UNIVERSITY OF CALIFORNIA

Ahmet Palazoglu
Chair, Assembly of the Academic Senate
Faculty Representative, UC Board of Regents

Academic Senate

Office of the President
1111 Franklin Street
Oakland, CA 94607

senate.universityofcalifornia.edu

CAMPUSES
Berkeley
Davis
Irvine
UCLA
Merced
Riverside
San Diego
San Francisco
Santa Barbara
Santa Cruz

MEDICAL CENTERS
Davis
Irvine
UCLA
San Diego
San Francisco

NATIONAL LABORATORIES
Lawrence Berkeley
Lawrence Livermore
Los Alamos

January 5, 2026

Academic Senate Division Chairs

Re: Report of the Academic Senate Workgroup on Artificial Intelligence (AI)

Dear Divisional Chairs:

On behalf of the Academic Council, I am pleased to transmit the attached report and recommendations of the UC Academic Senate Workgroup on Artificial Intelligence. **Please share the AI report with the appropriate Senate committees and other interested faculty in your division.**

The Council unanimously endorsed the report at its December 17, 2025 meeting. The Council views the report as timely and thoughtful in its identification of core principles and areas for further consideration across instruction, research, admissions, and data stewardship. The AI Workgroup anticipates this work will continue to evolve as technologies and institutional practices change. As such, the report is intended to serve as a framework for Senate-led engagement with issues related to artificial intelligence, rather than as a set of prescriptive policies.

The report is designed to stimulate broad discussion within the Academic Senate at both the systemwide and divisional levels, as well as ongoing dialogue with the administration and the UC Board of Regents. Many of the recommendations explicitly envision follow-up work by relevant Senate committees and collaborative consultation with administrative partners.

Members of the AI Workgroup have indicated their willingness to provide further consultation or context as Senate divisions and committees take up the report's recommendations and proceed with next steps.

Thank you for your leadership and engagement in these important AI-related issues.

Sincerely,

Ahmet Palazoglu
Chair, Academic Council

cc: President Milliken
Provost & Executive Vice President Newman
Academic Council
AI Workgroup Chair Steintrager
Senate Division Executive Directors
Senate Executive Director Lin

UNIVERSITY OF CALIFORNIA, IRVINE
BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO
SANTA BARBARA • SANTA CRUZ

Department of English

435 Humanities Instructional Building
Irvine, CA 92697-2650
(949) 824-6712
(949) 824-2916 fax

James A. Steintrager
Professor of English, Comparative Literature, and European Languages & Studies
University of California, Irvine

December 10, 2025

Dear Chair Palazoglu,

Please find attached the final report and recommendations of the Academic Senate AI Workgroup. The report itself, if it is to be useful, will need to be taken up by various Senate committees at the systemwide and divisional levels. Further, much of the work that will enable the University constructively to engage with and adapt to AI will require consultation and close collaboration with the Administration—and perhaps even the Board of Regents for some of the more ambitious recommendations. Although the workgroup considers its assigned endeavor done with the competition of the report, I would be happy to make myself available to answer questions and further engage with any of the groups just mentioned. Other workgroup members have also expressed their willingness in this regard.

The task of writing the report was challenging given the multifariousness of the charge and the rapidly changing nature and roles of AI technologies in higher education, research, and clinical practice. It was never an uninteresting task, however, and we hope that our Senate colleagues and others find our recommendations and insights helpful.

Yours,

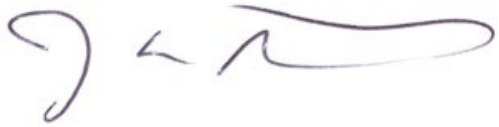

Jim Steintrager

Academic Senate AI Workgroup Report and Recommendations

December 10, 2025

Summary: Generative AI and machine learning are currently impacting most operations at the University of California, including how faculty fulfill the missions of instruction, research, clinical care, and service. This report identifies three key principles that should guide faculty and the university at large during this period of rapid change: agency, adaptability, and trustworthiness. Starting from these principles, the workgroup makes numerous recommendations in the specific areas of research, instruction, admissions, and data stewardship. These recommendations and the analyses on which they are based will in turn need to be considered by the appropriate Academic Senate committees at the systemwide and divisional levels. The Senate will then need to work cooperatively with the administration on implementation.

Background: In October 2024 the charge for the UC Academic Senate Workgroup on Artificial Intelligence was finalized and membership was established (see appendices 1 and 2 for charge and membership). We secured representation across the various divisions, representation in terms of relevant systemwide Academic Senate committees, as well as subject-matter expertise. The charge itself was developed based on a survey of systemwide Senate committees during the 2023-24 academic year and focused on the use of AI in admissions, teaching, and research, alongside issues of faculty workload and welfare.

In keeping with the charge, the workgroup surveyed Academic Senate divisional leaders on Senate engagement with AI in relation to the focus areas. UC faculty have long employed machine learning and related advanced computing in research—from foundational development to applied use—and are now integrating AI into workflows such as data analysis, grant writing, and literature review. Nevertheless, we found that broader institutional engagement with new technological developments has largely come from administrative IT units, often in conjunction with administrative units of undergraduate education, to provide information and tools to students, faculty, and staff. Faculty engagement through the Academic Senate has generally been minimal.[1] Moreover, in terms of AI-related policies in areas of faculty interest and purview such as academic integrity, personnel review, and ownership of data, little at present has been done. An important contributing factor in this regard is the difficulty of drafting or revising policies in a rapidly changing technological environment. This workgroup report and recommendations should help better position faculty via the Academic Senate to engage fruitfully and in an ongoing manner with AI in relation to our missions as a public research university.

We note that the charge envisioned the workgroup broadly surveying the faculty on their current and intended uses of AI and on their views of tradeoff considerations. It turned out that the

systemwide Provost had requested and received funding to have Tritonlytics, a UC San Diego-based survey and analytics group, develop and administer a faculty survey on AI use, opportunities, and concerns across the system. The Academic Senate workgroup worked with the Tritonlytics team in developing the survey instrument, providing feedback through focus groups on the instrument, and consulting on the process. The results of the survey, which closed in late June, were released in October 2025. The overall response rate was 6.85%, with the survey yielding ≈6k responses out of potential ≈92k respondents. Of these responses, ≈3.6k came from academic faculty and were unevenly divided across the ten campuses. While the survey's statistical power is limited, we nonetheless draw cautious inferences for topics germane to our report.

