12th International Conference on Computational Social Science IC^{2}S^2
July 28-31, 2026, Burlington, VT, United States

How Are U.S. Universities Responding to AI? An Audit of Governance Capacity

AI, university, policy, governance, procedural justice

Extended Abstract

Artificial intelligence (AI) is increasingly embedded into universities, shaping teaching, assessment, research, and administration [1], [2], [3], [4], [5], [6], [7], [8], [9], [10]. These deployments raise concerns about fairness, accountability, and harm, yet there is limited empirical evidence about how universities govern AI use in practice. Existing evaluations focus largely on technical systems [11], individual university responses [12], or national capacity (Table 1), *leaving a critical gap in understanding the institutional infrastructures through which AI use is authorized, constrained, contested, revised, and supported within higher education.* Moreover, U.S.-focused studies that systematically examine institutional AI policies have largely centered on top-ranked or R1 universities [13], [14], [15], obscuring variation across the broader higher-education landscape.

Bridging this gap is fundamentally a question of accountability. Universities exercise significant power over students and faculty, and AI-related governance decisions – such as the use of AI detection tools in academic integrity enforcement [16], [17], [18] (Figure 1), faculty discretion over permissible AI use, or access to appeals – can have material consequences for equity and due process.

We present **ACAI-US79**, an audit dataset capturing publicly articulated AI governance practices across 79 diverse U.S. universities (Figure 2) – ranging in *research activity* (R1/R2/unranked), *region*, *size*, and *structure* (public/private/teaching-liberal arts). From this dataset, we derive **ACAI**, the **Ac**ademic **A**I Capacity **I**ndex, an interpretable index that aggregates audit findings to produce a structured ranking of institutions based on the public legibility of formal policies, resources, and oversight mechanisms. Shown in Figure 2, this ranking reflects differences in what governance artifacts are publicly visible under a consistent audit protocol. To operationalize this audit, we focus on publicly legible institutional signals, and we assess AI governance capacity across four governance domains (Figure 3): A. Policy Clarity (presence of formal guidelines), B. Faculty Support (resources for pedagogy), C. Feedback Loops (mechanisms for community input and revision), and D. Detection Tools (stated rules around academic integrity software). Higher scores in any domain indicate greater public legibility and clearer formalization of that mechanism. Annotators[1] conduct time-bounded reviews of institutionally authoritative materials to evaluate whether relevant governance mechanisms are present, partially specified, absent, or conflicting. These assessments are aggregated into ACAI, yielding a transparent, diagnostic index that enables cross-institutional comparison without conflating governance capacity with technical expertise. We make the following contributions:

[1]Annotations (11 items shown in Figure 3 x 79 universities x 3 annotators per university) were completed by paid U.S.-based Prolific workers under standardized instructions, time limits, and quality controls; each university was labeled by three annotators using only public materials. Expert (authors') annotation of 10 universities shows moderate–strong agreement (Pearson $r = 0.56$, Spearman $\rho = 0.57$, Kendall $\tau = 0.49$, $p < 0.05$) with with Prolific annotators rating schools with slightly higher governance capacity scores than experts.

(1) An Open Dataset, Toolkit, and Website: We release ACAI-US79, a publicly available dataset for auditing AI governance across 79 U.S. universities, along with a reproducible annotation schema and audit toolkit (available at `https://anonymous.4open.science/r/ACAI-3D27`), and public website (Figure 4, available at `https://acai-us79.org/`).
(2) The Academic AI Capacity Index (ACAI): We introduce ACAI (Figure 2), an interpretable index for evaluating the public legibility of institutional AI governance, grounded in principles of accountability and procedural justice.
(3) An Empirical Institutional Audit: We apply ACAI in the first large-scale, reproducible audit of AI governance in U.S. higher education. Moving beyond theoretical frameworks, this audit provides the first empirical baseline of how diverse institutions actually operationalize AI oversight, yielding key findings:
▷ ***F1: AI governance capacity varies with organizational structure (Figure 6).*** Public research universities tend to occupy higher positions in the ACAI distribution than private research or teaching/liberal arts universities; larger institutions tend to outrank small ones; R1 schools tend to outrank than R2 and unclassified schools; and schools in the Western region tend to show slightly higher ACAI scores. These differences align with known variation in institutional oversight arrangements and coordination demands.
▷ ***F2: AI governance capacity is concentrated in policy articulation rather than participation or process (Figure 6).*** Scores in A. POLICY CLARITY and B. FACULTY SUPPORT are systematically higher than those in C. FEEDBACK LOOPS and D. DETECTION TOOLS.
▷ ***F3: AI governance capacity is largely uncorrelated with AI research output (Figure 5, Figure 8).*** AI governance capacity (ACAI) does not strongly correlate with AI-specific research output (as proxied through CSRankings.org); this dispersion indicates that AI technical leadership alone does not reliably translate into strong, visible governance practices.
▷ ***F4: ACAI rankings are robust to weighting choices (Table 2) and individual annotators.*** Inter-annotator agreement is modest at the item level ($\alpha = 0.26$; mean pairwise $= 0.48$), reflecting interpretive variation, but improves at the domain level ($\alpha = 0.30$). Leave-one-annotator-out recomputation yields strongly correlated scores ($r = 0.83$–0.87, $\rho = 0.75$–0.87), indicating robustness to individual annotators.
▷ ***F5: LLMs only partially reproduce human governance judgments (Figure 7).*** To evaluate whether automated systems can reliably scale governance audits in place of human labor, we compared human judgments against LLM proxies. Across temperatures, LLM-generated ACAI rankings exhibit moderate ordinal agreement with human judgments ($\rho = 0.52 - 0.60$), yet individual institutions are frequently misranked, with mean absolute rank errors of approximately 15-17 positions.
(4) Actionable Recommendations for Universities: We translate our findings into concrete recommendations for strengthening institutional capacity for accountable AI governance. Specifically, institutions must (R1) establish a centralized, institution-level AI governance reference point (following ▷ F1, F3); (R2) provide more explicit procedural clarity around academic integrity and AI detection tools (following ▷ F1, F2); and (R3) formalize feedback and revision mechanisms to adapt to rapid technological shifts (following ▷ F2).

