2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #52078

Integration of Social and Ethical AI Topics into a Computational Biomedical Engineering Course Increases Engagement Without Loss of Technical Learning

Dr. Alexander Phillip Clark, University of Virginia
Dr. Sarah Maddox Groves, University of Virginia
Dr. Jeffrey J Saucerman, University of Virginia
Kaitlyn Wintruba, University of Virginia

Integration of Social and Ethical AI Topics into a Computational Biomedical Engineering Course Increases Engagement Without Loss of Technical Learning

Abstract:

Introduction: There is increasing emphasis on the need for tightly integrated technical, social, and ethical training, particularly in areas with profound societal impact such as AI. However, whether social and ethical content integrated into a technical class infringes on technical learning remains unclear. In an undergraduate course on Computational Biomedical Engineering, this study examined whether partially replacing technical AI content with social-ethical AI content affects student engagement and technical learning. The hypothesis was that due to redistribution in class time, incorporating social and ethical topics would increase engagement but reduce technical learning outcomes.

Methods: A computational biomedical engineering course was taught in 2023 and 2024 with comparable student cohorts and the same instructor. In 2024, approximately 25% of class time was allocated to social and ethical topics, replacing equivalent technical content. Quantitative and qualitative data were collected through pre- and post-course surveys assessing student interest in engineering, biology, and medicine, as well as performance on problem-solving questions identical across years. Additional data sources included short essays, written reflections on ethical AI, and course evaluations.

Results: The hypothesis that increased proportion of time spent on social and ethical topics would reduce technical learning was not supported. Students in the 2024 cohort demonstrated comparable pre-to-post course gains in performance to those in 2023, despite reduced time on technical content. In 2024, average test scores increased from 40% pre-course to 95% post-course, exceeding the 2023 post-course average of 87%, suggesting that inclusion of social-ethical material was not associated with a measurable decrease in assessed technical learning. Student interest levels also remained high or improved: the proportion of students who strongly agreed with the statements *"I am interested in engineering," "I am interested in biology,"* or *"I am interested in medicine"* rose from 60% to 72%, 53% to 72%, and 53% to 72% in 2023 vs. 2024, respectively. Surveys and written reflections further indicated increased engagement and greater awareness of AI's role in society.

Conclusion: Incorporating social and ethical perspectives into computational coursework can strengthen engagement without diminishing technical learning. These findings suggest that balanced pedagogical approaches—combining technical skills with reflection on ethical and societal implications—prepare future engineers to navigate both the computational and human dimensions of biomedical innovation.

Introduction

The social, ethical, and technical impacts of AI across domains cannot be understated. The Artificial Intelligence Literacy Act clearly stated, "Americans of every age, in every district, and from every background will be impacted by AI, and therefore need AI literacy–an understanding of basic AI principles and applications, the skills to recognize when AI is employed, and awareness of its limits" [1]. In particular, users of generative AI must appreciate the social and ethical implications of its use. While academic integrity and AI in the classroom are an important domain, building AI literacy focuses on developing a student's understanding of how AI works and its broader impacts on society and the environment. Furthermore, students must learn how to use this new tool ethically, keeping in mind that AI readily hallucinates [2], shows bias [3], [4], and can have ethical and legal concerns [5].

Currently, much of the educational research on AI ethics in the classroom focuses on integration into ethics-based electives rather than technical courses. For example, a recently published protocol [6] was developed to investigate best practices for teaching postsecondary students about AI from a social and ethical perspective and to detail evaluation methods for educators. The researchers emphasize that there is little standardization of teaching AI literacy across disciplines. Furthermore, the protocol is focused on courses that focus on AI literacy and ethics, rather than those focused on technical content with integration of ethical discussions.

Even with the growing availability of such AI ethics courses, a majority of students feel they are not prepared by their academic courses to use AI [7]. An ability to use AI ethically and responsibly is important in every discipline and career, in or outside of academia. However, a large gap exists between the percentage of graduates who felt their postsecondary program prepared them to use AI tools (45%) and the percentage of employers who use and expect employees to use AI (73%). To develop critical thinking skills and practice ethical reasoning in an age of AI, students must be equipped with curricula that focuses on AI literacy and its socio-ethical impact. Such curricula should train AI specialists while also preparing the future workforce for using AI ethically. Long and Magerko [8] developed a conceptual framework for AI literacy, including competencies and design considerations (Reprinted in **Appendix A**). This framework enumerates five themes that should be addressed in any curriculum focused on ethical use of AI.

Since these competencies were published, curriculum focused on developing AI literacy has increased through courses on AI ethics alone. For example, of the 53 science or engineering courses being considered as part of a proposed AI minor at our institution, none integrate social and technical content by design. Rather social-ethical content related to AI are taught on their own in 11 courses. Other researchers have developed core competencies that should be addressed by integrating AI content into BME courses, specifically [9]. At the Fifth Biomedical Engineering Education Summit, strategies were developed for addressing the challenges of integrating AI into BME curricula, focusing on key skills such as mathematical foundations, programming proficiency, ethical awareness, and adaptability. Many of these skills therefore focus on understanding machine learning itself, rather than the social implications. Therefore, a significant curricular gap is the absence of AI literacy and socio-ethical content integrating into technical courses**.** Closing this gap would increase the reach of AI ethical content to a broader array of students, as some may feel that they do not have the availability for an AI ethics elective. Rather, integration into a core class would require socio-ethical content to replace some technical content. While Khojah [9] provides a framework for integration, it is currently unclear if such a replacement would infringe on technical learning.

In this manuscript, we use the term "AI" broadly to encompass machine learning, deep learning, and generative AI, consistent with the Long and Magerko [8] framework and the AI Literacy Act [1]. The technical course content focuses on numerical methods (e.g., optimization, regression) foundational to these learning algorithms in BME applications , while the social-ethical discussions address both traditional machine learning (ML) applications in biomedicine and generative AI systems that produce open-ended responses, with much wider implications that have great societal significance even beyond the realm of BME [8]. We recognize these categories have distinct implications for biomedical engineering practice and have noted this distinction where relevant.

