



THE UNIVERSITY *of* EDINBURGH
informatics

Applied Machine Learning (AML)

Class Starting at 4:10pm

Oisin Mac Aodha • Siddharth N.

Applied Machine Learning

Week 1: Introduction

*This slides will be made available on the project website after the class.
This session will be recorded.*

Course Instructors



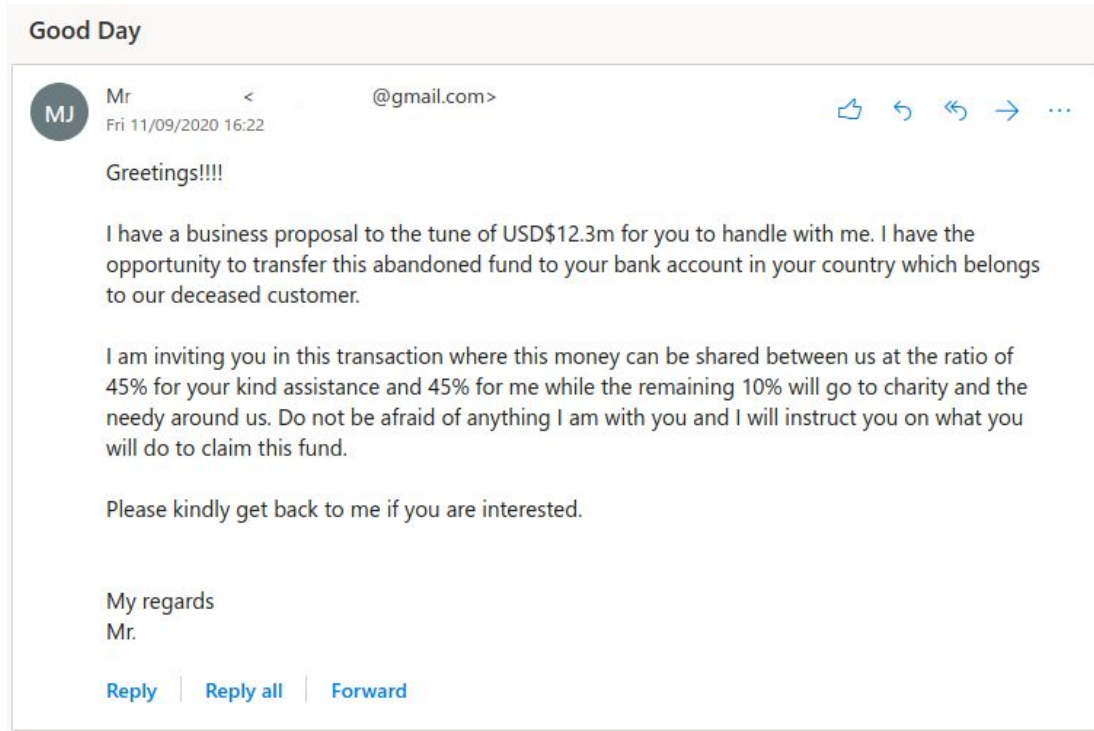
**Oisín
Mac Aodha**



**Siddharth
N.**

+ a big team including TA, lab demonstrators, and tutors helping out

How to Make a Spam Filter?



What is Machine Learning?

Machine Learning (ML) is the study and development of algorithms that learn from data in order to make predictions about new data.


What is Machine Learning?

Machine Learning (ML) is the study and development of algorithms that learn from data in order to make predictions about new data.

