Can Generative AI Write Effective Abstracts for the ETC Community?

Crista Mohammed, Azim Abdool and Janine Jaggernauth

April 2, 2026

Can Generative AI Write Effective Abstracts for the ETC Community?

Crista Mohammed
The University of the West Indies
Crista.Mohammed@uwi.edu

Azim Abdool
The University of the West Indies
Azim.Abdool@uwi.edu

Janine Jaggernauth
The University of the West Indies
Janine.Jaggernauth@uwi.edu

Abstract - The proliferation of generative AI (GAI) raises the question as to whether writing-from-sources tasks remain valid credit-bearing assessments of learning and writing achievement. This study probed whether select GAIs, using popular science articles as input texts, produce effective abstracts for the engineering, technology, and computing (ETC) discourse community. We input two articles into three GAIs: M365 Copilot (GPT-5 model); ChatGPT (GPT-4 model); and Perplexity (GPT-5.1 model). Using two prompts—a basic single-shot prompt and a single-shot prompt with a macro-structure for abstracts—we generated 12 abstracts. We present two sets of data: numerical scoring of the GAI-generated abstracts and outstanding discourse features of the abstracts. Of the 12 abstracts, all received a passing score of 50% or higher. More guidance in prompting did not lead to overall improved abstracts. The macro-structure in prompt two eroded the accuracy of abstracts produced by Copilot and ChatGPT. In terms of discourse quality, while the GAIs maintained an academic register, their credibility as authors was questionable. Bear in mind that authors of studies are understood to be the authors of the abstracts of those studies. All three GAIs were unable to inhabit the role of author in their abstracts by avoiding the first person; making overt references to third parties; and using active-voice constructions. While this study does not recommend that for-credit writing-from-sources tasks be retired, it does present evidence of the need to re-design this type of assessment.

Index Terms – Abstract writing, engineering writing, generative AI

INTRODUCTION

Writing from sources is an inescapable activity for university students. Two common writing-from-sources genres which have been the subject of writing research are abstracts [1-3], and literature reviews [4-6]. Instructors have relied on writing-from-sources activities to achieve several learning outcomes. Among these are having students learn new content [1, 2, 7]; develop argumentation skills [5, 8]; produce target genres [1, 2, 7]; and develop facility with disciplinary vocabulary and style [1, 2, 5].

The proliferation of free-to-use generative AI (GAI) raises the question as to whether writing-from-sources tasks remain valid assessments for testing learning and writing achievement. This question arose in the context of a long-standing communication intensive (CI) module that is embedded in a sophomore microprocessor course. The module requires students to write an abstract based on a third-party popular science article on microprocessor innovation. If students could offload the task to GAIs, then this for-credit assignment would need to be redesigned or retired.

This study probed whether select GAIs, using popular science articles as input texts, produce effective abstracts for the engineering, technology, and computing (ETC) discourse community.

While the study is motivated by an immediate concern about the validity of our abstract-writing task, these findings are of value to larger discussions on GAI in ETC classrooms; GAI and student assessment; GAI and academic integrity; and GAI and student writing.

RELATED WORK

Previous work [9, 10] identifies five types of empirical scholarship which investigates GAI in higher education contexts. These include studies that:

1) examine how GAI can assist in completing instructor tasks;
2) consider the affordances of GAIs and suggest how they can realize teaching and learning goals;
3) test the ability of GAI to complete assessments;
4) explore student use and experiences of GAI; and
5) investigate GAI use external to the classroom, where such use has implications for teaching and learning.

This study is located among scholarship identified in (3) above. It follows that our review draws on empirical literature that reports on testing the capabilities of GAIs.

And in this case, we focus on research that evaluates the ability of GAIs to either select content and or synthesize content using an input text: these are cognitive tasks essential to abstract writing.

Several studies tested only one GAI (see for example [11-13]). This paper makes the argument that in designing assessments, where students are permitted to use GAI, there is need for instructors to go beyond understanding a single GAI. We tested three, recognizing that testing all GAIs and all versions is nigh impossible. However, the intent is to work towards a nuanced understanding of how GAIs handle writing-from-sources tasks.

Our review of related work covers two types of empirical studies. First are those which evaluate the ability of GAI to extract text from an input text toward an academic end (content coding for example). And the second set of literature reports on experimentation with GAIs to write an abstract using an input text.

I. GAIs and content extraction

In examining ChatGPT's classification of rhetorical moves in engineering literature reviews, researchers [11] found that ChatGPT did not perform satisfactorily. However, the AI chatbot did offer feedback at the levels of content and organization. And this feedback can be supportive of student writing development.

Another study [13] found that with sufficient detail, specifically an eight-shot prompt, GPT4-Copilot accurately annotated abstracts in applied linguistics. The authors note that this affordance is best exploited by those with domain-specific linguistic knowledge that informs prompting. Though these scholars [13] did not comment on student use of GAIs to annotate abstracts, it is unlikely that learners are able to craft successful prompts for this type of task without a teaching intervention, given that students do not yet have the required genre knowledge identified by the authors.

II. GAIs and abstract writing

In terms of testing the ability to write abstracts, Babl et al [12] tested ChatGPT. They found that this GAI was able to produce an effective abstract. They generated a single abstract using a medical research article as the input. In wider testing of ChatGPT, other scholars [14] present a more complex picture. Inputting 30 papers drawn from the medical sciences, eight experts found that ChatGPT-generated abstracts were lower in quality than the actual abstracts in the published papers. And this poor performance was particularly notable with unstructured AI-generated abstracts. The authors [14] also note that beyond writing quality, the AI did not always accurately summarize the input texts, with ChatGPT getting conclusions wrong in one third of the corpus.

Our review of related work makes three significant points which give impetus to this current study.

Firstly, more testing of GAIs is needed if we are to re-design or retire our credit-bearing writing-from-sources assignments. While the literature abounds with experimentation with GAIs to complete many types of assessments, we found few empirical studies that tested the ability of GAIs to synthesize academic artefacts using input text or to extract text from source documents. See Table 1.

Secondly, where effective GAI-generated artefacts were realized these needed well-engineered prompts, drawing on human expertise [13]. So perhaps, GAIs may not generate accurate responses to writing tasks where linguistic knowledge, specifically genre knowledge, is implicated in the task. But this needs further investigation if educators are to comfortably retain extant writing-from-sources tasks as valid assessments.

Thirdly, of the studies reviewed here, only one used engineering texts as the input. That study [11] did not examine the ability of GAIs to write an abstract but rather to annotate literature reviews. Despite searching several databases (Google Scholar, IEEE Xplore, ProQuest, Ingenta Connect) we did not find any investigation on the ability of GAIs to write abstracts for the ETC community.

TABLE 1. SUMMARY OF RELATED WORK

Article	**Discipline**	**GAI testing**	**genre**
Issa, et al. [11]	engineering	identify rhetorical moves from input texts	literature review
Babl and Babl [12]	medical sciences	write with input text	abstract
Yu, et al. [13]	linguistics	identify rhetorical moves from input texts	abstract
Cheng, et al. [14]	medical sciences	write with input text	abstract

METHOD

We present two sets of data: numerical scoring of the GAI-generated abstracts and examples of outstanding discourse features of the GAI-generated abstracts.

In terms of scoring, two ETC scholars assessed our GAI-generated abstracts. We generated 12 abstracts in total. See Table 2. Three select GAIs were used: M365 Copilot (GPT-5-based chat model); ChatGPT (GPT-4 architecture); and Perplexity (GPT-5.1 model). We input two popular science articles [15, 16] per GAI. Two prompts per article were used:

- a basic single-shot prompt; and
- a single-shot prompt providing a macro-structure for content selection and sequencing.

The macro-structure in prompt two draws on Hartley's structured abstract framework [17]. Hartley's framework requires content which conveys, in this order, the background, aim, method, results, and conclusion.

To assess the quality of abstracts, we operationalized *effectiveness* to mean:

- accuracy of content, that is the GAI correctly represents content presented in the input text;
- relevance of content, that the text generated by the AI is salient in an abstract; and
- logical ordering of content, that is the GAI sequenced content so that it can be understood.

TABLE 2. OVERVIEW OF METHOD

Dimension of Method	Details
Evaluators	• Two (2) ETC Instructors • One (1) Language Instructor
GAIs	• M365 Copilot (GPT-5-based chat model) • ChatGPT (GPT-4 architecture) • Perplexity (GPT-5.1 model)
Input Texts	• Two popular science articles on microprocessor advances [15, 16]
Number of Abstracts	• six using prompt one • six using prompt two
Prompts	1. Using the attached article, write an informative abstract in 250 words or fewer. 2. Using the attached article, write an informative abstract in 250 words or fewer. Use this structure to select and sequence content: background; aim; method; results; and conclusion.

We used consensus marking—that is, in a single sitting both ETC scholars agreed on final scores. Each dimension of effectiveness—accuracy; relevance; sequence—was scored using a three-point Likert scale:

- 0-Unacceptable
- 1-Acceptable
- 2-Very Good

We also compared abstracts generated by prompt one with those generated by prompt two. We anticipated that prompt two would generate better abstracts as we gave the GAIs a macro-structure for extracting and sequencing content. As noted, domain-specific linguistic expertise can inform prompting, leading to better outputs [13].

In terms of qualitative assessment of the abstracts, we report on authorial stance and register. This element of our study was done by a language instructor.

RESULTS & DISCUSSION

Findings are presented using four subheadings:

- overall performance per GAI;
- scores per performance dimension—accuracy; relevance; and sequence
- scores per prompt
- discourse features

I. *Overall performance*

Abstracts were scored in three performance dimensions using a three-point Likert scale of zero to two. Therefore, the highest possible score is six and the lowest possible score is 0. Assuming that a score of three of six or 50% is adequate for a passing effort, then all GAIs generated passing efforts for both input texts and for each prompt. See Table 3.

Four abstracts received perfect scores and one earned the barest passing score of 50%. Overall, M365 Copilot had the best average score at 83.3%.

II. *Assessment per performance dimension*

In terms of accuracy, irrespective of text and prompt, M365 Copilot was most accurate. That is this GAI tended to better represent the ideas of the original input texts. Perplexity was the least accurate. Both M365 Copilot and ChatGPT tended to output more relevant text. This suggests that these two GAIs are more likely to output text that is of rhetorical value in an abstract meant for an ETC audience. Perplexity received the highest score for its ability to logically order content, and this held for both prompts. See Table 4.

TABLE 3. OVERALL SCORES

	Text	M365Copilot	ChatGPT	Perplexity
Prompt 1	Majumdar, et al. [15]	66.7 %	100 %	50 %
Prompt 1	Watts, et al. [16]	100 %	66.7 %	66.7 %
Prompt 2	Majumdar, et al. [15]	100 %	66.7 %	100 %
Prompt 2	Watts, et al. [16]	66.7 %	83.3%	66.7 %
average		83.3 %	79.2 %	70.8 %

TABLE 4. AVERAGE PERFORMANCE PER DIMENSION

Performance dimension	M365Copilot	ChatGPT	Perplexity
accuracy	1.75	1.5	1
relevance	1.5	1.5	1.25
sequencing	1.75	1.75	2
Key: 0 = unacceptable; 1 = good; 2 = very good			

III. *Scores per prompt*

It was expected that prompt two would generate higher scoring abstracts overall, because it offered a macro-structure for selecting and sequencing content. And if prompt two did not impact accuracy and relevance, it was reasonable to expect, at minimum, that sequencing scores would improve.

The macro-structure built into prompt two did improve the sequencing for M365 Copilot and ChatGPT, but it eroded accuracy. See Table 5. This loss in accuracy suggests that both articles of their own had sufficient discourse cues to support the GAIs' selection of text and or that the macro-structure provided in prompt two perhaps confounded the GAIs. This finding needs more probing.

TABLE 5. AVERAGE PERFORMANCE PER PROMPT

	Performance dimension	M365Copilot	ChatGPT	Perplexity
Prompt 1	accuracy	2	2	0.5
	relevance	1.5	1.5	1
	sequencing	1.5	1.5	2
Prompt 2	accuracy	1.5	1	1.5
	relevance	1.5	1.5	1.5
	sequencing	2	2	2
Key: 0 = unacceptable; 1 = good; 2 = very good				

While Perplexity gave more accurate content with prompt two, there was no reduction on its sequencing score. This GAI scored two with both prompts, having the best sequencing score irrespective of prompt.

IV. *Discourse features*

We comment on three language issues: authorial stance; wordiness; and register.

According to scoring by the ETC evaluators (Table 2), four abstracts earned perfect scores, whereas one abstract earned half the total score. Figures 1 and 2 provide the verbatim output of one high scoring [18] and the low scoring abstract [19] respectively.

> Modern chip design has reached extraordinary complexity, with system-on-chip (SoC) layouts involving hundreds of billions of transistors and combinatorial design spaces ranging from 10^{100} to 10^{1000} possibilities. Existing automation tools struggle with multi-constraint optimization, forcing human designers to intervene and slowing design cycles. Machine learning has shown promise, but current AI models fail on tasks requiring constraint satisfaction and backtracking. The article aims to determine whether AI alone can handle physical-design challenges—particularly floorplanning—and to explore how classical search techniques and AI can be combined to improve design effectiveness. The authors began by representing floorplans using B*-trees, enabling overlap-free layouts. They trained graph neural networks, diffusion models, and transformer-based architectures to place blocks relative to one another but found that sequential placement failed under realistic constraints. They then shifted to simulated annealing (SA), ultimately developing constraints-aware simulated annealing (CA-SA), which separates optimization from constraint-repair actions. This led to the creation of Parsac, a parallel CA-SA floorplanner. Parsac outperformed existing solvers, solving commercial-scale floorplanning tasks of ~100 blocks in under 15 minutes and beating all published benchmarks. The authors also released two new datasets—FloorSet-Lite and FloorSet-Prime—containing millions of realistic synthetic floorplans with practical design constraints. The study concludes that AI alone cannot yet solve major chip-design optimization problems. Instead, hybrid approaches—combining AI guidance with classical search—offer the most promising path, enabling faster, more efficient design workflows and supporting future chip complexity growth.

Source: M365 Copilot [18]

FIGURE 1. VERBATIM ABSTRACT OF MAJUMDAR, ET AL. [16] USING PROMPT TWO

While both these abstracts received a passing score from the ETC evaluators, and are grammatically correct, there is language that is somewhat problematic.

Consider the highlighted text. The problem is the position of the author [20-22]—in this case the GAI—with respect to the object of its reporting—an input research article for which the GAI is the putative author. In other words, the writing task required the GAI to take on the identity of author/s of the original input text. And this identity must be signaled through its language. The credibility of the GAI as author is eroded by avoiding the first person; overt references to a third party; and active-voice constructions.

> Chip design has become a massive combinatorial optimization challenge, where existing automated tools struggle under exploding design complexity and numerous hard constraints, forcing engineers into slow, manual iterations. This article argues that AI alone, especially deep learning, is ill-suited to such problems and that hybrid methods combining classical search and machine learning offer a more promising path. The authors first explored AI-only floorplanning using transformer-style models operating on B*-tree representations that guarantee nonoverlapping block placements. Trained on millions of optimal layouts, these models could optimize area and wirelength but failed once realistic hard constraints and larger block counts were introduced, largely because sequential prediction cannot backtrack effectively. The authors then developed constraints-aware simulated annealing (CA-SA), separating cost optimization from constraint repair, and instantiated it in Parsac, an open-source parallel floorplanner. Parsac solves commercial-scale floorplanning problems with around 100 blocks and practical placement and grouping constraints in under 15 minutes, outperforming prior simulated-annealing and learning-based approaches on standard benchmarks and extended constrained versions. To support future research, the team released FloorSet-Lite and FloorSet-Prime, large synthetic but statistically realistic datasets of constrained floorplans for training and testing machine-learning models. The article concludes that fully automated, end-to-end AI solutions for chip design are unlikely soon; instead, hybrid algorithms in which learned models guide but do not replace powerful search procedures are expected to yield faster design closure and more complex, energy-efficient chips.

Source: Perplexity [19]

FIGURE 2. VERBATIM ABSTRACT OF MAJUMDAR, ET AL. [16] USING PROMPT ONE

The nouns *authors* and *team* can appropriately substitute for the first-person. This can be read as a commitment to reporting in an academic register where the first-person is generally frowned on. However, these nouns displace the putative author from the text, giving the impression that the author/s of the abstract are not the same as those of the study.

Added to this is the effect of the third-person pronoun *they*, which suggests once again that the author of the abstract is an entity that is separate from the researcher/s. In one abstract there is even reference to the title of the input text, signaling very strongly that the abstract writer is different from the input text writer/s:

> In *"Lidar on a Chip Enters the Fast Lane,"* the authors examine the technological evolution of lidar systems for autonomous vehicles … [23].

The avoidance of passive constructions compounds this problem of authorial positioning. While the active voice is preferred for its concision and clarity [24], the passive voice has a distinct role in scientific writing. Passive voice has expository function [25] focusing on action not actor, where the prevailing understanding of the discourse community is that the author is the actor [26]. Here are examples:

> The authors began by representing floorplans using B*-trees, enabling overlap-free layouts [18].
>
> To address SA's limitations with hard constraints, the authors introduce a hybrid variant called constraints-aware simulated annealing (CA-SA), implemented in a parallel tool named Parsac [23].

These active-voice constructions ignore a taken-for-granted understanding in the ETC discourse community—that is, the reporting voice is implicated in the action. More natural sounding re-casts are:

> Floorplans were represented using B*-trees, enabling overlap-free layouts.
>
> To address SA's limitations with hard constraints, a hybrid variant called constraints-aware simulated annealing (CA-SA) was introduced and implemented in a parallel tool named Parsac.

Taken together—the first-person substitutes, the third-person pronoun, and the passive voice—the GAI-generated abstracts do not read like abstracts written by humans in the ETC community. These text features were found across all 12 abstracts.

A few possible correctives exist: the GAI could be instructed to inhabit the role of the author of the input text and to use the first-person perspective where needed. But the effectiveness of these correctives needs to be tested. The upshot of this unnatural language is that it may clue instructors in on unauthorized student use of GAI. But it is merely suggestive. Where charges of academic dishonesty are to be laid, more compelling evidence is needed.

The abstracts written by M365 Copilot have a few instances of wordiness or redundancies. See Table 6 for examples and rewrites. We found no such occurrence in the abstracts written by ChatGPT and Perplexity.

One welcome feature of all 12 abstracts was the absence of conversational language found in the input texts. Recall that the input texts are popular science articles, where breezy direct addresses of the reader are appropriate. Here are a few examples, which did not make their way into the GAI-generated abstracts:

> Doing so will generate two new frequencies: the sum and difference of the two frequencies *you* initially mixed [16].
>
> As *you might imagine*, FMCW lidar systems use a very different laser source than ToF systems do [16].

TABLE 6. WORDINESS AND REDUNDANCIES

GAI-generated sentence	**Re-cast, eliminating redundancies and unnecessary words**
Modern chip design has become increasingly complex, with today's systems-on-chip containing hundreds of billions of transistors and requiring engineers to navigate extraordinarily large combinatorial design spaces [18].	~~Modern~~ Chip design ~~has become~~ ***is*** increasingly complex, with today's systems-on-chip containing hundreds of billions of transistors and requiring engineers to navigate extraordinarily large combinatorial design spaces.
This article examines recent advancements in lidar technology that aim to enable safer and more affordable autonomous vehicles [27].	This article examines recent advancements in lidar technology that ~~aim to~~ enable safer and more affordable autonomous vehicles.

Other notable revisions to the popular-science register as found in the input texts [15] were the elimination of:

- figurative language: *hairy* optimization problems
- contractions: *here's* how it works
- colloquialisms: *back to the drawing board*; *a far cry.*

Another discourse feature which marked the GAI-generated abstracts as distinctively academic was complex punctuation. We consider complex punctuation to be that which we would not normally find in student writing. Consider here the use of the em-dash:

> While traditional lidar systems have been costly, bulky, and mechanically complex—hindering widespread adoption—new chip-scale designs promise to overcome these limitations [18].

Only M365 Copilot used the em-dash. It appeared in all four of its abstracts.

CONCLUSION

Our experimentation with three widely available GAIs—M365 Copilot (GPT-5-based chat model); ChatGPT (GPT-4 architecture); and Perplexity (GPT-5.1 model)—shows that they can produce adequate to exceptional abstracts for the ETC discourse community. Of the 12 abstracts none were found to be ineffective, as in failing to communicate germane, accurate content. These findings warrant rethinking writing-from-sources assessments, as it is possible that students can offload this type of task to a GAI.

The assumption that more context in prompts would lead to better output did not pan out. When two GAIs were provided with a macro-structure for extracting and sequencing text, there were gains in some dimensions of the writing performance but a reduction in other dimensions. In short, more prompting does not equal better output. This underscores the need for wider testing, with varying prompts and multi-shot prompting. And, there is need for domain specific expertise to inform prompting, as genre knowledge is needed in the processes of extracting and synthesizing content [13].

Some welcome findings were that without explicit prompting all GAI-generated abstracts were free of figurative, informal language. This supports the use of GAI as an editor for students who are learning to produce a formal academic register.

Our experimentation with three GAIs, using two input texts needs to be extended: there is need for testing with more and different types of texts, asking for the output of different academic genres, using more evaluators, and examining more performance dimensions. And while our study does not provide a definitive statement on the validity of for-credit writing-from-sources tasks, it does expose how tricky assessment design has become with the proliferation of GAI.

Indeed, the upshot of our limited testing of GAIs is a reminder that we are at the mercy of algorithmic calculations of the permutations and combinations of natural language use, as found in corpora of unknown composition and size. And this black box, into which our inputs vanish, and 'new' content emerges, is not going away or explaining itself anytime soon. Researchers will do well to study these technologies more. For with better understanding, we can make the best use of the technology without undermining our students' learning and development.

References

[1] C. Mohammed and C.-A. Radix, "Creating Rhetorical Complexity: Using Popular Science Articles to Teach Abstract Writing," in 2020 IEEE International Professional Communication Conference (ProComm), 2020, pp. 137-140: IEEE.

[2] C. A. Radix and C. Mohammed, "The efficacy of Hartley's "structured format" in the teaching and assessment of abstract writing," in 2017 IEEE Frontiers in Education Conference (FIE), 2017, pp. 1-6: IEEE.

[3] H. H. Feng, "Designing, implementing, and evaluating an automated writing evaluation tool for improving EFL graduate students' abstract writing: A case in Taiwan," Iowa State University, 2015.

[4] C. Liu and R. Houdek, "Teaching computer science graduate students scholarly literature review techniques," in Proceedings. Frontiers in Education. 36th Annual Conference, 2006, pp. 14-19: IEEE.

[5] M. K. Sastry and C. Mohammed, "The summary-comparison matrix: A tool for writing the literature review," in IEEE International Professional Communication 2013 Conference, 2013, pp. 1-5: IEEE.

[6] J. B. Lenart and C. G. Berdanier, "Optimizing a genre analysis framework to investigate engineering literature reviews," in 2017 IEEE International Professional Communication Conference (ProComm), 2017, pp. 1-6: IEEE.

[7] R. E. Cummings, S. M. Monroe, and M. Watkins, "Generative AI in first-year writing: An early analysis of affordances, limitations, and a framework for the future," Computers and Composition, vol. 71, p. 102827, 2024.

