Agenda

- 13:00: Welcome: Anne Katrine Windfeld, Grundfos
- 13:10: AI: from experimentation to delivering value: Jan Bosch,
  Chalmers University of Technology & Software Center
- 13.30 Trustworthiness in testing and continuous integration:
  Kristian Sandahl, Linköping University
- 13.50: Presentation of selected Software Center projects:
  Helena Holmström Olsson, Malmö University
  – Accelerating Digitalization Through Data
  – Strategic Ecosystem-Driven R&D Management
- 14.15: Interactive workshop on digitalization
- 15:45: Closing: Jan Bosch

Artificial Intelligence
from experimentation to delivering value

Jan Bosch

Director Software Center
www.software-center.se
Professor of Software Engineering
Chalmers University of Technology
Gothenburg, Sweden.

Overview

- Vem är jag? Wie ben ik? Who am I?
- Reflecting on AI
- AI Engineering
- Example research activities
- Conclusion

Academic Research

10X Software Center

Consultancy

Boards

Angel investing

Industry

Software Center

Mission: To significantly improve the ***digitalization*** capability of the European Software-Intensive industry

RESEARCH THEMES

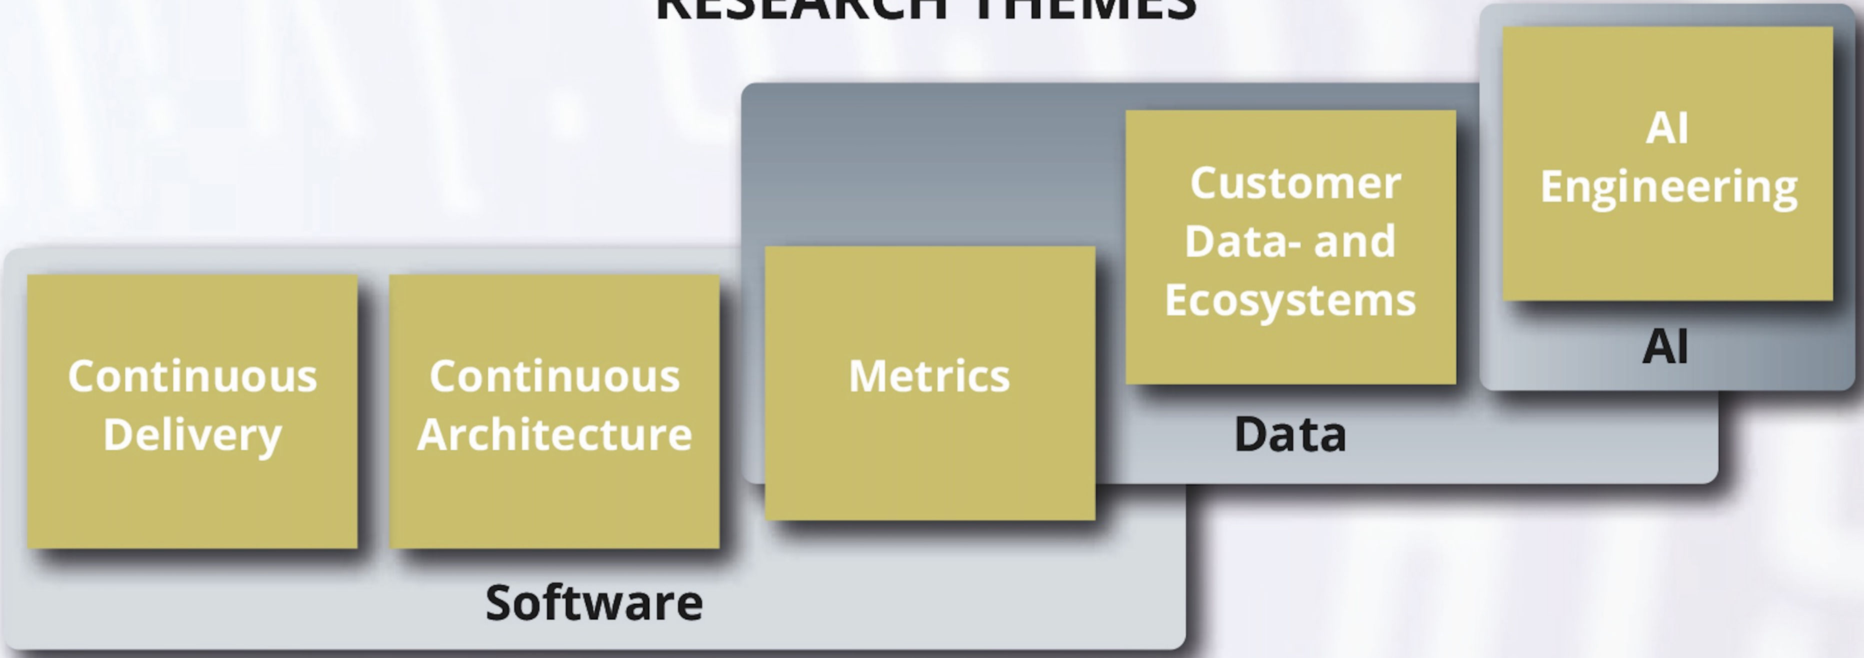

Some Online Companies

BOOKING.COM
online hotel reservations

Spotify

ROVIO

Klarna
Simplifying Buying

King

Microsoft

Advantages for Industry

- More consistent, integrated focus on your key change initiatives
- Holistic approach including technical, organizational and business aspects
- Value every 6 months
- Opportunity to steer projects frequently

Three Key Take-Aways

- Digital companies need to be world-class at software, data and **AI**
- Moving AI from prototyping and experimentation to deployment in production requires significant engineering effort
- AI engineering is concerned with software engineering for AI

Overview

- Vem är jag? Wie ben ik? Who am I?
- Reflecting on AI
- AI Engineering
- Example research activities
- Conclusion

The new spring in artificial intelligence is the most significant development in computing in my lifetime

— Sergey Brin, President ALPHABET

Artificial Intelligence

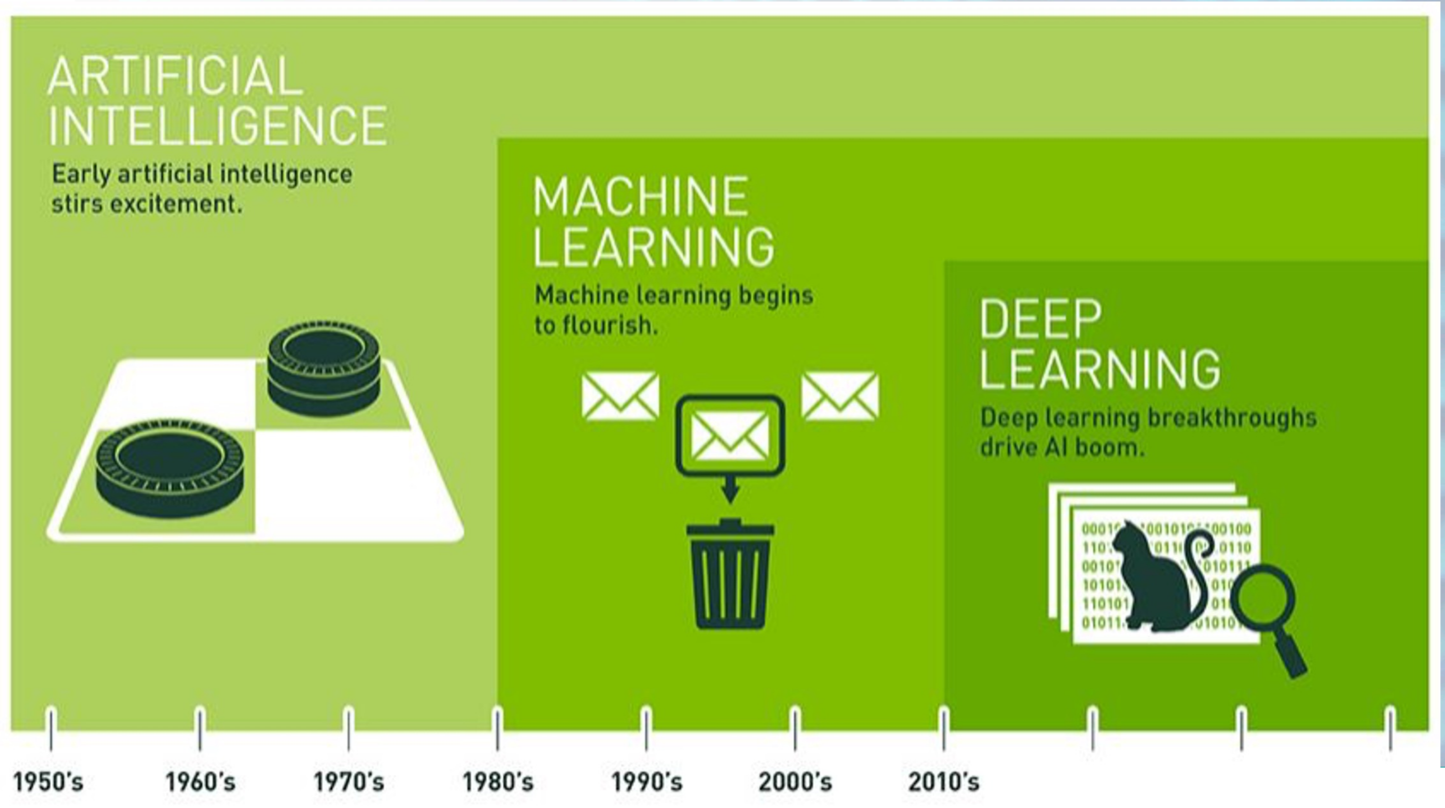

Why Machine Learning Revolution now?

INCREASE OF TRAINING DATA

Large datasets of labeled training data like ImageNet

INCREASE OF COMPUTING POWER

Especially GPUs & Distributed Computing.

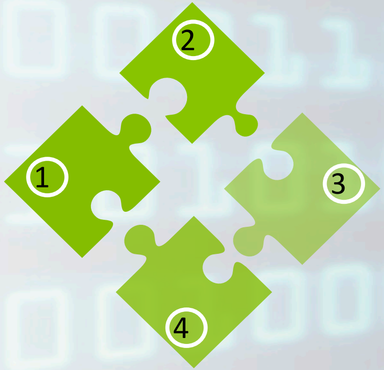

ALGORITHMIC IMPROVEMENTS

Especially in deep reinforcement learning but also normal DL (ReLUs [Dahl et al., 2013], Dropout [Srivastava et al., 2014], etc.)

BETTER TOOLING

Libraries like TensorFlow and Cloud Services

The Typical Problem

50%-85% of AI projects fail[1]

[1]https://www.pacteraedge.com/pactera-white-paper-reveals-85-percent-ai-projects-ultimately-fail-0


If you haven’t failed, you’re
not trying hard enough.
Jennifer Crusie
quotefancy

Why Software Engineering For AI?