$$y = f(\mathbf{x}; \theta)$$

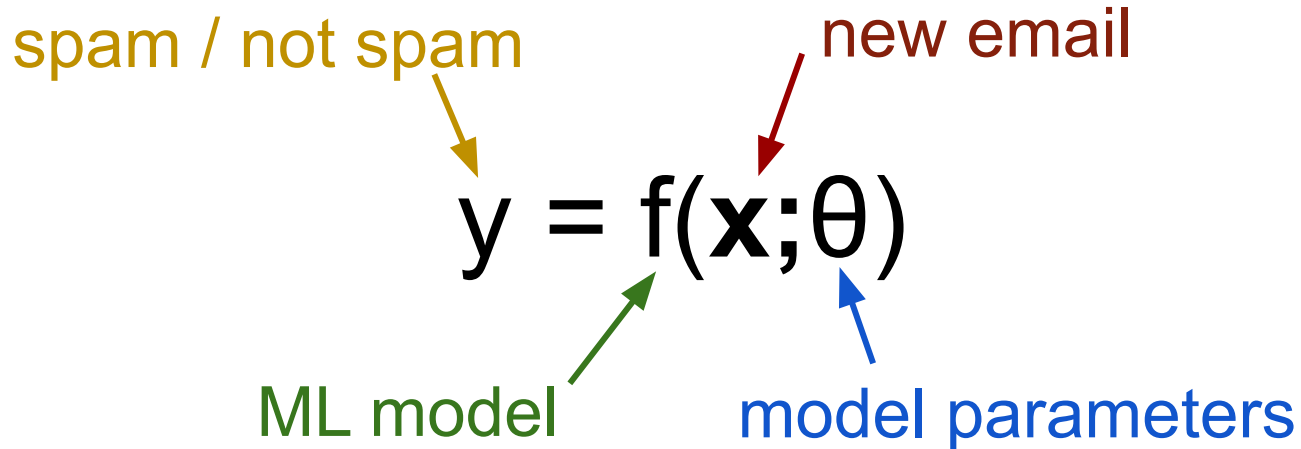
What is Machine Learning?

Machine Learning (ML) is the study and development of algorithms that learn from data in order to make predictions about new data.



What is Machine Learning?

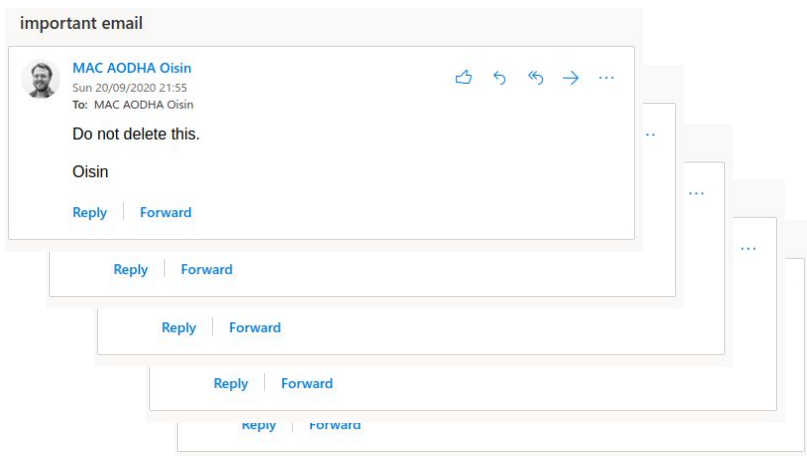
Machine Learning (ML) is the study and development of algorithms that learn from data in order to make predictions about new data.



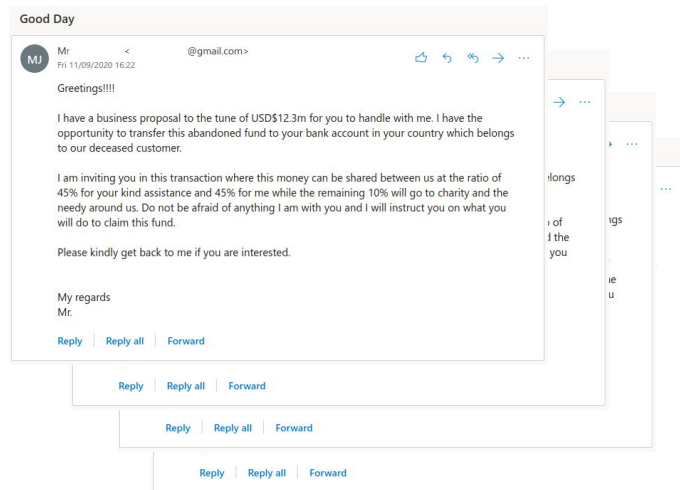
How to Learn the Model Parameters θ ?

We will use data to optimise for the model parameters that make the fewest mistakes on previously seen data.

