

# Quit Playing Language Games with My Heart; or, Conversational AI and Knowledge Sharing

LARRY S. McGRATH

*This paper shares guidance for designing conversational AI based on findings from linguistic and social analysis of core shifts introduced by large language models. When Duolingo and Babbel added AI features to their language learning apps, the games that users play on the apps was transformed, giving us a window into the structure of AI-driven user interactions in general. Dynamic, turn-taking dialogues resembling natural conversation now appear alongside multiple-choice and drag-and-drop vocabulary games. The transformation is not just technological: It also lays bare the shifting language games that we have come to play more broadly in our interactions with LLMs such as Chat-GPT, Claude, and Gemini. People engage with them in back-and-forth conversations in which different kinds of speech acts are at play. By contrast, traditional conversational agents based on decision-tree learning rely on discrete, finite, and close-ended commands. The paper argues that the future of AI conversation is one in which rule-based interactions co-exist with intention-based interactions powered by generative AI. Design elements on Duolingo, in particular, draw on diverse styles of communication and offer guidance to ensure that conversational AI platforms in general remain open to new forms of dialogue that enrich both artificial and natural conversation alike.*

“Here the term ‘language-game’ is meant to bring into prominence the fact that the speaking of language is part of an activity, or of a form of life.”

– Ludwig Wittgenstein, *Philosophical Investigations* §23

## 1. Introduction

Speaking a language is like playing a game. For Wittgenstein, learning a language – that is, grasping how to describe, question, command, and explain – involves a process of mastering the rules of the games that shape human interactions. In this paper, I take Wittgenstein’s words literally and examine language games in the learning apps, Duolingo and Babbel, which transformed in 2023-2024 from rule-based to AI-powered conversation platforms (Freeman 2023). The transformation is not just technological; it also lays bare shifting modes of interaction with chatbots. Unlike decision-tree chatbots, intention-based AI interfaces support communication through back-and forth dialogues in a dynamic manner not wholly unlike natural conversation. I suggest that studying language games in everyday conversation can enrich the design of language learning apps and ensure, more broadly, that

conversational AI platforms remain open to the diverse forms of interaction that already animate dialogue among humans.

Language games serve to demonstrate that words' meaning is better understood as part of a conversation among humans rather than as a representation of the world. For Wittgenstein, speakers of a language can understand each other because they partake in a shared activity. What speakers achieve in mutual understanding, however, comes at the risk of incomprehension, failure, and rejection. Much can go wrong in conversation. No facts – whether in the world or in our minds – guarantee the successful transmission of meaning. Such is also the case, as I show in this paper, when playing games with language learning apps. As Large Language Models (LLMs) grow in size and complexity, expanding conversational AI platforms' emulation of natural dialogue, so too do they replicate the fragility of their fallible creators. An important characteristic of language learning apps' sense of realism, I'm suggesting, is to be found in their potential for confusion. Duolingo and Babbel can bewilder. AI integrations that make the apps approximate human-like conversation also make them succumb to what Stanley Cavell describes as “the deceptions and temptations and dissatisfactions of the ordinary human effort to make oneself intelligible (to others and to oneself)” (Cavell 2004). As thinkers in Wittgenstein's wake have shown, language games can always run aground and land in misunderstanding. Our techniques for navigating confusion in quotidian conversation offer a guide to steering artificial conversation back on course.

Duolingo and Babbel exemplify the transformative impact of AI-powered conversation platforms. The apps use Natural Language Processing (NLP) to create personalized and adaptive learning experiences. Duolingo uses AI to tailor lessons to individual users' performance, providing each learner with appropriate levels of challenge and support (Duolingo 2023). Similarly, Babbel employs AI to recognize users' speech and provide real-time feedback on pronunciation and grammar (Babbel 2023). AI's integration not only enhances user engagement through realistic dialogue and instant feedback; a window also opens onto broader interaction patterns between humans and conversational user interfaces. Indeed, conversational AI is a partner in dialogue.

Conversational AI platforms are chatbots that people interact with by typing natural language, often over the course of back-and-forth interactions. There are two kinds of platforms. First, universal AI platforms such as Chat-GPT, Claude, and Gemini have a general scope of knowledge; their intended use is as standalone products. Second, product-specific AI chatbots (including Duolingo and Babbel) have a scope of knowledge limited to an organization with a specific service; their

intended use cases are as embedded features designed to address the organization's needs. I'm suggesting in this paper that language learning apps' AI features open a window onto the deeper structure of conversational interaction with the former platforms.

My own conversation is with scholars in Science and Technology Studies as well as Human-Computer Interaction, especially those who have moved beyond debates about AI platforms' epistemic capacities. An outdated current of thought remains captivated by AI's ability to mimic intelligence, reasoning, creativity, or consciousness akin to that of humans. On the one hand, technology enthusiasts contend that conversational AI realizes a theory of mind (Kosinski 2023; Summers-Stay 2023). Infamously, Google fired Blake Lemoine, an engineer, after he claimed that the Language Model for Dialogue Applications (LaMDA) had reached a level of human-like consciousness (Wakabayashi 2022). On the other hand, Noam Chomsky argues that nothing of the sort is possible. "The human mind is not, like Chat-GPT and its ilk, a lumbering statistical engine for pattern matching, gorging on hundreds of terabytes of data and extrapolating the most likely conversational response or most probable answer to a scientific question" (Chomsky 2023). Both sides of the debate, however, frame epistemic capacities as *possessions* of humans or AI systems. I want to suggest the ethnographers are well positioned, instead, to trace the behaviors and patterns that shape interactions *between* humans and AI platforms. Autonomy is bound up with automaticity. Neither are located on one side of the human-machine divide. A fresh scholarly focus on relations (rather than identities) guides recent inquiry into the natural conversation styles that precede and exceed conversational agents (Li 2023; Packer (2023). My article builds on that literature. As conversations with AI platforms approximate conversations among humans, I show how we can improve the former by drawing lessons from the latter.