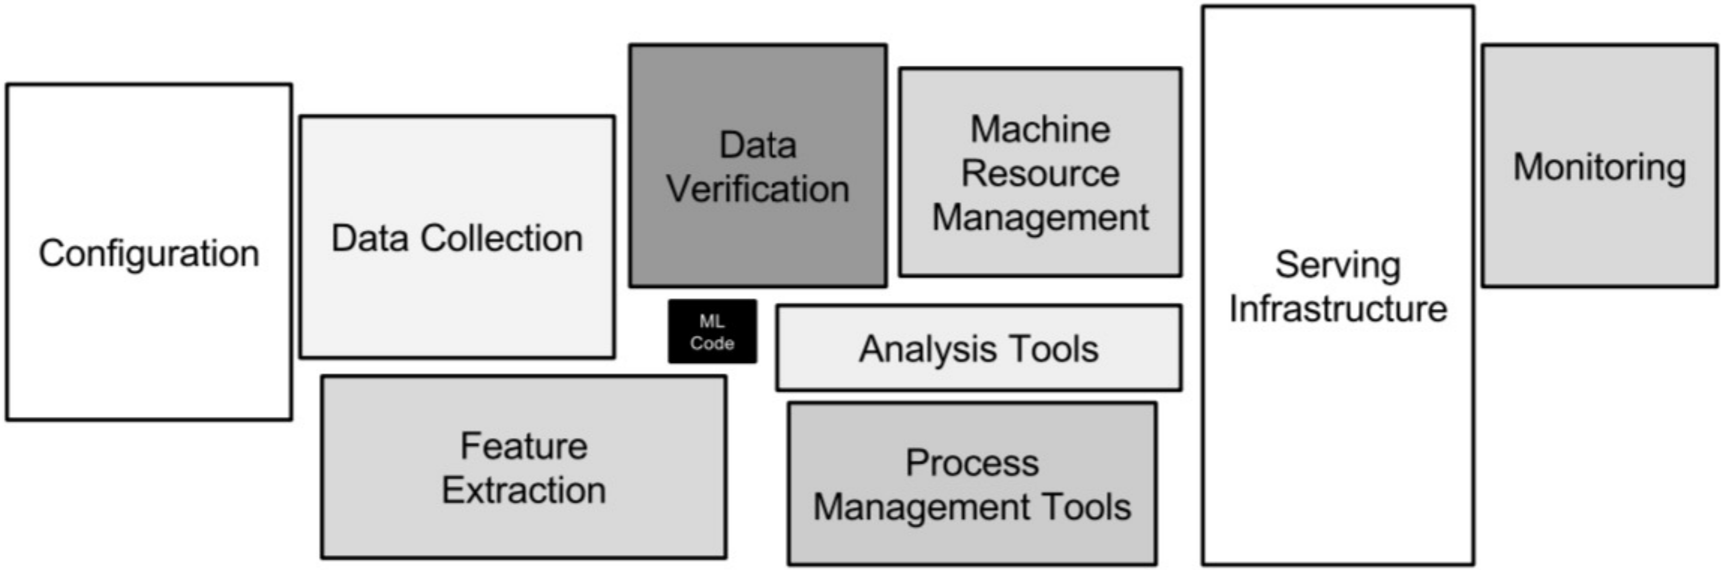

Google: D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, and D. Dennison, "Hidden technical debt in machine learning systems," in Advances in Neural Information Processing Systems, 2015, pp. 2503–2511

Where The Effort Goes …

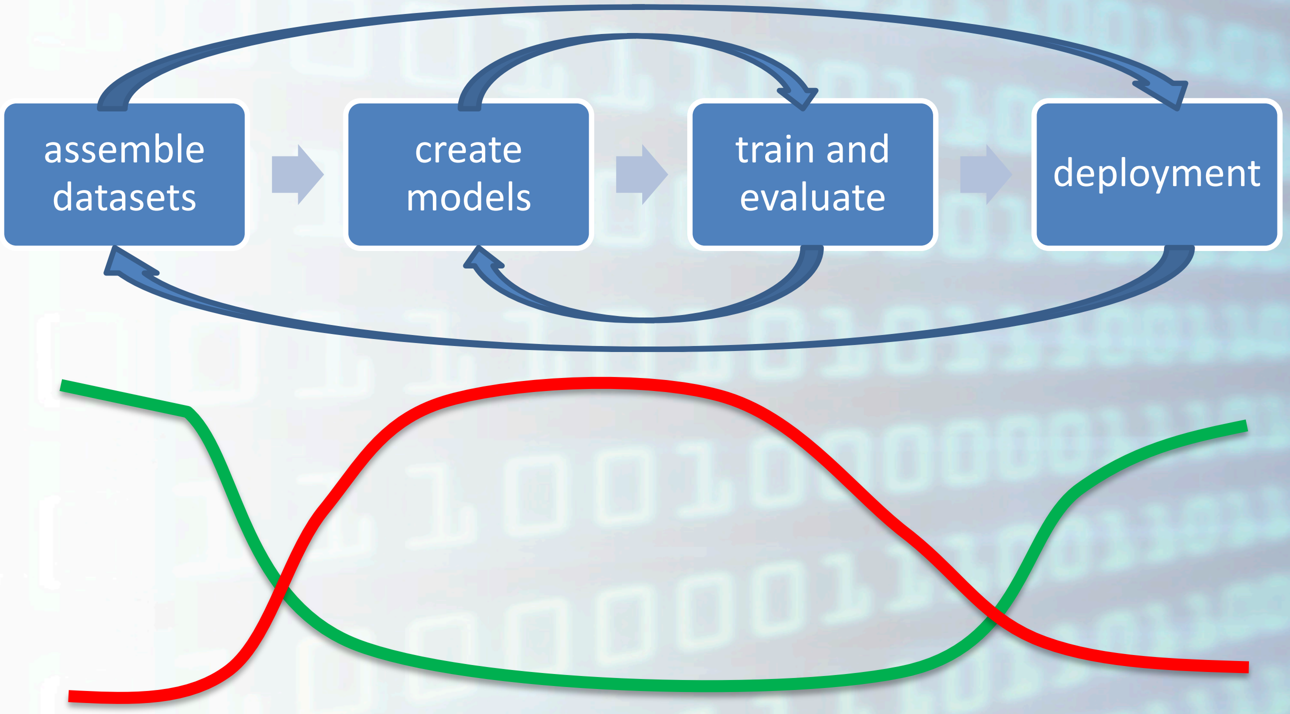

Amount of effort

Amount of attention

Quote from KPMG

“One of the biggest problems facing businesses is getting AI from pilot phase to scalable deployment”

KPMG Consulting

Five Most Important Jobs in AI

- **AI Architect –** These specialists look at individual business processes - as well as the big picture organization - and determine where they can inject and embed AI successfully. They are also responsible for measuring performance, and sustaining the AI model over time – ensuring it removes mundane tasks to optimize humans in the workforce. The lack of AI architects is a big reason why companies cannot successfully sustain AI initiatives.
- **AI Product Manager –** Working close … uct manager serves as a liaison across multiple business teams to en… …emented. They also work closely with these teams – as well as HR – to ident… to ensure optimal performance of both humans and machines.
- **Data Scientist –** With the ever-growing amount of data available to businesses, there is a shortage of experts with the skills to clean this data, and then design and apply the appropriate algorithms to glean meaningful insights.
- **Software Engineer –** One of the biggest problems facing businesses is getting AI from pilot phase to scalable deployment. Software engineers work hand-in-hand with data scientists to bring AI into production, blending business acumen with a deep understanding of how AI works.
- **AI Ethicist –** As ethical and social implications of AI continue to unfold, companies may need to create new jobs tasked with the critical responsibility of establishing AI frameworks that uphold company standards and codes of ethics. Initially, these roles could be fulfilled by existing leaders in an organization, but as the effects of AI fully take shape, it may need to be the responsibility of one person to ensure these guidelines are upheld.

Software Engineering

https://info.kpmg.us/news-perspectives/technology-innovation/top-5-ai-hires-companies-need-to-succeed-in-2019.html

https://qz.com/work/1517594/the-five-most-important-new-ai-jobs-according-to-kmpg/

AI: The New Electricity

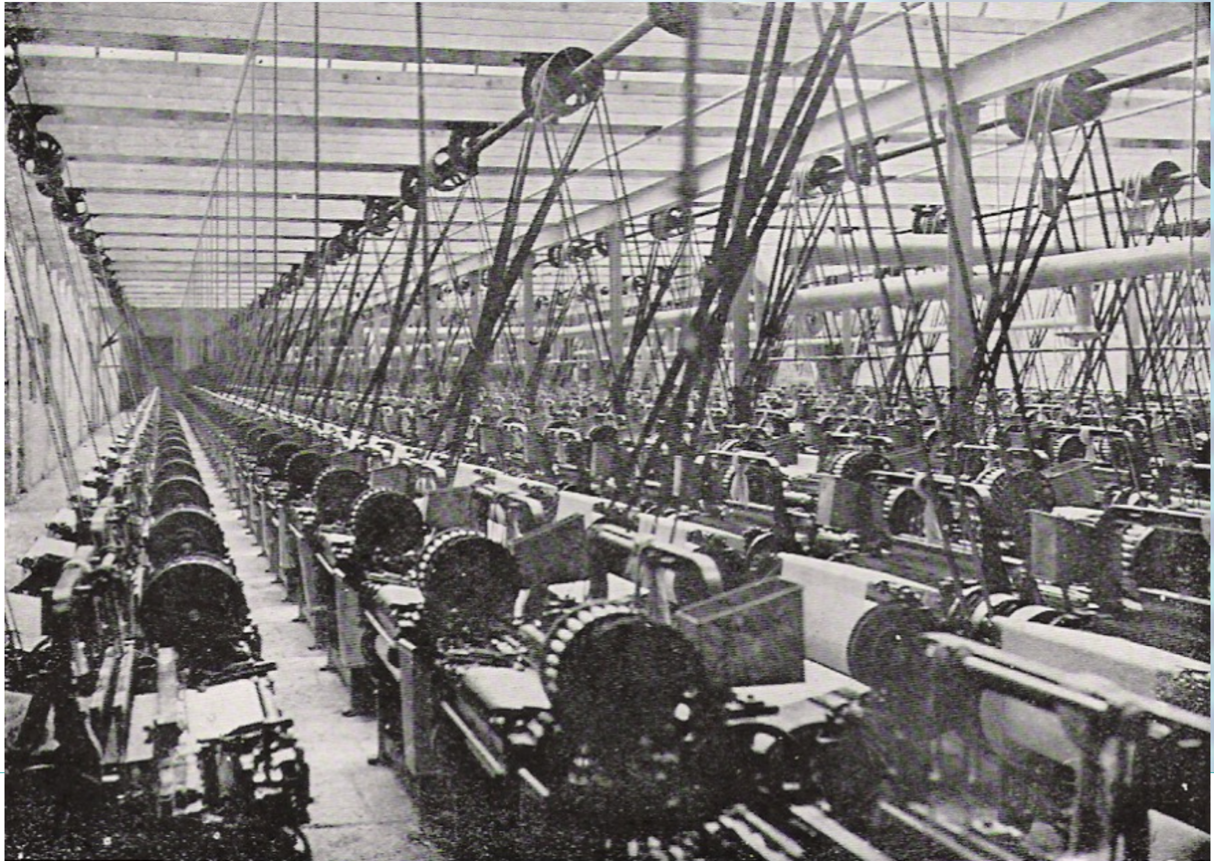

Holistic DevOps Framework

Holistic DevOps Framework

Requirements driven development
- Regulatory features
- Competitor parity features
- Commodity features

