

WAYNE STATE UNIVERSITY
ACADEMIC SENATE

PROCEEDINGS OF THE POLICY COMMITTEE
May 13, 2024

Present: L. Beale; S. Chrisomalis; L. Clabo; D. Donahue; r. hoogland; P. Khosla; J. Lewis; N. Rossi; S. Schrag

Absent with Notice: D. Aubert; B. Roth

Guest: Curtis Kratt, Sr. Dir., Service Management; Rob Thompson, AVP CIO

I. C&IT PLANS FOR UNIVERSITY AI

Thompson and Kratt were invited to Policy to discuss the plans for a university AI platform, which they recently presented at the Information Systems Management Committee (ISMC) meeting. We are all learning to navigate generative AI and what it means to campus in terms of policy and technology. It has become clear that we must wrestle with the problem of privacy and security of student data and information, especially with the use of ChatGPT and OpenAI in the classroom, as well as figure out how to navigate academic dishonesty and the re-working of the curriculum. Thompson has tasked Kratt with helping us navigate OpenAI- and ChatGPT-style generative AI chatbots, and figuring out how we can bring this technology to WSU in a way that is private and secure and provides instructors with the ability to have more granular control over how it is used in the classroom. A solution may be to deploy a generative AI toolset on campus, either on-premises (on-prem) or in the cloud, but one that WSU can control and apply our own policy, security and rules about how it is used in the classroom. This is still in the decision-making phase, so the Policy Committee's feedback is timely.

The goal is to provide faculty and students with a secure and private platform to utilize generative AI large language models (LLMs) with conversational chatbot-type experiences or the resource retrieval augmented where you are feeding it content (i.e., Canvas course content such as assignments, syllabi, et cetera) and having it respond to students as a helper. We would like it to integrate into Canvas and to give the faculty insight and control over usage within their course with the ability to turn it on or off, and to change the models for the students to use. The introduction will be fall 2024. Kratt went on to explain why we will not use the free and easy ChatGPT: the terms of service in the privacy policy tell us that their LLMs are trained on the information that our users enter. Therefore, as we upload Wayne State information (including research data or copyrighted material, et cetera), we allow their programs to acquire WSU's sensitive information.

The challenge with generative AI and LLMs is how quickly the entire discipline is advancing. These environments are expensive and doing this wrong could result in an extreme and useless expense. Research is undertaken to identify models and the functionality that they need (e.g., text-to-speech generation, conversational, summarization). What is necessary to provide that private ChatGPT-like experience? Which LLM should we use? How do we service it? We must be able to service this and have concurrency because we may have thousands of students using this at once. Additionally, now that we can talk to it, we will want a nice chat interface instead of asking our students and faculty to use Python or Jupyter Notebook to interact. Kratt noted they have identified models available to us, either open source with Apache or MIT licensing or at a nominal cost to buy that license from the provider.

If we buy a license from a provider, Linda Beale asked how we ensure that the provider is not taking the data from research or a plagiarism check. Do we have clear information about these that we are considering? Thompson assured Policy that if we use an external model, C&IT will go through the privacy agreements with a fine-tooth comb.

Beale also asked whether the expectation is to negotiate a single, customizable license that will do what we want. Thompson confirmed if we end up using an external model, there will be privacy protections as well. We do have options for an LLM that we could create with software packages that are out there. Beale suggested anything we create from scratch would not have enough data to perform like ChatGPT. Kratt explained in most cases the LLM software is pre-trained. Whatever data set they were originally trained on will have that benefit. ChatGPT 3.5 has scraped the internet since 2021. We can do our own fine tuning on top of that, but the goal would be an offline model, not a hosted model. That is one of the options that we can utilize, but it is costly because we pay per interaction. As students and faculty ask, the model responds, and there is a charge at both sides with tokens that can rack up quickly. U-M was quick to launch their own in-house generative AI model that costs about \$100,000 per month. Thompson noted the goal is to find tools that will get us most of the way there with less cost.