Not spam emails



Spam emails



Machine Learning Examples



Whats is machine learning?



Machine learning is a branch of artificial intelligence (AI) that focuses on the development of algorithms and models that enable computers to learn and make predictions or decisions without being explicitly programmed to perform

Large Language Models e.g. ChatGPT

Machine Learning Examples



Whats is machine learning?



Machine learning is a branch of artificial intelligence (AI) that focuses on the development of algorithms and models that enable computers to learn and make predictions or decisions without being explicitly programmed to perform

Large Language Models e.g. ChatGPT



Generative Image Models e.g. StableDiffusion, DALL·E

Machine Learning Examples



Whats is machine learning?




Machine learning is a branch of artificial intelligence (AI) that focuses on the development of algorithms and models that enable computers to learn and make predictions or decisions without being explicitly programmed to perform

Large Language Models e.g. ChatGPT



Protein Structure Prediction e.g. AlphaFold



Generative Image Models e.g. StableDiffusion, DALL·E

Machine Learning Examples




Whats is machine learning?




Machine learning is a branch of artificial intelligence (AI) that focuses on the development of algorithms and models that enable computers to learn and make predictions or decisions without being explicitly programmed to perform

Large Language Models e.g. ChatGPT



Protein Structure Prediction e.g. AlphaFold



Generative Image Models e.g. StableDiffusion, DALL·E

+ Many more examples:

- Recommendation systems
- Finance
- Image recognition
- Medicine
- ...

Ethical Considerations

- What are the implications of AI-based decision making?
- How has an AI model come to a particular decision?
- Do AI models exhibit biases?
- How was the training data sourced?
- ...

Importance of Fundamentals

Big advances in LLMs, but important to know how they work and their limitations.



Melanie Mitchell
@MelMitchell1



"OpenAI Srawberry"?

ChatGPT o1-preview Share MM

"Strawberry" has how many t's?

Thought for a couple of seconds

The word "strawberry" does not contain the letter "t." Therefore, it has **zero** "t's."

3:52 AM · Sep 13, 2024 · 26.1K Views

Overview for Today

- 1) Discussion of how AML is organised
- 2) Outline your tasks this for week

Have you taken a machine
learning course before?

What is AML?

- Introduction to machine learning methods and concepts
- How these algorithms work
- Which algorithm to apply for a given task and how to evaluate it
- Practical experience of applying machine learning to real problems
- Foundation for learning about more advanced machine learning methods in the future

Flipped Classroom

- There are **no “live” lectures** for AML
- Instead, you watch the lecture videos and read the slides in your own time and can engage in discussions on our forum (Piazza)
- In the tutorials, labs, and Q&A sessions we will focus on applying the knowledge you learned in the lectures and get **feedback**



ChatGPT Generated Image

Course Delivery

- Pre-recorded lectures - two topics per week
 - Watch in **week i**, Q&A in **week i+1**
- Live Q&A session - one per week (Tuesday at 4:10pm)
- Practical labs - four in total + one optional introduction lab
- Tutorials - four in total, pen and paper exercises

Course Website - <https://tinyurl.com/aml2024>

🔒 <https://groups.inf.ed.ac.uk/teaching/aml/> ☆

Applied Machine Learning (AML)

🔍 Search Applied Machine Learning (AML) DRPS Piazza Learn Labs Tutorials

Home

- Course Information
- Schedule
- Mini-Project
- Exam
- Staff

Applied Machine Learning (MSc)

Informatics (INFR11211), Semester 1, 2024



Oisín Mac Aodha
oisin.macaodha@ed.ac.uk



Siddharth N.
n.siddharth@ed.ac.uk



Week 0 Announcement

Sep 10 · 0 min read

Welcome Week!

Announcements

*URL also available on **Learn***

Make sure to refresh the page

Math Requirements

Linear algebra: vectors and matrices, e.g. scalar (dot) product, matrix multiplication, matrix inversion, eigenvectors, and determinants, ...

Special functions: e.g. logarithms and exponentials

Calculus: Differentiation of standard functions

Geometry: Basics of lines, planes, and hyperplanes

Probability theory: Discrete and continuous random variables, univariate and multivariate Gaussians, joint and conditional distributions, ...

