

In this issue:

- 4. Embracing Artificial Intelligence to Improve Self-Directed Learning: A Cybersecurity Classroom Study**
Jim Marquardson, Northern Michigan University
- 14. Student-Driven Programming Instruction: A Follow-Up Study**
Deepak Dawar, Miami University
- 30. Investigating FinTech Education and Training in Australian Universities**
Pak-Lok Poon, Central Queensland University
Santoso Wibowo, Central Queensland University
Srimannarayana Grandhi, Central Queensland University
Sau-Fun Tang, Central Queensland University
- 41. A Pragmatic Approach to Investigating the Digital Existence of Food Bank Users**
Yaprak Dalat Ward, Fort Hays State University
James G. Ward, Fort Hays State University
Li-Jen Lester, Sam Houston State University
- 53. An Eye Toward the Softer Side of CC2020 Computing Curricula: Professional, Legal, and Ethical Artificial Intelligence Issues**
Christine Ladwig, Southeast Missouri State University
Dana Schwieger, Southeast Missouri State University
- 89. Redesigning CS 100 in the Context of a Changing University Curriculum**
Mark Frydenberg, Bentley University
William (Bill) VanderClock, Bentley University

The **Information Systems Education Journal** (ISEDJ) is a double-blind peer-reviewed academic journal published by **ISCAP** (Information Systems and Computing Academic Professionals). Publishing frequency is five times per year. The first year of publication was 2003.

ISEDJ is published online (<https://isedj.org>). Our sister publication, the Proceedings of the ISCAP Conference (<https://iscap.us/proceedings>) features all papers, abstracts, panels, workshops, and presentations from the conference.

The journal acceptance review process involves a minimum of three double-blind peer reviews, where both the reviewer is not aware of the identities of the authors and the authors are not aware of the identities of the reviewers. The initial reviews happen before the ISCAP conference. All papers, whether award-winners or not, are invited to resubmit for journal consideration after applying feedback from the Conference presentation. Award winning papers are assured of a publication slot; however, all re-submitted papers including award winners are subjected to a second round of three blind peer reviews to improve quality and make final accept/reject decisions. Those papers that are deemed of sufficient quality are accepted for publication in the ISEDJ journal. Currently the target acceptance rate for the journal is under 35%.

Information Systems Education Journal is pleased to be listed in the Cabell's Directory of Publishing Opportunities in Educational Technology and Library Science, in both the electronic and printed editions. Questions should be addressed to the editor at editor@isedj.org or the publisher at publisher@isedj.org. Special thanks to volunteer members of ISCAP who perform the editorial and review processes for ISEDJ.

2024 ISCAP Board of Directors

Jeff Cummings
Univ of NC Wilmington
President

Amy Connolly
James Madison University
Vice President

Eric Breimer
Siena College
Past President

Jennifer Breese
Penn State University
Director

David Gomillion
Texas A&M University
Director

Leigh Mutchler
James Madison University
Director/Secretary

RJ Podeschi
Millikin University
Director/Treasurer

David Woods
Miami University
Director

Jeffry Babb
West Texas A&M University
Director/Curricular Items Chair

Tom Janicki
Univ of NC Wilmington
Director/Meeting Facilitator

Paul Witman
California Lutheran University
Director/2024 Conf Chair

Xihui "Paul" Zhang
University of North Alabama
Director/JISE Editor

Copyright © 2024 by Information Systems and Computing Academic Professionals (ISCAP). Permission to make digital or hard copies of all or part of this journal for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial use. All copies must bear this notice and full citation. Permission from the Editor is required to post to servers, redistribute to lists, or utilize in a for-profit or commercial use. Permission requests should be sent to Paul Witman, Editor, editor@isedj.org.

INFORMATION SYSTEMS EDUCATION JOURNAL

Editor Comment:

We are pleased to present our first issue of 2024, including papers on AI in and out of the classroom. We are also delighted to welcome Dr. Kevin Mentzer to the editorial team this year, joining Dr. Dana Schwieger as Associate Editor, along with Drs. Ira Goldstein and Michelle Louch as Cases & Exercises Co-Editors. Thank you to each of these individuals, to our reviewers, and to our Publisher, Dr. Tom Janicki.

Paul Witman, Senior Editor, ISEDJ

Editors

Paul Witman
Editor
California Lutheran
University

Thomas Janicki
Publisher
U of North Carolina
Wilmington

Dana Schwieger
Associate Editor
Southeast Missouri
State University

Kevin Mentzer
Associate Editor
Nichols College

Ira Goldstein
Teaching Cases & Exercises
Co-Editor
Siena College

Michelle Louch
Teaching Cases & Exercises
Co-Editor
Duquesne University

Donald Colton
Emeritus Editor
Brigham Young University
Hawaii

Jeffry Babb
Emeritus Editor
West Texas A&M
University

An Eye Toward the Softer Side of CC2020 Computing Curricula: Professional, Legal, and Ethical Artificial Intelligence Issues

Christine Ladwig
cladwig@semo.edu
Department of Marketing

Dana Schwieger
dschwieger@semo.edu
Department of Management

Southeast Missouri State University
Cape Girardeau, MO 63701, USA

Abstract

Hollywood screenwriters worry about Artificial Intelligence (AI) replacements taking over their jobs. Famous museums litigate to protect their art from AI infringement. A major retailer scraps a machine-learning based recruitment program that was biased against women. These are just a few examples of how AI is affecting the world of work, learning, and living. MIS and computer science students are among the professional groups who are embarking into careers with nebulous frontiers obscured by the outcroppings brought on by AI. Computer Science and Information System curriculum task forces have recognized the increasing ethical and professional implications developers' work can have beyond the scope of the programmers' code. In this article, the authors examine the professional, legal, and ethical implications of copyrights and algorithmic bias resulting from development of AI-enhanced applications and offer suggestions for addressing these topics in courses considering changes to the CC2020 and IS2020 Model Curriculum frameworks.

Keywords: Artificial Intelligence, Copyright infringement, Teaching strategies, Risk management, Programming bias, CC2020 Computing Curricula, IS2020 framework

Recommended Citation: Ladwig, C., Schwieger, D., (2024). An Eye Toward the Softer Side of CC2020 Computing Curricula: Professional, Legal, and Ethical Artificial Intelligence Issues. *Information Systems Education Journal*, 22(1), pp.53-88. <https://doi.org/10.62273/FFDJ1170>

An Eye Toward the Softer Side of CC2020 Computing Curricula: Professional, Legal, and Ethical Artificial Intelligence Issues

1. INTRODUCTION

Artificial Intelligence (AI) is changing the landscape of development, for the good, the bad, and the in-between (Bolukbasi, Chang, Jou, Saligrama, & Kalai, 2016; Hogan, 2019; Nelson & Reed, 2023). As awareness of the potential devastating effects that can result from AI enhanced systems expands, so does the need for protective measures. Governments and industrial leaders are calling for AI development guidelines; for example, in October 2022, the U.S. White House issued its *Blueprint for an AI Bill of Rights*, intended to provide guidance in the design and execution of machine learning systems (OSTP, 2022). On April 11, 2023, the Cyberspace Administration of China (CAC) revealed its draft generative artificial intelligence services policy (Ye, 2023). And on June 14, 2023, European lawmakers passed a landmark Artificial Intelligence Act that provides guardrails for AI development including classification levels for AI risk, "greater privacy standards, stricter transparency laws, and steeper fines for failure to cooperate" (Sharp, 2023, p. 1). While these guardrails are still evolving, it is important that future system developers be familiar with their underlying implications. Such rapid advancements generate various unknowns in many areas including: software development, human resources, privacy, security, ethics, regulatory implications, copyrights, education, etc. Due to the potential impact on the MIS and computer science professions, there are many critical factors — technical, ethical, regulatory, and professional — that should be considered when developing and updating course curricula.

Recent updates to the computing disciplines curriculum guidelines (CS2020; IS2020) have recognized the increasing importance of ethics and professionalism. For example, Section 6.5 of the Computing Curricula 2020 framework addresses the importance of incorporating professionalism and ethics, indicating that "it should be a permanent element of any computing curriculum" (CC2020 Task Force, 2020, p. 76). The framework suggests that concepts could be taught in dedicated courses either inside or outside the computing discipline or distributed over the body of curricula.

The Organizational Domain in the IS2020 framework has likewise been modified to include two required competency areas: "IS Management and Strategy/Ethics" and the "Use and Implications for Society" (IS2020 Task Force, 2020, p. 57). In making this change to the framework, the task force noted, "As IT is being deployed increasingly outside the traditional business organization context, and also incorporated closely to products and services for consumers, there are new ethical challenges to comprehend and address" (p. 57).

Therefore, in this paper, the authors discuss two nascent areas with significant ethical and professional considerations that could be efficiently highlighted in the MIS and/or computer science classroom: programming bias and copyright infringement protection for both programmer-developed and AI-developed works. The authors additionally provide exercises and an example of a CC2020 Competency Statement that may be used to incorporate these topics into computer science, information system, technology management, and/or business curriculums.

2. FRAMEWORKS FOR THE STUDY OF AI ISSUES

Two recent developments — one related to computer science education guidelines, and another connected to the current rise of Artificial Intelligence use — may be used by educators to incorporate important AI issues into the classroom:

- 1) The 2020 Computing Curriculum Task Force (CC2020) recently released a revised "template for specifying the subject matter of baccalaureate computing education" building upon the different educational frameworks provided in the IT2017 report (CC2020 Task Force, 2020, p. 47). The CC2020 guidelines move from knowledge-based learning to competency-based learning (Ormond, 2021) and define competency as the combination of knowledge (know-what), skills (know-how), and dispositions (know-why) (K-S-D) *in task* (CC2020 Task Force, 2020, p. 13). When the

K-S-D construct is tied to the performance of a task, CC2020 indicates that this “frames the skilled application of knowledge and makes dispositions concrete” (CC2020 Task Force, 2020, p. 48). The *task* therefore becomes the embodiment of the purpose for the competency, and “competency statements” — task descriptions that align with the relevant knowledge elements, skill level and disposition — become the learning delivery platform. The IS2020 task force adopted the CC2020 competency approach to developing their framework as well (IS2020 Task Force, 2020, p. 39). Thus, using the original CC2020 framework as a guide, the authors created a sample competency statement (Appendix A) for the risk management of legal and ethical issues in machine learning which may be used in conjunction with the exercises included in Appendices B and C.

