



GÖTEBORGS
UNIVERSITET



CHALMERS

DAT 550 / DIT 978

Advanced Software Engineering for AI/ML- Enabled Systems

Your teachers:

Hans-Martin Heyn, Universitetslektor,

Eric Knauss, Docent

Computer Science and Engineering Department, Göteborg University

What will you learn?

Advanced Software Engineering for AI/ML-Enabled Systems

- Processes and engineering practices for developing AI/ML-enabled systems, from requirements engineering to testing;
- Typical roles in software engineering of AI/ML enabled systems;
- Architectures and patterns for AI/ML-enabled systems;
- Existing techniques to verify and explain decisions made by AI/ML-enabled systems;
- Overview of recent research on SE for AI/ML-enabled systems.
- Find and read relevant research papers on software engineering for AI/ML-enabled systems;
- Present and demonstrate a software engineering approach for AI/ML-enabled systems;
- Assess new engineering knowledge for AI/ML-enabled system.

*What are your expectations?
What do you expect to learn?*

Why do you learn that?

Advanced Software Engineering for AI/ML-Enabled Systems

- Judge the extent to which an AI/ML component needs to be safe-guarded;
- Judge what verification methods are appropriate when developing an AI/ML-enabled system given the requirements of that system;
- Judge whether a model has systematic biases and discuss the consequences of these biases;
- Judge fairness and potential other ethical issues of an AI/ML-enabled system;
- Judge user's information needs to work with an AI-enabled system;
- Judge limitations of a state-of-the-art software engineering approach for AI/ML given evidence presented in research paper.

How important is it to learn that?

Your teachers

YOU!!!



Eric Knauss
Docent
Software Engineering
CSE



Hans-Martin Heyn
Senior Lecturer
Software Engineering
CSE

This is a seminar course

You are not only the students but also the teachers of this course

- No frontal class-room teaching...
- ...but discussions and interaction instead.
- We are learning new material together.
- Discussions are highly encouraged.
- You determine what you want to study in depth.
- But you will learn about a broad spectrum of topics through the presentations of the other students.



A round of introductions

As we are working together, we should introduce each other


- Your (preferred) name.
- Your background with AI / ML? Any goals or special interests?
- Your background with software engineering? Any goals or special interests?
- One topic you especially hope to learn more about in this class.
- Your favourite ChatGPT question or command.

What will you learn today?

Advanced Software Engineering for AI/ML-Enabled Systems

Introduction to SE for AI

The themes and topics of
this course



Administrative


What will you learn today?

Advanced Software Engineering for AI/ML-Enabled Systems

Introduction to SE for AI

The themes and topics of
this course

Administrative



What makes software engineering for ML challenging?



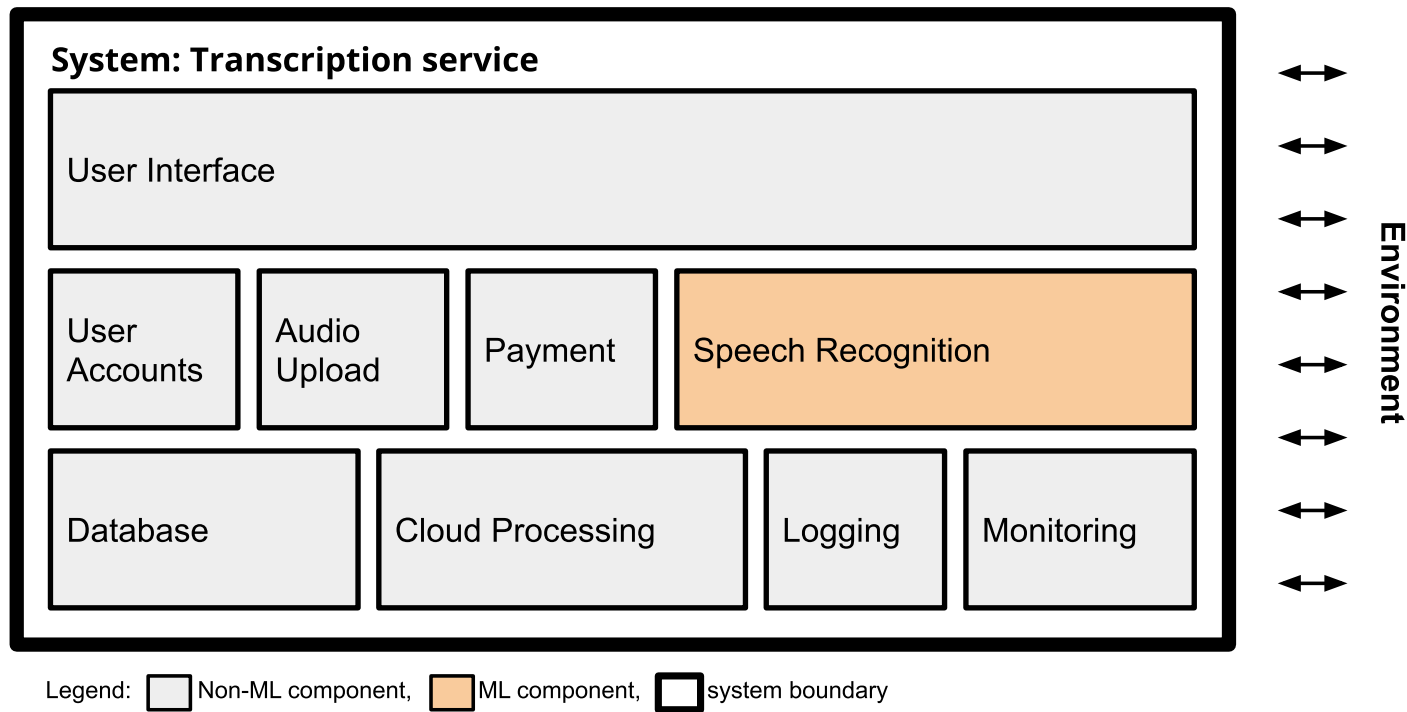
NeuralTalk2: A flock of birds flying in the air

