

DOI: https://doi.org/10.48009/1_iis_2026_130

How artificial intelligence is taught today: A mixed-methods analysis of introductory AI courses

Samuel Yudhistira, *Penn State University*, spy5259@psu.edu

Alan R. Peslak, *Penn State University*, arp14@psu.edu

Abstract

To gain a better understanding about how the rapidly growing field of artificial intelligence is currently being taught in universities, 45 introductory AI course syllabi were analyzed. A mixed-methods strategy using quantitative frequency count analysis and qualitative thematic coding analysis was primarily utilized. Among other findings, the results showed that machine learning was the most taught topic across courses, followed by neural networks and large language models. The most common pedagogy utilized across the various disciplines was lecture-heavy instruction. Ethics-related content was present in many syllabi but varied widely in depth and structure. These findings highlight the key trends and patterns in how AI is taught in universities and the growing need for AI-related skills.

Keywords: artificial intelligence, AI education, curriculum analysis, pedagogy, ethics, machine learning, syllabi

Introduction

The development of artificial intelligence, or AI for short, has grown exponentially in recent years. This has led to a widespread use and adoption of AI by industries across the world. Some notable fields that have incorporated and adopted AI are healthcare, business, education, and manufacturing. In business, 78% of organizations were reported to be using AI in 2024, which was a 55% increase from the prior year (Stanford HAI, 2025). AI has become increasingly more integrated in daily life due to these developments. This can be seen through various advancements like AI chatbots, assistants, and transportation. With the demand for AI currently being so pertinent, it is imperative that people know the skills to properly develop and maintain it. This is where education remains critical for AI.

In response to the demand for AI skills and knowledge, education has added a wide variety of AI-related courses. Most of these courses are in universities and colleges; however, high schools have also started to add AI courses as well (Ng, Su, Leung, and Chu, 2024). AI degrees and programs have even been added to a rapidly growing list of universities around the world. According to an AI degree report conducted by Degree Prospects (2026), there has been a total of 194 new AI master's programs in universities in the United States from 2022 to 2025, which is an increase of 167%. Additionally, bachelor's programs have more than doubled from 2024 to 2025, from 90 to 193 programs (Degree Prospects, 2026). Based on this response in education, it is evident that AI is still growing as a field and will continue to do so in the future.

Generally, AI courses can be separated into a few different categories. These categories are the tech, applied, and general categories. The tech category involves disciplines like computer science and

engineering and comprises the majority of AI courses. This is largely in part due to the coding and tech knowledge required to create and develop AI systems. The applied category involves the application of AI in disciplines like business, health, and humanities. These courses focus on how to use AI to solve problems in the specific disciplines. The last category is the general one which covers general education courses about AI. These courses aim to teach a basic understanding of the key concepts of AI.

The main objective of this study is to assess how AI is taught in contemporary university settings. Specifically, this is done by analyzing the syllabi of numerous introductory AI courses. The analysis is conducted using a determined coding framework to see if certain topics are taught and to determine what tools are used. The pedagogy and the mention of ethics are also included in the codebook. To accomplish this objective, three research questions were established:

RQ1: *What AI topics are most taught in university introductory AI courses? This question is integral in reaching a proper conclusion about how exactly AI is taught today due to directly covering the content of the courses.*

RQ2: *How do pedagogical approaches differ across disciplines? This question assesses in what way AI topics are presented and the relationships between different disciplines.*

RQ3: *How is ethics addressed in introductory AI courses? This question examines whether AI courses teach a significant value present in common problems with AI.*

Literature Review

In order to properly acknowledge and assess how AI is taught in universities, the growth of AI in education must first be examined. Yang, Qin, Lei, and Liu (2023) conducted a study on artificial intelligence courses and curriculum from a global view. They established that based on tracking the number of papers issued for AI courses from 2013 to 2022, there are three development stages of AI curriculum. These stages are the slow development stage (2013 ~ 2016), rapid development stage (2017 ~ 2020), and mature development stage (2021 ~ present). The papers issued during the mature development stage account for the majority of total papers, with 2021 accounting for 27.21% and 2022 accounting for 35.73% (Yang et al., 2023). This displays the global impact of AI in education around the world and precisely shows the growing demand for AI-related skills.

While it is important to learn about AI's growth in education, identifying the topics taught is crucial as well. Yang et al. (2023) additionally created a density diagram to identify the popular topics in AI courses based on keywords. The keywords with the most results were machine learning and education. The rest were then represented into lower color gradients, including deep learning, prediction, and curriculum. Surrounding these were more specific keywords such as neural networks, models, and reinforcement learning (Yang et al., 2023).

The dominance of machine learning as a keyword shows how focused AI's roots in technical topics are and how it is a vital component of AI. Zhou, Tong, Lan, Zheng, and Zhan (2021) have found similar results that reinforce that technical topics dominate. Particularly, it was core algorithms and typical applications they found that stood as the main content of AI education in massive open online courses (Zhou et al., 2021).

