AI for Good

The W^4H of AI for Public Good.

Anthony, Faustine

Data Scientist, and ML Researcher

sambaiga sambaiga@gmail.com

sambaiga.github.io/sambaiga/

Wait

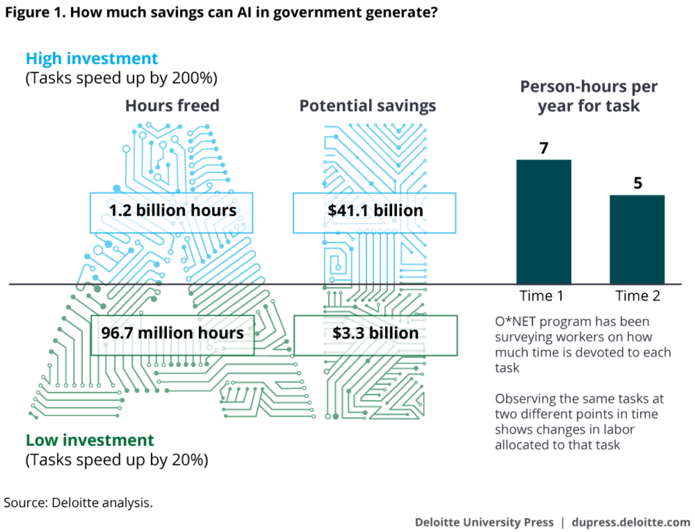

Figure 1: Credit: Deloite Analysis

The (W^1): What is Intelligence?

Intelligence: the ability to learn and perform suitable techniques to solve problem.

Figure 2: credit:Lindah Mavengere

A fully pre-programmed factory robot is flexible, accurate, and consistent but not intelligent.

The W^1: What is Artificial Intelligence

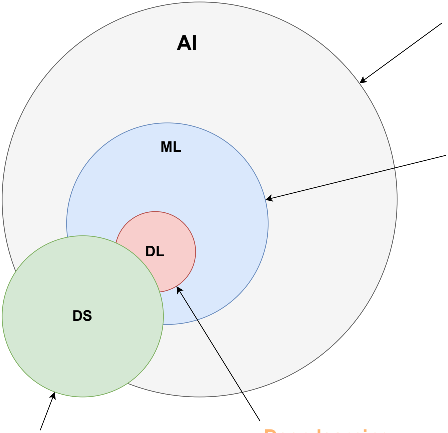

Artificial Intelligence

- the science and engineering of making intelligent machines.
- machines that can learn, at least somewhat as human beings do.

Machine Learning

- A subset of AI studying how computer agents can improve their perception, knowledge, thinking, or actions based on experience or data.

Supervised-ML

- A computer learns from input-output pairs in order to predict the output of the new input.

Unsupervised-ML

- Discover the pattern and learn the structure of unlabeled data.

Reinforcement-ML

- Agent learn action sequences that optimize its total rewards

Data science

- Understanding or making sense of data e.g visualization, exploratory analysis, data products etc

Deep learning

- A subset of ML that use cascade multi-layer (artificial) neural networks of non-linear preprocessing for feature learning and pattern recognition.

Figure 3: AI defined

The W^1: What is ML?

Machine learning (ML): the science (and art) of programming computers so they can learn from data.

Learn from data

- Automatically detect patterns in data and
- Build models that explain the world
- Use the uncovered pattern to understand what is happening (inference) and to predict what will happen(prediction).

This gives computers the ability to learn without being explicitly programmed.

The W^2: Why AI?

Consider how you would write a spam filter using traditional programming techniques.

The W^2:Why ML?

- Hard problems in high dimensions, like many modern CV or NLP problems require complex models $\Rightarrow$ difcult to program the correct behavior by hand.
- Machines can discover hidden, non-obvious patterns.
- A system might need to adapt to a changing environment.
- A learning algorithm might be able to perform better than its human programmers.

