Course Outline

School:	Eng. Tech. & Applied Science
Department:	Information and Communication Engineering Technology (ICET)
Course Title:	Unsupervised and Reinforcement Learning
Course Code:	COMP 257
Course Hours/Credits:	56
Prerequisites:	COMP 247, COMP 254
Co-requisites:	N/A
Eligible for Prior Learning, Assessment and Recognition:	Yes
Originated by:	Hao Lac
Current Semester:	Fall 2022
Approved by:	ppesikan c/o

Chairperson/Dean

Students are expected to review and understand all areas of the course outline.

Retain this course outline for future transfer credit applications. A fee may be charged for additional copies.

This course outline is available in alternative formats upon request.

Acknowledgement of Traditional Lands

Centennial is proud to be a part of a rich history of education in this province and in this city. We acknowledge that we are on the treaty lands and territory of the Mississaugas of the Credit First Nation and pay tribute to their legacy and the legacy of all First Peoples of Canada, as we strengthen ties with the communities we serve and build the future through learning and through our graduates. Today the traditional meeting place of Toronto is still home to many Indigenous People from across Turtle Island and we are grateful to have the opportunity to work in the communities that have grown in the treaty lands of the Mississaugas. We acknowledge that we are all treaty people and accept our responsibility to honor all our relations.

Course Description

In the first half of this course, students will be exposed to unsupervised learning (dimensionality reduction, k-means clustering, hierarchical clustering, Gaussian mixtures, autoencoders, and Kohonen Self-Organizing Map (SOM)).
In the second half of the course, students will be exposed to Reinforcement Learning (policy gradient, Markov Decision Processes, Q-Learning, and TF-Agents Library).
Students will gain hands-on experience by applying unsupervised learning and reinforcement learning techniques.

Program Outcomes

Successful completion of this and other courses in the program culminates in the achievement of the Vocational Learning Outcomes (program outcomes) set by the Ministry of Colleges and Universities in the Program Standard. The VLOs express the learning a student must reliably demonstrate before graduation. To ensure a meaningful learning experience and to better understand how this course and program prepare graduates for success, students are encouraged to review the Program Standard by visiting http://www.tcu.gov.on.ca/pepg/audiences/colleges/progstan/. For apprenticeship-based programs, visit http://www.collegeoftrades.ca/training-standards.

Course Learning Outcomes

The student will reliably demonstrate the ability to:

1. Understand the fundamental concepts of various unsupervised learning techniques.
2. Understand the fundamental concepts of Reinforcement Learning.
3. Analyze and implement appropriate intelligent solutions to problems by applying relevant unsupervised learning techniques or Reinforcement Learning.
4. Demonstrate the ability to work in a pair programming environment.
5. Code and test the application.
6. Produce both user and system documentation.
7. Document the programming aspects of a system.
8. Apply appropriate data structures and software engineering practices applicable to the goals of the system.
9. Present, explain decisions taken to solve the problem, and demonstrate the implemented system.

Essential Employability Skills (EES)

The student will reliably demonstrate the ability to*:

1. Communicate clearly, concisely and correctly in the written, spoken, and visual form that fulfills the purpose and meets the needs of the audience.
2. Respond to written, spoken, or visual messages in a manner that ensures effective communication.
3. Execute mathematical operations accurately.
4. Apply a systematic approach to solve problems.
5. Use a variety of thinking skills to anticipate and solve problems.
6. Locate, select, organize, and document information using appropriate technology and information systems.
7. Analyze, evaluate, and apply relevant information from a variety of sources.
9. Interact with others in groups or teams in ways that contribute to effective working relationships and the achievement of goals.
10. Manage the use of time and other resources to complete projects.
11. Take responsibility for one's own actions, decisions, and consequences.

**There are 11 Essential Employability Skills outcomes as per the Ministry Program Standard. Of these 11 outcomes, the following will be assessed in this course.*

Global Citizenship and Equity (GC&E) Outcomes

N/A

Methods of Instruction

Online

Text and other Instructional/Learning Materials

Text Book(s):

Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc

Online Resource(s):

https://scikit-learn.org

https://www.tensorflow.org

https://keras.io

https://gym.openai.com

https://www.javatpoint.com/hierarchical-clustering-in-machine-learning

https://www.jeremyjordan.me/autoencoders/

https://towardsdatascience.com/how-to-implement-kohonens-self-organizing-maps-989c4da05f19

https://towardsdatascience.com/kohonen-self-organizing-maps-a29040d688da

https://github.com/Kursula/Kohonen_SOM

Material(s) required for completing this course:

Access to a fast computer.

Custom Courseware:

Open source software.