In understanding AI education, it is also important to properly look at the teaching methods and pedagogy of how AI is taught in general. Findings based on massive open online courses found that the most used

methods were online procedural teaching at 47.56%, task-driven learning at 20.73%, and online collaborative learning at 17.89% (Zhou et al., 2021). The study also acknowledged that scenario-based learning and project-based learning are not fully adopted due to a more demanding teaching environment being needed to use them more effectively (Zhou et al., 2021). These findings suggest that learning AI from these courses is largely structured and task-oriented but can differ greatly from AI courses in universities.

For trends in another context of AI education, in an analysis of K-12 AI education, interactive learning comprised 100% of the lessons (Wu, Chen, Chen, and Liu, 2024). Inquiry-based learning and collaborative learning were next at 77.55% and 63.27% respectively (Wu et al., 2024). This emphasizes that the interactive and hands-on approach dominates K-12 AI education. While these are not representative of AI course pedagogy in universities, they still contribute as an insightful look into broad trends from different perspectives.

Building upon these broad trends, AI course pedagogy in universities has been shown to prioritize collaborative project-based and problem-based learning (Ng, Lee, Tan, Hu, Downie, and Chu, 2023). In articles reviewed from 2000 to 2020, Ng et al. (2023) discovered that 61.2% included collaborative curricula designs, and 85.7% had a project-based learning approach. These would incorporate technological teaching tools with robotics or game elements. The universities focused on developing students' interests, abilities, and social skills with these in mind (Ng et al., 2023). The results from this study indicate that in a university setting, students can actively engage in learning AI through these pedagogies.

Furthermore, this is supported by a study about pedagogical approaches in teaching AI courses conducted in Bangladesh (Rufai, Khan, and Hasan, 2024). This study was conducted through interviews with undergraduate students about their experiences with different teaching methods and pedagogies. Various pedagogies, such as a project-oriented approach and practical implementations of concepts, were asked about and considered effective. They concluded that effective AI instruction and pedagogy involve the integration of theory and practice, the value of feedback, self-paced learning, real-life problem connections, and future-oriented education (Rufai et al., 2024). These findings connect and contribute towards enhancing student engagement in learning AI.

In addition to topics and pedagogy, AI ethics has been extensively studied due to its implications for privacy and bias. AI can allow access to private information depending on the dataset, and biased algorithms can lead towards poor outcomes that reinforce negative decisions and inequalities (Fiesler, Garrett, Beard, 2020). Because of these risks, it is important to teach ethics in AI courses. However, a study of how AI ethics is taught in top U.S. CS universities revealed that ethics is heavily underrepresented. The study found that out of 955 AI-related courses, only 83 courses incorporated AI ethics content (Aljedaani and Parthasarathy, 2026). Most courses within the 83 gave one or two sessions to AI ethics topics such as algorithmic bias, data privacy, transparency, and social impact. The study concluded by providing suggestions and a framework for AI ethics in AI-related courses (Aljedaani and Parthasarathy, 2026). This indicates that while ethics is considered important, it is often treated as an add-on in AI education.

Altogether, current studies reveal that while the rapid growth of AI in education is responding to meet the demand of AI-related skills, core AI topics are still rooted in technical topics. Pedagogical approaches in universities focus on projects and collaboration, while other approaches generally vary across educational platforms and contexts. Ethics in AI-related courses appears to be neglected despite having a significant impact on the real world. The trends that these studies highlight provide a background and foundation to examine how exactly AI is taught across universities.

Methodology

This study utilized a mixed-methods strategy to assess how AI is taught in universities. To gather the data, 45 syllabi of introductory AI courses in universities across the world were collected and analyzed. The majority of these syllabi were publicly sourced and collected from official university websites or the course instructors' websites. The selected syllabi were also limited to terms within the past three years to get proper representation of the current state of AI education. An Internet search was performed to find as many detailed and publicly available syllabi as possible over a period of several months. The 45 syllabi included contained sufficient information for analysis. The courses found were composed of various disciplines with the majority being computer science introductory AI courses due to being the most prominent. Introductory AI courses offered in business and healthcare are from other disciplines that were gathered in lower quantities. Given the rapidly changing nature of artificial intelligence tools and methods, the syllabi may not fully represent current classroom practices and topics but reflected current practice at time of manuscript development.

Three research questions were established focusing on the topics taught, the pedagogy used, and the inclusion of ethics to accomplish the main objective of assessing AI education in universities. These questions shaped the foundation of the codebook that was used to code and analyze the syllabi.

