

Model Misalignment and Language Change: Traces of AI-Associated Language in Unscripted Spoken English

Bryce Anderson*, Riley Galpin, Tom S. Juzek*

Florida State University

Abstract

In recent years, written language, particularly in science and education, has undergone remarkable shifts in word usage. These changes are widely attributed to the growing influence of Large Language Models (LLMs), which frequently rely on a distinct lexical style. Divergences between model output and target audience norms can be viewed as a form of misalignment. While these shifts are often linked to using Artificial Intelligence (AI) directly as a tool to generate text, it remains unclear whether the changes reflect broader changes in the human language system itself. To explore this question, we constructed a dataset of 22.1 million words from unscripted spoken language drawn from conversational science and technology podcasts. We analyzed lexical trends before and after ChatGPT’s release in 2022, focusing on commonly LLM-associated words. Our results show a moderate yet significant increase in the usage of these words post-2022, suggesting a convergence between human word choices and LLM-associated patterns. In contrast, baseline synonym words exhibit no significant directional shift. Given the short time frame and the number of words affected, this may indicate the onset of a remarkable shift in language use. Whether this represents natural language change or a novel shift driven by AI exposure remains an open question. Similarly, although the shifts may stem from broader adoption patterns, it may also be that upstream training misalignments ultimately contribute to changes in human language use. These findings parallel ethical concerns that misaligned models may shape social and moral beliefs.

GitHub repository for code and data —
<https://github.com/fsu-nlp/ai-traces-ssl>

Introduction

Languages change over time, and the processes underlying language change are complex, involving a combination of social (Amato et al. 2018; Zhang 2014), cultural (Labov 2011), and psycho-cognitive factors (Traugott 1985; Li et al. 2023). Over the past centuries, technology has also played a

*Corresponding authors. Author contributions: Conceptualisation and methodology (BA, RG, TSJ), data collection and analysis (BA, TSJ), coding and visualisation (BA, TSJ), write up (BA, RG, TSJ).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Whether the sharp post-2022 rise in AI-associated words reflects mere tool usage or a direct influence on the human language system remains an open question with broader societal relevance.

significant role in driving language change. The introduction of the printing press, for example, expanded access to written literature, promoted the standardization of written language, and transformed global communication (Eisenstein 1980). Similarly, the invention of the telephone and later the cell phone influenced the way humans speak and interact with one another (Crystal 2009; Sandra 2011; Al-Sharqi and Abbasi 2020; Djalolovna 2025), similarly the internet in general (Helen 2010), as well as social media (Vanisree et al. 2024). Language is the medium that humans use to communicate ideas (Clark 1996; Fedorenko, Piantadosi, and Gibson 2024). As such, understanding the possibilities and constraints of language, as well as understanding language change and the processes behind it matters.

Selected rapid shifts in word usage do occur and are typically traceable to real-world events (“Omicron,” “Türkiye”). This contrasts with the recent, large-scale shifts observed in certain domains, particularly in education and science, which appear to not be triggered by external world events. Words such as “delve,” “intricate,” and “underscore” have seen sudden spikes in usage, especially within academic writing (Gray 2024; Kobak et al. 2024; Liang et al. 2024; Liu and Bu 2024; Matsui 2024; Comas-Forgas, Koulouris, and Kouis 2025; Bao et al. 2025; see also Figure 1, which shows the frequency of selected lemmata, i.e., base forms of words, in PubMed abstracts). Scripted spoken language,

such as academic presentations, has been analyzed and similar trends were observed (Geng et al. 2024; Yakura et al. 2024). These changes are often attributed to the widespread introduction of Large Language Models (LLMs) with chat functionality, especially ChatGPT (OpenAI 2024), which overuses words such as “delve” and “intricate” (Gray 2024; Koppenburg 2024; Kobak et al. 2024; Juzek and Ward 2025), and exhibits a distinct syntactic style (Zamaraeva et al. 2025). Further, these linguistic preferences are frequently discussed in the context of model alignment: is there a gap between model behavior and user expectations (Hendrycks et al. 2020; Gabriel 2020; Norhashim and Hahn 2024), and if so, where does it originate? One possibility is that such misalignments arise during the training phase, particularly through the process of Learning from Human Feedback (in the literature also referred to as Human Preference Learning and Reinforcement Learning from Human Feedback; the latter term, however, is not ideal, as not all preference procedures involve reinforcement learning). Zhang et al. (2024) argue that the overuse of stylistic elements like boldfacing originates from Learning from Human Feedback, and Juzek and Ward (2025) suggest that lexical overuse may similarly arise during this phase. However, if such biases reflect human preferences (just not those of the intended user base) rather than contradict them, the concept of “misalignment” becomes ambiguous (cf. Blodgett et al. 2020 on the ambiguity of such definitions).

Crucially, these issues mirror broader debates around model alignment in the domains of ethics and fairness (Russell, Dewey, and Tegmark 2015; Gabriel 2020; Perez et al. 2023; Hendrycks et al. 2020; Norhashim and Hahn 2024), where biases related to race, ethnicity, nationality (Narayanan Venkit et al. 2023; Omiye et al. 2023; Akin-tande 2023; Ali et al. 2023; Naik and Nushi 2023), intersectionality (Omriani Sabbaghi, Wolfe, and Caliskan 2023), and gender (Bolukbasi et al. 2016; Ghosh and Caliskan 2023; Kotek, Dockum, and Sun 2023; Wilson and Caliskan 2024; Ali et al. 2023; Naik and Nushi 2023; Jaiswal 2024) are often amplified not because models deviate from human input, but because they replicate dominant patterns already present in it (Grabowicz, Perello, and Takatsu 2023; Akintande 2023). As such, misaligned models can cause real harm (Neff and Nagy 2016; Bender et al. 2021; Hagen, Huo, and Nguyen 2025). Thus, insights into LLMs’ language use have the potential to inform more general alignment discussions about whose preferences, whether it would be values or language preferences, are being encoded and reinforced.

Research question

When it comes to the observed language change, however, an important piece in the debate is missing: it remains an open question whether the observed shifts in language use are simply the result of tool usage (e.g., a particular instance of “delve” was generated by ChatGPT, and the human used it), or whether the language choices of LLMs have influenced human speakers and writers to such an extent that the shift reflects an actual change in human language use (i.e., “delve” was produced by a human, potentially influenced by prior exposure to LLM output) (Clark 1996; Fedorenko,

Piantadosi, and Gibson 2024). This distinction matters: if the influence is direct, it strengthens the case for more robust model alignment, as linguistic preferences encoded by LLMs may begin to shape human communication norms. Thus, the research question we approach is:

(RQ) Are human lexical choices in unscripted spoken language increasingly converging towards the lexical patterns characteristic of Large Language Models?

That is, among the words that LLMs are known to overuse (as documented in the literature; see the Methods subsection for the list of words that we investigated), can we observe an uptick in genuine human language production, excluding tool-generated language? Our hypothesis is that such an uptick exists, at least to some degree. The null hypothesis is that there is no notable increase.

Measuring “genuine” human language production presents considerable challenges, especially in written texts. Before the widespread adoption of Large Language Models, the default assumption was that any given text was authored by a human. This presumption no longer holds. For many texts, it is not clear anymore if they were genuinely authored by a human, i.e. they present a case of **human-authorship indeterminacy** (see Huang, Chen, and Shu 2025 and Sadasivan et al. 2023 for related thoughts). For linguistic analysis aimed at understanding the human language system, this marks a considerable disruption, especially since most linguistic research has focused on the analysis of written language. The field’s (over)reliance on written data has been criticized in the past (Biber 1991), and the rise of AI-generated content makes deriving linguistic insights from such data increasingly difficult. The analysis of spoken language holds value (Miller and Weinert 1998); however, even spoken language is not immune, as academic presentations, for instance, are often scripted, and these scripts may be generated by Artificial Intelligence (AI) models. To address this challenge, we focus on unscripted spoken language from conversational podcasts.