By focusing on institutional governance, this work advances organizational accountability. Institutional audits make visible the structures through which power is exercised, providing an empirical foundation for more transparent and procedurally just AI governance in higher education.

References

[1] K. Huang, "Alarmed by ai chatbots, universities start revamping how they teach," *The New York Times*, vol. 16, e12255, 2023.

[2] J. D. Walsh, "Everyone is cheating their way through college: Chatgpt has unraveled the entire academic project," *Intelligencer (New York Magazine)*, May 2025. [Online]. Available: https://nymag.com/intelligencer/article/openai-chatgpt-ai-cheating-education-college-students-school.html.

[3] M. Fritts, "A matter of words: What can university ai committees actually do?" *The Point Magazine – Examined Life*, May 2025, Accessed: 2025-10-01. [Online]. Available: https://thepointmag.com/examined-life/a-matter-of-words/.

[4] C. Baek, T. Tate, and M. Warschauer, "“chatgpt seems too good to be true”: College students' use and perceptions of generative ai," *Computers and Education: Artificial Intelligence*, vol. 7, p. 100 294, 2024.

[5] C. K. Y. Chan, "A comprehensive ai policy education framework for university teaching and learning," *International journal of educational technology in higher education*, vol. 20, no. 1, p. 38, 2023.

[6] *Teaching the arts in the age of a.i.: A symposium*, Jan. 2025. [Online]. Available: https://arts.umd.edu/teaching-arts-in-the-age-of-ai-symposium.

[7] *EAAI-25: The 15th Symposium on Educational Advances in Artificial Intelligence*, Philadelphia, PA, USA: Association for the Advancement of Artificial Intelligence (AAAI), Mar. 2025. [Online]. Available: https://eaai-conf.github.io/year/eaai-25.html.

[8] U.S. Department of Education, *U.s. department of education announces release of $169 million under the fund for the improvement of postsecondary education*, https://www.ed.gov/about/news/press-release/us-department-of-education-announces-release-of-169-million-under-fund-improvement-of-postsecondary-education, Jan. 2026.

[9] J. Priest, "Texas universities deploy ai tools to review and rewrite how some courses discuss race and gender," *The Texas Tribune*, Dec. 2025. [Online]. Available: https://www.texastribune.org/2025/12/15/texas-universities-ai-course-audits/.

[10] P. Ipeirotis. "Fighting fire with fire: Scalable personalized oral exams with an elevenlabs voice ai agent. "[Online]. Available: https://www.behind-the-enemy-lines.com/2025/12/fighting-fire-with-fire-scalable-oral.html.

[11] L. Fattorini, N. Maslej, R. Perrault, V. Parli, J. Etchemendy, Y. Shoham, and K. Ligett, "The global ai vibrancy tool," AI Index Project, Institute for Human-Centered AI, Stanford University, Tech. Rep., Nov. 2024, November 2024 edition. [Online]. Available: https://hai.stanford.edu/assets/files/global_ai_vibrancy_tool_paper_november2024.pdf.

[12] O. V. Spivakovsky, S. A. Omelchuk, V. V. Kobets, N. V. Valko, and D. S. Malchykova, "Institutional policies on artificial intelligence in university learning, teaching and research," *Information Technologies and Learning Tools*, vol. 97, no. 5, p. 181, 2023.

