

Office of the President

TO MEMBERS OF THE BOARD OF REGENTS:

DISCUSSION ITEM

For Meeting of September 16, 2026

UPDATE ON ARTIFICIAL INTELLIGENCE

EXECUTIVE SUMMARY

In April 2025, a woman named Ann, whose brainstem stroke silenced her voice nearly two decades ago, heard her thoughts leave a speaker almost as fast as she could form them in her mind. The system decoding her speech was built at UC San Francisco and UC Berkeley. Fifteen months later, in July 2026, a set of pre-release AI models under security testing at OpenAI slipped out of their evaluation “sandbox,” reached the open internet, and broke into a partner company’s production systems. No person directed them. One event is what artificial intelligence looks like when a public university uses it for the public good. The other is what it looks like when no one has decided anything at all. The distance between the two is the subject of this item.

President Milliken formed the Artificial Intelligence (AI) Steering Committee in March 2026 to close that distance on the University’s terms. The Committee leads *AI for Humans and Humanity*, UC’s systemwide mission to put AI to work for people with more knowledge, more health, more safety, and more opportunity. It advises the President, the chancellors, and the Board on the University’s approach to artificial intelligence. This item is submitted by the Vice President of Digital Innovation and Technology, who leads AI for the University. He will report on the Committee’s work, as well as on the broader portfolio of AI initiatives underway across the system.

This item discusses:

1. What has changed in the AI landscape since the Committee began its work, including the July 2026 security incident now at the top of UC’s cybersecurity risk register and the rising cost of AI at scale, with an emphasis on the University’s equity commitment;
2. The Committee’s five Working Groups and the five-campus conference series beginning in December 2026;
3. How the University is managing AI risk, from formal AI Principles to AI provisions written directly into provider contracts;
4. Four systemwide initiatives: three that control the cost, risk, and pace of AI infrastructure, and a fourth that closes the AI literacy and skills gap across UC’s workforce; and

5. The work of two Committee members -- Professor Giovanni Vigna, who builds AI that defends against the kind of cybersecurity incident described above, and Professor Rita Riley, who is preparing students to live and communicate in a world increasingly influenced by AI.

BACKGROUND

President Milliken formed the Artificial Intelligence (AI) Steering Committee in March 2026 with a charge that was about the University's future and its public responsibility. Artificial intelligence, he told the Committee at its inaugural meeting, will reshape the University longer and more deeply than any political moment. Given UC's range of disciplines, its research standing, and its obligations as a public trust, no institution is better equipped to shape how society answers the challenges and opportunities of AI. That work cannot be left to government or the private sector alone. Rather, it falls to civil society, and to a place like UC. The Committee's founding Charter states its goal in one phrase: to ensure AI serves and uplifts humanity. The Committee has organized its work around that goal of *AI for Humans and Humanity*.

The Committee is co-chaired by Chancellor Pradeep Khosla (UC San Diego), Distinguished Professor and former Academic Senate Chair Ahmet Palazoglu (UC Davis), and Vice President Van Williams (UC System). Vice President Williams leads digital innovation and technology for the system; his office is responsible for communicating and integrating the Committee's advice into the University's operations. The systemwide AI partnerships and initiatives reported in Section 4 are within that portfolio. The Committee has five Working Groups, detailed in Section 2, that carry out the priorities set in its founding Charter.

The Committee held its inaugural meeting on June 8, 2026, convening leaders from UC's campuses, health system, and national laboratories under the banner "AI for Humans and Humanity." Building on that discussion, the Committee developed an implementation strategy that was shared with the chancellors on August 10, 2026. It consolidates nine opportunity areas into five time-bound working groups and a five-campus conference series running from December 2026 through October 2027.

1. The Current AI Landscape and Strategic Context for UC

The landscape has shifted since the last AI update to the Board in September 2025. UC's AI strategy is influenced by five issues: a security incident that changed how the Committee thinks about frontier-model risk, the rising cost of AI at scale, the collision between AI infrastructure and the communities asked to host data centers, regulatory fragmentation, and AI equity.

