Western University

Faculty of Engineering

Artificial Intelligence Systems Engineering Program

AISE4010/A – DEEP LEARNING FOR TIME SERIES DATA

Course Outline Fall 2025

COURSE DESCRIPTION:

This course introduces deep learning models for time series data. In this course, the students will become familiar with sequence models and their engineering applications. The students are introduced to Recurrent Neural Networks (RNNs) and commonly - used variants such as Gated Recurrent Units (GRUs) and Long Short-Term Memory (LSTMs). In addition, the course introduces transformer architectures and their engineering applications. Students will get hands- on experience with Deep Learning from a series of practical engineering case-studies.

ACADEMIC CALENDAR: AISE 4010/A

https://www.westerncalendar.uwo.ca/Courses.cfm?CourseAcadCalendarID=MAIN_030389_1&SelectedCalendar=Live&ArchiveID=

PRE OR COREQUISITES:

Data Science 3000A/B, AISE 3010A/B.

ANTIREQUISITES: N/A

CEAB ACADEMIC UNITS:

Engineering Science 65%, Engineering Design 35%.

CONTACT HOURS:

Lectures occur weekly. Laboratory sessions occur weekly.

LECTURE:	3 hours weekly
LAB:	2 lab and tutorial hrs /12 times during the term

All lectures and labs are in-person.

RECOMMENDED/ REQUIRED SOFTWARE: Python

RECOMMENDED RESOURCES/REFERENCES:

- Gharehbaghi, A. (2023), Deep Learning in Time Series Analysis. CRC Press.
- Zhang, A., et al., Dive into Deep Learning, Cambridge University press. [Available at: https://d2l.ai/d2l-en.pdf]
- Goodfellow,I., Bengio, Y., Courville, A. (2016), Deep Learning. MIT press.

- Chollet, F., (2021), Deep Learning with Python, Manning Publications.
- Brownlee, J. (2018), Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs and LSTMs in Python, [Available at: https://books.google.ca/books?id=o5qnDwAAQBAJ&printsec=frontcover&source=gbs_ge_summary_r&cad=0#v=onepage&q&f=false]
- Nielsen, A. (2020), Practical Time Series Analysis: Prediction with Statistics and Machine Learning, O'Reilly.
- Cerqueira, V. and Roque, (2024), L. Deep Learning for Time Series Cookbook: Use PyTorch and Python Recipes for Forecasting, Classification, and Anomaly Detection. Packt Publishing.
- Gridin, I. (2021), Time Series Forecasting using Deep Learning: Combining PyTorch, RNN, TCN, and Deep Neural Network Models to Provide Production-Ready Prediction Solutions. BPB Publications.

Other Required References: Students must check OWL BRIGHTSPACE (https://westernu.brightspace.com/d2l/home/) on a regular basis for news and updates. This is the primary method by which information will be disseminated to all students in the class. Students are responsible for checking OWL BRIGHTSPACE on a regular basis.

GENERAL LEARNING OBJECTIVES (CEAB GRADUATE ATTRIBUTES)

Knowledge Base	A	Engineering Tools	A	Impact on Society	
Problem Analysis	D	Individual & Teamwork		Ethics and Equity	D
Investigation		Communication		Economics and Project Mgmt.	
Design	A	Professionalism		Life-Long Learning	

Notation: x represents the content level code as defined by the CEAB. blank = not applicable; I = introduced (introductory); D = developed (intermediate) and A = applied (advanced).

Rating: I – The instructor will introduce the topic at the level required. It is not necessary for the student to have seen the material before. D – There may be a reminder or review, but the student is expected to have seen and been tested on the material before taking the course. A – It is expected that the student can apply the knowledge without prompting (e. g. no review).

COURSE MATERIALS: Weekly content and guides for the laboratories will be available on the course OWL site.

COURSE TOPICS AND SPECIFIC LEARNING OUTCOMES:

The following table summarizes the course learning outcomes along with CEAB GAIs where the GAIs in bold indicate ones to be measured and reported annually.

COURSE TOPICS AND SPECIFIC LEARNING OUTCOMES	(CAEB) Graduate Attribute
1. Time Series Overview, Pre-processing, Visualisation, and Statistical and Machine Learning Analysis At the end of this section, students will be able to: a. Define time series types, characteristics, trend and seasonality, and challenges. b. Define preprocessing techniques for handling missing data, normalization, noise reduction and applying rolling window both in words and mathematically. c. Understanding and identifying self-correlation d. Define different time series analysis (forecasting, classification, and anomaly detection). e. Statistical models for time series (ARIMA, SARIMA), their advantages and disadvantages. f. Feature engineering for time series and define concepts such as lag features, rolling statistics, and time-based features both in words and mathematically g. Machine learning for time series analysis h. Implement these methods in Python.	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
2. Deep Learning At the end of this section, students will be able to: a. Define in words and mathematically deep learning and its differences versus machine learning and statistical techniques. b. Define in words and mathematically deep neural networks, their learning process and categories c. Define in words and mathematically data splitting techniques (for different types of time series), optimization and loss functions, validation, and evaluation of deep learning models. d. Implement model training/validation/test e. Define in words and mathematically the signs of overfitting and methodologies to address that (regularization, dropout). f. Define in words and mathematically the signs of bias and methodologies to address that. g. List various applications of these models.	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**

