

Co-Adaptation in Learner–ChatGPT Dyadic Interaction: A Multi-Leveled Linguistic Analysis

Ashley Beccia, Sue Min Park, Jill Williams

Teachers College, Columbia University

INTRODUCTION

Despite the rapid uptake of large language models (LLMs) like ChatGPT in second language (L2) learning environments, the interactional dynamics of LLM–learner dyads remain under-examined. Existing research has primarily focused on the *products* of LLM–learner interactions, while the interactional *process* is rarely a central concern. For example, Sok and Shin (2025) compared learners’ task performance before and after interacting with ChatGPT, emphasizing the importance of outcome gains rather than the turn-by-turn exchanges with the LLM. Kusumaningrum et al. (2024) analyzed the degree of conceptual, lexical, and structural overlap between ChatGPT-generated text and learners’ final email drafts, focusing on learners’ appropriation of AI output rather than the dynamics of learner–ChatGPT interaction.

When interactional data *have* been examined, researchers tend to employ global metrics rather than fine-grained analyses of the ongoing linguistic coordination. For instance, Sim, Kim and Ku (2025) compared ChatGPT–learner and learner–learner interaction using pre/post gains and aggregate counts of pragmatic-related episodes and exposure to appropriate forms. As a result of this limited research scope, little is known about how learners and LLMs mutually adapt on a granular level over time.

Understanding these dynamics requires recognizing that LLMs represent a fundamentally different type of interlocutor. While humans co-adapt through embodied perception, affective attunement, and lived experience (Larsen-Freeman, 2020), LLMs generate language probabilistically based on large-scale pretraining and alignment procedures such as instruction-following and reinforcement learning from human feedback (Kulkarni et al., 2023; Ouyang et al., 2022). These systems display stylistic defaults, such as elaborated and expository responses, and typically operate within bounded context windows that limit durable, cross-session memory (Zhong et al., 2024). This creates an interactional ecology in which alignment may occur without symmetry, and responsiveness reflects pattern-matching rather than experiential attunement. Even so, research shows that humans can align linguistically with computational partners, producing priming effects comparable to human–human interaction (Branigan et al., 2010). Therefore, co-adaptation does not require shared human cognition or

intention, but the mechanisms and constraints that shape alignment in human–LLM interaction remain largely unexplored.

ANALYTIC APPROACH

To investigate how co-adaptation unfolds in a human–LLM dyad, we conducted a multi-level analysis of the seven-week interactional dataset. Guided by a Complex Dynamic Systems Theory (CDST) perspective, co-adaptation was conceptualized as the dynamic alignment that emerges as interlocutors modify their linguistic behavior over time. Because alignment can occur on multiple levels, the analysis focused on two dimensions: language (form) and content (meaning).

To capture how these adjustments unfolded moment-by-moment, we employed process tracing (Hiver & Al-Hoorie, 2019). Rather than correlating variables, process tracing reconstructs causal chains by iteratively locating observable manifestations of theorized mechanisms and examining their sequencing and co-occurrence (Thomas et al., 2024). In our context, this allowed for micro-level tracking of when a lexical item, syntactic pattern, or topic introduced by one party was adopted—or not—by the other.

To guide our analysis, we developed an analytic template that stratifies the evidence into two levels. First, the level of language captured alignment in the linguistic forms used in the interaction and comprises two sublevels: lexical synchrony, referring to the adoption of specific lexical items introduced by the participants; and syntactic synchrony, referring to the mirroring of clausal and sentential patterns. These sublevels allowed us to examine how linguistic forms were taken up, reshaped, or resisted across adjacent turns. Second, the level of content captured alignment in the subject matter of the interaction. At this level, we examined how topics were initiated, sustained, and shifted by the participants across turns and weeks, allowing us to trace the emergence and evolution of shared topical trajectories.

In a shared text file, the data were systematically coded for co-adaptation at the two levels. These codes were examined for patterns, especially the co-occurrence of codes in adjacent turns, which was taken as evidence for co-adaptation—that is, one participant adjusting their linguistic behavior to align with the other.

FINDINGS

Through iterative analysis of the data, co-adaptation was revealed on both the level of language and content. In the following sections, we present the lexical items, syntactic constructions, and content areas for which the learner and ChatGPT most clearly exhibited co-adaptation throughout their interactions.