The W^3: AI in Energy

Energy forecasting (Demand and supply)

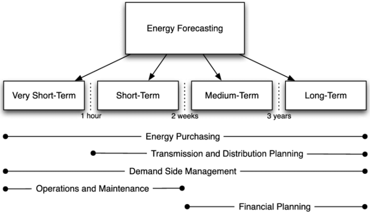

SDG-7 monitoring and mapping

Grid and microgrids planning and management

Improve Energy literacy and efficiency

Watt's up at Home? Smart Meter Data Analytics from a Consumer-Centric Perspective

by Benjamin Völker [1], Andreas Reinhardt [2], Anthony Faustine [3] and Lucas Pereira [4,*]

[1] Chair of Computer Architecture, University of Freiburg, 79110 Freiburg, Germany
[2] Department of Informatics, TU Clausthal, 38678 Clausthal-Zellerfeld, Germany
[3] Center for Artificial Intelligence (CeADAR), University College of Dublin, D04 V1W8 Dublin 4, Ireland
[4] ITI, LARSyS, Técnico Lisboa, 1049-001 Lisboa, Portugal
* Author to whom correspondence should be addressed.

The W^3: AI and EO for Public Goods

Urban Planning and Management

Improve Agriculture productivity and enhance Food Security

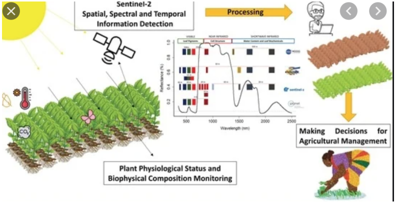

Natural resources monitoring and management

Improve National Statics and SDGs monitoring

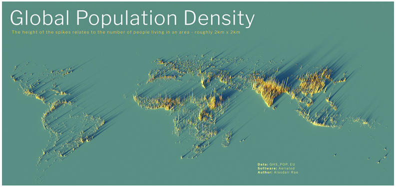

The W^4:When to apply AI

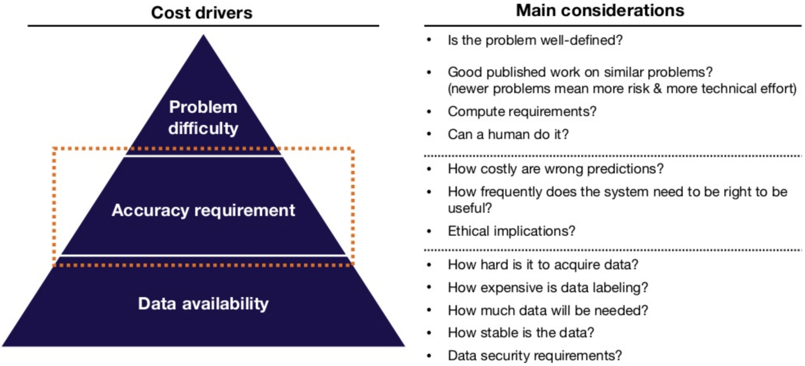

Figure 4: credit:Full stack DL

More reference:cs329s ML System Design Lecture #2 part 5

The H: Data and Resources

Data

- Data strategy that support the use of data-driven solutions.
- Data development that follow acceptable standards and best practices.
- Development of tools and platform for data sharing/hosting

Resources

- Invest in upgrading legacy IT systems.
- Invest on necessary computing resources for an AI project.
- Hire and retain AI talents (Juniors and Seniors).
- Engage research centers and universities.

The H:Research and Collaboration

Research

- Invest on Research and Development (R&D)
- Develop research-driven innovation.
- Encourage experimental-mindset.
- Invest on capacity and capabilities required to drive research-innovation agenda.