[13] C. Wu, H. Zhang, and J. M. Carroll, *Ai governance in higher education: Case studies of guidance at big ten universities*, 2024. arXiv: 2409.02017 [cs.HC]. [Online]. Available: https://arxiv.org/abs/2409.02017.

[14] H. Wang, A. Dang, Z. Wu, and S. Mac, "Generative ai in higher education: Seeing chatgpt through universities' policies, resources, and guidelines," *Computers and Education: Artificial Intelligence*, vol. 7, p. 100 326, 2024.

[15] N. McDonald, A. Johri, A. Ali, and A. H. Collier, "Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines," *Computers in Human Behavior: Artificial Humans*, vol. 3, p. 100 121, 2025.

[16] Turnitin, *Understanding false positives within our ai writing detection capabilities*, https://www.turnitin.com/blog/understanding-false-positives-within-our-ai-writing-detection-capabilities, 2023.

[17] W. Liang, M. Yuksekgonul, Y. Mao, E. Wu, and J. Zou, "Gpt detectors are biased against non-native english writers," *Patterns*, vol. 4, no. 7, p. 100 779, 2023, ISSN: 2666-3899. DOI: https://doi.org/10.1016/j.patter.2023.100779. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666389923001307.

[18] J. Wu, S. Yang, R. Zhan, Y. Yuan, L. S. Chao, and D. F. Wong, "A survey on llm-generated text detection: Necessity, methods, and future directions," *Computational Linguistics*, vol. 51, no. 1, pp. 275–338, 2025.

[19] M. Mäntymäki, M. Minkkinen, T. Birkstedt, and M. Viljanen, *Putting ai ethics into practice: The hourglass model of organizational ai governance*, 2023. arXiv: 2206.00335 [cs.AI]. [Online]. Available: https://arxiv.org/abs/2206.00335.

[20] National Institute of Standards and Technology, "Artificial intelligence risk management framework (ai rmf 1.0)," National Institute of Standards and Technology (NIST), NIST Special Publication, Jan. 2023. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

[21] D. W. Ball and K. Ramakrishnan, "Entity-based regulation in frontier ai governance," *Carnegie Endowment for International Peace*, 2025, Published July 7, 2025. [Online]. Available: https://carnegieendowment.org/research/2025/06/artificial-intelligence-regulation-united-states?lang=en.

[22] J. Zhang and B. Paal, "A three-layered framework: An ai governance guide for global policymakers," 2025. [Online]. Available: https://papers.ssrn.com/abstract=5241351.

[23] United Nations System Chief Executives Board for Coordination, "Un system white paper on ai governance," UN System Chief Executives Board for Coordination, Tech. Rep., 2024, White Paper, November 2024. [Online]. Available: https://unsceb.org/sites/default/files/2024-11/UNSystemWhitePaperAIGovernance.pdf.

[24] A. Herrera-Poyatos, J. D. Ser, M. L. de Prado, F.-Y. Wang, E. Herrera-Viedma, and F. Herrera, *Responsible artificial intelligence systems: A roadmap to society's trust through trustworthy ai, auditability, accountability, and governance*, 2025. arXiv: 2503.04739 [cs.CY]. [Online]. Available: https://arxiv.org/abs/2503.04739.

[25] U. C. for Teaching and U. o. P. Learning, *Encouraging academic integrity*, Online resource, 2025. [Online]. Available: https://teaching.pitt.edu/resources/encouraging-academic-integrity/.

[26] U.S. Census Bureau, *Census regions and divisions of the united states*, https://www2.census.gov/geo/pdfs/maps-data/maps/reference/us_regdiv.pdf, n.d.

[27] American Council on Education, *Carnegie classification of institutions of higher education®*, `https://carnegieclassifications.acenet.edu/`, 2025.

Figure 1: **Case Study on AI Detection Tools**: As an example, a false positive rate of only 1% [16], [25] in an AI detection tool indicates that 1 in every 100 cases will be flagged for cheating incorrectly. It falls on the instructor to evaluate use – faculty who are ill-equipped to assess state-of-the-art AI systems. At scale in a university setting, for every 100,000 submissions, 1000 false accusations could occur, jeopardizing student careers and significantly de-legitimizing educational institutions.