RQ1: *What AI topics are most taught in university introductory AI courses?*

RQ2: *How do pedagogical approaches differ across disciplines?*

RQ3: *How is ethics addressed in introductory AI courses?*

In order to make the analysis more streamlined, disciplines were also categorized into tech, applied, and general. Categories were developed inductively through content analysis of the syllabi by the authors. The resulting scheme is consistent with distinctions in the AI literacy literature, which differentiates curricula by their emphasis on technical mechanics versus conceptual and sociocultural understanding (Gu & Ericson, 2025); our 'applied' category further reflects the professional/vocational orientation distinction common in computing education (e.g., the IS-CS divide; Long & Magerko, 2020 on literacy competencies). Out of the 45 courses, 30 tech, 12 applied, and 3 general courses were collected and analyzed.

Specific data collected from the syllabi analyzed quantitatively were predefined content topics, pedagogy, tools, and ethics. Course description, topics covered, and learning objectives were analyzed qualitatively using thematic analysis.

To assess the statistical significance of observed differences, several inferential tests were applied to the coded data. For single-topic prevalence, exact binomial tests were used to evaluate whether each topic's frequency differed significantly from 50% ($H_0: p = 0.50$, two-tailed). Cochran's Q test assessed whether the nine content topics were covered at significantly different rates. For comparisons across the three discipline categories, chi-square tests of independence were computed; Fisher's exact test was substituted when expected cell counts fell below five. Effect sizes were estimated using Cramér's V. For topic breadth and qualitative theme scores, Kruskal-Wallis H tests were applied. Pairwise topic co-occurrence was evaluated using phi (ϕ) coefficients. All tests used $\alpha = .05$.

Results

As mentioned previously in the methodology section, 45 syllabi of introductory AI courses in universities of varying disciplines were analyzed. The disciplines were generalized into three main categories: tech,

applied, and general. The analysis primarily followed the created research questions about topics, pedagogy, and ethics to provide key insights as to how AI was taught in universities. This section was conducted using quantitative analysis.

To examine the first research question about the presence of topics in introductory AI courses, frequency counts were established. Table 1 shows the distribution of topics in the 45-course dataset.

Table 1. Frequency Counts of AI Topics

Topic	Frequency	Percentage
Machine Learning	41	91.1%
NLP/LLMs	29	64.4%
Knowledge Representation & Reasoning	28	62.2%
Ethics	27	60.0%
Neural Networks	27	60.0%
Search	26	57.8%
Applications	22	48.9%
Planning/Scheduling	20	44.4%
Bayesian Networks	10	22.2%

From Table 1, it can be seen that machine learning was the most frequent topic in these courses, with a noticeable gap compared to the other topics. This shows that machine learning is typically the core of introductory AI courses. The other topics were in a similar range together except Bayesian networks which appeared the least by a noteworthy margin. A visual depiction of the data is represented in Figure 1.




Figure 1. Frequency Counts of AI Topics (n = 45)

To assess whether these frequencies reflect statistically meaningful patterns, exact binomial tests were applied to each topic ($H_0: p = 0.50$). The results are presented in Table 1a. Cochran's $Q(8) = 49.53$, $p < .001$, confirming that the topics are taught at significantly unequal rates across the dataset.

Table 1a. Binomial Significance Tests for Topic Prevalence (n = 45)

Topic	n	%	Binomial p
Machine Learning	41	91.1%	$p < .001$ ***
NLP/LLMs	29	64.4%	$p = .073$ n.s.
Knowledge Representation & Reasoning	28	62.2%	$p = .135$ n.s.
Ethics	27	60.0%	$p = .233$ n.s.
Neural Networks	27	60.0%	$p = .233$ n.s.
Search	26	57.8%	$p = .371$ n.s.
Applications	22	48.9%	$p = 1.000$ n.s.
Planning/Scheduling	20	44.4%	$p = .552$ n.s.
Bayesian Networks	10	22.2%	$p < .001$ ***

Note. H_0 : $p = .50$, two-tailed. *** $p < .001$. n.s. = not significant. Cochran's $Q(8) = 49.53$, $p < .001$.

To help assist with the second research question about pedagogy, five different pedagogies were tracked. These results are shown in Figure 2.




Figure 2. Frequency Counts of Pedagogical Approaches (n = 45)

Based on Figure 2, it is evident that lectures dominate as a pedagogy in introductory AI courses. This may be due to the ease and effectiveness of lectures in large classroom settings. The next pedagogies, project-based and tool-driven, were also frequently used. Math-heavy and case-study-based followed them with the latter having the least.

To analyze the third research question about ethics, three separate coding categorizations were made. These considered whether the syllabi in the courses analyzed had ethics listed as a dedicated module and topic, only briefly mentioned it, or did not mention ethics at all. The results of this can be found in Table 2.

Table 2. Frequency Counts of Ethics Categories

Ethics Categories	Frequency	Percentage
Dedicated ethics module	20	44.4%
Not mentioned	18	40.0%
Brief mention	7	15.6%

From the results displayed in Table 2, it can be observed that there was an ethics module in approximately half of the courses analyzed. Similarly, with a comparable proportion, ethics was not mentioned anywhere in 40% of the syllabi. In approximately 15.6% of the course syllabi, ethics was only briefly mentioned.

