Artificial Intelligence in the CS2023 Undergraduate Computer Science Curriculum: Rationale and Challenges

Eric Eaton[1] **and Susan L. Epstein**[2,3]

[1]Department of Computer and Information Science, University of Pennsylvania
[2]Department of Computer Science, Hunter College of The City University of New York
[3]Department of Computer Science, The Graduate Center of The City University of New York
eeaton@seas.upenn.edu, sepstein@hunter.cuny.edu

Abstract

Roughly every decade, the ACM and IEEE professional organizations have produced recommendations for the education of undergraduate computer science students. These guidelines are used worldwide by research universities, liberal arts colleges, and community colleges. For the latest 2023 revision of the curriculum, AAAI has collaborated with ACM and IEEE to integrate artificial intelligence more broadly into this new curriculum and to address the issues it raises for students, instructors, practitioners, policy makers, and the general public. This paper describes the development process and rationale that underlie the artificial intelligence components of the CS2023 curriculum, discusses the challenges in curriculum design for such a rapidly advancing field, and examines lessons learned during this three-year process.

Introduction

Artificial intelligence (AI) now pervades and heavily influences our daily lives (Solaiman et al. 2023). Over the past decade, the term "AI" has become commonplace in business news articles and everyday conversation, driven largely by a series of high-impact machine learning (ML) applications (Littman et al. 2021). Terms such as "neural networks" and "deep learning," once relegated to advanced undergraduate or graduate computer science (CS) coursework, have entered the public lexicon. With access to sophisticated AI-driven tools, ordinary individuals can generate essays or poems from a prompt, produce artwork from a description, and create fake photographs and videos that depict real people. AI technologies now trade in stock markets alongside humans, curate our news and social media feeds, control vehicles on public highways, evaluate job applications, detect medical conditions, and influence prison sentences.

Meanwhile, AI and ML techniques have infiltrated many areas of CS. They are now used, for example, in automatic code generation, vulnerability prediction in software engineering, programming language translation, resource allocation in operating systems, load balancing in distributed computing, intrusion detection in networks, and database query optimization. Many disciplines, including biology, chemistry, physics, art, architecture, and finance, increasingly harness AI techniques to address problems within their fields.

AI's prevalence has made it a key focus in the education of computer scientists. For the past few decades, the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers – Computer Society (IEEE-CS) have jointly issued curriculum guidelines for undergraduate CS education. ACM released its first set of CS curricular guidelines in 1968, with a subsequent update in 1978 (Hemmendinger 2007). In 1991, IEEE-CS joined this process; together they published significant updates in 2001, 2008, and 2013. These curricula heavily influence CS major requirements and courses worldwide, and thereby guide the development of the next generation of researchers and practitioners.

ACM and IEEE-CS recognize that those who develop and deploy AI-enhanced technology must understand and consider both its potential and its impact. As a result, in 2020 they invited the Association for the Advancement of Artificial Intelligence (AAAI) to participate in the latest revision of the curricular guidelines for undergraduate CS that would include a focus on AI. Over the last three years, a steering committee of volunteer representatives from AAAI, ACM, and IEEE-CS have developed CS2023 (ACM/IEEE-CS/AAAI Joint Task Force 2023), the latest version of those guidelines. These representatives have been supported by volunteer subcommittees for each of the major CS areas (termed "knowledge areas" in CS2023 parlance) with frequent community feedback from academia, industry, government, and public stakeholders.

This paper describes how AI has been incorporated into CS2023. We first explore the major changes from the previous CS2013 curriculum (ACM/IEEE Joint Task Force 2013), and discuss the challenges to the development of an AI curriculum that could last a decade despite rapid advancement in the field. We then detail the AI knowledge area in the updated CS2023 curriculum and discuss the rationale and development process behind the recommendations. Finally, we examine the lessons learned during the development of the CS2023 AI curriculum.

