

The Essentials of AI for Life and Society: An AI Literacy Course for the University Community

**Joydeep Biswas*, Don Fussell, Peter Stone,
Kristin Patterson, Kristen Procko, Lea Sabatini, Zifan Xu**

The University of Texas at Austin
2317 Speedway, Stop D9500, Austin, TX 78712

Abstract

We describe the development of a one-credit course to promote AI literacy at The University of Texas at Austin. In response to a call for the rapid deployment of class to serve a broad audience in Fall of 2023, we designed a 14-week seminar-style course that incorporated an interdisciplinary group of speakers who lectured on topics ranging from the fundamentals of AI to societal concerns including disinformation and employment. University students, faculty, and staff, and even community members outside of the University, were invited to enroll in this online offering: *The Essentials of AI for Life and Society*. We collected feedback from course participants through weekly reflections and a final survey. Satisfyingly, we found that attendees reported gains in their AI literacy. We sought critical feedback through quantitative and qualitative analysis, which uncovered challenges in designing a course for this general audience. We utilized the course feedback to design a three-credit version of the course that is being offered in Fall of 2024. The lessons we learned and our plans for this new iteration may serve as a guide to instructors designing AI courses for a broad audience.

Introduction

The increased public availability of large language model-based chat tools prompted a burst of interest in artificial intelligence (AI) in the higher education community in 2023. Despite the wealth of information about AI available in the literature and through blog posts, most of it was geared toward a technical audience. News stories tended to focus on recent innovations with sweeping claims about how AI would change life and society, without providing background on how AI tools work and their limitations.

Responding to the interest in and concern about AI by faculty, staff, and students across disciplines at The University of Texas at Austin, the Dean of the College of Natural Sciences requested the rapid development of a new AI course targeting a non-technical audience. With a mere six weeks remaining until the start of the fall semester, we assembled a team that included three computer science faculty leads, a

project manager, and experts in instructional design and assessment to create a broadly accessible new course about AI. The institutional support and multidisciplinary team were critical to the swift course deployment.

In the years leading up to the development of this course, UT Austin had already built institutional structures that allowed rapid responses to emerging research and education needs. UT Austin recognized the need to understand and shape the sociotechnical challenges innate to developing and deploying AI-enabled systems in the real world. In 2019, UT Austin launched Good Systems¹, a research grand challenge to form human-AI partnerships that are beneficial to society. Good Systems comprises an interdisciplinary team of faculty, which promotes research in ethical AI use by leveraging its benefits while considering and mitigating harms or unintended consequences that may be caused by new technologies. Conveniently, Good Systems was well established when the course was proposed, and its members supported the computer science faculty in responding to the call. Another key source of institutional support was a college-level STEM education center, whose flexible structure, pedagogical and course-design expertise, and experience navigating complex institutional systems enabled rapid development and deployment of the course.

This class was designed to provide a gentle introduction to Artificial Intelligence: what it is, its capabilities, its limitations, and its potential implications for society. To promote accessibility, we proposed a 1-credit-hour course that students and faculty alike could easily fit into their schedules. The course was publicized through email, department newsletters, the campus newspaper, and social media. Undergraduate students enrolled in the course for credit, while graduate students, post-doctoral fellows, faculty, and staff participated as auditors.

The set of fourteen lectures on fundamental concepts for AI literacy, as well as the ethical and societal implications of AI technologies, were given by faculty from across the University. Recordings of these lectures are now available for anyone to enjoy and learn from at any time.² Only a curiosity and willingness to explore are needed - no prior technical knowledge is required.

*The author order includes the three co-instructors of the course in alphabetical order by last name, followed by the other co-authors, also in alphabetical order.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://bridgingbarriers.utexas.edu/good-systems>

²<https://www.cs.utexas.edu/~pstone/Courses/109fall23/>

Related Work

AI Literacy There have been a few attempts to create *AI Literacy* courses at universities. For example, Kong et al. (Kong, Cheung, and Zhang 2021) designed a seven-hour AI literacy course for university students from diverse backgrounds in Hong Kong.

