


AI Engineering Bootcamp

26 Week Program



WeCloudData

Contents

3	Introduction
4	Why WeCloudData
5	Who This Program Is For
6	Program Introduction
7	Program Details
8	Syllabus
12	Summary
13	Instructors
14	Admission Process
14	Tuition, Grants and Financial Options



Introduction

In Canada, Machine Learning Engineers (MLE) will earn around

\$120k

In the U.S., the average base pay for ML Engineers is around

\$130k



Placement rate of WeCloudData alumni

98%




Why WeCloudData

WeCloudData is the only bootcamp and accredited career college that focuses on the data field. We offer five distinctive programs - Data Science, Business Intelligence, Data Engineering, Artificial Intelligence, and Cloud Computing to prepare you for your dream career. The programs are instructor-led and project-based. You will learn advanced data techniques at a fraction of the cost from other institutes.

- Relevant to industry: Our programs cover the most up-to-date skills that are required by employers across the world. In addition to offering personal training, we also work with a wide range of well-known data-driven companies in North America to help them up-skill their employees for cutting-edge technologies. That being said, we know what qualifications employers are looking for in future candidates.
- Focus on practice: Our course design committee is comprised of data scientists from a range of organizations such as BMO, RBC, Dessa, Google, and Amazon which ensure real-world projects are built into each course curriculum. Our programs are rigorous and students will be expected to complete the assignments and projects to meet the course completion requirements.
- Service is our success: We understand the challenges of career transformation, and we're committed to helping each and every student gain the knowledge, skills, and experience that will help them launch a successful career. You can communicate with the instructors and TA's regularly through our online platform and communication app. We are here to ensure your positive outcome and many accomplishments from our programs.



Who This Program Is For



Data scientists, data engineers, software engineers, and AI researchers who are interested in ML model deployment in production.

Suitable for anyone who has prior experience with machine learning, data science, and has solid programming skills.



New grads from computer science (CS) or software engineering (SE) background.



Program Introduction

According to deeplearning.ai, only 22% of companies using machine learning have successfully deployed a model. The need for ML Engineers is growing exponentially as the industry moves towards Data-centric AI. ML Engineering (MLOps) is at the intersection of Machine Learning, DevOps, and Data Engineering. It is a critical role that makes sure the AI products get deployed in production in a scalable and reliable way.

In this advanced program, you will learn deep learning techniques, use cases in computer vision and NLP, and the engineering and operations side which allow you to master the entire lifecycle of ML projects from data preparation to model deployment and monitoring.



Program Details

Length

26 weeks

Schedule

Thur 6:30PM - 9:30PM

Sun 9:00AM - 5:00PM

Lab: Sun 1:00PM - :00PM

TA Office Hours: Evenings

Price

\$9,100 CAD (Part-time)

\$13,000 CAD (Full-time with Real Client project)

Course Structure

Overview of program



General Structure

Semester 1

- Machine Learning Foundations
- Deep Learning Foundations
- Project #1
- DL: Computer Vision
- DL: Natural Language Processing
- Project #2

Semester 2:

- Linux, Docker, and Github
- Intro to Cloud Computing (GCP)
- Infrastructure as Code (Terraform)
- Building Prediction Services
- Model Packaging & Deployment
- Model Versioning, Packing & Tracking with MLflow
- Model Monitoring
- Data Versioning Control
- CI/CD
- Project #3

Syllabus

Machine Learning Foundations

Content	What Will Be Covered & Tools Used
Review of Basic Concepts + Data Processing	
Parameter and Hyperparameter Optimization	
Supervised Learning	
Unsupervised Learning	<ul style="list-style-type: none">• Ensemble learning clustering

Deep Learning Foundations

Content	What Will Be Covered & Tools Used
Introduction to Neural Networks	<ul style="list-style-type: none">• Neural Network as a flexible system of learning• History of machine learning starting from perceptrons• What is deep learning?
Forward and Backward Propagations	<ul style="list-style-type: none">• Mathematics of Forward and Backward propagations• Forward and Backward propagation in PyTorch
Hyperparameter Optimization in Deep Learning	<ul style="list-style-type: none">• Important hyperparameters in Neural Network modeling• Intuition behind each hyperparameter• Playing with nnet hyperparameters in PyTorch
Intro to PyTorch	<ul style="list-style-type: none">• Neural Network architecture specification and initialization, Forward and Backward propagation in PyTorch• Review of PyTorch tutorials
Convolutional Neural Networks	<ul style="list-style-type: none">• Theoretical aspects of Convolutional Neural Networks• CNN for image classification• CNN for image segmentation• Introduction to transfer learning
Recurrent Neural Networks	<ul style="list-style-type: none">• Introduction to Recurrent Neural Networks• Application of RNN in NLP, speech recognition, etc. widely used units in RNN such as GML and LSTM
Transformers and Attention Mechanism	<ul style="list-style-type: none">• Introduction to attention• Multi-head attention in NLP• Transformers