The CS2023 Curriculum

The CS2023 curriculum is divided into 17 *knowledge areas* (KAs), each of which roughly equates to a major subfield of CS. Each KA is in turn further subdivided into *knowledge units* focused on individual topics or smaller subfields. No-

AI	Artificial Intelligence
AL	Algorithmic Foundations
AR	Architecture and Organization
DM	Data Management
FPL	Foundations of Programming Languages
GIT	Graphics and Interactive Techniques
HCI	Human-Computer Interaction
MSF	Mathematical and Statistical Foundations
NC	Networking and Communication
OS	Operating Systems
PDC	Parallel and Distributed Computing
SDF	Software Development Fundamentals
SE	Software Engineering
SEC	Security
SEP	Society, Ethics, and the Profession
SF	Systems Fundamentals
SPD	Specialized Platform Development

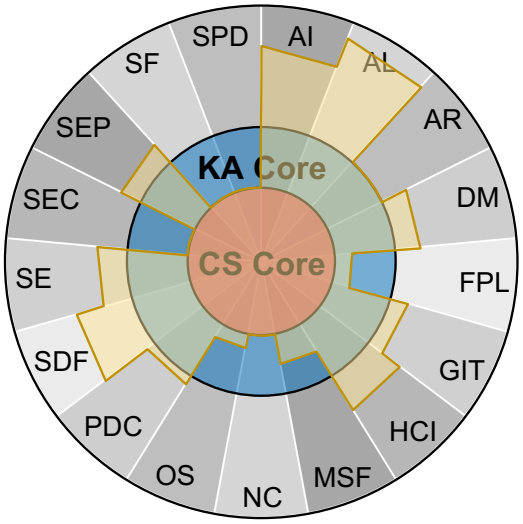

Figure 1: (left) The 17 knowledge areas that compose the CS2023 curriculum. (right) An instantiation of the curriculum at a college or university as a course of study would follow a sunflower model, including the full set of CS Core topics (in red) and selected other knowledge areas (the blue core plus additional elective topics in gray). An example instantiation is shown as the full red CS Core with yellow highlights drawn from the KA Cores and elective topics.

tably, KAs are not equivalent to courses—one course might cover one or more KAs, or a KA might be addressed by multiple courses. Figure 1 depicts the CS2023 KAs as a sunflower model from which to instantiate a curriculum from the guidelines. Each KA identifies a set of *CS Core* topics required for every CS major. In addition, each KA designates significant additional topics essential for those who study the knowledge area in depth (known as the *KA Core*), and a much broader set of optional elective topics.

Crucially, CS2023 must cover CS programs of all sizes, from community colleges to small liberal arts programs to large research universities. Thus, the full CS Core and all the KA Cores minimize their classroom contact hours to allow smaller programs to implement CS2023 successfully. Although CS2023 describes *what* should be taught, not *how* it should be taught, it does include some suggested course packaging. This demonstrates how faculty can flexibly combine topics and knowledge units across KAs into courses in order to fit their needs and yet ensure coverage of the essential topics in the CS Core and the KA Cores.

Summary of Changes from 2013 to 2023

To reflect its recent growth and societal impact, AI in CS2023 has been revised substantially from the previous CS2013 curriculum.

Updated name for the KA. The knowledge area's name has been changed to "Artificial Intelligence" in CS2023. The previous CS2013 included AI topics under the KA name of "Intelligent Systems." CS2023 updates the name to reflect the most common terminology used for these topics within CS and its more widespread use outside CS.

AI requirements in the CS Core. The curriculum now specifies 16 classroom hours of study in AI for all CS majors as part of the CS Core. This ensures that all CS undergraduates will have some minimal foundation in AI. Given the breadth of interconnections between AI and other KAs in practice, the goal is to develop basic AI literacy and critical thought in every CS student. By comparison, CS2013 required no hours at all of AI instruction in its equivalent of the CS Core (called "Tier 1" in CS2013) .

Neural networks and generative models. To reflect recent technological advances, there is now an increased emphasis on neural networks, representation learning, and generative models. The impact of deep neural networks and large language models necessitates that all CS majors be familiar with these topics. Students in AI courses will be provided with greater depth.