AI in K-12 Education While there is broad interest in inclusion of AI at the K-12 level, there is a lack of consensus on what AI literacy means (Ng et al. 2021b), particularly for K-12 students. Early attempts by Ng et al. (Ng et al. 2021a) define “AI Literacy” as consisting of four components: 1) know and understand, 2) use, 3) evaluate, and 4) ethical issues. To provide guidance on how to create AI literacy curricula for K-12 students, Ng et al. (Ng et al. 2023) conducted a systematic review of the literature on AI literacy, producing pedagogical models, teaching tools, and challenges as a result. Williams (Williams 2023) also conducted a review of AI literacy in K-12 curricula, with a focus on assessments, yielding a set of recommendations for approaches to use, and existing gaps in assessment.

AI in University Curricula AI for computer science or engineering majors is a common offering at universities at the undergraduate and graduate levels. There have also been several recent AI majors introduced recently at the undergraduate (Carnegie Mellon University 2024) and graduate level (University of Texas at Austin 2024). There have also been several instances of integrating ethical considerations in AI courses (Vekhter and Biswas 2023; Saltz et al. 2019). *CPSC 170 – AI For Future Presidents* (Scassellati 2023) at Yale University is the closest previous analogue to our course, designed for all students, with no pre-requisites, and requiring no math or programming background.

Relation to Previous Work In this paper, we present the design and retrospective analysis from offering an AI literacy course to all members of a large public university, without assumptions of technical background. Unlike previous courses, this offering is simultaneously *technical* — it covers the main technical sub-disciplines that comprise AI; while including *ethical* considerations; and is aimed at a *broad audience, including non-technical majors*.

Course Design

The ambition of the course was to introduce AI to *anyone* with interest. To reach this large target audience, accessibility was crucial. Accordingly, this online course was delivered with support from the Liberal Arts Instructional Technology Services (LAITS) at UT Austin.

Curricular Objectives

The main learning objective of the course was to improve participants’ AI literacy. To support students’ achievement of this AI literacy objective, we introduced the types of AI and their applications, highlighted the risks and benefits of various AI technologies, discussed the societal impacts of AI technologies, and aimed to improve learners’ ability to

distinguish AI science from science fiction. The course faculty determined that background on concepts that are foundational to AI would be included in the first five lectures, and the foundational concepts would be connected to specific applications (computer vision, robotics, and natural language processing, including large language models) in the following three presentations (Table 1). The remaining lectures provided historical context, followed by an exploration of some of AI’s inherent dangers and limitations. Lecturers were primed to expect a wide range of backgrounds and experiences among the students, and were asked to assume no prior technical knowledge — as a consequence, where necessary, lecturers included introductory technical primers in the lecture content to support understanding.

Course Structure

Students were expected to attend weekly one-hour lectures synchronously, and recordings were made available afterwards to provide flexibility for faculty and staff whose real-time attendance may be challenging. A limited number of students could attend the lectures in-person in the studio.

The course requirements for students were to complete readings, attend the online lectures, and respond to in-class quiz questions on the course material. Weekly written reflections on the reading and lecture were also required; because we were developing this course concurrently with its delivery, we used these reflections as an opportunity to evaluate the course and make corrections. Grades were based on attendance (30 percent), correctness of responses on real-time quizzes (40 percent), and completion of weekly reflections (30 percent). Assignments were designed to be an incentive to participate, rather than to be evaluative — we determined early on that it would be acceptable, indeed desirable, for every student to earn an A.

Each lecture included about 30-40 minutes of presentation time, where one of the course faculty or a guest lecturer presented on a designated topic. We recruited experts from computer science to introduce the foundational concepts we had thoughtfully curated, and in the latter half of the course we identified guest speakers from diverse departments across UT Austin. The guest lecturers were instructed to tie their lectures back to a working definition of AI, but were free to structure the talk as they liked. Each speaker selected a required reading assignment to accompany their lecture, and was welcome to provide additional optional readings for students who wished to explore a topic more deeply.