3. Convolutional Neural Networks (CNNs) for Time Series At the end of this section, students will be able to: a. Define in words and mathematically the 1D CNN and TCN architectures b. Analyze time series using these architectures (Classification and Forecasting) c. Implement these architectures in test/validation/test steps d. Define their pros/cons versus other deep learning techniques for time series applications	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
4. Recurrent Neural Networks (RNNs) At the end of this section, students will be able to: a. Define in words and mathematically the RNN model architecture and its variants (LSTM, GRU) b. Analyze time series using these architectures c. Implement these architectures in test/validation/test steps d. Define their pros/cons versus other deep learning techniques for time series applications	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
5. Transformers At the end of this section, students will be able to: a. Define in words and mathematically concept of attention mechanism in deep learning b. Define in words and mathematically transformer architectures c. Define in words and mathematically applications of attention in time series forecasting d. Analyze time series using these architectures e. Implement these architectures in test/validation/test steps f. Define their pros/cons versus other deep learning techniques for time series applications.	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
6. Autoencoders for Anomaly Detection At the end of this section, students will be able to: a. Define in words and mathematically time series anomalies b. Define in words and mathematically autoencoder structure c. Define in words and mathematically reconstruct time series and detect anomalies d. Implement these algorithms in practical examples in various domains	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**

7. Probabilistic Forecasting, Uncertainty, Fairness, and Hyperparameter Optimization At the end of this section, students will be able to: a. Explain probabilistic forecasting methods and distinguish aleatoric vs epistemic uncertainty. b. Implement models that provide uncertainty estimates. c. Identify fairness issues in time series and apply mitigation strategies. d. Compare and apply hyperparameter optimization methods (grid search, Bayesian, Hyperband, ASHA). e. Evaluate trade-offs between accuracy and computational cost in real-world applications.	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
8. Advanced Representations and Hybrid Approaches At the end of this section, students will be able to: a. Transform time series into alternative domains (spectrograms, Gramian fields, recurrence plots) and apply CNNs. b. Describe and compare state space models (e.g., HMMs and Kalman filters) with deep models. c. Combine classical models with deep learning for noisy or irregular data. d. Assess advantages and limitations of hybrid and representation-based approaches.	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**

EVALUATION:

Assessment Type	**Weight**	**CEAB GAs ASSESSED**
Laboratory Assignments (tentatively 3 assignments)	15%	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
Group Project	40%	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**
Participation Quizzes	5%	EE4, D1, **D2, KB3, KB4, PA2, PA3, ET2**
Final Examination	40%	EE4, D1, **D2, D3, KB3, KB4, PA2, PA3, ET2, ET3**

Marks will be assigned on the basis of the method of analysis and presentation, correctness of solution, clarity and neatness.

To obtain a passing grade in the course, a mark of 50% or more must be achieved on the final examination as well as on the laboratory assignments. A final examination or laboratory assignment mark < 50% will result in a final course grade of 48% or less.

COURSE POLICIES:

Laboratory Assignments:

There will be multiple group assignments. All assignments have equal weights. Assignments are programming-based. Python programming language and various Python-based machine learning packages will be used in this course. Assignments will be released on OWL BRIGHTSPACE with due dates. Assignments must be demonstrated by all group members in the lab and will be graded by the TA/Instructor during lab hours. **The laboratory assignment with the lowest grade will be dropped for every student. Therefore, no late laboratory assignments will be accepted, and special accommodation for missed laboratory assignments will be denied**. A grade of 50% or higher on the laboratory assignments is required to pass the course.

Quizzes: Each student will need a laptop to complete the **in-person** quizzes. Participation pop quizzes will be conducted during randomly selected classes throughout the semester. A quiz may include a concept as well as a practical component. **The lowest quiz mark will be dropped for every student. Therefore, special accommodation for missed quizzes will be denied**.

Group Project:

Groups must consist of **3 or 4 students in the same section of the course.**

There are four project deliverables. The weightings below are relative to the overall final grade.

1. List of group members and identification of the dataset to work with (Mid-September; 0%)
2. Progress Report (1 page) (Mid-November; 8%)
3. Project Presentation (The week before the end of class ~ Dec 3)
   - Final Group Presentation (an 8-minute presentation in the lab; 15%)
4. Code Demo in the lab + GitHub repository submission (1% + 1%; The week before the end of class ~ Dec 3)
5. Project Report and Code (Dec 9)
   - Final Report (4 pages + optional appendix if desired) (15%)

Project Presentation: All project group members must present the project's slides in the lab and will be graded by the TA/Instructor during lab hours. Students who miss the lab for the presentation will be graded 0 for the Project Presentation.

Project Code: The project code must be written in Python. The GitHub repository of the project's code must be submitted and shared with the course Instructor and TAs, and the code must be demonstrated by all project group members in the lab and will be graded by the TA/Instructor during lab hours. Students who miss the lab for the demo will be graded 0 for the Code Demo.

Final Examination: The final exam date will be scheduled by the Office of the Registrar. The final exam will be in-person, in "**closed-book**" structures, and cover concepts from the entire course. Each student may bring one (1) 8.5" by 11" or A4 sheet of paper into the final exam as a "cheat sheet". A grade of 50% or higher on the final examination is required to pass the course.

Late Submission Policy: No late submissions will be accepted; late submissions will be graded with 0%.

Use of English: In accordance with Senate and Faculty Policy, students may be penalized up to 10% of the marks on all assignments, tests, and examinations for improper use of English. Additionally, poorly written work (except examinations) may be returned without grading. If resubmission of the work is permitted, it may be graded with marks deducted for poor English and/or late submission.

Attendance: Attendance is mandatory for all lectures and labs. Any student who, in the opinion of the instructor, is absent too frequently from class, laboratory, or tutorial periods will be reported to the Dean (after due warning has been given). On the recommendation of the department, and with the permission of the Dean, the student will be debarred from taking the regular final examination in the course.