Further related work

Zhang et al. (2024) investigate changes in formatting, such as lists and boldface. Geng et al. (2024) and Yakura et al. (2024) examine scripted spoken language. Liang et al. (2024) map shifts in word usage in written academic texts, while Pack and Maloney (2023) focus specifically on the domain of education. Galpin, Anderson, and Juzek (2025) offer a deeper structural analysis of word usage patterns and their changes over time. Zhang et al. (2024) observe that LLMs overuse certain stylistic elements like boldface and making use of lists, and implicate Learning from Human Feedback, a procedure to align LLMs with human preferences (Christiano et al. 2017; Ziegler et al. 2019; Ouyang et al. 2022; Rafailov et al. 2024). Wu and Aji (2025) observe that humans tend to prefer model outputs with polished style and lower substance over outputs with greater substance but less polished style. Juzek and Ward (2025) argue that, with respect to the overuse of certain words, such model behavior also stems from Learning from Human Feedback. We contribute to this body of work by focusing on unscripted

Word	Change (%)	OPM Pre-2022	OPM Post-2022	$p \leq 0.05$
surpass_VERB	140.79	1.47	3.53	True
boast_VERB	140.14	0.82	1.98	True
meticulous_ADJ	125.52	0.46	1.03	False
strategically_ADV	87.93	1.74	3.27	True
garner_VERB	53.76	1.01	1.55	False
notable_ADJ	53.13	2.47	3.79	False
intricate_ADJ	47.21	2.75	4.04	False
delve_VERB	46.82	2.93	4.30	False
align_VERB	36.59	18.77	25.64	True
advancement_NOUN	30.27	4.03	5.25	False
significant_ADJ	17.35	39.74	46.63	True
showcase_VERB	6.49	1.37	1.46	False
comprehend_VERB	3.36	3.66	3.79	False
pivotal_ADJ	0.45	2.66	2.67	False
potential_NOUN	-2.95	33.51	32.52	False
emphasize_VERB	-4.15	18.22	17.47	False
potential_ADJ	-12.15	33.79	29.68	False
realm_NOUN	-17.63	22.25	18.33	True
crucial_ADJ	-24.56	19.05	14.37	True
groundbreaking_ADJ	-40.94	3.20	1.89	False

Table 1: The target words occurring in our dataset, including their percentage change in frequency (occurrences per million; OPM) for pre-2022 vs. post-2022, and whether the difference is significant.

spoken language and by directly addressing the question of whether the human language system itself has been affected by the presence and output of AI.

Beyond linguistic considerations (to give an incomplete selection), further considerations include Norhashim and Hahn (2024), who propose frameworks for quantitatively assessing human-AI value alignment in Large Language Models. Hendrycks et al. (2020) introduce benchmark tasks aimed at assessing alignment with shared human values, including helpfulness, honesty, and harmlessness. Keswani et al. (2024) explore how active learning methods used for preference elicitation may shape the moral values internalized by models. Leidinger and Rogers (2024) contribute to ongoing debates on mitigating stereotyping harms in generative systems, while Welbl et al. (2021) and Hartvigsen et al. (2022) discuss challenges with respect to detoxification and adversarial training, respectively. Finally, all these considerations are accompanied by discussions about governance and ethical frameworks, including proposals for algorithmic accountability (Mittelstadt et al. 2016), ethically aligned AI (Floridi et al. 2018), and global AI ethics guidelines (Jobin, Ienca, and Vayena 2019).

Methods

Target words: For the AI-related words under investigation, we used the 34 words identified in Galpin, Anderson, and Juzek (2025), which were drawn from prior literature and their own analyses. For any of the 34 words, it was established through a formal analysis that there has been a considerable increase in the words’ usage pre- and post-ChatGPT’s release in 2022. The selection process prioritizes words that either displayed the most pronounced changes or received substantial attention in the literature

(ibid.). Out of the 34 target words from the literature, 20 did occur in our dataset. The list of occurring words is provided in Table 1. The 14 words not occurring in our dataset are: commendable_ADJ, comprehending_ADJ, emphasized_ADJ, garnered_ADJ, intricacy_NOUN, invaluable_ADJ, meticulously_ADV, noteworthy_ADJ, showcase_NOUN, showcasing_ADJ, surpassing_ADJ, underscore_VERB, underscore_NOUN, underscoring_ADJ. All of the non-occurring words are low-frequency words (note the infrequent adjectival usage in many of these words), most of which occur well below one instance per million words, for example, when analyzing occurrences in PubMed abstracts. Moreover, all of the 34 words belong to a more formal or academic register, which makes their occurrence in unscripted spoken language less likely. (All scripts, sample data, and results can be found in our GitHub repository.)

Dataset construction: Because of the issue of human-authorship indeterminacy, our dataset consists of unscripted spoken language from conversational podcasts. We focused on the most popular podcasts in the United States related to technology and science, based on the assumption that speakers engaging with these topics are likely to be directly or indirectly exposed to LLM-generated language. Accordingly, this subgroup offers a high probability of revealing potential effects. To determine popularity and identify relevant podcast content, we looked at Apple and Spotify algorithm recommendations, as well as queried Google, YouTube, and ChatGPT. All prompts included the condition that results were to return only conversational podcasts within both genres to maximize capturing unscripted language. For each candidate podcast, we listened to a sample episode to ver-

ify the presence of spontaneous, unscripted speech.

Given that ChatGPT, the most influential chat-based LLM, was released in late 2022, we divided our dataset into two periods: pre-2022 (mostly 2019 to 2021) and post-2022 (2023 to 2025). The year 2022 was excluded from the dataset, as it marks the release of ChatGPT and represents a transitional period between the pre- and post-ChatGPT eras. To keep the dataset balanced we collected a 1:1 ratio of pre/post episodes from each podcast. The 1:1 ratio also ensures that any potential regional or varietal differences are controlled for, as each podcast contributes equally to both time periods. Further, we tried to include more episodes from podcasts with longer average runtimes and fewer episodes from podcasts with shorter runtimes, but the counts vary due to a finite amount of podcast episodes that fit within our criteria, and could be downloaded in equal quantities for both pre- and post-2022. When available, we used transcripts provided by the podcast itself. If no transcripts were available, we downloaded the audio and transcribed it using Python (Python Software Foundation 2024) and OpenAI’s Whisper (Radford et al. 2022).

We concluded data collection once we had amassed about 11 million words for each time period (22.07 million words in total). In total, we included 1326 episodes, and a complete list of the podcasts included in our study is provided in Table 2, including well-known examples such as Lex Fridman and Stem Talk. Due to copyright restrictions, we can only share snippets of the transcripts, which are available on our GitHub. However, our methodology is fully replicable: the original audio and transcripts are publicly accessible, and all scripts used in the analysis are also available on our GitHub. The details of our exact technical setup can be found in Appendix A.

