Lessons Learned from Building Products with ML

Dr. Mikio L. Braun, YOW! Data 2021, May 12, 2021

1. Some cool ML algorithm
2. ???
3. Profit!!

Lessons learned - from what?

- 15 years ML research (4 Ph.D., 11 PostDoc)
- 5 years “startup-on-the-side”
- 5 years industry (e-commerce, Zalando, GetYourGuide)
- 6 months consulting
- You might’ve seen jblas.org (531 stars on github)

Understand how ML fits into products

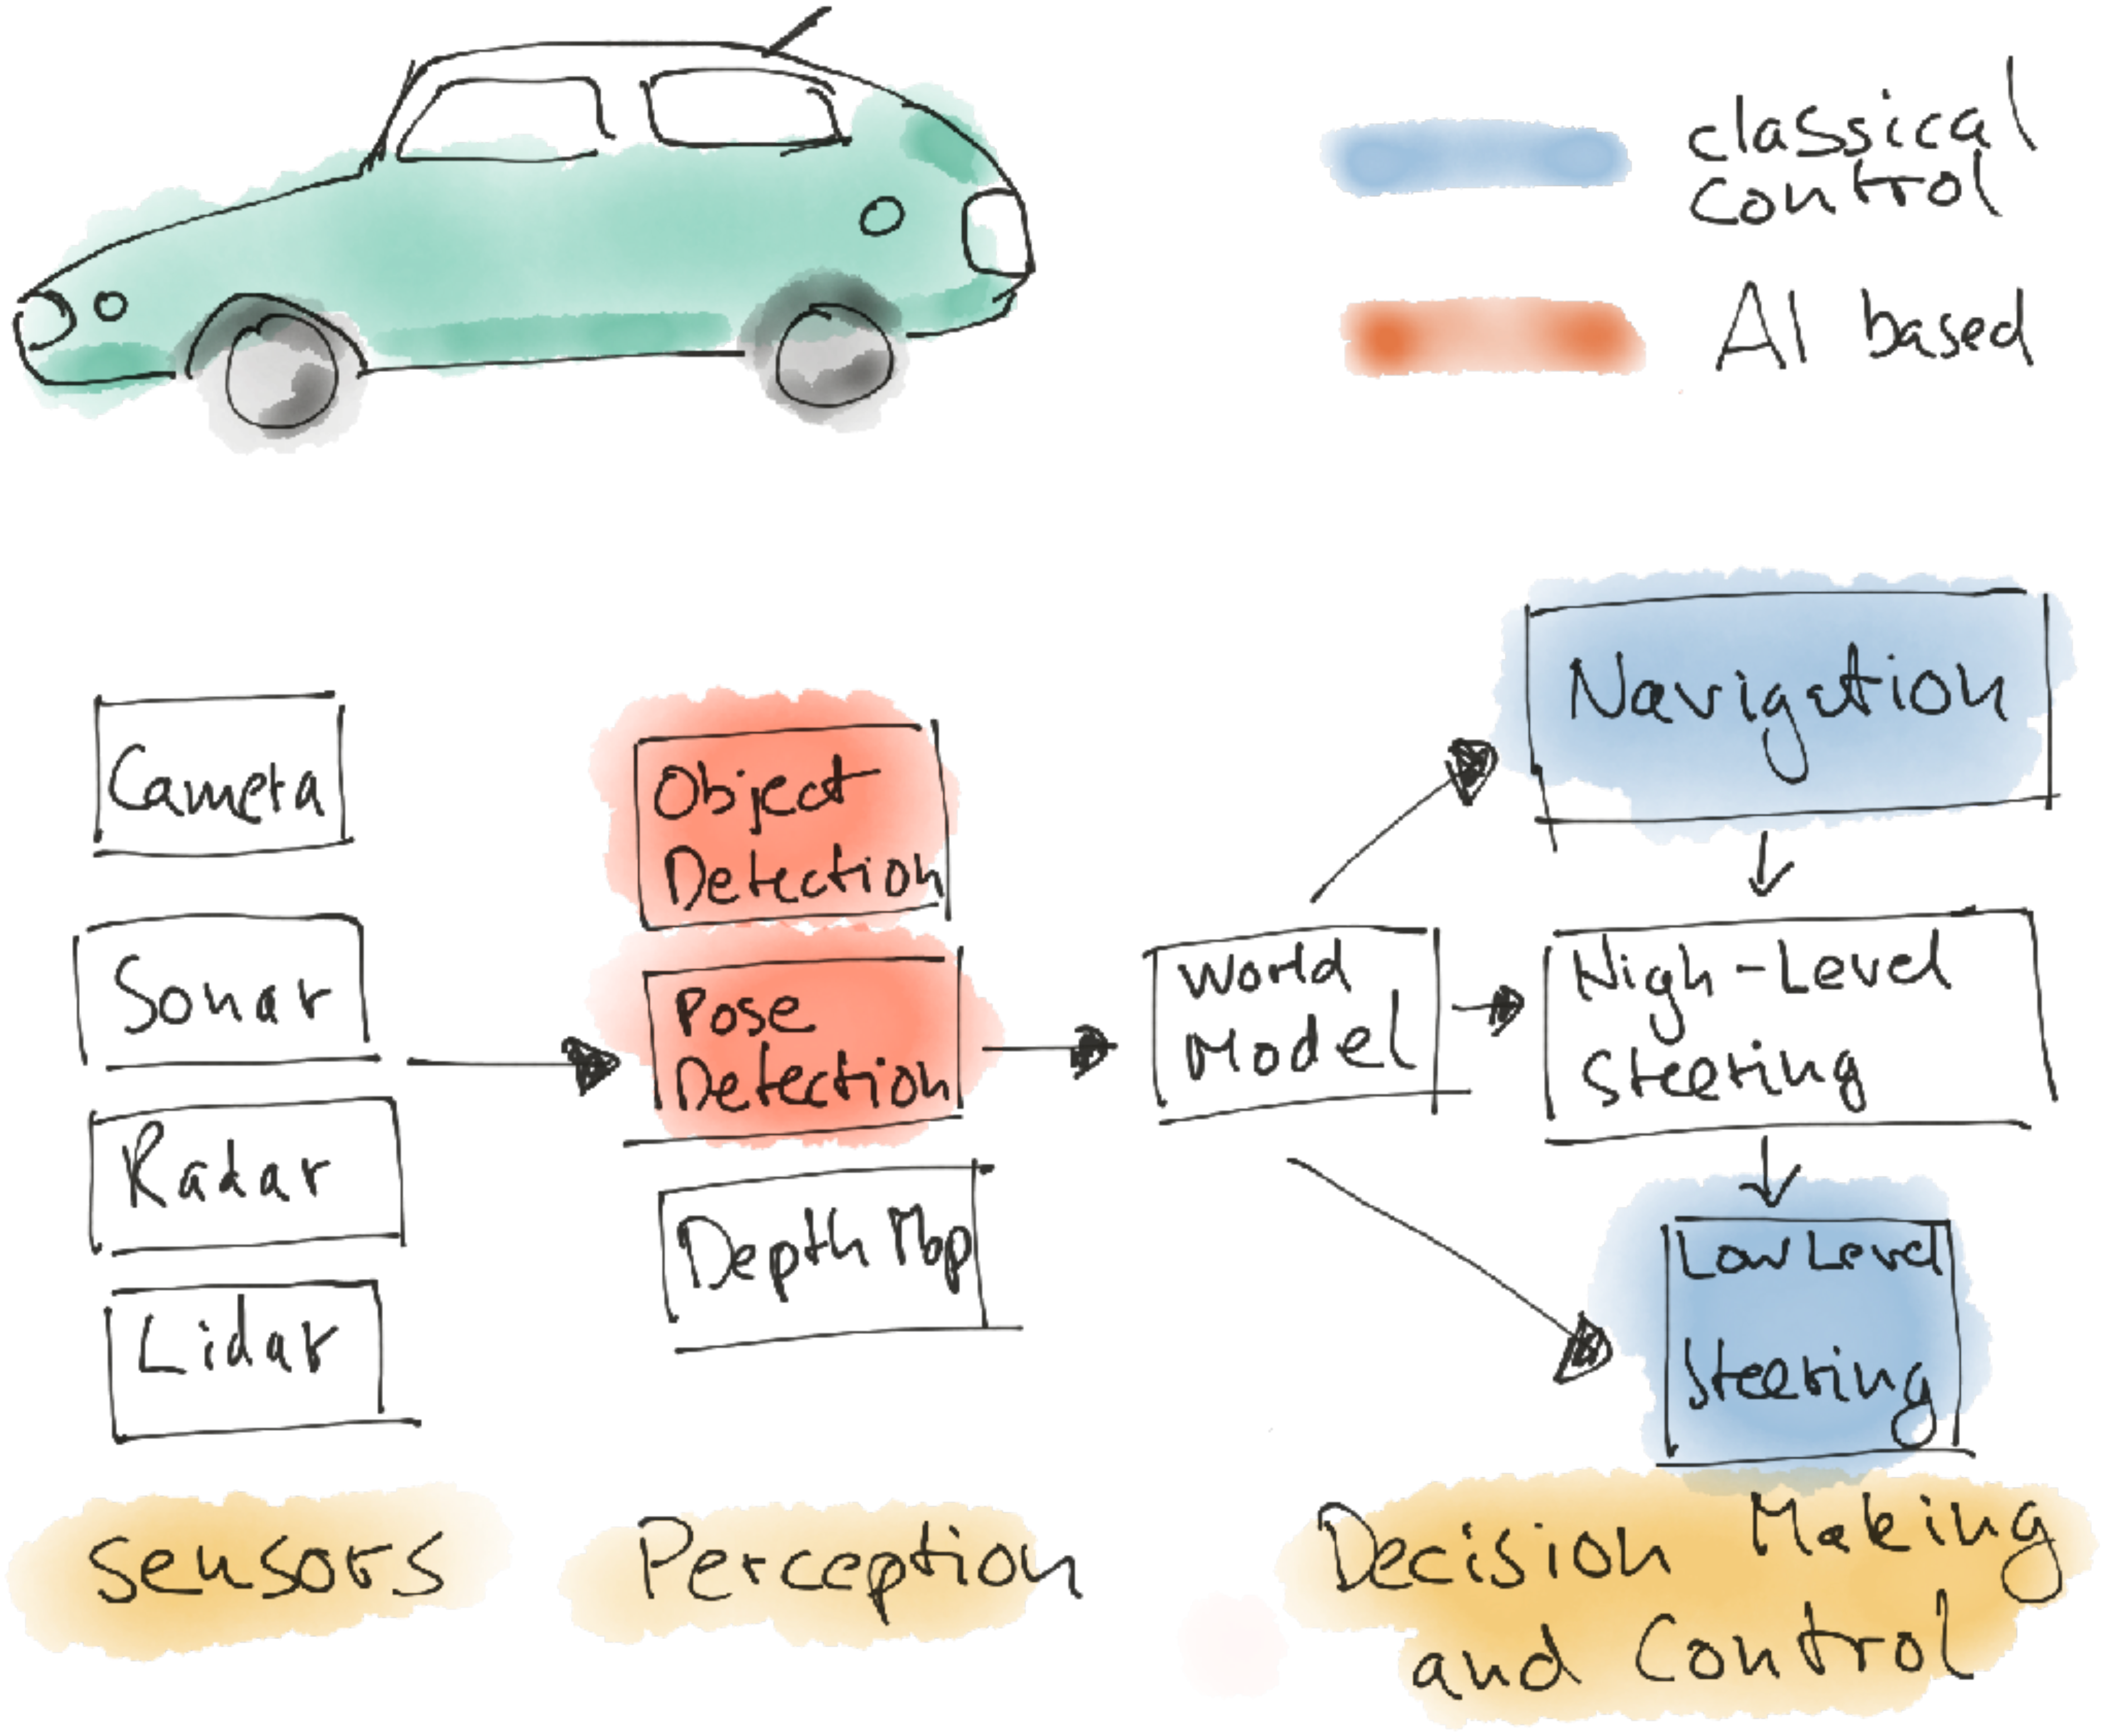

classical control
AI based
Camera
Sonar
Radar
Lidar
Object Detection
Pose Detection
Depth Map
World Model
Navigation
High-Level Steering
Low Level Steering
Sensors
Perception
Decision Making and Control

Data-Driven Approach To Autonomous Driving

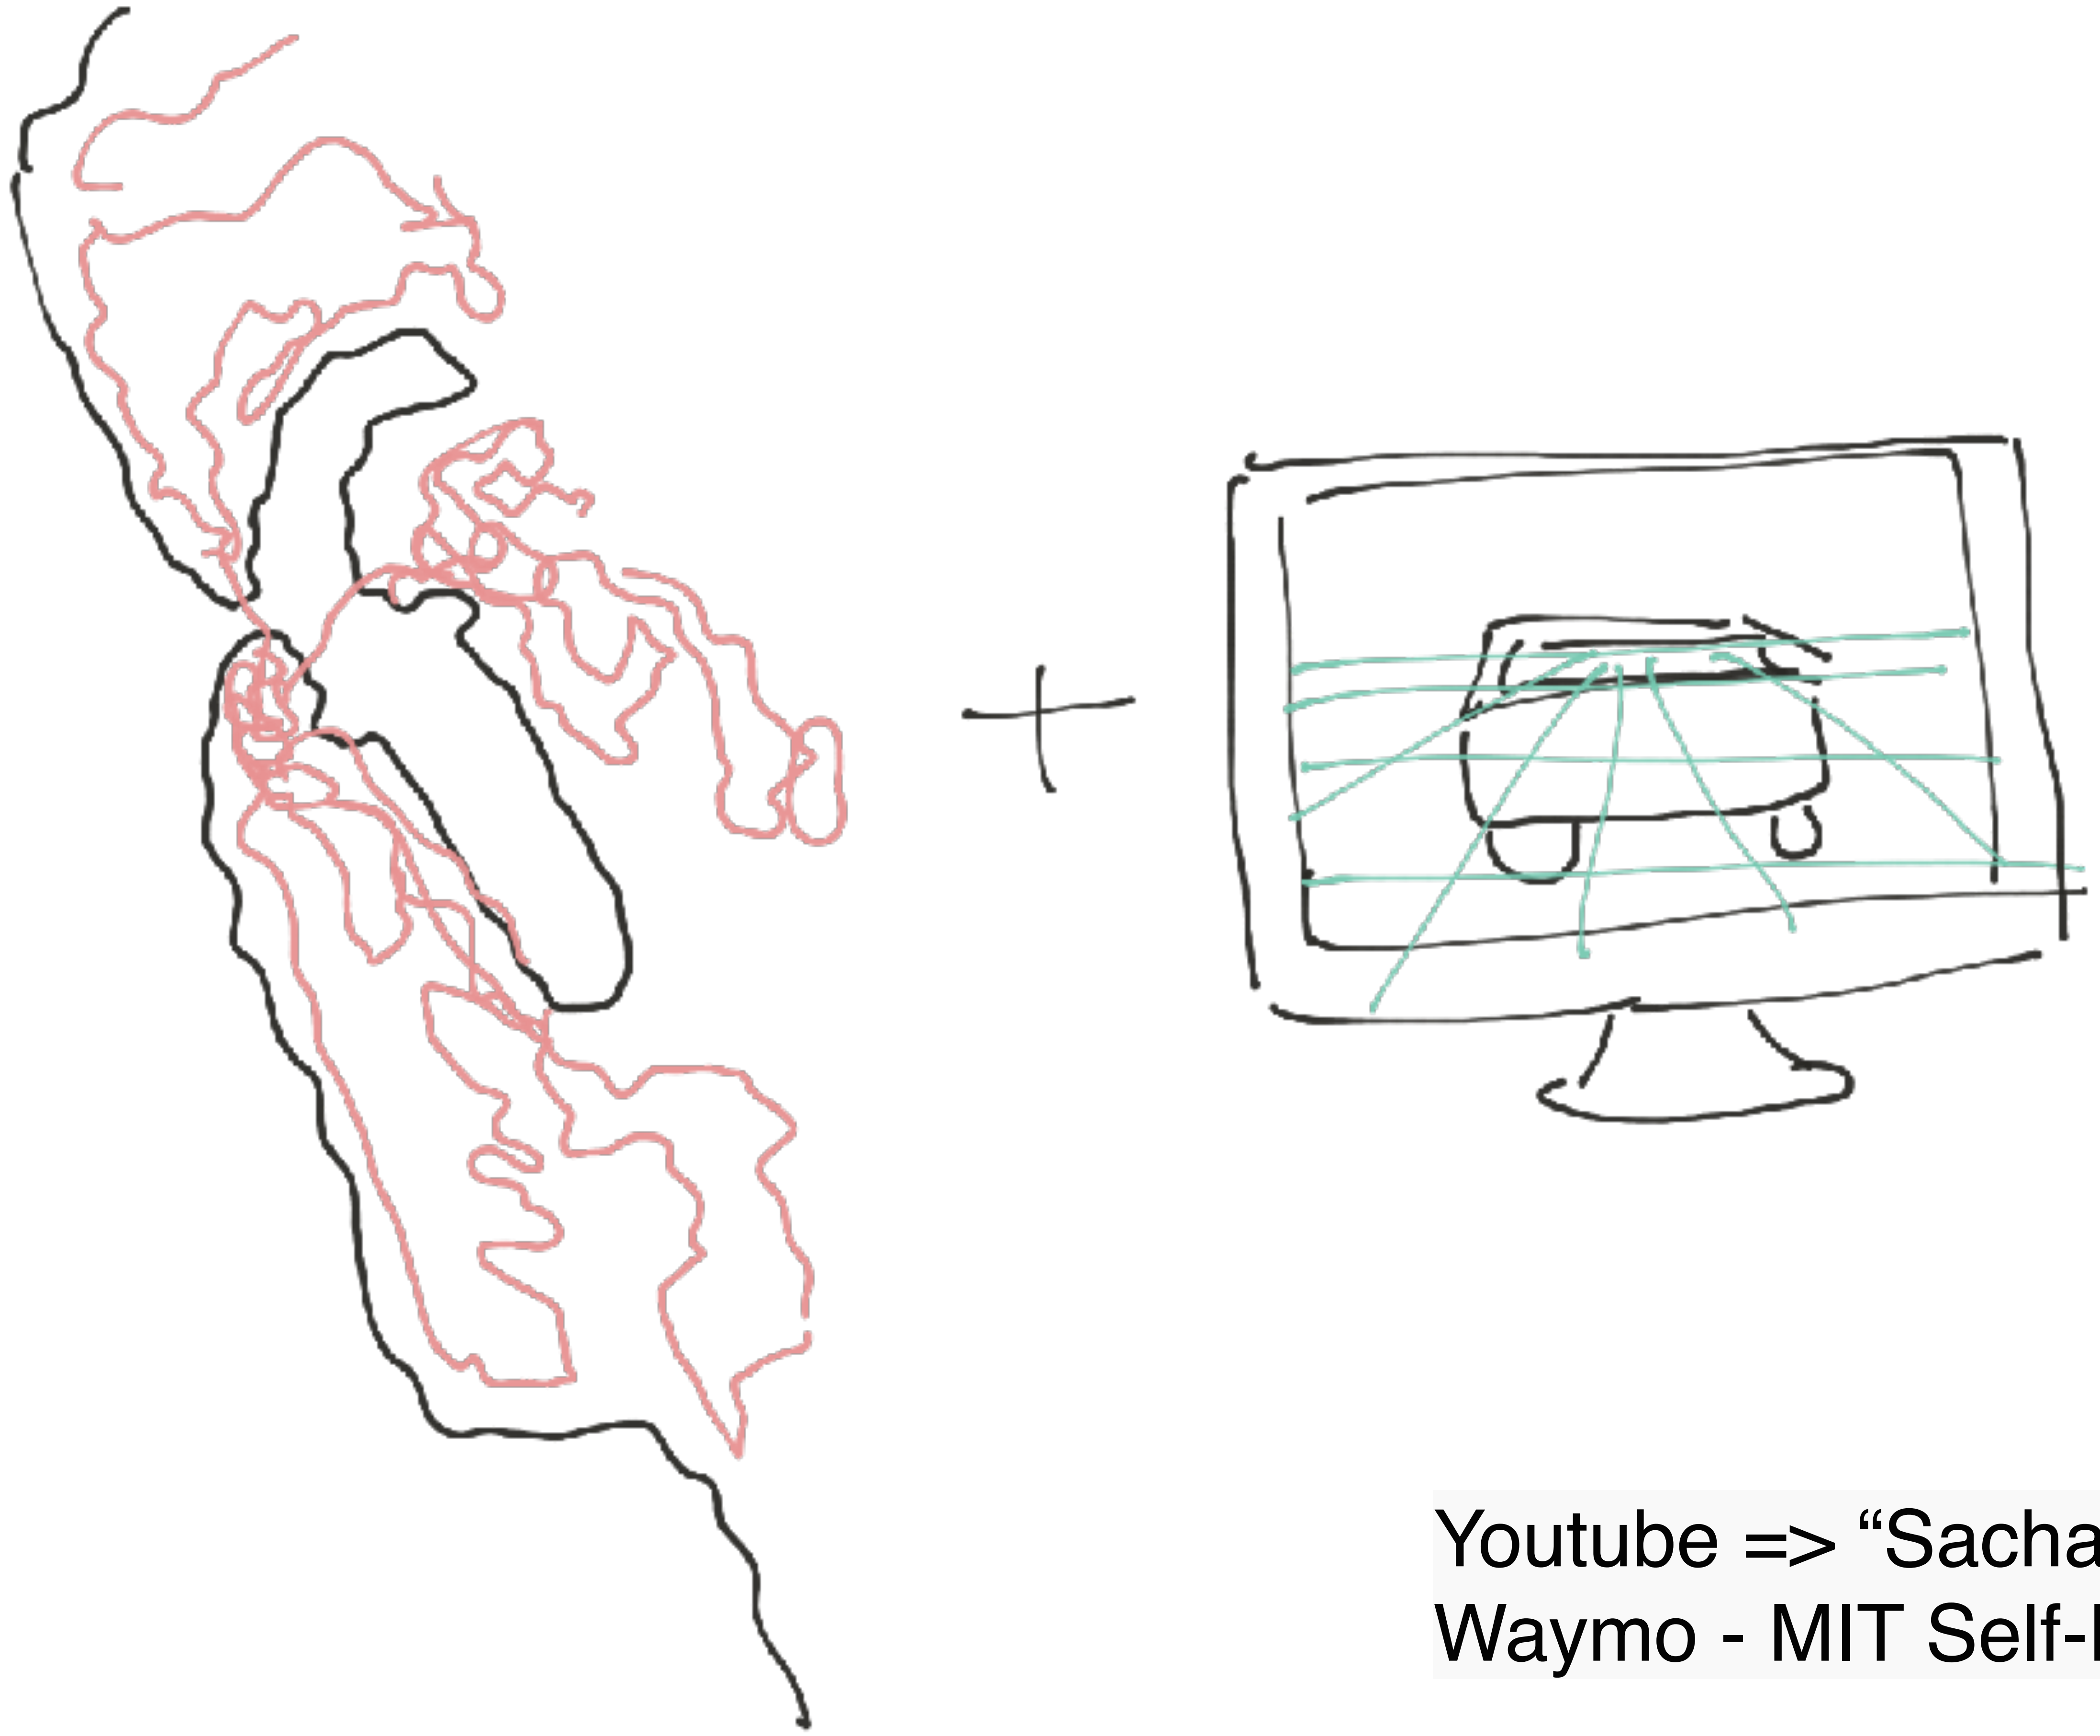

Youtube => “Sacha Arnoud, Director of Engineering, Waymo - MIT Self-Driving Cars”

