
DTU

Nicki Skafte Detlefsen, Associate Professor
Section for Cognitive Systems, DTU Compute

Introduction to MLOps

The Journey from Machine Learning Model to Machine Learning Product

Agenda

- 9-9:45: What is MLOps?
- 9:45-10:00: Break
- 10:00-10:30: Practical example + rambling
- 10:30-10:50: Trustworthy AI
- 10:50-11:00: Break (dance?)
- 11:00-11:15: Designing Machine Learning Systems
- 11:15-12:00: CANVAS time!

Who am I

Who am I

- 💡 Associate Professor at DTU Compute
- 💡 Research in Data-centric ML, MLOps, Efficient Machine Learning, Machine Learning Development
- 💡 Open-source developer
- 💡 Part-time ML Engineer at https://lightning.ai/
- 💡 Technical advisor for 4 DTU startups

= A good mix between the academic and industrial worlds

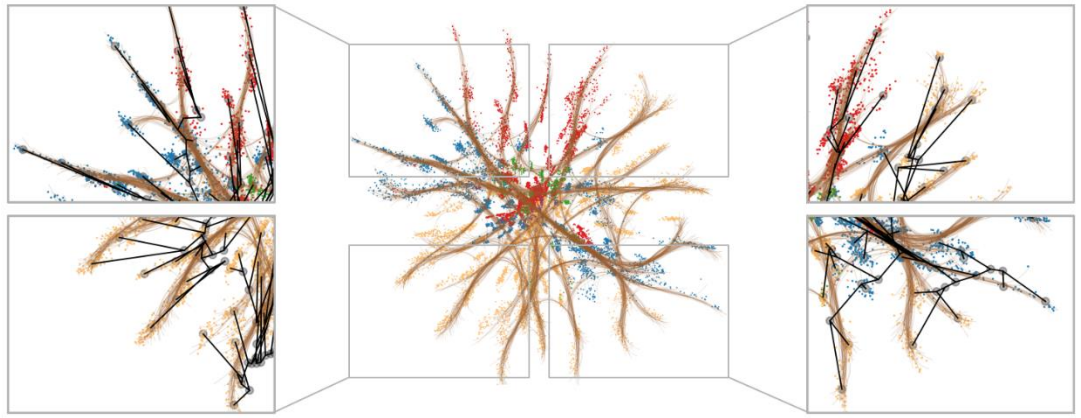

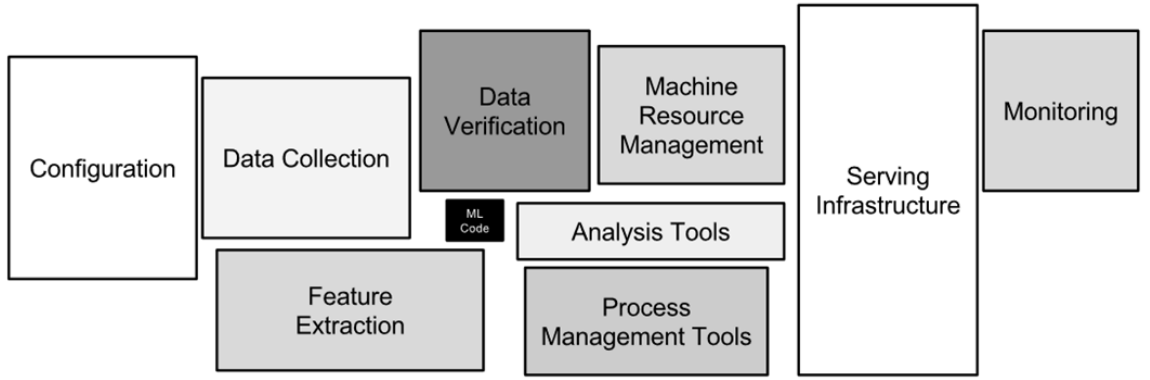

My secret identity

💻 Eager open-source contributor

- Numpy
- Scikit-learn
- Pytorch
- Matplotlib
- …

⚡ ML Engineer at https://lightning.ai/

- Pytorch-lightning
- Torchmetrics
- LitData
- LitGPT

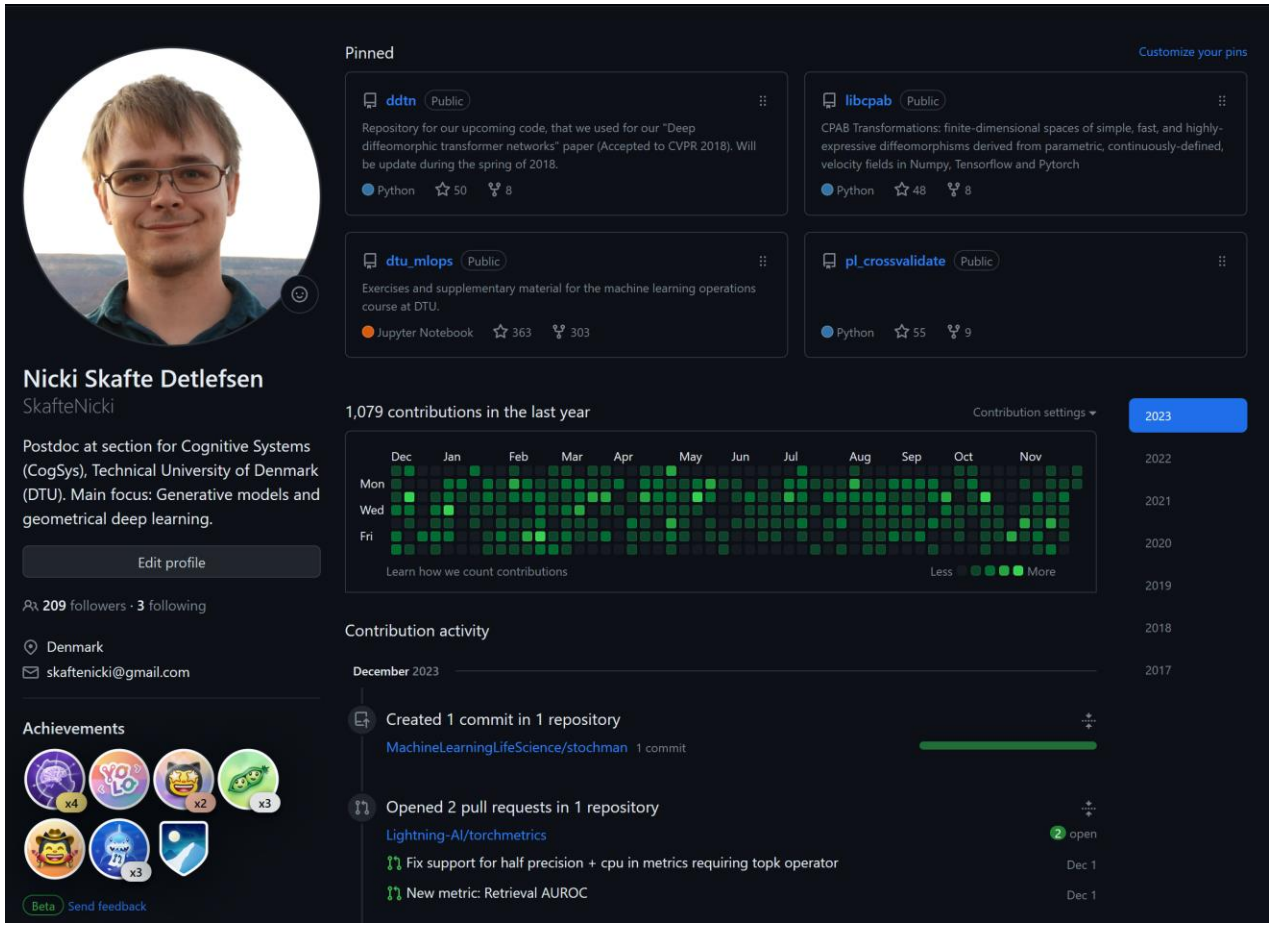

Mr. MLOps

https://skaftenicki.github.io/dtu_mlops/

✅ I been consulting companies for the last couple of years on how to implement AI in their businesses

✅ I done 20+ master project that have implemented some sort of ML I companies

✅ Technical advisor on 4 startups from DTU

Course information

- Course responsible
  - Postdoc Nicki Skafte Detlefsen, nsde@dtu.dk
  - Professor Søren Hauberg, sohau@dtu.dk
- 5 ECTS (European Credit Transfer System), corresponding to 140 hours of work
- 3 week period of January 2023
- Master course

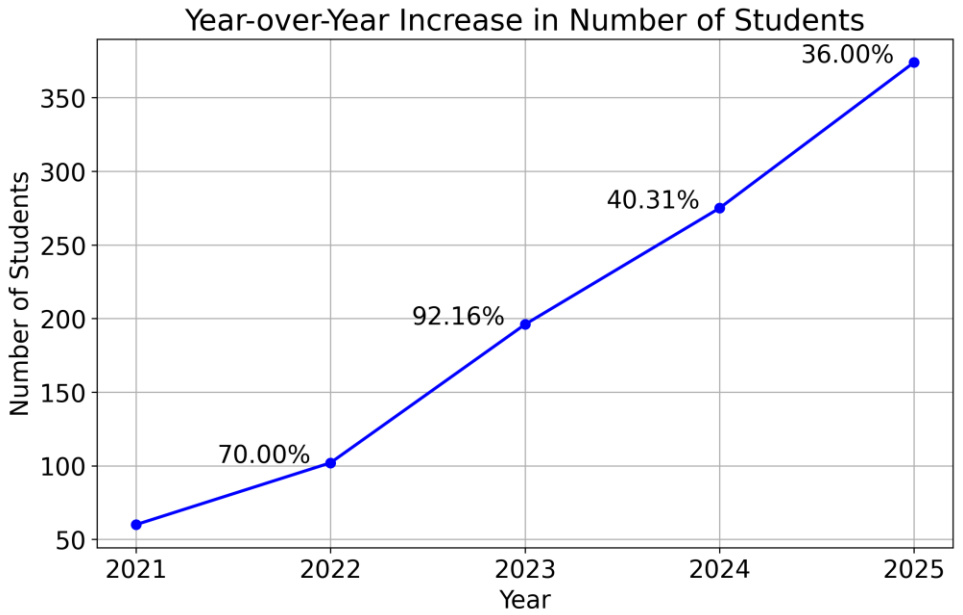

A definition of MLOps

What is AI?

What is AI?

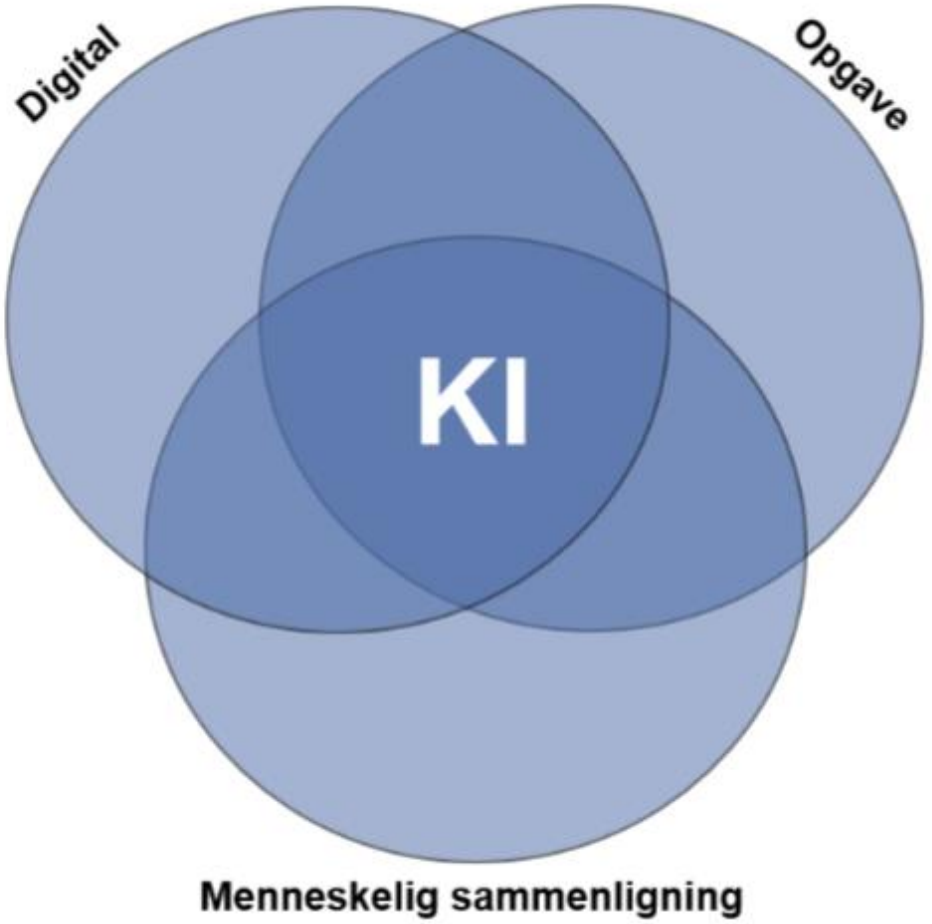

[1] Thomas Ritter, Kunstig Intelligens eller Kaotisk Idioti

[1] Thomas Ritter, Kunstig Intelligens eller Kaotisk Idioti

Let's agree that ML/AI is fantastic

Chatbot

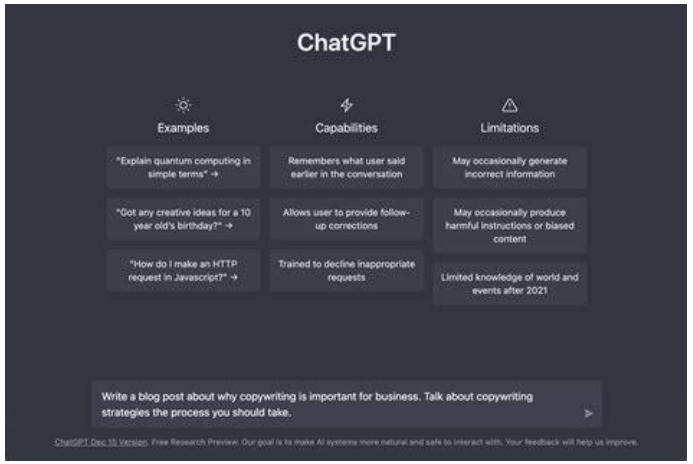

Self-driving cars

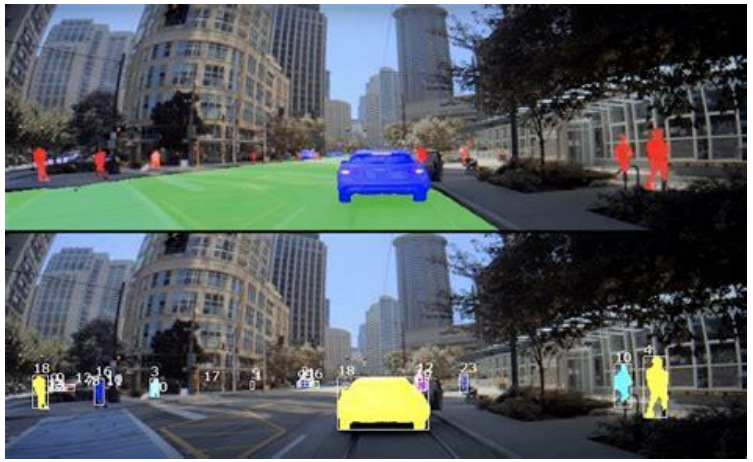

Efficient diagnoses

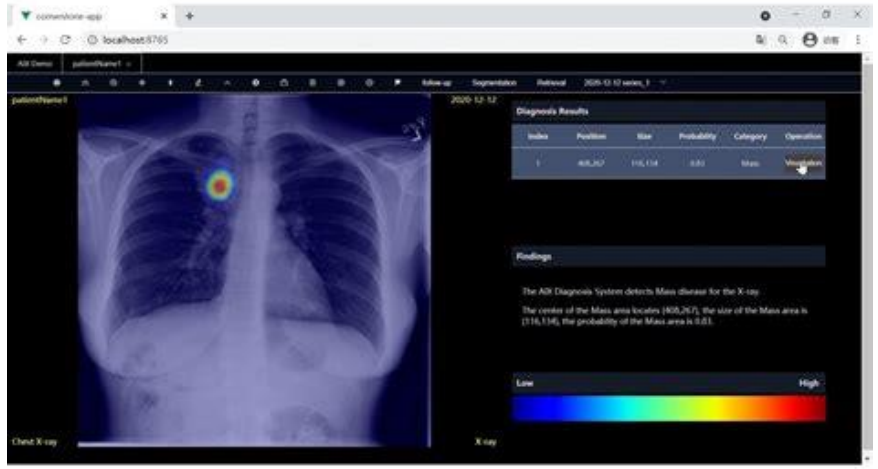

- 💡 AI is a key component of what we call industry 5.0
- 💡 It can solve problems on unprecedented scales

[1] https://chatgpt.com/
[2] https://blogs.nvidia.com/blog/drive-labs-panoptic-segmentation/
[3] https://www.innovationhub.hk/article/aixchest-x-ray-screening-with-ai

But errors does happen

Chatbot

Self-driving cars

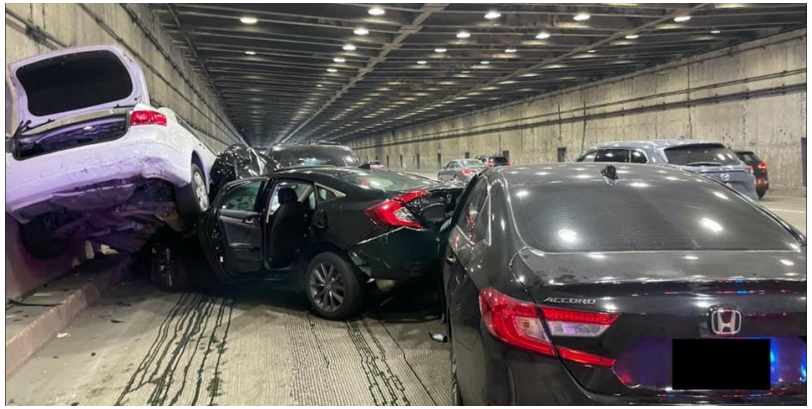

Efficient diagnoses

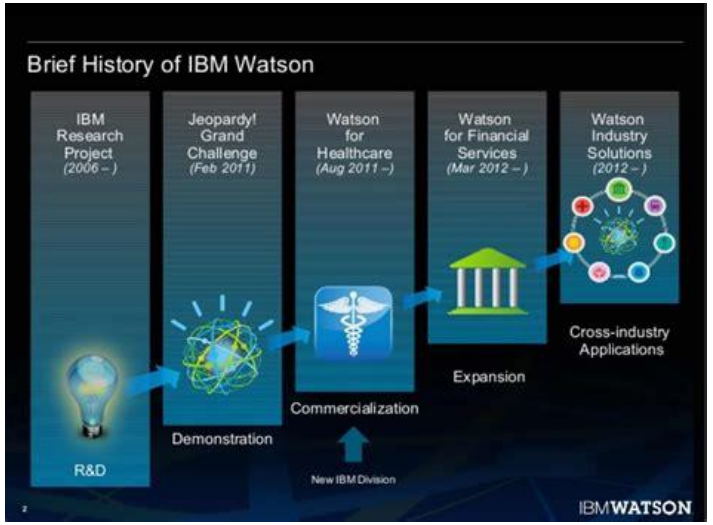

- 💡 AI, similar to humans, sometimes takes the wrong decision
- 💡 But due to its unpresented scale, things can quickly go very wrong

[1] https://dailywireless.org/internet/what-happened-to-microsoft-tay-ai-chatbot/
[2] https://www.businessinsider.nl/video-shows-8-car-pileup-after-a-tesla-allegedly-using-full-self-driving-stopped-in-a-highway-tunnel/
[3] https://www.nextplatform.com/2019/02/14/ibm-mashes-up-powerai-and-watson-machine-learning-stacks/

The duality of AI

Developing AI is easy – Running AI in production is hard

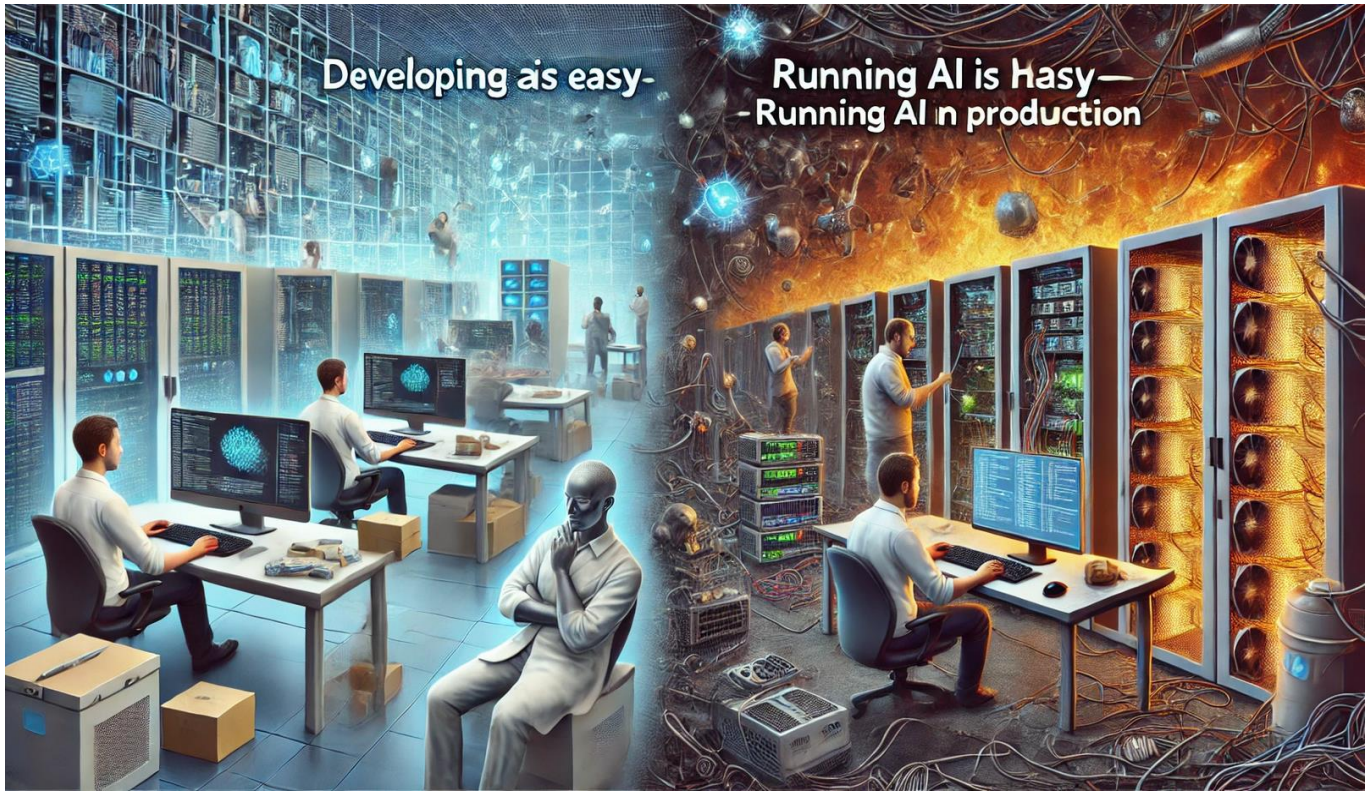

AI-Generated

This is not a new

Why is this?

💡 Because AI is just software at the end of the day

Being an AI developer <-> Being a software developer

✅ We have been developing software for 30+ years

✅ We have a lot of tools for software development

❌ Software can break

❌ Software needs to be maintained

Let's therefore try to understand AI in the context of software development

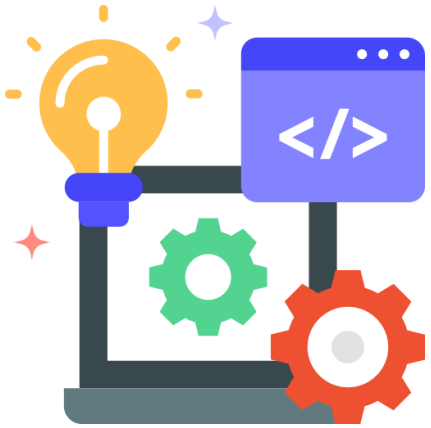

Let's start where it all began

Machine learning in production is fantastic

BUT

Massive technical debt is incurred if not careful

Hidden Technical Debt in Machine Learning Systems

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips
{dsculley,gholt,dgg,edavydov,toddphillips}@google.com
Google, Inc.

Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, Dan Dennison
{ebner,vchaudhary,mwyoung,jfcrespo,dennison}@google.com
Google, Inc.

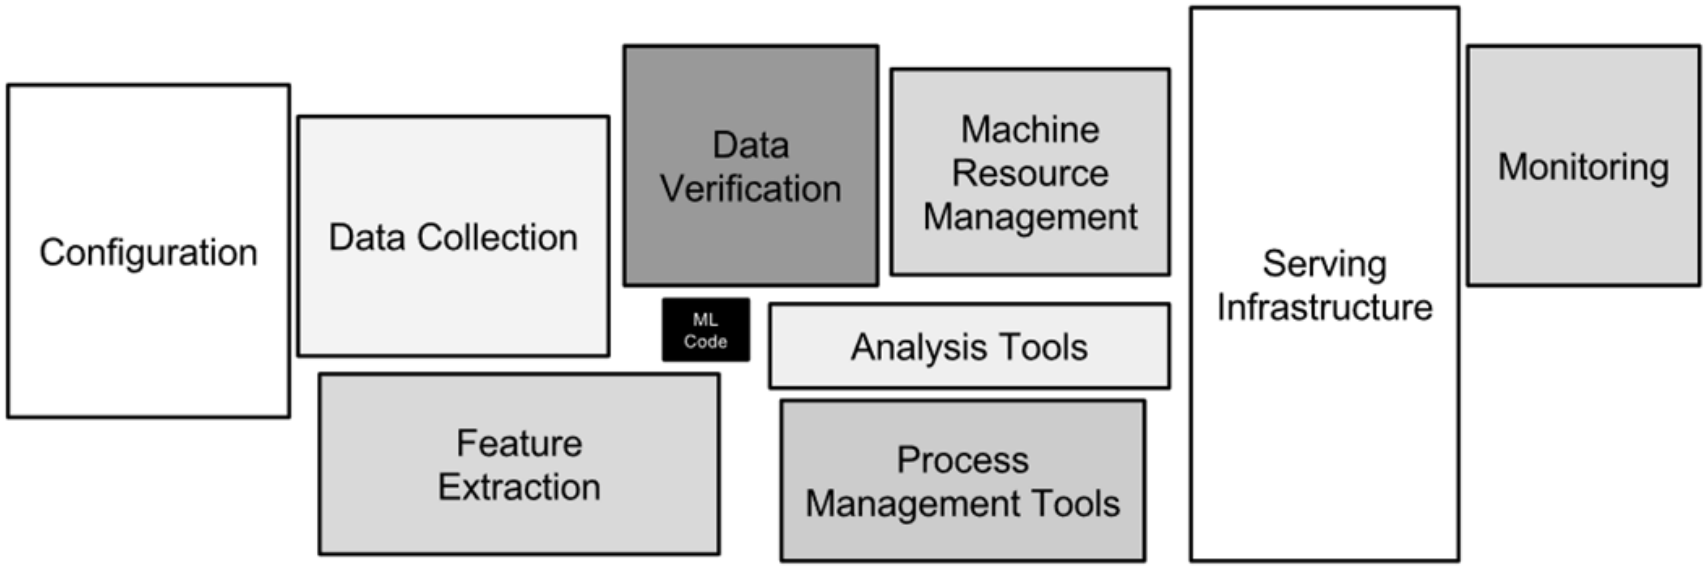

[1] Hidden Technical Debt in Machine Learning Systems, https://proceedings.neurips.cc/paper/2015/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf

Why do we focus on modelling?

Because we teach people it!

Courses / Projects are linear in nature

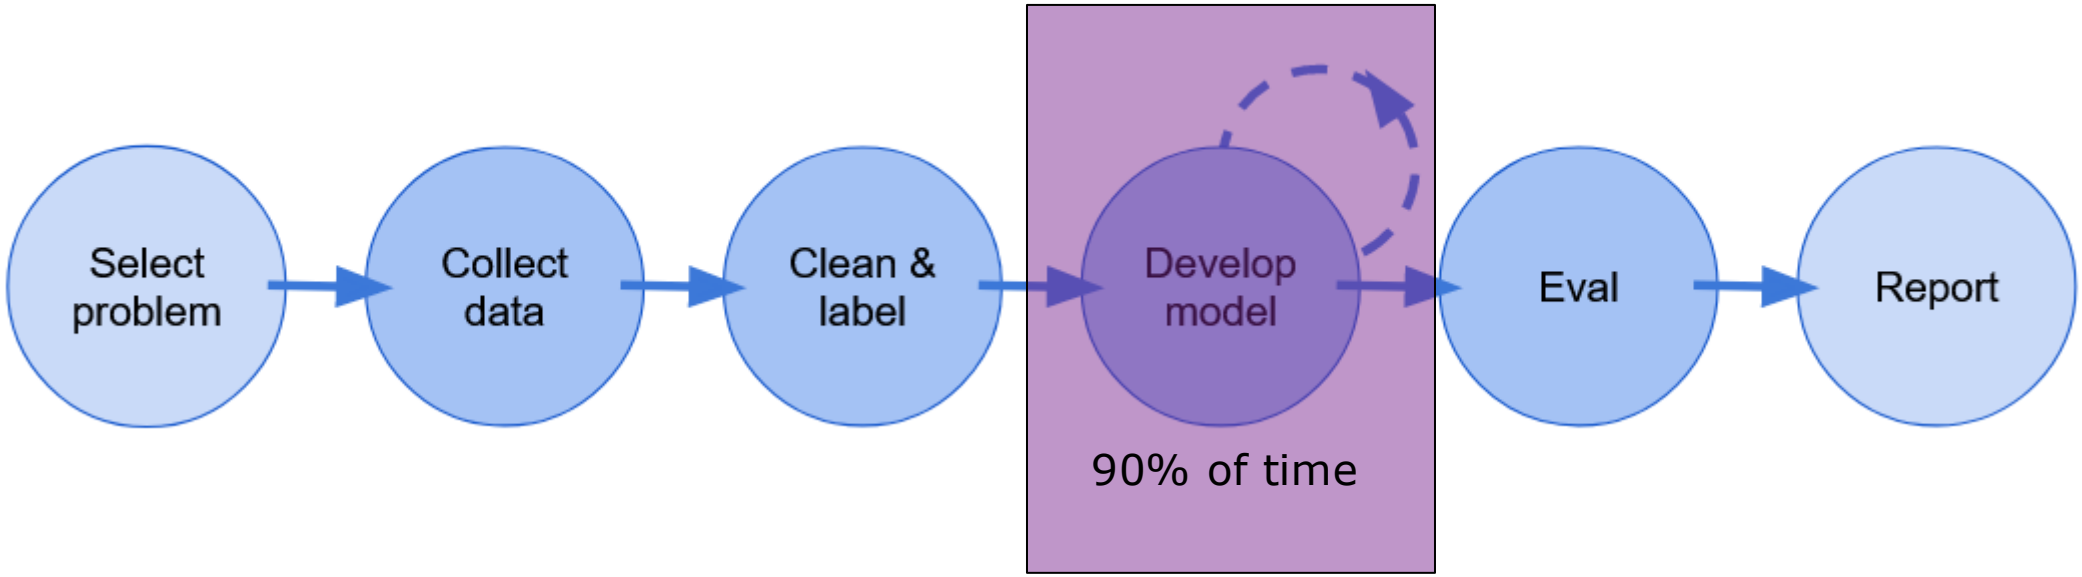

Feedback is grades / funding

[1] Full Stack Deep Learning course 2022, https://fullstackdeeplearning.com/course/2022/lecture-1-course-vision-and-when-to-use-ml/

Machine Learning in the real world

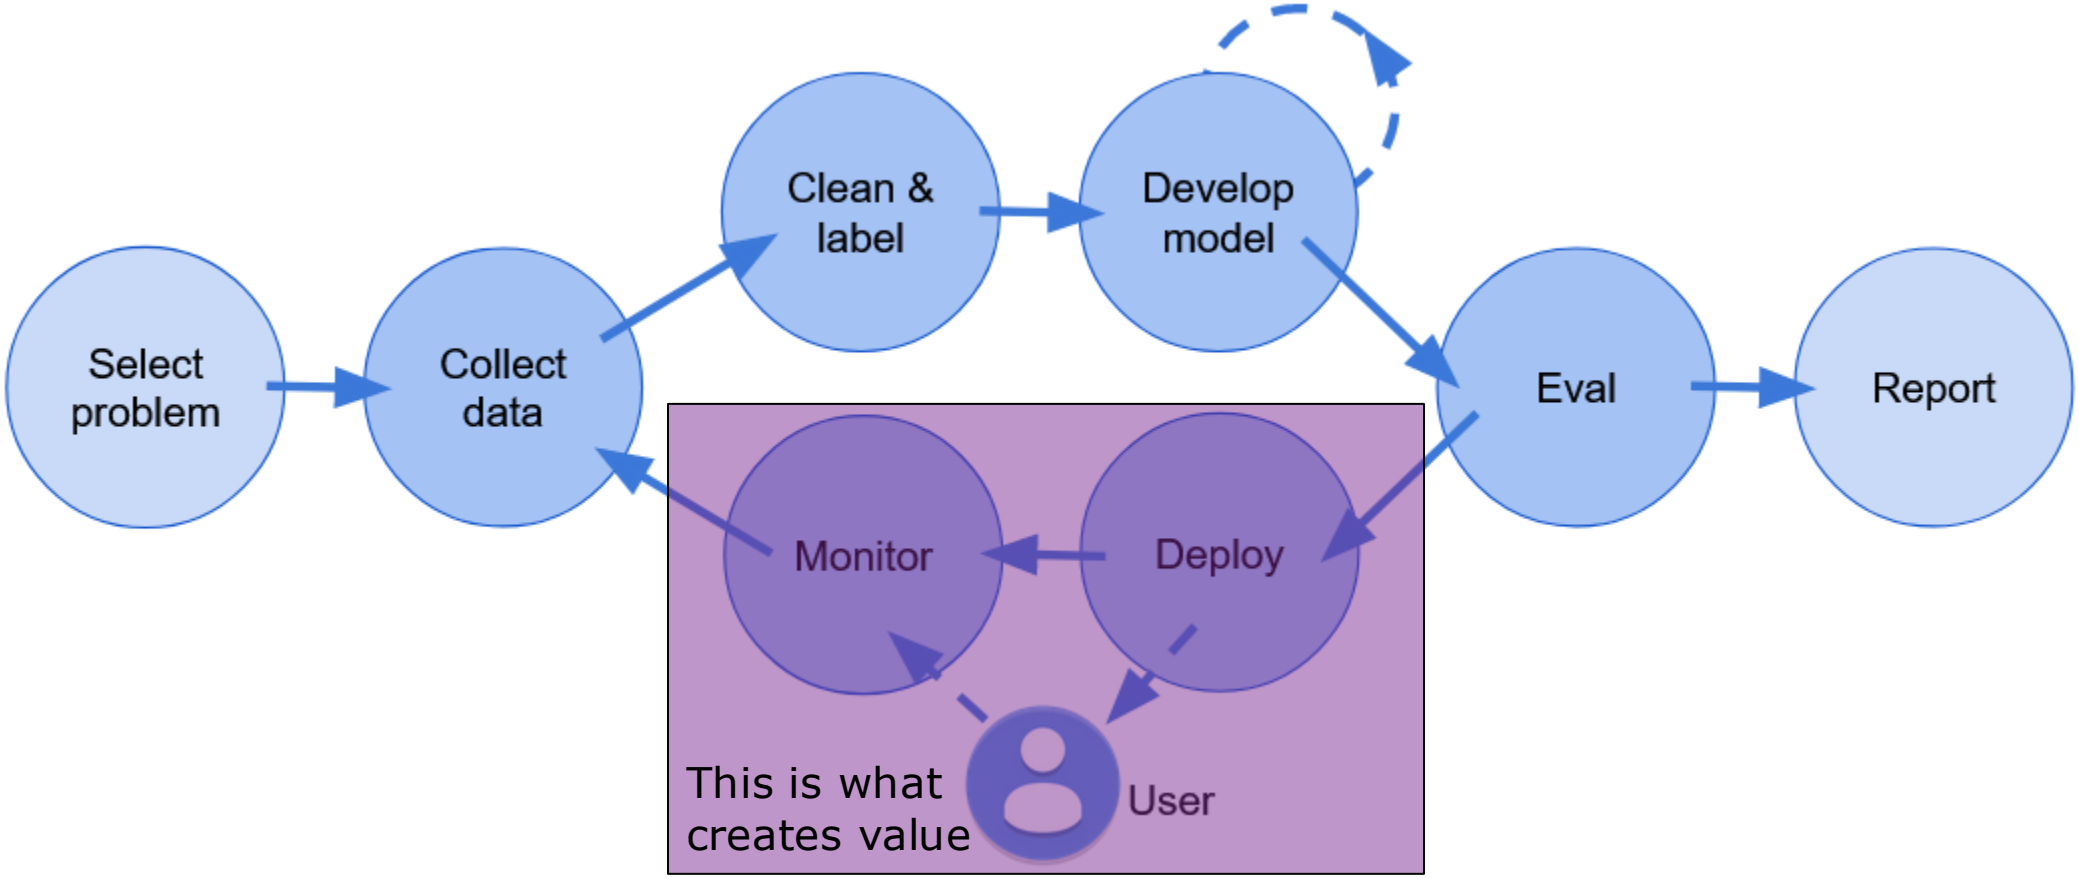

[1] Full Stack Deep Learning course 2022, https://fullstackdeeplearning.com/course/2022/lecture-1-course-vision-and-when-to-use-ml/

