WASHINGTON STATE UNIVERSITY

To: Elizabeth Chilton, Provost and Chancellor, WSU-Pullman

From: Karen Metzner, Assistant Dean of Students & Director of Center for Community Standards

Bill Davis, Interim Vice Provost

Date: August 1, 2023

RE: Final report from the WSU AI Strategy and Policy Task Force

Dear Dr. Chilton,

On behalf of the members of the task force listed below, please find attached to this memo the final report for the WSU AI Strategy and Policy Task Force that you commissioned on June 13, 2023. There are five recommendations we are providing to you based upon our work and we have provided our overview of what we found to be best practices in higher education as well as areas where there is currently a lack of consensus or guidance.

We would be happy to meet with you to review this report at your convenience.

Members of the Task Force:

Kristina Peterson-Wilson, Greg Crouch, Erika Offerdahl, Erin Gordon, Eric Shelden, Levi O'Loughlin, Kim Christen, Sharyl Kammerzell, Jacqueline Southwick, Michael Walters

Summary

Artificial Intelligence (AI) will continue to be an evolving topic of conversation for the foreseeable future. After receiving its charge, the WSU Artificial Intelligence (AI) taskforce worked throughout June and July of 2023 to explore the impacts of AI on four different areas: i) academic policies, ii) instructional pedagogy, iii) intellectual property and research policy, and iv) data security. After reviewing materials from the federal government, state governments, professional societies, and other institutions of higher education, the AI Task Force endorses the following recommendations.

1. Improve Communication. WSU should provide system wide communication to all its community members regarding existing policies and how they are applicable to generative AI platforms.
2. Establish a new AI Website. WSU should establish an AI website and populate it with resources and communication to stakeholders throughout the university.
3. Expand Faculty Training. Through workshops and communities of practice, WSU should continue to support instructors as they modify their course designs to account for the impacts of generative AI on classroom assessment tools.
4. Expand Training for Research-Active faculty. Through workshops, guides, and smaller focused working groups, WSU should support researchers across disciplines to account for the varied impacts of generative AI on research data collection, analysis, reporting, and publications.
5. Establish a continuing AI Council. WSU should establish an AI Council to continue to review and establish best practices within instruction and research, as well as continuously monitor the global trends and developments of AI.

Introduction

AI denotes the emulation of human intelligence by computational systems. These systems, composed of computers or machines, are programmed to execute tasks that typically necessitate human intelligence for completion. There are many different applications of AI, including machine learning, natural language processing, computer vision, robotics, and more. Such AI systems acquire knowledge from data, experience, and feedback to continuously enhance their capabilities. We include this broader definition of AI because future task forces or workgroups at WSU may need to explore these areas of AI.

The AI Task Force limited its focus to just one narrow area within AI, Generative AI. Generative AI produces new data (e.g., text, images, and video) based on user prompts and existing data within a defined training set such as internet accessible information. Generative AI does not include spelling and grammar checking tools.

As one example of generative AI, Chatbot GPT (Generative Pre-trained Transformer), a natural language processing software developed by Open AI, originally launched in June 2018. During its initial launch, it had significant limitations. In June 2020, ChatGPT-3 was released and brought artificial intelligence into the national and international spotlight. Its successor ChatGPT-4 now generates output that is difficult, or sometimes impossible, to detect as being created by a non-human writer. ChatGPT is just one example of language development AI, however, it has received a significant amount of attention because it is widely available on the internet for free, or a minimal fee.

Generative AI presents both opportunities and challenges for institutions of higher education, and WSU is no exception. During the Spring 2023 semester, there were several documented conversations about data security concerns related to generative AI platforms. In the realm of academics, the Center for Community Standards (CCS) received its first academic integrity report relating to the use of ChatGPT in a WSU course in late April 2023. Throughout May 2023, CCS continued to receive questions and concerns from faculty regarding the permissibility of AI in coursework and how existing academic integrity university policies apply to the use of AI in academic settings. Currently, CCS has received approximately ten academic integrity reports resulting from suspected use of AI to complete coursework. Four of those cases were appealed and are awaiting review by the Academic Integrity Hearing Board. Currently, the Academic Integrity Hearing Board has yet to review an appeal resulting from the use of AI.