Deep Learning -- Computer Vision

Content	What Will Be Covered & Tools Used
Literature Review/Discussion on the Latest Papers	<ul style="list-style-type: none">• How to search for relevant literature• How to evaluate sources and find relevant concepts• Identify themes, debates and gaps• Different areas of ML and how to get started
CNN Architectures + Visualizing and Understanding CNN models	<ul style="list-style-type: none">• CNN architectures; Modern Convolutional Neural Networks; Visualizing CNN
Object Detection	<ul style="list-style-type: none">• Detecting a single/multiple object• Bounding Box Predictions• Sliding Windows• Anchor Boxes• YOLO Algorithm for Object Detection
Semantic Segmentation	<ul style="list-style-type: none">• What is semantic segmentation?• Semantic Segmentation before and after Deep Learning• Constitutional Neural Network for Segmentation such as UNet
Ingesting Data, Dealing with Data, Image Processing	<ul style="list-style-type: none">• Introduction to Image Pre-processing• Image Processing using OpenCV
Image Augmentation & Fine-Tuning	<ul style="list-style-type: none">• Common image augmentation methods and strategies; transfer learning via fine-tuning
Industry Use Cases	<ul style="list-style-type: none">• CV in industry applications

Deep Learning -- NLP

Content	What Will Be Covered & Tools Used
Introduction to NLP (Fundamentals + RNN; How can we solve problems with NLP?)	
Traditional and DNN-based Methods Overview (topic modelling, n-grams, etc.)	<ul style="list-style-type: none">• Sklearn: Traditional algorithms• Gensim: LDA
[Basic Models] Word Vectors: Word2Vec (Skip-gram), Glove	<ul style="list-style-type: none">• Gensim: word2vec (skip-gram, CBOW Glove)
[Basic Models] BERT - Modelling, Tips and Applications, Attention Mechanisms, HuggingFace	<ul style="list-style-type: none">• HuggingFace: BERT• Fairseq: RoBERT
Generative Models - GPT	
Embedding Use Cases (search engines, recommendations, ranking)	<ul style="list-style-type: none">• ElasticSearch - Search Engines• DPR - Ranking• SentenceTransformers - Recommendations/Ranking
Classification Use Cases #1 (e.g. token, sentence, text) + Introduction of Datasets	<ul style="list-style-type: none">• HuggingFace

Lecture	What Will Be Covered & Tools Used
Classification Use Cases #2	<ul style="list-style-type: none"> • HuggingFace
Seq2seq Models for Machine Translation	<ul style="list-style-type: none"> • HuggingFace
Differentiating Between Models	<ul style="list-style-type: none"> • HuggingFace: BERT, RoBERT, alberta, T5, DistilBERT

MLOps

Lecture	What Will Be Covered & Tools Used
Linux, Docker, and Github	<ul style="list-style-type: none"> • Linux • Github • yaml files • Containers • Docker
Intro to Cloud Computing (GCP)	<ul style="list-style-type: none"> • Intro to GCP managed services specifically tailored for data and ML • Tools: Vertex AI
Infrastructure as Code (IaC with Terraform)	<ul style="list-style-type: none"> • Intro to Terraform scripts • Standing up and tearing down infrastructure • Tools: Terraform
Building Prediction Services #1	<ul style="list-style-type: none"> • Building ML models • Offline and online prediction • Serving • Tools: Python
Building Prediction Services #2	<ul style="list-style-type: none"> • Building and deploying models • REST endpoints • Tools: Python, FASTAPI/gunicorn
Model Packaging & Deployment	<ul style="list-style-type: none"> • Yaml files • Containers • Docker, Docker images • Intro to Kubernetes • Docker run command • Tools: Docker, kubectI
Model Versioning, Packaging & Tracking with MLflow	<ul style="list-style-type: none"> • ML models, versions, project structure, MLflow • Tools: MLflow, Python
Kubernetes	<ul style="list-style-type: none"> • Intro to Kubernetes • Kubernetes basics, running docker images on Kubernetes • Container orchestration • Tools: Kubernetes, kubectI

