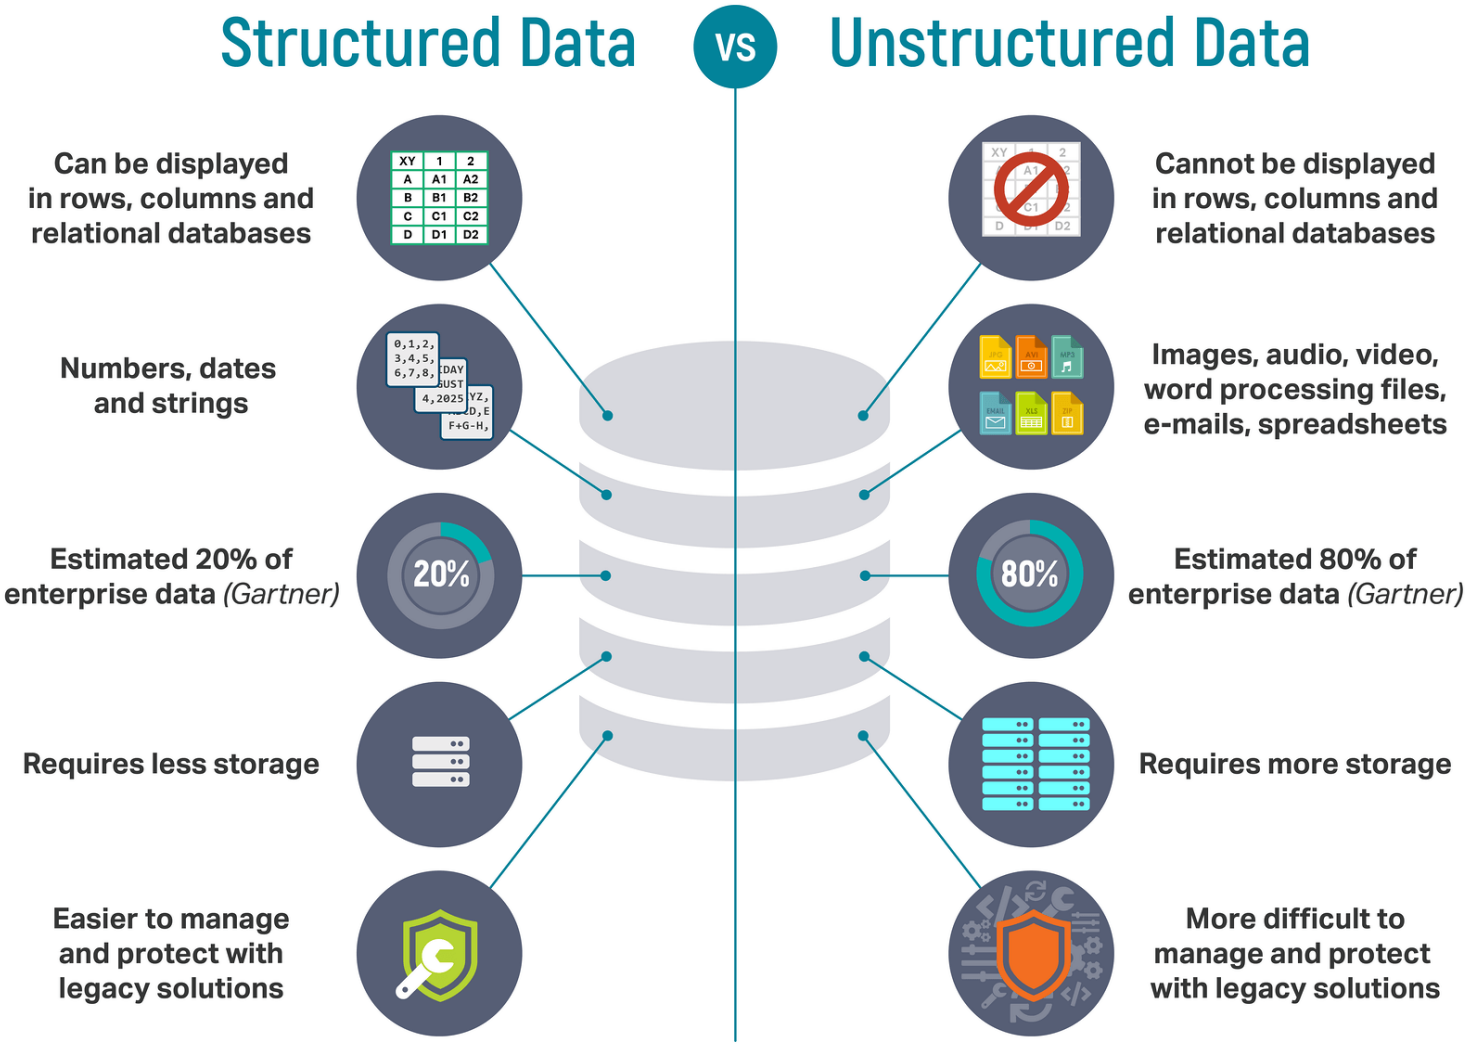

Source

The Full Human Data Stack

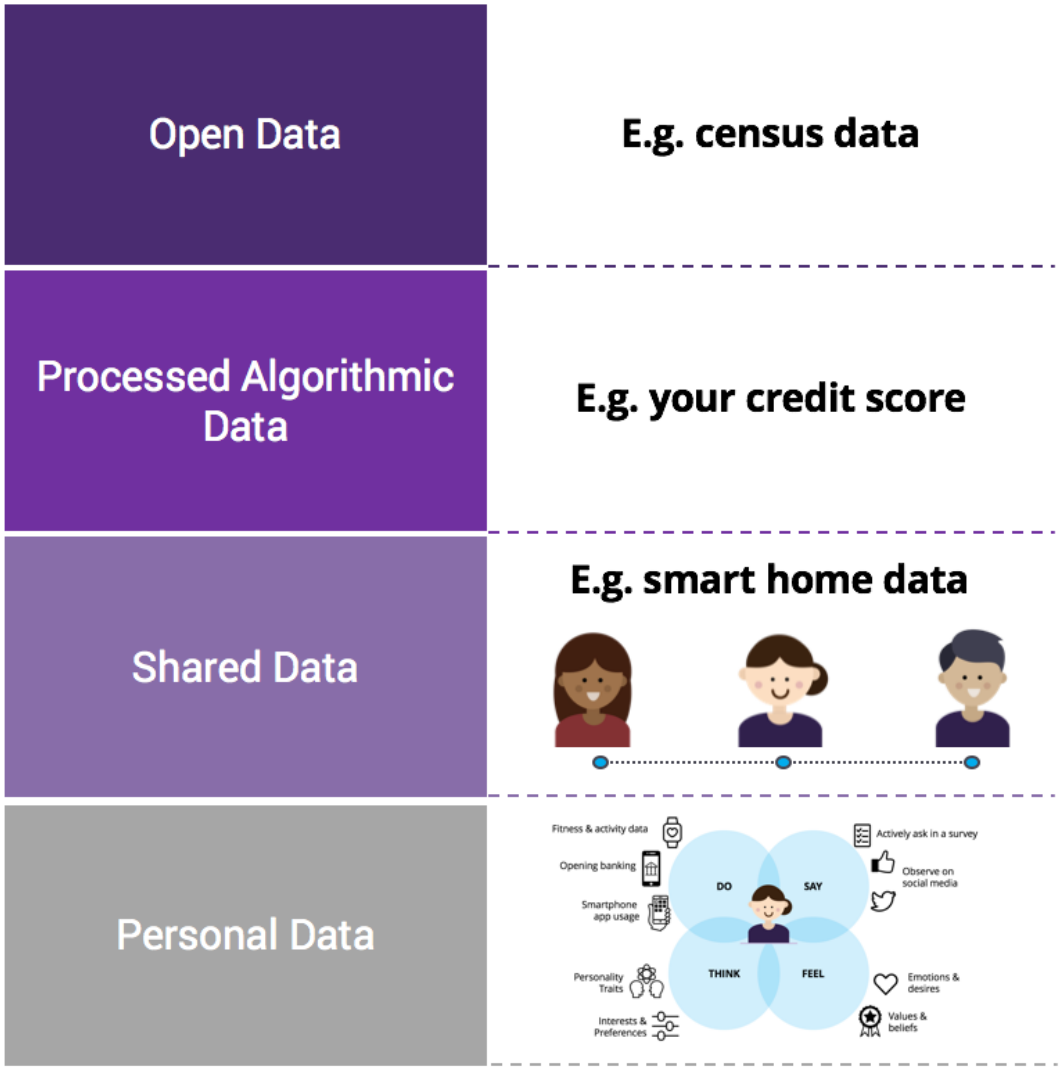

https://www.citizenme.com/human-data/