In June 2023 the Provost established a system-wide AI Strategy and Policy Task Force, co-chaired by Bill Davis, Interim Vice Provost for Academic Engagement and Student Achievement and Karen Metzner, Assistant Dean of Students and Director, Center for Community Standards. The charge of the task force was to chart an institutional strategy for addressing the impact of AI in four main areas: academic policies, pedagogy (including faculty training), intellectual property and research policy, and data security.

The taskforce initially met on June 13, 2023, and continued its work in subcommittees throughout June and July. As mentioned previously, the Task Force quickly narrowed its focus to just generative AI. As the subcommittees assessed the current landscape at WSU and across the nation, it became evident that there was significant overlap in the findings of many of the subgroups. Therefore, this report consolidates the findings and recommendations from all four subcommittees into one summary. The rest of this summary will present the benefits of generative AI, the risks of generative AI, the impacts of generative AI tools on current WSU policies, ethical concerns arising from the use of generative AI, and recommendations for future action.

Benefits of Generative AI

AI is quickly becoming a transformational force in higher education. By integrating AI systems such as ChatGPT into the university environment, we can reshape the educational and research landscape, fostering a more personalized and dynamic experience for both students and faculty.

1. Generative AI will likely boost the productivity of researchers and employees since it can perform routine tasks and thus automate routine work. According to McKinsey, half of today's work activities may be automated as early as 2030 (reference).
2. Given the potential impact of generative AI in higher education and research and development, post-secondary educational programs will need to incorporate generative AI into their programs of study to prepare the workforce of the future in how to manage

and use these tools in both a practical and an ethical manner. As an innovative educational institution, WSU is well situated to be a leader in this area.

3. Generative AI has the potential to alleviate financial barriers faced by students, staff, and faculty in their pursuit of education and research. By offering no or low-cost research assistance, editing, and tutoring support, generative AI, if used ethically and effectively, can significantly increase equity and access for students, staff, and faculty.

Student Perspective

From a student's standpoint, the introduction of AI in their academic journey offers many advantages.

1. Personalized Learning: AI-based educational tools can provide students with a tailored learning experience. AI can adapt to individual learning approaches, pace, and progress, providing customized feedback, recommendations, and resources.
2. Efficient Study Tools: AI, like ChatGPT, can serve as an intelligent tutor if used within the appropriate scope. It can assist in learning new concepts, exploring topics, revising content, or even practicing problem-solving skills. *Used in combination with developing a student's ability to critically evaluate output,* AI can strengthen learning.
3. Career Guidance: AI tools can analyze data and predict trends, helping students make informed decisions about their career paths and future job markets. This is especially helpful when used in conjunction with academic advising.

Faculty Perspective

Faculty can also benefit from AI integration into the educational sphere. Although generative AI is in its infancy, potential benefits from its incorporation can be envisioned in educational settings. An important consideration will be making sure that humans stay in the loop in the implementation and introduction of AI in educational settings (reference). The US Department of Education (reference) has identified several potential benefits for instructors, including:

1. Automated Administrative Tasks: AI can assist in grading assignments, tracking student progress, and managing class schedules, significantly reducing administrative workload, and allowing more time for interaction with students.
2. Better Student Assessment: As AI continues to improve, it may eventually provide detailed analytics on each student's strengths and weaknesses, helping educators understand where a student might be struggling and tailor their teaching accordingly.
3. New forms of learning: AI is expected to create engaging, interactive learning environments and simulations that can aid in teaching complex concepts. It will be especially beneficial if these environments are built to enhance the strengths of individual learners while retaining critical aspects of social learning.
4. Continual Learning: AI itself is becoming a vast field of study. Faculty can use AI tools as a part of their research, opening new possibilities for investigation, and contributing to the advancement of this exciting field.

When used judiciously, and with the acknowledgement that current AI still lacks nuance based on differences in race, class, gender and other socio-economic factors, AI in the university setting does not replicate human intelligence; it augments our capabilities, making the educational process more efficient, engaging, and personalized. As we continue to explore and integrate AI technologies like ChatGPT into our university life, we are paving the way for a more enriched and comprehensive learning and teaching experience.