Evaluation Scheme

- Dimensionality Reduction using PCA (Assignment 1): Reducing data dimensionality via PCA.
- K-Means Clustering (Assignment 2): K-Means Clustering for unlabelled data.
- Hierarchical Clustering (Assignment 3): Structured Hierarchical Clustering for unlabelled data.
- Gaussian Mixture Models (Assignment 4): Probabilistic approach to clustering using Gaussian Mixture Models.
- Autoencoders (Assignment 5): Object reconstruction using autoencoders.
- Reinforcement Learning (Assignment 6): Agent training using TF-Agents.
- Final Project: Final project based on real-world dataset.

Evaluation Name	CLO(s)	EES Outcome(s)	GCE Outcome(s)	Weight/100
Dimensionality Reduction using PCA (Assignment 1)	1, 3, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 10, 11		10
K-Means Clustering (Assignment 2)	1, 3, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 10, 11		10
Hierarchical Clustering (Assignment 3)	1, 3, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 10, 11		10
Gaussian Mixture Models (Assignment 4)	1, 3, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 10, 11		15
Autoencoders (Assignment 5)	1, 3, 4, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 9, 10, 11		15
Reinforcement Learning (Assignment 6)	2, 3, 4, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 9, 10, 11		15
Final Project	1, 3, 4, 5, 6, 7, 8, 9	1, 2, 3, 4, 5, 6, 7, 9, 10, 11		25
Total				100%

If students are unable to write a test they should immediately contact their professor or program Chair for advice. In exceptional and well documented circumstances (e.g. unforeseen family problems, serious illness, or death of a close family member), students may be able to write a make-up test.

All submitted work may be reviewed for authenticity and originality utilizing Turnitin®. Students who do not wish to have their work submitted to Turnitin® must, by the end of the second week of class, communicate this in writing to the instructor and make mutually agreeable alternate arrangements.

When writing tests, students must be able to produce official Centennial College photo identification or they may be refused the right to take the test or test results will be void.

Tests or assignments conducted remotely may require the use of online proctoring technology where the student's identification is verified and their activity is monitored and/or recorded, both audibly and visually

through remote access to the student's computer and web camera. Students must communicate in writing to the instructor as soon as possible and prior to the test or assignment due date if they require an alternate assessment format to explore mutually agreeable alternatives.

Student Accommodation

The Centre for Accessible Learning and Counselling Services (CALCS) (http://centennialcollege.ca/calcs) provides programs and services which empower students in meeting their wellness goals, accommodation and disability-related needs. Our team of professional psychotherapists, social workers, educators, and staff offer brief, solution-focused psychotherapy, accommodation planning, health and wellness education, group counselling, psycho-educational workshops, adaptive technology, and peer support. Walk in for your first intake session at one of our service locations (Ashtonbee Room L1-04, Morningside Room 190, Progress Room C1-03, The Story Arts Centre Room 285, Downsview Room 105) or contact us at calcs@centennialcollege.ca, 416-289-5000 ext. 3850 to learn more about accessing CALCS services.

Use of Dictionaries

- Any dictionary (hard copy or electronic) may be used in regular class work.

Program or School Policies

N/A

Course Policies

Meeting with the instructor in a team based assignment/project requires all members to be present to promote transparency.

College Policies

Students should familiarize themselves with all College Policies that cover academic matters and student conduct.

All students and employees have the right to study and work in an environment that is free from discrimination and harassment and promotes respect and equity. Centennial policies ensure all incidents of harassment, discrimination, bullying and violence will be addressed and responded to accordingly.

Academic Honesty

Academic honesty is integral to the learning process and a necessary ingredient of academic integrity. Forms of academic dishonesty include cheating, plagiarism, and impersonation, among others. Breaches of academic honesty may result in a failing grade on the assignment or course, suspension, or expulsion from the college. Students are bound to the College’s AC100-11 Academic Honesty and Plagiarism policy.

To learn more, please visit the Libraries information page about Academic Integrity https://libraryguides.centennialcollege.ca/academicintegrity and review Centennial College's Academic

Honesty Module:
https://myappform.centennialcollege.ca/ecentennial/articulate/Centennial_College_Academic_Integrity_Module_%202/story.html

Use of Lecture/Course Materials

Materials used in Centennial College courses are subject to Intellectual Property and Copyright protection, and as such cannot be used and posted for public dissemination without prior permission from the original creator or copyright holder (e.g., student/professor/the College/or third-party source). This includes class/lecture recordings, course materials, and third-party copyright-protected materials (such as images, book chapters and articles). Copyright protections are automatic once an original work is created, and applies whether or not a copyright statement appears on the material. Students and employees are bound by College policies, including AC100-22 Intellectual Property, and SL100-02 Student Code of Conduct, and any student or employee found to be using or posting course materials or recordings for public dissemination without permission and/or inappropriately is in breach of these policies and may be sanctioned.