Programming Requirements

Should be familiar with programming in a modern object-oriented language, ideally **Python** which is the course language

Comfortable with the material in Lab 0 (more on this later)

Course Structure

Four main themes

1. Supervised Learning
2. Data Exploration and Evaluation
3. Unsupervised Learning
4. Ethics

Course Structure

W1 L1: Intro to ML

W1 L2: Intro to Classification

W2 L3: Naive Bayes Classification

W2 L4: Logistic Regression

W3 L5: Linear Regression

W3 L6: Decision Trees

W4 L7: Representing Data

W4 L8: Exploratory Data Analysis

W5 L9: Optimization

W5 L10: Generalization

W6 L11: Evaluation

W6 L12: Model Selection

W7 L13: Clustering

W7 L14: Non-Linear Dim Reduction

W8 L15: Recommender Systems

W8 L16: Neural Networks

W9 L17: Ethics and Fairness

W9 L18: Further Topics

Prohibited Course Combinations

Machine Learning and Pattern Recognition (MLPR)

You should **not** be taking both AML and MLPR **INFR11130**

Introductory Applied Machine Learning (IAML)

You should **not** have taken any of **INFR10069**, **INFR11182**, **INFR11130**, or **INFD11005** in the past or future

Other courses

You should **not** have taken Machine Learning (MLG) **INFR10086** or Data Analysis and Machine Learning 4 **ELEE10031**

Make sure you are eligible to take the course If not, contact your student advisor / personal tutor.

What is AML not?

AML is **not** an advanced Machine Learning course, instead try:

Machine Learning and Pattern Recognition (MLPR INFR11130) Sem 1

Machine Learning Theory (MLT INFR11202) Sem 2

AML is **not** a Deep Learning course, instead try:

Machine Learning Practical (MLP INFR11132) Sem 1&2

AML vs MLPR

AML

Covers some of the same material as MLPR, but with more focus on practical application of existing methods, than the mathematical construction and principles of the methods

MLPR

Also an introductory course, however, there are enough materials to stretch most of those with some machine learning background

<https://mlpr.inf.ed.ac.uk/2024/faq.html>

Make this decision soon (i.e. before end of week 2). Contact your student advisor / personal tutor if you need to change

Auditing AML

Some of you may want to take “audit” the course, i.e. not take it for credit

All the necessary information (e.g. lecture videos and slides) will be on the course webpage

No need to email us to ask for permission to audit

Students auditing **cannot** attend the labs or tutorials or get the coursework marked

Lectures - Videos

Pre-recorded - you watch on your own in advance of the Q&A sessions.

Watch in week i , Q&A in week $i+1$


Links on course Webpage -> **Schedule**


Home
Course Information
Schedule
Mini-Project
Exam
Staff

Schedule

DRPS Timetable

Week 1

Introduction to ML	Playlist • Slides • Handout
Introduction to Classification	Playlist • Slides • Handout
 Q&A Session	



Lectures - Videos

Each topic has ~1-3 videos per playlist. Watch them all, **in the correct order**