In the pages that follow, I offer an armchair anthropology primarily of Duolingo (and secondarily of Babbel) guided by the philosophy of language. Wittengstein's insight that meaning is a form of action in a community orients my study of the forms of interaction between human and computer made possible by AI-powered language games. In the second section, I present AI features in Duolingo and Babbel and explain their leap beyond rule-based language games. In the third section, I show how two kinds of language games – traditional rule-based games and dynamic conversation games – engender distinct varieties of learning. Learners commit errors of knowledge when playing traditional games. With dynamic conversations, however, errors of acknowledgement take place when conversations follow unintended directions. In the fourth section, I elaborate on the theme of acknowledgment and

argue that it's pivotal not just in language learning apps but also in emergent interaction patterns with universal AI platforms. As technological progress brings those platforms' models closer to natural conversation, they also bear the risks of uncertainty and skepticism that haunt our all-too-human efforts to interpret others' intentions in everyday language games. In the final section, I show how language learning apps come to terms with these risks by combining rule-based interactions with open-ended conversations. With Duolingo in particular, the combination offers an elegant solution to the challenges that conversational AI platforms confront as they emulate natural conversation.

## 2. From Rigid Rules to Dynamic Conversations

Babbel was founded in 2007 with the mission of “creating mutual understanding through language” (Babbel 2024). Duolingo followed in 2011 the method as its centerpiece: “learn by doing” (Freeman 2023). Both mobile apps still feature language games in 2024 whose purpose is to learn languages and facilitate human conversation via speaking, writing, reading, and listening. Duolingo employs brief 2-5 minute exercises such as matching paired words, translating sentences, speaking sentences, and multiple-choice questions. Babbel employs longer scenarios in the form of dialogues, fill-in-the-blank exercises, and pronunciation practice. What makes the exercises *games* is both the progressive structure of the exercises, which advance by levels, as well as engaging features such as points, badges, rewards, and challenges (Saleem 2022). Although Duolingo explicitly gamifies exercises with playful designs, a cartoon bird mascot, and a streak system that rewards daily usage, games are central to both apps in the form of progress tracking and certificates for finishing courses. Like a board game, users can take out their mobile device, play briefly with Babbel or Duolingo, and return the apps to their pocket.

The apps initially relied on decision-tree algorithms to personalize language games for users' language skills. Decision trees are used in machine learning to make predictions based on a finite data set with predefined categories. The apps used Simple Adaptive Learning to match the difficulty of a game with the user's ability. Like a tree, the model begins with a root node (the user's initial language assessment after signing up and subsequent exercise levels while using the app); decision nodes split the dataset on the basis of defined conditions (such as the user's performance on an exercise and the time passed since the prior exercise); leaf nodes generate predictions (such as correcting the user's response and presenting feedback). Predetermined rules guide the decision nodes of the tree. They include, among

others, progression rules (if the user completes a lesson with at least 80% accuracy then the next level unlocks), exercise rules (if the user struggles with writing then more typing games appear), and error rules (if the user makes the same error three times then detailed explanations follow). Although both apps use supervised decision trees to personalize the games presented to users, Duolingo came to deploy more sophisticated machine learning models. By 2023, a recurrent neural network internally called Birdbrain supported game selection on the basis of a user's entire history on the app (Freeman 2023). A long short-term memory model translates users' performance into 40 vectors. Completing a game updates the vectors. Individual games are tailored to each user's learning level. Yet, the internal elements of any game remain fixed. A finite number of routes through the decision tree circumscribe the bounds of possible interactions with the traditional language games on Duolingo and Babbel.

Rule-based algorithms underlie traditional, rule-based language games on the apps. The user is introduced to the meaning of new words by using the words – that is, by playing games. The app takes turns by displaying a prompt; the user responds by answering. The user learns by doing in the sense that rules are embedded in the games; he or she does not grasp a rule apart from the words' use. “For a large class of cases,” Wittgenstein wrote, “in which we employ the word ‘meaning’ it can be defined thus: the meaning of a word is its use in the language” (Wittgenstein 1953, §43). When a user performs well, and correctly applies a rule, the algorithm advances to more advanced topics. When the user fails to apply the rule, the algorithm revisits weaker areas. The rule-based games reflect the contractual terms of language. As John McDowell puts it, “to learn the meaning of a word is to acquire an understanding that obliges us subsequently – if we have occasion to deploy the concept in question – to judge and speak in certain determinate ways, on pain of failure to obey the dictates of the meaning we have grasped” (McDowell 1984, 325). The linguistic rules used to evaluate users' turns depend on specific vocabularies and grammatical structures, the categorical values that function as decision nodes in the apps' decision trees. The limited range of input prevents overfitting (when a machine learning model predicts accurately from training data but not from new data). To-and-fro, the app and user take turns playing language games.

Turns are finite, discrete, and close-ended. In the apps' traditional language games, a finite set of clickable individual words appear. For pronunciation and free-form response games, the apps recognize only predefined answers. Each of the user's turns constitutes an individual unit in the game. The units are close-ended in the sense that they follow a binary logic of validation. Answers are correct or

incorrect. Each individual turn in the game either applies or violates grammatical rules and, as a result, triggers the algorithmic rules in the decision tree to issue the next turn.