For more information on these and other policies, please visit www.centennialcollege.ca/about-centennial/college-overview/college-policies.

Students enrolled in a joint or collaborative program are subject to the partner institution's academic policies.

PLAR Process

This course is eligible for Prior Learning Assessment and Recognition (PLAR). PLAR is a process by which course credit may be granted for past learning acquired through work or other life experiences. The PLAR process involves completing an assessment (portfolio, test, assignment, etc.) that reliably demonstrates achievement of the course learning outcomes. Contact the academic school to obtain information on the PLAR process and the required assessment.

Topical Outline (subject to change):

Week	Topics	Readings/Materials	Weekly Learning Outcome(s)	Instructional Strategies	Evaluation Name and Weight	Evaluation Date
1	-Review of basic data pipeline used in Machine Learning systems. -Setup development environment. -Data preprocessing (data cleaning, feature scaling, transformation).	"Week 1: Introduction and Review" notes. Chapter 2, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand the basic data pipeline in Machine Learning systems. -Understand how to setup an isolated environment for development. -Understand basic data cleaning, feature scaling, and the transformation pipeline.	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 1
2	-The curse of dimensionality. -Data dimension reduction using projection, manifold learning, and PCA.	"Week 2: Dimensionality Reduction" notes. Chapter 8, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand the concept of the curse of dimensionality in real-world data. -Understand the concept of projection and manifold learning. -Understand the various approach based on PCA including randomized PCA, incremental PCA, and kernel PCA.	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 2
3	-Clustering. -Applications of clustering. -K-Means clustering and the K-Means algorithm. -Centroid initialization methods. -The K-Means++ algorithm. -Finding the optimal number of clusters.	"Week 3: Clustering and K-means" notes. Chapter 9, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand the benefits of clustering. -Understand K-Means and its algorithm. -Understand K-Means++ and its algorithm. -Understand centroid initialization methods and its use for avoid local optimum. -Understand how to find the optimal number of clusters.	Lecturing, Demos, Practice Problems, Individual work outside class.	Dimensionality Reduction using PCA (10%)	Friday of Week 3 by 11:59pm EST
4	-Hierarchical Clustering. -Agglomerative Hierarchical	"Week 4: Hierarchical Clustering" notes.	-Understand Hierarchical Clustering. -Understand Agglomerative Hierarchical Clustering (AHC) and Divisive Hierarchical Clustering (DHC).	Lecturing, Demos, Practice Problems, Individual work outside class.	K-Means Clustering (10%)	Friday of Week 4 by 11:59pm EST

Week	Topics	Readings/Materials	Weekly Learning Outcome(s)	Instructional Strategies	Evaluation Name and Weight	Evaluation Date
	Clustering (AHC) and Divisive Hierarchical Clustering (DHC). -Clusters and dendrograms. -Distance measure between two clusters (single linkage, complete linkage, average linkage, centroid linkage). -Metrics for measure similarity (Euclidean distance, Manhattan distance, Minkowski distance, Jaccard Similarity Coefficient/Jaccard Index, Cosine similarity, Gower's similarity coefficient).		-Clusters and dendrograms. -Understanding different rules of distance measure between two clusters (single linkage, complete linkage, average linkage, centroid linkage). -Understand different metrics for measure similarity between instances (Euclidean distance, Manhattan distance, Minkowski distance, Jaccard Similarity Coefficient/Jaccard Index, Cosine similarity, Gower's similarity coefficient).			
5	-Gaussian Mixture Model (GMM). -Expectation-Maximization (EM) algorithm. -Limitations of GMM. -Bayesian information criterion (BIC) and the Akaike information criterion (AIC). -Bayesian Gaussian Mixture Models (BGMM). -KL divergence.	"Week 5: Gaussian Mixtures" notes. Chapter 9, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand Gaussian Mixture Model (GMM) and how it addresses the limitations of K-Means. -Understand the use of the Expectation-Maximization (EM) algorithm to perform soft cluster assignments. -Understand that limitations of GMM. -Understand the use of metrics Bayesian information criterion (BIC) and the Akaike information criterion (AIC) for selecting the number of clusters. -Understand Bayesian Gaussian Mixture Models (BGMM) for clustering. -Understand KL divergence minimization to obtain good approximations to p(z|X).	Lecturing, Demos, Practice Problems, Individual work outside class.	Hierarchical Clustering (10%)	Friday of Week 5 by 11:59pm EST
6	-Autoencoders. -Undercomplete autoencoders. -Sparse autoencoders.	"Week 6: Autoencoders" notes.	-Understand autoencoders as encoding and decoding networks. -Understand undercomplete autoencoders as more powerful nonlinear generalization of PCA.	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 6