Lexical Items

Looking longitudinally across the seven weeks, there was modest evidence of lexical synchrony. For example, the learner introduced the lexical item *paper* to the ecosystem in Week 1 when explaining scientific writing and data commentaries: “Data commentary is an essential skill for writing a good *paper*.” After some discussion between the two, ChatGPT employed the word *paper* three times toward the end of the conversation as it made suggestions about the length of data commentary paragraphs and offered resources for examples.

Though the lexical item *paper* was not mentioned in Week 2, it reappeared at the end of the interaction in Week 3 when the learner asked:

in your opinion, which terms are more common in a paper: “in this study, “in this paper”, and “in this work?

Subsequently, ChatGPT quoted the learner’s inquiry to answer the question, adopting “in this paper” verbatim.

The word *paper* continued to appear frequently in subsequent weeks as the dyad took turns using this item, especially in Week 4 when the learner initiated the item when asking about the IMRAD style of paper and its various sections: “in Imrad style paper, which section is the most common to write first?” ChatGPT responded by detailing the introduction section and its importance until the learner explained the original intention of the inquiry by clarifying that it is for a *draft* of the paper.

In Week 5, the learner also initiated the synchrony by stating that their class was learning more about the introduction section of a paper. In ChatGPT’s subsequent turn, it asked the learner to be more specific about writing introductions for scientific papers. From there, the two participants took turns detailing the structure and contents of the introduction, both using the lexical item *paper*, but ChatGPT on several more occasions than the learner.

The item *paper* was initiated once again by the learner in Week 6 when asking about the abstract, title, and acknowledgement components of a scientific paper. Similar to before, ChatGPT engaged with this inquiry by asking for more specifics about what the learner gleaned from class about these topics. In the learner’s third turn in Week 6, they asked “how long is the abstract in a paper?”, in which ChatGPT detailed the guidelines to some scientific papers. From this turn, the learner continued the inquiry by asking how to choose a title for a paper and later about a typical outline for a survey paper. The dyad continued to use the lexical item *paper* in synchrony in this question-response style throughout the remainder of Week 6.

Lastly, in Week 7, the learner began the synchrony in the third turn by asking about how to determine the order of authors in a paper. ChatGPT matched the learner verbatim: “the order of authors in a paper” in its response and provided details to answer this question. The chatbot continued the synchrony by using this item in its seventh and eighth turns when the topic

changed from papers to presentations, despite the learner not continuing the synchrony in both their previous and subsequent turns during Week 7.

Overall, with the exception of Week 2, the lexical item *paper* appeared consistently throughout the dataset as a content word, with the participants synchronizing in their use of the lexical item when engaging in the question–response process. The learner always initiated this specific lexical item, and ChatGPT co-adapted by implementing this term in its responses and questions back to the learner, even when the learner discontinued their use of the term in subsequent responses each week.

Unlike the lexical item *paper*, however, most of the lexical synchrony either occurred locally, such as within one conversation, or appeared inconsistently across the weeks. For instance, the word *paragraph* was introduced to the ecosystem by the learner during Week 1 with the inquiry, “What is the typical organizational flow in the data commentary paragraph?” From there, ChatGPT started using this term frequently in its subsequent response. The next turn taken by the learner included the question, “How long does usually data commentary paragraph have?” ChatGPT responded using a plural variation of the term: “There isn’t a set length for data commentary paragraphs [emphasis added]....” This term reappeared in Week 2 when the learner began the inquiry with: “Can you help to spot some punctuations or grammars error in my paragraph?” and ChatGPT responded with a turn containing *paragraph*. This lexical item did not reappear for the remainder of the dataset.

A final example of lexical synchrony is the item *scientific writing*. Once again, the learner introduced the lexical item to the ecosystem in Week 1: “Today I learned about data commentary for scientific writing...” Though ChatGPT did not use *scientific writing* on its subsequent turn, it did on its third, fourth, sixth and seventh turns. The learner, however, did not use this item again during this conversation. ChatGPT continued to employ this item in Weeks 3, 4, and 5 without any reciprocation from the learner until Week 6, when the learner employed this item once more. This time, the opposite occurred: ChatGPT did not reciprocate by using this term.

In all, there was moderate evidence of lexical synchrony in the dataset. Most of the evidence was identified within particular conversations, rather than across the conversations week by week. There was also evidence for the lack of co-adaptation as one participant continued to use a lexical item that the other participant did not.

Syntactic Constructions

Evidence for syntactic synchrony was limited, emerging primarily in two interrogative–declarative pairings. In the following, these syntactic constructions are introduced and defined, and evidence for co-adaptation in the participants’ use of these constructions (or the lack thereof) is presented.