Personal Data

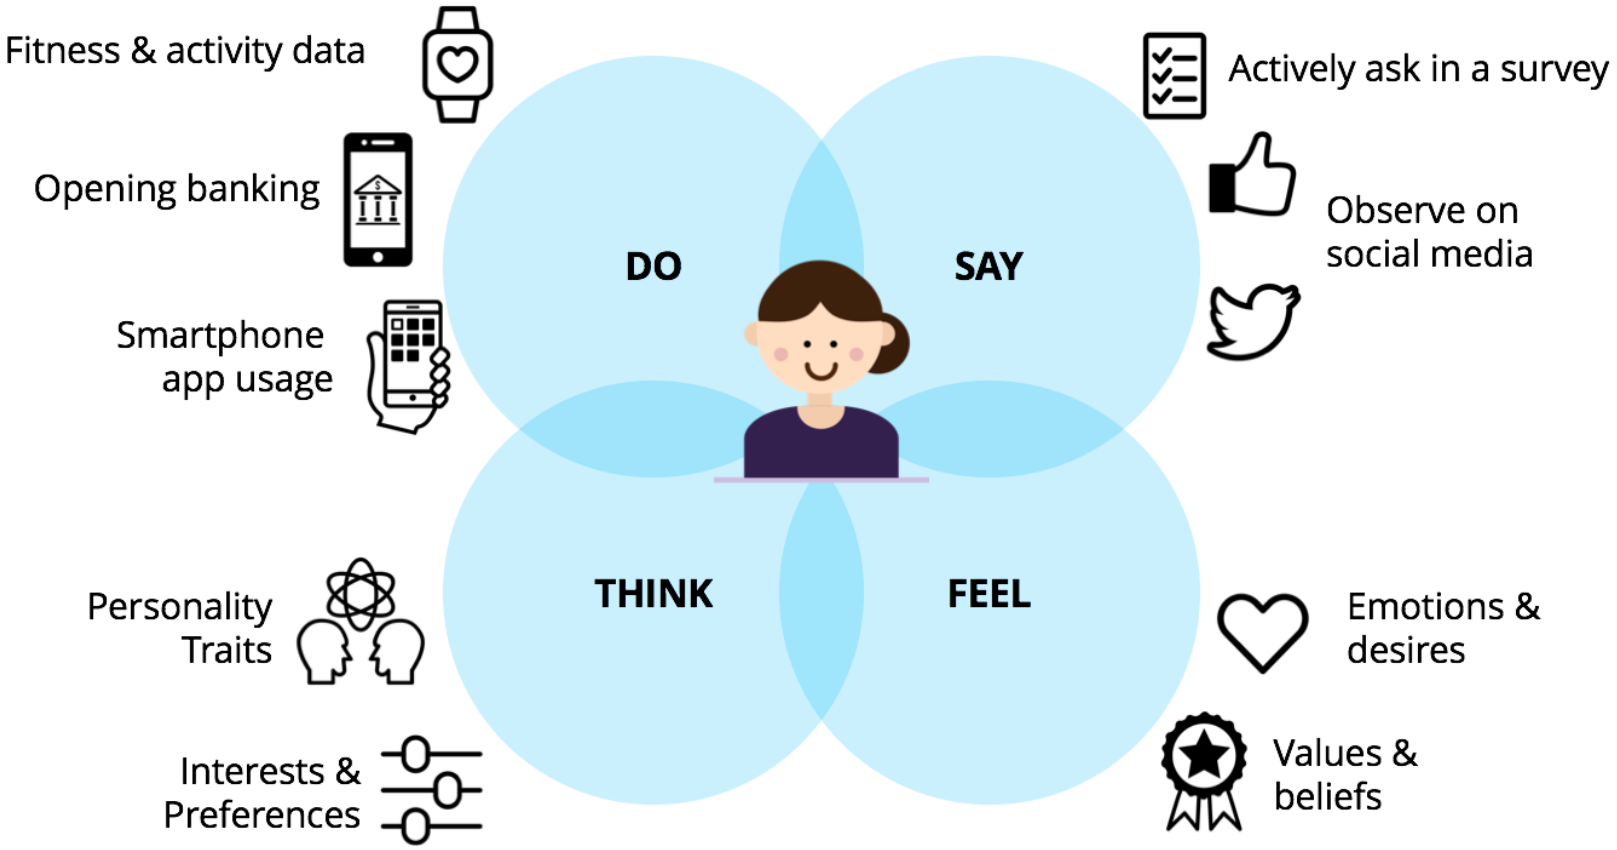

https://www.citizenme.com/human-data/

What type of data does machine learning need?