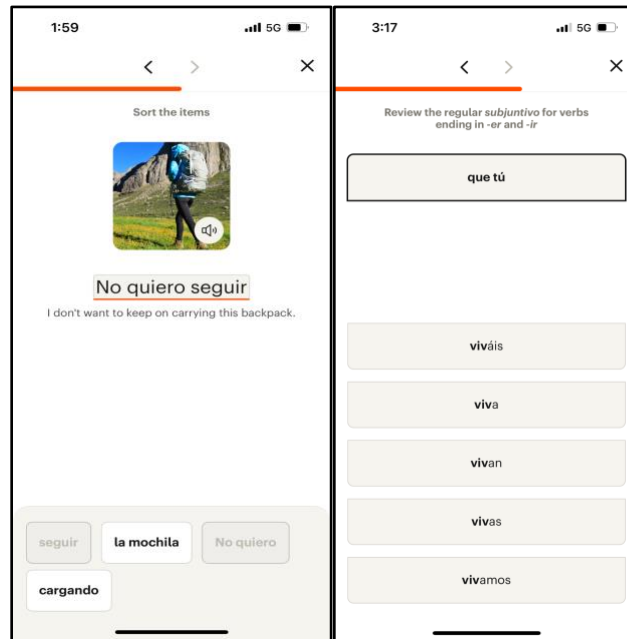


Figure 1 a–b. Babel – Rule-based language games

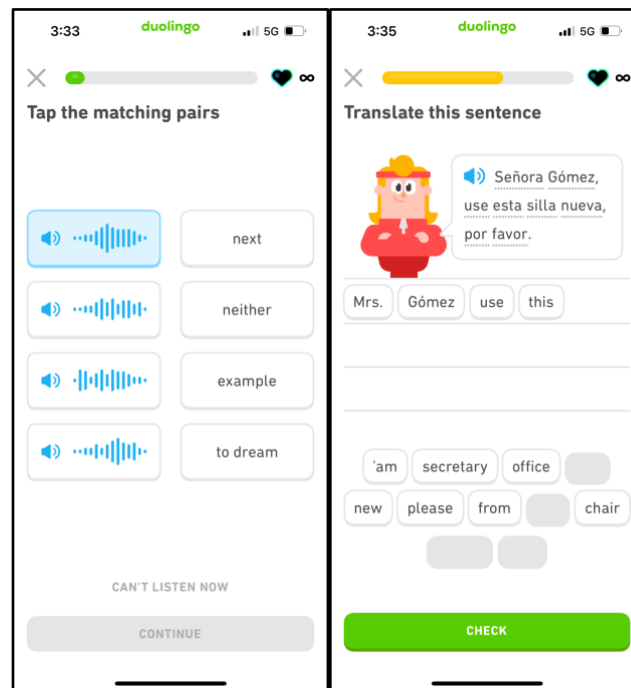


Figure 2 a–b. Duolingo – Rule-based language games

Yet, the traditional language games in Duolingo and Babbel lack the continuous character of natural conversation that unfolds between speakers of a shared language. That is why the apps have demonstrated mixed results in studies. When it comes to effective rule learning, Duolingo and Babbel help reinforce knowledge of vocabulary and grammar. In one study, 54 college students improved their Spanish scores by .7 of 5 levels on the standardized exam of the American Council on the Teaching of Foreign Languages after using Duolingo for 15 minutes per day over 12 hours (Jiang 2021). A review found that app-based vocabulary learning can be as effective as traditional instruction methods for vocabulary. However, students have had to go elsewhere to learn conversation skills (Tommerdahl 2022). Fluid turn-based conversation, as students would enjoy in classrooms, has been a shortcoming of traditional app-based language games.

Babbel and Duolingo added AI-powered features in 2023 to make their language games feel more like natural interactions with humans. Babbel added AI-enhanced speech recognition to “Everyday Conversations” for French, German, Italian, and Spanish. Duolingo’s “Roleplay” surpasses Babbel’s AI integration in its sense of realism and fluidity (currently available only for French and Spanish). Whereas Babbel’s users read aloud predetermined lines, Duolingo’s users can express anything that they’d like in open-ended interactions. “Roleplay” comes far closer to the back-and-forth dynamics of natural conversation.




Figure 3a–b. Duolingo – “Roleplay” AI feature

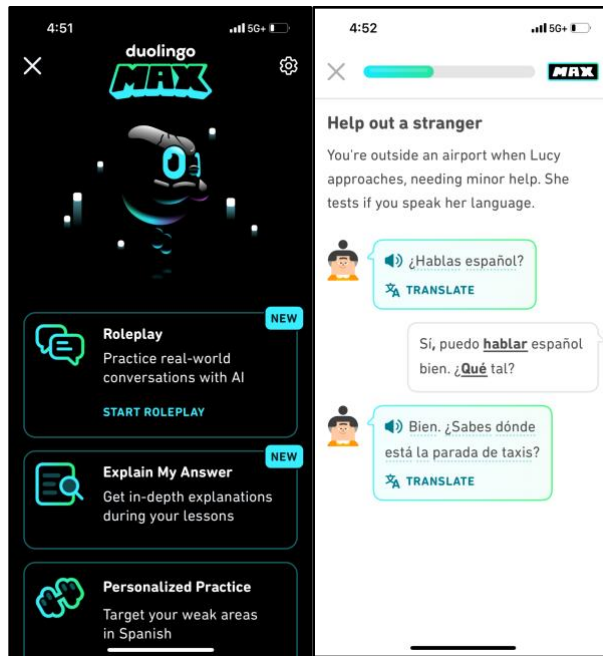


Figure 4 a–b. Babbel – “Everyday Conversations” AI feature

Continuous conversation comes to the fore in Duolingo’s “Roleplay.” Users do not carry out finite, discrete, and close-ended exercises. Rather, dialogue is open-ended. The uncertainty of the chatbot’s response hangs on every utterance. Users can engage literally or add commentary; they can elaborate by injecting color, making revisions, or taking exceptions. Points are rewarded for turns with more words. Akin to a fluid conversation between humans, words are not set in stone but confront the pushback of a fellow conversant. Although Duolingo is unmistakably non-human, “Roleplay” depends on an interaction pattern that stands out from those of the app’s traditional language games. Back-and-forth dialogue endures through time as each turn between chatbot and user spills into the unforeseen utterances of the next turn.

Neural networks facilitate the continuous character of dialogue in “Roleplay” by registering the various ways that users articulate the same utterance. Instead of parsing individual entries word by word, the networks deploy hidden layers between a user’s input and Duolingo’s output. Each node in the layers applies a set of attention mechanisms to the input and passes the result through an activation function before sending the output to the next layer. The attention mechanisms function as weights and biases; they assess all prior states of text according to learned measures of context and relevance. “Roleplay” integrates a Generative Pre-Trained Transformer (GPT), which applies weighted mean reductions to interpret input in the light of the collective input of users’ interactions with the app’s language games



(Duolingo 2023b). (Babbel has not disclosed which model facilitates “Everyday Conversations.”) The result in Duolingo is a conversation whose continuity extends across alternating turns between chatbot and user.

