



Institute of Technology Entrepreneurship and Commercialization

TIMG 5103–P, Prompt Engineering in Business [0.5 credits]

Fall 2025

Time and Place

Thursdays, 6:05 pm to 8:55 pm, Nicol Building NI-4050

September 4th, 2025 - December 4th, 2025

Course offered online at Brightspace (Prompt Engineering in Business – TIMG 5103P – Fall 2025).

Updates to this outline (version 1.0) will be made as necessary.

Instructor

Daniel Cardenas

Course materials

Access to online course sessions, course materials and recorded videos will be provided through the new CU Brightspace system: <https://carleton.ca/brightspace/>. To access Brightspace, you should use your CU credentials and select the TIMG 5103–P, Prompt Engineering in Business (LEC), Fall 2025.

Office hours

The instructor could be contacted via email. Email is the preferred mode of communication because it provides a record of the content exchanged. Depending on availability, the instructor will be available during class and for online meetings by appointment.

Target audience

This course best suits students in the Master of Advanced Business Analytics or the Master of Digital Transformation and Entrepreneurship. Students in other programs are welcome to attend this course, depending on space availability. However, all students must meet the high academic standards of the TIM program.

Calendar description

Prompt Engineering in Business – TIMG 5103–P

Generative AI applications. Artificial Intelligence (AI), Machine Learning (ML), and Large Language Models (LLMs) concepts. Prompt Engineering best practices. Retrieval-Augmented Generation (RAG) and Fine-Tuning optimization techniques. Agents and agentic frameworks. LLM benchmarks and evaluation.

Prerequisites: None.

Prompt Engineering in Business course description

The course examines the principles and applications of Generative Artificial Intelligence (Gen AI) in the context of business enterprises, small and large, and how to understand the business value of Large Language Models (LLMs). Participants will learn the concepts that support most LLMs and how to design clear and impactful prompts that guide LLMs to generate accurate and relevant responses. Topics covered include prompt optimization, API integrations, Retrieval-Augmented Generation (RAG) and Fine-Tuning optimization techniques. The second part of the course covers LLM-based Agents and agentic frameworks, as well as LLM benchmarks and evaluation methods. The course emphasizes the importance of prompt engineering in enhancing user experiences, streamlining workflows, and making informed business decisions. Real-world case studies are examined to highlight the impact and risks of LLMs on customer engagement, data collection, and personalized interactions.

Objectives / Learning Outcomes

Students will be expected to:

- Understand the capabilities and potential business value of using Generative AI from exemplary real-life cases of LLM applications in specific business domains.
- Describe where Generative AI and Large Language Models fall within the large Artificial Intelligence domain.
- Acquire an understanding of the foundational models that power most current LLMs.
- Learn how to create optimized prompts and use them with API integration mechanisms.
- Get familiar with the complexities of creating custom LLM-based agents.
- Propose an integrated business system solution that leverages the capabilities of Generative AI.
- Understand the need for basic principles to guide the responsible and risk-mitigated use of LLMs.

Class sessions

The course sessions will be broadcast live and recorded on Zoom. This is an alternative mode of attending class if you are travelling, sick, or otherwise unable to attend class. Since this course includes interactive exercises using Google Colab or a similar platform, it is best experienced in person. Please see the detailed tutorials on participating in an online classroom using Zoom at <https://carleton.ca/online/online-learning-resources/zoom-guide-for-students/>.

Paul Menton Centre

Students with disabilities requiring academic accommodations in this course are encouraged to contact a coordinator at the Paul Menton Centre (PMC) for Students with Disabilities to complete the necessary letters of accommodation. After registering with PMC, schedule an appointment to meet with your instructor and discuss your needs at least two weeks prior to needing accommodations for assignments or presentations. This is necessary to ensure sufficient time to make the necessary arrangements.

Course assignments

The course comprises four group-based assignments due in weeks 5, 6, 10, and 14, plus two individual exams (weeks 6 and 12). Assignments #2, #3, and #4 are summative towards the term project and comprise a deliverable (document and/or proof of concept) and a presentation. Due to its potentially complex nature, the term project should not be completed individually. What you submit for the term project must reflect your work, so each document must include the names of all participants and each person's contribution in percentage. Grading is platform- and model-agnostic, so individual platform-specific attributes will not be evaluated. However, students will receive additional credit for effectively leveraging platform-specific features to overcome challenges in their solution. This is not a programming course, and no coding knowledge is required. Nevertheless, basic coding workshops could be offered to assist students without a coding background, ensuring everyone can engage in the technical aspects of the project. Groups will have opportunities for feedback from the instructor during the scheduled presentations to ensure they are on the right track and can address any major issues early on the project.