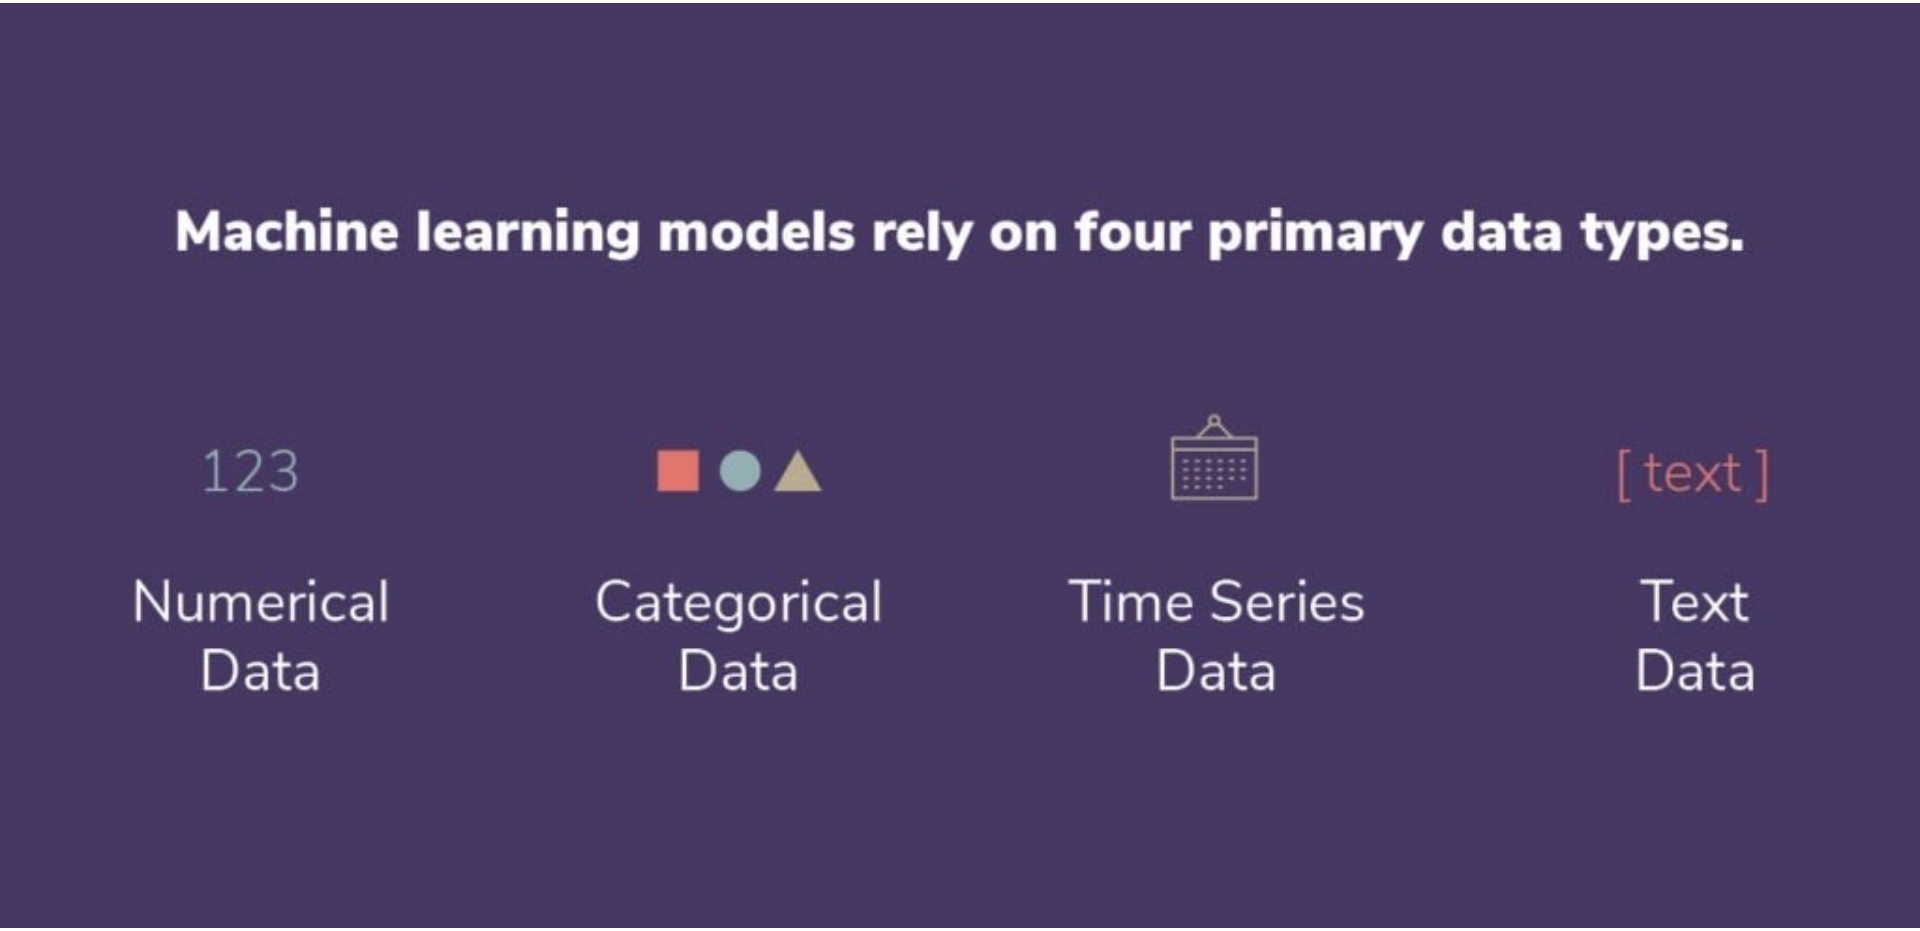

Source

Categorical Data

Color	Digitized
Red	0
Green	1
Blue	2

Hometowm	
Guangdong	0
Hunan	1
Fujian	2
…	…

Gender	
Female	0
Male	1
…	…

Numerical data

O9

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	DATA						RANK					TREND		
2	Sales Pers	May	June	July	Aug		May	June	July	Aug		June	July	Aug
3	1	84	138	72	45		35	2	42	60		↑	↓	↓
4	2	98	57	122	129		27	52	13	9		↓	↑	↑
5	3	83	108	107	107		36	19	20	20		↑	↓	→
6	4	91	135	120	56		30	3	14	54		↑	↓	↓
7	5	133	61	47	62		6	49	59	47		↓	↓	↑
8	6	73	80	86	113		41	37	33	18		↑	↑	↑
9	7	57	98	66	117		52	27	44	16		↑	↓	↑