Can PRO V...? / PRO can(not) V...

The *Can PRO V...?* construction refers to an interrogative beginning with the modal “can” followed by a pronoun and a verb. This syntactic pattern, along with the parallel declarative construction *PRO can(not) V...*, was identified throughout the dataset, especially in Weeks 1, 2, 3, and 5.

ChatGPT introduced the interrogative pattern at the start of each sentence with its fixed opener (“**Can you** please **tell** me...”). In Week 1, the learner did not immediately use a *Can PRO V...?* construction. However, during their fifth turn, the learner asked, “where **can I find** some example of good data commentaries including their data examples?” Although this interrogative begins with a WH- question word, the learner still adopted the *can PRO V...?* pattern introduced by ChatGPT. Then, in the subsequent turn, ChatGPT continued employing this pattern, opting for the *PRO can V* declarative form on two occasions to respond to the learner’s question. On the first occasion, ChatGPT even preceded the pattern with a WH- question word, emulating the learner.

In Week 2, co-adaptation occurred more immediately. Once again, ChatGPT initiated the conversation with its opening message containing the *Can PRO V...?* construction. In response, the learner mirrored ChatGPT’s language use to ask, “**Can you help** to spot some punctuations or grammars error in my paragraph?” Then, on their third turn, the learner produced another *Can PRO V...?* interrogative construction.

In Week 3, the learner produced a *Can PRO V...?* interrogative construction on four occasions in successive turns (see Appendix A). These constructions functioned as indirect commands (e.g., asking the AI chatbot to do something) or means for seeking information about ChatGPT’s capabilities (i.e., what the AI chatbot can and cannot do). It is this latter function that prompted ChatGPT to employ the declarative form of the construction in the subsequent turn:

Yes, as an AI language model, I am able to conduct some of the same analyses that AntConc can do. For example, **I can generate** concordances and analyze word frequencies within a given text...

Here, ChatGPT produced a *PRO can(not) V...* construction to inform the learner about the functionalities it shares with AntConc. Similarly, in the subsequent turn, when the learner asked ChatGPT if it can replace AntConc, ChatGPT replied with a *PRO can(not) V...* construction to delimit what it can do, followed by a negated version of the declarative construction to clarify what it cannot do.

In the remainder of the weeks, ChatGPT continued to use *Can PRO V...?* constructions, not only in its fixed opening but also in turns such as the following from Week 5, “Great! **Can you tell** me more about what you specifically learned about writing introductions for scientific papers?” However, the learner almost entirely discontinued their use of the construction, producing a *Can PRO V...?* interrogative only once more towards the end of Week 5. As a result,

there is no further evidence of co-adaptation in the participants' use of *Can PRO V...? / PRO can(not) V...* constructions.

COP SUBJ X? and SUBJ COP (not) X

The second syntactic pattern for which there is evidence of co-adaptation in the dataset is the interrogative *COP SUBJ X?* and its declarative counterpart *SUBJ COP (not) X*. In these constructions, the “COP” slot is filled with a copular verb (e.g., “is” or “are”) and the “SUBJ” slot is filled with a subject (e.g., “there” or “it”). The “X” slot may be filled with an adjective, such as “important,” or a noun phrase, such as “good resources.” The negator “not” is optionally placed before the “X” slot (e.g., “Is is not important...”).

Though these constructions only appeared in Weeks 1, 2, 4, and 7, co-adaptation was evident in each of these conversations. In Week 1, as the participants were discussing data commentaries, the learner asked ChatGPT on their second turn, “**Is it important** to position ourselves as author in the data commentary?” In response, ChatGPT used the declarative *SUBJ COP X* on two occasions:

Yes, **it is important** to clearly identify yourself as the author in the data commentary. This helps to establish your authority and credibility as the researcher or scholar who conducted the study or analyzed the data. **It's also essential** that you use language and tone that is appropriate for scientific writing, keeping in mind the need for objectivity and clarity.

The first occurrence above mirrors the learner's language use more closely than the second, as ChatGPT used the uncontracted copular *to be* (i.e., “it is”) and same adjective (i.e., “important”). Nevertheless, the second occurrence also mimics the learner's language use, as ChatGPT employed a dummy subject (i.e., “it”), copular *to be*—albeit contracted—and semantically related adjective (i.e., “essential”) to expand on its response to the learner's question.