Figure out which problems to tackle with ML

AI Innovation Landscape

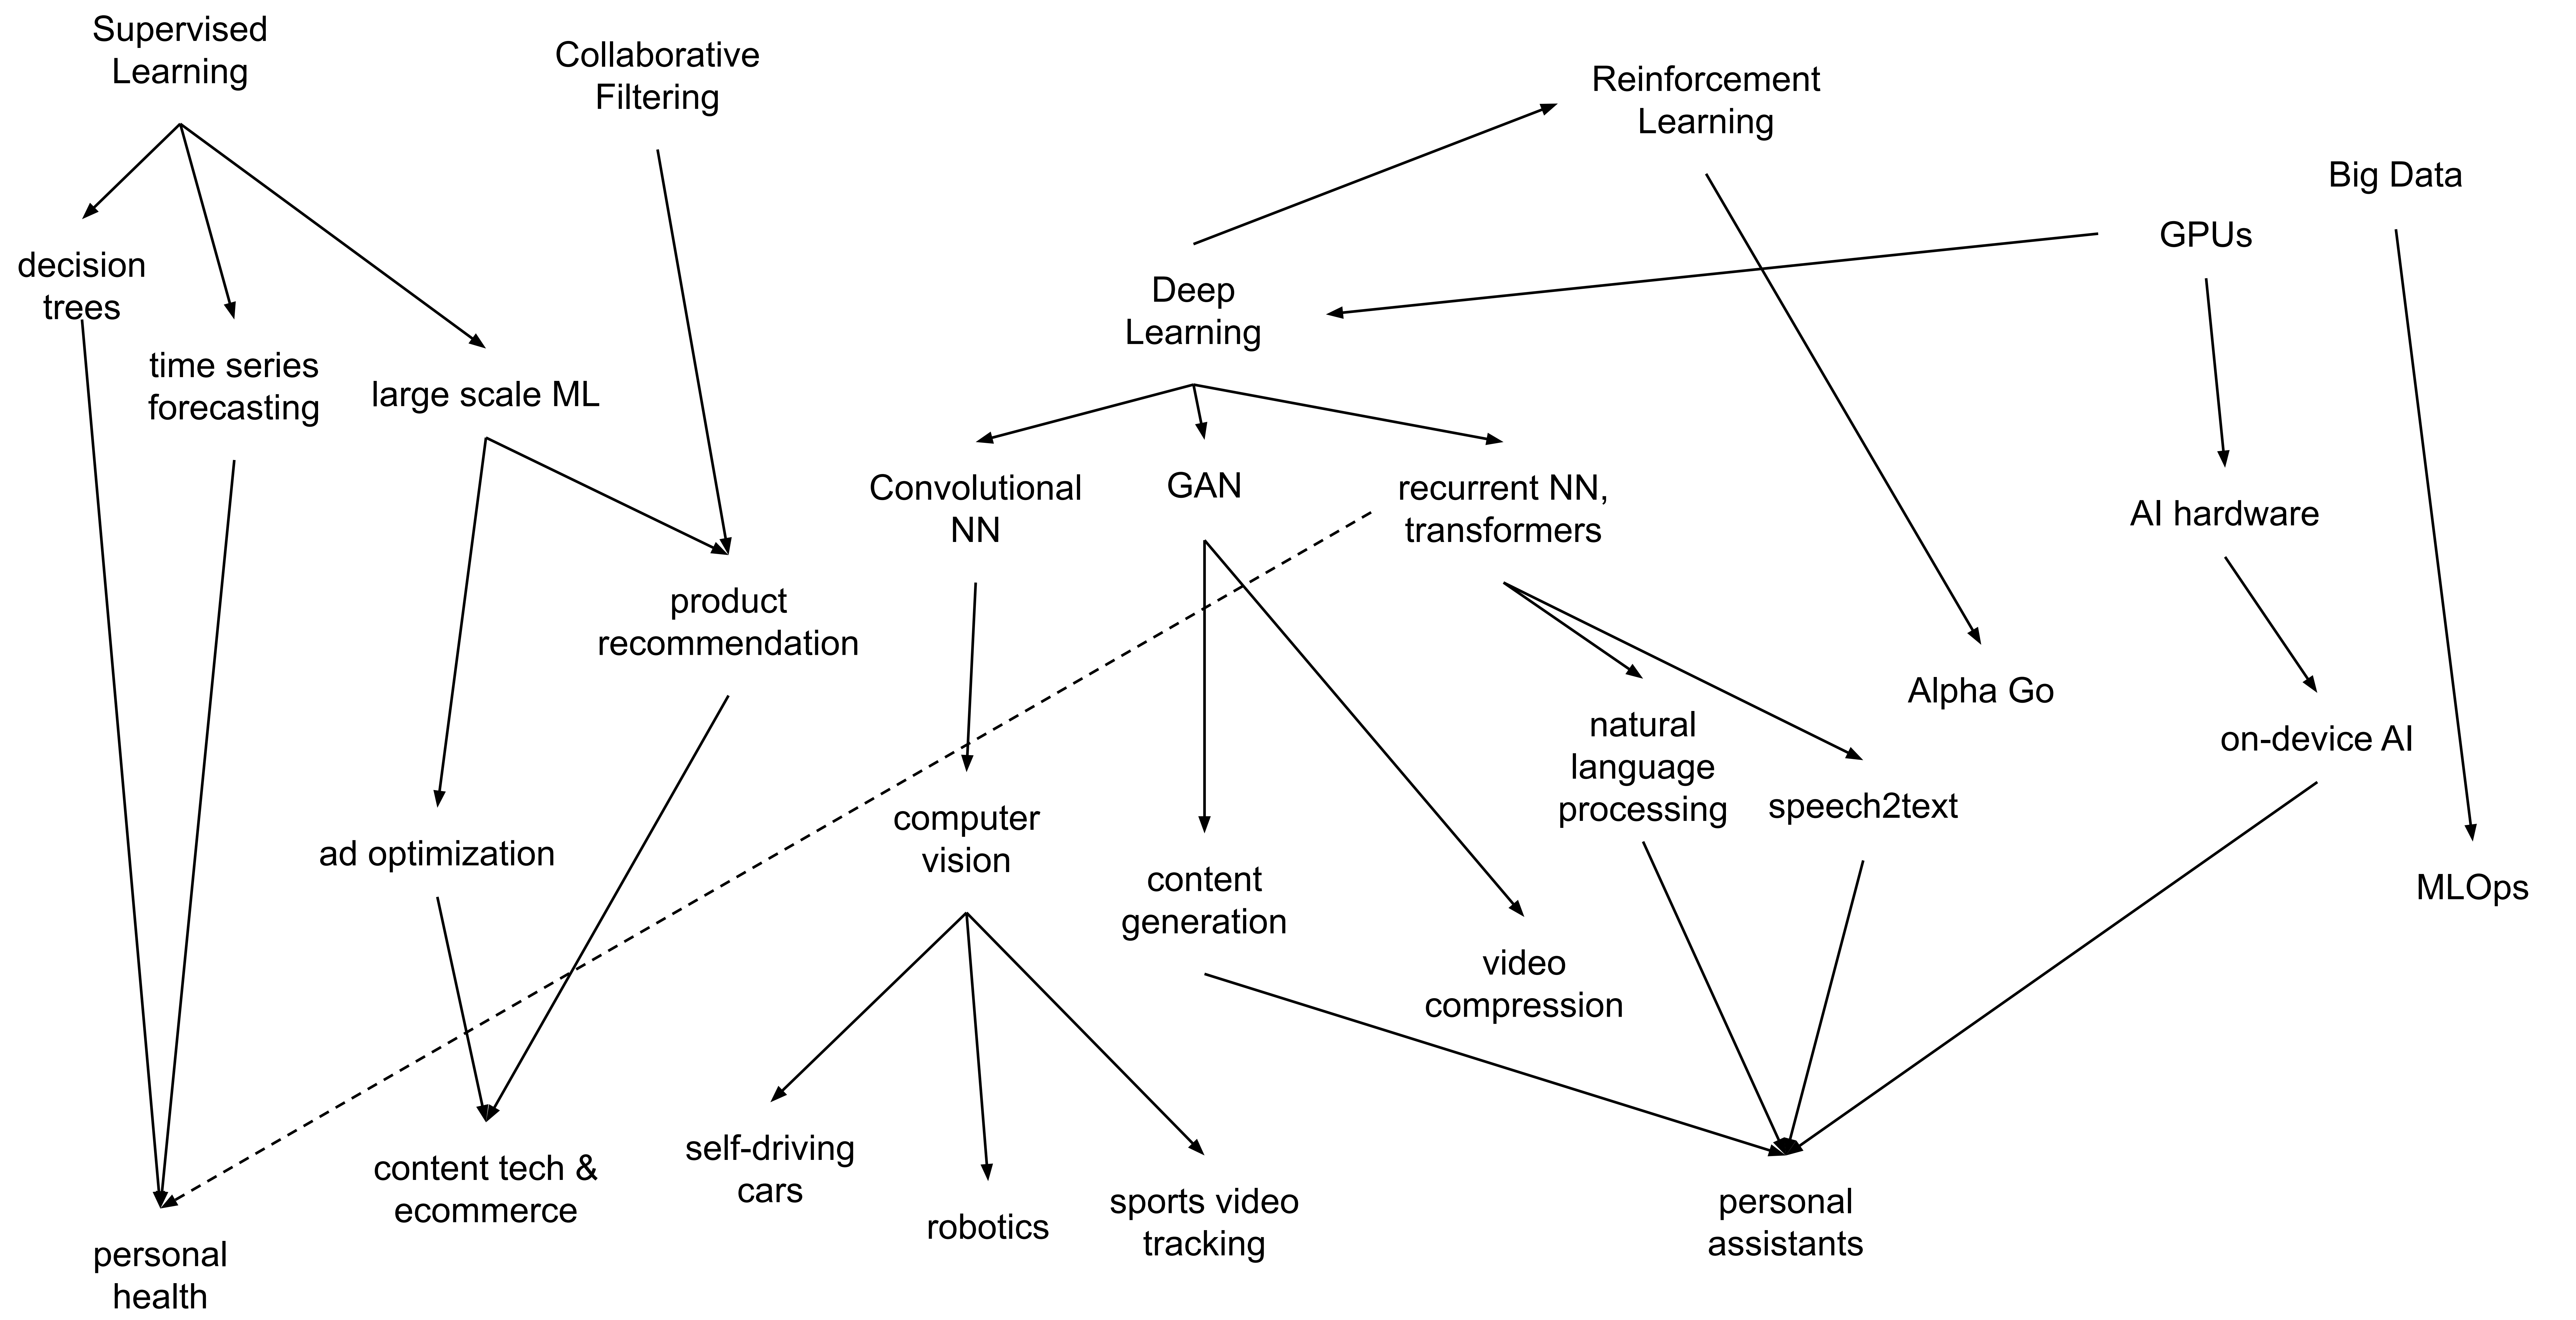

Dr. Mikio Braun

Research To Product

AI Startup Challenges

(a) Technology => Product?	(e.g. Alpha Go)
(b) Academia => Real World?	(e.g. Computer Vision)
(c) Complex solution from existing technology	(e.g. Self-driving cars)
(d) Data	(e.g. Health)
(e) Technology Challenges (performance, resources)	(e.g. GANs)