Key observations

1. Machine Learning in production is much more than doing ML modelling

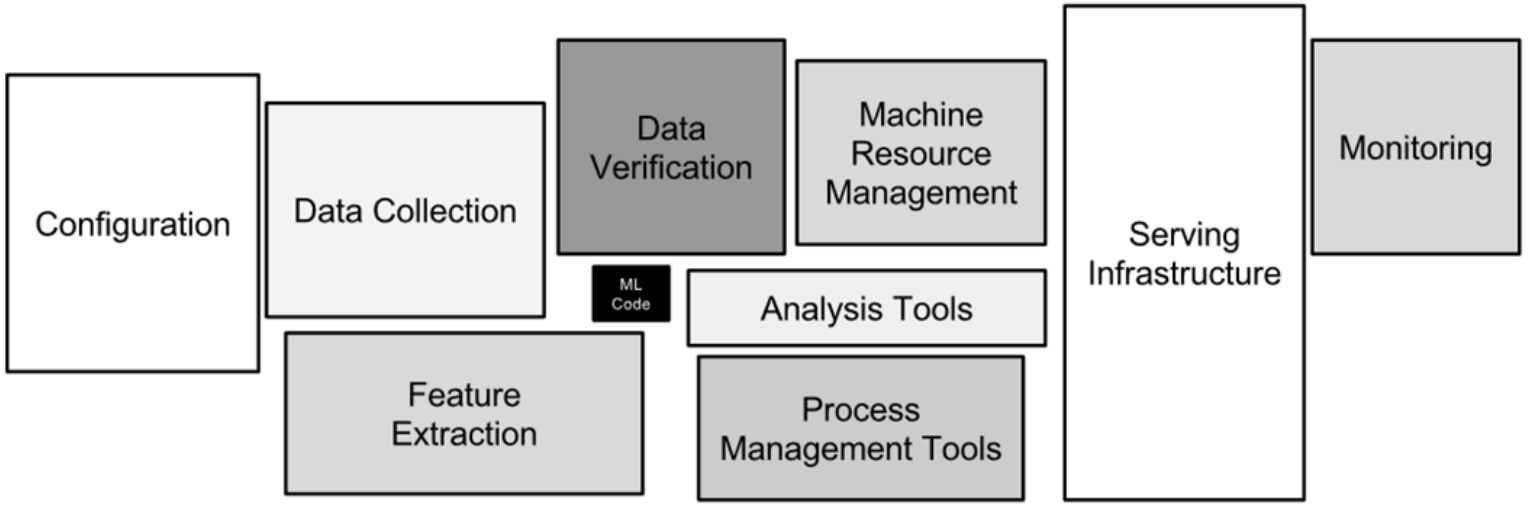

2. Machine Learning in production is a cycle

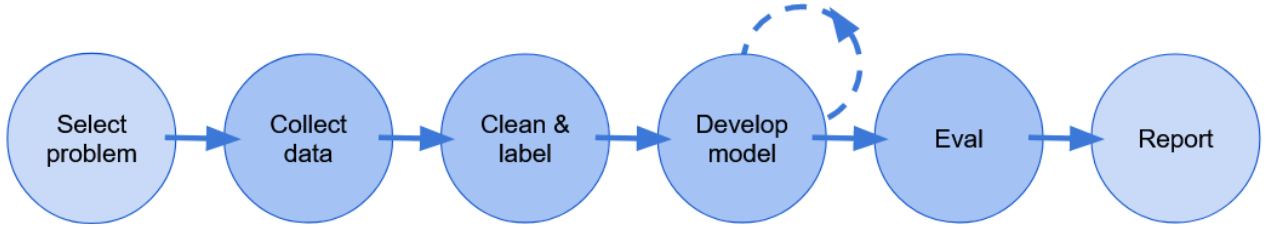

vs

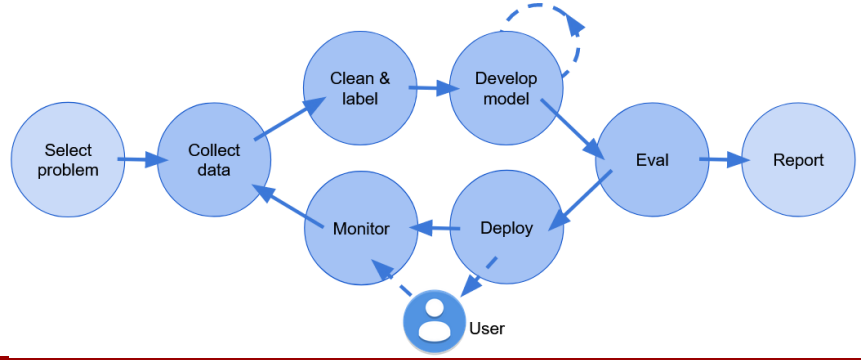

The other stuff is DevOps

DevOps = Developer operations

- Dates back to the late 1980s and early 1990s
- Gained popularity around 2007/2008 to remove the separation between software development and IT operations

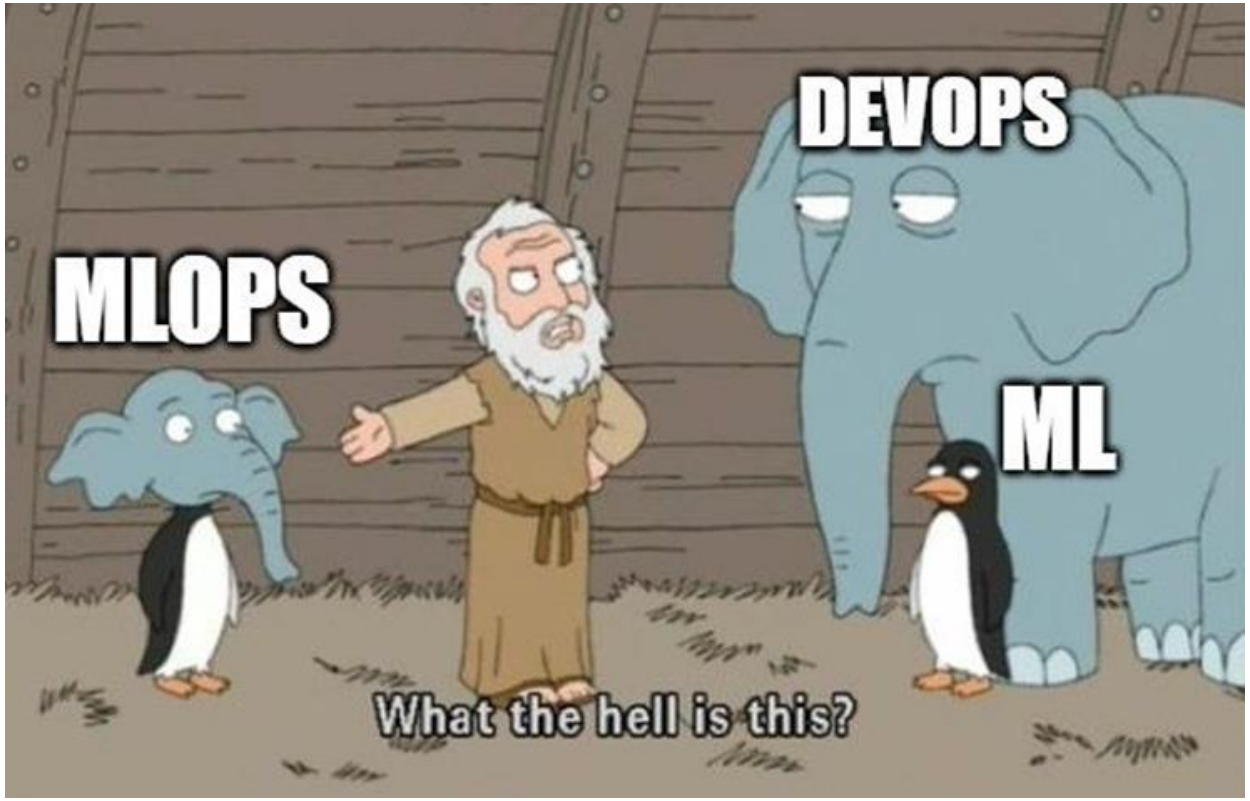

Hvad var det for et problem man identificeret?

There are two teams in software development:

- 💡 Dev team = focuses on developing and improving software
- 💡 Ops team = focuses on infrastructure and operationalizing the software

If these two don’t communicate, Dev might develop software that Ops can’t deploy — and Ops might set up the wrong infrastructure for what Dev is building.

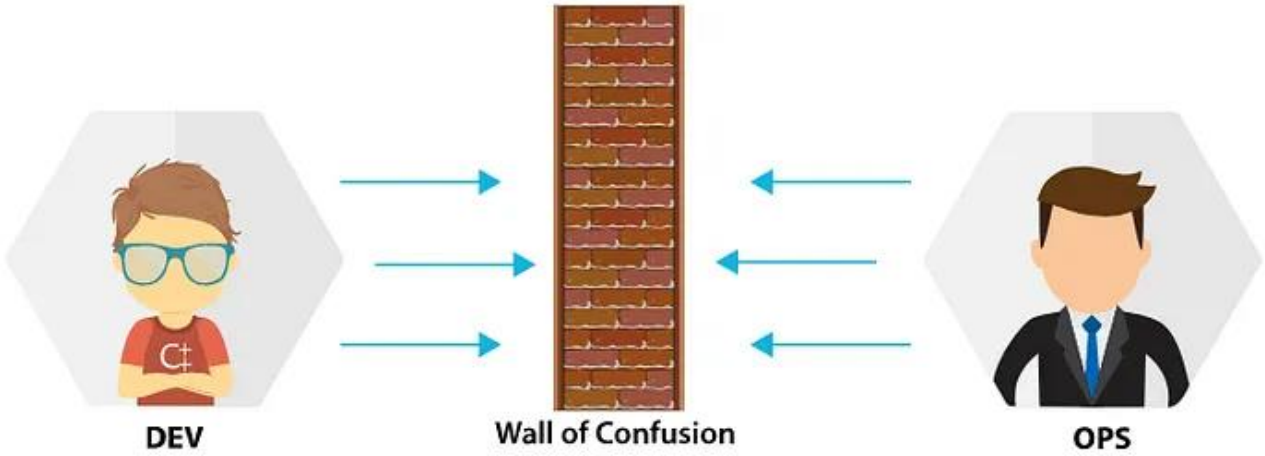

So, what is DevOps?

This is the closest to a definition that I could find:

DevOps is a set of practices that combines software development (*Dev*) and IT operations (*Ops*). It aims to shorten the systems development life cycle and provide continuous delivery with high software quality. It's an combination of human mindset, processes and technologies that continuously creates value.

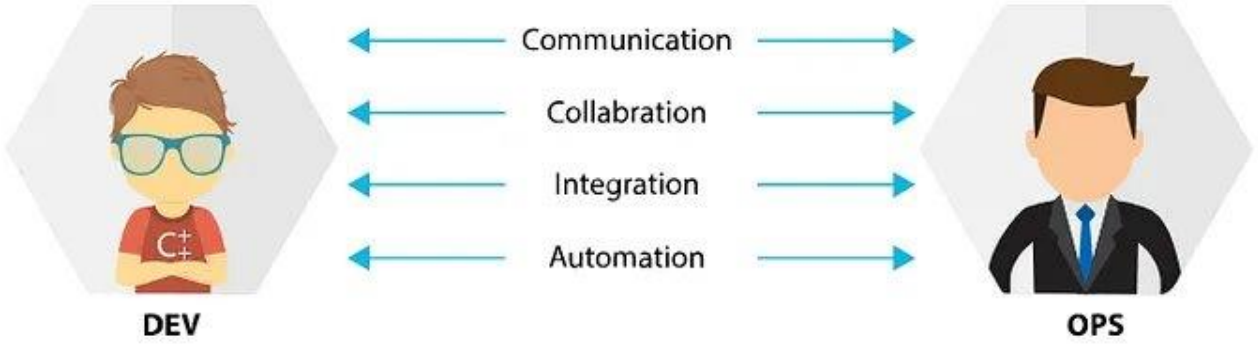

So, what is DevOps?

This is the closest to a definition that I could find:

DevOps is a set of practices that combines software development (*Dev*) and IT operations (*Ops*). It aims to shorten the systems development **life cycle** and provide continuous delivery with high software quality. It's a combination of human **mindset**, **processes** and **technologies** that continuously creates **value**.

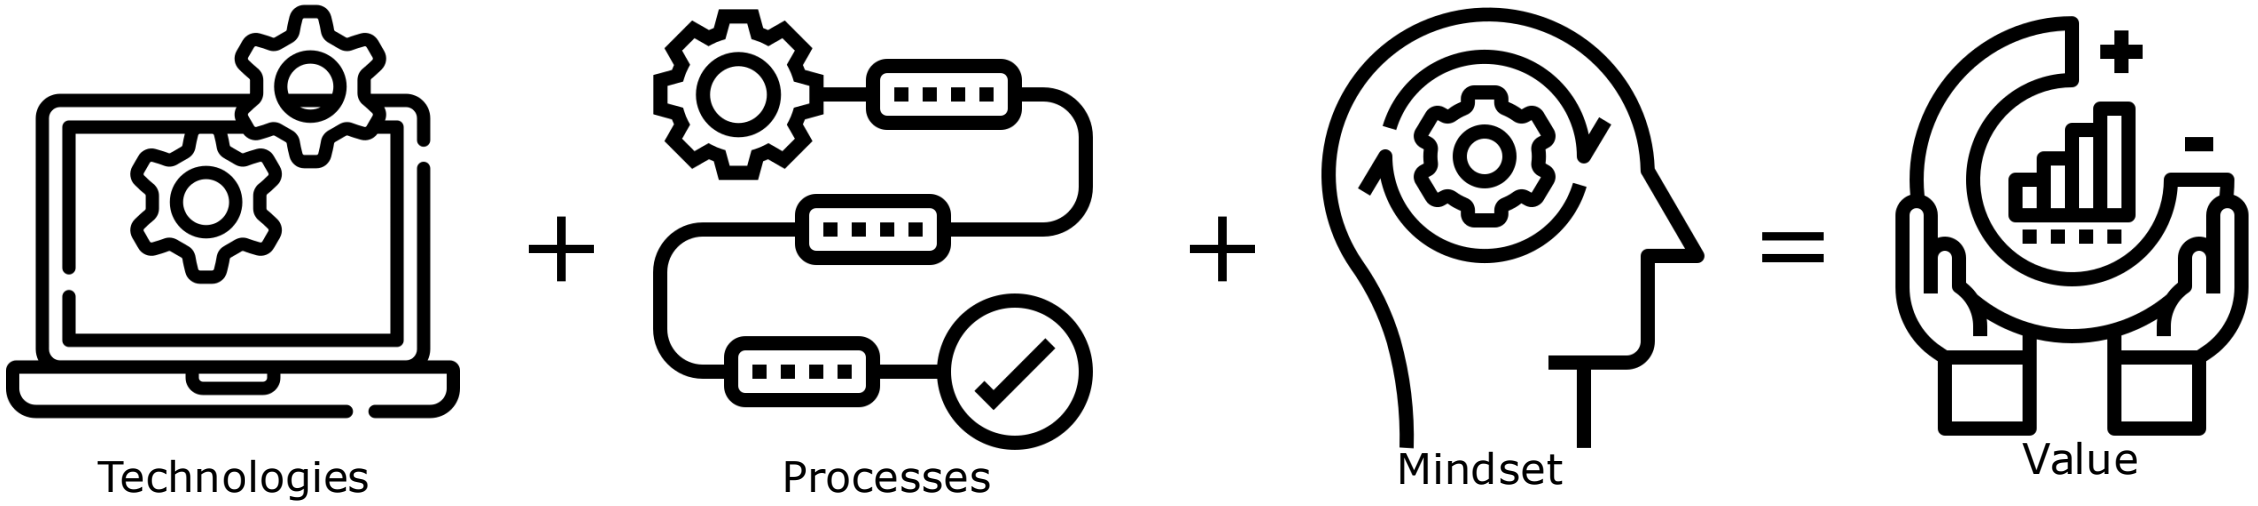

Technology, Processes, Mindset

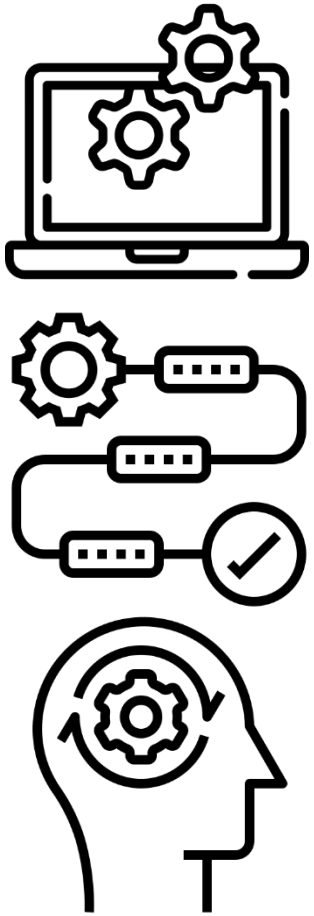

Use technologies that support the different parts of the lifecycle

Implement processes to make sure everyone is in sync about the lifecycle

Always consider all part of the lifecycle, not just its parts

But then MLOps must be…

Is a set of **tools**, **processes**, and **mindset** that aim to make **ML Lifecycle** *reproducible, trackable, testable and maintainable* to continuously create value.

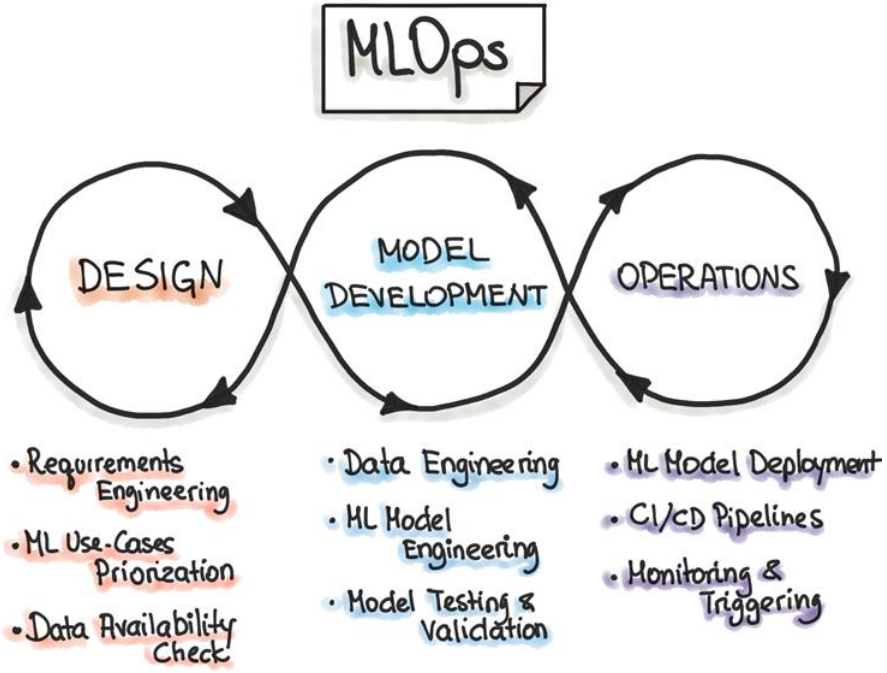

Define:

- 💡 Reproducible
- 💡 Trackable
- 💡 Testable
- 💡 Maintainable

Let's look at the different phases of MLOps

Data phase

- Business understanding
- Data understanding
- Designing the ML-powered software

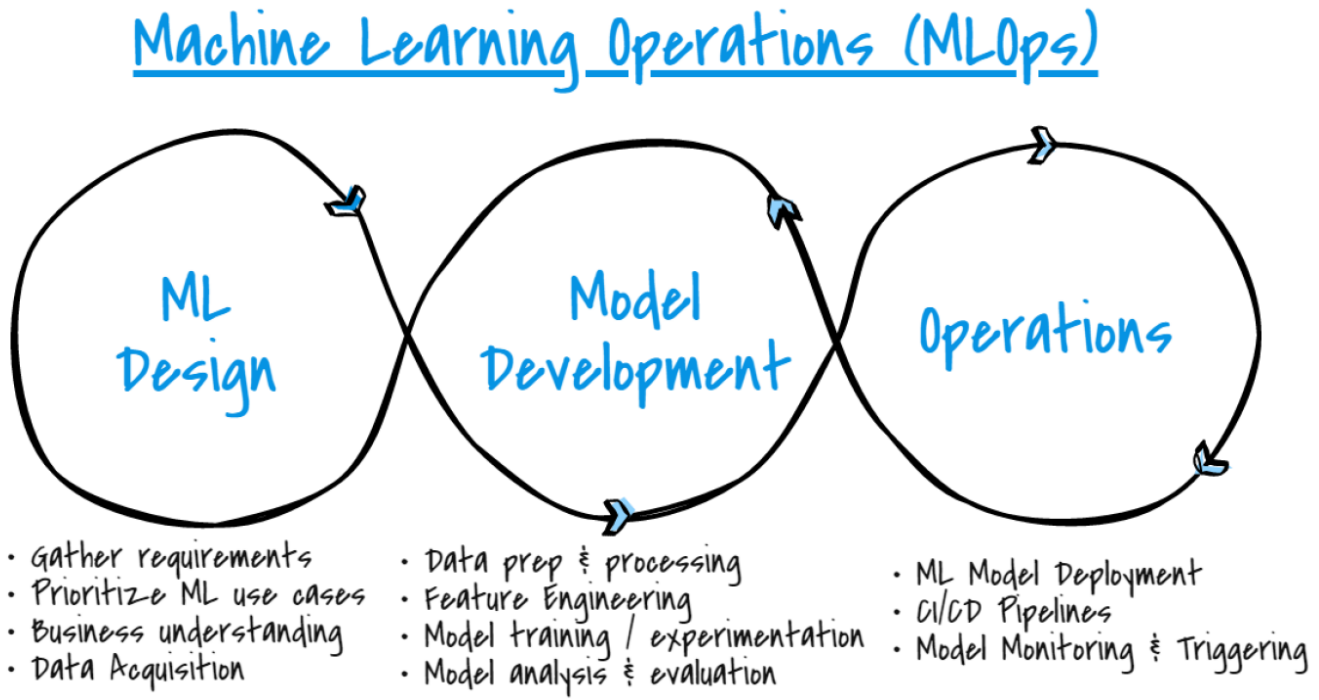

Model phase

- Model engineering
- Data engineering
- Deliver a stable quality ML model that we will run in production

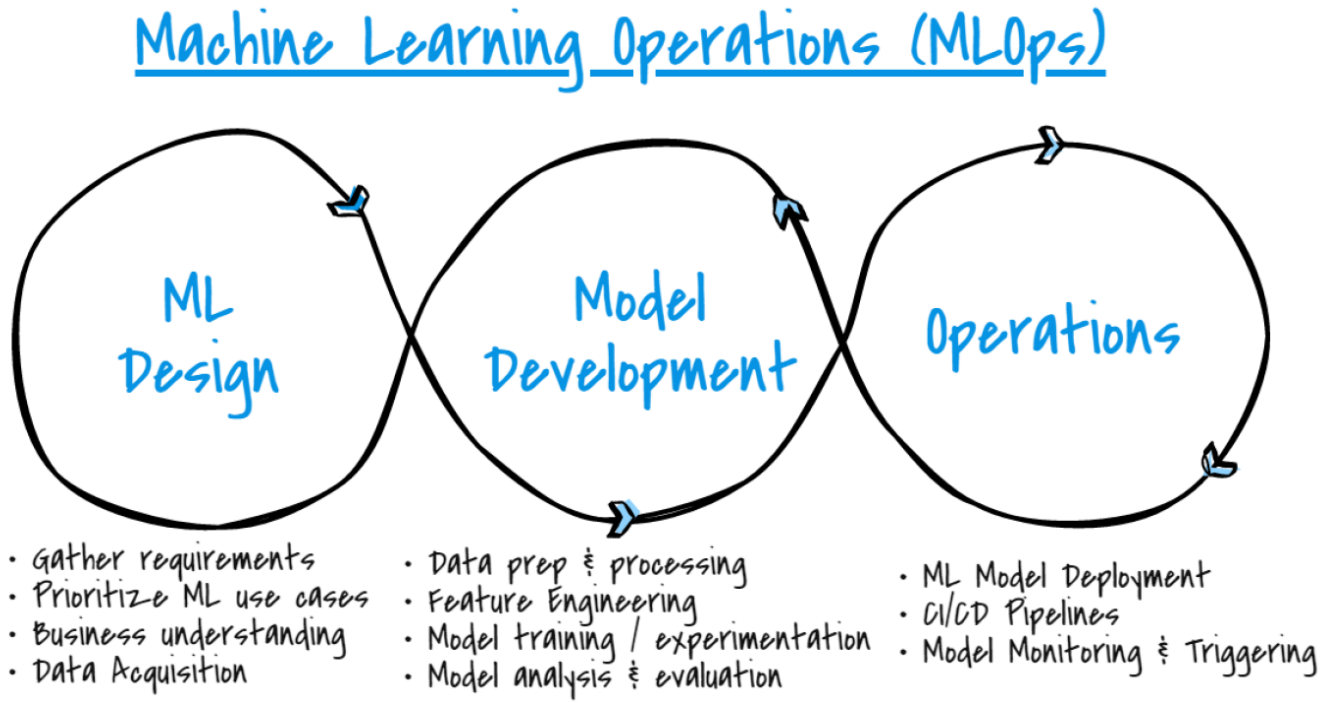

Operations phase

- Deliver the previously developed ML model in production
- Testing, versioning, continuous delivery, and monitoring

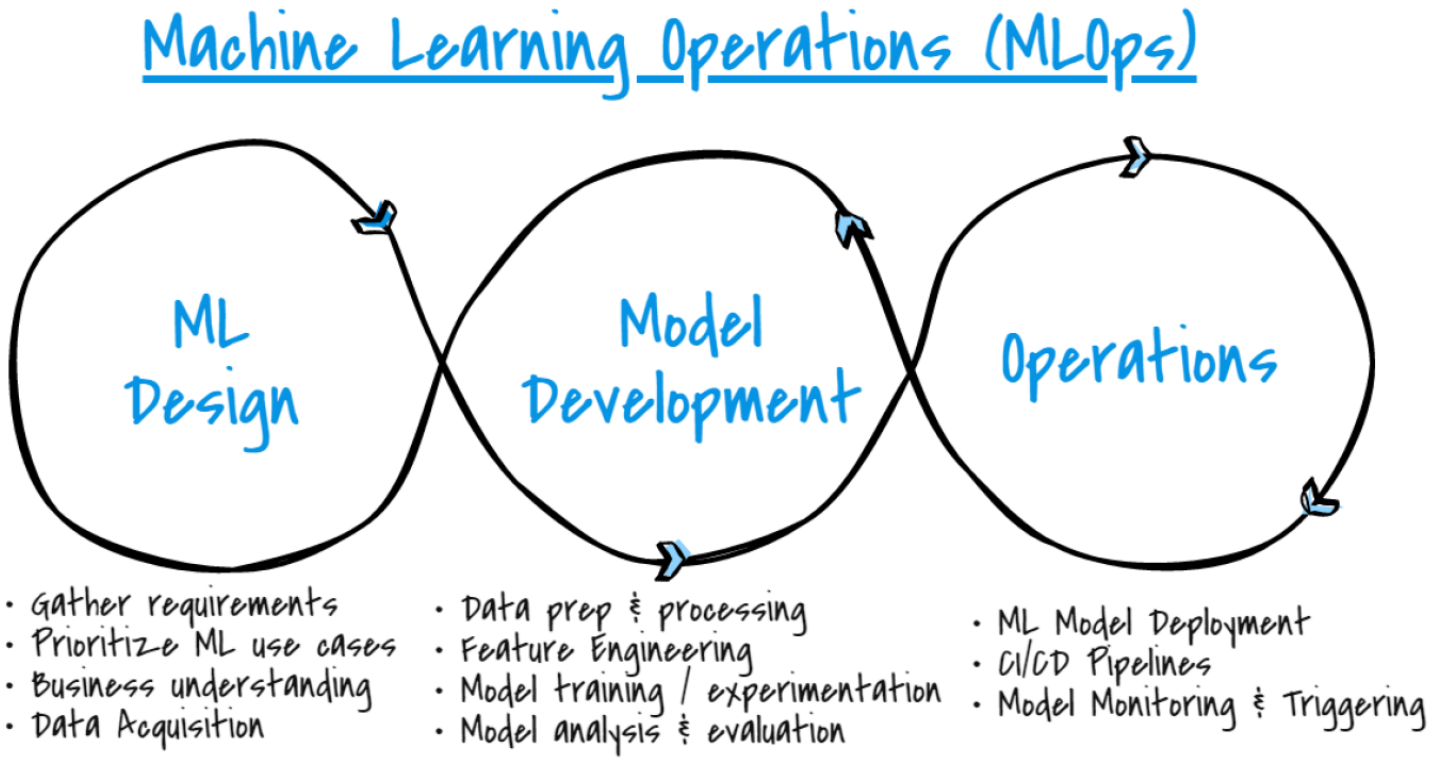

If DevOps exist, then why do we need MLOps?

Because data changes everything

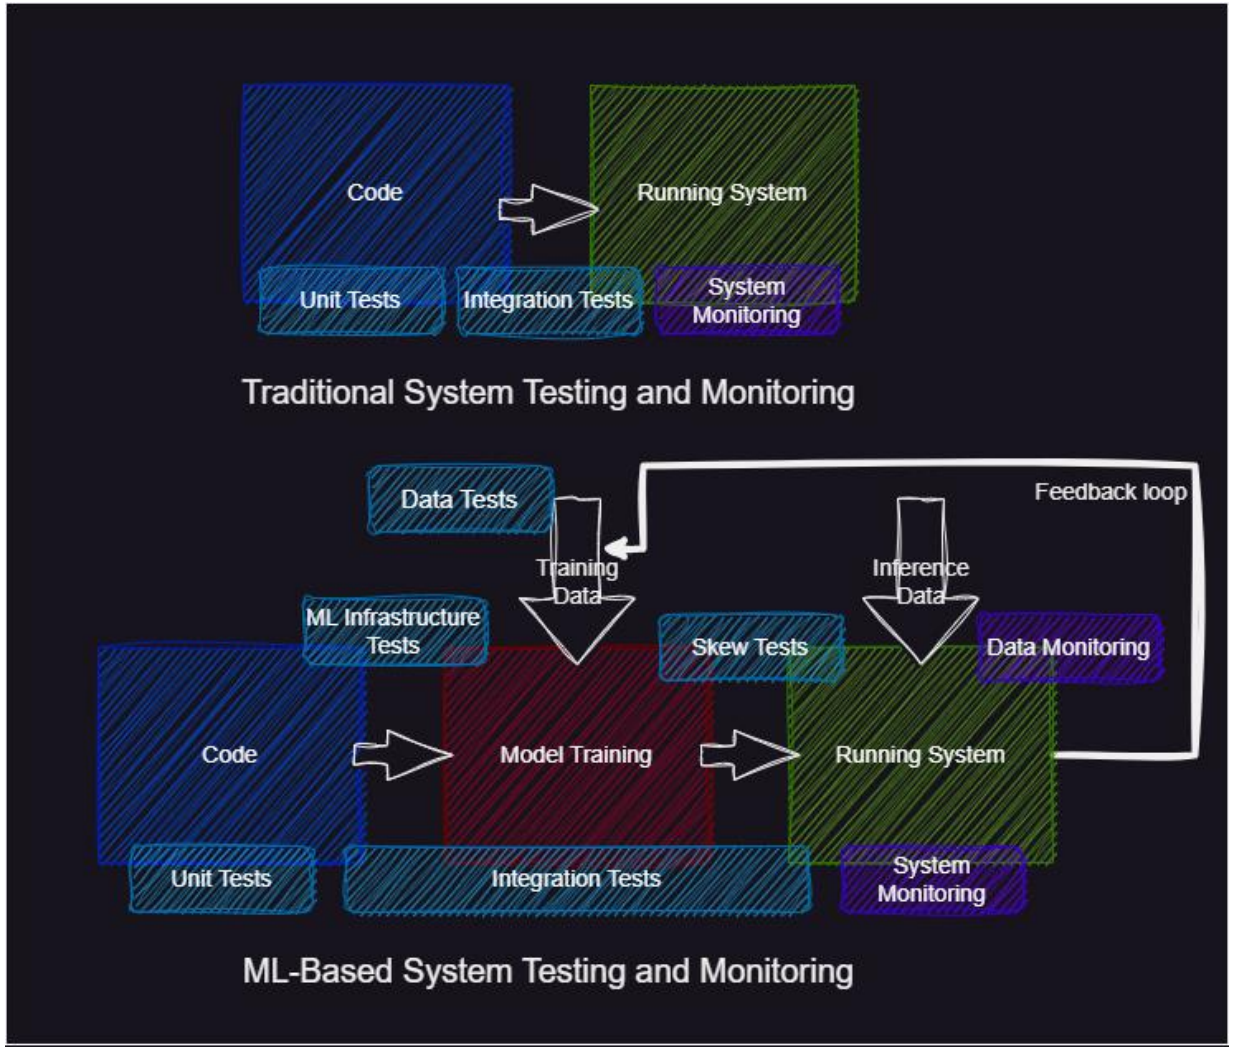

What is an MLOps engineer?

What makes an MLOps engineer?

A mix of

- Software developing
- Machine Learning
- Data engineering

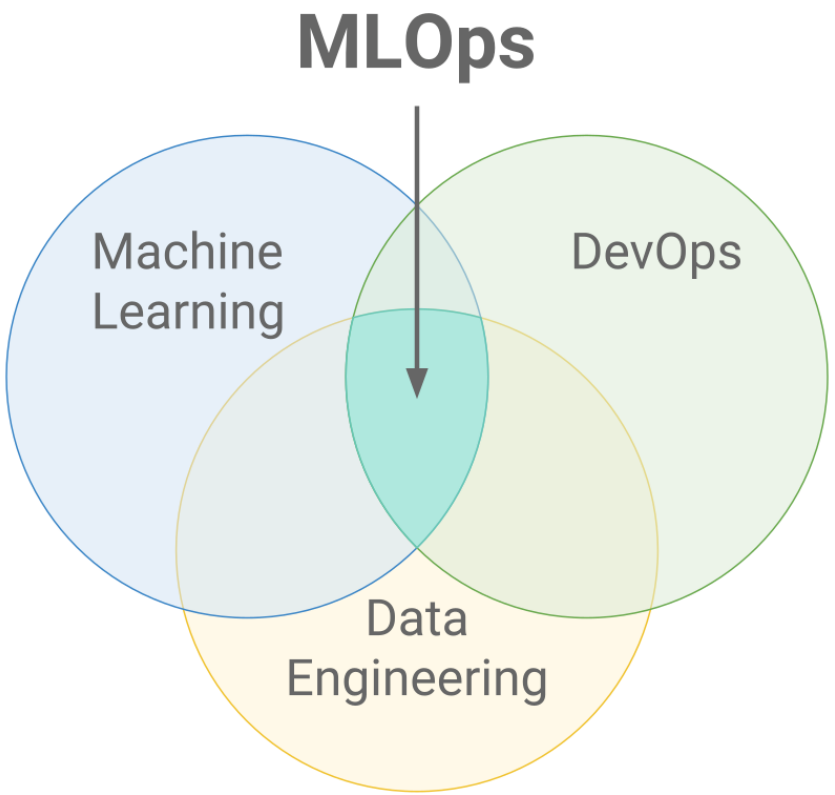

Where's waldo?

