ENGE3940

AI and Language

Term 2, 2026–27 · Tom van Nuenen

1. Course code	ENGE 3940
2. Title in English	AI and Language
3. Title in Chinese	人工智慧和語言

4. Course description

This course introduces students to the relationship between artificial intelligence and human language, with a focus on everyday writing and communication. Students will learn basic principles of how AI systems process and generate language, then explore their capabilities and limitations through readings, discussion and guided practical activities. Topics include how wording and context shape meaning, how language changes for different audiences, and how social and cultural assumptions appear in AI-generated texts. Working primarily with English examples, students will learn to organise and compare small sets of language data and support their interpretations with evidence. A final poster project, developed through structured activities and feedback, will explore a question arising from the course. No programming background is required.

5. Learning outcomes

By the end of the course, students should be able to:

1. Explain how a chat model produces a response, including the difference between what it learns during pretraining and the behaviour it develops during post-training.
2. Analyse a chatbot response closely by quoting its wording, explaining its effects, and identifying its assumptions about the reader and the situation.
3. Use concepts from the study of language, meaning and social interaction to explain what a chatbot response does and how people interpret and rely on it.
4. Use AI tools, including embeddings and model-based labelling, to examine a small set of texts, and check the results through close reading.
5. Design a small comparison that changes one factor, such as wording, context or model, and assess whether the resulting differences reflect that change or variation between runs.
6. Evaluate a chatbot response using stated criteria, support the evaluation with textual evidence, explain disagreements between readers or models, and draw an evidence-backed conclusion.

6. Course syllabus

Topic	Contents and fundamental concepts
How language models produce text	Symbolic and connectionist AI; large language models (LLMs); tokens and next-token prediction; pretraining and post-training; base and chat models.
How people understand AI	Changing definitions of “AI”; the ELIZA effect; Turing’s imitation game; anthropomorphism and metaphors for AI.
Close reading of chatbot responses	Positivism and constructionism; close reading; word choice, hedging, forms of address and structure; substitution and comparison tests.

Topic	Contents and fundamental concepts
Prediction and patterns in language	Information and prediction (Shannon); linguistic value (Saussure); the distributional hypothesis and collocation (Firth); word embeddings; debates about form and meaning, including the octopus test, the Chinese Room and stochastic parrots.
Meaning in context	Speech acts (Austin); implicature and the cooperative principle (Grice); system prompts.
Language in social interaction	Register and audience; footing (Goffman); politeness and face; the assistant persona.
Authorship and varieties of English	AI-mediated communication; authorship; standard and non-standard Englishes; cultural defaults; model collapse.
Evaluating AI language	Evaluation criteria; blind judging; controlled comparison; variation between runs.
Using AI to make judgments	LLM-as-judge methods; sycophancy; hallucination; cognitive offloading.
Analysing and presenting findings	Combining counts, embedding maps and close reading; reporting findings and limitations.

7. Course components

Teaching mode	Percentage of class time
All classes on site and face to face	
Lectures, two hours a week, not hybrid	67%
Interactive tutorials in small groups, one hour a week, not hybrid	33%

How the classes work

This is an introductory course. You do not need programming skills or previous study of AI. Lectures include demonstrations and short experiments. Tutorials take place in small groups led by the instructor or a teaching assistant. You work in pairs or groups, record what you tried, and discuss the results. Please bring a laptop to tutorials. The class also runs a shared experiment throughout the term. You collect phrases and messages from Hong Kong life. Each week, a pair of students tests them on several chatbots and briefly reports any changes in the responses.

Tools

We mainly use a chatbot available through CUHK, such as Google's Gemini Enterprise for Education. We may use DeepSeek or Qwen for comparisons. Some activities also use free browser tools, which are introduced in class. You do not need a paid account or a VPN.

8. Assessment types, weights and criteria

Requirement	Weight	When
Participation and contribution through reading cards	20%	Most tutorials
Group presentation	20%	After reading week

Requirement	Weight	When
Individual analysis check	20%	After reading week
Group research poster, with AI use allowed	40%	Last week of term

Assessment criteria and guidance

Work receives a letter grade based on the following criteria: (1) intellectual capacity; (2) analytical power and reasoning; (3) knowledge of the subject; (4) imagination and originality; (5) writing and argumentation; (6) ability to conduct scholarly research; (7) capacity for independent work; and (8) reliability and responsibility. All assessment tasks use procedures practised in class. Marks reflect the quality of your analysis rather than the polish of your English. Short phrases, labels and spoken explanations are acceptable.