Jobs-To-Be-Done

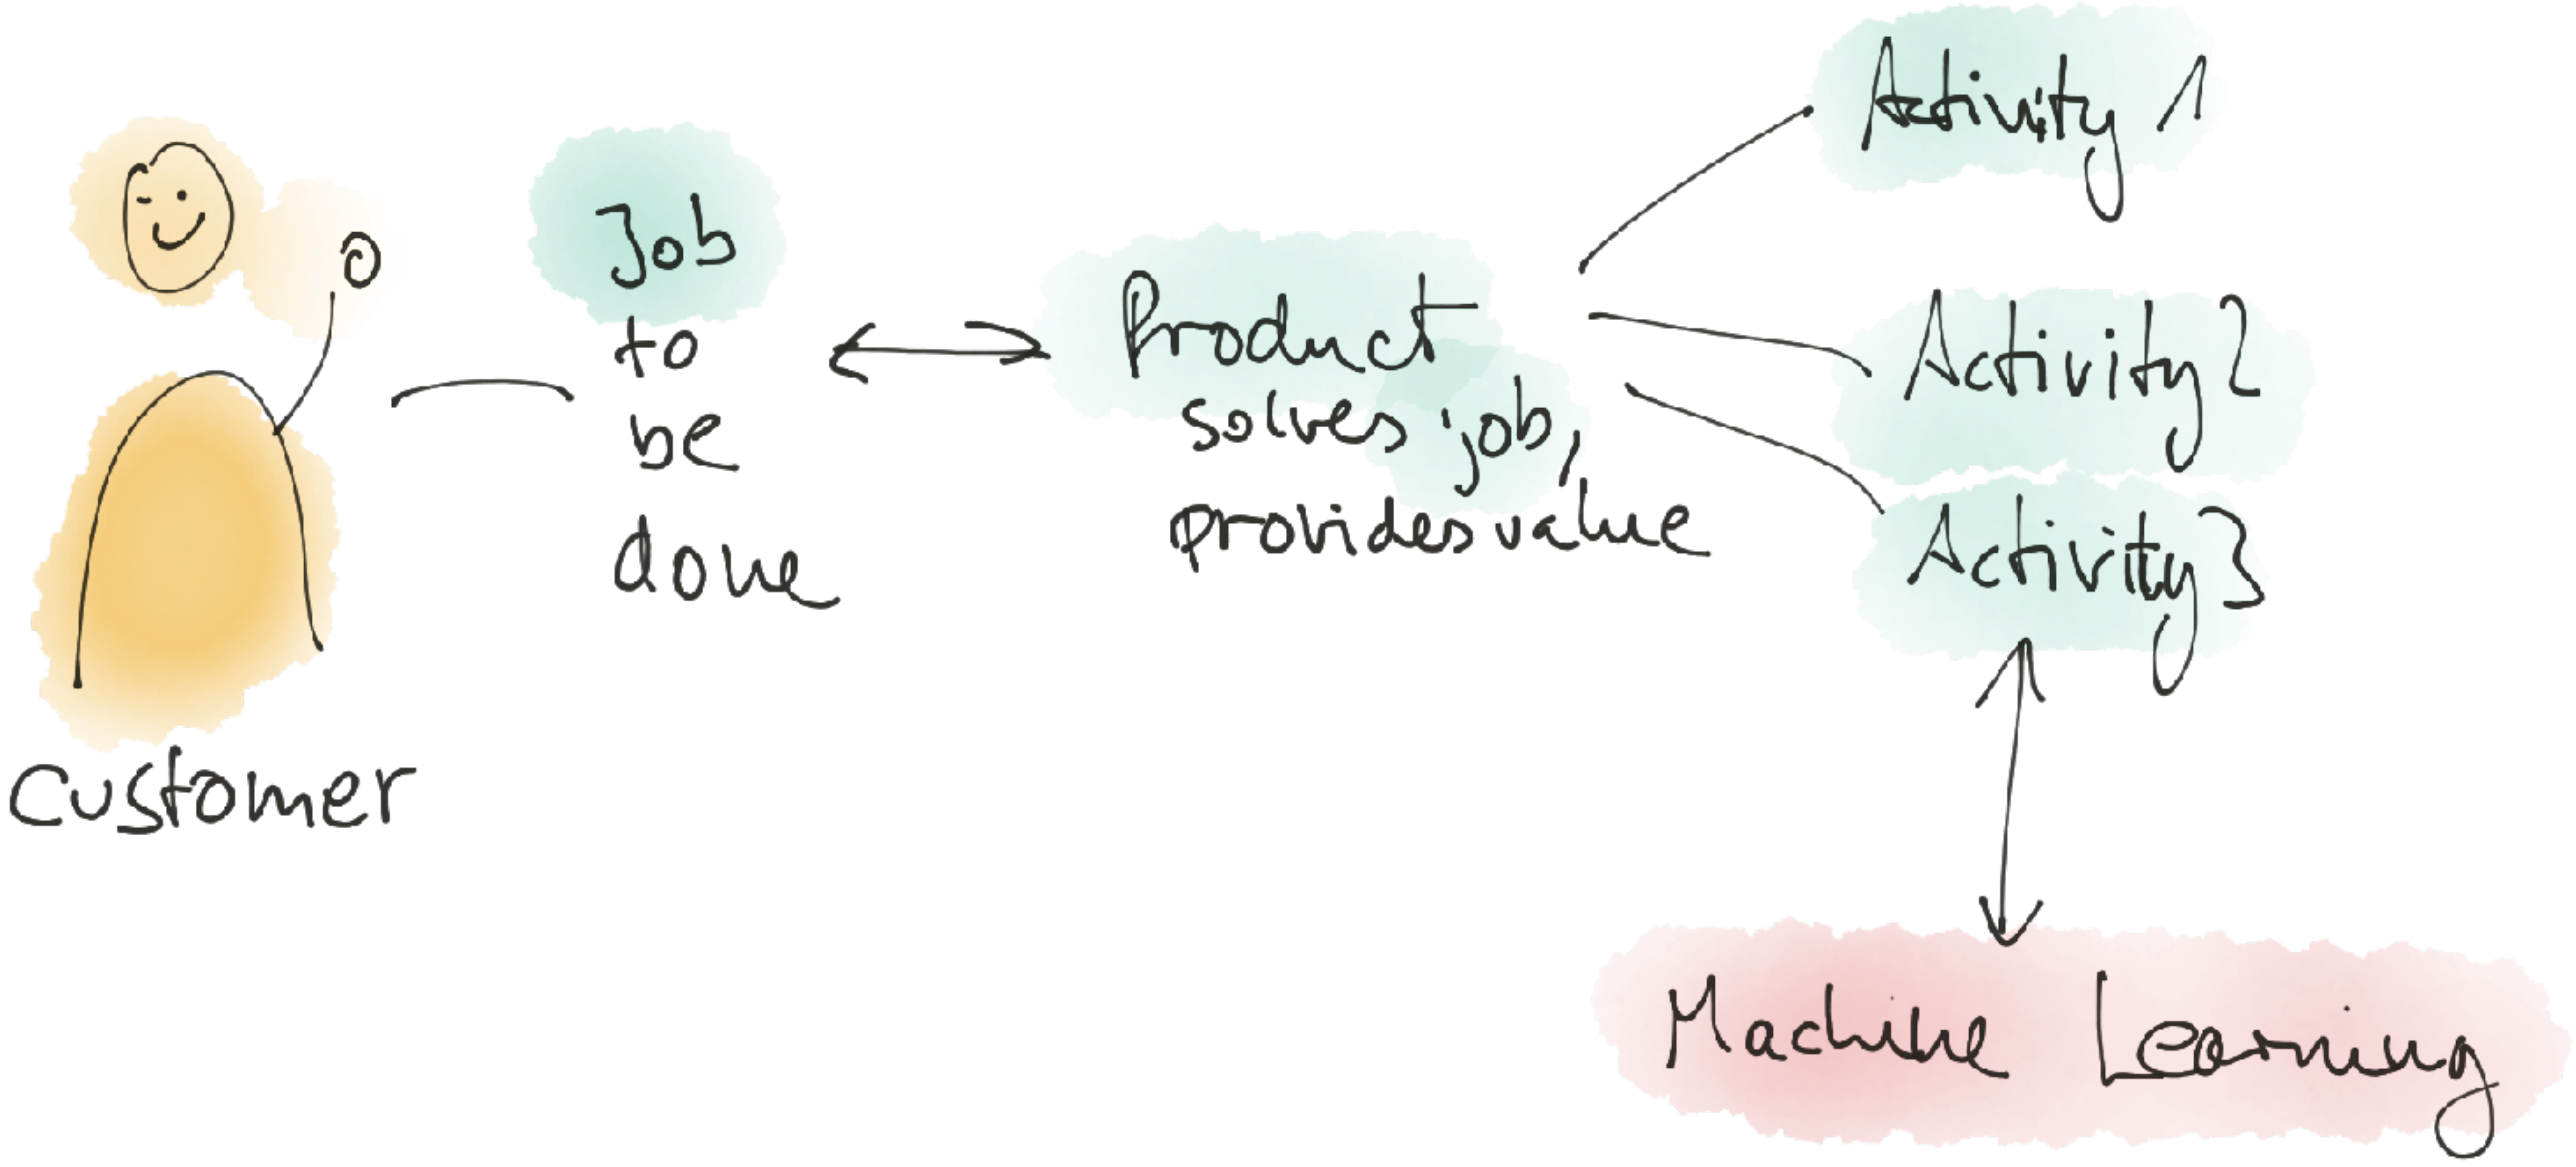

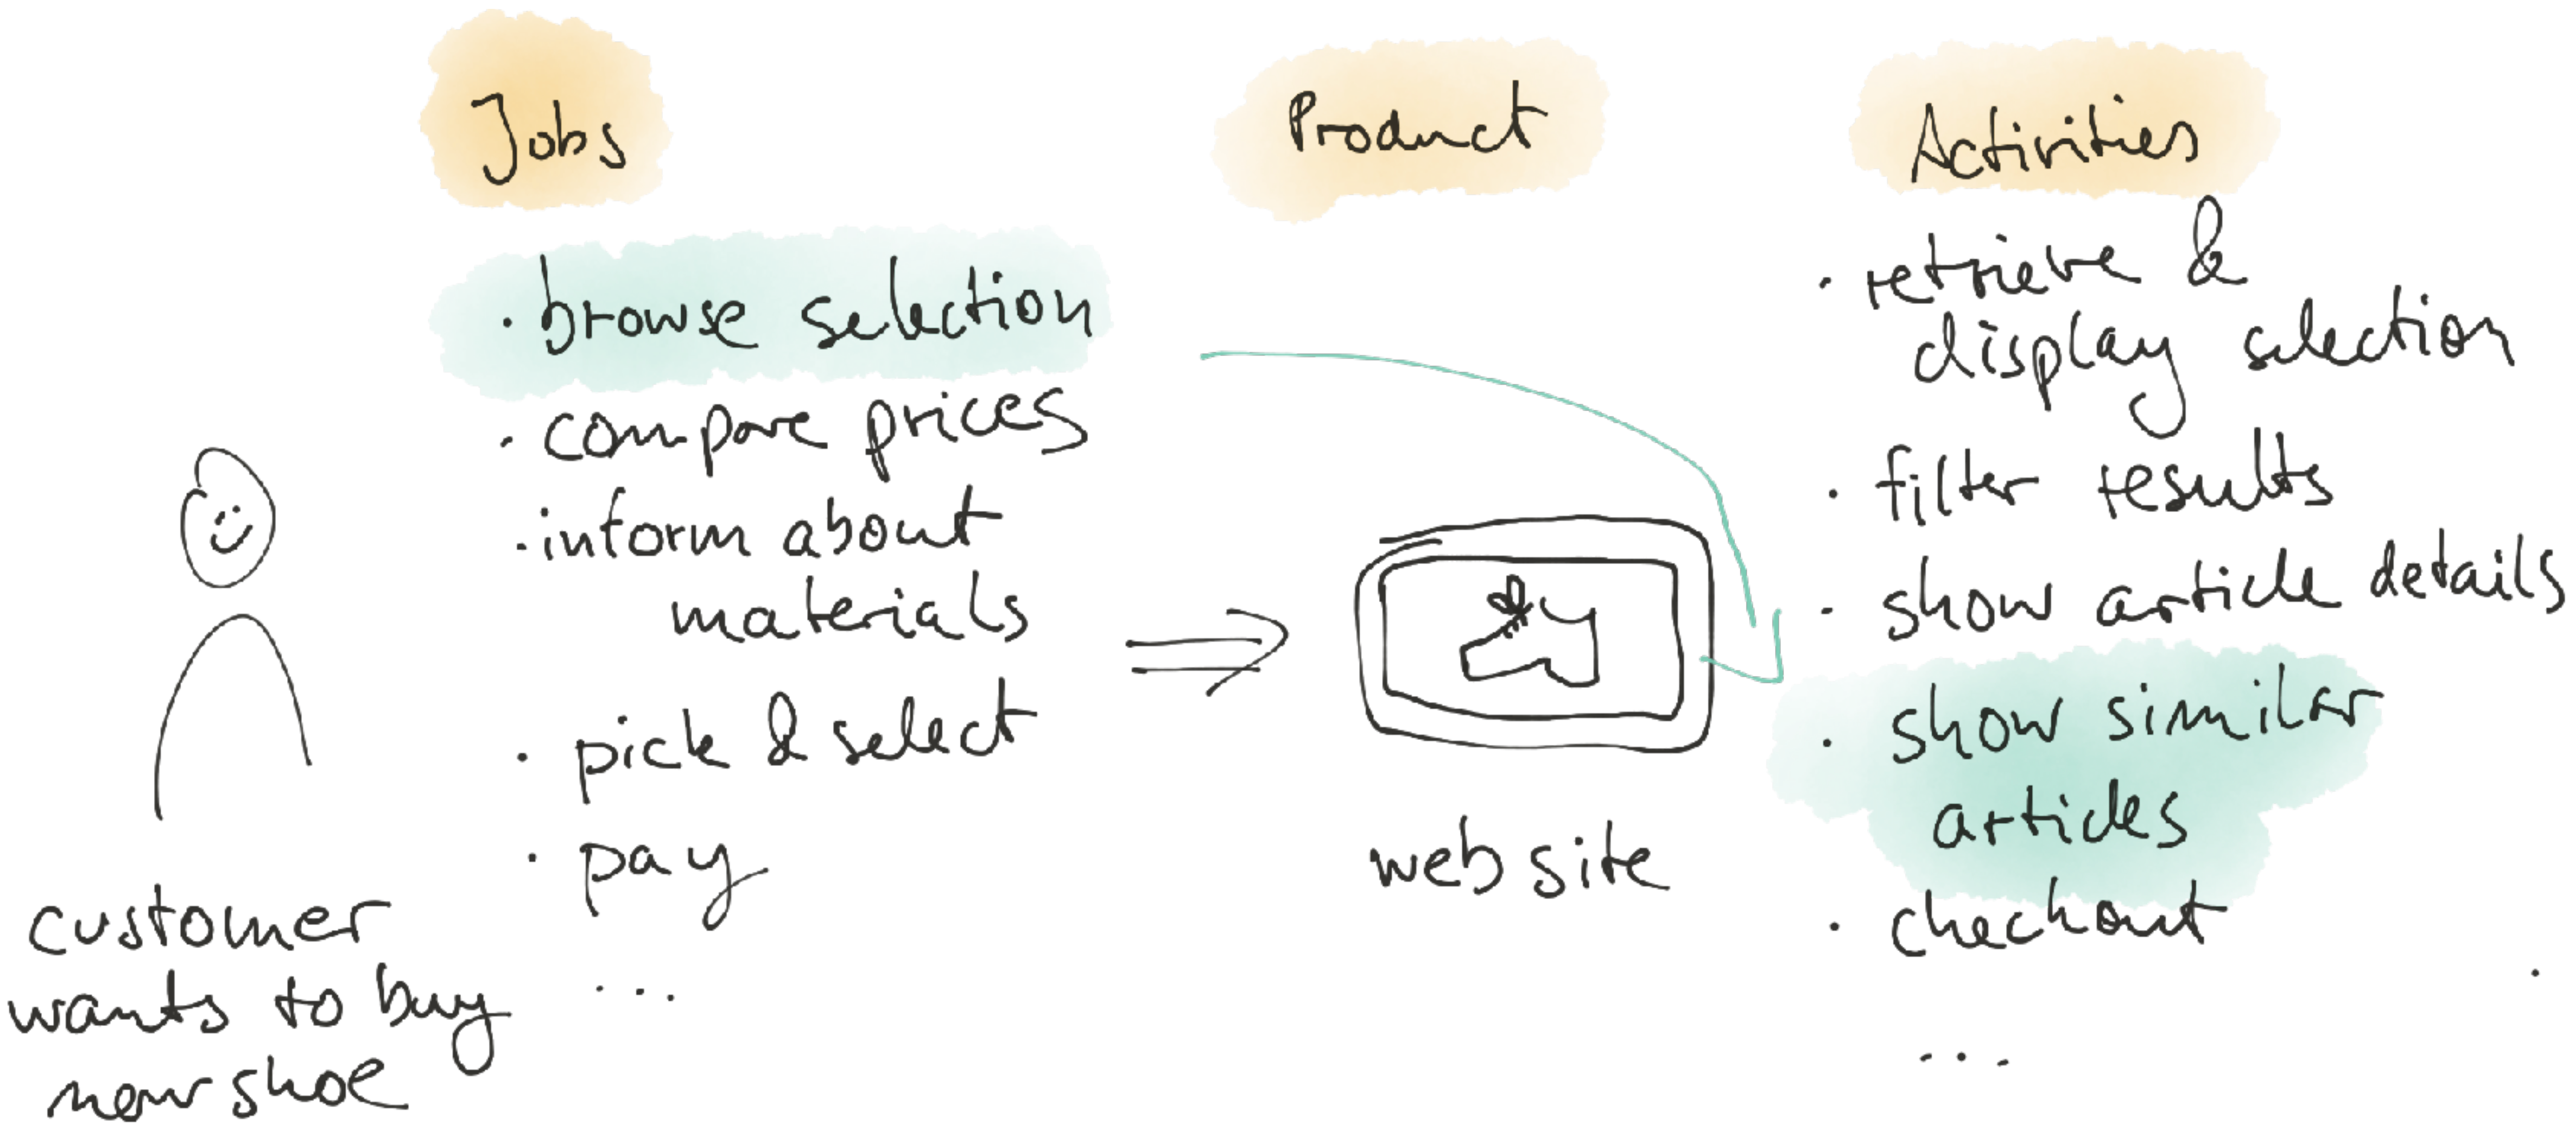
Jobs
· browse selection
· compare prices
· inform about materials
· pick & select
· pay
...
customer wants to buy new shoe
Product
web site
Activities
· retrieve & display selection
· filter results
· show article details
· show similar articles
· checkout
...

Typical problem for machine learning:

- Hard to specify what exactly means “similar.”
- A lot of example data is available.
- Recommendations have to change based on new articles frequently.

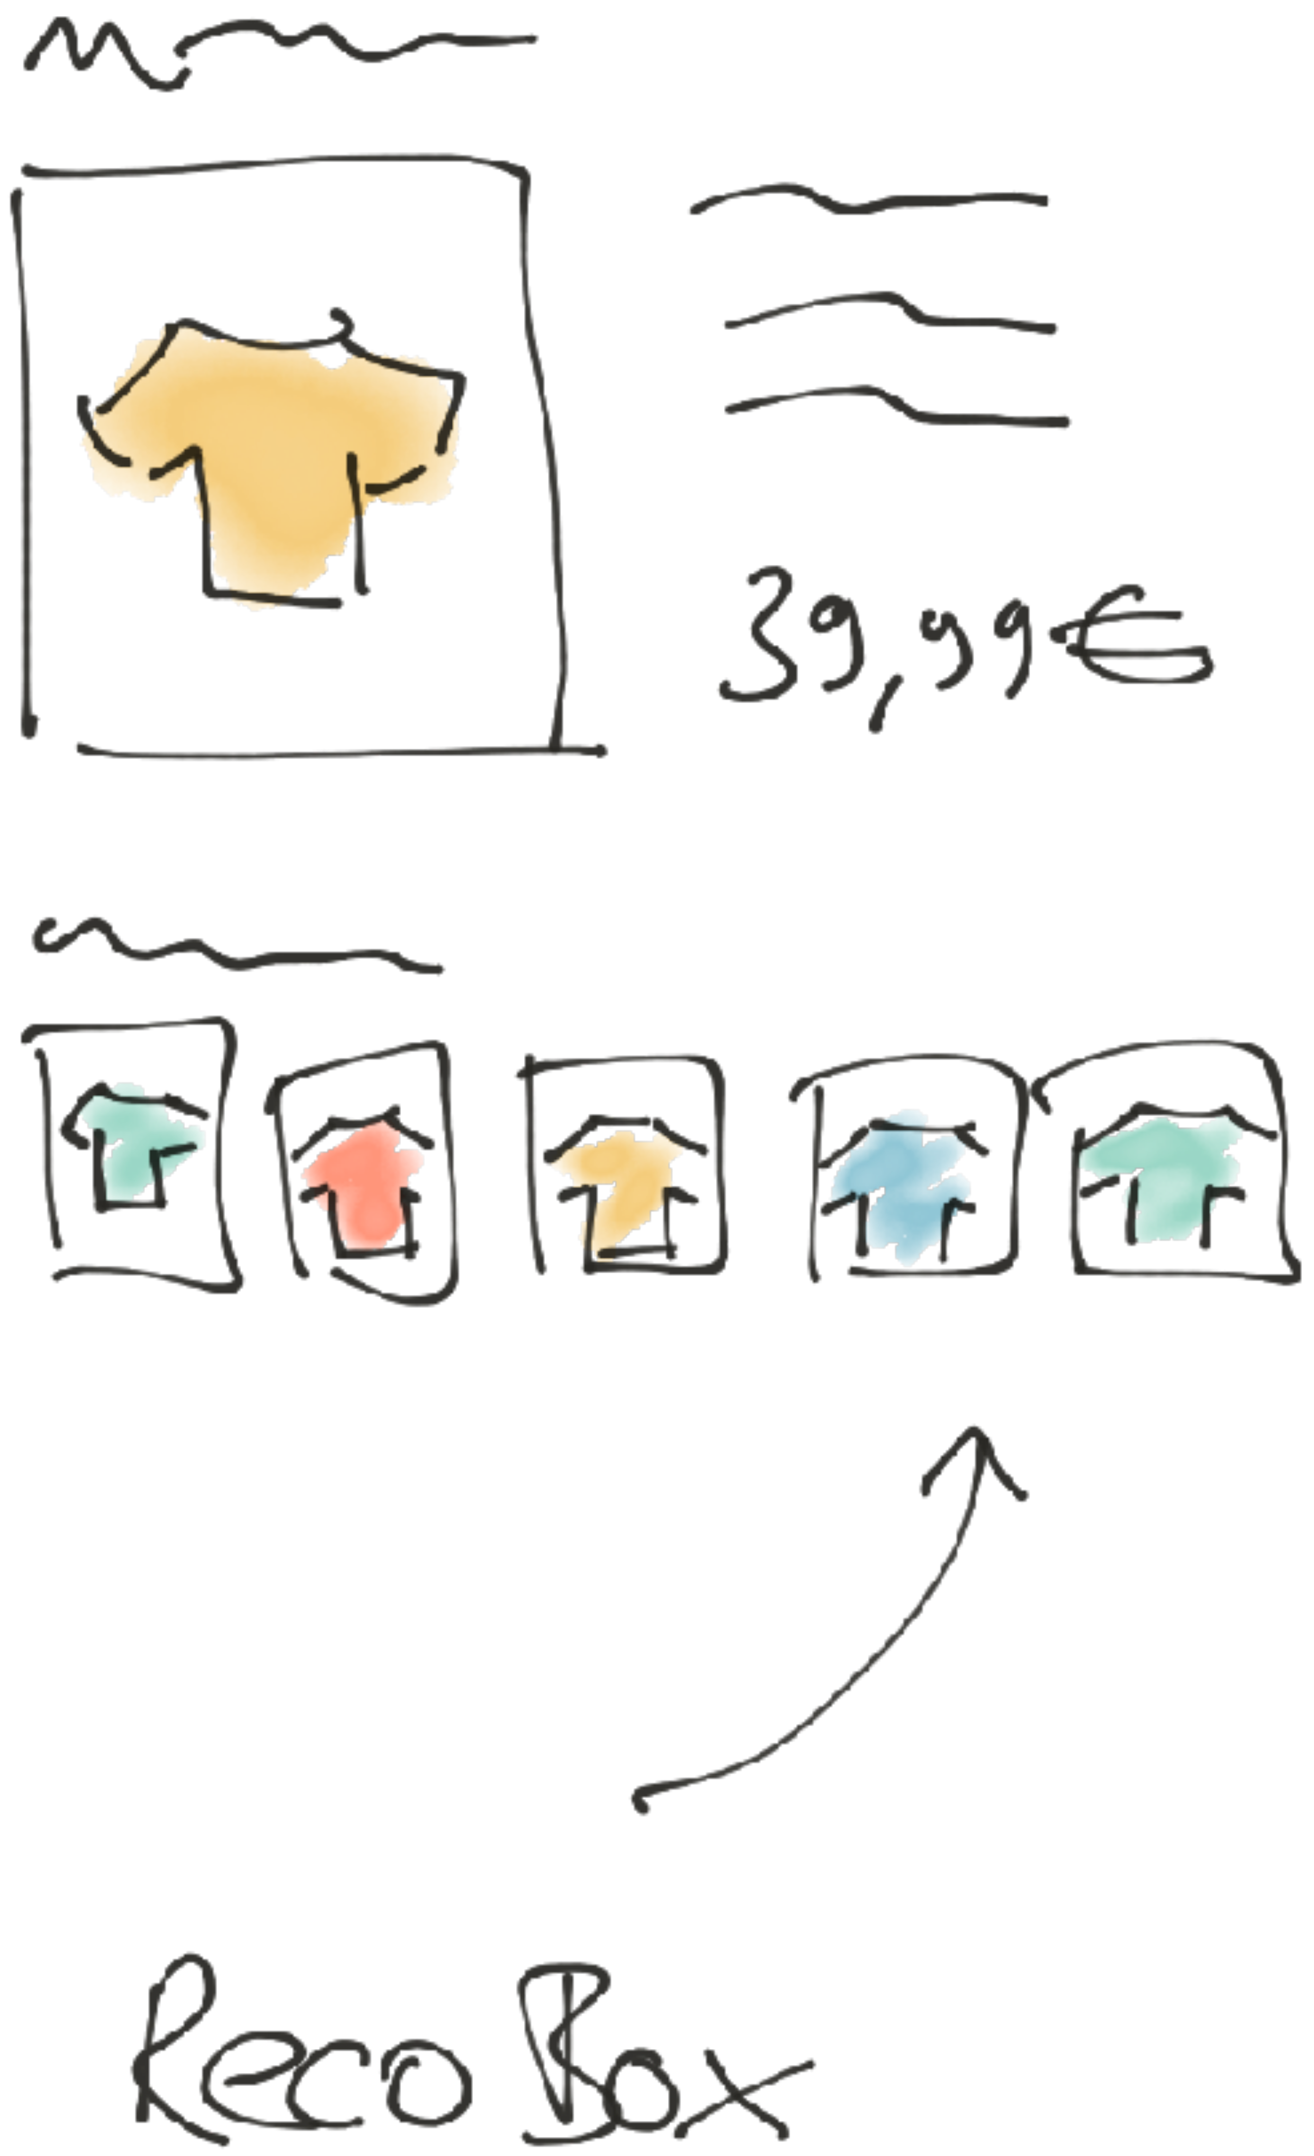

Learn how to gather your data

1. Data quality is more important than models. If you have an okay model, invest time to improve data quality than tweak the model.

2. Don't spend years building up data infrastructure first. Data is important, but you also need to learn what kind of data you need, learn by doing small ML projects.

3. Don't just look for problems best suited for ML, but also for the most important business problems. Solve important customer problems.

4. There is a "small data regime" where you can look at individual points and discuss labels, etc.

Andrew Ng: Forget about building an AI-first business. Start with a mission.