Evaluations and Learning Assessments

Students were asked to complete the associated reading assignments prior to watching the lectures. To promote engagement in the synchronous online lecture, we incorporated live quiz questions. We originally interspersed the questions throughout the presentation via a polling system to encourage attentiveness in the largely remote audience; however, midway through the course we opted to change the format to provide a single integrated quiz following the lecture. This mitigated some technical difficulties and gave the speaker an opportunity to transition to a seat beside the host faculty member to engage in 5-10 minutes of Q&A

Topic	Lecturer (Department)
Introduction - AI100 study(Littman et al. 2022)	Peter Stone (Computer Science)
Planning and Search	Joydeep Biswas (Computer Science)
Probabilistic Modeling	Roberto Martín-Martín (Computer Science)
Machine Learning Fundamentals	Adam Klivans (Computer Science)
Machine Learning Paradigms	Ray Mooney (Computer Science)
Computer Vision	Kristen Grauman (Computer Science)
Intelligent Robotics	Luis Sentis (Aerospace Engineering /Engineering Mechanics)
Natural Language Processing (Large Language Models)	Greg Durrett (Computer Science)
Philosophical Foundations and Relation to Computing	Don Fussell (Computer Science)
AI and Mis/disinformation	Matt Lease (School of Information)
Bias and Fairness in AI Models / Elargethical Datasets	S. Craig Watkins (Journalism and Media)
Impacts on Workplace and Economics	Sherri Greenberg (LBJ School of Public Affairs)
AI Alignment and Existential Threats	Scott Aaronson (Computer Science)
Current and Future Directions	Joydeep Biswas, Don Fussell, Peter Stone (Computer Science)

Table 1: Course lecture schedule

about the presentation. Course participants submitted questions through a chat feature, with priority given to questions from students in the live studio audience.

Student Reactions

We analyze the effectiveness of the course in terms of 1) the background of the enrolled students, 2) responses to weekly surveys, and 3) a final course survey.

Enrollment

The course was successful in enrolling a broad audience. A total of 788 individuals signed up for the course, including 132 undergraduates who signed up for credit, 631 auditors from the university, and 25 external participants. Among the 631 university individuals that signed up to audit, all 17 colleges from UT Austin were represented. University-affiliated auditors included faculty, emeritus faculty, staff, postdoctoral fellows, graduate students, and alumni. Table 2 lists the college affiliations and standings of the university individuals who enrolled for the course.

Response to Weekly Surveys

The weekly reflections were a required assignment for students that was graded based on completion; auditors were encouraged, but not required, to complete weekly reflections. Each week, we collected over 100 student responses and usually 10–20 reflections from course auditors. Following a brief free-form essay reflecting on the course material, we integrated several rating scales to gauge the amount of time students required to complete the reading, their level of engagement with the topic, and prior understanding. While designing the course, we anticipated offering a 3-credit option in the future. Accordingly, we included one open-ended targeted question: “If this topic were expanded, what would you like to learn more about? In the space below, list any relevant subtopics you would like more information about.” In addition to using these responses to identify future subtopics to expand upon, we uncovered challenges students faced with some readings and course material. In general, student responses indicated a desire for more examples connecting the early lectures on fundamentals to specific applications

UNIT	UG	GS	P	F	S	T
Engineering	3	22	3	9	6	40
Education	1	3	0	28	33	64
Fine Arts	2	1	0	10	7	18
Liberal Arts	36	12	1	43	11	67
Natural Sciences	80	24	3	67	51	145
Pharmacy	0	1	1	1	0	3
Medical School	0	2	1	2	25	30
Geosciences	0	0	0	0	1	1
Public Affairs	0	0	0	2	0	2
Business	9	1	0	10	13	24
Communication	1	4	1	16	2	23
Architecture	0	1	0	9	6	16
Information	0	22	3	0	1	26
Law	0	1	0	4	1	6
Nursing	0	2	0	0	1	3
Social Work	0	0	0	3	7	10
UG College	0	0	0	0	8	8
Other	0	5	0	12	129	146
Total of Type	132	101	13	216	302	763

Table 2: Course enrollment by unit and standing: undergraduate students (UG), graduate students (G), postdocs (P), faculty (F), staff (S), and unit total (T).

in AI, which was largely limited by the single, brief lecture period. They also requested current news articles to connect topics to the most recent advances in the field. Some weeks, responses to the open-ended question revealed the reading selected by the guest speaker or the lecture content presented was too high-level for the broad audience.