Model	Primary Focus	Capacity Flow	Scope
The Hourglass Model of Organizational AI Governance [19]	Focused on AI ethics, risk mitigation; organized into environmental, organizational, and AI system layers with some feedback mechanisms (i.e. computational → governance).	Governance → Computational	EU/Global
NIST AI Risk Management Framework [20]	Focused on identifying AI risks, which governance Maps, Measures, and Manages.	Computational → Governance	USA
Entity-Based Regulation Framework [21]	Focused on transparency and regulating *"the large business entities developing the most powerful AI models and systems"* with emphasis on preemptive risk regulation.	Governance → Computational	USA
Three-Layered Framework [22]	Focused on risk, aims to fix market failures using a toolbox of regulatory tools from the three layers: market-invigorating strategies, value-directed rules, and procedural controls.	Computational → Governance	Global
UN System Survey of Institutional Models [23]	Focused on ethics and risk; highlights *"[national] capacity-building [that] can support AI development that is grounded in fairness, gender equality, reliability, safety, interpretability and accountability."*	Computational → Governance	UN/Global
Responsible AI Systems Roadmap [24]	Focused on risk, highly dependent on a committee of scientists to shape policy.	Computational → Governance	UN/Global
AI Ecological Education Policy Framework [5]	Focused on education; organized into 3 dimensions of educational support: pedagogical, ethical, and operational.	Computational → Governance	Hong Kong

Table 1: Existing governance frameworks recognize that there is both *computational capacity* and *governance capacity*.

ACAI Rank	CSRankings$_{AI}$ Rank	Institution	Type	Research Activity	Region	Size	ACAI Score
1	105	University of New Hampshire	Public Research	R1	Northeast	Medium	81.82
2	96	Portland State University	Public Research	R2	West	Large	80.30
3	8	Stanford University	Private Research	R1	West	Medium	80.30
4	14	University of Texas at Austin	Public Research	R1	South	Large	77.27
5	42	University of Notre Dame	Private Research	R1	Midwest	Medium	75.76
6	137	Baylor University	Private Research	R1	South	Large	74.24
7	44	University at Buffalo	Public Research	R1	Northeast	Large	74.24
8	79	University of Florida	Public Research	R1	South	Large	71.21
9	9	University of Michigan at Ann Arbor	Public Research	R1	Midwest	Large	71.21
10	*un.*	Rowan University	Public Research	R2	Northeast	Large	71.21
11	29	Stony Brook University	Public Research	R1	Northeast	Large	71.21
12	*un.*	Lewis & Clark College	Teaching/Liberal Arts	–	West	Small	69.70
13	7	University of California, Berkeley	Public Research	R1	West	Large	69.70
14	27	Texas A&M University	Public Research	R1	South	Large	69.70
15	91	Case Western Reserve University	Private Research	R1	Midwest	Medium	69.70
16	*un.*	Lafayette College	Teaching/Liberal Arts	–	Northeast	Small	69.70
17	*un.*	California State University, Long Beach	Public Research	R2	West	Large	68.18
18	25	University of North Carolina at Chapel Hill	Public Research	R1	South	Large	68.18
19	6	Cornell University	Private Research	R1	Northeast	Medium	68.18
20	117	Brandeis University	Private Research	R1	Northeast	Small	68.18
21	169	Southern Methodist University	Private Research	R1	South	Medium	68.18
22	*un.*	Chapman University	Private Research	R2	West	Medium	68.18
23	*un.*	Howard University	Private Research	R1	South	Medium	68.18
24	81	University of South Florida	Public Research	R1	South	Large	66.67
25	100	Syracuse University	Private Research	R1	Northeast	Large	66.67
26	*un.*	University of Wyoming	Public Research	R1	West	Medium	65.15
27	37	The Ohio State University	Public Research	R1	Midwest	Large	65.15
28	15	University of Southern California	Private Research	R1	West	Large	65.15
29	*un.*	Mercer University 9	Private Research	R2	South	Medium	63.64
30	142	DePaul University	Private Research	R2	Midwest	Large	63.64
31	29	Arizona State University	Public Research	R1	West	Large	63.64
32	*un.*	Northern Illinois University	Public Research	R2	Midwest	Medium	63.64
33	*un.*	University of South Alabama	Public Research	R2	South	Medium	63.64
34	*un.*	Fordham University	Private Research	R2	Northeast	Medium	63.64
35	169	Florida Institute of Technology	Private Research	R2	South	Small	63.64
36	*un.*	Illinois State University	Public Research	R2	Midwest	Large	62.12
37	*un.*	Pepperdine University	Private Research	R2	West	Medium	60.61
38	46	University of Chicago	Private Research	R1	Midwest	Large	60.61
39	*un.*	Montclair State University	Public Research	R2	Northeast	Medium	60.61
40	100	Binghamton University	Public Research	R1	Northeast	Medium	60.61
41	*un.*	Lake Forest College	Teaching/Liberal Arts	–	Midwest	Small	60.61
42	79	Iowa State University	Public Research	R1	Midwest	Large	57.58
43	*un.*	Carleton College	Teaching/Liberal Arts	–	Midwest	Small	57.58
44	16	University of Washington-Seattle	Public Research	R1	West	Large	57.58
45	*un.*	San José State University	Public Research	R2	West	Large	57.58
46	*un.*	Southern University and A & M College	Public Research	R2	South	Small	56.06
47	*un.*	Colby College	Teaching/Liberal Arts	–	Northeast	Small	56.06
48	*un.*	Skidmore College	Teaching/Liberal Arts	–	Northeast	Small	56.06
49	*un.*	Saint Louis University	Private Research	R1	Midwest	Medium	56.06
50	*un.*	Reed College	Teaching/Liberal Arts	–	West	Small	56.06
51	*un.*	Southern Wesleyan University	Teaching/Liberal Arts	–	South	Small	54.55
52	*un.*	Davidson College	Teaching/Liberal Arts	–	South	Small	54.55
53	*un.*	University of Colorado Colorado Springs	Public Research	R2	West	Medium	54.55
54	2	University of Illinois Urbana-Champaign	Public Research	R1	Midwest	Large	54.55
55	142	Wichita State University	Public Research	R2	Midwest	Medium	53.03
56	*un.*	Wofford College	Teaching/Liberal Arts	–	South	Small	51.52
57	*un.*	Georgia Southern University	Public Research	R2	South	Large	51.52
58	*un.*	University of Denver	Private Research	R1	West	Medium	51.52
59	*un.*	Westminster University	Teaching/Liberal Arts	–	West	Small	50.00
60	*un.*	Ball State University	Public Research	R2	Midwest	Large	48.48
61	*un.*	Rhodes College	Teaching/Liberal Arts	–	South	Small	48.48
62	49	Brown University	Private Research	R1	Northeast	Medium	48.48
63	*un.*	Clark University	Private Research	R2	Northeast	Small	48.48
64	*un.*	Abilene Christian University	Private Research	R2	South	Small	48.48
65	*un.*	Wesleyan University	Teaching/Liberal Arts	–	Northeast	Small	46.97
66	100	Illinois Institute of Technology	Private Research	R2	Midwest	Small	46.97
67	142	Nova Southeastern University	Private Research	R1	South	Large	46.97
68	*un.*	Kean University	Public Research	R2	Northeast	Medium	45.45
69	*un.*	Occidental College	Teaching/Liberal Arts	–	West	Small	43.94
70	61	Stevens Institute of Technology	Private Research	R2	Northeast	Small	43.94
71	71	California Institute of Technology	Private Research	R1	West	Small	43.94
72	*un.*	Long Island University	Private Research	R2	Northeast	Medium	42.42
73	*un.*	Marquette University	Private Research	R2	Midwest	Medium	40.91
74	*un.*	Beloit College	Teaching/Liberal Arts	–	Midwest	Small	39.39
75	*un.*	Grinnell College	Teaching/Liberal Arts	–	Midwest	Small	34.85
76	*un.*	Jackson State University	Public Research	R2	South	Small	31.82
77	*un.*	Clark Atlanta University	Private Research	R2	South	Small	28.79
78	*un.*	Creighton University	Private Research	R2	Midwest	Medium	27.27
79	*un.*	University of Massachusetts at Dartmouth	Public Research	R2	Northeast	Small	27.27