An AI pioneer reflects on how companies can use machine learning to transform their operations and solve critical problems.

by **Karen Hao** March 26, 2021

JEREMY PORTJE

https://www.technologyreview.com/2021/03/26/1021258/ai-pioneer-andrew-ng-machine-learning-business/

Bootstrapping Your Data

- Iteration 0: Manually created heuristic, focus on UI.
- Iteration 1: Model trained from hand labelled data.
- Iteration 2: First model trained on feedback data.
- Iteration 3: First model trained on click data.
- …

Put ML in production

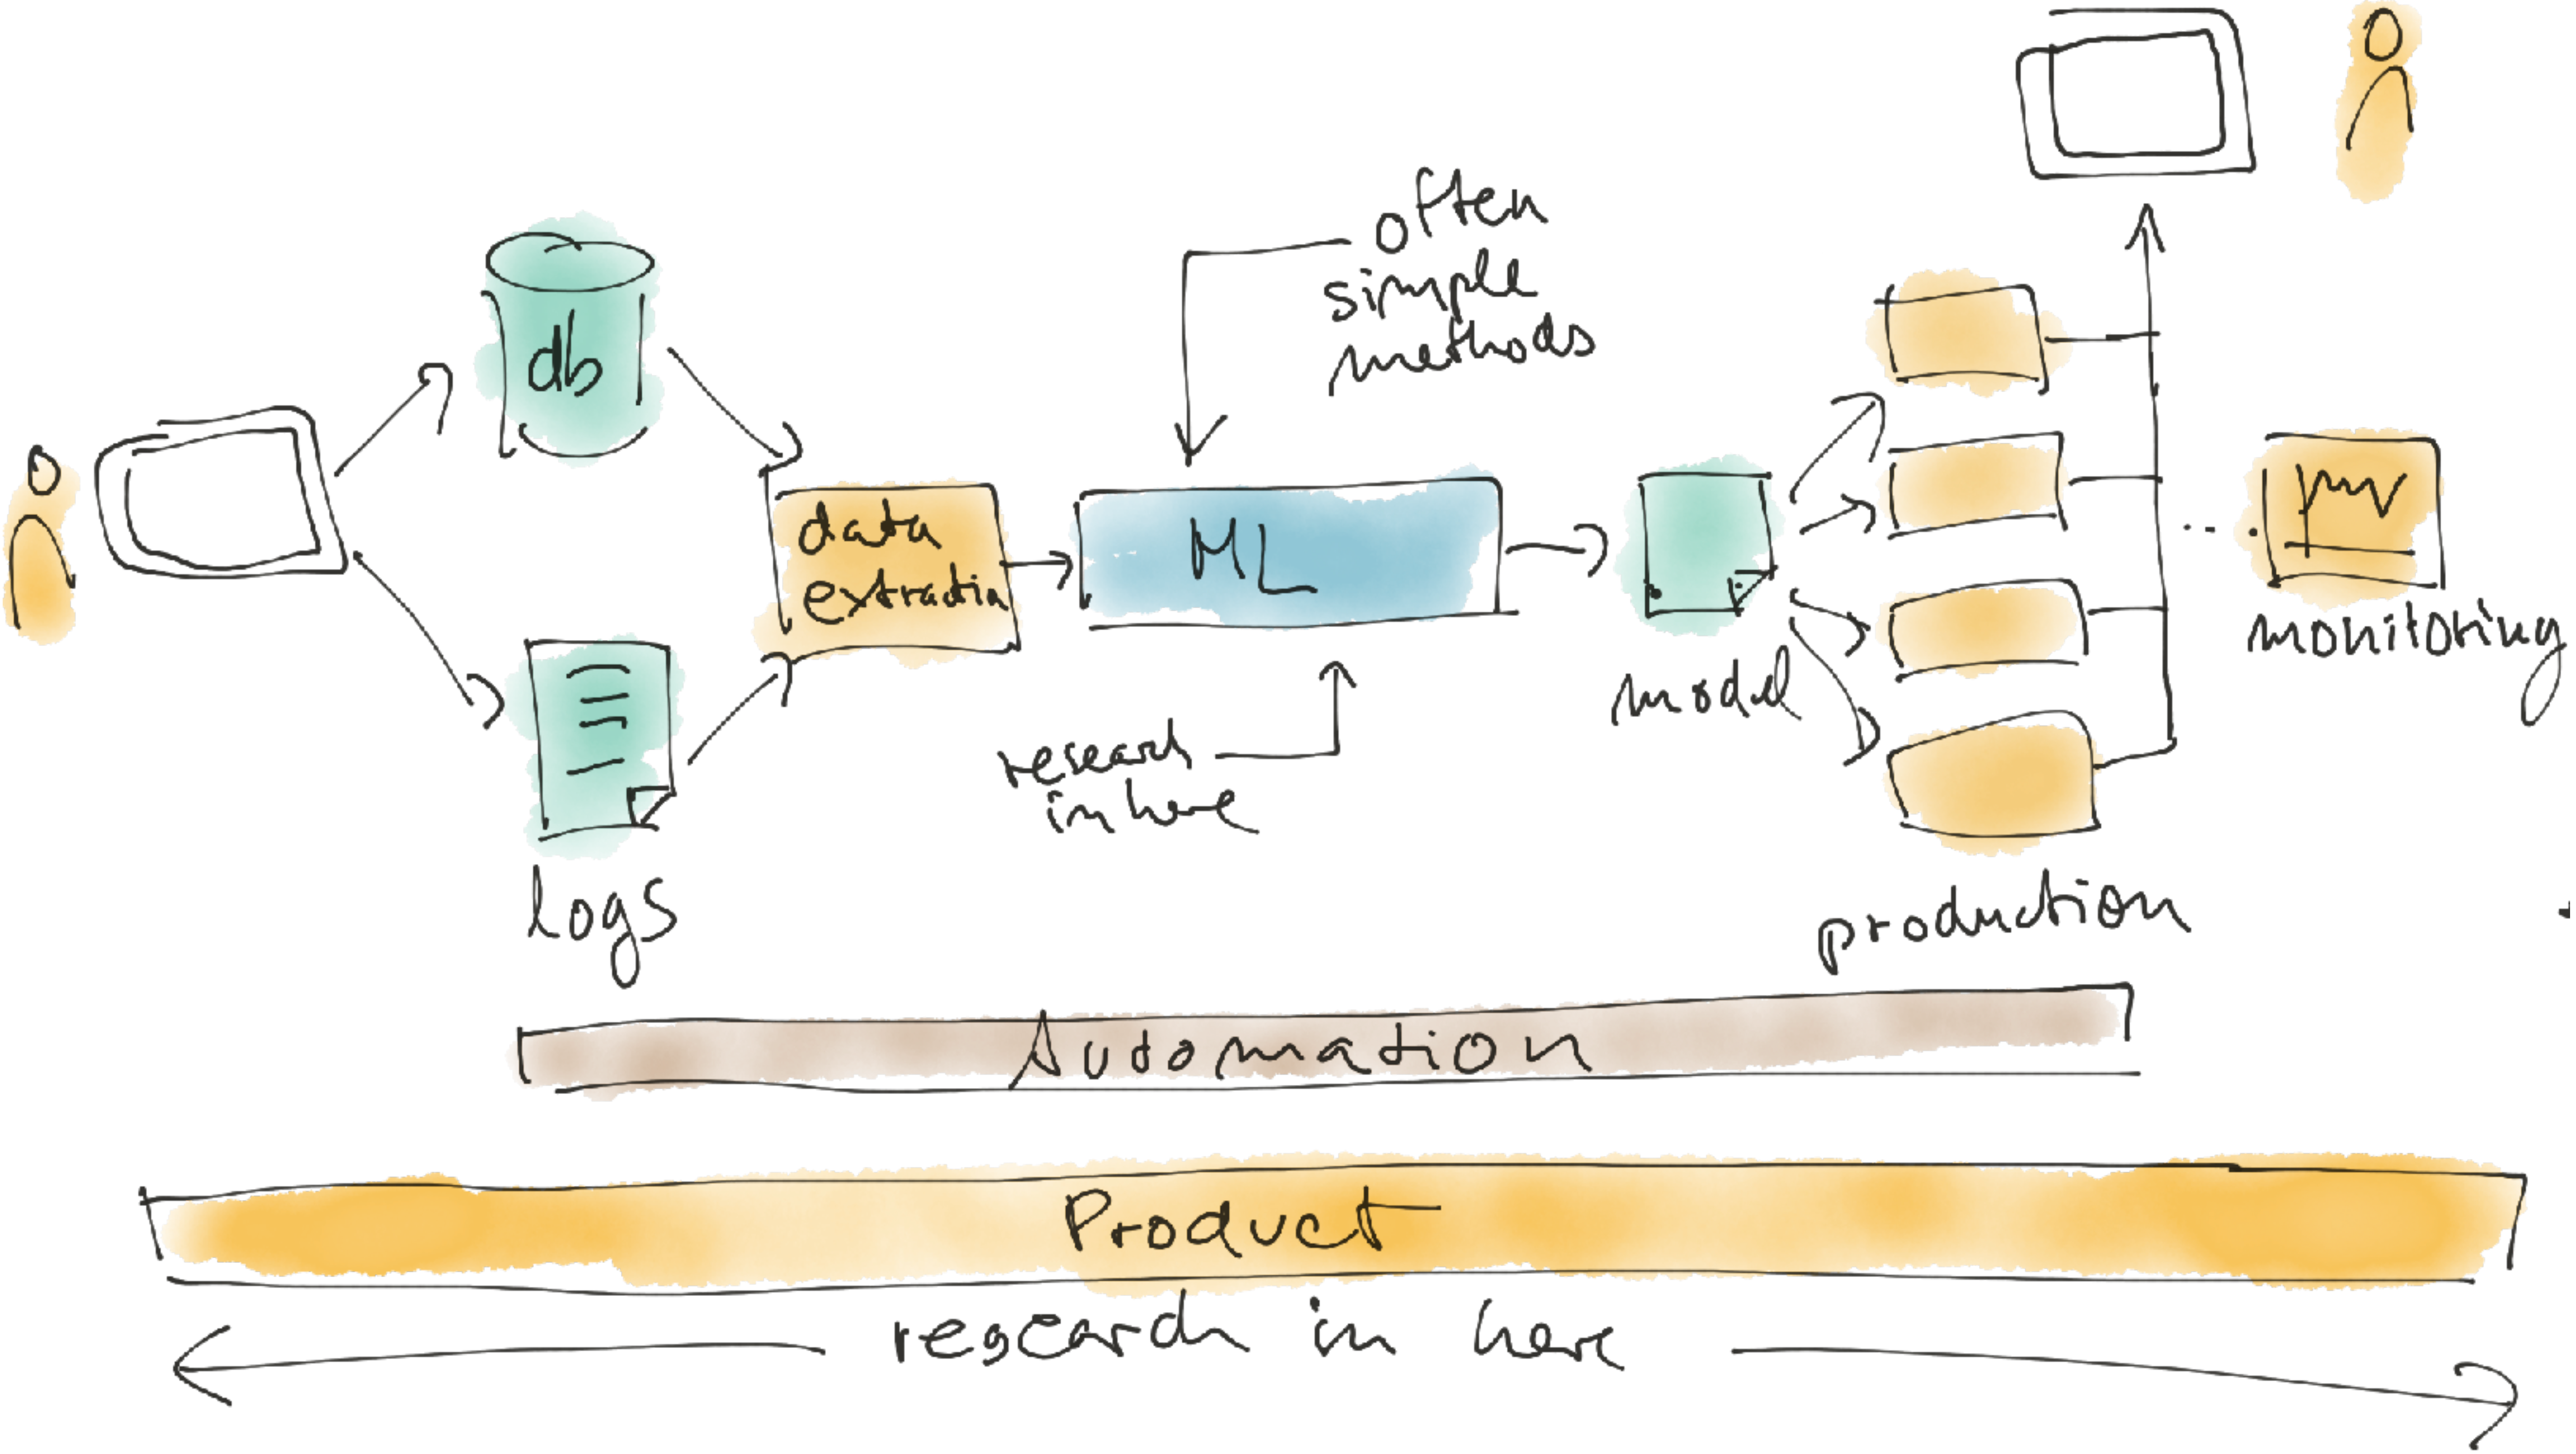

Dr. Mikio Braun

Design Patterns for AI Architecture

Core Machine Learning
—how to train, evaluate, etc.