This study examined student engagement and technical learning in a core undergraduate course (Computational Biomedical Engineering, BME2315) with integrating AI ethical discussions. Each week, students engaged in discussions covering several literacy competencies and wrote opinion papers on the topics. We hypothesized that incorporation of such activities would increase engagement but reduce technical learning outcomes. Instead, we found no decrease in the technical learning that we assessed alongside an increase in interest in the course material and the biomedical engineering discipline. These results suggest that incorporating social and ethical AI content in a core, technical course is an important way to increase AI literacy across engineering students without decreasing technical learning.

Methods

Integration of social-ethical content

Social-ethical AI content to increase AI literacy was incorporated into a core Biomedical Engineering course, BME2315. This course introduces second-year students to foundational numerical methods, including mathematical background and computational applications of the methods, such as regression, interpolation, optimization, solving systems of linear equations, and solving systems of ordinary differential equations (ODEs).

In Fall 2024, 25% of class time was allocated to discussion of social/ethical implications of AI rather than technical content. Each unit was reduced by an average of 25% to accommodate such discussions, and several topics (including one full unit) were entirely removed (Table 1). Table 1 specifies which topics were removed (bolded) and the number of lectures dedicated to each topic in 2023 and 2024. In Fall 2023, the course did not include social-ethical AI discussions and is included in this study as a control. Social-ethical AI content replaced technical content on specific numerical methods that are more advanced or do not scale well to larger or more complex data (**Table 1**). A pre- and post-course quiz (see below) assessed knowledge of topics covered in both years; omitted topics (root finding, golden section search, stiff ODEs, adaptive time stepping) were not tested.

Table 1: Reductions to technical content from Fall 2023 to Fall 2024 by topic. Bolded topics were not covered in 2024.

Unit	**Fall 2023 Technical Content (Topics and Time Spent)**	**Fall 2024 Technical Content (Topics and Time Spent)**
Python Coding	IDEs, Basic Python Code, Unit Tests, Jupyter Notebooks (2 classes)	Equivalent content (1.5 classes)
Large Language Models	AI, Large Language Models, Neural Networks (1 class)	Equivalent content (0.5 class)
Error	Roundoff error, Truncation error, Error propagation (2 classes)	Equivalent content (1.5 classes)
Numerical Differentiation & Integration	Divided-difference approximations, Newton-Cotes, Trapezoidal Rule, Simpson's Rules (2 classes)	Equivalent content (2.5 classes)
Linear Systems	Linear algebra basics, Gaussian elimination (3 classes)	Equivalent content (2.5 classes)
Root Finding	**Bisection method, Newton-Raphson method** (2 classes)	Not covered (0 classes)
Interpolation	Linear, quadratic, polynomial, B-spline (1 class)	Equivalent content (0.5 class)
Optimization	**Golden section search**, gradient ascent, Newton's method, Multivariable optimization (2 classes)	Gradient ascent, Newton's method, Multivariable optimization (2 classes)
Regression	Linear least-squares, gradient descent & Newton's method for nonlinear regression (3 classes)	Equivalent content (2.5 classes)
ODEs	Euler's Method, Midpoint Method, 4th order RK, Oscillators,	Euler's Method, Midpoint Method, 4th order RK, Oscillators (2 classes)

	Stiff ODEs, Adaptive time stepping (3 classes)	
MATLAB	Intro to MATLAB (2 classes)	Equivalent content (2 classes)
DISTRIBUTION OF CLASS TOPICS	Technical content = 22 classes Intro/Review days = 4 classes Exams = 3 classes TOTAL = 29 classes	Technical content =17.5 classes AI content = 4.5 classes Intro/Review days = 3 classes Exams = 3 classes TOTAL = 28 classes

In Fall 2024, seven AI social-ethical topics (**Table 2**) that span the core competencies for AI literacy (**Appendix A**) were integrated into the course with a reading or viewing assignment, response paper, and discussion. Students completed the reading assignment and initial ½-page response paper outside of class. In the following class period, the students engaged in a discussion about the AI topic with their peers and instructor. The structured discussions began with small group discussions for about 10 minutes to identify areas of consensus and disagreement. Students then shared with the rest of the class their common and distinct group views and individual views, and their peers could reflect and respond to their points. The instructor emphasized the need for these discussions to remain respectful as sensitive topics could be discussed.

Data collection approach

During each semester, an instructor-developed assessment in the form of pre- and post-test quizzes were required for all students and graded by participation. Both tests included 1) a subjective assessment including questions to gauge interest in major themes of the course (engineering, biology, and medicine), perceived usefulness of the course, confidence in coding skills and technical knowledge; and 2) a technical proficiency assessment (**Appendix D**). Official course evaluations from both years provide an avenue for students to provide comments in response to 3 questions: 1) Please tell us briefly how any of the above learning activities (or other activities not included above) contributed to your learning in this course; 2) What would you like the instructor and university administrators to know about your experience in this course?; and 3) Please give specific examples as to how [the instructor] created an environment that respected difference and welcomed diverse perspectives.

Student discussions were not quantitatively scored; however, participation was encouraged during in-class activities and reflected in qualitative course evaluations. Written reflections were completed individually prior to class and then expanded after peer discussion; reflections were not peer-reviewed or graded by classmates.

The concept knowledge assessment was administered as a pre-test during the first week of the course prior to AI-related instruction and as a post-test on the day of the final exam.