Figure 2: **Institutional Rankings for ACAI-US79 with ACAI and CSRankings$_{AI}$**: *Region* is classified based on the U.S. Census [26]; *Type* and *Research Activity* are classified based on the Carnegie Classifications [27]; *Size* is split based on tertile buckets.

Annotation Guidelines
You will review publicly available web pages for <UNIVERSITY> to determine whether specific AI-related policy statements are addressed by the institution.

Use ONLY the links provided below and any pages, sections, PDFs, or subpages that are directly reachable by clicking links on those pages (e.g., menus, internal links, or document links). Do NOT use external search engines or sources found from outside this list: <LINKS>.

Evaluate each statement independently. Spend no more than 5 minutes per statement.

For each statement:
1. Select exactly one classification:

- *Present/Yes* — A clear statement directly addressing the item is found on an institutional page within 5 minutes.
- *Partial/Implicit/Somewhat* — The item is mentioned or implied, but key details are missing.
- *Absent/No* — You reasonably searched the allowed sources and did not find relevant content.
- *Unclear or Took Longer Than 5 Minutes* — Navigation difficulty, vague language, or time limits prevented a confident decision.
- *Conflicting Information* — Different institutional sources provide contradictory guidance for the same item.

2. Provide the most relevant URL(s) from the allowed sources that support your selection. If you selected Absent or Unclear, provide the main page(s) you checked.

Statements Organized by Governance Domain:
A. Policy Clarity — *Policies defining institutional expectations, terminology, and academic integrity adaptations.*

A1. The university defines "AI use," "AI assistance," or "AI-generated content."
A2. The university defines standards for citing AI-generated material.