Serving
—access predictions in real-time

Data Preprocessing and Features
—how to deal with preprocessing

Automation & Monitoring
—making it more production ready

Machine Learning Integration
—how to fit it into a larger picture

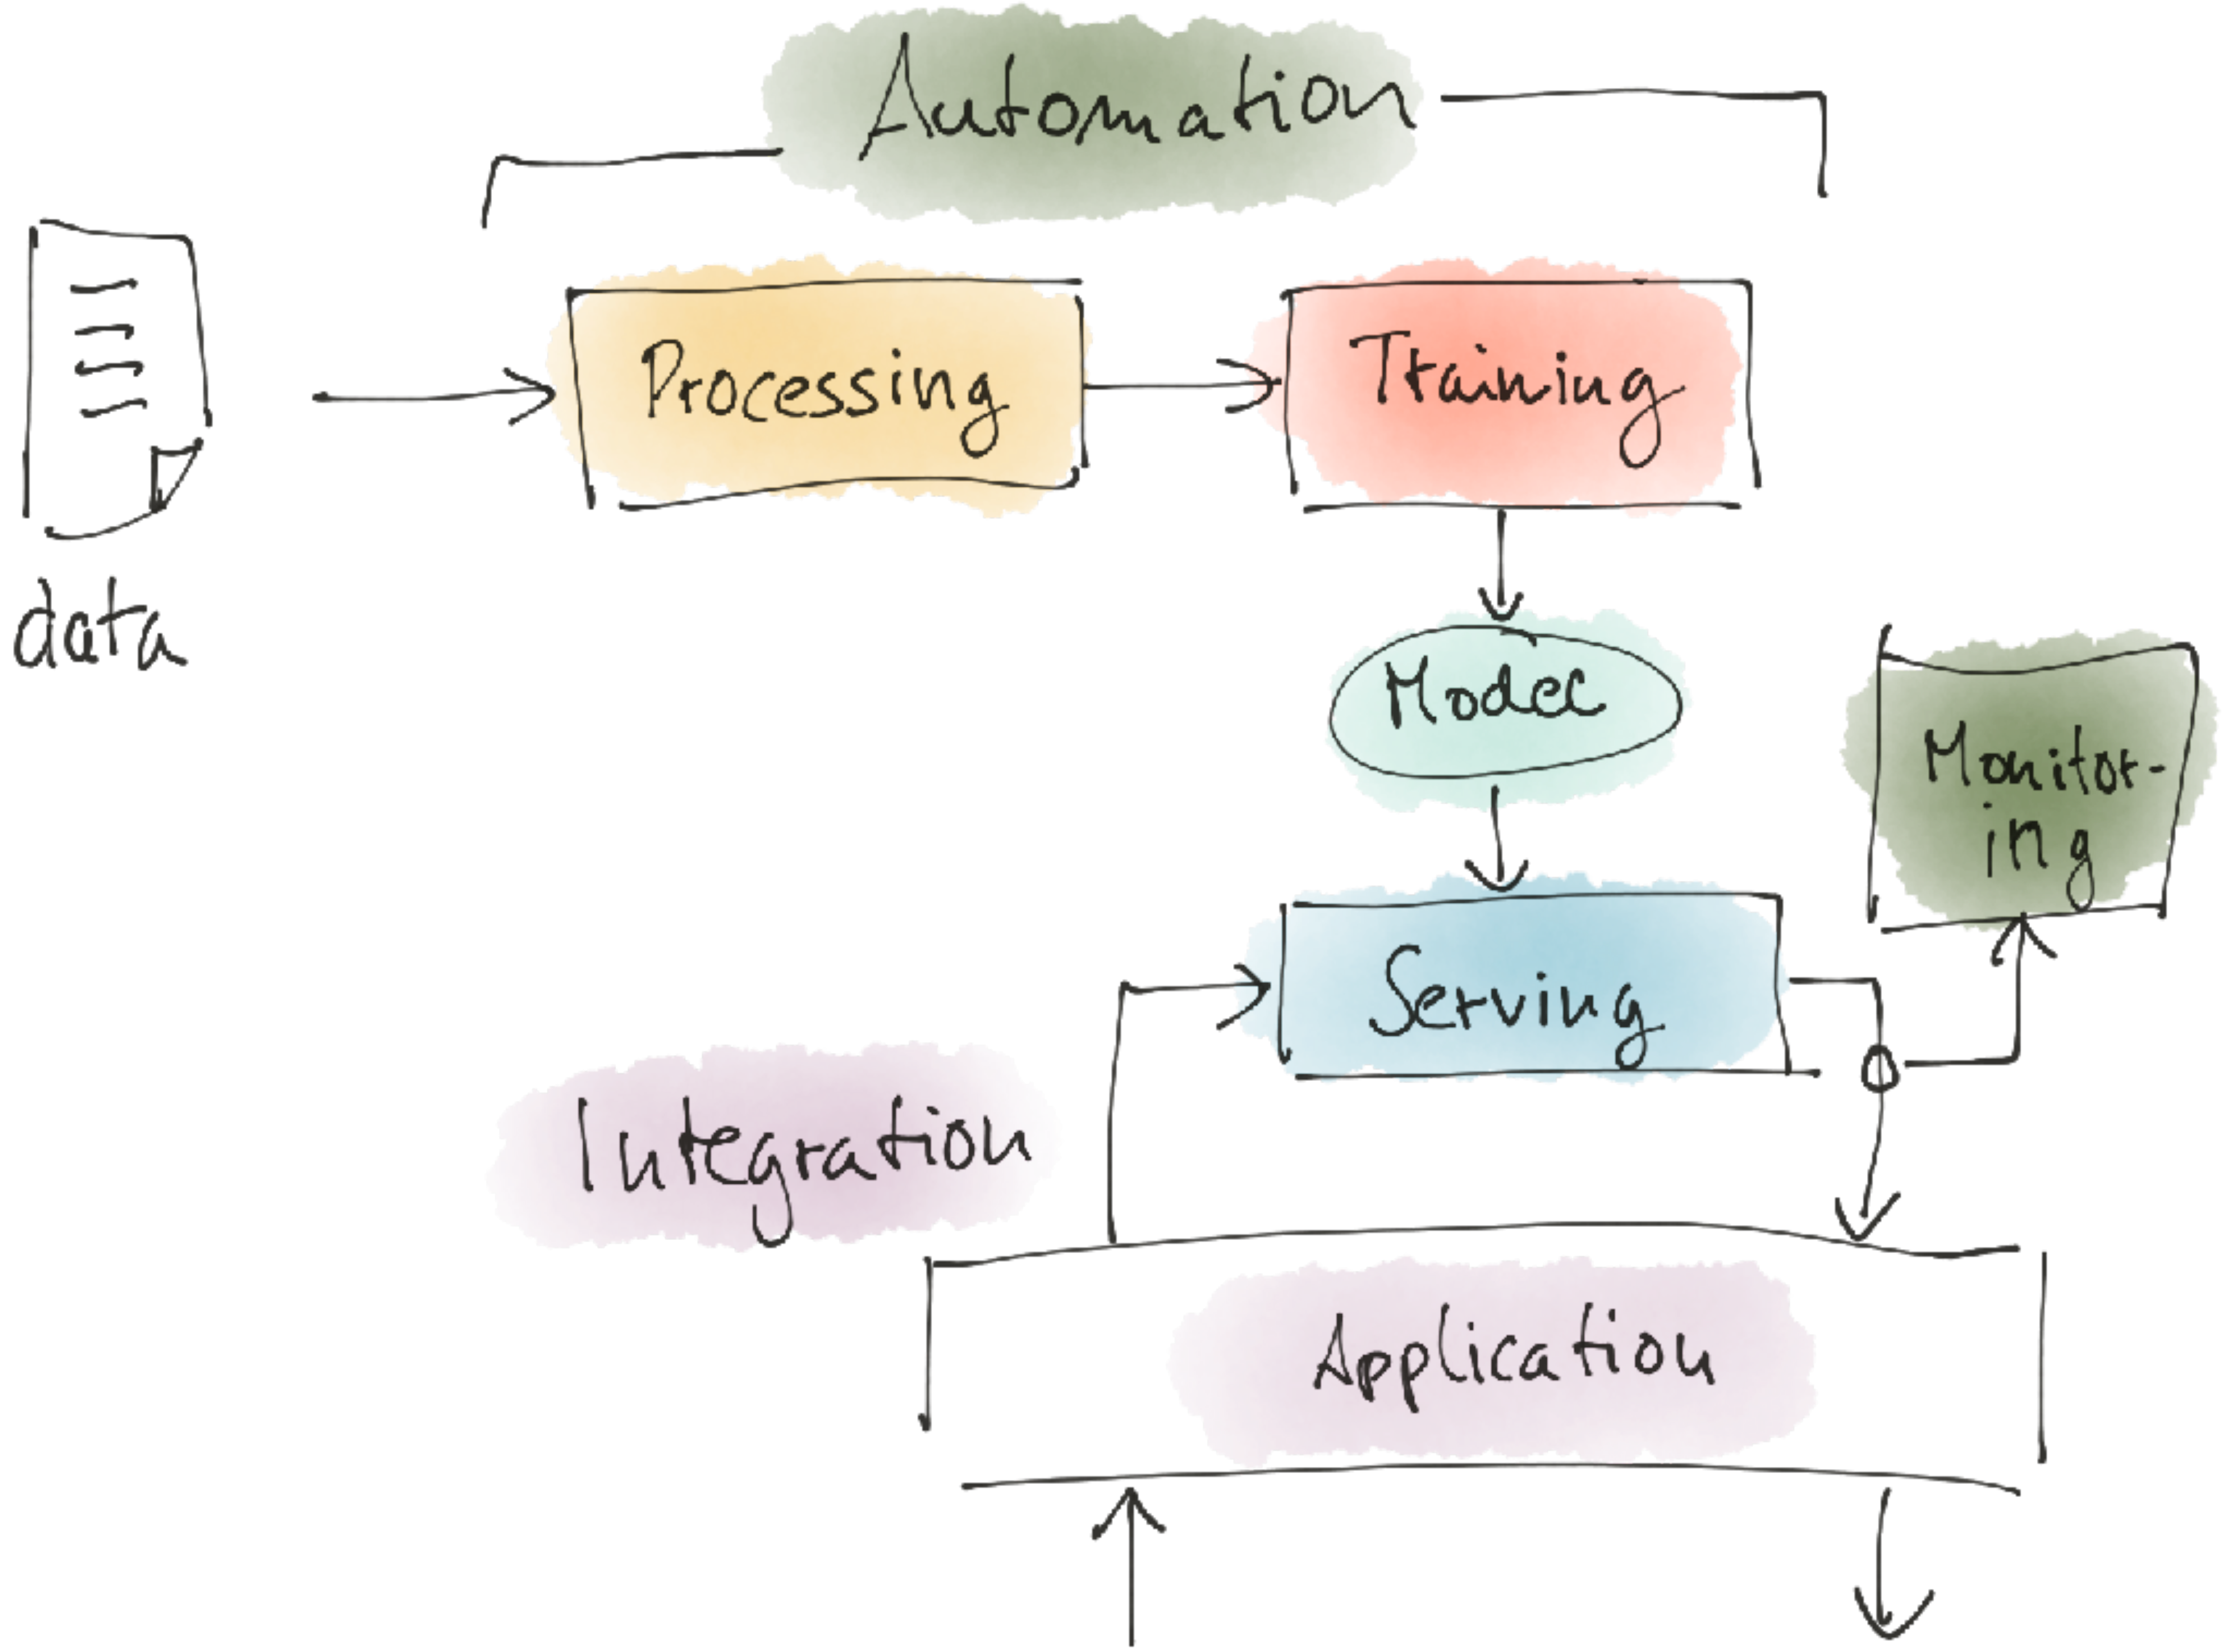

Dr. Mikio Braun

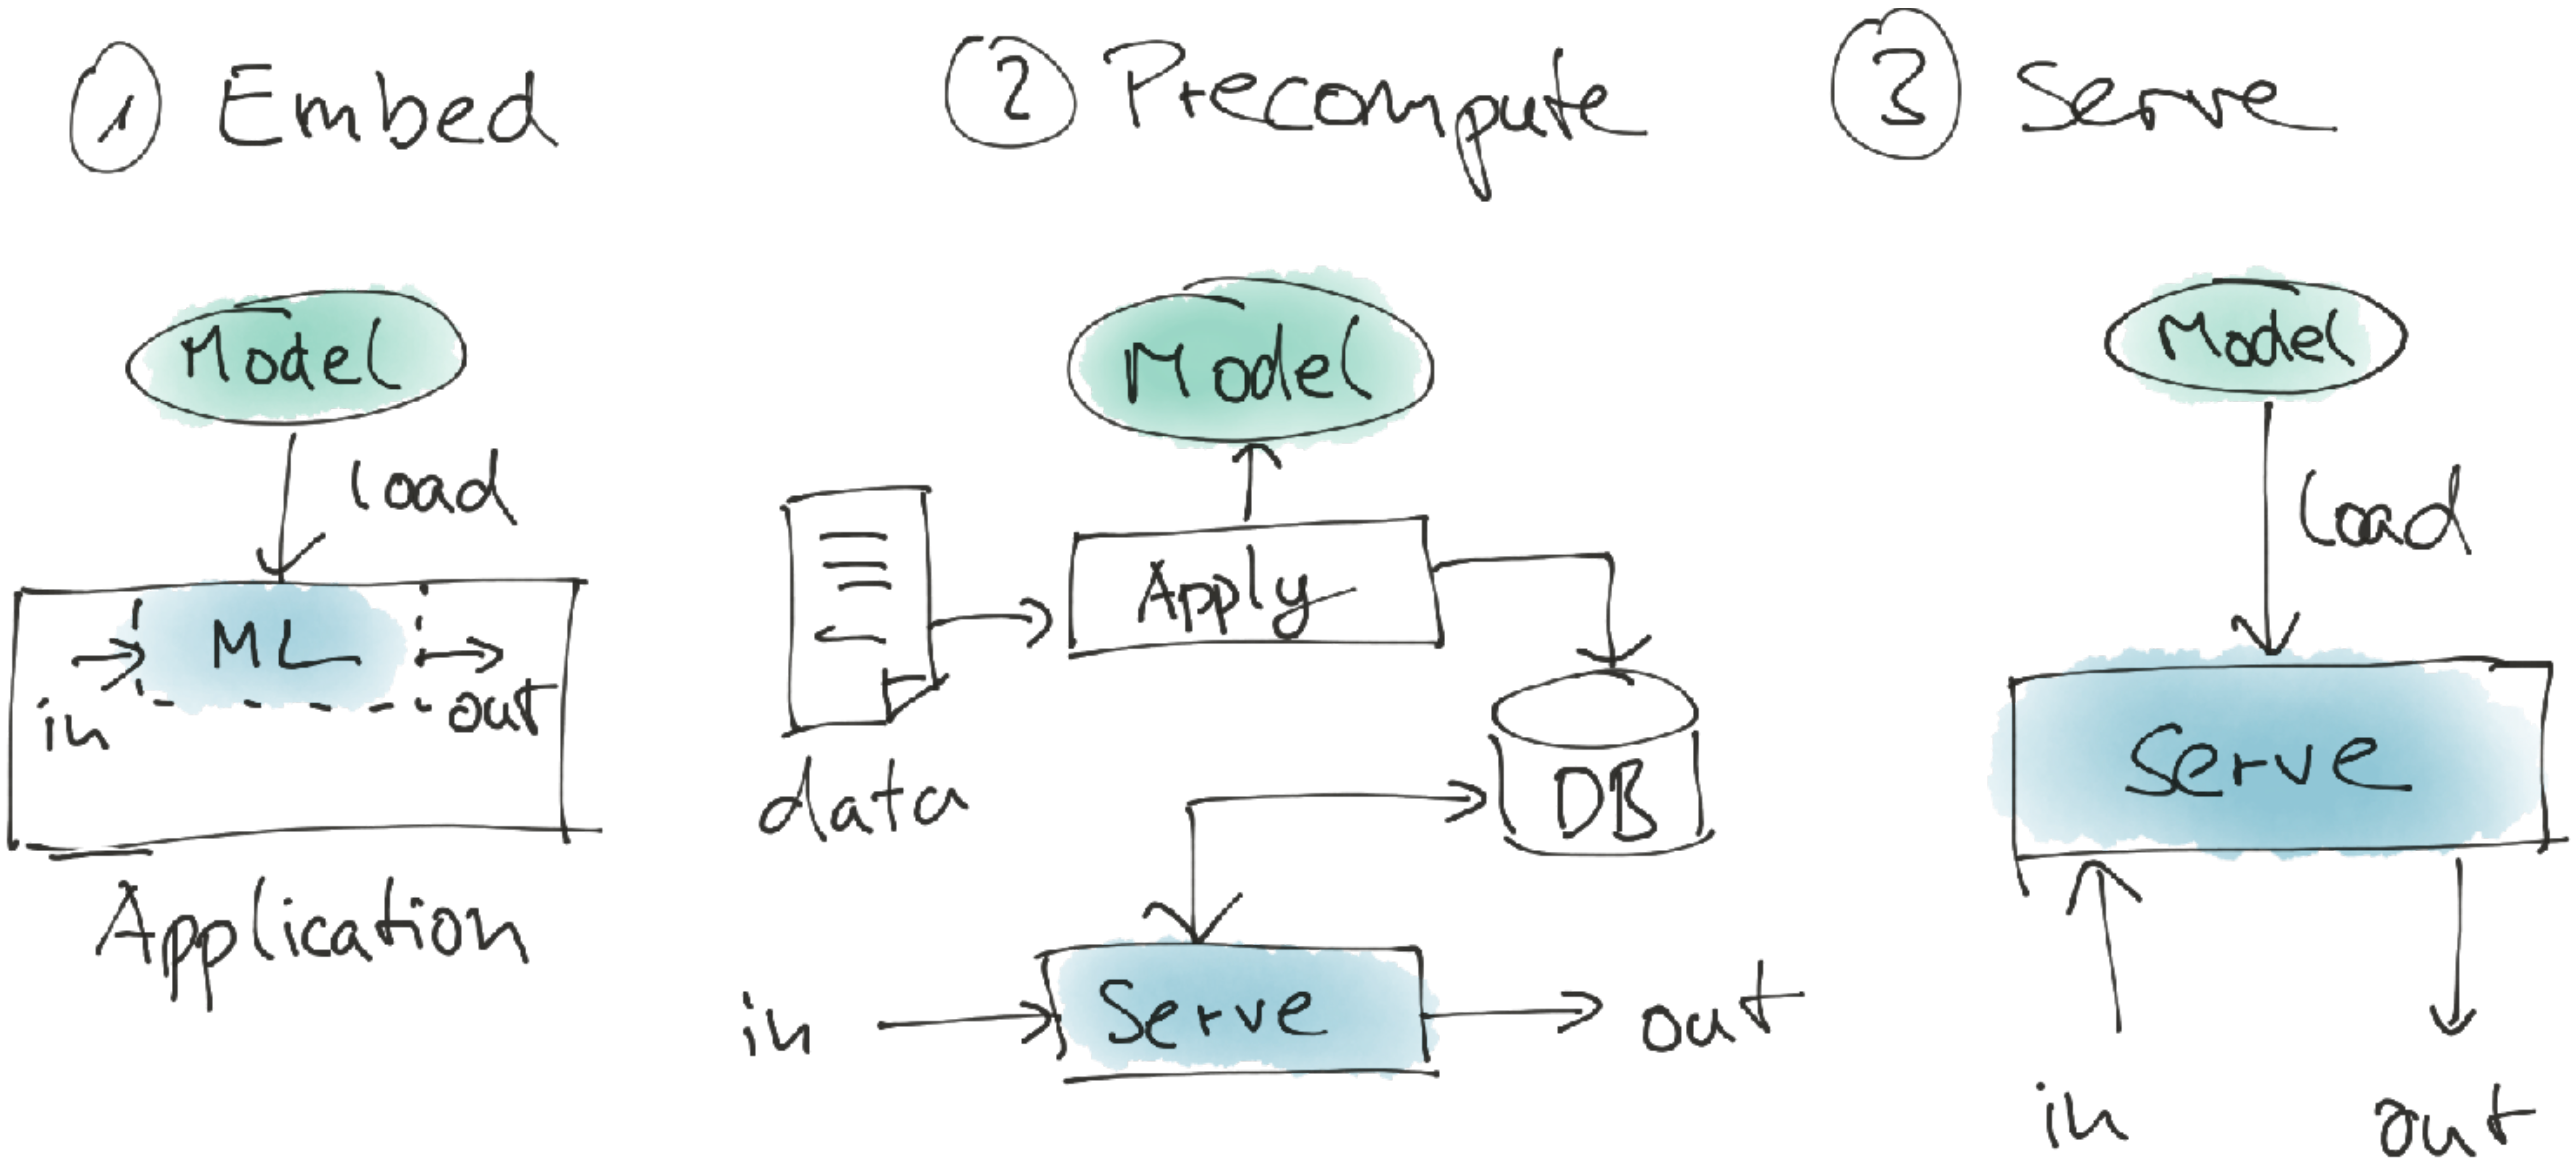

1 Embed
Model
load
ML
in
out
Application
2 Precompute
Model
Apply
data
DB
in
Serve
out
3 Serve
Model
Load
Serve
in
out

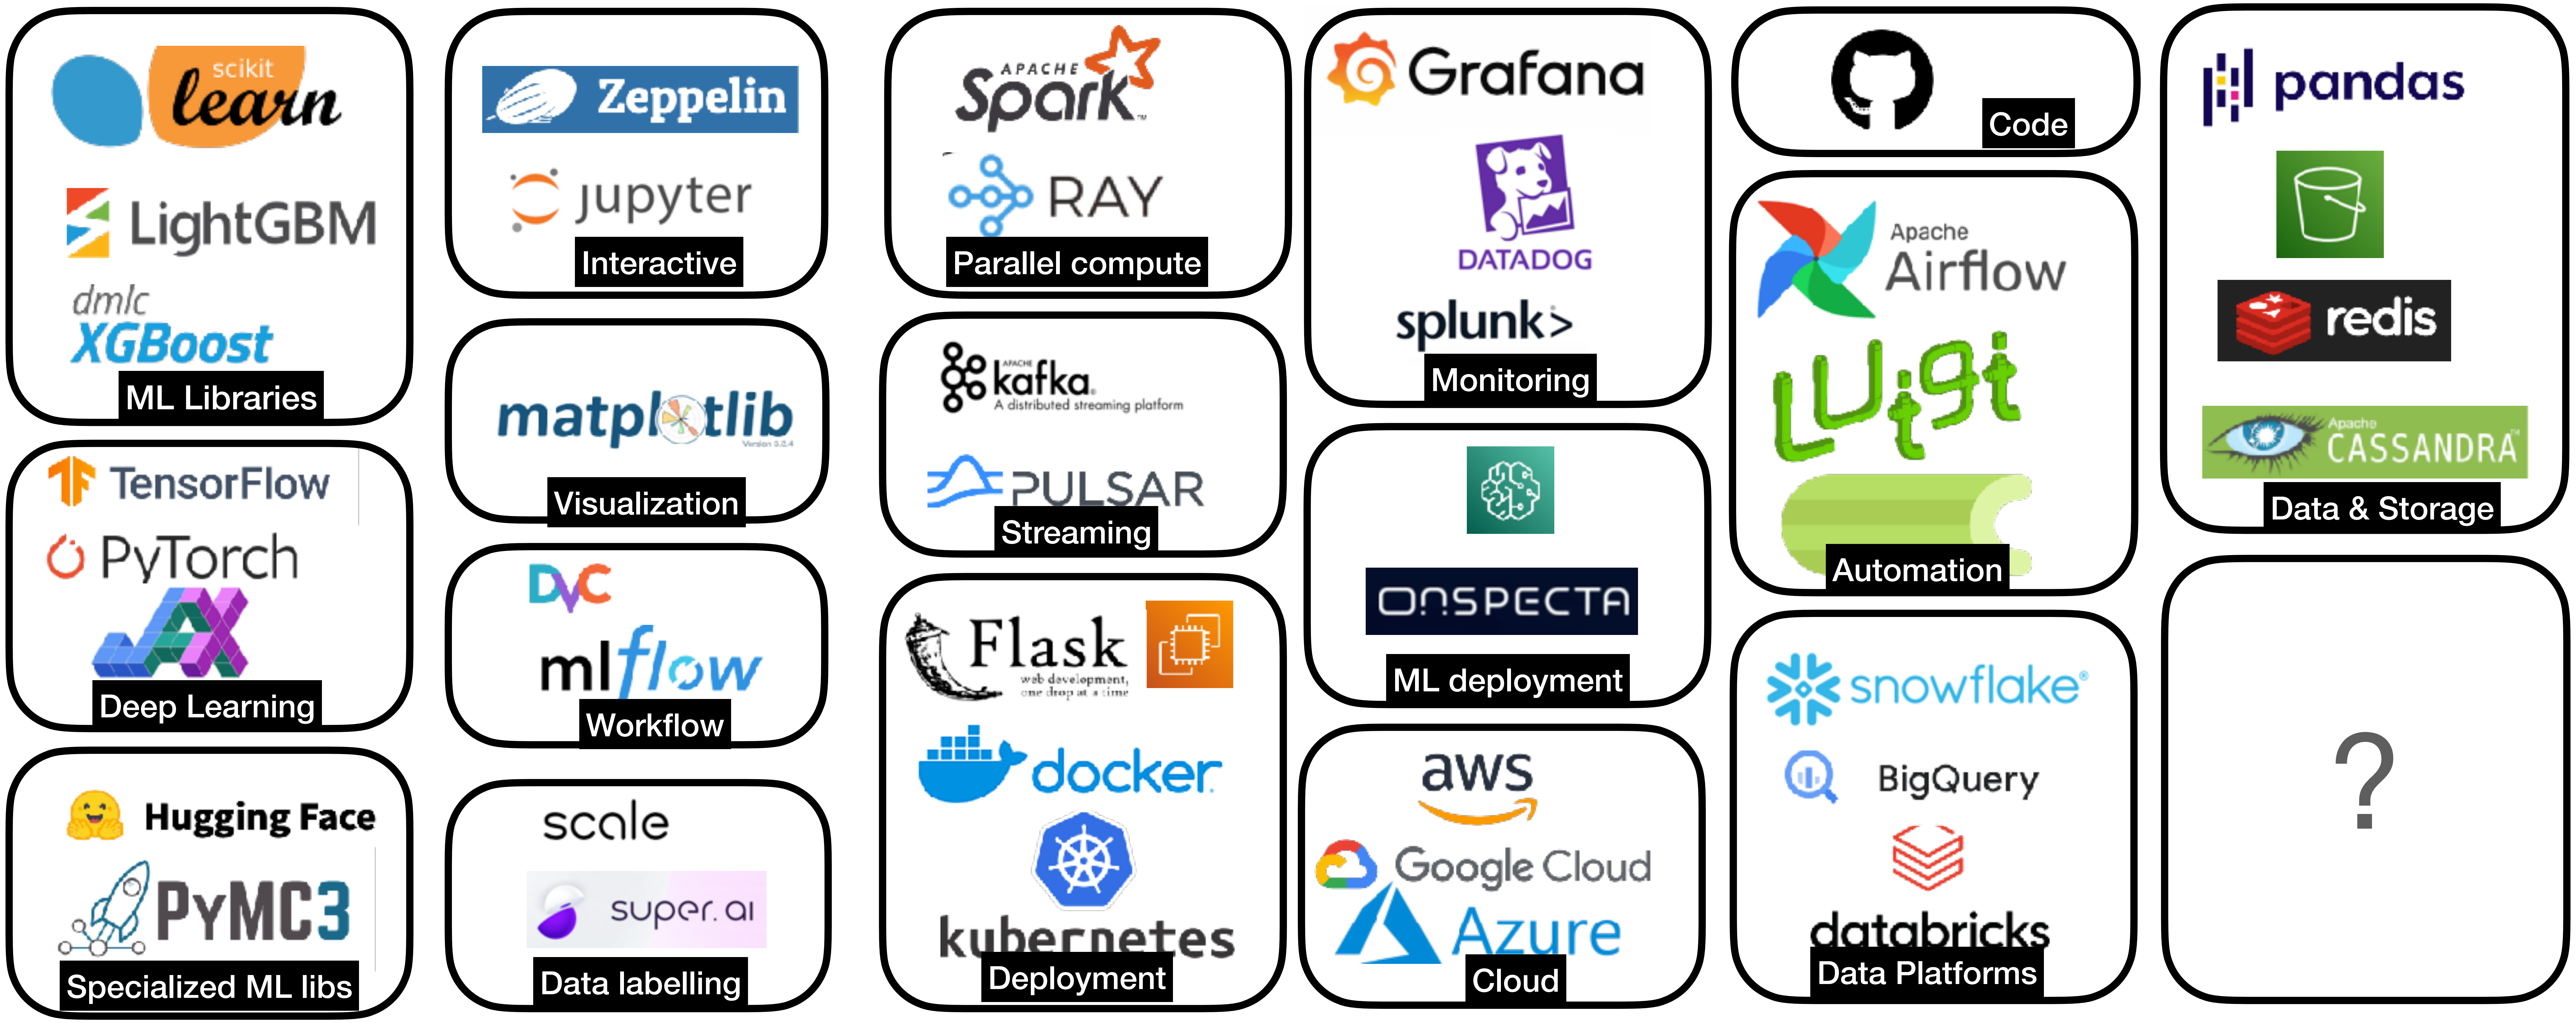
scikit learn
LightGBM
dmlc XGBoost
ML Libraries
TensorFlow
PyTorch
JAX
Deep Learning
Hugging Face
PyMC3
Specialized ML libs
Zeppelin
jupyter
Interactive
matplotlib
Visualization
DVC
mlflow
Workflow
scale
super.ai
Data labelling
APACHE Spark
RAY
Parallel compute
APACHE kafka
A distributed streaming platform
PULSAR
Streaming
Flask
web development, one drop at a time
docker
kubernetes
Deployment
Grafana
DATADOG
splunk>
Monitoring
ONSPECTA
ML deployment
aws
Google Cloud
Azure
Cloud
Code
Apache Airflow
Luigi
Automation
snowflake
BigQuery
databricks
Data Platforms
pandas
redis
Apache CASSANDRA
Data & Storage
?