Week	Topics	Readings/Materials	Weekly Learning Outcome(s)	Instructional Strategies	Evaluation Name and Weight	Evaluation Date
	-Sparsity constraint. -Denoising autoencoders.		-Understand that sparse autoencoders as alternative to reducing hidden unit numbers. -Understanding sparsity constraints to achieve sparse autoencoders. -Understand that denoising autoencoders as generalizable models for slightly corrupt input data with uncorrupted data as target output.			
7	-Kohonen Self-Organizing Map (SOM). -Structure of SOM. -How SOM learns. -The SOM learning algorithm.	"Week 7: Kohonen Self-Organizing Map (SOM)" notes.	-Understand the fundamentals of Kohonen Self-Organizing Map (SOM). -Understand how the structure of SOM naturally displays the inherent clustering of data. -Understand how SOM learns. -Understand the SOM learning algorithm.	Lecturing, Demos, Practice Problems, Individual work outside class.	Gaussian Mixture Models (15%)	Friday of Week 7 by 11:59pm EST
8	-Rewards in Reinforcement Learning. -Policy search strategies.	"Week 8: Introduction to Reinforcement Learning (RL)" notes. Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand rewards in Reinforcement Learning. -Understand different policy search strategies.	Lecturing, Demos, Practice Problems, Individual work outside class.	Autoencoders (15%)	Friday of Week 8 by 11:59pm EST
9	-OpenAI Gym. -Neural Network policies. -Evaluating actions and the credit assignment problem. -Policy gradients.	"Week 9: Environment and Policy Gradients" notes. Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand OpenAI Gym as the agent's environment and how to install it. -Understand Neural Network policies for learning better actions resulting in higher rewards. - Understand how to evaluate actions and challenges due to the credit assignment problem. -Understand policy gradients for selecting higher rewards action paths.	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 9
10	-Markov chains. -Markov Decision Processes.	"Week 10: Markov Decision Processes" notes.	-Understand the concept of Markov chains. -Understand Markov Decision Processes. -Learn how to estimate optimal state value	Lecturing, Demos, Practice Problems, Individual work outside		Week 10

Week	Topics	Readings/Materials	Weekly Learning Outcome(s)	Instructional Strategies	Evaluation Name and Weight	Evaluation Date
	-Estimaing optimal state value. -Temporal Difference Learning.	Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	using the Bellman Optimality Equation, Dynamic Programming, and the Q-Value Iteration algorithm. -Understand the Temporal Difference Learning algorithm when the agent has partial knowledge.	class.		
11	-Q-Learning. -Exploration policies. -Approximate Q-Learning and Deep Q-Learning. -Deep Q-Network (DQN).	"Week 11: Q-Learning" notes. Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Understand Q-Learning as an off-policy as opposed to policy gradients' on-policy. -Understand exploration policies. -Understand approximate Q-Learning and Deep Q-Learning. -Understand Deep Q-Network (DQN).	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 11
12	-TF-Agents library. -Installation of TF-Agents. -Training an agent to play Breakout. -TF-Agents environment and specifications. -Environment wrappers and Atari preprocessing.	"Week 12: The TF-Agents Library" notes. Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly Media, Inc.	-Basic knowledge of TF-Agents library. -Be able to install of TF-Agents. -Understand the code on how to train an agent to play Breakout. -Understand the TF-Agents environment and specifications. -Understand environment wrappers and Atari preprocessing.	Lecturing, Demos, Practice Problems, Individual work outside class.		Week 12
13	-Training architecture for TF-Agents for the Breakout arcade game. -The Deep Q-Network (DQN) for Breakout. -The DQN agent for Breakout. -Replay buffer and its observer. -Training metrics for	"Week 13: Training An Agent How To Play Breakout" notes. Chapter 18, Géron, Aurélien. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, SecondEdition, 2019, Published by O'Reilly	-Understand the training architecture for TF-Agents for the Breakout arcade game. -Understand the Deep Q-Network (DQN) for Breakout. -Understand the DQN agent for Breakout. -Understand the replay buffer and its observer. -Understand the training metrics for Breakout. -Understand the collect driver for Breakout. -Understand how to create the dataset and training for Breakout.	Lecturing, Demos, Practice Problems, Individual work outside class.	Reinforcement Learning (15%)	Friday of Week 13 by 11:59pm EST

Week	Topics	Readings/Materials	Weekly Learning Outcome(s)	Instructional Strategies	Evaluation Name and Weight	Evaluation Date
	Breakout. -The collect driver for Breakout. -Creating the dataset and training Breakout.	Media, Inc.				
14	N/A	N/A	N/A	N/A		Week 14