As perhaps expected, the survey revealed that many faculty have concerns about academic integrity and interest in gauging the impacts, positive and negative, of AI on learning. There are also concerns about intellectual property and copyright. A subset of faculty, particularly in the humanities and arts, are concerned about discipline-specific impacts of AI on, e.g., reading and writing skills and creative production. While the survey analysis as presented by Tritonlytics claims that the vast majority (89%) of UC faculty are "AI-engaged advocates" whereas the remaining minority are "disengaged/skeptical" (elsewhere described as "cautious"), we would counter that the survey itself does not support such a robust claim and that the realities are more complex. For example, one can be excited about and fruitfully use AI tools in the context of research while worrying about the impacts of AI tools on pedagogy—and this within the same discipline. "Engaged" and "skeptical" or "cautious," moreover, are not antonyms and may in fact be complementary. In the academic context, we would argue that they should be. Nonetheless, the survey results do provide useful information for faculty and administrators alike, and we recommend that future surveys of the faculty on work climate, satisfaction, and other areas of concern, conducted either by the Academic Senate or through Institutional Research and Planning (IRAP) in the UC Office of the President, include questions on the use of and adaptation to AI.

Introduction and Statement of Core Principles

The report is divided into four sections that we found best covered the main areas of the charge: research, instruction, admissions, and data stewardship.[2] The workgroup settled on three core principles that provide overarching guidance for our recommendations in each of these areas:

- **Agency**: Fostering faculty agency and shared decision-making is essential to incorporating AI into the university's teaching, research, and public service missions.

  The University of California will successfully address the challenges and opportunities of AI only if faculty are meaningfully consulted and participate in decision-making by the administration, and only if faculty assert their rights and responsibilities through shared governance. The Academic Senate and its various divisional and systemwide committees are the mechanisms through which faculty at UC exercise their authority as delegated by the Board

of Regents over admissions, instruction, and degree requirements; make recommendations on merit-based step increases and promotions, which is essential to maintaining academic freedom; and advise on university budgetary and other matters (see Regents Bylaw 40: https://regents.universityofcalifornia.edu/governance/bylaws/bl40.html; and APM 200: https://ucop.edu/academic-personnel-programs/_files/apm/apm-200.pdf). Throughout this report, we will indicate which Academic Senate committees are best positioned to take up a particular concern or opportunity related to AI, although we would hasten to add that these indications, much like the concerns and opportunities, are not exhaustive.

- **Adaptability**: Given the impact of AI across the university's missions and the rapidly changing nature of AI itself, faculty will need to be adaptable. Further, administrative structures will also need to maximally provide faculty with the tools for adaptation and flexibility.

  During the workgroup's discussions, two key points quickly emerged. First, AI, itself multifarious, has already profoundly impacted—and will continue to reshape—how faculty carry out research, instruction, and service to the university, as well as most, if not all, administrative functions of the university. Second, given the speed of AI's deployment and development, the institution will need to approach adaptation as an exploratory, intellectually challenging, and ongoing process.

  To take one example, addressed at greater length below, in the wake of ChatGPT's release in November 2022, many institutions of higher learning recommended instructional policies of an on-off nature (either allow or do not allow students to use large language models [LLMs] in a given class). The ubiquity and integration of LLMs into word-processing software, search engines, and learning management systems such as Canvas has made such initial policy recommendations appear quaint and largely infeasible. We assume that some of the recommendations made in this report may likewise become obsolete, which is why we insist up front and repeatedly on the importance of ongoing adaptation and on the structures of shared governance to guide such adaptation.

  We further recognize that departments and interdisciplinary units will—and should—develop their own policies and protocols in response to AI's impact on their broad disciplines and specific fields, including research, clinical, and pedagogical practices. We recommend a center-periphery approach: campuses should provide a central framework that is flexible and capacious enough to accommodate discipline-specific practices emerging from divisions and departments. Such practices may range from active and thorough incorporation of AI tools to their outright prohibition, depending on pedagogical goals and disciplinary norms. This structure allows for innovation and experimentation at the departmental and individual faculty levels while ensuring coherence and adherence to shared principles across the institution.

- **Trustworthiness**: Responsible data stewardship means trustworthiness, with sustained attention to transparency, data privacy, data ownership, and incorporating the voices of various university constituencies.

A central goal of data stewardship by UC faculty and, we would hope, by the university as a whole is that data should be managed so that those who are interested in this data and affected by its use will trust that such management is clear and appropriate. While trustworthiness has always been important to data stewardship, the emergence and widespread adoption of generative AI tools have complicated the picture and made the principle both more important and harder to implement. The risks posed by LLMs include fundamental technical challenges, such as the possibility of providing incorrect information or so-called "hallucinations," and various forms of systematic bias baked into results depending on the training data sets, particular prompts, and model architecture. Furthermore, the unauthorized use and appropriation of copyrighted material in training sets introduces significant legal and intellectual property risks that must be addressed through, e.g., transparent licensing agreements and audit mechanisms.

While the principle of trustworthiness applies broadly to all data stewardship, implementing it will take different forms depending on the type of data involved, different constituencies, and different uses. Trustworthiness includes transparency about the ways in which data are gathered, stored, and used, but transparency itself must be tailored to the particular data in question and the uses to which those data are being put. Access to some data is restricted by law (e.g., FERPA and HIPAA restrictions), by the need for confidentiality in personnel processes, and so forth. In many cases, however, the level and nature of transparency is more nuanced and contingent. For example, an admissions office noting on its website that it routinely collects and uses applicant-submitted data in certain ways might be deemed sufficient transparency; with clinical data, on the other hand, it might be more appropriate to contact patients individually to let them know that their health data is being collected, how it will be used, and who will have access. Transparency is not an absolute principle, although trustworthiness of data stewardship is or should be.