Assignment # 1: A case study of a Generative AI application (10%, Week # 5: Sunday, October 05)

- Find a solution that uses LLMs/Agents as a core component.
 - Describe the context in which the problem occurs, identifying the main pain points.
 - Perform a literature review to assess the feasibility of applying LLMs in the problem context.
 - Propose metrics to assess the business value obtained from implementing the Generative AI solution to that problem.
 - Explain the solution architecture: components and proposed process/workflow.
 - Identify the models, tools, and techniques used in the solution.
 - Describe the results obtained and identify the potential limitations and risks of the solution.
- Critical thinking skills: Synthesize.

Assignment # 2: Term Project – Problem definition & solution proposal (20%, Week # 6: Sunday, Oct 12)

- Problem Definition (40 %)
 - Environment: Describe the context where the problem occurs.
 - Current process and data: Describe the workflow and data used at each stage.
 - Identification of pain points: Describe and quantify (measure) the main problems, clearly stating why AI (or alternatives) could help.
- Solution Proposal (60 %)
 - Perform a literature review to assess the feasibility of applying LLMs in the problem domain.
 - General description: How LLMs can help overcome the previously identified main pain points.
 - Proposed process: Describe the new workflow, emphasizing the differences between the baseline and the new process.
 - Data requirements: Source, quantity, frequency, currency and limitations.
 - Component identification: High-level description of each module's purpose. Identify each component's input(s) and output(s).
 - Knowns/Unknowns: Identify what is known and what is not known about the proposed solution.
- Critical thinking skills: Synthesize.

Assignment # 3: Term Project – Solution design (25%, Week # 10: Sunday, Nov 09)

- Solution Design (60 %)
 - Solution Architecture:
 - Provide static and dynamic diagrams showing system components, data flow, and processes.
 - Ensure platform-specific features are highlighted and analyze these features' advantages.
 - Data Management:
 - Describe the data preparation process, offering detailed steps for cleaning, structuring, and aligning data with each system's component.
 - Suggest alternatives for data limitations (scalability, storage, accessibility, bias, etc.).
 - Risk and Mitigation Plan:
 - Identify principal risks (data accuracy, AI outputs, scalability, ethical issues, potential bias, etc.) directly related to the components used in the solution and propose mitigation strategies.
 - Rank risks based on impact and likelihood.
 - Validation Proposal:
 - Perform a literature review to identify potential ways to evaluate the performance of the proposed solution.
 - Propose how the system's outputs will be validated (performance, reliability, etc.).
 - Present examples of potential good and bad outputs from the system that can be used to assess its performance.
- Prototype (40 %)
 - Data collection, cleaning and preparation: Describe the steps required to prepare the data.
 - LLMs: Platform, model(s), prompt definitions, expected outcomes. A basic prototype or mockup should demonstrate how the proposed solution could work.
- Critical thinking skills: Recommend.

Assignment # 4: Term Project – Solution & Proof of Concept (30%, Week # 14: Friday, Dec 05)

- Solution Design (40 %)
 - Revisit and improve the sections outlined in Assignment #3 based on feedback.
 - Evaluate the solution's problems and limitations found during development.
 - Compare the expected outputs defined in Assignment # 3 with those obtained from the PoC.
 - Identify areas of improvement and provide recommendations.
- Presentation & Proof of Concept (60 %)
 - Solution architecture: main components and process/workflow used in the PoC.
 - Describe the data used in the PoC: quantity, attributes, source (real vs. synthetic).
 - Describe the PoC use case: roles, scenarios, assumptions.
 - Highlight the techniques being applied in the PoC: Prompts Engineering, RAG, Fine Tuning, and Agents.
 - Execute the Proof of Concept.
- Critical thinking skills: Generate and execute.

Methodology based on Plummer (2019). A short guide to building your team's critical thinking skills. Harvard Business Review.

<https://hbr.org/2019/10/a-short-guide-to-building-your-teams-critical-thinking-skills>.