Microsoft Azure: A group of giraffe standing next to a tree

Image: Fred Dunn, <https://www.flickr.com/photos/gratapictures> - CC-BY-NC

- It is often hard to even describe what goes wrong in software with ML.
- One root cause problem is probably the lack of suitable specifications:
 - How do you define the desired context of a ML model?
 - How do you specify the necessary data?
- What is a correct output of a ML model? Is 99% correctness acceptable? 99.2%? 99.8%?
- How do we judge the consequences of a wrong decision?

What makes software engineering for ML challenging?



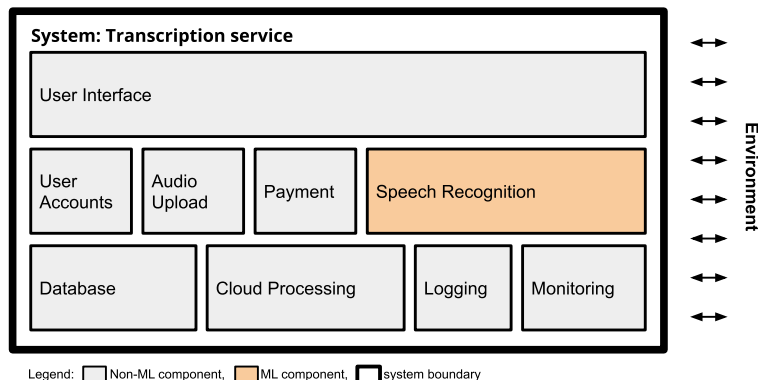
© Eunsuk Kang, Chrisitan Kästner, Machine Learning in Production

Not everything is a new problem

Many problems exist also with conventional software systems

- Building safe software with unreliable components;
 - Systems interacting with the environment (aka Cyberphysical systems);
 - Big data systems, cloud systems;
 - "Good enough" and "fit for purpose" != correct.
-
- However, ML is getting widely used now, so the problems are being put into the spotlight of attention.

Can we specify ML models?



© Eúnsuk Kang, Chrisitan Kästner, Machine Learning in Production

```
/**
 * Return the text spoken within the audio file
 * ???
 */
String transcribe(File audioFile);
```

- With ML, it is hard to derive clear specification. Why?
- Well, we use ML precisely because we do not know how to specify and implement a specific task.
- For example speech recognition:
 - We want a function that converts spoken language into machine readable text.
- But can we be more specific?
- How do we derive test cases from vague specifications?