Statistical tests and power analysis

All study procedures were reviewed and approved by the Institutional Review Board for Social and Behavioral Sciences (IRB-SBS) at our institution. Likert-scale survey responses were analyzed as ordinal data. For visualization, response distributions were summarized as proportions across the five Likert categories (Strongly Disagree to Strongly Agree) for each survey item and time point. For statistical testing, Likert responses were converted to numeric scores (Strongly Disagree = 1 to Strongly Agree = 5), and nonresponses were excluded. For statistical analysis, responses were analyzed using rank-based nonparametric tests without assuming equal spacing between categories; numeric coding (1–5) was used only to determine rank ordering. Nonresponses were excluded. Pre- and post-course responses from the same students were compared using a two-sided Wilcoxon signed-rank test paired by student identifier (testing whether the median within-student change differed from zero). Comparisons between independent student cohorts (2023 vs 2024) were performed using a two-sided Mann–Whitney U (Wilcoxon rank-sum) test. These analyses evaluated whether responses shifted toward higher or lower agreement rather than testing individual category counts.

Post hoc power analyses were performed for Mann–Whitney U tests using effect sizes estimated from observed rank-biserial correlations. For each comparison, the rank-biserial correlation was computed from the Mann–Whitney U statistic and used to approximate statistical power under a normal approximation to the rank-sum distribution. Required sample sizes to achieve target power levels were estimated by iteratively increasing group sizes under an equal-group assumption until the calculated power exceeded 80%. These calculations account for the ordinal nature of the data and were conducted to quantify the sample sizes necessary to detect distributional shifts of the observed magnitude under the same testing framework.

Concept survey scores were analyzed separately from Likert survey responses because they represent continuous percentage measurements rather than ordinal categories. For within-student pre- and post-course comparisons, differences in quiz scores were evaluated using a two-sided paired t-test. Comparisons of post-course performance between independent cohorts (2023 vs 2024) were evaluated using a two-sample t-test with unequal variances (Welch's t-test). These parametric tests were used because score differences were approximately continuous and symmetric, and inference concerned mean performance rather than rank ordering. Scores are reported as mean ± standard deviation.

Table 2: Social-ethical AI Topics integrated into BME2315 in Fall 2024

Topic #	Discussion & Opinion Paper Topic	Core AI Literacy	Effective and enduring skills

		Competencies (Long and Magerko, [8])	**for AI education** (Khojah [9])
1	**AI in Sci-Fi**: What is the role of AI in a sci-fi movie of your choice? What are the most promising and most concerning implications of AI? How close are we to this form of AI?	**1, 3, 6, 14, 15**	Adaptability, interdisciplinary collaboration
2	**What is AI, anyway?** Review of TED talks by Mustafa Sulerman (CEO of Microsoft AI) and Sebastian Thrun (co-founder of Udacity, Google X). In your opinion, what is AI? In what ways do you agree with the views of Mustafa Sulerman or Sebastian Thrun? In what ways are their explanations of AI similar/dissimilar to the AI seen in sci-fi?	**2, 3, 7, 8, 9, 12, 14**	Mathematical foundations, machine learning techniques
3	**AI Hallucinations**: Identify and explain one example of a prominent hallucination/failure of a computational or AI prediction. How good is good enough for predictions made by a computational model? Is the threshold the same for an AI model? Why or why not?	**2, 5, 8, 10, 12, 13, 16**	Ethical awareness and decision making
4	**Environmental Impacts:** Provide a brief overview of the current and projected environmental costs of AI, including quantitative estimates of scale and time. Make the case for one particular mitigation strategy and its potential to reduce these environmental costs.	**5, 16**	Ethical awareness and decision making
5	**Algorithmic bias in AI:** Identify a specific, prominent real-world example of algorithmic bias or discimination involving AI. Explain the source of the algorithmic bias, its implications for use of this AI, and suggest practical strategies for mitigating this bias.	**2, 5, 8, 10, 13, 16**	Ethical awareness and decision making, Model evaluation & bias awareness
6	**Biomedical Applications of AI:** Identify a specific major biomedical AI advance in 2024, providing a link or citation. Write a paragraph summarizing its goal, the numerical methods we've discussed that are most relevant to how it works and evaluate its long-term promise.	**1, 3, 10, 13, 14, 15, 16**	Interdisciplinary collaboration, open science practices, mathematical foundations

7	**Generative AI for Biomedical Visualization:** Use a generative AI of your choice to visualize a biomedical concept or computational approach of your choice. Describe both the prompt you used and your interpretation of the result. How might generative AI change how you make such visualizations in the future? How does it make you think differently about art?	**1, 6, 17**	Interdisciplinary collaboration

Results

Students report the value of integrating social-ethical AI content

To identify student opinions on the value of integrating social-ethical AI content, we examined responses to three questions from the course evaluation questions:

1) "tell us briefly how any of the learning activities… contributed to your learning in this course",

2) "what would you like the instructor… to know about your experience in the course", and

3) "what constructive suggestions do you have to help [the instructor] improve this course for future students".

Based on these post-course evaluations, overall sentiment to the integration of social-ethical AI content was positive. Out of 70 comments in response to these three course evaluation prompts, 10 mentioned the AI discussions and only two were slightly negative (**Appendix B**, Discussion). Several students reported that AI discussions increased their enjoyment and interest in the class, such as one student who said, "I really liked this class, especially the opinion papers and discussion to supplement the more technical learning." The AI discussions were seen as "helpful" for learning, "very interesting," and "engaging."

The opinion papers displayed a wide range of themes for which the students gained appreciation. For example, submissions for the first opinion paper, which focused on AI in sci-fi (**Table 2**), spanned themes of human emotion, dystopian fears, practical utility, and ethical concerns. Through this first assignment, students began to appreciate that AI is efficient and can provide value to healthcare and the medical field, but at the expense of high amounts of [environmental] waste, bias, and limited moral reasoning. In particular, students noted that human oversight is necessary for ethical use of AI. In the second assignment ("What is AI, anyway?"), several students described the tension between unique human traits that remain irreplaceable and the necessity for the workforce to evolve. Students emphasized the need for individuals to learn to use AI effectively and be open to change, with most students believing that AI is a powerful tool to augment human capabilities (Core competency #10: Human role in AI). In this early assignment, students already began to uncover ethical implications of AI, including data privacy, potential bias in algorithms (aligning with Topic 5), and the need for higher-level (governmental) regulation. Regarding AI in the classroom, student sentiment was that the educational system needs to adapt to prepare future generations, teaching students to think critically and work with AI.