- 2) In October 2022, the U.S. White House Office of Science and Technology Policy issued a *Blueprint for an AI Bill of Rights* consisting of “five principles that should guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence” (OSTP, 2023, p. 1). The Blueprint notes that AI technologies “can drive great innovations, like enabling early cancer detection or helping farmers grow food more efficiently”, but conversely, the same advancements are “too often developed without regard to their real-world consequences and without the input of the people who will have to live with their results” (OSTP, 2023, p. 1). Using the AI Bill of Rights as a guide, the authors suggest additional content and strategies for addressing the “danger zone” issues associated with the meteoric rise of AI system use.

In Sections 3 and 4, the authors describe two (of the many) professional, ethical, and legal considerations — software copyrights and algorithmic bias — that raise significant issues for industry as the rapid advancement of AI continues.

3. IT RISK MANAGEMENT ISSUE: SOFTWARE COPYRIGHTS

As students prepare to enter the computer science and MIS industries, it is important that they understand intellectual property and the legal conventions protecting it, especially as AI begins to play a role in the generation of products and content. Software is an example of intellectual property (IP), or an intangible original work that has value for its developer. It is

essential that entities safeguard their IP investments by applying legal protections such as copyrighting and patenting. Because software is considered a “literary work” under U.S.C. §101 of the Copyright Act, a program’s written code can be protected from being copied and used without the developer’s permission.

The U.S. Copyright Act (1976) allows copyright holders to:

- Make copies of their software,
- Distribute the work (sell it)
- Make “derivative works”, and
- Share or perform the work in public.

As soon as the software is expressed as an operational program, it is automatically copyrighted with the developer as its owner. However, officially registering the software for a copyright with the U.S. Copyright Office and the Library of Congress is a good idea, as it gives the developer more legal protection in case someone tries to steal the software or use it without the developer’s permission (also known as “infringement”). A copyright lasts for the lifetime of the author, plus 50 years in many countries, and for 70 years in others, including the United States. Obtaining a copyright is relatively inexpensive (\$45-\$125) and usually only takes about three months. Copyrightable software includes a vast array of programs and components, such as graphical user interfaces, mobile phone apps, video conferencing software, animated graphical sequences, soundtracks or sound effects, and social media platforms.

Although software is primarily protected by copyrights, patents can be used to protect the underlying “ideas, procedures, and operational/computing methods” behind the software (Rouse, 2013). However, patents are more costly and the application process more complex.

Copyright Exceptions: First Sale Doctrine & Fair Use

Under the First Sale Doctrine, the user/owner of a purchased or legally obtained copy of a work is entitled to sell that copy without infringing copyright. The legal owner is not allowed, however, to make and sell copies of the owned copy — or even to give them away. Copyright law will sometimes allow a legal owner to make archival copies for personal use. For instance, you are permitted to make a single copy of a legally obtained computer software program in case the original is lost or damaged, but that back-up copy must be destroyed or transferred if you ever sell the original copy.

The Fair Use rule is a part of copyright law that allows for the limited use of copyrighted material without permission from the rights holder. In *Galoob v. Nintendo*, the 9th Circuit Court held that modification of copyrighted software for personal use was fair (Farris, 1992). At issue in *Galoob* was use of the *Game Genie* accessory, a device that could alter the output of video games in the Nintendo Entertainment System. The Court determined that the *Genie* did not make derivative works of Nintendo's games (which would violate their software copyright) and qualified as non-commercial fair use. In *Sega v. Accolade*, the 9th Circuit held that making copies during reverse engineering is a fair use, when it is the only way to get access to the "ideas and functional elements" in the copyrighted code, and when "there is a legitimate reason for seeking such access" (Reinhardt, 1992). More recently in 2021, The U. S. Supreme Court ruled in *Google LLC v. Oracle America, Inc.* (2021) that the reuse of application programming interfaces (APIs), including representative source code, can be transformative and fall within fair use. However, they did not rule on whether such APIs are copyrightable.

Copyright Infringement and AI

If not complex enough, the legal environment of software and program copyright now faces a new era of issues as developer products are being used to support artificial intelligence or create new works generated by AI. Lawsuits have been launched regarding copyrights potentially infringed by AI-enhanced programs, either in the use of copyrighted training materials or the expression of substantially similar copyrighted content in AI generated works. On March 16, 2023, the U.S. Copyright office launched an Artificial Intelligence Initiative to examine copyright law issues created by the rise of machine learning. The study is intended to examine "the scope of copyright in works generated using AI tools and the use of copyrighted materials in AI training" (U.S. Copyright Office, 2023).

With sophisticated AI technologies training on vast quantities of human-created content and producing expressive materials, the danger of copyright infringement is omnipresent. There are a number of concerning issues evolving from the current AI landscape, including: (1) Is AI-generated content copyrightable? (2) If AI-generated content is copyrightable, who owns that copyright? (3) Does AI-generated content infringe copyrighted materials, such as those used for training the program? (4) If AI-

generated content does infringe copyrighted material, who is liable?

In addressing the first and second questions, the Copyright Office has stated that "it is well-established that copyright can protect only material that is the product of human creativity" (U.S. Copyright Office, 2023, p. 2). Works generated through an AI prompt would be made by a non-human machine, and therefore not be eligible for registered copyright protection. But, what about human manipulation of that AI generated work? How much human intervention is required to move the AI content over the threshold of including human authorship? And what about the developers of the complex algorithms and code that comprise the AI enhanced program — should these individuals be recognized as the authors of the AI generated works?

In addressing the third question, there is already evidence that AI enhanced programs are potentially infringing copyrighted material by training on — and then using in generated works — existing copyrighted material. In 2021, Microsoft, its subsidiary GitHub, and its business partner OpenAI were sued in a class action lawsuit which alleged that "the companies' creation of AI-powered coding assistant GitHub Copilot relies on 'software piracy on an unprecedented scale.'" (Vincent, 2023). It is well known that many text-to-image AI, like the open-source program Stable Diffusion, are also created by scraping copyrighted material from the web. (Scraping is a technical approach to extract text and images from a web page for use as raw material for training.) Although AI firms contend that these actions are a copyright fair use exception, experts suggest that this is far from settled law (Vincent, 2023).

Regarding the fourth question — liability for AI infringement — plaintiffs in legal actions are assigning the companies behind the AI enhanced programs with this responsibility. Author Ellen Glover (2023) details the lawsuit filed by Getty Images (of the Getty Museum) against Stability AI (Stable Diffusion) for "copying and processing millions of [Getty] images that are protected by copyright, as well as their associated metadata, without getting permission or providing compensation." According to Ben Zhao, a computer science professor at the University of Chicago, "The large majority of independent artists make their living through commissioned works. And it is essential for them to keep posting samples of their art. But the websites they post their work on are being scraped by AI-enhanced

programs in order to learn and then mimic that particular style. Artists are literally being replaced by models that have been trained on their own work" (Glover, 2023, p. 3).

Similarly, Hollywood writers were on strike in April 2023, with a primary complaint of potential copyright infringement of their existing works. Their fears were based in studios taking their prior generated scripts and using AI to generate new stories and writings for film and television — without the human writers' involvement. Because the studios hired the writers to create the material originally (known as "works for hire"), they could train AI on these prior scripts without potentially infringing the copyrights that they hold. The ethics of these actions, however, are highly questionable. So, although the state of the law is currently dynamic regarding the questions of copyrightability and infringement, it is vital that programmers maintain an understanding of the vital role this issue has, and will continue to have, as the AI generation of content accelerates in scope. And in addition to copyright issues, they must also consider the potential for inherent bias unintentionally introduced, and magnified, in AI supported programs by algorithm training data.

4. IT RISK MANAGEMENT ISSUE: PROGRAMMING AND/OR TRAINING DATA BIAS

The prevalence of AI algorithms in areas in which life-altering decisions may occur, e.g., healthcare, transportation, job placement, school admission, loans, etc., is escalating. Although great efforts have been made to develop accurate AI algorithms to assist decision-makers in making high-quality consistent decisions, that is not always the case (Parikh, 2021). Whether intended or not, bias may be incorporated into a program due to the nature of the data used to develop and/or train the system (Sparkes, 2022). Stories of racial and gender bias have been known for some time (Caliskan, Bryson, & Narayanan, 2017; Sparkes, 2022) and embedded gender bias has even been found in general text references and correlations between pronouns and roles on Internet news searches (Bolukbasi, et al., 2016.)

Programming Bias in Employment and Hiring

Researcher Hogan (2019, p. 2) found that "most hiring algorithms will drift toward bias by default" and that "deeper disparities" in predictive algorithms must be addressed. For instance, the author discussed recruiters using algorithmic ad

generators and job boards to advertise to relevant potential applicants. These services are interested in attracting the most clicks for their clients' dollars but may be "delivered in a way that reinforces gender and racial stereotypes" (Hogan, 2019, p. 2). The author also noted that some personalized job boards automatically learn patterns in recruiters' preferences as they correspond with job seekers and dynamically adjust algorithms to solicit similar applicants. Thus, by directing ads to potential candidates matching the dynamically adjusted algorithm, dissimilar potential applicants are inadvertently excluded. In addition, applicant screening tools often model past hiring decisions which may, in turn, further support an unintended bias. Amazon reportedly canceled further development of an AI program intended to help human resources vet resumes. The machine learning tool was trained on observing patterns in resumes submitted to the company over a 10-year period that inherently reflected the dominance of males in the tech industry (Dastin, 2018).