Overall Course Survey

The UT Course Evaluation Survey (CES) includes a 1-5 Likert scale for the overall course, and the course instructor(s). Our course ratings were 4.23 and 4.33, respectively. For the same semester, the average ratings for all courses in the College of Natural Sciences were 4.00 and 4.20. In addition to the standard CES, we conducted an extensive evaluation to assess the course’s effectiveness at meeting its curricular objectives, which we present next. Unfortunately, unlike the CES, there are no baselines to compare the detailed evaluation results to, but we hope our results will serve as baselines for evaluations of future iterations at UT and/or elsewhere.

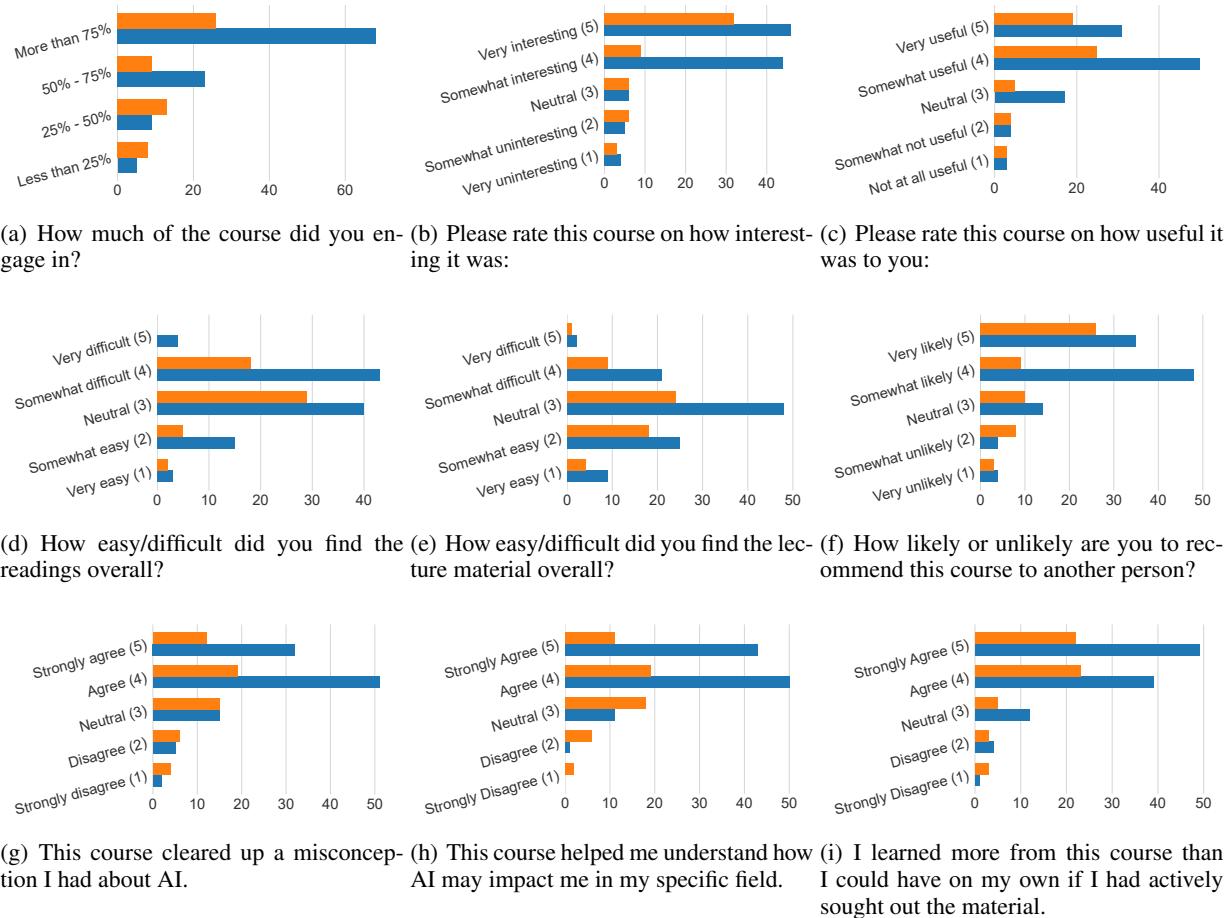


Figure 1: Summary of survey responses on engagement and interest, difficulty of course materials, and overall understanding. The blue and orange bars represent responses from students and auditors, respectively. Subtitles indicate the prompts corresponding to each survey question.