In all, the opinion papers and reflections were evidence that students gained better understanding of ethical AI use and the impact of AI on humanity and the environment. Students began to recognize when AI is successful and when it fails. Overall, evidence of AI literacy improvement was demonstrated in the opinion papers, and subjective responses on course evaluations show perceived learning and engagement with these topics.

Pre- to post-course surveys in 2024 revealed an impact on student interest in the technical but not social interest in AI **(Figure 1).** Interest in the technical aspects of AI increased significantly from pre- to post-survey ($p = 0.031$). Prior to the course, most students reported being undecided or moderately interested (Undecided = 5, Agree = 11, Strongly Agree = 3), whereas after the course responses on technical interest in AI shifted toward higher enthusiasm, with a marked increase in Strongly Agree responses (Undecided = 2, Agree = 6, Strongly Agree = 9). In contrast, interest in the social and ethical aspects of AI did not change significantly over the course. However, there was a modest redistribution in responses, with students moving from undecided toward agreement (Undecided decreased from 4 to 0; Agree increased from 7 to 9), while the proportion of Strongly Agree responses remained relatively stable (8 pre vs. 7 post). Together, these results suggest that the course increased student interest in AI's technical dimensions, reinforcing pre-existing interest in its social and ethical considerations.

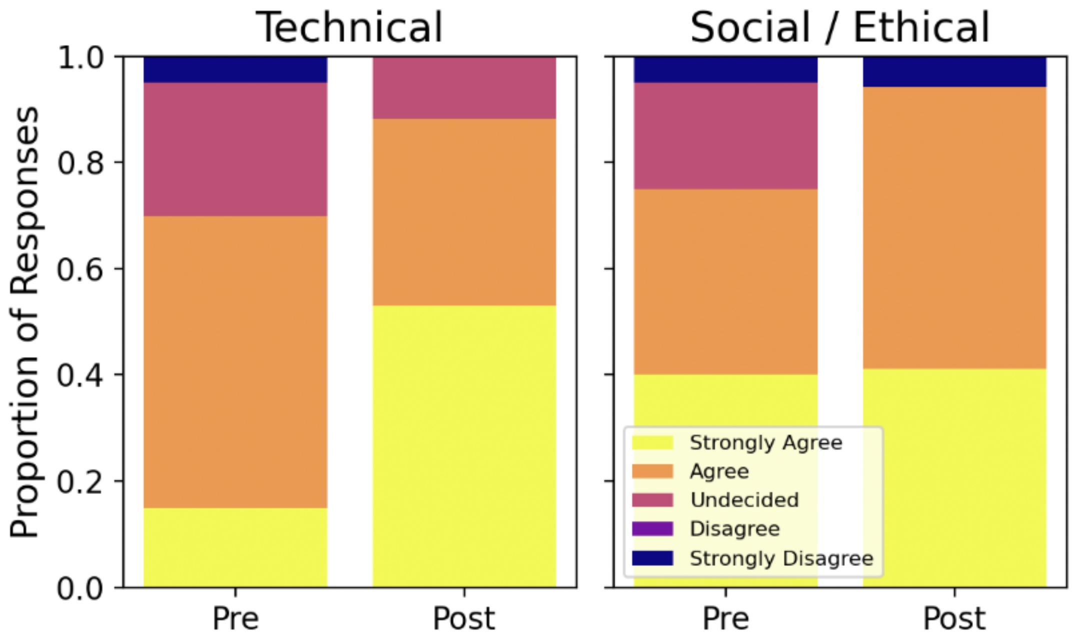

Figure 1: Change in student interest in technical and social / ethical aspects of AI from pre- (n=20) to post-course (n=17) surveys in 2024. For students who completed pre- and post-course surveys, there was a significant increase in interest in technical aspects of AI ($p = 0.031$) but not social / ethical ($p=0.80$), via Wilcoxon paired signed-rank.

Student engagement in technical material remains high with social-ethical AI content

Student engagement with core disciplinary content remained high over the course of instruction, including in the presence of social–ethical AI material. In 2024, pre- to post-course surveys assessing interest in engineering, biology, and medicine showed no statistically significant changes (engineering $p = 0.53$, biology $p = 0.08$, medicine $p = 1$, **Figure 2**). The proportion of students reporting strong interest increased for engineering (from 60% pre-course to 72% post-course) and biology (from 65% to 72%), while strong interest in medicine decreased modestly over the same period (from 80% to 72%). These results indicate that, although changes did not reach statistical significance, student interest in engineering and biology was maintained or modestly strengthened, while interest in medicine remained high overall despite a small downward shift.