Tool usage was also tracked for further insights. Table 3 represents the results for the presence of specific tools in intro to AI courses.

Table 3. Frequency Counts of Tools in AI Courses

Tools	Frequency	Percentage
Python	22	48.9%
ChatGPT/LLM APIs	12	26.7%
Other tools	12	26.7%
No tools listed	10	22.2%
Scikit-learn	3	6.7%
TensorFlow/PyTorch	3	6.7%

As shown in Table 3, the most frequently used tool in these courses was Python. ChatGPT and other LLM APIs were the next most common, representing generative AI as a tool for these courses. A portion of courses notably did not have tools mentioned. The least mentioned tools were Python libraries scikit-learn, TensorFlow, and PyTorch.

To further develop and answer the research questions, cross-tabulations were conducted comparing discipline with both topics and pedagogy of the courses. The cross-tabulation of disciplines and topics is shown in Table 4, while the cross-tabulation between disciplines and pedagogies is shown in Table 5.

Table 4. Cross-Tabulation of Disciplines and Topics

Topic	Tech	Applied	General
Machine Learning	28	10	3
NLP/LLMs	16	11	2
Knowledge Representation & Reasoning	25	2	1
Neural Networks	19	8	0
Search	25	0	1
Ethics	13	11	3
Applications	7	12	3
Planning/Scheduling	18	1	1
Bayesian Networks	10	0	0

The data from Table 4 exhibits the difference in topics taught based on the discipline. Technical courses prioritize teaching AI topics such as machine learning, knowledge representation, and search. While applied courses also have machine learning as a core topic, they have a higher emphasis on teaching ethics and real-world applications. General courses lack more technical topics but still cover key AI topics. Notably, the majority of search and Bayesian networks topics coverage were in the technical courses.

To determine which of these observed differences are statistically significant, chi-square tests of independence were applied to each topic. Table 4a presents these results, sorted by effect size.

Table 4a. Chi-Square Tests for Discipline × Topic (n = 45)

Topic	Tech %	Applied %	General %	Sig.
Applications	23.3%	100.0%	100.0%	p < .001 ***
Search	83.3%	0.0%	33.3%	p < .001 ***
Knowledge Rep. & Reasoning	83.3%	16.7%	33.3%	p < .001 ***

Topic	Tech %	Applied %	General %	Sig.
Ethics	43.3%	91.7%	100.0%	p = .001 **
Planning/Scheduling	60.0%	8.3%	33.3%	p = .004 **
NLP/LLMs	53.3%	91.7%	66.7%	p = .046 *
Bayesian Networks	33.3%	0.0%	0.0%	p = .019 *
Neural Networks	63.3%	66.7%	0.0%	n.s.
Machine Learning	93.3%	83.3%	100.0%	n.s.

Note. *** $p < .001$, ** $p < .01$, * $p < .05$, n.s. = not significant. Fisher's exact test used when expected cell counts < 5.

Table 4a confirms that five of the nine topics differ significantly across discipline categories. The largest effects are for Applications ($\chi^2 = 23.52$, $p < .001$) and Search ($\chi^2 = 24.09$, $p < .001$), representing near-complete disciplinary segregation. Knowledge Representation and Reasoning ($\chi^2 = 17.06$, $p < .001$), Ethics ($\chi^2 = 10.42$, $p = .001$), and Planning/Scheduling ($\chi^2 = 8.82$, $p = .004$) also show significant differences. Machine Learning and Neural Networks show no significant differences, reflecting their near-universal inclusion regardless of discipline.

Table 5. Cross-Tabulation of Disciplines and Pedagogies

Pedagogy	Tech	Applied	General
Lecture-heavy	30	11	3
Project-based	23	12	2
Tool-driven	23	10	3
Math-heavy	25	1	0
Case-study-based	1	10	2

The cross-tabulation from Table 5 reveals that lectures are one of the primary pedagogies in all the disciplines gathered. In technical courses, the project-based, tool-driven, and math-heavy pedagogies also play a prominent role. Applied courses are primarily project-based with the other pedagogies next to them except math-heavy. General courses follow a similar pattern with the applied courses.

Chi-square tests were also applied to the pedagogy cross-tabulation. As shown in Table 5a, the results reveal a sharp disciplinary divide for two pedagogies while confirming the remaining three are statistically universal.

Table 5a. Chi-Square Tests for Discipline $\times$ Pedagogy (n = 45)

Pedagogy	Tech %	Applied %	General %	χ^2 / Sig.
Case-Study-Based	3.3%	83.3%	66.7%	$\chi^2 = 28.61$, $p < .001$ ***
Math-Heavy	83.3%	8.3%	0.0%	$\chi^2 = 24.09$, $p < .001$ ***
Lecture-Heavy	100.0%	91.7%	100.0%	n.s.
Project-Based	76.7%	100.0%	66.7%	n.s.
Tool-Driven	76.7%	83.3%	100.0%	n.s.