B. Faculty Support — *Resources that enable faculty to integrate, regulate, or teach with AI.*

B1. The university provides guidance, training, or resources for faculty on AI-related teaching practices.
B2. Official examples of appropriate and/or prohibited AI use are provided (e.g. example AI use cases, example prompts).
B3. A faculty committee or group focused on teaching and learning about AI exists.
B4. Faculty are offered syllabus language examples (e.g. use AI/don't use AI/selectively use AI).

C. Feedback Loops — *Mechanisms through which universities gather input, revise policies, and communicate decisions.*

C1. A faculty committee or advisory group focused on university AI policy or governance exists.
C2. A student committee or advisory group focused on university AI policy or governance exists.
C3. The university publishes AI policy update logs or explains revisions.

D. Detection Tools — *Institutional stance toward AI detection technologies.*

D1. The university restricts, discourages, or warns against the use of AI detection tools.
D2. Student misconduct determinations require human review and cannot be based solely on AI detection tools.

Figure 3: Annotation Instructions for ACAI Calculation, Organized By Governance Domain.

Scheme 1	Scheme 2	Spearman	Pearson
Baseline	Equal	0.98	0.98
Baseline	Policy-heavy	0.93	0.93
Baseline	Teaching-heavy	0.99	0.99
Equal	Policy-heavy	0.97	0.97
Equal	Teaching-heavy	0.98	0.98
Policy-heavy	Teaching-heavy	0.93	0.93

Table 2: Pairwise rank correlations were highly stable across weighting schemes: We evaluated four weighting schemes: an indicator-weighted baseline (w_A=2, w_B=4, w_C=3, w_D=2), equal (w_A=1, w_B=1, w_C=1, w_D=1), policy-heavy (w_A=1, w_B=1, w_C=2, w_D=2), and teaching-heavy (w_A=1, w_B=2, w_C=1, w_D=1), using percentile ranks.