10	8	86	52	134	132		33	56	5	8				
11	9	53	99	48	106		55	25	58	22				
12	10	96	80	59	69		29	37	50	43				
13	11	78	102	104	116		40	24	23	17				
14	12	133	119	90	89		6	15	31	32				
15	13	79	127	128	124		39	11	10	12				
16	14	49	66	64	62		57	44	46	47				
17	15	58	135	99	141		51	3	25	1				

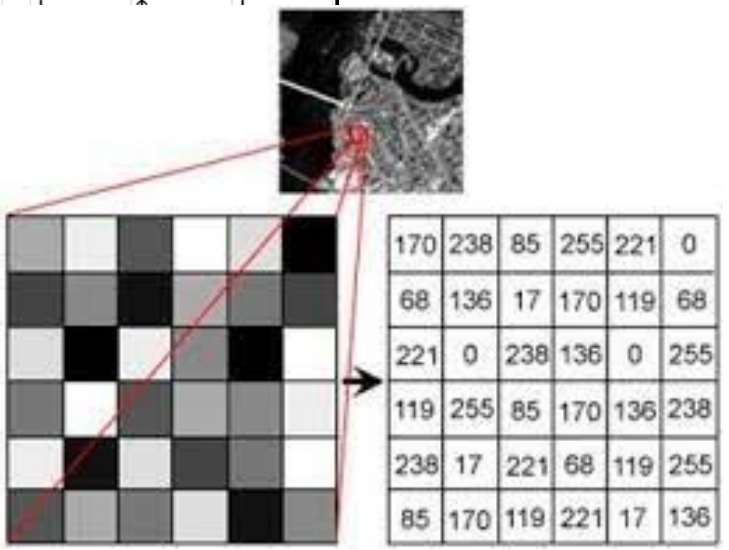

170	238	85	255	221	0
68	136	17	170	119	68
221	0	238	136	0	255
119	255	85	170	136	238
238	17	221	68	119	255
85	170	119	221	17	136

Time series data

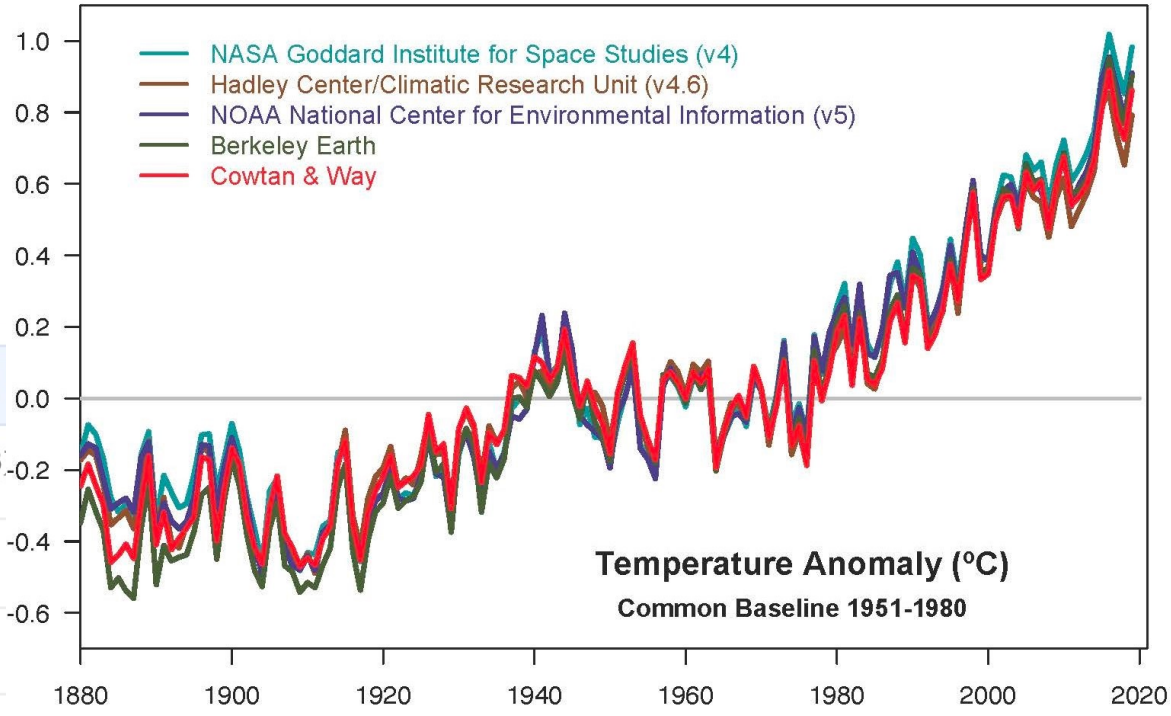

Text data

Source

Where Do We Get Data for ML?