AML: Linear Regression - Introduction to Regression

0 0 0

Related Media

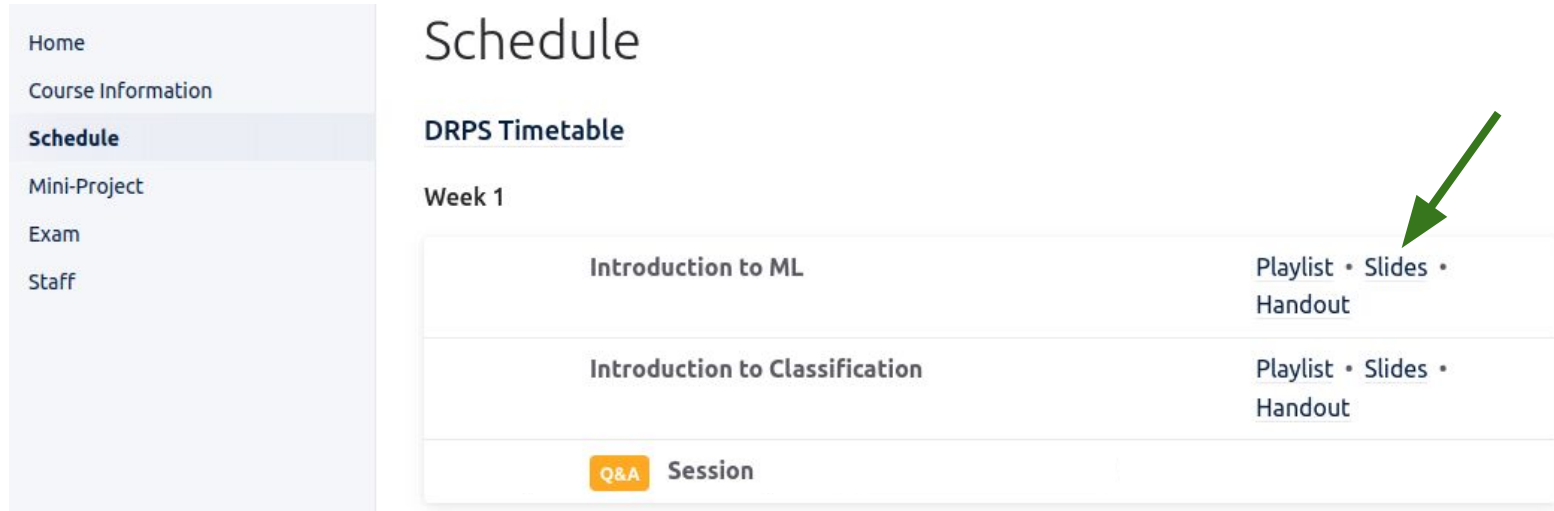
Make sure you are logged in correctly <https://media.ed.ac.uk>




Lectures - Slides

The slides for the lectures are available on the course website.

The class session slides will also be available there.



The screenshot shows a course website interface. On the left is a vertical navigation menu with the following items: Home, Course Information, **Schedule** (highlighted), Mini-Project, Exam, and Staff. The main content area is titled "Schedule" and contains a section for "DRPS Timetable". Underneath, it lists "Week 1" with three rows of lecture information:

Introduction to ML	Playlist • Slides • Handout
Introduction to Classification	Playlist • Slides • Handout
 Session	

A green arrow points to the "Slides" link in the first row of the schedule table.

Lectures - Recommendation

Note a time in your calendar each week to watch the videos e.g.

Wednesday at 2pm - Topic 1

Thursday at 2pm - Topic 2

Q&A Sessions

You are in/watching the class meeting right now!

Once a week: **Tues @ 4:10pm**

We will go over additional examples and questions

Opportunity to ask questions (best to do this on Piazza in advance - more on this later)

Q&A Sessions

Q&A sessions will be recorded so you can watch them offline later

Click here to find them on Learn - after some delay e.g. 1 day



Lecture Recordings

Access to lecture recordings for this course (Opens in a new window).

Labs

Weeks: 3, 5, 7, and 9

Python using Jupyter Notebooks in **Notable**

Lab sessions will be ~one hour long and conducted in Appleton Tower

Labs (and tutorials) are only for those taking the course for credit

Labs are not assessed

Numpy exercises

The following short exercises test your understanding of simple numpy functions and objects. Make sure you can complete them and feel free to reference the official [documentation](#) should you need it.

You may need to google some solutions

===== Question 1 =====

Print your numpy version.

In []: *# Your Code goes here:*

===== Question 2 =====

Create a zero vector of size 5.

In []: *# Your Code goes here:*

===== Question 3 =====

Create a zero vector of size 5 of type integer. Set the third element to 1.

In []: *# Your Code goes here:*

===== Question 4 =====

Create a vector ranging from 0 to 9.

In []: *# Your Code goes here:*

===== Question 5 =====

Create a vector ranging from 10 to 29.

In []: *# Your Code goes here:*

Labs

Your Lab groups assignment will be performed by ITO

Will be available in your timetable soon (MyEd -> Studies -> Timetable)

There are twelve lab sessions per week LAB01:LAB12

Only go to one of them, i.e. the one you have been assigned to

What programming
experience do you have?