Risks of Generative AI

AI systems, like any other technological advancement, can introduce new cyber risks and vulnerabilities. Here are some key risks associated with generative AI:

1. Data Breaches: AI systems rely on vast amounts of data, and if not properly protected, this data can be targeted by cybercriminals. Breaches can result in unauthorized access to sensitive information, including personal data, financial records, or research data.
2. Adversarial Attacks: Adversarial attacks exploit vulnerabilities in AI models to manipulate their behavior. By making subtle modifications to input data, attackers can deceive AI systems, leading to incorrect outputs or decisions. Adversarial attacks can be particularly concerning in critical areas such as autonomous vehicles or medical diagnosis.
3. Model Poisoning: Model poisoning involves manipulating the training data used to train AI models. By injecting malicious data or subtly modifying existing data, attackers can compromise the integrity and performance of AI systems. This can lead to biased outcomes, unauthorized access, or disruption of services.
4. Privacy Risks: AI systems often process large amounts of personal data, raising concerns about privacy. If not properly secured, AI systems can become targets for hackers seeking to gain unauthorized access to personal information or engage in identity theft.
5. Malicious Use of AI: AI technology can be harnessed for malicious purposes, such as creating highly sophisticated phishing attacks, generating realistic deepfake content, or automating social engineering techniques. This poses significant risks to individuals, organizations, and society as a whole.
6. Lack of Transparency: Some AI models, such as deep learning neural networks, can be complex and difficult to interpret. This lack of transparency makes it challenging to understand how AI systems make decisions or identify the causes of errors or biases, potentially hindering accountability and making it difficult to detect malicious behavior.
7. Supply Chain Attacks: AI systems often rely on various software libraries, frameworks, and external APIs. If these dependencies are compromised or maliciously altered, it can lead to vulnerabilities in the AI system, enabling unauthorized access or control by attackers.
8. Social Engineering: AI can be used to automate and enhance social engineering attacks, where attackers manipulate individuals to gain unauthorized access to systems or divulge sensitive information. AI-powered chatbots or voice assistants can be programmed to deceive users, making social engineering attacks more sophisticated.

Impacts of Generative AI on Existing WSU IT and Educational Policies

There is no national or local consensus on whether the use of ChatGPT or other AI text generation platforms should be allowed or outlawed writ large in classrooms, research settings, or business procedures. As one example from teaching and learning, there are several instructors known to the task force who have already chosen to incorporate the use of these AI tools as a part of their pedagogy in brainstorming, to inform revisions, or for the generative steps of a classroom assignments. Other instructors allow students to use ChatGPT or other generative AI tools only if students disclose their use of such tools. Finally, there are some instructors who wish to ban its use on all their class assignments. As is always the case, WSU supports instructors in their choices regarding the most appropriate pedagogy for their discipline and specific course context and to allow or forbid the use of any generative AI tool.

There are existing avenues available to instructors who wish to disallow the use of AI text generation platforms within their courses. If an instructor states clearly within their syllabus that these tools are prohibited, then, with evidence that indicates it is more likely than not that an academic integrity violation occurred, a conduct case should be submitted to the Center for Community Standards for a student who used one of these platforms. This is consistent with and under Section 2(b) “Use of sources beyond those authorized by the instructor in writing papers, preparing reports, solving problems, or carrying out other assignments” or Section 2(d)(i) “Plagiarism”. Therefore, it’s important to be clear what the policies are for any class, specific assignment, and/or activity if the guidelines will vary. After an instructor meets with a student to discuss the concern, a report can be submitted.

Existing WSU policies, including Executive Policy 8, already prohibits the inclusion of legally protected or regulated data (e.g., proprietary, personally identifiable information, HIPAA, FERPA) in queries provided to generative AI platforms like ChatGPT. However, it is not clear that all stakeholders are aware that the current policy applies to generative AI tools. Therefore, the possibility exists that there may be protected or regulated information being shared with AI platforms.

While current university policy already addresses the use of AI in the classroom and the utilization of third-party platforms for regulated data, WSU should take the initiative to establish explicit and uniform communication with all stakeholders, including faculty, staff, and students. This proactive approach will solidify the applicability of existing policy to the use of AI tools and better ensure everyone is complying.

