Western Science

Computer Science

Course Outline
CS 9840: Probabilistic Generative Artificial Intelligence

1 Course Summary

Course information The following are the key details about the course:

- **Course Title:** Probabilistic Generative Artificial Intelligence
- **Course Code:** CS 9840
- **Academic Term:** Winter 2026
- **Lectures:** Tuesdays, 10 AM–12 PM and Thursdays, 1–2 PM; Room MC 320
- **Laboratory:** There will be no laboratory
- **Course website:** https://probgenai.kristia.de and OWL Brightspace

Course summary From chatbots to image generators, Generative AI (GenAI) is reshaping the world. What algorithms and probabilistic principles power this revolution? This course introduces the theoretical and algorithmic foundations of GenAI, covering large language models (LLMs) and diffusion models. We will also explore applications in AI for Science, such as drug discovery and weather forecasting. Because GenAI is inherently probabilistic, prior knowledge of probability/statistics, machine learning, and linear algebra will be helpful.

Topics This course will cover the following topics:

- **Latent-space models:** deterministic autoencoders, variational autoencoders
- **Implicit and explicit likelihood models:** generative adversarial networks, normalizing flows
- **Diffusion models:** forward & backward diffusion processes, score matching, guided diffusion
- **LLMs:** attention, transformers, finetuning, alignment, decoding, reasoning
- **Applications:** latent-space Bayesian optimization, climate modeling, de novo molecular discovery

Prerequisites There are no formal prerequisites for this course. However, students are expected to know about the following topics:

- **Probability & statistics:** sum, product, and Bayes' rules, probability distributions (especially Gaussian), maximum likelihood.
- **Linear algebra:** vectors and matrices; matrices/vectors products, matrix inversion,
- **Multivariable calculus:** partial derivatives, gradients, Jacobians, Hessians, multivariable integration.
- **Machine learning:** regression, classification, neural networks, gradient descent, familiarity with deep learning libraries (code examples are in PyTorch).

2 Instructor Information

- **Instructor:** Dr. Agustinus Kristiadi
- **Email:** akristi@uwo.ca
- **Office:** MC 362
- **Website** https://agustinus.kristia.de
- **Office Hours:** Thursdays, 2:00–3:30 PM (in-person)

3 Course Syllabus

Generative AI (GenAI) is crucial for today's and future AI landscape, from personal assistants to scientific discoveries. The key theoretical and algorithmic aspects of GenAI are built upon probabilistic principles. Hence, to develop and harness GenAI well, one must understand it deeply from a probabilistic point of view.

This course covers a comprehensive set of topics, crucial for gaining intuition and understanding of the current state-of-the-art GenAI models, including large language models (LLMs) and diffusion models. Moreover, a wide range of applications of GenAI will be discussed, with a focus on AI for Science, such as drug discovery and climate modeling.

On successful completion of this course, students will be able to:

- Demonstrate a solid understanding of various foundational, probabilistic principles used by recent GenAI models.
- Be comfortable in reading and understanding the latest GenAI research papers.
- Implement various GenAI models using standard deep learning libraries.
- Utilize GenAI models in various applications, especially for scientific decision-making.
- Write clear and stylistic scientific reports about GenAI.
- Create clear and engaging presentations on technical GenAI topics.

Important academic dates

- **Classes Begin:** January 5, 2026.
- **Winter Study Break/Reading Week:** February 14–22, 2026.
- **Holiday:** February 16, 2026 is Family Day; Western will be closed.
- **Classes End:** April 9, 2026.
- **Study Days (not applicable to this course):** April 10—11, 2026.
- **Final Examinations (not applicable to this course):** April 12-–30, 2026.

Planned course schedule

This course is a 12-week course, with a two-hour lecture session on Tuesdays, 10 AM–12 PM, and a one-hour tutorial session on Thursdays, 1–2 PM. The tutorial sessions will be used for reviewing the prerequisite knowledge (probability theory, linear algebra, and calculus), homework discussions, questions and answers, and tutorials on selected topics. Please refer to Table 1 for the planned weekly course schedule and topics.

4 Course Materials

There is no official textbook for this course. However, the following textbooks are highly recommended:

- Bishop, Christopher M., "Deep Learning: Foundations and Concepts", Springer, 2024.
  - Freely available online at https://www.bishopbook.com.

Table 1: Weekly course schedule and topics.