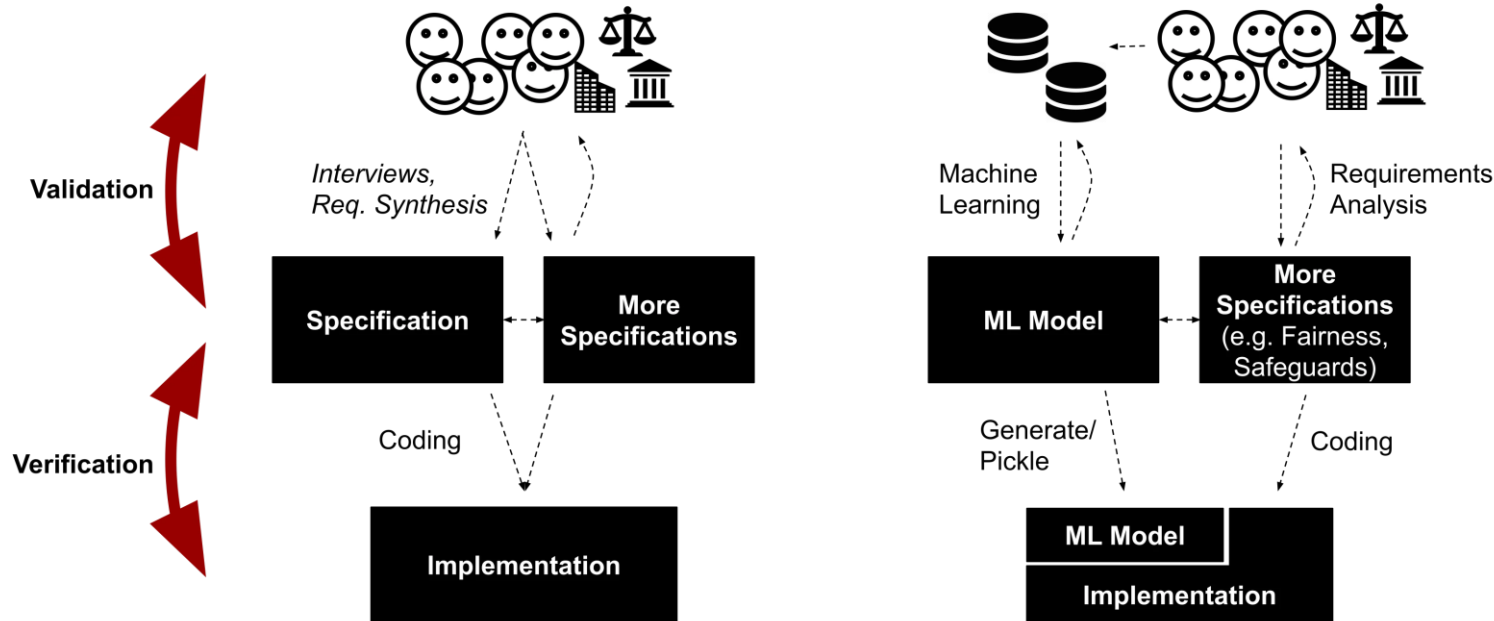
ML models and specifications

Excerpt from: Machine Learning is Requirements Engineering — On the Role of Bugs, Verification, and Validation in Machine Learning, Christian Kästner

Talking about specifications in machine learning is difficult. We have data, that somehow describes what we want, but also not really: We want *some sort of generalization* of the data, but also don't want a precise generalization of the training data — it's quite okay if the model makes wrong predictions on some of the training data if that means it generalizes better. Maybe, we can talk about some implicit specification derived from some higher system goals (e.g., thou shall best predict the stock market development, thou shall predict what my customer wants to buy), but its unclear where such specification comes from or how it could be articulated. *This vague notion of specification is confusing (at least to me) and makes it very hard to pin down what “testing”, “debugging”, or “bug” mean in this context.*

Identify relevant and representative data = identify representative stakeholders

How large should a training dataset be Is it better to select data, or just “throw all data on the problem”?



Would you hire a software engineer or data scientist?

Software Engineer


- Builds a product
- Concerned about cost, performance, stability, release time
- Identify quality through customer satisfaction
- Must scale solution, handle large amounts of data
- Detect and handle mistakes, preferably automatically
- Maintain, evolve, and extend the product over long periods
- Consider requirements for security, safety, fairness

Data Scientist

- Often fixed dataset for training and evaluation (e.g., PBS interviews)
- Focused on accuracy
- Prototyping, often Jupyter notebooks or similar
- Expert in modelling techniques and feature engineering
- Model size, updateability, or implementation stability typically does not matter

What is a Software Engineering for AI?