Disclosure has emerged as a key term for addressing transparency and trustworthiness with respect to AI use. Concerning research, most discussions of what constitutes necessary disclosure of AI use is taking place at the disciplinary level. This is as it should be. Departments and divisions, however, will likely need to adapt their personnel review practices in light of prevailing disciplinary norms. Further, faculty should know when and if AI is being used in the preparation of their review files and whether and how the administration is using AI in the academic personnel process. Regarding pedagogy, where the failure of students to inform instructors of AI use has been the focus of attention, we would add that students themselves should be informed when and how faculty are using AI for teaching and evaluation purposes. Again, the ongoing integration of AI into a variety of software platforms and internet-based tools will complicate our understanding of meaningful and useful disclosure. Is, for example, the trustworthiness of this report enhanced if we disclose that AI-generated summaries of the group's Zoom meetings were employed or that some research underlying the report included Internet searches with AI-generated overviews? Answering these and similar, if more important,

questions across the range of the university's missions will require frank discussion and debate.

Research

UC faculty have long used machine learning, high-performance computing, and other AI-adjacent technologies, so the recent UC-administered survey's findings are unsurprising: faculty in the main hold positive views about the impact and potential of AI for their research. AI has already transformed research practices by simplifying and supporting many key tasks. In the recent systemwide survey, "literature review" and "coding and thematic analysis of qualitative data" show up as the top two current and intended uses of AI in research (survey pp. 8-12). Beyond these, the workgroup identified and discussed additional applications not captured in the survey, including automated programming; synthetic data generation for experimental and clinical work; efficient and sophisticated search capacities, allowing researchers to navigate large databases and data sets efficiently; and automation of secondary tasks (although survey data suggests automation of "administrative tasks" in relation to research is not a faculty priority; see Q.2 and Q.3 responses, p.8). A particularly promising area that could bridge disciplinary silos and enhance collaboration is AI's improving translation capabilities, whether for cross-linguistic research or terminological alignment between fields.

To take advantage of AI tools as they enhance the university's research mission, while adhering to this report's core principles, we make the following recommendations.

Research Recommendations:

1. **The UC should invest in scalable AI/machine learning infrastructure to support research across all campuses**. While some have already made significant investments in high-performance computing (e.g., UCSF), infrastructure remains inadequate systemwide. To anecdotally illustrate the scale of the need: we have heard of individual research labs requiring GPU capacity fifty times greater than their entire campus currently provides (based on informal workgroup interviews). Campus administrative units supporting research should, if they have not already done so, conduct comprehensive assessments of computing infrastructure needs in consultation with faculty and with particular attention to GPU availability and high-performance computing capacity for AI applications.

   Given the financial constraints under which UC operates, we recommend creative and sustainable investment strategies that will avoid zero-sum game scenarios (that is, where investment in computing infrastructure would be made at the expense of other university priorities). Considering the State of California's interest in developing AI research, UC could seek a funding earmark that would not impact the State's base budget for the university. We could also pursue systemwide private-public partnerships along the lines of the non-

profit Allen Institute for AI's partnership with Nvidia and the NSF (the Open Multimodal AI Infrastructure to Accelerate Science [OMAI] project). Such partnerships can provide access to cutting-edge infrastructure while distributing costs and expertise across multiple institutions. However, the UC should prioritize collaborations that co-develop open, interoperable resources (e.g., joint research labs, shared datasets, or open-source tooling) over proprietary, closed ecosystems (e.g., subscription-based models like OpenAI's API, which restrict data portability and interoperability). To safeguard long-term flexibility and cost control as AI evolves, partnerships should embed open standards, joint decision-making, and exit clauses into their agreements.

2. **Address AI's environmental costs and societal risks.** Generative AI's resource intensity—high energy use, carbon emissions, and water consumption for cooling—conflicts with the University of California's commitments to carbon neutrality and to addressing California's climate and water challenges. Consequently, the UC has a responsibility to track and mitigate its own AI-related environmental footprint while supporting faculty research into energy-efficient AI approaches and sustainable computing.

Beyond environmental harms, AI poses social, economic, and ethical challenges including algorithmic discrimination, workforce displacement, misinformation (e.g., deepfakes and hallucinations), and intrusive surveillance. The University should actively support interdisciplinary research addressing these concerns through dedicated funding mechanisms and support for cross-disciplinary programs and initiatives. This research should encompass economic impacts (especially labor transformations), ethical issues (algorithmic bias, fairness, accountability), legal frameworks (regulation, IP rights), and safety concerns (adversarial attacks, misuse prevention). UC must remain a leader in understanding and addressing AI's broader societal implications.

3. **UC should not take a "one-model-fits-all" or "one application-fits-all" approach**. The machine-learning space is heterogenous. It encompasses a broad spectrum of techniques, tools, and applications. No single vendor or proprietary ecosystem can adequately serve the full range of UC's research needs, from historians analyzing medieval manuscripts to biologists modeling protein folding. Resources dedicated by the administration and technical implementations should thus facilitate the exploration and exploitation of different models and do so in close collaboration with the faculty.

Further, **all AI features and implementations with regard to research applications should be "opt-in" rather than "opt-out."** Here we affirm the principle of the right of refusal, noting that faculty have quite divergent views on how AI should be incorporated into research workflows.

4. **Departments should codify discipline-specific guidelines for the use of AI in research.** As with tenure guidelines and general research expectations and disciplinary

norms for merit-based career advancement, departments should come to an understanding of and explicitly articulate guidelines for the legitimate use of AI in research and protocols for disclosure.

The University Committee on Academic Computing and Communications (UCACC) and the University Committee on Research Policy (UCORP), along with their divisional counterparts are best positioned to further consider, endorse, communicate, and coordinate with the administration the implementation of the first three principles and recommendations above. The fourth is more in the purview of the University Committee on Academic Personnel (UCAP), which should work with divisional Academic Senate personnel advisory committees on encouraging departments to develop clear, discipline-specific guidelines for the use of AI in research in relation to personnel review and advancement.