Five of the most popular ML dataset resources:

Google → Google's Dataset Search

Microsoft → Microsoft Research Open Data

aws → Amazon Datasets

UCI → UCI Machine Learning Repository

ĐATA.GOV → Government Datasets

Why is data important for ML?

- Big data provides ample amounts of raw material from which machine learning systems can derive insights.
- To truly understand how machine learning works, you must also understand the data by which it operates.

Why is ML important to Data?

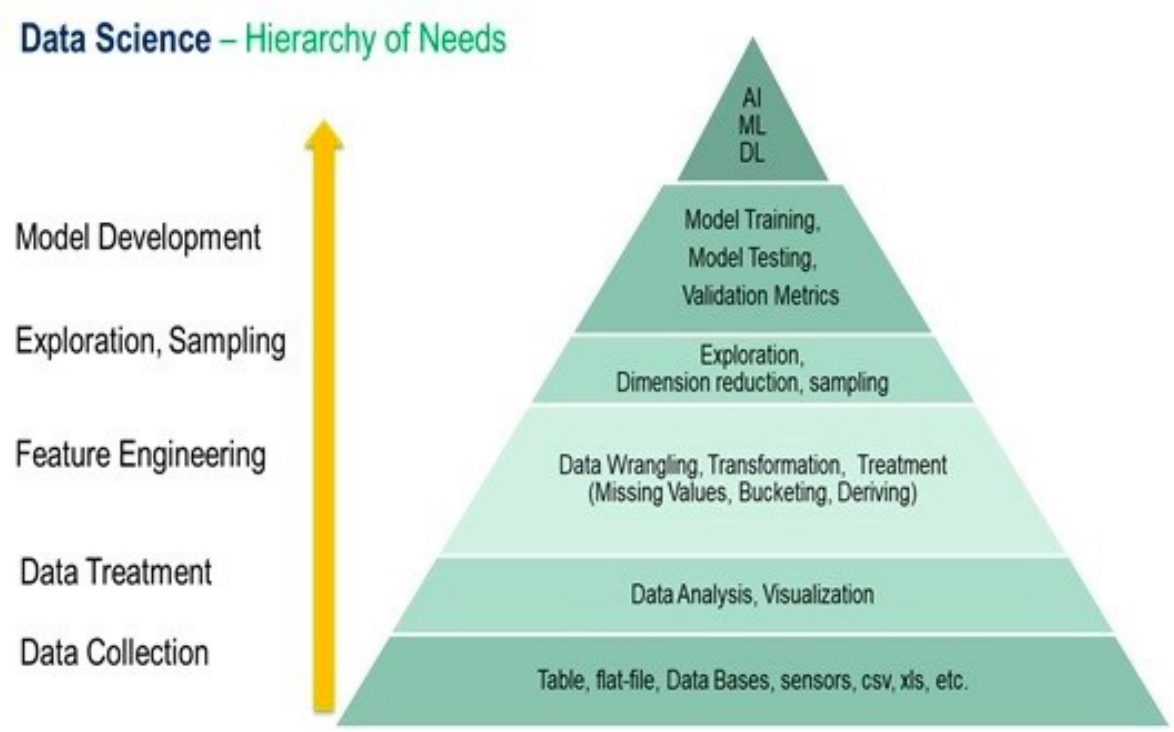

Exercise: Play, Observe, and understand Data

Wan Fang
Southern University of Science and Technology

Exercise

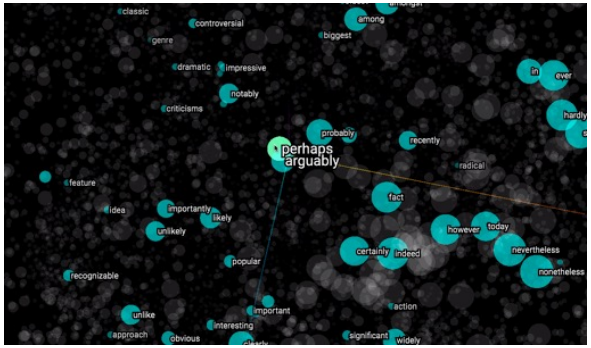

- Pluto notebook: Image Processing
- Explore 2 open dataset
  - https://www.kaggle.com/datasets
  - Visualizing High-Dimensional Space from Google
- Share with us the datasets you found: what.

Activity:
User-centered problem solving

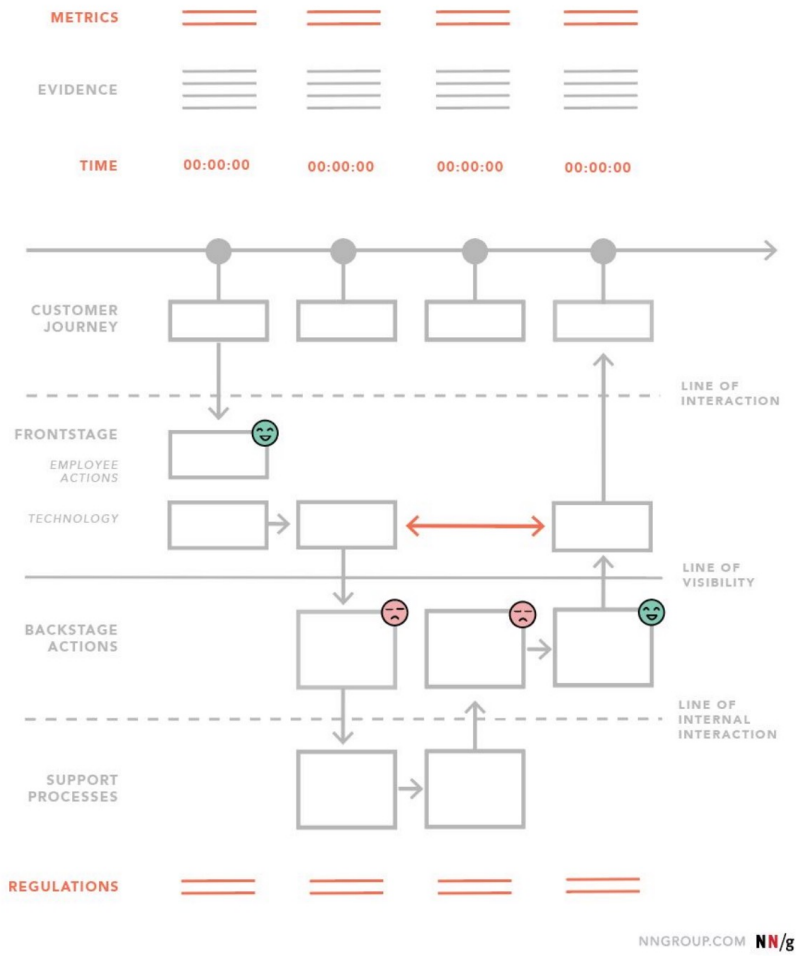

Image by Norman Nielsen Group

Before we move into more tech-driven approaches and spotting (latent) opportunities, we first align on user personas, needs, and context. In this step, we do the groundwork for everything else and state obvious and less-obvious user problems to solve.