Kratt discussed the ongoing research to determine how the pieces fit together. One piece is the model: Hugging Face is one of the largest language model providers and supporters. They aggregate and have a rating system. It is the most popular system to benchmark open-source models against. The second piece is serving that model. We will need to be able to talk to it remotely. Hugging Face Text Generation Inference (TGI) and Virtual LLM (vLLM) are open-source technologies available to us. The third piece is the need for a chat interface, not just a command line or terminal. We need a chat UI such as HuggingChat Chat UI or Open WebUI (previously called LlamaUI).

The next steps are to determine which to use. One option is using cloud generative AI services like U-M uses (e.g., Azure OpenAI, Amazon Bedrock, Google Gemini). It is easy to set up with no heavy metal on-prem. However, it has a high transactional cost (pay as it is used) and it will be difficult to estimate use. This could be extremely popular, and we might struggle putting a boundary on that.

Another option is the cloud-hosted/rented GPU (graphics processing units such as EC2, Vast.AI, Runpod) or big video cards that drive LLMs. They are faster than CPUs and can essentially provide instant responses. A rental GPU is a good start as we better understand our needs and investigate a full on-prem solution. That will entail bringing in Dell or Exact hardware. Campus researchers are buying some of this now, so we may have some engagement with our grid team for high performance computing when our LLM system is slow, and perhaps those extra resources can be part of the grid to help our researchers.

Beale assumed that for the on-prem model, we decide in advance how much to pay for, but individual use is not charged separately. Kratt confirmed that on prem—the estimate of expected use will determine the amount of hardware we set up, with the hope that this quickly changing discipline does not change so quickly that our hardware becomes obsolete. Thompson noted this is an advantage of first renting GPUs: it will be slightly more expensive for the short period of time used to determine need. It will make a switch to an on-prem system easier. The advantage of cloud computing is the ability to spin up as many resources as needed: if it is only for a short period of time, it is not too expensive.

Beale raised another aspect of the privacy concern. Is it clear that the internal is not accessible when renting? Thompson confirmed it is not: GPU rentals in the cloud are raw computing resources over which we have full control.

Noreen Rossi noted that AI uses high energy. Does WSU have infrastructure to support it? Thompson explained we have extra capacity in our new data center. It is true some places, such as U-M, struggle with the GPU use of electricity, but we expect to have sufficient capacity for what we are planning right now. Beale questioned what costs are projected for the next years. Thompson estimated we are at 50% capacity now, so if AI continues expanding, we may need to consider a build out in six to eight years with DTE.

Our next steps are a proof of concept. Kratt noted C&IT will utilize our Amazon EC2 (elastic compute cloud) to spin up a smaller, hosted-GPU vLLM, install a local model and an Open WebUI/ChatUI interface, and begin testing internally to see if we can meet our needs with that information. We can then estimate sizing up before we open it more broadly. We would probably ask the ad hoc subcommittee on AI members like Richard Pineau (CLAS) or Bob Reynolds (ENG) to test it.

Beale asked about the process if a faculty member were to decide this summer to have their students do particular assignments using generative AI. Will that interface connect with whatever the pilot rental system is? Thompson confirmed it will connect. The plan is for instructors to have individual control over enabling this in their Canvas course, just like they can with other tools. Beale suggested the default should have the tool turned off so that users are encouraged to be aware of various concerns. Steve Chrisomalis asked whether text output only or image generation is expected. Kratt confirmed we will start with text and send surveys to faculty to better understand whether other models are needed. Image generation may not be needed across the board.

Chrisomalis understood that this approach will handle privacy concerns by keeping it local, but he remained concerned about addressing the issues around student code of conduct and academic dishonesty. For instance, there is a lot of discourse right now on watermarking images: it is much harder to watermark text because you can just change text. Are we going to have any capacity, given that AI detection is terrible? Are there going to be any means that we can use to ensure detection? Another concern is faculty being able to see what their students have put into the tool as a prompt. What if the student is just using this casually? There are still concerns about privacy just at the institutional level. Thompson explained they have discussed whether it makes sense to try to provide instructors with information about how their students are using the tool. The capability is there, but it is potentially problematic. We should be transparent with our students, but if we are transparent about logging all their activity, they are just going to use ChatGPT and not use our tool. Beale agreed it is problematic if faculty do not know, and it is also problematic in terms of student privacy; yet the detection problem and not knowing how it is used is the biggest concern after privacy concerns are resolved. ChatGPT logs much more than just their activity, which is something students (and faculty) need to know.