Collaborate

- Collaborate with domains experts (local and international).
- Research institute-govt partnership project.
- Collaborate with research centres and companies leading in AI (local and international)

The H:Strategies

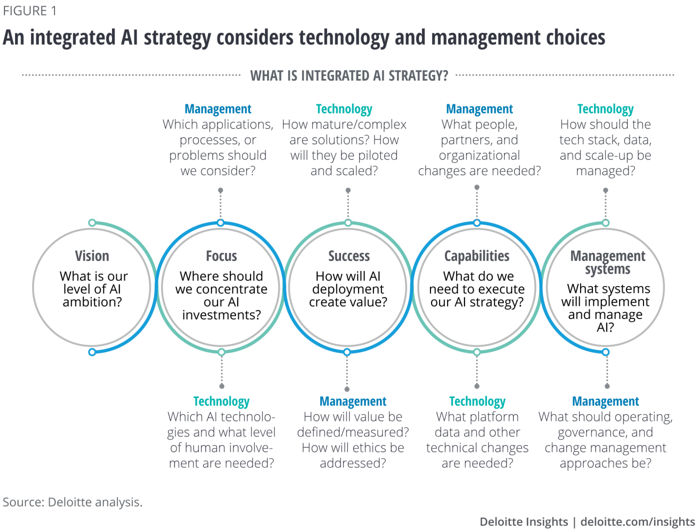

Figure 5: Credit:Deloite:Crafting an AI strategy for government leaders

The H:AI Policy

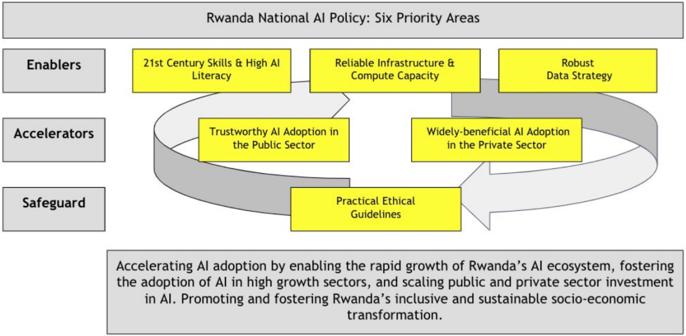

Figure 6: credit:RwandaICT

More reference: National AI Policy Making

The H: AI readiness

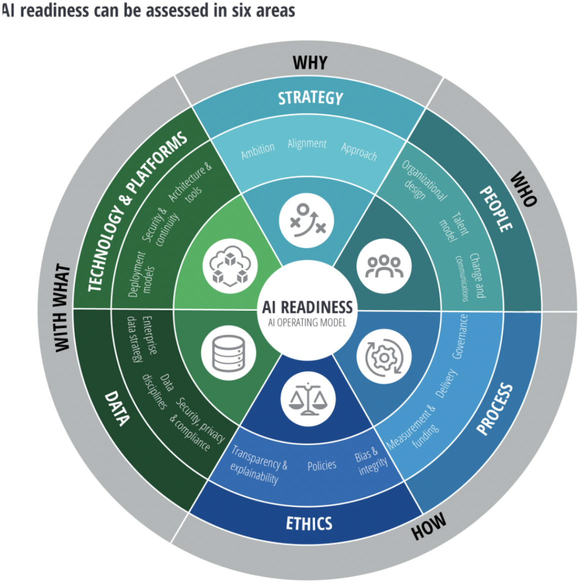

Figure 7: credit:deloite:AI readiness for government

ML life-cycle

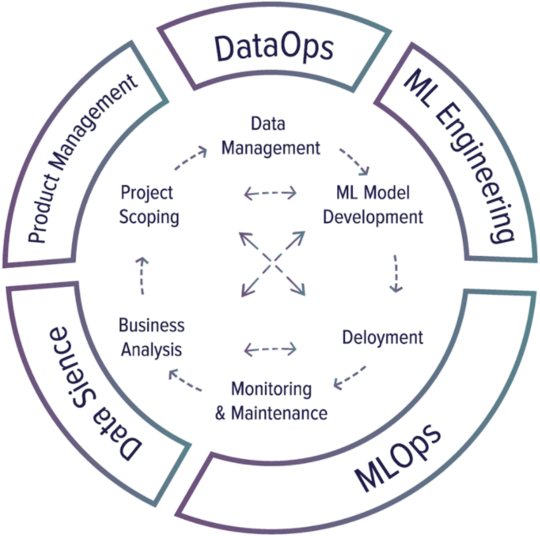

Figure 8: credit: CS 329S: Machine Learning Systems Design

Where can I contribute?