To address the increasing use of AI in employment and the potential risks for bias and discrimination, the Equal Employment Opportunity Commission ("EEOC") published guidance in 2022 regarding artificial intelligence and employer obligations under the Americans with Disabilities Act ("ADA"). Then in May 2023, the EEOC issued similar guidance (the "Recent Guidance"), this time regarding employers' use of AI in their "selection procedures" (e.g., hiring, promotion, and termination) and the potential for disproportionate adverse effects (i.e., "disparate impact") on applicant groups who are protected under Title VII of the Civil Rights Act of 1964 ("Title VII"). The Recent Guidance explains that the "Uniform Guidelines on Employee Selection Procedures" from 1978 still apply — even though the technology has changed significantly — and will help employers understand how to use AI in hiring and avoid legal violations (Nelson & Reed, 2023).

The EEOC also confirms that "employers may be held responsible for algorithmic decision-making tools that create a disparate impact, even if the tools are designed or administered by another entity, such as a software vendor." Therefore, an employer using AI to make hiring decisions may be liable under Title VII if "the AI discriminates on a protected basis, such as gender or race, even if an outside vendor developed the AI." The Recent Guidance encourages employers who learn that an AI tool is creating a disparate impact to "take steps to reduce the impact or select a different

tool in order to avoid engaging in a practice that violates Title VII.”

Algorithmic Bias in Healthcare

Unintended bias may also have life threatening results (Feiner, et al., 2007; Jamalia, et al., 2022; Larrazabal, et al., 2020; Sparkes, 2022). Medical researchers have found that algorithmic bias may not only be inadvertently programmed into an AI system but may possibly be amplified as well (Larrazabal, et al., 2020; Zou & Schiebinger, 2018). Studies have found that some pulse oximeters have a tendency to overestimate oxygen levels for people having lower oxygen saturation levels and darkly pigmented skin (Feiner, et al., 2007; Jamalia, et al., 2022).

In a large-scale study conducted by Larrazabal, et al., (2020) on medical imaging datasets, the authors ran multiple different gender-imbalanced ratios of training data on the AI system and found that training deep learning-based CAD medical imaging systems on gender-imbalanced datasets had the potential to affect pathology results in minority groups (Larrazabal, et al., 2020). When the authors used diverse and balanced datasets to train the AI system, the system performed the best on test data (Larrazabal, et al., 2020).

In their 2019 paper “Artificial intelligence and algorithmic bias: implications for health systems,” researchers Panch, Mattie, and Atun first defined “algorithmic bias” in healthcare as the “application of an algorithm that compounds existing inequities in socioeconomic status, race, ethnic background, religion, gender, disability, or sexual orientation and amplifies inequities in health systems.” This definition suggests that some forms of bias have been active in medical information and decisions even before the application of AI systems. “Algorithmic bias is not just a technical issue” say Panch, Mattie, and Atun, “teams developing algorithms should be explicitly aware of the specificities of the health system context for which they are developing algorithms, by considering differential needs of different groups--best achieved through multi-disciplinary data science teams and by appropriate regulation and evaluation of algorithms and the data science process itself.” The authors suggest creating a “human-in-the-loop” system to counteract algorithmic bias; program outputs can thereby be vetted with the human as the ultimate decision maker (Panch, et al., 2019).

5. INCORPORATING PROFESSIONAL SKILLS AND CONTENT IN ALIGNMENT WITH COMPETENCY FRAMEWORKS

AI and its intrinsic effects have sprung into the educational limelight with lightning speed. Many educators are struggling to incorporate rapidly advancing current AI topics, such as algorithmic bias and copyright issues, into skill-heavy course syllabi. However, the overarching CC2020 Computing Curriculum guidelines and subsequent program directed models (e.g., IS2020, IT2020, etc.), provide a list of competencies in which these concepts can be addressed.

In Appendix A, the authors provide an example of a competency statement, based upon the CC2020 guidelines, that faculty can modify and implement to address these subsequent issues. The following section provides content and strategies supporting the *AI Bill of Rights* and the proposed competency statement to address the professional and ethical aspects of copyrights and algorithmic bias. Table 1 illustrates the alignment of the IS2020 competency realms (IS2020, 2020, p. 51) with the U.S. White House’s *Blueprint for an AI Bill of Rights*. Additional classroom suggestions and exercises, including the proposed competency statement, are provided in Appendices A, B, and C.

Tactics to Address the IS Competency Realms

Listed below are suggested skills and strategies that faculty can incorporate into their courses to address bias and copyrights in relation to the IS2020 competency realms.

IS Foundations Realm: This competency, representing IS as a whole, is usually addressed in an introductory course covering general IS concepts. Faculty can ensure that concepts such as AI, data management, data governance, data bias, ethics, copyrights, data privacy, data security, and legal implications are addressed. The authors provide a Moral Machine exercise in Appendix B that faculty can assign to generate awareness in students of their own biases. In addition, the White House’s *Blueprint* as well as other frameworks can be presented.

IS2020 Realm	White House Blueprint
IS Foundations	Overview
Data (Data & Info Mgt.)	Data Privacy & Protections Algorithmic Discrimination Protections
Technology (IT Infrastructure)	Safe & Effective Systems
Development (Design, Devpt. & Programming)	Algorithmic Discrimination Protections
Organizational Domain (Ethics, society, IS Mgt & Strategy)	Notice & Explanations Human Alternatives Consideration & Fall Back
Integration (Project Mgt & IS Practicum)	All

Table 1: White House Blueprint Mapped to IS2020 Competency Realms

Data Realm (Data/Information Management/Business Analytics): This competency may be covered in multiple courses and may include addressing data preprocessing techniques such as data cleaning, dimension reduction, variable selection, data sampling, data transformation, and balancing data (Azevedo, 2022). Discussions of the different types of data bias (e.g., propagated current state, inaccurate data focus, under-represented populations, sampling, selection, analytics, confirmation, etc.) can be included (Lawton, 2020). Techniques and tools available to reduce those biases can also be introduced such as: subpopulation analysis, comparing model results over time, IBM's open source AIF360 toolkit, IBM Watson OpenScale, Google's What-IF Tool (Dilmegani, 2022; IBM Developer Staff, 2018), Microsoft's Fairlearn, and TensorFlow's open-source standardized data sets and data tools. Faculty can incorporate discussion about bias and the importance of developing good algorithms and training data.

Technology Realm (IT Infrastructure): Discussions about the current use and future direction of AI could be incorporated into multiple courses to address the optional competency aspect of "Emerging Technologies". In addition, discussion about copyrights, copyrighted material and AI training data, and the outputs generated by AI algorithms can be addressed as well.

Awareness of Copyrights: Along with a discussion on the content on copyrights provided in Section 3, faculty can incorporate an exercise on code theft and copyright infringement provided in Appendix C.

Development Realm: In the systems analysis and design and application development courses, faculty can incorporate a mindset of quality assurance by eliminating bias and copyright infringement at the start of the development process. As students being to develop the requirements for their systems, they could be asked to list the types of issues that would occur, where and how those issues would arise, delineate the steps they would take to eliminate the identified issues, and describe how they would audit the system for those issues. Quality assurance procedures and methodologies could be addressed when having students evaluate the results generated by AI algorithms may provide an opportunity to incorporate quality control procedures and methodologies into the classroom. The algorithm must first be tested to determine whether or not the ideal results are returned for the real problem of interest or if the results address a similar, but different, problem (Bembeck, et al., 2021). Once the ideal result is returned, then a "blind taste test" can be run on the algorithm to test for bias (Uzzi, 2020). With the blind test, the AI algorithm is first trained on the entire data set and then trained again on the same data set with data suspected of introducing bias removed. If the results are different, then the suspected variable should be evaluated to determine if the variable provides a valid explanation for the model's performance or introduces bias into the decision-making (Uzzi, 2020).

Organizational Domain Realm (Ethics, society, IS Management and Strategy) Governance Policies: *The Economics and Regulation of Artificial Intelligence and Emerging Technologies* published by the Brookings Institute, recommended that, to reduce inherent bias built into AI, policymakers should define bias in respect to its real-world results, use the definition to assist the industry in investigating biased algorithms, and insist that organizations develop internal accountability structures to prevent bias before it happens. This methodology could be incorporated into faculty discussions on IT governance policies and procedures, accountability structures, documentation, and the

establishment of structures for preventing bias (Bembeck, et al., 2021).

Need for Compliance with Industry Regulations: Future developers may develop software in a variety of industries. It is impossible to address the specific regulations and requirements of each industry. However, by reviewing some of the requirements outlined by a regulatory body affecting almost every industry, students are made aware that additional factors outside the software's specific functional requirements may need to be assessed and compliance issues addressed.

AI has been used in workplace employment decisions to assist with such activities as advertising job openings, screening applicant resumes, determining salary offers, establishing terms and conditions of employment, monitoring worker performance, and making promotion decisions (Bogen, 2019; EEOC, 2022; Nelson & Reed, 2023). U.S. employers "can be held liable for using procedures that overly favor a certain group of applicants" (Hogan, 2019, p. 2). In response to the growing use of AI in human resource decision-making, the EEOC published a technical assistance document to provide guidance (2022). EEOC guidelines indicate that decision-making tools could be considered unlawful for:

- **Not providing reasonable accommodations** for fair and equitable treatment of all employees.
- **Screening out individuals**, either intentionally or unintentionally, for not meeting a particular standard resulting from algorithmic decisions made resulting from a disability. For example, the disability may affect the accuracy of the algorithm's assessment of the individual or prevent or hinder the individual from participating in the screening process (EEOC, 2022)
- **Violating ADA restrictions on disability-related inquiries.** For instance, if an applicant for a position has a significant gap in their employment history due to their disability, the algorithm may screen the employee out of the process if the algorithm is not programmed to handle such employment abnormalities (EEOC, 2022).
- **EEOC compliance by vendor-purchased software** must be ensured to prevent companies using the software from being held liable for violations.