Integration of real-world applications, ethics, and societal issues. There is an increased emphasis on practical applications of AI and their social and ethical ramifications. AI's enormous potential for societal impact demands attention and consideration to those aspects from practitioners who apply AI. Consequently, students must study real-world AI applications and their implications. This encompasses explicit discussion of tools that employ deep generative models (e.g., ChatGPT, LLaMA, DALL-E, Midjourney) and are now in widespread use, including high-level coverage of how these models work, how they can be used, and the shortcomings and pitfalls in their application. The CS2023 curriculum highlights how important it is to understand and assess the broader societal impacts and implications of AI methods and applications, including issues in AI ethics, fairness, trust, and explainability. Study of the AI techniques and applications that have the most impact, particularly deep generative models, is intimately tied to examination of these societal and ethical issues.

Relationship to data science. Data science was among the most important topic areas in feedback surveys for CS2013 (ACM/IEEE-CS/AAAI Joint Task Force 2023). As a direct result, CS2023 includes a suggested Data Science packaging for AI topics.

Challenges and Important Considerations

There were two major challenges in the creation of the AI curriculum: stability and size. The first challenge was to devise guidelines that would remain current for the ten-year expected lifespan of the CS2023 document, even while AI itself continues to develop rapidly. Indeed, in the 18 months between the first draft and the final version of CS2023, several new techniques (e.g., generative networks and large language models) became widely used. They were essential to incorporate into both the CS Core and several KA Cores. Such rapid change in AI is likely to continue for the foreseeable future. These forthcoming AI advances must be considered when faculty use the CS2023 AI curriculum. Until the next revision, those who teach AI will have to gauge which advances should be included to keep the curriculum current.

The second challenge was minimizing the set of CS Core and KA Core topics to address the needs of smaller CS programs. This requirement constrained the number of CS Core and KA Core hours that could be required, and therefore restricted the topics the Cores include. The CS Core and KA Core hours should therefore be viewed in that context, as the minimum amount of time to be devoted to these topics. Programs with more flexibility are encouraged to devote additional classroom hours to those topics, at the discretion of individual instructors.

The AI Knowledge Area in CS2023

This section provides details on the structure of the AI KA and an overview of its topics.

Structure. The CS2023 Artificial Intelligence knowledge area is divided into 12 knowledge units:

- Fundamental Issues
- Search
- Fundamental Knowledge Representation and Reasoning (KRR)
- Machine Learning
- Applications and Societal Impact

CS Core, KA Core, Elective

- Probabilistic Representation and Reasoning
- Planning
- Logical Representation and Reasoning
- Agents and Cognitive Systems
- Natural Language Processing
- Robotics
- Perception and Computer Vision

Elective only

Topics in the first five knowledge units are partitioned among the CS Core, the KA Core, and electives (all topics outside of the CS Core and the KA Core). The remaining seven knowledge units contain elective topics only. Figure 2 provides an overview of the topics within the first five units that compose the CS and KA Cores.

Table 1 compares the recommended classroom hours in the CS2023 AI curriculum for the CS Core and KA Core knowledge units with equivalent hours in CS2013. CS2023 places significant emphasis on AI compared to CS2013, with 16 hours of AI instruction now required for all CS majors, and 18 additional hours for those who focus on AI. CS2013 had zero required hours of AI instruction in its Tier 1 (analogous to CS2023's CS Core) and only ten hours of AI instruction in its Tier 2 (analogous to CS2023's KA Core and required of approximately 80% of all CS majors). The total CS Core, across all 17 knowledge areas, is 270 hours, so AI topics constitute 5.93% of the CS2023 CS Core curriculum.

Note that the recommended CS and KA Core hours for AI represent the minimum requirements, with many CS programs offering more hours of instruction on these topics. As stated previously, the CS and KA Cores and their recommended hours are deliberately kept as small as possible to meet the needs of smaller CS programs.