Lab 0: Introduction

Lab 0 “**00 - Introduction.ipynb**” in GitLab page

Getting setup on Notable - online Jupyter Notebook. We only support Notable

Introduction to Python and some of the core libraries we will use e.g. numpy

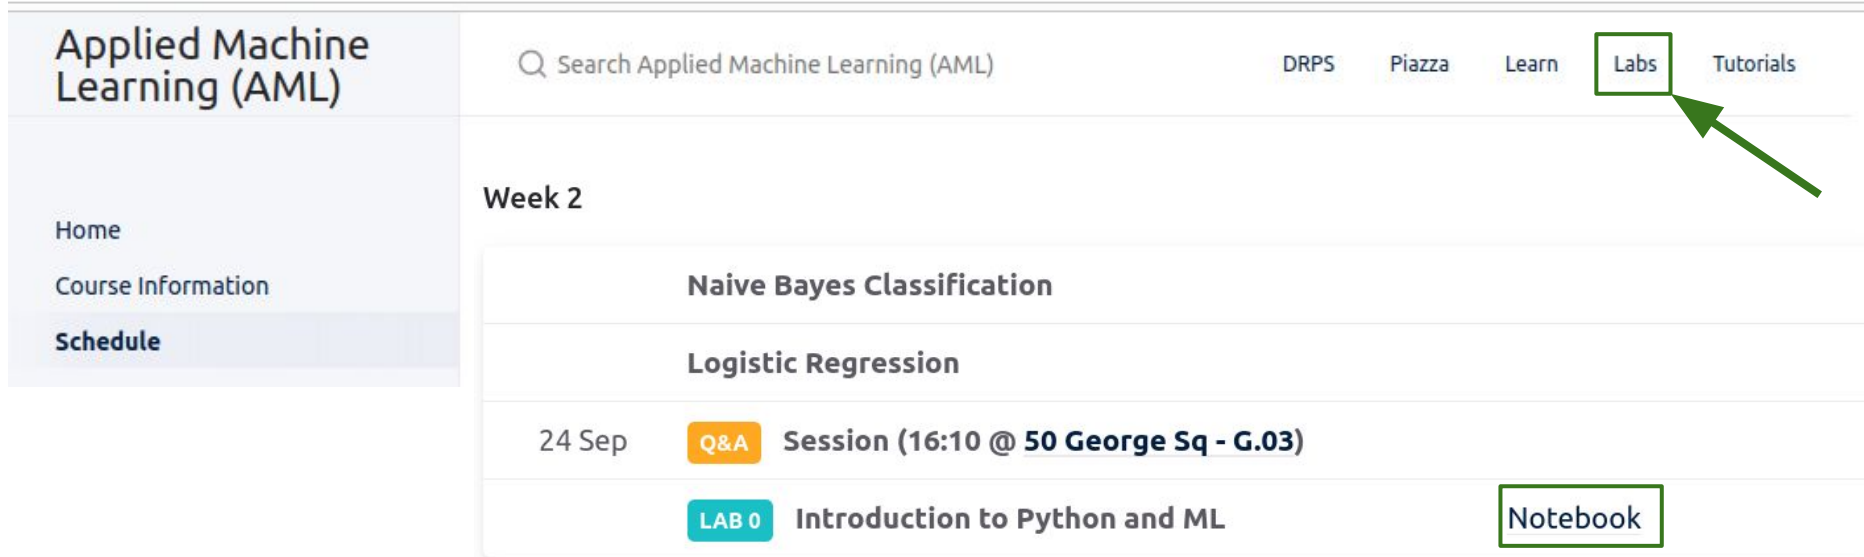
Important to be comfortably with this so that you can do the labs and courseworks in future weeks. Get started now!

