

From Plausibility to Practice: Generative AI and the Changing Research Environment

Faculty Contributor
Dr. J. Elliott Casal

Three years ago, “generative AI” meant ChatGPT: a surprising, sometimes impressive, often wrong, meticulously hyped, and generally concerning technology. Many also saw tremendous promise. Three years later, this all still largely applies, although we have a broader set of so-called ‘GenAI’ tools and are somewhat less surprised by new affordances as they roll out to the public. As I see it, the real impact of GenAI on research and teaching within my field (and beyond) is not merely that machines can produce language, but that they are reshaping the conditions under which language, expertise, and knowledge circulate.

First impressions

The first thing I ever did with ChatGPT back in late 2022 was ask it to write a poem about itself in the style of Emily Dickinson. The second was to ask it to give me home-care advice for hallux rigidus while recreating the persona of Mike Tyson from the cartoon Mike Tyson Mysteries (everyone else was trying pirates). Both tests were strange, uneven, and compelling.

Focusing on my first example, my early impression was that because responses were instantaneous and had continuity, there was a sort of interactive persistence between prompts. Context accumulated through refinement in a way that was distinct from the trial-and-error logic of traditional coding. I had to teach and model slant rhyme, I had to define and model Dickinson’s unique wit, I had to work on a core metaphor for the poem, but we landed on the verb trace to describe what GenAI does. The resulting poem was surprisingly good for an early public version of ChatGPT, and I felt an eerie sense of ownership and involvement. The first stanza reads:

*Upon the Web I met a Mind — A
Wit without a Face —
I asked if it knew Humankind,
“It’s you,” it said, “I trace.”*

When I was first asked to speak on ChatGPT at our university, my charge was to address faculty and student thoughts and concerns about the new technology. In preparation, I asked faculty across campus about GenAI, and there was a clear focus on immediate institutional disruption via plagiarism and academic dishonesty. Undoubtedly, this remains a primary issue in higher education. When I asked undergraduate students about their views, a consistent existential feeling of uncertainty emerged: What could students dedicate themselves to ‘now’ that would still be valued in four years?

At the time, conversations about what this technology meant for research (and research ethics) in my field (and in general) were limited, and so, on the last day of AAAL 2023 my colleague Matt Kessler (University of South Florida) and I sought to jumpstart such discussions. Using a long-form prompting protocol, we investigated the extent to which 72 editorial board members from top Applied Linguistics journals could distinguish AI-generated research article abstracts from human-written ones and what criteria formed the basis of their decisions. To the surprise of many, the findings of our paper (Casal & Kessler, 2023) suggest that despite employing multiple rationales to judge texts, reviewers were largely unsuccessful in identifying AI versus human writing in this context, with an overall positive identification rate of only 38.9% even at that time. The key was that we fed entire research papers into the system, rather than providing a generic prompt to ‘write an abstract for a research paper on X topic’. Like my Emily Dickinson poem experiment, this highlighted not simply that ChatGPT was already capable of producing lucid text, but that users could collaboratively shape genre-appropriate discourse through iterative input and refinement, and they could do so with remarkable speed.

Suddenly, everything has changed

Early questions were often framed around plausibility: *could ChatGPT/GenAI make texts that were ‘good enough?’* However, a short time later we are in a radically changed linguistic ecology. Doubts of plausibility have given way to rapidly evolving practices with surprisingly little policy. In “Taking the reins as researchers and steering its use toward ethical outcomes” (Kessler & Casal, 2024), my co-author and I argue that researchers must take an active role in shaping how these technologies are used rather than treating them as inevitable or external to academic practice. As they emerge, policies tend to be frameworks for permissiveness rather than tools for making principled practical decisions. Admittedly, this applies to my own work contributing to GenAI policies as well; the landscape is changing quickly, and there is widespread disagreement.

Clearly, AI-generated language is no longer marginal or experimental. Recent reports claim that substantial portions of online text (at least 30% of text on active websites), and particularly new online content (up to 75%), is now produced or assisted by generative AI tools (Law et al., 2025; Spennemann, 2025). Researchers across scientific disciplines increasingly report using these systems in the preparation of scholarly writing as well (Liang et al., 2025). The ethics surrounding these trends are murky and context-specific, and some communities are overrun with ‘slop’ too.

