Psych 390/650: Artificial Intelligence (AI) Methods for Psychological Science

Meeting Time: Tuesdays/Thursdays, 11:30 a.m. – 12:50 p.m.

Location: PAS 4032

Term: Spring **Year:** 2026

Units: 0.5 **Section:** 2

Instructor

Name	Dr. Trent Cash – feel free to call me Trent or Dr. Cash!
Contact Info	tcash@uwaterloo.ca
Office location	PAS 4042
Office hours	Tuesdays (10:30am – 11:15am); Thursdays (1:00pm – 2:30pm)

Teaching Assistant

Name	Caroline Simpson
Contact Info	Caroline.Simpson@uwaterloo.ca
Office location	PAS 4044
Office hours	By Appointment Only

390 Course Description (650 description N/A)

390: **Natural Science Advanced Research Methods Topics**
Topics explore natural science research methods in Psychology. Consult departmental listings for upcoming topics, which may include methods for studying psychophysiology, neuroanatomy, and/or cognitive processes (e.g., perception, attention, memory, information-processing, decision-making). Activities may include research proposals, group and/or individual projects (e.g., lab experience, data collection), research reports, critiques of published/proposed research, and student presentations.

Course Topic Description

In this course, students will be introduced to the rapidly growing scientific study of the psychology of Artificial Intelligence (AI) and have the opportunity to become producers of AI Psychology research. Given their popularity, the course will focus primarily on one type of AI: Large Language Models (LLMs), such as ChatGPT, Gemini, and Claude. The course will proceed through three sections.

In the first section of the course, students will survey the literature on the capabilities of LLMs – what they're good at, what they're bad at, and how LLM performance is empirically evaluated via psychological benchmarks. In the second section of the course, students will be introduced to AI tools that can be used to augment their own research skills. Then, in the final section of the course, students will explore the literature on how humans think about and interact with AI in their daily lives, and how this can be experimentally investigated. The course will primarily cover AI-relevant studies within the domain of cognitive psychology. Throughout the course, students will also participate in labs designed to teach specific research skills that are necessary to become producers of AI Psychology research. Students will also have the opportunity to design and run a research study in which they leverage AI to better understand a research area of interest to them.

Prior Knowledge

This course is designed to be accessible to undergraduate and graduate students in psychology. Familiarity with basic psychological concepts – such as those taught in a typical introductory psychology course – will be helpful, but no specific content knowledge is required. Participation in this course does not require any programming knowledge, but students with advanced programming skills are encouraged to leverage them in the design of their final projects.

Learning Objectives

- **Learning Objective 1**: Students will develop a working knowledge of the rapidly growing research domain of AI Psychology (APA Goal 1.2)
- **Learning Objective 2**: Students will be able to use AI to interpret, design, and evaluate psychological research (APA Goal 2.2)
- **Learning Objective 3:** Students will be able to exercise scientific reasoning to investigate questions at the intersection of AI and psychology (APA Goal 2.1)
- **Learning Objective 4:** Students will be able to employ ethical standards – including transparency and open science – when conducting research on AI Psychology (APA Goal 3.1)
- **Learning Objective 5:** Students will be able to augment their psychological literacy using artificial intelligence (APA Goal 4.3)
- **Learning Objective 6:** Students will be able to leverage artificial intelligence to improve their written and spoken communication skills (APA Goal 4.2)

Learning Resources

All the materials/readings for this course will be available on the course LEARN page. Assignments that require the usage of Large Language Models (LLMs) can be completed using Microsoft CoPilot, which can be accessed using your WatIAm login. If you wish to use a paid version of a different LLM (e.g., ChatGPT, Gemini, Claude), you will be responsible for the cost of the subscription. Lab #3 will require the use of the website Open Scholar, which offers a free trial. If students wish to use Open Scholar beyond what is available through the free trial, they will be responsible for the cost of additional usage tokens. A full reading list for this course is available on LEARN and at the end of this syllabus.

Attendance

Attendance and participation in this class is mandatory and will count for 10% of your final grade (10 points out of 100). Your first two absences will have no impact on your grade. Each subsequent absence will subtract 0.5 points from your attendance grade (minimum 0/10). So, if you attend 22/24 classes, you will earn full credit. To receive credit for attending a class, you should be present for the entire class. Attendance will be tracked via the completion of "Exit Tickets" that will be given at random times in each class.