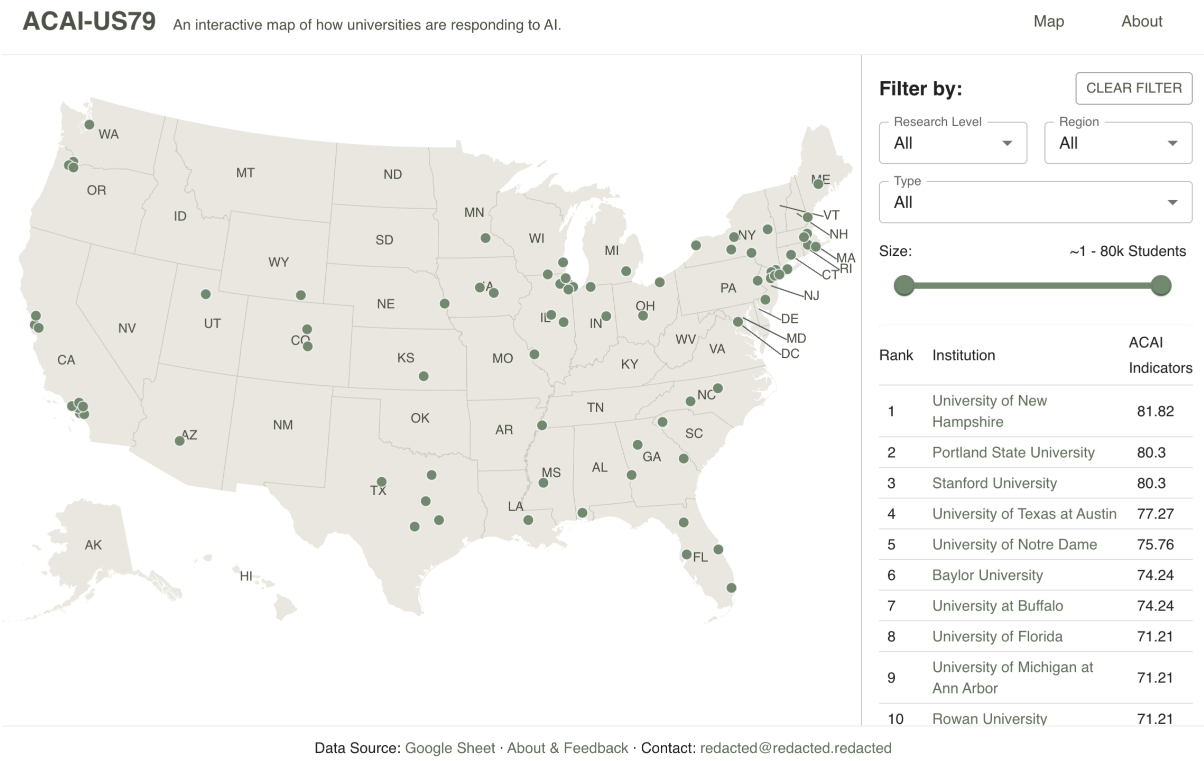

Figure 4: **Interactive map of ACAI-US79 at** `https://acai-us79.org/`, visualizing the 79 U.S. universities and describing their publicly articulated governance capacity. Institutions are shown as clickable markers and ranked by ACAI score, with filters enabling comparison across research activity, institutional type, region, and size. Selecting an institution reveals its score and links to the publicly available policies, guidance, and governance materials reviewed in the audit, supporting traceability and independent inspection of how governance capacity is publicly articulated.

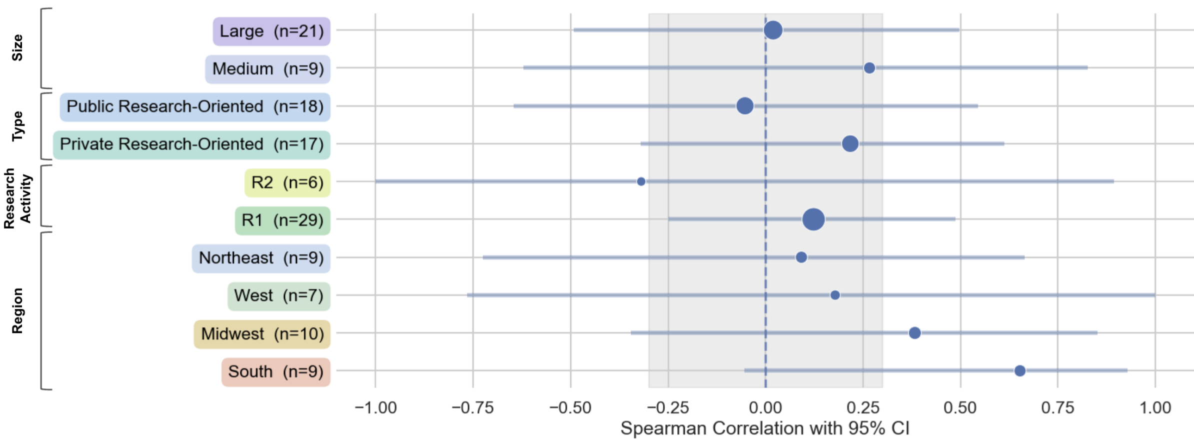

Figure 5: Spearman ρ rank correlations between ACAI ranks and CSRankings_{AI} ranks across institutional subgroups show that AI governance capacity and AI research output are almost entirely uncorrelated: All CIs cross zero, unranked CSRankings_{AI} universities are excluded, and categories with $n \leq 5$ are dropped.

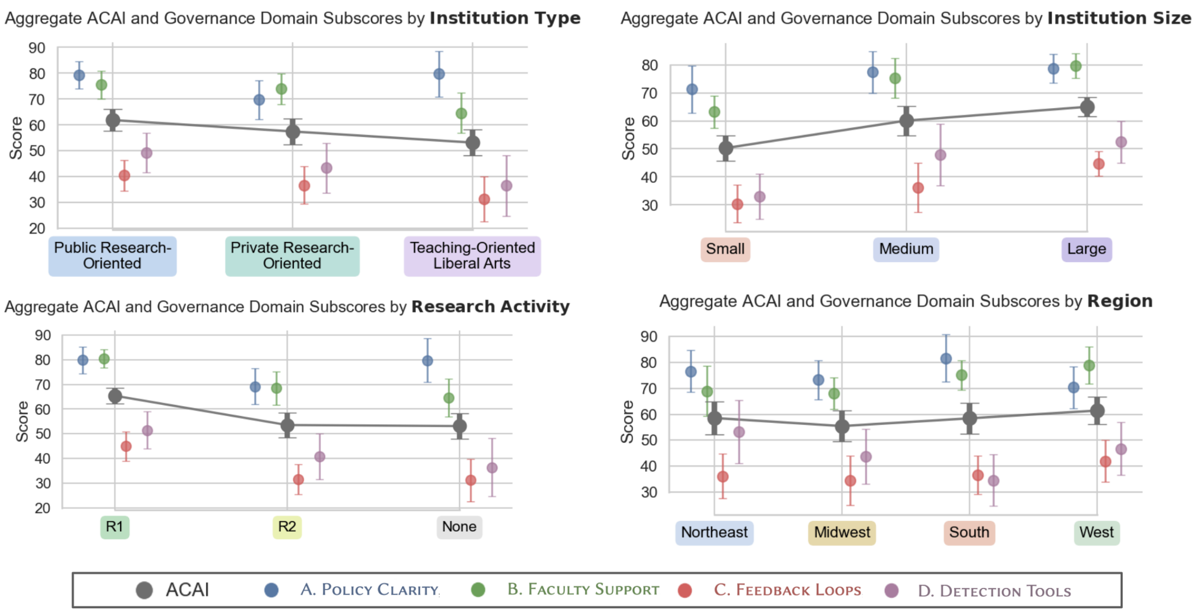

Figure 6: Aggregate ACAI and governance domain subscores (A,B,C,D) show that AI governance capacity varies with organizational structure, and that governance participation (C) and due process (D) are undersupported.