Participation and contribution through reading cards (20%). At the end of most tutorials, you complete a short handwritten reading card. You copy a phrase from a chatbot response and explain what it does. Cards are checked for completion rather than graded for quality. Missing one or two cards does not reduce your mark.

Group presentation (20%). After reading week, each group presents its plan for the poster study. The presentation introduces the research question, a relevant language concept and some early examples. Every group member speaks. The presentation is assessed on: (1) a clear question; (2) use of a relevant language concept; and (3) explanation of the examples.

Individual analysis check (20%). This is a short individual task completed in class, on paper and without devices, after reading week. You analyse chatbot responses you have not seen before using methods practised in tutorials.

Group research poster (40%). Groups submit their poster and a record of their data before the poster session in the last week of term, and present the poster at the session. You may use AI tools for any part of the poster, provided you record how you used them. The poster is assessed on: (1) a clear question and a relevant concept; (2) a sound comparison; (3) analysis supported by quoted examples; (4) use of AI tools checked against your own reading; and (5) clear presentation and a conclusion supported by the evidence. Individual marks may be adjusted through peer evaluation under the department's rules for group work.

The poster project

Groups form before reading week, with members drawn from the same tutorial group. Each group chooses a small question about how chatbots handle language. The question may come from something recorded in your tutorial logs. To investigate the question, compare two versions of a prompt, run each version more than once, and read the responses closely. When groups form, you receive a short project brief with example topics and instructions. Projects do not require private messages, scraped data or participants. Ask the instructor before using any of these.

Auditors are expected to participate fully in all classes and complete the presentation assignment. This allows them to contribute to and benefit from the collaborative work.

9. Required and recommended readings

Most weeks have either one short reading or a handout containing short passages from classic works on language. Questions guide your reading, and the passages are discussed in class. There is no textbook to buy. All readings are available through the University Library or on Blackboard.

Required readings

- Shanahan, Murray. 2024. "Talking about Large Language Models." *Communications of the ACM* 67 (2): 68–79.
- Bender, Emily M., and Alexander Koller. 2020. "Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data." *Proceedings of ACL*, 5185–5198. The assigned selection is the octopus thought experiment.

- Shanahan, Murray, Kyle McDonell, and Laria Reynolds. 2023. “Role Play with Large Language Models.” *Nature* 623: 493–498.
- Liang, Weixin, et al. 2023. “GPT Detectors Are Biased against Non-Native English Writers.” *Patterns* 4 (7): 100779.
- Hicks, Michael Townsen, James Humphries, and Joe Slater. 2024. “ChatGPT Is Bullshit.” *Ethics and Information Technology* 26: 38.

Handouts

Handouts include short passages from classic works on language discussed in class, along with short guides to close reading, the poster project and keeping a record.

Recommended reading

- Jurafsky, Daniel, and James H. Martin. 2026. *Speech and Language Processing*. Third edition, online draft. This is a reference for the technical parts of the course. We use selected explanations and figures only.

10. Feedback

You receive feedback on the individual analysis check and at several stages of the poster project. Feedback comes from your tutor and other groups. Common problems are discussed in class. You can also email the instructor or arrange a meeting.