Outcome/data driven development
- Value hypothesis
- New ”flow” features
- Innovation

AI driven development
- Minimize prediction errors
- Many points in data set
- Combinatorial explosion of alternatives

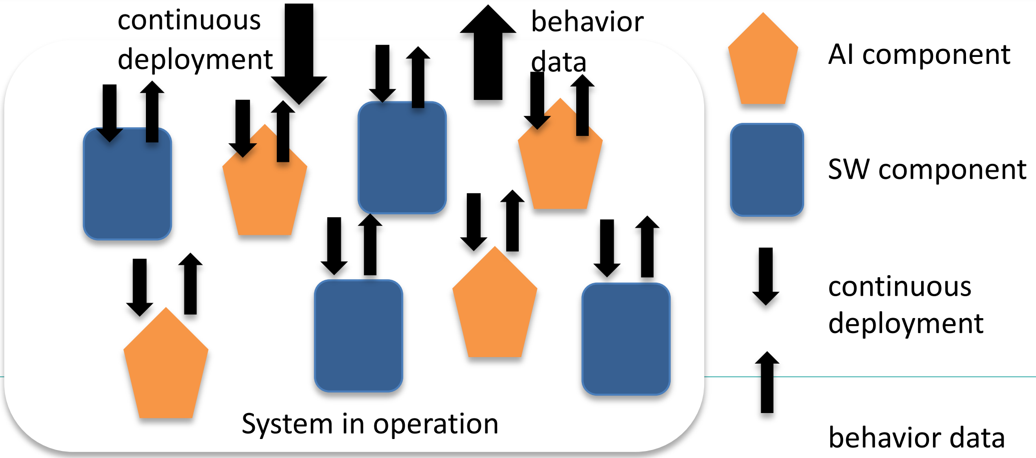

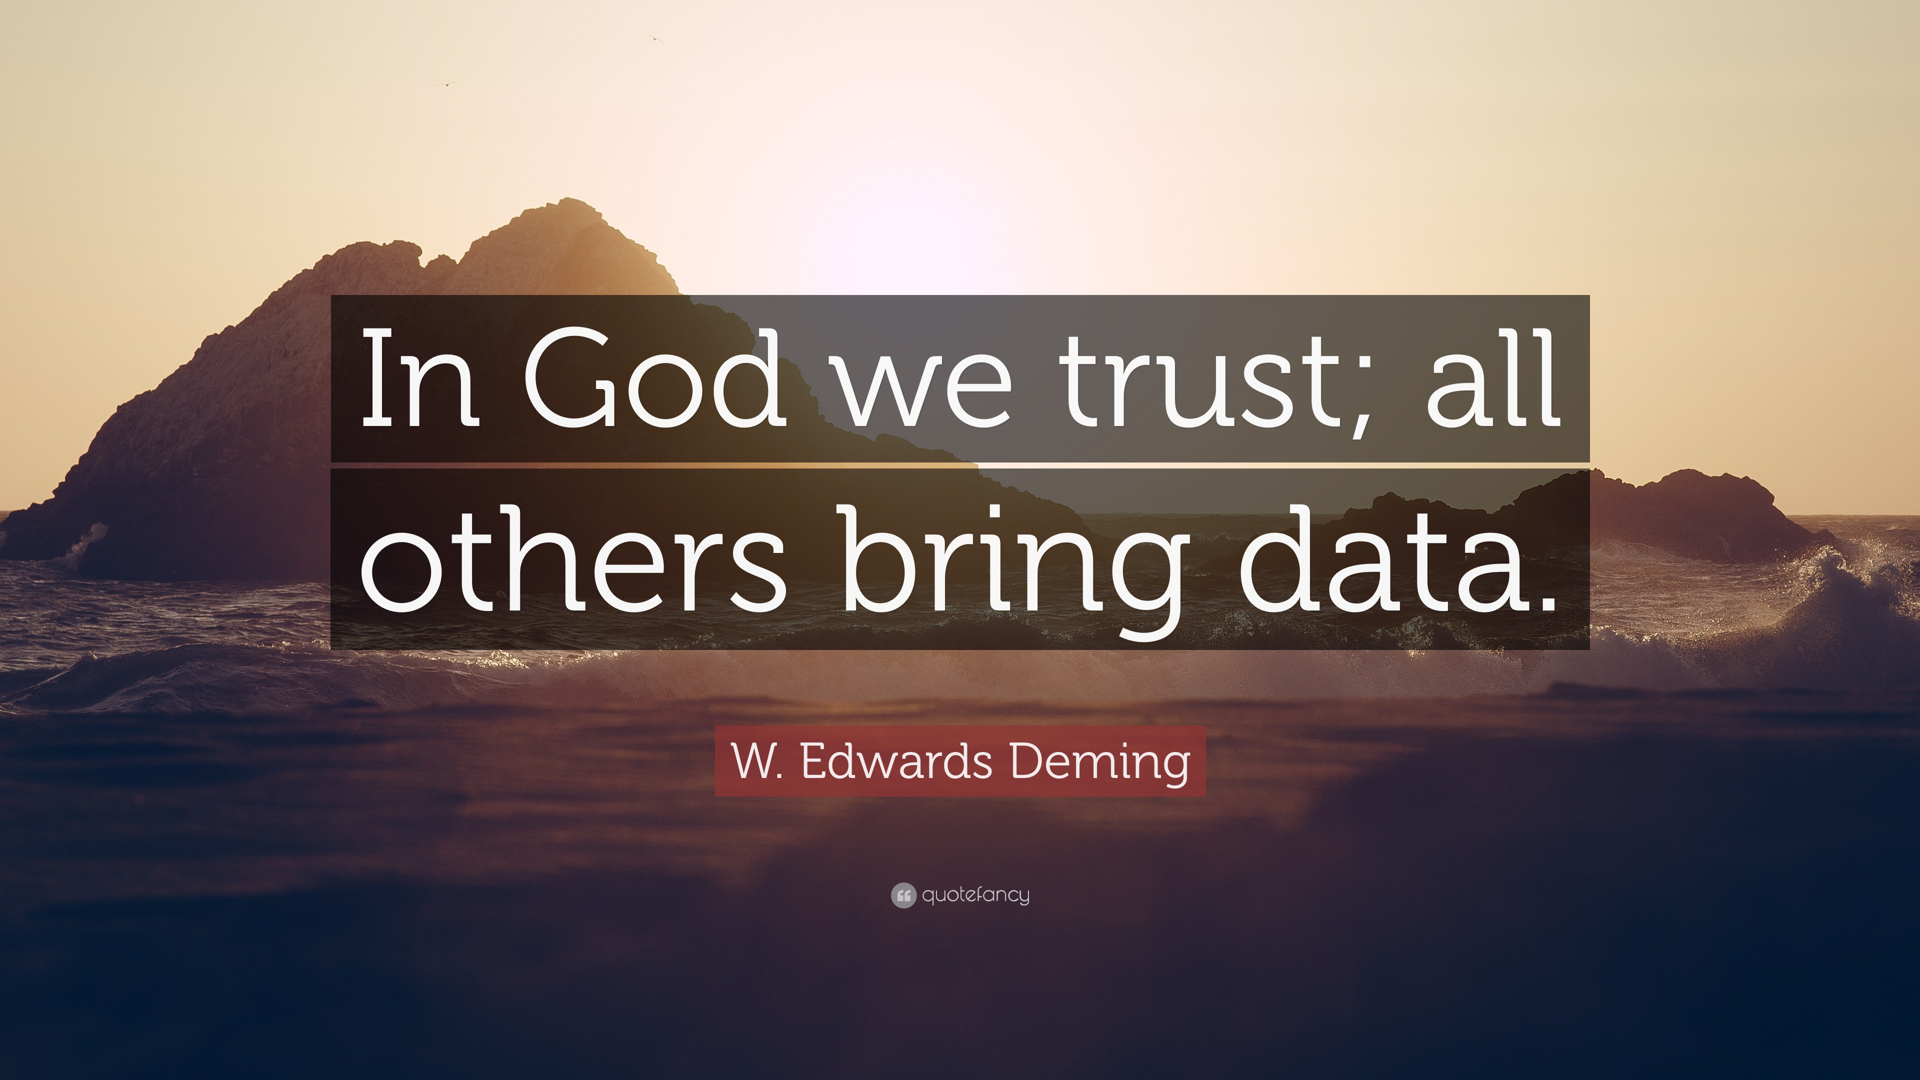
In God we trust; all
others bring data.
W. Edwards Deming
quotefancy

Overview

- Vem är jag? Wie ben ik? Who am I?
- Reflecting on AI
- AI Engineering
- Example research activities
- Conclusion

AI Engineering

- AI engineering is concerned with applying (software) engineering principles to the creation, deployment and evolution of AI driven systems
- AI engineering focuses on the business, architectural, process and organizational implications of using ML/DL models in industrial systems, including data pipelines, monitoring & logging, quality attributes, etc.