Ethical Concerns arising from the use of Generative AI

The use of generative AI raises several ethical concerns. Here are some key considerations:

1. Bias and Discrimination: Generative AI models can inadvertently learn biases present in the training data, which can perpetuate or amplify existing biases and discrimination. This can lead to biased outcomes in areas like language generation, image synthesis, or decision-making systems.

2. Misinformation and Deepfakes: Generative AI can be misused to create highly realistic fake content, including deepfake videos, images, or text, which can be used to spread misinformation, deceive individuals, or manipulate public opinion.
3. Intellectual Property and Plagiarism: Generative AI has the potential to generate content that infringes upon intellectual property rights, leading to issues of plagiarism and unauthorized use of copyrighted materials.
4. Privacy and Data Security: Generative AI models often require large amounts of data for training, raising concerns about privacy and data security. In some cases, these models can inadvertently reveal sensitive information present in the training data. Unintended Consequences: The use of generative AI can have unforeseen consequences, especially when deployed in critical domains such as healthcare, finance, or autonomous systems. Ensuring the safety, reliability, and accountability of AI-generated outcomes is a significant ethical challenge.
5. Impact on Human Labor: Generative AI has the potential to automate tasks traditionally performed by humans, potentially leading to job displacement and socioeconomic inequalities. Ethical considerations should be given to the impact on employment and the need for retraining or upskilling.

Addressing these ethical concerns requires a comprehensive approach that includes careful dataset curation, robust evaluation of models, transparency in AI systems, responsible deployment, and ongoing monitoring and regulation to mitigate potential risks and harms.

Task Force Recommendations to the WSU Provost

Enact a Risk Mitigation Plan

To mitigate the risks arising from generative AI, it is crucial to implement robust security measures, conduct thorough vulnerability assessments, and ensure ongoing monitoring and updates of AI systems. Regular training and awareness programs can help educate users about potential threats and best practices for AI security. Collaboration between AI researchers, cybersecurity experts, and policymakers is also essential to develop effective defense mechanisms against evolving cyber risks associated with AI.

To mitigate the cyber risks associated with AI, here are several important strategies and best practices:

1. Data Security and Privacy:
    - Implement strong data protection measures, including encryption, access controls, and secure storage of sensitive data.
    - Conduct regular data privacy impact assessments to identify and address potential vulnerabilities.
    - Adhere to data protection regulations and industry best practices when collecting, storing, and processing personal or sensitive information.

2. Secure Development Practices:

- Follow secure coding practices and conduct thorough security testing throughout the AI system's development lifecycle.
- Regularly patch and update AI frameworks, libraries, and dependencies to address known vulnerabilities.
- Implement secure software development practices, such as code review, static code analysis, and secure configuration management.

3. Adversarial Attack Mitigation:

- Employ robust testing and validation techniques to identify and mitigate vulnerabilities to adversarial attacks.
- Implement techniques like adversarial training and robust model architectures to improve the resilience of AI systems against adversarial manipulations.
- Regularly evaluate and improve the security of AI models against emerging attack vectors.

4. Model Validation and Testing:

- Perform comprehensive testing, including input validation and model performance evaluation, to identify potential weaknesses or vulnerabilities.
- Implement model validation techniques to ensure the accuracy, reliability, and fairness of AI systems.
- Conduct rigorous testing against real-world scenarios, considering various edge cases and potential attack vectors.

5. Transparency:

- Promote the development of interpretable and explainable AI models, allowing users to understand how decisions are made.
- Conduct audits and establish mechanisms for transparency and accountability in AI systems' decision-making processes.
- Implement methods to detect and explain biases in AI systems, ensuring fairness and preventing discriminatory outcomes.

6. Employee Training and Awareness:

- Educate employees, including developers, data scientists, and end-users, about potential cyber risks associated with AI.
- Train employees on secure coding practices, data handling, and the responsible use of AI systems.
- Raise awareness about common attack vectors, social engineering techniques, and best practices for cybersecurity.

7. Continuous Monitoring and Response:

- Implement robust monitoring systems to detect and respond to cyber threats in real-time.