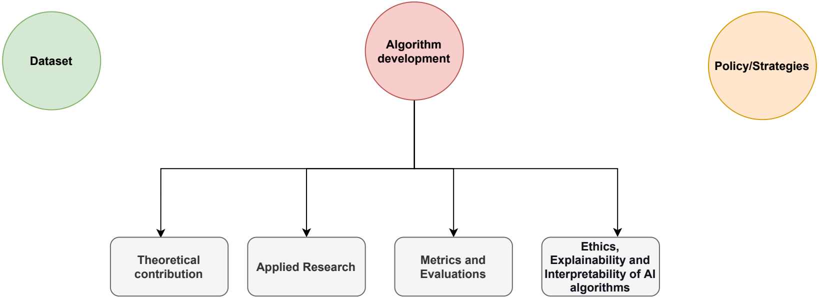

Figure 9: Contributions

Identify open and specific problem

Establish hypothesis about the problem.

Review literature

- Identify open questions that need answers.
- Learn about common methods, datasets and libraries.
- Lay out goals & objectives, constraints, and evaluation criteria.

Identify datasets

Identify data-set(s) to benchmark your solution.

At least one dataset that appeared in related prior work.

Where to find datasets

- Build them.
- Scrape them.
- Find them (contact authors).
- Generate them (artificial data).

Take time to understand your data: exploratory analysis.

Establish baseline

Define your baseline:Any publishable performance with simplest approach.

Write code quickly

- data-pipeline.
- Training-evalution-pipeline.
- Analysis-pipeline.

First get a baseline running $\Rightarrow$ this is good research practise.

Run Experiments

Track experiments: Take notes of what each experiment was meant to test.

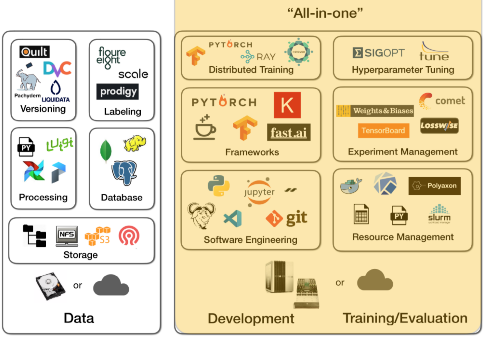

Figure 10: credit:Full stack DL

- Use recommended best practices for managing and monitoring ML experiments.
- Version code, data and ML experiments.
- Use existing frameworks and tools.

More references: CS329: Versioning & experiment tracking

Experiment Evaluation

When the cook tastes the soup that is formative; when the customer testes that is summative.

Formative evaluation

- They guide further investigations.
- Compare design option A to B, tune hyper-parameters etc.

Summative evaluation

- compare your approach to previous approaches,
- compare different major variants of your approach.

Don't save all your qualitative evaluation for the summative evaluation.

Experiment Evaluation

Follow prior work precisely in how to choose and implement main evaluation metric.

Quantitative evaluation

- Show metric as many variants of your model as you can.
- Test for statistical significance (for highly variable models or small difference performance).
- If your results are not significant. say so and explain what you found.

Experiment Evaluation

Conduct a thorough analysis

Qualitative-evaluation

- convince reader for your hypothesis $\Rightarrow$ look to prior work to get started.
- Show examples of system output.
- Plot how your model performance varies with the amount of data.
- Present error analysis.

Conclusion

References

- PATTERNS, PREDICTIONS, AND ACTIONS: A story about machine learning
- Probabilistic Machine Learning - a book series by Kevin Murphy
- Model-based Machine Learning, John Winn
- Foundations: How to design experiments in NLU