1

Start with exploring, defining and framing user needs as usual. Enrich insights and (in)validate assumptions through qualitative user research and, where available, quantitative data.

For a refresher, check out IDEO's HMW, d.school's POV, or IBM's Hills on how to formulate a helpful challenge statement.

IDEO's HMW:
How might we .....?

d.school's POV:
[user] needs a way to [verb] because [insight]

IBM's Hills:
[who] [what] [wow]

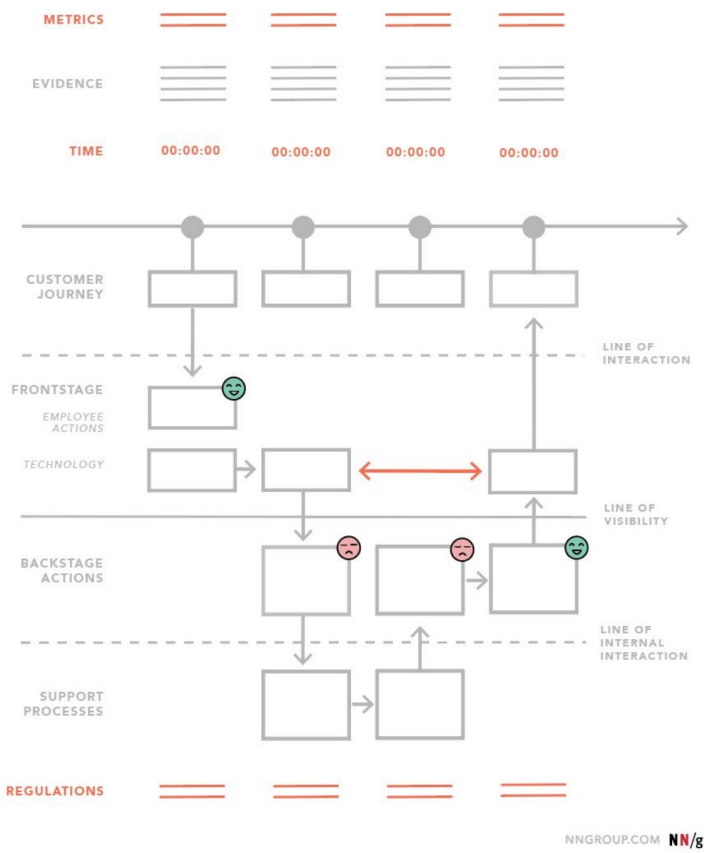

Image by Norman Nielsen Group

2

For each of the needs, challenge statements, and pain points, **ask if and how AI might help solve or fulfill x in a new / unique / better way?**

It could be in the form of a user experience, interaction, channel, process augmentation, or even data insight. It can help to refer back the common ML types described in the intro.

You can think about machine learning as **3 key capabilities**:
detection
prediction
generation

3

Map out your existing user journey or service blueprint. AI is particular useful for certain tasks, so look specifically for pain-points in the journey that:

.. are delayed or sequenced
.. are repetitive
.. are labor-intensive
.. are overwhelming
.. are emotionally sensitive
.. are context-dependent
.. are unintuitive or ill-fit to screens
.. are generic or impersonal

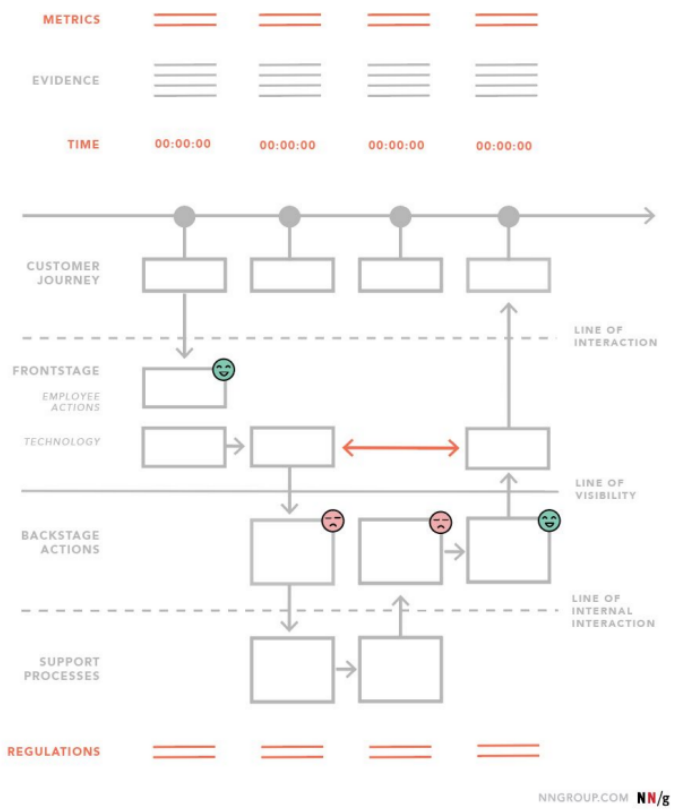

Image by Norman Nielsen Group

Activity:
Tech-driven opportunity spotting

Embracing machine learning as a design material, there is value in tech-driven innovation. Not every problem is an AI problem. Not everywhere we can use AI, we should. Understanding the type of problems that make for great machine learning candidates can help us spot new opportunities.