[8] N. N. Spivey, "Transforming texts: Constructive processes in reading and writing," Written communication, vol. 7, no. 2, pp. 256-287, 1990.

[9] C. Mohammed and W. Sarjusingh, "Productive uses of GAI: Preliminary findings from a capstone course," in 2024 IEEE Frontiers in Education, Washington, D.C., 2024: IEEE, 2024.

[10] C. Mohammed, "Work in Progress: A Rapid Review of the Scholarship on Generative AI in Engineering Workplaces—Implications for Engineering Education," in 2025 IEEE Engineering Education World Conference (EDUNINE), 2025, pp. 1-4: IEEE.

[11] M. Issa, M. Faraj, and N. AbiGhannam, "Exploring ChatGPT's ability to classify the structure of literature reviews in engineering research articles," IEEE Transactions on Learning Technologies, vol. 17, pp. 1819-1828, 2024.

[12] F. E. Babl and M. P. Babl, "Generative artificial intelligence: Can ChatGPT write a quality abstract?," Emergency Medicine Australasia, vol. 35, no. 5, pp. 809-811, 2023.

[13] D. Yu, M. Bondi, and K. Hyland, "Can GPT-4 learn to analyse moves in research article abstracts?," Applied Linguistics, 2024.

[14] S. Cheng et al., "Comparisons of quality, correctness, and similarity between ChatGPT-generated and human-written abstracts for basic research: cross-sectional study," Journal of Medical Internet Research, vol. 25, no. 1, p. e51229, 2023.

[15] S. Majumdar, U. Mallappa, and H. Mostafa, "AI alone isn't ready for chip design: A combination of classical search and machine learning may be the way forward," IEEE Spectrum, vol. 61, no. 12, pp. 38-43, 2024.

[16] M. R. Watts, C. Poulton, M. Byrd, and G. Smolka, "Lidar on a chip enters the fast lane: Sensors for self-driving cars and robots will be tiny, reliable, and affordable," IEEE Spectrum, vol. 60, no. 9, pp. 38-43, 2023.

[17] J. Hartley, "Improving the clarity of journal abstracts in psychology: the case for structure," Science Communication, vol. 24, no. 3, pp. 366-379, 2003.

[18] M365 Copilot (GPT-5-based chat model), "Using the attached article (Majumdar et al 2024) write an informative abstract in 250 words or fewer. Use this structure to select and sequence content: background; aim; method; results; and conclusion. ," ed, 2026.

[19] Perplexity (GPT-5.1 model), "Using the attached article (Majumdar et al 2024) write an informative abstract in 250 words or fewer.," ed, 2026.

[20] J. D. Ford, "Knowledge transfer across disciplines: Tracking rhetorical strategies from a technical communication classroom to an engineering classroom," IEEE Transactions on Professional Communication, vol. 47, no. 4, pp. 301-315, 2004.

[21] K. Hyland, "Disciplinary discourses: Writer stance in research articles," in Writing: Texts, processes and practices, C. N. Candlin, K.Hyland., Ed. London, 1999.

[22] P. K. Matsuda and C. M. Tardy, "Voice in academic writing: The rhetorical construction of author identity in blind manuscript review," English for Specific Purposes, vol. 26, no. 2, pp. 235-249, 2007.

[23] ChatGPT (GPT-4 architecture), "Using the attached article (Watts et all 2023), write an informative abstract in 250 words or fewer," ed, 2026.

[24] S. Conrad, K. Lamb, and T. J. Pfeiffer, "Where grammar, content, and professional practice meet: The case of the passive voice," in 2018 ASEE Annual Conference & Exposition, 2018.

[25] K. Riley, "Passive voice and rhetorical role in scientific writing," Journal of technical writing and communication, vol. 21, no. 3, pp. 239-257, 1991.

[26] D. D. Ding, "The passive voice and social values in science," Journal of Technical Writing and Communication, vol. 32, no. 2, pp. 137-154, 2002.

[27] M365 Copilot (GPT-5-based chat model), "Using the attached article (Watts et al 2023) write an informative abstract in 250 words or fewer.," ed, 2026.