Run data science projects and teams

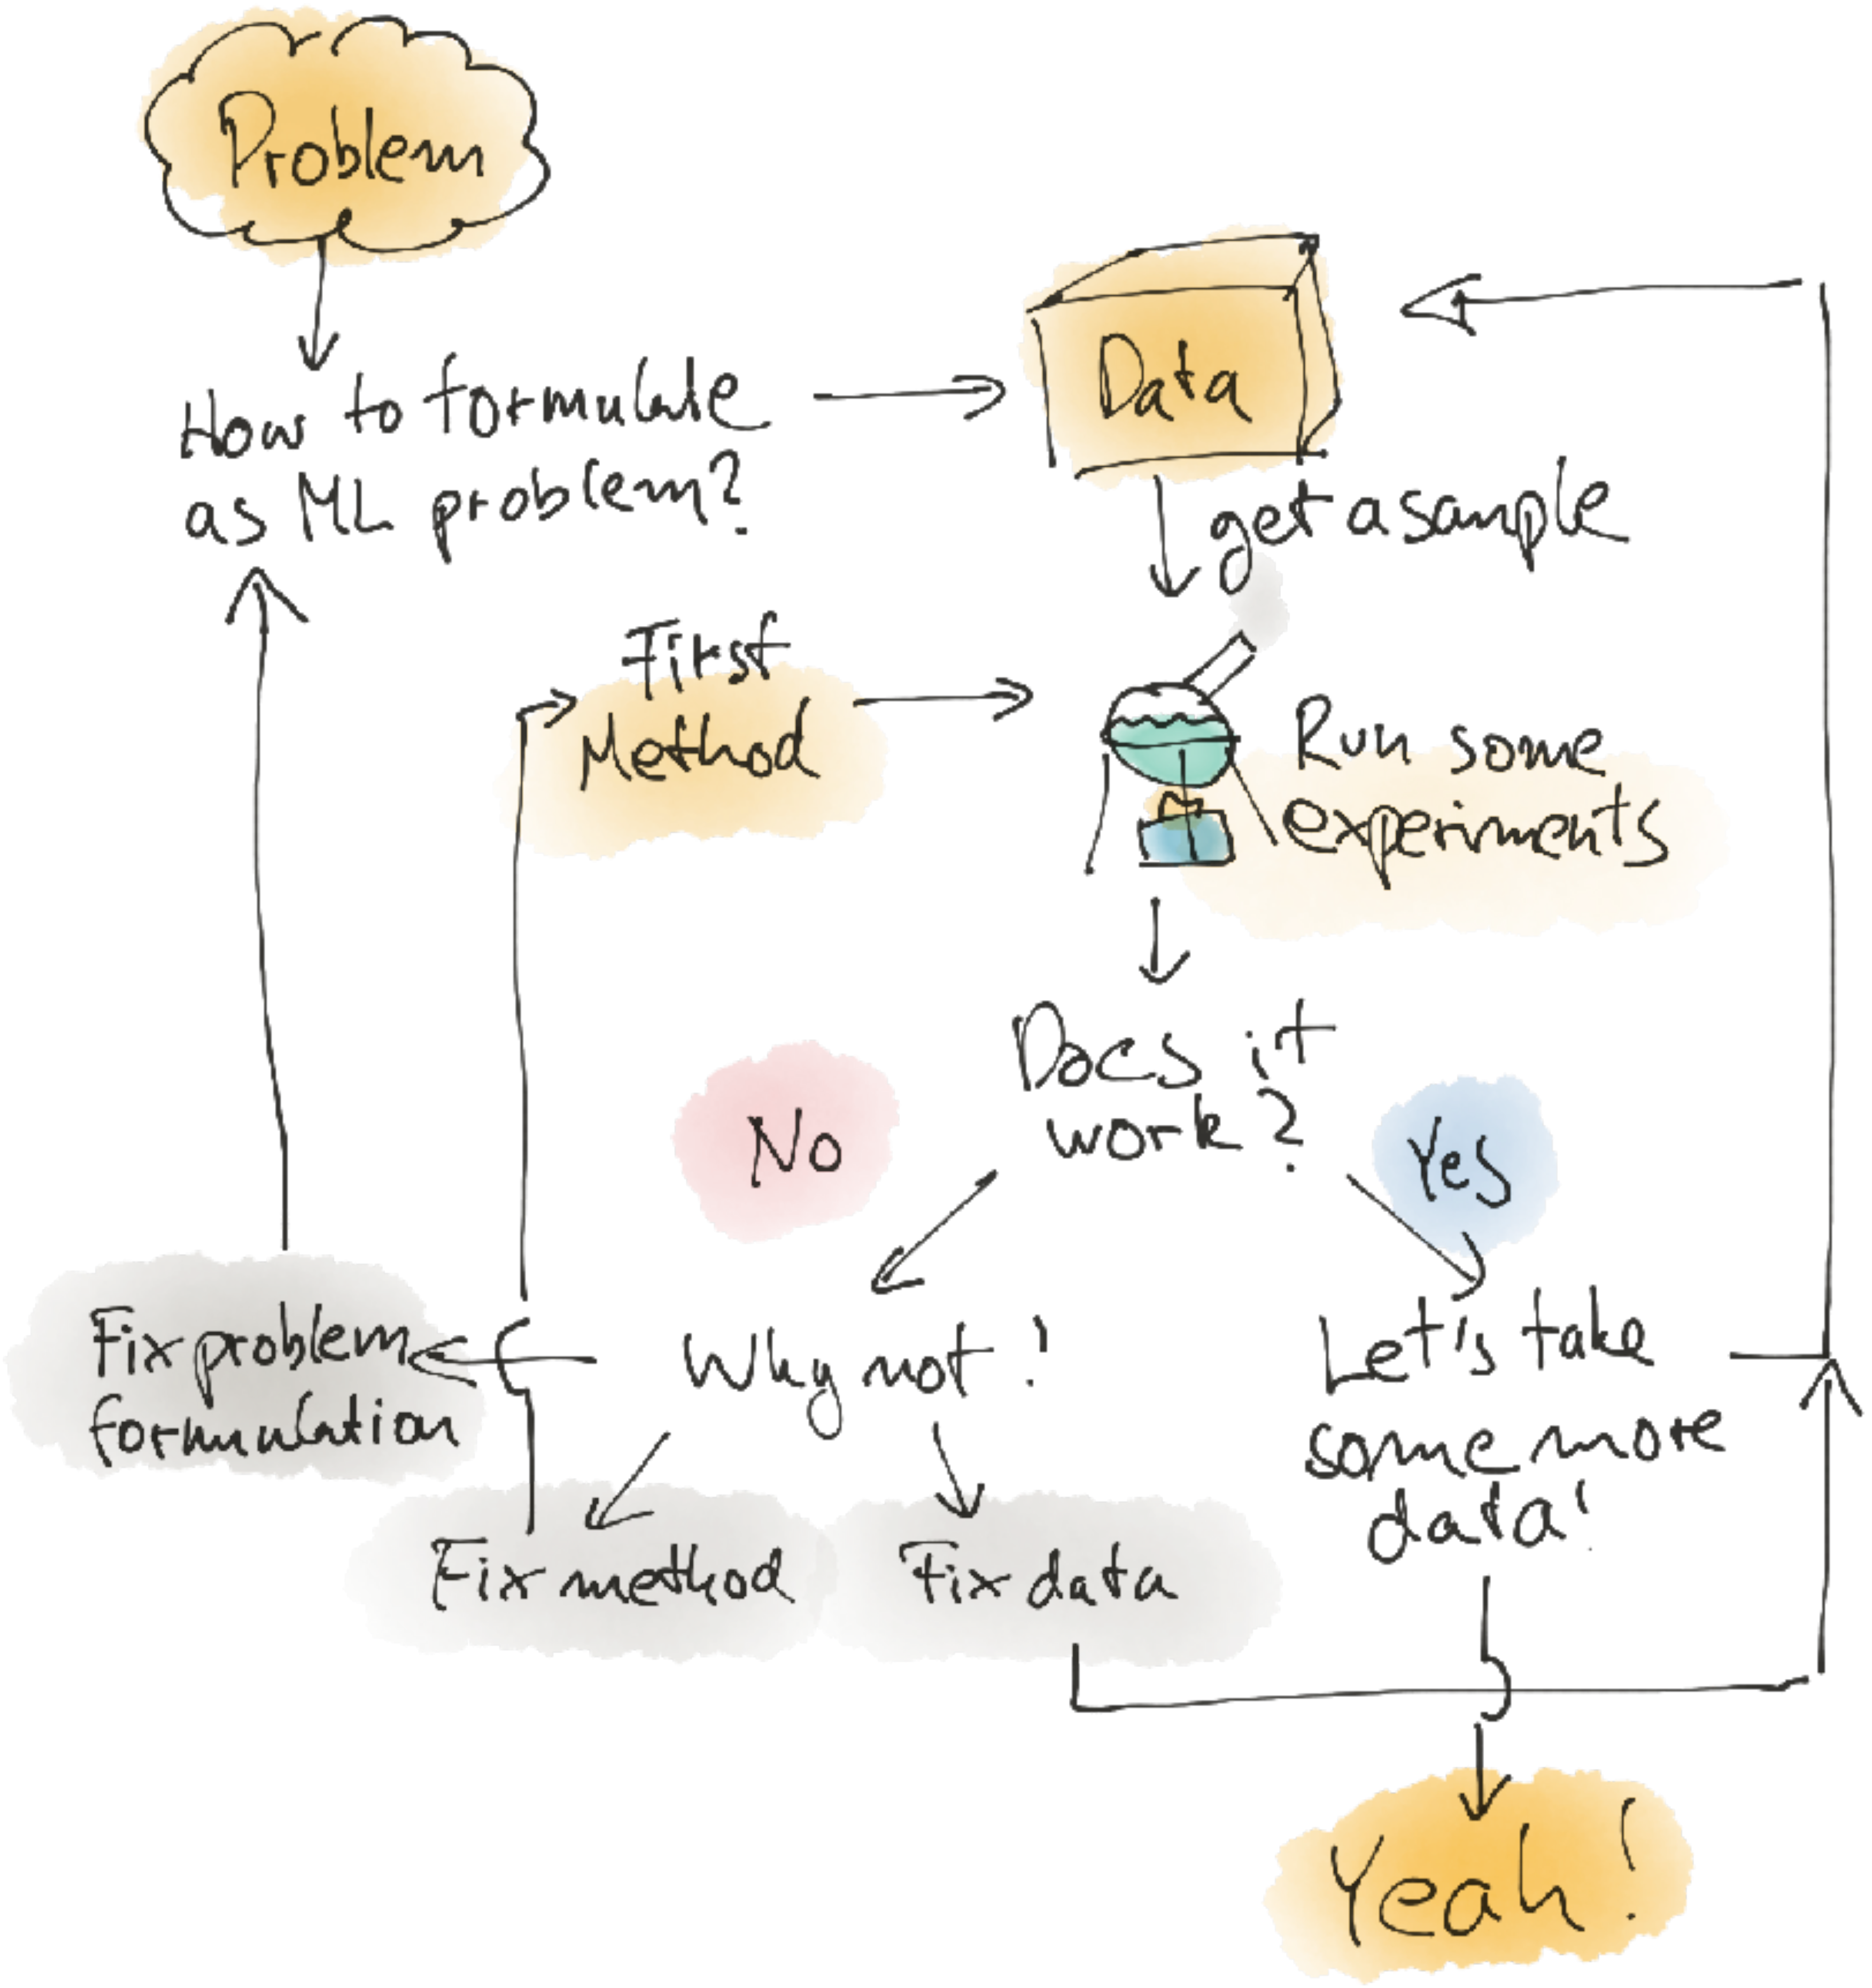
Problem
How to formulate as ML problem?
Data
get a sample
First Method
Run some experiments
Does it work?
No
Yes
Why not!
Fix problem formulation
Fix method
Fix data
Let's take some more data!
Yeah!

Minimize risk & time box

https://book.mlinpractice.com/people-and-processes/data-science-projects

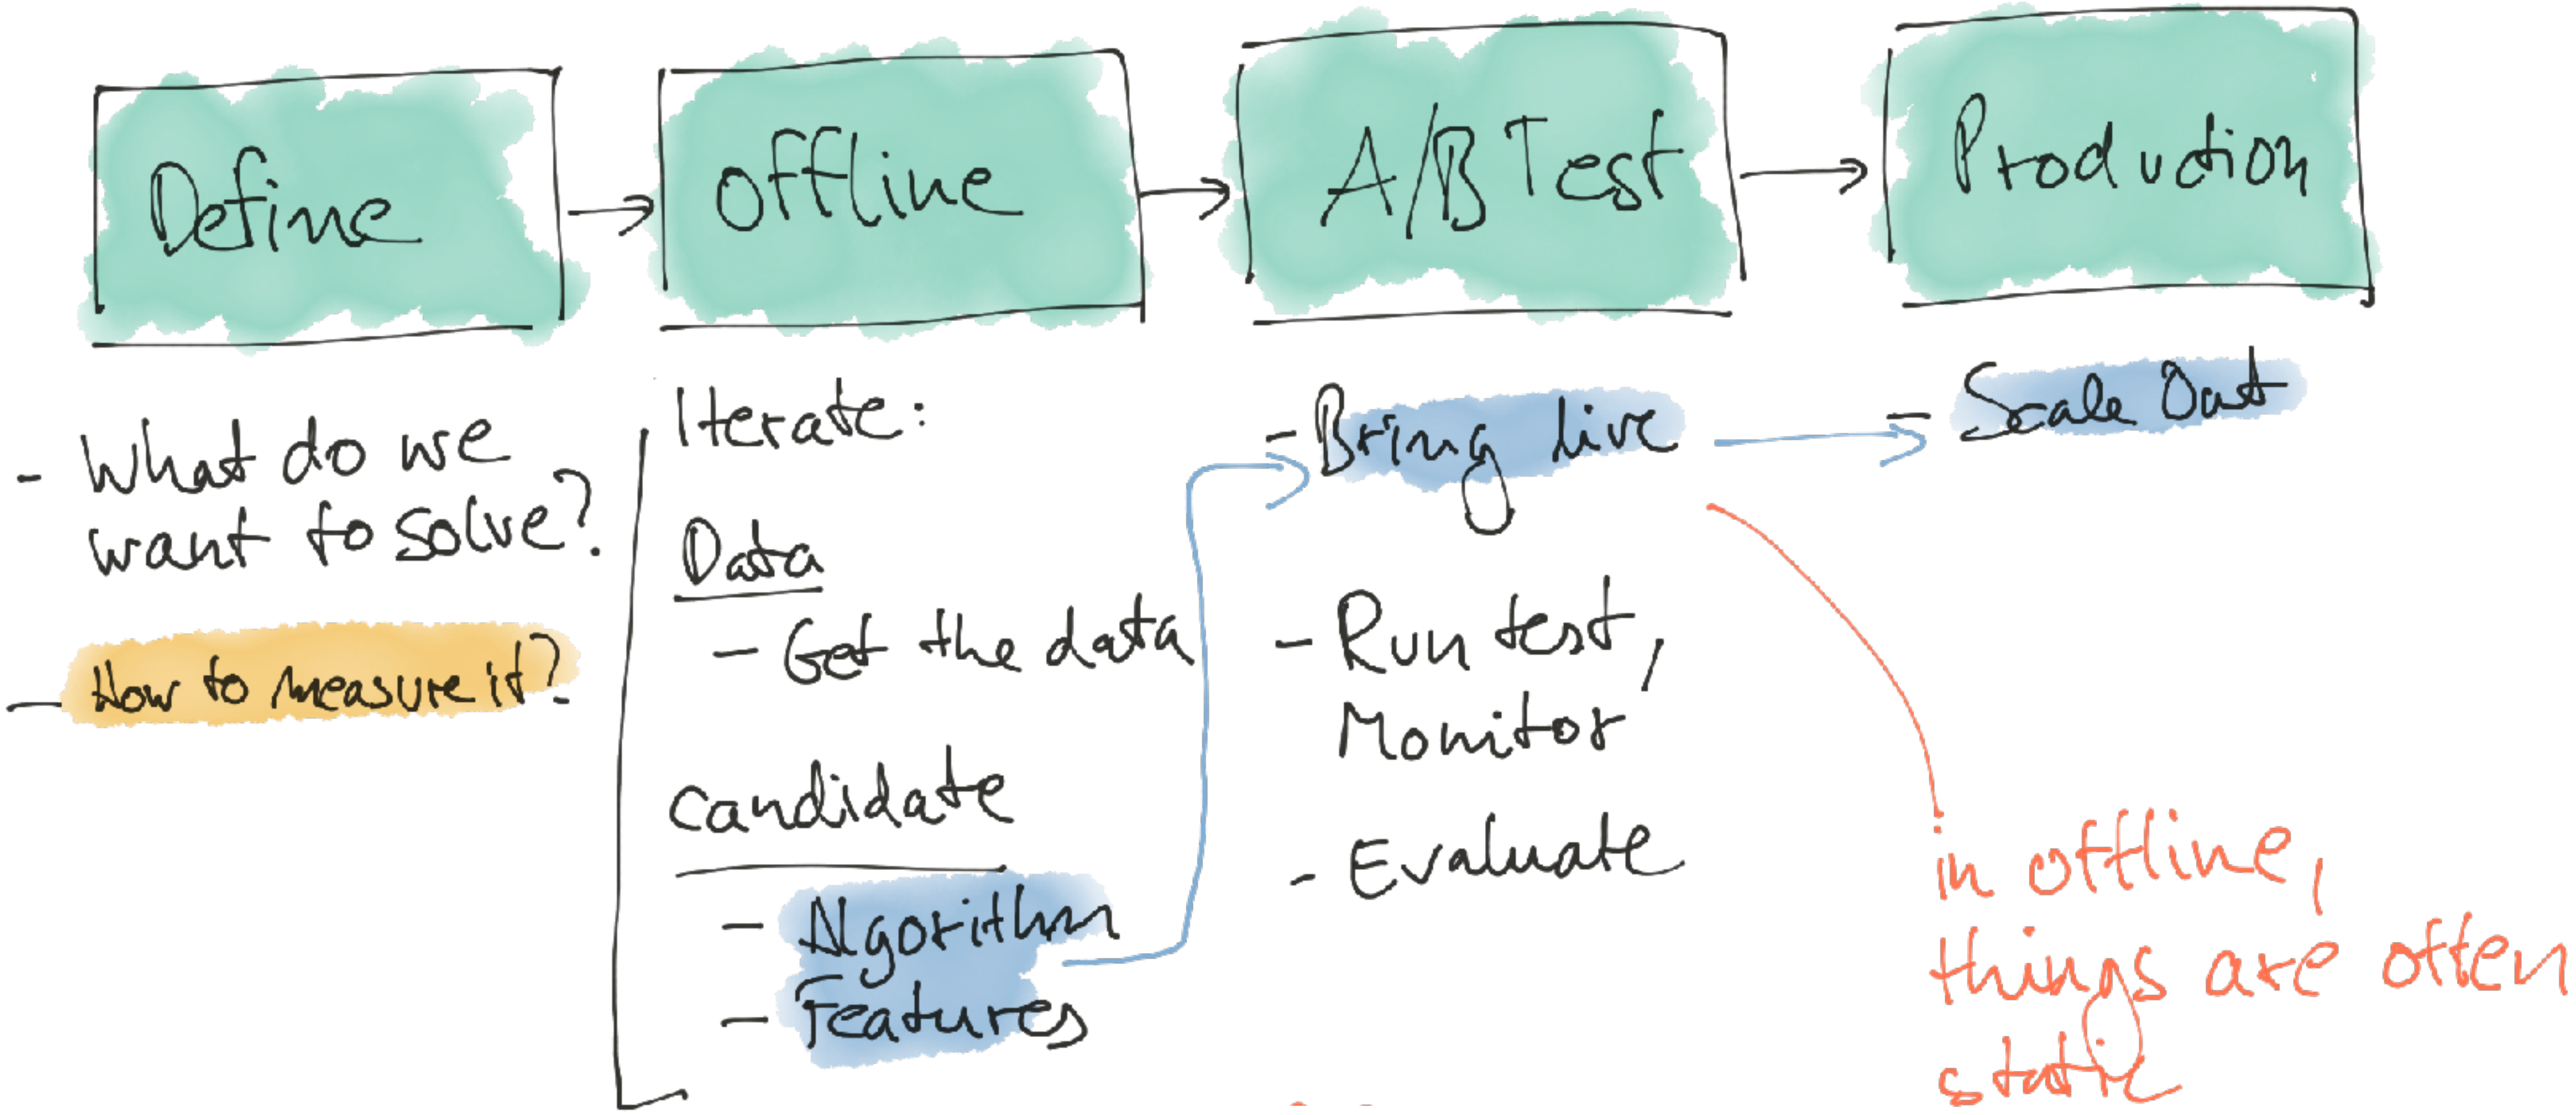
Define
Offline
A/B Test
Production
- What do we want to solve?
- How to measure it?
Iterate:
Data
- Get the data
Candidate
- Algorithm
- Features
- Bring live
- Run test, Monitor
- Evaluate
- Scale Out
in offline, things are often static

