

# What Changes When the Interlocutor Is an AI? Interactional Fluency and Linguistic Uptake in L2 Spoken Dialogue

Russell Scheinberg  
Portland State University  
rschein2@pdx.edu

Kalab Kahsay  
Portland State University  
kalabk@pdx.edu

Ameeta Agrawal  
Portland State University  
ameeta@pdx.edu

Nina Vyatkina  
University of Kansas  
vyatkina@ku.edu

Tetyana Sydorenko  
Portland State University  
tsydorenko@pdx.edu

Griet Boone  
University of Antwerp  
griet.boone@uantwerpen.be

## ABSTRACT

Voice-based AI systems are increasingly used for L2 speaking practice, but evaluations rarely characterize the interactional processes they create. We analyze 78 university learners of German across four sites completing a counter-balanced spot-the-difference task with both a human peer and a real-time AI partner. From diarized ASR transcripts, we extract measures of interactional fluency, linguistic uptake, and learner experience. Human dialogue was faster and more balanced, with many short turns; AI dialogue resembled supported monologue, with fewer, longer turns, reduced learner floor share, and greater within-turn fluency. The AI's verbose, syntactically regular input was associated with greater short-term uptake and stronger syntactic priming after controlling for input volume. Attitudes toward AI improved after the task, and satisfaction was predicted by production fluency rather than uptake. The results show complementary affordances for AI and human dialogue in L2 practice.

## Keywords

interactional fluency, L2 speech, voice-based AI, turn-taking, educational data mining

## 1. INTRODUCTION

Voice-based conversational AI is increasingly positioned as a scalable speaking partner for second language (L2) learners. Yet these systems do not simply provide more practice, they change the *nature* of the interaction [1, 2, 29, 11, 22]. AI-based interlocutors, especially using large language models (LLMs), produce longer turns, respond at atypical moments, and lack the backchannels and floor-negotiation strategies of human partners [22, 29, 11, 30, 5]. If interaction structure shapes learning opportunities, then characterizing *how* AI-mediated dialogue differs from human dialogue is essential

Russell Scheinberg, Ameeta Agrawal, Tetyana Sydorenko, Kalab Kahsay, Nina Vyatkina, and Griet Boone. What Changes When the Interlocutor Is an AI? Interactional Fluency and Linguistic Uptake in L2 Spoken Dialogue. In Anthony Botelho, Maria Mercedes T. Rodrigo, Adish Singla, Hiroaki Ogata, Hyojeong So, and Young Hoan Cho (eds.) Proceedings of the 19th International Conference on Educational Data Mining, Seoul, Republic of Korea, June, 2026, pp. 376–387. International Educational Data Mining Society (2026).

© 2026 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.21039907>

for informed system design and deployment.

Most evaluations of dialogue-based computer-assisted language learning (CALL) focus on outcome measures: pre/post speaking scores, perceived confidence, or learner satisfaction. Meta-analytic evidence confirms positive effects overall but highlights substantial heterogeneity across systems, tasks, and evaluation criteria [2, 17, 10]. What remains underexplored is the *interactional process* itself—how learners and systems jointly construct spoken interaction during practice [19].

We address this gap through three complementary lenses. The first is *interactional fluency*: measurable properties of conversational coordination—response latency, turn-length distributions, floor sharing, and within-turn pausing—that reflect how speakers manage participation in real-time dialogue. While the distinction between monologic and dialogic fluency is well established in SLA research [38], less is known about how these dynamics shift when the interlocutor is a machine.

The second is *linguistic uptake*. Conversational alignment theory holds that interlocutors converge on lexical and syntactic forms during interaction [28], a process proposed as a mechanism supporting L2 learning [6, 23]. When partners differ fundamentally in how they generate language, it is an open question whether they afford similar uptake opportunities. We operationalize this as direction-sensitive lexical and morphosyntactic uptake: the extent to which a learner reuses words or constructions introduced by the interlocutor in the preceding context.

The third is *learner experience*. If AI and human dialogue create different interactional conditions, learners may perceive these differences, and the behavioral features of the interaction (e.g., production fluency, floor balance) may predict satisfaction independently of what learners linguistically adopt.

These three lenses motivate the following research questions:

**RQ1:** How do **interactional fluency** measures differ between human–human and human–AI dialogue?

**RQ2:** Do learners differ in **lexical and morphosyntactic uptake** from a human partner versus an AI-interlocutor, and do these differences vary across linguistic levels?

**RQ3:** How do learner **attitudes toward AI** change after direct experience, and what behavioral features of the interaction predict satisfaction?

To investigate these questions, we conducted a controlled, multi-site study with second- and third-year L2 German learners ( $N = 78$ ) at four universities in the United States and Europe. Participants completed a spot-the-difference task under four within-speaker conditions: two monologic scene descriptions and two dialogues (one with a human peer, one with a real-time AI partner). All measures are extracted from diarized, timestamped ASR transcripts without manual coding. Results reveal two distinct interaction modes: human dialogue is rapid and balanced, whereas AI dialogue resembles a supported monologue, with fewer but longer turns, slower responses, reduced learner floor share, and greater within-turn fluency. AI verbosity and syntactic regularity afford more uptake opportunities, and learners show stronger short-term syntactic priming with the AI partner even after controlling for input volume. Attitudes toward AI improve after the task, with satisfaction tied to production fluency rather than linguistic uptake.

This paper makes three contributions: (1) we operationalize interactional fluency and direction-sensitive uptake for spoken AI practice using metrics extractable from diarized ASR output without manual linguistic annotation; (2) we leverage a within-speaker design across four instructional sites to isolate how the same learner’s interaction changes with a human versus AI partner; and (3) we provide a reusable analysis pipeline for rapid, scalable evaluation of conversational AI in educational settings. The remainder of the paper describes the study design (§3), methods and results for interactional fluency (§4.1), uptake (§4.2), and survey analyses (§5), followed by discussion and design implications (§6).

## 2. RELATED WORK

Prior work establishes the promise of conversational practice for language learning, but also suggests that who (or what) the interlocutor is can reshape the interaction itself.

*Conversational practice and transcript-based measurement.* Dialogue-based computer-assisted language learning (CALL) has a long history motivated by the role of interaction in second language development. Meta-analytic evidence suggests that dialogue-based CALL systems can yield positive learning outcomes overall, while also reporting substantial heterogeneity across system types, tasks, and evaluation measures [2]. Recent reviews of conversational agents for language learning echo both pedagogical promise and persistent concerns about reliability, pedagogical alignment, and evaluation practices [17, 10]. As voice-based LLM systems become plausible speaking partners, a central question is not only whether learners improve, but what kinds of interaction these systems create and how those interactional properties shape practice opportunities [19].

A parallel development in learning analytics is the growing treatment of dialogue as analyzable behavioral trace data. This shift enables process-level evaluation of learning-support interactions, but it also raises a practical challenge for spoken settings: extracting reliable measures without extensive manual annotation. Recent work demonstrates that transcript-based pipelines can extract educationally meaningful structure from naturalistic collaborative speech, enabling scalable, fine-grained analysis [39]. This line of work motivates our use of diarized, timestamped transcripts to quantify interaction dynamics and linguistic behavior across interlocutor conditions and sites.