Within linguistics, modern usage-based approaches highlight that language users are sensitive to the distributional patterns of linguistic resources in their environments (e.g., Ellis et al., 2014; Ellis & Ferreira-Junior, 2009). In other words, people learn language from the language they are exposed to, building linguistic knowledge as they pay attention over their lifetimes to the ways that different meanings are expressed in context. Language produced by generative AI tools is now a prominent part of the linguistic ecology, as it increasingly appears in search results, institutional communications, student writing, public and online discourse, entertainment and media content, and even scholarly publication itself. As the presence of AI-generated language in these (and other) environments become normalized, significant questions are raised about how these systems may recursively shape future language use, genre expectations, and communicative norms.

To me, there is plenty of room for compelling research on these issues. Do we attend to and learn from linguistic input differently if we are aware a language source is AI generated? How strongly are community conventions reshaped by changes in patterns of AI language across models? What are the implications of students reading AI-generated summaries rather than primary sources on their literacy skills and depth of cultural knowledge? Does the flattening of lexical, phrasal, syntactic, and rhetorical conventions in AI-generated or assisted text push us to prize creativity further?

Let's say the human part out loud

Instead, much discussion in research, educational, and technology worlds focuses on 'AI capability,' as though outputs emerge independently of interaction. Discussions of AI in linguistics often focus on whether systems possess linguistic competence, what features they do or do not make use of 'like humans', or how effectively they contribute to language teaching and learning. To some degree and depending on one's theoretical perspective, aspects of these questions may indeed be valuable.

However, AI outcomes are often treated as properties of the models themselves rather than emergent properties of human-model interaction, including iterative refinement practices, user expertise and disposition, and discourse contexts. In many cases, the early discussions of 'the human in the loop' have already faded into set-it-and-forget-it approaches for prompt engineering that reinforce a static view of AI potential, rather than an extended cognition model of human computer interaction.

My own experiences with ChatGPT and other tools, including collaboratively refining an Emily Dickinson poem line-by-line, suggest something quite different. The most compelling outputs often emerge not from well-conceived isolated prompts, but from sustained interactions that leverage my own creativity, evolving goals, capacity for learning, knowledge and expertise, and playful disposition. AI-generated language is now abundant in our usage environments, but the best AI outcomes emerge through recursive interaction empowered by the things that make humans good at what they do.

The students were right about the uncertainty, and today the question is clearly no longer whether generative AI can produce plausible discourse. It is how research, teaching, and language itself change once plausible discourse becomes abundant and cumulatively embedded within the environments from which future discourse practices emerge.

References

- Casal, J. E., & Kessler, M. (2023). Can linguists distinguish between ChatGPT/ AI and human writing?: A study of research ethics and academic publishing. *Research Methods in Applied Linguistics*, 2. <https://doi.org/10.1016/j.rmal.2023.100068>
- Ellis, N. C., & Ferreira-Junior, F. (2009). Constructions and their acquisition: Islands and the distinctiveness of their occupancy. *Annual Review of Cognitive Linguistics*, 7(1), 188–221. <https://doi.org/10.1075/arcl.7.08ell>
- Ellis, N. C., O'Donnell, M. B., & Römer, U. (2014). The processing of verb-argument constructions is sensitive to form, function, frequency, contingency and prototypicality. *Cognitive Linguistics*, 25(1), 55–98. <https://doi.org/10.1515/cog-2013-0031>
- Kessler, M., & Casal, J. E. (2024). (Socio)linguistics and generative AI: Taking the reins as researchers and steering its use toward ethical outcomes. *Journal of Sociolinguistics*, 28(5), 31-35. 10.1111/josl.12682
- Law, R., Guan, X., & Suolo, T. (2025, May 19). 74% of New Webpages Include AI Content (Study of 900k Pages). SEO Blog by Ahrefs. <https://ahrefs.com/blog/what-percentage-of-new-content-is-ai-generated/>
- Liang, W., Zhang, Y., Wu, Z., Lepp, H., Ji, W., Zhao, X., Cao, H., Liu, S., He, S., & Huang, Z. (2025). Quantifying large language model usage in scientific papers. *Nature Human Behaviour*, 1–11. <https://doi.org/10.48550/arXiv.2404.01268>
- Spennemann, D. H. (2025). Delving into: The quantification of Ai-generated content on the internet (synthetic data) (arXiv:2504.08755). arXiv. <https://doi.org/10.48550/arXiv.2504.08755>