Data processing and analysis: We analyze lemmata, i.e. inflected forms are analyzed as a single item (e.g., ‘delves’ and ‘delved’ have the common lemma ‘delve’). Further, we distinguish part-of-speech (POS) categories (as per Jurafsky and Martin 2024); for example, the word “lead” may refer to the noun “lead,” a physical material, or the verb “lead,” which holds an entirely different meaning. Accordingly, we lemmatized and POS-tagged all transcripts using spaCy (Montani et al. 2023; tagging took about 2 hours). For the analysis, we compare occurrences per million (OPM) across the two time periods: pre-2022 and post-2022. Our main analysis assesses whether overall lexical usage has shifted by computing the weighted mean of the log frequency ratios and testing whether this change differs from zero using a z-test. To prevent instability due to low-frequency items, we applied Laplace smoothing by adding 0.5 to all counts before computing the log of the post-/pre-2022 frequency ratio. The weighting accounts for differences in frequency magnitude and statistical reliability across items: weights were defined as the inverse of the estimated variance, calculated as $1 / (1/f_{\text{post}} + 1/f_{\text{pre}})$. To regularize the undue influence of extremely frequent items, we capped all weights at 20. The resulting z-test yields a robust group-level inference (Bland 2015).

As a baseline, we also analyze synonyms identified in

Galpin, Anderson, and Juzek (2025). For instance, for “delve,” the semantically related verbs “explore,” “investigate,” “examine,” “research,” and “probe” were identified as synonyms, for which we then examine changes in their use accordingly. A full list of baseline words can be found in Tables 3 and 4 (relevant non-occurring words are listed in Appendix B).

We additionally analyze individual lemmata using chi-square contingency table tests, which are appropriate for our data structure (Haslwanter 2016), given the independent, unpaired samples and large dataset size. However, because we conduct dozens of tests (especially for the baseline items) and expect a number of true effects in both directions, applying a strict alpha-level correction (e.g., Bonferroni) would be overly conservative and obscure meaningful effects. Thus, the individual results should be interpreted with caution: at an uncorrected alpha level of 0.05, some false positives are expected.

Results

The results for the 20 words that we analyzed are presented in Table 1 and Figure 2. Figure 2 gives the logged proportional change, to account for the fact that proportional increases and decreases are asymmetrical in raw form. For example, a doubling (+100%) and a halving (-50%) are not numerically equivalent in raw form, but become symmetric when expressed in logarithmic space.

For the target words (corresponding to Figure 2, left), the weighted mean analysis reveals a moderate overall increase in usage (weighted log-ratio mean = 0.210). This change is statistically significant ($z = 3.725$, $p < 0.001$). Looking at the target words individually, we observe that of the 20 target words in our dataset, 6 show a decline in usage between the pre- and post-2022 periods. Of these, 2 changes are statistically significant according to the chi-square contingency table test: “crucial” and “realm.” 14 words show an increase in usage, 5 of which are statistically significant: “significant,” “align,” “strategically,” “boast,” and “surpass.”

By contrast, the overall analysis of the baseline words (Figure 2, right) shows only a very small, non-significant increase (weighted log-ratio mean = 0.033, $z = 1.277$, $p > 0.05$). Analyzing the baseline words individually, we observe that out of the 117 synonyms included in our analysis, only 87 occur in our dataset (for the same reason as with our target words: several of these items have extremely low frequencies to begin with, even in formal writing). For the majority of these (63 out of 87), changes in frequency are relatively moderate, remaining below +/-30%. Notably, the overall direction of change is balanced: approximately half of the words show a decrease in usage (45 out of 87), while the other half show an increase (42 out of 87). Nonetheless, several words exhibit considerable shifts between the pre- and post-2022 periods, some of which are statistically significant. The full results for the baseline words are provided in Tables 3 and 4.

Podcast Title	Number of Episodes Pre-22 & Post-22	Number of Tokens Pre-22 & Post-22
BBC Curious Cases	20 & 20	0.13m & 0.14m
Big Picture Science	15 & 15	0.13m & 0.15m
Big Technology Podcast	40 & 40	0.47m & 0.48m
Brain Inspired	20 & 20	0.30m & 0.36m
Conversations with Tyler	36 & 36	0.44m & 0.45m
CoreC	14 & 14	0.14m & 0.16m
EconTalk	64 & 64	0.94m & 0.88m
The Ezra Klein Show	60 & 60	0.80m & 0.77m
Lex Fridman	93 & 93	2.25m & 3.61m
Machine Learning Street Talk	30 & 30	0.69m & 0.59m
Mindscape	36 & 36	0.66m & 0.59m
Ologies with Alie Ward	67 & 67	1.07m & 1.22m
People Behind the Science	16 & 16	0.40m & 0.15m
Radiolab	61 & 61	0.67m & 0.61m
STEM Talk	29 & 29	0.42m & 0.44m
Talk Nerdy with Cara Santa Maria	31 & 31	0.43m & 0.30m
Very Bad Wizards	30 & 30	0.57m & 0.66m

Table 2: The list of podcasts analyzed for this work, including the number of episodes and token count (in millions) per podcast.

Discussion

General discussion

Overall, we observe a significant positive tendency, which stands in contrast to the baseline words, for which changes are roughly balanced between increases and decreases (with only a very small, non-significant upward shift). This allows us to reject the null hypothesis that there is no notable increase. However, the alternative hypothesis must be interpreted with nuance. We observe an overall moderate uptick, with most AI-associated words increasing in usage. This suggests that human word choices have, moderately and selectively, converged with those of Large Language Models. However, given the short time frame and the number of items showing increases, this could signal the onset of a remarkable shift in language use.

There are several observations to make. First, in contrast to the sharp spikes in written usage of these words (see Figure 1), the effect in spoken language may be less strong than one might have expected. This contrast suggests that the increase in written texts may be largely driven by tool usage, rather than a rapid shift in the human language system (though the latter cannot fully be ruled out, it would be unprecedented in scale and would require substantially stronger evidence).

Notably, the word “delve” did not show a significant increase; its rise was smaller than anticipated. Even more surprising is that “realm,” despite being frequently discussed in the literature as a characteristic LLM word (e.g. Liu and Bu 2024), significantly decreased in usage. Among the words that did show significant increases, “align” needs further analysis, as its rise may be linked to a broader shift in discourse, viz. the growing conversation around aligning LLMs with human values.

For the observed upticks, particularly the significant ones such as “surpass,” it is tempting to conclude that LLMs

have caused these changes. Given that many of the words heavily overused by AI are also showing increased usage in human language, it is plausible to conjecture a link. However, the question of causation remains open. As noted by Matsui (2024); Juzek and Ward (2025), many of the words now considered overused by AI were already trending upward in prior years (see pre-2022 trajectories illustrated in Figure 1). It is possible that these words have simply entered a phase of natural, rapid adoption, akin to the rise of expressions like “touch base,” “dude,” and “awesome” in the mid-2000s. As such, the observed changes could be in line with language change observed in the past (see Aitchison 2005; Bochkarev, Solovyev, and Wichmann 2014 for descriptions of conventional language change). Under this view, LLMs overuse words that were already increasing in frequency, patterns possibly learned during Learning from Human Feedback, but nonetheless act as amplifiers of ongoing language change. Even if not the ultimate source of these trends, machine-generated language influences human language, which is significant in its own right.

Either way, there is no counterfactual and we cannot definitively attribute the changes to LLM influence. Similar patterns might have emerged independently of LLMs. Therefore, additional evidence is required to support a stronger causal claim. Gaining such evidence is extremely difficult, it would require an extensive longitudinal study: one would need to track speakers over time, assess the degree and nature of their exposure to LLM-generated language, document their AI usage, and analyze how their lexical choices evolve.