Instruction

The rapid emergence and evolution of generative AI is shaping teaching and learning across the University of California. We believe that it is crucial not to constrain faculty agency when adopting and/or adapting to AI use in the instructional context and that any broad recommendations in this regard should take care not to flatten the disciplinary differences that characterize UC instruction.[3] In keeping with the overarching principles above, we would stress the following points:

1. **Flexibility, faculty agency, and shared decision-making are necessary conditions for engaging AI responsibly in the classroom.** There is no single model for incorporating AI or, for that matter, not incorporating AI in instruction—nor should there be. Effective policy must reflect the reality that AI interacts differently with different disciplines, pedagogies, and student populations. Preserving faculty autonomy and involving instructors in decision-making—especially around tool adoption, curricular change, and enforcement—is not only good governance but also a condition for pedagogical coherence.

2. **Responsible data stewardship in instruction requires ongoing attention to transparency, privacy, and data ownership.** Even in instructional settings, AI raises significant questions about student data, authorship, and consent. Many widely used tools send user inputs to proprietary systems with opaque data practices. The burden of disclosure often falls to instructors without institutional support. Any policy or practice related to instructional AI must engage these concerns directly.

3. **The adoption and use of instructional AI tools must be grounded in discussions about their intended contexts and communities.** AI tools reflect specific pedagogical priorities and institutional values. Their effects are not neutral, and their adoption should not be automatic. Any systemwide or campus-level tool procurement should include consultation

with instructors who will use—or decline to use—those tools, with special attention to accessibility, equity, and long-term reliance on commercial vendors.

At present, the Cal State system has bought into Open AI's ChatGPT Edu for supposed chatbot enhancement of instruction and the California Community College system has partnered with Google to similar ends. Within California's tripartite public higher education system, UC has not or not yet partnered with—or bought into—major private AI/LLM providers on the instructional front and is unlikely to do so given the autonomy that individual campuses generally enjoy. Some campuses have developed proprietary, customized chatbots designed primarily for students that are constructed from big tech products.[4] We would stress that it is not too late for administrators in the UC system and on individual campuses to consult proactively and work cooperatively with faculty through the Academic Senate on the procurement and development of AI products and software aimed at enhancing instruction.

With these three points in mind, we offer the following recommendations concerning classroom practice, curriculum design, and academic integrity. Accompanying these recommendations, we have articulated instruction-specific principles and rationales, along with suggested guidance and best practices.

Instruction Recommendations:

1. Support the development of a systemwide repository of AI-informed assignments, annotated by discipline and instructional goal.
2. Fund campus-based faculty development programs focused on assignment design in AI-rich environments.
3. Encourage integration of critical AI literacy into early coursework or general education programs.
4. Support partnerships between faculty and learning centers to scaffold research and argumentation skills in AI-informed classrooms.
5. Encourage the development of non-punitive educational responses for first-time AI misuse.
6. Create shared faculty resources for AI use and AI-related integrity in the classroom.
7. Work toward consistency between instructional support units and student conduct processes, especially on campuses where academic integrity is handled outside the teaching and learning infrastructure.
8. Have campus administrative units overseeing undergraduate instruction catalog and concisely describe, preferably on a regularly updated webpage, what AI instructional tools are currently available to students and instructors across a given campus (e.g., any AI-driven data analysis tools built into LMSs like Canvas; any AI-driven personalized tutoring tools; etc.).

Recommendations 1 and 2 flow from the following instruction-specific principle: **Learning is an active and demanding process, even when AI is present.** Generative tools can support exploration and synthesis, but they cannot replace the intellectual work that learning requires. Moreover, we recognize that faculty may have legitimate reasons to decline incorporating AI into their classrooms and that an individual faculty member's control over pedagogical methods is an essential feature of academic freedom.[5] When instructors do choose to integrate AI into their courses, such integration should enhance student engagement rather than shortcut it.[6] The education sector is just beginning to assess the impacts of AI on, e.g., critical thinking and numeracy, and faculty will want to carefully monitor and respond to emerging research in this area. Moreover, the design of assignments, assessments, and feedback mechanisms will require sustained attention as the capabilities of AI evolve.

Guidance: Assignments that emphasize iteration, reflection, or comparative analysis can help students use AI to explore ideas without outsourcing their thinking. When students are asked to revise AI-generated works, critique their structure, or compare them to peer work, they begin to see where their own judgment and voice matter.

Best Practices:

- Provide specific examples of appropriate and inappropriate AI use in syllabi and major assignments.
- Create low-stakes opportunities for students to reflect on how and why they are using AI.
- Design assignments that foreground process, such as draft-and-revision workflows or critical analyses of machine-generated outputs.

Recommendations 3 and 4 flow from the following instruction-specific principle: **The presence of AI in the instructional context offers an opportunity to reaffirm what makes learning at a research university distinctive.** Generative AI invites questions that lie at the heart of academic life: What counts as knowledge? What does it mean to develop an argument, evaluate evidence, or make a claim? Rather than responding defensively, instructors can use AI as a prompt for these questions—reaffirming the value of research-based learning in the process.[7]

The systemwide AI survey indicates strong agreement that "discussion of the appropriate use of AI-powered tools and technologies should be a part of all courses for students" and also indicates that "promoting AI literacy to avoid over-reliance" on the technology appears the top challenge (pp. 24 and 27). Regardless of what an individual faculty member thinks about the use of AI in their classroom, it is important to be able to discuss the topic with their students. Since there is no one-size-fits-all solution to the problem of whether or how to incorporate AI into the classroom, students look to faculty to understand what it means to engage with the technology ethically and appropriately. Students left to decide that for themselves will generate a wide range of responses that seem reasonable to them but unacceptable to their instructor. Cultivating classroom dialogue about AI helps to provide students with the language and guardrails needed for them to engage with the technology effectively. To enhance their ability to address AI with students, faculty should

have access to and input into critical AI literacy programs aimed at developing their understanding of AI use and misuse in the instructional context. Libraries are likely to be good partners in this regard.