- *Synthesize*: Students should be able to synthesize the information, drawing connections between Generative AI capabilities and business objectives. This will be done through discussions or written assignments requiring students to analyze and interpret the results of Generative AI systems.
- *Recommend*: Students should be able to recommend appropriate uses of LLMs. This could involve strategic decision-making where students must consider various factors, such as ethical implications, business goals, and technological constraints, to make informed proposals.
- *Generate*: Students should be able to generate new ideas or strategies based on their understanding of Generative AI's potential. This will be part of the capstone project, where students design innovative AI applications or propose improvements to existing systems.
- *Execute*: Students should be able to execute real-world tasks using Generative AI tools. This will involve practical labs or simulations to interact with LLM models to understand their capabilities and limitations.

Student evaluation and assignment grading

The final grade will be assigned using the following mark allocation:

Assignment/Exam	Description	Date/Week	Grade %
Assignment #1: Case study of a Gen-AI application.	<ul style="list-style-type: none"> Analyze a Generative AI solution: Problem, solution architecture, tools/models, limitations, and results. 1500 – 2500-word document. 	Sunday, Oct 05 (Week # 5)	10
Exam #1.	<ul style="list-style-type: none"> Concepts weeks 1-5 Individual, 20 minutes. 	Thursday, Oct 09 (Week # 6)	5
Assignment #2: Term Project – Problem definition.	<ul style="list-style-type: none"> Context and problem description High-level solution proposal 1000 – 2000-word document. 	Sunday, Oct 12 (Week # 6)	20
Assignment #3: Term Project – Solution design.	<ul style="list-style-type: none"> Architecture design, requirements, and limitations. 2000 – 3000-word document. 	Sunday, Nov 09 (Week # 10)	25
Exam #2.	<ul style="list-style-type: none"> Concepts weeks 1-12 Individual, 90 minutes. 	Thursday, Nov 20 (Week # 12)	10
Assignment #4: Term Project – Solution & PoC.	<ul style="list-style-type: none"> Proof of Concept (presentations on weeks 13 & 14). 2500 – 3500-word document. 	Friday, Dec 5 (Week # 14)	30
			100

Plagiarism

Plagiarism, including copying and handing in someone else's work for credit, is a serious instructional offence that will not be tolerated. Please refer to the section on instructional offences in the Graduate Calendar for additional information. A case of plagiarism will be referred to the Chair of the department and the Carleton University Ethics Committee. The instructor will not deal with the matter directly. The University has clear processes in place to deal with students suspected of plagiarism.

Group work and freeloaders

Group work is an important component of this course. Students will be randomly assigned to groups. Group conflicts are to be dealt with by the group in a way that is fair, fast and without personal attacks. The instructor does not settle group disputes. The instructor may dissolve a group that is late submitting an assignment. Freeloaders are not welcome anywhere, and this course is no exception: there is zero tolerance for freeloaders. The best way to deal with freeloaders is not to include their names on the group assignments' first page (or workbook's cell). If a student's name does not appear in an assignment submitted by his or her group, the student must submit his or her assignment. Failure to do so will result in the student receiving a zero grade for the assignment.

Class Schedule

Week/ Date	Topic	Assigned readings & details
1/ Thursday, Sep 4	Presentation & Overview - Course structure: Objectives, assignments, important dates - Basic concepts - Business use cases	- LLM-based agent exercise Recommended reading: - Ozdemir, S. (2023). <i>Quick start guide to large language models strategies and best practices for using CHATGPT and other LLMs</i> . Chapter 1 - Hadi, M. U., Al-Tashi, Q., Qureshi, R., Shah, A., Muneer, A., Irfan, M., ... & Shah12, M. LLMs: A Comprehensive Survey of Applications, Challenges, Datasets, Models, Limitations, and Future Prospects. - Zao-Sanders, M., & Ramos, M. (2023). A framework for picking the right generative AI project. <i>Harvard Business Review</i> . - Agrawal A., Gans J., & Goldfarb A. (2022). <i>Power and Prediction</i> . Harvard Business Review Press. Chapter 2. Korinek, A. (2024). LLMs Level Up—Better, Faster, Cheaper: June 2024 Update to Section 3 of “Generative AI for Economic Research: Use Cases and Implications for Economists,” Published in the. <i>Journal of Economic Literature</i> , 61(4), 1-38.
2/ Thursday, Sep 11	Prompt Engineering I - LLM's landscape - Generative AI economic impact - Prompt Engineering basic techniques and frameworks	- Basic prompt engineering exercises. Recommended reading: - Perrault, R., & Clark, J. (2024). Artificial Intelligence Index Report 2024. - Chui, M., Hazan, E., Roberts, R., Singla, A., & Smaje, K. (2023). The economic potential of generative AI. - White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., ... & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with chatgpt. <i>arXiv preprint arXiv:2302.11382</i> . - Ye, Q., Axmed, M., Pryzant, R., & Khani, F. (2023). Prompt engineering a prompt engineer. <i>arXiv preprint arXiv:2311.05661</i> . - Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., ... & Resnik, P. (2024). The Prompt Report: A Systematic Survey of Prompting Techniques. <i>arXiv preprint arXiv:2406.06608</i> . - Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. <i>Advances in neural information processing systems</i> , 35, 22199-22213.
3/ Thursday, Sep 18	Prompt Engineering II - Prompt Engineering advanced techniques - Context Engineering	- Meta prompting exercises and advanced prompts exercises. Recommended reading: - Ozdemir, S. (2023). <i>Quick start guide to large language models strategies and best practices for using CHATGPT and other LLMs</i> . Pearson Education (US). Chapter 3.