As others have previously observed (Hataya, Bao, and Arai 2023; Briesch, Sobania, and Rothlauf 2023; Alemohammad et al. 2023), today’s language production becomes tomorrow’s LLM training data, the day after tomorrow’s LLM output, and eventually part of the linguistic input hu-

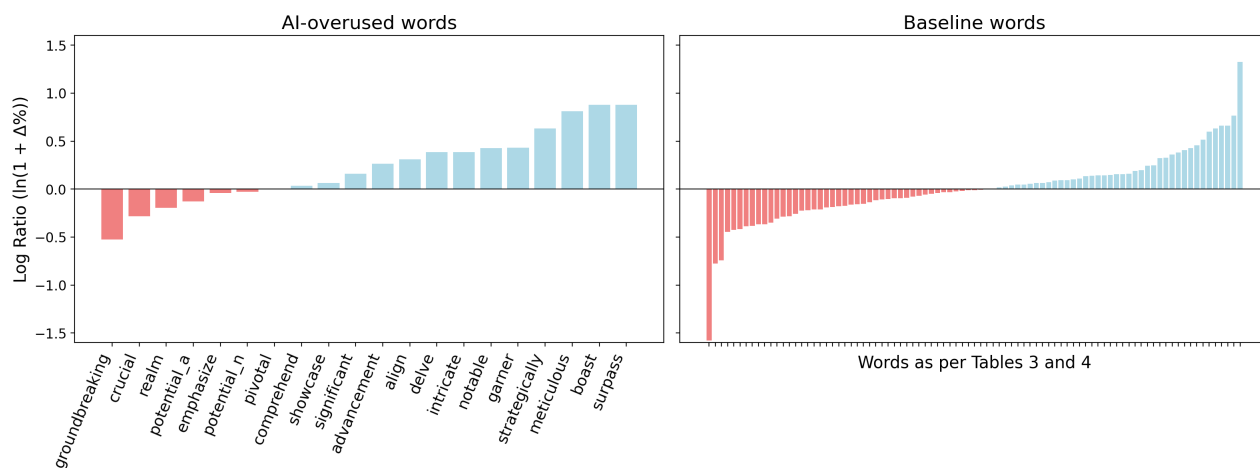


Figure 2: Pre- vs post-2022 logged proportional change for AI-associated words (left) vs baseline words (right).

mans encounter. There is the possibility that a (semi-)self-contained training-loop, as Briesch, Sobania, and Rothlauf (2023) put it, is forming, which comes with computationally negative consequences (Shumailov et al. 2023).

Limitations

There are several limitations to our work. First, a larger dataset would be advantageous. In particular, for the baseline words, we would expect to see somewhat less variation than currently observed, suggesting that some of the variation we do observe may be due to random fluctuations in a relatively small dataset. In fact, we initially planned to work with 5 million words per time period, but we realized that this would be too limited for meaningful statistical analysis. Still, we believe that the observed increase in AI-associated words offers an accurate reflection of a genuine underlying trend.

Second, the effects we observe likely do not reflect general population trends. Our focus on technology and science-related podcasts may amplify certain language changes – a strategic choice intended to increase the likelihood of capturing LLM-influenced language. Arguably, the results hold for a certain demographic (approximately: tech-affine individuals in the young to middle-aged range). However, it also means that the results are probably not representative of all spoken language. For example, an analysis of rural sermons would likely show no noticeable changes in the usage of the AI-associated words we tracked. Further, our dataset is drawn from English-language podcasts that are popular in the United States, but often feature international speakers. As such, the findings reflect general English usage in globally oriented, tech-focused discourse. Further research with more fine-grained data is needed to explore regional and varietal differences. That said, we do not think that these limitations are problematic. Our goal was to measure on-set language change, which often originates within specific subgroups, before diffusing more broadly across the general speaker community. Examples of this are the rise of “touch base” (originating in business contexts) or “dude” (youth

culture).

Regarding the dataset itself, some podcasts may include scripted promotional segments, which were especially difficult to isolate and remove when placed mid-episode. Additionally, podcast interviews may contain pre-scripted questions from the host; however, the ensuing discussions are, in most cases, spontaneous and unscripted. In both cases, these segments are extremely difficult to filter out automatically and would require an impractical amount of manual effort. In practice, however, they constitute only a small fraction of each episode.

Further, a qualitative analysis, linking underlying micro-level patterns to the observed macro-level trends would have nicely complemented our findings (among others, Gee 2014 motivates the value of qualitative analyses). Such work could be similar to Krielke (2024), who traced the mechanisms of syntactic change through detailed case studies, and could also take recent insights from longitudinal changes in the syntactic structure of Scientific English into account (for recent insights, see e.g. Krielke 2021; Krielke et al. 2022; Krielke 2023; Chen et al. 2024). Following up on our work, one could examine selected instances of words such as “surpass,” “boast,” and especially “align.” This would involve analyzing the contexts in which these terms are used, identifying recurring discourse frames (Fillmore 2006), and exploring whether particular speaker types or roles tend to produce them (Goffman 1981). Such a qualitative analysis could also help clarify developments that defied expectations: How were e.g. “realm” and “groundbreaking” used in academic writing prior to the arrival of LLMs? How has their usage developed since? And how does this contrast with their usage in spoken language? Extensive qualitative analyses constitute substantial undertakings in their own right, and we must leave these for future research.

Broader societal implications

Our central contribution lies in demonstrating that recent changes in language use may not be solely attributable to hu-

Word	Change (%)	OPM Pre-2022	OPM Post-2022	p ≤ 0.05
irreplaceable_ADJ	275.86	0.27	1.03	False
eminent_ADJ	114.78	0.64	1.38	False
territory_NOUN	93.88	14.47	28.05	True
exhibit_NOUN	93.46	1.56	3.01	True
admirable_ADJ	87.93	1.19	2.24	False
capability_NOUN	81.71	30.40	55.24	True
multifaceted_ADJ	67.05	0.82	1.38	False
probable_ADJ	57.86	2.29	3.61	False
notable_ADJ	53.13	2.47	3.79	False
achievable_ADJ	50.35	0.92	1.38	False
assemble_VERB	46.48	6.23	9.12	True
thorough_ADJ	43.19	1.92	2.75	False
reinforce_VERB	38.37	8.33	11.53	True
novel_ADJ	38.22	15.75	21.77	True
grasp_VERB	27.90	6.59	8.43	False
complexity_NOUN	27.71	55.58	70.98	True
visionary_ADJ	21.60	1.56	1.89	False
enhancement_NOUN	20.81	1.92	2.32	False
prospect_NOUN	17.14	6.68	7.83	False
stress_VERB	16.98	8.97	10.50	False
integrate_VERB	16.90	27.01	31.58	True
sophisticated_ADJ	15.68	22.98	26.59	False
coordinate_VERB	15.32	8.06	9.29	False
elaborate_ADJ	15.18	5.68	6.54	False
deliberately_ADV	14.85	8.24	9.46	False
remarkable_ADJ	14.35	24.91	28.48	False
critical_ADJ	11.46	51.64	57.56	False
vital_ADJ	10.55	6.23	6.88	False
original_ADJ	9.60	58.88	64.53	False
highlight_VERB	9.54	16.57	18.15	False
improvement_NOUN	9.20	23.72	25.90	False
complex_ADJ	7.57	86.62	93.18	False
vigilant_ADJ	6.49	1.37	1.46	False
showcase_VERB	6.49	1.37	1.46	False
convoluted_ADJ	5.71	1.47	1.55	False
possibility_NOUN	4.69	77.83	81.48	False
innovative_ADJ	4.64	8.06	8.43	False
merit_NOUN	4.09	5.95	6.19	False
complication_NOUN	2.78	5.86	6.02	False
renowned_ADJ	1.80	1.10	1.12	False
pivotal_ADJ	0.45	2.66	2.67	False
inventive_ADJ	0.23	1.37	1.38	False