How AI Evolves In Industry

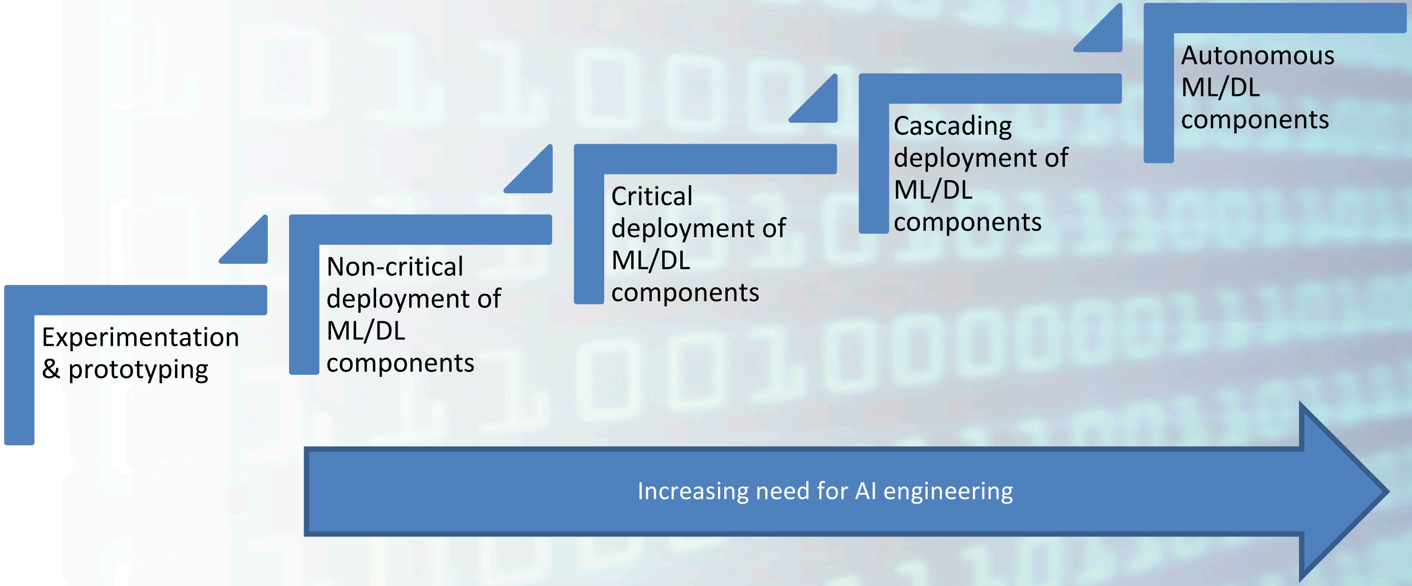

AI Engineering Research Agenda

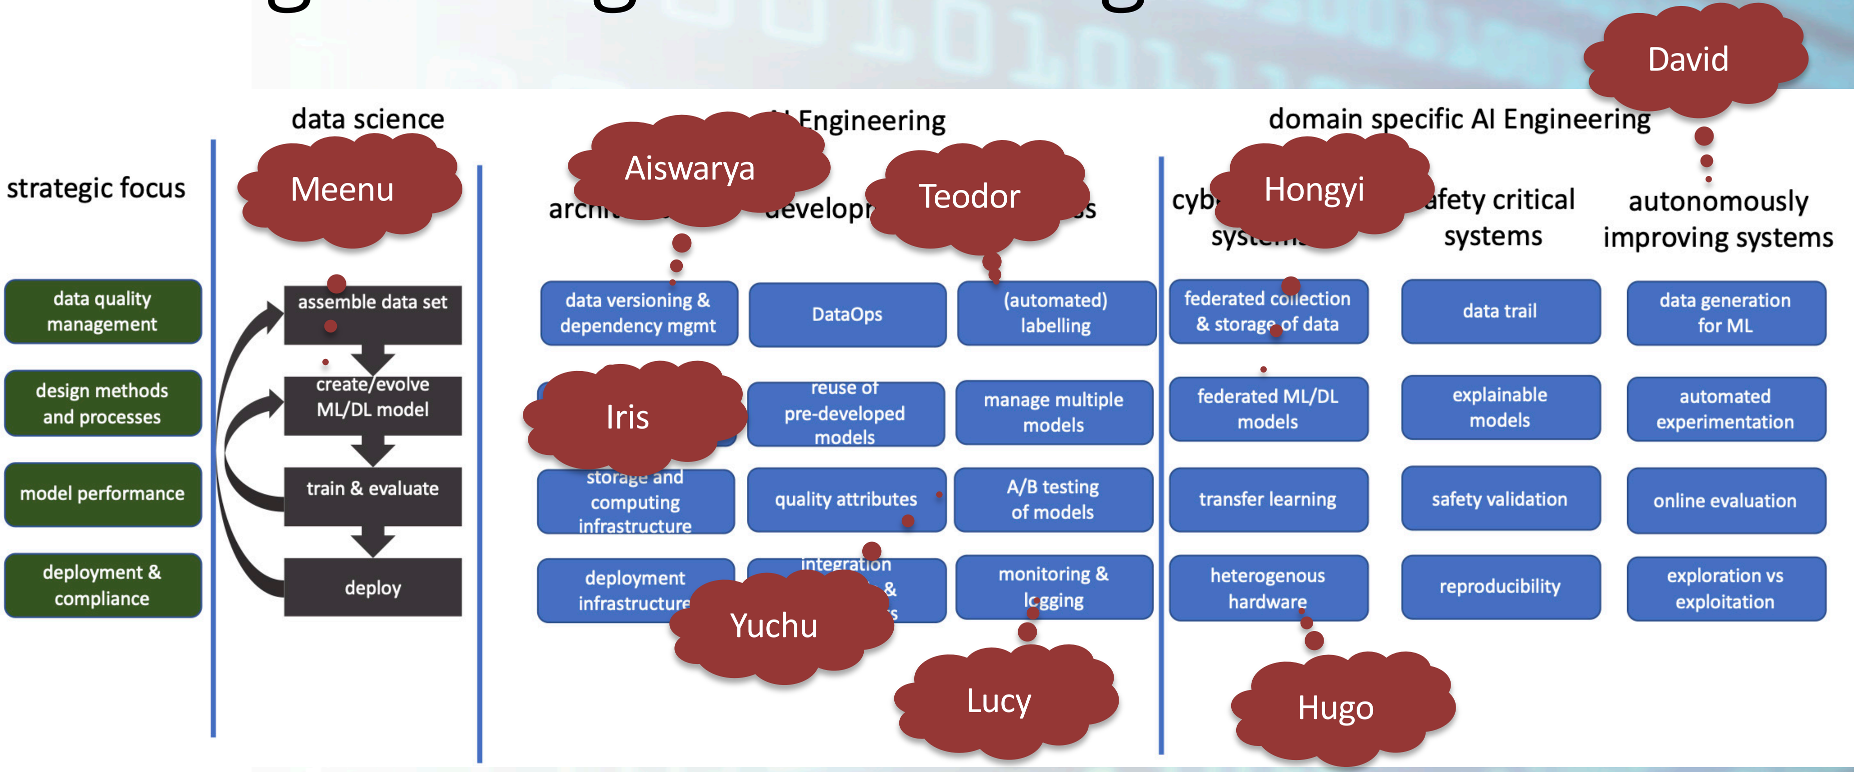

Bosch, Jan, Helena Holmström Olsson, and Ivica Crnkovic. "Engineering ai systems: A research agenda." Artificial Intelligence Paradigms for Smart Cyber-Physical Systems. IGI Global, 2021. 1-19.

Overview

- Vem är jag? Wie ben ik? Who am I?
- Reflecting on AI
- AI Engineering
- Example research activities
- Conclusion

Study #1

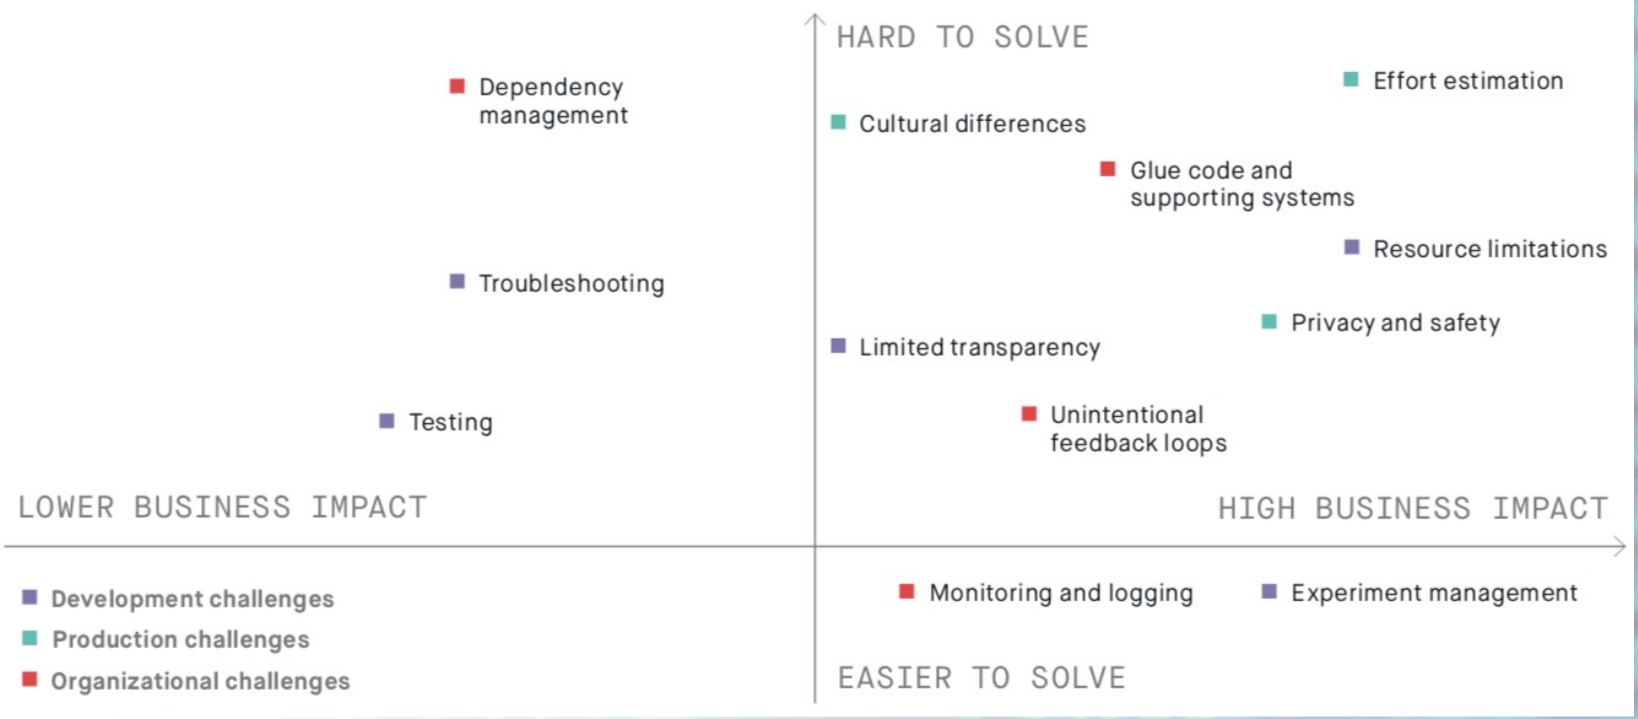

Arpteg, Anders, Björn Brinne, Luka Crnkovic-Friis, and Jan Bosch. "Software engineering challenges of deep learning.” In *2018 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pp. 50-59. IEEE, 2018.

Study #2

	Experiment Prototyping	Non-critical deployment	Critical deployment	Cascading deployment
assemble dataset	Issues with problem formulation and specifying desired outcome	Data silos, scarcity of labelled data, imbalanced training set	Limitations in techniques for gathering training data from large-scale, non-stationary data streams	Complex and effects of data dependencies
create model	Use of non-representative dataset, data drifts	No critical analysis of training data	Difficulties in building highly scalable ML/DL pipeline	Entanglements causing difficulties in isolating improvements
train and evaluate model	Lack of well-established ground truth	No evaluation of models with business-centric measures	Difficulties in reproducing models, results and debugging DL models	Need of techniques for sliced analysis in final model
deploy model	No deployment mechanism	Training-serving skew	Adhering to stringent serving requirements e.g., of latency, throughput	Hidden feedback-loops and undeclared consumers of the models

Table 2. Challenges in the evolution of use of ML/DL components in software-intensive systems

Lwakatare, Lucy Ellen, Aiswarya Raj, Jan Bosch, Helena Holmström Olsson, and Ivica Crnkovic.
"A taxonomy of software engineering challenges for machine learning systems: An empirical investigation."
In *International Conference on Agile Software Development*, pp. 227-243. Springer, Cham, 2019.