Note. *** $p < .001$, n.s. = not significant.

Case-study-based instruction shows the largest effect in the entire dataset ($\chi^2 = 28.61$, $p < .001$), with 83.3% of Applied courses using it versus only 3.3% of Tech courses. Math-heavy instruction shows an equally strong and opposite pattern ($\chi^2 = 24.09$, $p < .001$): present in 83.3% of Tech courses but only 8.3% of Applied courses. Lecture-heavy, project-based, and tool-driven pedagogies show no statistically significant differences across disciplines, confirming their role as universal instructional strategies.

To examine the distribution of ethics treatment across discipline categories, the three-category ethics coding was cross-tabulated with discipline and tested for independence. Table 6 presents the full contingency table.

Table 6. Ethics Treatment by Discipline Category (n = 45)

Discipline	Ethics Module	Brief Mention	Not Mentioned	Total
Tech (n=30)	9 (30.0%)	4 (13.3%)	17 (56.7%)	30
Applied (n=12)	9 (75.0%)	3 (25.0%)	0 (0.0%)	12
General (n=3)	2 (66.7%)	0 (0.0%)	1 (33.3%)	3
Total	20 (44.4%)	7 (15.6%)	18 (40.0%)	45

Note. $\chi^2(4) = 13.53$, $p = .009$, Cramér's $V = .39$.

The chi-square test is statistically significant ($\chi^2(4) = 13.53$, $p = .009$, $V = .39$, medium-to-large effect). Notably, 56.7% of Tech courses make no mention of ethics at all, while 75.0% of Applied and 66.7% of General courses include a dedicated ethics module.

Following the quantitative section, a thematic analysis with thematic coding was conducted to assess the syllabi qualitatively. Several key themes were identified by analyzing the course description, topics list, and learning objectives. The broad themes were concluded to be technical, application, hands-on/practical, conceptual, and ethics. Courses were given a value of 0, 0.5, or 1 in order to evaluate how prominent the themes were in the evaluated text. These results are shown in Figure 3 below.




Figure 3. Theme Prominence by Discipline Category (n = 45)

The most prominent theme found was the technical theme, found in essentially every course. The next most significant theme was the hands-on or practical one, centered on opportunities to learn content through projects, exercises, and coding assignments. Conceptual was the following theme, involving teaching AI theory, history, and core concepts. Common phrases for courses to be defined as conceptual were “foundational principles of AI” and “history of AI.” The theme of application focused on AI in real-world contexts, represented by phrases like “real-world application.” Finally, the theme of ethics was identified as the least common; however, it was still present in 60% of the courses. This theme was focused on

considerations regarding ethical practices, identified by indicators such as “bias,” “fairness,” or “responsible AI.”

Kruskal-Wallis tests were applied to each of the five theme scores to assess whether theme prominence differed significantly across discipline categories. Results are presented in Table 7.

Table 7. Kruskal-Wallis Tests for Theme Scores by Discipline Category

Theme	Tech Mean	Applied Mean	Gen. Mean	Kruskal H	p / Sig.
Technical	1.000	0.917	1.000	5.46	p = .065 n.s.
Hands-On/Practical	0.759	1.000	0.667	4.198	p = .123 n.s.
Conceptual	0.810	0.792	1.000	0.864	p = .649 n.s.
Application	0.638	1.000	0.333	7.873	p = .020 *
Ethics	0.448	0.833	1.000	7.266	p = .026 *

Note. Theme scores: 0 = absent, 0.5 = partial, 1 = present. * $p < .05$, n.s. = not significant.

Two of the five themes differ significantly across discipline categories. The Application theme is significantly higher in Applied courses (mean = 1.000) than in Tech (mean = 0.638) or General (mean = 0.333) courses ($H = 7.87$, $p = .020$). The Ethics theme follows the same pattern ($H = 7.27$, $p = .026$). The Technical, Hands-On/Practical, and Conceptual themes show no statistically significant differences across categories, consistent with their near-universal presence across all course types.

To examine structural relationships among content topics, phi (ϕ) coefficients were computed for all pairwise topic combinations and tested for significance. Fourteen of the 36 pairs showed statistically significant co-occurrence patterns, presented in Table 8.