In Week 2, synchronization in the participants' use of the *COP SUBJ X?* construction occurs later on in the exchange. On their penultimate turn, the learner asked, “**is it okay** to have number of percentage at the beginning of the sentence?” In its response, ChatGPT employed the declarative counterpart, “...**It is** therefore generally **better** to start the sentence with the subject or the information being described...” Though ChatGPT inserted two adverbs, the use of the syntactic pattern *SUBJ COP X*, particularly with an adjective in the “X” slot, demonstrates co-adaptation.

In the following week, neither participant employed a *SUBJ COP X* construction whatsoever. It is not until the learner's fifth turn during Week 4 that the construction appeared again. As the participants discussed writing the methods and results sections of an empirical report, the learner inquired, “**is there any good resource** to practice writing methods and results sections?” In this case, the learner employed the interrogative *COP SUBJ X* with the dummy

subject “there” in the “SUBJ” slot and noun phrase “any good resource” in the “X” slot. Mirroring these linguistic choices, ChatGPT responded by stating:

Yes, **there are several resources** you can use to practice writing methods and results sections. ... **There are** also **several online resources** that offer writing prompts and exercises specifically designed to help you practice writing methods and results sections...

In this reply, ChatGPT opted for the same *COP SUBJ X* construction, though employing different adjectives in the noun phrase, and on the second occasion, inserting an adverb between the copula and subject.

The final instance of co-adaptation in the use of the *COP SUBJ X?* construction can be found in Week 7. As the participants discussed the authorship of scholarly publications, the learner asked, “but sometimes, some advisor want to be listed as the first author even though the student has the most contribution. **is it okay?**” ChatGPT then emulated the learner’s language use by responding, “No, **it is not okay** for an advisor to be listed as the first author if the student has made the most significant contribution to the research project...” Aside from the addition of negators, ChatGPT began its response with the same syntactic pattern that the student employed as an interrogative in the preceding turn.

In brief, synchronization in the participants’ use of two interrogative–declarative pairings emerged sporadically across the dataset.

Content

At the content level, the dyad exhibited a stable, learner-led topical pattern across Weeks 1–7. The learner consistently initiated or redirected topics through brief, task-oriented prompts. Subsequently, ChatGPT responded with substantially more elaborate turns that offer definitions, step-by-step procedural guidance, genre conventions, and outline templates. This separation of roles resulted in synchrony in topical focus but non-convergence in production length. Even as ChatGPT’s responses closely tracked the student’s queries, its turns remained expansive in comparison to the student’s concise contributions. The difference in turn length was visually evident in the transcripts, producing an enduring asymmetry in the distribution of discourse across the database. Despite this asymmetry, the dyad displayed a coherent trajectory of topical development, progressing from micro-level writing concerns (e.g., editing, paragraph organization) to macro-level research practices (e.g., IMRAD sequencing, authorship order, research dissemination). In a nutshell, at the content level, the dyad showed a lack of adaptation in turn expansion but clear alignment in content progression. The summary of the weekly content highlights is displayed in Table 1.

TABLE 1
Week-by-Week Content Highlights

Week	Topic	Details
Week 1	Data Commentary	Definition, purposes, organization/flow, typical paragraph lengths; curtailed exemplars
Week 2	Editing	Punctuation/grammar fixes; concision trade-offs; guidance on sentence-initial percentages
Week 3	AntConc	Operations (concordance, collocates); capability boundaries relative to ChatGPT
Week 4	IMRAD	Drafting sequence; tense conventions, Results vs. Discussion separation; practice resources
Week 5	Introduction	“Three moves” negotiated and corrected; include-all-moves guidance; outline provided
Week 6	Abstract & Survey	Abstract organization and word counts; title heuristics; empirical vs. synthesis survey outlines
Week 7	Presentation & Authorship	Authorship order norms; oral vs. poster design, visuals, and selection

Content synchrony was achieved through student micro prompts (e.g., “What is...?”, “Can you...?”, “Which section...?”) that cued ChatGPT to deliver macro-level elaborations that extend beyond the immediate question to include rationale, exemplification, and models. This pattern was evident from the earliest exchanges on data commentary through later weeks on introductions, abstracts, survey papers, and presentation design. The topical scope expanded systematically (see Appendix B). Structural concerns surfaced repeatedly in Week 1 (data commentary), Week 4 (IMRAD), Week 5 (introductions), and Week 6 (abstract/survey). Style resurfaced as a cross-cutting theme in Week 2 (editing for correctness and concision), Week 4 (tense/passive), and Week 6 (title selection heuristics). Tools came to the fore in Week 3 (AntConc features and limits relative to ChatGPT), while Week 7 added professional norms (authorship order) and delivery considerations (oral versus poster presentations). In sum, the discourse advanced from local textual decisions to the broader ecology of scholarly communication.