The EEOC (2022) offered several recommendations to software developers to reduce the chances of making biased decisions resulting from the use of an algorithmic tool. Instructors can fashion discussion around governance policy development associated with the following topics:

- **Inquiring of vendors** about whether or not the AI algorithms were developed with individuals with disabilities in mind.
- **Involving experts** on various types of disabilities in the algorithm and software development process.
- **Addressing as many different kinds** of disabilities as possible to minimize bias.
- **Ensuring that user interfaces are accessible and/or alternative formats** are made available.
- **Testing the system** to ensure that individuals with disabilities are not disadvantaged by the algorithm.
- **Providing clear instructions for requesting accommodations** to disabled individuals using the system.
- **Providing an explanation of the algorithm in use** and the data assessed so that accommodations can be requested if necessary.
- **Limiting utilization of algorithms** in decision making. Utilizing an algorithmic decision-making tool only for making decisions that will not be affected by a disability.

The U.S. White House's *Blueprint for an AI Bill of Rights*

The White House Office of Science and Technology Policy issued the *Blueprint for an AI Bill of Rights* consisting of "five principles that should guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence" (OSTP, 2023, p. 1). Table 1 illustrates the alignment of these statements with the IS2020 required competency realms. Faculty can use the five principles, summarized below, as talking points for classroom discussion with the sections listed below or independently:

1. **Safe and Effective Systems** designed and developed to address potential risks, prevent algorithmic bias, and proactively protect individuals from unintended, yet foreseeable use, inappropriate use of data, and compounded harm of the algorithm's reuse. Systems should be developed with consultation from diverse communities,

evaluated and audited by independent parties, and the steps taken to mitigate potential harm should be reported.

2. **Algorithmic Discrimination Protections** against unfair treatment arising from “race, color, ethnicity, sex, religion, age, national origin, disability, veteran status, genetic information or any other classification by law” (OSTP, 2022, p. 1). Proactive and continuous steps should be taken to guard against discrimination including equity assessments, representative datasets, ensuring accessibility, and organizational oversight.
3. **Data Privacy and Protections** should be built into systems by default and data collected that is strictly necessary for use in the context intended. Enhanced protective measures should be enacted for sensitive domains including data collected about health, work, education, criminal justice, finance, and youths.
4. **Notice and Explanations** should be clearly articulated regarding the overall automated system function, the parties responsible for the automated system, explanations of the outcomes, and how and why an outcome that impacts an individual was determined.
5. **Human Alternatives, Consideration and Fallback** should be provided as appropriate. Human alternatives to the automated system should be provided to ensure accessibility. Individuals should have access to “timely human consideration and remedy by a fallback or escalation process” (p. 1). People who interact with AI systems should receive appropriate training for ensuring consistent fair treatment and addressing equity issues. Governance processes developed to carry out these processes should be developed and made publicly available.

6. CONCLUSION

Artificial Intelligence is poised to change the way in which organizations operate, job roles are carried out, programmers develop software, and faculty teach classes. While the field develops at a rapid pace, faculty should prepare MIS and computer science students for an AI-enhanced world with multiple risks and unknowns. In this paper, the authors address a few “danger zones” of the AI-enhanced world by reviewing relevant literature on software copyrights and programming bias and suggest course content strategies for addressing these issues using the frameworks provided by the CC2020 Task Force,

the White House’s *Blueprint for an AI Bill of Rights*, and the EEOC’s recommendations for addressing possible bias introduced by AI supported technologies.

7. REFERENCES

- Azevedo, A. (2022). Data Preprocessing: 6 Techniques to clean data. *Scalable Path*. Retrieved August 19, 2023, from <https://www.scalablepath.com/data-science/data-preprocessing-phase>
- Bembeck, E., Nissan, R., & Obermeyer, R. (2021). To stop algorithmic bias, we first have to define it. *Brookings*. Retrieved on June 13, 2023, from <https://www.brookings.edu/research/to-stop-algorithmic-bias-we-first-have-to-define-it/>
- Bogen, M. (2019). All the ways hiring algorithms can introduce bias. *Harvard Business Review*. Retrieved on May 29, 2023, from <https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias#:~:text=Many%20hope%20that%20algorithms%20will%20help%20human%20decision-makers,points%20like%20university%20attendance%20or%20performance%20evaluation%20scores>
- Bolukbasi, T. Chang, K.W., Jou, J., Saligrama, V., & Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Advances in Neural Information Processing Systems*, (29). Retrieved on May 29, 2023, from <https://arxiv.org/pdf/1607.06520v1.pdf>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186. <https://doi.org/10.1126/science.aal4230>
- CC2020 Task Force (2020). Computing Curricula 2020 (CC2020). Retrieved on May 29, 2023, from <https://www.acm.org/binaries/content/assets/education/curricula-recommendations/cc2020.pdf>. DOI: 10.1145/3467967
- Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. Retrieved June 1, 2023, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- Dilmegani, C. (2022). Bias in AI: What it is, Types, Examples & 6 Ways to Fix it in 2023.

- AI Multiple*. Retrieved August 19, 2023, from <https://research.aimultiple.com/ai-bias/>
- Farris, J. (1992) U.S. Court of Appeals, Ninth Circuit. *Lewis Galoob Toys, Inc. v. Nintendo of America*. (964 F 2d 965). Retrieved on May 29, 2023, from <https://h2o.law.harvard.edu/cases/4459>
- Feiner, J. R., Severinghaus, J. W., & Bickler, P. E. (2007). Dark skin decreases the accuracy of pulse oximeters at low oxygen saturation: the effects of oximeter probe type and gender. *Anesthesia and analgesia*, 105(6 Suppl), S18–S23.
<https://doi.org/10.1213/01.ane.0000285988.35174.d9>
- U.S. Equal Employment Opportunity Commission (EEOC). (2022). The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees. ADA, 29 CFR Part 1630 & app. Retrieve on June 8, 2023, from <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>
- Glover, E. (2023). AI-generated content and copyright law: What we know. *Builtin*. Retrieved June 14, 2023, from <https://builtin.com/artificial-intelligence/ai-copyright>
- Hogan, M. (2019). All the ways hiring algorithms can introduce bias. *Harvard Business Review*. Retrieved June 2, 2023, from <https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias>
- IBM Developer Staff (2018). AI Fairness 360. Open Source@IBM. . Retrieved on August 19, 2023, from <https://www.ibm.com/opensource/open/projects/ai-fairness-360/>
- IS2020 Task Force (2020). IS2020 A Competency Model for Undergraduate Programs in Information Systems. Retrieved on August 2, 2023, from <https://www.acm.org/binaries/content/assets/education/curricula-recommendations/is2020.pdf>
DOI: 10.1145/3460863
- Jamalia, H, Castillo, L. T., Morgan, C. C., Coult, J., Muhammad, J. L., Osobamiro, O. O., Parson, E. C., & Adamson, R. (2022). Racial disparity in oxygen saturation measurements by pulse oximetry: Evidence and implications. *Annals of the American Thoracic Society*. 19 (12). 1951-1964.
<https://10.1513/AnnalsATS.202203-270CME>
- Larrazabal, A. J., Nieto, N., & Peterson, V. (2020) Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *PNAS*, 117(23) 12592-12594.
<https://doi.org/10.1073/pnas.1919012117>
- Lawton, G. (2020). 8 Types of bias in data analysis and how to avoid them. *TechTarget: Business Analytics*. Retrieved on August 19, 2023, from <https://www.techtarget.com/searchbusinessanalytics/feature/8-types-of-bias-in-data-analysis-and-how-to-avoid-them>
- Nelson, S. M. & Reed, M. (2023). EEOC issues guidance on use of AI in employment decisions. *The National Law Review*, 13(159), Retrieved on June 8, 2023, from <https://www.natlawreview.com/article/eeoc-issues-guidance-use-ai-employment-decisions>
- Ormond, J. (2021). ACM and IEEE-CS release Computing Curricula 2020, Global Guidelines for Baccalaureate Degrees in Computing. *Association for Computing Machinery News Release*. Retrieved on June 8, 2023, from <https://www.acm.org/media-center/2021/march/computing-curricula-2020>
- Office of Science and Technology Policy. (2022). *Blueprint for an AI Bill of Rights*. Retrieved on June 14, 2023, from <https://www.whitehouse.gov/ostp/ai-bill-of-rights/#safe>
- Panch T, Mattie H, & Atun R. (2019). Artificial intelligence and algorithmic bias: implications for health systems. *Journal of Global Health*. 9(2):010318. doi: 10.7189/jogh.09.020318. PMID: 31788229; PMCID: PMC6875681.
- Parikh, N. (2021). Understanding bias in AI-enabled hiring. *Forbes*. Retrieved on June 1, 2023, from <https://www.forbes.com/sites/forbeshumanresourcescouncil/2021/10/14/understanding-bias-in-ai-enabled-hiring/?sh=737a60c17b96>
- Reinhardt, S. (1992). U.S. Court of Appeals, Ninth Circuit. *SEGA Enterprises Ltd. v. Accolade, Inc.* (977 F 2d 1510). Retrieved on May 29, 2023, from <https://h2o.law.harvard.edu/cases/4486>
- Rouse, M. (2013). Software patent. *Techopedia*. Retrieved on May 29, 2023, from <https://www.techopedia.com/definition/22199/software-patent>