Assessments

Final course grades will be calculated using the following categories:

Assessment	**Maximum Points (Undergraduates)**	**Maximum Points (Graduate Students)**
Discussion Board Posts	10	5
Discussion Board Responses	N/A	5
AI-Assisted Essays (x2)	10	10
Labs (x6)	30	30
Research Project Proposal	10	10
Research Project Final Paper	20	20
Research Project Presentation	10	10
Attendance	10	10
Research Participation	2 Extra Credit Points	N/A
Total	100 (+2 Extra Credit)	100

- **Discussion Board Posts:** For each class, students will be assigned 1-2 readings that they are expected to read before class. Each week (except the first and last week), students will be expected to choose one of the readings and write a brief (~ 100 word) discussion post about that reading. In this post, students can choose to ask a question about the reading, highlight something they found interesting about the reading, or propose a follow-up study that would build upon the reading. These Discussion Board Posts will be due at **5:00 p.m. on the Monday before the class in which that reading is discussed**.

- **Discussion Board Responses (Graduate Students Only):** In addition to making Discussion Board posts, graduate students will be required to respond to at least two posts made by other students (undergraduate or graduate). These responses should be of similar length to the original posts (~100 words) and will be **due each week by 5:00 p.m. on Friday.** Credit will only be earned if at least 2 responses are posted.

- **AI-Assisted Essays:** Students will be tasked with writing two brief opinion essays about various topics related to AI Psychology. Students will then upload their essays to an LLM and revise their essays based on the LLM's feedback. Details about each essay assignment will be made available on the course LEARN page.

- **Labs:** Students will be tasked with completing six lab assignments throughout the semester in which they practice a research skill related to AI Psychology. These labs will be started in class and finished outside of class. Details about each lab will be made available on the course LEARN page.

- **Research Project:** The main goal of this course is for students to practice designing, running, analyzing, and summarizing the results of a research project in which they answer a question related to AI Psychology. Undergraduate students will complete this assignment in groups of 2-3 students. Graduate students will complete this project independently. This project is meant to be completed over the course of the semester, with three key deliverables: Proposal; Final Paper; and Presentation. Details about each deliverable will be made available on the course LEARN page. All data collected as part of this research project will be for educational purposes only.

- **Research Participation**: Undergraduate students will have the opportunity to participate in studies through SONA. One point of extra credit in this course will be offered per SONA credit, up to a maximum of 2 credits/points. This extra credit opportunity will not be available to graduate students.

Graduate Student Additional Requirements

Because this course is cross listed for undergraduate and graduate students, some assignments will be altered to be more challenging for graduate students. Specifically:

1. **Discussion Board Responses**: Graduate students are expected to respond to at least two posts made by other students (undergraduate or graduate).

2. **Labs**: Each lab assignment will have additional questions that only the graduate students will be required to complete.

3. **Research Project**: Graduate students will complete the final research project independently.

Weekly Course Flow

Each week of the course except for Week 1 and Week 12 will adhere to the following structural flow:

- **Monday:** By 5 p.m., students should complete the readings for the upcoming week and submit one Discussion Board post.

- **Tuesday/Thursday:** Students should attend class. Classes will be a mix of lectures, in-class activities, labs, and group work time.

- **Friday**: By 5 p.m, graduate students should complete their Discussion Board responses.

- **Sunday**: If an assignment is due (see Course Schedule), submit it by 10 p.m.

A Few Tips for Success in This Course:

- **Tip 1:** Complete the readings and Discussion Board Posts early! In case something comes up, it's better to know that you already have it done.

- **Tip 2:** When reading for this course, focus on the big picture, not the small details. I can promise that I will never ask you to pull the small details from a paper out of thin air. Most academics refer to papers when they need to remember these kinds of things.

- **Tip 3:** Make friends in the class! Maybe even start a group chat. Having people to talk about the assignments with can be a lot of fun – especially when you find a paper that changes how you think.

- **Tip 4:** Start the final paper early! You will do a much better job if you work on it over two months, rather than two weeks.

Grading Policies

- **Late-work policy**: Because of their time-sensitive nature, late-work will NOT be accepted for Discussion Board Posts/Responses or the Final Paper Presentation. For all other assignments, 10% of the points for the assignment will be removed for each day the assignment is late. For this policy, a day is defined as a 24-hour period beginning one minute after the deadline. So, if an assignment is due at 10 p.m. on Sunday, then an assignment turned in at 10:01 p.m. on Sunday is considered one day late, and an assignment turned in at 10:01 p.m. on Monday is considered two days late.

- **Make-up work policy**: Make-up work will not be offered for any assignments. In the case of extenuating circumstances, please see me during office hours. Arrangements will be made at my discretion.