Study #3

MAPPING BETWEEN DATA MANAGEMENT CHALLENGES AND USE CASES

Phase	Challenge	Use cases of DL components					
		RS[1]	WPP[2]	HPP[3]	MD[4]	FFD[5]	MS[6]
Data Collection	Lack of metadata	X	X	X	X	X	X
	Data Granularity		X	X			
	Shortage of diverse samples		X	X	X	X	
	Need for sharing and tracking techniques	X	X	X	X	X	
	Data Storage	X					
Data Exploration	Statistical Understanding		X			X	X
	Deduplication Complexity	X	X	X	X	X	X
	Heterogeneity in data	X	X	X	X	X	
Data Preprocessing	Dirty data	X	X	X	X	X	X
	Managing sequences in data					X	X
	Managing categorical data				X	X	
Dataset Preparation	Data Dependency	X	X	X	X	X	X
	Data Quality	X	X	X	X	X	X
Data Testing	Tooling	X	X	X	X	X	X
	Expensive Testing	X			X		X
Deployment	Data Extraction Methods	X	X	X	X	X	X
	Overfitting				X	X	
Post Deployment	Data sources and Distribution	X	X	X			
	Data drifts	X	X	X			
	Feedback loops	X					

[1] Recommender System [2] Wind Power Prediction
[3] House Price Prediction [4] Melanoma Detection
[5] Financial Fraud Detection [6] Manufacturing Systems

Munappy, Aiswarya, Jan Bosch, Helena Holmström Olsson, Anders Arpteg, and Björn Brinne. "Data Management Challenges for Deep Learning." In *2019 45th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pp. 140-147. IEEE, 2019.

Conceptual Model for Data Pipelines

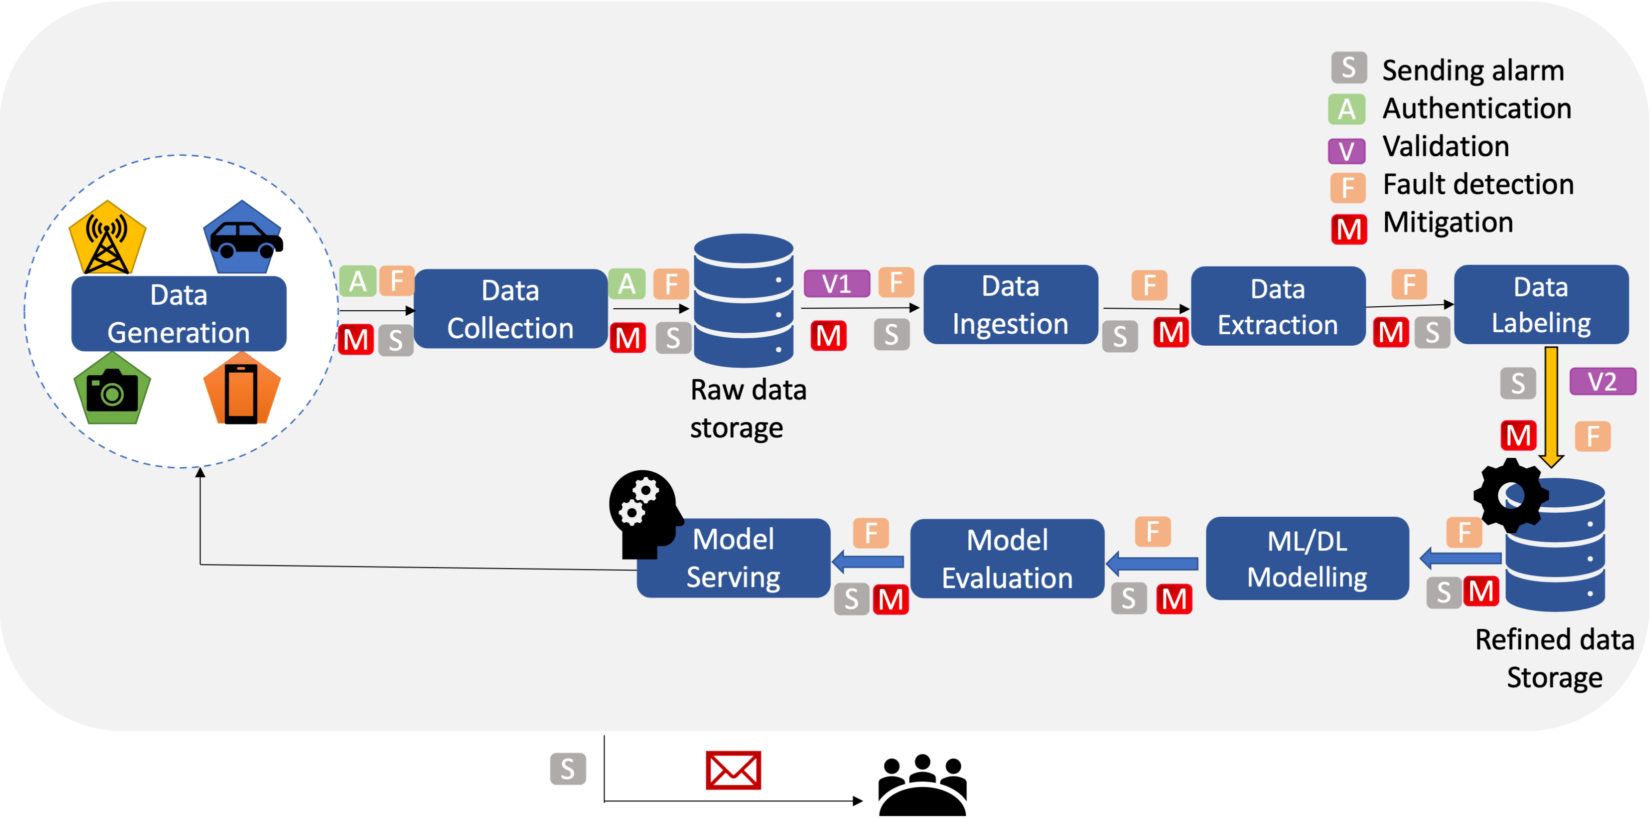

Automatic data validation

- Irrespective of ML algorithms, data errors (e.g., training-serving skew) can adversely affect the quality of generated model

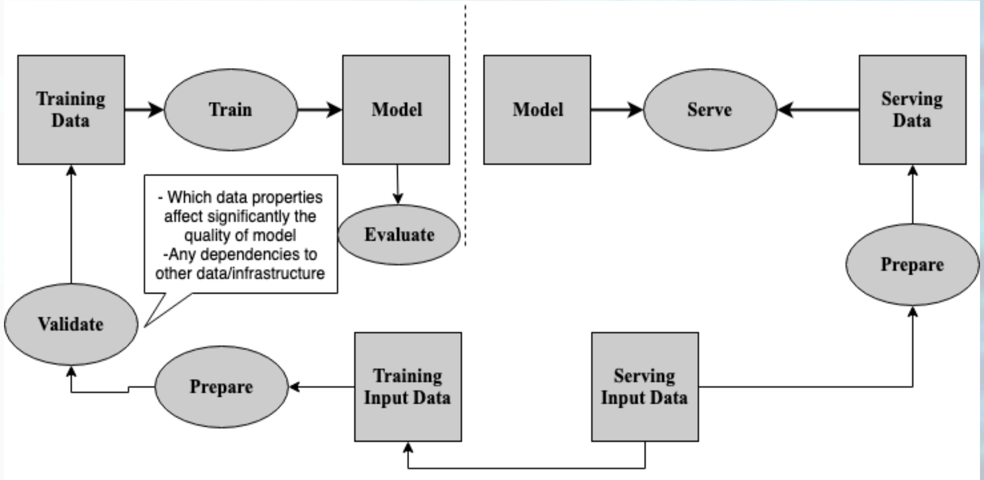

Need to ensure data errors are caught early in the context of continuous and frequent ML model training and serving