Based on other institutions that are using it and given the speed with which this whole field is developing, Pramod Khosla questioned how updates will be handled. Will we need new tools within six months or is that difficult to predict? Kratt agreed it is difficult to predict. New models come out almost daily, with triple the parameters of training and those bigger models require more gigabytes of storage on the host server. An 8 billion parameter model that is 10 gigabytes big will fit on a common 16 gigabyte card, but now OpenAI and Llama are suggesting 500 billion parameter models. We will determine what is sufficient on a per-semester basis. Our base majority use will be more structured, and it will not update as quickly, so we may keep a separate environment on the cloud-rented GPU where we can increase capacity for a more specialized group of users.

Beale asked about the relationship with research on grants in terms of indirect cost recovery. Could there be a per-account charge that takes AI use into account to make it fair (e.g., so that a huge user in one lab is not creating higher costs paid by the indirect cost recovery funds of a grant in another lab)? Thompson responded that there will be a different approach for AI research use. Discussion with VPR Ezenenari Obasi supports adding a Google cloud-hosted subscription service to their tool sets. The research piece takes two forms: (i) using generative AI and the AI tools to accelerate grant writing and grant application processes and (ii) the use of generative AI in the research itself. Because of the pace of research and the speed at which we turn around grants, the cloud-hosted model makes more sense. That is going to be part of this Google-focused subscription that they are working on. This is thus a two-part approach: (i) for research, the goal is to do that within the Google cloud; and (ii) for use by instructors and students in the classroom, the focus is the model presented today. Provost Clabo noted that allows segregation of costs, and use of the ICR/F&A funds appropriately to support research.

Beale asked about the ability to control the data in the Google Gemini cloud: is it really trustworthy? C&IT is working with them before signing anything, especially when it comes to research data. Beale also wondered what the result will be when many students just play with it a lot. Thompson confirmed that is the reason for eventually bringing this on-prem to eliminate per-interaction costs. We want to encourage instructors and students to experiment with the technology, rather than have a system that requires us to disincentivize use to avoid a limit.

Beale returned to the question about what faculty know about what students have done. One could assume that there could be an accounting for each individual student's course time using the system, without logging the specific requests and responses. Thompson responded that is likely possible. It may not be real time, but we can build it from log data in Canvas. It will identify the course in the section and what feature that they were using. Chrisomalis added that a student can still use ChatGPT. Imagine a faculty member designing an AI-based assignment that requests a list of prompts used, and the faculty member can confirm that students were in fact using the tool during the assignment. That does not raise a privacy concern. That use of account information will be helpful, but it will not solve the problem of a student's intent on academic dishonesty.

Policy discussed a range of the issues connected with AI detection and non-WSU AI data scraping. Beale noted one consideration should be how to help faculty recognize the kinds of ways they could do this kind of checking. Clabo agreed ongoing education of faculty around capacity and limitations, regardless of what we do, is going to be a big part of this process. We have to start somewhere because doing nothing is the equivalent of sticking our heads in the sand, and that is not offering students or faculty the best opportunity to learn in new ways. Beale agreed. We want to share the ad hoc subcommittee on AI's report with the faculty as a cautionary statement about the terms of service of publicly available models. There are summer courses for which syllabi are being written. When Policy talked about the terms of service issue, most members thought that faculty probably do not read them; they have no idea what they are asking students to give up. Chrisomalis asked about the legal implications. For instance, there are going to be faculty who will unknowingly commit IRB violations by putting material into non-WSU AI without realizing that data then becomes accessible to Amazon, Azure and Google. Thompson suggested the more dangerous companies are the ones that are mostly free or inexpensive that commodify everything. C&IT has posted some guidelines for the use of AI for administrative purposes, listing the ways that we can safeguard information. It is a constantly evolving document: we will learn as we progress and publish guidance for the campus community.