11. Course schedule

	Date	Topic	Reading	Tutorial
1	12 Jan	Experimenting with chatbots	Close-reading guide (in class)	Test chatbot robustness by changing a detail in a familiar riddle or saying. Examine responses to user correction by challenging an answer.
2	19 Jan	From rules to prediction	Shanahan	Compare next-token predictions from a small language model with a chat model's responses to the same input.
3	26 Jan	Prediction and word choice	Handout	Use word substitution to examine how lexical choice affects sentence meaning and the outputs of text and image models.
4	2 Feb	Word meaning and patterns of use	Bender & Koller	Compare a word's contexts in training data, chatbot sentence completions and nearest neighbours in an embedding space.
	9 Feb	*Lunar New Year holiday*		*No class*
5	16 Feb	Meaning in context	Handout	Compare interpretations of an indirect refusal across conversational contexts. Each group prepares a hidden system instruction for the following tutorial.
6	23 Feb	Register, audience and the assistant's voice	Shanahan et al.	Infer another group's hidden system instruction from a chatbot's responses. Review tutorial logs to identify a possible poster question.
7	2 Mar	AI writing and varieties of English	Liang et al.	Analyse how AI editing changes the register and style of your own writing. Form groups, agree on a poster question with your tutor, and prepare the presentation over reading week.
	9 Mar	*Reading week*		*No class*
8	16 Mar	Evaluating a chatbot response	Worked example	Present the planned group study and answer questions.
9	23 Mar	Prompt changes and variation between runs	Record guide	Complete the individual analysis check in the lecture. In the tutorial, define evaluation criteria, design a controlled prompt comparison, collect repeated responses and establish a data record.
10	30 Mar	Using AI to make judgments	Hicks et al.	Compare human and model-based labelling of the same responses, including the reasons given. Repeat the model's evaluation with a different response order to test for order effects.

	Date	Topic	Reading	Tutorial
11	6 Apr	Analysing responses with counts, maps and close reading	Worked example	Combine counts, embedding maps and close reading to complete the analysis. Select examples for the poster, including one that does not fit the pattern.
12	13 Apr	Reviewing poster drafts	Poster guide	Exchange full poster drafts for peer review. Check claims against the evidence and plan revisions with your tutor.
13	20 Apr	Poster session	—	Present findings in a poster gallery and vote for the most surprising finding.

Meeting details

1 Experimenting with chatbots

12 January 2027

We discuss the changing meanings of the term “AI” across the decades. Short experiments with poetry and altered riddles, through a lens of close reading, expose our expectations of chatbots and our interpretation of their responses. We introduce the syllabus, and consider how models generate conversational language.

Reading: No advance reading. After class, read the one-page close-reading guide and a short introductory selection from Jurafsky and Martin.

2 From rules to prediction

19 January 2027

We trace the development of AI from rule-based programs to machine learning and large language models. ELIZA is introduced as a historical example exposing the tendency to attribute understanding to conversational systems. The lecture covers tokens, next-token prediction, variation between responses, and the distinction between pretraining and post-training. The class begins collecting phrases for the shared experiment.

Reading: Shanahan (2024).

3 Prediction and word choice

26 January 2027

We examine information and word choice through Shannon’s guessing game and short passages by Shannon and Saussure. Activities compare the class’s next-word predictions with a model’s predictions and test the effects of word substitutions in text and image prompts.

Reading: Handout with short passages from Shannon and Saussure.

4 Word meaning and patterns of use

2 February 2027

We introduce Firth’s account of meaning through patterns of use and examine how language models represent words and sentences through their contexts. We discuss Bender and Koller’s octopus thought experiment, and a response to it, to frame the discussion on patterns and meaning. We use a research tool to illustrate internal model representations, including words absent from the final response.

Reading: The octopus thought experiment in Bender and Koller (2020), and a handout containing Firth’s passage on “the company it keeps”.

9 February: Lunar New Year holiday. No class.

5 Meaning in context

16 February 2027

We examine how context shapes meaning through Austin's account of speech acts and Grice's account of implicature. Discussion covers indirect refusals, whether models can promise or apologise, and how system prompts guide chatbot responses.

Reading: Handout with short passages from Austin and Grice.

6 Register, audience and the assistant's voice

23 February 2027

We examine how language varies with audience and situation, drawing on Aristotle's advice on style and studies comparing chatbot register with human writing. Goffman's distinction between the speaker, writer and person responsible for an utterance provides a way to analyse the chatbot's use of the first person.

Reading: Shanahan, McDonell and Reynolds (2023), and a handout with short passages from Aristotle and Goffman.

7 AI writing and varieties of English

2 March 2027

An experiment with repeated AI retellings of a local story introduces changes in language and cultural detail. We discuss AI-mediated communication, standard and non-standard Englishes, and the effects of AI polishing on personal and local writing styles. Groups form for the poster project.

Reading: Liang et al. (2023).