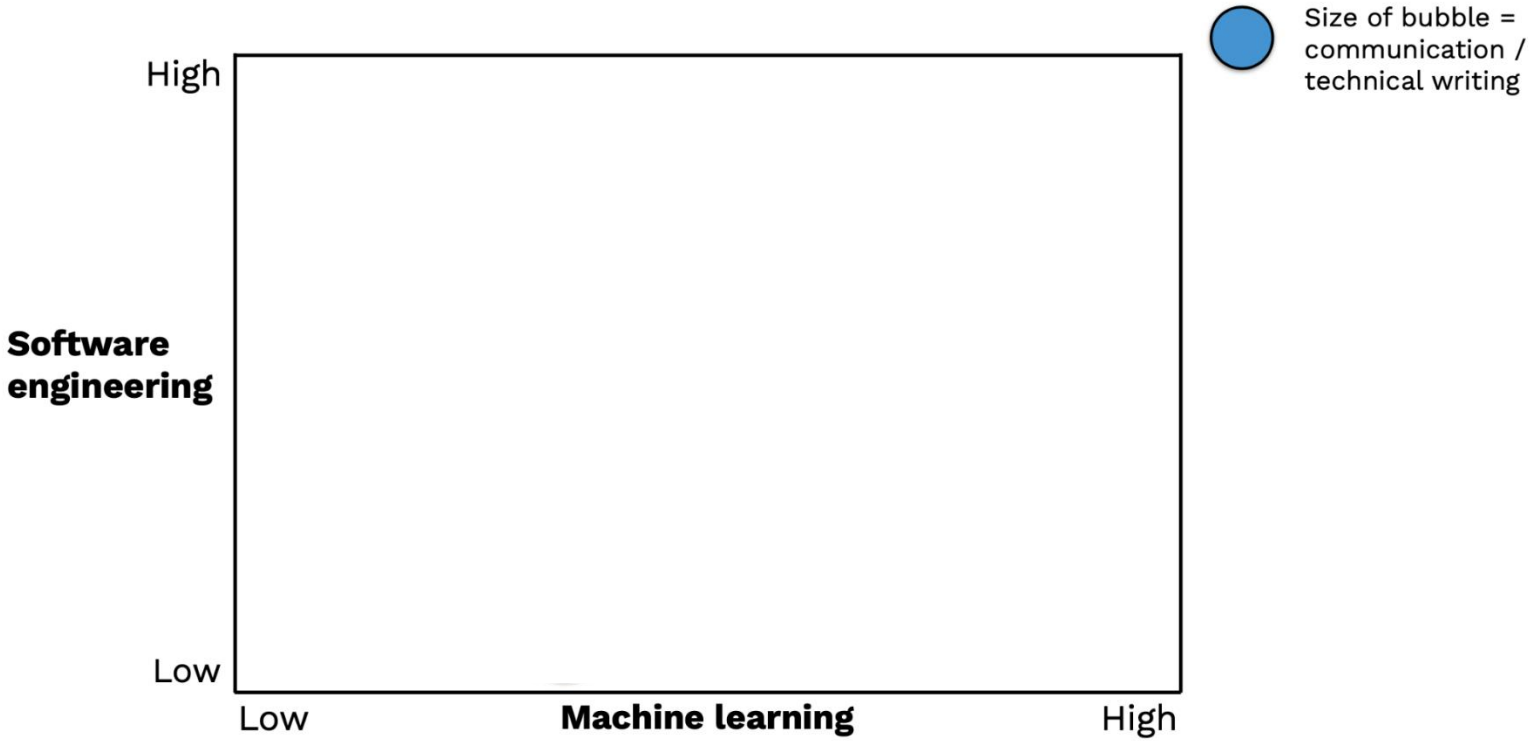

Where should the different positions be?

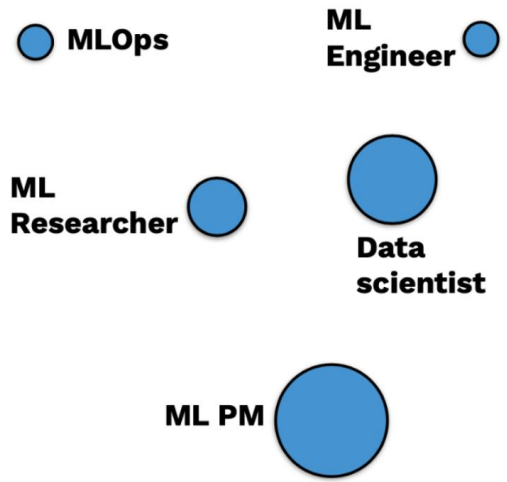

[1] Full Stack Deep Learning course 2022, https://fullstackdeeplearning.com/course/2022/lecture-8-teams-and-pm/

According to stable diffusion

https://stablediffusionweb.com/

Prompt:

“Machine Learning operations engineer”

A MLOps engineer is

- Buff
- Locked in
- Many screens

Its not completely wrong

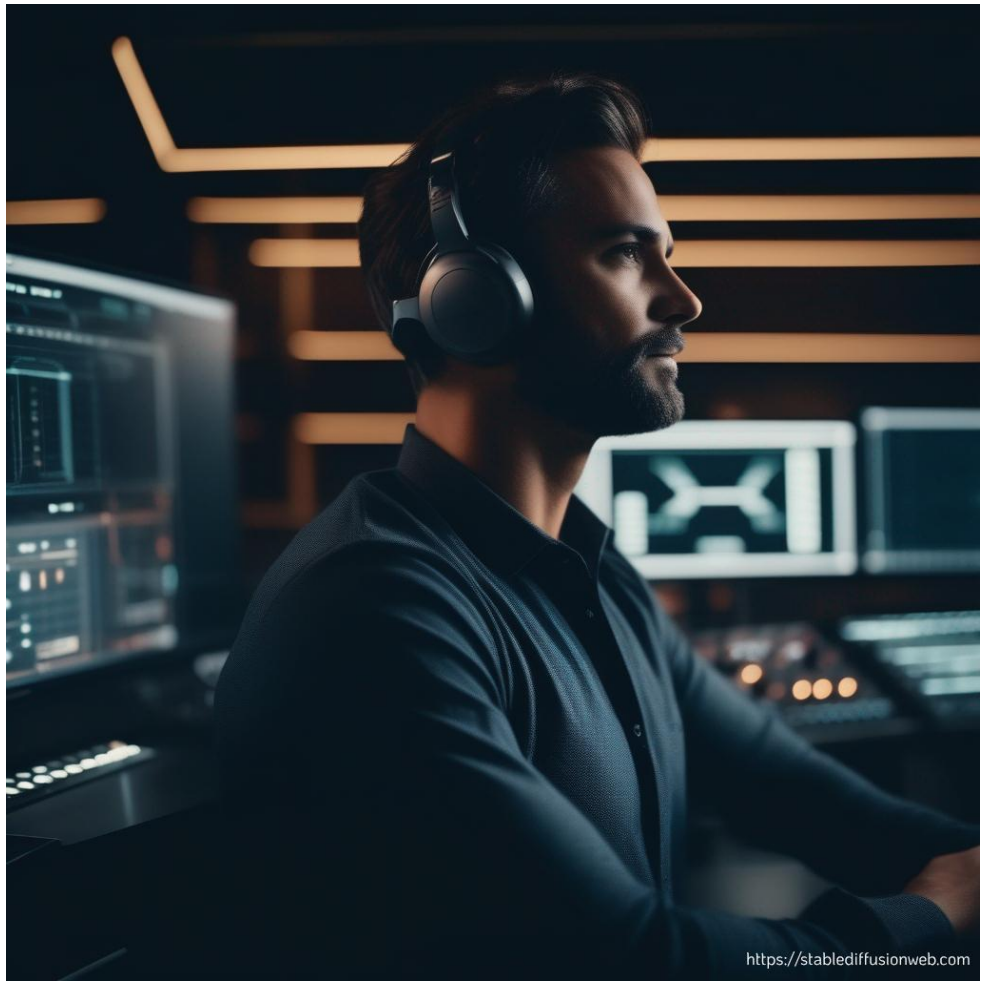

Why is MLOps hot?

Why does companies care about MLOps

Having automated model deployed with errors can cost A LOT of money:

”A famous example of the dangers here was Knight Capital´s system losing $465 millions in 45 minutes, apparently because of unexpected behavior from obsolete experimental codepaths” – Hidden Technical debt in Machine Learning Systems

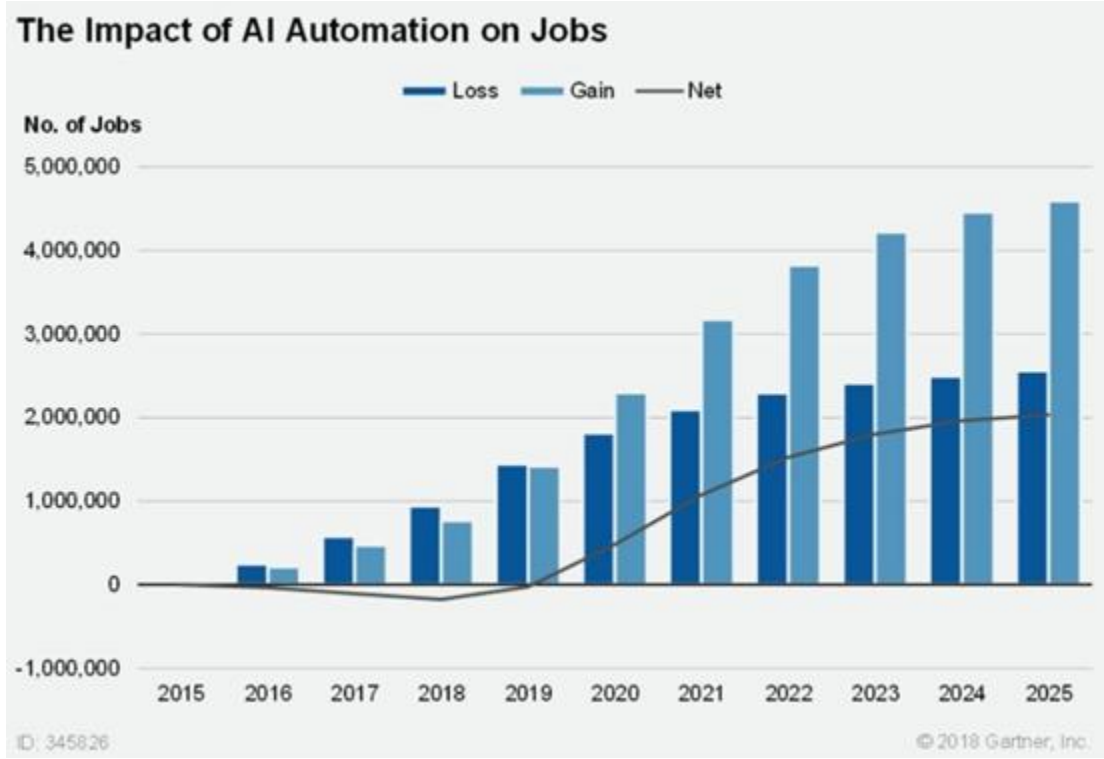

Let's ask ChatGPT

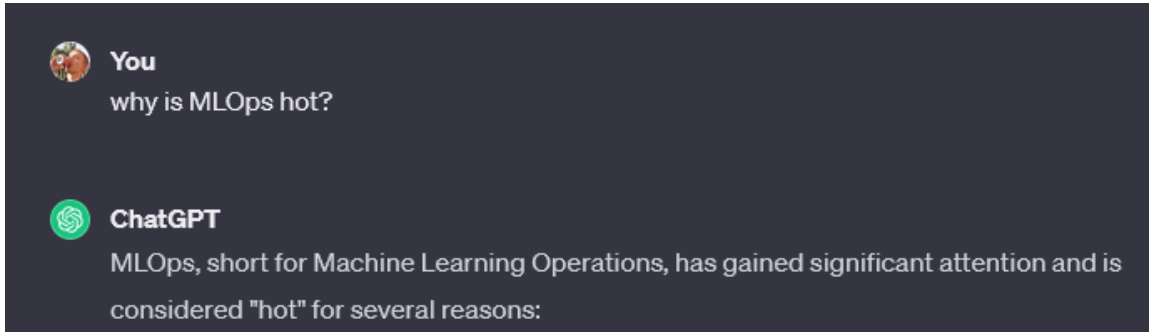

1. Growing Adoption of Machine Learning (ML)
2. Complexity of ML Workflow
3. Bridge between Development and Operations
4. Need for Collaboration
5. Ensuring Model Governance and Compliance
6. Automation and Scalability
7. Continuous Integration and Continuous Deployment (CI/CD)
8. Infrastructure Orchestration
9. Adoption of Cloud Services
10. Business Impact

Open AI study

🔎 What is the contributes to the success of OpenAI?

- 💡 Funding
- 💡 Data
- 💡 People
- 💡 Compute
- 💡 Service contract with Microsoft

Microsoft fired 10,000 workers in the same breath as the invested $10B in OpenAI

Trends in MLOps

MLOps has been trending for a couple of
Tools have been the main priority

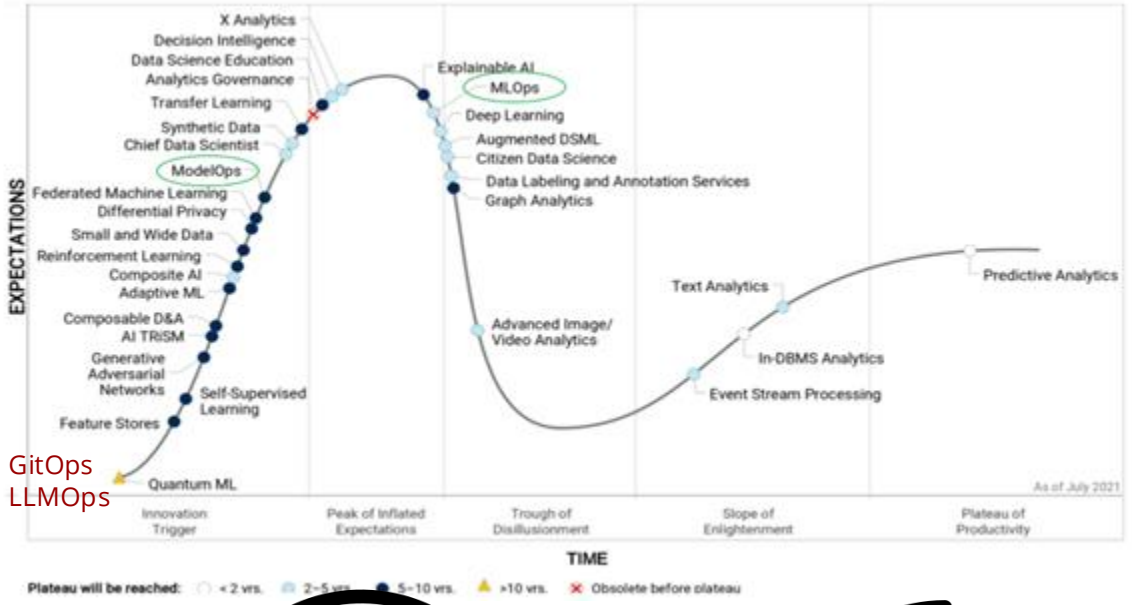

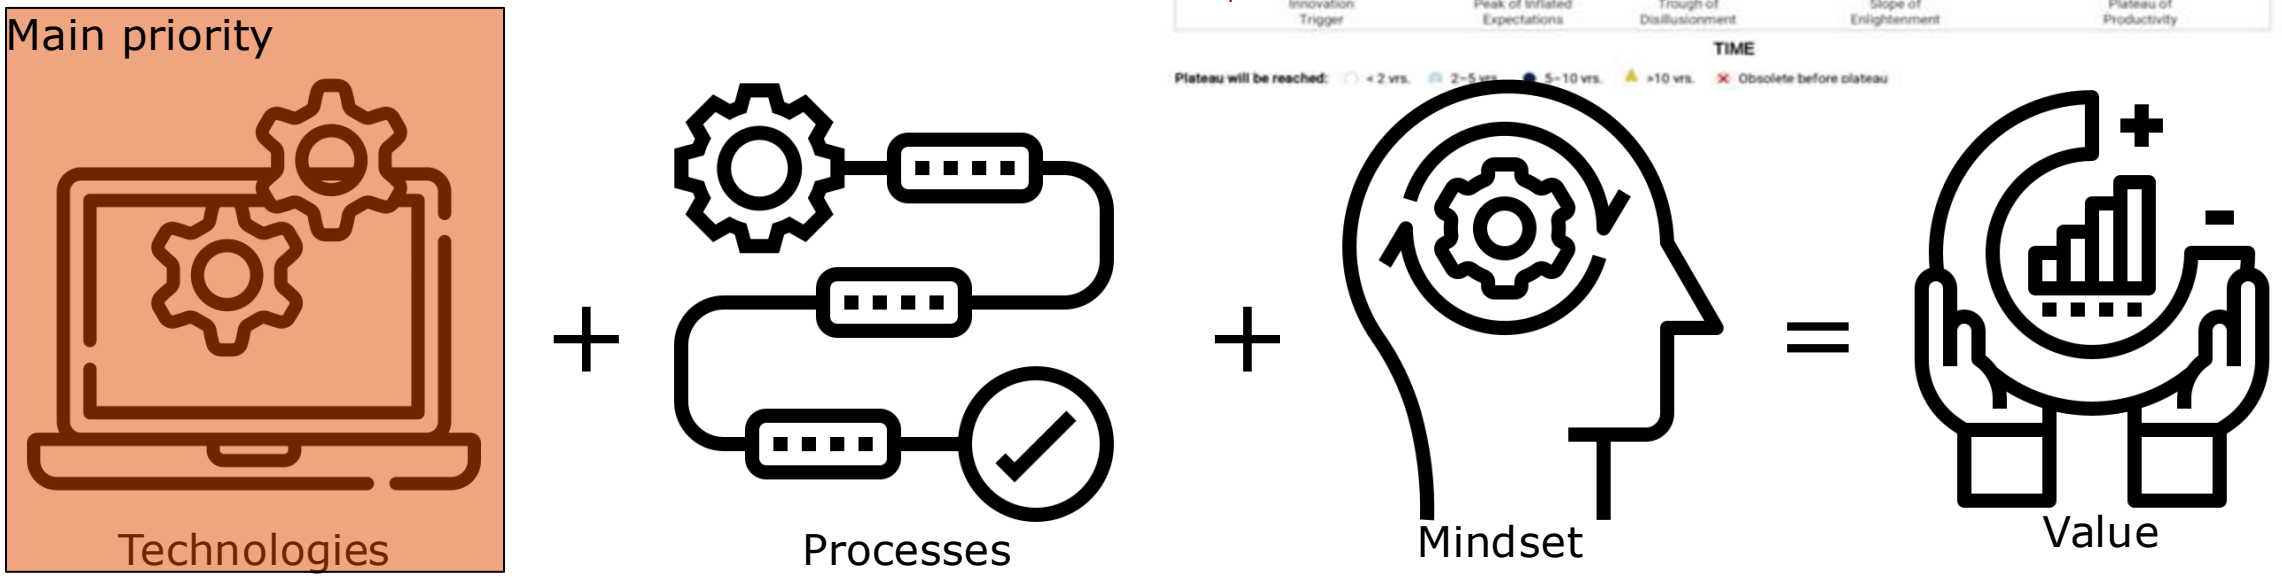

Choosing the right tool for the job

Looking back

MLOps around 2006 = write everything from scratch

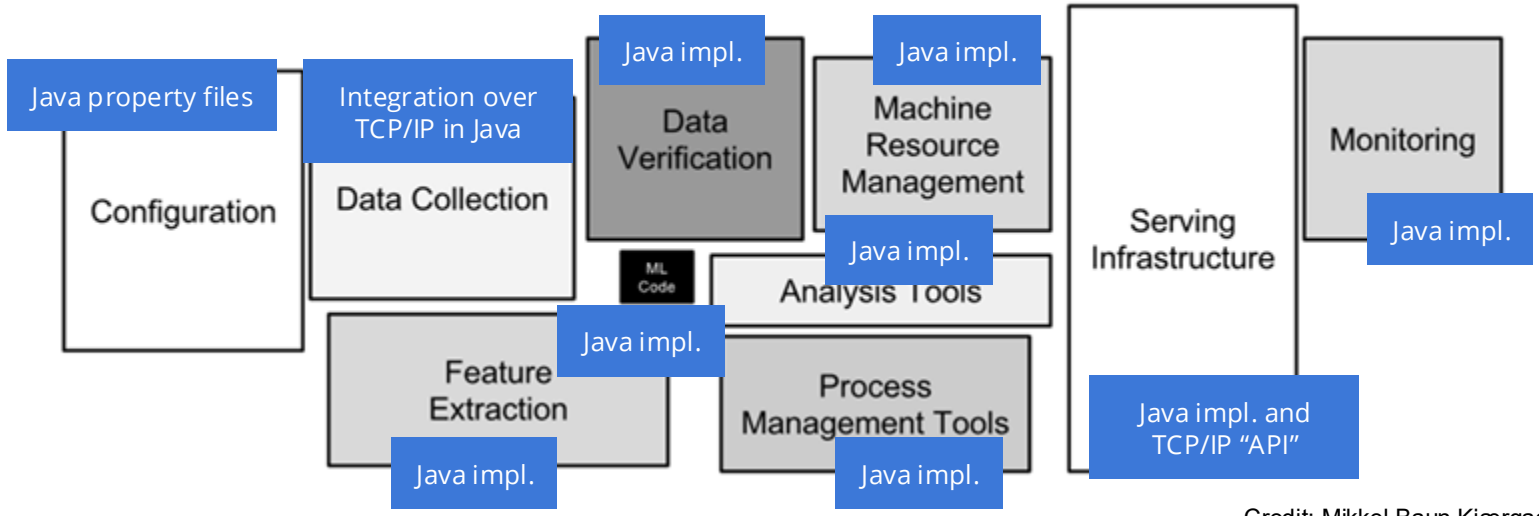

Credit: Mikkel Baun Kjærgaard

Pros:

+ Full control

Cons:

- Slow to iterate
- Hard to maintain
- Lot of manpower per project

Today we have options

There is a tool for everything you need

Pros:

+ Easy to get started
+ Easy to iterate

Cons:

− Framework integration can be really hard
− Hard to compare frameworks

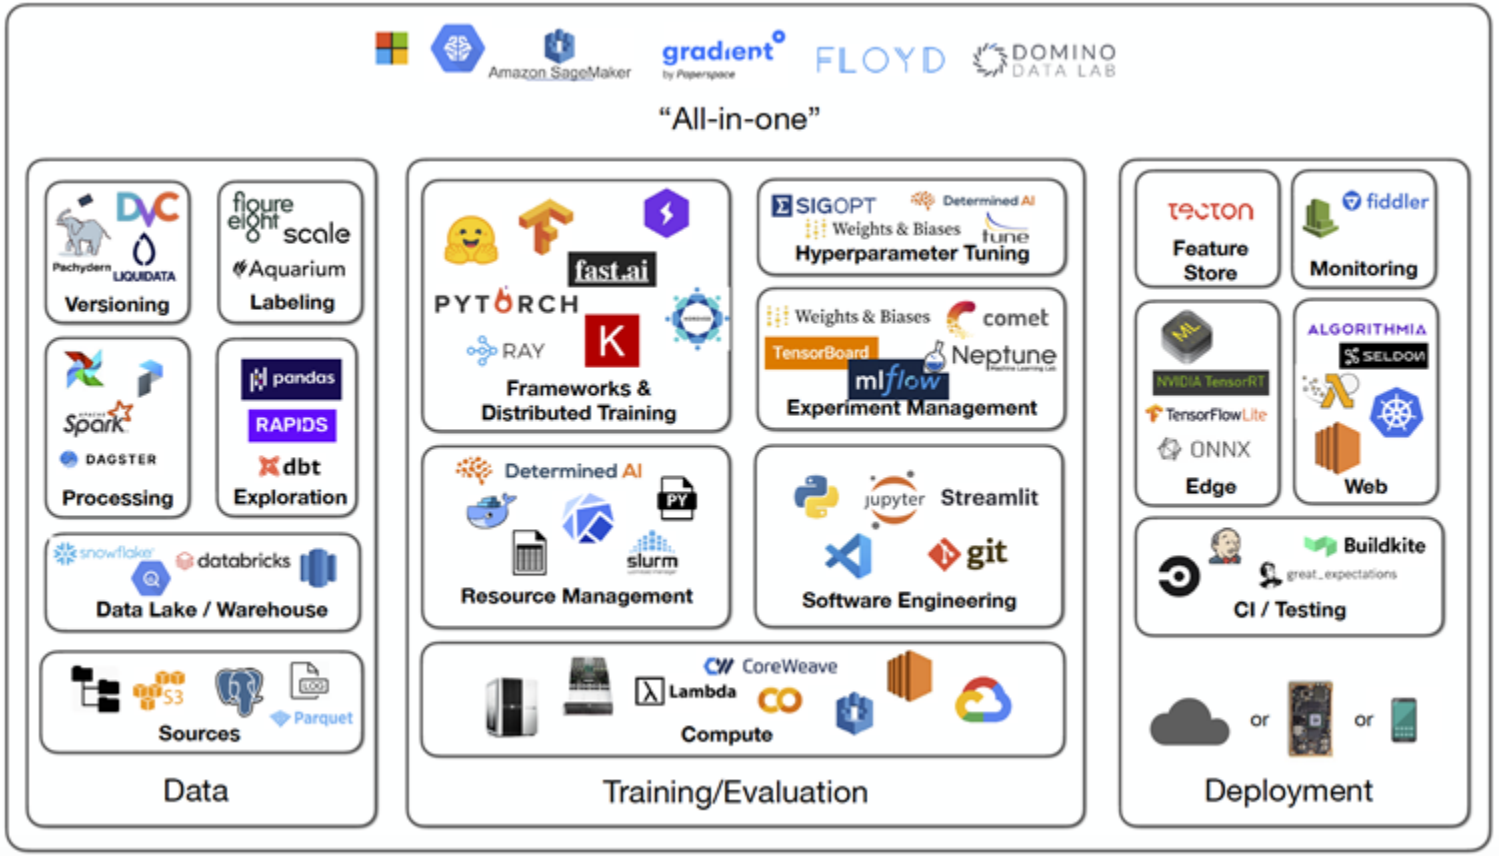

MLOps now

Pick a *stack* of tools

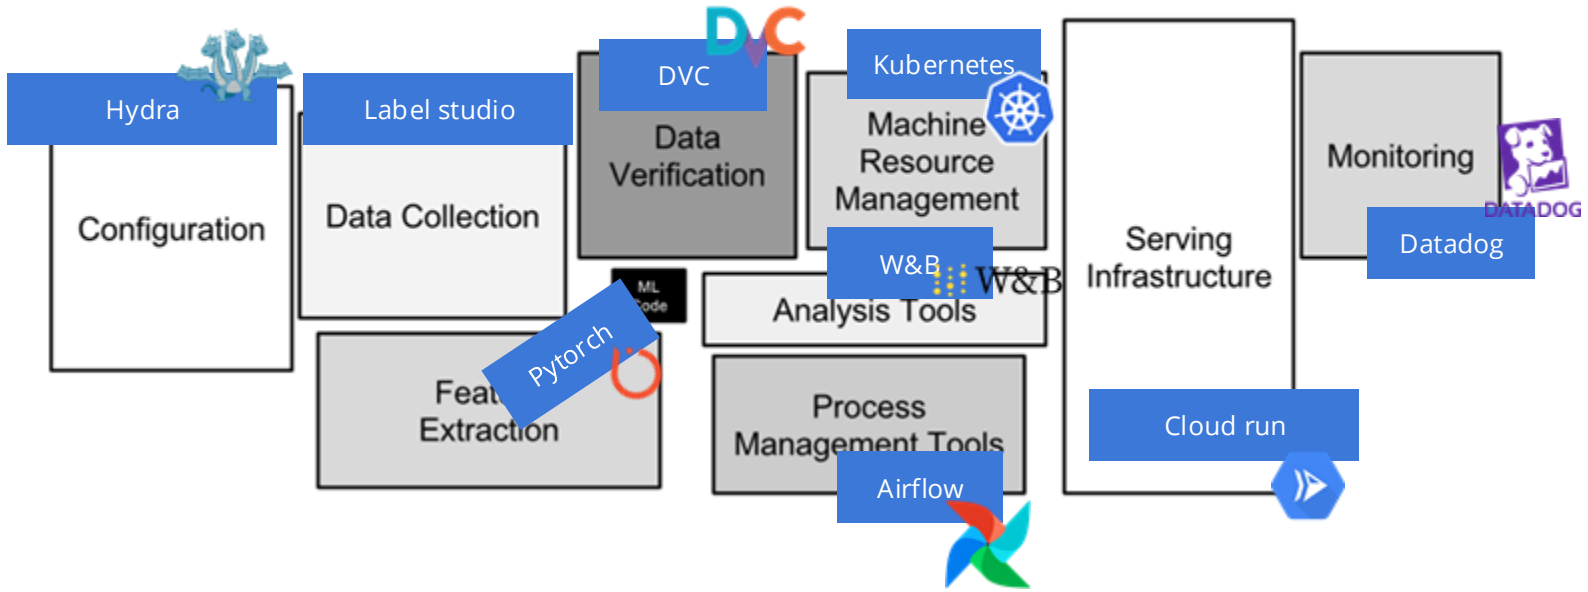

Not picking the right tool leads to TD

In a nutshell MLOps is about dealing reducing technical debt

Technical debt is the implied cost of future reworking required when choosing an easy but limited solution instead of a better approach that could take more time

MLOps is full stack

In MLOps we embrace the full stack of problems that comes from the full lifecycle. Especially integration problems.

Criteria for what goes into the stack (4Cs):

- Cost
- Coverage
- Complexity
- Community

Amazon SageMaker

Whenever we need to pick one tool over the other, we need to consider these 4 criteria.

And most time this is not possible without actually trying to use both.

MLOps as a process

MLOps has and is too tool centric

Why the Focus on Tools?

💡 **Tech-Driven Hype** – MLOps emerged alongside cloud AI services, containerization, and orchestration tools, making it feel like a tooling problem rather than a process and culture problem.

💡 **Vendor Influence** – Companies push their own MLOps stacks, leading to fragmented ecosystems that emphasize tool adoption rather than best practices.

💡 **ML Engineers' Backgrounds** – Many ML practitioners come from research backgrounds and are more familiar with coding than system design

Why Process Matters More

If we **strip away the tools**, the core of MLOps is about establishing **robust workflows** that ensure:

💡 Models are reproducible and traceable (versioning, experiment tracking).

💡 CI/CD practices extend beyond software to **model development** (automated validation, deployment gates).

💡 Models are **continuously monitored and retrained** as data distributions change.

💡 Organizations have **clear roles and responsibilities** (who builds, tests, deploys, and maintains models).

1. Maturity model

💡 The MLOps model is intended as a tool to help companies get an understanding of where they are the MLOps process and where to go next

💡 Do note that not all should aim for level 4, due to regulation, risk etc.

Level	Description	Highlights	Technology
0	No MLOps	• Difficult to manage full machine learning model lifecycle • The teams are disparate and releases are painful • Most systems exist as "black boxes," little feedback during/post deployment	• Manual builds and deployments • Manual testing of model and application • No centralized tracking of model performance • Training of model is manual
1	DevOps but no MLOps	• Releases are less painful than No MLOps, but rely on Data Team for every new model • Still limited feedback on how well a model performs in production • Difficult to trace/reproduce results	• Automated builds • Automated tests for application code
2	Automated Training	• Training environment is fully managed and traceable • Easy to reproduce model • Releases are manual, but low friction	• Automated model training • Centralized tracking of model training performance • Model management
3	Automated Model Deployment	• Releases are low friction and automatic • Full traceability from deployment back to original data • Entire environment managed: train > test > production	• Integrated A/B testing of model performance for deployment • Automated tests for all code • Centralized tracking of model training performance
4	Full MLOps Automated Operations	• Full system automated and easily monitored • Production systems are providing information on how to improve and, in some cases, automatically improve with new models • Approaching a zero-downtime system	• Automated model training and testing • Verbose, centralized metrics from deployed model

[1] https://learn.microsoft.com/en-us/azure/architecture/ai-ml/guide/mlops-maturity-model

2. Machine learning is experimental at its core

- This means starting with simple solutions first and then improving them.
- Everything should be thought of as a cyclical process.
- This means allocating time in the future to fix, improve, and develop your solution.

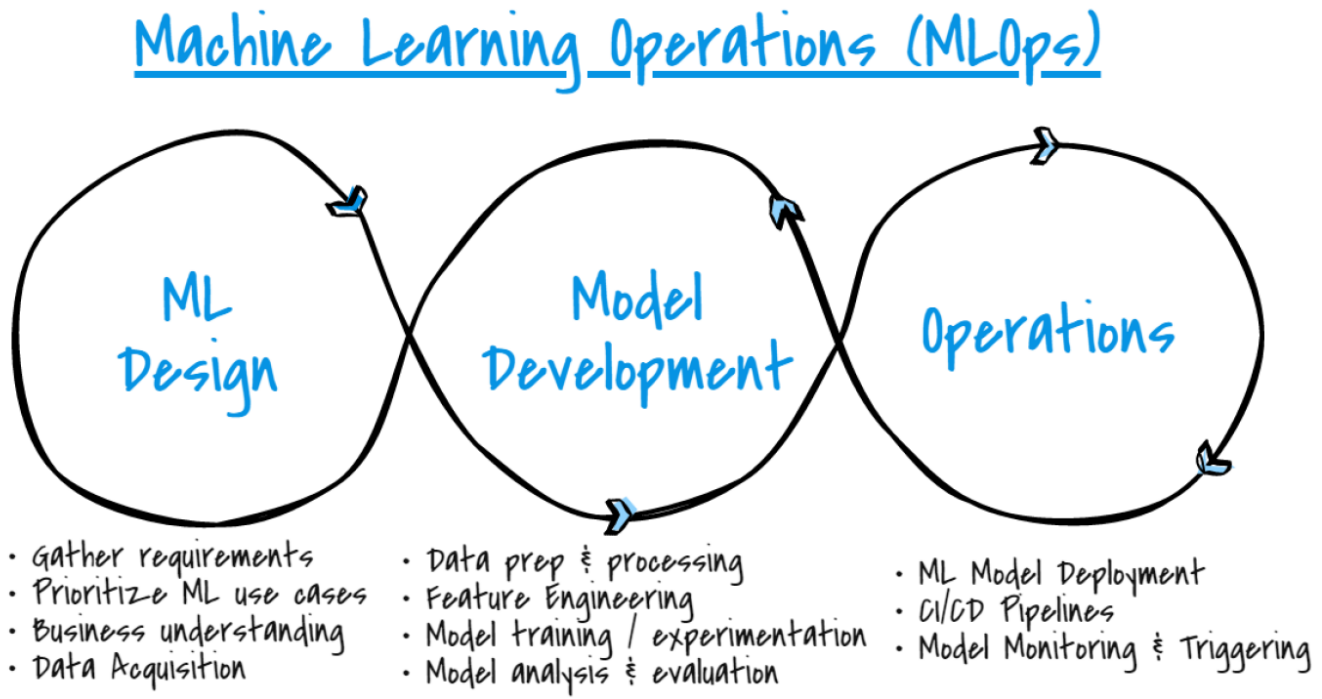

3. Focus on one problem at the time

We want an

ML lifecycle that is *reproducible, traceable, testable, and maintainable.*

Or in other words: we want to reduce future technical debt.

However, this is not possible on the first attempt. So, focus on one aspect, iterate until it is resolved, and then move on to the next—while ensuring:

💡 **Reproducibility**: Ensure that every step can be reproduced by a colleague. If not, it may mean having to start over at times.

💡 **Traceability**: Ensure that all data is stored throughout the pipeline. If not, parts of experiments/models may need to be redone.

💡 **Testability**: Implement proper CI for all parts of the pipeline. If not, there are no guarantees for robustness when making future changes.

💡 **Maintainability**: Ensure the project does not become outdated. If not, updating it in the future will take a long time.

4. Focus on what is important for your case

Velocity:

💡 ML is experimental by nature. The faster you can prototype and iterate on ideas, the better.

Validation:

💡 Errors become more expensive to fix once users encounter them. Early validation helps eliminate bad ideas and catch mistakes as soon as possible.

Versioning:

💡 Ensures reproducibility of the production pipeline. Minimizes downtime by quickly switching from a faulty model to another.