AI literacy is increasingly a fundamental skill that students need not only during their time as students but also on joining the workforce and to be responsible citizens. Properly critical AI literacy is not limited to proficiency but includes understanding how AI interacts with learning, positive and negative impacts on critical thinking and skill development, ethical concerns, and so forth. Some aspects of AI literacy are discipline specific, while others are more general. These differences can and should be reflected in the curriculum. Campuses may want to consider developing an AI general education/breadth requirement or incorporating critical AI literacy into extant general education and breadth requirements.

Guidance: With respect to discipline-specific AI literacy, instructors might use AI-generated text to teach disciplinary norms, expose bias, or frame discussions of method. Students can learn to ask where a model's answers come from, what assumptions they rely on, and how those assumptions differ from those of a scholar, a practitioner, or a peer.

Best Practices:

- Begin the quarter or semester with a conversation about how learning happens in the discipline, how AI tools interact with those methods, and the rules and standards that will govern AI use (if any) in the course.
- Use AI-generated responses as raw material for critique, synthesis, or counterargument.
- Invite students to consider how knowledge is made and what role they play in that process.

Recommendations 5, 6, and 7 assume that **academic integrity requires clarity, conversation, and trust—not just enforcement.** As AI tools become harder to detect and easier to misuse, students are encountering a patchwork of unclear expectations. Many students misunderstand what the tools do or assume that prior approval in one course applies in another. Faculty, in turn, can lack the training and knowledge needed to develop functional AI policies and speak openly with their students about the role that AI plays in their classroom.[8]

Staff handling complaints of academic integrity have developed a blend of human and technical means of AI detection, often through painstaking research of their own.[9] Our interviews with academic integrity officers made it abundantly clear that handling AI-related violations of academic integrity within offices of judicial affairs (as opposed to stand-alone units on academic integrity that handle the issue as a matter of teaching and learning) can lead to very substantial problems of staff morale and create unsustainable institutional processes.

Guidance: Academic integrity in the AI era is not primarily a matter of catching violations. It is about helping students understand how ideas are developed, what it means to take responsibility for

one's work, and why certain norms—authorship, attribution, collaboration—exist. These questions require conversation, not just policy.

Best Practices:

- Offer assignment-specific guidelines on how AI can be used and why.
- Discuss the ethical and epistemic stakes of using AI—not just the rules.
- Frame academic integrity issues related to AI as opportunities for students to reflect on the kind of learners they want to become.

Recommendation 8 simply ensures that instructors know what AI tools are available to them and to their students. This is particularly important in a rapidly changing instructional environment.

At the systemwide level, the University Committee on Educational Policy (UCEP) is best positioned to take up these various recommendations and help coordinate discussion across the divisions. The Coordinating Committee on Graduate Affairs (CCGA) should also consider these recommendations in the context of graduate education.

Admissions

UC faculty are delegated by the Regents authority over undergraduate and graduate admissions policy in the UC system. The implementation of undergraduate admissions policy is largely carried out by admissions offices at each campus. By Regental policy, undergraduate admissions procedures must be comprehensive (reliant on multiple factors) and holistic (no single factor is given a fixed weight).[10] Graduate admissions is carried out by professional schools and academic divisions and departments. In the case of graduate admissions, direct faculty involvement in implementation is the norm. The Board of Admissions and Relations with Schools (BOARS) is the Academic Senate committee with oversight of admissions policies and practices across the University of California system, while campus-level committees provide additional oversight of undergraduate admissions at the local level.[11] Local Senate committees approve admissions policies and procedures that adhere to UC policies and campus objectives prior to each admissions cycle. The rapidly evolving use of AI by applicants and admissions offices across the country highlights the pressing need for Senate committees to evaluate the use of AI in admissions at its campuses and to consider AI-related policies.

AI is already playing a significant role in admissions at institutions across the country, with college and university admissions offices adopting AI tools in efforts to improve outcomes and streamline workloads. The potential uses of AI tools in admissions are numerous. These include targeting students and high schools for recruiting efforts, identifying applicants who are most likely to be eligible for financial aid and scholarships, generating customized communications to applicants, streamlining the application process, evaluating domestic and international transcripts, reviewing essays, predicting student success, detecting fraud, and various other forms of analysis.[12] There

are additional uses in the context of graduate admissions, often conducted by professional schools and academic divisions and departments, such as aiding in matching of applicants to faculty based on research interests and external funding opportunities. For applicants, AI tools can be used to learn about colleges and universities, provide application and financial aid advice, and assist in writing admissions essays.

The important relationships between Senate admissions policy committees and admissions offices at each campus highlight the need for a clear process of consultation between these constituencies. The varied nature of admissions practices across campuses at the undergraduate level, as well as across professional schools and academic departments at the graduate level, highlights the need for autonomy and adaptability of policies and implementation. Further, the sensitive nature of admissions highlights the importance of careful stewardship, not just for data security but also to ensure that new practices are auditable and do not introduce bias. Finally, evaluating AI tools and adopting related policies and practices require a clear grasp of the goals of the admissions process. Some guiding principles are clear, such as reducing bias in the admissions process and ensuring data security. However, a more sophisticated consideration of how AI might aid in admissions and of potential pitfalls requires understanding what comprehensive and holistic review are meant to achieve and how they are being implemented across the system.

Admissions Recommendations:

1. Incorporate review of the use of AI by admissions offices into the regular approval cycle and ensure that the appropriate Academic Senate admissions committees are consulted about and approve of the procurement and use of new AI tools.

Admissions cycles require regular communication between Senate faculty who set policy and admissions offices responsible for implementation. Because AI is rapidly changing the admissions landscape, the university requires a clear framework for regular communication and consultation about AI use, including the need for any policy changes, between admissions administrators, the Senate policy committees at each campus, and BOARS. Senate approval should be sought prior to the adoption of new AI tools. This is particularly important regarding core aspects of admissions such as transcript and essay evaluation and applicant selection.