Week	Lecture (Tue)	Tutorial (Thu)	Resources
1	Introduction & review of fundamentals	Review of fundamentals	Bishop 2024, Chapters 2 & 3; D2L, Chapter 2
2	Variational autoencoders	Group formation deadline & review of fundamentals	Bishop 2024, Chapter 19; D2L, Chapter 2
3	Explicit and implicit models	Homework 1	Bishop 2024, Chapters 17 & 18
4	Application: *De novo* drug discovery	Homework 2	Gómez-Bombarelli, et al., 2018
5	Diffusion Models I: Introduction	Project info session	Bishop 2024, Chapter 20
6	Diffusion Models II: Advanced Topics	Project proposal deadline and Project Q&A	Bishop 2024, Chapter 20
7	Application: Climate modeling	Homework 3	
8	LLMs I: Architectures	Stylistic scientific writing 101	Bishop 2024, Chapter 12; D2L Chapter 11
9	LLMs II: Decoding and reasoning	Homework 4	Bishop 2024, Sec. 12.3.2; D2L Sec. 10.8; Lightman et al., 2023
10	LLMs III: Finetuning and alignment	Selected topics in GenAI	Bishop 2024, Chapter 12; Steinnon et al, NeurIPS 2020; Hu et al., ICLR 2022
11	Project presentations	Project presentations	
12	Project presentations	Project presentations	

- Aston Zhang, Zachary Lipton, Mu Li, and Alex Smola. "Dive into Deep Learning (D2L)".
  - Freely available online at https://d2l.ai.

Additionally, results from the following research papers will be discussed in the course:

- Gómez-Bombarelli, et al., "Automatic chemical design using a data-driven continuous representation of molecules", ACS Central Science 4(2), 2018.
- Price, et al., "Probabilistic weather forecasting with machine learning", Nature 637(8044), 2025.
- Stiennon, et al., "Learning to summarize with human feedback", NeurIPS, 2020.
- Lightman, et al., "Let's verify step by step", ICLR, 2023.
- Hu, et al., "LoRA: Low-rank adaptation of large language models", ICLR, 2022.

Lecture notes will be posted on both the external course website (https://probgenai.kristia.de) and OWL Brightspace (https://westernu.brightspace.com). Homework assignment sheets will be available on OWL Brightspace.

Students should check OWL Brightspace on a regular basis for news and updates. If students need assistance with the course OWL site, they can seek support on the OWL Brightspace Help page. Alternatively, they can contact the Western Technology Services (WTS) Helpdesk. They can be contacted by phone at 519-661-3800 or ext. 83800.

5 Methods of Evaluation

The grading scheme for this course is presented in Table 2.

Table 2: Course evaluation components.

Type	Weight	Count	Sub-Total
Homework	10%	4	40%
Project Proposal	15%	1	15%
Project Final Presentation	20%	1	20%
Project Final Report	25%	1	25%
Total			**100%**

Students will be asked to form a group of 2–3 members in week one. The groups will be fixed for the entire duration of the course—no group switching will be allowed. All homework and project submissions must be done in groups. Please note that detailed guidelines on homework submission, project proposal, final presentation, and final report will be discussed in class.

The following are tentative deadlines for the evaluation components. Please note that all deadlines, except for Project Presentation Slides and Project Final Report, are on the night before the tutorial session on Thursdays, i.e., on Wednesdays, 11:59 PM EST.

- **Group formation deadline:** Due Wednesday, January 14, 2026, 11:59 PM.
- **Homework 1:** Due Wednesday, January 21, 2026, 11:59 PM
- **Homework 2:** Due Wednesday, January 28, 2026, 11:59 PM.
- **Project Proposal:** Due Wednesday, February 11, 2026, 11:59 PM.
- **Homework 3:** Due Wednesday, February 25, 2026, 11:59 PM.
- **Homework 4:** Due Wednesday, March 11, 2026, 11:59 PM.
- **Project Presentation Slides:** Due Monday, March 23, 2026, 11:59 PM.
- **Project Final Report:** Due Thursday, April 9, 2026, 11:59 PM.

Submission guidelines All submissions must be done electronically via OWL Brightspace. They must be written in LaTeX and must be uploaded as PDF files. This includes the presentation slides. LaTeX templates will be provided.

Grading scheme Each homework consists of several questions, each graded in a pass-fail or sufficient-insufficient manner. The total score of each homework is the fraction of passed/sufficient answers. For the Project Proposal, Final Presentation, and Final Report, they will be evaluated based on (i) Problem Formulation, (ii) Analysis, and (iii) Clarity. Each rubric will be graded on a scale of A (corresponds to 100% mark) to E (corresponds to 20% mark), then they will be averaged with equal weighting. The breakdown is in Table 3.

Absences and late submissions Attendance at lectures and tutorials is strongly encouraged, though it will not be counted. Since both the homework and project are intended to be done in a group, the submission deadlines are strict: Late submission will not be accepted, except for valid reasons as per the university's academic consideration policies.