How do we find the unicorn?




- What knowledge is needed for Software Engineering?
- What additional knowledge is needed for AI?
- Can one person do the job? Would you rather hire a Software Engineer or a Data Scientist? Why?



"I-shaped"
Expert at one thing



Generalist
Capable in a lot of things
but not expert in any



"T-shaped"
Capable in a lot of things
and expert in one of them

Introduction to SE for AI

Fundamentals of Engineering AI-Enabled Systems

Holistic system view: AI and non-AI components, pipelines, stakeholders, environment interactions, feedback loops

Requirements:

System and model goals
User requirements
Environment assumptions
Quality beyond accuracy
Measurement
Risk analysis
Planning for mistakes

Architecture + design:

Modeling tradeoffs
Deployment architecture
Data science pipelines
Telemetry, monitoring
Anticipating evolution
Big data processing
Human-AI design

Quality assurance:

Model testing
Data quality
QA automation
Testing in production
Infrastructure quality
Debugging

Operations:

Continuous deployment
Contin. experimentation
Configuration mgmt.
Monitoring
Versioning
Big data
DevOps, MLOps

Teams and process: Data science vs software eng. workflows, interdisciplinary teams, collaboration points, technical debt

Responsible AI Engineering

Provenance,
versioning,
reproducibility

Safety

Security and
privacy

Fairness

Interpretability
and explainability

Transparency
and trust

Ethics, governance, regulation, compliance, organizational culture

Want to read more?

Machine Learning is Requirements Engineering — On the Role of Bugs, Verification, and Validation in Machine Learning


<https://medium.com/analytics-vidhya/machine-learning-is-requirements-engineering-8957aee55ef4>

Building Intelligent Systems – A Guide to Machine Learning Engineering by Geoff Hulten, Apress / Springer, **freely available through Chalmers's library**

<https://link.springer.com/book/10.1007/978-1-4842-3432-7>

- Chapter 1: Introducing Intelligent Systems (And the Internet Toaster)

Taking a break



Introduction to SE for AI

Fundamentals of Engineering AI-Enabled Systems

Holistic system view: AI and non-AI components, pipelines, stakeholders, environment interactions, feedback loops

Requirements:

System and model goals
User requirements
Environment assumptions
Quality beyond accuracy
Measurement
Risk analysis
Planning for mistakes

Architecture + design:

Modeling tradeoffs
Deployment architecture
Data science pipelines
Telemetry, monitoring
Anticipating evolution
Big data processing
Human-AI design

Quality assurance:

Model testing
Data quality
QA automation
Testing in production
Infrastructure quality
Debugging

Operations:

Continuous deployment
Contin. experimentation
Configuration mgmt.
Monitoring
Versioning
Big data
DevOps, MLOps

Teams and process: Data science vs software eng. workflows, interdisciplinary teams, collaboration points, technical debt

Responsible AI Engineering

Provenance,
versioning,
reproducibility

Safety

Security and
privacy

Fairness

Interpretability
and explainability

Transparency
and trust

Ethics, governance, regulation, compliance, organizational culture

This is a seminar course

- No frontal class-room teaching...
- ...but discussions and interaction instead.
- We are learning new material together.
- Discussions are highly encouraged.
- You determine what you want to study in depth.
- But you will learn about a broad spectrum of topics through the presentations of the other students.

The key to this course is that you will pick one specific topic in SE for AI and prepare a lecture, example, and introduction video about it.

You will also write a report on 2-3 other topics in SE for AI that interest you.

It is your choice!

Standing of the shoulder of giants

- You don't have to start from zero. We prepared a set of themes and topics, including starting literature for you.
- For each theme, we identified together with our colleagues several topics that you can choose to study.
 - Some of our colleagues also volunteered to help you finding suitable literature.
- Each topic contains a list of suggested starting literature:

Topic: Risk analysis for AI/ML

Starting Literature:

Hulten, G. (2018). Building intelligent systems: a guide to machine learning engineering. Apress. Chapter 24

[Towards Risk Modeling for Collaborative](#) ↓ [AI.pdf](#) ↓


[Dataset Fault Tree Analysis for Systematic Evaluation of Machine Learning Systems.pdf](#) ↓

What will you learn today?