- Sharp, A. (2023). EU lawmakers pass landmark AI regulation Bill. *FP*. Retrieved on June 14, 2023, from <https://foreignpolicy.com/2023/06/14/eu-ai-act-european-union-chatgpt-regulations-transparency-privacy/>
- Sparkes, M. (2022). How can we prevent AI from being racist, sexist, and offensive? *New Scientist*, 255(3302). p. 14.
- Supreme Court of the U.S. (2021). *Google LLC v. Oracle America, Inc.* Retrieved on May 29, 2023, from https://www.supremecourt.gov/opinions/20pdf/18-956_d18f.pdf
- Upcounsel. (2020). What is a software patent? *Upcounsel*. Retrieved on May 29, 2023, from <https://www.upcounsel.com/software-patent>
- Uzzi, B. (2020). A simple tactic that could help reduce bias in AI. *Harvard Business Review*. Retrieved on May 29, 2023, from <https://hbr.org/2020/11/a-simple-tactic-that-could-help-reduce-bias-in-ai>
- U.S. Copyright Office. (2023). *Artificial Intelligence Initiative; Issue No. 1004* - March 16, 2023. Retrieved on June 14, 2023, from <https://copyright.gov/newsnet/2023/1004.html>
- Vincent, J. (2022). The lawsuit that could rewrite the rules of AI copyright. *The Verge*. Retrieved on June 14, 2023, from <https://www.theverge.com/2022/11/8/23446821/microsoft-openai-github-copilot-class-action-lawsuit-ai-copyright-violation-training-data>
- Ye, J. (2023). Generative AI services. *Reuters*. Retrieved on June 14, 2023, from <https://www.reuters.com/technology/china-releases-draft-measures-managing-generative-artificial-intelligence-2023-04-11/>
- Zou J, Schiebinger L. (2018). AI can be sexist and racist - it's time to make it fair. *Nature*. 559(7714):324-326. doi: 10.1038/d41586-018-05707-8. PMID: 30018439.

Editor's Note:

This paper was selected for inclusion in the journal as a 2023 ISCAP Conference Distinguished Paper. The acceptance rate is typically 7% for this category of paper based on blind reviews from six or more peers including three or more former best papers authors who did not submit a paper in 2023.

APPENDIX A

Sample Competency Statement

In 2020, the Association for Computing Machinery (ACM) and the IEEE Computer Society (IEEE-CS), with input from several other groups including the Education Special Interest Group of Information Systems and Computing Academic Professionals (EDSIG/ISCAP), published the Computing Curriculum 2020 (CC2020). The CC2020 report summarizes and brings together “the current state of curricular guidelines for academic programs that grant baccalaureate-level degrees in computing, as well as propose a vision for future curricular guidelines” (CC2020, 2020 p. 12).

The central theme of the report focuses on the inclusion of competencies in computing education. According to CC2020, “A competency is a collection of specific components of knowledge, skills, and dispositions. The knowledge dimension of competency encompasses concepts that are technical (computing concepts), foundational and professional (indicative of a workplace), and domain specific (the task setting).” CC2020 believes that “competency statements” — task descriptions that align with the relevant knowledge elements, skill level and disposition — are a key method of expressing a model of knowledge aligned with skills and professionalism (CC2020, 2020, p.48).

The CC2020 provided several component tables to help guide instructors in writing competency statements. In Table 4.1 of the report (shown below), the elements of computing knowledge are featured. (CC2020, 2020, p.49).

Table 4.1 illustrates thirty-four abbreviated knowledge areas partitioned into an ordered sequence of six categories. While the table is incomplete, it does provide an example of high-level vocabulary for computing knowledge rooted in the collective wisdom of different computing communities. This summary of computing knowledge areas represents a well-understood and consistent vocabulary from which computing competency statements can evolve.

Table 4.1. Elements of Computing Knowledge

Users and Organizations	Systems Modeling	Systems Architecture and Infrastructure	Software Development	Software Fundamentals	Hardware
Social Issues and Professional Practice Security Policy and Management IS Management and Leadership Enterprise Architecture Project Management User Experience Design	Security Issues and Principles Systems Analysis & Design Requirements Analysis and Specifications Data and Information Management	Virtual Systems and Services Intelligent Systems (AI) Internet of Things Parallel and Distributed Computing Computer Networks Embedded Systems Integrated Systems Technology Platform Technologies Security Technology and Implementation	Software Quality, Verification and Validation Software Process Software Modeling and Analysis Software Design Platform-Based Development	Graphics and Visualization Operating Systems Data Structures, Algorithms and Complexity Programming Languages Programming Fundamentals Computing Systems Fundamentals	Architecture and Organization Digital Design Circuits and Electronics Signal Processing

In Table 4.2 of the report, the elements of foundational and professional knowledge are recorded. “The thirteen elements of foundational and professional knowledge listed in Table 4.2 represent a subset of the professional listings derived from the IT2017 report and subsequently from Appendix D in this report. Computing professionals are commonly expected to demonstrate high levels of skill in applying this knowledge which deserves explicit attention in baccalaureate programs” (CC2020, 2020 P49).

Table 4.2. Elements of Foundational and Professional Knowledge

Knowledge Elements	Meaning
Analytical and Critical Thinking	A mental process of simplifying complex information into basic parts and evaluating results to make proper decisions
Collaboration and Teamwork	Apportion challenging tasks into simpler ones and then work together to complete them efficiently
Ethical and Intercultural Perspectives	Ethical perspectives of the different viewpoints someone uses to view a problem in the context of individual human values
Mathematics and Statistics	Use of numbers and theories abstractly especially in the collection and analysis of numerical data
Multi-Task Prioritization and Management	Processing several issues or tasks at once while arranging them according to importance to do specific one first
Oral Communication and Presentation	Conveying a message orally using real-time presentations with visual aids related audience interests and goals
Problem Solving and Trouble Shooting	A logical and orderly search for the source of a unit problem and making the unit operational again
Project and Task Organization and Planning	A process to provide decisions about a project concerning organization and planning to achieve a successful result
Quality Assurance / Control	Use of techniques, methods, and processes to identify and prevent defects according to defined quality standards
Relationship Management	A strategy to maintain an ongoing level of engagement usually between a business and its customers or other businesses
Research and Self-Starter/Learner	Someone who begins or undertakes work or a project without needing direction or encouragement to do so
Time Management	An ability to use a person's time in an effective or productive manner to work efficiently
Written Communication	Use of a written form of interaction between people and organizations that provides an effective way of messaging

"As CC2020 defines skill — the proficient applying of knowledge — Table 4.3 summarizes an ordered sequence of six cumulative levels of skill (cognitive skill) together with abbreviated definitions. These levels correlate with Bloom's taxonomy that permits the adoption of a commonly agreed vocabulary as described in the 2001 revisions to Bloom's taxonomy of educational objectives. The table lists the cognitive skills as verbs" (CC2020, 2020, p. 50).

Table 4.3. Levels of Cognitive Skills Based on Bloom's Taxonomy

Remembering	Understanding	Applying	Analyzing	Evaluating	Creating
Exhibit memory of previously learned materials by recalling facts, terms, basic concepts, and answers.	Demonstrate understanding of facts and ideas by organizing, comparing, translating, interpreting, and giving descriptions.	Solve problems in new situations by applying acquired knowledge, facts, techniques, and rules in a different way.	Examine and break information into parts by identifying motives or causes; make inferences and find evidence to support solutions.	Present and defend opinions by making judgments about information, validity of ideas, or quality of material.	Compile information together in a different way by combining elements in a new pattern or by proposing alternative solutions.

"Dispositions define the third dimension of competency. Table 4.4 displays eleven prospective dispositions derived from the literature. Disposition, as an intrinsic component of competency, represents the opportunity to express institutional and programmatic values expected in the workplace. Dispositional expectations enrich the description/assessment of competency and/or the related pedagogy. Ascribing a disposition to a competency indicates a clear commitment to self-reflection and examination that distinctly distinguishes a competency from a learning outcome" (CC2020, 2020, p. 50).

Table 4.4. Prospective Elements of Dispositions

Element	Elaboration	Element	Elaboration
Adaptable	Flexible; agile, adjust in response to change	Professional:	Professionalism, discretion, ethical, astute
Collaborative:	Team player, willing to work with others	Purpose-driven:	Goal driven, achieve goals, business acumen
Inventive:	Exploratory. Look beyond simple solutions	Responsible:	Use judgment, discretion, act appropriately
Meticulous:	Attentive to detail; thoroughness, accurate	Responsive:	Respectful; react quickly and positively
Passionate:	Conviction, strong commitment, compelling	Self-directed:	Self-motivated, determination, independent
Proactive:	With initiative, self-starter, independent		

The authors have taken the focus of this article — professional, legal and ethical issues in machine learning — and created a competency statement which may be used to orientate the task of identifying the salient issues and creating awareness of potential resolutions.

Competency Title: Risk Management of Legal and Ethical Issues in Machine Learning	
Competency Statement Analyze machine learning scenarios and identify legal and ethical issues that would place business organizations at risk; propose solutions/resolutions/safeguards.	
Knowledge Element [Table #4.1 & 4.2](CC2020, 49; 50) (CC2020, 51)	Skill Level [Table 4.3]
Social Issues and Professional Practice	Analyzing
Data and Information Management	Applying
Intelligent Systems (AI)	Analyzing
Software Design	Applying
Analytical and Critical Thinking	Analyzing
Ethical and Intercultural Perspectives	Applying
Written [and/or Oral] Communication	Creating
Disposition(s) [Table 4.4] (CC2020, 51)	
Adaptable	Inventive Professional Responsible Responsive

Knowledge Elements:

(1) Social Issues and Professional Practice: It is important for computer scientists to understand the relevant social, ethical, and professional issues that surround their activities. According to the ACM Code of Ethics and Professional Conduct, a computing professional should “contribute to society and to human well-being, acknowledging that all people are stakeholders in computing” and “avoid harm” (ACM Code, 1.1 and 1.2). The Code also requires the computer scientist to “be fair and take action not to discriminate” (ACM Code 1.4). By analyzing the issues in the two exercises included in Appendices B (The Moral Machine) and C (Not OK: AI Copyright Infringement), students will gain awareness of the ethical issues of programming bias and intellectual property theft, and work through potential solutions to these issues.