Beale asked what groups have been involved beyond C&IT in this plan. Thompson spoke at the AI subcommittee for the Academic Senate's Facilities, Support Services and Technology subcommittee and received good feedback. He had been participating in OVPR committees, although they stopped during the transition to the new VPR. They have talked at the IT governance committees, the ISMC committee and the Research Technology Advisory Council. Beale noted former interim VPR Tim Stemmler set up a big data/AI group that met once, but has not met since the new VPR was appointed. There was interest in establishing an AI research center.

Damecia Donahue suggested it may be helpful having user test groups work with some of the different models. It might provide more insight about privacy, ethics, copyrights et cetera if we are actively engaged with the models. There is a learning curve and much education is needed. It would be helpful if people are actively involved with testing, understanding the different prompts and knowing what is happening behind the scenes. Clabo agreed. Once we have a system, the education of faculty, staff and students will be as large a project as the implementation. This can also be worked in to New Faculty Orientation, as well as some sort of welcome-back plan for faculty across the board. Beale added that the Senate office will send out the ad hoc AI subcommittee's report to all faculty and academic staff. Once the decision on which system is made, the C&IT announcement can be framed around some of the issues that were raised in the report.

II. APPROVAL OF POLICY PROCEEDINGS

The approval of the proceedings of April 29, 2024 will be approved by email, to ensure all have had a chance to review them.

III. REPORT FROM THE CHAIR

Senior Vice President for Business Affairs and Finance search: An email is going out to the search committee today announcing the generation of a sufficient pool of candidates with CVs. A discussion of candidates will be scheduled for next Monday as well as a plan for moving forward.

Permanent College of Nursing dean: Clabo discussed with Policy a plan for a permanent dean of the College of Nursing.

Enrollment: Clabo shared Monday enrollment numbers. New and continuing undergraduate enrollment is up 2% and master's is up overall 7.9% with new master's enrollment up 4.2%. Overall enrollment is up 5.06%, representing 12,814 new and continuing enrollments as compared to 12,197 this time last year.

IV. REPORT FROM THE SENATE PRESIDENT

Promotion & Tenure reception: Although something came up so that she could not attend, Beale heard the May 6 Promotion & Tenure reception was a good event. Clabo noted she was grateful to see so many members from university P&T show to support promoted faculty. People took selfies and had a fun evening.

Policy interviews with AVP Tech Commercialization and AVP SPA candidates: Policy is meeting with AVP tech commercialization candidates on May 20 and May 22 from 12 to 1 p.m. in Maccabees. Policy will also meet with AVP SPA candidates. So far only one candidate has been scheduled on May 22 from 11 to 12 p.m. in Maccabees (immediately before the AVP tech commercialization candidate).

Budget Planning Council summary document: Beale alerted Policy that she still has not seen the summary document. Clabo confirmed it is coming.

ACLU FOIA request: Beale learned that the ACLU FOIA request asked for certain parts of the proceedings of the Academic Senate plenary with the free speech discussions. That is obviously fine, since those are public documents.

Advocacy Day cancellation: An advocacy day was planned in Lansing on May 7 but was canceled due to procedural votes resulting in a quick change of schedules. Clabo pointed out no meetings will take place over the summer, but there will be a fresh start next year. Beale is hopeful that Policy gets involved in what is rescheduled in the fall.

Statement about AI terms of service: Beale shared with Policy a shorter statement versus a longer statement drafted by the ad hoc subcommittee on AI to send out with their AI report to all faculty and staff, along with the faculty and student guidelines. Policy preferred the longer statement. Jennifer Lewis suggested including the syllabus guidance distributed earlier.