Advanced Software Engineering for AI/ML-Enabled Systems

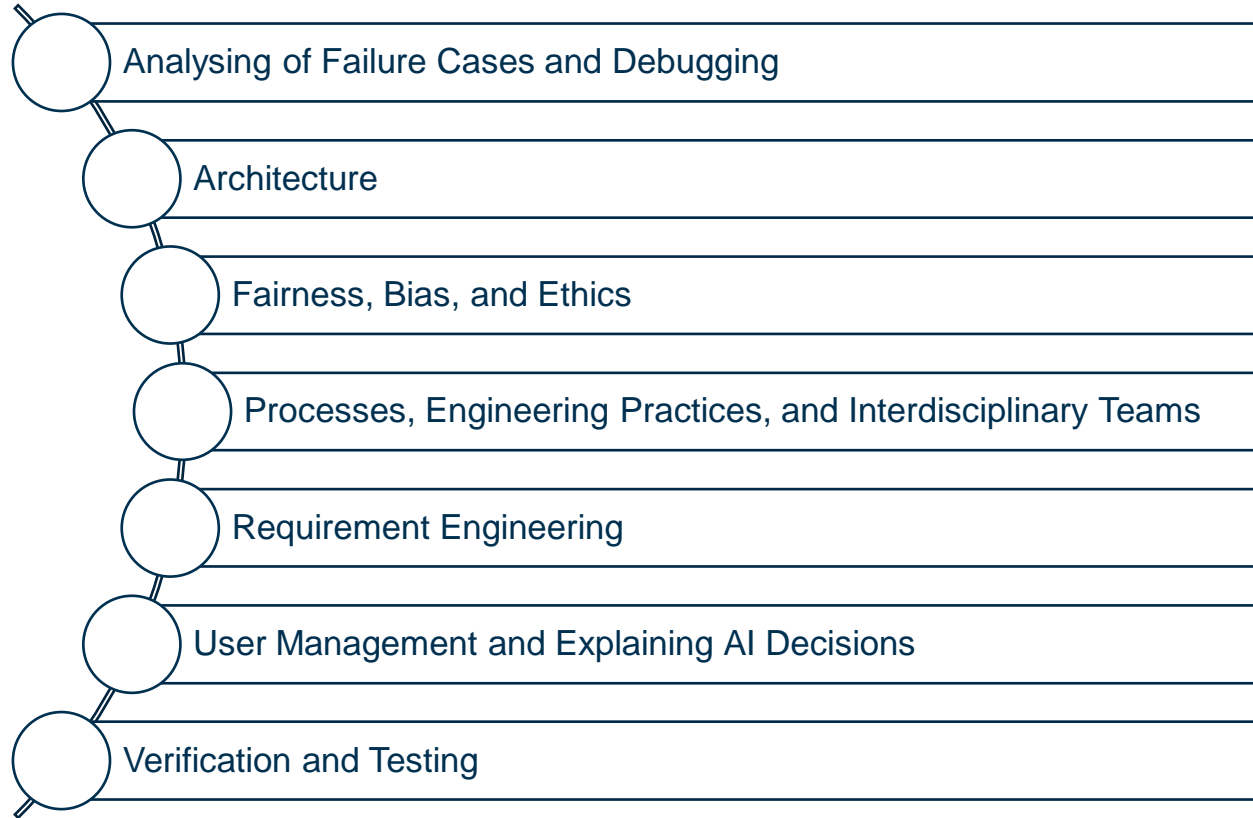
Introduction to SE for AI

The themes and topics of
this course



Administrative

Themes of the course



Analysing of Failure Cases and Debugging

Topics for study:

- Risk analysis for AI/ML
- Explanatory debugging
- Data Distribution Debugging
- Model misprediction diagnosis

Architecture

Topics for study:

- Software Architecture for ML
- Safe-guarding strategies and security
- Technical Dept. Code Smells, Patterns, and Anti-Patterns

Fairness, Bias, and Ethics

Topics for study:

- Accountability, Responsibility, and Transparency
- Fairness strategies and Fair-SMOTE
- Bias Mitigation Methods (Fairea)
- Fairness metrics and testing

Processes, Engineering Practices, and Interdisciplinary Teams

Topics for study:

- Challenges when Engineering AI/ML systems in industry / practice
- Process & Agile for AI/ML
- Engineering Practices
- MLOps
- Interaction and collaboration of roles

Requirement Engineering

Topics for study:

- Data Quality Requirements and Data Smells
- Functional and non-functional requirements specifications for ML

User Management and Explaining AI Decisions

Topics for study:

- Expectations / Expectation Setting Methods
- Tools for Interpretability A: Overview
- Tools for Interpretability B: Tools
- Design Practices of Explainable AI

Verification and Testing

Topics for study:


- Quality vs. Confidence and typical bug types and sources

What will you learn today?