Many campus committees approve the policies for the upcoming admissions cycle during the prior academic year (e.g., the policies and process for fall applications might be reviewed during the spring quarter or semester of the prior academic year). This ensures time for proper reader training and implementation of policy changes. The same timeline may be conducive to reviewing plans for the adoption of AI tools or changes in how such tools are used. Fine-tuning the specific parameters of AI use might occur with faculty input during the admissions cycle. However, we note that consultation further in advance is likely necessary in the case of the adoption of new AI-based admissions products that require more careful and significant evaluation. Having clear plans and

expectations in place for communication and approval is likely to be important for the adoption of new tools.

2. Have BOARS regularly examine and coordinate use of AI in admissions across the divisions.

At the systemwide level, BOARS is uniquely situated to evaluate the use of AI in admissions across campuses and to propose overarching AI-related policies. The committee's membership includes Senate faculty from each campus, and it regularly consults with admissions and enrollment management administrators. BOARS produces the "Annual Report on Undergraduate Admissions Requirements and Comprehensive Review," which summarizes the admissions process at each campus, and could be updated in future years to reflect the extent to which AI tools are used in evaluation and selection.

3. Reassert the purview of the Academic Senate over the admissions process.

Regular consultation may facilitate an increased role for faculty in the details of the admissions process. If and when admissions offices adopt AI tools, faculty will have the opportunity to evaluate whether their use is consistent with the intended admissions policies. University faculty also have expertise in computing and AI that might be especially useful when evaluating sophisticated tools, creating the potential for a strong link between faculty and admissions offices. Further, BOARS could serve a vital role in assessing the use of AI across campus admissions offices and policies across divisional Senate committees, thereby coordinating best practices and policies. Such evaluations might consider which AI tools and practices promote access and equity and reduce bias.

4. Recognize and encourage divisional differences, flexibility, and autonomy within the constraints of Regental and systemwide policy on UC admissions.

UC campuses differ significantly in their resources, selectivity, and the students they serve. They therefore need and enjoy significant autonomy in determining their admissions objectives, policies, and practices. This autonomy should extend to AI tools that individual campuses might adopt and how they would be used. The use of AI tools for admissions should not be required and policies relating to the adoption of AI tools should provide the flexibility necessary for campuses to achieve their particular goals.

Campus Academic Senate admissions committees are well situated to consider the specific needs and goals of their institutions and the policies necessary to achieve them, especially regarding undergraduate admissions. At the graduate level, professional schools and academic divisions and departments oversee admissions across a diverse set of programs. Such programs and their application processes differ substantially in their admissions criteria and should be granted significant autonomy regarding if, and how, AI tools might be adopted in the evaluation and selection of candidates.

5. Ensure that AI tools used in admissions are secure, auditable, and free of bias.

Several considerations are crucial for the responsible stewardship of AI in admissions. The first is ensuring the security of applicant data. Students entrust the UC with personally identifying information, transcripts, test scores, and financial records, as well as essays revealing sensitive personal experiences. The use of AI tools must not be allowed to jeopardize data security. For example, any adoption of third-party commercial AI products must be done in such a way that there is no chance that data from UC applicants is shared or used by outside parties. Second, AI tools have the potential to identify and reduce bias in admissions, but also the potential to introduce systematic bias. For example, AI tools may incorporate bias against certain groups because of the tool's programming or the data on which it is trained, and such bias may not be apparent without careful evaluation of the tool, its programming, and its performance. Thus, tools used in the admissions process (e.g., for recruitment, communication, evaluation, or selection) should be thoroughly vetted to ensure that they do not introduce bias, much as human readers are trained with the same objective. Finally, the use of AI tools in admissions should be done in such a way that it is auditable. The admissions process should be replicable and documented for oversight and potential legal review. Each of these factors (data security, bias, and auditability) should be considered explicitly during the evaluation of AI in admissions and during consultations between Senate committees and admissions offices.

6. Monitor and evaluate admissions-related workload and work distribution changes

At the administrative level, the adoption of AI tools promises, among other things, greater efficiency. As some tasks are automated (basic transcript evaluation, compliance work, etc.), we are likely to see a reduction in overall workload in admissions offices. We are also likely to see changes in work distribution, such as a reduction in the need for outside, part-time readers but, potentially, an increase in the need for internal data analysts. These AI-related changes should be tracked, evaluated, and reported to Senate admissions committees.

At the Senate faculty level, BOARS is likely to face an increase in workload in understanding how AI is being used and perhaps evaluating potential policies. For campus-level faculty committees, there will be additional work associated with adopting AI tools and refining their use for each admissions cycle. This work is likely to be quite intensive, especially during the transitional period, and require committee members with relevant skills. The adoption of AI may also allow Senate faculty to become more directly involved in the implementation of undergraduate admissions (which has been shifted over time to administrative units because of scale). Temporary and long-term changes in faculty service workload and work distribution should also be tracked and evaluated. If admissions-related service exceeds normal expectations for Senate faculty as part of their regular duties, compensation is appropriate (research stipend; partial or full course release; etc.).

Data Stewardship

Throughout this report, we have mentioned the need for appropriate data stewardship with relation to the adoption and use of AI tools. We take data stewardship to be the disciplined oversight and management of data generated across and held by the institution. Data stewardship covers ownership, governance, security, accessibility, sustainability, and ethical use in an academic setting to foster trust, innovation, and research advancement. Data generated at UC can be categorized in various ways. We have adopted the taxonomy below because practical considerations and implications flow clearly from it and because it delineates and emphasizes the rights and responsibilities of Senate faculty as teachers, researchers, clinicians, and as employees and/or administrators.

- **Instructional Data**. Data generated by faculty, lecturers, and students in the course of instruction.
- **Research Data**. Data generated in the course of research conducted by faculty, as well as staff and students.
- **Clinical Data**. Data generated in the conduct of faculty members' clinical practice.
- **Administrative Data.** Data such as payroll and other employment records, institutional financial data, data on, e.g., course enrollment and related matters, overseen by registrar's offices, etc.

The emergence of AI has underlined that data have value. Given the nature and missions of the University of California, data are not only a strategic asset for the institution but also part of **a public trust** and should be used to further **the public good**, and particularly the good of the people of California. Fundamental to trustworthy data stewardship is that all university constituencies, including faculty and the public at large, have sufficient insight into the university's policies around and handling of data that they believe these essential criteria have been met.