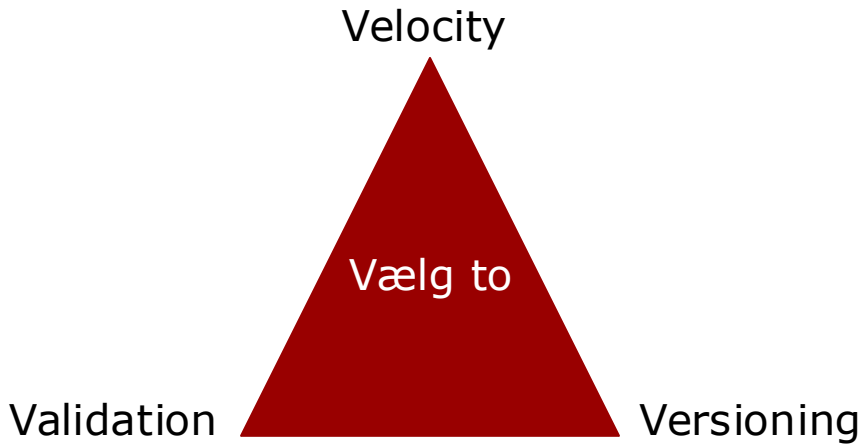

[1] Operationalizing Machine Learning: An Interview Study, https://arxiv.org/abs/2209.09125

Case study 1

Improving a pipeline

Danish company for cardiac monitoring

- Core business is hardware
- Slowly transforming into a software as an service company

A pain point analysis

Semi-structured interviews with stakeholders

- C-suite
- Data scientists
- Engineers

To identify restrictions and pain points of current workflow

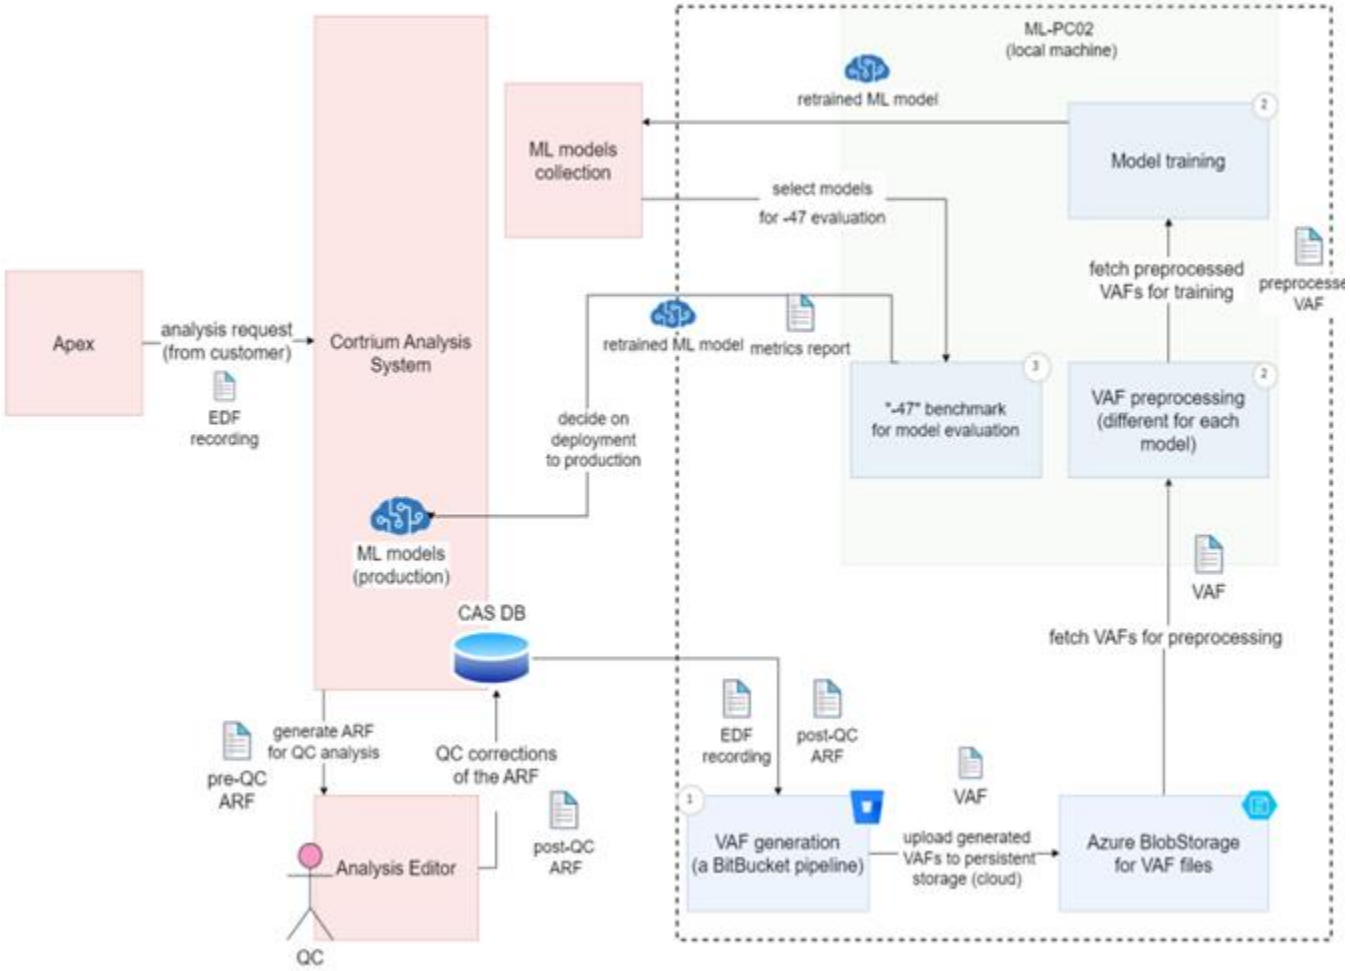

Restrictions and sectors to improve

Restrictions:

- Minimum intrusion on current workflow
- Minimum cost

Note:

I can guarantee that this will lead to technical debt down the line

Sectors to improve:

1. Data generation
2. Model training
3. Model evaluation

The pain points

Data generation
- Data is not persisted
- Manually extraction of data
- Manual analysis was not persisted

Model training
- Remote connection to machines
- No centralization of jobs config
- No consistency in training schedule

Model evaluation
- Manual executing on validation pipeline
- Hard to configure to a new model

-Versioning, -Velocity, +Validation

= Overall an orchestration platform was needed to run jobs

Choosing the right tool for the job

The search for a tool to solve most of the pain points was conducted

Coverage	5/5	4/5
Complexity	1/5	5/5
Cost	5/5	3/5
Community	4/5	2/5

The pain points

Data generation

- Data is ~~not~~ persisted
- ~~Manually~~ Automatic extraction of data
- ~~Manual~~ Automatic analysis ~~was not~~ that persisted and auto reported to stakeholders

Model training

- Web service to lunch jobs ~~Remote connection to machines~~
- ~~No~~ centralization of jobs config and requirements
- No consistency in training schedule (only through paid service)

Model evaluation

- *Manual executing on validation pipeline*
- ~~Hard to configure to a new model~~ Simple one line change to run on new model

Case study 2

Detection of flaws in products

Danish medical company want to automatically detect errors in viels

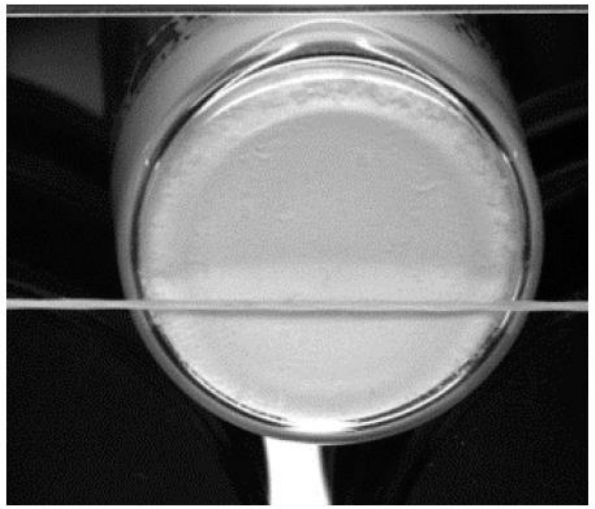

(a) Good Vial - No Defect

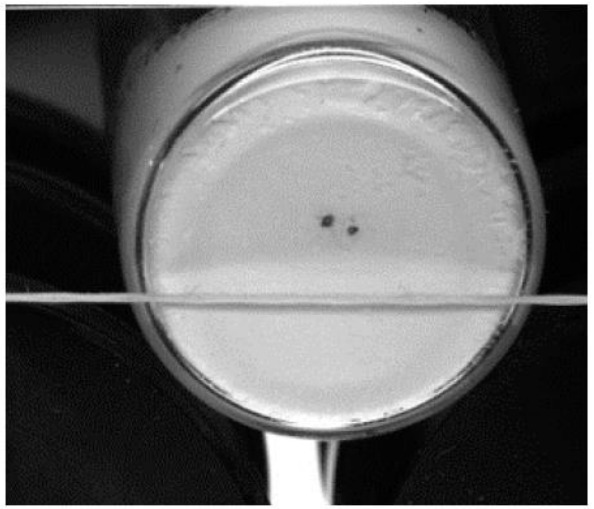

(b) Particle Defect

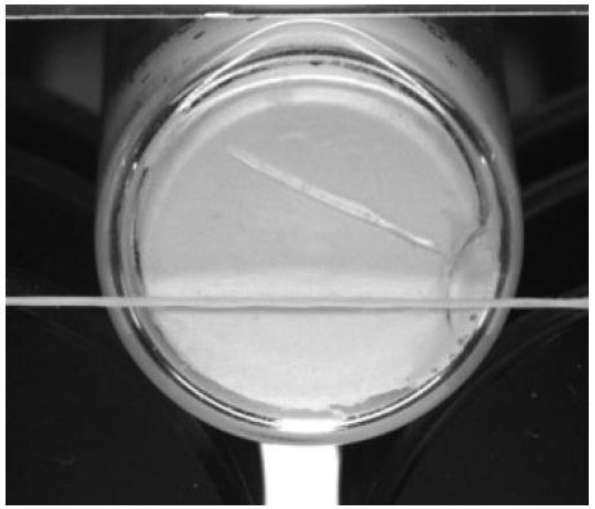

(c) Chips & Cracks Defect

Requirements specification

Model requirements

1. **Must** outperform human baseline
2. **Must** not approve flawed vials -> no risk to patient safety
3. *Should* have a low false rejection rate, minimizing product loss
4. *Should* outperform existing model

Model serving requirements

1. **Must** have a serving latency of less 3 sec pr vial
2. *Should* have a serving latency of less than 0.6 seconds per vial

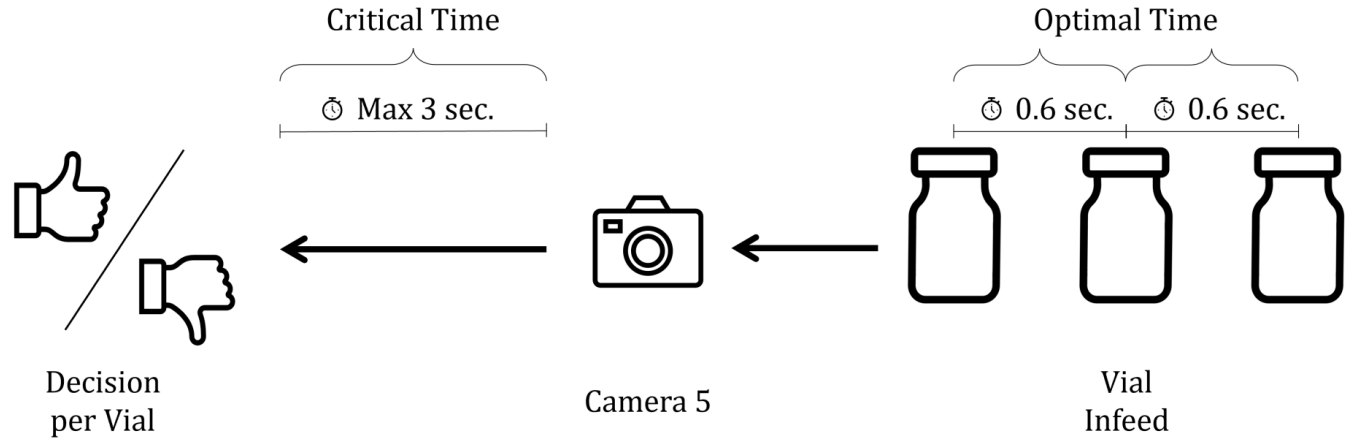

Pipeline

- Containerization
- Training
- Inference

Make sure pipeline is can deal with change in data and model specification

Containerization

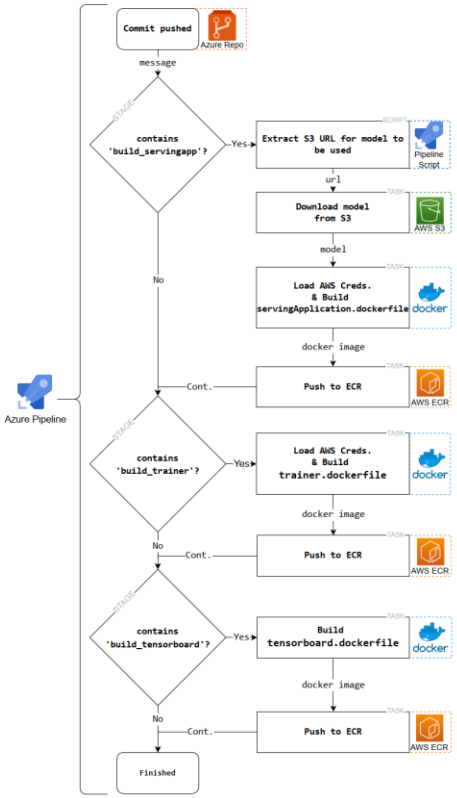

Training

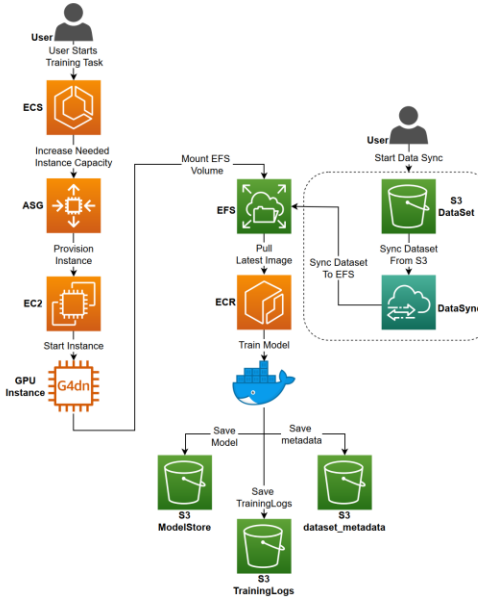

Inference

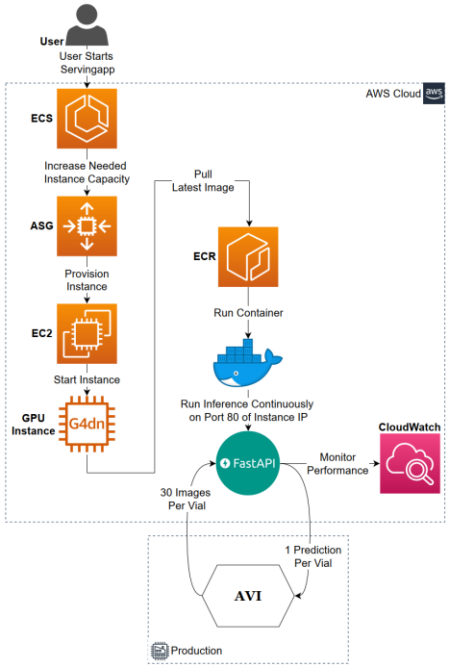

Did it work?

Model requirements

1. **Must** outperform human baseline ✅
2. **Must** not approve flawed vials -> no risk to patient safety ✅
3. *Should* have a low false rejection rate, minimizing product loss ✅
4. *Should* outperform existing model ✅

Model serving requirements

1. **Must** have a serving latency of less 3 sec pr vial ✅
2. *Should* have a serving latency of less than 0.6 seconds per vial ❌

		Predictions						
		NN Model			`Efficientnet_b0`			
		Good	**Particle**	**CC**	**Good**	**Particle**	**CC**	Total
Ground Truth	**Good**	648	13	22	682	0	1	683
	Particle	0	17	8	0	25	0	25
	CC	10	0	36	0	0	46	46

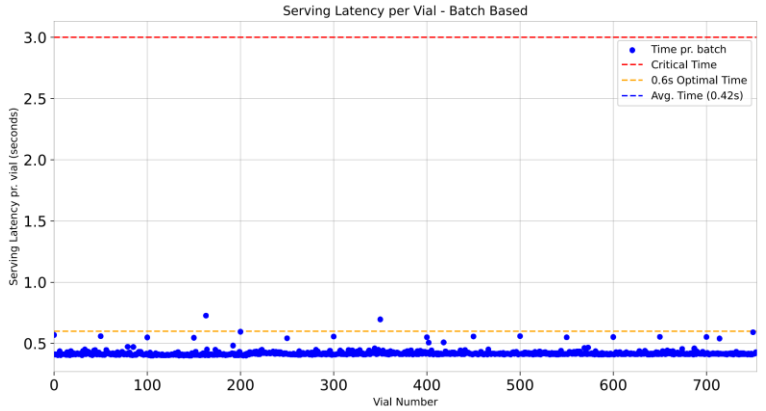

Let's get down to business, to defeat the bugs 🐛

💻 Practical live coding example from ML model to deployed model

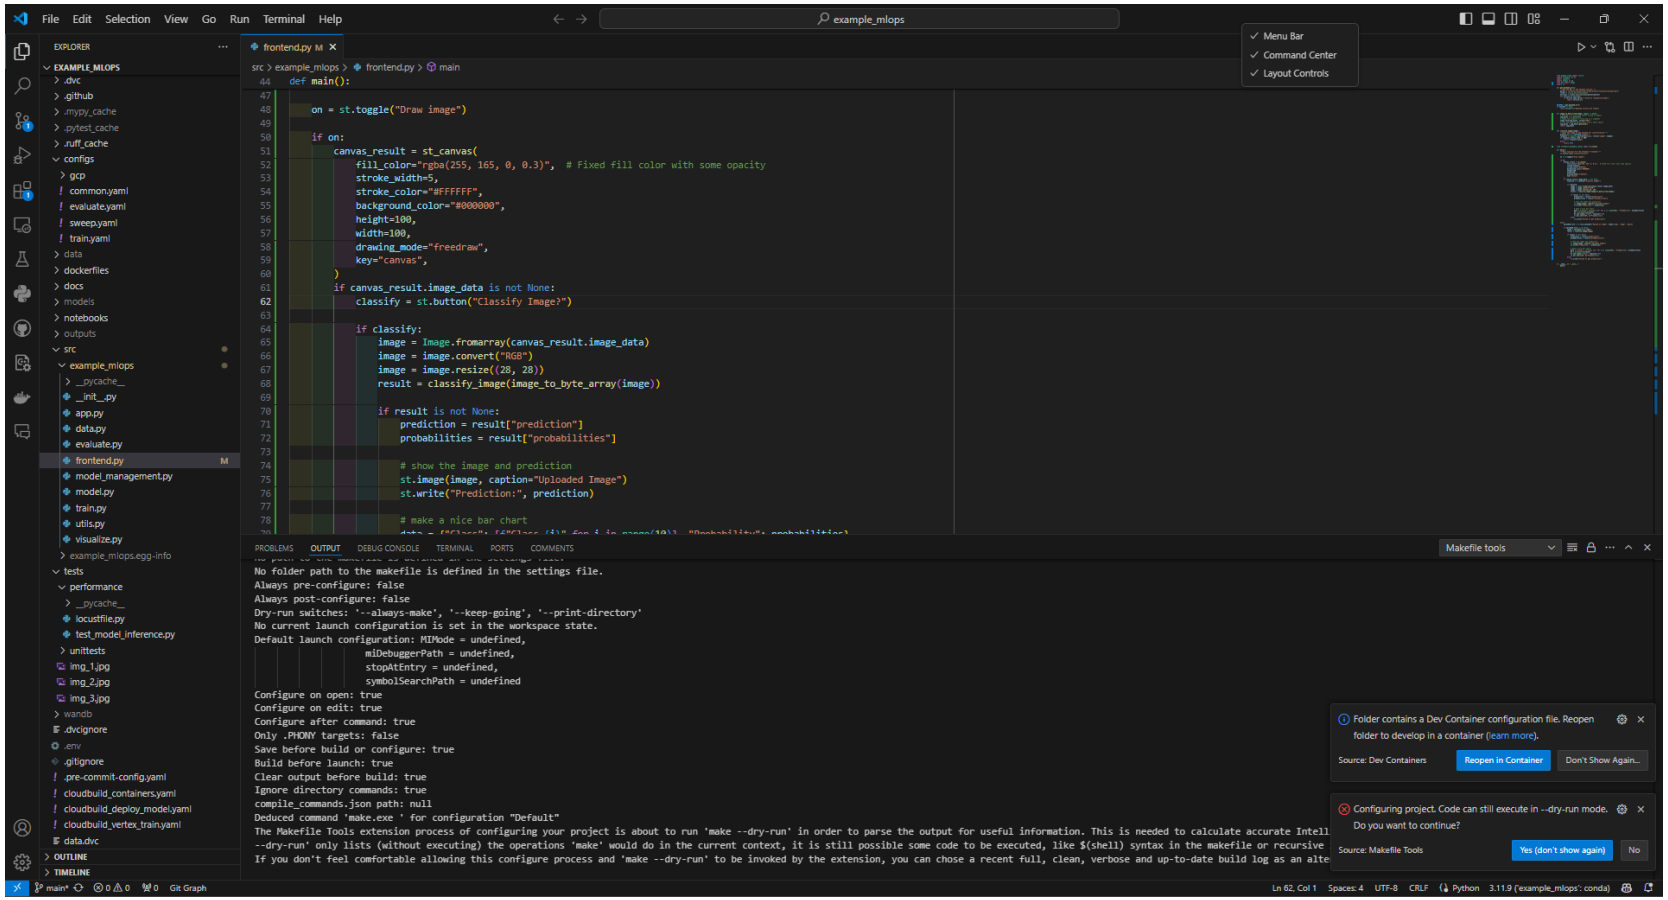

Trustworthy AI

What is trustworthy AI?

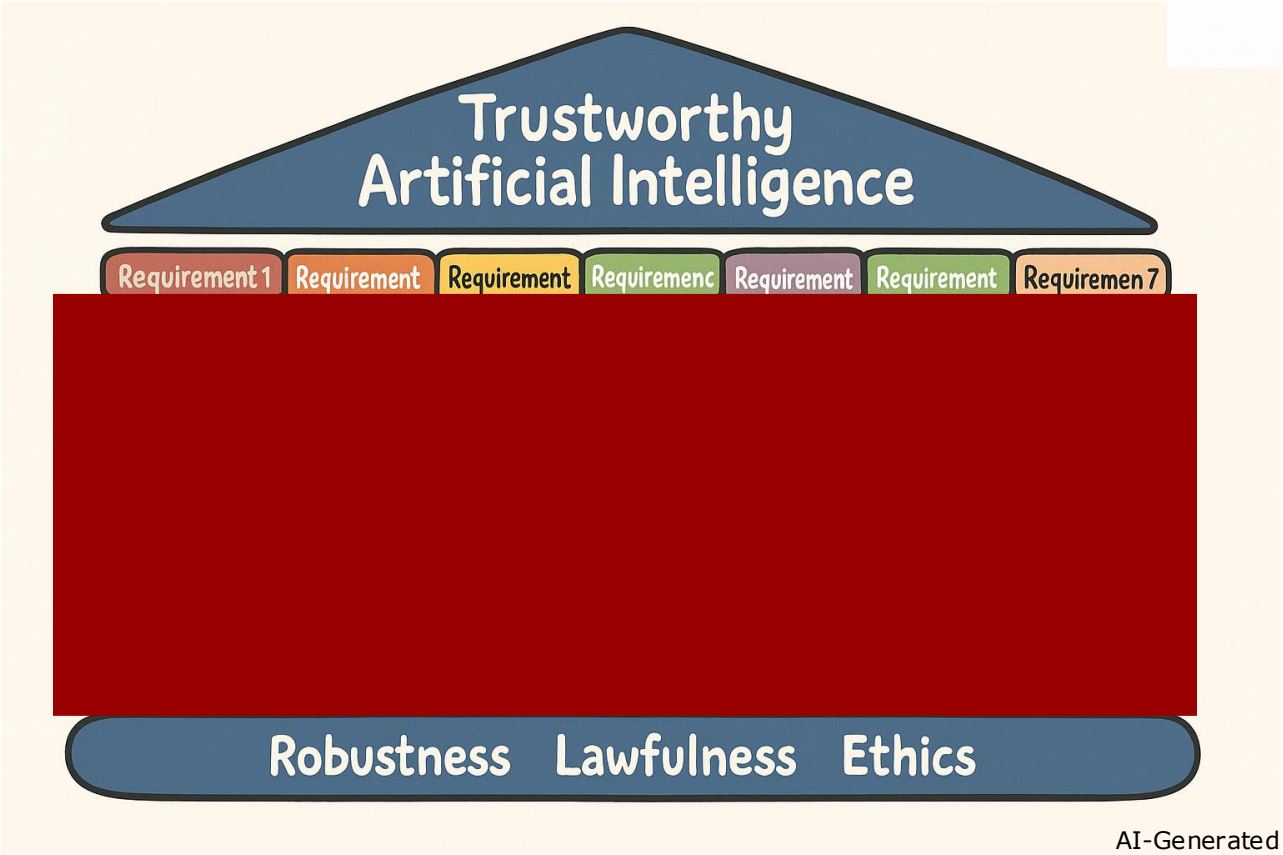

AI-Generated

[1] https://montrealethics.ai/connecting-the-dots-in-trustworthy-artificial-intelligence-from-ai-principles-ethics-and-key-requirements-to-responsible-ai-systems-and-regulation/

What is trustworthy AI?

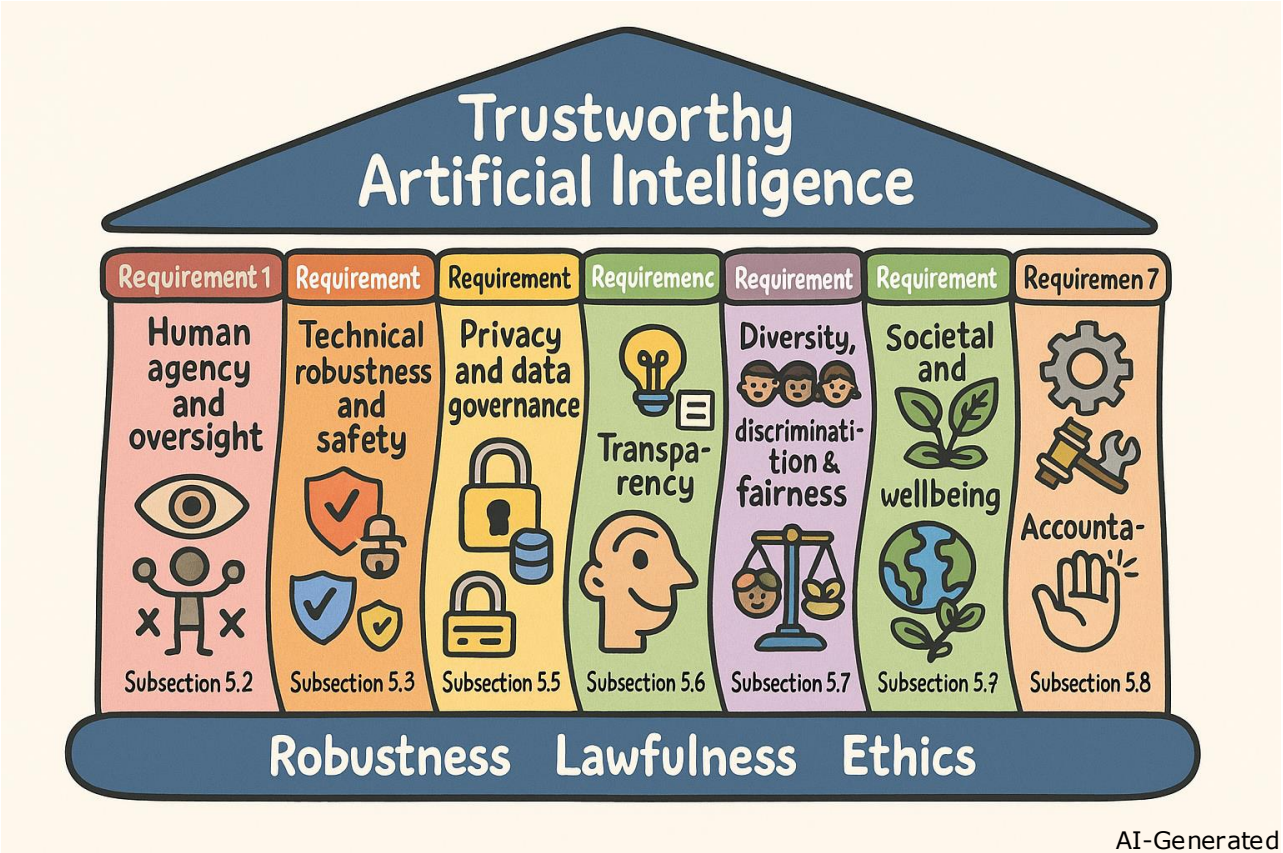

AI-Generated

[1] https://montrealethics.ai/connecting-the-dots-in-trustworthy-artificial-intelligence-from-ai-principles-ethics-and-key-requirements-to-responsible-ai-systems-and-regulation/

[2] Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation, https://www.sciencedirect.com/science/article/pii/S1566253523002129

OECD AI principles

Values-based principles

The OECD AI Principles promote use of AI that is innovative and trustworthy and that respects human rights and democratic values. Adopted in May 2019, they set standards for AI that are practical and flexible enough to stand the test of time.

- Inclusive growth, sustainable development and well-being +
- Human rights and democratic values, including fairness and privacy +
- Transparency and explainability
- Robustness, security and safety
- Accountability

Recommendations for policy makers

- Investing in AI research and development +
- Fostering an inclusive AI-enabling ecosystem +
- Shaping an enabling interoperable governance and policy environment for AI +
- Building human capacity and preparing for labour market transformation +
- International co-operation for trustworthy AI +

Regulation is coming

Think risk from the beginning to not get hit in the end

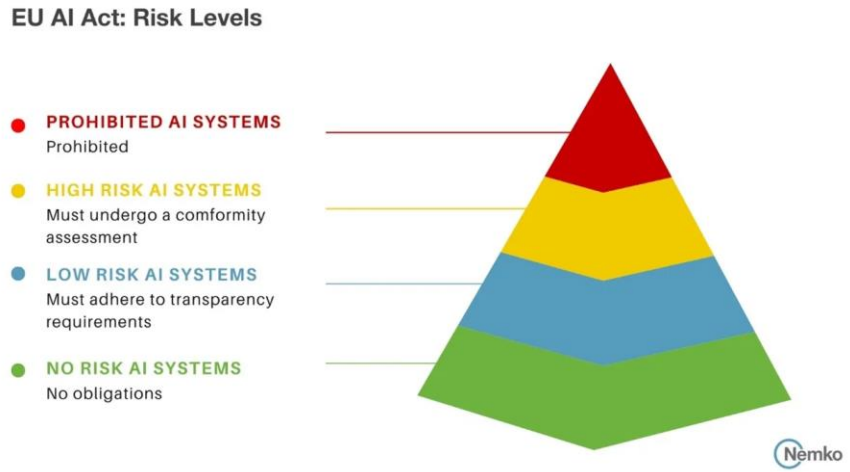

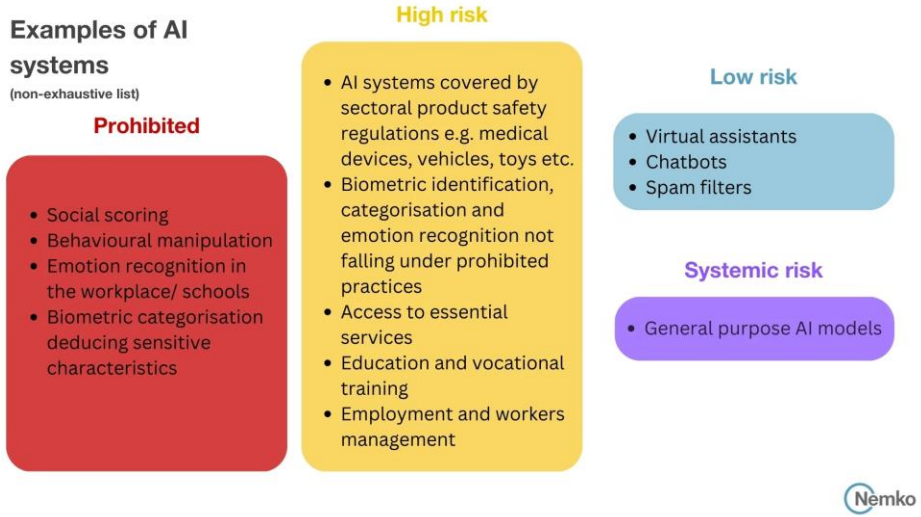

Nemko

[1] https://www.nemko.com/blog/a-quick-dive-into-the-eu-ai-act

Two ways to think AI in the future

Category	US USA	EU EU
⚖️ Regulation	🤝 Voluntary, flexible	📜 Strict, binding (AI Act)
🏛️ Legal Authority	🖋️ Executive orders	🏢 EU law with penalties
🚀 Innovation Focus	⚡ Fast, market-driven	🛡️ Balanced, precautionary
🔍 Risk Management	🧪 Voluntary standards	🧭 Mandatory assessment & oversight
🔒 Data Privacy	🧩 Fragmented (state-based)	🛡️ GDPR + AI-specific rules
🧭 Ethics	📘 Non-binding principles	📜 Mandated fairness & transparency
🌍 Global Influence	🥇 Compete to lead	🌐 Shape global standards
👷 Enforcement	🪶 Light-touch, decentralized	🔨 Strong, centralized

An optimistic view on EU's direction

"The winner of the AI race will not be the one with the best model, but the one that best deals with the social changes"
- LinkedIn post I can't find anymore

💡 Open-source models have (almost) caught up with closed-source models

💡 If everyone has access to the best model, then what really matters?

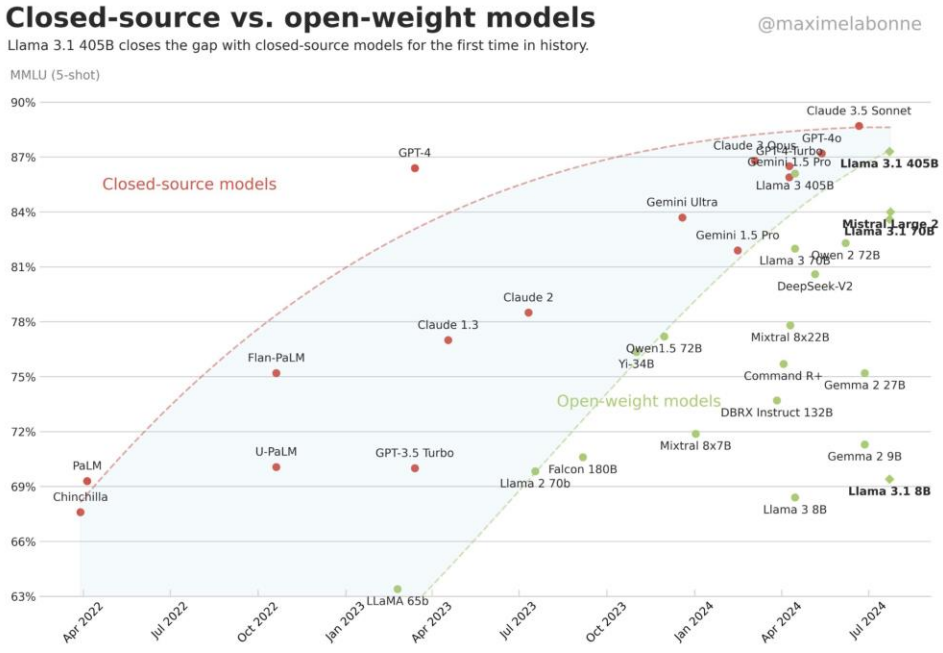

[1] https://the-decoder.com/googles-rumored-gemini-2-0-launch-in-december-could-support-llm-stagnation-thesis/

Consequences as they are right now

TLDR: Treat AI like the software it actually is!