The granularity of inquiry also intensified over time. Early prompts were definitional (e.g., “What is a data commentary?”), whereas later questions required fine-grained genre decisions (e.g., typical word counts for abstracts, outlines for research-synthesis surveys). In response, ChatGPT narrowed its topical scope and operationalized guidance by providing targeted templates, word-count ranges, and genre-sensitive rationales, thereby aligning with the student’s specific aims. In parallel, writing support evolved from generic references to bespoke

scaffolds. Week 1 curated exemplars and guides (e.g., Nature, Duke Science Writing Resource, UNC Writing Center, Elsevier), while later weeks delivered practice strategies (reading and summarizing Methods/Results; OWL exercises; workshops), customized outlines for introductions and survey papers, abstract length ranges, and title selection rules. The overall pattern demonstrates an increased level of writing support sophistication, aligning with the student's more discriminating request.

Two episodes illustrate topical co-adaptation particularly clearly. In Week 5, the student repeatedly corrected the description of the introduction's "three moves," and ChatGPT responded by revising and clarifying its account, ultimately consolidating an outline consistent with the student's formulation. In Week 6, the student re-specifies "survey paper" as a research synthesis paper, prompting ChatGPT to reframe the outline accordingly. In both cases, ChatGPT resupplied its content in real-time in direct response to student feedback, demonstrating sensitivity to the learner's topical framing while preserving its expository depth. At the same time, these adjustments did not affect response magnitude. That is, despite topical alignment, the dyad retained a persistent length asymmetry, with the student's prompts remaining brief and single-function while ChatGPT's turns remained long and multi-clause. In short, co-adaptation was strong in what is addressed but absent in the amount produced, yielding synchrony in topic selection without convergence in turn expansion.

Viewed holistically, content-level analysis revealed a student-led, ChatGPT-scaffolded interactional pattern. Brief student prompts determined the syllabus for each exchange. At the same time, extended ChatGPT responses delivered information with a few moments of correction-driven refinement, keeping the content aligned with evolving learner goals. The longitudinal exchange, therefore, traced a recognizable development from textual micro-skills (editing, paraphrasing, tense) to scholarly practices (IMRAD planning, authorship ethics, and presentation design), documenting topical co-adaptation alongside a lack of co-adaptation in turn expansion.

DISCUSSION AND CONCLUSION

Employing process tracing guided by a multi-leveled analytic template, we found that co-adaptation in this human-LLM dyad manifested unevenly across linguistic and content levels. In terms of lexical items, synchrony was present but mostly local, with a few items recurring longitudinally after the learner's introduction of the items to the ecosystem in Week 1. In terms of syntactic coordination, synchrony was concentrated in two constructions. In one case (*Can PRO V...?*), ChatGPT led synchrony through its recurring opening message; in the other (*COP SUBJ X?*), the learner initiated interrogatives that ChatGPT echoed in declarative responses. Yet, across weeks, both constructions showed inconsistent uptake, with stretches of alignment followed by stretches of non-use. This selectivity echoes findings from Beccia et al. (2024), who

similarly observed co-adaptation emerging robustly in some linguistic subsystems while largely bypassing others.

By contrast, content-level synchrony was markedly more stable and cumulative. Across all seven conversations, the learner reliably set the agenda through succinct prompts, while ChatGPT responded with extended expository turns containing definitions, explanations, heuristics, and procedural guidance. This division of labor produced a consistent interactional pattern: brief learner initiations followed by ChatGPT's elaboration, yielding high topical alignment alongside persistent asymmetry in turn length. Over time, this coordination supported an expansion of thematic scope from micro-level concerns (e.g., paragraph structure, scientific style) to broader scholarly practices (e.g., IMRAD organization, authorship conventions). Several motifs (e.g., grammar, research tools) resurfaced across conversations, contributing to the system's recurrent topical trajectory.