Users of “Roleplay” play a turn-taking language game. They learn not just the grammatical rules of meaning but also partake in the social rules that structure a form of life shared with others. Conversation is the conduit of community. The conclusion of the chatbot’s turn solicits the user’s to begin. It’s up to him or her to form a complete thought and bring it to a close in a manner that completes the turn and invites the chatbot’s response. Conversations with “Roleplay” last seven to 10 turns. In their study of turn-taking, Harvey Sacks, Emanuel Schegloff, and Gail Jefferson wrote, “For socially organized activities, the presence of ‘turns’ suggests an economy, with turns for something being valued and with means for allocating them, which affect their relative distribution” (Sacks 1974, 696). Each turn is not a wholly isolated unit (although Duolingo does deploy feedback in response to incorrect utterances). The chatbot and user articulate turns in an organized sequence, what the pioneers of conversation analysis call an “economy” whose distribution of social roles has at once “general abstractness and local particularization potential” (Sacks 1974, 700). Users’ words serve both to respond to the chatbot’s individual utterances and to advance the next turn in a structured conversation.

The transition from language learning apps’ rule-based decision trees to dynamic conversation supported by neural networks brought about a shift in interaction patterns from close-ended and finite actions to open-ended and continuous dialogue. Much like a natural conversation, the prior turn in “Roleplay” weaves with the subsequent turn. The chatbot’s warp crosses the user’s weft such that neither are entirely separate from their dialogical tapestry.

### **3. Learning a Language Via Two Varieties of Knowledge**

Learning a language depends on making mistakes. Indeed, opening one’s mouth, putting pen to paper, typing on a keyboard – all entail risks when sending new words into the world. The mistakes of schoolchildren who baffle their teacher as they stumble about trying to apply seemingly obvious rules appear again and again in Wittgenstein’s writings. “Let us now examine the following kind of language-game,” he writes, “when A gives an order B has to write down a series of signs according to a certain formulation rule... At first perhaps we guide his hand in writing out the series 0 to 9; but then the *possibility of getting him to understand* will depend on his going on to write down independently” (Wittgenstein, 1953, §143). On his own, B writes,

0, 1, 2, 3, 4... and suddenly, the series goes awry. B writes ... 5, 7, 6, 9. The scenario illustrates how rules are communal phenomena that bond us to their future application. B's mistakes do not arise because he fails to share with A a mental picture – perhaps a formula – of ordinal numbers. The student's understanding is, ultimately, to be found in applying the rule over time in a shared community. Wittgenstein concludes, “And hence also ‘obeying a rule’ is a practice. And to think one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule ‘privately’: otherwise thinking one was obeying a rule would be the same thing as obeying it” (Wittgenstein, 1953, §202). To speak a language is a public act. And coming to understand the rules of a language involves a long and frustrating process of unexpected errors in the face of errors. That's why errors are key elements of Duolingo and Babbel.

Error messages appear in the apps when users make mistakes. For the traditional language games, decision trees register errors across various aspects of language – e.g., errors of grammar, vocabulary, pronunciation. Error messages take a few forms: incorrect answer messages display the correct answer and explain words' meaning; grammar messages display the correct tense, word order, or article; typos qualify as correct responses accompanied by clarifications of words' spelling; missing diacritics (in the case of Latin languages) also qualify as correct responses with messages that display correct accents.

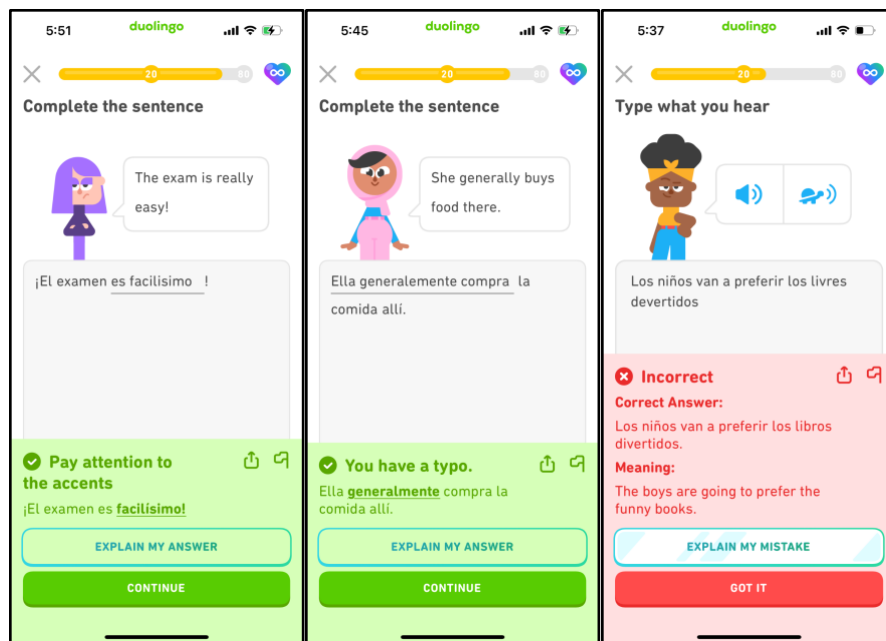


Figure 5 a–c. Error messages in Duolingo

The error messages respond to users' *errors of knowledge*. A user fails to correctly apply a rule of language. The messages are designed to guide users back to the rule's correct application. We're reminded of Wittgenstein's students, whose misapplications of a rule evince that they have yet to understand it. The error messages' goal is to restate the rule, encourage its correct applications in future instances, and thereby grow users' knowledge of a language.