Table 8. Significant Topic Co-Occurrence Patterns (ϕ Coefficients, $n = 45$)

Topic Pair	ϕ (phi)	Significance
KRR $\times$ Search	+0.726	p < .001 ***
Search $\times$ Applications	-0.694	p < .001 ***
KRR $\times$ Applications	-0.613	p < .001 ***
Search $\times$ Planning	+0.584	p < .001 ***
NLP $\times$ Ethics	+0.531	p < .001 ***
Planning $\times$ Applications	-0.517	p < .001 ***
KRR $\times$ Planning	+0.512	p < .001 ***
Bayesian $\times$ Ethics	-0.436	p = .003 **
Ethics $\times$ Applications	+0.436	p = .004 **
ML $\times$ Neural Networks	+0.383	p = .010 *
NLP $\times$ Applications	+0.355	p = .017 *
Bayesian $\times$ Search	+0.349	p = .019 *
Search $\times$ Ethics	-0.331	p = .027 *
Bayesian $\times$ Applications	-0.309	p = .038 *

Note. $\phi > 0$ indicates topics appear together more than chance; $\phi < 0$ indicates mutual exclusivity. *** $p < .001$, ** $p < .01$, * $p < .05$.

The co-occurrence results reveal two distinct curricular clusters. A classical AI cluster in which Knowledge Representation & Reasoning, Search, and Planning co-occur strongly (all $\phi \geq .51$, $p \leq .001$), and an applied/modern cluster in which NLP, Ethics, and Applications co-occur ($\phi \geq .36$). The two clusters are negatively correlated with each other, confirming two structurally distinct approaches to introductory AI curriculum. A visual depiction of the cluster structure is shown in Figure 4.




Figure 4. Significant Topic Co-Occurrence Patterns: Two Distinct Curricular Clusters (n = 45)

Note. Solid lines indicate positive co-occurrence ($\phi > 0$); dashed lines indicate mutual exclusivity ($\phi < 0$). Line weight is proportional to $|\phi|$. *** $p < .001$, ** $p < .01$, * $p < .05$.

A final analysis examined topic breadth, defined as the number of distinct content topics covered per course. Table 9 presents the mean and median breadth for each discipline category along with the Kruskal-Wallis test result.

Table 9. Topic Breadth by Discipline Category (n = 45)

Group	n	Mean Topics	Median	Kruskal-Wallis
Tech	30	5.37	6.0	H = 3.12
Applied	12	4.58	5.0	p = .21, n.s.
General	3	4.67	4.0	
Overall	45	5.11	5.0	

Note. Mann-Whitney U post-hoc comparisons: Tech vs. Applied $U = 241.5$, $p = .081$; Tech vs. General $U = 55.0$, $p = .542$; Applied vs. General $U = 20.0$, $p = .818$. All non-significant.

The Kruskal-Wallis test is not statistically significant ($H = 3.12$, $p = .21$), indicating that discipline category does not predict the number of topics a course covers. Tech courses average 5.37 topics, Applied courses 4.58, and General courses 4.67, with no pairwise comparison reaching significance. This is a substantively important null finding: while the preceding analyses demonstrate that Tech and Applied courses cover fundamentally different sets of topics, they do so at comparable levels of breadth. In other words, discipline predicts which topics are taught, not how many.

Overall, the results from the quantitative and qualitative analyses provided insights into how AI is taught in universities through different methods. In the following discussion section, these results are interpreted based on the previously stated research questions and current AI education.

Discussion

This study primarily aimed to assess and examine how AI is taught across universities in the current times through analyzing various aspects of introductory AI course syllabi. The analysis relied on a mixed-methods approach to be able to properly gather results and reach a conclusion. In particular, the results revealed key insights and patterns guided by the research questions. This discussion section interprets these results to answer the results in relation to the research questions and the overall objective of the study.

RQ1: What AI Topics Are Most Taught in University Introductory AI Courses?

In relation to the first research question about the exact topics that were taught in intro to AI courses, courses were coded based on relevant and determined content topics. The quantitative analysis results indicate that machine learning has an overwhelming majority in these courses. This indicates that machine learning has become the dominant focus of AI education today, reflecting its widespread use in real industries. This corroborates with the study from Yang et al. (2023) which found that machine learning was the most popular topic in AI courses. The binomial test confirms that machine learning is the only topic significantly above the 50% threshold ($p < .001$).

Following this, natural language processing and large language models were also highly represented in the analyzed courses. This suggests an increasing emphasis on AI systems that can understand and respond to regular human language, lining up with the recent developments and surge of popularity in generative AI. Chatbots are a sizable example of these tools, with growing use of AI voice-powered chatbots for fast food drive-throughs recently. Additionally, knowledge representation and reasoning (KRR) remained a consistent topic included in over half of the courses. This denotes that courses still give importance to the foundations and structure of AI systems before getting into their creation and utilization. It reflects a focus on the mechanics of intelligent systems and is considered a more traditional area of AI.

Ethics and neural networks were topics present in many of the course syllabi. The notable presence of neural networks represents deep learning in the modern AI landscape; considered a core component for advanced AI systems, neural networks comprise the backbone for tasks such as NLP and image recognition. Ethics being so widely included suggests that universities are increasingly recognizing the importance of ethics in AI. Based on the cross-tabulation and chi-square test, out of 30 courses in a tech-adjacent discipline, only 13 contained the topic of ethics, and this difference is statistically significant ($\chi^2 = 10.42$, $p = .001$), confirming that the underrepresentation of ethics in Tech courses is a reliable pattern, not an artifact of sample size.