V. REPORT OF PROVOST'S AD HOC COMMITTEE ON PHD EDUCATION and PETITION TO GRAD COUNCIL TO APPROVE MINIMUM 60 CREDIT HOURS

Beale reported the ad hoc committee met and concluded it would be better to reduce minimum PhD thesis credits to 18 and minimum overall credits to 60. Although programs can require more credits, the change

will influence how programs think about it. It is clear that some programs such as medicine will not want to reduce to 60.

Clabo emphasized that the committee followed process by petitioning the Graduate Council to modify the motion to include reducing the number of required dissertation credit hours for PhD from 30 block credits to 18 variable and add a motion to reduce the minimum number of total credits required to 60 instead of 90 (60 is 42 didactic plus 18 dissertation credits), with both these changes to apply to current students (if they select) as well as new students. This petition was submitted Friday to the Graduate Council.

Chrisomalis' concerns are all about process. The first piece was that the subcommittee of the PhD report that was responsible for the tuition credit hours was comprised solely of administrators, and there was concern that this was not coming from faculty. Second, there is a concern that one of the side effects of this will be that administrators will be encouraging programs to offer fewer graduate courses in order to increase the number of instructors teaching undergraduate courses. In other words, currently people who are teaching PhD seminars will not do so because there will be fewer didactic courses, and those faculty will be shunted into larger undergraduate classes as a cost-saving mechanism for the university. This is worrisome, since Jennifer Wareham (associate dean in CLAS) has said openly things like "all the programs have too many credit hours. You should all be cutting your credit hours." On the one hand the change sets 60 as a minimum (allowing departments to have more if needed in the field), but on the other hand an administrator who is one of the key authors of the petition puts pressure on departments in the downward direction.

Clabo suggested no committee member is a key author: the petition to the Graduate School is from the PhD program committee as a whole. Beale agreed that there is not a sole author; however, she noted that in the course of the very few committee meetings held, Wareham was a prominently outspoken person on that issue. She should not be viewed as having drafted this, but it is important that there be a clear pushback against any expectation that every program will cut courses. People do worry that CLAS Dean Steffi Hartwell will push PhD programs to offer fewer courses because Hartwell has actually said in a variety of meetings, including here at Policy, that she does not think PhD programs are all that important. Policy members agreed.

Chrisomalis added this change will help self-paying PhD students—and there are quite a few. Humanities and social science programs basically make up directed study courses for which they get no credit. The average number in Anthropology of didactic directed study courses for our PhD students is nine. They are taking nine of their 60 credits because we do not have enough faculty and resources to teach a big curriculum, but also because we got into the habit of saying, "well, you are preparing for your qualifying exams, when you need some credits, go take these courses." We could just say to the students, "here is a list of 50 books. You want a PhD? Read 50 books." We do not need a directed study, which is going to be funded off a GTA graduate professional scholarship. Nobody has made that case as powerfully as could be done.

Clabo's concern is that WSU's programs having a large number of non-essential didactic credits results in potential candidates going to places like OU where they get an inferior PhD education. Inflation of the number of credits artificially, in an R1 institution, creates an access issue.

Rossi commented on a potential problem for existing PhD students midway through a program. There may be required courses for that PhD that they have not taken, but they may already have 60 credits. Those required courses may be important and still need to be taken even if above 60 credits. Beale suggested this change would not allow a student to avoid taking required courses for the degree. If the program changes the number of credits, whatever remains required must be taken, even if that means a student will have taken more than the new minimum credits. Clabo agreed: it depends on what the required courses are in the reduced curriculum. When we change a PhD student's program of study, it requires the consent of their committee.

Lewis pointed out another side of the argument. A PhD recipient is acknowledged as a public intellectual and a true scholar, but we are now playing The Hunger Games in education: when other universities offer the degree for fewer credits, we compete by reducing ours. Those didactic credits are important: she has her students take things outside the field so they will be truly educated people (e.g., learn another language, take a music history course, take a mathematics course). Clabo agreed within limits. This is about teaching what our colleagues consider to be essential in the discipline to earn the doctorate. If we all took more courses every day, we would be brighter, but we do not necessarily keep people in school for multiple years to force them to take more.