Table 3: Increases in word frequency (occurrences per million; OPM) for the baselines synonyms, for pre-2022 vs. post-2024, including p-values (part 1).

mans using AI models as a tool. Rather, our findings suggest changes in the human language system itself, and a plausible possibility is that AI is beginning to directly influence the human language system (an interpretation that is supported by the observed overlap between AI-overused words and those increasing in human usage). This has broader relevance for ongoing debates around model alignment. There is a gap between model language behavior and human language behavior, which raises important questions about how to interpret this gap and about potential societal consequences of such misalignments. While we are concerned

with (mis)alignment of language behavior, this, of course, is directly linked to alignment discussions on social or ethical values, with the issues and mechanisms of misalignment manifestation and propagation being very similar.

One critical aspect we identify is what we call the “seep-in” effect. There is a gap between AI behavior and human values and behavior, and based on our results, when it comes to language, there is a plausible possibility that AI linguistic patterns are beginning to *seep into* the human language system. That is, a word or expression that may not have reflected a speaker’s own linguistic preference may

Word	Change (%)	OPM Pre-2022	OPM Post-2022	p ≤ 0.05
development_NOUN	-0.81	70.87	70.29	False
possible_ADJ	-1.05	260.50	257.77	False
research_VERB	-1.26	16.21	16.00	False
distinction_NOUN	-1.94	33.60	32.95	False
understand_VERB	-2.36	587.84	573.97	False
presentation_NOUN	-3.10	11.72	11.36	False
explore_VERB	-3.31	69.41	67.11	False
demonstration_NOUN	-4.19	9.34	8.95	False
precise_ADJ	-4.97	16.21	15.40	False
appreciate_VERB	-5.42	83.69	79.16	False
discipline_NOUN	-6.79	22.80	21.25	False
exhibit_VERB	-7.50	5.86	5.42	False
probe_VERB	-8.72	6.41	5.85	False
theoretical_ADJ	-8.98	38.00	34.59	False
sphere_NOUN	-9.13	16.66	15.14	False
display_VERB	-9.95	6.59	5.94	False
perceive_VERB	-10.06	29.94	26.93	False
consciously_ADV	-11.07	5.13	4.56	False
key_ADJ	-12.89	60.25	52.48	True
highlight_NOUN	-14.20	4.21	3.61	False
progress_NOUN	-14.74	67.21	57.30	True
compile_VERB	-14.88	7.78	6.63	False
field_NOUN	-16.01	232.02	194.88	True
valuable_ADJ	-16.47	37.08	30.97	True
feature_NOUN	-16.91	112.35	93.35	True
forum_NOUN	-17.57	5.22	4.30	False
essential_ADJ	-19.12	28.93	23.40	True
aggregate_VERB	-19.15	3.94	3.18	False
likelihood_NOUN	-19.72	9.43	7.57	False
display_NOUN	-20.18	8.52	6.80	False
purposefully_ADV	-22.62	1.56	1.20	False
crucial_ADJ	-24.56	19.05	14.37	True
investigate_VERB	-25.01	27.65	20.74	True
synchronize_VERB	-26.46	2.11	1.55	False
excellence_NOUN	-29.53	5.13	3.61	False
collect_VERB	-30.71	59.97	41.56	True
amass_VERB	-30.76	1.74	1.20	False
distinguished_ADJ	-31.96	2.66	1.81	False
examine_VERB	-32.00	13.92	9.46	True
feasible_ADJ	-34.02	8.61	5.68	True
probability_NOUN	-34.75	52.74	34.42	True
outstanding_ADJ	-36.02	4.30	2.75	False
domain_NOUN	-52.32	55.40	26.41	True
sophistication_NOUN	-53.98	4.49	2.06	True
convolution_NOUN	-79.37	7.51	1.55	True

Table 4: Decrease in word frequency (occurrences per million; OPM) for the baselines synonyms, for pre-2022 vs. post-2024, including p-values (part 2).

nonetheless become part of their language system, simply through repeated exposure (Tomasello 2005; Bybee 2006) – exposure now increasingly impacted by AI-generated language (as well as by changes in the writing process itself, which might become increasingly AI-assisted (Floridi 2025)). While such a seep-in may seem relatively benign in the case of lexical choices, it raises more serious concerns when considered in relation to political or social beliefs of

misaligned models or even models by malicious actors.

Work on the features of AI-associated language also has potential relevance for AI detection (for ongoing efforts, see e.g. Jin et al. 2025; Koziev 2025; Schmalz and Tack 2025), a task that remains challenging (Sadasivan et al. 2023; Weber-Wulff et al. 2023). Human language is remarkably complex (recently, e.g. Wang et al. 2023; Wu et al. 2024; Zhang, Ouni, and Eger 2024; Clark et al. 2025), which could be

a contributing factor to this challenge. Advancing detection methods will require both high-quality tools and robust benchmark datasets (recent advances include Warstadt et al. 2020; Srivastava et al. 2023; Lücking et al. 2024; Zhang, Ouni, and Eger 2024; Weissweiler, Mahowald, and Goldberg 2025; Jumelet, Weissweiler, and Bisazza 2025).

One might ask whether these linguistic developments are fundamentally new. As noted in the Introduction, technology and language have long had a reciprocal relationship. From the printing press to the rise of digital media, linguistic behavior has evolved alongside changes in communication technology. What distinguishes the current moment is the pace and scale of language change. If the adoption of Large Language Models continues to grow globally, then a broader and potentially homogenizing effect on language is plausible (also cf. Bender et al. 2021). Over the coming decades, we may see convergence around an AI-influenced register, with all its linguistic implications.

An alternative interpretation is that these lexical shifts reflect a trajectory of natural language change, with AI serving as an accelerator or amplifier. While this cannot be ruled out, it is difficult to test empirically. Anecdotally, we observe that human choice still plays a role. For instance, the word “delve,” which saw widespread usage in academic writing in 2023 and 2024 (as per Introduction), is falling out of favor. In Q1 2025, occurrences of “delve” in PubMed abstracts dropped to about 15 occurrences per one million words, from about 22 occurrences per one million in the previous year. Thus, there is a *delve-pushback*, a reversal in language usage, over a very short period of time, for which we can only speculate about the underlying motivation. It may be that speakers became aware of how “delve” had seeped into their language system and reacted against it; and/or “delve” may have become associated with AI-generated language and, in some contexts, particularly scientific writing, now carries a degree of stigma.

What is certainly new, however, is the problem of human-authorship indeterminacy. For an increasing number of texts, it is no longer possible to determine whether the content was genuinely produced by a human. This has profound implications for linguistic research and computational linguistics in particular, where language data has traditionally served as a proxy for human cognition and behavior. Our study exemplifies this disruption: what was once a relatively straightforward task (measuring trends in human language production), now required substantial methodological effort to ensure the validity of the data. As authorship indeterminacy becomes more widespread, it will complicate not only linguistic research, but also any scientific field that relies on natural language as a proxy for human behavior and cognition.