- **Re-grade policy**: Re-grade requests should be submitted to me via email within 48 hours of the graded assignment being made available. Students should indicate precise reasons for which they feel they deserve a different grade on a specific aspect of the assignment. A re-grade is not a guarantee that your grade will be changed. If your grade is changed, the new grade will be retained even if it's lower than the original grade.

A Note on Work-Life Balance:

Being a student is no easy task. It is important to me to remind you that classes are only one part of your university experience. While you're here at UW, you should be growing not only academically, but personally and socially as well. If you feel that this course is inhibiting your ability to maintain a good work-life balance, please let me know. While I am excited to teach this course on a topic about which I am deeply passionate, I care even more about your well-being. If I can make changes to this course that would make your life better without impacting the quality of the learning environment, I'm all ears. As an instructor, my goal is to help you grow as a scholar – but as a human, my job is to help you grow as a person as well. With that said, please know that my door (or Zoom room) is always open.

Course Schedule

Week	Date	Theme/Topic	Required Readings	Assignments Due (see note below)*
Section 1: Evaluating AI Capabilities and Tendencies				
1	5/12	Introduction to Course & Logistics	None	
	5/14	History of Artificial Intelligence and LLMs	Copeland (2026) Roose (2023)	
2	5/19	AI Thinking + Reasoning	Binz & Schulz (2023) Hickman et al. (2024)	
	5/21	AI Emotions + Personality	Salecha et al. (2024)	
3	5/26	Lab 1: Empirically Testing AI Performance	Reuel et al. (2024)	
	5/28	AI Biases	Bai et al. (2025) Tao et al. (2024)	Lab #1 (5/31 @ 10 p.m.)
4	6/2	Final Project Work Day	None	
	6/4	AI Metacognition and Wisdom	Cash et al. (2025) Johnson et al. (2026)	AI-Assisted Essay #1 (6/7 @ 10 p.m.)
Section 2: Using AI As a Research Tool				
5	6/9	Lab 2: Prompt Engineering	Geromienko (2025) – Chapter 2	
	6/11	(Open) Science in the Age of AI	Van Noorden & Perkel (2023) Grossman et al. (2023)	Lab #2 (6/14 @ 10:00 p.m.)
6	6/16	Lab 3: AI for Literature Review	Asai et al. (2026)	
	6/18	AI for Hypothesis	Banker et al. (2024)	Lab #3

		Generation + Testing	Cui et al. (2025)	(6/21 @ 10 p.m.)
7	6/23	Lab 4: AI for Statistical Analysis	Lee (2025) Sun et al. (2025) (Note: Sun et al. (2025) gets a bit complicated – don't get caught up in the details)	
	6/25	AI for Stimulus Generation + Stimulus Testing	Lechuga & Karle (2026) Behrend & Landers (2025)	Research Proposal (6/28 @ 10 p.m.) Lab #4 (6/28 @ 10 p.m.)
8	6/30	Lab 5: AI for Coding Data + Text Analysis	Gilardi et al. (2023) Bojic et al. (2025)	
	7/2	Final Project Work Day	None	Lab #5 (7/5 @ 10 p.m.)
Section 3: Studying Humans in an AI World				
9	7/7	Learning with AI vs. Offloading to AI	Bastani et al. (2025)	
	7/9	Lab 6: Detecting AI Respondents	Westwood et al. (2025) Rodriguez & Oppenheimer (2024)	Lab #6 (7/12 @ 10 p.m.)
10	7/14	Final Project Work Day	None	
	7/16	AI Persuasion and Belief Formation	Costello et al. (2024)	AI-Assisted Essay #2 (7/19 @ 10 p.m.)
11	7/21	Human Beliefs about AI – What is it? Is it conscious?	Colombatto & Fleming (2024)	
	7/23	Student Presentations	None	
12	7/28	Student Presentations	None	
	7/30	Student Presentations	None	
Final Paper Due: Sunday, August 2nd @ 10:00 p.m.				

***Note: To avoid clutter, the "Assignments Due" column does not include weekly Discussion Posts (All Students) or Responses (Graduate Students Only)**

Assignment Screening

No automated assignment screening will be utilized within this course.

Generative AI

Generative artificial intelligence (GenAI) trained using large language models (LLM) or other methods to produce text, images, music, or code, like Chat GPT, DALL-E, or GitHub CoPilot, may be used under specific conditions in this course with proper documentation, citation, and acknowledgement. Permitted uses of and expectations for using GenAI will discussed in class and outlined on assignment instructions.