(2) Data and Information Management: Computer science and MIS students may be responsible for managing and mining significant amounts of data, and the way in which they use that data to create products and algorithms is becoming increasingly important. Students as computing professionals

engaged in system development will have a personal responsibility for programs and data use that may adversely affect the public. Again, by analyzing the issues in the two exercises included in Appendices B (The Moral Machine) and C (Not OK: AI Copyright Infringement), students will gain awareness of how proper data management is critical to protecting individual rights and preventing bias and discrimination. Once aware of the potential issues, students can apply this knowledge in designing safeguards which may be utilized in the design of data management and use.

(3) Intelligent Systems (AI): Another professional responsibility of computing professionals is embodied in the ACM Code 2.5: "Give comprehensive and thorough evaluations of computer systems and their impacts, including analysis of possible risks." This Code note especially highlights that "extraordinary care should be taken to identify and mitigate potential risks in machine learning systems." Therefore, this knowledge component in the competency focuses on awareness of the potential risks that AI enhanced programs pose through their operation, and the computer professional's responsibility to establish preventative measures to protect users. The content of this article, and the two exercises included here — one on programming bias and one on copyright infringement — detail the dangers that may be associated with AI enhanced programs, and ask students to critically think through potential remedies and safeguards.

(4) Software Design: Students should be able to identify issues in software and program design that may be problematic for users and their respective organizations. Programming bias in AI algorithms, and copyright infringement by AI enhanced programs, are significant issues for organizations that either create products or use AI to generate content. Using the two exercises in Appendices B and C, as well as the White House's Blueprint for an AI Bill of Rights, students should be able to identify risks in software design and apply this awareness to creating safeguards that will mitigate these risks.

(5) Analytical and Critical Thinking: Using the exercises on programming bias and copyright infringement, students will analyze the potential legal and ethical risks inherent in AI enhanced programs use and propose potential solutions and safeguards. By creating awareness and understanding of these complex issues, students will be able to evaluate risks and apply this knowledge to make proper decisions. Knowledge of intellectual property law and employment law, as well as current regulatory guidelines as included in this article's content, should aid students in identifying whether their organization's practices are aligned with compliance — and if not, what steps could be taken to refocus a company's operations to ensure that ethical and legal tenants are adhered to.

(6) Ethical and Intercultural Perspectives: Once again using the exercises included in Appendices B and C, students may apply their knowledge of law and ethical issues regarding AI programming and copyright issues in ensuring that policies and decisions include safeguards against discrimination and IP violations. This will be especially important as the AI industry sees more and more applications in narrow artificial intelligence, and critical as we may someday realize the development of strong AI and general artificial intelligence.

(7) Written Communication: This Competency Statement requires the student to identify legal and ethical issues within the topics of programming bias and copyright infringement, and propose solutions/resolutions/safeguards. In the Moral Machine Exercise (Appendix B), students are asked to share their opinions of the reasons for cultural bias found by the MIT team by posting to collaboration boards or otherwise submitting written and/or oral communication on the focused issues. In the Not OK: AI Copyright Infringement Exercise (Appendix C), students are also asked to analyze and present, either in writing or verbally, the ethical issues in the "taking" of computer science professor Tim Davis's copyrighted code. Students are also asked to propose solutions and/or safeguards to these complex issues.

Appendix B

Moral Machine Exercise

As originally envisioned, AI was thought of to eliminate bias and promote Diversity, Equity & Inclusion. Increasing use of weak AI has certainly led to positives and efficiencies in modern society, yet it has also uncovered a darker side — including use characterized by perpetuating stereotypes, incorporating bias, and failing to eliminate discriminatory practices.

The focus of this exercise is to create awareness in computing students with regard to the insidious way bias can creep into their programming, coding, and product designs. Such discrimination is not only harmful to users, but it can generate significant professional and legal liability for program developers and businesses.

In the Moral Machine Experiment, MIT researchers sought to test how people around the world would decide significant moral dilemmas. The project gathered data on millions of humans' moral decisions, and this data was used to train machine-learning algorithms. The intent — built around autonomous vehicle decisions — was to determine what may potentially influence human machine learning programming. The results surprised the social psychologists involved in the experiment; significant cultural biases were unearthed that could be regionalized, and conversely certain trends were globally apparent. There were findings that aligned with developed vs. developing countries; the extent of economic inequality in a region; and individualist and collectivist cultures. The data revealed that programmers should be aware of how bias can influence outcomes when they are creating products or services — especially when human lives are on the line.

The authors have created a lesson around the Moral Machine Experiment. Although the data gathering phase of the project ended in 2020, the exercise is still available for use by individuals, and may be found at the following link: <https://www.moralmachine.net>. One of the authors uses the Moral Machine lesson in an ethics course, so we have created a Nearpod Lesson that may be accessed in self-paced mode by students anywhere in the world. Nearpod is a hybrid learning tool that combines multimedia learning with digital assessments; the program is highly interactive and may be used in self-paced or a live-action mode. Students can access the self-paced lesson by going to www.nearpod.com, and enter the access code under "Join a lesson" when prompted. They can re-enter the lesson anytime using the same code. Instructors who would like an editable copy of the lesson can sign up for a free Nearpod account, and then access the editable link below and add the lesson to their library. Once in your library, an instructor can change the lesson, set up their own self-paced version, or add a "live" session to their classroom. There is also a significant reporting feature in Nearpod. The slides from the exercise are additionally included (following the links) here if the instructor would prefer to use the exercise resources outside the Nearpod app.

Nearpod Editable Link for Educators:

https://np1.nearpod.com/sharePresentation.php?code=b59aba4c7b642572f5b4bf9f05446e00-1&oc=user-created&utm_source=link



Lesson: Moral Machine Self-Paced Nearpod Lesson

1/25



The MIT Moral Machine Experiment

- The Moral Machine is an online platform developed by Iyad Rahwan's Scalable Cooperation group at the Massachusetts Institute of Technology (MIT).
- The project tested human responses in over 200 countries to moral dilemmas and collected information on the decisions that these people made between two destructive outcomes.
- The project has gathered data on millions of humans' moral decisions, and this data was used to train machine-learning algorithms.
- As artificial intelligence plays an increasingly significant role in autonomous driving technology, research projects like Moral Machine are intended to help find solutions for challenging life-and-death decisions that will face self-driving vehicles.

Lesson: Moral Machine Self-Paced Nearpod Lesson

2/25



The MIT Moral Machine Experiment

On the next slide, please view the introduction on how to participate in the Moral Machine experiment.


The directions may also be viewed at the following link: <https://www.moralmachine.net/> or on YouTube: <https://www.youtube.com/watch?v=XCO8ET66xE4>.

Following the video, the instructions for participation are posted on the next slide. These directions may also be accessed by pressing the "View Instructions" button at the website.

View Instructions

Lesson: Moral Machine Self-Paced Nearpod Lesson

3/25



Instructions

This website has three main functional interfaces that can be accessed from the menu bar:

Judge

You will be presented with random moral dilemmas that a machine is facing. For example, a self-driving car, which does not need to have passengers in it. The car can sense the presence and approximate identification of pedestrians on the road ahead of it, as well as of any passengers who may be in the car.

The car also detects that the brakes failed, leaving it with two options: keep going and hit the pedestrians ahead of it, or swerve and hit the pedestrians on the other lane. Some scenarios will include the case of a non-empty car; in those cases, one of the two lanes have a barrier that can be crashed into, affecting all passengers. One or two pedestrian signals may also be included in a given scenario, changing the legality of a pedestrian's position on their respective lane.

You are outside the scene, watching it from above. Nothing will happen to you. You have control over choosing what the car should do. You can express your choice by clicking on one of the two choices in front of you. In each of the two possible outcomes, the affected characters will be visually marked with the symbol of a skull, a medical cross, or a question mark to signal what will happen to this character, corresponding to death, injury, or an uncertain outcome, respectively.

You may proceed from scenario to scenario by selecting the outcome you feel is most acceptable by you. This can be done by clicking the outcome of your choice, which will be highlighted when you hover your cursor over it. A button below each outcome depiction will let you toggle the display of a textual summary of the outcome that you can read. A counter at the top right will let you know your progress in the sequence of scenarios.

Upon finishing all the scenarios, you will see a summary of the aggregated trend in your responses in the session of the game you just played, compared to the aggregated responses of other players, along several different dimensions. You may "Share" or "Link" your results using the corresponding buttons, and/or play another scenario sequence by clicking "More".

Design

Additionally, you may try to create a new scenario yourself. You will first be asked to choose whether you want to have the dilemma be between two sets of pedestrians, or, if between pedestrians and passengers, whether the self-driving car will have to swerve to save the passengers or pedestrians. You can then choose to add legal complications in the form of a pedestrian signal.

Finally, you can choose characters to add to each possible location in the scenario. The default fate for the characters who crash or are hit is death, but you can change this using the dropdown for each location. Note that the fates of impacted characters can be set independently for each character, even within the same location.

You can reset the scenario creation interface at any time by clicking the "Start Over" button on the left. Once you are done creating a scenario, and have given it a creative title, you can submit it by clicking the "Submit" button on the right. Once you do, your scenario will be added to the database of scenarios created by users of the platform.


Browse

This interface lets you view the scenarios you and other users of this platform have created. The scenarios are arranged chronologically; you can click "C" or "P" to move ahead or back in this arrangement. Alternatively, you can click the "Random" button to take you to any random scenario in the timeline. As with the Judge interface, you are able to toggle the display of textual descriptions for each outcome using the button below the respective outcome.

You can show your appreciation for particularly interesting scenarios by clicking the "Like" button, and can "Share" or "Link" such scenarios using the corresponding buttons. A discussion thread for each scenario is displayed below the depiction of the scenario; we encourage you to participate.