- ⭐ Documentation and governance
- ⭐ Model cards
- ⭐ Bias audits
- ⭐ Usage logs
- ⭐ Data-model lineage

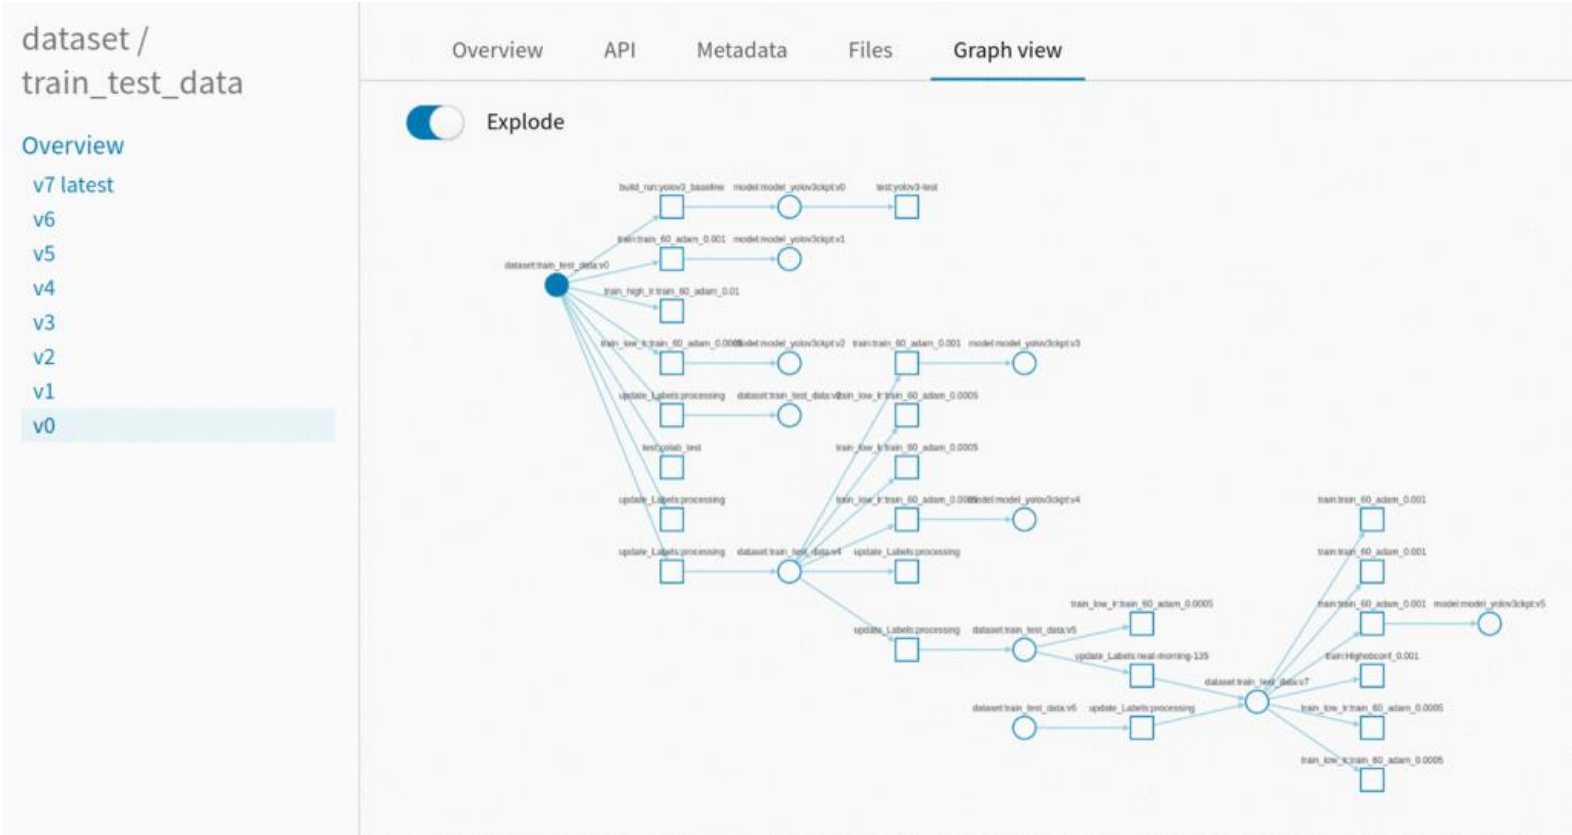

[1] https://docs.wandb.ai/guides/core/registry/lineage/

Z-inspection

💡 Proven and tested framework for trustworthy AI

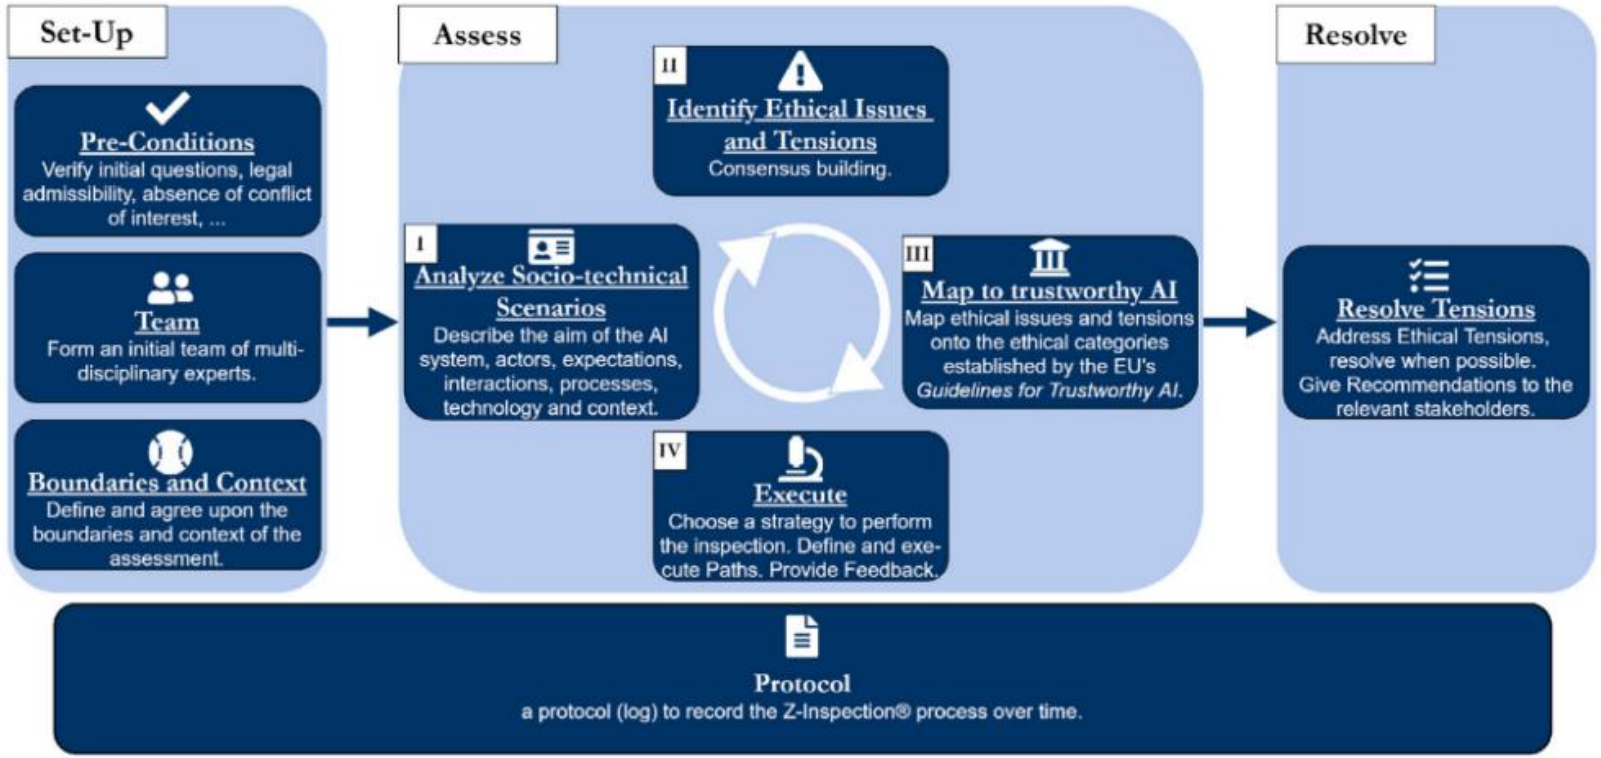

Fig. 1. Z-Inspection process in a nutshell.

[1] https://oecd.ai/en/catalogue/tools/z-inspection

[2] Getting Ready for the EU AI Act in Healthcare, https://arxiv.org/abs/2505.07875

Power analysis

A personal take on science and society

World view

By Pratyusha Kalluri

Don't ask if AI is good or fair, ask how it shifts power

Definitions:

Value is something that causes future actions

Power predicts impact of actions

Through the lens of power, it's possible to see why accurate, generalizable and efficient AI systems are not good for everyone. In the hands of exploitative companies or oppressive law enforcement, a more accurate facial recognition system is harmful. Organizations have responded with pledges to design 'fair' and 'transparent' systems, but fair and transparent according to whom?

❓ Are Meta behaviors aligned with their values ❓

❓ Are Meta behaviors aligned with your values ❓

❓ Is the Meta-you relation "power symmetric" – how are EULA/terms negotiated ❓

Design of Machine Learning Projects

Machine Learning Operations Canvas

Anecdotally, as many as 90% of machine learning models never make it to production. The reasons vary—from misaligned objectives to operational bottlenecks—but the message is clear: building a model is only one part of the journey.

🔥 *A machine learning model is only as impactful as the system that supports it—strategy, scalability, and accountability form its true foundation.*

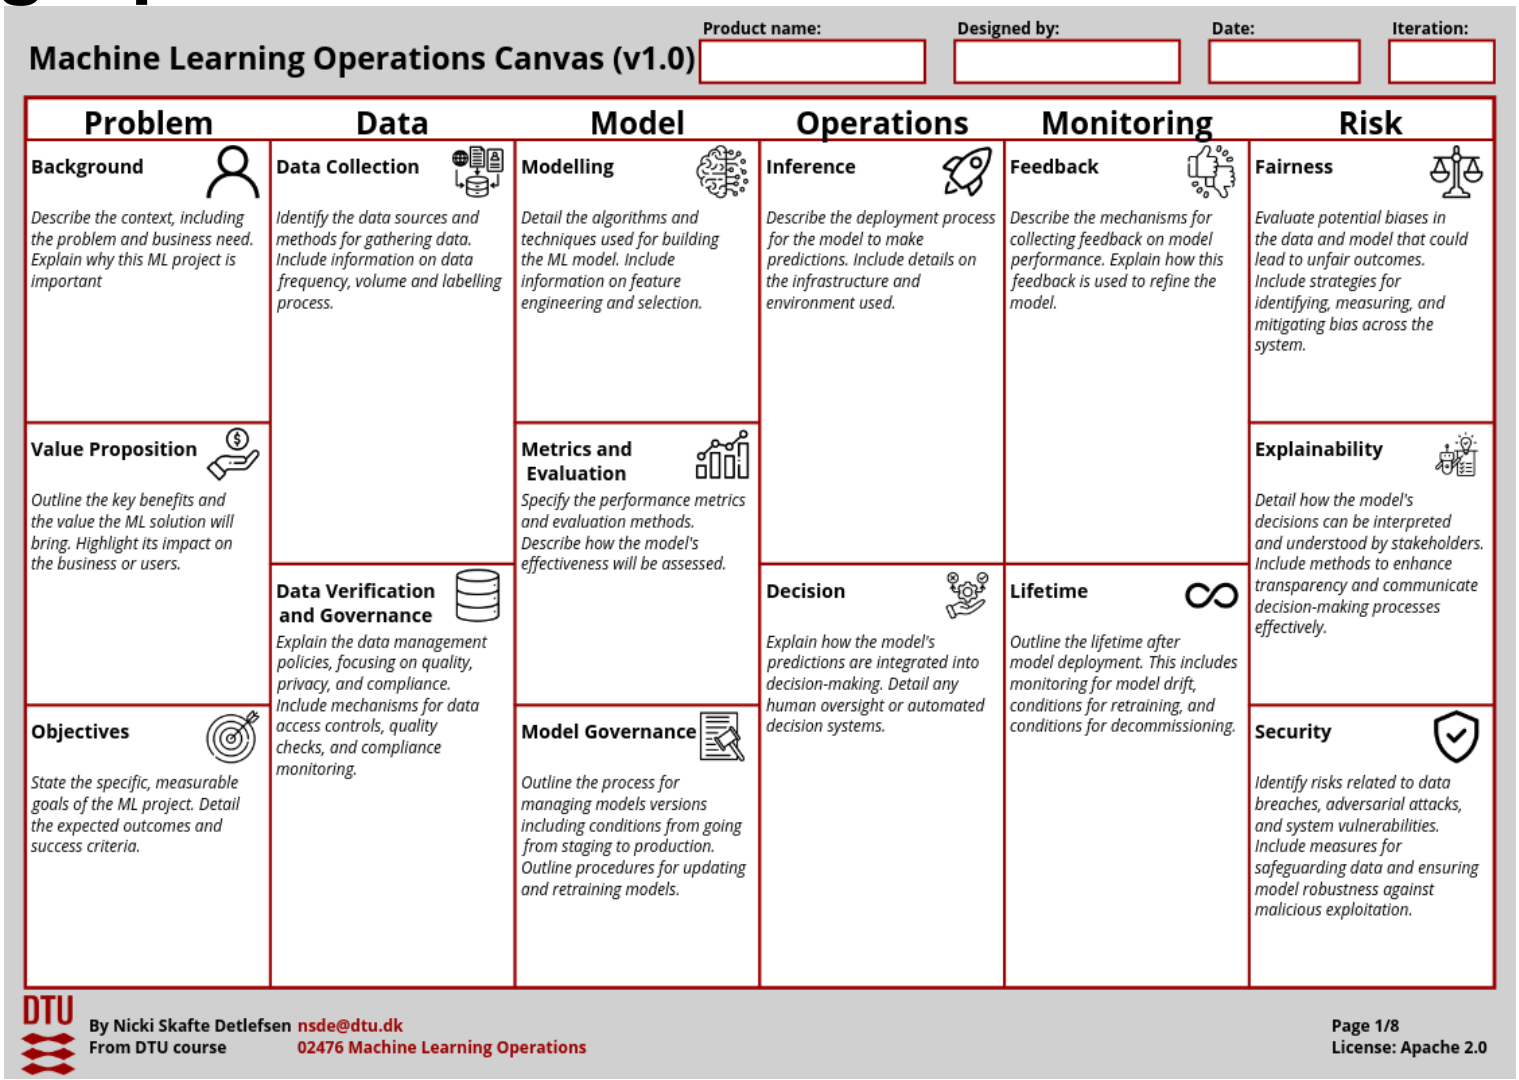

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Machine Learning Operations Canvas

MLOps Canvas (Machine Learning Operations) is a structured framework designed to assist with planning, execution, and management of machine learning projects.

Inspired by the widely used Business Model Canvas, the MLOps Canvas aims to help teams, consisting of both technical and non-technical members, collaborate effectively on machine learning projects.

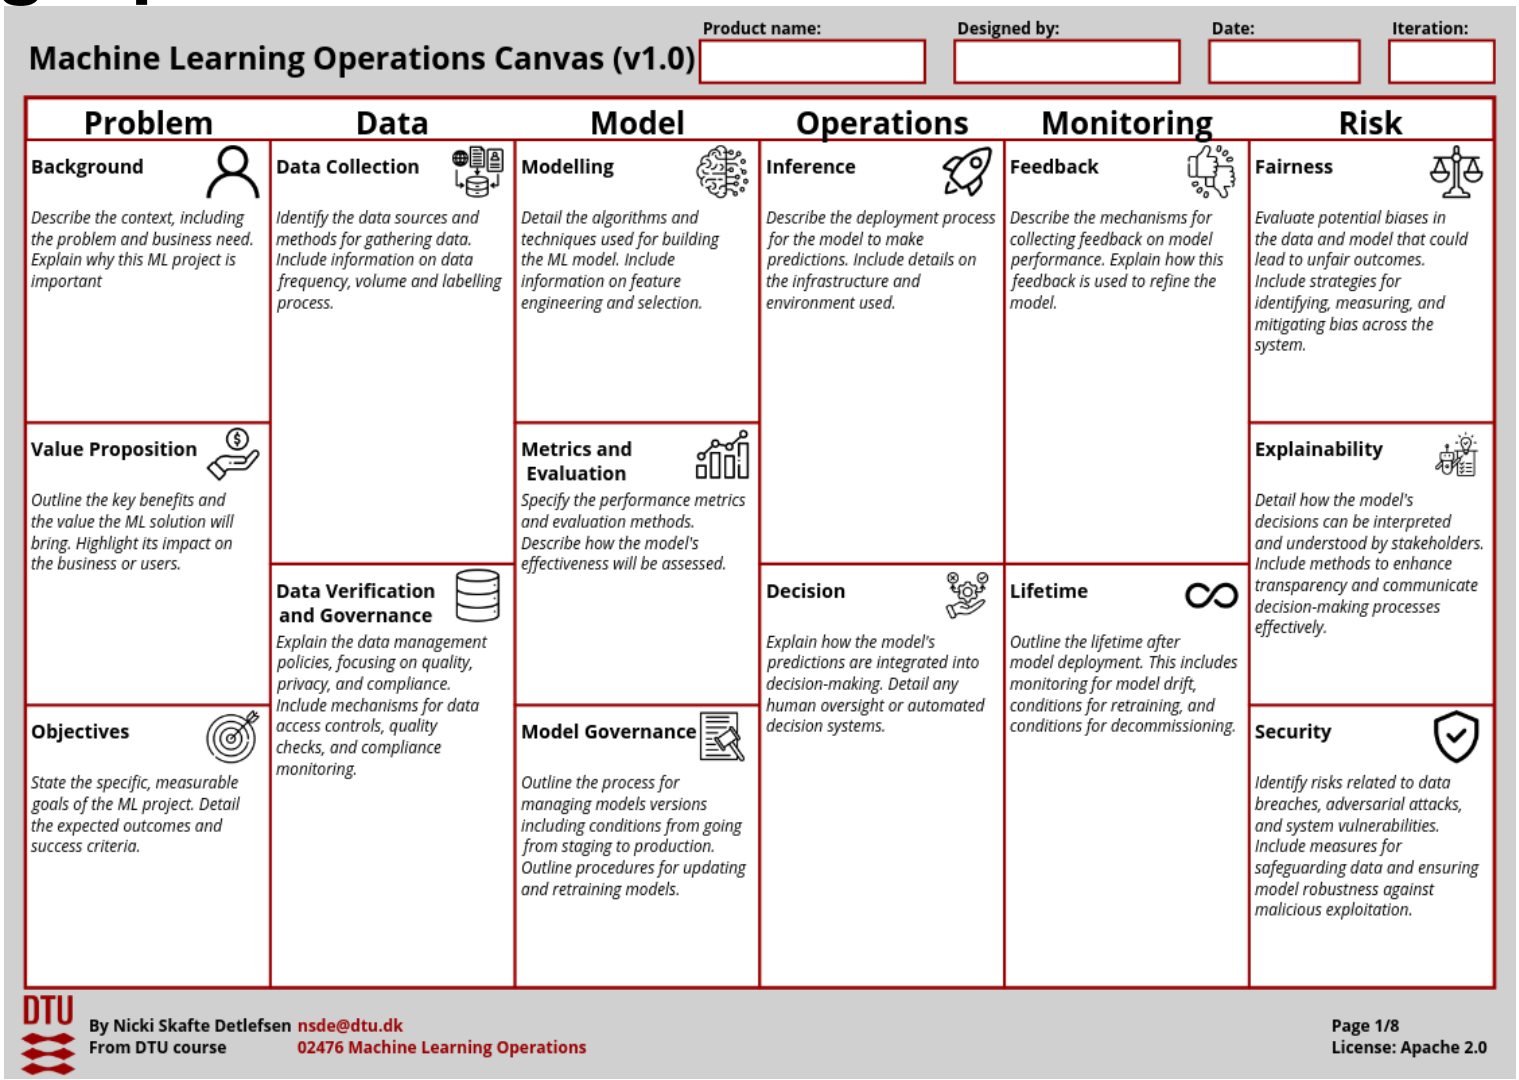

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU
By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations
Page 1/8
License: Apache 2.0

Problem - Background

- What is the context or environment where this problem exists?
- Why is this problem worth solving?
- Who are the stakeholders affected by this problem?
- What is the current state of the problem?
- Are there any constraints or assumptions that need to be considered?

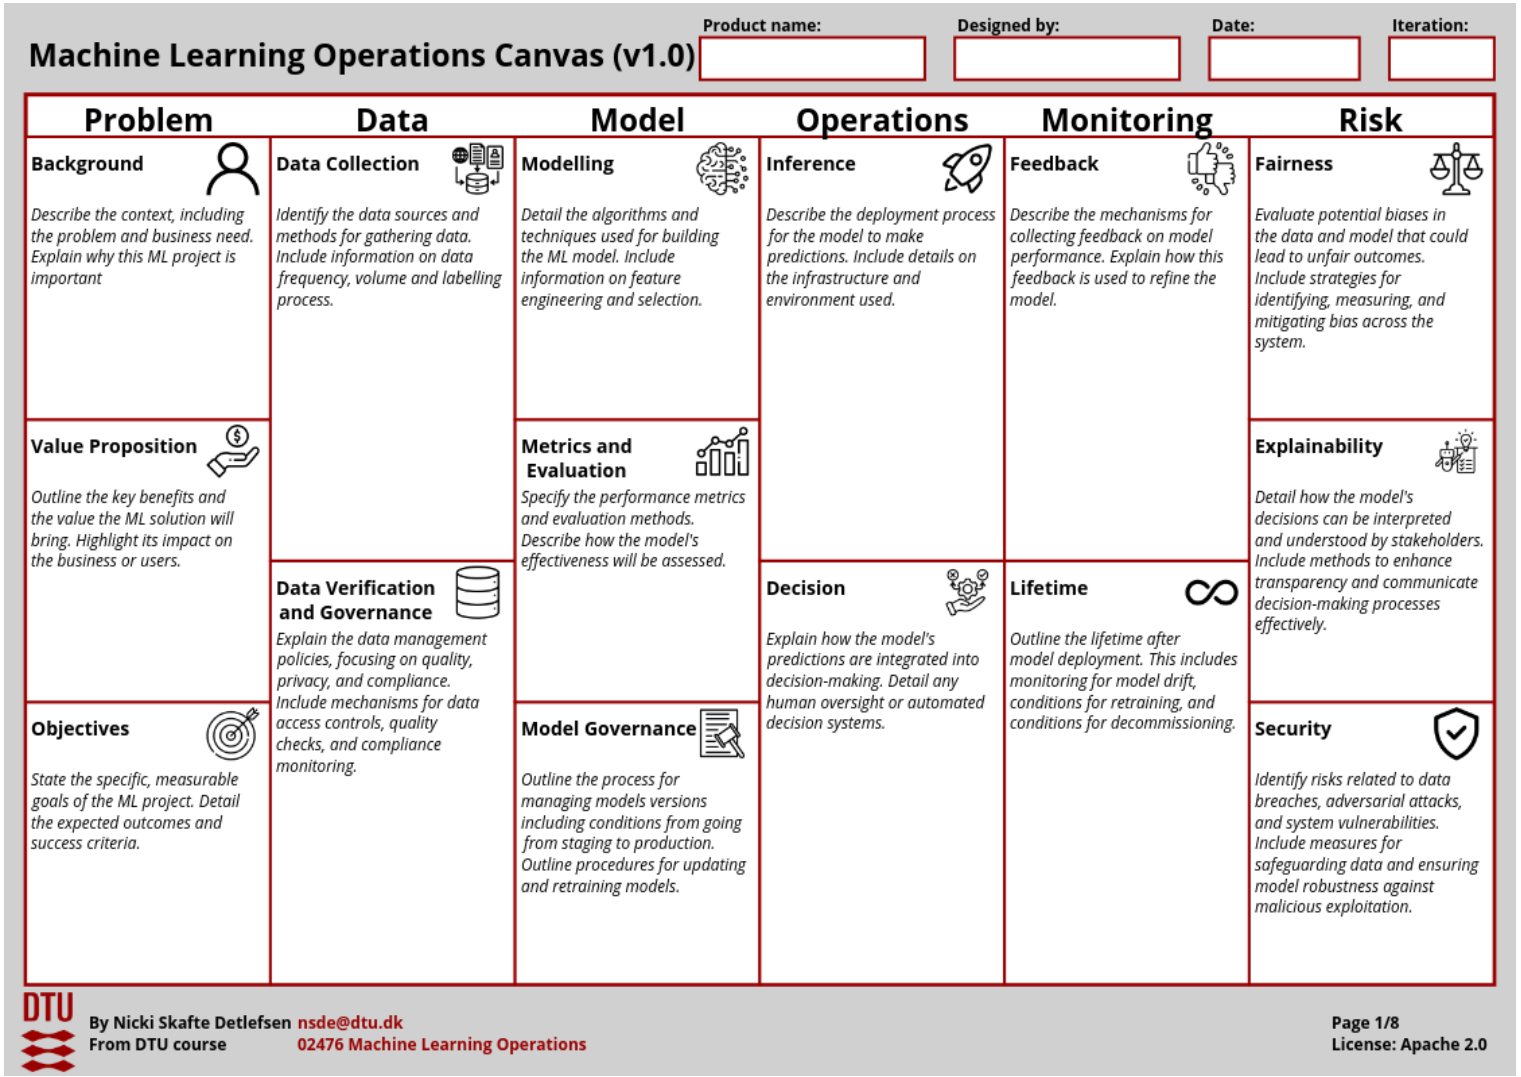

Problem – Value Proposition

- What value will the solution provide to the stakeholders?
- How does this solution improve existing alternatives?
- What specific needs or pain points does it address?
- Who benefits the most from solving this problem?
- How will success be measured in terms of delivered value?

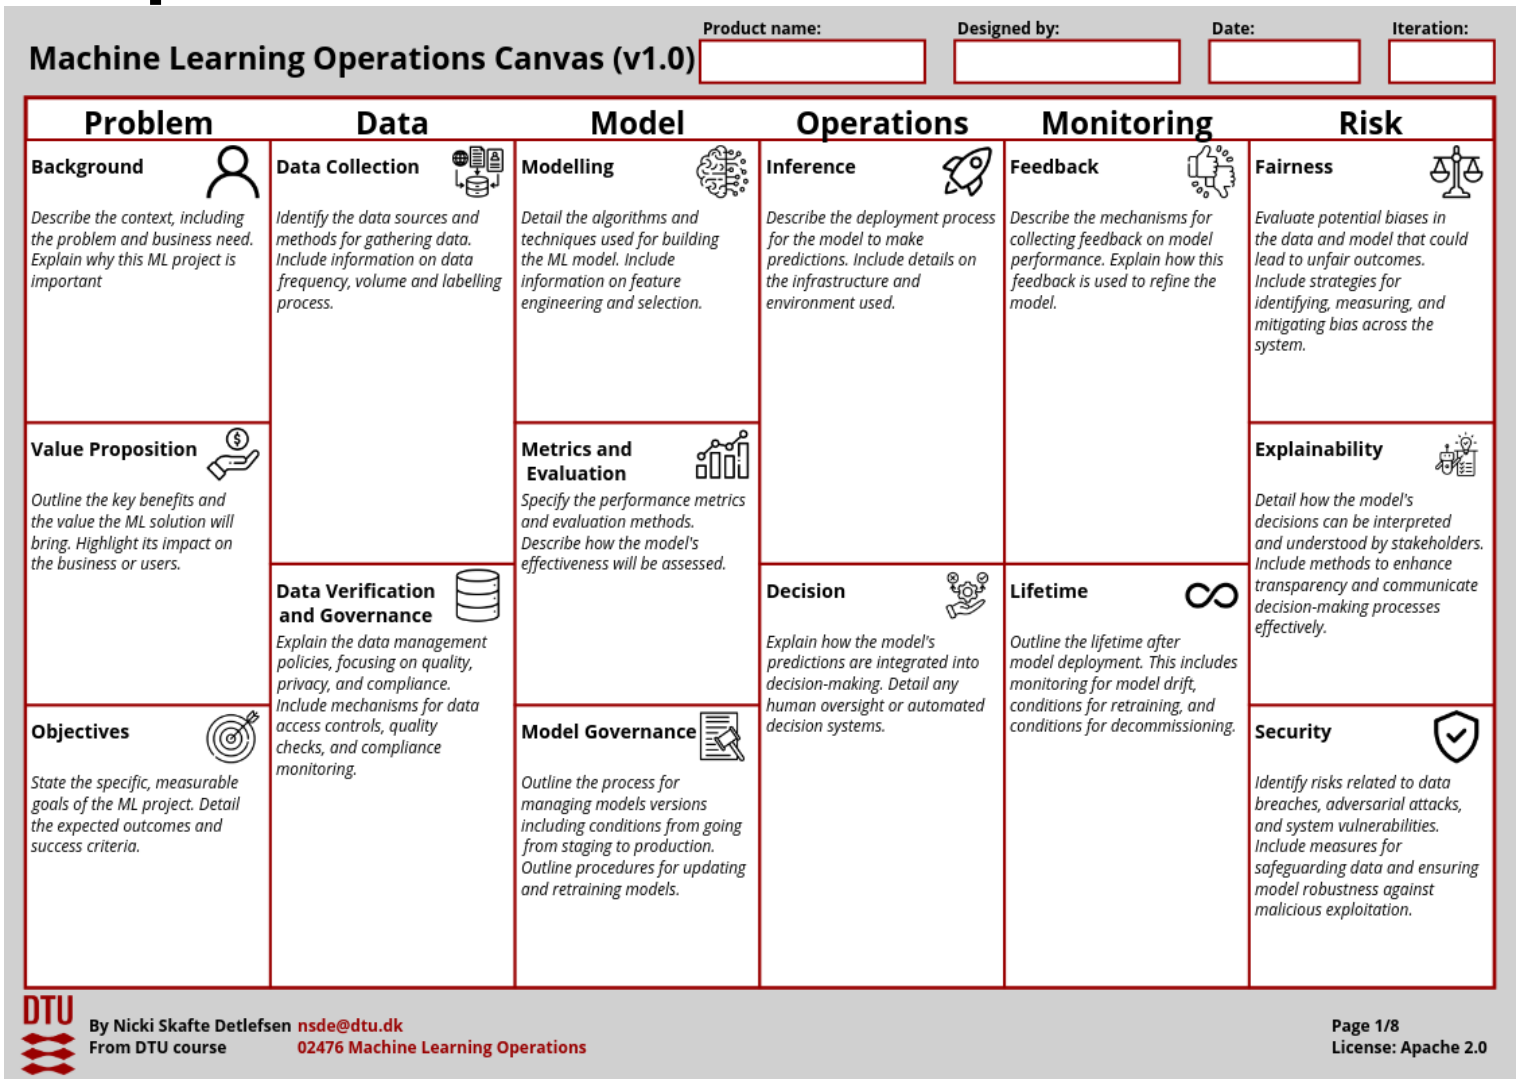

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU
By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Problem - Objectives

- What are the measurable goals for the machine learning system?
- How do these goals align with the overall business or project objectives?
- What are the short-term and long-term goals?
- How will you prioritize conflicting goals, if any?
- Are there clear performance limits or benchmarks to achieve?

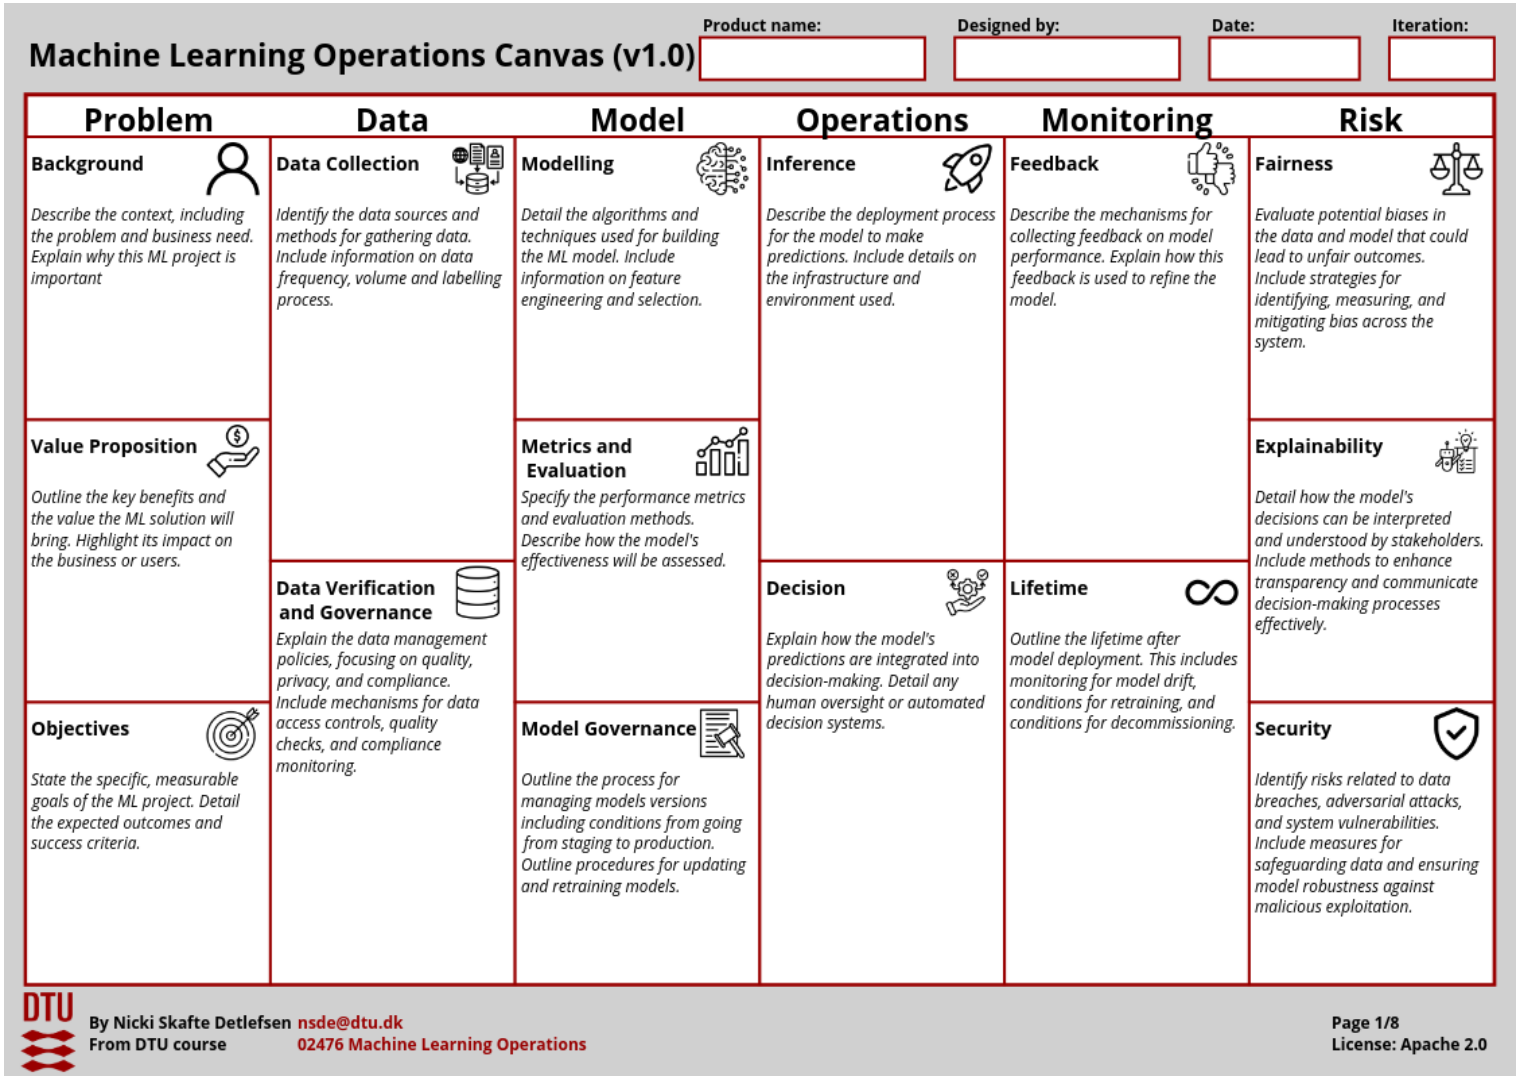

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Data – Data Collection

- What are the primary data sources, and how will the data be accessed?
- Is the data collected at a fixed frequency, in real-time, or as a one-time batch?
- How much data is needed (volume), and how will it be labeled or annotated?
- Are there license, ownership, or copyright issues associated with the data?
- Does the data contain diverse and representative samples to effectively address the problem?