Succeeding these topics were search, applications, and planning/scheduling. Search and planning/scheduling are traditional AI methods. The results show that while these methods are still being taught to a significant degree, they are slightly less emphasized than modern machine learning topics.

Lastly, Bayesian networks were the least frequently covered topic by a large margin (22.2%), and the binomial test confirms this frequency is significantly below chance ($p < .001$). The data from the cross-tabulation supports this with all instances of Bayesian networks being mentioned in the tech discipline category. The co-occurrence analysis confirms two structural clusters: classical AI topics (KRR, Search, Planning; $\phi = .51-.73$) and applied topics (Ethics, Applications, NLP; $\phi = .36-.53$), with the two clusters negatively correlated, reflecting two fundamentally distinct curricular philosophies. Additionally, the pairwise post-hoc comparisons provide a notable refinement: the Applications topic differs significantly for both Tech vs. Applied ($p < .001$) and Tech vs. General ($p = .022$), indicating that the absence of real-world application content is specific to Tech courses rather than a general feature of smaller course categories. By contrast, KRR differs significantly only for Tech vs. Applied ($p < .001$) but not for Tech vs. General ($p =$

.107), suggesting the classical algorithmic divide is primarily a CS-versus-non-CS phenomenon rather than a broader formal-versus-informal split.

RQ2: *How Do Pedagogical Approaches Differ Across Disciplines?*

With reference to the second research question about how pedagogical approaches differ across disciplines, the quantitative analysis reveals that lectures were the primary pedagogical method by a vast amount, only not being included in one course according to the syllabi. This indicates that lecture-based teaching is one of the best and most effective ways to teach introductory AI content, as lectures have the ability to provide students with structured explanations of content like algorithms and models. It is also essential for teaching a rapidly growing field to large groups of students. Chi-square tests confirm that lecture-heavy, project-based, and tool-driven pedagogies show no statistically significant differences across disciplines (all $p > .23$), confirming their discipline-neutral status.

Project-based learning typically involves assignments where students need to apply what they have learned, suggesting that educators see the importance of execution rather than just theory. Tool-driven instruction also had a high frequency. Python was the most frequently listed tool (48.9%), while ChatGPT and LLM APIs (26.7%) signal the growing presence of generative AI in course toolsets. Math-heavy instruction was heavily concentrated in technical courses over any other discipline ($\chi^2 = 24.09$, $p < .001$). In contrast, case-study-based instruction shows the strongest disciplinary divide in the dataset ($\chi^2 = 28.61$, $p < .001$), primarily found in applied and general courses.

Case studies allow students to explore the decisions and challenges that go into applying AI realistically. The qualitative Kruskal-Wallis tests confirm that the Application theme is significantly more prominent in Applied courses ($H = 7.87$, $p = .020$), while Technical, Hands-On, and Conceptual themes are near-universal. Pairwise post-hoc comparisons add further nuance to the pedagogy findings: Math-heavy instruction differs significantly for both Tech vs. Applied ($p < .001$) and Tech vs. General ($p = .010$), and Case-study-based instruction likewise differs for both Tech vs. Applied ($p < .001$) and Tech vs. General ($p = .017$). This means the pedagogical divide is not simply a matter of course size or formality; it holds consistently whether comparing Tech courses to the larger Applied group or the very small General group of three courses.

RQ3: *How Is Ethics Addressed in Introductory AI Courses?*

Regarding the third research question about ethics and how it's addressed in introductory AI courses, a little under half of all the courses contained a dedicated ethics module or section. This indicates a growing emphasis on ethics but does not currently take priority over other core AI topics, as these modules often take place over the course of one or two sessions. The chi-square test of the full three-way ethics cross-tabulation confirms that the disciplinary disparity in ethics coverage is statistically significant ($\chi^2(4) = 13.53$, $p = .009$, $V = .39$, medium-to-large effect).

Before discussing the courses that do not mention ethics at all, the brief mention ethics category should be discussed. This category was only coded when ethics was mentioned briefly in some capacity, usually as a small mention as an objective in the syllabi. The presence of this specific category shows that while ethics is important, some courses only touch on it very lightly in an unstructured way.

The Kruskal-Wallis test on the qualitative Ethics theme score further supports the disciplinary divide ($H = 7.27$, $p = .026$). Notably, 56.7% of Tech courses make no mention of ethics at all, while 75.0% of Applied courses include a dedicated ethics module. Considering that Tech students are most likely to build AI systems, this gap raises significant concerns about the preparation of future AI practitioners.

Conclusion

Through the conducted mixed-methods research and analysis, meaningful insights were attained regarding the main research objective of assessing how exactly AI is taught in universities. The three research questions did this by examining the topics, pedagogical approaches, and ethical considerations within the courses. Together, these areas provided sufficient data to understand the situation within introductory AI courses.