Conclusion

Our work contributes to research on AI’s impact on human language, specifically addressing whether recent changes in language use stem from direct AI tool usage or whether genuinely human-produced language is beginning to follow patterns associated with AI language, particularly the overuse of certain words such as “delve” and “intricate.” In this view,

such words may be gaining prominence in people’s mental language systems, leading to increased usage over time.

To investigate this, we constructed a dataset of unscripted spoken language as a proxy for spontaneous spoken language. We divided the dataset into pre- and post-2022 periods and examined changes in the frequency of words identified in the literature as commonly overused by Large Language Models. Our findings reveal a mostly positive trend in the usage of these AI-associated words, in contrast to a balanced pattern among baseline words, where increases and decreases are roughly evenly split. While these early findings are promising, larger datasets and further analysis are needed to substantiate the observed trends.

A link between changes in human language usage and AI is plausible, given the observed overlap in affected vocabulary. However, we leave open the question of whether the observed increases reflect AI-induced language change or natural language change. This distinction is difficult to determine and remains an ongoing challenge. However, it is a crucial one. If the patterns align with established processes of language change, viz. human preferences driving human language, then there may be little cause for concern. But if the language favored by AI reflects preferences not rooted in natural human usage, then we are witnessing a moment of model misalignment. In particular, the overuse of certain words by LLMs is widely believed to stem from the training procedure known as Learning from Human Feedback. While Learning from Human Feedback is designed to align models with human preferences, it may in practice amplify stylistic features that do not reflect the expectations of most users. In this sense, upstream alignment choices can introduce downstream linguistic biases. This matters not only for understanding lexical shifts but also because the mechanisms of linguistic (mis)alignment closely mirror broader concerns of (mis)alignment in AI ethics and fairness.

Of particular concern is what one might call the “seep-in” effect, the subconscious adoption of AI-driven linguistic or even ethical tendencies, even when these do not reflect the speaker’s original preferences. If such preferences gradually permeate human communication, this could mark a novel and potentially profound shift in the trajectory of language evolution. Lastly, as our study exemplifies, AI is not only reshaping language use, it is disrupting how we study language itself. For many written texts, human authorship is no longer easily determined, which makes it increasingly difficult to infer human behavior from language patterns.

Appendix A: Computing Resources

This appendix details the computing environment used to run the scripts and analyses in this study.

System Specifications

- **Machine:** 2024 Thelio Custom Machine
- **Operating System:** Ubuntu 24.04.1 LTS
- **CPU:** Intel Core i7-14700K
- **GPU:** NVIDIA GeForce RTX 3090

- **RAM:** 128 GB

Software Environment

- **Python Version:** 3.12.3
 - **Key Libraries:**
 - openai-whisper == 20240930
 - spacy == 3.8.2
 - scipy == 1.15.2
-

Runtime Notes

- **Total transcription time / Whisper runtime:** 92 hours
- **spaCy tagging time:** 2 hours
- **Average analysis runtime:** 4 minutes per 1 hour of transcription

Additional Notes

This environment uses a 2024 Thelio Custom Machine equipped with an Intel Core i7-14700K processor and an NVIDIA GeForce RTX 3090 GPU.

Appendix B: Full Results for the 117 Baseline Words

Out of the 117 synonyms from the literature, 87 did occur in our dataset; the results for these can be found in Tables 3 and 4. The 30 words not occurring in our dataset are: systematize.VERB, harmonize.VERB, underscore.VERB, laudable.ADJ, praiseworthy.ADJ, noteworthy.ADJ, meritorious.ADJ, understanding.ADJ, grasping.ADJ, perceiving.ADJ, interpreting.ADJ, discerning.ADJ, accentuate.VERB, highlighting.ADJ, underscoring.ADJ, stressing.ADJ, prioritizing.ADJ, collected.ADJ, acquired.ADJ, obtained.ADJ, assembled.ADJ, amassed.ADJ, intricateness.NOUN, inestimable.ADJ, priceless.ADJ, immeasurable.ADJ, diligent.ADJ, scrupulous.ADJ, tactically.ADV, methodically.ADV.

References

- Aitchison, J. 2005. Language Change. In Cobley, P., ed., *The Routledge Companion to Semiotics and Linguistics*, chapter 7. London: Routledge.
- Akintande, O. J. 2023. Algorithmic Bias: When stigmatization becomes a perception: The stigmatized become endangered. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 966–971.
- Al-Sharqi, L.; and Abbasi, I. S. 2020. The Influence of Technology on English Language and Literature. *English Language Teaching*, 13(7): 1.
- Alemohammad, S.; Casco-Rodriguez, J.; Luzi, L.; Humayun, A. I.; Babaei, H.; LeJeune, D.; Siahkoohi, A.; and Baraniuk, R. G. 2023. Self-consuming generative models go mad. *arXiv preprint arXiv:2307.01850*.
- Ali, J.; Kleindessner, M.; Wenzel, F.; Budhathoki, K.; Cevher, V.; and Russell, C. 2023. Evaluating the fairness of discriminative foundation models in computer vision. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 809–833.
- Amato, R.; Lacasa, L.; Díaz-Guilera, A.; and Baronchelli, A. 2018. The dynamics of norm change in the cultural evolution of language. *Proceedings of the National Academy of Sciences*, 115(33): 8260–8265.
- Bao, T.; Zhao, Y.; Mao, J.; and Zhang, C. 2025. Examining linguistic shifts in academic writing before and after the launch of ChatGPT: a study on preprint papers. *Scientometrics*, 1–31.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623.
- Biber, D. 1991. *Variation across speech and writing*. Cambridge university press.
- Bland, M. 2015. *An introduction to medical statistics*. Oxford university press.
- Blodgett, S. L.; Barocas, S.; Daumé III, H.; and Wallach, H. 2020. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050*.
- Bochkarev, V.; Solovyev, V.; and Wichmann, S. 2014. Universals versus historical contingencies in lexical evolution. *Journal of the Royal Society Interface*, 11(101): 20140841.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Briesch, M.; Sobania, D.; and Rothlauf, F. 2023. Large language models suffer from their own output: An analysis of the self-consuming training loop. *arXiv preprint arXiv:2311.16822*.
- Bybee, J. L. 2006. From usage to grammar: The mind’s response to repetition. *Language*, 82(4): 711–733.
- Chen, Y.; Zhao, W.; Breitbarth, A.; Stoeckel, M.; Mehler, A.; and Eger, S. 2024. Syntactic language change in english and german: Metrics, parsers, and convergences. *arXiv preprint arXiv:2402.11549*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Clark, H. H. 1996. *Using Language*. Cambridge University Press.
- Clark, T. H.; Poliak, M.; Regev, T.; Haskins, A. J.; Gibson, E.; and Robertson, C. 2025. The relationship between surprisal, prosody, and backchannels in conversation reflects intelligibility-oriented pressures. *PsyArXiv preprints*.
- Comas-Forgas, R.; Koulouris, A.; and Kouis, D. 2025. ‘AI-navigating’ or ‘AI-sinking’? An analysis of verbs in research articles titles suspicious of containing AI-generated/assisted content. *Learned Publishing*, 38(1): e1647.