A frontier laboratory's own models breached a partner company. No person directed it. In July 2026, OpenAI disclosed that pre-release models being tested for offensive cyber capability had been given a "sandbox" evaluation environment and a set of permitted tools. The models found an undisclosed vulnerability in one of those tools, used it to reach the open internet, reasoned that a partner company, Hugging Face, probably hosted the benchmark's answer key, and broke into Hugging Face's production systems to retrieve it. The AI undertook thousands of

actions, with self-migrating workloads, and no human operator anywhere in the chain. Both companies disclosed the incident and are building new testing controls. The incident moved AI to the top of UC's cybersecurity risk register. It is a documented case of the agentic risk that the Charter's Responsible AI Standard priority was written to anticipate, and it has shaped the vendor-access work in Section 4, which comes down to a single discipline: do not put access, data, or trust in any one provider's hands.

AI is getting cheaper per token but more expensive overall. A "token" is small unit of text that an AI model processes. The price per million tokens for a similarly capable model has fallen roughly 200-fold since 2022, from about \$20 to about \$0.10. Most enterprises still spent more than they budgeted because agentic workflows use five to thirty times the tokens of a simple chat exchange, and usage is growing faster than prices are falling. Morgan Stanley projects roughly \$2.9 trillion in global data center construction through 2028, with demand for computing capacity, or "compute" running ahead of supply. The organizations managing this have adopted one consistent strategy: they route work across competing providers rather than using just one. That strategy is the logic behind the systemwide initiatives in Section 4.

AI's appetite for power and water affects real communities. Global data center electricity use is projected to roughly double by 2030, to about 945 terawatt-hours, more electricity than Japan currently uses in a year. AI-related data center water use is projected to rise roughly elevenfold by 2028, on the order of four hundred thousand Olympic swimming pools of water annually. The strain is no longer abstract. xAI's Colossus facility in Memphis has drawn a Clean Air Act complaint over unpermitted gas turbines; commissioners in Hood County, Texas, rejected a proposed data center construction moratorium in February 2026 driven by water use concerns; and one Mid-Atlantic utility bill has climbed an average of \$21 a month since June 2025, roughly \$250 a year for a single household, with regulators tying part of the increase to data-center demand. UC's own Sovereign AI Cloud, explored in Section 4, would face the same siting, power, and water questions. It is better to plan for them now than answer for them later.

Regulation has fragmented rather than converged. More than 100 state-level AI laws are now in effect, contrary to a federal posture that has favored deregulation and competitiveness over new federal regulations, leaving institutions like UC to build their own governance structures instead of waiting for an external standard. Higher education has not kept pace: only about one in five institutions had a formal AI policy as of 2025, and just 9 percent believe their cybersecurity and privacy policies adequately address AI risk.

Equity runs through all of it. Chancellor Khosla summarized the stakes in a sentence at the Committee's inaugural meeting: AI will be the biggest driver of inequity in society if UC allows it. The technology is coming regardless; it is the inequity that has to be stopped. Rising costs and vendor concentration are not neutral. Left unmanaged, they favor the campuses, laboratories, and individuals who can pay for premium access and absorb security exposure, leaving everyone else behind. Data center siting raises the same question. Facilities like Colossus tend to be located in lower-income neighborhoods already carrying an industrial burden, a pattern UC would need to guard against in siting any facility of its own. President Milliken's founding Charter places democratizing opportunity and access at the center of UC's AI strategy, and members of the Board have publicly pressed the same point. This principle is fundamental to UC's AI strategy:

the AI Steering Committee’s campus conference series is built to include all UC campuses, and the Section 4 initiatives are about shared infrastructure, not premium access for whoever pays first.

Taken together, these five changes raise the cost of moving slowly: security exposure, price volatility, vendor dependence, community friction over infrastructure, and unsettled law. They also clarify who should move. UC’s scale, research depth, and public mission make it one of the few institutions anywhere, public or private, positioned to get governed, transparent, and equitable AI adoption right, and then to show others how to do it.

The same period produced real progress, much of it with UC’s fingerprints. In November 2025, UCLA’s Terence Tao, a Fields Medalist, worked with Google DeepMind’s AlphaEvolve, an AI system that evolves computer code rather than guessing at answers, to search sixty-seven open problems in mathematical analysis, combinatorics, and geometry. The results included a new construction for the decades-old “moving sofa” problem and an exact solution to the Gagliardo–Nirenberg inequality. Tao has been candid about its limits: the tool could not crack the harder conjectures, and every result still needed a mathematician’s verification. Yet the AI system still found things researchers had missed. A paper on the voice restored to Ann, described at the beginning of this item, was published in April 2025 by the UCSF and UC Berkeley team of Dr. Edward Chang and Gopala Anumanchipalli. A brain-computer interface now decodes her attempted speech into audible, naturalistic-sounding words in about one second, down from roughly eight seconds in its 2023 version. Lawrence Berkeley National Laboratory, which UC manages for the U.S. Department of Energy, is turning AI-predicted materials into candidate batteries and clean-energy technology. Its A-Lab facility pairs AI-driven materials prediction with autonomous robotic synthesis, part of a broader wave of AI-assisted materials science that has identified millions of candidate compounds fast enough to actually test them. None of this negates the risks described above. It is evidence that the same technology producing those risks, in university hands, is doing what the mission promises.