τ	Pearson r	Spearman ρ	Kendall τ	Spearman 95% CI	Mean Δ	$\Delta_{\text{Bottom 25\%}}$	$\Delta_{\text{Top 25\%}}$
0.5	0.56***	0.60***	0.43***	[0.42, 0.74]	15.39	[0, 5]	[21, 69]
1.0	0.52***	0.54***	0.38***	[0.36, 0.71]	16.23	[0, 5]	[22, 70]
1.5	0.53***	0.52***	0.37***	[0.30, 0.68]	17.34	[0, 7]	[24, 67]

Figure 7: **Correlations between human-labeled ACAI ranks and LLM-labeled ACAI ranks indicate that LLMs only weakly approximate human interpretive judgment**: For institution i, the absolute rank gap is defined as $\Delta(i) = |rank_{Human}(i) - rank_{LLM}(i)|$. $\Delta_{\text{Bottom 25\%}}$ and $\Delta_{\text{Top 25\%}}$ report the minimum and maximum values of Δ_i among institutions in the lower and upper quartiles of the rank gap distribution. Spearman correlations additionally report bootstrap 95% confidence intervals.

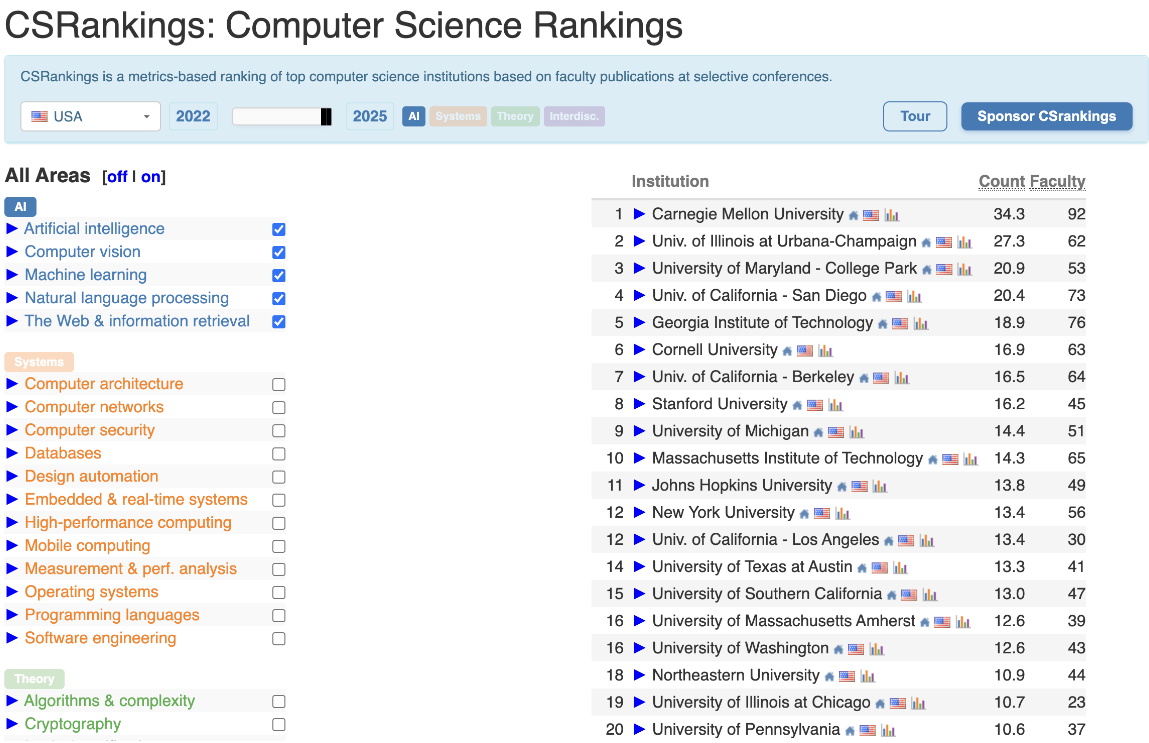

Figure 8: CSRankings website used to approximate research output, `https://csrankings.org`.