	CS2023		CS2013
Knowledge Unit	**CS Core**	**KA Core**	**Tier 2**
Fundamental Issues	2	1	1
Search	5	6	4
Fundamental KRR	2	2	3
Machine Learning	4	6	2
Applications and Societal Impact	3	3	-
Total	16†	18	10

Table 1: CS2023 recommended core hours for AI. For comparison, CS2013 hours are shown in gray. CS2013 required zero hours of AI instruction in its Tier 1 and ten hours of AI in its Tier 2, roughly the equivalent of CS2023's CS Core and KA Core, respectively. †Included in the total are four CS Core hours that are shared with other KAs: three hours with Algorithms (uninformed search) and one hour with Mathematical and Statistical Foundations (probability).

Topics. The chosen topics in the AI KA differ considerably from CS2013. CS2023 requires the study of neural networks, generative models, AI applications, societal impact, and working with data—topics not included in CS2013's recommendations. Given its key role throughout AI, search is still emphasized in both the CS and KA Cores. CS2023 places more emphasis, however, on understanding subsymbolic methods and learned representations than on symbolic methods. It still retains knowledge-based and symbolic approaches within the AI curriculum because they offer unique capabilities, are commonly used in practice, ensure a broad education, and support interpretability. In addition, the CS2023 AI recommendations anticipate increased future attention to neurosymbolic approaches that integrate both learned and symbolic representations.

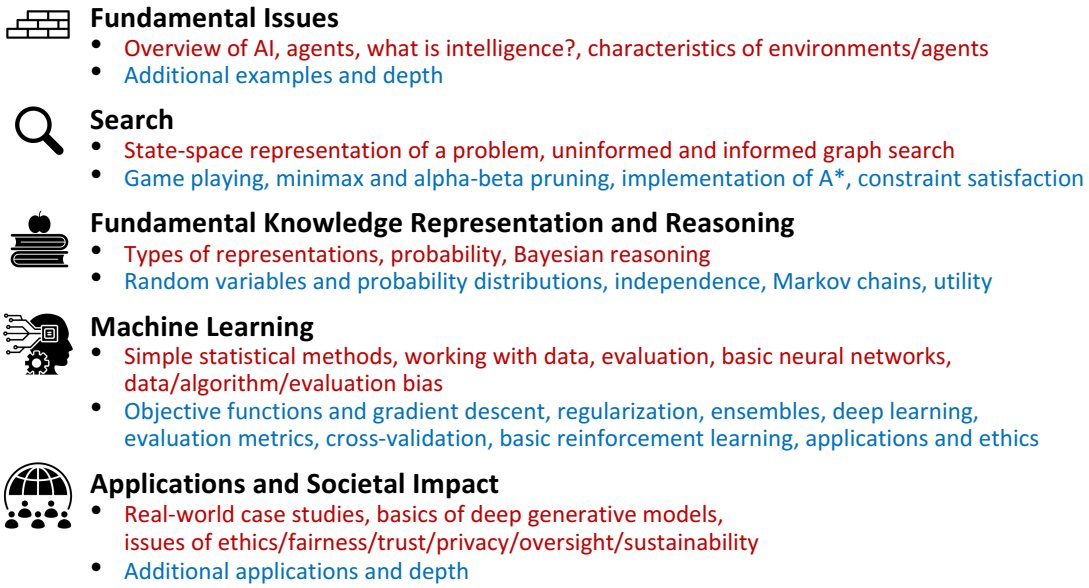

Figure 2: Summary of the CS Core (in red) and KA Core (in blue) topics in each of the first five AI knowledge units. Note that these lists of topics are not comprehensive and should not be relied upon to develop a curriculum. See the full CS2023 curriculum document for a complete list of all CS Core and KA Core topics under these knowledge units.

The topics selected for the AI CS Core and KA Core, along with their recommended classroom hours, were chosen to ensure that all CS students acquire basic literacy in AI. All CS students should know:

- How to use uninformed and informed graph search to solve a problem
- The basic principles of probabilistic reasoning, simple statistical models, and neural networks
- How to work with and clean data
- The sources of bias in machine learning
- The basics of deep generative models
- How these methods have been applied in real case studies
- The associated societal issues of ethics, fairness, privacy, oversight, and sustainability.