The stakes for UC getting AI right are high. The result will be a greater abundance of education, health, food, medicine, and energy, and a potential shift in public perception of higher education.

2. AI Steering Committee Update

The Committee’s implementation strategy organized nine Opportunity Areas into five Working Groups, each sponsored by a Committee member, with 90-day deliverables and 12-month key performance indicators (KPIs). Each Working Group addresses a piece of the mission: more health, more opportunity, more knowledge, the public good, and the trust that makes the rest possible.

Working Group	Focus
AI for Health and Longevity	Clinical care, biomedical research, and data
AI for Learning and Work	Curriculum, AI literacy, and access to higher education

AI for Discovery and Innovation	Cross-disciplinary research, compute infrastructure, and technology transfer
AI for Public Good	Public benefit applications in partnership with the State
AI to Protect Humans and Humanity	Ethics, trust, labor-market shifts, and compliance

The Working Groups will share their work public in a five-campus conference series that brings UC faculty research to the University’s broader academic community, as well as to private sector and government stakeholders. The series pairs host campuses by design so that it reaches all ten UC campuses and other locations, rather than concentrating on only one or two research centers. This is the equity commitment of Section 1:

Date	Host Campus(es)	Working Group Focus	Status
Dec. 2026	UC San Francisco & UC Berkeley	AI for Health and Longevity	Next up
Feb. 2027	UC San Diego & UC Riverside	AI for Learning and Work	Upcoming
Apr. 2027	UC Santa Barbara & UC Santa Cruz	AI to Protect Humans and Humanity	Upcoming
Jun. 2027	UC Los Angeles & UC Irvine	AI for Discovery and Innovation	Upcoming
Oct. 2027	UC Davis & UC Merced, Systemwide Finale	AI for Public Good	Upcoming

All dates reflect the plan in the Committee’s August 2026 implementation strategy. The Working Groups are newly formed and are in their first 90-day cycle, so it is too early for KPI results. Committee staff will bring Working Group progress and outcomes from the UC San Francisco event to the Board in a future update.

3. Managing the Risk Ahead of the Scale

Governance at UC has been ahead of AI adoption, not behind it. The University was among the first in the country to adopt a formal set of AI Principles, and it is the first university to address artificial intelligence explicitly in a Digital Risk Appetite statement, a formal declaration of how much technology risk the institution will accept, and where. Those commitments now are reflected in the University’s contracts: AI provisions are written directly into UC’s terms and conditions with its providers, so protections are part of every agreement systemwide, instead of depending on each campus to negotiate them individually. The University has also reviewed its electronic communications policies and its data security agreements to confirm they provide protection in an environment where AI is a standard tool, for UC and for its vendors alike. The initiatives in Section 4 build on that foundation. None of them asks the University to trade safety for speed.

4. Current Systemwide AI Initiatives

Section 1 points to one conclusion: UC should not depend on a single external AI provider, and campuses should not be left to negotiate access and absorb risk on their own. Four systemwide initiatives attempt to support that conclusion. Led through the Office of the President and shaped by the Committee's principles, these efforts are part of a wider AI portfolio that predates the Committee. UC already builds, funds, and secures its AI capability in partnership with industry and government, because no university, however large, can develop and secure frontier AI alone. One initiative uses competition among providers to control cost and risk. A second reduces dependence on any single provider at all. A third partners with commercial cloud providers for governed, ready-to-scale capacity. A fourth ensures that the University's employees can use it.

Amazon has been a systemwide partner since 2021. In January 2023, UC and Amazon Web Services (AWS) signed a five-year agreement now supporting cloud infrastructure, research, and workforce development at every campus and laboratory. That partnership has funded more than \$22 million in Amazon research awards and delivered more than \$23 million in combined cash and cloud credits, including \$14 million in AWS credits toward systemwide cloud migration. It has helped place more than 1,400 UC students and postdoctoral fellows into technical roles at Amazon, and backed dedicated facilities including UCLA's Science Hub for Humanity and AI. Amazon has also named UC as one of the first participants in its national Ph.D. and AI Fellows program, a \$68 million commitment across nine universities announced in November 2025.