Recognizing and guiding the exploitation of the value of data produced within UC and across all categories of the taxonomy is, overall, not currently covered by university- or campus-based policies. A key starting point for addressing data stewardship is applying the taxonomy and then **identifying who, what, or whether anyone owns the data** at issue. However, the fact that the university or a faculty member owns a given data set does not entail that it should be exploited in any legally allowable manner. This becomes a vexed issue with the possibility of monetizing access to data sets (a possibility akin to, but not the same as, to monetization of IP). Moreover, while the university as a public institution is committed to transparency and open access, it is not clear or obvious that all university- or faculty-owned data of whatever type should be freely accessible to private companies that can exploit such access for gain (even where, arguably, such exploitation ultimately benefits the public economically or otherwise).

The university ought to demonstrate adaptability to using AI but should not do so without guardrails in place. AI tools can currently be used internally to analyze UC-generated data across the

taxonomy, and we assume that more opportunities lie ahead. At present, trustworthy data stewardship entails **holding AI tools and utilization to a higher standard** than established methods of data analysis because of the novelty of these tools, potential for various sorts of bias, lack of algorithmic transparency, and potential for large-scale impact. The choice of which AI tool to use to process certain university data is potentially of more significant consequence now than, for example, the choice of which spreadsheet tool to use to process university financial data. The results of AI tool use should also be auditable, and audits be routinely carried out. Even if this higher-standard metric and routine audits will offset some of the efficiencies of AI, these are necessary transitional costs to establish trust in the university's use of AI.

In keeping with the principle of agency, **the voice of constituents in decision-making around AI use** should be heard and incorporated, and this is particularly true insofar as faculty form a constituency regarding many uses and aspects of administrative and other types of data at the university. We would add, however, that identifying constituencies and incorporating constituent voices beyond the faculty is a crucial feature of data stewardship for the public good. For example, student voices should be heard, as well as those of patients in the university's clinical settings. We would also suggest gathering input from those in K-12 education, especially with respect to potential applicants to UC and the admissions process. This principle touches on the broader issue of what has been termed "data dignity": transparency about data collection, informed and affirmative consent to data collection, clear opt-out or opt-in procedures, privacy protections, and, according to strong models of data dignity, compensation for those who provide their data. While incorporating the voices of constituents and attending to data dignity are complex matters, the university should be at the forefront of these discussions.

While we believe that the data taxonomy presented above and suggested for adoption will lead to greater clarity about various data stewardship issues, the application of the taxonomy is not always self-evident. In fact, the taxonomy may help identify and bring into focus various complications of data stewardship. These complications are not necessarily new with AI but addressing them now takes on an even greater urgency. The following example may help elucidate this issue. A wealth of data is produced within and by Learning Management Systems (LMSs). Senate faculty instructors by policy own copyright to their course materials unless such copyright is *explicitly* waived or superseded (e.g., when the university expends "exceptional university resources" on course development, as may be the case with the development of online courses - see University of California Policy on Ownership of Course Materials: https://policy.ucop.edu/doc/2100004/CourseMaterials). LMSs such as Canvas, however, are far more than repositories for course materials; they also include data about individual assignments and grades, as well as an array of "student engagement" data. Who owns such data and who has access to them? Are they faculty owned or the property of the institution? Are they or should they be available to faculty researchers at the instructor's home campus? To researchers at other campuses? To what extent are students the owners of the data they produce in their interactions with LMSs? To what extent should LMS data be available to companies outside that see them as a valuable resource? Who profits from and controls the sale of such data? We assume that the

institution can make a strong claim to instructional data that extend beyond course materials as administrative data, but we would also insist, given the Senate purview over instruction, that Senate faculty: need to be involved in the work of line drawing (i.e., determining which data count as course materials and instructional data over which faculty maintain ownership and which do not); need to understand whether and how data generated in the course of instruction are accessed and being used; and need to be involved in making decisions about the questions raised above and in adapting current policies and drafting of new policies related to these questions. In other words, LMSs, which have become a ubiquitous feature of the instructional landscape, touch on all the principles articulated above and raise a host of complex questions about data ownership and use.

Data Stewardship Recommendations:

1. **Ensure that faculty, via the Academic Senate, are members of administrative data stewardship committees and equivalents.** At present, several campuses have administrative data stewardship committees. We recommend not only that all campuses have such committees and that the Academic Senate should encourage their creation where they do not exist, but that there should be a systemwide group providing coordination and sharing insights across the campuses. In all cases, participation of faculty will be crucial to the success of such committees, even when their focus is on administrative data (e.g., payroll and other employment records, institutional financial data, course enrollments and related data overseen by registrar's offices). While administrative data is often seen as belonging to the administration, clear governance on usage, transparency, monetization, etc., is needed and faculty need to be consulted and meaningfully incorporated into decision making.

2. **Develop clear policy guidance on the monetization of data.** In the course of research and clinical practice, faculty, as well as students and staff, generate data that have the potential for monetization. Existing policies such as the systemwide Policy on Inventions, Patents, and Innovation Transfer (https://policy.ucop.edu/doc/2500493/PatentPolicy) and UCSF's various data governance policies and guidelines (https://data.ucsf.edu/compliance) are incomplete when it comes to AI, data stewardship, and monetization. These issues should be taken up by systemwide Academic Senate committees and divisional counterparts in conjunction with the administration. The most appropriate systemwide committees in this regard are the University Committee on Planning and Budget (UCPB) and the University Committee on Research Policy (UCORP). Of crucial importance in such considerations will be how monetization of data fits or does not fit with the university's commitment to the public good and to pursuing its publicly funded missions.

3. **Develop clear policy guidance on the sharing of university-generated data.** Similarly, more policy guidance is needed about the sharing of data both within and without

the university. For example, we are concerned that the university's commitment to open access regarding research results is in tension with the ability of private companies to mine such research for gain and not necessarily for the public good. We also think that better guidance needs to be provided for access to and sharing of data sets produced by faculty research and in the course of clinical practice within and across divisions. These are issues that are relevant to systemwide Academic Senate committees including UCORP and the University Committee on Library and Scholarly Communication (UCOLASC).

