

Machine Learning Engineering

As Artificial Intelligence becomes increasingly sophisticated and ubiquitous, the role of Machine Learning Engineering grows increasingly important. Machine Learning Engineers operate at the intersection of Software Engineering and Machine Learning and are responsible for **integrating AI systems into production systems**. Without Machine Learning Engineers, cutting-edge AI products wouldn't become a reality.

What Machine Learning is:

Machine Learning Engineers design, develop and maintain robust and scalable data science and machine learning pipelines. MLE encompasses duties such as:

- **Model training:** MLEs are responsible for preparing data sets, data scraping and storage for production-ready machine learning systems.
- **Model Deployment:** Implementing models in production environments, ensuring they are scalable, efficient, and can handle real-world data.
- **Monitoring and Maintenance:** Continuously monitoring model performance, handling model drift, retraining models, and ensuring they remain accurate and relevant.
- **Collaboration:** Working closely with data science, data engineers and software engineers to integrate machine learning solutions into broader systems.

What Machine Learning isn't:

- **Pure data science/machine learning:** MLEs are not solely focused on hypothesis testing, or theoretical aspects of data science.
- **Traditional software engineering:** While closely related, MLEs focus on the unique challenges of integrating non-deterministic systems into production environments.
- **Research and development:** MLEs typically implement existing models, or models developed by data science teams.
- **Exploratory model development:** MLEs go beyond training models for a specific project, it involved continuous management, retraining and updates as data improves.

About the career

With an average salary of USD 153,252 [1], Machine Learning Engineers come in at number eight in the 2023 Indeed Top Jobs survey, which considers the number of opportunities, compensation, and job flexibility over the previous three years [1]. In 2019, Indeed listed Machine Learning Engineering as its #1 job of the year [2]. Machine Learning Engineers are increasingly sought after by employers like OpenAI [3], Canva [4], CSIRO, and Real Estate Australia.

Pre-requisites

- COMP1531: Software Engineering Fundamentals
- ML?

Course Learning Outcomes

CL01: Develop machine learning models using pre-designed algorithms, datasets, and techniques

CL02: Analyse and sanitise datasets for machine learning applications, including data scraping, and storage

CL03: Monitor, test and maintain machine learning pipelines in cloud environments, including writing unit tests for ML models, designing dashboards to monitor performance and accuracy

CL04: Contribute to existing production-quality systems

Syllabus

Week	Content
1	Data Collection <ul style="list-style-type: none">• Dataset scraping• Storage (data lakes vs databases)• Pipelines
2	Data Analysis and Evaluation <ul style="list-style-type: none">• Raw and Tidy Data• Data sanitisation• Training and Holdout sets
3	Types of common machine learning goals <ul style="list-style-type: none">• Classification• Generation
4	Model Deployment <ul style="list-style-type: none">• Static deployment vs dynamic deployment• On-device deployment
5	Model Deployment <ul style="list-style-type: none">• Cloud deployment
6	Model Serving, Monitoring and Maintenance <ul style="list-style-type: none">• Security• Developer experience
7	Model Serving, Monitoring and Maintenance <ul style="list-style-type: none">• Ease of Recovery• Logging and Dashboards (Grafana)
8	Contributing to Open-Source Systems <ul style="list-style-type: none">• Guest lecture
9	Final Project <ul style="list-style-type: none">• Guest lecture



Weekly contact hours

- 1x2h Lecture
- 1x3h Laboratory

Assessment Structure

Assessment 1: Portfolio of tasks –60%

A range of 2-3 tasks per week, with formative feedback provided.

Tasks due two weeks after task start-date.

- Data collection, sanitisation, and storage: 20%
- Model training, deployment: 20%
- Ongoing deployment and robustness: 20%

Assessment 2: Take-home final project – 40%

Individual or group-based final project.

Due at end of term.

- If individual: 40%,
- If group:
 - 30% project delivery,
 - 10% peer report

References:

[1] <https://www.indeed.com/career-advice/news/best-jobs-of-2023>

[2] <https://indeed.com/lead/best-jobs-2019>

[3] <https://openai.com/careers/machine-learning-engineer-applied-ai/>

[4] <https://www.lifeatcanva.com/en/jobs/743999982220573/graduate-machine-learning-engineer-anz/>



UNSW
SYDNEY