### *Interactional fluency and turn-taking with AI partners.*

Within second language acquisition (SLA) research, fluency is often decomposed into speed, breakdown, and repair, and dialogic fluency work argues that monologic temporal measures can miss key interactional phenomena in conversation [38]. Overlap, interruptions, and between-turn gaps can materially affect fluency assessments, motivating measurement at the level of turns and turn transitions. Pauses, repetitions, and stalling may likewise function as interactional resources rather than purely as disfluency [25]. These perspectives build on classic accounts of turn-taking norms that describe strong pressure to minimize gaps and avoid overlap, with response latency as a central coordination variable [33, 37, 20].

When one interlocutor is a machine, these norms can shift systematically. Studies of voice interfaces document altered turn design and recurring breakdown and repair patterns [29, 21], and users may perform additional interactional work to maintain progressivity when systems are slow or insufficiently contingent [11]. With LLM-based agents, system contributions become more content-rich, yet interface timing and modality still constrain turn-taking and repair; users often scaffold the interaction through reformulations and related strategies [30]. Emerging work in language learning contexts suggests that L2 learners’ turn-taking practices may shift when interacting with conversational AI [5]. Interaction structure is also increasingly treated as consequential for learning processes, with dialog acts in AI-supported peer tutoring linked to learning rates [3]. Together, these strands motivate our analysis of interactional fluency as a way to characterize how human–AI dialogue differs from human–human dialogue in the interactional conditions it affords.

### *Lexical and morphosyntactic uptake as alignment in dialogue.*

Psycholinguistic theories of alignment propose that interlocutors tend to converge on lexical and syntactic representations through priming mechanisms, supporting efficient coordination [28]. Related work on lexical entrenchment and conceptual pacts shows how partners establish shared referring expressions as part of building common ground [4]. In SLA, alignment has been proposed as a mechanism that can support learning, since repeated exposure to and reuse of forms in interaction may increase accessibility and entrenchment [6, 23], and empirical work suggests that alignment patterns vary with interactional context [6, 18]. In the present one-shot design, we therefore treat uptake as short-term mining or alignment, not as evidence of durable

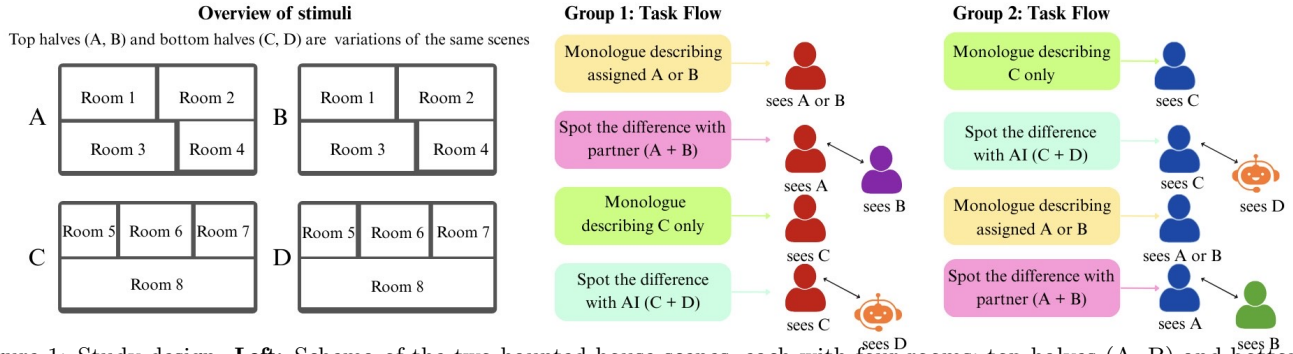


Figure 1: Study design. **Left:** Schema of the two haunted-house scenes, each with four rooms; top halves (A, B) and bottom halves (C, D) are slight variations of the same scenes. In the human condition each partner sees a different version (A or B), while in the AI condition the participant describes C and the AI is provided a description of D. **Right:** Task flow for the two counterbalanced groups. Each participant completes four tasks: two monologues and two dialogues (one with a human partner, one with the AI). Group 1 completes the human block first; Group 2 completes the AI block first.

learning; repeated alignment may provide affordances for learning over time, but our data measure immediate reuse opportunities rather than acquisition.

Measuring alignment at scale remains challenging. However, computational dialogue research offers measures ranging from global similarity to targeted accounts of repetition and construction reuse [32]. Work in human computer interaction further shows that people align their language with spoken dialogue systems, and that the strength and form of this alignment can vary with system framing and design [7, 35]. Educational natural language processing has also moved toward turn-linked measures that capture responsiveness beyond surface similarity [9]. Recent work similarly demonstrates the feasibility of automated, language-based measurement in educational settings, including LLM-based detection of learning-relevant constructs in verbal protocols [41]. At the same time, these approaches foreground familiar concerns about construct validity and fairness in model-based inference [36]. Building on these threads, we operationalize lexical and morphosyntactic uptake as a direction-sensitive form of alignment, measuring the extent to which learners reuse lexical items and morphosyntactic patterns introduced by the interlocutor in the immediately preceding context.

*Learner affect and attitudes.* CALL studies of chatbot and LLM-based speaking practice frequently include self-report measures—perceived usefulness, enjoyment, comfort, willingness to communicate—since these perceptions shape sustained engagement [17, 12]. Recent work shows that some conversational AI designs can increase willingness to communicate and reduce speaking anxiety [40], and that perceived enjoyment and satisfaction predict intentions to continue using AI tools for speaking practice [16]. However, these attitudinal outcomes are rarely connected to fine-grained interaction traces, leaving open which moment-by-moment behaviors drive learner perceptions—a gap we address by relating survey responses to transcript-derived interaction metrics.

### 3. DATA COLLECTION

### 3.1 Task and Stimuli

Participants completed a spot-the-difference task, a cooperative information-gap activity widely used in communicative language teaching [27]. A human artist created a “haunted house” scene containing eight rooms, each drawn in two matched versions with systematic differences in object placement (e.g., repositioned ghosts, added or removed objects). Participants communicated in the L2 language to identify these differences. The eight rooms were roughly matched for complexity and divided into two sets of four; one set was used for the human dialogue condition and one for the AI dialogue condition. Figure 1 provides an overview of the study design.

### 3.2 Conditions and Procedure

The study used a within-participant design with four speaking conditions along two dimensions: (i) task format (monologue vs. dialogue) and (ii) partner type (human peer vs. AI model).

Participants were assigned to a counterbalanced design, with one group completing the human–human condition first and the other completing the human–AI condition first. In the human–human condition, participants first completed a monologue describing either image A or image B, and then engaged in a dialogue about the same room set with a human partner who had previously described the complementary version. In the human–AI condition, all participants first described image C, followed by a dialogue with the AI system, which was provided a textual description of image D.