Errors of knowledge also afflict LLMs. For its public debut, Google's Bard (now known as Gemini) incorrectly stated that the James Webb Space Telescope took the first image of a planet outside our solar system. In fact, the European Southern Observatory's Very Large Telescope already did so in 2004. Such "hallucinations" are false predictions that take the form of factually incorrect output (Ji 2023). Google's LaMDA model trained Bard on the basis of more than 1.56 trillion words and deployed 137 billion parameters to predict erroneously that the James Webb Space Telescope took the first image of an exoplanet.

Interactions with LLMs can also bring about varieties of error other than incorrect facts. *Errors of acknowledgment* occur when the course of conversation goes astray, takes unforeseeable turns, and results in misunderstandings. Words fail to act on each other with the appropriate force. And the resultant turns in a back-and-forth dialogue prove dissatisfying. Divergent interactions or shifting contexts in Duolingo's "Roleplay" are examples of acknowledgement errors that can arise artificial and natural conversation alike, even when conversants might utter factually correct sentences.

Take as an example when conversation runs aground. Duolingo's "Roleplay" organizes dialogues according to topics, such as "Help out a stranger" and "Mingle at a party." The topics serve as points of departure but – like a meandering dialogue – they need not be the final destination. Users can stray from the topic on account of confusion or an effort to steer the conversation toward their own ends. Because the AI-powered conversation unfolds continuously through turns, each turn has the potential for either conversant to lead the next down errant paths, even when chatbot and user contribute empirically correct content. Below, a user of French "Roleplay" responds to a conversation about parks in Paris by instead suggesting that they go to a park in Montréal. The conversation takes an errant turn to discuss a North American country – not a European country. Chatbot and user are not on the same wavelength. Although the sentences are meaningful, the conversation exemplifies an error of acknowledgment due to misaligned intentions.

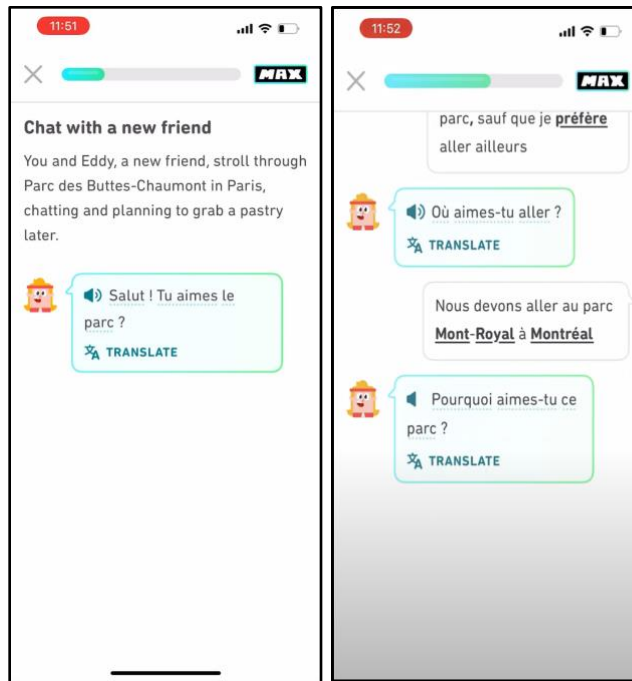


Figure 6 a–b. “Roleplay” in Duolingo

In the second example, new contexts lead to ambiguity and misunderstanding when words’ meaning shifts. Learners of a language grasp new meanings of words by using them in different situations. Creativity and flexibility are key. Consider a user of Spanish “Roleplay” who meets a friend in Barcelona’s Parque Güell. The user mentions that he gets lost in the designs (“Siempre me pierdo en sus diseños”). “Pierdo” can mean getting lost physically or metaphorically, as if to be captivated by the designs. For a Spanish language learner familiar with just the word’s literal meaning, the idea that his friend might be lost could feel unsettling; it would certainly fail to convey his sense of amazement stirred by Gaudi’s architecture. The error of acknowledgment reflects the same words being used differently by chatbot and user.

Acknowledgment is a peculiar species of knowledge. I acknowledge another when my utterances convey not just facts but also an appeal to respond. The word stems from 15th- century Middle English, *aknow*, meaning an admission of one’s knowledge. I confess my words, as it were, and make claims on another. As Cavell put it, “acknowledgment is more than knowledge, it includes an invitation to action” (Cavell 1969, 259). In the least, that action enjoins another to respond in kind and continue the conversation. At the most, my words’ action makes an ethical appeal. When I feel pain in front of a friend, for instance, I might wince, groan, and grab the hurt part of my body. My speech does not convey a desire that my pain be known.

(This is the sense in which Wittgenstein remarked, “It can’t be said of me at all (except perhaps as a joke) that I *know* I am in pain” (Wittgenstein, 1953, §246). I seek my friend’s acknowledgement: that he responds attentively, asks how I feel, and gestures to help.

The passage from rule-based to AI-powered language games, viewed from the vantage of linguistics, hinges on the different speech acts at play. Games of knowledge assess users’ constative utterances: words whose description of a state of affairs is true or false. These are generally statements (“He attends the wedding”) with meaning and lexical qualities that are assessed according to their truth. With AI-powered conversations, Duolingo empowers users also to make performative utterances, which, as J.L. Austin showed, do not describe but bring about a state of affairs (Austin 1962). These are statements (to use Austin’s favorite example, “I take thee to be my lawfully wedded wife”) whose force implicates another. They enjoin the other’s acknowledgment in order to create a new truth. In the case of “Roleplay,” performative utterances impel the chatbot to continue the conversation in an intended direction.

As AI-powered conversation approximates natural conversation, conversations between users and chatbots fall victim to the infelicities that already haunt our struggles to acknowledge each other. In natural conversation, my pronouncement of marriage might fall apart when spoken in the wrong context (perhaps I lack the authority of a rabbi, priest, or civil servant) or when the other leaves my words unacknowledged (and never says “I do”). Even when the utterances spoken might be empirically valid, unforeseen turns and shifting contexts can equally lead conversation with a chatbot to errors of acknowledgement.