Recommendations for how to cite generative AI in student work at the University of Waterloo may be found through the Library.

Please be aware that generative AI is known to falsify references to other work and may fabricate facts and inaccurately express ideas. GenAI generates content based on the input of other human authors and may therefore contain inaccuracies or reflect biases.

To protect the privacy and security of any data entered, students should use the University's version of Co-Pilot and login with their UW ID. Data entered into other systems can be added to training sets, monitored, geolocated and even reproduced as output which may share private personal information or result in intellectual property breaches.

In addition, you should be aware that the legal/copyright status of generative AI inputs and outputs is unclear. Exercise caution when using large portions of content from AI sources, especially images. More information is available from the Copyright Advisory Committee.

You are accountable for the content and accuracy of all work you submit in this class, including any supported by generative AI. You should be able to readily demonstrate your knowledge of your submissions. To demonstrate your learning, you should keep your rough notes, including sources, research notes, brainstorming, drafting notes and prompts. You may be asked to submit these notes along with earlier drafts of your work, either through saved drafts or saved versions of a document.

Administrative Policy

Accommodating religious and spiritual observances
[Applies only to students taking undergraduate courses]
The University of Waterloo has a duty to accommodate religious, spiritual, and other creed-based beliefs and practices under the Ontario Human Rights Commission (2015) Policy on preventing discrimination based on creed. Students may seek accommodations for missed course components on religious, spiritual, or other creed grounds. In such cases the students should please consult the instructor within two weeks of the announcement of the due date for which the accommodation is being sought. Students also may request accommodations for temporary absences from classes or other course-related activities to engage in prayer or other daily spiritual practices. Instructors are expected to make reasonable arrangements to respectfully accommodate such requests. The University of Waterloo has a number of multi-faith spaces that students can use for faith-based practices.

Declaring absences
[Applies only to students taking undergraduate courses]
Please see the Academic Regulations section of the Undergraduate Calendar for more details. Regardless of the process used to declare an absence, it is the student's responsibility to contact their instructor so that the instructor can determine how to accommodate the missed work.

Absences due to religious, creed, and spiritual observances
Beginning Fall 2024, students can submit a Religious Observance Self-Declaration Form in Quest. This form allows the student to register dates of potential conflicts with coursework. Students are expected to record their absence within two weeks of announcement of the due date or scheduled examination date for which academic accommodation is being sought. Submission through Quest will notify your instructors of your absence.

Absences due to verified condition
For absences due to illness, students need to obtain a Verification of Condition (VOC) Form and submit the VOC form through the Absence Database. Once the verified condition is registered the system your instructors will be notified of your absence. Students should not submit their VOC form or any other medical documentation to your instructors, teaching assistants, or other course personnel. To request and arrange accommodations, it is the student's responsibility to contact the instructor within 48 hours of the date of any missed course component, or as soon as the student is able to given their health condition.

Absences due to extenuating circumstances (e.g., bereavement)
For absences due to extenuating circumstances (e.g., bereavement, serious family illness), students submit a request for absence due to extenuating circumstances to the Absence Database for processing. Students use the form to submit supporting documentation, which will be considered in a comparable manner to a VOC. Official documentation is necessary before any action can be taken. Examples are plane/train/bus tickets, court documents, police reports, death certificates/obituaries/letters from funeral directors. To request and arrange accommodations, it is the student's responsibility to contact the instructor within 48 hours of the date of any missed course component, or as soon as the student is able to given their circumstances.

Self-declared short-term absences for any reason
Students can self-declare one short-term absence per term for any reason. A short-term absence covers two calendar days, and no documentation is required. The student has only one self-declared short-term absence per term and this absence can only be used during the formal lecture period. Self-declared absences cannot be applied to Laboratory (LAB), Clinic (CLN), or Studio (STU) course components. Submitting the self-declaration through Quest will notify your instructors of your absence. It is the student's responsibility to contact their instructors within the first 24 hours after submitting their self-declaration to discuss accommodations for missed course components.

Recording lecture
Use of recording devices during lectures is only allowed with explicit permission of the instructor of the course. If allowed, video recordings may only include images of the instructor and not fellow classmates. Posting of videos or links to the video to any website, including but not limited to social media sites such as: Facebook, Twitter, etc., is strictly prohibited.

University Policy

Mental Health: At the University of Waterloo, we are dedicated to supporting your mental and emotional well-being. Our Counselling Services offer confidential support, including individual counselling, workshops, and crisis intervention. If you're struggling, please reach out for help at 519-888-4096 or visit their website for more information.