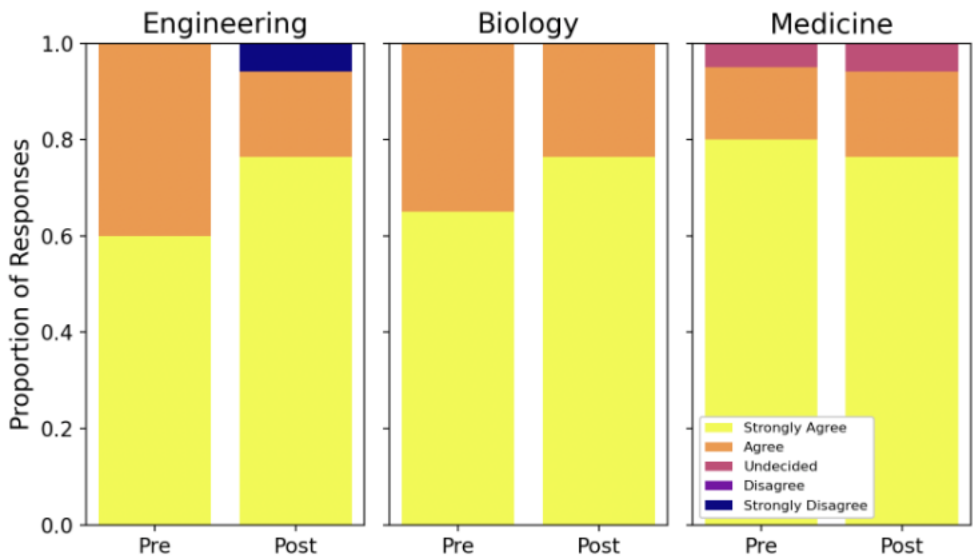

Figure 2: Proportion of students expressing interest in engineering, biology, or medicine in 2024 pre- (n=20) vs post-course (n=17) surveys. The plot shows a pre-to-post increase in Strongly Agree proportion for engineering and biology but not medicine, with no statistically significant difference (engineering $p = 0.53$, biology $p = 0.08$, medicine $p = 1$) via Wilcoxon paired signed-rank.

Comparison of post-course surveys between 2023 and 2024 revealed a consistent pattern of higher reported strong interest in 2024 across all three domains, though again there is no significant difference between groups (engineering $p = 0.37$, biology $p = 0.09$, medicine $p = 0.14$, **Figure 3**). The proportion of students selecting "Strongly Agree" increased from 2023 to 2024 for all three categories: engineering (61% to 72%), biology (52% to 72%), and medicine (52% to 72%). Together, these findings suggest that incorporating social and ethical AI content did not negatively impact student engagement with engineering, biology, or medicine and may be associated with broad, though statistically non-significant, increases in strong disciplinary interest across course iterations.

To contextualize these non-significant results, we performed post hoc power analyses based on the observed Mann–Whitney U effect sizes. For the biology interest comparison, detecting an effect of the observed magnitude with conventional power would require approximately 210 students per group. Required sample sizes for the other domains were even larger, on the order of several hundred students per group. These high estimates are likely driven by pronounced ceiling effects and a large number of tied responses on the 5-point Likert scale, which substantially reduce statistical power. Consequently, the absence of statistically significant differences should be interpreted with caution and does not preclude small but meaningful shifts in student interest.

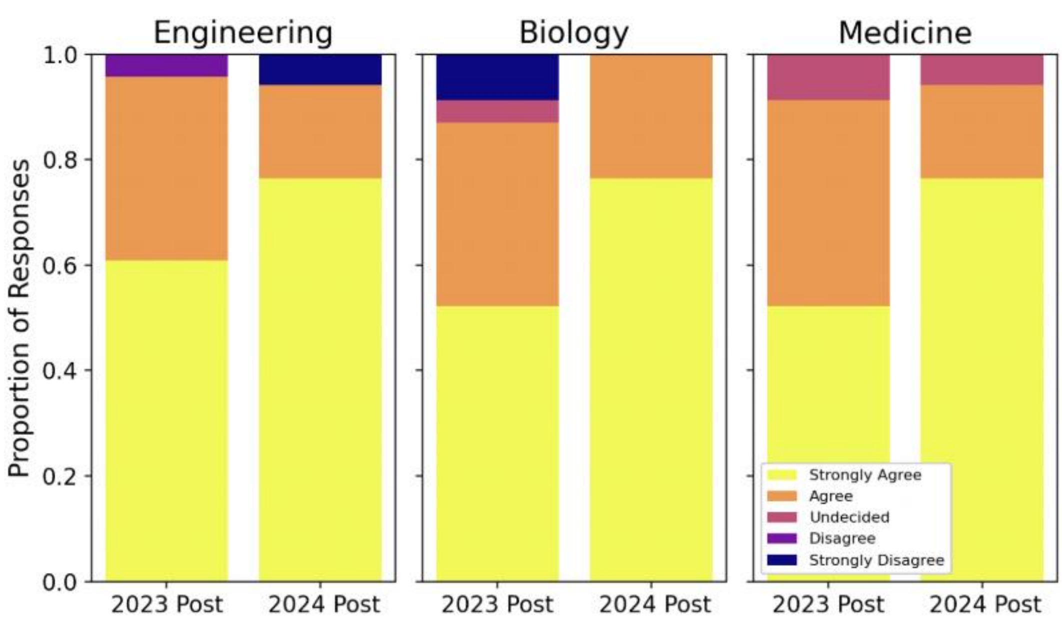

Figure 3: Proportion of students expressing interest in engineering, biology, or medicine in 2023 (n=23) vs 2024 (n=17) post-course surveys. This comparison shows a qualitative increase in Strongly Agree proportion in 2024 vs 2023 for all three areas. There was no significant difference (p_{eng}=0.37, p_{bio}=0.09, p_{med}=.14) via two-sided Mann-Whitney test.

Decrease in technical content by 25% shows no evidence of decrease in knowledge of assessed technical concepts

Despite a 25% reduction in technical content in the 2024 course relative to 2023, technical learning on the topics assessed in a post-course quiz was comparable across cohorts **(Appendix C).** In contrast, student performance on this concept survey improved substantially over the duration of the 2024 course **(Figure 4).** Average quiz scores increased from 40% ± 24% on pre-course tests (n = 20) to 95% ± 8% on post-course tests (n = 17), with a highly significant improvement observed in paired pre- to post-course comparisons (paired t-test $p = 3.75 \times 10^{-7}$). Moreover, post-course performance in 2024 was comparable to, and slightly exceeded, performance post-course scores in 2023 (Quiz 3: 87% ± 13%, n = 23), despite the reduced technical emphasis. Together, these results indicate that decreasing the volume of technical material did not compromise—and may have supported—student learning of core technical concepts.