Errors of acknowledgment could also be the result of willful manipulation. Take as an example Air Canada’s customer service chatbot, which invented a bereavement refund policy thanks to a conniving user’s back-and-forth conversations with the bot (Matsakis 2023). He claimed that he was entitled to a free ticket because he was traveling to attend a grandparent’s funeral. After Air Canada rescinded the reimbursement, a Canadian tribunal ruled that the company had to enforce the policy. The customer kept a screenshot and demanded the refund that had been promised. We find humans interacting with AI platforms as if they were humans – testing or coaxing chatbots. The history of computing is rife with malfeasance. Humans hack systems to steal information or falsify documents using text and image processors. But with AI platforms, we might imagine that users no longer break into a bank; a manipulative user instead persuades the guards to unlock the safe. In this

case of Air Canada, the bereavement policy was not lying in wait. It was the product of perlocutionary utterances between chatbot and user.

The errors of knowledge that attract much attention in popular and academic accounts of AI-powered conversation should not occlude the errors of acknowledgement that arise not simply due to mistakes of engineering. Acknowledgment remains a delicate struggle in everyday interactions. Neither human nor machine are entirely to blame. If solutions to such errors of acknowledgement have proved elusive, whether by modifications of interface design or back-end engineering, that is in large part because the infelicities of human interaction magnify in artificial interaction. Intentions can misalign. Either the chatbot misinterprets the user's intentions or the user confounds those of the chatbot. Perhaps, still, a mischievous user flippantly beguiles an AI-powered conversation. The potential for acknowledgement to go awry knows no bounds.

#### **4. Language Games at The Heart of Our Intentions**

Acknowledging another's intention is rarely instantaneous. There is a back-and-forth interaction by which we express our thoughts and solicit the other's response. To be sure, an AI-powered chatbot does not acknowledge a user's intentions *tout court*. A machine cannot *re-cognize* a user's words, tone, or facial expression with the emotional immediacy that humans of a common culture perceive in each other. AI models make predictions about users' intentions on the basis of learned associations found in training data. Duolingo's "Roleplay" nonetheless registers the linguistic signs of users' intentions thanks to Natural Language Understanding (NLU). Deployed by the transformer model, NLU concatenates patterns of intentional statements by tokenizing inputs across training data and embedding the tokens in vectors representing their relationships to other words. Context management finalizes the prediction. NLU tracks each turn in a dialogue to maintain conversations' coherence over time. Although AI-powered conversation relies on prediction, it cannot be said that predicting the intentions in users' natural language input is a small feat.

The intentional interface in Duolingo's "Roleplay" marks a shift from the command-based interfaces of modern computing (Norman 2023). Command-based interaction began in 1963 with the Teletype Model 33 ASR teleprinter, which registered users' lines of text via keyboard and punch tape. Originally designed for the US Navy, the teleprinter was soon used to send and receive messages in company offices. Command-based interactions became digital thanks to Graphic

User Interfaces (GUIs). Alto by Xerox PARC enabled users in 1973 to issue commands by clicking visual elements. Apple's Lisa (1983) and Macintosh (1984) inaugurated a world in which users click a mouse to engage windows, icons, and menus. AI changed everything. Instead of executing tasks via discrete commands, LLMs facilitate interactions by which users input what they want in the form of natural language. Duolingo's "Roleplay" deploys the technology to weigh mean reductions and interpret individual users' input against the collective input of user interactions. As a result, the app's underlying GPT predicts users' intentions.

Turn-taking with an AI-powered intentional interface such as Duolingo unfolds through time. Much the same, natural conversations advance through time when conversants speak intentionally. That is, we share more than bare facts with each other. Hardly innocent, the nature of intentionality is the subject of longstanding debates in the philosophy of language. If we think of users' back-and-forth conversations with chatbots as anything like natural conversations – rife with the hazards of knowledge and acknowledgment – then everything hinges on how we think of intentions. That is especially the case if natural conversation is to yield lessons that can improve the design of intention-based interfaces, as I have suggested. After all, intentionality establishes the conditions of a conversation's success or failure.

Consider three approaches to intentionality. For Aristotle, intention involves purposeful reasoning. All animals experience desire (appetency) but humans stand apart for the capacity to rationally deliberate (*βούληση*) about the means to fulfill our desires (Aristotle 1999). With *prohairesis* (*προαίρεσις*) we act intentionally and make choices that lead to a desired end. This line of thinking established in Western thought the principle that intention consists of rational action.

A second approach considers intention to be a mental state. When psychology took shape in the late nineteenth century as an experimental and clinical science, Franz Brentano contended that intentionality is the "mark of the mental" (Brentano 1874). Beliefs and desires direct the mind toward the world in a way that objects do not; they lack the intentional relation to believe or desire a state of affairs. Advancements in brain science in the twentieth century drove John Searle to suggest that intentionality is the mind's causal power to direct actions (Searle 1980). Although the second approach shares much with the folk concept that intentions motivate decision-making, it's difficult to apply to user interactions with conversational AI because computers lack the ability to access the thoughts and desires anterior to the words that users type in an interface. All that is available for natural language processing is the user's input.

A third approach is that intention depends on the order of reasoning. Elizabeth Anscombe suggested that intentions are not mental states that precede our actions (Anscombe 1957). We act intentionally when we are prepared to answer the question, “why?” Anscombe took as an example a man moving his arm up and down while holding a pump. We might ask, “why is he moving his arm?” We arrive at his intentions by explaining that the man acts *because* (and not just that) he is operating the pump, which he’s doing because (and not just that) he’s pumping water to the house. The order of the arm’s actions up and down really exists in the world; it is not the sum of merely physical events nor anterior thoughts or desires. *This* order of actions, connected by giving and taking reasons, constitutes the reasoning involved in intention.