The KA Core expands these topics to provide a more in-depth education typical of an undergraduate AI course:

- Techniques for adversarial reasoning and game-playing
- Details on how to implement informed search
- Constraint satisfaction
- In-depth probabilistic reasoning and utility
- Mathematical formulations of machine learning, more advanced evaluation methodology, basic reinforcement learning, deep learning, and ML ethics
- In-depth explorations of applications and societal issues of AI, with a focus on deep generative models.

Knowledge Units in the CS Core and the KA Core

We now briefly examine each of the knowledge units with topics that are included in the CS and KA Cores (Figure 2).

Fundamental Issues. This knowledge unit includes topics typically covered in introductory lectures on AI. For the CS Core, students should receive a general introduction to the field of AI, its notable successes and applications, the concept of agents and their properties, and how to analyze the characteristics of problems that can be solved by AI. The KA Core extends this coverage with additional examples, discussions, and depth, especially on applications.

Search. The Search knowledge unit addresses how to frame and solve problems using graph search, including both uninformed and informed (e.g., A*) search in the CS Core. Consequently, it partially overlaps with topics covered under the Algorithms KA in CS2023. These topics form the foundation for solving practical search problems and thinking of ML optimization as a search problem. The KA Core extends these topics with an exploration of adversarial games, including minimax with alpha-beta pruning, details for implementing informed search, and constraint satisfaction.

Fundamental Knowledge Representation and Reasoning This knowledge unit, for the CS Core, covers different kinds of representations (e.g., symbolic vs subsymbolic, learned vs specified) and graphical models. It also reviews basic probability concepts and how to use them for basic Bayesian reasoning. The KA Core takes a more rigorous approach to provide the foundations for understanding prediction, uncertainty in machine learning, and generative models. It covers random variables and probability distributions, independence and conditional independence, Markov chains, and utility and decision making.

Machine Learning. Because ML underlies most recent AI advances, the ML knowledge unit has received a substantial increase in recommended core hours over those in

CS2013. Since many CS students will be called upon to apply ML models, topics in the CS Core provide a basic overview of machine learning, working with data, and the pitfalls of prediction. Topics include the kinds of ML problems (e.g., supervised, unsupervised, semi-supervised), a simple statistical classifier (e.g., linear regression or decision trees), and the concepts of overfitting and representation. The AI CS Core also includes basic neural networks, how to work with data (including preprocessing, missing values, and how human choices affect ML outcomes), basic evaluation procedures, and ML ethics with a focus on real-world applications and societal issues.

The KA Core provides significantly more depth on current techniques and mathematical formulations, including how to frame an ML problem as an objective function and optimize it (e.g., via gradient descent), the basics of ensemble models, deep learning and basic reinforcement learning, more advanced evaluation methods and analysis, and in-depth examinations of ML case studies and ethics, including fairness and privacy.

Applications and Societal Impact. This material spans all AI knowledge units. It examines one or more applications of AI techniques to a broad set of problems and diverse fields, such as medicine, health, sustainability, social media, economics, education, and robotics. Within the context of a specific application, students should study how to formulate a task as an AI problem, dealing with ill-posed problems, data availability and cleanliness, and biases in data sets, algorithms, and evaluation. These aspects are particularly important to bridge between the study of AI methods and their practical use (Gautam, Akgun, and Mitra 2023).

The societal impact of an application is key here—every application raises questions about ethics, fairness, trust, explainability, privacy and the use of training data, human autonomy and oversight, and sustainability. This knowledge unit also includes discussions of deployed deep generative models with reference to state-of-the-art AI methods. For the CS Core, coverage is expected to focus on one specific application for this analysis. For the KA Core, this knowledge unit should include additional applications to provide more breadth, and substantially greater depth on deep generative models and state-of-the-art AI methods.