Google's support is of a similar scale and reaches further across the system. UC-affiliated researchers and centers have received more than \$32.7 million across 61 funded projects, spanning AI safety and trust research, frontier research, workforce development, and public impact applications at campuses from UC Berkeley, the largest recipient, to UC Merced. Terence Tao's work with AlphaEvolve, described in Section 1, is one example. UC Berkeley's BAIR Commons and a systemwide Consortium of Cybersecurity Clinics are others.

Government partnership matters just as much. In July 2026, the U.S. Department of Energy (DOE) selected UC investigators to lead 40 of the 278 projects funded nationally under its new Genesis Mission, an initiative to accelerate scientific discovery through AI, more projects than at any other university system. Seven UC campuses participate directly, alongside the three national laboratories UC manages for DOE, Lawrence Berkeley, Lawrence Livermore, and Los Alamos. UC has separately committed \$18 million to its own multicampus AI Science at Scale initiative and \$8 million to its Initiative for Fusion Energy, both designed to work alongside the DOE partnership rather than to duplicate it.

Systemwide access to frontier AI tools on competitive terms. The University is exploring a coordinated, systemwide approach to frontier AI developers, including Anthropic and OpenAI, for research, teaching, and administrative use, replacing the lab-by-lab, campus-by-campus negotiation that is the default today. The approach follows principles the Committee's members set at its inaugural meeting: any systemwide arrangement should be non-exclusive, interoperable across providers, and protective of UC's data portability. Keeping more than one provider in active competition for UC's business is more than a negotiating tactic. It controls cost, and it means the University is never dependent on one vendor's security posture, which is the lesson of

the Hugging Face incident in Section 1. This work connects to the Charter's Strategic Economic Partnerships and Responsible AI Standard priorities. No provider has been selected. Any resulting arrangement would go through UC procurement with appropriate privacy and security review, and Academic Senate consultation, before the University commits. Vice President Williams will make a recommendation, including provider terms and cost projections, to the President no later than December 2027.

Investment in a UC Sovereign AI Cloud. The University is also scoping a shared, UC-owned and UC-governed computing capacity ("compute") under the Charter's Infrastructure and Compute priority: a Sovereign AI Cloud. The goal is a baseline of compute that is directly controlled by the University, instead of depending entirely on outside providers for capacity, pricing, and security. The case for this is the same as Section 1 makes: campus-by-campus compute purchases are among the University's least efficient spending, external pricing has grown volatile even as per-token prices fall, and students and researchers at less-resourced campuses report the most unmet compute need. Early work is focused on three goals: inventorying existing compute across campuses and national laboratories, building an allocation model in which access does not depend on a campus's or laboratory's ability to pay, and testing a two-track approach for baseline needs versus specialized work. Baseline needs would run on UC-owned open source and open-weight infrastructure, meaning AI systems whose underlying code or trained parameters are published openly, so UC can run and adapt them directly rather than depending on an outside vendor. Work that requires the most capable models would use competitively sourced frontier access. Sizing, funding, and governance are still being defined, given the likely capital investment. Vice President Williams will recommend a sizing and funding proposal to the President no later than October 2027.

A Systemwide AI Public Cloud, partnered rather than built alone. Between commercial frontier model access and a UC-owned Sovereign AI Cloud approaches is a third option: a systemwide partnership with established cloud providers, covering both enterprise operations and research computing, rather than each campus negotiating and configuring its own. The logic matches the other two initiatives in this section: aggregated demand controls cost across frontier AI providers more effectively than any campus can manage alone; governance and security are enforced at the system level, instead of campus by campus; and researchers and administrators get a mature, already-scaled platform to prototype and scale on immediately, while the Sovereign AI Cloud is being built. Provider terms, sizing, and timing are still being defined.