In human dialogues, participants were paired with a classmate and completed the task face-to-face or on Zoom; no constraints on turn-taking or strategy were imposed. AI dialogues were conducted through a web application with audio input. This modality asymmetry is representative of how each mode would typically be deployed, but co-presence in the human condition provides non-verbal coordination cues unavailable in the AI condition. Each session lasted approximately 90 minutes including a background survey, vocabulary assessment, the four speaking tasks, and a post-task survey. Data were collected between October and December 2025.

All procedures were approved by the institutional review boards at the participating sites. Before completing the background survey and speaking tasks, participants provided informed consent for audio recording and transcript-based analysis; recordings without valid consent were excluded.

### 3.3 AI Interlocutor System

We developed the AI system using OpenAI’s `gpt-realtime` model<sup>1</sup>, selected for its native speech-to-speech capability (eliminating separate ASR/TTS latency), support for a plurality of languages, and real-time turn-taking via server-side voice activity detection. Because direct image input produced excessive hallucination during piloting, we provided the AI with manually verified English-language text descriptions of each room, written by the research team and checked against the images for accuracy. This choice kept the AI grounded in the task materials but creates a stimulus-representation asymmetry: human partners viewed images directly, whereas the AI received textual descriptions of the rooms.

At session initialization, all room descriptions were included in the system prompt<sup>2</sup>; at each room transition, the current room’s description was re-sent. The AI retained full conversation history within a session.

The system parameters were set as follows: voice `alloy`, temperature 0.8, max response tokens 150. Turn-taking used server-side VAD with an open microphone (prefix padding 300 ms, silence threshold 200 ms). After pilot testing, we found it useful to instruct the AI model to use informal register, take short turns (one to two sentences), allow time for learner responses, and avoid excessive praise. Despite these constraints, the AI consistently occupied a larger share of the conversational floor than human partners.

### 3.4 Participants

All participants were enrolled in second- or third-year university German courses at one of four sites: two Flemish universities in Belgium ( $n = 24, 34$ ), a Midwestern US university ( $n = 14$ ), and a Pacific Northwest US university ( $n = 14$ ), totaling  $N = 86$ . After excluding non-consented, failed, and incomplete recordings,  $N = 78$  participants remained, yielding 78 AI dialogues and 39 human-human conversations (one per dyad). Per-analysis sample sizes vary by minimum-turn thresholds and are reported with each analysis. Approximately 50% of participants listed Dutch as their sole L1, English was the second largest group (34.5%), and the remaining 15% comprised speakers of Romanian, Vietnamese, Spanish, Albanian, Hindi, and Ndebele, several of whom reported bilingual combinations with Dutch.

### 3.5 Transcription

We analyze diarized transcripts rather than raw audio, a deliberate choice favoring scalability: transcript-based measures avoid prosodic confounds across L1 backgrounds and

recording environments, require no manual coding, and yield an easily replicable pipeline.

All sessions were transcribed using ElevenLabs `scribe_v2` with automatic speaker diarization (language: `deu`), producing word-level timestamps and speaker attribution. German-speaking authors reviewed all transcripts, verifying speaker labels, trimming non-task speech at recording boundaries, and confirming coherence with assigned room content. Word-level ASR errors were not systematically corrected. In three AI-dialogue transcripts, mid-stream speaker label swaps were detected via an automated disfluency-rate classifier (the AI produces zero filled pauses) and corrected manually.

## 4. METHODS

From the diarized, timestamped transcripts we extract two complementary sets of metrics. *Interactional fluency* (§4.1) captures how speakers coordinate in real time – how long they hold the floor, how quickly they respond, and how much they hesitate within turns. *Linguistic uptake* (§4.2) captures what speakers do with the language they hear – whether they adopt vocabulary or reproduce morphosyntactic patterns from the interlocutor’s speech. The first captures *how* learners interact; the second captures *what language* they take away from the interaction.

### 4.1 Interactional Fluency

Dialogue involves continuous coordination over who speaks, for how long, and when transitions occur. These properties can differ systematically between human and AI partners, shaping the kind of practice an interaction affords. We capture this through seven features extracted from word-level timestamps, organized into two groups.

*Conversation structure.* Four features describe the overall shape of a dialogue. *Conversation length*: total session duration in seconds. *Turn count*: number of participant turns. *Mean turn duration*: average length of a participant turn in seconds. *Words per turn*: average words per participant turn. Together, these indicate whether the dialogue consists of many short exchanges or fewer, longer stretches of speech.

*Turn-taking dynamics.* Three features capture how speakers manage transitions and share the floor. *Response latency*: the gap between the partner’s last word and the participant’s first word; only gaps within  $[0, 10]$ s are retained. *Floor ratio*: the proportion of session time the participant is speaking. *Within-turn pause ratio*: the proportion of a participant’s turn spent in silence ( $\geq 0.3$ s gaps between words), reflecting online planning difficulty. Figure 2 illustrates how within-turn pauses are distinguished from between-turn gaps.

Session-level metrics (conversation length, turn count) were compared using paired *t*-tests ( $n = 78$ ). Per-turn metrics use sessions with  $\geq 8$  participant turns, yielding unbalanced samples; these were analyzed using linear mixed-effects models with condition as a fixed effect and participant as a random intercept.

<sup>1</sup>Specifically, `gpt-realtime-2025-08-28`

<sup>2</sup>An anonymized repository containing the full LLM prompt is available at [https://osf.io/9kxsu/overview?view\\_only=ff7b5f764f1647a788c8541bc0b3fc31](https://osf.io/9kxsu/overview?view_only=ff7b5f764f1647a788c8541bc0b3fc31).

**Within-turn pause ratio: only gaps between consecutive words within the same speaker's turn are counted**




Figure 2: Illustration of temporal metrics extracted from word-level timestamps. Blue blocks are participant words; orange blocks are partner words. *Within-turn pause ratio* counts only gaps  $\geq 0.3$ s between consecutive words within the same speaker’s turn (green). Gaps between speakers’ turns (gray) are excluded from the pause ratio but are used to compute *response latency* (purple). Turn boundaries are reconstructed by grouping consecutive words from the same speaker.

## 4.2 Linguistic Uptake

Interactive alignment theory [28] predicts that interlocutors converge on shared linguistic representations during dialogue, from lexical choices to syntactic structures. We operationalize this as linguistic uptake: whether L2 learners adopt linguistic material from their interlocutor’s speech. Figure 3 presents an illustration of how uptake is estimated. We measure uptake at two complementary levels:

**Lexical uptake** tracks whether learners adopt vocabulary items (lemmas) introduced by the partner, capturing breadth of vocabulary activation across the full session.

**Syntactic uptake** tracks whether learners reproduce morphosyntactic patterns (e.g., determiner–noun agreement, prepositional case government) from the partner’s prior speech, capturing structural alignment within a dialogue.

Both measures impose a **first-use constraint**, such that each linguistic unit is scored only when it first appears in the learner’s speech. This operationalization follows an input-mining paradigm in L2 research, where learners’ later production is interpreted relative to prior independent production and intervening input [15]. In our design, the monologue preceding each dialogue provides a learner-specific baseline for lexical uptake, while syntactic uptake applies the same conservative first-use logic within dialogue by crediting templates only after prior partner use within a bounded recency window. We use uptake/mining to mean short-term reuse of available input, not acquisition; repeated task performance and the shared visual context may also contribute to some cases.

### 4.2.1 Lexical Uptake

The lexical uptake algorithm uses a **global timeline** approach, tracking each lemma across all four tasks in sequential order. The counterbalanced design creates two orderings:

- **Group 1** ( $n = 46$ ): Monologue A/B  $\rightarrow$  Human dialogue

$\rightarrow$  Monologue C  $\rightarrow$  AI dialogue

- **Group 2** ( $n = 32$ ): Monologue C  $\rightarrow$  AI dialogue  $\rightarrow$  Monologue A/B  $\rightarrow$  Human dialogue

The participant’s first monologue establishes a **baseline vocabulary**. Words are tokenized, filtered (removing determiners, 49 pre-task vocabulary items, 153 English code-switching tokens, and words  $\leq 2$  characters), and lemmatized using spaCy’s `de_core_news_sm` model. Words appearing on the Goethe-Institut A1 wordlist [26], approximately 714 high-frequency items, are also excluded as uninformative adoption candidates. A lemma counts as **adopted** if: (1) the interlocutor produced it before the participant, (2) it is not in the participant’s baseline, and (3) it has not already been attributed to another source (**first-source-wins** rule). We use “adopted” rather than “learned” because we are measuring activation through interactive exposure, not acquisition.

Adoption is measured cumulatively without a recency window, consistent with accounts of lexical entrainment as a *historical* phenomenon in which interlocutors maintain partner-specific conceptual precedents beyond immediate adjacency [4], and with L2 alignment work that omits distance constraints under an interaction-as-learning framing [24]. Because the AI interlocutor produces substantially more speech than human partners, we report both raw adoption counts and adoption rates controlled for verbosity by dividing by the number of unique novel lemmas each interlocutor introduced beyond the participant’s baseline (**pool size**). To test whether interaction type predicts adoption beyond input volume, we fit a mixed-effects regression with condition, pool size, and learner word count as fixed effects and participant as a random intercept.

### 4.2.2 Syntactic Uptake

Syntactic uptake measures whether learners reproduce morphosyntactic patterns from the interlocutor’s speech **within**

## How linguistic uptake is measured

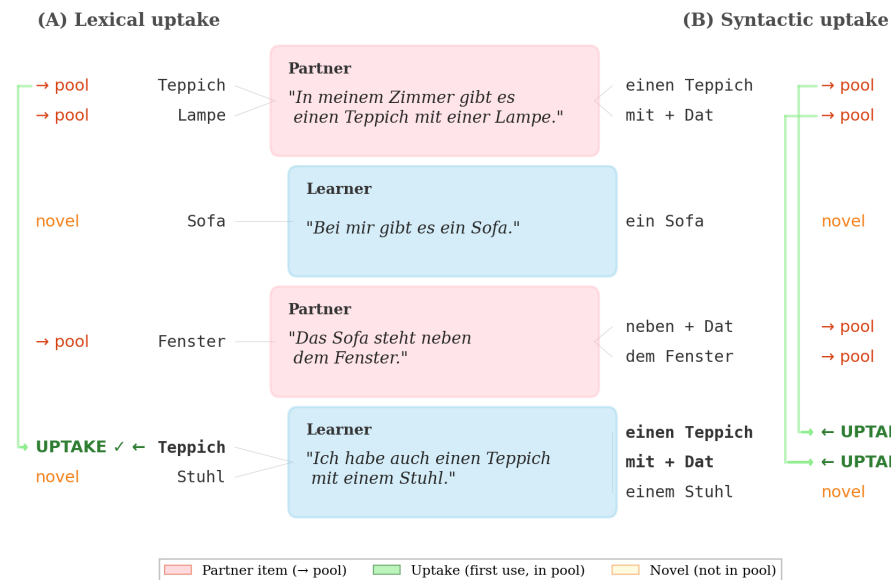


Figure 3: Illustration of linguistic uptake measurement at two levels, applied to the same example dialogue ( $N_{AI} = 78$ ,  $N_{Human} = 78$  in the full dataset). **(A)** Lexical uptake: content words used by the partner accumulate in a pool; when the learner first produces a matching word, it is counted as uptake. **(B)** Syntactic uptake: morphosyntactic templates—determiner–noun agreement (e.g., *einen Teppich*, masculine accusative) and preposition case government (e.g., *mit + dative*)—are tracked analogously. Green arrows trace each uptake event to its source in the partner’s speech. Only the learner’s first use of each item counts (first-use constraint); repeated uses are not double-counted. For syntactic templates, grammatical correctness is additionally validated against a noun-gender lexicon (e.g., *Teppich* → masculine).

Table 1: Target morphosyntactic patterns for syntactic uptake.

Pattern	Example	What it captures
DET-NOUN	<i>der Tisch</i> (Masc/Nom/Sg)	Det–noun gender/case/number agreement
ADP-CASE	<i>mit dem Hund</i> (mit → Dat)	Preposition case government
ADJ-NOUN	<i>großes Gemälde</i> (Neut/Nom/Sg)	Attributive adjective agreement
CONTRACTION	<i>im</i> (= <i>in + dem</i> , Dat)	Preposition–article fusion

**a single dialogue.** We extract four pattern types that constitute the core of German nominal morphology (see Table 1): determiner–noun agreement, prepositional case government, attributive adjective agreement, and preposition–article contractions. These are pervasive in the data ( $M = 86$  learner-produced template instances per AI dialogue, 113 per human dialogue) and target well-documented areas of L2 German difficulty.

**Extraction pipeline.** Each turn is lightly preprocessed (removing filled pauses such as *äh* and *ähm*) and parsed with the Stanza German UD pipeline [31]. Each extracted template has an **anchored** representation that includes the noun lemma, surface form, and morphological features (e.g., DET-NOUN|*tisch*|*dem*|Masc|Dat|Sg), tying uptake to specific lexical items rather than abstract morphological categories.

**Correctness checking.** For DET-NOUN and ADJ-NOUN templates, we verify gender agreement against a lexicon of 91,270

German nouns. For ADP-CASE, correctness is checked for strict-case prepositions (e.g., *mit* → dative); two-way prepositions (e.g., *auf*, *in*), which can take dative or accusative case, are included in uptake matching but excluded from correctness analysis, as case selection depends on semantic context that cannot be determined from syntax alone. Contractions are inherently correct.

**Uptake scoring.** On the participant’s first production of a template, we check whether it appeared in the source’s prior speech. Unlike lexical adoption, syntactic uptake uses a recency window, reflecting the lag-sensitivity of structural priming documented in both L1 [14] and L2 task-based interaction [8], where prime–target contingency is typically coded within a bounded span [24]. Because AI turns are longer and syntactically denser than human turns, a fixed turn-count window captures unequal amounts of input across conditions. We therefore compare three window operationalizations: **turn-based** (last 3 source turns), **time-based** (last 15 seconds of source speech), and **template-count** (last 15 source




Figure 4: Conversation structure: AI vs. Human partner. (A) Participant turn count and (B) conversation length use all sessions ( $n = 78$  per condition). (C) Words per turn and (D) mean turn duration use sessions with  $\geq 8$  participant turns ( $n_{AI} = 63$ ,  $n_{Human} = 72$ ). Box plots show median and interquartile range; individual sessions are overlaid as points. \*\*\*  $p < .001$ .

Turn-taking dynamics (sessions with  $\geq 8$  participant turns; AI  $n=63$ , Human  $n=72$ )




Figure 5: Turn-taking dynamics: AI vs. Human partner (sessions with  $\geq 8$  participant turns). (A) Overlapping histograms of mean response latency show clear separation between conditions. (B) Floor ratio: learners hold the floor for a smaller share of session time with the AI. The dashed line marks equal floor time (0.5). (C) Within-turn pause ratio: learners pause *less* within turns when speaking with the AI. \*\*\*  $p < .001$ .

templates, explicitly equalizing input volume). All results are tested via mixed-effects regression with condition, source pool size, and learner unique template count as fixed effects and participant as a random intercept.

The algorithm also detects **error-correction sequences**: the participant produces an incorrect form, the source subsequently models the correct form, and the participant later produces the corrected version.

## 5. RESULTS

We present results in three parts. Interactional dynamics (§4.1) characterize how dialogue structure and turn-taking differ between AI and human partners. Lexical and syntactic uptake (§4.2) examine whether and how learners adopt language from each interlocutor type. Survey results (§5.3) connect interaction behavior to learner attitudes and satisfaction.

### 5.1 Interactional Dynamics

We organize the findings into three domains: conversation structure, turn-taking dynamics, as well as consistency of

results across four sites.

#### 5.1.1 Conversation Structure

Human-human dialogues were significantly longer ( $M = 764$  s vs.  $608$  s,  $d = -0.44$ ,  $p < .001$ ) and contained nearly three times as many turns ( $M = 66.5$  vs.  $22.2$ ,  $d = -1.35$ ,  $p < .001$ ). Human partners engaged in rapid back-and-forth exchange, while AI dialogues were characterized by fewer, longer turns. With the AI, participants produced significantly longer turns ( $M = 10.8$  s vs.  $7.7$  s,  $d = 0.39$ ,  $p < .001$ ) containing more words per turn ( $M = 18.0$  vs.  $12.1$ ,  $d = 0.51$ ,  $p < .001$ ). Learners pack more content into each turn with the AI, but the threefold reduction in turn count means that total learner output is substantially lower than in human dialogue.

Figure 4 illustrates this trade-off: human dialogues produce many more turns (A), but learners produce more words per turn with the AI (C).

#### 5.1.2 Turn-Taking Dynamics

Site consistency (sessions with  $\geq 8$  participant turns)

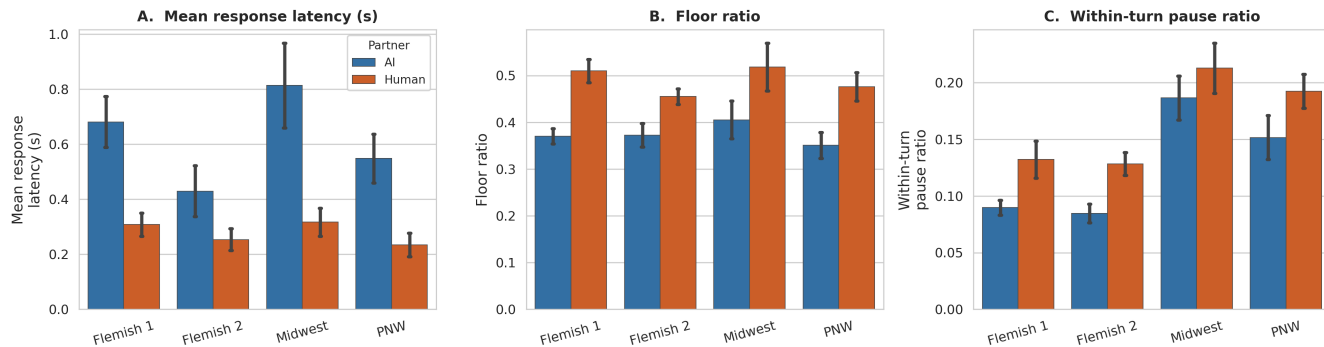


Figure 6: Site consistency: the three turn-taking metrics broken down by site and partner type (sessions with  $\geq 8$  turns). Error bars show  $\pm 1$  SE. At every site, response latency is higher with the AI, floor ratio is lower, and within-turn pause ratio is lower. No site reverses the direction of any effect.

Three interrelated measures reveal qualitatively different turn-taking patterns across conditions. Figure 5 visualizes the three turn-taking metrics, showing clear separation between conditions for response latency (A) and consistent differences in floor ratio (B) and within-turn pause ratio (C).

Participants took roughly twice as long to begin speaking after the AI ( $M = 0.61$  s) as after a human partner ( $M = 0.32$  s;  $d = 0.90$ ,  $p < .001$ ), indicating substantially longer response latencies in AI dialogue. Most human-partner latencies cluster below 0.3 s, while AI latencies spread between 0.2–1.0 s, likely reflecting the AI’s longer, more information-dense turns.

Learners held the floor for a smaller proportion of session time when interacting with the AI ( $M = 0.37$ ) than with a human partner ( $M = 0.48$ ;  $d = -1.14$ ,  $p < .001$ ), representing one of the largest effects observed in the dataset. The AI’s longer turns account for much of this difference, as they occupy a greater share of the conversational time.

Learners paused *less* within their own turns when speaking with the AI ( $M = 0.12$  vs. 0.15;  $d = -0.53$ ,  $p < .001$ ), indicating greater within-turn fluency. When learners do hold the floor with the AI, they tend to speak more fluently within each turn.

Taken together, these measures show a consistent shift in interactional organization: AI dialogue features slower turn transitions and reduced learner floor time, but more fluent production when learners hold the floor.

### 5.1.3 Site Consistency

Figure 6 shows the contrasts were consistent in direction across all four sites. Response latency was higher with the AI at every site; floor ratio was lower; within-turn pause ratio was lower. The magnitudes varied—Midwest showed the largest response-latency difference, while Flemish 2 showed the smallest floor-ratio gap—but no site reversed the direction of any effect.

## 5.2 Linguistic Uptake

### 5.2.1 Lexical Uptake

The AI produced over twice as many novel lemmas as human partners (pool size:  $M = 79.6$  vs. 36.4), yielding a higher descriptive adoption rate for human partners (16.2% vs. 9.3% of pool adopted). However, a mixed-effects regression controlling for pool size and learner word count found no significant condition effect ( $\beta = -1.77$ ,  $SE = 0.95$ ,  $z = -1.87$ ,  $p = .062$ ). Pool size ( $\beta = 0.10$ ,  $p < .001$ ) and learner word count ( $\beta = 0.005$ ,  $p < .001$ ) were the significant predictors: after accounting for input volume and learner output, condition does not reliably predict adoption.

To disentangle partner type from task order, we compared adoption at the same sequential position. Table 2 presents the results of this analysis. In position 2—where the depleted vocabulary pool provides a more sensitive test—AI participants adopted significantly more words ( $M = 4.6$  vs. 2.6;  $d = 0.72$ ,  $p = .003$ ). In position 1, the direction is consistent but nonsignificant ( $d = 0.42$ ,  $p = .079$ ).

### 5.2.2 Syntactic Uptake

A total of 156 participant-level dialogue records were analyzed: 78 AI dialogue records, and 78 human-partner records derived from 39 human-human conversations. The results are summarized in Table 3. Learners’ cumulative uptake rate was higher with the AI ( $M = .16$ ,  $SD = .11$ ) than with human partners ( $M = .08$ ,  $SD = .06$ ;  $d = 0.93$ ), but as with lexical adoption, the AI also provided a substantially larger source pool ( $M = 155$  vs. 90 unique templates;  $d = 1.35$ ). After controlling for pool size and learner unique template count, cumulative uptake showed a positive but nonsignificant trend ( $\beta = 1.55$ ,  $SE = 0.92$ ,  $z = 1.69$ ,  $p = .090$ ). Recency-windowed uptake told a different story: the AI advantage was significant across all three operationalizations—turn-based ( $\beta = 3.18$ ,  $p < .001$ ), time-based ( $\beta = 3.04$ ,  $p < .001$ ), and template-count ( $\beta = 2.84$ ,  $p = .001$ )—including the window that explicitly equalizes input volume. This confirms that the AI’s short-term priming advantage reflects genuine structural priming, not merely a larger input pool.

Correctness rates were comparable: AI  $M = .82$  ( $SD = .13$ )

Table 2: Lexical adoption by sequential position and partner type. Group 1 ( $n = 46$ ) had the human partner first; Group 2 ( $n = 32$ ) had the AI first. Position-matched comparisons use between-subjects Welch’s  $t$ -tests.

Position	Partner type	$n$	$M$	$SD$	$t$	$p$	$d$
First dialogue	Human (Group 1)	46	8.6	5.5	1.79	.079	0.42
	AI (Group 2)	32	11.2	6.8			
Second dialogue	Human (Group 2)	32	2.6	2.8	3.12	.003**	0.72
	AI (Group 1)	46	4.6	2.8			

Table 3: Syntactic uptake by condition. Descriptive statistics (Welch’s  $t$ -tests) and mixed-effects regression results ( $\beta$  = condition coefficient, AI = 1; models control for source pool size and learner unique templates with participant random intercept).

Measure	AI $M$ ( $SD$ )	Human $M$ ( $SD$ )	$t$	$d$	$\beta$	$SE$	$z$	$p$
Cumulative uptake (rate)	.16 (.11)	.08 (.06)	5.84	0.93	—	—	—	—
Cumulative uptake (count)	7.3 (5.6)	5.4 (4.6)	2.28	0.36	1.55	0.92	1.69	.090
Recency: 3 turns (count)	6.3 (5.2)	2.9 (2.9)	4.98	0.80	3.18	0.86	3.69	<.001***
Recency: 15 s (count)	5.7 (5.0)	2.7 (2.7)	4.65	0.74	3.04	0.84	3.61	<.001***
Recency: 15 templates (count)	5.2 (5.3)	3.2 (2.8)	3.00	0.48	2.84	0.87	3.27	.001**

Table 4: Key baseline attitudes: AI vs. Human peers ( $N=55$ ). All items rated on a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree).

Item	$M$	$SD$
Practicing speaking is valuable	4.56	0.69
AI effectiveness	3.22	0.76
AI enjoyment	2.73	0.89
Relaxation: AI	3.25	1.02
Relaxation: Proficient speakers	2.82	1.19
Access to proficient speakers	2.67	1.11

vs. Human  $M = .80$  ( $SD = .13$ );  $p = .316$ ,  $d = 0.16$ . Higher uptake with the AI does not come at the cost of morphological accuracy. Correction sequences were numerically more frequent with human partners (total: 38 vs. 27), but per-participant rates did not differ significantly ( $M = 0.49$  vs. 0.35;  $p = .215$ ).

### 5.3 Survey Analysis

Of the 78 participants, 55 completed matched pre- and post-surveys; 52 had complete data for the directional sentiment composite used in the paired  $t$ -test below. These data allow us to assess how learner attitudes toward AI shift after direct experience and whether behavioral features of the interaction predict satisfaction.

**Baseline attitudes.** Participants entered with predominantly skeptical attitudes toward AI: 60% described themselves as skeptical or unsure, and only 14.5% as enthusiastic adopters. AI-specific Likert items scored near or below neutral (perceived effectiveness  $M = 3.22$ , enjoyment  $M = 2.73$ ), while general speaking practice was highly valued ( $M = 4.56$ ). Notably, relaxation with AI ( $M = 3.25$ ) already exceeded relaxation with proficient speakers ( $M = 2.82$ ), and the lowest-scoring item, having enough opportunities to speak with proficient speakers ( $M = 2.67$ ), points to a perceived access gap (see Table 4).

**Post-experience shift.** Attitudes improved significantly after the study: a paired  $t$ -test on directional sentiment scores

showed a mean increase of +0.48 ( $t(51) = -4.58$ ,  $p < .0001$ ,  $d = 0.64$ ), with 75% of participants becoming more positive. Post-study, 67.3% liked the AI interaction, 81.8% agreed AI can be useful for L2 speaking practice, and 34.5% reported learning new vocabulary.

**Behavioral predictors of satisfaction.** To connect attitudinal outcomes to interaction behavior, we fit OLS regressions predicting post-satisfaction and attitude change from behavioral features of the AI dialogue. Table 5 presents the results of survey regressions, fit using ordinary least squares (`statsmodels`; [34]) with site as a categorical covariate.

Pre-study attitude predicted post-satisfaction ( $\beta = 0.36$ ,  $p = .006$ ), but the strongest behavioral predictor was within-turn pause ratio ( $\beta = -4.78$ ,  $p = .008$ ): learners who paused more during AI turns, reflecting greater online planning difficulty, reported lower satisfaction and smaller attitude gains ( $\beta = -5.91$ ,  $p = .006$ ). Floor ratio was not significant, suggesting that *how fluently* learners spoke mattered more than *how much* they spoke. Neither lexical adoption rate nor syntactic uptake rate predicted satisfaction ( $p > .45$ ), indicating that the subjective experience of AI interaction is shaped by production fluency rather than by receptive alignment.

## 6. DISCUSSION

The fluency, uptake, and survey results converge on one conclusion: AI and human dialogue create distinct but complementary conditions for L2 practice.

**Monologue-in-dialogue.** AI dialogue structurally approximates supported monologue. Learners produce sustained output with fewer, longer turns, slower responses, and less floor competition, while producing more fluent speech within each turn. The AI’s patient, non-competitive style shifts learners from reactive turn management to more deliberate, content-focused production. Human dialogue, by contrast, is rapid and interactive: frequent turns, fast responses, balanced floor time, and more within-turn hesitation. The interactional work of peer conversation (collaborative reference-building, repair negotiation, and affiliative floor manage-

Table 5: OLS regressions predicting post-interaction satisfaction and attitude change from behavioral features of the AI dialogue ( $n = 55$ ). Site is dummy-coded with Flemish 1 as the reference category. \* $p < .05$ , \*\* $p < .01$ .

Predictor	Satisfaction			Attitude change			Satisfaction (uptake)		
	$\beta$	$SE$	$p$	$\beta$	$SE$	$p$	$\beta$	$SE$	$p$
Intercept	4.01	0.54	<.001**	2.10	0.46	<.001**	3.34	0.51	<.001**
Pre-AI attitude	0.36	0.13	.006**				0.31	0.14	.030*
Within-turn pause ratio	-4.78	1.71	.008**	-5.91	2.07	.006**			
Floor ratio	-0.95	0.92	.307	-1.16	1.12	.306			
Lexical adoption rate							1.22	1.62	.454
Syntactic uptake rate							-0.39	1.08	.717
Site: Flemish 2	-0.53	0.26	.046*	-0.22	0.31	.482	-0.58	0.27	.038*
Site: Midwest	-0.44	0.35	.209	-0.18	0.42	.667	-0.98	0.32	.004**
Site: PNW	-0.34	0.26	.196	-0.04	0.31	.907	-0.60	0.28	.037*
$R^2$ (Adj. $R^2$ )		.43 (.36)			.24 (.16)			.33 (.25)	
$F$		6.03**			3.10*			3.98**	

ment) is largely absent from AI dialogue. This is not a failure of the AI to “be human” but a description of what the interaction actually is.

**What the AI scaffolds.** The AI’s verbose, syntactically regular speech, in this specific model-prompt-task configuration, creates favorable conditions for short-term structural alignment. Overall, AI dialogue exposes learners to more input and elicits more uptake in absolute terms. The deeper finding is that the AI’s syntactic priming advantage survives all controls for input volume, pointing to a qualitative difference: complete, well-formed constructions without the fragments and restarts of peer speech yield cleaner targets for short-term reuse. At the lexical level, the AI provides greater exposure but not more efficient uptake per unit of input ( $p = .062$  after controls). The AI’s distinctive pedagogical contribution is thus specifically structural. This is particularly notable given that meaning-focused tasks like spot-the-difference typically direct learner attention toward vocabulary rather than grammar [13]; the syntactic priming observed here operates implicitly, suggesting a potential affordance of AI interaction worth investigating in future work.

**What human partners scaffold.** Human dialogue exercises a different set of capacities. The rapid turn exchange, floor competition, and negotiation sequences characteristic of peer interaction create conditions for developing interactional competence, skills not exercised in AI dialogue. The information-gap task produces its richest interactional outcomes with a partner who *cannot* resolve gaps efficiently. Whether human dialogue also promotes more durable lexical learning through interactional grounding remains a question our data suggest but cannot confirm.

**Learner experience.** Attitudes toward AI improved significantly after the task, but satisfaction was predicted by within-turn production fluency, not by uptake. Learners who paused more within their turns reported lower satisfaction, suggesting that the subjective quality of AI interaction hinges on whether the learner can produce fluent speech in the space the AI provides, not on how much language they absorb from the AI’s input.

## 7. CONCLUSIONS

AI and human dialogue scaffold L2 production through complementary mechanisms: structural priming through regular

input in one case, interactional competence through coordinative demand in the other. The design challenge is not to make AI dialogue more like human dialogue, but to exploit what each mode does well.

This points to concrete principles for integrating voice-based AI into language curricula: For curricula, this suggests four design principles. First, AI tasks are well suited to focused morphosyntactic practice, where regular input can support priming of constructions such as dative prepositions and gender agreement. Second, AI practice should complement rather than replace peer speaking, since human dialogue exercises turn competition, clarification requests, repair, and collaborative floor management. Third, AI verbosity should be calibrated: shorter, less complete AI turns may reduce floor dominance and create more space for learner-initiated language. Finally, uptake metrics can support adaptive feedback on vocabulary and constructions adopted, missed, or produced incorrectly from ASR transcripts.

**Limitations.** Several limitations apply. Uptake measures cannot fully separate priming from parallel activation driven by shared task context; causal claims require manipulating interlocutor language. NLP pipelines trained on Standard German may degrade on disfluent L2 speech; word-level ASR errors were not corrected, although German-speaking authors verified speaker labels, task boundaries, and broad coherence. First-use counting captures breadth rather than depth of uptake, and the lexical first-source-wins rule creates a task-order confound ( $d = 1.71$ ), addressed with lower-powered position-matched comparisons. The study cannot determine whether short-term priming becomes durable learning; delayed post-tests are needed. Transcript-based analysis also excludes prosody, pronunciation, stress, intonation, and segmental accuracy. The comparison involved stimulus and modality asymmetries: human partners viewed images directly and interacted face-to-face or by Zoom, whereas the AI received English text descriptions and interacted through audio only. Some temporal effects may therefore partly reflect interface latency, VAD behavior, and visual coordination cues rather than interlocutor type alone. Finally, the AI condition reflects one system configuration (model, voice, prompt, VAD setting, and task interface) with university German learners, mostly Dutch- or English-dominant, performing one task type; future work should replicate across AI systems, L1 backgrounds, proficiency levels, target languages, and tasks.

## Acknowledgments

We would like to thank all of the professors, teachers and students who worked with great dedication to plan for, collect and process all of the data for these experiments.

## References

- [1] S. Bibauw, T. François, and P. Desmet. Discussing with a computer to practice a foreign language: Research synthesis and conceptual framework of dialogue-based call. *Computer Assisted Language Learning*, 32(8):827–877, 2019.
- [2] S. Bibauw, W. Van den Noortgate, T. François, and P. Desmet. Dialogue systems for language learning: A meta-analysis. *Language Learning & Technology*, 26(1):1–24, 2022.
- [3] C. Borchers, K. Yang, J. Lin, N. Rummel, K. R. Koedinger, and V. Alevan. Combining dialog acts and skill modeling: What chat interactions enhance learning rates during ai-supported peer tutoring? In B. Paassen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 117–130, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.
- [4] S. E. Brennan and H. H. Clark. Conceptual pacts and lexical choice in conversation. *Journal of experimental psychology: Learning, memory, and cognition*, 22(6):1482, 1996.
- [5] S. Choi and S. Oh. Developing L2 turn-taking with ChatGPT: A longitudinal conversation analytic study. *System*, 138:103959, 2026.
- [6] A. Costa, M. J. Pickering, and A. Sorace. Alignment in second language dialogue. *Language and Cognitive Processes*, 23(4):528–556, 2008.
- [7] B. R. Cowan and H. Branigan. Does voice anthropomorphism affect lexical alignment in speech-based human-computer dialogue? In *Proc. Interspeech 2015*, pages 155–159, 2015.
- [8] P. Dao, P. Trofimovich, and S. Kennedy. Structural alignment in l2 task-based interaction. *ITL - International Journal of Applied Linguistics*, 169(2):293–320, 2018.
- [9] D. Demszky, J. Liu, Z. Mancenido, J. Cohen, H. Hill, D. Jurafsky, and T. B. Hashimoto. Measuring conversational uptake: A case study on student-teacher interactions. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*, pages 1638–1653, 2021.
- [10] J. Du and B. K. Daniel. Transforming language education: A systematic review of ai-powered chatbots for english as a foreign language speaking practice. *Computers and Education: Artificial Intelligence*, 6:100230, 2024.
- [11] J. E. Fischer, S. Reeves, M. Porcheron, and S. Sharples. Progressivity for voice interface design. In *Proceedings of the 1st International Conference on Conversational User Interfaces*, CUI '19. Association for Computing Machinery, 2019.
- [12] L. Fryer and R. Carpenter. Bots as language learning tools. 2006.
- [13] M. d. P. García Mayo and A. Azkarai. Efl task-based interaction: Does task modality impact on language-related episodes? In *Peer interaction and second language learning: Pedagogical potential and research agenda*, pages 241–266. John Benjamins Publishing Company, 2016.
- [14] R. J. Hartsuiker, S. Bernolet, S. Schoonbaert, S. Speybroeck, and D. Vanderelst. Syntactic priming persists while the lexical boost decays: Evidence from written and spoken dialogue. *Journal of Memory and Language*, 58(2):214–238, 2008.
- [15] H. Hoang and F. Boers. Re-telling a story in a second language: How well do adult learners mine an input text for multiword expressions? *Studies in Second Language Learning and Teaching*, 6(3):513–535, 2016.
- [16] F. Huang and B. Zou. English speaking with artificial intelligence (ai): The roles of enjoyment, willingness to communicate with ai, and innovativeness. *Comput. Hum. Behav.*, 159(C), Oct. 2024.
- [17] W. Huang, K. F. Hew, and L. K. Fryer. Chatbots for language learning—are they really useful? a systematic review of chatbot-supported language learning. *Journal of Computer Assisted Learning*, 38(1):237–257, 2022.
- [18] Y. Kim and M. Michel. Linguistic alignment in second language acquisition: A methodological review. *System*, 115:103007, 2023.
- [19] W. Y. W. Lai and J. S. Lee. A systematic review of conversational ai tools in elt: Publication trends, tools, research methods, learning outcomes, and antecedents. *Computers and Education: Artificial Intelligence*, 7:100291, 2024.
- [20] S. C. Levinson and F. Torreira. Timing in turn-taking and its implications for processing models of language. *Frontiers in Psychology*, 6:731, 2015.
- [21] E. Luger and A. Sellen. Like having a really bad PA: The gulf between user expectation and experience of conversational agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 5286–5297, San Jose, CA, USA, 2016. Association for Computing Machinery.
- [22] E. Mayor, L. M. Bietti, and A. Bangerter. Can large language models simulate spoken human conversations? *Cognitive Science*, 49(9):e70106, 2025.
- [23] K. McDonough. Interaction and syntactic priming: English l2 speakers' production of dative constructions. *Studies in Second Language Acquisition*, 28(2):179–207, 2006.

- [24] M. Michel, C. Appel, and S. Cipitria. Lexical and syntactic alignment during english-spanish teletandem meetings. *System*, 111:102930, 2022.
- [25] P. Peltonen. Temporal fluency and problem-solving in interaction: An exploratory study of fluency resources in L2 dialogue. *System*, 70:1–13, 2017.
- [26] M. Perlmann-Balme. Goethe-zertifikat A1: Start Deutsch 1 — Wortliste. Technical report, Goethe-Institut, 2012.
- [27] T. Pica, H.-S. Kang, and S. Sauro. Information gap tasks: Their multiple roles and contributions to interaction research methodology. *Studies in Second Language Acquisition*, 28(2):301–338, 2006.
- [28] M. J. Pickering and S. Garrod. The interactive-alignment model: Developments and refinements. *Behavioral and Brain Sciences*, 27(2):212–225, 2004.
- [29] M. Porcheron, J. E. Fischer, S. Reeves, and S. Sharples. Voice interfaces in everyday life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 1–12, Montreal, QC, Canada, 2018. Association for Computing Machinery.
- [30] O. Pütz and E. Esposito. Performance without understanding: How chatgpt relies on humans to repair conversational trouble. *Discourse & Communication*, 18(6):859–868, 2024.
- [31] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, and C. D. Manning. Stanza: A python natural language processing toolkit for many human languages. In A. Celikyilmaz and T.-H. Wen, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online, July 2020. Association for Computational Linguistics.
- [32] D. Reitter and J. D. Moore. Alignment and task success in spoken dialogue. *Journal of Memory and Language*, 76:29–46, 2014.
- [33] H. Sacks, E. A. Schegloff, and G. Jefferson. A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4):696–735, 1974.
- [34] S. Seabold and J. Perktold. Statsmodels: Econometric and statistical modeling with python. *SciPy 2010*, 2010.
- [35] L. Spillner and N. Wenig. Talk to me on my level – linguistic alignment for chatbots. In *Proceedings of the 23rd International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '21. Association for Computing Machinery, 2021.
- [36] F. Stinar, H. Lee, C. Belitz, N. Nasiar, S. Fancsali, S. Ritter, H. Almoubayyed, R. Baker, J. Ocumpaugh, and N. Bosch. Fairness of bayesian knowledge tracing for math learners of different reading ability. In C. Mills, G. Alexandron, D. Taibi, G. L. Bosco, and L. Paquette, editors, *Proceedings of the 18th International Conference on Educational Data Mining*, pages 170–181, Palermo, Italy, July 2025. International Educational Data Mining Society.
- [37] T. Stivers, N. J. Enfield, P. Brown, C. Englert, M. Hayashi, T. Heinemann, G. Hoymann, F. Rossano, J. P. de Ruiter, K.-E. Yoon, and S. C. Levinson. Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26):10587–10592, 2009.
- [38] P. Tavakoli. Fluency in monologic and dialogic task performance: Challenges in defining and measuring L2 fluency. *International Review of Applied Linguistics in Language Teaching*, 54(2):133–150, 2016.
- [39] V. Venkatesha, A. Nath, I. Khebour, A. Chelle, M. Bradford, J. Tu, J. Pustejovsky, N. Blanchard, and N. Krishnaswamy. Propositional extraction from natural speech in small group collaborative tasks. In B. Paassen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 169–180, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.
- [40] Y. Wang, Y. Wang, K. Crace, and Y. Zhang. Understanding attitudes and trust of generative ai chatbots for social anxiety support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, page 1–21. ACM, Apr. 2025.
- [41] J. Zhang, C. Borchers, V. Alevan, and R. S. Baker. Using large language models to detect self-regulated learning in think-aloud protocols. In B. Paassen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 157–168, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.