Week/ Date	Topic	Assigned readings & details
		<ul style="list-style-type: none"> - Suzgun, M., & Kalai, A. T. (2024). Meta-prompting: Enhancing language models with task-agnostic scaffolding. <i>arXiv preprint arXiv:2401.12954</i>. - Bsharat, S. M., Myrzakhan, A., & Shen, Z. (2023). Principled instructions are all you need for questioning llama-1/2, gpt-3.5/4. <i>arXiv preprint arXiv:2312.16171</i>. - Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in Neural Information Processing Systems</i>, 35, 24824-24837. - Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., ... & Zhou, D. (2022). Self-consistency improves chain of thought reasoning in language models. <i>arXiv preprint arXiv:2203.11171</i>.
4/ Thursday, Sep 25	RAG Fundamentals <ul style="list-style-type: none"> - Foundational Models - Embeddings - Retrieval-Augmented Generation (RAG) basics 	<ul style="list-style-type: none"> - Basic RAG application exercise. Recommended reading: <ul style="list-style-type: none"> - Ozdemir, S. (2023). <i>Quick start guide to large language models strategies and best practices for using CHATGPT and other LLMs</i>. Pearson Education (US). Chapter 2. - Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. <i>Advances in neural information processing systems</i>, 26. - Muennighoff, N., Tazi, N., Magne, L., & Reimers, N. (2022). MTEB: Massive text embedding benchmark. <i>arXiv preprint arXiv:2210.07316</i>.
5/ Thursday, Oct 02	Advanced RAG <ul style="list-style-type: none"> - Advanced RAG applications - Modular RAG - Vector databases 	<ul style="list-style-type: none"> - RAG application using a vector database. - Assignment #1: Case study of a Generative AI application (Sunday, Oct 05) Recommended reading: <ul style="list-style-type: none"> - Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., ... & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. <i>arXiv preprint arXiv:2312.10997</i>.
6/ Thursday, Oct 09	Term Project evaluation	<ul style="list-style-type: none"> - Exam # 1 (content: weeks 1-5) - Assignment #2: Term Project - Problem definition. Group Presentation (all groups). - Document due date: Sunday, October 12
7/ Thursday, Oct 16	LLMs optimization <ul style="list-style-type: none"> - Fine-Tuning - Model Alignment - Reinforcement Learning from Human Feedback (RLHF) - Fundamentals of Agents - Agent architectures 	<ul style="list-style-type: none"> - Fine-tuning a foundational model exercise Recommended reading: <ul style="list-style-type: none"> - Ozdemir, S. (2023). <i>Quick start guide to large language models strategies and best practices for using CHATGPT and other LLMs</i>. Pearson Education (US). Chapter 9. - Kaufmann, T., Weng, P., Bengs, V., & Hüllermeier, E. (2023). A survey of reinforcement learning from human feedback. <i>arXiv preprint arXiv:2312.14925</i>. - Liu, Y., Yao, Y., Ton, J. F., Zhang, X., Cheng, R. G. H., Klovchov, Y., ... & Li, H. (2023). Trustworthy LLMs: a Survey and Guideline for Evaluating Large Language Models' Alignment. <i>arXiv preprint arXiv:2308.05374</i>. - Zhi-Xuan, T., Carroll, M., Franklin, M., & Ashton, H. (2024). Beyond preferences in AI alignment. <i>Philosophical Studies</i>, 1-51. - Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., ... & Hadfield-Menell, D. (2023). Open problems and fundamental limitations of reinforcement learning from human feedback. <i>arXiv preprint arXiv:2307.15217</i>.
8/ Thursday, Oct 23	Fall Break	