We replaced the final reflection essay with a targeted course survey which was part of the students' course grade. To recruit auditors to complete the survey, we sent several targeted messages to this group through the learning management system. We encouraged them to respond to the survey regardless of the amount of the course they had completed. We collected responses from 145 course participants, the majority (70 percent) of whom were undergraduate students. The remaining respondents were auditors: 12 percent were faculty, 11 percent were staff, and the remaining 7 percent were graduate students or indicated other roles at UT Austin. Among those internal to UT Austin, the majority (54 percent) were affiliated with the College of Natural Sciences, with an additional 24 percent from the College of Liberal Arts and 8 percent from the McCombs School of Business. Figure 1 summarizes the survey responses.

We first evaluated the data for course engagement. Gratifyingly, the majority of respondents (73 percent) indicated that they were likely to recommend the course to another person. Regarding how interesting participants found the course, 81 percent of respondents indicated that the course was "very interesting" or "somewhat interesting" (average

rating of 4.14 on a 5-point Likert scale where 1 = very uninteresting and 5 = very interesting). Perhaps unsurprisingly, students who were earning a grade in the course were more engaged than auditors: 87 percent of students reported that they engaged with more than half of the material, compared to 63 percent of auditors. We also evaluated engagement through attendance logs, which peaked in the first week at 387 live participants and in later weeks dropped to 170.

Open-ended questions were also used to evaluate the course. Response to the question: "What aspect(s) of the course did you find most useful?" were thematically coded (Auerbach and Silverstein 2003; Saldaña 2021) to uncover themes in participant responses. Of the 122 open-ended responses collected, 22 percent described the variety of guest speakers and broad overview. The lectures stood out as a useful learning tool with 37 percent of open-ended responses mentioning them. We also sought critical feedback through the open-ended question: "What aspect(s) of the course needs improvement?" Most of the content-related comments (19 percent of 115 open-ended responses) focused on the level and length of the reading, for example: "In my opinion, some of the reading was mathematically chal-

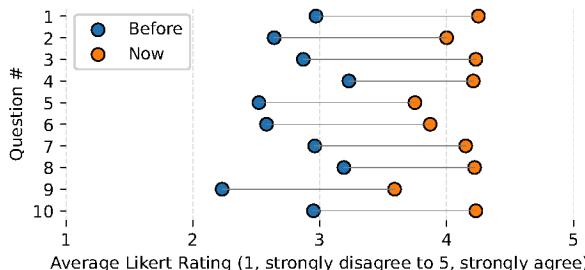


Figure 2: Retrospective pre-/post-survey questions. Each result is statistically significant, with a p -value < 0.01 , $n = 151$ for all questions, except for Q5 where $n = 150$.

lenging, so I would recommend adding prerequisite courses for students.” This was corroborated by quantitative analysis, where survey respondents were asked to rate the question: “How easy/difficult did you find the readings overall?” on a five-point Likert scale (1=very easy to 5=very difficult). Of the 159 responses, a majority (51 percent) selected “somewhat difficult” or “very difficult”. In an analogous question about the reading, only 30 percent rated the lecture material as difficult. Other course structure critiques emerging in the open-ended feedback (13 percent) focused on interactivity and specific assignments: “I would like to see an active student community, either making the lectures way more interactive, or implementing new and fun ideas for students to participate in.”

Our overarching goal for this course was to improve AI literacy for a broad audience. We utilized retrospective pre-/post survey (Geldhof et al. 2018; Howard and Dailey 1979) questions to examine this; through the course survey given at the end of the course, participants asked to rate their skill level before the course and now, after taking the course on a 5-point Likert scale (1 = strongly disagree and 5 = strongly agree). For each of ten questions probing participants’ AI literacy, we saw an increase in the rating after taking the course. The list of questions follows; in parenthesis following the question the change in rating is reported. Each reported change is statistically significant with a P -value < 0.01 and $n = 151$ for all questions (except Q5 where $n = 150$). Figure 2 summarizes the average Likert rating for each question in the survey.