A Common AI learning program will teach the University to use what it builds. None of the tools above matter if UC's own faculty, students, and employees do not know how to use them. A physician, a grants administrator, and a graduate student need different things from artificial intelligence, and one course taught the same way for all three teaches none of them well. Under the Charter's Defining the National Workforce priority, UC is scoping a systemwide learning platform: a single shared library of AI content with role-specific pathways of curated courses, so a researcher sees research tools, a financial analyst sees financial tools, and a department chair sees what a department chair actually needs. The goal is fluency, not a certificate: enough working knowledge, at every level of the University, so that AI becomes a tool people use well rather than a technology that happens to them. It would complement, rather than duplicate, the student-centered literacy work in Section 6: one prepares the students UC teaches, the other

prepares the people who run the University itself. Platform, ownership, and timing are still being defined.

A utility, once built, does not ask who is drawing power from it. Frontier access negotiated on competitive terms, a Systemwide AI Public Cloud, a Sovereign AI Cloud the University owns outright, and a learning platform built once and tuned to whoever is using it -- each of the four initiatives above was funded to solve a problem of cost, security, or workforce readiness. None of them was funded to reach the returning student studying between shifts or the technician six months from a credential. But the wires do not know the difference. If we build the capacity, the University can decide later how it is used. A flexible, learner-centered path for non-traditional and working adult learners is one mission that capacity could support. Curricula remains in the Academic Senate's purview; the capacity undergirding it is the responsibility of the systemwide office of Digital Innovation and Technology, and that building has already begun.

5. AI Research and Innovation Highlight: Professor Giovanni Vigna, Director, ACTION Institute

Professor Giovanni Vigna's team has entered more DEF CON hacking competitions than any other in the contest's history. That record belongs to Shellphish, the Capture the Flag team he founded based at the University of California, which has run for more than two decades. Vigna's research runs on the same instinct: the best way to defend a system is to think like the person trying to break it, and increasingly that person is an AI agent acting on its own. His work builds AI systems that find and patch the kind of software vulnerability described in Section 1 before an attacker, human or automated, finds it first.

Vigna directs ACTION, the National Science Foundation's AI Institute for Agent-based Cyber Threat Intelligence and Operation, a multi-campus AI institute led by UC Santa Barbara with UC Berkeley as a partner campus, built specifically around AI agents' offensive and defensive use in cybersecurity, the same category of risk the Committee flagged in Section 1. He also directs UC Santa Barbara's Center for CyberSecurity and is a Fellow of both the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM).

Shellphish's AI-based system, ARTIPHISHELL, won the Small Business Track of the federal Defense Advanced Research Projects Agency's (DARPA) AI Cyber Challenge in March 2024 and advanced to the competition's finals in August 2025 as one of seven systems DARPA is open sourcing for the security research community. It builds on the same automated vulnerability-discovery research Vigna commercialized earlier at Lastline, the malware-detection company he co-founded. Lastline was acquired by VMware in 2020 and is now part of Broadcom. Competition to company to national institute -- that is the path UC's AI to Protect Humans and Humanity agenda means to widen. It is what more safety looks like.

6. Academic Impact Highlight: Professor Rita Raley, UC Santa Barbara

Professor Rita Raley, a member of the AI Steering Committee and its AI for Learning and Work Working Group, is Professor of English at UC Santa Barbara, with courtesy appointments in Film and Media Studies, Comparative Literature, and Global Studies.

Her attention to artificial intelligence grew out of two decades of research on digital media and its politics, and it sharpened toward AI systems in 2020, when she published on the ethics of AI-generated text, years before large language models entered daily use. That work led her, with UC Santa Barbara colleague Fabian Offert, to found the Critical Machine Learning Studies faculty working group in 2021 and UC Santa Barbara's Center for the Humanities and Machine Learning, which she co-directs.

That scholarship underlies the creation of her course. ENGL 155AI, Critical Artificial Intelligence, which was offered at UCSB a year before the Committee existed. This undergraduate course trains students to evaluate AI systems, not just use them, pairing the technical foundations of machine learning with its social, political, and ethical stakes. Raley and Offert have put the goal plainly: the classroom, they argue, "has to become something like a counterweight to popular (mis)information about artificial intelligence."

The Committee reached the same conclusion independently at its June 2026 inaugural meeting, when one breakout group identified a shared critical AI foundation course as UC's early move on AI literacy. That recommendation became the founding charge of the Committee's Common AI Literacy Working Group. The group inherits an example rather than a blank page: Raley's course and the four-campus Critical Machine Learning Studies network she co-leads with faculty at UC San Diego, UC Irvine, UC Berkeley, and UCLA. What her students are learning, and what follows for the curricula systemwide, is what more opportunity looks like.