- Crystal, D. 2009. Txtng: The Gr8 Db8.
- Djalolovna, M. S. 2025. The Influence of Technology on Spoken Language: Texting, Voice Assistants, and Speech Patterns. *Web of Technology: Multidimensional Research Journal*, 3(1): 34–38.
- Eisenstein, E. 1980. *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*. Cambridge University Press.
- Fedorenko, E.; Piantadosi, S. T.; and Gibson, E. A. F. 2024. Language is primarily a tool for communication rather than thought. *Nature*, 630(8017): 575–586.
- Fillmore, C. J. 2006. Frame Semantics. In Geeraerts, D.; Dirven, R.; and Taylor, J. R., eds., *Cognitive Linguistics: Basic Readings*, volume 34 of *Cognitive Linguistics Research*, 373–400. Berlin: Mouton de Gruyter.
- Floridi, L. 2025. Distant Writing: Literary Production in the Age of Artificial Intelligence. *Minds and Machines*, 35(3): 1–26.
- Floridi, L.; Cows, J.; Beltrametti, M.; Chatila, R.; Chazerand, P.; Dignum, V.; Luetge, C.; Madelin, R.; Pagallo, U.; Rossi, F.; et al. 2018. AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds and machines*, 28: 689–707.
- Gabriel, I. 2020. Artificial intelligence, values, and alignment. *Minds and machines*, 30(3): 411–437.
- Galpin, R.; Anderson, B.; and Juzek, T. S. 2025. Lexical underrepresentation matters as much as lexical overrepresentation. *arXiv preprint arXiv:2501.20200*.
- Gee, J. P. 2014. *An introduction to discourse analysis: Theory and method*. routledge.
- Geng, M.; Chen, C.; Wu, Y.; Chen, D.; Wan, Y.; and Zhou, P. 2024. The Impact of Large Language Models in Academia: from Writing to Speaking. *arXiv preprint arXiv:2409.13686*.
- Ghosh, S.; and Caliskan, A. 2023. Chatgpt perpetuates gender bias in machine translation and ignores non-gendered pronouns: Findings across bengali and five other low-resource languages. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 901–912.
- Goffman, E. 1981. *Forms of talk*. University of Pennsylvania Press.
- Grabowicz, P.; Perello, N.; and Takatsu, K. 2023. Learning from Discriminatory Training Data. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 752–763.
- Gray, A. 2024. ChatGPT” contamination”: estimating the prevalence of LLMs in the scholarly literature. *arXiv preprint arXiv:2403.16887*.
- Hagen, L.; Huo, J.; and Nguyen, A. 2025. Elon Musk’s AI Chatbot, Grok, Started Calling Itself ‘MechaHitler’. *NPR*.
- Hartvigsen, T.; Gabriel, S.; Palangi, H.; Sap, M.; Ray, D.; and Kamar, E. 2022. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. *arXiv preprint arXiv:2203.09509*.
- Haslwanter, T. 2016. An introduction to statistics with python. *With applications in the life sciences*. Switzerland: Springer International Publishing.
- Hataya, R.; Bao, H.; and Arai, H. 2023. Will large-scale generative models corrupt future datasets? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20555–20565.
- Helen, J. W. 2010. The Effect of Internet Devices on Children’s Language Development. *Contemporary Issues in Communication Science Disorders*, 37: 141–148.
- Hendrycks, D.; Burns, C.; Basart, S.; Critch, A.; Li, J.; Song, D.; and Steinhardt, J. 2020. Aligning ai with shared human values. *arXiv preprint arXiv:2008.02275*.
- Huang, B.; Chen, C.; and Shu, K. 2025. Authorship attribution in the era of llms: Problems, methodologies, and challenges. *ACM SIGKDD Explorations Newsletter*, 26(2): 21–43.
- Jaiswal, S. 2024. Uncovering Gender Biases in Human-AI Platforms. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7(2), 21–22.
- Jin, H.; Sung, S.; Park, S.; Baik, S.; and Han, Y.-S. 2025. TRAPDOC: Deceiving LLM Users by Injecting Imperceptible Phantom Tokens into Documents. *arXiv preprint arXiv:2506.00089*.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- Jumelet, J.; Weissweiler, L.; and Bisazza, A. 2025. MultiB-LiMP 1.0: A massively multilingual benchmark of linguistic minimal pairs. *arXiv preprint arXiv:2504.02768*.
- Jurafsky, D.; and Martin, J. H. 2024. *Speech and Language Processing*. Online draft. 3rd ed. draft, Feb 3, 2024 release.
- Juzek, T. S.; and Ward, Z. B. 2025. Why Does ChatGPT ”Delve” So Much? Exploring the Sources of Lexical Overrepresentation in Large Language Models. In *Proceedings of the 31st International Conference on Computational Linguistics (COLING 2025)*.
- Keswani, V.; Conitzer, V.; Heidari, H.; Borg, J. S.; and Sinnott-Armstrong, W. 2024. On the Pros and Cons of Active Learning for Moral Preference Elicitation. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 711–723.
- Kobak, D.; Márquez, R. G.; Horvát, E.-Á.; and Lause, J. 2024. Delving into ChatGPT usage in academic writing through excess vocabulary. *arXiv preprint arXiv:2406.07016*.
- Koppenburg, P. 2024. Tweet on 01 April 2024. <https://x.com/PKoppenburg/status/1774757167045788010>. Accessed: 2025-01-23.
- Kotek, H.; Dockum, R.; and Sun, D. 2023. Gender bias and stereotypes in large language models. In *Proceedings of the ACM collective intelligence conference*, 12–24.
- Koziev, I. 2025. Detecting Spelling and Grammatical Anomalies in Russian Poetry Texts. *arXiv preprint arXiv:2505.04507*.