LLM usage policy LLMs such as ChatGPT may be used only for fine-tuning the report (e.g., grammar, formatting, or clarity). They must not be used for content generation or to answer questions in homework.

Table 3: Project evaluation.

	Problem Formulation	Analysis	Clarity
Proposal	Motivation, potential impacts, feasibility	Quality of hypothesis and literature review, feasibility of plan	Styles and formatting of the paper; presence of errors/typos
Presentation	Communication of motivation, examples	Presentation depth and Q&A performance	Slides quality and their styles
Final Report	Motivation, potential impacts	Methods/analysis novelty and quality	Styles and formatting of the paper; presence of errors/typos

6 Accommodation and Accessibility

Religious Accommodation When a course requirement conflicts with a religious holiday that requires an absence from the University or restricts participation in certain activities, students should request accommodation in writing at least two weeks before the holiday. Requests should be submitted to the course instructor and/or the Academic Counselling Office of the student's Faculty of Registration. Students are encouraged to check out the diversity calendar posted on the university's EDID website for the recognized religious holidays: https://www.edi.uwo.ca. More information can be found here: https://www.uwo.ca/univsec/pdf/academic_policies/appeals/accommodation_religious.pdf.

Accommodation Policies Students with disabilities are encouraged to contact Accessible Education, which provides accommodation recommendations based on medical documentation or psychological and cognitive assessments. The University's Policy on Academic Accommodation for Students with Disabilities can be found here: https://academicsupport.uwo.ca/accessible_education/graduate_students.html. More details about graduate academic consideration policies can be found here:
https://www.uwo.ca/univsec/academic_policies/grad-accommodation-consideration.html.

7 Academic Policies

The website for Registrar Services is https://www.registrar.uwo.ca. In accordance with policy, https://www.uwo.ca/univsec/pdf/policies_procedures/section1/mapp113.pdf, the centrally administered e-mail account provided to students will be considered the individual's official university e-mail address. It is the responsibility of the account holder to ensure that emails received from the University at their official university address are attended to in a timely manner.

Scholastic Offences are taken seriously. Students are directed to read the appropriate policy, specifically, the definition of what constitutes a Scholastic Offence, at the following: https://www.uwo.ca/univsec/pdf/academic_policies/appeals/scholastic_discipline_grad.pdf.

All required reports/papers may be subject to submission for textual similarity review using the commercial plagiarism detection software under license to the University for the detection of plagiarism. All papers submitted for such checking will be included as source documents in its database solely for the purpose of detecting plagiarism in future submissions. Use of this service is subject to the licensing agreement currently between the University of Western Ontario and Turnitin (https://www.turnitin.com).

Remote proctoring software may be used in this course in the event of a health-related lockdown. By enrolling in this course, you consent to the use of this software and acknowledge that you may be required to provide personal information, including biometric data, and that your session will be recorded. Comple-

tion of this course requires a reliable internet connection and a device that meets the software's technical requirements. More information, including technical specifications, is available on Western's remote proctoring website at https://remoteproctoring.uwo.ca.

8 Support Services

Graduate students should direct any questions regarding course registration, academic consideration, leaves of absence, or appeals to their Graduate Program Office and/or the School of Graduate and Postdoctoral Studies (SGPS) at https://grad.uwo.ca.

Students experiencing emotional or mental distress are encouraged to reach out to Mental Health@Western (https://uwo.ca/health) for a complete list of available resources and support options.

Western is committed to preventing gender-based and sexual violence and providing compassionate support to anyone affected by such experiences. If you have experienced sexual or gender-based violence, whether recently or in the past, you can find information about support services for survivors, including emergency contacts, at: https://www.uwo.ca/health/student_support/survivor_support/get-help.html. To connect with a case manager or set up an appointment, please contact support@uwo.ca.

Please contact the course instructor if you require lecture or printed materials in an alternate format, or if other arrangements could help make this course more accessible to you. You may also contact Accessible Education for additional support and resources at http://academicsupport.uwo.ca/accessible_education/index.html, if you have any questions.

Learning-skills counsellors at the Student Development Centre (https://learning.uwo.ca) are available to help you improve your academic and learning skills. They offer presentations on topics for time management, exam preparation/writing, reading strategies, and more. Individual support is available during the Fall/Winter terms through the drop-in Learning Help Centre, and year-round through individual counselling.

Western University is committed to fostering a thriving campus community through both virtual and in-person learning experiences. Students are encouraged to visit the Student Experience website to access resources that support academic success, well-being, and campus life.

Additional student-run support services are offered by the USC: https://westernusc.ca/services.