“AI shines in problems where the goals are understood, but the means aren’t”

- Yonatan Zunger

“AI shines in problems where the goals are understood, but the means aren’t”
- Yonatan Zunger

Activity: Tech-driven opportunity spotting

1

Capability-inspired
To find problems where AI is uniquely positioned to help, begin to look for parts of your process and interactions where:

People generally agree on what a correct answer or outcome looks like + People struggle to agree on how to arrive there, how it’s done or how it works

And there is a visible action or consequence as a result

Or: they agree but the task is repetitive or inconvenient for a human to perform

Or: they agree but writing out all the rules or processing all the data is unfeasible

Keep in mind the key capabilities of machine learning (detection, prediction, generation) and its key functions (clustering, regression, classification).

Refer back to the **crash course in AI/ML** in the beginning of this booklet for a refresher.

2

Research-inspired
Most machine learning progress comes from the academic space. Keep an eye out for new capabilities and models, then see if you can think of valuable applications and use cases for them.

Look at paperswithcode.com /sota for **recent research** developments.

3

Industry-inspired
You might’ve seen an AI feature or API offering somewhere else and see value in adopting a similar solution into your product or service.

For **industry inspiration** on what’s already happening, use the AI x Design Prompt Cards attached to this toolkit.

Embracing machine learning as a design material, there is value in tech-driven innovation. Not every problem is an AI problem. Not everywhere we can use AI, we should. Understanding the type of problems that make for great machine learning candidates can help us spot new opportunities.

Activity:
Tech-driven opportunity spotting

1

Capability-inspired

To find problems where AI is uniquely positioned to help, begin to look for parts of your process and interactions where:

People generally agree on what a correct answer or outcome looks like

+

People struggle to agree on how to arrive there, how it's done or how it works

And there is a visible action or consequence as a result

Or: they agree but the task is repetitive or inconvenient for a human to perform

Or: they agree but writing out all the rules or processing all the data is unfeasible

Keep in mind the key capabilities of machine learning (detection, prediction, generation) and its key functions (clustering, regression, classification).

Refer back to the **crash course in AI/ML** in the beginning of this booklet for a refresher.

Activity:
Tech-driven opportunity spotting

2

Research-inspired

Most machine learning progress comes from the academic space. Keep an eye out for new capabilities and models, then see if you can think of valuable applications and use cases for them.

Look at paperswithcode.com /sota for **recent research** developments.

3

Industry-inspired

You might've seen an AI feature or API offering somewhere else and see value in adopting a similar solution into your product or service.

For **industry inspiration** on what's already happening, use the AI x Design Prompt Cards attached to this toolkit.

Activity: Data-driven opportunity spotting

Like cars run on fuel, machine learning runs on data. Looking at the data that is available to us, and how it might help our users, can help us spot new opportunities.

1

Data availability
What data do you have? What data do your direct partners or collaborators own? Which other data sources do you have access to?

If you own data, count yourself lucky. If they are large sets of (semi-)organized data, count yourself in.

Which insights might you be able to draw from this data? Who and how could this help?

If you don't have any data available, let's get creative with your data acquisition.

2

Creative data acquisition
Look for public and relevant data sets in your industry or region. Ask which and how they might help serve your users?

Imagine if you can scrape the data required from the internet, Wiki, news sites, and social media?

Think what data you could collect or label through user interactions? When CAPTCHA asks people to select the images with traffic lights, you are actually labelling objects in their data set of images.

3

Buying data
Although often costly, data is for sale. You can probably buy the data you need either as a product, selling one-off data sets, or as a service, with subscription-based models.

If your data is unlabelled or unstructured, resort to **data preparation** services on Amazon Mechanical Turk.

Places to look for **public data sets** include:
Kaggle
Socrata OpenData
data.gov (US)
UCI Repository
Academic Torrents

If you want to become more **data-driven** as an organization and facilitate a team session around this, check out the free resources from *data.world*.

Tool:
AI prompt card deck for ideation

Attached to this toolkit you'll find a card deck. The card deck includes over 20 what-if statements to help idea generation based on machine learning capabilities that are possible and feasible today.

24 what-if prompt cards are organized into 6 categories - each symbolizing a new area of opportunity for user experiences.

You can use the cards for ideation and brainstorming sessions, as elements for a force fitting exercise, for communication between design and engineering, to learn, to spark critical discussions around technology, and more.

Each card includes:

Tool:
AI prompt card deck for ideation

Attached to this toolkit you'll find a card deck. The card deck includes over 20 what-if statements to help idea generation based on machine learning capabilities that are possible and feasible today.

24 what-if prompt cards are organized into 6 categories - each symbolizing a new area of opportunity for user experiences.

You can use the cards for ideation and brainstorming sessions, as elements for a force fitting exercise, for communication between design and engineering, to learn, to spark critical discussions around technology, and more.

Each card includes:

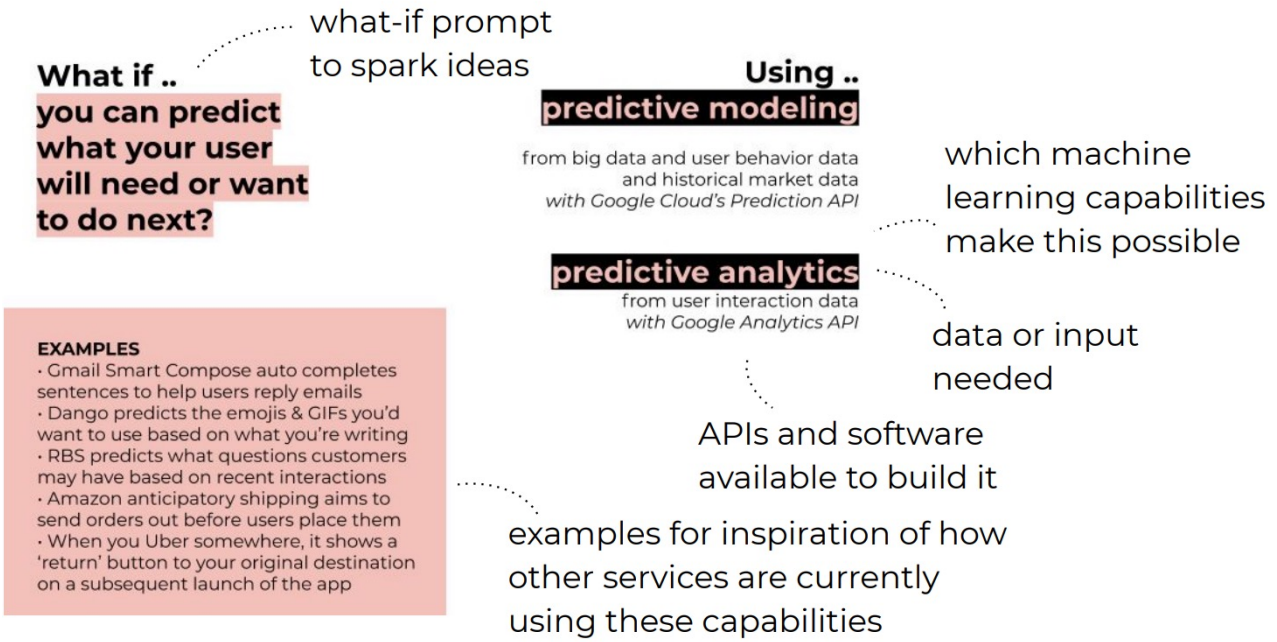

Concept development + idea selection

You've generated a bunch of ideas and now you're wondering which to move forward with. You need to roughly assess feasibility, viability, and desirability of your concepts and their AI elements to inform your selection.

How feasible, viable, and desirable are the ideas? How to select which to move forward with? How to develop concepts in more detail?

In this chapter you will find:

Impact matrix for idea selection
to map ideas according to desirability and feasibility

Value proposition design *
to better understand what value you're offering your user

Assessing feasibility
to roughly assess feasibility of your AI ideas without expert knowledge

Framing your task + Plotting your model *
to begin thinking about your task in a computational way

Activity:
Impact matrix for idea selection

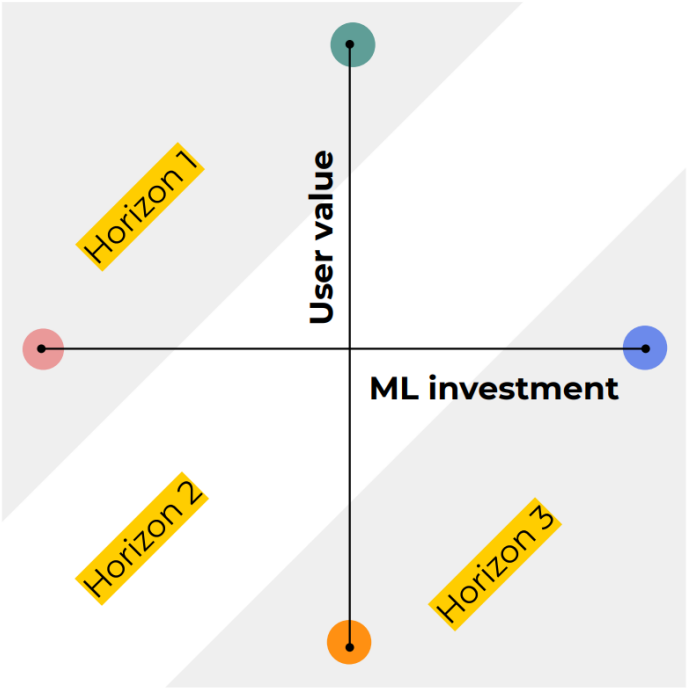

1

Plot your generated ideas on an impact matrix. The x-axis represents the estimated user value (based on desirability + responsibility), the y-axis the ML investment required (based on feasibility + viability). Use sticky notes so you can shift them as you learn.

2

The following tools and questions will help you assess desirability, responsibility, feasibility, and viability to fill out the matrix:

Desirability
Use the assessing desirability worksheet, value proposition design, and user testing practices attached to this toolkit to help determine user value.

Responsibility
Consider unintended consequences, edge cases, misuse. What ethical concerns arise at first glance?

Feasibility
Use the assessing feasibility flowchart attached to this toolkit to help determine feasibility.

Viability
Is data available? Which building blocks are available? Is it a custom or out-of-the-box model? What is the time-to-model? How does it return value on investment?

3

Evaluate your ideas and decide which to move forward with.

Horizon 1	Horizon 2	Horizon 3
= high value	= medium / unknown value	= low value
+ low effort	+ medium / unknown effort	+ high effort
Yas! Get to work on these ideas	Google away and talk to your team to evaluate in more detail	It's not even worth it

To ensure you're designing for human values and well-being, use the Design for Happiness Deck from the Delft Institute of Positive Design.

1

Draft your value proposition statement with the madlib on the right and iterate it as you learn more about your user and solution.

Data - Which input will you use to inform the model?

AI capability - What are you looking to do to turn data input into a valuable output?

Persona - Who is your user?

Job to be done - What user need does it solve or fulfill?

Gain - What does your user gain in using this solution compared to what they're currently doing?

Value - Why is this important to the user or humanity at large?

_____
[concept name]

using .....
[data]

to .....
[AI capability]

we can help .....
[persona]

with a better way to .....
[job to be done]

with/without .....
[gain/pain]

because/so that .....
[value]

2

Now go out and (in)validate your value proposition with user research.

In the next chapter you'll find best practices and things to consider for feedback, prototyping, and testing AI concepts with users.

3

Iterate and refine the statement, or even pivot or discard your idea, based on your learnings.

Worksheet: Value proposition design for assessing desirability

Before we continue, it's time to check in with our users and cross-check if we're solving a real need, if we're solving it in a unique and helpful way.

Worksheet:
Value proposition statement

[concept name]

using .....

[data]

to .....

[AI capability]

we can help .....

[user persona]

with a better way to .....

[job to be done]

with/without .....

[gain/pain]

because/so that .....

[value]

Draft your value proposition statement with the madlib on the right and iterate it as you learn more about your user and solution.

Data
Which input will you use to inform the model?

AI capability
What are you looking to do to turn data input into a valuable output?

Persona
Who is your user?

Job to be done
What user need does it solve or fulfill?

Gain/pain
What does your user gain in using this solution compared to what they're currently doing?

Value
Why is this important to the user or humanity at large?

Now go out and (in)validate your value proposition with user research.

Iterate and refine the statement, or even pivot or discard your idea, based on your learnings.

Day 02
AI Meets Design I

Thank you~

Wan Fang
Southern University of Science and Technology