Advanced Software Engineering for AI/ML-Enabled Systems

Introduction to SE for AI

The themes and topics of
this course




Administrative

Administrative

All information can be found on Canvas



☰ DAT550 / DIT978 > Syllabus



lp4 VT23

Home

Syllabus

Modules

People

Grades

DAT550 / DIT978 Advanced Software Engineering for AI/ML-Enabled Systems

Jump to
today

Course-PM: Welcome to DAT550 / DIT978 Advanced Software Engineering for AI/ML-Enabled Systems lp4 VT23 (7.5 hp)

Artificial intelligence and machine learning are more and more used in practice. However, the introduction of AI/ML components into a software system comes with new challenges and needs and changes the way the software system is engineered. This course introduces processes, practices and techniques for engineering AI/ML-enabled software systems.

Contact details

Examiner: Eric Knauss ([mail](#))

Lecturers: Hans-Martin Heyn ([mail](#)), Eric Knauss ([mail](#))

Organisation

This course is offered to students from the University of Gothenburg and from Chalmers. The course homepage is located in Canvas.

The course is provided in the form of a literature seminar, which combines reading papers, student presentations, and discussions. Students will explore one

Course Design

| Week Number | Tuesday (08:15 - 10:00) | Friday (08:15 - 11:45) | Task Due Dates |
|---------------------------------------|------------------------------------------------------------------------------|------------------------------------------------------|--------------------------------------------------------------|
| 12 (21.03 / 24.03) | Introduction, General Overview of Themes (Link to Slides ↓) | How to read & research scientific papers. | Fr. 24.03 17:00 Task A: Topic Selection |
| 13 (28.03 / 31.03) | Example of system architecture for ML enabled systems | Responsible AI Engineering | |
| 14 | Self-Studies: Preparation of Presentations / Eastern | | |
| 15 (14.04) | Self-Studies: Preparation of Presentations / Eastern | Guest Lecture: AI Engineering in Volvo Trucks | |
| 16 (18.04 / 21.04) | Self-Studies: Preparation of Presentations | | |
| 17 (25.04 / 28.04) | Student Presentations | Student Presentations | |
| 18 (02.05 / 05.05) | Student Presentations | Student Presentations | |
| 19 (09.05 / 12.05) | Student Presentations | TBD | |
| 20 (16.05 / Wednesday, 17.05) | Student Presentations | OBS! Wednesday 17.05 Student Presentations | Fr. 19.05 17:00 Task D: Teach-me video |
| 21 (23.05 / 26.05) | Student Presentations | Teach-me video presentations and summary | |
| 22 Exam Week | Exam Week, no lectures | | Fr. 02.06 17:00 Task F: Individual Report |

The course has four phases:








- Introduction lectures (Week 12&13)
- Preparation phase (Week 14-16)
- **Presentation phase (Week 17-21)**
- Individual report (Week 22)

Examination

| Task | Name | Description | Team | Individual | Due Date |
|------|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------------|--------------------------------------|
| A | Topic Selection | Topic Selection and Assignment: Students should form teams of 2 and submit their topic preferences | X | | Fri. 24. March 2023 17:00 |
| B1 | Presentation and Pedagogic example, exercise, or tutorial | The presentation will happen on one of the following days. The order of presentations will be determined by topic. Presentation days are: Tu. April 25, Fr. April 28, Tu. May 2, Fr. May 5, Tu. May 9, Fr. May 12, Tu. May 16, and Tu. May 23. The exact times will be shown in Canvas. The Pedagogic example, exercise, or tutorial is to be delivered as part of the presentation. | X | | Individual |
| B2 | Submission of material | Presentation slides and materials from the pedagogic example, exercise, or tutorial must be submitted one week after the presentation via Canvas. | X | | Individual |
| C | Leading a discussion | Leading a discussion will happen directly after the corresponding presentation. | | X | Individual |
| D | Teach-me video | Should be submitted on Monday of the week after the presentation of the team. | X | | Fri. 19. May 2023, 17:00 |
| E | Attendance | Attendance to all student presentations (see dates above) is mandatory. A student may miss one of the dates. | | X | All lectures |
| F | Report | A final report | | X | Fri. 02. June 2023, 17:00 |

How do I pass?