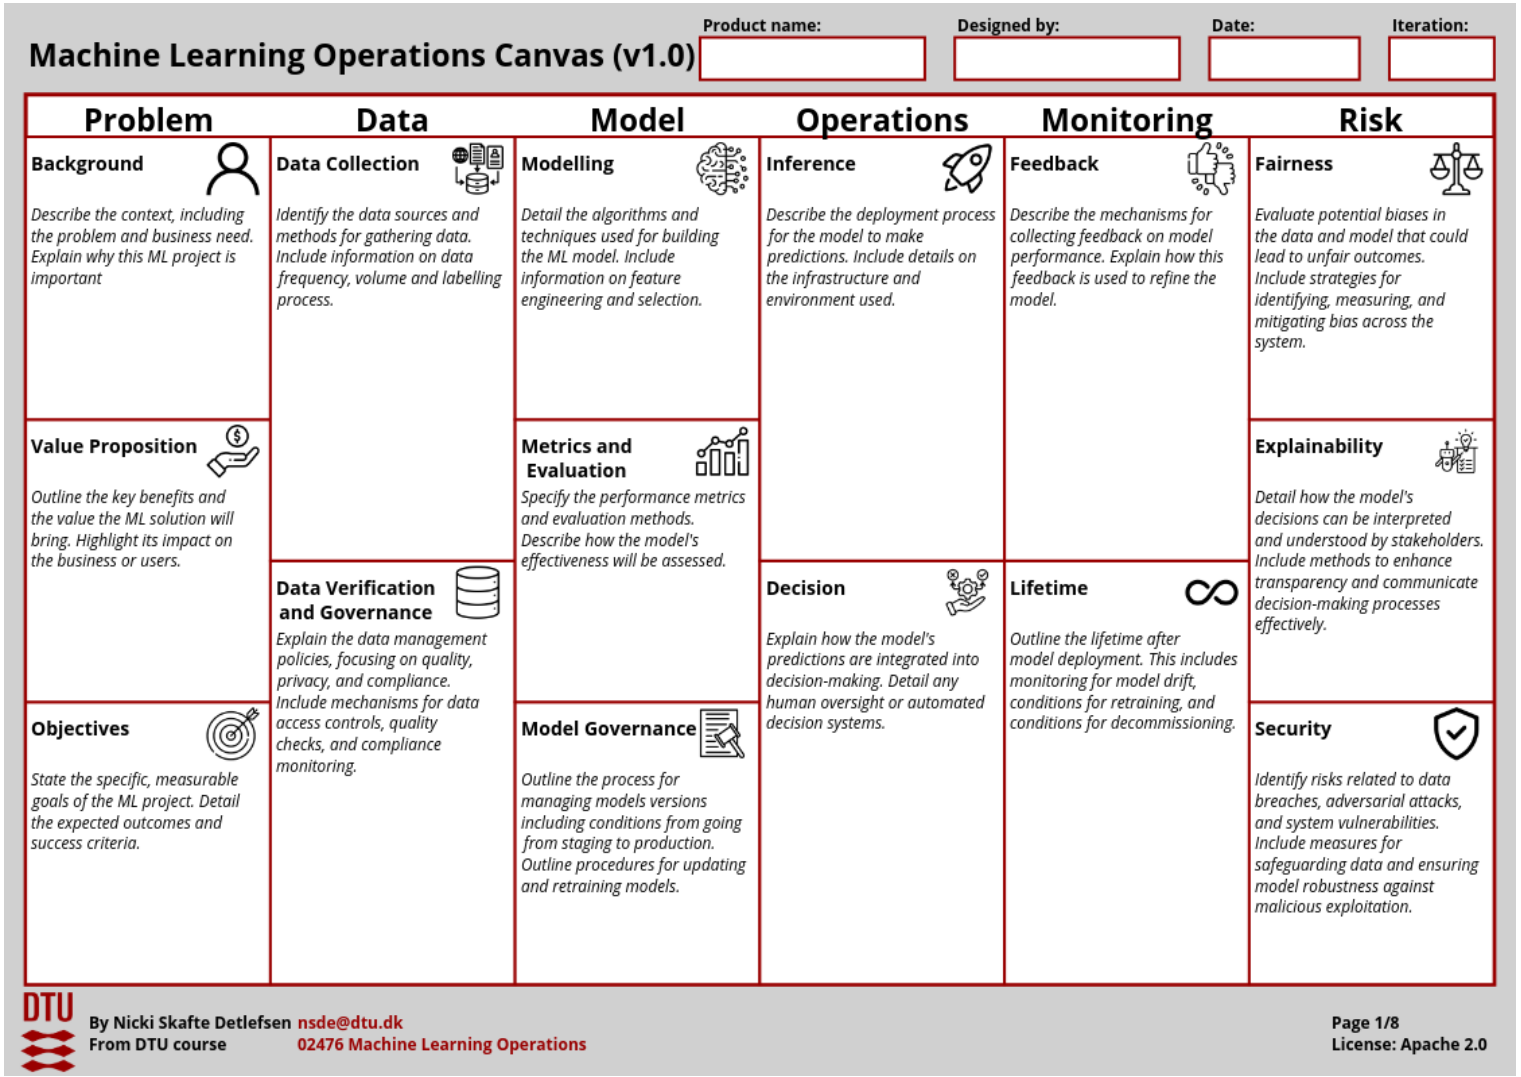

Data – Data Verification and Governance

- What steps will you take to ensure data quality (e.g., handling missing, duplicate, or incorrect data)?
- How will you validate that the data accurately represents the real-world problem?
- Are there policies in place to manage the storage, access, and use of the data?
- How will you handle sensitive data or ensure compliance with privacy laws (e.g., GDPR)?
- What methods will you use to detect and mitigate potential biases in the data?

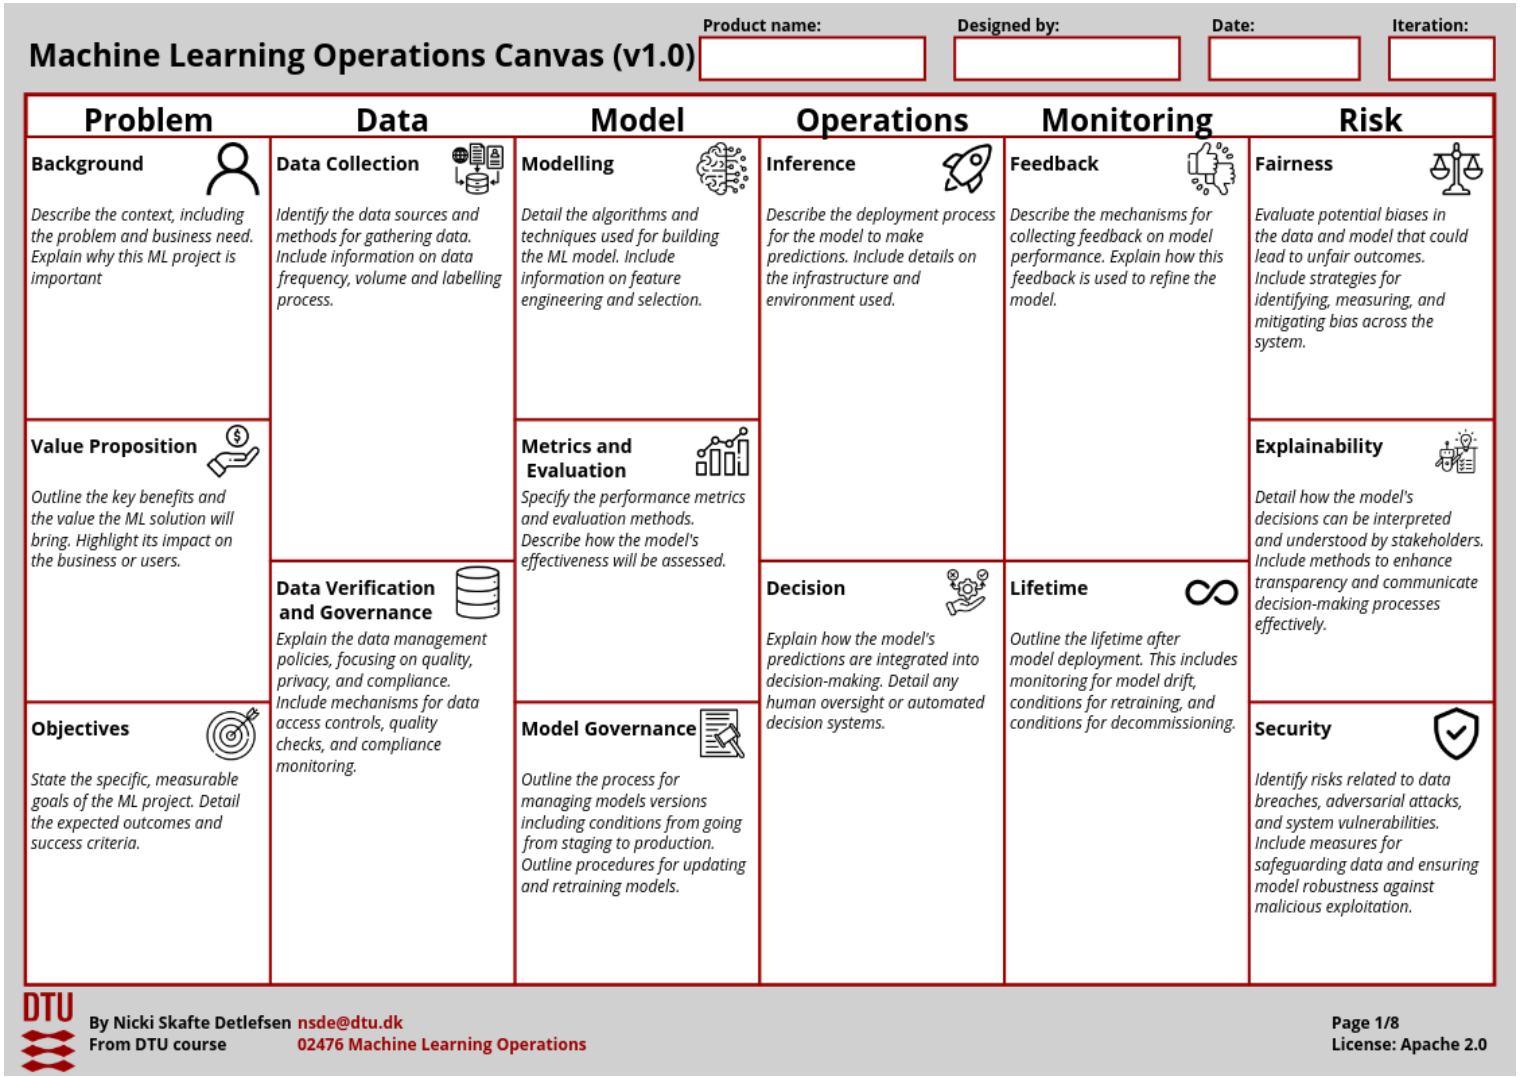

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Model - Modelling

- Which algorithms or techniques will be used to build the model?
- How will you decide which features to include or exclude?
- What is the expected complexity of the model, and is it justified?
- What tools and frameworks will you use to develop the model?
- What trade-offs (e.g., accuracy vs. interpretability) will you need to consider?

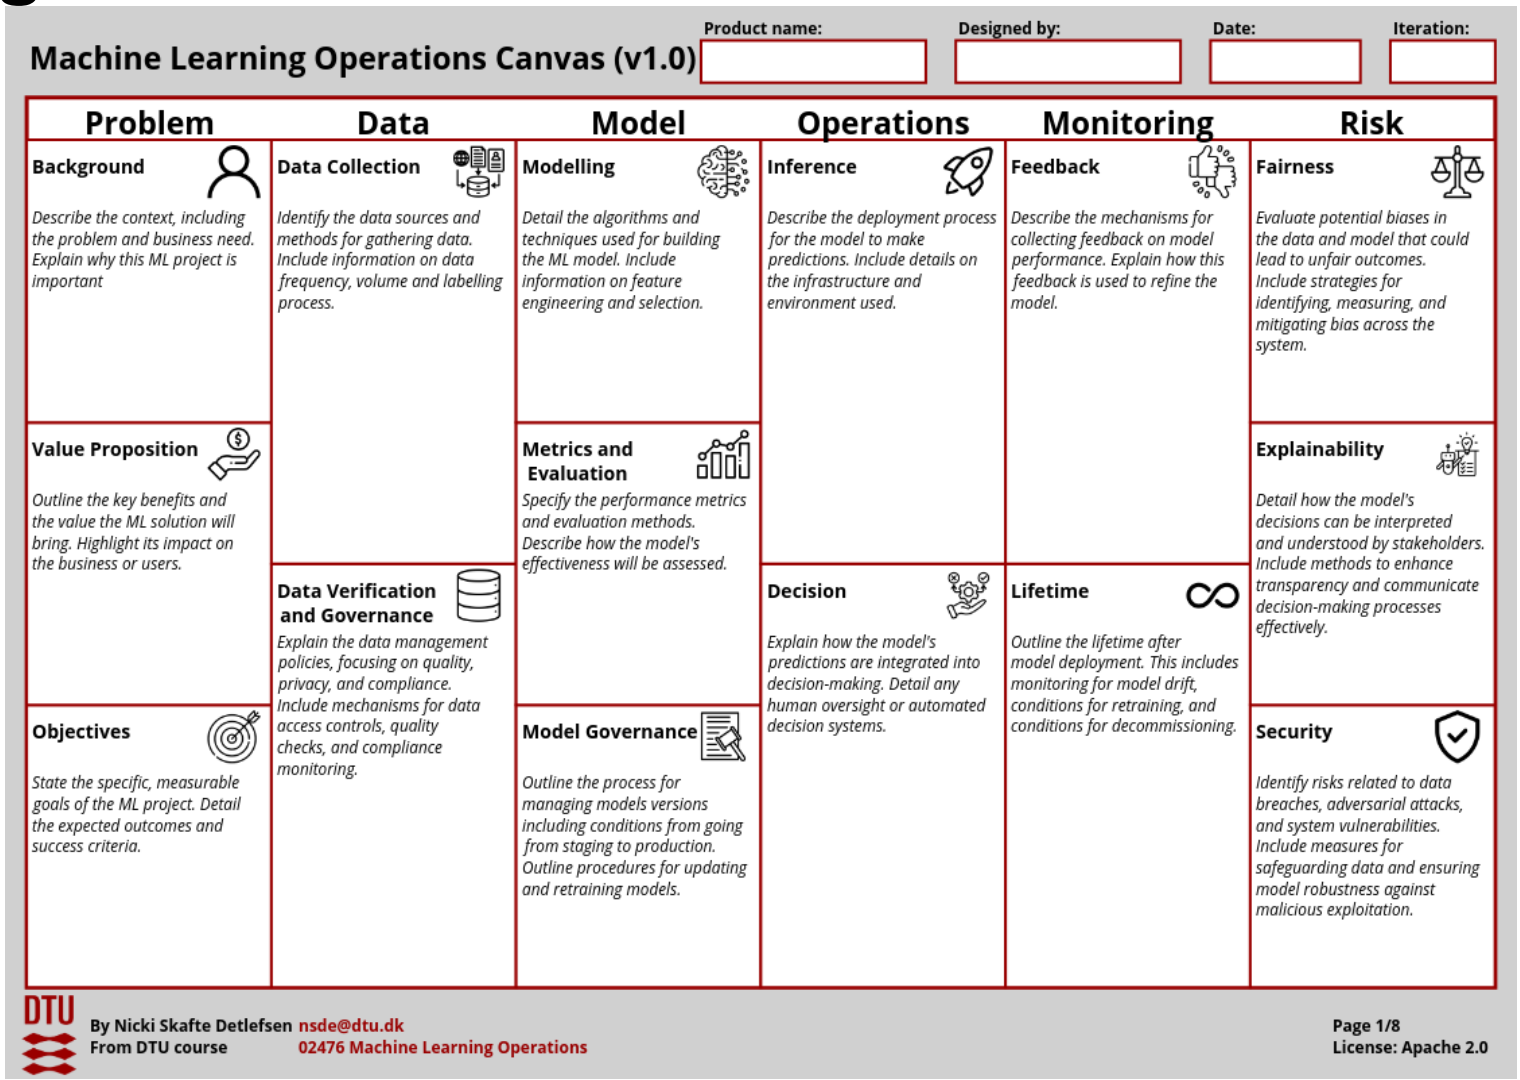

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Model – Metrics and evaluation

- Which metrics will be used to evaluate the model's performance?
- How will you determine if the model meets its objectives?
- What baseline or benchmarks will you compare the model with?
- How will you evaluate the model's performance across different datasets?
- How will you validate the model's robustness to edge cases and noise?

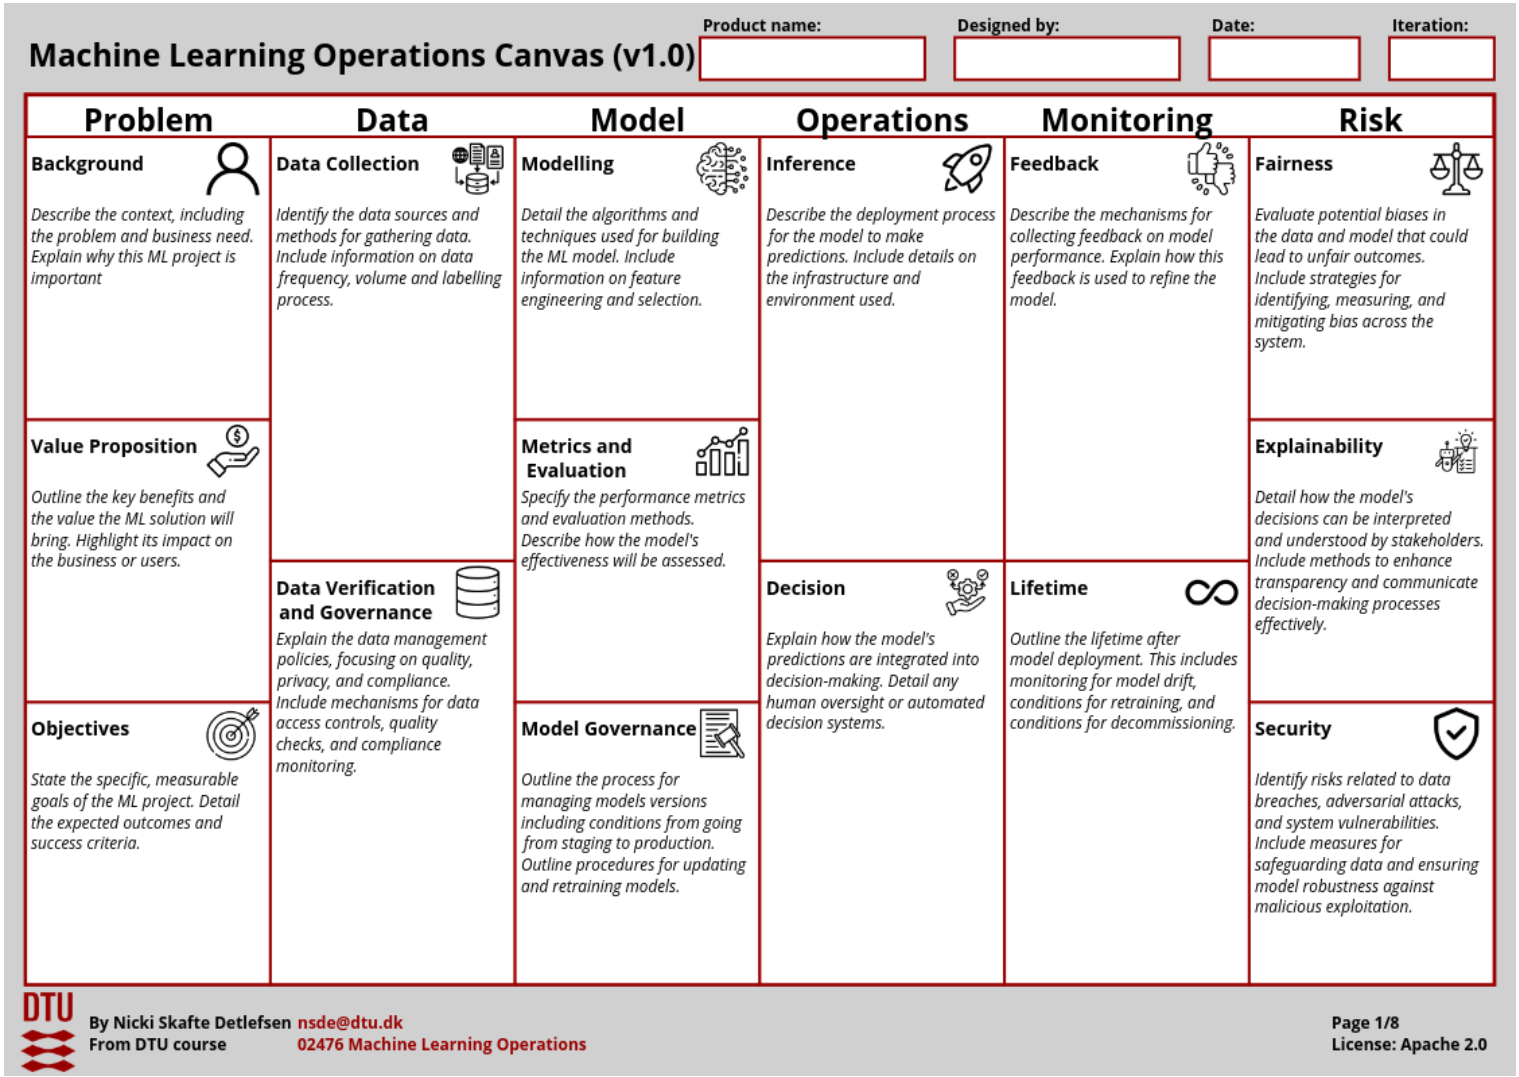

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem

Background

Describe the context, including the problem and business need. Explain why this ML project is important

Value Proposition

Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.

Objectives

State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.

Data

Data Collection

Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.

Data Verification and Governance

Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.

Model

Modelling

Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.

Metrics and Evaluation

Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.

Model Governance

Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.

Operations

Inference

Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.

Decision

Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.

Monitoring

Feedback

Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.

Lifetime

Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.

Risk

Fairness

Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.

Explainability

Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.

Security

Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Model – Model governance

- What processes are in place to document the model's development lifecycle?
- How will you ensure accountability for the model's predictions?
- What security measures are in place to monitor ethical compliance?
- How will you track changes and versions of the model over time?
- Are clear roles and responsibilities defined for maintaining the model?

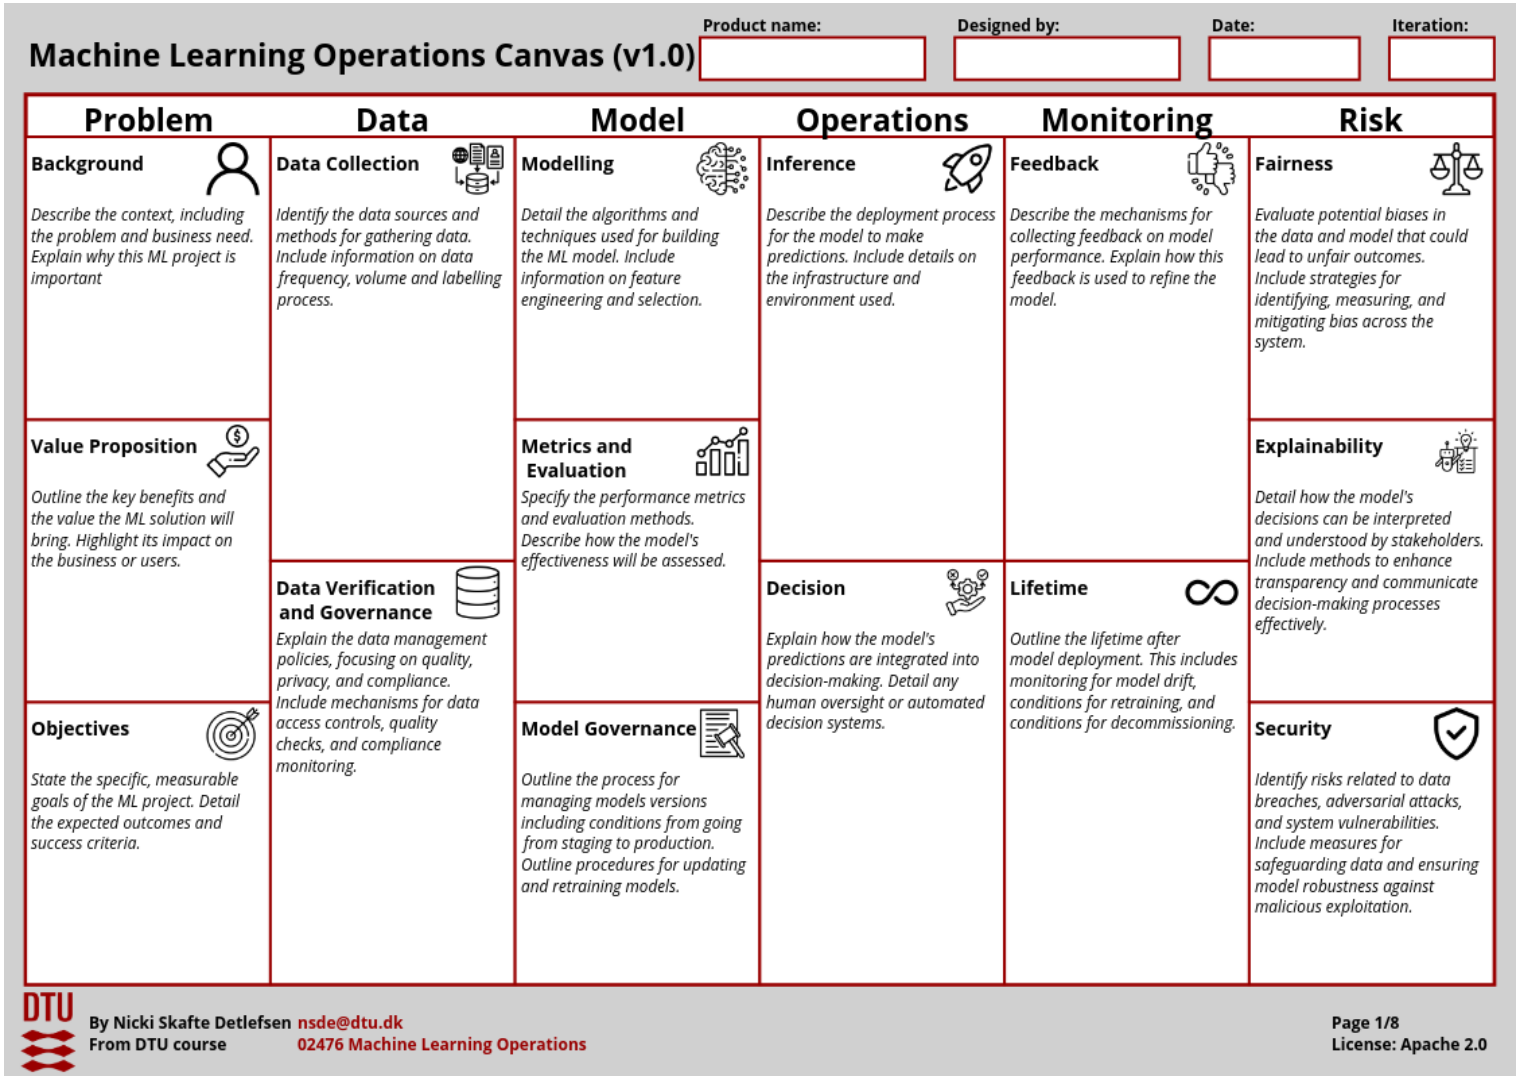

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Operations - Inference

- How will the model's predictions be generated in production?
- What infrastructure is necessary to support inference at scale?
- How will latency and throughput requirements be met?
- What methods will be used to handle errors or failed predictions?
- What format will the predictions be delivered in for further use?

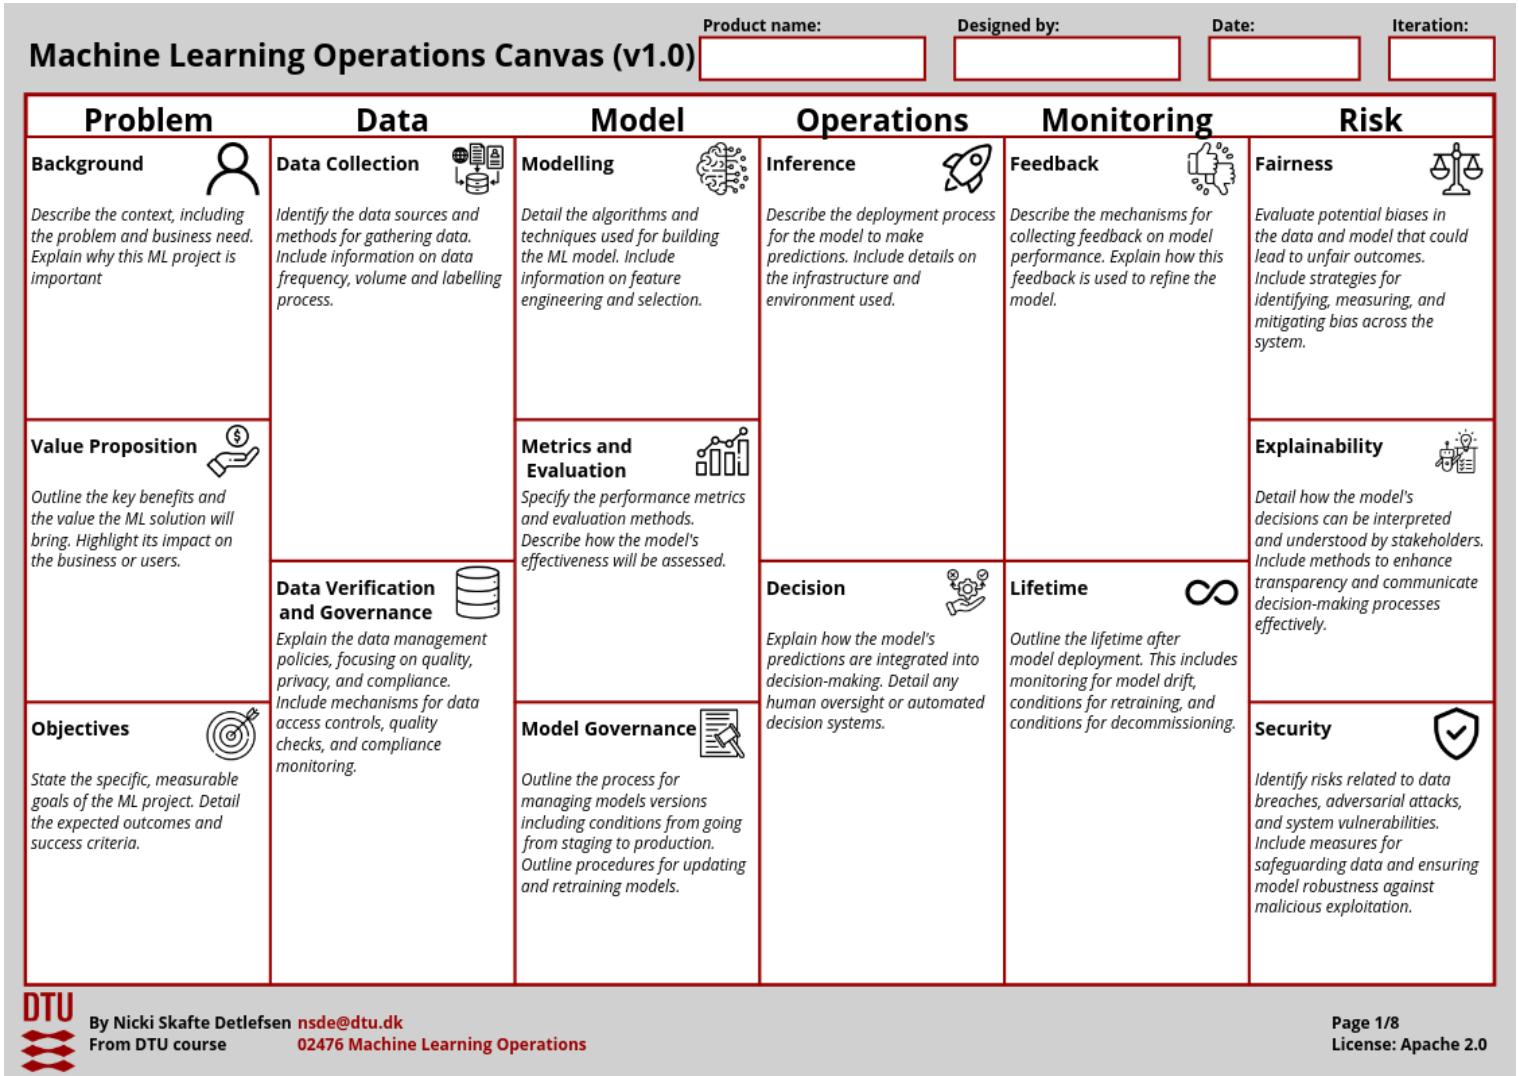

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Operations - Decision

- How will predictions be integrated into decision-making processes?
- Who or what will act on the model's predictions?
- Are there automatic or manual checks in place for critical decisions?
- What are the expected consequences of these decisions for stakeholders?
- How will the model's output be communicated to end-users or stakeholders?

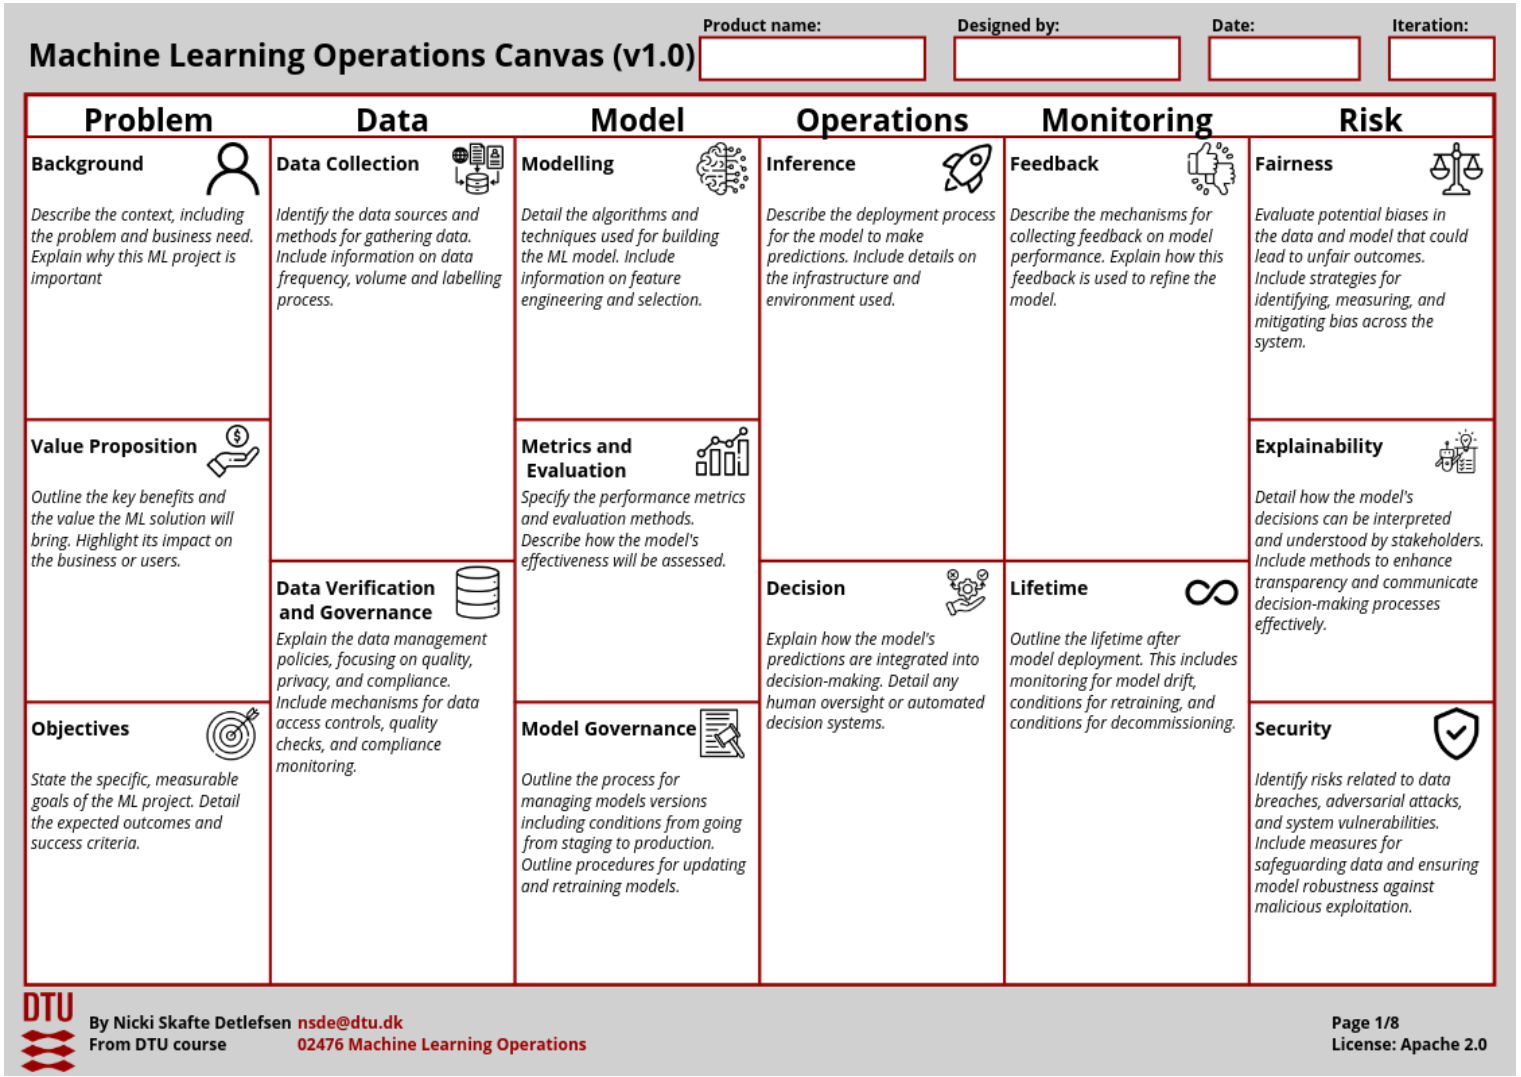

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU
By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations
Page 1/8
License: Apache 2.0

Monitoring - Feedback

- What mechanisms are in place to collect feedback on the model's predictions?
- How will feedback be used to improve the model over time?
- Who will be responsible for analyzing and acting on feedback?
- What channels will you use to collect user or system feedback?
- How will you handle negative feedback or identified performance issues?