So too does the order of reasons between chatbot and user make it possible to achieve a mutual acknowledgment of each other’s intentions. That the chatbot’s output follows my input, and in turn, that I extend the interaction by building subsequent input on prior output, makes a successful conversation possible. For H.P. Grice, we implicate each other. Conversations consist of more than “a succession of disconnected remarks...They are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, at least a mutually accepted direction” (Grice 1975, 45). The order of reasons, as Anscombe considered them, can be understood in Grice’s remark to take the shape of a “mutually accepted direction.” Errors of acknowledgment arise when intentions misalign; conversants express what might be otherwise factual claims without the shared connections expected in an intentional exchange.

Errors of acknowledgment are not a bug but a feature of everyday conversation. When conversants fail to connect – when we fall short of cooperative efforts and our words crumble into disconnected remarks – the urge ensues to flee or fasten the conversation. On the one hand, I might seek an escape from the frustration of elusive intentions (perhaps yours appear opaque or my own prove difficult to convey). On the other hand, I might seek iron-clad knowledge of the other’s intentions; to peer behind my interlocutor’s words, as it were, and grasp the underlying mental picture. Yet, the wish for a perfect communication, a complete transparency between minds, is a fantasy; it masks the reality that language makes our community possible as much as it divides us. Deprived of acknowledgment in a shared space, we seek knowledge as consolation in the recesses of another’s mind. What we seek are linguistic rules: answers to our doubts that are independent, finite, and discrete. Confronted with a deficit of acknowledgement, we might wish that our



broken conversation could be fixed if what stood between us were merely a deficit of knowledge. For Anscombe, Austin, Cavell and other thinkers who wrote in the wake of Wittgenstein’s theory of language games, the precarity of natural conversation drives in us the ineluctable desire to step beyond an exchange of reasons, to attain rigid – yet illusory – rules and thereby regain our bearings in the face of uncertainty.

Users can feel a similar urge to regain their bearings when they encounter errors of acknowledgment in AI-powered language games. Conversations in “Roleplay” can stray from the intended topic or expose new contexts that bring words’ meaning into question. As technological progress brings AI-powered conversation platforms into closer proximity with natural conversation, they confront the uncertainties that cause human interaction to slip and leave us desiring words’ traction. In Duolingo, we find a mix of language games – games of knowledge and acknowledgment – that evince the possibilities and pitfalls of natural conversation.

## **5. Designing AI Chatbots to Deepen Conversation**

Until 2023, users played only rule-based games on Duolingo and Babbel. Close-ended exercises would transport users from the dynamism of everyday dialogue in order to make moves that are finite and discrete. With the integration of GPT-4 on Duolingo, continuous conversation launched “Roleplay” into the orbit of natural turn-taking communication. Yet, desire for close-ended features with finite and discrete designs persisted. Clickable individual words as well as drag-and-drop vocabulary offer relief from the uncertainties and infelicities that can arise when users feel lost in continuous conversation. “Roleplay” includes design features that provide footing when the grounds of the game slip away. The language game’s interface offers a mix of rigid rules and dynamic interactions.

Facets in “Roleplay” appear beneath the dialogue frame and provide filtered navigation when a user is unsure how to proceed. The feature displays a finite set of words to select and add to the dialogue. The facets are spatially positioned outside the conversation, providing rule-based relief, as it were, when users are unsure what to say. When errors of acknowledgment arise in natural conversation, we might desire a script to know which moves to make. The urge, although unattainable in natural conversation, remains a real and insistent impulsion that Duolingo visualizes in the form of faceted conversation. Roleplay thereby offers knowledge when users face errors of acknowledgment.



Figure 7. Facets in Duolingo's "Roleplay"

A second design feature is translation overlays, which appear when users hover over foreign words in the dialogue frame. At times when the user is unsure of words' meaning, especially when multiple meanings might be at play, he or she can make the English translation appear. It opens an escape hatch, an opportunity to flee the uncertainty of translation, and momentarily return to a language (game) whose rules are known. Unsure of another's intentions in natural conversation, we might desire telepathic powers to know what he or she *really means* behind the veil of language. "Roleplay" accommodates the desire with translation overlays.

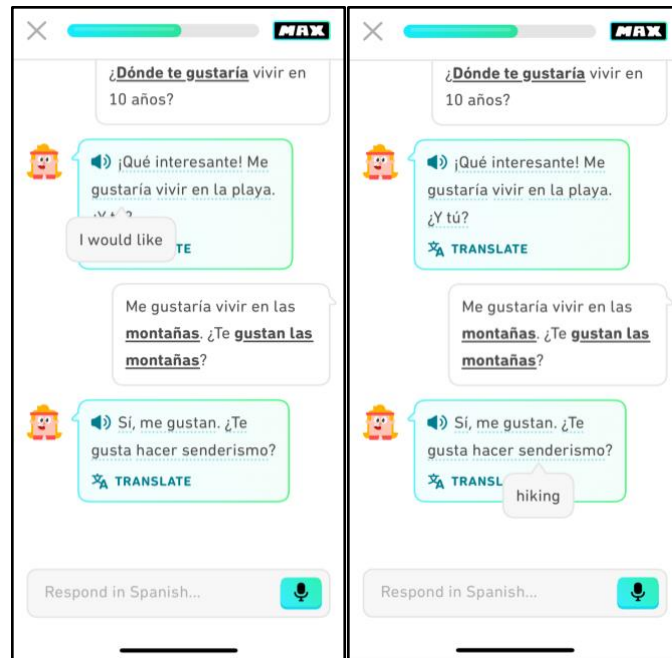


Figure 8 a–b. Translation overlays in Duolingo’s “Roleplay”

Both design features make visible what language hides. They respond to the temptations of a transparent communication liberated from the opacities and illusions of language. Whereas misunderstanding and incomprehension can leave us groping to secure acknowledgment in day-to-day interactions, AI interactions secure traction for understanding and comprehension. Facets and translation overlays make text clickable. The errors of acknowledgment implicit to natural conversation appear as knowledge explicit in artificial conversation. As a result, the features fulfill the illusory – yet inescapable – urge for a transparent conversation devoid of precarity. Language games with chatbots rise above the games we always already play in language.