Lecture	What Will Be Covered & Tools Used
Kubeflow	<ul style="list-style-type: none"> • Kubeflow pipelines • Pipeline components • DAG, experiments • Output artifacts and ML metadata • Tools: Kubeflow
Seldon for Deployment	<ul style="list-style-type: none"> • Seldon concepts • Intro to Seldon Core • Seldon ML inference servers • Intro to Seldon Deploy and Alibi • Tools: Seldon, kubectl
ML Feature Store	<ul style="list-style-type: none"> • Intro to Feature Store and Concepts • Materializing features on feature store • Use for online and offline serving of models • Tools: Feast, Vertex AI feature store
ML Pipelines with TFX and Apache Airflow	<ul style="list-style-type: none"> • Intro to TFX for Model Pipelines • Comparison with Kubeflow • Orchestration using Airflow • Tools: Python, TFX, Airflow/GCP Cloud Composer
Model Monitoring	<ul style="list-style-type: none"> • Model Monitoring basics • Intro to Model Monitoring MLflow, and other tools like Neptune, Sagemaker, Weights and Biases etc.
Data Versioning Control (DVC)	<ul style="list-style-type: none"> • Data Versioning • Touch on DVC, Pachyderm, etc.
CI/CD 101	<ul style="list-style-type: none"> • CI/CD - Github
CI/CD Jenkins	<ul style="list-style-type: none"> • CI/CD - Jenkins



Summary

Less theory-based, more on practical aspect + use cases

- Focus on how to put things into production
- Strong learning support in program

Projects students work on will be modeled based on real-life use cases/projects (<https://weclouddata.com/courses/ai-engineering-bootcamp/>)

Possible project examples:

- Object detection, image processing using DL
semantic segmentation, autonomous driving, IoT -
creating a robot

Career mentorship and support at the end of the program

Option to join our real-client projects



Instructors



Ali Madani

Director of Machine Learning, Cyclica

Ali is the head of machine learning at Cyclica Inc. and leads the team to further improve Cyclica's deep learning technology for predicting interaction between ligands and target proteins. As a computational biologist and machine learning specialist, Ali has worked on a series of scientific articles in high impact scientific journals and international conferences covering such fields as transfer learning, dimensionality reduction and unsupervised clustering. He earned a Ph.D degree from the University of Toronto, and a master of a mathematics degree from the University of Waterloo.

Farnoosh Khodakarami

Machine Learning Researcher, Cyclica

Experienced computer scientist with a demonstrated history of working in the research industry. Skilled in application development with experience in machine learning applications. Strong research professional with a Doctor of Philosophy (Ph.D.) focused in Computer Science from the Amirkabir University of Technology - Tehran Polytechnic. Creative, self-motivated and committed to working with a team-player attitude, great problem-solving skills and ability to quickly grasp new concepts.



Yi Zhang

Senior AI Applied Research Scientist, ServiceNow ATG

Dr. Yi Zhang is a Computer Scientist specializing in NLP, Machine Learning, and AI, working as a Senior AI Applied Research Scientist at ServiceNow Advanced Technology Group. He has years of academia and industry experience and has architected solutions to convert the latest technology to robust Business Intelligent. Prior to his previous experience as the Head of Engineering at NLPmatics, he also assumed team lead role at Tencent and Research Scientist at Chan Zuckerberg Institute.

Indrani Gorti

Director of Platform and ML engineering, Loblaw Digital

Indrani Gorti is a polyglot with experience in data engineering, data science, and solving business problems using ML and AI. She has worked in major industries like Finance, Telecommunications, Technology, Retail, and Startups. Her work has involved end-to-end projects from data ingestion, data analysis to the deployment of models in various setups using user transaction and interaction logs.

Indrani is a Director of Platform and ML engineering at Loblaw Digital. She holds a BS and MS in Computer Science.




Admission Process



Book a Counseling Meeting

Our program advisor will have a 1-on-1 meeting with you to see if the program is a good fit for you. There is no application fee. It takes about 30 minutes



Pass the Challenge

There will be a technical test and an interview. Applicants spend up to 2 hours on the challenge.



Reserve the Spot & Sign the Contract

Our admissions officer will work with you directly to help you fill out a contract, pay a \$500 deposit and assist you in applying for any kind of grant and finance options.



Start your Pre-bootcamp Learning

Preparing ahead will help you gain more experience and competence. Research shows that preview and preparation account for 73.7% success in academic achievement of university students.

Tuition, Grants and Financial Options

FULL-TIME

\$13,000
CAD

PART-TIME

\$9,100
CAD

As an Ontario registered private career college, you can apply for student line of credit from BMO with lower interest rate.

Scholarship

WeCloudData is offering student scholarships. Please contact our advisor for more information.

Grant

With Ontario Second Career grant you may be eligible for up to \$28,000 for costs including: tuition, books, manuals, transportation, basic living allowance, child care.

Contact Us

Contact our program advisor, Amir, for more information
info@weclouddata.com