Concluding Remarks

As we noted at the outset, the university's engagement with AI, very much including faculty engagement, will need to be ongoing and adaptive. There are no simple, singular or definitive responses, for example, to how AI might be used most effectively in the classroom, research, or clinical practice. For these reasons, most of our recommendations take the form of assigning more work to the faculty through the Academic Senate. Here we would stress that AI has and will continue to have profound workload implications for faculty, whether we are eager adopters, highly skeptical, or somewhere in between. Although the tech industry and many in the IT world have highlighted AI's potential for introducing efficiencies and other labor-saving aspects of the technology, the very need to adapt to AI will tend to increase workload for faculty in the short term to middle term and perhaps beyond. All instructors, for example, regardless of whether they use AI in the classroom, will have to figure out how to adapt their pedagogical approaches to its presence, experiment with different modes of assessment, and evaluate and address the impacts of AI on learning. Many researchers will want to take advantage of AI's capabilities but will also need to work to identify and mitigate bias, identify and correct for "hallucinations," and work to not allow AI to too narrowly determine research agendas and protocols. (Indeed, while AI may be extremely good at optimizing explicitly targeted metrics, this often occurs at the expense of more implicit goals that would usually be byproducts of the explicit goal.) Clinicians will need to actively figure out how to balance the beneficial aspects of face-to-face care with time-saving AI-driven communication and how to take advantage of AI's diagnostic capabilities while maintaining complementary and crucial practitioners' skills.

Finally, we would underline the need for AI procurement decisions in the domain of academic affairs to be informed, and ultimately driven by, faculty input. The administration and administrators should not expect faculty to readily adopt AI tools for which they have not asked and that may be ill-suited to their goals in teaching, research, and clinical settings. There exist already numerous Academic Senate committees at the systemwide and divisional levels that are appropriate forums for procurement conversations and coordination among IT and other administrative units and the faculty. Nonetheless, individual Senate divisions may wish to establish standing committees with the specific charge of coordinating AI-related procurement at the campus level. Any procurement or establishment of collaborative relationships with AI industry

partners at the systemwide level should be done in close coordination with the systemwide University Committee on Academic Computing and Communication, with Academic Council and other committees, and with the need to protect academic freedom front and center.

Notes:

[1] The problem appears general in the higher education sector: lack of meaningful consultation with faculty on the procurement and deployment of AI tools, especially those that impact instruction, is a leitmotif of the AAUP's recent and survey-based report "Artificial Intelligence and Academic Professions" (https://www.aaup.org/reports-publications/aaup-policies-reports/topical-reports/artificial-intelligence-and-academic).

[2] To facilitate and focus our work, subgroups were established in each of these main areas. Subgroup leads were Lisa Yeo (Data Stewardship); Rita Raley (Research); Gerardo Diaz (Instruction); and George Bulman (Admissions). The workgroup as a whole considered, discussed, and ultimately endorsed recommendations made at the subgroup level.

[3] We would note that, in arriving at recommendations discussed and endorsed by the workgroup as a whole, the instruction subgroup combined review of pedagogical research, structured conversations within the subgroup, and interviews with academic integrity officers at UC Davis and UC San Diego. These interviews underscored the practical challenges of enforcement, the unevenness of institutional resources, and the difficulty of separating technological change from structural strain.

[4] UC Irvine's ZotGPT, for example, uses Microsoft Azure AI and Amazon Web Services [AWS] for infrastructure and accesses data from various big tech LLMS; UC Davis's GauchoGPT is built on Open AI's ChatGPT. Meanwhile, UCSD's TritonGPT is built out of various open-source components, including Llama 3, an open-source LLM developed by Meta.

[5] As section 010 of the UC Academic Personnel Manual, which covers academic freedom, states: "The faculty has authority for all aspects of the course, including content, structure, relevance of alternative points of view, and evaluations" (https://www.ucop.edu/academic-personnel-programs/_files/apm/apm-010.pdf). At the same time, it is the faculty's responsibility to ensure that instruction adheres to equitable and professional standards.

[6] On the relationships between AI and human effort, see Jose Antonio Bowen and C. Edward Watson, *Teaching with AI: A Practical Guide to a New Era of Human Learning* (Johns Hopkins University Press, 2024).

[7] For more on exploring disciplinary conventions with AI, see Joseph Dumit and Gerardo Con Diaz, "Co-Writing with AI (Academic Style)," in N. Begus (ed.), *Writing with AI* (University of Michigan Press, forthcoming).

[8] Tricia Bertram Gallant, Director of Academic Integrity at UC San Diego, has written extensively on this topic. See Tricia Bertram Gallant and David Rettinger, *The Opposite of Cheating: Teaching for Integrity in the Age of AI* (University of Oklahoma Press, 2025).

[9] "Current Trends in Academic Misconduct," presentation by Marilyn Derby, Associate Director for Student Conduct, Office of Student Support, University of California, Davis.

[10] Campuses may and do implement comprehensive and holistic review variously, but the within the guidelines determined by the Regental Policy 2102 (https://regents.universityofcalifornia.edu/governance/policies/2102.html).

[11] Campus-level Academic Senate admissions committees include the Admissions, Enrollment, and Preparatory Education Committee (UCB), Committee on Admissions and Enrollment (UCD), Council on Enrollment Management and Admissions (UCI), Committee on Undergraduate Admissions and Relations with Schools (UCLA), Admissions and Financial Aid Subcommittee (UCM), Undergraduate Admissions Committee (UCR), Academic Senate has a Committee on Admissions (UCSD), Committee on Admissions, Enrollment, and Relations with Schools (UCSB), and Committee on Admissions and Financial Aid (UCSC).

[12] For an example of one way that a UC campus is already incorporating AI into admissions, UCLA has recently launched a chatbot that provides guidance on financial aid and the admissions process: https://dts.ucla.edu/newsroom/ai-in-action-transforming-uclas-enrollment-management-support-with-bruinchat.