Week/ Date	Topic	Assigned readings & details
9/ Thursday, Oct 30	Agent-based applications <ul style="list-style-type: none"> - Agentic frameworks - Multi-agent applications Artificial Intelligence & LLMS <ul style="list-style-type: none"> - Neural Networks - Positional Encoding - Transformer Model 	<ul style="list-style-type: none"> - Agent-based application exercises. Recommended reading: <ul style="list-style-type: none"> - Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. <i>Advances in neural information processing systems</i>, 30. - Kaddour, J., Harris, J., Mozes, M., Bradley, H., Raileanu, R., & McHardy, R. (2023). Challenges and applications of large language models. <i>arXiv preprint arXiv:2307.10169</i>.
10/ Thursday, Nov 06	Term Project evaluation	<ul style="list-style-type: none"> - Assignment #3: Term Project – Solution design. Group Presentation (all groups). - Document due date: Sunday, November 9th.
11/ Thursday, Nov 13	Benchmarking LLMs <ul style="list-style-type: none"> - General-purpose and domain-specific LLM evaluations - RAG evaluations - Generative AI Risks 	Recommended reading: <ul style="list-style-type: none"> - Ooi, K. B., Tan, G. W. H., Al-Emran, M., Al-Sharafi, M. A., Capatina, A., Chakraborty, A., ... & Wong, L. W. (2023). The potential of Generative Artificial Intelligence across disciplines: perspectives and future directions. <i>Journal of Computer Information Systems</i>, 1-32. - Dalvi, F., Hasanain, M., Boughorbel, S., Mousi, B., Abdaljalil, S., Nazar, N., ... & Alam, F. (2023). LLMebench: A Flexible Framework for Accelerating LLMs Benchmarking. <i>arXiv preprint arXiv:2308.04945</i>. - Kaddour, J., Harris, J., Mozes, M., Bradley, H., Raileanu, R., & McHardy, R. (2023). Challenges and applications of large language models. <i>arXiv preprint arXiv:2307.10169</i>.
12/ Thursday, Nov 20	- Model Context Protocol	<ul style="list-style-type: none"> - MCP exercise. - Exam # 2 (content: weeks 1-12)
13/ Thursday, Nov 27	Artificial Intelligence & LLMS <ul style="list-style-type: none"> - Neural Networks - Positional Encoding Term Project Presentations.	- Assignment #4: Term Project – Solution. Group Presentation (part I).
14/ Thursday, Dec 04	Artificial Intelligence & LLMS <ul style="list-style-type: none"> - Transformer Model Term Project Presentations.	<ul style="list-style-type: none"> - Assignment #4: Term Project – Solution. Group Presentation (part II). - Document due date for all groups: Friday, December 5th

Recommended books

Agrawal A., Gans J., & Goldfarb A. (2022). *Power and Prediction*. Harvard Business Review Press.

Caelen, O. & Blete, M.A. (2023). *Developing apps with GPT-4 and ChatGPT: Build intelligent chatbots, content generators, and more*. O'Reilly Associates.

Ozdemir, S. (2023). *Quick start guide to large language models strategies and best practices for using CHATGPT and other LLMs*. Pearson Education (US).

Recommended articles

Bsharat, S. M., Myrzakhan, A., & Shen, Z. (2023). Principled instructions are all you need for questioning llama-1/2, gpt-3.5/4. *arXiv preprint arXiv:2312.16171*.

Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., ... & Hadfield-Menell, D. (2023). Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.

Chui, M., Hazan, E., Roberts, R., Singla, A., & Smaje, K. (2023). The economic potential of generative AI.

Dalvi, F., Hasanain, M., Boughorbel, S., Mousi, B., Abdaljalil, S., Nazar, N., ... & Alam, F. (2023). LLMeBench: A Flexible Framework for Accelerating LLMs Benchmarking. *arXiv preprint arXiv:2308.04945*.

Hadi, M. U., Al-Tashi, Q., Qureshi, R., Shah, A., Muneer, A., Irfan, M., ... & Shah12, M. LLMs: A Comprehensive Survey of Applications, Challenges, Datasets, Models, Limitations, and Future Prospects.