Elective topics. Although elective topics are not required by either the CS Core or the KA Core, courses that instantiate CS2023 will likely include many of them. All five of the knowledge areas listed above include elective topics. In addition, there are seven elective knowledge units within the AI KA: agents and cognitive systems, logical representation and reasoning, probabilistic representation and reasoning, planning, natural language processing, robotics, and perception and computer vision. This allows instructors to customize their individual courses and still meet the CS2023 requirements. In larger programs, most knowledge units could reasonably appear as full courses.

Integration of AI with Other Knowledge Areas

While it is possible to cover all the AI CS Core topics by a required AI course, it may be more realistic to integrate them into existing required courses. Such an approach would help emphasize the applications of AI to other CS areas. For example, courses on data structures or algorithms could incorporate problem solving with uninformed and informed graph search, courses on the mathematical foundations of CS could cover probabilistic reasoning, and a senior-level capstone course could include data analytics via machine learning. This allows CS departments to package the required CS Core topics and knowledge units into courses flexibly—a necessity for smaller CS programs. The required topics could then be complemented by an AI elective that combines the KA Core with additional topics.

AI's interconnections with several other knowledge areas in CS2023 enables the creation of crosscutting courses. For example, a robotics course could combine the search, ML, and robotics knowledge units under AI with knowledge units from the Specialized Platform Development KA, which includes robotic platforms, design considerations, and middleware (e.g., ROS, the Robot Operating System). As another example, a course on data science could combine knowledge units from the Data Management, Graphics and Interactive Techniques, and the Human-Computer Interaction KAs with the ML knowledge unit from AI.

Mathematics and Statistics Requirements

CS2023 outlines specific mathematics and statistics requirements. It became abundantly clear during discussions within the CS2023 steering committee that there is a need for a CS-centric prerequisite mathematics course that combines topics in math and statistics with programming and applications. This is in contrast to outsourcing the math and statistics education of CS students to a Mathematics department, which may induce a separation between topic coverage and its practical use within CS. Many universities have already implemented such a course within their CS programs.

For AI, CS2023 recommends requiring the following mathematics and statistics topics:

- discrete mathematics on sets, relations, and functions
- linear algebra, including linear functions, matrix operations and algebra, and basis sets
- probability and statistics, including basic probability theory, conditional probability, independence, Bayes theorem and its applications, expected value, distributions, basic descriptive statistics, and significance testing.

Additionally, it is desirable that students who study AI have a background in calculus-based probability and statistics, optimization, and other statistical topics such as hypothesis testing, data resampling, and experimental design.

Lessons Learned

This section briefly discusses some of the key lessons learned during the development of the AI curriculum.

Widely differing opinions. There is no universal agreement on the ideal contents of an AI curriculum. Even among the volunteers who developed CS2023 and its AI curriculum, there is some disagreement as to what should or should not be included in the CS Core and KA Core for AI. This

divide only broadened as other stakeholders, including academics from all sizes of institutions, industry, government, and the public, offered feedback on different drafts of the curriculum. For example, when we solicited feedback from the AAAI general membership on a near-final version of the curriculum in October 2023, largely the one described above, we received widely varied opinions, such as:

- It shortchanges classical AI techniques
- It looks archaic, without proper coverage of modern deep learning methods that students need to find current jobs
- It lacks sufficient emphasis on ethical implications of AI.

Further details on these and other comments, and the committee's response to each, is available in the revision report included with the CS2023 curriculum. In each case, the committee has tried to take a balanced approach in determining what to include in the curriculum. Indeed, the feedback above led, respectively, to inclusion of constraint satisfaction in the KA Core for search, increased hours and depth on deep learning and deep generative models, and additional hours devoted to ethics and the societal implications of AI.

Core size. Smaller academic programs require a smaller Core while larger academic programs and industry want a larger one. There is pressure both to expand and to reduce the set of Core topics and hours. In particular, the CS Core must be achievable for smaller programs while including all essential topics. Consequently, the final CS2023 AI curriculum is admittedly one-size-fits-none. Nonetheless, we believe that it is achievable for every undergraduate CS program, regardless of its size. Although it is a compromise across all stakeholders, it includes those AI topics currently necessary for a modern CS undergraduate education.

