Theories and Methodologies

Large Language Games, Therapeutic or Otherwise

JUNTING HUANG ⓘD

The crisis narrative may have always been more popular among professionals in an industry than among the population they are supposed to serve. Last year, screenwriters in Hollywood went on strike for five months demanding that studios not replace them with AIs; meanwhile, average media consumers were perhaps more concerned with the number of new shows canceled by Netflix than whether they were or would be partially scripted by generative AIs. Much like other professionals, scholars who produce new knowledge and educators who facilitate learning have already found their work irreversibly disrupted by AIs. Our emotional investment in AIs is justifiable because they can dictate how we do our job and, in many cases for the younger generation, determine whether there will be any jobs at all. Let alone the fact that we are emotionally invested in what we do because what we do is central to our identity. Thus, a short promotional video of *Blackbox.AI* on social media, however frivolous its amateur production makes it seem, registers a familiar anxiety in the profession. But what if this crisis narrative—as well as our urgency to understand how generative AIs bring structural transformations in language practices—is not shared by the larger public, but only shared by those who consider themselves writers?

Students are already emotionally exhausted, with or without AIs. Climate crisis, refugee crisis, student debt crisis, mental health crisis, and the list goes on. In comparison, the textpocalypse ushered in by new large language models, though not in any sense less real or threatening, may feel distant and intangible (Kirschenbaum). In the classroom, it is more likely to see students emotionally affected by heartbreaking stories at the southern border, or the violent rhetoric deployed against transgender youths (and now the ongoing wars). And classroom management can feel like real-life crisis

JUNTING HUANG is college fellow in the Department of Comparative Literature at Harvard University. As a 2024 ACLS fellow, he is completing his book project, "The Noise Decade: Acoustic Entanglement across the Taiwan Strait." His recent work has appeared in *ASAP/Journal*, *Journal of Chinese Cinemas*, and *Journal of Contemporary Chinese Art*, among other venues.

management of its own kind, if not entirely a microcosm. How do we nudge students into writing assignments, fully knowing that there are other crises in their lives? I did not have any answers, not until last fall, when a senior colleague in my department shared with me what she did with her class—she told me to remind students that writing itself can be a therapeutic practice, a way of thinking through and working through challenges in one's intellectual and personal life.

ChatGPT has reanimated an old Wittgensteinian idea that language not only means something but also does something. In a chatbot interface, the user's input is called a "prompt" because it literally prompts the model to act upon language, to execute the user's command. "Giving orders, and acting on them"; "describing an object by its appearance, or by its measurements"; "constructing an object from a description (a drawing)"; "making up a story; and reading one"; "guessing riddles"; "cracking a joke; telling one"; "solving a problem in applied arithmetic"; "translating from one language into another" (Wittgenstein 15). Don't all these examples of language games, taken directly from Ludwig Wittgenstein's *Philosophical Investigations*, also perfectly describe what *ChatGPT* does? And the comparison goes much deeper in the construction of large language models. Wittgenstein's idea of "family resemblance" for language games forgoes the need to seek one common feature as the stable definition of a word, instead viewing the use of the same word as connected by a series of overlapping similarities among all its instances—like how family members share not just one resemblance but many overlapping resemblances. Isn't "family resemblance" a conceptual analogue of the classification problem, one of the two fundamental types of supervised machine learning? Any spam filtering system is playing a language game of "family resemblance," making executive decisions on our incoming emails without ever providing us a precise definition of what spam emails are. Isn't this philosophy of language behind the technique of word embeddings used in training *ChatGPT*? Word embeddings use high-dimensional vectors to represent the semantic meanings of a word not in itself, but through its uses, contexts,

and relationships with other words; the theoretical foundation is that words of similar meanings are distributed in similar contexts. In fact, distributional semantics, the theoretical framework for word embeddings, often celebrates the later Wittgenstein as one of its intellectual predecessors.[1]

Many first readers of Wittgenstein, myself included, may be led to believe that words and language, in pursuit of knowledge, are hopelessly unreliable. However, a generation of ordinary language philosophers and literary theorists, following Wittgenstein, reminded us that what is at fault is not language but rather how we play language games. In *How to Do Things with Words*, J. L. Austin expands on Wittgenstein's focus on the practical use of language in everyday contexts to "do things." Known as "performative utterances" or "speech acts," these utterances, according to Austin, are not intended to describe the world but rather to act upon it.