Scale data science

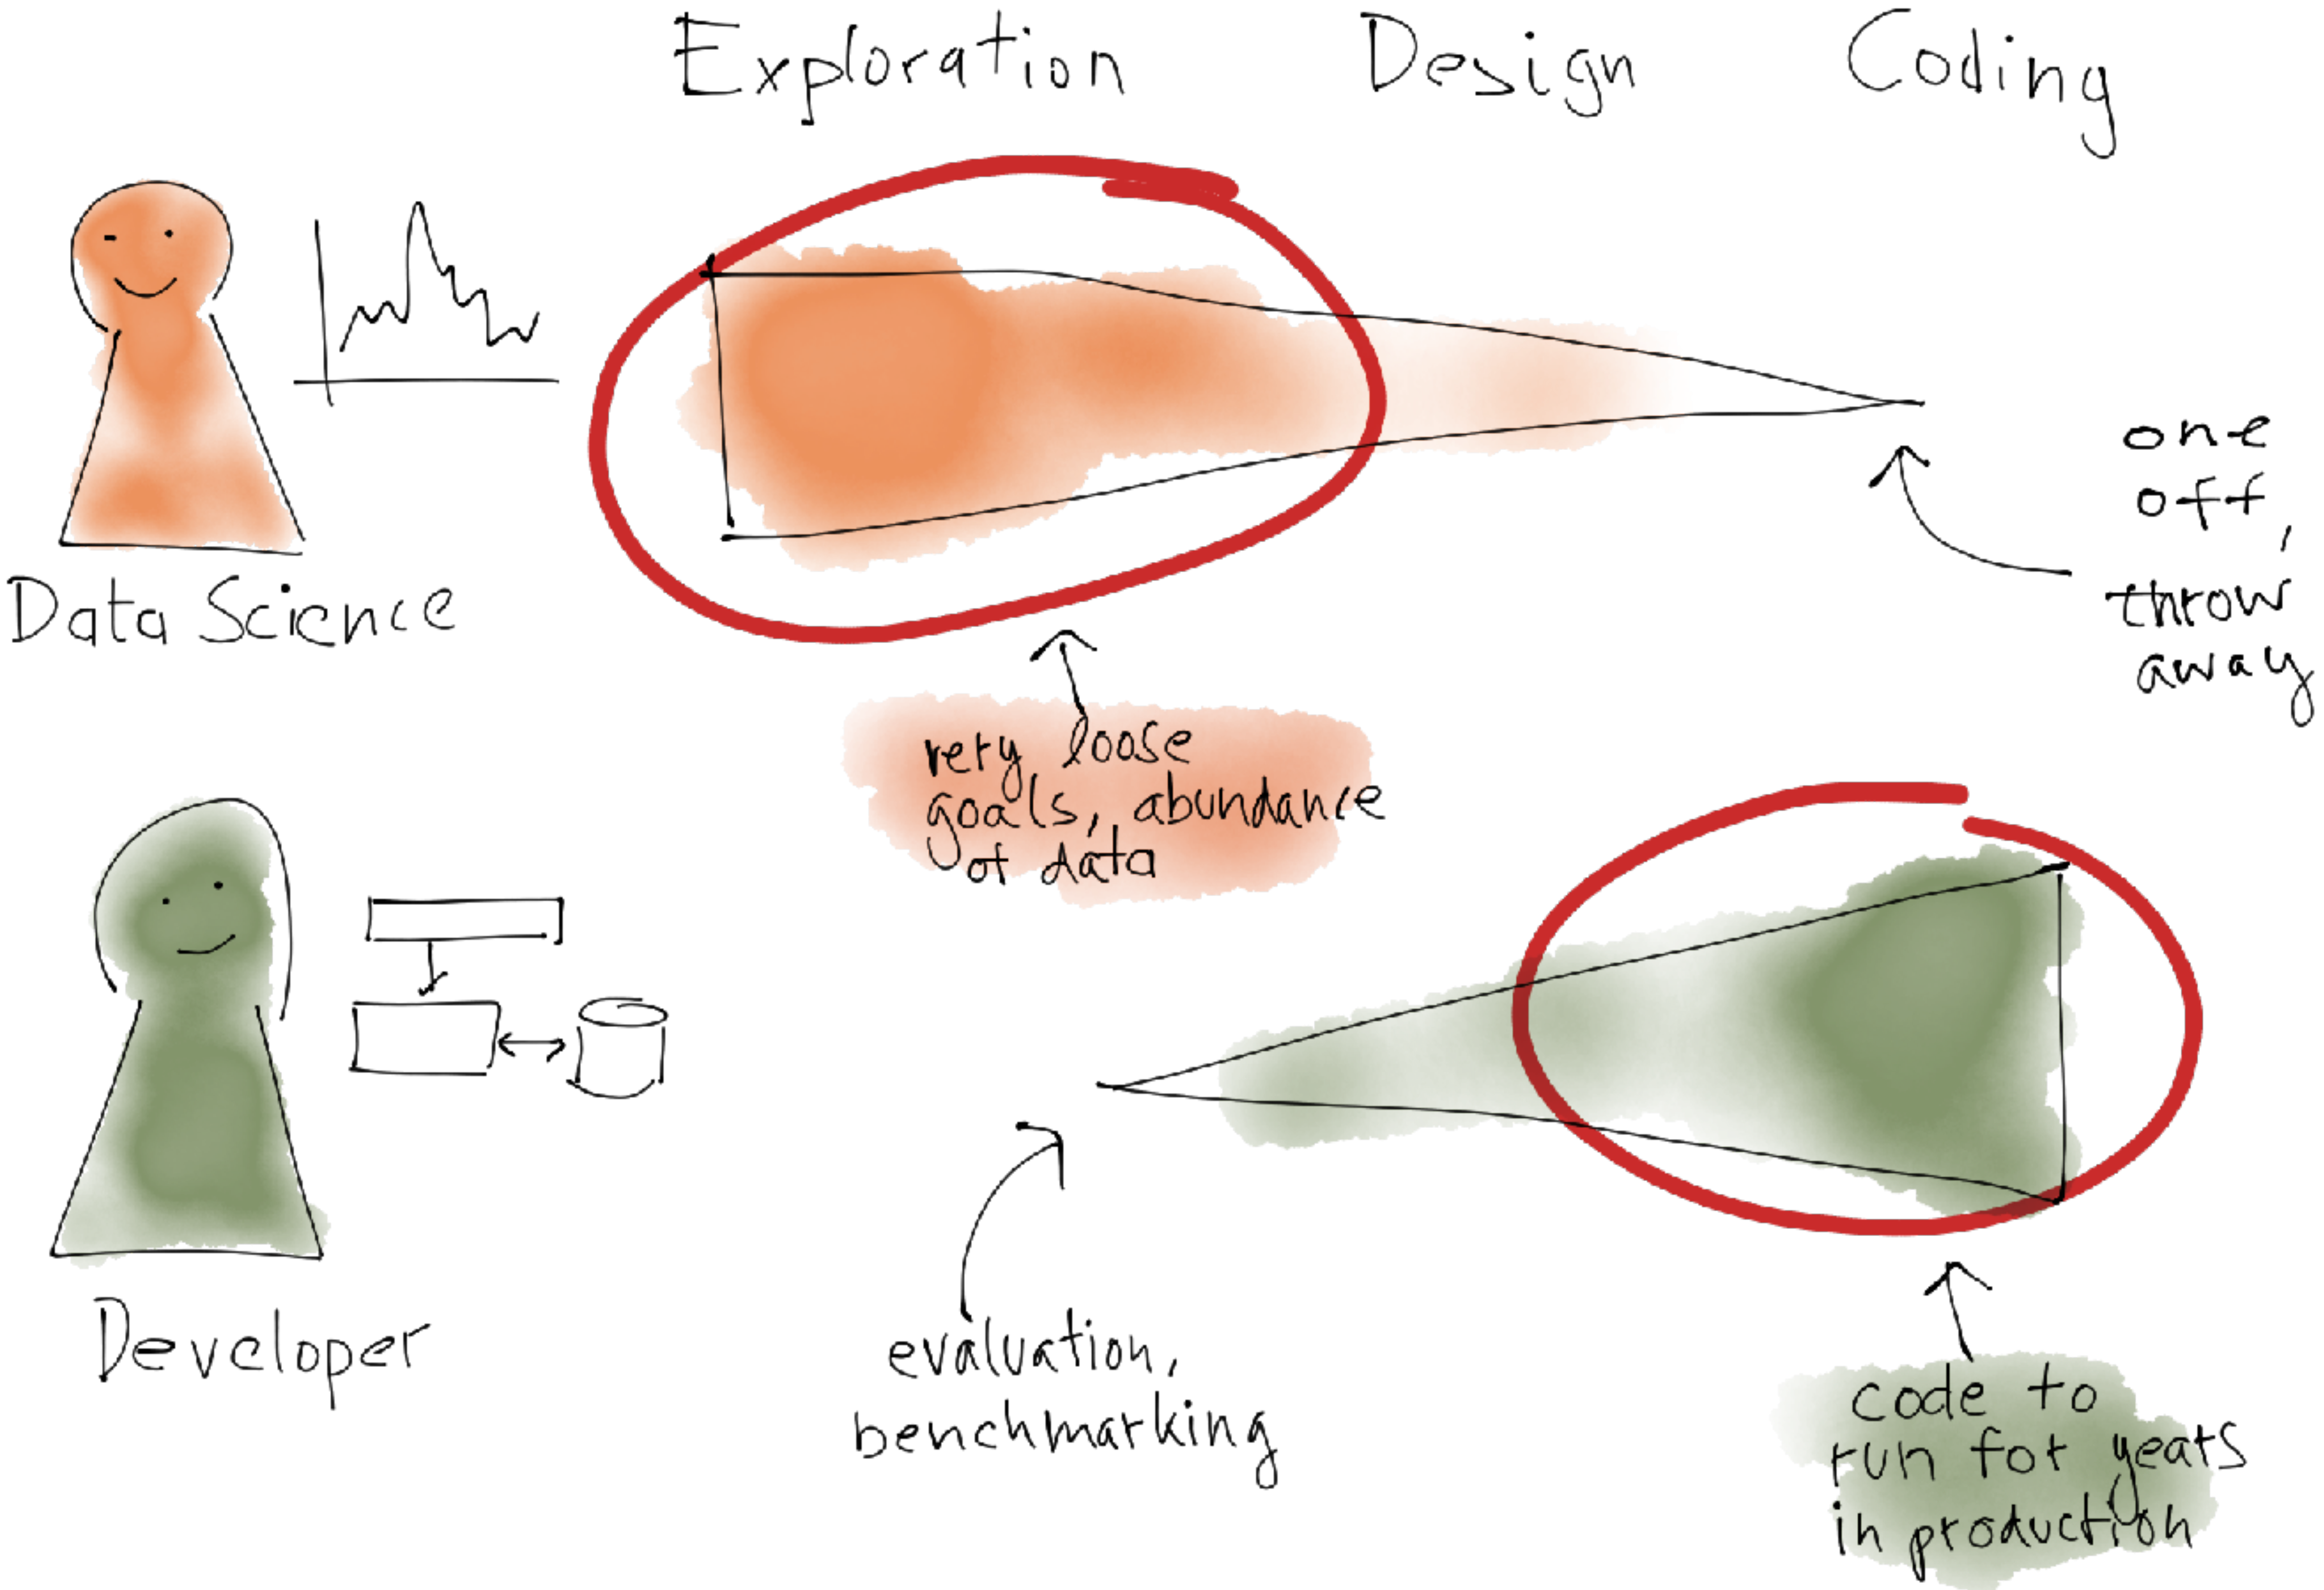
Exploration
Design
Coding
Data Science
one off, throw away
very loose goals, abundance of data
Developer
evaluation, benchmarking
code to run for years in production

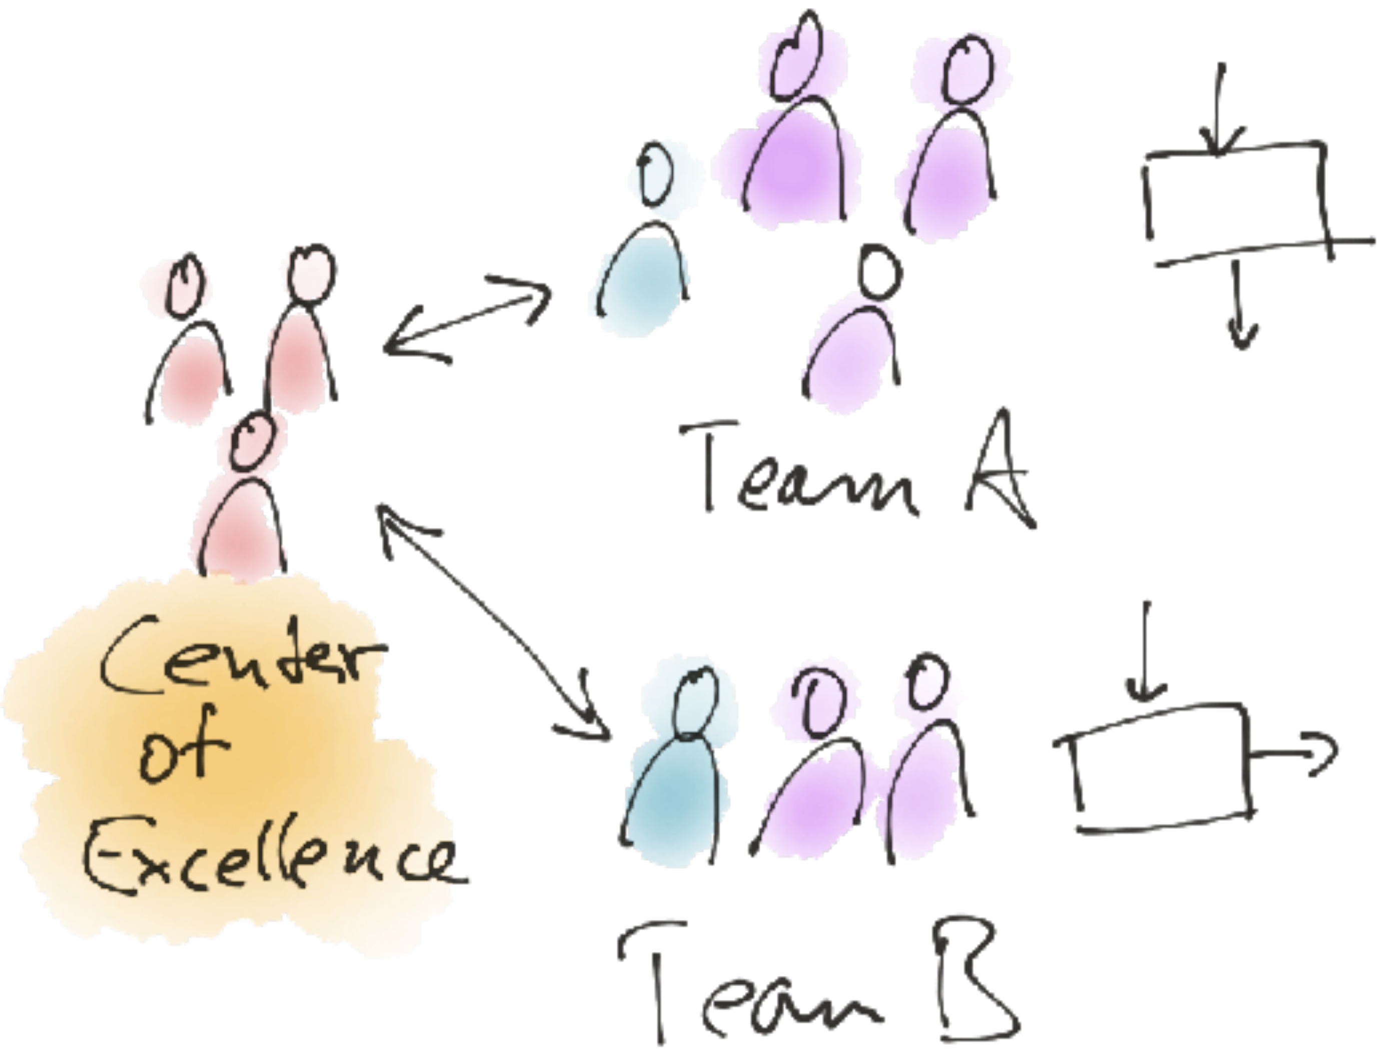

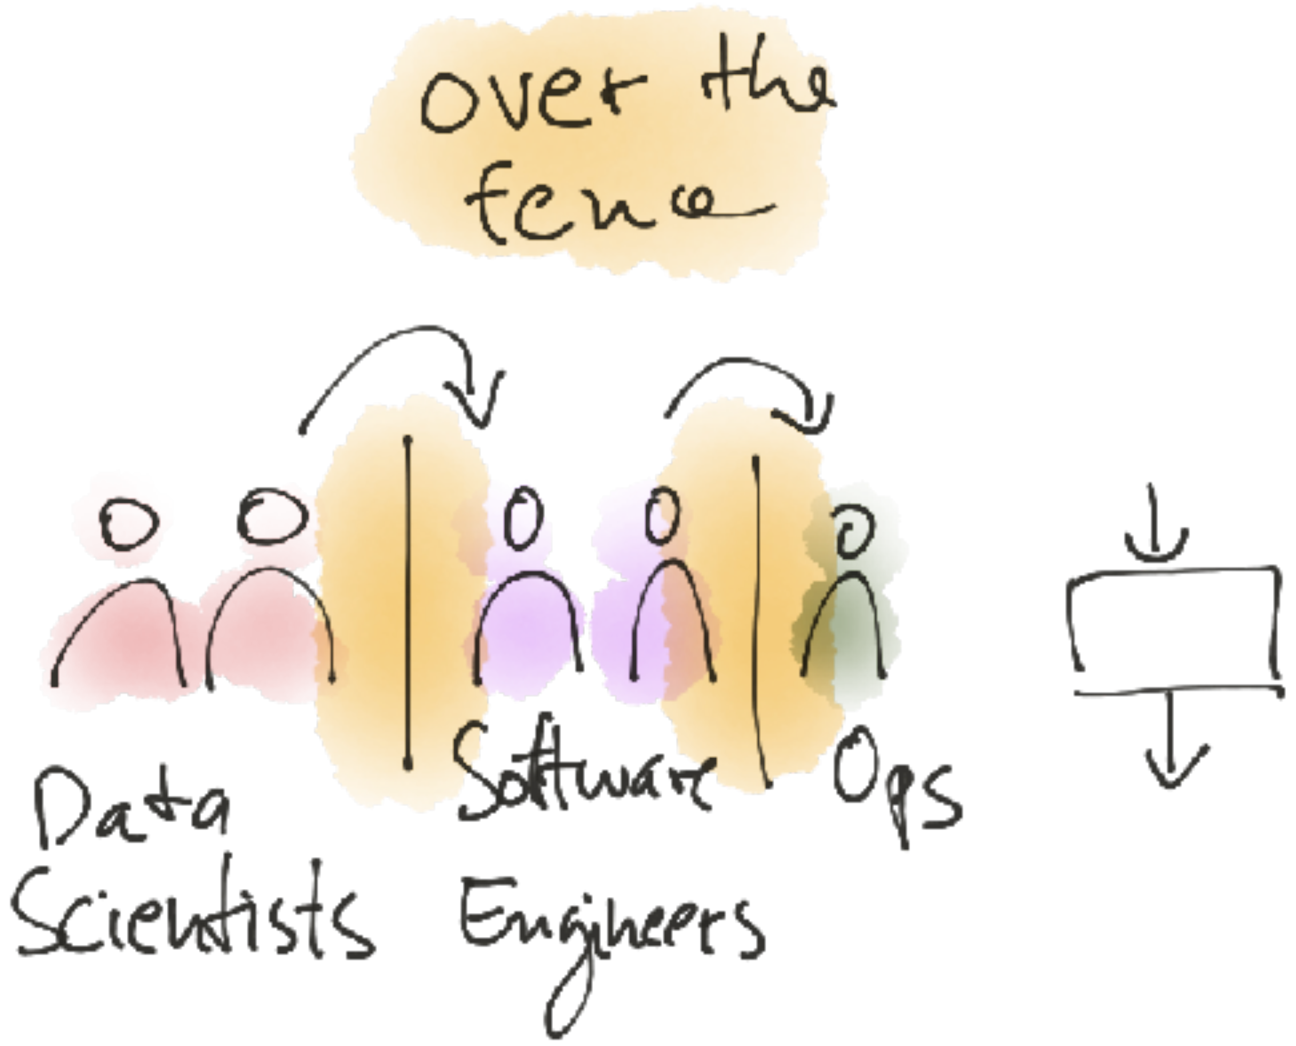

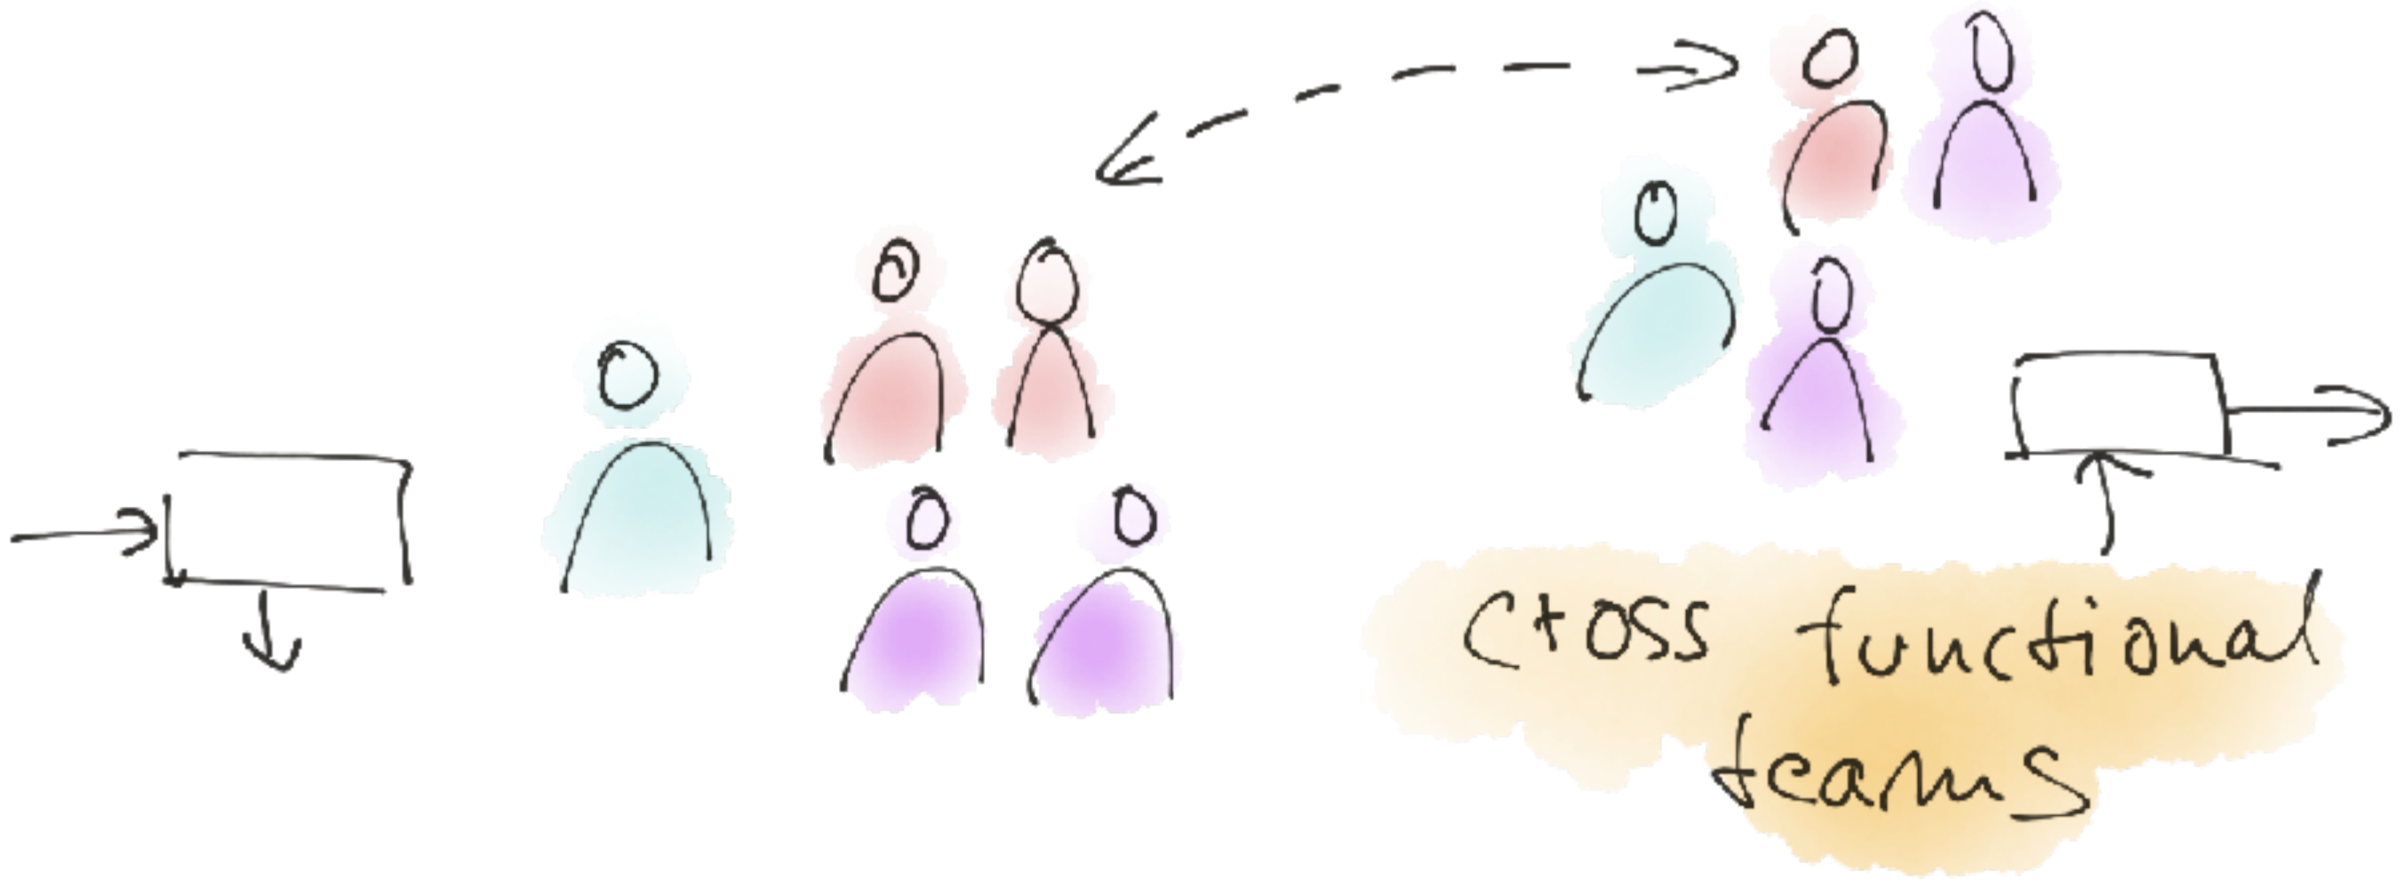

Dr. Mikio Braun

Machine Learning:
The High-Interest Credit Card of Technical Debt

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov,
Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young
{dsculley,gholt,dgg,edavydov}@google.com
{toddphillips,ebner,vchaudhary,mwyoung}@google.com
Google, Inc