Automatic data validation

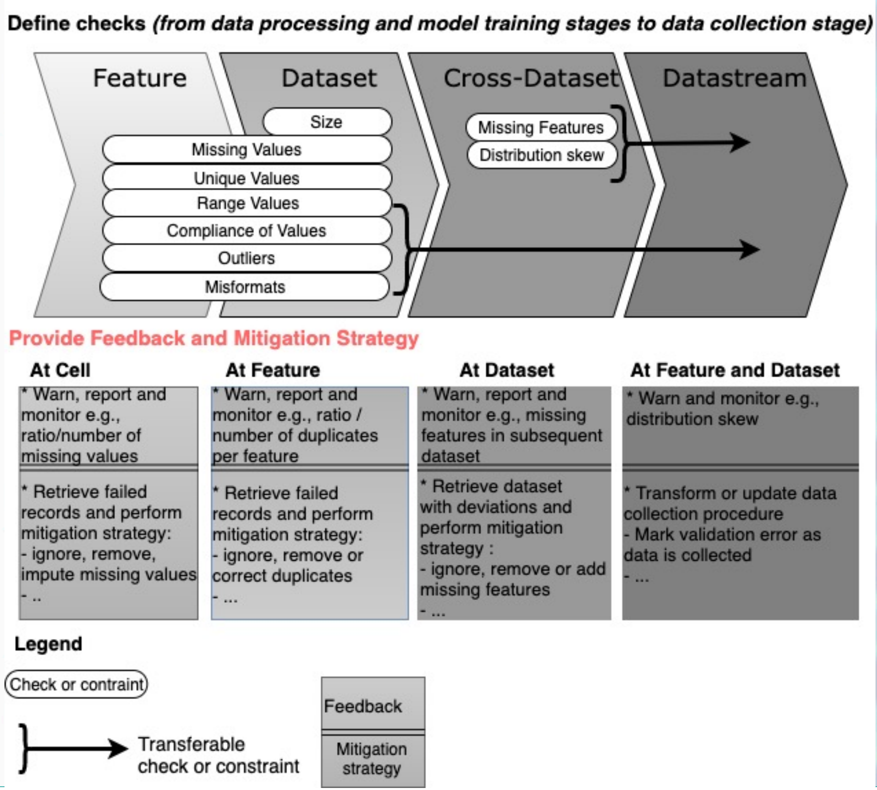

Automatic Labeling

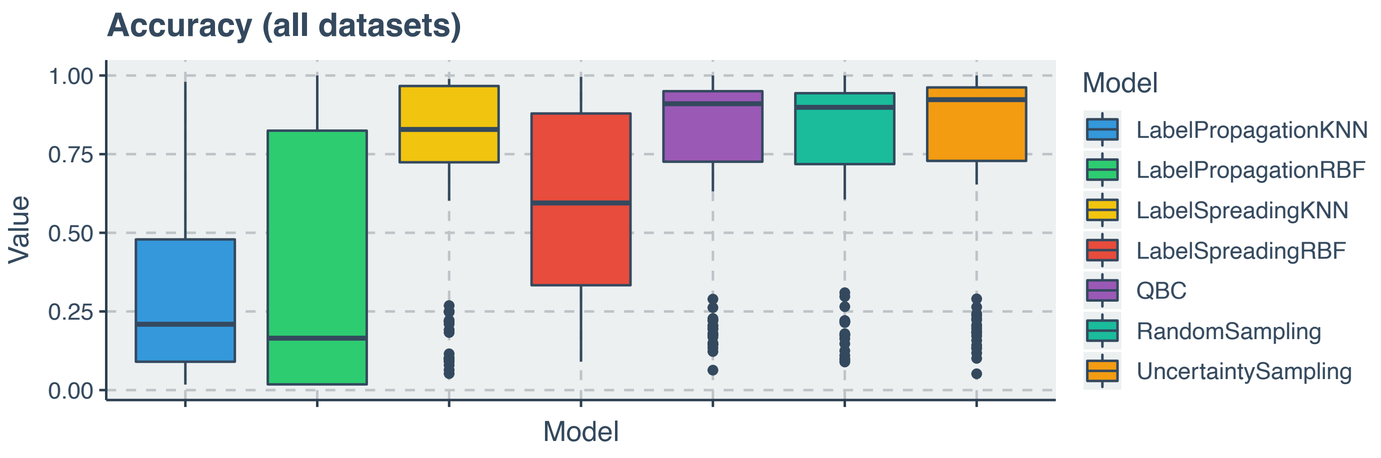

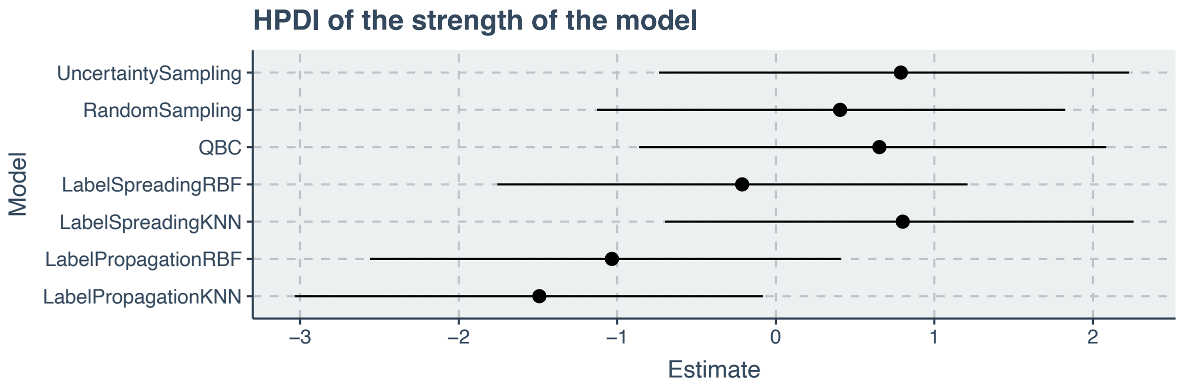

Deep Semi-Supervised Learning

- The two most common learning paradigms are Supervised learning and Unsupervised learning.
  - Supervised learning utilizes labelled data.
    - The more labelled data you have the better, especially for deep learning algorithms.
    - In industrial settings labelled data is often scarse.
  - Unsupervised learning utilizes unlabelled data
    - Learns patterns from unlabelled data.
    - Unlabelled data is easy to find
- Question: Can we learn from both labelled and unlabelled data and can this improve supervised classification?
- Yes we can!
  - Semi-Supervised Learning