Academic integrity: In order to maintain a culture of academic integrity, members of the University of Waterloo community are expected to promote honesty, trust, fairness, respect and responsibility.

Grievance: A student who believes that a decision affecting some aspect of their university life has been unfair or unreasonable may have grounds for initiating a grievance. Read Policy 70, Student Petitions and Grievances, Section 4. When in doubt, please be certain to contact the department's administrative assistant who will provide further assistance.

Discipline: A student is expected to know what constitutes academic integrity to avoid committing an academic offence, and to take responsibility for their actions. A student who is unsure whether an action constitutes an offence, or who needs help in learning how to avoid offences (e.g., plagiarism, cheating) or about "rules" for group work/collaboration should seek guidance from the course instructor, academic advisor, or the undergraduate associate dean. For information on categories of offences and types of penalties, students should refer to Policy 71, Student Discipline. For typical penalties, check Guidelines for the Assessment of Penalties.

Appeals: A decision made or penalty imposed under Policy 70, Student Petitions and Grievances (other than a petition) or Policy 71, Student Discipline may be appealed if there is a ground. A student who believes they have a ground for an appeal should refer to Policy 72, Student Appeals.

Note for students with disabilities and disabling conditions: The University of Waterloo recognizes its obligations under the Ontario Human Rights Code to accommodate students with known or suspected disabilities and disabling conditions (e.g. medical conditions, injuries, impacts of trauma such as from violence or discrimination) to the point of undue hardship. To support this obligation, AccessAbility Services (AAS) collaborates with all academic departments and schools to facilitate academic accommodations for students with disabilities and disabling conditions without compromising the academic integrity of the curriculum. If you believe you may require academic accommodations (e.g., testing accommodations, classroom accommodations), register with AAS as early in the term as possible by completing the online application. Students already registered with AAS must activate their accommodations for each of their courses at the beginning of each term using AAS' online system. If you require assistance, contact AAS by phone (519-888-4567 ext. 35082), email (access@uwaterloo.ca) or in-person (Needles Hall North, 1st Floor, Room 1401).

Turnitin.com: Text matching software (Turnitin®) may be used to screen assignments in this course. Turnitin® is used to verify that all materials and sources in assignments are documented. Students' submissions are stored on a U.S. server, therefore students must be given an alternative (e.g., scaffolded assignment or annotated bibliography), if they are concerned about their privacy and/or security. Students will be given due notice, in the first week of the term and/or at the time assignment details are provided, about arrangements and alternatives for the use of Turnitin in this course.

It is the responsibility of the student to notify the instructor if they, in the first week of term or at the time assignment details are provided, wish to submit alternate assignment.

Full Reading List References + Optional Readings

Week 1: History of Artificial Intelligence and LLMs:

Required:

Copeland, B. J. (2026). History of artificial intelligence. *Encyclopedia Brittanica. https://www.britannica.com/science/history-of-artificial-intelligence*

Roose, K. (2023): How does ChatGPT really work? *The New York Times*. https://www.nytimes.com/2023/03/28/technology/ai-chatbots-chatgpt-bing-bard-llm.html

Optional:

Wolfram, S. (2023). *What is ChatGPT doing… and why does it work?* https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/

Hoffman, C. H. (2022). Is AI intelligent? An assessment of artificial intelligence, 70 years after Turing. *Technology in Society, 101893.* https://doi.org/10.1016/j.techsoc.2022.101893

Week 2: AI Thinking & Reasoning

Required

Binz, M., & Schulz, E. (2023). Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences, 120*(6), e2218523120. https://doi.org/10.1073/pnas.2218523120

Hickman, L., Dunlop, P. D., & Wolf, J. L. (2024). The performance of large language models on quantitative and verbal ability tests: Initial evidence and implications for unproctored high-stakes testing. *International Journal of Selection and Assessment, 32*, 499-511. https://doi.org/10.1111/ijsa.12479

Optional:

Trott, S., Jones, C., Chang, T., Michaelov, J., & Bergen, B. (2023). Do large language models know what humans know? *Cognitive Science, 47*(7), e13309. https://doi.org/10.1111/cogs.13309

Week 2: AI Emotions + Personality

Required:

Salecha, A., Ireland, M. E., Subrahmanya, S., Sedoc, J., Ungar, L. H., & Eichstaedt, J. C. (2024). Large language models display human-like social desirability biases in Big Five personality surveys. *PNAS Nexus, 3*, 533. https://doi.org/10.1093/pnasnexus/pgae533