1. I can define artificial intelligence (AI). (+1.27)
2. I have the necessary vocabulary to discuss AI. (+1.36)
3. I can list five examples of AI. (+1.37)
4. I can describe how AI affects my daily life. (+0.97)
5. My understanding of AI makes me well-equipped to apply it to my future professional work. (+1.24)
6. I am prepared to weigh in on the deployment of AI in products that affect me. (+1.29)
7. I can evaluate news stories about AI. (+1.19)
8. I can differentiate AI science from science fiction. (+1.03)
9. I am literate about the technical components of AI. (+1.36)
10. I am literate about the societal implications of AI. (+1.28)

Lessons Learned

Analysis of course feedback shed light on key considerations when designing an AI literacy course for a broad audience. Course participants appreciated the variety of speakers and connection to examples. Indeed, participants expressed a seemingly limitless desire for additional examples in the weekly reflection feedback. Interest in the integration of popular news articles connected to the weekly topics was noted frequently in the open-ended reflections as well. These types of articles could mitigate the main course flaw, which was that speaker-selected readings were often challenging for our non-technical audience.

Another key takeaway is that a broad audience from different backgrounds can be difficult to satisfy. Auditors reported less course engagement and some described having different expectations for the course (e.g., wanting to learn how to use AI tools, a greater focus on incorporating AI in teaching). Some staff members with limited experience using the University learning management system reported difficulty navigating and finding course materials. Gratifyingly, despite these challenges, the audience that participated in the final course survey improved their AI literacy.

Summary Of Lessons Learned

Engaging students and a wide variety of speakers can pose challenges for an AI literacy course. Enrollment in the course included 131 students and 584 auditors, which included faculty, staff, postdoctoral fellows, graduate students, and participants external to The University of Texas at Austin. All UT Austin colleges were represented by auditors that enrolled in the course. Surveyed auditors reported less engagement in the course than students and some expressed having different expectations for the course.

Course participants appreciated the variety of speakers and connection to examples—the more the better! 22% of the 122 open-ended responses described the variety of guest speakers and broad overview. “I loved that we heard from different guest speakers every week; it was great to get to see so many experts in so many different fields. I also found the course content to be sufficiently reflective of the instructors’ thesis that the AI problem is an interdisciplinary problem, not just a computer science problem as is commonly thought of.” Some also highlighted value of connections to examples: “Anytime the material was tied into real life examples or scenarios I found it very helpful and applicable. I also liked hearing from speakers with knowledge of different areas of AI.” The lectures stood out as a useful learning tool with 37% of open-ended responses mentioning them: “I liked the lectures and content about how AI works at a fundamental level and what we need to be aware of as we start integrating its use in our real lives.”

The readings were the most challenging part of the course for our non-technical audience. In response to the open-ended survey item: “What aspect(s) of the course needs improvement?”, most of the content-related comments (19% of 115 open-ended responses) focused on the level and length of the reading: “In my opinion, some of the

reading was mathematically challenging, so I would recommend adding prerequisite courses for students.” We provide our reading list online,³ but based on the feedback recommend revising it to include more introductory materials.

Creating an engaged student community in a one-credit course was challenging. Other course structure critiques (13%) focused on interactivity and specific assignments. We address our plans for increased interactivity below.

Future Plans

Based on the feedback from this first version of the course, we further developed it to a 3-credit-hour course, first offered in fall of 2024⁴. The expanded course covers the same set of topics and remains geared toward a broad audience, with no technical background. Due to the challenges in satisfying the expectations of both students and a broad pool of auditors, we designed this course exclusively for students.

We examined the readings from the first iteration of the course, with non-computer science course staff reviewing them and making recommendations for modification and/or replacement. We replaced several technical readings from journals with approachable blogs and popular news articles. The readings were read and discussed by groups of students using a collaborative annotation tool to foster asynchronous discussion and a sense of community.