Chrisomalis noted a weakness in the original PhD report was the lack of the 2019-20 data generated by Sharon Lean (CLAS) and the Academic Standards Committee of Graduate Council showing that our credit hour requirements were much higher than most other schools. The average number of didactic credit hours, both median and mode, across this sample of about 25 programs was 42, and the average number of dissertation credit hours was 18—so we are now moving to the national median and mode for those two things (compared to our current 60 didactic and 30 thesis, which puts us as an extreme outlier). It is likely that PhD education was wrongly considered a revenue-generating program here. He agrees that students can be given good guidance about what they ought to be doing, and any department can retain its current credit hour requirements. Nothing stops a student from taking additional courses, but they may not want to do so because they want to finish. Telling students they have to take a range of courses not necessary for the degree because it is “good for them to be better educated” is not a good approach. He remains concerned that the case was not made to the faculty at large in the PhD Committee process, although this change is clearly the right thing to do.

Policy members voted unanimously in support of the change to a total credit hour minimum of 60, including at least 42 didactic and at least 18 thesis credits.

VI. REPORT FROM STANDING COMMITTEE LIAISONS AND NEW BUSINESS

Campus tour: It was brought to Beale’s attention by a staff member that a potential first-year student who had planned to attend WSU and live on campus had a negative experience on a campus tour. Upon arrival at WSU’s Student Center Building (SCB), the tour guide did not greet the potential student and her family or direct them where to go, rather a security guard directed them. The guide showed no enthusiasm and read a PowerPoint verbatim. The potential students were directed to a table with t-shirts to help themselves (only XX sizes were offered). Participants of the tour asked for a bag to hold their WSU materials, but none were available. The tour took about an hour and consisted of viewing a few buildings. The only dorm toured was Chatsworth: when asked if there were other dorm options for freshmen, the guide was not able to answer. They toured UGL, but the guide noted there were no books in the library, only briefly mentioning there is another library on campus for books. When another potential student on the tour asked to see the Mike Ilitch School of Business, the guide said they would have to go on their own and warned them they may not be able to get inside the building. When the tour ended, the guide said the group was free to go and offered to show them back to the SCB but did not walk them to the parking garage.

In comparison, EMU mailed packets to potential students in advance with a parking pass. Upon arrival at EMU, guests received a bag full of swag and were offered t-shirts in a variety of sizes. The mother of the potential student explained she was not only frustrated that they had to pay for parking, but they had to figure out how to get back to the parking structure on their own using the campus map, unlike the tour at EMU where guests were walked back to the parking structure by the tour guide. The potential student was extremely disappointed with the way WSU was presented. When she arrived home later that afternoon there was a hand-written thank-you note from the enthusiastic EMU tour guide in the mail. Only a follow-up survey was emailed to the potential student after the WSU tour. The student ultimately decided to enroll at EMU.

Clabo suggested having secret shoppers participate in campus tours. There needs to be quality control measures in place. This information will be shared with Vice Provost for Enrollment Management Charles Cotton when he arrives on June 1.

Year of Focus Committee: Donahue reported the Year of Focus Committee met but there is uncertainty about the committee's task. Clabo has asked the co-chairs of the committee to present at the University Leadership Council meeting tomorrow. It is already May, and she would like to make an announcement to campus about the year of focus.

Upcoming Policy agenda items: Policy members discussed upcoming agenda items including the Humanities Center's Detroit Humanities Hub, the Campus Police Oversight Committee and the Foreign Influence Policy and documents.

Academic Senate statement on academic freedom and free speech: Policy members discussed the Senate speaking on academic freedom and free speech. Beale suggested including this in the charges for the Student Affairs, DEI and Faculty Affairs committees in the fall. Additionally, Policy could work on a draft statement over the summer and then send it to those three committees for consideration in the fall with the idea that it comes to the plenary in December.

Approved as revised via email on June 25, 2024.