FORUM

The Unattended University

TO THE EDITOR:

In Matthew Kirschenbaum and Rita Raley's "AI and the University as a Service," one sentence struck me: "Let us bring the entire tradition of critical judgment—as it has been honed, debated, and theorized over the centuries—to bear on the problem of how to qualitatively evaluate model output" (vol. 139, no. 3, May 2024, p. 511). I agree with Meredith Martin that "[a]ll the interventions" they "propose could use better circulation beyond the pages" of *PMLA*, and especially this one: our critical work has a history that has remained relevant over centuries ("Command Lines for the Humanities"; vol. 139, no. 3, May 2024, p. 543). The questions the two ask of current AI products are topical and pertinent: "Are we sure we know what 'good' output is? Can the basis for evaluation be concretely identified and scored, or is it a tacit hermeneutic contract?" (511). Their essay, however, doesn't do much to connect the "entire tradition of critical judgment" to their own questions. It recognizes that we are now dealing with a "technolinguistic situation" and its "epistemic and epistemological" consequences (506). It exhorts us to pick up the tools we have to hand, but doesn't explain how these tools have been, and still might be, used. Two years down the road, I'd like to invite Kirschenbaum and Raley to elaborate on their claims.

The piece seems to be operating on several timescales at once. It's good to know that the MLA is responding to the advent of AI and that throughout its history our organization has been "think[ing] and work[ing] through structural transformations in language practices." Kirschenbaum and Raley direct us to the pages of *PMLA*, which, "since its inception in 1884," has been a platform for such debates (505). Behind these "agonistic disputes," they say, there has been running a "coherent project on the back end": figuring out "*language*" (506). If you peruse the back issues of this journal, you *might* see something like a "paradigm" (as we learned to call it) sustained over time. But you would certainly see the obsolescence of whole schools of thought, each with its attendant ethos, values, and structures of feeling, as well as the ghostly survival and permutation of those schools into something "else." (And often it's just the practices that survive: "close reading.")

PMLA invites members of the association to submit letters that comment on articles in previous issues or on matters of general scholarly or critical interest. The editor reserves the right to reject or edit Forum contributions and offers *PMLA* authors discussed in published letters the opportunity to reply. Submissions of more than one thousand words are not considered. The journal omits titles before persons' names and discourages endnotes and works-cited lists in the Forum. Letters should be e-mailed to pmlaforum@mla.org.

© 2026 The Author(s). Published by Cambridge University Press on behalf of the Modern Language Association of America
PMLA 141·3 (2026), doi:10.1632/S0030812926102090

https://doi.org/10.1632/S0030812926102090 Published online by Cambridge University Press

What continuities—other than a willingness to "think and work" on these matters—do they trace between the disputes that marked our shared past and those of our contested present? Is "*language*" a category too capacious (or else too definite) to subsume all that controversy? Even Kirschenbaum and Raley's account in the "profession's journal of record" does not clarify how the journal's archive informs the argument (505).

It's a question of timing—and of scaling. Some of what Kirschenbaum and Raley say is explicitly keyed to the "exigencies of the moment" (508)—for instance, their remarks on prompt engineering: it is "the design and development of methods by which to instruct, cue, or otherwise initiate the process of language generation so as to result in 'good' output" (507). And some of it hints at epochal change: "Perhaps" subword tokenization means that "the leading sciences and technologies are no longer operated and mediated by language *as such*. And perhaps we truly are at the beginning of something new" (508). Such timescale shifts can be dizzying, and I'm not sure that zooming way in or way out like this does that much to bring the AI picture into focus. Prompt engineering, to take their example, is now so 2024, and context engineering, in 2026, has displaced it. (But who knows what 2027 will bring? Semiautonomous "agents," probably.) AI cognoscenti often tell us how hard it is to keep up with their field from week to week, even day to day. Some new model release or market projection or research study is always dropping, while, at the same time, the arrival of something truly new, something that will change everything, is always about to happen. (When will "superintelligence" manifest itself? Is it five years from now, or ten, or never?) To me, it seems unlikely that a useful AI-adjacent literary criticism—"'good' output"—could in practice be either extremely presentist or extremely futurist. The specifics of AI knowledge age very quickly, and the field itself is volatile and portentous, a battle space for hypers and p(doomers). Which makes it all the more important for us to locate a "foundation model" (509), a disciplinary staging ground defined by the hermeneutic skill set that critics now have, and also, necessarily, by the "entire tradition of critical judgment" that subtends our skill set—that is, to do exactly what Kirschenbaum and Raley ask us to do.

What we're told, though, is that "*language*" is and long has been the ground of our discipline and also, simultaneously, that "language *as such*" may no longer be operating down at the technological "foundation." It's not obvious, at least to me, what the object of our criticism should then be, or by what protocols its readings would proceed. It's true we cannot predict how, or even whether, our interventions would play out "beyond the pages" of *PMLA*. Kirschenbaum and Raley are right to say that incorporating LLMs into our workflows might provide "domain-specific" benefits for "literary critical and historical research," but still not offer "exceptional status for even the most digitally forward humanists" (508). Isn't this way of thinking part of the problem, though? In her response to their essay, Katherine Elkins recalls that "I've had [AI] model developers defer to me in conversations, insisting, 'You tell me: you're the language expert.'" And "[i]ndeed, we are," she says, "it's our field as much as theirs" ("A(I) University in Ruins: What Remains in a World with Large Language Models?"; vol. 139, no. 3, May 2024, p. 560). Maybe so; I hope so. If such an entente is to come about at scale, it must be clear—first to literary critics themselves—just what their theory and their criticism have to offer. And so: what is that? What has the "entire tradition of critical judgment . . . honed, debated, and theorized over the centuries," taught us?