- All submissions happen in Canvas

| | |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ⋮ | ▼ Team Assignments |
| ⋮ |  Task A: Topic Selection Due 24 Mar at 17:00 |
| ⋮ |  Task B1: Presentation and Pedagogic example, exercise, or tutorial |
| ⋮ |  Task B2: Submission of presentation material |
| ⋮ |  Task D: Teach-me video Due 19 May at 17:00 |
| ⋮ | ▼ Individual Assignments |
| ⋮ |  Task C: Leading a discussion |
| ⋮ |  Task E: Attendance |
| ⋮ |  Task F: Final Report Due 2 Jun at 17:00 |

Task A: Topic Selection (and team building)

Finding a team partner

- Each team consist of two students.
- You need to find a team partner (we are 40 students, so it should work out). If not, we can have one team with three student.
- If you want me to assign you randomly into a team with another student, please send me an e-mail today to heyn@chalmers.se .

Selecting a topic

- Under Assignment Task A: Topic Selection you find a link to a Microsoft Form survey.
- Here you can specify 3 most desired and 2 least desired topics out of the list of topics we just saw.
- **We will try our best to allocate one of the three choices to you** for which you will prepare a presentation, pedagogic example and teach-me video.

The topic selection needs to be done by Friday, 24th March, 17:00.

Task B1: Presentation and Pedagogic example

Each presentation shall be...

- ... between 30-40 minutes long;
- ... motivate the topic you selected (Why is it important? Why are you interested in it?);
- ... show an overview of relevant papers on the topic (min. 5) (What is happening in current research? Why?);
- ... identify interesting research gaps or topics that you think are interesting for further studies (Why?);
- ... provide one pedagogical example (see below);

Pedagogical examples could be:

- Mentimeters / Kahoots
- Interactive Jupyter Notebooks that the audience / student can try out
- Exercises (e.g., calculations on the blackboard)
- Tool / Code introductions / walk-throughs
- ...

The presentation will happen at one of the lectures in week 17-21.

After each presentation you will engage in a discussion with the discussion leaders and the other students for about 10-15 minutes.

Task B2: Presentation Materials

You need to upload all material from your lecture and example in one .zip file in Canvas under Assignment Task B2.

Task B2: Submission of presentation material

 Published  Edit 

Here you upload your team's presentation slides and material, including the pedagogic example, exercise (incl. solution) or tutorial.

You should have uploaded all materials latest one week after your presentation.

Please collect all material in one .zip file which you upload here.

Points 0

Submitting a file upload

File types zip

| Due | For | Available from | Until |
|-----|----------|----------------|-------|
| - | Everyone | - | - |

+ Rubric

Task C: Leading a discussion

The idea of the course is also to trigger discussions among us.

- Each student will act as a discussion leader for another presentation together with another student from another team.
- The assignment will be random.
- The discussion will take about 10-15 minutes.

As a discussion leader you, together with your co-leader, shall...

- ... prepare 1-2 questions during the presentations (or before, if you are familiar with the topic / material);
- ... engage other students in joining the discussion;
- ... moderate the discussion, ensure that everyone gets an equal chance to participate;
- ... provide constructive feedback to the presentation / lecture given.

Task D: Teach-me video

The videos will be uploaded to Canvas and available to everyone in the course

The video shall...

- ...be between 7 - 10 minutes;
- ...both team members shall present sections of the video;
- ...provide a motivation of the topic you selected (1-2 minutes);
- ...recap one specific takeaway from your studies in an engaging way (5-6 minutes);
- ...answer: if you were to propose a master thesis on the presented research topic, what could be a possible research question or research focus (1-2 minutes);

The videos are due on Friday, 19th May at 17:00.

Task E: Attendance

Having discussions only works if we have people to discuss with!

- You should attend at least(!!) 7 out of 10 presentation days.
- If you are sick or otherwise cannot join, please inform me before(!) the lecture by e-mail at heyn@chalmers.se
- We trust you, so we do not use roll-call lists or similar. But we need you to be honest and excuse yourself before the lecture if you cannot join.
- There is a mechanism to compensate with an extra brief written report if you are sick / unavailable more than 3 times, see Canvas Assignment Task E.

Task F: Final Report

To go beyond the chosen topic.

The report shall...