Automated Experimentation

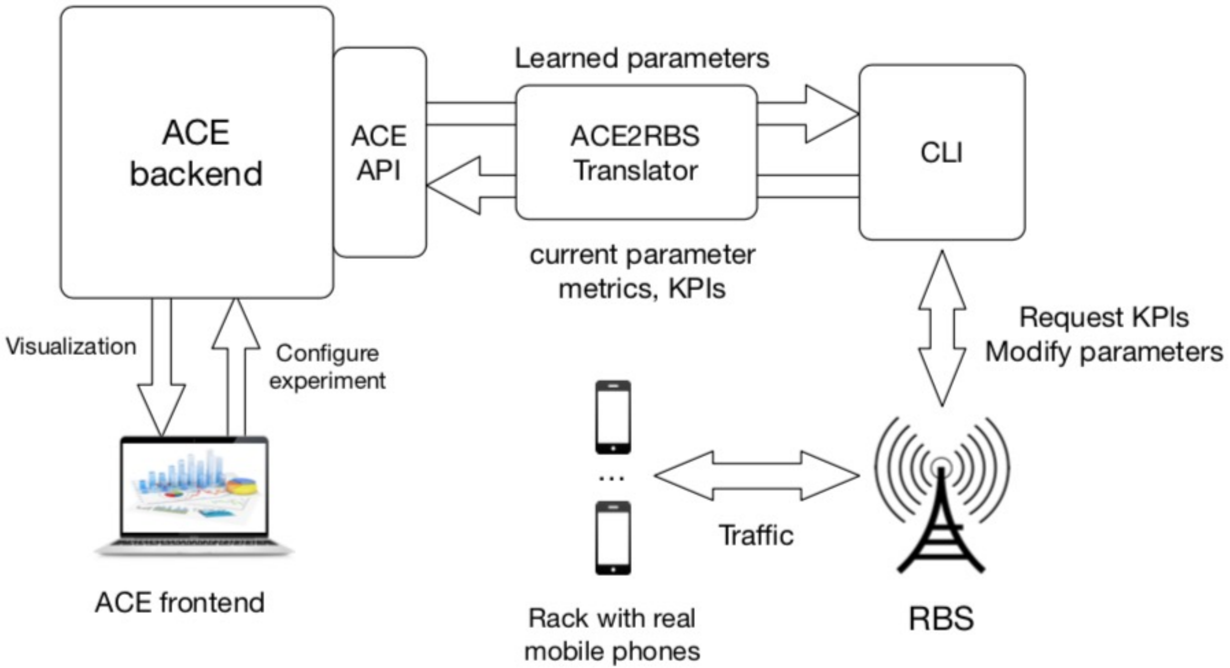

Multi-armed bandit algorithms

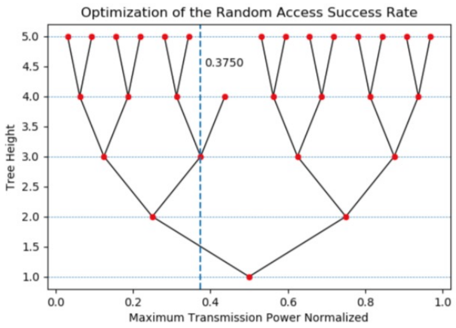

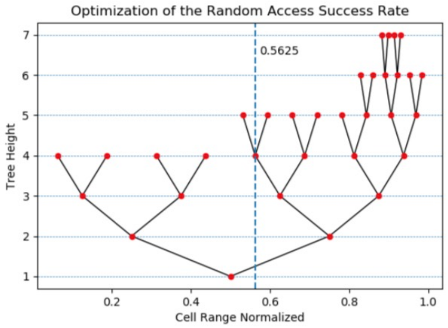

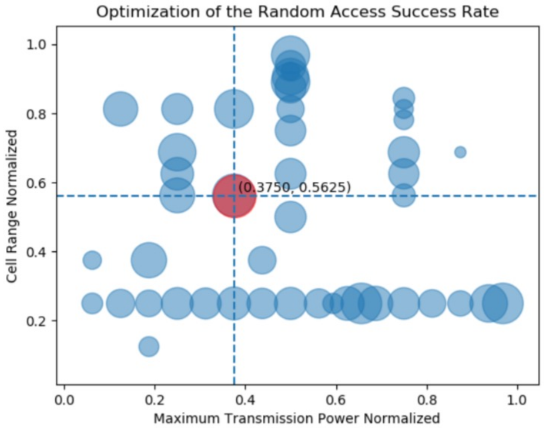

Federated Learning

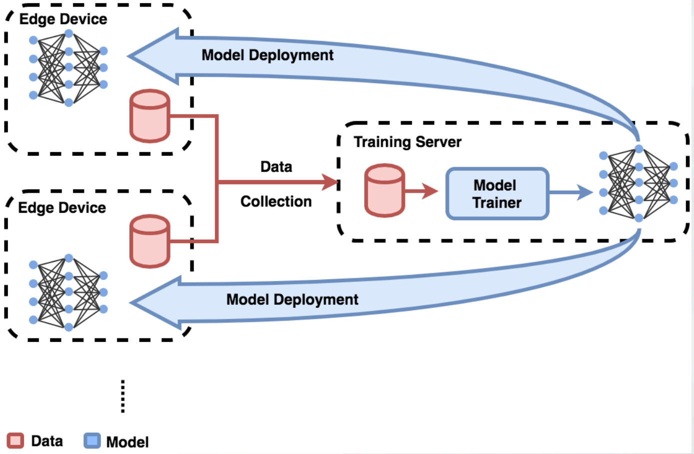

(a) Traditional Machine Learning system

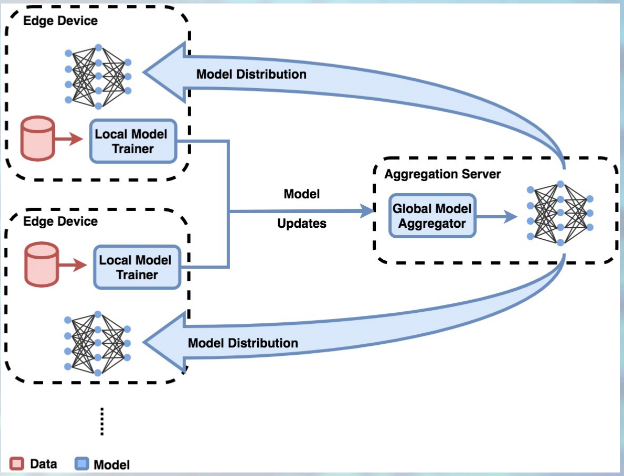

(b) Federated Learning system

Aynchronous Federated Learning

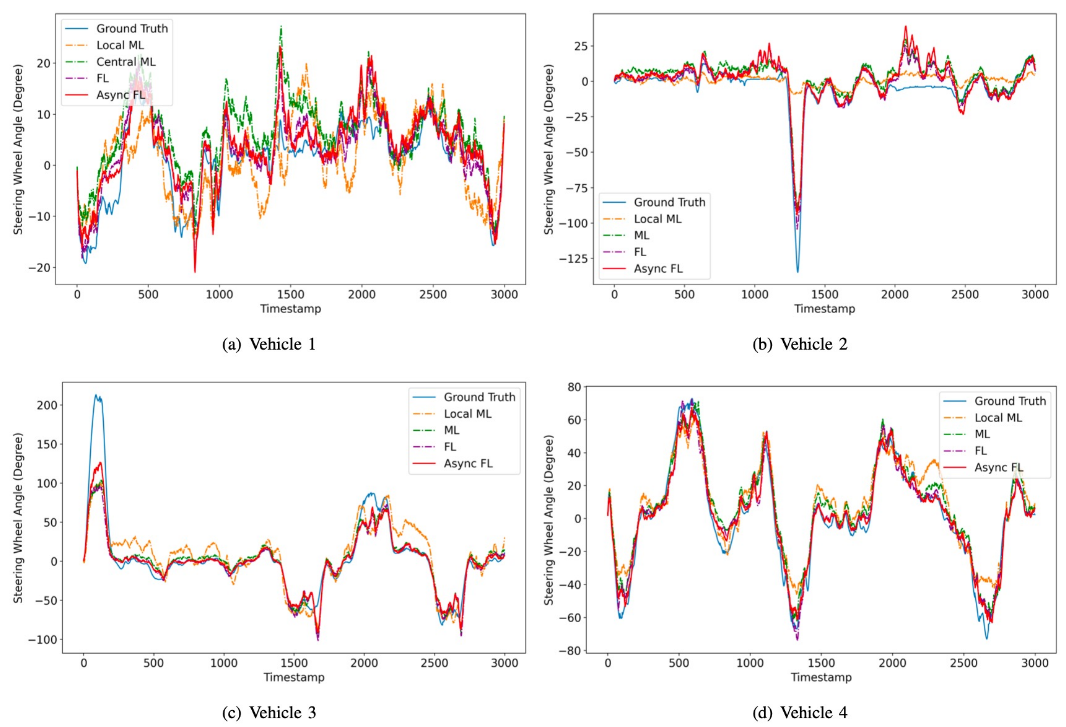

(a) Vehicle 1

(b) Vehicle 2

(c) Vehicle 3

(d) Vehicle 4

Decentralized Reinforcement Learning - System Scenario

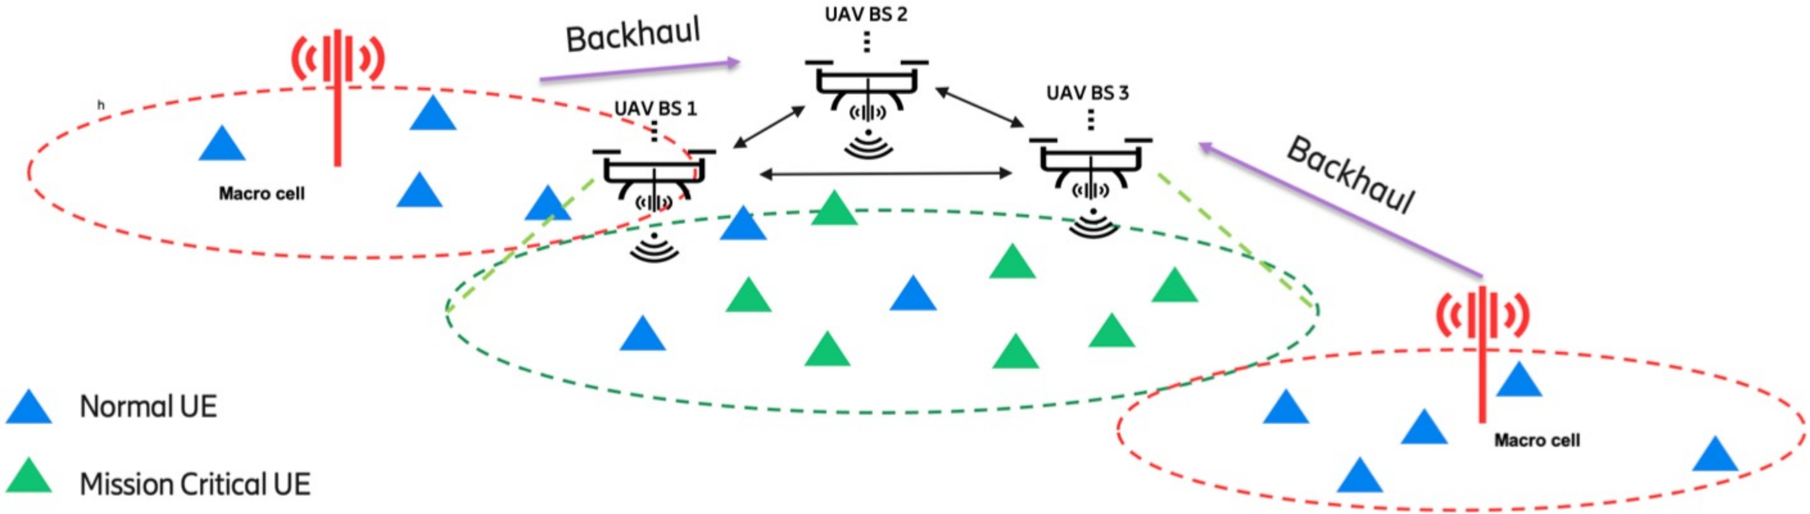

Fig: UAV-assisted system design: Mission-critical scenario enabled by three wireless drones

Decentralized Reinforcement Learning

The system has two different channels for exchanging information:

- The brown line indicates the sensor data channel, which will communicate to nearby drones connection performance and UAV-BS location data.

- The local model of each UAV-BS will be shared with its neighbours via the model data channel, which is indicated by the green line.

- Each UAV-BS will share their learning experiences as a result, and the others can gain information from the experiences of the others. After multiple training epochs, the UAV-BS can swap their model with their neighbours under the control of a frequency parameter.

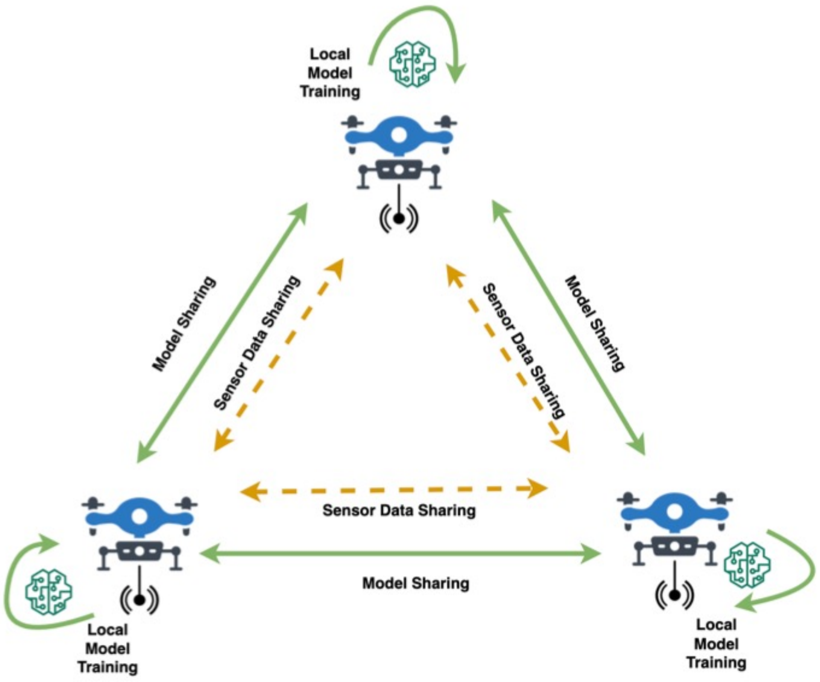

- To assist companies to determine whether to deploy AI models in the cloud or on the edge, we evaluate architectural alternatives for deploying AI on the edge based on the key factors.
- Involves experimentation with 3 architectural alternatives:
  - Centralized architecture
  - Federated architecture
  - Decentralized architecture
- Validate these architectures against key factors
  - *Mass customization* - Extent to which a local model can be customized and how it impacts overall architecture
  - *Scalability* – Measures how effectively the local models scale from one-to-many instances
  - *Model Performance* – Indicates the overall quality of the architecture
- Implementation:
  - Using SNIC (2 servers and 3 edge instances)
  - Three datasets - MNIST, CIFAR-10, ECMWF
  - Model - CNN

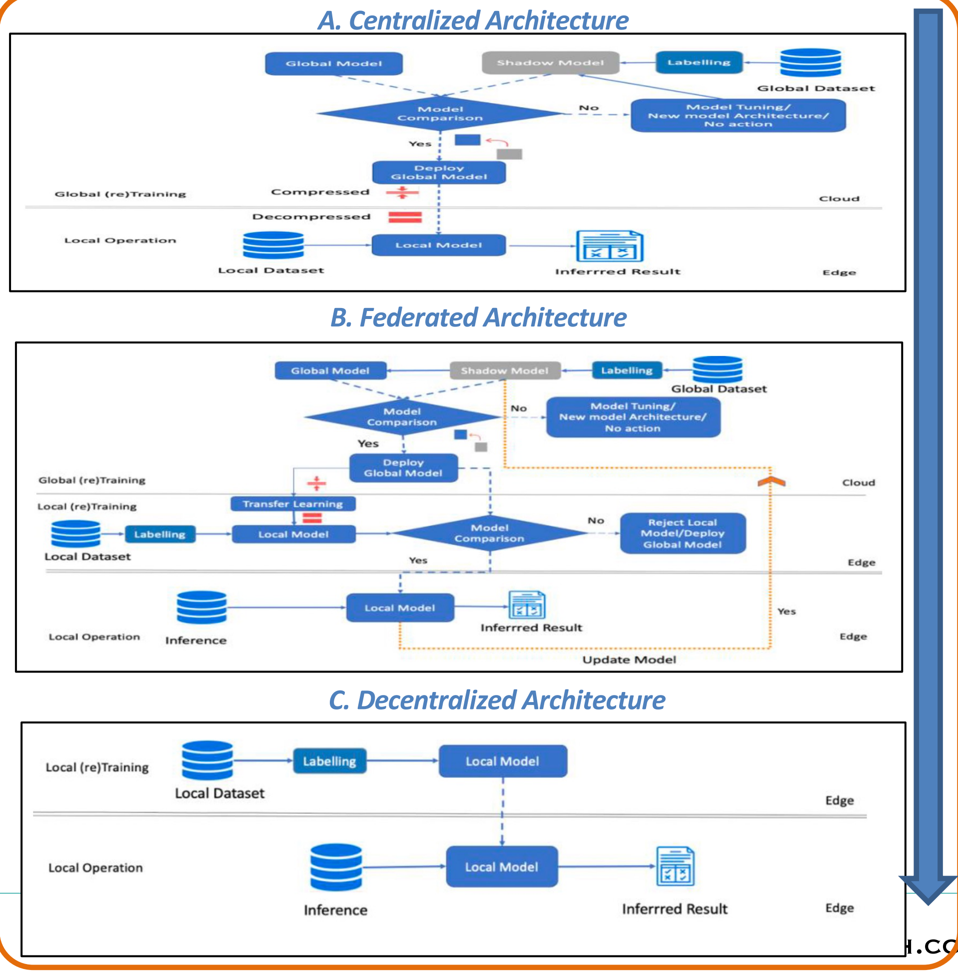

Why monitoring is important??

- To **track the performance of ML/DL models in production** and ensure that business objectives are met, model monitoring is a key step
- Even though ML/DL practitioners may benefit from existing lessons and approaches from software monitoring, ML/DL model monitoring presents **unique challenges that necessitate the use of specialized techniques, processes, and tools**
- To address these challenges, we propose to conduct in-depth studies **focusing on the specifics of monitoring models** after they have been operationalized in embedded systems.

Research Questions:

RQ. How to monitor models to ensure that it performs well once it is operational?

RQ1. In software-intensive embedded systems companies, what are the related activities and challenges that practitioners experience while monitoring ML/DL models and what are the the potential solutions to these challenges?

RQ2. What are the different ways to compare results from a newer version(s) of the model against the in-production version(s)?

RQ3. How to build an infrastructure around the ML/DL model for monitoring to make sure it performs well once it is operational?