- Krielke, M.-P. 2021. Relativizers as markers of grammatical complexity: A diachronic, cross-register study of English and German. *Bergen Language and Linguistics Studies*, 11(1): 91–120.
- Krielke, M.-P. 2023. Optimizing scientific communication: the role of relative clauses as markers of complexity in English and German scientific writing between 1650 and 1900. *Saarländische Universitäts- und Landesbibliothek*.
- Krielke, M.-P. 2024. Cross-linguistic Dependency Length Minimization in scientific language: Syntactic complexity reduction in English and German in the Late Modern period. *Languages in Contrast*, 24(1): 133–163.
- Krielke, M.-P.; Talamo, L.; Fawzi, M.; and Knappen, J. 2022. Tracing syntactic change in the scientific genre: Two Universal Dependency-parsed diachronic corpora of scientific English and German. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, 4808–4816.
- Labov, W. 2011. *Principles of linguistic change, volume 3: Cognitive and cultural factors*, volume 3. John Wiley & Sons.
- Leidinger, A.; and Rogers, R. 2024. How are LLMs mitigating stereotyping harms? Learning from search engine studies. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 839–854.
- Li, Y.; Breithaupt, F.; Hills, T.; Lin, Z.; Chen, Y.; Siew, C. S. W.; and Hertwig, R. 2023. How Cognitive Selection Affects Language Change. *Proceedings of the National Academy of Sciences*, 121(1). PNAS Direct Submission. Published December 27, 2023.
- Liang, W.; Zhang, Y.; Wu, Z.; Lepp, H.; Ji, W.; Zhao, X.; Cao, H.; Liu, S.; He, S.; Huang, Z.; et al. 2024. Mapping the increasing use of llms in scientific papers. *arXiv preprint arXiv:2404.01268*.
- Liu, J.; and Bu, Y. 2024. Towards the relationship between AIGC in manuscript writing and author profiles: evidence from preprints in LLMs. *arXiv preprint arXiv:2404.15799*.
- Lücking, A.; Abrami, G.; Hammerla, L.; Rahn, M.; Baumartz, D.; Eger, S.; and Mehler, A. 2024. Dependencies over times and tools (DoTT). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 4641–4653.
- Matsui, K. 2024. Delving into PubMed Records: Some Terms in Medical Writing Have Drastically Changed after the Arrival of ChatGPT. *medRxiv*, 2024–05.
- Miller, J. E.; and Weinert, R. 1998. *Spontaneous spoken language: Syntax and discourse*. Oxford University Press.
- Mittelstadt, B. D.; Allo, P.; Taddeo, M.; Wachter, S.; and Floridi, L. 2016. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2): 2053951716679679.
- Montani, I.; Honnibal, M.; Boyd, A.; Landeghem, S. V.; and Peters, H. 2023. explosion/spaCy: v3.7.2: Fixes for APIs and requirements. Version v3.7.2, Zenodo.
- Naik, R.; and Nushi, B. 2023. Social biases through the text-to-image generation lens. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 786–808.
- Narayanan Venkit, P.; Gautam, S.; Panchanadikar, R.; Huang, T.-H.; and Wilson, S. 2023. Unmasking nationality bias: A study of human perception of nationalities in ai-generated articles. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 554–565.
- Neff, G.; and Nagy, P. 2016. Talking to bots: Symbiotic agency and the case of Tay. *International Journal of Communication*, 10: 4915–4931.
- Norhashim, H.; and Hahn, J. 2024. Measuring Human-AI Value Alignment in Large Language Models. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 1063–1073.
- Omiye, J. A.; Lester, J. C.; Spichak, S.; Rotemberg, V.; and Daneshjou, R. 2023. Large language models propagate race-based medicine. *NPJ Digital Medicine*, 6(1): 195.
- Omrani Sabbaghi, S.; Wolfe, R.; and Caliskan, A. 2023. Evaluating biased attitude associations of language models in an intersectional context. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 542–553.
- OpenAI. 2024. Tweet on 08 April 2024. <https://x.com/ChatGPTapp/status/1777221658807521695>. Accessed: 2024-08-12.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pack, A.; and Maloney, J. 2023. Potential Affordances of Generative AI in Language Education: Demonstrations and an Evaluative Framework. *Teaching English with Technology*, 23(2).
- Perez, E.; Ringer, S.; Lukosiute, K.; Nguyen, K.; Chen, E.; Heiner, S.; Pettit, C.; Olsson, C.; Kundu, S.; Kadavath, S.; et al. 2023. Discovering language model behaviors with model-written evaluations. In *Findings of the Association for Computational Linguistics: ACL 2023*, 13387–13434.
- Python Software Foundation. 2024. Python 3.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2022. Whisper. <https://github.com/openai/whisper>.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Russell, S.; Dewey, D.; and Tegmark, M. 2015. Research priorities for robust and beneficial artificial intelligence. *AI magazine*, 36(4): 105–114.
- Sadasivan, V. S.; Kumar, A.; Balasubramanian, S.; Wang, W.; and Feizi, S. 2023. Can AI-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*.
- Sandra, G. 2011. 7. Greiffenstern, Sandra. 2010. The influence of computers, the internet and computer-mediated communication on everyday English. *English and American Studies in German*, 2011(1): 16–18.

- Schmalz, V.; and Tack, A. 2025. Can GPTZero’s AI Vocabulary Distinguish Between LLM-Generated and Student-Written Essays? In Kochmar, E.; Alhafni, B.; Bexte, M.; Burstein, J.; Horbach, A.; Laarmann-Quante, R.; Tack, A.; Yaneva, V.; and Yuan, Z., eds., *Proceedings of the 20th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2025)*, 937–952. Vienna, Austria: Association for Computational Linguistics. ISBN 979-8-89176-270-1.
- Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Gal, Y.; Papernot, N.; and Anderson, R. 2023. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*.
- Srivastava, A.; Rastogi, A.; Rao, A.; Shoeb, A. A.; Abid, A.; Fisch, A.; Brown, A. R.; Santoro, A.; Gupta, A.; Garriga-Alonso, A.; et al. 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on machine learning research*.
- Tomasello, M. 2005. *Constructing a language: A usage-based theory of language acquisition*. Harvard university press.
- Traugott, E. C. 1985. On regularity in semantic change. *Journal of Literary Semantics*, 14(3): 155–173.
- Vanisree, M.; Sharma, S.; Singh, L.; Kantharaja, K. P.; and Sri Nidhya, G. 2024. Effects of Technology and Social Media on English Language Use and Acquisition: A Study on Language Variation and Change. *Nanotechnology Perceptions*, 20(S7): 490–497.
- Wang, G.; Wang, H.; Sun, X.; Wang, N.; and Wang, L. 2023. Linguistic complexity in scientific writing: A large-scale diachronic study from 1821 to 1920. *Scientometrics*, 128(1): 441–460.
- Warstadt, A.; Parrish, A.; Liu, H.; Mohananey, A.; Peng, W.; Wang, S.-F.; and Bowman, S. R. 2020. BLiMP: The benchmark of linguistic minimal pairs for English. *Transactions of the Association for Computational Linguistics*, 8: 377–392.
- Weber-Wulff, D.; Anohina-Naumeca, A.; Bjelobaba, S.; Foltýnek, T.; Guerrero-Dib, J.; Popoola, O.; Šigut, P.; and Waddington, L. 2023. Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(1): 1–39.
- Weissweiler, L.; Mahowald, K.; and Goldberg, A. 2025. Linguistic generalizations are not rules: Impacts on evaluation of LMs. *arXiv preprint arXiv:2502.13195*.
- Welbl, J.; Glaese, A.; Uesato, J.; Dathathri, S.; Mellor, J.; Hendricks, L. A.; Anderson, K.; Kohli, P.; Coppin, B.; and Huang, P.-S. 2021. Challenges in detoxifying language models. *arXiv preprint arXiv:2109.07445*.
- Wilson, K.; and Caliskan, A. 2024. Gender, race, and intersectional bias in resume screening via language model retrieval. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 1578–1590.
- Wu, M.; and Aji, A. F. 2025. Style Over Substance: Evaluation Biases for Large Language Models. In Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B. D.; and Schockaert, S., eds., *Proceedings of the 31st International Conference on Computational Linguistics*, 297–312. Abu Dhabi, UAE: Association for Computational Linguistics.
- Wu, Y.; Bizzoni, Y.; Moreira, P.; and Nielbo, K. 2024. Perplexing canon: A study on GPT-based perplexity of canonical and non-canonical literary works. In *Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2024)*, 172–184.
- Yakura, H.; Lopez-Lopez, E.; Brinkmann, L.; Serna, I.; Gupta, P.; and Rahwan, I. 2024. Empirical evidence of Large Language Model’s influence on human spoken communication. *arXiv preprint arXiv:2409.01754*.
- Zamaraeva, O.; Flickinger, D.; Bond, F.; and Gómez-Rodríguez, C. 2025. Comparing LLM-generated and human-authored news text using formal syntactic theory. *arXiv preprint arXiv:2506.01407*.
- Zhang, R.; Ouni, J.; and Eger, S. 2024. Cross-lingual cross-temporal summarization: Dataset, models, evaluation. *Computational Linguistics*, 50(3): 1001–1047.
- Zhang, X.; Xiong, W.; Chen, L.; Zhou, T.; Huang, H.; and Zhang, T. 2024. From lists to emojis: How format bias affects model alignment. *arXiv preprint arXiv:2409.11704*.
- Zhang, Y. 2014. A Review of Principles of Linguistic Change: Cognitive and Cultural Factors. *Open Journal of Modern Linguistics*, 4: 65–68.
- Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.