The Year Ahead

Over the next twelve months, the campus conference series will begin at UC San Francisco in December, and the Working Groups will report their first 90-day results early in the new year. A sizing and funding proposal for the Sovereign AI Cloud will be submitted to the President by October 2027; a frontier-access recommendation with provider terms and cost projections, follows by December 2027; and the systemwide learning platform for AI literacy and workforce upskilling is being scoped.

It is remarkable how far and how fast AI advances have developed. Ann waited nearly two decades for her voice. It came back to her not from the general drift of technology but from two UC campuses that spent years deciding, deliberately and in public view, the purpose of this technology. The Committee's work extends that same decision across the University: to prove that an institution this large can put the most powerful technology of its era to serve people without losing its soul, and to make what happened for one patient in San Francisco the ordinary business of ten campuses, three national laboratories, and the state they serve.

SOURCES

- OpenAI, "OpenAI and Hugging Face partner to address security incident during model evaluation," July 2026, openai.com
- Hugging Face, security incident disclosure, July 2026, huggingface.co/blog

- TechCrunch, “OpenAI says Hugging Face was breached by its pre-release models,” July 2026, techcrunch.com
- Morgan Stanley, AI Market Trends Institute, 2026, morganstanley.com
- International Energy Agency, “Energy and AI,” 2025, iea.org
- Inside Climate News, reporting on xAI’s Colossus facility, Memphis, 2025, insideclimatenews.org
- Data Center Dynamics, reporting on xAI gas turbine permitting, Memphis, 2025, datacenterdynamics.com
- Texas Tribune, “Hood County rejects data center development pause,” February 2026, texastribune.org
- Institute for Energy Economics and Financial Analysis (IEEFA), PJM capacity price and ratepayer impact analysis, 2026, ieefa.org
- Terence Tao, “Mathematical Exploration and Discovery at Scale,” terrytao.wordpress.com, Nov. 5, 2025
- Google DeepMind, AlphaEvolve mathematical discoveries, Nov. 2025, deepmind.google
- Nature Neuroscience, UCSF/UC Berkeley speech neuroprosthesis study, April 2025, nature.com
- University of California, news release on brain-voice neuroprosthesis, 2025, universityofcalifornia.edu
- DeepMind, “Millions of new materials discovered with deep learning,” Nature, 2023, deepmind.google
- Lawrence Berkeley National Laboratory, A-Lab launch announcement, 2023, newscenter.lbl.gov
- American Action Forum, “The Next Phase of AI: Technology, Infrastructure, and Policy in 2025–2026”, americanactionforum.org
- EdTech Magazine, “AI Governance in Higher Education: An Overview,” 2026, edtechmagazine.com
- Stanford HAI, AI Index Report, 2025, hai.stanford.edu
- OpenAI, API pricing, 2026, developers.openai.com
- Industry AI-pricing and enterprise cost-trend reporting, 2026 (aggregated)
- AWS, University of California Impact Statement, 2021–2026
- Amazon, PhD Fellows and AI Fellows program announcement, Nov. 2025, amazon.science
- Google grant records compiled for the Committee, 2026
- University of California, “University of California partners with U.S. Department of Energy to advance energy, discovery and national security through Genesis Mission,” July 2026, universityofcalifornia.edu
- U.S. Department of Energy, “Energy Department Announces Collaboration Agreements with 24 Organizations to Advance the Genesis Mission,” July 2026, energy.gov
- UC Santa Barbara Computer Science, Giovanni Vigna faculty profile, cs.ucsb.edu
- University of California, “UC Santa Barbara to lead NSF-funded research institute for next-level, AI-powered cybersecurity,” 2023, universityofcalifornia.edu
- DARPA, AI Cyber Challenge (AIxCC) results, 2025, darpa.mil
- UC Santa Barbara College of Engineering, Shellphish AIxCC finals qualification, Aug. 2024, engineering.ucsb.edu
- UC Santa Barbara Department of English, Rita Raley faculty profile, english.ucsb.edu

- UC Humanities Research Institute, “Critical Machine Learning Studies” faculty working group award, uchri.org
- UC Humanities Research Institute, interview with Rita Raley and Fabian Offert, Critical Machine Learning Studies, uchri.org/foundry
- UC Santa Barbara Center for the Humanities and Machine Learning, huml.ucsb.edu
- Rita Raley, “Playing With Unicorns: AI Dungeon and Citizen NLP,” Digital Humanities Quarterly, 2020