Kaddour, J., Harris, J., Mozes, M., Bradley, H., Raileanu, R., & McHardy, R. (2023). Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*.

Kaufmann, T., Weng, P., Bengs, V., & Hüllermeier, E. (2023). A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*.

Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35, 22199-22213.

Korinek, A. (2024). LLMs Level Up—Better, Faster, Cheaper: June 2024 Update to Section 3 of “Generative AI for Economic Research: Use Cases and Implications for Economists” Published in the *Journal of Economic Literature*, 61(4), 1-38.

Liu, Y., Yao, Y., Ton, J. F., Zhang, X., Cheng, R. G. H., Klochkov, Y., ... & Li, H. (2023). Trustworthy LLMs: A survey and guideline for evaluating large language models' alignment. *arXiv preprint arXiv:2308.05374*.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.

Muennighoff, N., Tazi, N., Magne, L., & Reimers, N. (2022). MTEB: Massive text embedding benchmark. *arXiv preprint arXiv:2210.07316*.

Ooi, K. B., Tan, G. W. H., Al-Emran, M., Al-Sharafi, M. A., Capatina, A., Chakraborty, A., ... & Wong, L. W. (2023). The potential of Generative Artificial Intelligence across disciplines: perspectives and future directions. *Journal of Computer Information Systems*, 1-32.

Perrault, R., & Clark, J. (2024). Artificial Intelligence Index Report 2024.

Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., ... & Resnik, P. (2024). The Prompt Report: A Systematic Survey of Prompting Techniques. *arXiv preprint arXiv:2406.06608*.

Suzgun, M., & Kalai, A. T. (2024). Meta-prompting: Enhancing language models with task-agnostic scaffolding. *arXiv preprint arXiv:2401.12954*.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., ... & Zhou, D. (2022). Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., ... & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.

Zao-Sanders, M., & Ramos, M. (2023). A framework for picking the right generative AI project. *Harvard Business Review*.

Zhi-Xuan, T., Carroll, M., Franklin, M., & Ashton, H. (2024). Beyond preferences in ai alignment. *Philosophical Studies*, 1-51.

Additional resources

- OpenAI
 - Documentation: <https://platform.openai.com/docs/introduction>
 - API Reference: <https://platform.openai.com/docs/api-reference>
- LangChain
 - Introduction: <https://python.langchain.com/docs/introduction/>
 - High-Level Conceptual Guide: <https://langchain-ai.github.io/langgraph/concepts/>
- Hugging Face: <https://huggingface.co/docs>
- Meta Llama: <https://www.llama.com/>
- Video: What are the risks of generative AI? - The Turing Lectures with Mhairi Aitken: <https://www.youtube.com/watch?v=si1jcl7UFqU>
- Google Colab Tutorial: <https://colab.research.google.com/drive/16pBJQePbqkz3QFV54L4NikOn1kwpuRrj>
- Video: What is generative AI and how does it work? – The Turing Lectures with Mirella Lapata: https://www.youtube.com/watch?v=6R7Ym6Vy_I
- The Microsoft Responsible AI Standard: <https://www.microsoft.com/en-us/ai/principles-and-approach>
- Patterns for Building LLM-based Systems & Products: <https://eugeneyan.com/writing/llm-patterns/>
- LLM Benchmarks: <https://livebench.ai/#/>

Appendix: ADDITIONAL INFORMATION

Group work

The Spratt School of Business encourages group assignments in the school for several reasons. They provide you with opportunities to develop and enhance interpersonal, communication, leadership, followership and other group skills. Group assignments are also good for learning integrative skills for putting together a complex task. Your professor may assign one or more group tasks/assignments/projects in this course. Before embarking on a specific problem as a group, it is your responsibility to ensure that the problem is meant to be a group assignment and not an individual one.

In accordance with the Carleton University Graduate Calendar, the letter grades assigned in this course will have the following percentage equivalents:

A+ = 90-100	B+ = 77-79	C+ = 67-69	D+ = 57-59
A = 85-89	B = 73-76	C = 63-66	D = 53-56
A - = 80-84	B - = 70-72	C - = 60-62	D - = 50-52
F = Below 50			

<https://calendar.carleton.ca/grad/gradregulations/administrationoftheregulations/#10>

A grade of B- or better must normally be obtained in each course credited towards the master's degree. A candidate may, with the support of the departmental graduate supervisor/associate chair (graduate affairs) and the approval of the Dean of the Faculty of Graduate and Postdoctoral Affairs, be allowed a grade of C+ in 1.0 credit. Some programs do not permit the C+ option and apply a B- minimum rule.