AI as a moving target. The fast pace of AI developments makes it challenging to craft a curriculum that will remain current for some years to come. Consequently, we focused the curriculum on the essential techniques and ideas within AI, while including freedom for instructors to choose current AI applications of those techniques.

Society and ethics. Societal and ethical issues are intertwined with applications. Rather than separate them out into their own knowledge unit, which could encourage their isolation into separate, dedicated class sessions, these AI curricular guidelines insist that students learn about societal and ethical issues in the context of AI applications. Indeed, the same underlying AI technique may provoke different issues depending on the context of its application (e.g., consider using a large language model to generate video game dialogue versus having it summarize legal texts to craft a trial defense). The study of societal and ethical issues in the context of an application both grounds the discussion and allows detailed examination of the ways technical choices in the underlying AI techniques impact those issues. CS2023 does include some knowledge-unit-specific societal and ethical issues (e.g., the ethics and privacy of data in the Machine Learning knowledge unit). However, most societal and ethical issues apply across all knowledge units and are specified for examination across all applications.

Closing Thoughts

The CS2023 AI curriculum will require adaptation as AI progresses. Thus instructors should use their judgement as this curriculum ages. As of this writing, however, we believe that the ACM/IEEE/AAAI curriculum is a current assessment of what matters and how it might be delivered effectively. We hope that it proves useful to other instructors and institutions as they revise their own CS curricula to provide a modern and exciting education and to produce thoughtful, responsible professionals.

Acknowledgments

Both authors are the AAAI-designated representatives to the CS2023 Task Force and members of the CS2023 AI subcommittee. This paper is based on a near-final version of the CS2023 curriculum, but it is possible that the final published CS2023 curriculum may differ slightly. We gratefully acknowledge the contributions of the other members of the CS2023 AI subcommittee: Zachary Dodds, Laura Hiatt, Amruth N. Kumar, Peter Norvig, Meinolf Sellmann, and Reid Simmons. We also wish to acknowledge the contributions of Nate Derbinsky, Eugene Freuder, Ashok Goel, and Claudia Schulz to the AI curriculum, and the many anonymous reviewers who provided valuable feedback on each AI curriculum draft version. Finally, we are grateful to the EAAI 2024 reviewers, whose feedback improved this paper.

References

ACM/IEEE-CS/AAAI Joint Task Force. 2023. Computer Science Curricula 2023. https://csed.acm.org/. Accessed: 2023-09-10.

ACM/IEEE Joint Task Force. 2013. Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science.

Gautam, S.; Akgun, M.; and Mitra, P. 2023. Exploring the Challenges of AI Experts to Inform AI Curriculum. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 2*, SIGCSE 2023, p.1338.

Hemmendinger, D. 2007. The ACM and IEEE-CS Guidelines for Undergraduate CS Education. *Communications of the ACM*, 50(5): 46–53.

Littman, M.; Ajunwa, I.; Berger, G.; Boutilier, C.; Currie, M.; Doshi Velez, F.; Hadfield, G.; Horowitz, M.; Isbell, C.; Kitano, H.; Levy, K.; Lyons, T.; Mitchell, M.; Shah, J.; Sloman, S.; Vallor, S.; and Walsh, T. 2021. Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report. Stanford University, Stanford, CA. http://ai100.stanford.edu/2021-report.

Solaiman, I.; Talat, Z.; Agnew, W.; Ahmad, L.; Baker, D.; Blodgett, S. L.; Daumé III, H.; Dodge, J.; Evans, E.; Hooker, S.; Jernite, Y.; Luccioni, A. S.; Lusoli, A.; Mitchell, M.; Newman, J.; Png, M.-T.; Strait, A.; and Vassilev, A. 2023. Evaluating the Social Impact of Generative AI Systems in Systems and Society. arXiv:2306.05949.