If there is demand, we will arrange drop in labs next week

Lab 0: Introduction

URL for Lab 0 is available on the course webpage -> schedule

Links to solutions will also be provided 1 week later



The screenshot shows the top navigation bar of the Applied Machine Learning (AML) course website. The left sidebar contains links for Home, Course Information, and Schedule. The main navigation bar includes a search bar, DRPS, Piazza, Learn, Labs (highlighted with a green box and arrow), and Tutorials. The main content area displays the course schedule for Week 2, listing Naive Bayes Classification, Logistic Regression, and a Q&A session on 24 Sep. The Lab 0 entry, 'Introduction to Python and ML', is highlighted with a blue box, and a 'Notebook' link is highlighted with a green box.

Applied Machine Learning (AML)		DRPS	Piazza	Learn	Labs	Tutorials
Home	Week 2					
Course Information	Naive Bayes Classification					
Schedule	Logistic Regression					
	24 Sep	Q&A	Session (16:10 @ 50 George Sq - G.03)			
	LAB 0		Introduction to Python and ML			Notebook

Tutorials

Weeks: 4, 6, 8, and 10

Exam style questions

They will help you further understand the material so that you are better prepared for the exam

Tutorials

Weeks: 4, 6, 8, and 10

There are seven tutorial sessions per week

WORKSHOP01:WORKSHOP07

Only go to one of them, i.e. the one you have been assigned to

Tutorials are not assessed, but it is strongly recommended that you complete them

Lab and Tutorial Solutions

They will be available on the course website the week after a given event

e.g. the solution for Tutorial 1 (held in week 4), will be on the course website in week 5

Lab and Tutorial Groups

Note

You may not have been allocated to a group yet

MyEd -> Studies -> Timetable

	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
Monday		Applied Machine... OPT 12		Applied Machine... OPT 10		Applied Machine... OPT 07	Applied Machine... OPT 11	
Tuesday			Applied Machine... OPT 01	Applied Machine... OPT 02				Applied Machine Learn... Q&A SESSION APPLIED MACHINE LEARNING - Q&A SESSION 16:10 - 17:00 ▾ Course code INFR11211 Open course Sem 1 Week 1-10 Lecture Theatre G.03 - 50 George Square SEM 1 1 2 3 4 5 6 7 8 9 10 11
Wednesday			Applied Machine... OPT 02	Applied Machine... OPT 03				
		Applied Machine... OPT 03	Applied Machine... OPT 04			Applied Machine... OPT 04	Applied Machine... OPT 05	
				Applied Machine... OPT 06				

Lab and Tutorial Groups

Changing Groups

Please stick to your group, but if you notice a conflict now you may be able to request a group change. No need to contact the course organisers.

MyEd -> Studies -> Timetable

	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
Monday		Applied Machine... OPT 12		Applied Machine... OPT 10		Applied Machine... OPT 07	Applied Machine... OPT 11	
Tuesday			<div data-bbox="629 473 867 934" style="border: 2px solid red; padding: 5px;"><p>Applied Machine... OPT 01</p><p>TUTORIAL</p><p>APPLIED MACHINE LEARNING - WORKSHOP/01</p><p>11:10 - 12:00 ▾</p><hr/><p>Course code</p><p>INFR11211</p><p>Open course</p><p>Sem 1 Week 4, Sem 1 Week 6, Sem 1 Week 8, Sem 1 Week 10</p><p>5.04 - Teaching Studio - Appleton Tower</p><p>SEM 1</p><p>1 2 3 4 5 6 7 8 9 10 11</p></div>	Applied Machine... OPT 02				Applied Machine Learn... Q&A SESSION APPLIED MACHINE LEARNING - Q&A SESSION 16:10 - 17:00 ▾ <hr/> <p>Course code</p> <p>INFR11211</p> <p>Open course</p> <p>Sem 1 Week 1-10</p> <p>Lecture Theatre G.03 - 50 George Square</p> <p>SEM 1</p> <p>1 2 3 4 5 6 7 8 9 10 11</p>
Wednesday				Applied Machine... OPT 03	Applied Machine... OPT 04	Applied Machine... OPT 04	Applied Machine... OPT 05	
				Applied Machine... OPT 06				

Changing Lab/Tutorial Group

If your assigned time slot clashes with your other classes, you can submit a **Group Change Request Form**.