Abstract

Machine learning offers a fantastically powerful toolkit for building complex systems quickly. This paper argues that it is dangerous to think of these quick wins as coming for free. Using the framework of *technical debt*, we note that it is remarkably easy to incur massive ongoing maintenance costs at the system level

https://research.google/pubs/pub43146/

What's your ML Test Score? A rubric for ML production systems

Eric Breck, Shanqing Cai, Eric Nielsen, Michael Salib, D. Sculley
Google, Inc.
{ebreck, cais, nielsene, msalib, dsculley}@google.com

Abstract

Using machine learning in real-world production systems is complicated by a host of issues not found in small toy examples or even large offline research experiments. Testing and monitoring are key considerations for assessing the production-readiness of an ML system. But how much testing and monitoring is enough? We present an ML Test Score rubric based on a set of actionable tests to help quantify these issues.

https://research.google/pubs/pub45742/

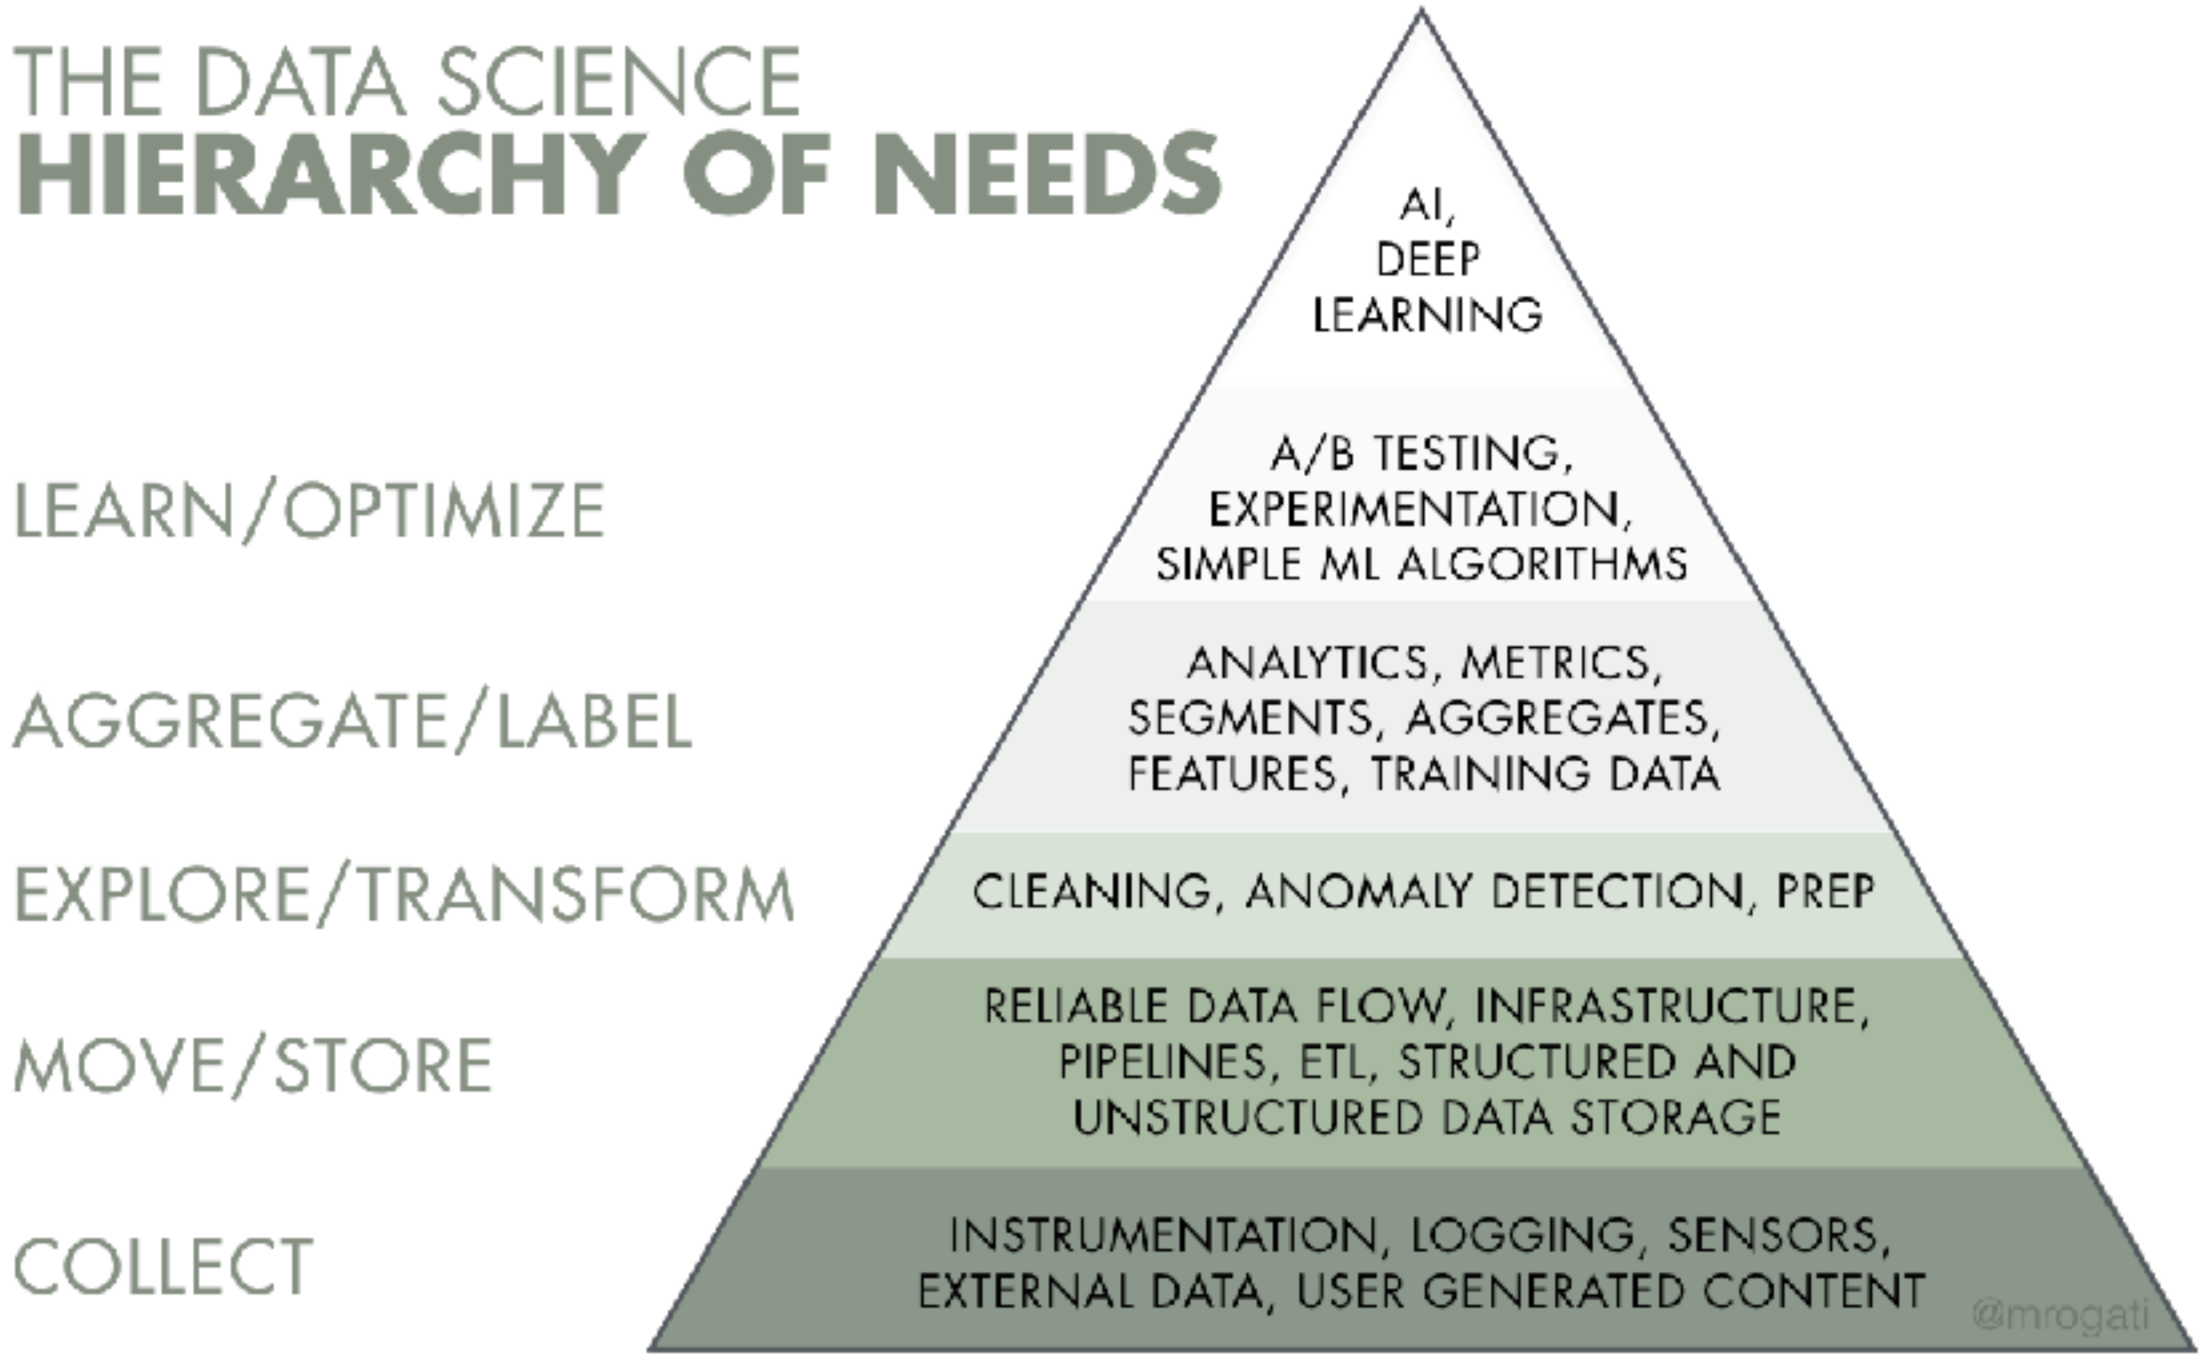

https://hackernoon.com/the-ai-hierarchy-of-needs-18f111fcc007

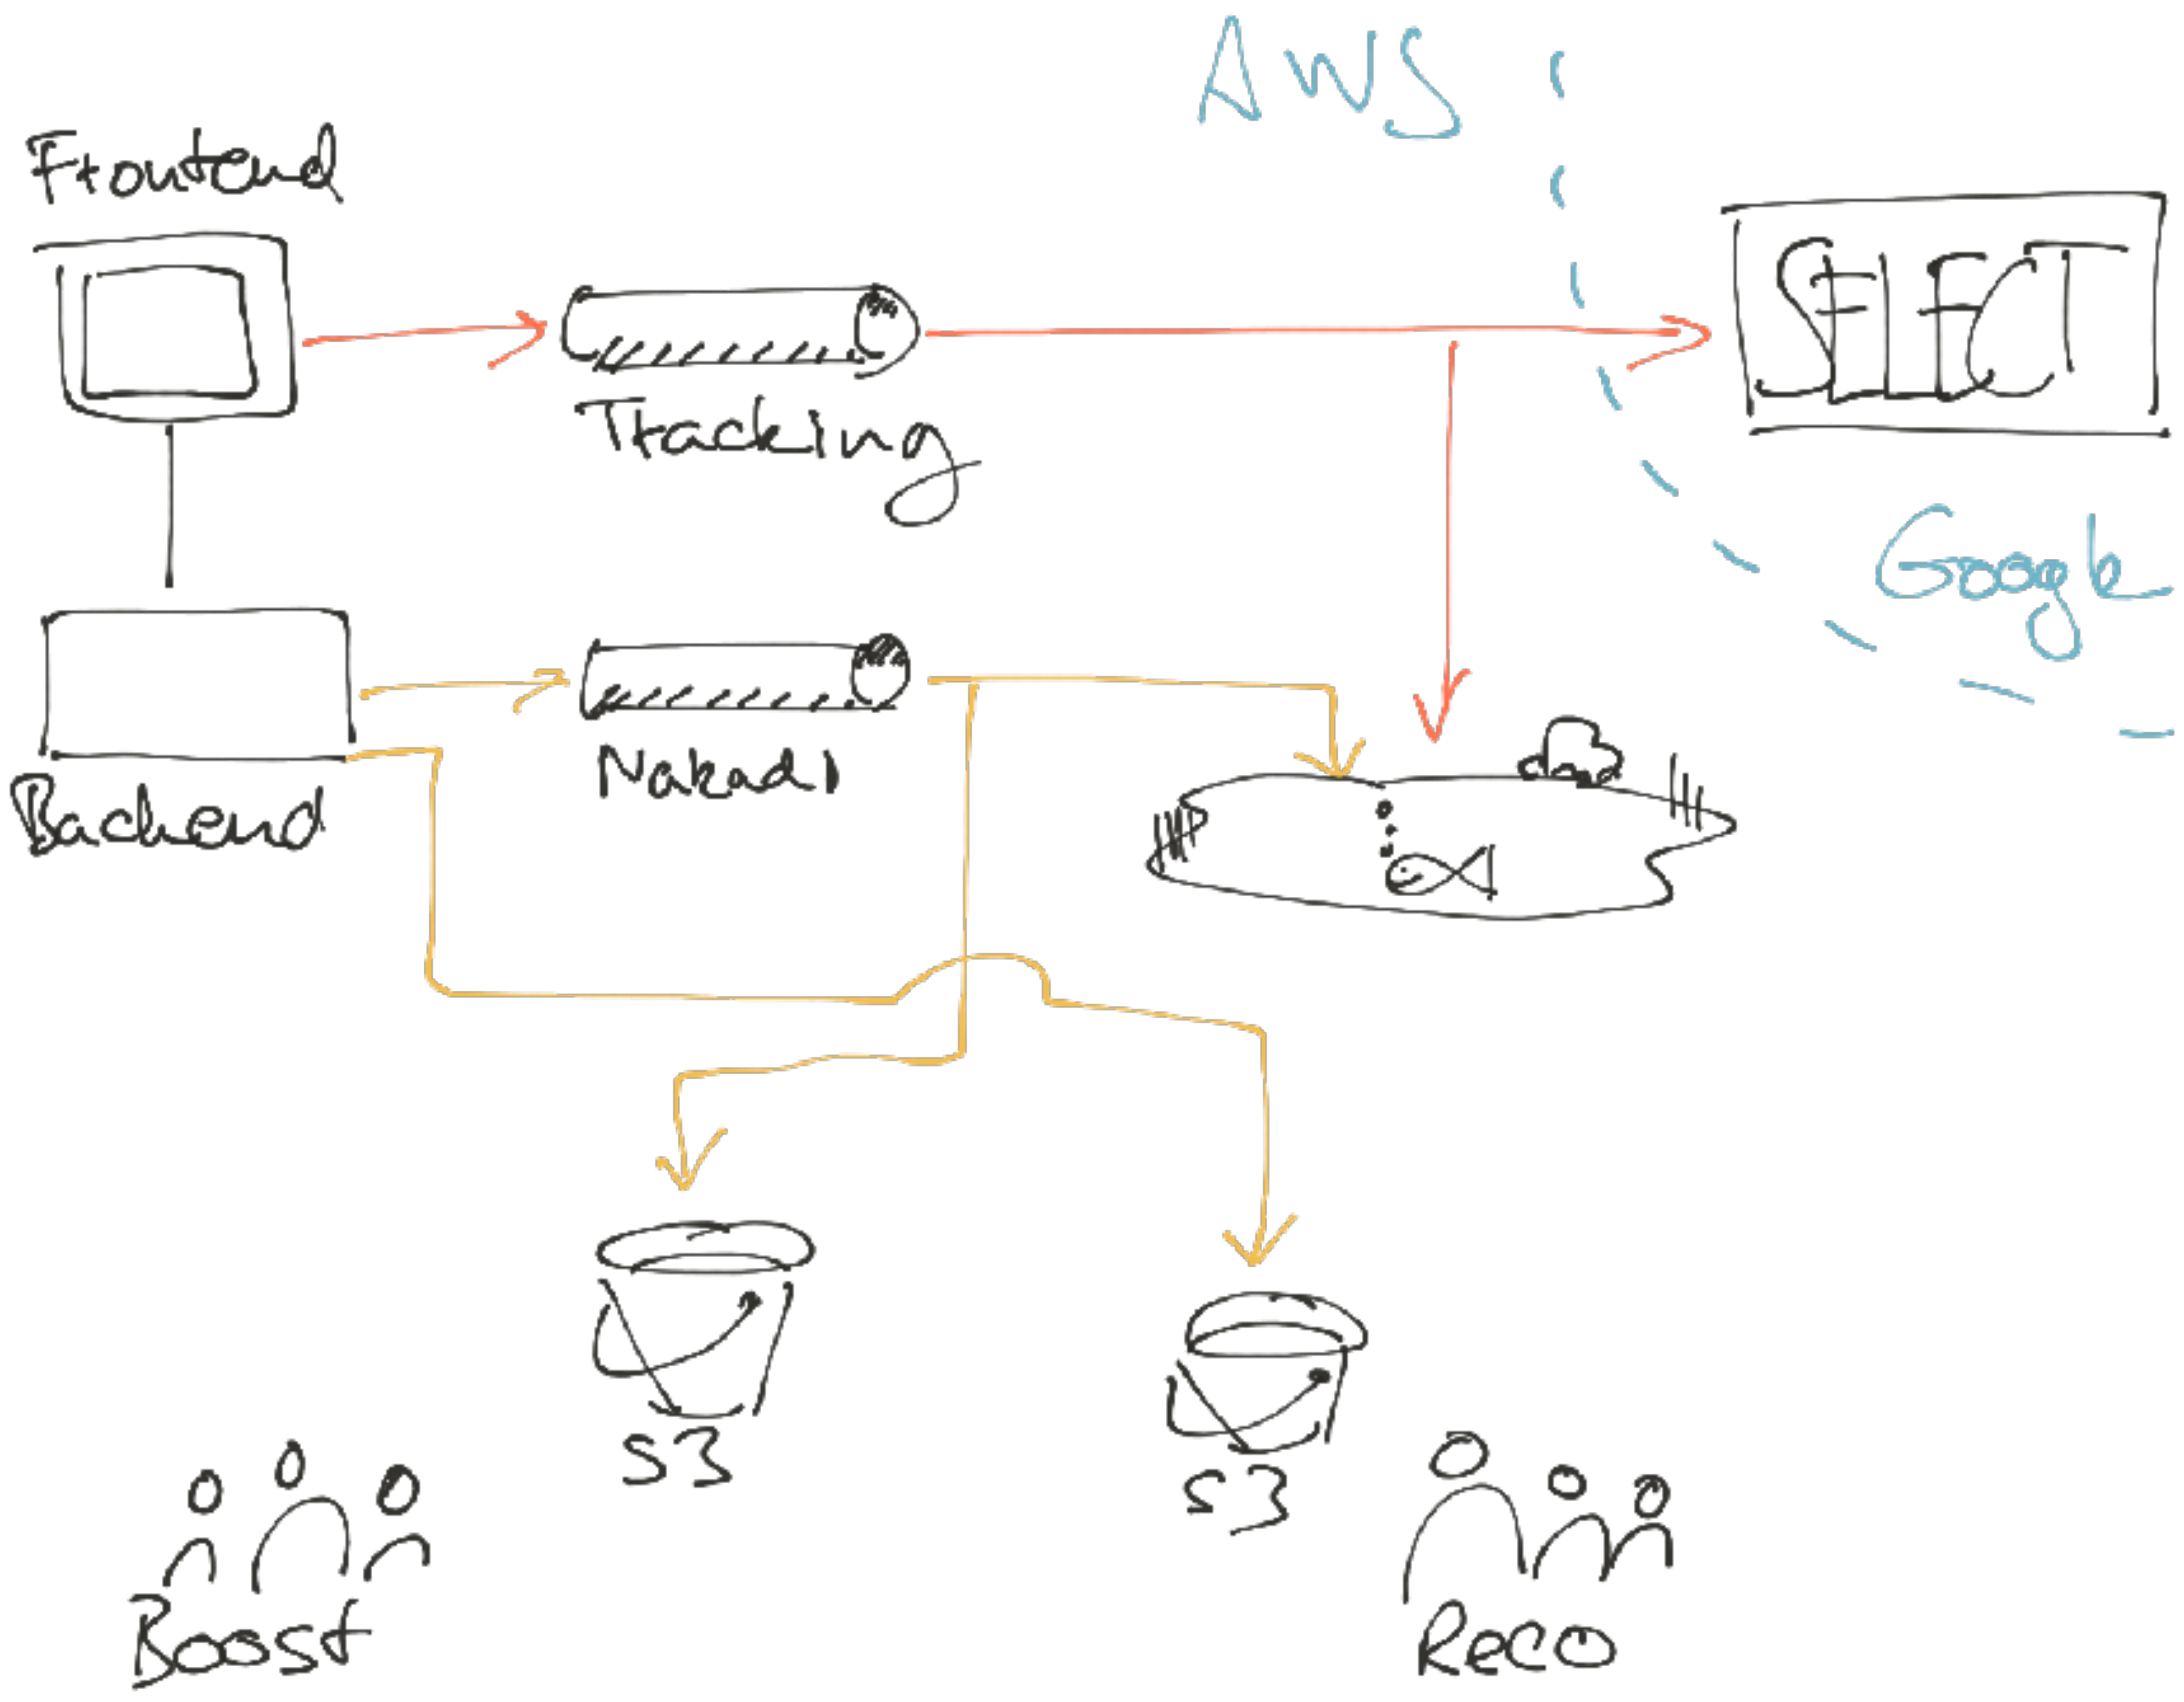

Thank you!