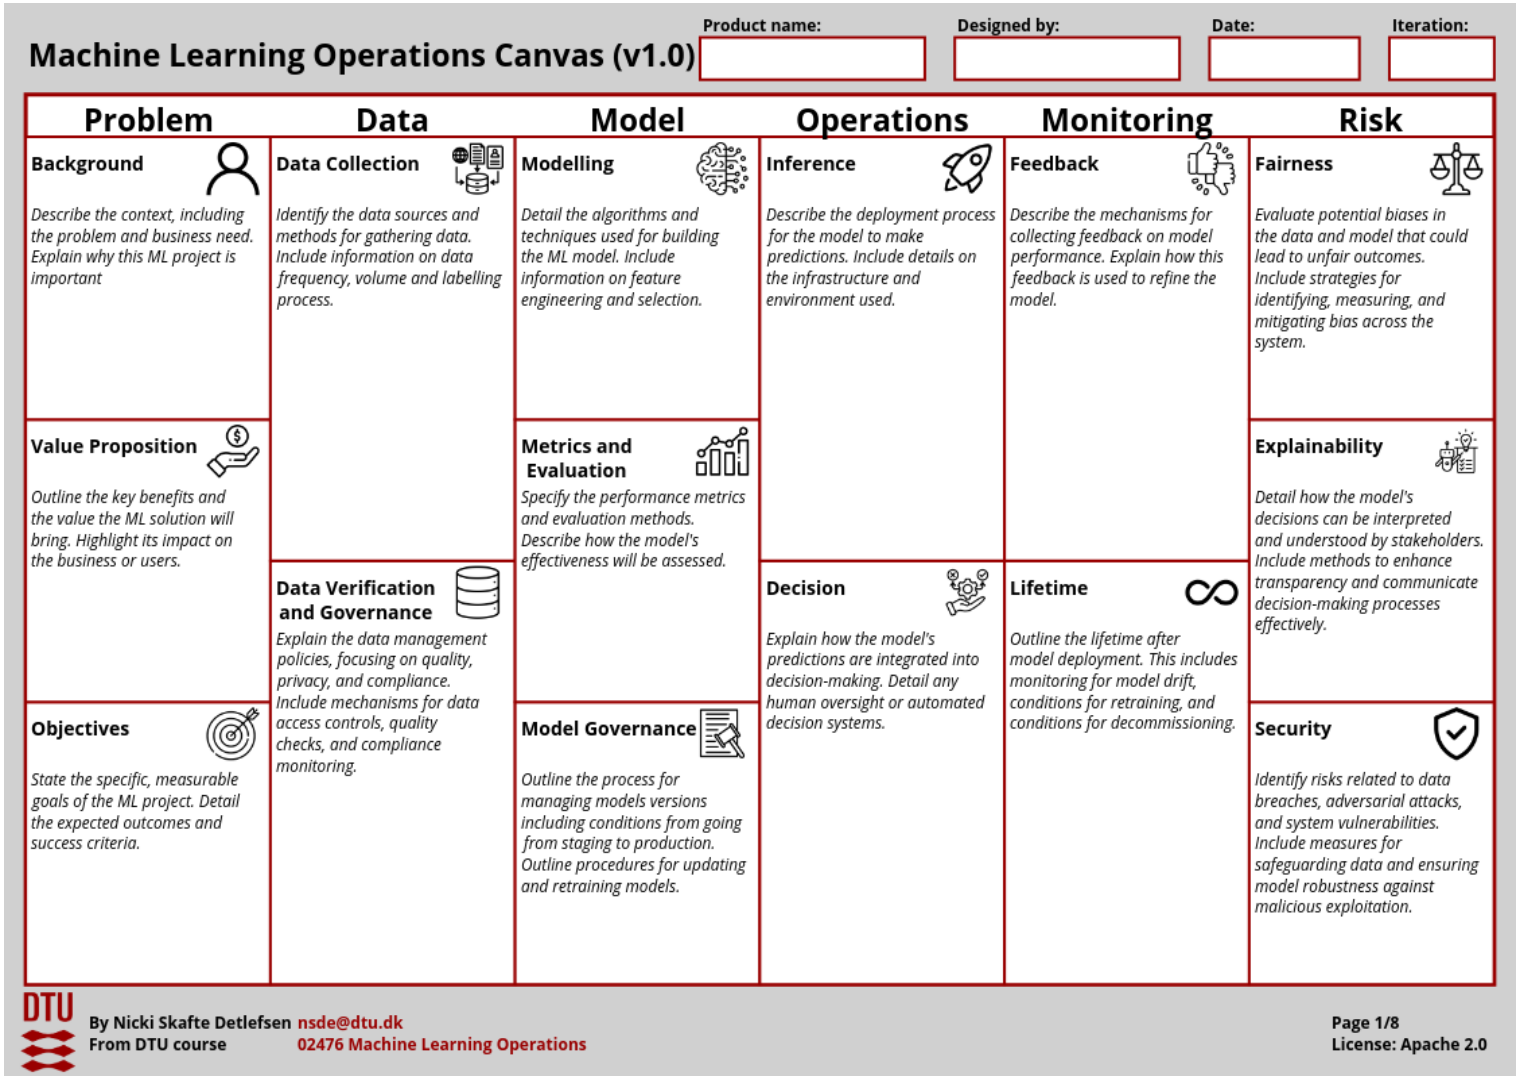

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU
By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations
Page 1/8
License: Apache 2.0

Monitoring - Lifetime

- What is the expected lifespan of the model before it needs to be retrained or replaced?
- How will you monitor for model drift or performance degradation?
- What triggers will indicate the need for model updates?
- What processes are in place to phase out outdated models?
- How will you handle dependencies and compatibility over the model's lifespan?

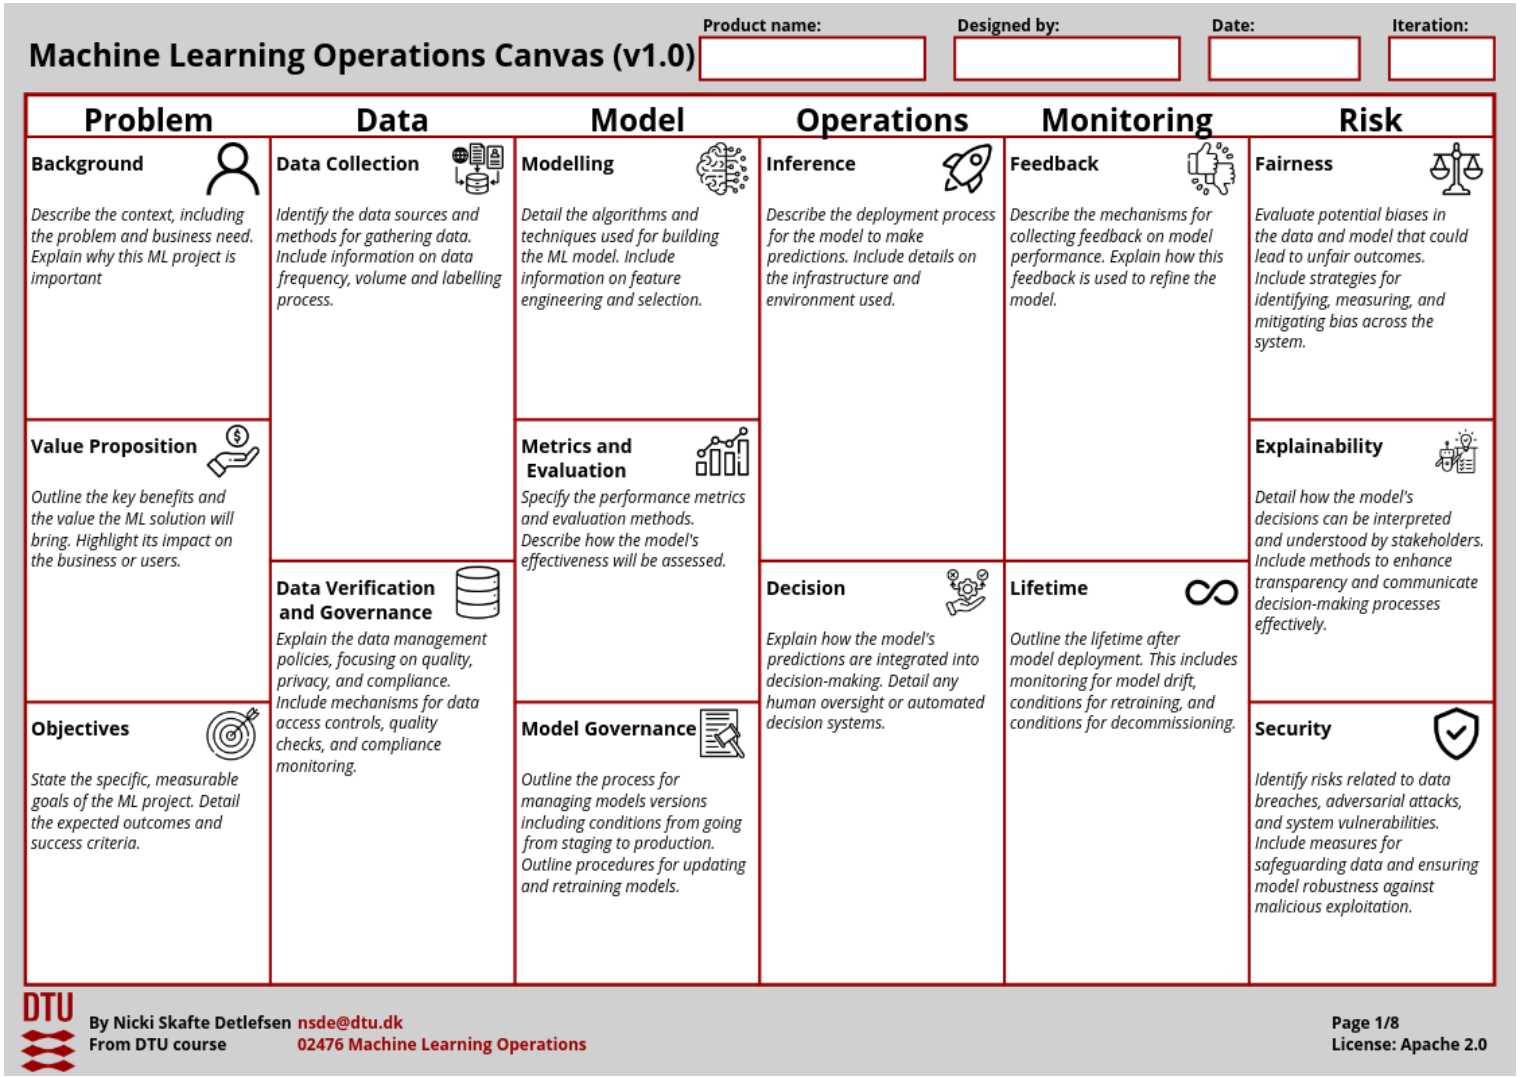

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Risk - Fairness

- What potential biases could arise in the data or model, and how will you identify them?
- How will you ensure fair outcomes for all stakeholder groups?
- What metrics or methods will you use to measure fairness?
- What steps will you take to mitigate bias during development and deployment?
- Are there any groups or scenarios that could be disproportionately affected?

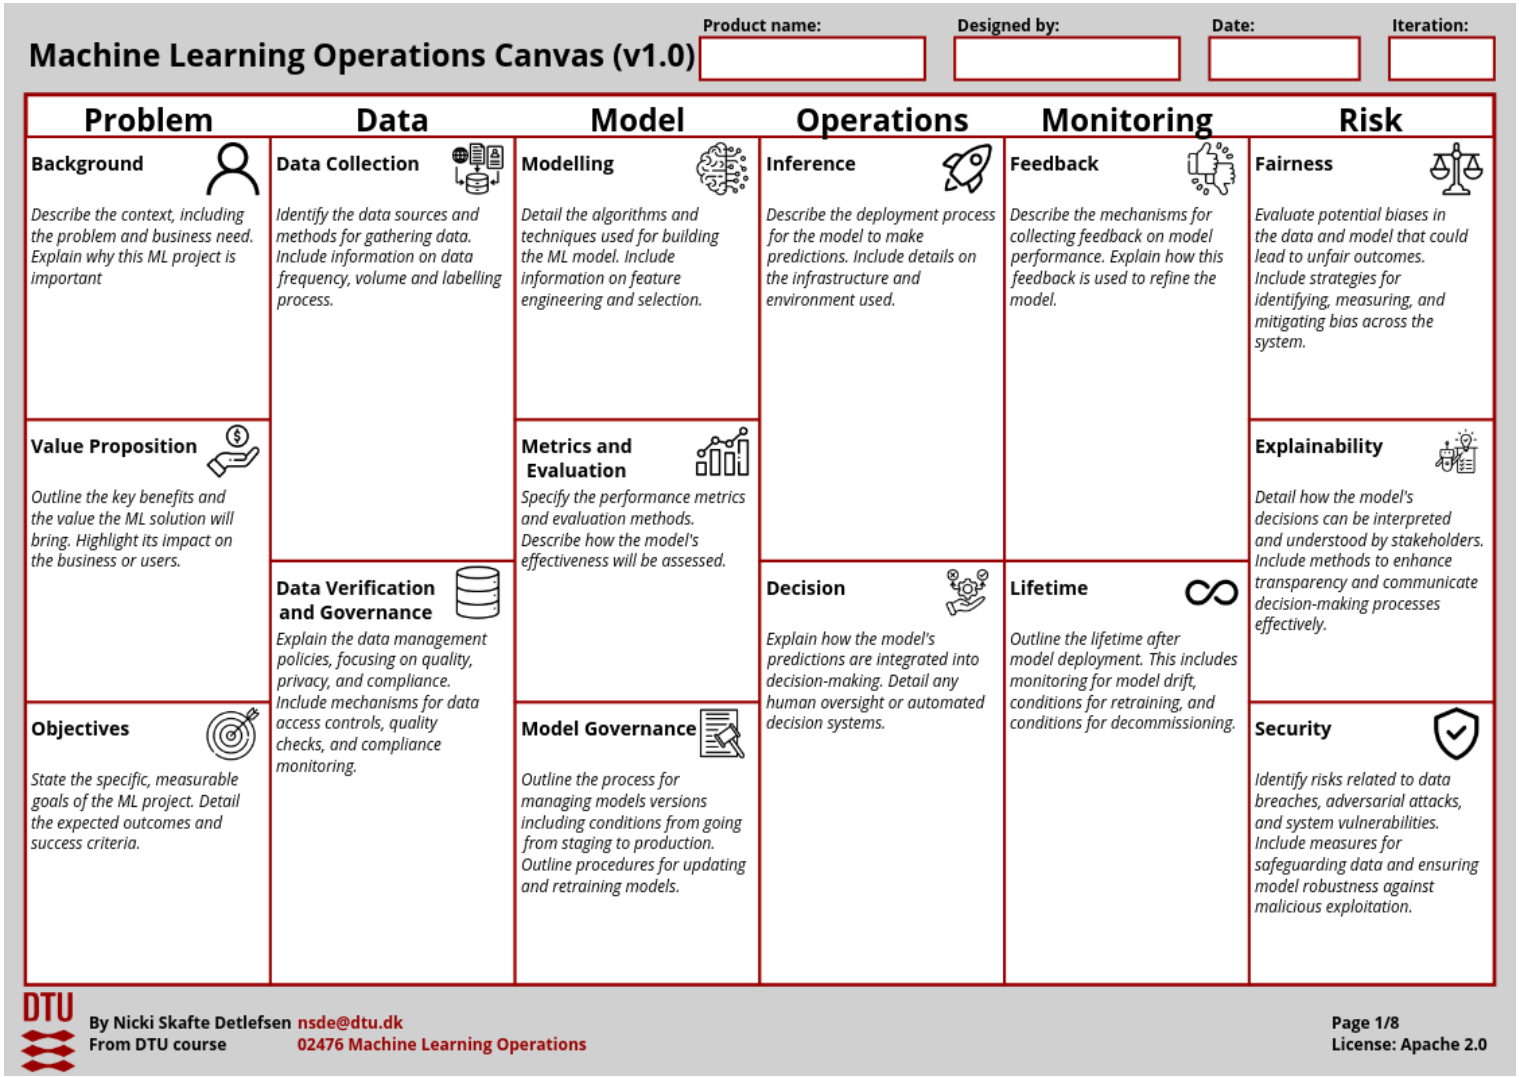

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Risk - Explainability

- How will you ensure that the model's decisions can be interpreted by stakeholders?
- What tools or techniques will you use to improve explainability?
- How will you communicate the model's decision-making process to non-technical audiences?
- Are there trade-offs between explainability and performance to consider?
- What level of transparency is required for regulatory or ethical purposes?

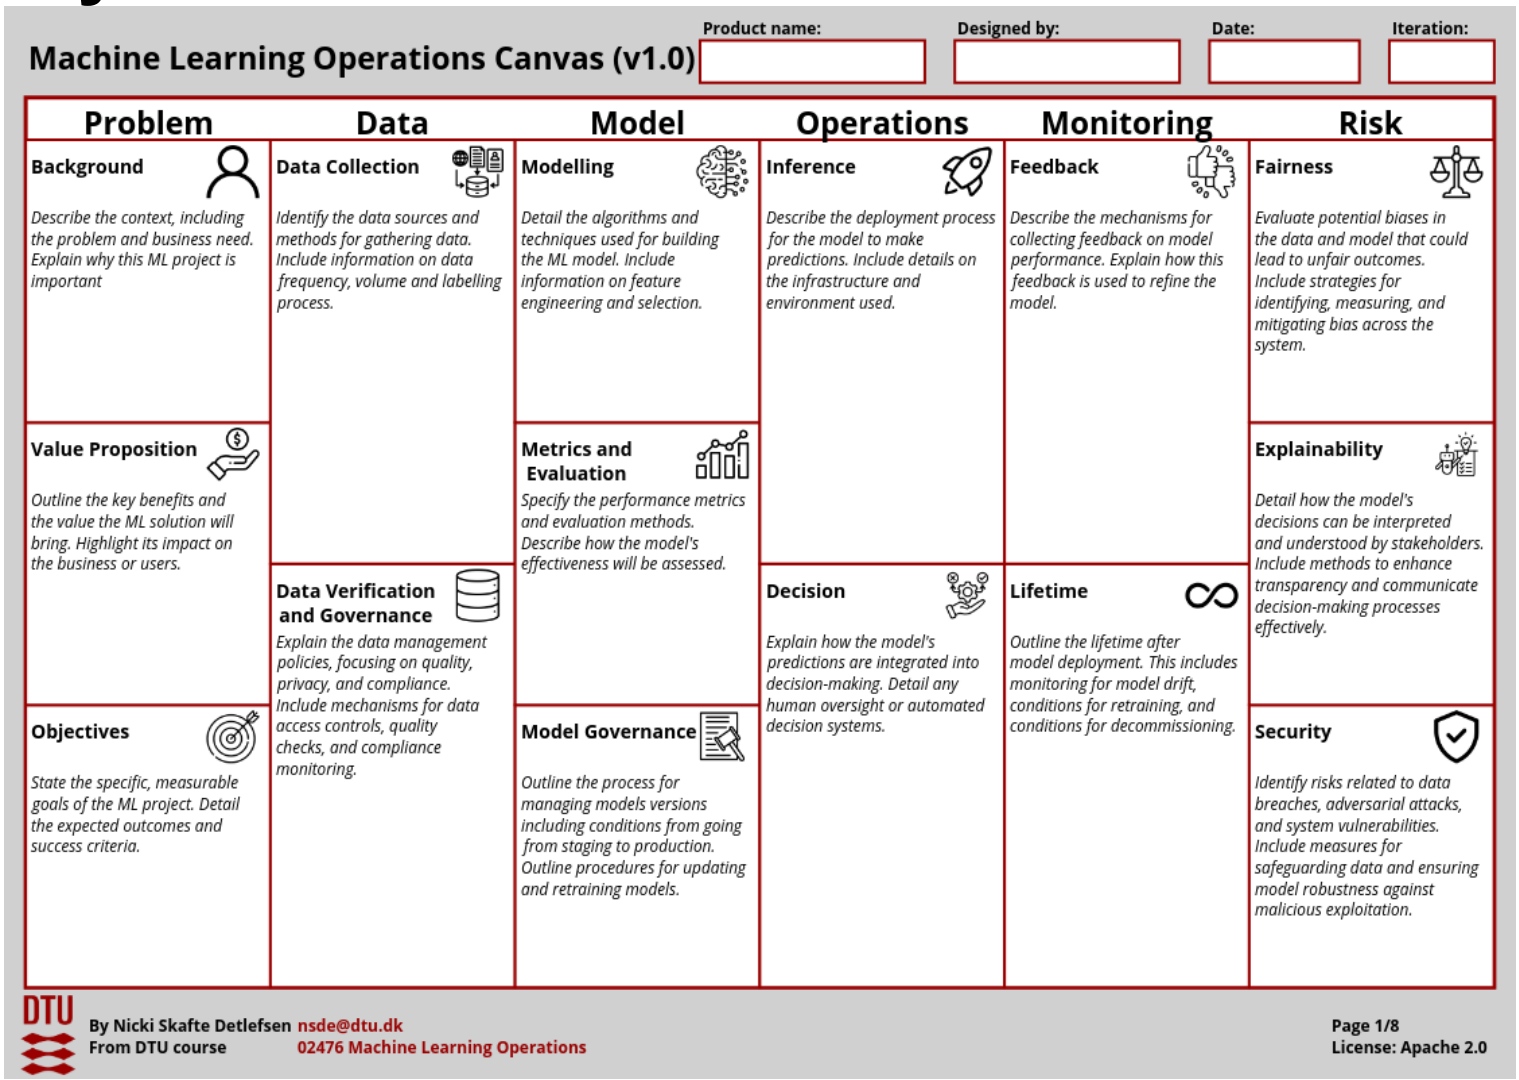

Risk - Security

- What risks exist for data breaches or leaks, and how will you mitigate them?
- How will you protect against adversarial attacks on the model?
- What processes are in place to ensure data protection during and after model usage?
- How will you monitor and address vulnerabilities in the implemented system?
- Are there compliance requirements related to security that need to be addressed?

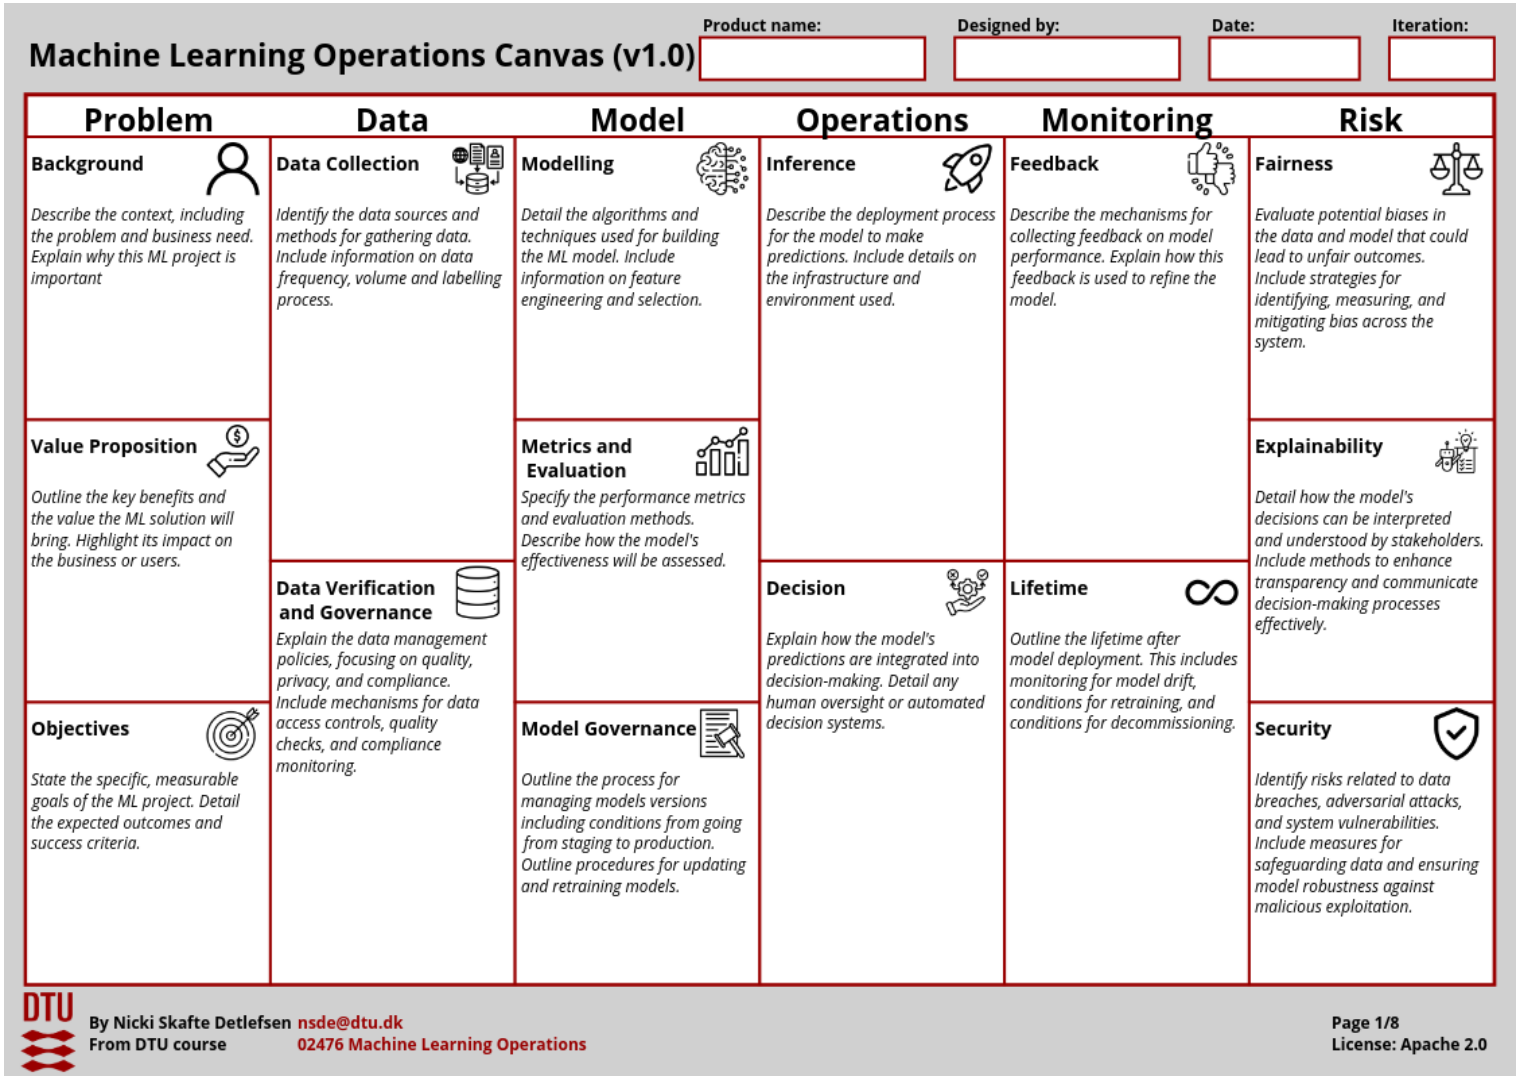

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

DTU
By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations
Page 1/8
License: Apache 2.0

Machine Learning Operations Canvas

How do I answer these questions?

- Discussion with stakeholders
- Research
- Iterative updates

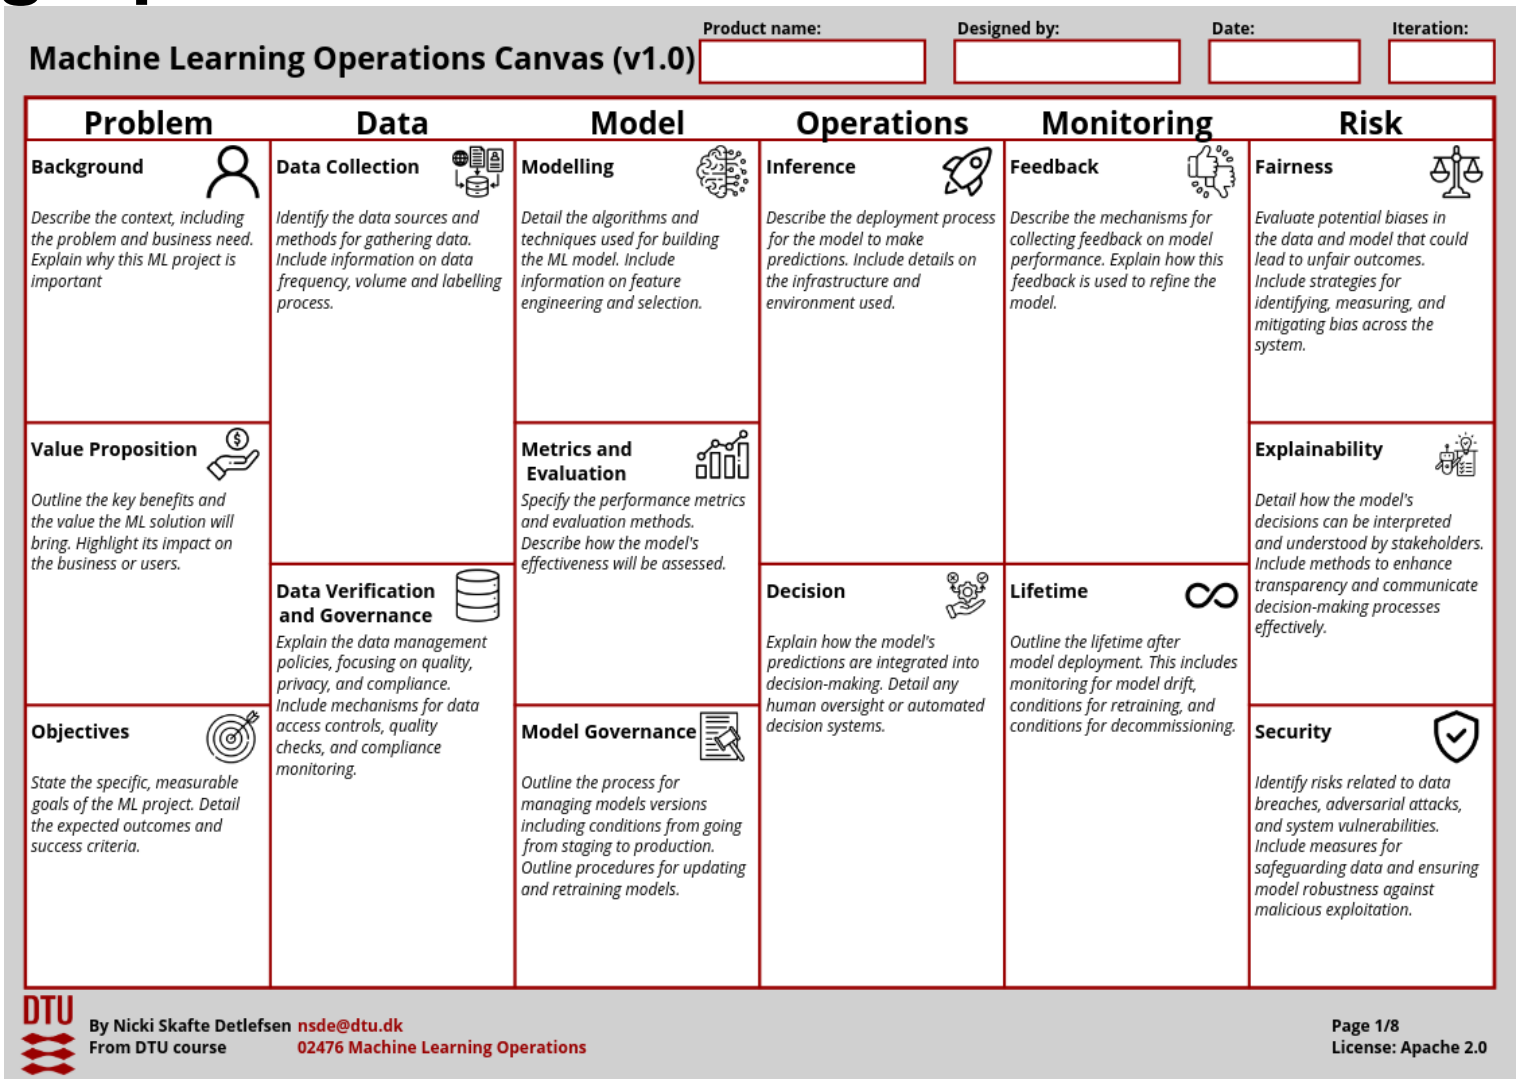

Case description

Description:

Company X operates in the pharmaceutical industry and produces life-saving medicine for its customers. Medicine X is sold in small bottles. In very rare cases, small particles may enter the bottles along with the medicine, contaminating it and posing a potential risk to customers.

Focus:

- Patient safety
- The solution must run in real-time

What did we do?