We might quit playing language games with our heart, and start playing language games with chatbots, when AI-powered conversation combines rule-based, clickable options with the dynamic realism of back-and-forth conversation. The future of AI conversation, I’m suggesting, is one in which the rule-based interactions of traditional decision trees co-exist with intention-based interactions powered by generative AI. We are just beginning to witness AI platforms’ capacity to mimic the continuity of natural conversation; progress is to be made by learning from the pitfalls of human communication as well. With AI platforms as our interlocutors, companionship takes many forms. Users might learn a skill with Claude and treat the

platform as an educator. Others might test a thesis with Chat-GPT and use it as a debate coach. Perhaps we explore a niche topic with Gemini and interact with it as we would with a librarian. Across diverse language games, conversational AI should bring to light the ineluctable desires churning beneath the surface of natural conversation; to thereby suture artificially what our natural words, as the Backstreet Boys sang, can break asunder:

I should have known from the start,  
You know you have got to stop,  
You are tearing us apart,  
Quit playing (language) games with my heart.

## References Cited

- Anscombe, G.E.M. 1957. *Intention*. Basil Blackwell.
- Aristotle. 1999. *Nicomachean Ethics*, translated by Terence Irwin, 2nd ed. Hackett Publishing Company.
- Austin, J.L. 1962. *How to Do Things with Words*, edited by J.O. Urmson. Oxford University Press.
- Babbel. 2024. "About Us." Babbel, September 16. <https://uk.babbel.com/about-us>
- Babbel. 2023. "Learn with Your Own Voice: Babbel Launches Two New Speech-Based Features." Babbel, May 5. <https://www.babbel.com/press/en-us/releases/learn-with-your-own-voice-babbel-launches-two-new-speech-based-features-us#>
- Brentano, Franz. 1874. *Psychology from an empirical standpoint*, edited by A. C. Rancurello, D. B. Terrell, & L. L. McAlister (1995). Routledge.
- Cavell, Stanley. 2004. *Cities of Words: Pedagogical Letters on a Register of the Moral Life*. Harvard University Press.
- Cavell, Stanley. 1969. *Must We Mean What We Say?* Cambridge University Press.
- Chomsky, Noam. 2023. "The False Promise of ChatGPT." *The New York Times*, March 8, 2023. <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>
- Duolingo. 2023a. "Large Language Models + Duolingo = More Human-Like Lessons." Duolingo Blog, April 12, 2023. <https://blog.duolingo.com/large-language-model-duolingo-lessons/>
- Duolingo. 2023b. "Introducing Duolingo Max, a Learning Experience Powered by GPT-4." Duolingo Blog, March 14, 2023. <https://blog.duolingo.com/introducing-duolingo-max-powered-by-gpt-4>

- Freeman, Cassie, Audrey Kittredge, Hope Wilson, and Bożena Pająk. 2023. *The Duolingo Method for App-Based Teaching and Learning*. Duolingo Research. <https://duolingo-papers.s3.amazonaws.com/reports/duolingo-method-whitepaper.pdf>
- Grice, H.P. 1975. "Logic and Conversation," *Syntax and Semantics 3: Speech Acts*. Academic Press.
- Ji, Ziwei, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii. 2023. "Survey of Hallucination in Natural Language Generation." *ACM Computing Surveys* 55 (12): Article 248. <https://doi.org/10.1145/3571730>
- Jiang, Xiangying, Jenna Rollinson, Luke Plonsky, Erik Gustafson, and Bożena Pająk. 2024. "Evaluating the Reading and Listening Outcomes of Beginning-Level Duolingo Courses." *Foreign Language Annals* 54 (4): 974-1002. <https://doi.org/10.1111/flan.12600>
- Kosinski, Michal. 2023. "Theory of Mind May Have Spontaneously Emerged in Large Language Models." *arXiv Preprint arXiv:2302.02083*. <https://doi.org/10.48550/arXiv.2302.02083>
- Li, Xinyuan. 2023. "'There's No Data Like More Data': Automatic Speech Recognition and the Making of Algorithmic Culture." *Osiris* 38 (1): 165–82
- Matsakis, Louise. 2023. "Air Canada's Chatbot Soars Out of Control, Invents Refund Policy," *Wired*, February 6. <https://www.wired.com/story/air-canada-chatbot-refund-policy/>
- McDowell, John. 1984. "Wittgenstein on Following a Rule," *Synthese* 58 (3): 325-363.
- Norman, Donald. 2023. "AI: First New UI Paradigm in 60 Years," Nielsen Norman Group, June 18. <https://www.nngroup.com/articles/ai-paradigm/>
- Packer, J.D., and James Reeves, eds. 2023. *Abstractions and Embodiments: New Histories of Computing and Society*. Baltimore: Johns Hopkins University Press
- Sacks, Harvey, Emanuel A. Schegloff, and Gail Jefferson. 1974. "A Simplest Systematics for the Organization of Turn-Taking for Conversation." *Language* 50, (4): 696-735
- Saleem, Awaz, Narwin Noori, and Fezile Ozdamli. 2022. "Gamification applications in E-learning: A literature review." *Technology, Knowledge, and Learning* 27: 139–159. <https://doi.org/10.1007/s10758-020-09487-x>
- Searle, John. 1980. "Intention and action," *Philosophical Perspectives* 4: 79-105.
- Summers-Stay, David, Christopher R. Voss, and Scott M. Lukin. 2023. "Brainstorm, Then Select: A Generative Language Model Improves Its Creativity Score." In *The AAAI-23 Workshop on Creative AI Across Modalities*.

Tommerdahl, Jodi, A. Olsen, and C. Dragonflame. 2022. "A Systematic Review Examining the Efficacy of Commercially Available Foreign Language Learning Mobile Apps." *Computer Assisted Language Learning* 37 (1): 1-30. <https://doi.org/10.1080/09588221.2022.2035401>

Wakabayashi, Daisuke. 2022. "Google Engineer Says AI Has 'