Studies in Applied Linguistics & TESOL at Teachers College, Columbia University, Vol. 25, No. 2, pp. 61-66
The Forum

Co-Adaptation in Human–Machine Interaction: What We Learned and Where To Go Next

Abby Massaro & Ashley Beccia
Teachers College, Columbia University

SUMMARY OF FINDINGS

This forum set out to explore the question: How does co-adaptation unfold in an ecosystem of learner–ChatGPT dyadic interaction? Working with one graduate-level EFL learner from a seven-week interactional dataset from RECIPE4U, the three contributions approached this question from complementary perspectives: communicative naturalness and stylistic synchrony; formal and topical alignment; and cognitive-psychological trajectories indexed by LIWC and grounded in interactional moves.

Massaro and Wicker showed that the dyad began with a relatively high degree of stylistic synchrony: both interlocutors were formal, polite, and grammatically careful. Over time, however, repeated violations of Voss and Waring's (2024) naturalness principles on ChatGPT's part—its "perfect," non-egalitarian paragraphs; explicitness without anaphoric tying; and limited recipient design and contiguity—coincided with a striking divergence in style. The learner's contributions became shorter, more informal, and less carefully punctuated; closings disappeared, and discourse markers crept in. ChatGPT, by contrast, remained locked in a consistently upbeat, didactic, fully sentential, "neutral" register. The end state was a dyad that was functionally coordinated but stylistically bifurcated.

Beccia, Park, and Williams zoomed in on lexical, syntactic, and topical coordination. Lexically, they found modest but clear pockets of synchrony (e.g., *paper*, *paragraph*, *scientific writing*), usually initiated by the learner and taken up, and sometimes even over-extended, by ChatGPT. Syntactically, synchrony clustered around two interrogative–declarative pairings (*Can PRO V…? / PRO can(not) V…* and *COP SUBJ X? / SUBJ COP (not) X*), with leadership alternating: ChatGPT seeded one construction through its recurring opener, whereas the learner seeded the other through interrogatives that ChatGPT echoed in declaratives. Yet for both lexical items and constructions, alignment tended to be local and discontinuous, appearing within sessions and then fading.

By contrast, content-level synchrony was robust and longitudinal. Week after week, the learner set the agenda through brief, task-oriented prompts, while ChatGPT responded with long, expository turns that provided definitions, guidelines, templates, and rationales. Topically, the dyad moved in a coherent arc: from micro writing concerns (data commentary, editing) to tools (AntConc), then to genre and structure (IMRAD, introductions, abstracts, survey papers), and

finally to professional practices (authorship norms, presentations). Throughout, topical alignment remained high, even as turn length remained asymmetrical.

Chen and Ostolaza examined cognitive and psychological dimensions of co-adaptation, as well as the underlying interactional moves. Using LIWC, they found that ChatGPT consistently produced more analytic language——logical, structured, and academic—while the learner tended toward more authentic language marked by openness, emotional expressivity, and spontaneity. Both indices fluctuated over time, and the authors identified periods, particularly around Week 5, when the participants' scores converged, suggesting brief moments of cognitive and psychological alignment.

Their micro-interactional analysis clarified how these convergences emerged. Interactional episodes in Weeks 5–7 unfolded through recurring cycles of mismatch, correction, reformulation, and confirmation. In Week 5, the learner repeatedly challenged ChatGPT's account of the "three moves" of a scientific introduction, driving a recursive repair sequence until ChatGPT aligned with the course-taught version. In Week 6, the learner introduced a disciplinary distinction between empirical and synthesis-based survey papers, prompting ChatGPT to reframe its explanation accordingly. In Week 7, the learner raised an ethically charged authorship scenario, eliciting a shift in ChatGPT's stance toward an explicitly ethical position. Across these episodes, the learner increasingly assumed epistemic leadership, while ChatGPT adapted through uptake, reformulation, and stance adjustment. Some of these repair sequences corresponded with local peaks or temporary convergence in LIWC scores, indicating that surface-level linguistic alignment sometimes unfolded alongside psychological and/or cognitive alignment.

DISCUSSION

The three sets of analyses shed light on how co-adaptation unfolds in an ecosystem of human–machine dyadic interaction. They reveal co-adaptation as a multi-layered, uneven, and role-dependent process that emerges reliably in some dimensions, sporadically in others, and sometimes not at all.

Co-Adaptation is Multi-Layered and Uneven

Across the contributions, one throughline is clear: co-adaptation did not occur uniformly across different levels of analysis, nor did alignment at one level imply symmetry across the interactional system.

At the level of content, co-adaptation was the most robust and longitudinal. Week after week, the dyad returned to a consistent pattern in which the learner set the agenda through brief prompts and ChatGPT expanded, elaborated, and reformulated. There was also a clear topical

progression over the course of the exchange, with the dyad moving from narrow, writing-specific subject areas to broader, domain-level ones.

At the level of linguistic form, co-adaptation was selective and fragile: lexical and syntactic synchrony occurred in brief, localized bursts tied to specific constructions or task phases, but these patterns dissipated across weeks. This contrasts with human–human interaction, where interlocutors tend to align simultaneously across multiple linguistic levels (Pickering & Garrod, 2004). In the present dyad, pervasive alignment was not possible because ChatGPT lacked durable memory; once the interaction window reset, it could not carry forward lexical or syntactic constructions into subsequent weeks. As Kim and Michel (2023) note, linguistic synchrony between humans relies on persistent representational states in both interlocutors. In LLM–human interaction, however, only the human retains these states.

At the stylistic level, synchrony was the least stable. Violations of communicative naturalness led to stylistic divergence, with the learner gradually abandoning formality while ChatGPT remained anchored to a "neutral," didactic register. Meanwhile, at cognitive and psychological levels, alignment was intermittent, emerging occasionally during interactional repair, especially in the latter half of the dataset.

In short, co-adaptation was not monolithic. Depending on the level of analysis, the same dyad appeared highly aligned, partially aligned, or starkly misaligned.

The Human is the Primary Driver of Co-Adaptation

Across all three sets of analyses, the learner tended to lead the co-adaptation process. They introduced the lexical items that became sites of synchrony, initiated the topical agenda for every exchange, and drove repair cycles through which alignment deepened. Even the rare moments of stylistic coordination early in the dataset reflected the learner's uptake of ChatGPT's initial tone rather than the reverse. Further, as the dyad failed to show stylistic convergence, this, too, was driven primarily by the learner as they increasingly digressed from ChatGPT's fixed formal register.

As discussed in the forum introduction, humans are biologically and socially primed for coordination. Decades of interdisciplinary research shows that people spontaneously align with their interlocutors on multiple levels, including linguistically. Our forum shows that the human predisposition for co-adaptation appears to extend, at least in part, to interactions with machines. This aligns with prior research showing that human interlocutors often adopt language introduced by computer partners (e.g., Branigan et al., 2011; Pearson et al., 2006).

In our dataset, however, alignment was skewed in one direction. Rather than generating new pathways for alignment, ChatGPT most often responded by adopting words and phrases introduced to the ecosystem by the learner, revising content when corrected, or adjusting stance when prompted. ChatGPT produced language according to architectural defaults: fully sentential responses, didactic expansion, polite formality, and bounded context windows. This asymmetry

does not preclude co-adaptation, but it does imply that co-adaptation unfolded within the learner's interactional frame, not the LLM's.

Co-Construction of Meaning Functions as an Impetus for Synchrony

Another finding that emerged from the three pieces is that major shifts in synchrony could be found across dimensions in the latter weeks of the dataset, especially in Week 5. Chen and Ostolaza found greater patterns of psychological alignment were brought about through recurrent cycles of "correction, clarification, reformulation, and confirmation" (p. 37); these interactional moves began in earnest in Week 5. Beccia, Park, and Williams also found greater topical alignment in Week 5 as ChatGPT refined and revised its understanding of the elements of an introduction in academic writing. Massaro and Wicker, meanwhile, found that stylistic divergence started to become more pronounced this week of the interaction, triggered by repeated failure of ChatGPT to index corrections made by the learner.

Unlike prior research on co-adaptation, where alignment is facilitated by shared understanding (e.g., Clark & Brennan, 1991), here, co-adaptation emerged even as the interlocutors held opposing understandings of the topic under discussion—and alongside jostling for didactic authority. What seemed essential to the manifestation of co-adaptation in this dyad was engagement in the ongoing process of meaning construction. While alignment across more core communicative domains such as lexis and content was needed to co-construct meaning, synchrony in style and syntax was not. The perceived interpersonal and epistemic distance between interlocutors may have attenuated the learner's willingness to align in these more accessory domains, while the pressure to maintain mutual intelligibility may have "overrode" this distance in domains critical to conveying meaning.

In fact, this perceived distance may have *triggered* greater alignment. Prior research has shown that humans are *more* inclined to show synchrony with computers than humans—and further, are more inclined to align (at least at the lexical level) with basic computers relative to advanced computers (Branigan et al., 2011; Pearson et al., 2006). Brantigan and her team (2011) say that this increased inclination to synchronize may be due to the fact that human interlocutors find it difficult to know what a computer knows and doesn't know; using the language that the machine has introduced, even when the linguistic choices are less common, might be a way of "supporting" the computer and helping ensure that communication is successful (Pearson et al., 2006). In the context of the present dataset then, as the distance between the machine and human increases—marked by the human's growing awareness of how ChatGPT failed to index what was taught in the classroom and its inability to index the corrections offered by the learner—the more the learner may have felt it necessary to adapt to the machine. Style and syntax, as mentioned above, not being essential to conveying meaning, may have fallen by the wayside.

SURFACED QUESTIONS

Several questions surface from this forum. A few that appear especially pressing include:

1. Within a given human–machine dyadic interaction (e.g., conversation), why does co-adaptation emerge more robustly at some levels than others?

2. What is the role of memory in human–machine co-adaptation? Would formal or stylistic alignment become more sustained if the LLM retained information across sessions?

3. To what extent is co-adaptation task-dependent? How might co-adaptation differ in other academic tasks (e.g., narrative writing), as well as non-academic tasks (e.g., gaming)?

4. Does human–machine co-adaptation result in L2 development?

FUTURE DIRECTIONS

This forum points to avenues for future SLA research on human–LLM co-adaptation. First, there is a need to broaden the empirical base. Examining multiple learners, LLM architectures, including models with persistent memory, and instructional contexts would help determine which co-adaptive patterns are learner-driven, model-driven, or emergent from particular dyads. Doing so at multiple levels of analysis as we did in this forum would shed light on why co-adaptation emerges more robustly at some levels than others, even within one conversational interaction.

Second, future studies may expand beyond a single task domain. Because scientific writing foregrounds informational and analytic goals, it may naturally privilege content-level alignment. Other academic tasks and non-academic activities may elicit different balances of stylistic, linguistic, and affective alignment. Cross-task comparisons will clarify the task-dependence of co-adaptation.

Third, a central SLA priority is to investigate whether and how co-adaptation contributes to L2 development. This requires conducting longitudinal studies that measure independent outcomes, such as changes in complexity, accuracy, fluency, or grammatical knowledge, and employ analytic tools capable of capturing dynamic, multi-level change.

While the scope of this forum was quite narrow, it represents a small step towards understanding the extent to which a known human phenomenon—namely, the synchrony of interlocutors in an ecosystem of dyadic interaction—extends to contexts with AI conversation partners.

REFERENCES

Branigan, H., Pickering, M., Pearson, J., McLean, J., & Brown, A. (2011). The role of beliefs in lexical alignment: Evidence from dialogs with humans and computers. *Cognition, 121*(1), 41–57. https://doi.org/10.1016/j.cognition.2011.05.011

Clark, H. H., & Brennan, S. A. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). APA Books. https://doi.org/10.1037/10096-006

Kim, Y.-J., & Michel, M. (2023). Linguistic alignment in second language acquisition: A methodological review. *System, 115*(103007). https://doi.org/10.1016/j.system.2023.103007

Pearson, J., Hu, J., Branigan, H., Pickering, M., & Nass, C. (2006). Adaptive language behavior in HCI: How expectations and beliefs about a system affect users' word choice. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM Digital Library,* 1177–1180. https://doi.org/10.1145/1124772.1124948

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*(2), 169–226. https://doi.org/10.1017/S0140525X04000056

Voss, E., & Waring, H. Z. (2024). When ChatGPT can't chat: The quest for naturalness. *TESOL Quarterly*, *59*(2). https://doi.org/10.1002/tesq.3374

Abby Massaro is a doctoral student in Applied Linguistics at Teachers College, Columbia University. Her research interests include heritage language acquisition, instructed SLA, and language attrition and reacquisition.

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.