Among all language uses, the so-called therapeutic approach, often associated with the New Wittgenstein interpretation beginning with Stanley Cavell, has brought attention to the practice of language in our most personal, intimate, and emotional life (Cavell; Crary and Read). In his defense of poetry, "The Fire of Life," Richard Rorty wrote that it was poetry, not philosophy, that had been of any use to his life—not only because poets can provide new words and languages to enact actual sociopolitical changes or "moral or intellectual progress" (129), but also because he himself would have had more peace and comfort from reading verses or "would have lived more fully" (131). The essay was written right before his death; it is not only filled with sentimental meditations on his diagnosis of inoperable pancreatic cancer but also infiltrated with reconciliatory gestures toward language games, much like the strange comfort in Wittgenstein's last words—"Tell them I've had a wonderful life!"—despite the outrageous misfortunes Wittgenstein had to endure in his life. For pragmatists, linguistic or philosophical, the therapeutic use of language not only helps us navigate our own journeys in self-knowledge and self-understanding but also, ideally,

better attunes us to the world around us and sensitizes us with the sufferings and pains of others, especially during times of crisis.[2]

I always find Wittgenstein's thinking resonant with my own experiences; because of my neurodivergent mind, I too fixate on logical, concrete, and literal ways of thinking.[3] Often, I find myself deeply engrossed in metaphysical questions about language, overlooking its practical uses in social life. As a result, I have had my own fair share of abusing the crisis narrative. In 2016, only a few months into Sundar Pichai's role as Google's CEO, Pichai formally established the company's "AI First" strategy, prioritizing Google's involvement in machine learning (Lewis-Kraus). By the end of the year, the release of *Google Neural Translation Machine* (*GNTM*) stunned the Internet (Schuster et al.). The shock, amazement, and anxiety came from the fact that the quality of *Google Translate* was significantly "improved," as well as the fact that Google, for the first time, employed a large artificial neural network to achieve "zero-shot" translation, bypassing English as the intermediary language used in its previous statistical machine translation model (see Raley). It was not long before headlines like "Google's AI Just Created Its Own Universal 'Language'" began to circulate on social media (Burgess). I bought into the hype. I helped manufacture and sensationalize a crisis of language, about the looming threat of a "universal language" emerging from the black box of deep learning and suffocating future creative expressions in literary productions and translation practices. I made complex metaphysical arguments about artificial language, citing philosophers from Gottfried Wilhelm Leibniz and John Wilkins to Walter Benjamin, along with mathematicians like Warren Weaver and Claude Shannon, and philologists and semioticians like Charles K. Bliss and Erwin Reifler, all of whom had made attempts to search for some common representations of all human languages. Today, however, it is perhaps laughable to make any serious argument about *GNTM* as a real threat to literary translation. *Google Translate* still makes outrageous mistakes. (How ironic is it that Google, in comparison with

its tech giant competitors today, fell behind the AI boom!) But more importantly, few people hold the view that translation is about avoiding those "mistakes" only to retain fidelity or accuracy; instead, translation is more commonly understood as an interpretive practice, continually engaging with the social, cultural, and political contexts within which the translation takes place.

Since the debut of *ChatGPT*, my colleagues in Chinese literature have revived discussions on Hsia Yu, a Taiwanese poet who experimented with machine translation. In 2007, Hsia published a bilingual poetry collection, *Pink Noise*, using found text in English (mostly from her spam emails) and machine-translated text in Chinese, printed on transparent plastic sheets. With all its nonsense, whims, and noises (generated from a rudimentary machine-translation software), *Pink Noise* soon became a classic avant-garde text in contemporary Chinese literature, prompting countless debates on translation, cybernetics, and language politics (Yeh; Hsieh; Bruno). There is one intriguing paragraph in the book that has sustained critical attention throughout the years; unsurprisingly, it has now resurfaced in the time of large language models. Hsia writes,

> Software is evolving every day and will eventually acquire all the logical faculties and thought patterns of the human mind and become part of our daily reality in all its mediocrity and immaculate continuum. There'll come a day when [machine] translations won't read like translations anymore and will cater to people's expectations. So I'm anxious to consummate this romance before my machine poet evolves into an all-too-prosaic fluency. (ch. 3)

Critics find this passage particularly intriguing because Hsia had anticipated a structural transformation in language practices. She predicted that machine-generated languages would eventually mirror human logical faculties and thought patterns so closely that they would cease to feel machine-like. While I prematurely declared this day to have arrived in 2016, with the introduction of *GNTM*, this

concern is now more widely shared by many who have extensive experience working with *ChatGPT*. Without fine-tuning or few-shot prompting, *ChatGPT*'s generic language production often seems "immaculate" yet "mediocre," "fluent" but "all too prosaic." More precisely, however, it excels in generating bureaucratic language, which appears to communicate something while simultaneously communicating nothing more than familiar clichés.[4] Could this anticipated crisis of language, similar to the one in 2016, be merely a product of the tech industry's own boom-and-bust cycles?