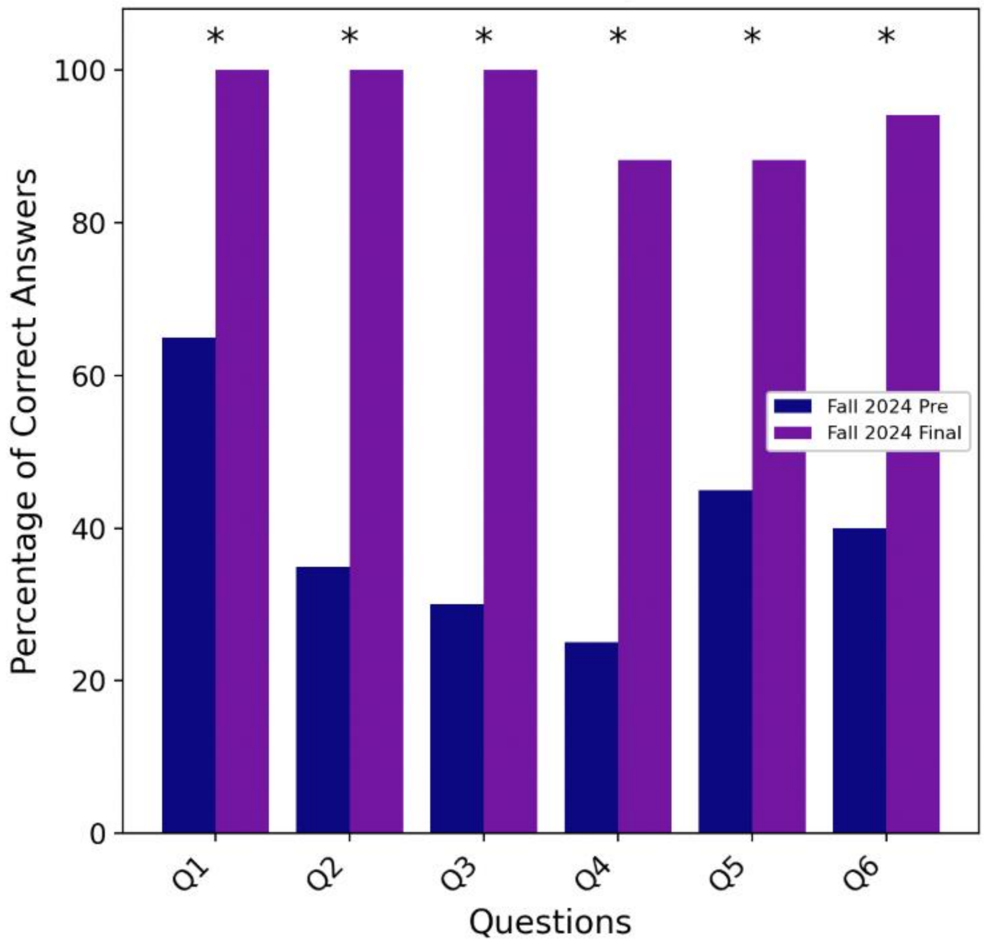

Figure 4: Student changes in knowledge of technical concepts, compared via pre- vs. post-course surveys in 2024. There was a statistically significant increase from pre-test (n = 20) to post-test (n = 17) on all six questions, via paired t-test ($p = 3.75 \times 10^{-7}$). Technical concept survey questions are provided in **Appendix D**.

Discussion

This study provides preliminary evidence for the educational value of integrating social-ethical AI content into a technical course. Even when 25% of the technical content was replaced with discussions and assignments focused on the social impact of AI, technical learning of assessed content in a core BME course did not decrease. Furthermore, students gained AI literacy skills and self-reported that the incorporated content was engaging and interesting. Interest in core biomedical engineering domains also remained high, with more students strongly agreeing with the statements "I am interested in engineering" and "I am interested in biology." Together, these results suggest that incorporation of social-ethical AI content into a technical BME course showed positive effects on engagement, learning, and interest and no negative effect on technical proficiency.

Because this was a retrospective study, it had some notable limitations. Firstly, sentiment towards AI content integration was gauged using general course evaluation comments, rather than survey questions specifically geared towards uncovering sentiment. Because students were not specifically asked about the integration of new content, only 10 responses mentioned AI or opinion papers. Similarly, the pre- and post-test questions were kept consistent between years, which meant that questions specific to the removed content were in some cases not tested. Because the quiz tested only topics retained in both course iterations, we cannot draw conclusions about student learning of omitted material. The quiz used in this study was developed by the course instructor and has not undergone formal psychometric validation. While pre-to-post gains and cross-cohort comparisons are internally consistent, future work should include expert review, item analysis, and reliability testing to strengthen the evidentiary basis. While a limitation of our study, the foundational knowledge developed and tested in the quiz, as well as evaluations throughout the semester, suggest that students, with the tool set we gave them, would be capable of self-teaching the removed content. In future iterations of this study, we would like to specifically include questions regarding the sentiment towards AI content integration, and we would test the assumption that the removed content would be buildable with self-learning. Furthermore, formal validation such as expert review of the assessment, test-retest reliability, and psychometric validation of a concept inventory would strengthen future iterations of this study.

Secondly, our study was underpowered to adequately detect differences in self-reported interest in core bioengineering domains. Power analysis showed hundreds to thousands of students would be needed to detect such differences. Additionally, analyses relied on rank-based nonparametric tests rather than hierarchical ordinal regression models. Larger multi-year studies could apply mixed-effects ordinal models to jointly estimate cohort and within-student changes.

Thirdly, two students provided critical feedback related to the replacement of technical content with social-ethical AI content, stating that the course "felt a little too broad," or "too fast," and "it felt a bit disconnected at times." These statements point to possible improvements to the integration strategy by aligning the AI topics directly with the technical content topic so that the course feels more cohesive. For example, the optimization technical unit would align well with a discussion of how AI and LLMs work (and what is being optimized). The section on error would be an appropriate time to talk about AI bias and hallucinations.