<https://www.ed.ac.uk/timetabling-examinations/timetabling/personalised-timetables/student-timetables>

<https://registryservices.ed.ac.uk/timetabling-examinations/timetabling/personalised-timetables/group-change-request>

However, we encourage you to only do this if you have a conflict. Otherwise, you should stay with the group you have been assigned.

Coursework

Coursework is worth 40%

Released in week 4 with a submission deadline of 21st of November

Group project of size three/four. We will use Piazza to allow you to create groups

Actively engaging with the labs, and tutorials will be very helpful preparation for the courseworks

Don't worry about this now. We will discuss more in the coming weeks

Coursework - How to submit

You will submit a PDF group report containing text and plots

There will be one report per group

A Latex template will be given to you

Don't worry about this now. We will discuss more in the coming weeks

Coursework Feedback

Marks and comments will be provided ~2 weeks after the submission deadline

Exam

Worth 60% - will take place in December (sometime between 9th and 20th)

Timetabling will provide the exact date later in the semester and form of the exam - **course organisers do not know**

Past solutions for all previous exams will not be made available, i.e. instructors will also not provide solutions on Piazza

Past Exam Papers

This is the third year AML (INFR11211) has run.

Past papers for IAML (INFR11182) are also very relevant

<https://exampapers.ed.ac.uk>

Note, in 2020 and 2021, the exam for IAML was “open book” and held online. Those exam papers are less relevant.

Common Marking Scheme

The University of Edinburgh uses a Common Marking Scheme (CMS) for taught student assessment. If you are unfamiliar with it, please take a look at this webpage:

<https://web.inf.ed.ac.uk/infweb/student-services/taught-students/information-for-students/information-for-all-students/your-studies/common-marking-scheme>

>70% is considered “excellent”

AML Piazza

You should have access via Learn if enrolled in AML

Your opportunity to ask questions in advance of class sessions

Use the relevant topics e.g. **topic_classification**

Ask questions in advance of the Q&A sessions i.e. latest by Monday afternoon for Tuesday's class. Earlier the better!

Learn

 [Piazza](#)

Access the Piazza discussion boards for the course.

AML Website

[DRPS](#)

[Piazza](#)

[Learn](#)

[Labs](#)

[Tutorials](#)

AML Piazza

You can also answer each other questions - peer learning

We will monitor Piazza and respond to questions

We do not provide solutions to all past exam questions (some sample solutions will be provided)

No discussion of the coursework across **different** groups

AML Piazza - Guidelines

- Check if someone has already asked/answered your question
- Questions should be relevant to the course
- Be respectful to each other and instructors
- If needed you can post an anonymous question, however the course organisers and TAs will be able to see your identity

Office Hours

No dedicated office hours

Instead, make use of the Piazza and these Q&A sessions to ask questions about the course content

Labs and tutorials are also helpful for discussing the respective pieces of work

Recommended Texts

Probabilistic Machine Learning: An Introduction, *Kevin Murphy*

<https://probml.github.io/pml-book/book1.html>

Pattern Recognition and Machine Learning, *Christopher Bishop*

<https://www.microsoft.com/en-us/research/people/cmbishop/prml-book>

Useful for extra background and detail, but not essential.

Note, there many more topics in these books that what we cover.

Best of Luck!

What do you want to get out of AML? Use this to guide your learning

Advice

- Monitor email and course webpage (but refresh) for announcements
- Watch the lectures before the class sessions
- Do not leave the coursework until the last minute
- Engage in the labs and tutorials
- Start on the labs and tutorials on your own before the sessions

Contacts

General Questions

- Piazza or Q&A Sessions

Administrative

- Course Secretary: Lindsay Seal

Otherwise

- Contact course organisers - Oisin and Sid
 - Best to use private message on Piazza, if confidential

Week 1: Your tasks for this week

- 1) Watch the (i) **Intro to ML** and (ii) **Intro to Classification** lectures
 - a) We will discuss these in the class session next week
- 2) Ask questions on Piazza about the lectures if stuck - earlier the better
- 3) Start “Lab 0 - Intro to Python”.

 Week 1

Introduction to ML	Playlist • Slides • Handout
Introduction to Classification	Playlist • Slides • Handout
Q&A Session	

Questions?

Course Website

<https://tinyurl.com/aml2024>