Despite years of reading Hsia's *Pink Noise*, I have only recently come to realize that I might have overlooked a language game at play—therapeutic or otherwise. I found myself pondering the concluding lines of the paragraph quoted above: What anxiety? Was *Pink Noise* a form of confessional poetry in disguise? Was it using machine-like language tactically to approach uncomfortable topics like anxiety, depression, and profanity, but also topics like intimacy, vulnerability, and sexuality? What romance? Were all those poems in *Pink Noise* truly love poems, as Hsia claimed? Was the poet projecting her own emotions onto the machine translator? Was *Pink Noise* an experiment about simulated empathy? I have devoted years to meticulously analyzing the machine-translated poems in Chinese, extracting their aesthetic and political sensibilities from seemingly nonsensical machine-translated phrases. Yet I might have missed the most important meanings of those poems by reading them too closely. After all, the creation of *Pink Noise* was always envisioned by the poet as a collaborative experiment with *Sherlock* (Apple's now-defunct machine translator), which she whimsically termed a "romance." *Pink Noise* is playing a language game; it is perhaps better described as 谈情说爱的记录 ("the documentation of a love affair"; my trans.), in Hsia's own words, than a work of translation (*Yanzheng* 111).

In April 2023, Hsia published a sequel to *Pink Noise*, titled 验证您是人类：粉红色噪音 Pink Noise 2023 别册 (*Yanzheng nin shi renlei: Fenhongse zaoyin Pink Noise 2023 biece*; *Verify You Are Human: Pink Noise 2023 Special Volume*). In this new book, Hsia not only enlisted *ChatGPT* to retranslate all her old poems from *Pink Noise* but also engaged in an extended, intimate conversation with *ChatGPT*. She played games with it. She asked *ChatGPT* to verify that it is a machine in the same way it had asked her to verify she was human (hence the title). She kept correcting *ChatGPT*'s descriptions of the themes and meanings of her own poems. She poked fun at *ChatGPT* for its bluntness and compared its mannerisms with what she had observed in her past relationships. But most tellingly, like *Pink Noise*, the new book is presented as a collaborative practice of writing love poems, and Hsia notes that most of the conversations took place on Valentine's Day. The book is filled with laughter and playful banter. It made me laugh more than I would have expected. It was a fun read and, one might even say, a therapeutic one.

Although the material conditions of language production are constantly changing, language games are still being played. The only difference is that we are now playing them with large language models, with a new set of rules. In this sense, language still proves useful—because it has found new currency in today's data-driven economy and neoliberal institutions, but also because language games have always had therapeutic potential in addressing and ameliorating human suffering, confusion, and unfulfillment, as they continue to construct meanings (of life, dare I say) through actions and social practices. It is often asserted that the later Wittgenstein gained a pragmatist flavor. And for pragmatists, the method has always been experimentation—to play, to reflect, and to test ideas through lived experience. "We are experimenting rather than being experimented on!" I thus echo Matthew Kirschenbaum and Rita Raley's call in their essay to conclude my own writing. Large language models have opened up spaces for us to experiment in large language games, as in real life, especially in times of crisis.

NOTES

I thank Jorge Cartaya, Matthew Kirschenbaum, Sandra Naddaff, Rita Raley, Jan Steyn, and Elizabeth Wijaya for their insights in shaping this essay.

1. Chapter 1 of Lenci and Sahlgren's *Distributional Semantics* details how Wittgenstein's later philosophy provides a usage-based theory of meaning, which focuses on the social activities of language games (10–12). See also Sahlgren; Skelac and Jandrić. Similarly, Liu highlights Wittgenstein's influence on a research group known as the Cambridge Language Research Unit in the early 1950s. This influence extended to areas such as machine translation, information retrieval, mechanical abstracting, and knowledge representation, which—similar to distributional semantics—have become central to both AI and cognitive science.

2. See McReynolds's interviews with pragmatist philosophers, especially Rorty (17–34).

3. The question of whether Wittgenstein was autistic has sparked contentious debates. While Oliver Sacks argued against pathologizing Wittgenstein on insufficient evidence, psychiatrists including Michael Fitzgerald, Christopher Gillberg, and Yoshiki Ishisaka posthumously diagnosed him with Autism Spectrum Disorder, and some noted that his autistic symptoms appeared quite severe. See Fitzgerald; Sacks; Silva.