- Establish incident response plans specific to AI systems and conduct regular drills to test their effectiveness.
- Monitor AI systems for anomalous behavior, unauthorized access attempts, or data breaches, and respond promptly to mitigate potential damage.

8. Collaboration and Knowledge Sharing:

- Foster collaboration between AI researchers, cybersecurity experts, and policymakers to share knowledge, best practices, and emerging threats.
- Engage in industry initiatives, standards bodies, and academic collaborations to stay up to date with the latest cybersecurity advancements.

Remember, cybersecurity is an ongoing process, and it's essential to continuously evaluate and improve the security measures in place to adapt to evolving cyber threats and advancements in AI technology.

Educate the WSU Community about the impacts of generative AI on copyright and patentability of its products

The following concepts need to be shared widely with WSU students, staff, and faculty.

1. Copyright Infringement is a risk arising from the use of generative AI. The use of copyrighted materials in generative AI input may result in liability for copyright infringement. Faculty, staff, and students alike are responsible for ensuring that all work they present is their own. Blaming a technology tool for copyright infringement will not persuade courts that the infringement did not occur. Copyright covers specific expression, not ideas, but WSU will also continue to hold individuals to the highest standards of academic integrity. Sloppy or negligent use of technological tools does not excuse inaccuracy, plagiarism, or failure to provide credit for ideas to their human source.
2. Material created solely by generative AI currently cannot be copyrighted. Copyright law requires human authorship, though the U.S. Copyright Office recognizes that computers can be “assisting instruments” so long as “traditional elements of authorship . . . (literary, artistic, or musical expression or elements of selection, arrangement, etc.) were actually conceived” by the human author (reference). Courts will develop further guidance regarding use of generative AI with a human-in-the-loop as fact-specific cases are presented, but generally for expressive works subject to copyright protection, courts require a “minimal creative spark" of human effort. In other contexts, this threshold creativity has been demonstrated through creative and non-mechanical transformation.
3. Depending on the terms of use, a user of a generative AI tool may not own its input or output. In some cases, output is owned by the organization and individual who uses the tool. In others the toolmaker may own the output, or ownership may be shared.

4. Generative AI tools likely record every query and conversation logged by users and so any confidential or proprietary information in a query or conversation may end up in the output provided to third parties.
5. Patentable inventions require a human inventor, and conception is the key to inventorship. [reference] WSU researchers are reminded that patent applications require disclosures which enable other specialists in the field to practice (reproduce) the invention. As such, proofs of concept, method steps, and parameters should be specified in detail. Generative AI promises to be a useful tool in brainstorming ideas and testing various technologies, but its output alone will almost certainly be deemed insufficient for the essential purpose of patents, which reward human effort and genius to encourage useful innovation.

Provide continuing education and training to the WSU Community about generative AI

The taskforce envisions future trainings for the following groups at the university.

1. Employee training: Discuss with HRS employee training needs and opportunities, and review current trainings for updates and added examples, such as cybersecurity and ethics.
2. Faculty/Instructor training: The Transformational Change Initiative has already planned professional development opportunities (Aug 16 and Sept 7, 2023) for faculty. The subgroup recommends additional support be given to TCI and other areas within the Provost's Office and across the different WSU campuses to continue to establish and publicize Zoom events relating to AI monthly. These events would be focused solely on teaching and learning with generative AI and will provide much-needed opportunities for faculty to interact with and learn from one another and members of the WSU learning innovations team.
3. Establish a Community of Learners: Create a virtual space for faculty to share resources. We recommend a shared file or folder on OneDrive/SharePoint with permissions to allow anyone with a WSU NID, WSU Employees, or WSU Faculty to log in. This space would be used to diagnose emergent solutions and problems related to AI in the classroom and create "just-in-time" faculty professional development opportunities. Accompanying this website will be the creation of small communities of learners, working together on shared problems, to identify and create effective solutions.
4. Student training: In consultation with ASWSU, the Graduate School, and other parties (residence halls, student organizations, etc.), trainings related to generative AI should be provided to students to ensure that they are informed about best practices and potential risks arising from generative AI. Given the need for consistent messaging and coordination of this training, it was not immediately clear to the task force who / what group in the WSU system should have primary responsibility for this activity.