These findings show that co-adaptation at the level of language was transient, while co-adaptation at the level of content was recurrent. The fleeting nature of linguistic synchrony likely reflects ChatGPT's lack of cross-session memory. Because the LLM cannot retrieve prior conversations, lexical items or syntactic patterns introduced earlier in the longitudinal sequence cannot serve as sustained anchor points for alignment. Consequently, synchrony often emerged within a given conversation but dissipated in subsequent ones. Content synchrony, however, did not require long-term memory of specific forms; rather, it was supported by the learner's consistent prompting practices.

In CDST terms, the dyad's content-level co-adaptive behavior appeared to settle into an attractor state, or a stable pattern repeatedly approached over time (Hiver, 2014). The learner consistently occupied the agenda-setting role, while ChatGPT consistently took on the elaborative role, regardless of the specific topic. This repeated co-enactment suggests self-organization within the dyadic ecosystem, as the interaction settled in a configuration that fit the participants' goals, capabilities, and constraints (Larsen-Freeman & Cameron, 2008).

The findings also shed light on who leads co-adaptation and under what conditions. In terms of lexical items and subject areas, the learner frequently led the synchrony by introducing content words and setting topical directions. ChatGPT then incorporated these elements into its responses, demonstrating alignment. In terms of syntactic coordination, leadership was construction-specific. ChatGPT repeatedly initiated the *Can PRO V...?* construction, while the learner initiated *COP SUBJ X?* interrogatives that ChatGPT mirrored syntactically. These patterns suggest that leadership in co-adaptation is not uniformly human- or machine-driven, but emerges from the affordances of particular constructions and roles within the interaction.

Overall, the findings contribute to emerging work on co-adaptation as a mechanism for L2 development by exploring how the mechanism might operate in ecosystems of dyadic interaction when one interlocutor is an LLM. Though L2 development was not measured, the analysis paves the way for future research linking patterns of co-adaptation to independent measures of development.

Several limitations should be noted. Because the dataset represents a single dyad, the findings are not generalizable to all human–LLM interactions, particularly those involving models with durable memory. The focus on synchronous written interaction also leaves open questions about co-adaptation in oral modalities. Finally, without independent learning measures, developmental implications cannot be inferred. Future research should examine human–LLM co-adaptation across diverse tasks, learners, modalities, and LLM architectures, and incorporate learning measures to determine whether and how co-adaptation contributes to L2 development.

REFERENCES

- Beccia, A., Lew, W. M. A., & Han, Z. (2024). Exploring co-adaptation in an ecosystem of dyadic interaction. *Language Teaching Research Quarterly*, 39, 125–144.
<https://doi.org/10.32038/ltrq.2024.39.10>
- Branigan, H. P., Pickering, M. J., Pearson, J., & McLean, J. F. (2010). Linguistic alignment between people and computers. *Journal of Pragmatics*, 42(9), 2355–2368.
<https://doi.org/10.1016/j.pragma.2009.12.012>
- Hiver, P. (2014). 3. Attractor states. In Z. Dörnyei, P. MacIntyre & A. Henry (Eds.), *Motivational dynamics in language learning* (pp. 20–28). Multilingual Matters.
<https://doi.org/10.21832/9781783092574-005>
- Hiver, P., & Al-Hoorie, A. (2019). *Research methods for complexity theory in applied linguistics* (1st ed.). Multilingual Matters. <https://doi.org/10.21832/9781788925754>
- Kulkarni, A., Shivananda, A., Kulkarni, A., & Gudivada, D. (2023). *Applied generative AI for beginners: Practical knowledge on diffusion models, ChatGPT, and other LLMs* (1st ed.). Apress. <https://doi.org/10.1007/978-1-4842-9994-4>
- Kusumaningrum, W. R., Lim, H. J., Challis, K., & Beck, J. (2024). Exploring how English language learners use ChatGPT in their email writing processes. In C. A. Chapelle, G. H. Beckett, & J. Ranalli (Eds.), *Exploring artificial intelligence in applied linguistics* (pp. 42–58). Iowa State University Digital Press. <https://doi.org/10.31274/isudp.2024.154.04>
- Larsen-Freeman, D. (2020). Chapter 8. Complexity Theory: Relational systems in interaction and in interlocutor differences in second language development. In L. Gurzynski-Weiss (Ed.), *Cross-theoretical explorations of interlocutors and their individual differences* (pp. 189–208). John Benjamins Publishing Company. <https://doi.org/10.1075/lllt.53.08lar>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in Neural Information Processing Systems* 35 (pp. 27730–27744). Neural Information Processing Systems Foundation, Inc.