- ... use the IEEE conference template (it's on Canvas);
- ... not be longer than 2 pages (without references);
- ... reflect on two of the student lectures you thought were most interesting (why did you find them most interesting? what did interest you most? did you miss something in the presented lectures?);
- ... propose and motivate one additional topics not mentioned in the course OR argue for the importance / significance of a topic that is already part of the course and you think should remain in the course. Try to support your argumentation with references to literature.

The report is due on Friday, 02nd June at 17:00.

Grading Requirements

3 The student gives a presentation that represents the theme correctly and to a reasonable degree. The presentation includes a correct pedagogical example, exercise, or tutorial. The student acts as a discussion leader and submits a sound individual written report that addresses the main aspects asked for in the report.

4 The student gives a presentation that represents the theme correctly and to a large degree. The presentation includes a correct and illustrative pedagogical example, exercise, or tutorial. The student acts as a discussion leader, asking questions that show a good understanding of the presented topic. The student submits a teach-me video that is correct in presented content and of sufficient quality so that it can be shown to the class. The student submits a sound individual written report that addresses the main aspects asked for in the report and shows the ability to critically reflect on the discussed topic.

5 The student gives a presentation that represents the theme correctly, to a large degree, and shows a depth of knowledge on the topic. The presentation includes a correct and illustrative pedagogical example, exercise, or tutorial that encourages participants to reflect critically on the topic and its limitations. The student acts as a discussion leader, asking questions that show a good understanding of the presented topic and critical examination of the presented information. The student submits a teach-me video that is correct in presented content, of sufficient quality so that it can be shown to the class and illustrates the information in an engaging way. The student submits a sound individual written report that addresses all aspects asked for in the report and shows the ability to critically reflect on the discussed topic.

What if I miss a deadline?

Missed Deadlines and Revisions

Handling of missed deadlines and revisions depends on the grading component.

- Presentation: If a student fails to show up for their own presentation, a substitute presentation can be given on the 17th of May, 2023. Re-submissions of the presentation slides or materials can be done on the following dates: August 21st, 2023, October 23rd, 2023, and January 15th, 2024
- Pedagogic example, exercise, or tutorial: If a student fails to present a pedagogic example, exercise, or tutorial, as part of their presentation, materials for this task can be resubmitted on the following days: August 21st, 2023, October 23rd, 2023, and January 15th, 2024
- Teach-me video: The course can be passed without submission of a teach-me video. A re-submission is not possible.
- Discussion Leader: If a student fails to show up and act as a discussion leader for the assigned presentation, the student has one chance to compensate for that by submitting a 5-page report that discusses open questions and gives constructive feedback on the presentation slides, pedagogic example, exercise, or tutorial, and the final report of the group who gives that presentation. The deadline for this substitution task is August 21st, 2023.
- Individual written final report: The individual report can be resubmitted on three dates: August 21st, 2023, October 23rd, 2023, and January 15th, 2024
- Attendance: If a student fails to attend at least 8 of the 10 presentation days (including teach-me video presentations), the student has the chance to compensate for that by submitting a report including 1 page per lacking presentation day (i.e., if a student attended 8-X presentation days, they should pick X of the non-attended presentations). The report should discuss open questions and give constructive feedback on the presentation slides and pedagogic example, exercise, or tutorial, of the non-attended presentation. The deadline for this substitution task is August 21st, 2023.

Welcome to our course 😊

Fundamentals of Engineering AI-Enabled Systems

Holistic system view: AI and non-AI components, pipelines, stakeholders, environment interactions, feedback loops

Requirements:

System and model goals
User requirements
Environment assumptions
Quality beyond accuracy
Measurement
Risk analysis
Planning for mistakes

Architecture + design:

Modeling tradeoffs
Deployment architecture
Data science pipelines
Telemetry, monitoring
Anticipating evolution
Big data processing
Human-AI design

Quality assurance:

Model testing
Data quality
QA automation
Testing in production
Infrastructure quality
Debugging

Operations:

Continuous deployment
Contin. experimentation
Configuration mgmt.
Monitoring
Versioning
Big data
DevOps, MLOps

Teams and process: Data science vs software eng. workflows, interdisciplinary teams, collaboration points, technical debt

Responsible AI Engineering

Provenance,
versioning,
reproducibility

Safety

Security and
privacy

Fairness

Interpretability
and explainability

Transparency
and trust

Ethics, governance, regulation, compliance, organizational culture



GÖTEBORGS
UNIVERSITET



CHALMERS