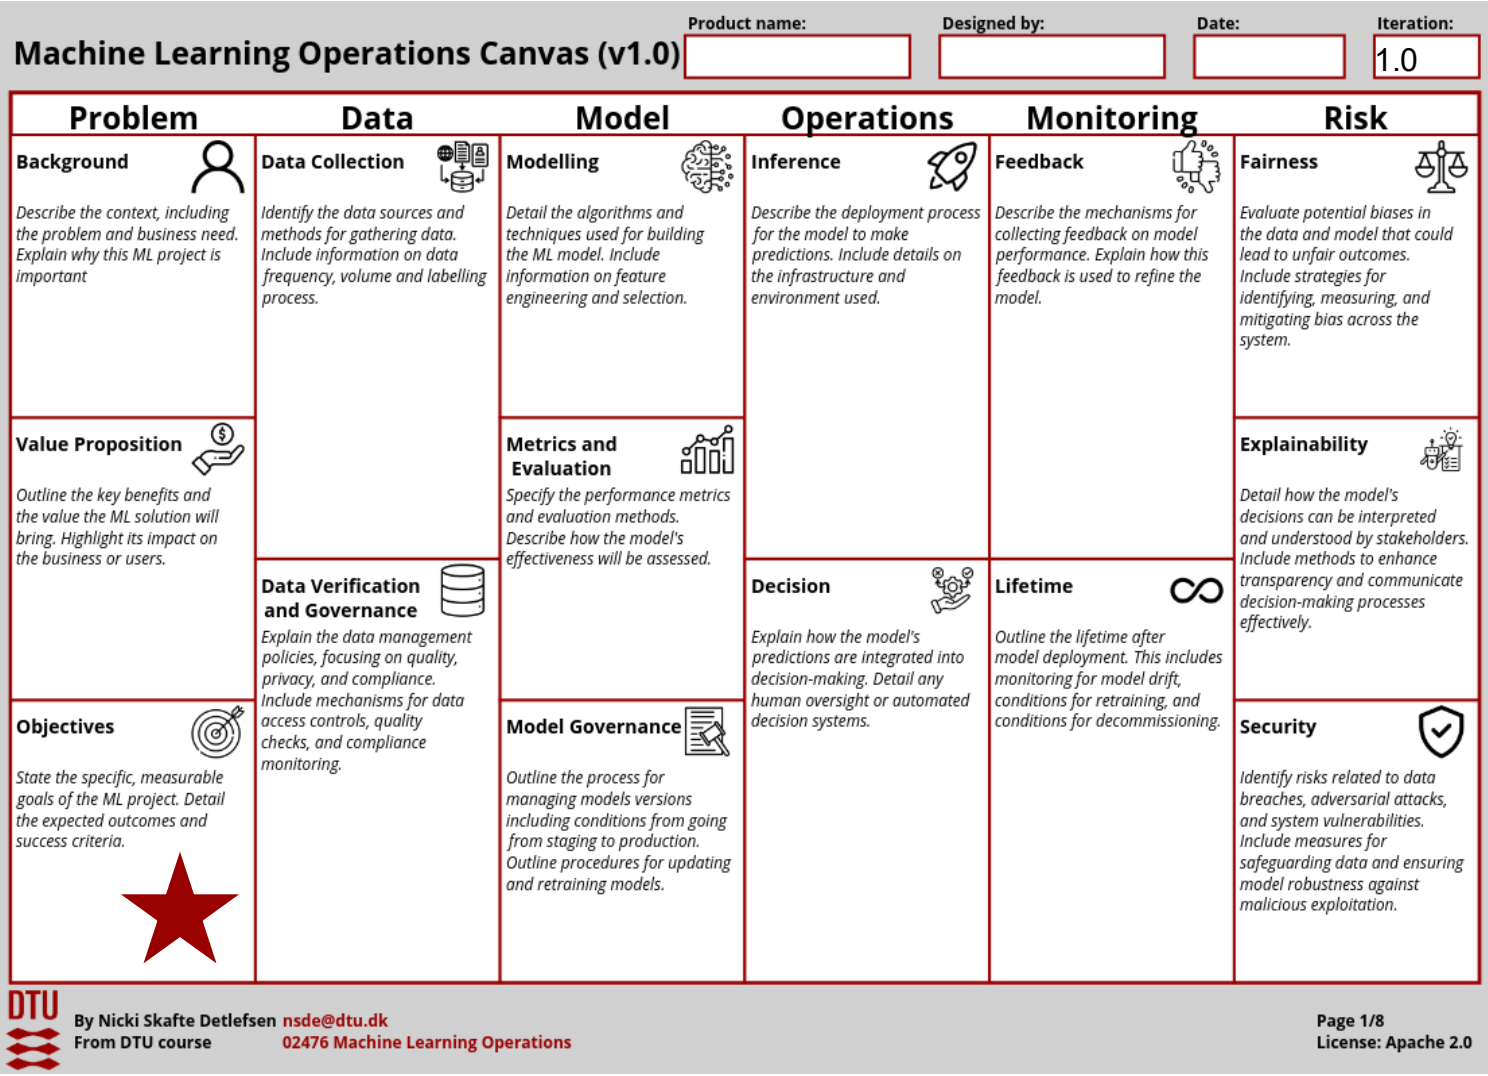

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration: 1.0

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Identified main criteria

Model

1. No contaminated product should go to the user
2. Reduce product discarded

Operations

1. Inference needs to run faster 1sec
2. Inference should preferably run faster than 0.4sec

What did we do?

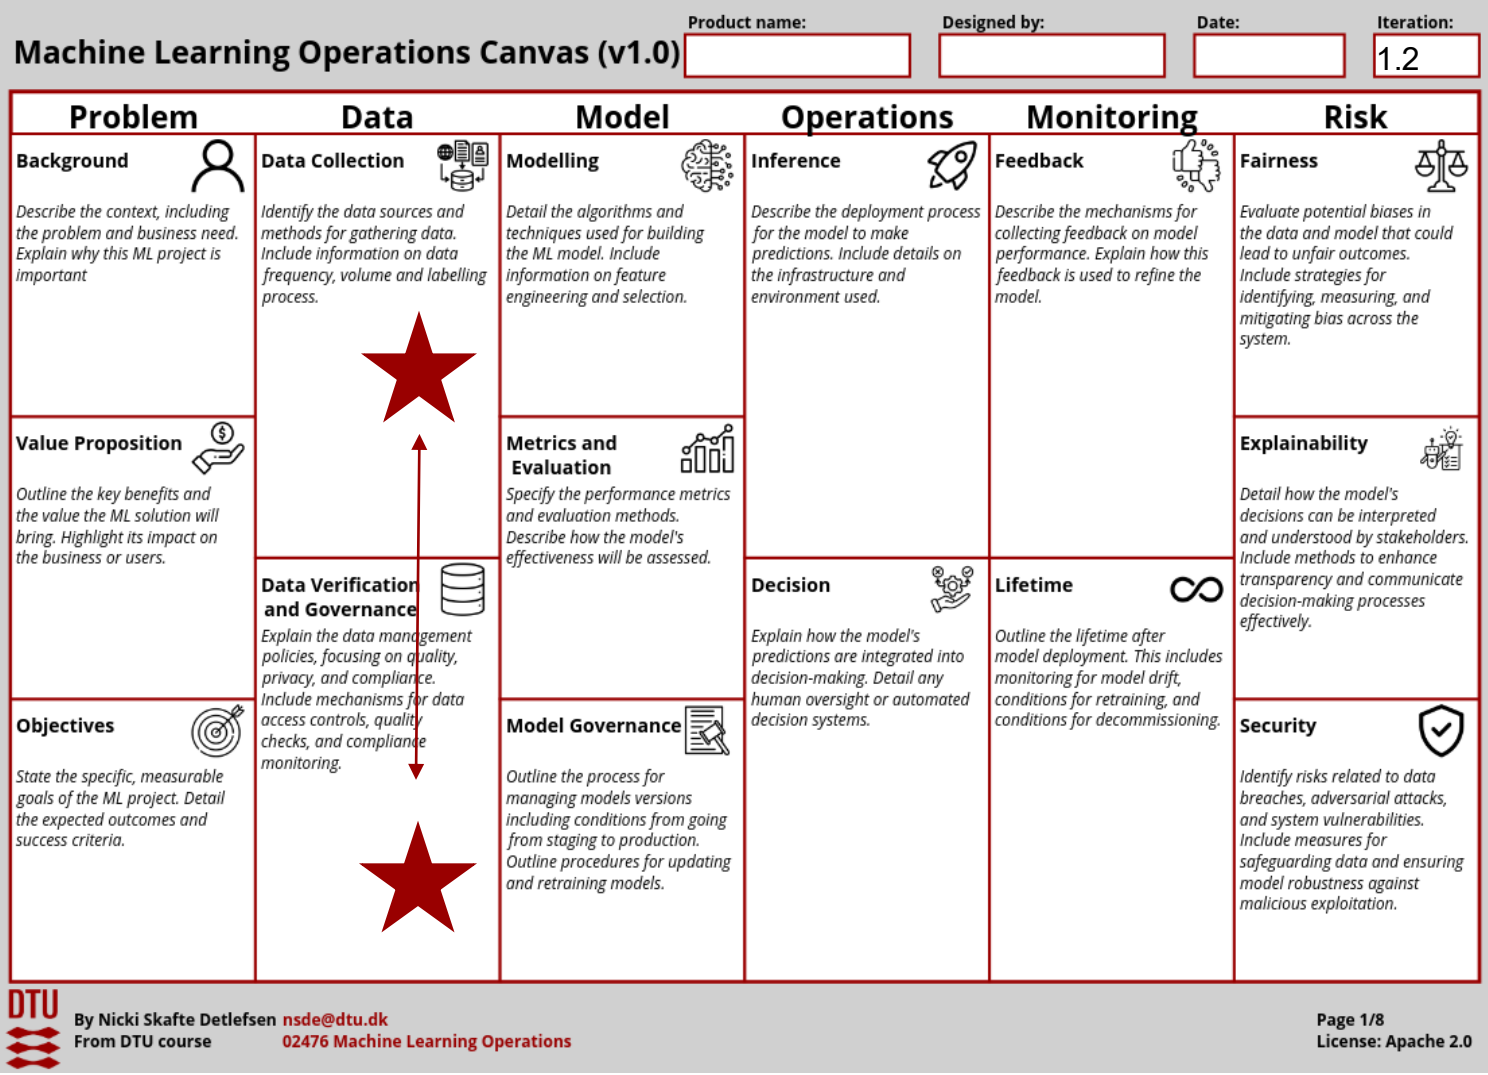

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration: 1.2

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Install single camera at some point during conveyor belt

Because rare event:

1. Add contaminant to bottles
2. Run bottles through setup
3. Automatic labelling of images

Because the contaminant is in a liquid it was not visible on all images -> change data collection setup to include more angels and repeat

What did we do?

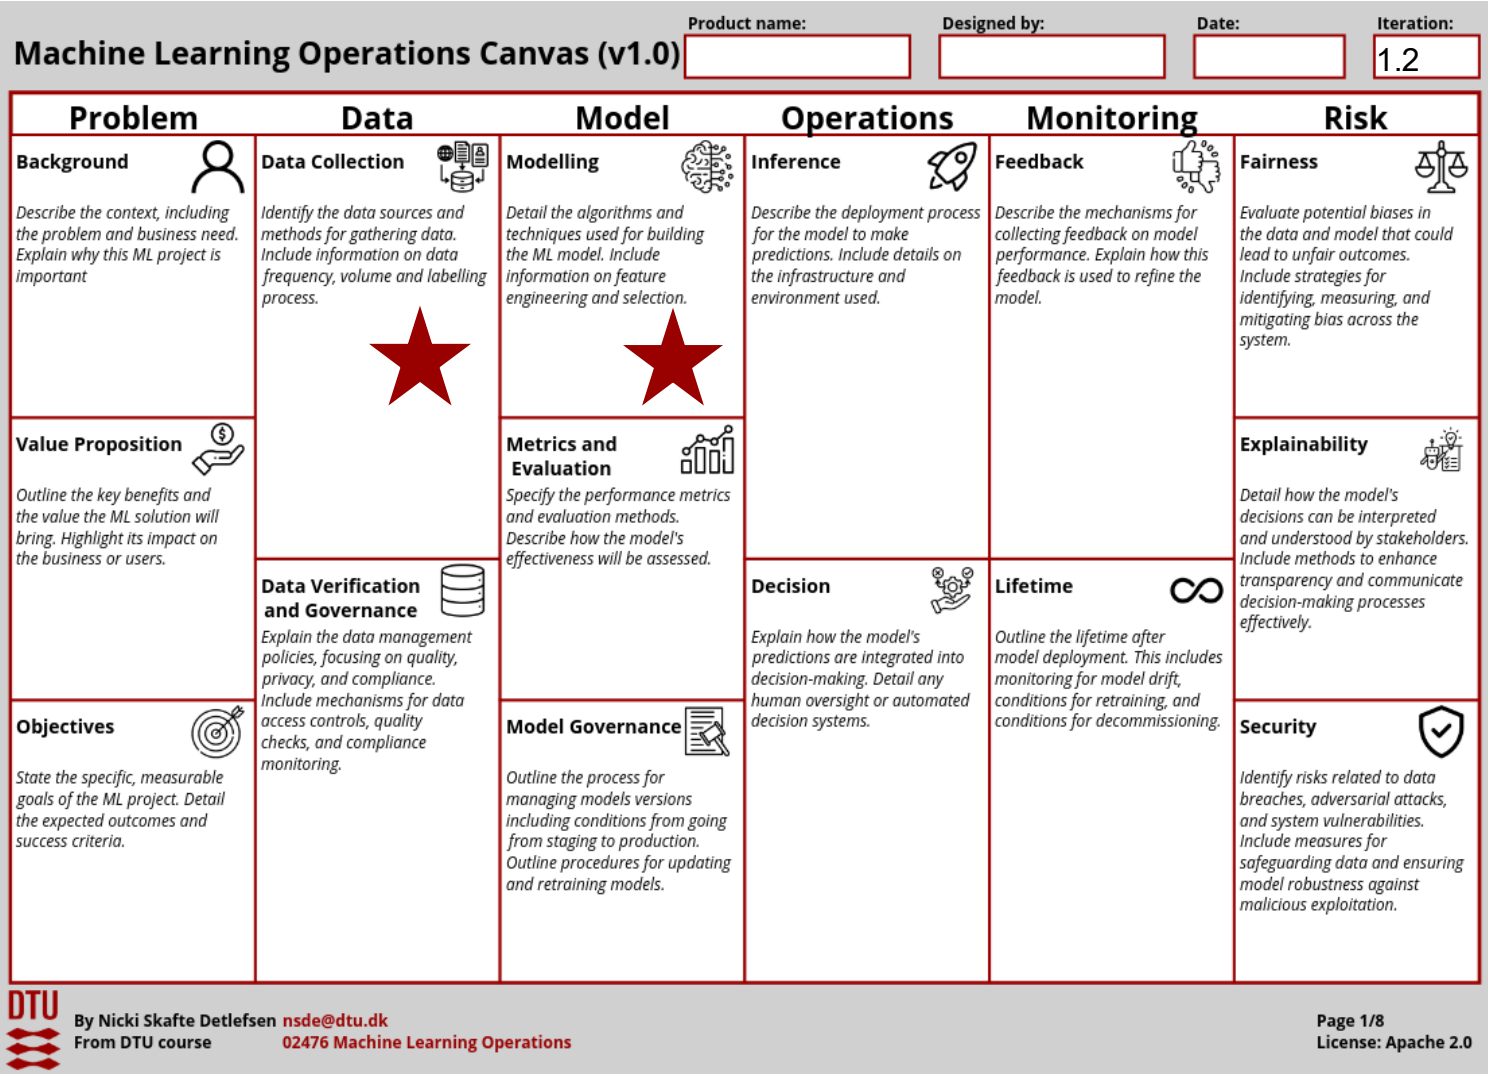

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration: 1.2

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Development of model

1. Doing classification per image -> good but not good enough performance

2. Doing classification per set of images (multiple angels) -> no bad vials where approved

What did we do?

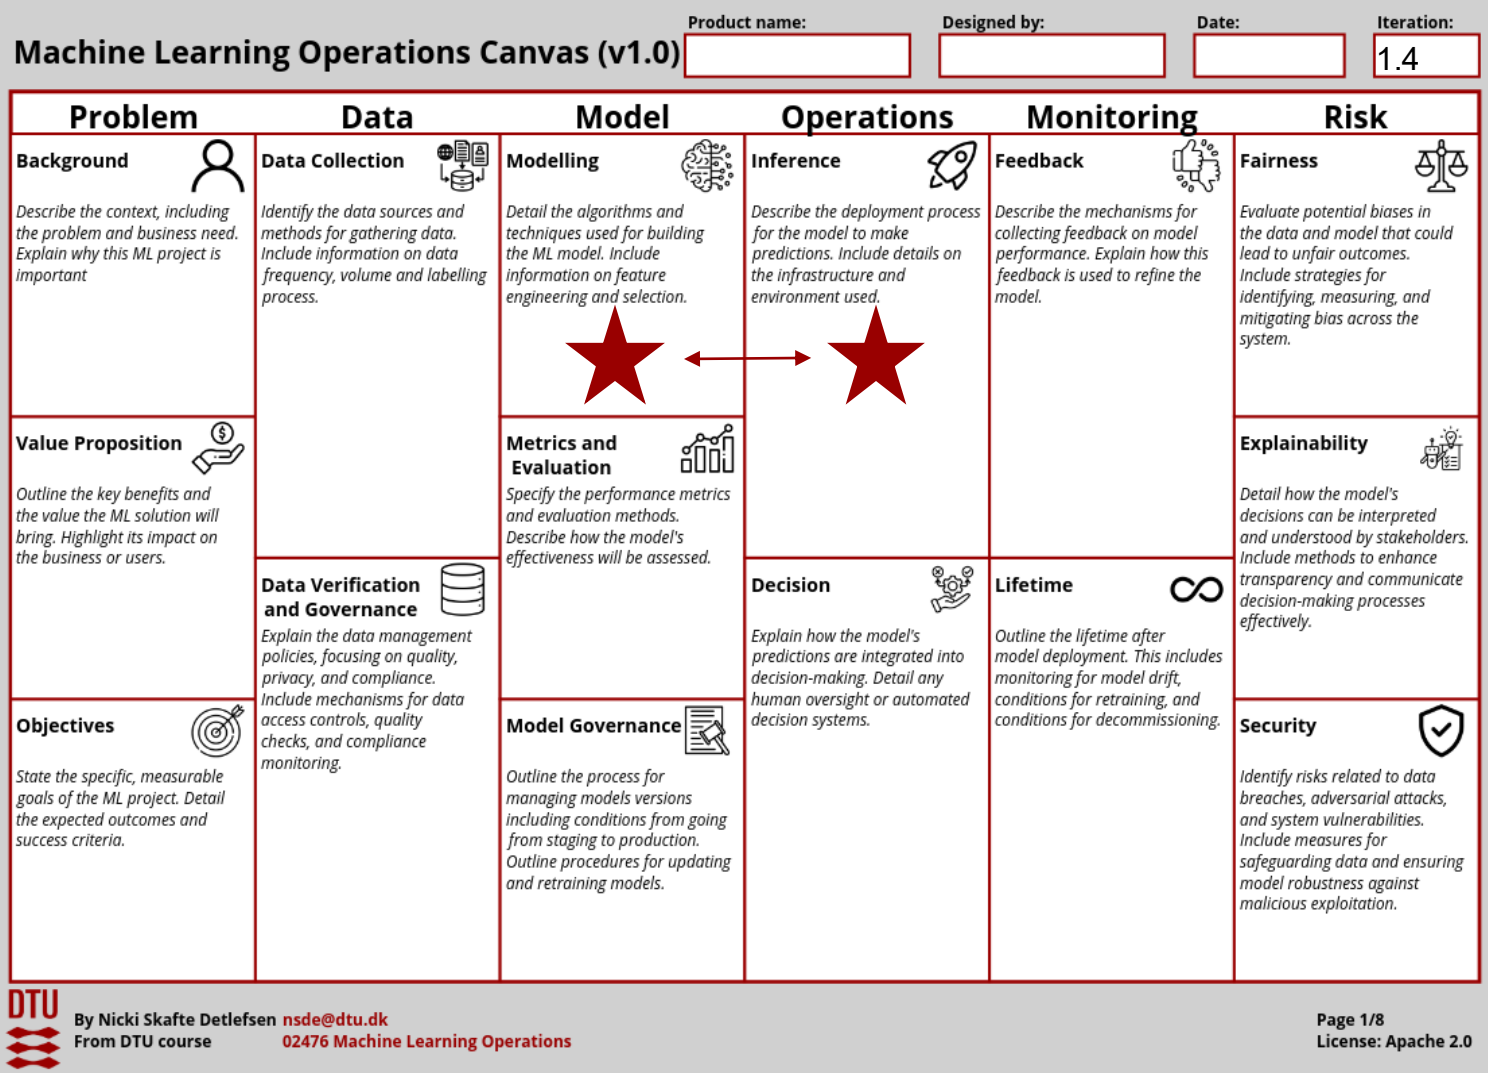

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration: 1.4

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*		**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*			**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Inference of model

1. Deployment of model to company cloud provider -> average inference time of 0.8sec but with some above 1sec
2. Optimization of model architechture for faster inference
3. Deployment of new model to cloud -> average inference of 0.6sec, with no spikes above 1sec
4. Deployment of new model to local laptop connect to visual setup -> all inference below 0.4 sec

What did we do?

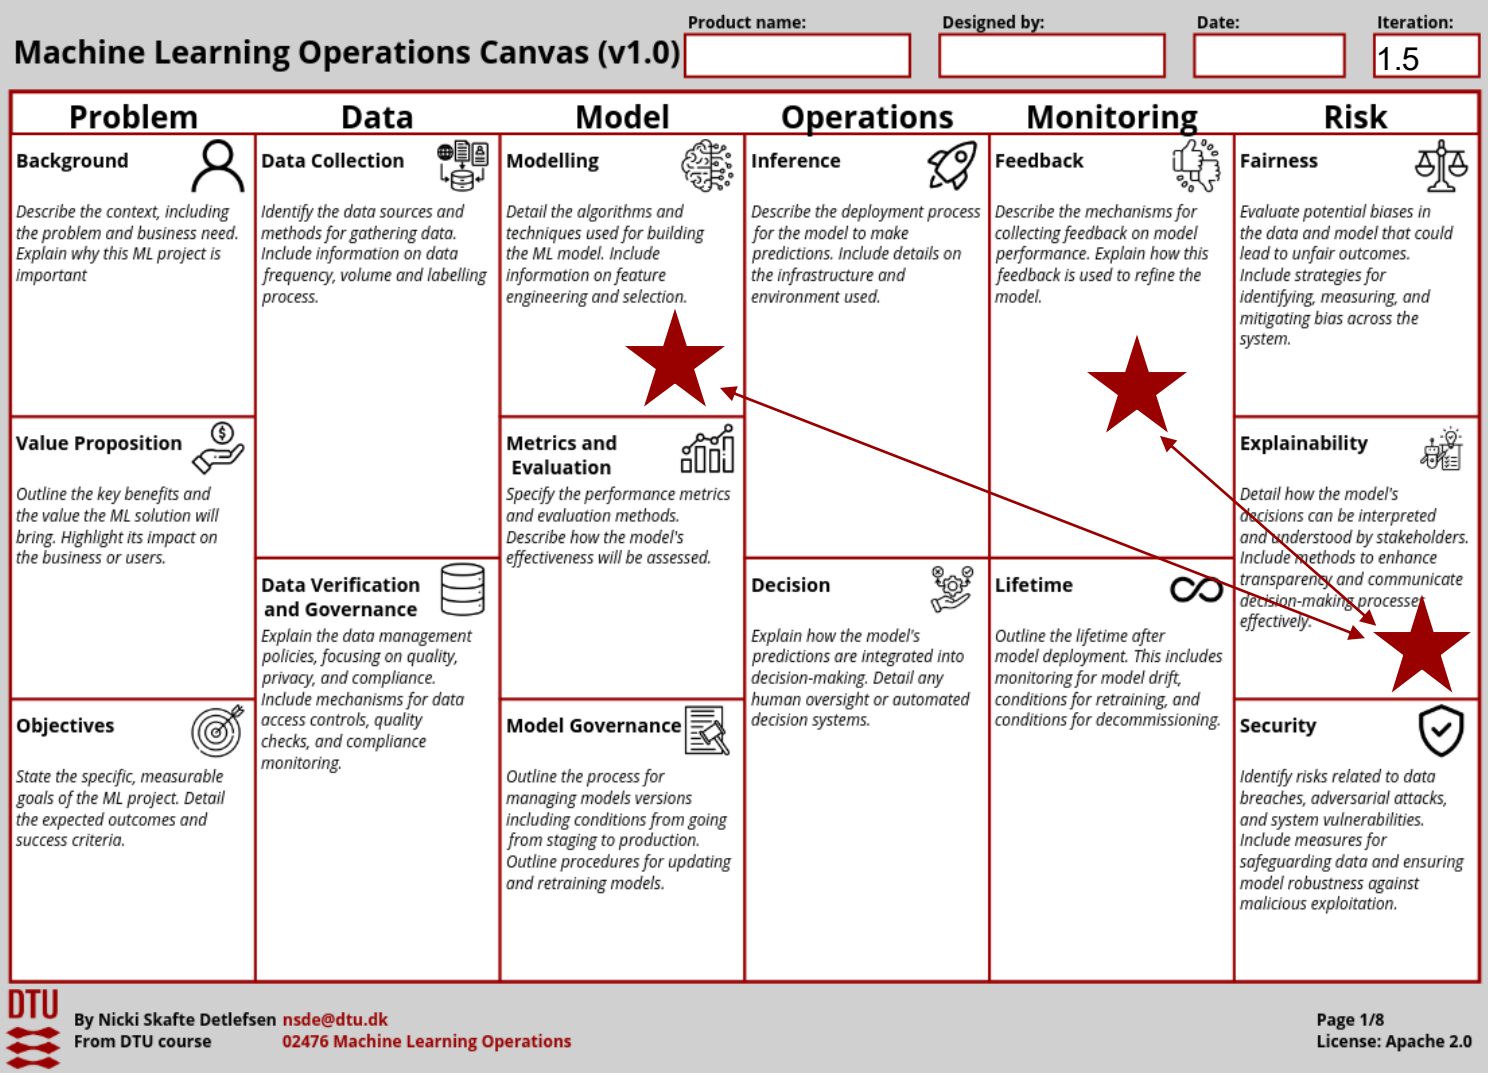

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration: 1.5

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Page 1/8
License: Apache 2.0

Because medical companies needs to comply with strict regulation

1. Small adjustment to model architecture to make it fits with explainability frameworks
2. Human-in-the-loop feedback procedure established that evaluates predictions, data setup and explainability on regular basis

CANVAS time!

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

Summary

Take aways

Just like machine learning models:

Garbage in, garbage out 🗑️

💡 If the fundamental problems and requirements are not defined early, any machine learning pipeline will accumulate technical debt.

💡 We can reduce the debt we have to pay by considering the entire pipeline.

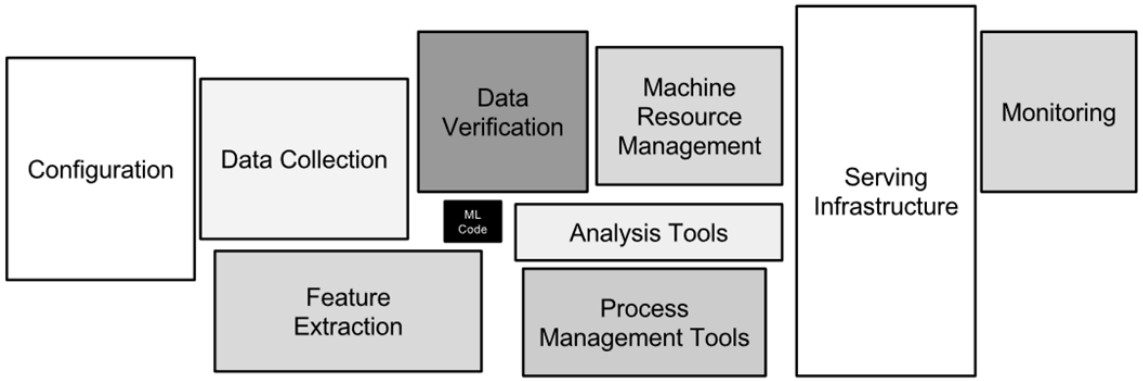

MLOps is really just...

...delivers value to the business 💸

...considers the entire pipeline, not just data and model 🤔

...takes long-term goals into account from the start 🥅

…that includes thinking about the ethical risks ⚠️

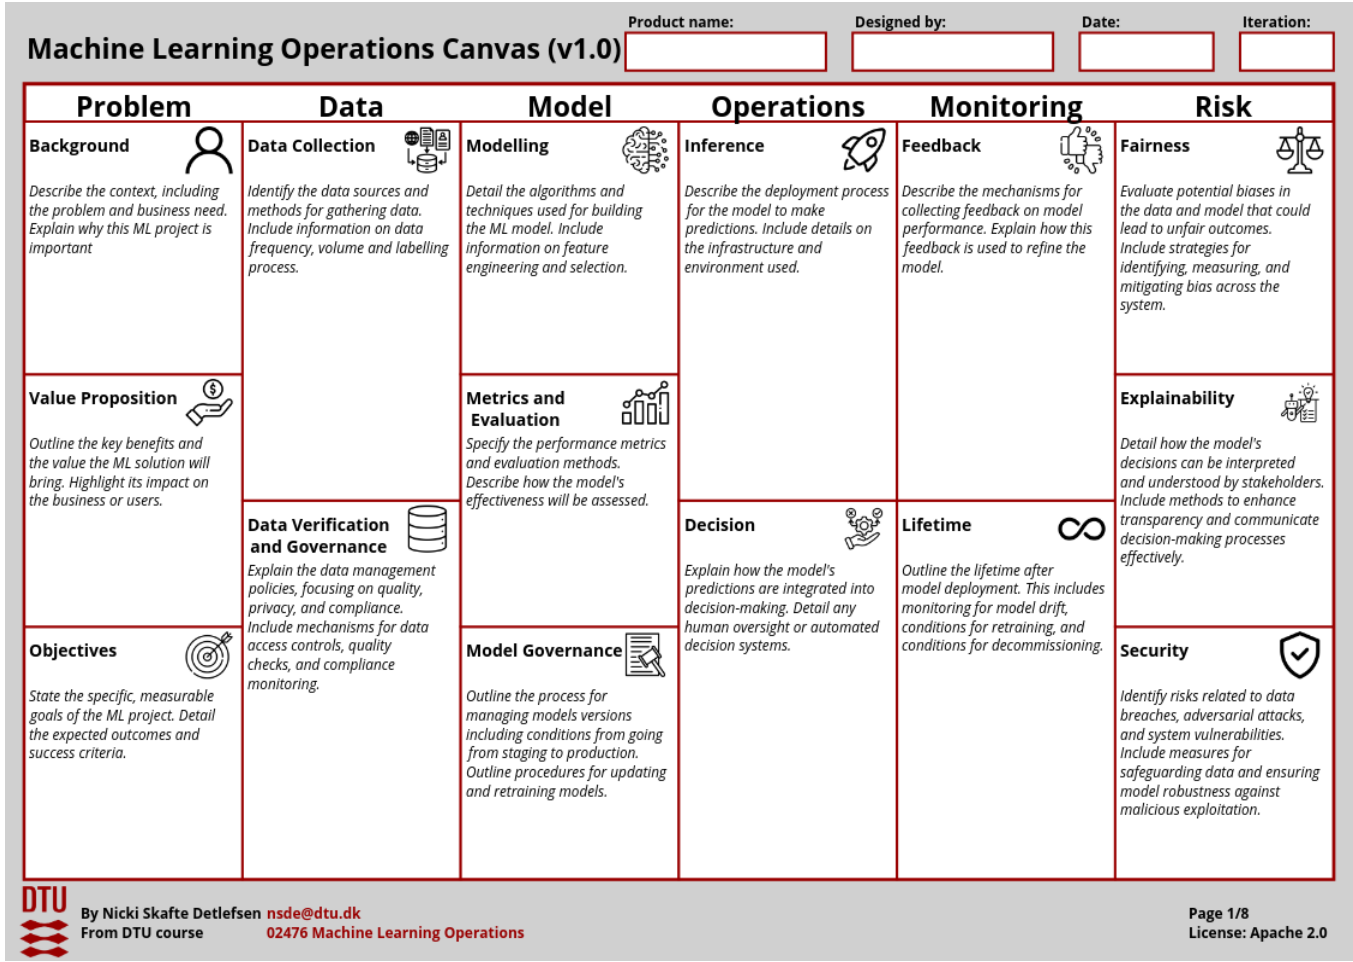

Machine Learning Operations Canvas (v1.0)

Product name: | Designed by: | Date: | Iteration:

Problem	Data	Model	Operations	Monitoring	Risk
Background *Describe the context, including the problem and business need. Explain why this ML project is important*	**Data Collection** *Identify the data sources and methods for gathering data. Include information on data frequency, volume and labelling process.*	**Modelling** *Detail the algorithms and techniques used for building the ML model. Include information on feature engineering and selection.*	**Inference** *Describe the deployment process for the model to make predictions. Include details on the infrastructure and environment used.*	**Feedback** *Describe the mechanisms for collecting feedback on model performance. Explain how this feedback is used to refine the model.*	**Fairness** *Evaluate potential biases in the data and model that could lead to unfair outcomes. Include strategies for identifying, measuring, and mitigating bias across the system.*
Value Proposition *Outline the key benefits and the value the ML solution will bring. Highlight its impact on the business or users.*		**Metrics and Evaluation** *Specify the performance metrics and evaluation methods. Describe how the model's effectiveness will be assessed.*			**Explainability** *Detail how the model's decisions can be interpreted and understood by stakeholders. Include methods to enhance transparency and communicate decision-making processes effectively.*
Objectives *State the specific, measurable goals of the ML project. Detail the expected outcomes and success criteria.*	**Data Verification and Governance** *Explain the data management policies, focusing on quality, privacy, and compliance. Include mechanisms for data access controls, quality checks, and compliance monitoring.*	**Model Governance** *Outline the process for managing models versions including conditions from going from staging to production. Outline procedures for updating and retraining models.*	**Decision** *Explain how the model's predictions are integrated into decision-making. Detail any human oversight or automated decision systems.*	**Lifetime** *Outline the lifetime after model deployment. This includes monitoring for model drift, conditions for retraining, and conditions for decommissioning.*	**Security** *Identify risks related to data breaches, adversarial attacks, and system vulnerabilities. Include measures for safeguarding data and ensuring model robustness against malicious exploitation.*

By Nicki Skafte Detlefsen nsde@dtu.dk
From DTU course 02476 Machine Learning Operations

A canvas is "just" a template

- Maybe the ML canvas I presented doesn't fit your company.
- That's completely fine—create a version that suits you better.

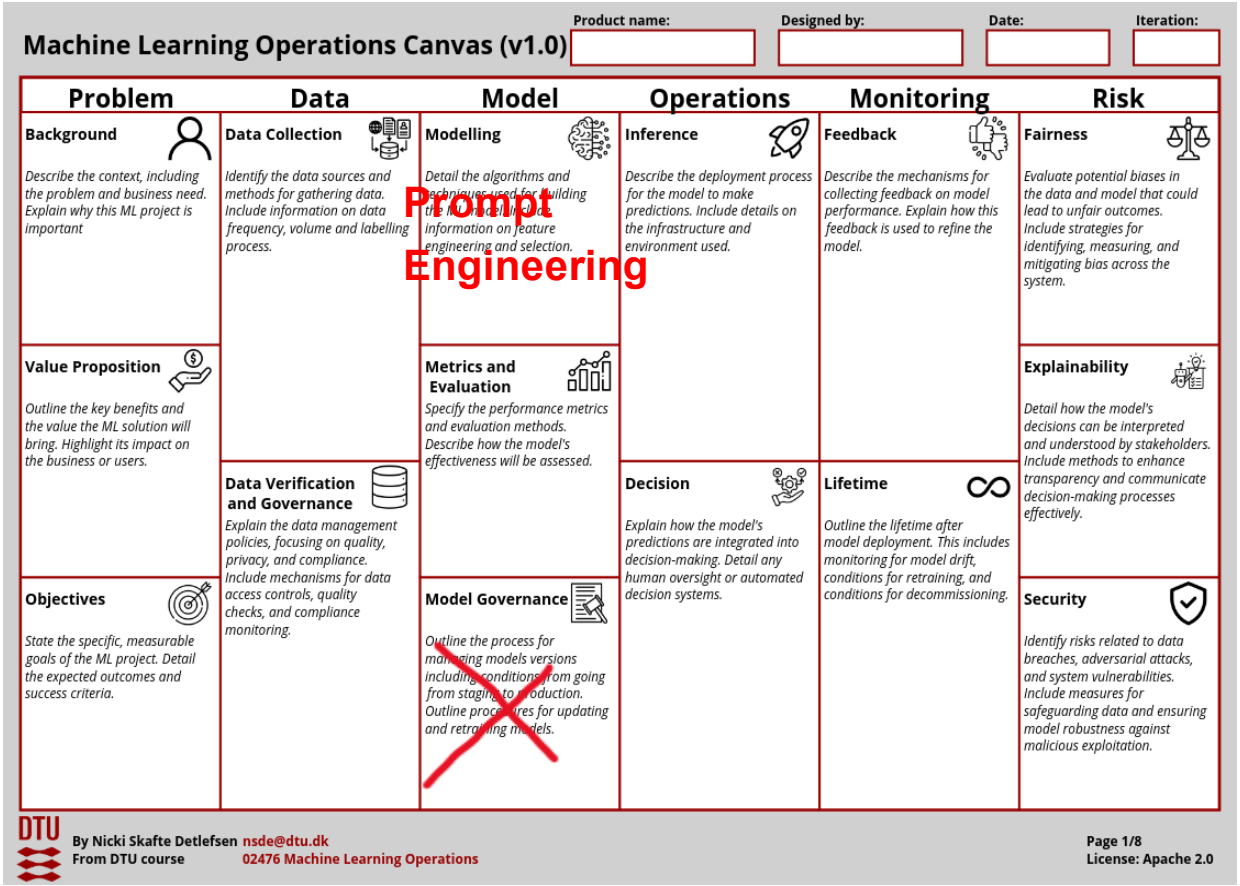

Links

💡 3 weeks MLOps course: https://skaftenicki.github.io/dtu_mlops/

💡 Practical MLOps pipeline example: https://github.com/SkafteNicki/example_mlops

Thanks for your attention 🤗

Feel free to reach out to me if you want to collaborate

nsde@dtu.dk

AI-Generated