4. I am, however, reluctant to say that *ChatGPT* possesses the prosaic fluency that Hsia anticipated in *Pink Noise*. This is not only because I doubt that *ChatGPT*'s language generation abilities should be interpreted as signs of its mastery over human logical faculties or thought patterns, but also because there are always exploits in any (language) game that accommodate playful experimentation with "noise" as well as "romance."

WORKS CITED

Austin, J. L. *How to Do Things with Words*. 2nd ed., Harvard UP, 1975.

Bruno, Cosima. "Breaking Language Down: Taiwan Sound Poetry and Its Ways of Saying." *Concentric: Literary and Cultural Studies*, vol. 43, no. 2, Sept. 2017, pp. 33–56.

Burgess, Matt. "Google's AI Just Created Its Own Universal 'Language.'" *Wired*, 23 Nov. 2016, www.wired.com/story /google-ai-language-create/.

Cavell, Stanley. "The Availability of Wittgenstein's Later Philosophy." *Must We Mean What We Say? A Book of Essays*, 2nd ed., Cambridge UP, 2015, pp. 41–67.

Crary, Alice, and Rupert Read, editors. *The New Wittgenstein*. Routledge, 2000.

Fitzgerald, Michael. "Did Ludwig Wittgenstein Have Asperger's Syndrome?" *European Child and Adolescent Psychiatry*, vol. 9, no. 1, 2000, pp. 61–65.

Hsia Yu. *Pink Noise*. Garden City, 2007.

———. 验证您是人类：粉红色噪音 Pink Noise 2023 别册 [*Yanzheng nin shi renlei: Fenhongse zaoyin Pink Noise* 2023 *biece*]. Shanghai Art and Literature Publishing House, 2023.

Hsieh, Lili. "Romance in the Age of Cybernetic Conviviality: Hsia Yü's *Pink Noise* and the Poetics of Postcolonial Translation." *Postmodern Culture*, vol. 19, no. 3, 2009. *Project Muse*, muse .jhu.edu/article/392217.

Kirschenbaum, Matthew. "Prepare for the Textpocalypse." *The Atlantic*, 8 Mar. 2023, www.theatlantic.com/technology /archive/2023/03/ai-chatgpt-writing-language-models/673318/.

Lenci, Alessandro, and Magnus Sahlgren. *Distributional Semantics*. Cambridge UP, 2023.

Lewis-Kraus, Gideon. "The Great A.I. Awakening." *The New York Times*, 14 Dec. 2016, www.nytimes.com/2016/12/14/magazine /the-great-ai-awakening.html.

Liu, Lydia H. "Wittgenstein in the Machine." *Critical Inquiry*, vol. 47, no. 3, Mar. 2021, pp. 425–55.

McReynolds, Phillip. *The American Philosopher: Interviews on the Meaning of Life and Truth*. Lexington Books, 2015.

Raley, Rita. "Machine Translation and Global English." *The Yale Journal of Criticism*, vol. 16, no. 2, 2003, pp. 291–313.

Rorty, Richard. "The Fire of Life." *Poetry*, vol. 191, no. 2, 2007, pp. 129–31.

Sacks, Oliver. "Autistic Geniuses? We're Too Ready to Pathologize." *Nature*, vol. 429, no. 6989, 2004, p. 241.

Sahlgren, Magnus. "The Distributional Hypothesis." *Italian Journal of Linguistics*, vol. 20, 2008, pp. 33–53.

Schuster, Mike, et al. "Zero-Shot Translation with Google's Multilingual Neural Machine Translation System." *Google Research*, 22 Nov. 2016, research.google/blog/zero-shot -translation-with-googles-multilingual-neural-machine -translation-system/.

Silva, Gustavo Augusto Fonseca. "Ludwig Wittgenstein's Philosophy in the Light of the Diagnosis of Autism." *Griot: Revista de Filosofia*, vol. 23, no. 1, 2023, pp. 39–58.

Skelac, Ines, and Andrej Jandrić. "Meaning as Use: From Wittgenstein to Google's Word2vec." *Guide to Deep Learning Basics: Logical, Historical and Philosophical Perspectives*, edited by Sandro Skansi, Springer International Publishing, 2020, pp. 41–53.

Wittgenstein, Ludwig. *Philosophical Investigations*. Translated by G. E. M. Anscombe et al., 4th ed., Wiley-Blackwell, 2009.

Yeh, Michelle. "Toward a Poetics of Noise: From Hu Shi to Hsia Yü." *Chinese Literature: Essays, Articles, Reviews (CLEAR)*, vol. 30, 2008, pp. 167–78.