Optional:

Wang, Y., Zhao, J., Ones, D. S., He, L., & Xu, X. (2025). Evaluating the ability of large language models to emulate personality. *Scientific Reports, 15*, 519. https://doi.org/10.1038/s41598-024-84109-5

Chen, S., Gao, M., Sasse, K., Hartvigsen, T., Anthony, B., Fan, L., Aerts, H., Gallifant, J., & Bitterman, D. S. (2025). When helpfulness backfires: LLMs and the risk of false medical information due to sycophantic behavior. *Npj Digital Medicine, 8*, 605. https://doi.org/10.1038/s41746-025-02008-z

Week 3: Empirically Testing AI Capabilities

Required:

Reuel, A., Hardy, A., Smith, C., Lamparth, M., Hardy, M., Kochenderfer, M. J. (2024). What makes a good AI benchmark. *Human-Centered Artificial Intelligence: Policy & Society*. https://hai.stanford.edu/assets/files/hai-policy-brief-what-makes-a-good-ai-benchmark.pdf

Optional:

Eriksson, M., Purificato, E., Noroozian, A., Vinagre, J., Chaslot, G., Gomez, E., Fernandez-Llorca, D. (2025). Can we trust AI benchmarks? An interdisciplinary review of current issues in AI evaluation. *Proceedings of the Eighth AAAI/ACM Conference on AI, Ethics, and Society*. https://doi.org/10.48550/arXiv.2502.06559

Week 3: AI Biases

Required:

Bai, X., Wang, A., Sucholutsky, I., & Griffiths, T. L. (2025). Explicitly unbiased large language models still form biased associations. *Proceedings of the National Academy of Sciences, 122*(8), e2416228122. https://doi.org/10.1073/pnas.2416228122

Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus, 3*, 346. https://doi.org/10.1093/pnasnexus/pgae346

Optional:

Omiye, J. A., Lester, J. C., Spichak, S., Rotemberg, V., & Daneshjou, R. (2023). Large language models propagate race-based medicine. *Npj Digital Medicine, 6*, 195. https://doi.org/10.1038/s41746-023-00939-z

Peng, T.-Q., Yang, K., Lee, S., Li, H., Chu, Y., Lin, Y., & Liu, H. (2026). Beyond partisan leaning: A comparative analysis of political bias in large language models. *Journal of Information Technology & Politics,* advanced online publication. https://doi.org/10.1080/19331681.2026.2646990

Week 4: AI & Metacognition

Required:

Cash, T. N., Oppenheimer, D. M., Christie, S., & Devgan, M. (2025). Quantifying uncert-AI-nty: Testing the accuracy of LLMs' confidence judgments. *Memory & Cognition, 54*(2), 375-400. https://doi.org/10.3758/s13421-025-01755-4

Johnson, S. G. B., Karimi, A.-H., Bengio, Y., Chater, N., Gerstenberg, T., Larson, K., Levine, S., Mitchell, M., Rahwan, I., Scholkopf, B., & Grossman, I. (2026). Imagining and building wise machines: The centrality of AI metacognition. *Trends in Cognitive Sciences,* advanced online publication. https://doi.org/10.1016/j.tics.2026.01.002

Optional:

Steyvers, M., Tejeda, H., Kumar, A., Belem, C., Karny, S., Hu, X., Mayer, L. W., & Smyth, P. (2025). What large language models know and what people think they know. *Nature: Machine Intelligence, 7*, 221-231. https://doi.org/10.1038/s42256-024-00976-7

Steyvers, M., & Peters, M. A. K. (2025). Metacognition and uncertainty communication in humans and large language models. *Current Directions in Psychological Science,* advanced online publication. https://doi.org/10.1177/09637214251391158

Week 5: Prompt Engineering

Required:

Geroimenko, V. (2025). Chapter 2: Key Principles of Good Prompt Design. In *The Essential Guide to Prompt Engineering.* https://doi.org/10.1007/978-3-031-86206-9_2

Optional:

Giray, L. (2023). Prompt engineering with ChatGPT: A guide for academic writers. *Annals of Biomedical Engineering, 51*, 2629-2633. https://doi.org/10.1007/s10439-023-03272-4

Geroimenko, V. (2025). Chapter 3: Key techniques for writing effective prompts. In *The Essential Guide to Prompt Engineering.* https://doi.org/10.1007/978-3-031-86206-9_3

Week 5: (Open) Science in the Age of AI

Required:

Grossman, I., Feinberg, M., Parker, D. C., Christakis, N. A., Tetlock, P. E., & Cunningham, W. A. (2023). AI and the transformation of social science research. *Science, 380*(6650), 1108-1109. https://doi.org/10.1126/science.adi1778

Van Noorden, R., & Perkel, J. M. (2023). AI and science: What 1,600 researchers think. *Nature, 621,* 672-675. *https://doi.org/10.1038/d41586-023-02980-0*

Optional:

Resnik, D. B., & Hosseini, M. (2025). The ethics of using artificial intelligence in scientific research: New guidance needed for a new tool. *AI and Ethics, 5*, 1499-1521. https://doi.org/10.1007/s43681-024-00493-8

Tyner, A. H. et al. (2026). Investigating the replicability of the social and behavioural sciences. *Nature, 652*, 143-150. https://doi.org/10.1038/s41586-025-10078-y

Week 6: AI for Literature Review

Required:

Asai, A. et al. (2026). Synthesizing scientific literature with retrieval-augmented language models. *Nature, 650*, 857-863. https://doi.org/10.1038/s41586-025-10072-4

Optional:

Wagner, G., Lukyanenko, R., & Pare, G. (2022). Artificial intelligence and the conduct of literature reviews. *Journal of Information Technology, 37*(2), 209-226. https://doi.org/10.1177/02683962211048201

Bolaños, F., Salatino, A., Osborne, F., & Motta, E. (2024). Artificial intelligence for literature reviews: Opportunities and challenges. *Artificial Intelligence Review, 57,* Article 259. https://doi.org/10.1007/s10462-024-10902-3

Week 6: AI for Hypothesis Generation + Testing

Required:

Banker, S., Chatterjee, P., Mishra, H., & Mishra, A. (2024). Machine-assisted social psychology hypothesis generation. *American Psychologist, 79*(6), 789-797. https://doi.org/10.1037/amp0001222

Cui, Z., Li, N., & Zhou, H. (2025). A large-scale replication of scenario-based experiments in psychology and management using large language models. *Nature Computational Science, 5*, 627-634. https://doi.org/10.1038/s43588-025-00840-7

Optional:

Hermida Carrillo, A., Stachl, C., & Talaifar, S. (2024). A workflow for human-centered machine-assisted hypothesis generation: Commentary on Banker et al. (2024). *American Psychologist, 79*(6), 800-802. https://doi.org/10.1037/amp0001256

Ke, L., Tong, S., Cheng, P., & Peng, K. (2025). Exploring the frontiers of LLMs in psychological applications: A comprehensive review. *Artificial Intelligence Review, 58*, 305. https://doi.org/10.1007/s10462-025-11297-5

Sohail, A., & Zhang, L. (2025). Using large language models to facilitate academic work in the psychological sciences. *Current Psychology, 44*, 7910-7918. https://doi.org/10.1007/s12144-025-07438-2

Week 7: AI for Statistical Analysis

Required

Lee, W. (2025). LLMs-powered data analysis: 5 transformative use cases for data workflow. *Medium. https://medium.com/@whee.2013/llms-powered-data-analysis-5-transformative-use-cases-for-data-workflow-3bd0043a4abb*

Sun, M., Han, R., Jiang, B., Qi, H., Sun, D., Yuan, Y., & Huang, J. (2025). A survey on large language model-based agents for statistics and data science. *The American Statistician,* advanced online publication. https://doi.org/10.1080/00031305.2025.2561140

Optional

Wang, Z., Zhou, Z., Song, D., Huang, Y., Chen, S., Ma, L., & Zhang, T. (2025). Towards understanding the characteristics of code generation errors made by large language models. *arXiv*. https://doi.org/10.48550/arXiv.2406.08731

Sebastian, J., Moon, J.-M., & Camburn, E. (2025). Enhancing quantitative analysis in social sciences with large language models (LLMs): A methodological case study in educational research. *AERA Open, 11*(1), 1-18. https://doi.org/10.1177/23328584251389621

Week 7: AI for Stimulus Generation + Stimulus Testing

Required

Behrend, T. S., & Landers, R. N. (2025). Participant interactions with artificial intelligence: Using Large Language Models to generate research materials for surveys and experiments. *Journal of Business and Psychology, 40*, 1275-1297. https://doi.org/10.1007/s10869-025-10035-6

Optional

Lechuga, J., & Karle, N. N. (2026). Generating experimental text stimuli for psychological research using ChatGPT. *Advances in Methods and Practices in Psychological Science, 9*(1), 1-18. https://doi.org/10.1177/25152459251407631