David J. Baker
University of North Carolina, Chapel Hill

Reply:

David J. Baker identifies a real tension in our essay—the oscillation between the immediate exigencies of the moment and the epochal shift of the technolinguistic condition—and asks, in effect, for a foundation model of criticism: a disciplinary staging ground, grounded in the "entire tradition of critical judgment" we ourselves invoked, that might stabilize practice against the volatility of the AI landscape. The desire is understandable, even necessary, for a profession long anchored in the relative persistence of the text, the archive, the interpretive act. But in the context of the University as a Service (UaaS), this demand for stable critical ground is structurally isomorphic with the service model itself. Standardization, interoperability, deployability—the modular virtues that render critique itself legible, packageable, measurable, and therefore optimizable—are equally the governing terms of what we first set out to diagnose. The foundation model of criticism and the foundation model

https://doi.org/10.1632/S0030812926102090 Published online by Cambridge University Press

of the LLM converge not in their content but in their logic. Both promise an invariant core from which downstream applications can be spawned.

Since we wrote, the landscape has shifted in ways that make the question less tractable. The shift is not away from prompting but rather from generation to delegation, from representation to *execution*—a word whose computational, juridical, and performative senses are here impossible to disentangle. A tool like *NotebookLM* could still be contained within a logic of representation—ingesting the *Canvas* bookshelf, transcribing lectures, producing summaries on demand—presupposing, however minimally, a student who engaged with the output and still tended a "notebook." The emergence of agentic protocols has tilted the spectrum toward execution: software architectures that allow AI systems not merely to produce text but to act across applications and platforms, (note)booking, submitting, authenticating, deploying, completing. The scene of writing we described in our essay presumed a human at a keyboard prompting a model and receiving output. The agentic turn dispenses with even that residual scene. We still have to prompt, if only once; what we delegate is not comprehension but presence.

The conditions for this were latent in the platform architecture all along; what has changed is that their operational expression has become trivial. Earlier this year a graduate-school dropout launched a short-lived application called *Einstein*, inviting students to hand over their *Canvas* credentials so that autonomous agents could complete coursework on their behalf. The credential itself becomes sufficient; the system executes assignments, discussion posts, assessments, the entire modular performance of learning. The innovation, such as it is, lies not in *Einstein* itself—the underlying capabilities are running rampant through your email inbox—but in what its brief existence revealed about how little stands between the logic of the service model and its full operational expression. Instructure, which owns *Canvas*, issued a cease and desist; so did the Hebrew University of Jerusalem, which owns Albert Einstein's name. The site was dark in four days. In March, on a timetable announced the previous summer, Instructure shipped its own agentic tool, *IgniteAI Agent*, designed to automate what it calls "low-value tasks" for faculty members—rubric generation, content alignment, discussion review (Kathryn Palmer; "Canvas Unrolls AI Teaching Agent"; *Inside Higher Ed*, 23 Mar. 2026, insidehighered.com/news/tech-innovation/artificial-intelligence/2026/03/23/canvas-unrolls-ai-teaching-agent). The agent expelled through the front door walked right back in through the vendor contract. The agentic horizon recedes as fast as one tries to enforce it.

If *Einstein* and *IgniteAI* complete the platform from below and within—the student's agency delegated to a script, the instructor's labor rationalized by the vendor—Arizona State University's *Atomic* completes it from above: faculty lectures scraped from *Canvas*, chopped into short clips stripped of classroom context and scholarly framing, and sold at five dollars a month, noncredit, to people the university has not enrolled. The university harvests its own intellectual labor and repackages the slop as personalized education. *Canvas* does not know, and cannot know, whether the entity completing the assignment is a student or a script; neither, it seems, does the university know—or care—when its faculty members are reduced to raw material for its content platform. One might say the agent *completes* the institution, in both senses: it fulfills its logic and finishes it off. *Einstein* takes the credential without the content; *Atomic* sells the content without the credential. The university's remaining function is to certify outcomes it no longer witnesses.

And here Baker's question returns with a different valence. What would the foundation model of criticism evaluate? There is no text to read. There is no output to interpret. There is a transaction between two automated systems—agent and platform—producing a data record operationally indistinguishable from learning. The object of criticism has not become probabilistic or opaque. It has been replaced by a service call: the performative utterance stripped of its human speaker and carried out as automated process. A service call is not a text, in this class or any other. It is not even a bad text. It is an execution.

The UaaS thesis, then, is no longer only about the outsourcing of infrastructure (lecture capture, transcription, course management) but about the outsourcing of agency itself: the student delegates action to the agent, the instructor delegates assessment to the platform, the institution delegates infrastructure to the vendor, and the vendor delegates governance to the model. What is outsourced at each level is not labor but the capacity to act, to judge, to

https://doi.org/10.1632/S0030812926102090 Published online by Cambridge University Press

decide—the capacities the humanities have claimed as their province. What remains is the form of the institution without its function: not the university in ruins but the university running, as it were, unattended. Excellence without content, and now without even the pretense of human administration. None of this stops at the campus; the university is only where the outsourcing of judgment, the capacity on which institutional legitimacy depends, first becomes legible as such. Will the same agentic architectures now being tested on syllabi and discussion boards be deployed to manage kill chains? The interoperability protocols do not change with the domain and neither, one has to assume, will the integration timelines, the compliance regimes, the change-management plans, or the executives who will describe the transition as an opportunity to reduce human capital costs.

We did not have a program to offer before. We have less of one now. The agent is not the crisis. The agent is the invoice.

Matthew Kirschenbaum
University of Virginia

Rita Raley
University of California, Santa Barbara

https://doi.org/10.1632/S0030812926102090 Published online by Cambridge University Press