The three-credit course includes weekly asynchronous modules with readings, videos, and assignments before an additional synchronous session each week. During the synchronous class, the instructor provides additional context, and leads discussions about the material. Each student attends class in the studio at least once during the semester, serving as the spokesperson for a group of remote students with whom they engage in discussion through an online chat. A series of (non-programming) assignments familiarizes students with the capabilities and limitations of some AI-based tools and technologies. A significant writing assignment focused on the societal and ethical implications of AI guides students to consider its real-world applications. Accordingly, the course satisfies an ethics requirement for UT Austin.

Acknowledgments

This study was evaluated and considered exempt from the IRB of the University of Texas at Austin and was conducted in accordance with relevant ethical guidelines. This work was supported in part by the National Science Foundation (CAREER-2046955, NRT-2125858, FAIN-2019844), and UT Austin’s Good Systems grand challenge. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors. Peter Stone serves as the Chief Scientist of Sony AI and receives financial compensation for that role. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

³<https://www.cs.utexas.edu/~pstone/Courses/109fall23/assignments.html>

⁴<https://www.cs.utexas.edu/~pstone/Courses/309fall24/>

References

Auerbach, C.; and Silverstein, L. B. 2003. *Qualitative data: An introduction to coding and analysis*, volume 21. Carnegie Mellon University. 2024. B.S. in Artificial Intelligence. <https://www.cs.cmu.edu/bs-in-artificial-intelligence/>. Accessed: 2024-09-16.

Geldhof, G. J.; Warner, D. A.; Finders, J. K.; Thogmartin, A. A.; Clark, A.; and Longway, K. A. 2018. Revisiting the utility of retrospective pre-post designs: the need for mixed-method pilot data. *Evaluation and program planning*, 83–89.

Howard, G. S.; and Dailey, P. R. 1979. Response-shift bias: A source of contamination of self-report measures. *Journal of Applied Psychology*, 64(2): 144.

Kong, S.-C.; Cheung, W. M.-Y.; and Zhang, G. 2021. Evaluation of an artificial intelligence literacy course for university students with diverse study backgrounds. *Computers and Education: Artificial Intelligence*.

Littman, M. L.; Ajunwa, I.; Berger, G.; Boutilier, C.; Currie, M.; Doshi-Velez, F.; Hadfield, G.; Horowitz, M. C.; Isbell, C.; Kitano, H.; Levy, K.; Lyons, T.; Mitchell, M.; Shah, J.; Sloman, S.; Vallor, S.; and Walsh, T. 2022. Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report. arXiv:2210.15767.

Ng, D. T. K.; Lee, M.; Tan, R. J. Y.; Hu, X.; Downie, J. S.; and Chu, S. K. W. 2023. A review of AI teaching and learning from 2000 to 2020. *Education and Information Technologies*, 28(7): 8445–8501.

Ng, D. T. K.; Leung, J. K. L.; Chu, K. W. S.; and Qiao, M. S. 2021a. AI literacy: Definition, teaching, evaluation and ethical issues. *Proceedings of the Association for Information Science and Technology*, 58(1): 504–509.

Ng, D. T. K.; Leung, J. K. L.; Chu, S. K. W.; and Qiao, M. S. 2021b. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*.

Saldaña, J. 2021. The coding manual for qualitative researchers.

Saltz, J.; Skirpan, M.; Fiesler, C.; Gorelick, M.; Yeh, T.; Heckman, R.; Dewar, N.; and Beard, N. 2019. Integrating Ethics within Machine Learning Courses. *ACM Trans. Comput. Educ.*, 19(4).

Scassellati, B. 2023. AI for Future Presidents (CS170). <https://zoo.cs.yale.edu/dsac/blog/2023/12/19/cpsc-170/>. Accessed: 2024-09-16.

University of Texas at Austin. 2024. M.S. in Artificial Intelligence. <https://cdso.utexas.edu/msai>. Accessed: 2024-09-16.

Vekhter, J.; and Biswas, J. 2023. Responsible robotics: a socio-ethical addition to robotics courses. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15877–15885.

Williams, R. 2023. A Review of Assessments in K-12 AI Literacy Curricula. https://randi-c-dubs.github.io/K12-AI-ed/Constructionist_AI_Assessments.pdf. Accessed: 2024-09-16.