Adhikari, D. M., Hartland, A., Weber, I., & Cannanure, V. K. (2025). Exploring LLMs for automated generation and adaptation of questionnaires. *arXiv*. https://doi.org/10.48550/arXiv.2501.05985

Week 8: AI for Coding Data + Text Analysis

Required

Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences, 120*(30), e2305016120. https://doi.org/10.1073/pnas.2305016120

Bojic, L., Zagovora, O., Zelenkauskaite, A., Vukovic, V., Cabarkapa, M., Veseljevic Jerkovic, S., & Jovancenic, A. (2025). Comparing large language models and human annotators in latent content analysis of sentiment, political leaning, emotional intensity, and sarcasm. *Scientific Reports, 15*(1), 11477. https://doi.org/10.1038%2Fs41598-025-96508-3

Optional

Bermejo, V. J., Gago, A., Galvez, R. H., & Harari, N. (2025). LLMs outperform outsourced human coders on complex textual analysis. *Scientific Reports, 15*, 40122. https://doi.org/10.1038/s41598-025-23798-y

Week 9: Learning with AI vs. Offloading to AI

Required

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabacki, O., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Optional

Melumad, S., & Yun, J. H. (2025). Experimental evidence of the effects of large language models versus web search on depth of learning. *PNAS Nexus, 4*(10), pgaf316. https://doi.org/10.1093/pnasnexus/pgaf316

Lira, B., Rogers, T., Goldstein, D. G., Ungar, L., & Duckworth, A. L. (2025). Coach not crutch: Evidence that AI can improve writing skill despite reducing effort [Preprint]. *ArXiv.* https://doi.org/10.48550/arXiv.2502.02880

Oppenheimer, D. M., Cash, T. N., & Connell Pensky, A. E. (2025). You've got AI friend in me: LLMs as collaborative learning partners. *International Journal of Artificial Intelligence in Education, 35*(6), 3896-3921. https://doi.org/10.1007/s40593-025-00521-6

Week 9: Detecting AI Respondents

Required

Westwood, S. J. (2025). The potential existential threat of large language models to online survey research. *Proceedings of the National Academy of Sciences, 122*(47), e2518075122. https://doi.org/10.1073/pnas.2518075122

Rodriguez, C., & Oppenheimer, D. M. (2024). Creating a bot-tleneck for malicious AI: Psychological methods for bot detection. *Behavior Research Methods, 56*, 6258-6275. https://doi.org/10.3758/s13428-024-02357-9

Optional

Peer, E., Rothschild, D., Gordon, A., Evernden, Z., & Damer, E. (2022). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods, 54*, 1643-1662. https://doi.org/10.3758/s13428-021-01694-3

Week 10: AI Persuasion and Belief Formation

Required

Costello, T. H., Pennycook, G., & Rand, D. G. (2024). Durably reducing conspiracy beliefs through dialogues with AI. *Science, 385*(6714), eadq1814. https://doi.org/10.1126/science.adq1814

Optional

Huang, G., & Wang, S. (2023). Is artificial intelligence more persuasive than humans? A meta-analysis. *Journal of Communication, 73*(6), 552-562. https://doi.org/10.1093/joc/jqad024

Matz, S. C., Teeny, J. D., Vaid, S. S., Harari, G. M., & Cerf, M. (2024). The potential of generative AI for personalized persuasion at scale. *Scientific Reports, 14*, 4692. https://doi.org/10.1038/s41598-024-53755-0

Week 11: Human Beliefs about AI – What is it? Is it conscious?

Required

Colombatto, C., & Fleming, S. (2024). Folk psychological attributions of consciousness to large language models. *Neuroscience of Consciousness, 1*, niae013. https://doi.org/10.1093/nc/niae013

Optional

Anthis, J. R., Pauketat, J. V. T., Ladak, A., & Manoli, A. (2025). Perceptions of sentient AI and other digital minds: Evidence from the AI, morality, and sentience (AIMS) survey. *arXiv. https://doi.org/10.48550/arXiv.2407.08867*

Kang, B., Kim, J., Yun, T., Bae, H., & Kim, C.-E. (2026). Identifying features that shape perceived consciousness in LLM-based AI: A quantitative study of human responses. *Computers in Human Behavior Reports, 21*, 100901. https://doi.org/10.1016/j.chbr.2025.100901

Sebo, J., & Long, R. (2025). Moral consideration for AI systems by 2030. *AI and Ethics, 5*, 591-606. https://doi.org/10.1007/s43681-023-00379-1