While social-ethical content regarding AI is not usually added to technical courses, some studies have shown the impact of adding such discussions regarding machine learning (ML) in general. For example, Saltz [10] Discussed the need for ethical consideration within ML courses. In response to this need, the researchers developed a framework for such incorporation. Specific to BME courses, Wang [11] developed an approach to focus on biases in ML, particularly AI-enabled medical devices. Their framework focused on similar topics as our discussion topics 4 and 5, while our framework aimed to more broadly cover the implications of AI.

This initial study on socio-technical integration with the subject of Computational Biomedical Engineering tested the integration of seven topics on the social and ethical implications of AI, which was limited by the semester-long course. In future courses, based on the core competencies of AI literacy, other possible topics include: 1) content ownership and intellectual property: who owns the results of AI when it is trained on content from a huge variety of sources? (Competencies 5 and 16); 2) How does AI actually learn from data, and what if that data is biased? (Competencies 12, 13, and 16); and 3) How can and should AI be used in medical decisions? (Competencies 14, 15, and 16). Future work should also be done to incorporate more students into a follow-up study in order to test the underpowered associations we show here. Overall, this study provides evidence that integration of social-ethical AI content into a technical course is an important avenue for developing AI literacy in biomedical engineering students without reducing technical proficiency.

References

[1] *Artificial Intelligence Literacy Act of 2023*. 2023. Accessed: Apr. 01, 2026. [Online]. Available: https://www.congress.gov/bill/118th-congress/house-bill/6791

[2] S. Farquhar, J. Kossen, L. Kuhn, and Y. Gal, "Detecting hallucinations in large language models using semantic entropy," *Nature*, vol. 630, no. 8017, pp. 625–630, Jun. 2024, doi: 10.1038/s41586-024-07421-0.

[3] B. Hofmann, "Biases in AI: acknowledging and addressing the inevitable ethical issues," *Front. Digit. Health*, vol. 7, p. 1614105, Aug. 2025, doi: 10.3389/fdgth.2025.1614105.

[4] P. Radanliev, "AI Ethics: Integrating Transparency, Fairness, and Privacy in AI Development," *Applied Artificial Intelligence*, vol. 39, no. 1, p. 2463722, Dec. 2025, doi: 10.1080/08839514.2025.2463722.

[5] Y. Qian, K. L. Siau, and F. F. Nah, "Societal impacts of artificial intelligence: Ethical, legal, and governance issues," *Societal Impacts*, vol. 3, p. 100040, Jun. 2024, doi: 10.1016/j.socimp.2024.100040.

[6] C. Hillis *et al.*, "Teaching postsecondary students about the ethics of artificial intelligence: A scoping review protocol," *PLoS One*, vol. 20, no. 7, p. e0329020, Jul. 2025, doi: 10.1371/journal.pone.0329020.

[7] Cengage Group, "2024 Graduate Employability Report: Preparing Students for the GenAI-Driven Workplace," Cengage Group, Boston, MA, Jul. 2024. Accessed: Apr. 01, 2026. [Online]. Available: https://cengage.widen.net/s/bmjxxjx9mm/cg-2024-employability-survey-rep

[8] D. Long and B. Magerko, "What is AI Literacy? Competencies and Design Considerations," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, Apr. 2020, pp. 1–16. doi: 10.1145/3313831.3376727.

[9] R. Khojah, A. Werth, K. W. Broadhead, L. W. Dobrucki, C. Geiger, and D. A. Rubenstein, "Integrating Generative Artificial Intelligence Tools and Competencies in Biomedical Engineering Education," *Biomed Eng Education*, vol. 5, no. 2, pp. 135–151, Jul. 2025, doi: 10.1007/s43683-025-00175-9.

[10] J. Saltz *et al.*, "Integrating Ethics within Machine Learning Courses," *ACM Trans. Comput. Educ.*, vol. 19, no. 4, pp. 1–26, Dec. 2019, doi: 10.1145/3341164.

[11] X. Wang, T. Chan, and A. Tamura, "BOARD # 28: Work in progress: Preparing Biomedical Engineers to Tackle Biases in Machine Learning," in *2025 ASEE Annual Conference & Exposition Proceedings*, Montreal, Quebec, Canada: ASEE Conferences, Jun. 2025, p. 55643. doi: 10.18260/1-2--55643.

Appendix A: Reprint of Core Competencies for AI Literacy from Long and Magerko [8]. These competencies are referenced in Table 2.

Theme	**Competencies**
What is AI?	1 (Recognizing AI) Distinguish between technological artifacts that use and do not use AI
	2 (Understanding Intelligence) Critically analyze and discuss features that make an entity "intelligent", including discussing differences between human, animal, and machine intelligence.
	3 (Interdisciplinarity) Recognize that there are many ways to think about and develop "intelligent" machines. Identify a variety of technologies that use AI, including technology spanning cognitive systems, robotics, and ML.
	4 (General vs Narrow) Distinguish between general and narrow AI.
What can AI do?	5 (AI's Strengths and Weaknesses) Identify problem types that AI excels at and problems that are more challenging for AI. Use this information to determine when it is appropriate to use AI and when to leverage human skills.
	6 (Imagine Future AI) Imagine possible future applications of AI and consider the effects of such applications on the world.
How does AI work?	7 (Representations) Understand what a knowledge representation is and describe some examples of knowledge representations.
	8 (Decision-Making) Recognize and describe examples of how computers reason and make decisions.
	9 (ML Steps) Understand the steps involved in machine learning and the practices and challenges that each step entails.
	10 (Human Role in AI) Recognize that humans play an important role in programming, choosing models, and fine-tuning AI systems.
	11 (Data Literacy) Understand basic data literacy concepts
	12 (Learning from Data) Recognize that computers often learn from data (including one's own data).
	13 (Critically Interpreting Data) Understand that data cannot be taken at face-value and requires interpretation. Describe how the training examples provided in an initial dataset can affect the results of an algorithm.
	14 (Action & Reaction) Understand that some AI systems have the ability to physically act on the world. This action can be directed by higher-level reasoning (e.g. walking along a planned path) or it can be reactive (e.g. jumping backwards to avoid a sensed obstacle).
	15 (Sensors) Understand what sensors are, recognize that computers perceive the world using sensors, and identify sensors on a variety of

	devices. Recognize that different sensors support different types of representation and reasoning about the world.
How should AI be used?	16 (Ethics) Identify and describe different perspectives on the key ethical issues surrounding AI (i.e. privacy, employment, misinformation, the singularity, ethical decision making, diversity, bias, transparency, accountability).
How do people perceive AI?	17 (Programmability) Understand that agents are programmable.