- Sim, S., Kim, Y-J., & Ku, K. (2025). Exploring generative AI as a roleplay interlocutor in L2 task-based pragmatics learning: Comparing ChatGPT-learner and technology-mediated peer interactions. *TESOL Quarterly*, 59(S1), S86–S116.
<https://doi.org/10.1002/tesq.70010>
- Sok, S., & Shin, H.W. (2025). Do interactions with ChatGPT influence L2 learners' oral speaking ability, summarization ability, and perceptions of generative AI tasks? *TESOL Quarterly*, 59(S1), S19–S51. <https://onlinelibrary.wiley.com/doi/abs/10.1002/tesq.70001>
- Thomas, N., Hultgren, A. K., Zuaro, B., Yuksel, D., Wingrove, P., Nao, M., & Beach, D. (2024). Process tracing for applied linguistics. *Research Methods in Applied Linguistics*, 3(2), 100118. <https://doi.org/10.1016/j.rmal.2024.100118>
- Zhong, W., Guo, L., Gao, Q., Ye, H., & Wang, Y. (2024). MemoryBank: Enhancing large language models with long-term memory. *Proceedings of the AAAI conference on artificial intelligence*, 38(17), 19724–19731. <https://doi.org/10.1609/aaai.v38i17.29946>

APPENDIX A

“Can you” Constructions Produced by the Learner in Week 3

Turn	ChatGPT (before)	Learner (during)	ChatGPT (after)
2	Yes, I'm familiar with AntConc. It's a corpus analysis software that can be used to analyze large amounts of text data. Did you find the demonstration of AntConc that we did in class helpful?	Yes, it showed the basic use of it. <u>Can you elaborate</u> more on how to use antconc?	Certainly, AntConc is a powerful tool for analyzing text data and discovering patterns in large amounts of text...
3	... I hope this gives you a better idea of what AntConc is and how to use it. Please let me know if you have any other questions or concerns.	<u>can you give</u> specific examples of you meant by concordance and collocates?	Certainly! Concordance and collocates are two of the main features in AntConc, and they allow you to analyze your text data in different ways...
4	... Overall, both concordances and collocates can help you to gain a better understanding of your text data and identify interesting patterns or trends that you might not have noticed otherwise.	<u>can you do</u> what antconc does?	Yes, as an AI language model, I am able to conduct some of the same analyses that AntConc can do. For example, <u>I can generate</u> concordances and analyze word frequencies within a given text...
5	... However, I may not be able to run some of the more complex analyses that require specialized software like collocation analysis, keyword analysis, or n-gram analysis.	<u>can you replace</u> antconc?	As an AI language model, <u>I can perform</u> some of the same functions as AntConc such as searching for word usage patterns in a large text corpus, but <u>I cannot replace</u> AntConc altogether...

APPENDIX B

Content Matrix by Theme and Week

Theme	Week1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7
Scientific Writing Structure	Data commentary	-	-	IMRAD	Introduction	Abstract & Survey	Presentation structure
Academic Style/ Grammar	-	Editing	-	Tense, passive	Statement length	Title selection heuristics	-
Research Tools	-	-	AntConc	-	-	-	-
Research Ethics/ Norms	-	-	-	-	-	-	Authorship order
Writing Support Resources	Curated exemplars (Nature, Duke, SciComm, UNC Writing Center, Elsevier)	ChatGPT's revision and length adjustments	Conceptual exemplification	Practice strategies	Model outline for introductions	Abstract word-count ranges; title rules; refined survey-paper outline	Presentation & poster guidance
Presentation Guidance	-	-	-	-	-	-	Oral vs. poster

Note. Cells summarize the dominant content moves per week; some items (e.g., resources) extend across weeks as the discourse evolves.

Ashley Beccia is a doctoral candidate in Applied Linguistics at Teachers College, Columbia University. Her research interests include child SLA, task-based language teaching, and complex dynamic systems theory. Correspondence should be sent to Ashley Beccia, E-mail: ajb2301@tc.columbia.edu.

Sue Min Park is a doctoral candidate in Applied Linguistics at Teachers College, Columbia University. Her primary research interest is digital language learning, particularly the affordances of digital games for incidental vocabulary acquisition.

Jill Williams is a doctoral student in Applied Linguistics at Teachers College, Columbia University. Her research interests include foreign language learning anxiety, complex dynamic systems theory, and second language acquisition in secondary and higher education.