Academic Regulations

University rules regarding registration, withdrawal, appealing marks, and most anything else you might need to know can be found on the university's website, here:

<https://calendar.carleton.ca/grad/gradregulations/>

Requests for Academic Accommodation

You may need special arrangements to meet your academic obligations during the term. For an accommodation request, the processes are as follows:

Pregnancy obligation

Please contact your instructor with any requests for academic accommodation during the first two weeks of class, or as soon as possible after the need for accommodation is known to exist. For more details, visit the Equity Services website: <https://carleton.ca/equity/>

Religious obligation

Please contact your instructor with any requests for academic accommodation during the first two weeks of class, or as soon as possible after the need for accommodation is known to exist. For more details, visit the Equity Services website: <https://carleton.ca/equity/>

Academic Accommodations for Students with Disabilities

If you have a documented disability requiring academic accommodations in this course, please contact the Paul Menton Centre for Students with Disabilities (PMC) at 613-520-6608 or pmc@carleton.ca for a formal evaluation or contact your PMC coordinator to send your instructor your Letter of Accommodation at the beginning of the term. You must also contact the PMC no later than two weeks before the first in-class scheduled test or exam requiring accommodation (if applicable). After requesting accommodation from PMC, meet with your instructor as soon as possible to ensure accommodation arrangements are made. carleton.ca/pmc

Survivors of Sexual Violence

As a community, Carleton University is committed to maintaining a positive learning, working and living environment where sexual violence will not be tolerated, and its survivors are supported through academic accommodations as per Carleton's Sexual Violence Policy. For more information about the services available at the university and to obtain information about sexual violence and/or support, visit carleton.ca/sexual-violence-support

Accommodation for Student Activities

Carleton University recognizes the substantial benefits, both to the individual student and for the university, that result from a student participating in activities beyond the classroom experience. Reasonable accommodation must be provided to students who compete or perform at the national or international level. Please contact your instructor with any requests for academic accommodation during the first two weeks of class, or as soon as possible after the need for accommodation is known to exist.

<https://students.carleton.ca/course-outline/#accommodation-for-student-activities>

For more information on academic accommodation, please contact the departmental administrator or visit students.carleton.ca/course-outline

Academic Integrity

Violations of academic integrity are a serious academic offence. Violations of academic integrity – presenting another’s ideas, arguments, words or images as your own, using unauthorized material, misrepresentation, fabricating or misrepresenting research data, unauthorized co-operation or collaboration or completing work for another student – weaken the quality of the degree and will not be tolerated. Penalties may include a grade of Failure on the submitted work and/or course; academic probation; a refusal of permission to continue or to register in a specific degree program; suspension from full-time studies; suspension from all studies at Carleton; expulsion from Carleton, amongst others. Students are expected to familiarize themselves with and follow the Carleton University Student Academic Integrity Policy, which is available, along with resources for compliance at <https://carleton.ca/registrar/academic-integrity/>.

Sprott Student Services

The Sprott student services office, located at 710 Dunton Tower, offers academic advising, study skills advising, and overall academic success support. If you are having a difficult time with this course or others, or need some guidance on how to successfully complete your Sprott degree, please drop in any weekday between 8:30am and 4:30pm. Our advisors are happy to discuss grades, course selection, tutoring, and concentrations and will ensure that you get connected with the resources you need to succeed! <https://sprott.carleton.ca/current-students/undergraduate-students/academic-advising/>

Centre for Student Academic Support

The Centre for Student Academic Support (CSAS) is a centralized collection of learning support services designed to help students achieve their goals and improve their learning both inside and outside the classroom. CSAS offers academic assistance with course content, academic writing and skills development. Visit CSAS on the 4th floor of MacOdrum Library or online at <https://carleton.ca/csas/>.

Important Information:

- Students must always retain a hard copy of all submitted work.
- All final grades are subject to the Dean’s approval.

- For us to respond to your emails, we need to see your full name and CU ID, and the email must be written from your valid CARLETON address. Therefore, to respond to your inquiries, please send all email from your Carleton CMail account. If you do not have or have yet to activate this account, you may wish to do so by visiting <https://carleton.ca/its/>.