Publish and maintain a WSU Artificial Intelligence Website

The taskforce has envisioned and has already started to generate content for a new WSU AI website that would have the following broad content areas:

1. What is artificial intelligence?
2. Benefits of AI
3. Challenges of AI
4. Data Security and Privacy
    a. Completed Security Reviews
5. Intellectual Property and Research
6. Teaching Recommendations
    a. Detecting and Reporting Academic Misconduct and Research Misconduct
    b. Policies and Setting Classroom Expectations
7. Professional Development
    a. Additional training opportunities
    b. Discussion Forum
    c. Calendar of Events
8. FAQ for Faculty
9. FAQ for Students

Implement a university Communication Plan related to generative AI

1. A comprehensive communication plan that reaches all WSU stakeholders with applicable information including security, privacy, and ethical considerations needs to be established. The communication plan needs to be comprehensive, with messaging through HTML, Provost/Chancellor newsletters, Deans/Chairs, and a website. There should be monthly communication during AY 2023-24 and reevaluated summer 2024. Communications will highlight the need for transparency around the use of generative AI in developing creative or inventive material.
2. Initial communication to all faculty with teaching recommendations and setting classroom expectations should be provided shortly after faculty are back on contract. This communication needs to be sent as early as possible so that instructors can update their syllabus statements as appropriate prior to the start of the Fall 2023 term. The key item to include in this communication is information about setting expectations regarding the use of generative AI in their courses.
3. For students, communication should be sent at the beginning of classes, reminding students of the importance of academic integrity, and encouraging them to be thoughtful and deliberate about their use of generative AI when completing coursework. Students should be reminded to cite their sources, including the use of generative AI, appropriately. Students should also be reminded that individual course expectations regarding the use of generative AI may vary and that they need to be responsible for understanding the expectations of all their instructors. An additional communication should be sent around midterms and around finals week with recommended best practices for students as they prepare for exams and term papers.

Establish an AI Committee

1. An institutional level Committee needs to be established to continue to develop material and monitor the development of generative AI.
2. Tasks for the AI Committee include:

- Risk Mitigation
    1. IT security expert review and identify where we have compensating controls.
    2. Have an ethics advisor review for feedback on ethics elements and where we already have applicable policy.
    3. Policy Development: Policies for research are still very much in development nationally. Given this, WSU should form a yearlong AI Research Policy Working Group made up of a select group of ADRs/VCRs and faculty who direct large research groups and faculty across colleges. This working group should be representative of the range of research interests at WSU across all disciplines. This group should gather relevant information from national organizations (including APLU's Council on Research and the National Organization of Research Development Professionals, the National Science Foundation, the National Institute of Health, Council on Government Relations, and other relevant organizations).
- Continuous Monitoring
    1. In the area of IP and AI, there needs to be continuous monitoring of guidance coming from federal agencies around the copyright and trademark status of AI-generated content. Additional training for faculty, graduate, postdoctoral, and undergraduate researchers is going to be needed, perhaps aligned, or added to existing Responsible Conduct in Research training. Information about the impacts of AI, especially generative AI on IP should be communicated through the new AI website and updated on an annual basis.
    2. ITS security: internal monitoring, security controls that are in place will be monitored, and monitoring of gap analysis and plan of action to address risks
    3. As AI software tools continue to develop, AI generated text detection tools will also continue to develop. Continued monitoring will be needed to assess the efficacy of these tools and their use in detecting AI generated content.
    4. Continued monitoring is also needed as industry standards and regulations change. As an institution of higher education, it is imperative that we stay up to date and communicative as additional industry standards and recommendations are established or updated.
- Professional Development, Curriculum, and Pedagogy
    1. This subgroup will need to interface with Faculty Senate, the TCI, the teaching academies, and associate deans to continuously diagnose the faculty experience with AI so that responsive and timely support can be provided to teaching staff.
    2. Given the potential impact of Generative AI in higher education and Research and Development, post-secondary educational programs will need to incorporate generative AI into their programs of study to prepare

the workforce of the future in how to manage and use these tools in both a practical and an ethical manner.