Based on the data collected, the overall conclusion for how AI is taught in universities can be separated into a few main findings. For the topic side, current AI courses have a large focus on teaching machine learning, since machine learning is considered the staple of modern AI. In terms of pedagogy, lecture-heavy instruction remains as the most used one across every discipline, complemented by project-based and tool-driven approaches. Between disciplines, technical ones include math-heavy as a pedagogy over case-study-based while it is the opposite for applied and general disciplines, with both differences confirmed as highly significant ($p < .001$). For ethics, they are addressed in a good number of courses but are not addressed in many as well, and the disciplinary gap is statistically confirmed ($p = .009$, $V = .39$).

These overall key findings point towards the conclusion that current AI education and syllabi have a strong emphasis on modern AI topics. This prepares students more for the future and current AI use in the workforce. The results from topics and pedagogies also represent a fine balance between theory and practice, making sure students can learn both. The main distinction between disciplines is that technical ones generally prioritize algorithms while applied and general ones emphasize real-world use. In conclusion, these findings and research reflect the current state of how AI is taught in universities according to many introductory AI course syllabi.

Limitations

Regarding limitations of this study, it is recognized that a bigger syllabi sample size over 45 could have influenced the results and provided a different outcome. Additionally, this study relied mostly on publicly available syllabi, which may not represent the contents of every AI course. The distribution of disciplines being heavily in favor of technical ones is another notable limitation. A significant limitation is that course syllabi had varying amounts of detail included. The coding for the thematic qualitative analysis also had a level of subjectivity involved. These directly affected the coding process.

Regarding the inferential tests, the small size of the General category ($n = 3$) required collapsing it with Applied in some pairwise comparisons. The absence of a Bonferroni correction means that results at the $p < .05$ threshold should be interpreted cautiously for pairwise and co-occurrence analyses, though the largest effects are robust given their $p < .001$ thresholds and large effect sizes.

Acknowledgements

During the preparation of this work the author(s) used Anthropic Claude in order to assist in statistics, graphics development, and content review. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

- Aljedaani, W., & Parthasarathy, P. D. (2026). How AI ethics is taught: Insights from a syllabus-level review of U.S. Computing courses. *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.1*, 31–37. <https://doi.org/10.1145/3770762.3772583>
- Degree Prospects. (2026). 2025 AI degree report: The landscape of AI degrees in the United States. *MastersInAI.org*. <https://www.mastersinai.org/ai-degree-report/>
- Garrett, N., Beard, N., & Fiesler, C. (2020). More than “if time allows”: The role of ethics in AI education. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 272–278. <https://doi.org/10.1145/3375627.3375868>
- Gu, X., & Ericson, B. J. (2025, August). AI literacy in K-12 and higher education in the wake of generative AI: An integrative review. In *Proceedings of the 2025 ACM Conference on International Computing Education Research V. 1* (pp. 125–140).
- Long, D., & Magerko, B. (2020, April). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16).
- Ng, D. T. K., Lee, M., Tan, R. J. Y., Hu, X., Downie, J. S., & Chu, S. K. W. (2023). A review of AI teaching and learning from 2000 to 2020. *Education and Information Technologies*, 28(7), 8445–8501. <https://doi.org/10.1007/s10639-022-11491-w>
- Ng, D. T. K., Su, J., Leung, J. K. L., & Chu, S. K. W. (2024). Artificial intelligence (AI) literacy education in secondary schools: A review. *Interactive Learning Environments*, 32(10), 6204–6224. <https://doi.org/10.1080/10494820.2023.2255228>
- Rufai, A. A., Hossain Khan, M. S., & Hasan, M. (2024). An exploration of pedagogical approaches in teaching artificial intelligence courses. *Social Sciences & Humanities Open*, 10, 101075. <https://doi.org/10.1016/j.ssaho.2024.101075>
- Stanford Human-Centered Artificial Intelligence. (2025). *2025 AI index report*. Stanford University. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- Wu, D., Chen, M., Chen, X., & Liu, X. (2024). Analyzing K-12 AI education: A large language model study of classroom instruction on learning theories, pedagogy, tools, and AI literacy. *Computers and Education: Artificial Intelligence*, 7, 100295. <https://doi.org/10.1016/j.caeai.2024.100295>
- Yang, Y., Qin, J., Lei, J., & Liu, Y. (2023). Research status and challenges on the sustainable development of artificial intelligence courses from a global perspective. *Sustainability*, 15(12), 9335. <https://doi.org/10.3390/su15129335>
- Zhou, X., Tong, Y., Lan, X., Zheng, K., & Zhan, Z. (2021). AI education in massive open online courses: A content analysis. *2021 3rd International Conference on Computer Science and Technologies in Education (CSTE)*, 80–85. <https://doi.org/10.1109/CSTE53634.2021.00023>