Appendix B: Anonymous course evaluations in Fall 2024 related to AI and opinion paper integration.

The opinion papers were very interesting and I truly enjoyed learning about AI bias, and more.
I enjoyed the group discussions on our opinions on AI as well as the lectures.
Real time lecture on both course material and AI news were very helpful.
Great class and I really enjoyed the weekly discussions on AI.
Definitely a useful course, but sometimes felt a little too broad. We went over a lot of very different topics (AI, coding, math based skills), and so I'm walking away with a lot of useful tools but it felt a bit disconnected at times.
I loved how the opinion papers allowed us as engineers to broaden our perspective on diverse topics of the engineering field.
The course was very interesting and I like the intertwining of AI discussions but sometimes the topics taught were discussed too fast or the methods (for example) were not discussed enough.
I really liked this class, especially the opinion papers and discussions to supplement the more technical learning.
Lectures were super engaging, especially with the opinion papers.
I think that the lectures helped the most for this course; it is where I learned the most. I also think that the opinion papers were also very interesting where we learned about AI's implications and talked about that with the class!

Appendix C:

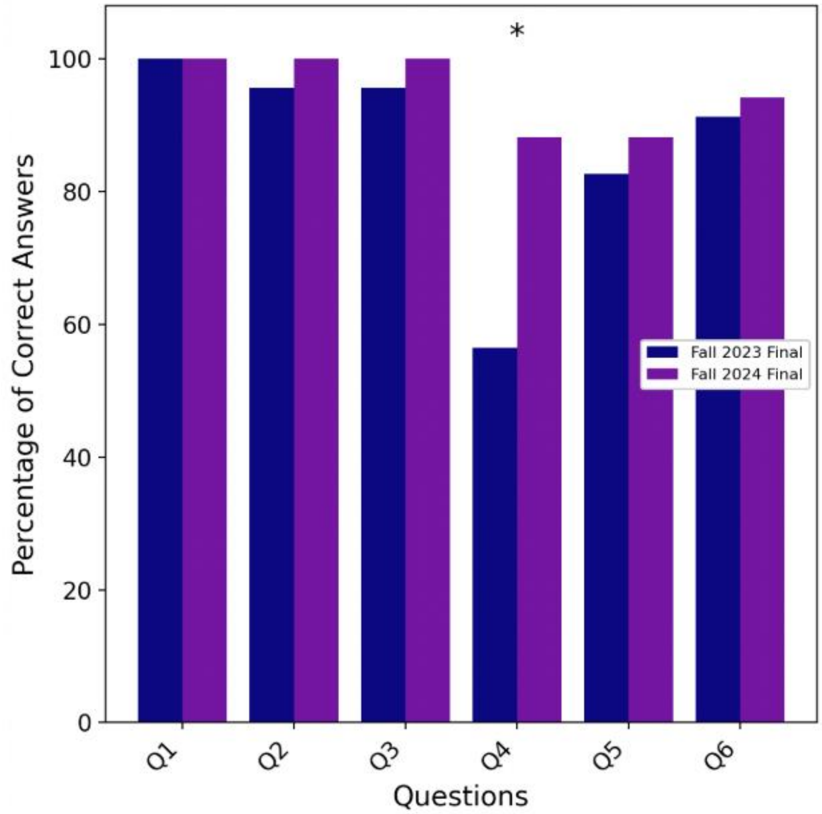

Student changes in knowledge of technical concepts, compared via 2023 and 2024 post-course surveys. There was a statistically significant increase via unpaired ttest ($p = 0.02$) from 2023 ($n = 23$, mean=87%) to 2024 ($n = 17$, mean=95%) on overall scores, largely driven by improvement on Q4. Technical concept survey questions are provided in **Appendix D**.

Appendix D: Concept quiz questions for Computational Biomedical Engineering course. Correct answers are shown in *italics*.

1. What is the goal of nonlinear regression on a function f(t,p) with parameter p?
    a. To maximize f(t,p)
    b. To minimize f(t,p)
    c. To find a value of p that is a root of the sum of squares error
    d. To find a value of p that maximizes the sum of squares error
    e. *To find a value of p that minimizes the sum of squares error*
2. In the vicinity of a function's maximum, which optimization method converges the fastest?
    a. Golden section
    b. Steepest ascent
    c. *Newton's method*
    d. Newton-Raphson method

3. What type of error causes unexpected results when you type 1 - 3*(4/3 - 1) in Python?
    a. Truncation error
    b. *Roundoff error*
    c. Approximate error
    d. Syntax error
4. Which of the following approximation methods is NOT derived from a Taylor Series expansion?
    a. Newton's method
    b. *Gaussian Elimination*
    c. Interpolation
    d. Simpson's 1/3 rule
5. Which of the following is the binary version of the smallest non-negative number that can be represented in an 8-bit floating point number (sign, 3 bits for the signed exponent, and 4 bits for the mantissa)?
    a. *11101000*
    b. 01111000
    c. 01101001
    d. 11101001
6. Which of the following methods can be applied to data points?
    a. Adaptive Euler’s method
    b. *Linear interpolation*
    c. Finite difference method for PDEs
    d. Newton-Raphson method