Lesson: Moral Machine Self-Paced Nearpod Lesson

4/25



The MIT Moral Machine Experiment

You can browse the various scenarios, like the one seen here, by pressing the "Browse Scenarios" button.

When you are ready to participate in the experiment, press the "Start Judging" Button.

Ride or die

Share Like 0 Likes 3C Random

In this case, the self-driving car with sudden brake failure will continue ahead and drive through a pedestrian crossing ahead. This will result in ...

Dead:

- 1 male executive
- 1 elderly woman
- 1 criminal

Note that the affected pedestrians are abiding by the law by crossing on the green signal.

Hide Description Show Description

Lesson: Moral Machine Self-Paced Nearpod Lesson 5/25 

The MIT Moral Machine Experiment

After judging the scenarios, be sure to save your results. You can share your results and opinions through the polls and collaboration boards here in this self-paced Nearpod Lesson, or in class/online through another forum.

Results

Most Saved Character

Most Killed Character

Saving More Lives

Between Not Matter Others Matters a Lot


Lesson: Moral Machine Self-Paced Nearpod Lesson 6/25 

Moral Machine Directions at MIT



Lesson: Moral Machine Self-Paced Nearpod Lesson

7/25



Ethical Issues?

What are some ethical issues w/ the programming of self-driving cars?

Instructions

Teacher

In addition to text, you can add images, video, or audio as a part of your response.

0 comments

Write a comment




Collaborate Board

Ethical Issues?

Lesson: Moral Machine Self-Paced Nearpod Lesson

8/25






Lesson: Moral Machine Self-Paced Nearpod Lesson 9/25 



In your moral machine survey, how did you rate: Saving More Lives?


- ☐ Does not matter
- ☐ Matters a lot
- ☐ In the middle

Lesson: Moral Machine Self-Paced Nearpod Lesson 10/25 

Poll

Lesson: Moral Machine Self-Paced Nearpod Lesson

11/25




In your Moral Machine survey what was your preference in gender?

- ☐ Males
- ☐ Females
- ☐ No preference

Lesson: Moral Machine Self-Paced Nearpod Lesson

12/25





Lesson: Moral Machine Self-Paced Nearpod Lesson

13/25



In the Moral Machine Survey, what preference did you have for age?

- ☐ Younger
- ☐ Older
- ☐ No preference

Lesson: Moral Machine Self-Paced Nearpod Lesson

14/25



Poll

Lesson: Moral Machine Self-Paced Nearpod Lesson

15/25



In the Moral Machine Survey, what preference did you have for fitness?

- ☐ Fit people
- ☐ Unfit people
- ☐ No preference

Lesson: Moral Machine Self-Paced Nearpod Lesson

16/25



Poll

Lesson: Moral Machine Self-Paced Nearpod Lesson

17/25



In the Moral Machine Survey, what was your social value preference?

- ☐ Higher social status
- ☐ Lower social status
- ☐ No preference

Lesson: Moral Machine Self-Paced Nearpod Lesson

18/25




Cultural Bias in the Moral Machine Experiment

Please watch the video about cultural bias on the next slide. This video may also be accessed on YouTube: <https://www.youtube.com/watch?v=jPo6bby-Fcg>


Then respond to the questions in the collaboration boards in this Nearpod Lesson, or during a discussion in class or online.

Lesson: Moral Machine Self-Paced Nearpod Lesson

19/25




The Moral Machine How Culture Affects Decisions



Lesson: Moral Machine Self-Paced Nearpod Lesson

20/25





Why do you believe Far East countries select to save elders?

Instructions



Collaborate Board

Why do you believe Far East countries select to save elders?

Lesson: Moral Machine Self-Paced Nearpod Lesson

21/25



Why do French-affiliated countries prefer to save females?

Instructions




Collaborate Board

Why do French-affiliated countries prefer to save females?

Lesson: Moral Machine Self-Paced Nearpod Lesson


22/25





Why would affluent countries prefer to save higher status?

Instructions




Collaborate Board

Why would affluent countries prefer to save higher status?

Lesson: Moral Machine Self-Paced Nearpod Lesson

23/25



Programming Bias in the Moral Machine Experiment


Think about the following statement: "People like to assume that computers and machines are neutral in their results. In reality, machine-learning software that uses datasets to "train" software often amplifies existing social biases. Companies are currently relying heavily on software that learns by sorting piles of data. This has led to computers taking on unsavory biases from both the programmer and society in general."

Pun, I., Evrim, B. and Cheng, Y. (2019). «Are we programming computers to be biased?» Ethics in the Built Environment, <https://sites.psu.edu/visionaryarchitecture/2019/06/17/are-we-programming-computers-to-be-biased/>.

How can programmers protect against bias in their algorithm, coding and product design?

Lesson: Moral Machine Self-Paced Nearpod Lesson

24/25





Programming Bias?

How can programmers protect against bias in their algorithm, coding and product design?

Instructions



Collaborate Board

Programming Bias?

Lesson: Moral Machine Self-Paced Nearpod Lesson

25/25

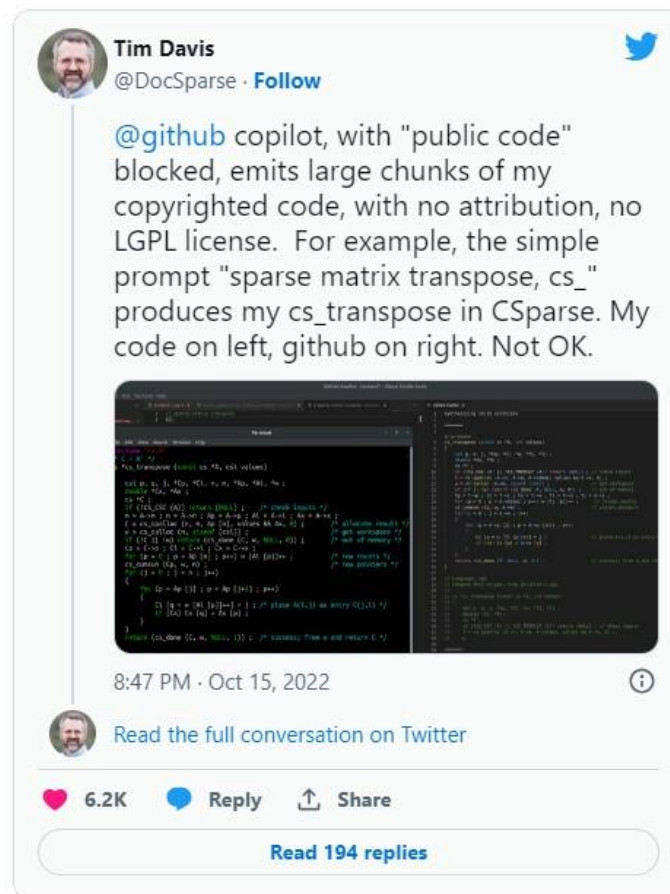


Appendix C

Not OK: AI Copyright Infringement Exercise

In 2021, Microsoft and its computer code-sharing website GitHub, as well as artificial intelligence firm OpenAI, were sued in California in a class-action lawsuit. The complaint (J. DOE 1, et al., Plaintiffs, v. GITHUB, INC., et al., Defendants) claimed that the companies' AI-powered programming tool *Copilot* infringed copyright by using millions of lines of human-written code without proper attribution. According to *NewScientist*, this is the "first big copyright lawsuit over AI and potential damages could exceed \$9 billion" (Wilkins, 2022).

Texas A&M University Computer scientist Tim Davis claimed on Twitter that the Microsoft-owned AI programming assistant "emits large chunks of my copyrighted code, with no attribution, no LGPL license." The "LGPL" Davis mentions is a type of Open-Source use permission — the Lesser General Public License — which makes the code available for use to anyone *if* they adhere to the license requirements, such as attribution (which identifies the copyright holder of the work being reused — in this case developer Davis).



The LGPL is considered to be a weak "copyleft" license, and it typically applies to a narrow set of code. If a user modifies and distributes code covered by a weak copyleft license — as GitHub may have done with Davis's code — they would need to release the modified version under the same license as the original. Davis contends that the AI generated code he's seeing does not include the license or the attribution required by the LGPL terms. Individuals who use an Open Source Software (OSS) component are legally responsible for complying with the terms of the license. When that "individual" is AI, the question becomes "who is responsible if the license terms are violated?"

Copilot works by translating natural language to suggestions for lines of code, and trains on Open Source Software material. With the OSS license clearly attached to Tim Davis's code, he alleges that the GitHub generated code is violating his copyright.



In his Twitter discussions, Davis correctly points out that while algorithms are not generally copyrightable, the expression of the code can be protected. Even though Davis has placed his copyrighted code more or less in the public domain through the use of OSS, he is still entitled as the copyright holder, to set the terms of use including: attribution to him as the original developer, notice of his copyright, and inclusion of a copy of the LGPL license.

The monetization of the work by Microsoft/GitHub/Copilot is one of the primary irritants to developers like Davis. GitHub users pay \$10 per month or \$100 per year for access to Copilot. "Copilot's goal is to replace a huge swath of open source by taking it and keeping it inside a GitHub-controlled paywall," the complaint against the companies said. "It violates the licenses that open-source programmers chose and monetizes their code despite GitHub's pledge never to do so." The complaint also cites unlawful competition by the defendants in the "passing-off" of licensed code as their own creation, and alleges GitHub has been unjustly enriched by their unlawful conduct.