9 March: reading week. No class.

8 Evaluating a chatbot response

16 March 2027

We distinguish accuracy, instruction following and appropriateness to the situation as criteria for evaluating chatbot responses. Activities practise supporting judgments with textual evidence and approaches to make judgments more reliable.

Reading: Handout with a worked example of a comparison.

9 Prompt changes and variation between runs

23 March 2027

We examine how to distinguish the effects of prompt changes from variation between runs. Activities repeat part of a published experiment and retest examples from Meeting 1. We discuss the popularity of "magic" prompting tips. The individual analysis check takes place during the lecture.

Reading: The one-page guide to keeping the record.

10 Using AI to make judgments

30 March 2027

We examine the use of models to label and evaluate text, including ways to check their judgments. We also discuss how agreement, flattery and model-generated explanations can affect human judgment, returning to the correction activity from Meeting 1.

Reading: Hicks, Humphries and Slater (2024).

11 Analysing responses with counts, maps and close reading

6 April 2027

Groups combine simple counts, maps of their responses and close reading. We examine how examples that do not fit the overall pattern can help make conclusions more precise.

Reading: Revisit the worked example from Meeting 8. No new reading.

12 Reviewing poster drafts

13 April 2027

We review how to organise a poster around a research question and supporting examples, with readable figures, short captions and a conclusion supported by the study. Each group reviews its evidence with the instructor.

Reading: The poster guide and a sample poster.

13 Poster session

20 April 2027

Groups present and discuss their posters in a gallery session. The class votes for the most surprising finding and the best example. We review the shared class experiment and reconsider the experiments from Meeting 1 using concepts developed during the course.

12. Teacher contact details

Name	Tom van Nuenen
Office location	To be confirmed
Email	To be confirmed
Teaching venue	To be confirmed
Website	Not applicable

13. Course website

Blackboard is used to share course resources and submit work.

14. Academic honesty and plagiarism

You must follow the University's policies and regulations on academic honesty. These include the disciplinary guidelines and procedures for breaches of academic honesty. Details are available at http://www.cuhk.edu.hk/policy/academichonesty/.

Each assignment must include a signed declaration confirming that you are aware of these policies, regulations, guidelines and procedures. The declaration is available at https://www.cuhk.edu.hk/policy/academichonesty/Eng_htm_files_(2013-14)/p10.htm.

- For group projects, every member should sign the declaration. Every member is responsible and may face disciplinary action, regardless of whether they signed it or contributed directly or indirectly to the problematic content.
- For assignments that are mainly text-based computer-generated documents submitted through VeriGuide, the system issues the declaration as a receipt when you upload the assignment.
- Your work may be investigated using AI content detection software to assess its originality.
- You must understand the course's approach to AI use. Where AI tools are allowed, acknowledge and cite their use as instructed by the course teacher.

Assignments without a properly signed declaration will not be graded. Submit only the final version through VeriGuide.

Submitting the same work, or part of it, for more than one purpose without declaring this is an undeclared multiple submission. This includes submitting work to meet the requirements of two different courses. Reusing a phrase or a sentence or two from your own work is common and acceptable, but substantial reuse is problematic. Obtain agreement from the relevant course teacher or teachers before submitting reused work.

CUHK owns the copyright in teaching materials produced by its staff and teachers, including lecture notes, assignments and examination questions. You may download these materials from University learning management systems, such as Blackboard, for your own educational use. You must obtain permission from the relevant staff member or teacher before distributing, sharing or copying them for a third party.

15. Use of generative AI in teaching, learning and assessment

Approach 2 — Use only with prior permission. You may use AI tools in some activities and assessments but not in others. This section sets out where AI use is allowed and how to acknowledge it.

Chatbot responses used as data. Chatbot responses are also material for analysis in this course. When you use them as data for an activity or assignment, keep the exact prompts and responses in your record. This record serves as your acknowledgement for those responses. Acknowledge other uses of AI, such as drafting or editing, as described below.

AI use in assessed work. Reading cards, the individual analysis check and the group presentation assess your own understanding. You may not use AI tools to help with these tasks. You may use AI tools for any part of the group poster, provided you record how you used them.