RQ4. Once the model has been deployed to different clients in the field, how can we identify subgroups of clients with similar characteristics that underperform?

- **Methodology:**
  - Conduct the research either on-site or online --- If on-site, we propose action research (work in close collaboration with experts in the company), Otherwise, in-depth case study.
  - The overall research involves interviews, workshops, focus groups, prototyping /demonstration, experimentation and validation.

For more information, please contact meenu-mary.john@mau.se, helena.holmstrom.olsson@mau.se or jan.bosch@chalmers.se

M

Overview

- Vem är jag? Wie ben ik? Who am I?
- Reflecting on AI
- AI Engineering
- Example research activities
- Conclusion

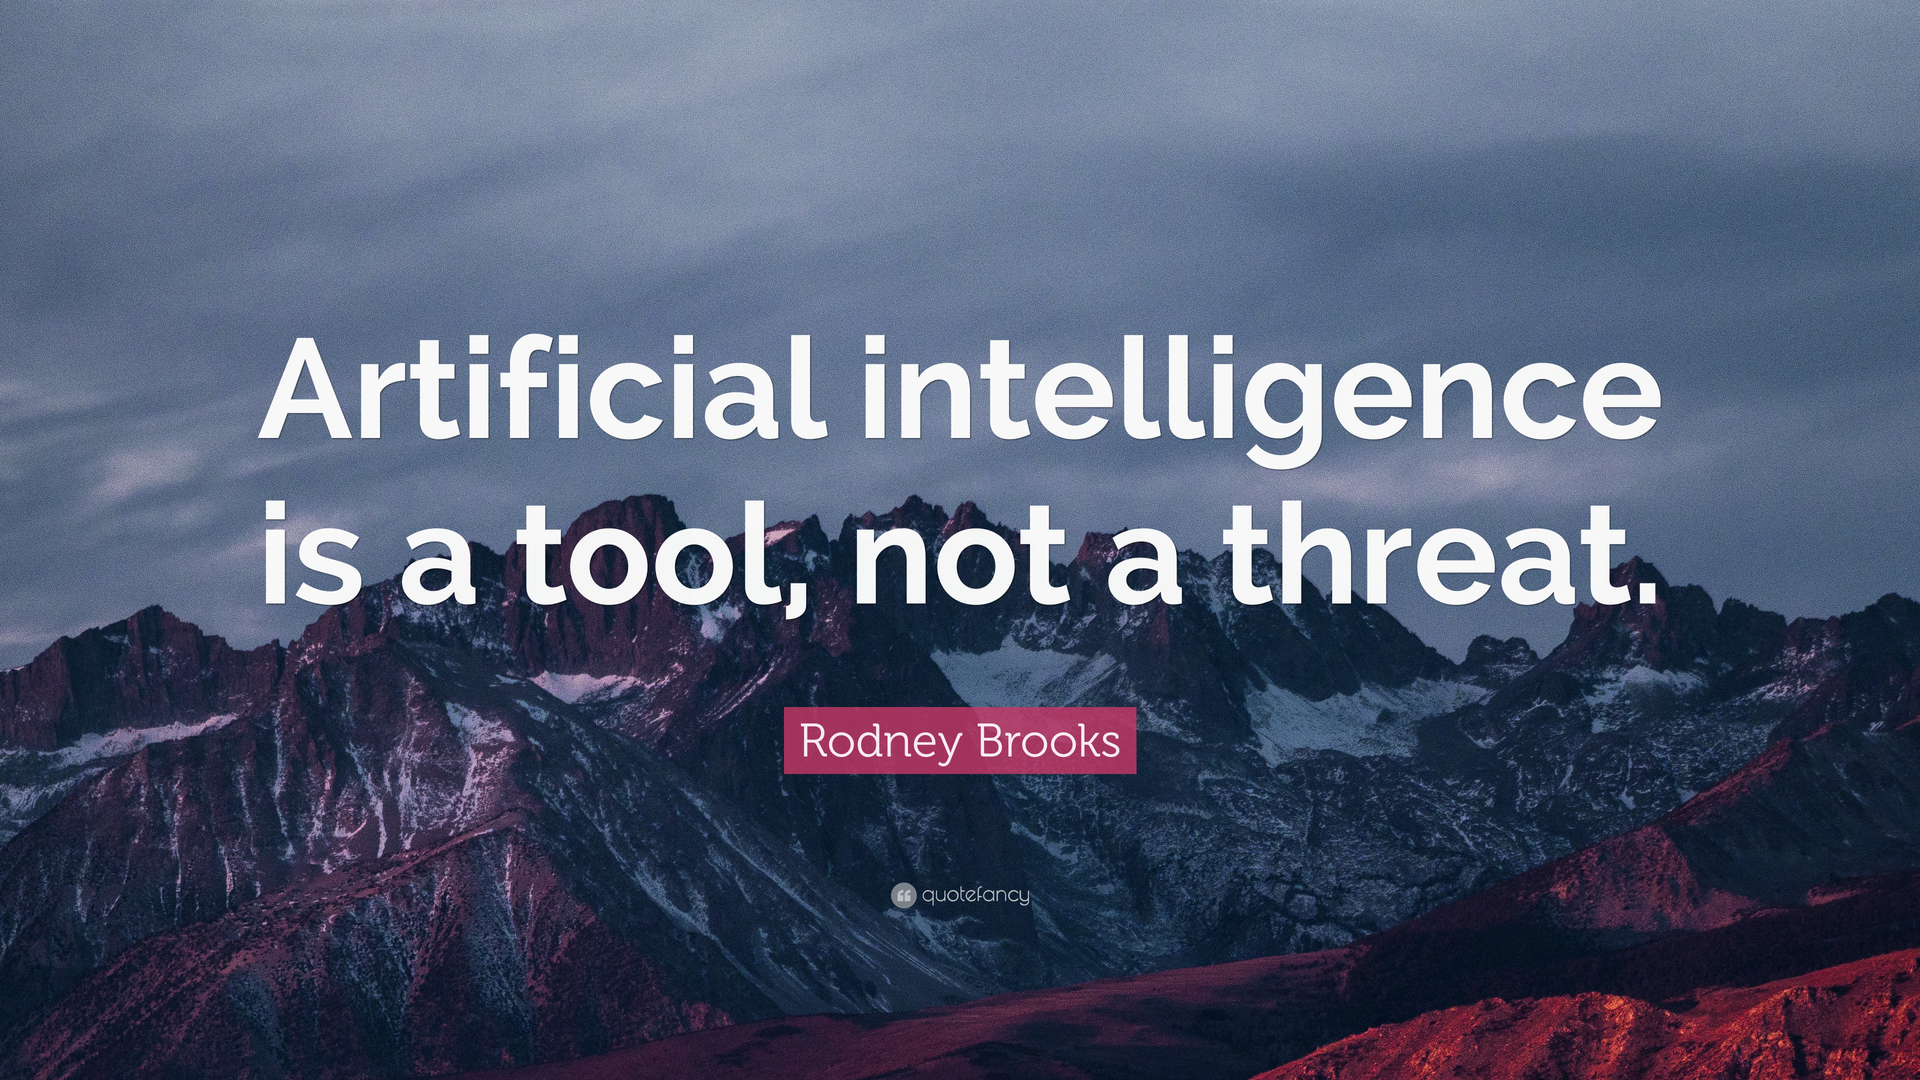
Artificial intelligence is a tool, not a threat.
Rodney Brooks
quotefancy

Conclusion

- Digital companies need to be world-class at software, data and **AI**
- Moving AI from prototyping and experimentation to deployment in production requires significant engineering effort
- AI engineering is concerned with software engineering for AI

Learn More?

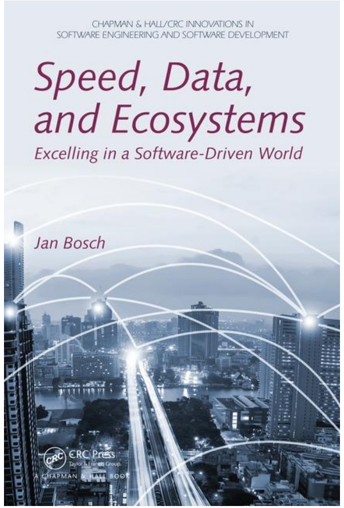
CHAPMAN & HALL/CRC INNOVATIONS IN
SOFTWARE ENGINEERING AND SOFTWARE DEVELOPMENT
Speed, Data, and Ecosystems
Excelling in a Software-Driven World
Jan Bosch
CRC Press

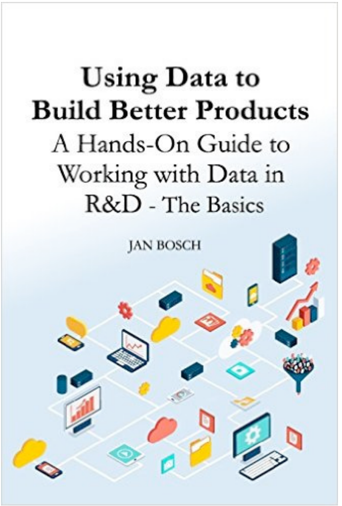
Using Data to
Build Better Products
A Hands-On Guide to
Working with Data in
R&D - The Basics
JAN BOSCH

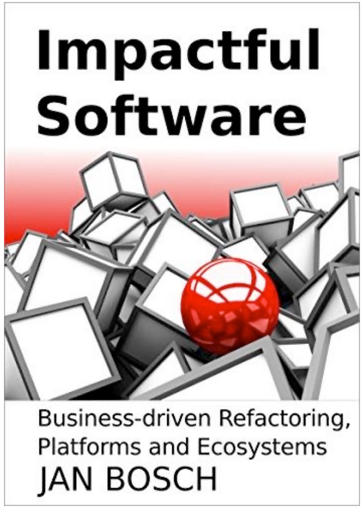
Impactful
Software
Business-driven Refactoring,
Platforms and Ecosystems
JAN BOSCH

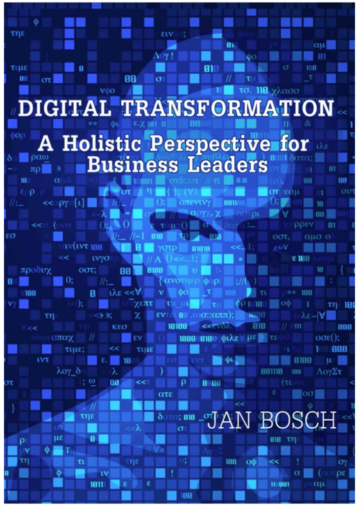
DIGITAL TRANSFORMATION
A Holistic Perspective for
Business Leaders
JAN BOSCH

Software Center

www.janbosch.com
jan@janbosch.com

Follow me on LinkedIn, Twitter (@JanBosch) or www.janbosch.com/blog

Closing thoughts

- Professional development
- Echo chamber
- Role
- Thank you!
    - Grundfos for hosting
    - DEIF and Grundfos for being members
    - All of you for participating!
- What’s next?
    - Software Center reporting workshop – December 15 (and June 15 next year?)
    - Join Software Center! And if you’re a member, join a(nother) project!