The defendants responded to the copyright allegations in early 2023 by evoking the doctrine of fair use. The companies cited the 2021 Supreme Court of the United States (SCOTUS) decision in *Google LLC v. Oracle America, Inc.* where the court determined that the reuse of application programming interfaces (APIs), including representative source code, can be transformative and fall within fair use. Although the court has yet to address the fair use defense in the case, Judge Jon Tigar ruled in May 2023, on the defendants' motion to dismiss (throw out the case), and a number of claims were dismissed; the plaintiffs were given leave to amend most of these claims; however, and they did so on June 8, 2023. Featured predominantly in the amended complaint is the allegation that "Codex (powering Copilot)

outputs copyrighted materials without following terms of the applicable licenses” as seen in the excerpt below:

"52. Below is an explanation of how Codex functions. When Codex is prompted with:

*function isEven(n) {
it assumes this is the beginning of a function written in the JavaScript language that will test whether a number is even.*

53. Based on this assumption, Codex will then provide Output meant to complete the rest of the function. Based on the given prompt, it produced the following response:

```
function isEven(n) {  
  if (n == 0)  
    return true;  
  else if (n == 1)  
    return false;  
  else if (n < 0)  
    return isEven(-n);  
  else  
    return isEven(n - 2);  
}  
console.log(isEven(50));  
// → true  
console.log(isEven(75));  
// → false  
console.log(isEven(-1));  
// → ??
```

54. The function itself occupies the first ten lines. Six additional lines follow the function, beginning with "console.log(isEven(50))". One possible explanation for Codex's inclusion of these lines is to test the "isEven" function. Though not part of the function itself, the lines will confirm the function works for certain values. In this case, the code implies that "isEven(50)" should return the value "true", and "isEven(75)" should return "false". Those answers are correct.

55. The penultimate line indicates "isEven(-1)" should return "??". This is an error, as "isEven(-1)" should return "false".

56. Codex cannot and does not understand the meaning of software code or any other Licensed Materials. But in training, what became Codex was exposed to an enormous amount of existing software code (its "Training Data") and — with input from its trainers and its own internal processes — inferred certain statistical patterns governing the structure of code and other Licensed Materials. The finished version of Codex, once trained, is known as a "Model."

57. When given a prompt, such as the initial prompt discussed above — "functionisEven(n) {" — Codex identifies the most statistically likely completion, based on the examples it reviewed in training. Every instance of Output from Codex is derived from material in its Training Data. Most of its Training Data consisted of Licensed Materials.

58. Codex does not "write" code the way a human would, because it does not understand the meaning of code. Codex's lack of understanding of code is evidenced when it emits extra code that is not relevant under the circumstances. Here, Codex was only prompted to produce a function called "isEven". To produce its answer, Codex relied on Training Data that also appended the extra testing lines. Having encountered this function and the follow-up lines together frequently, Codex extrapolates they are all part of one function. A human with even a basic understanding of how JavaScript works would know the extra lines are not part of the function itself.

59. Beyond the superfluous and inaccurate extra lines, this "isEven" function also contains two major defects. First, it assumes the variable "n" holds an integer. It could contain some other kind of value, like a decimal number or text string, which would cause an error. Second, even if "n" does hold an integer, the function will trigger a memory error called a "stack overflow" for sufficiently large integers. For these reasons, experienced programmers would not use Codex's Output.

60. Codex does not identify the owner of the copyright to this Output, nor any other — it has not been trained to provide Attribution. Nor does it include a Copyright Notice nor any License Terms attached to the Output. This is by design — Codex was not coded or trained to track or reproduce such data. The Output in the example above is taken from Eloquent JavaScript by Marijn Haverbeke."

And then comparing Copilot's response to that of Codex:

72. When we give Copilot the same prompt discussed above in Paragraph 52, "function isEven(n) {", it interprets the prompt as the beginning of a function written in the JavaScript language that will test whether a number is even, same as Codex.

73. However, the Output of Copilot in response to the prompt is different than Codex, namely:

```
function isEven(n) {  
    return n % 2 === 0;  
}
```

74. This function is much closer to what a human programmer might use as compared to Codex's older, inaccurate offering. It handles all values and types of "n" correctly. It does not cause a stack overflow for larger values of "n" like the Codex Output.

75. Copilot's Output, like Codex's, is derived from existing code. Namely, sample code that appears in the online book *Mastering JS*, written by Valeri Karpov. Like Codex's Output, Copilot's is also based upon copyrighted educational material. *Mastering JS* is a set of educational exercises for programmers. Like *Eloquent JavaScript*, there are many copies of Karpov's exercise stored in public repositories on GitHub. Programmers working through *Mastering JS* store their answers there.

76. If Copilot is prompted with the name of a function that will test whether a number is prime (that is, a number that can only be evenly divided by 1 and itself), namely "function isPrime(n) {", it returns:

```
function isPrime(n) {  
    if (n < 2) {  
        return false;  
    }  
    for (let i = 2; i < n; i++) {  
        if (n % i === 0) {  
            return false;  
        }  
    }  
    return true  
}
```

77. Though this function will work, it contains an error often made by beginner programmers that makes it much slower than it could be. Namely, the loop in the middle, which checks possible divisors, does not need to check every divisor smaller than "n," only the divisors smaller than the square root of "n". As with Codex, Copilot has no understanding of how the code works. It knows that more functions called "isPrime" contain the portion that checks for all divisors smaller than "n", so that is what it offers. It does not return what it "thinks" is best, it returns what it has seen the most. It is not writing, it is reproducing (i.e., copying).

78. Like the other examples above — and most of Copilot's Output — this output is nearly a verbatim copy of copyrighted code. In this case, it is substantially similar to the "isPrime" function in the book *Think JavaScript* by Matthew X. Curinga et al, which is:

```
function isPrime(n) {  
  if (n < 2) {  
    return false;  
  }  
  for (let i = 2; i < n; i++) {  
    if (n % i === 0) {  
      return false;  
    }  
  }  
  return true;  
}
```

79. As with the other examples above, the source of Copilot's Output is a programming textbook. Also like the books the other examples were taken from, there are many copies of Curinga's code stored in public repositories on GitHub where programmers who are working through Curinga's book keep copies of their answers.

80. The material in Curinga's book is made available under the GNU Free Documentation License. Although this is not one of the Suggested Licenses, it contains similar attribution provisions, namely that "You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License."

81. As with Codex, Copilot does not provide the end user any attribution of the original author of the code, nor anything about their license requirements. There is no way for the Copilot user to know that they must provide attribution, copyright notice, nor a copy of the license's text. And with regard to the GNU Free Documentation License, Copilot users would not be aware that they are limited in what conditions they can place on the use of derivative works they make using this copyrighted code. Had the Copilot user found this code in a public GitHub repository or a copy of the book it was originally published in, they would find the GNU Free Documentation License at the same time and be aware of its terms. Copilot finds that code for the user but excises the license terms, copyright notice, and attribution. This practice allows its users to assume that the code can be used without restriction. It cannot.

The amended complaint also contains a significant number of examples similar to those of Tim Davis's comparisons of his code with Copilot's generated code. Although the litigation is still in early stages, the result will be an important step in clarifying federal protections for software and understanding the liability associated with AI programming infringement.

Discussion Questions

1. In a customer support message, GitHub stated the following:

Training machine learning models on publicly available data is considered fair use across the machine learning community . . . OpenAI's training of Codex is done in accordance with global copyright laws which permit the use of publicly accessible materials for computational analysis and training of machine learning models, and do not require consent of the owner of such materials. Such laws are intended to benefit society by enabling machines to learn and understand using copyrighted works, much as humans have done throughout history, and to ensure public benefit, these rights cannot generally be restricted by owners who have chosen to make their materials publicly accessible.

Is this claim that *training machine learning models on publicly available code is widely accepted as fair use* accurate? Why or why not?

2. Review the following Twitter except from Tim Davis in response to a user query:



Why is the expression of code, but not an algorithm, copyrightable? What types of issues would be created if you could protect an algorithm as intellectual property?

3. Review the plaintiffs' arguments for demonstrating that Copilot (and Codex) has violated the terms of applicable licenses by generating copyrighted code. Do you agree with their evidence? Why or why not?

4. Almost immediately after Copilot began generating code, there were concerns about how the AI was trained. According to GitHub, it was trained on "billions of lines of code" in dozens of programming languages; this included code on GitHub itself, which is a common tool used for open-source developers.

- How could you design Copilot's training regimen to avoid violating the copyright (and copyleft) licenses of other programmers and developers?
- How could you ensure that any AI generated code you may use is not violating copyright?

5. According to Jonathan Bailey in the article "The Ethical and Legal Challenges of GitHub Copilot",

GitHub made much of its name and brand due to open-source developers. Before its purchase by Microsoft, it was the leading tool for such development. Now with Copilot, GitHub launched a product that many developers feel is against the ethos of open-source development by exploiting open-source code but omitting the attribution and licensing requirements that come with it. The response to Copilot's use of open-source code has been, overall, negative from the community.

- If you were the developer who created and trained Copilot, what would be your response to the software and programming community?
- What would be a viable argument to support how Copilot operates?

6. Review the following excerpt from Tim Davis's Twitter thread on the issue of the alleged copyright violations:



- What is your opinion of the use of blockchain (mentioned above) to provide remuneration to code developers for use of their copyrighted content?
- Do you have other ideas about how copyright holders could be compensated for their holdings?

Commentator Kevin Fischer mentions above that we should rethink our fundamentals around IP ownership. There have been many arguments for a number of years that believe we should forego intellectual property rights and protections, because they are hindering research and development efforts.

- What is your opinion of IP protections?
- What are some advantages of legal protections such as copyright licenses? Disadvantages?

Bailey, J. (2022). The Ethical and Legal Challenges of GitHub Copilot, PlagiarismToday. Retrieved on June 18, 2023 at <https://www.plagiarismtoday.com/2022/10/19/the-ethical-and-legal-challenges-of-github-copilot/>.

Wilkins, A. (2022). Microsoft's Copilot code tool faces the first big AI copyright lawsuit. NewScientist. Retrieved on June 17, 2023 at <https://www.newscientist.com/article/2346217-microsofts-copilot-code-tool-faces-the-first-big-ai-copyright-lawsuit/>.