Responsibility for submitted work. You are the author of the work you submit and are responsible for every claim it contains, regardless of which tools helped you produce it.

Acknowledging AI support. Acknowledge every functional use of a generative AI tool. Cite the tool whenever you paraphrase, quote or incorporate content it generated, including text, images, data or other formats.

- An acknowledgement may take this form: I acknowledge the use of [name of AI tool, for example Gemini] to [describe the support, for example plan my essay, generate ideas or suggest examples of data collection instruments].
- When including AI-generated text, record the prompt and identify the tool and displayed model: Prompt: [Insert the prompt.] Tool and displayed model: [Insert the name.] Generated text: [Insert the text.]

If you are unsure whether a use is allowed, ask the course teacher and obtain permission where necessary. See the CUHK Library guide for advice on using AI information responsibly in assignments.

16. Use of AI tools for grading and provision of feedback on student work

Key Point: Teachers may use AI tools for grading and providing feedback on student work. When doing so, the following guidelines should be noted:

- teachers using AI tools to grade and/or provide feedback on student work in a course should allow their students to make use of AI tools in the same course, i.e. adopt Approaches 2 – 4;
- for summative assessment tasks, teachers are not permitted to use AI tools to grade and/or provide feedback on them, except for closed-ended answers or solutions, irrespective of the approach adopted;
- for formative assessment tasks, teachers may use AI tools to assist them to grade and/or provide feedback on those tasks (including open-ended and close-ended answers or solutions) if they allow students to make use of AI tools in learning activities and/or assessments, i.e. adopt Approaches 2 – 4 for the course. Teachers should scrutinize the AI-generated outputs before finalizing the grades and/or feedback to be given to students;
- teachers should inform students in advance that trusted AI tools will be used to provide feedback on and/or grade student work;
- clear assessment rubrics should be provided to the AI tools to ensure the quality of the AI-generated feedback; and
- if teachers choose to adopt AI tools not supported by the University in grading and providing feedback, students' written consent will be required before their work is input to these tools.

Use of AI tools for grading and provision of feedback on student work

Not applicable. AI tools will not be used to grade or provide feedback on student work in this course.

17. Course grade descriptors

Grade	Descriptor
A	Outstanding performance on all learning outcomes. The work combines and applies the course's principles or skills at a level beyond normal expectations, reaching standards typical of higher levels of study. This grade is reserved for exceptional work that exceeds the standard expected of most students and is expected to be achieved by only a small minority.
A-	Generally excellent performance on almost all learning outcomes. The work combines and applies the course's principles or skills in a way that fully meets expectations at this level and occasionally reaches standards typical of higher levels of study.
B+	Very good performance on all learning outcomes, or very good performance on some that compensates well for slightly weaker performance on others. Overall performance is substantial. The work applies the course's principles or skills well and comprehensively, comfortably meeting expectations at this level.
B	Good performance on all learning outcomes, or good performance on some that compensates for slightly weaker performance on others. Overall performance is substantial. The work applies the course's principles or skills in a more comprehensive way that meets expectations at this level.
B-	Somewhat good performance on almost all learning outcomes, or somewhat good performance on some that compensates for slightly weaker performance on others. Overall performance is substantial. The work applies the course's principles or skills comprehensively enough to meet expectations at this level.
C+	Very satisfactory performance on most learning outcomes. The work shows a somewhat sustained ability to apply the course's principles or skills and meets the basic requirements at this level.
C	Satisfactory performance on most learning outcomes. The work applies some of the course's principles or skills and meets the basic requirements at this level.
C-	Somewhat satisfactory performance on a number of learning outcomes. The work shows some ability to apply the course's principles or skills and meets the minimum basic requirements at this level.
D+	Barely satisfactory performance on a few learning outcomes. The work addresses the course inadequately, meeting basic requirements in only some areas. Other areas receive minimal attention and may contain only loosely relevant material.
D	Almost barely satisfactory performance on very few learning outcomes. The work addresses the course inadequately, meeting basic requirements in very few areas. Other areas receive very little attention and may contain only loosely relevant material.
F	Unsatisfactory performance on a number of learning outcomes, or failure to meet specified assessment requirements. The work fails to address the course and suggests that its requirements have not been understood. It misses the point of the task.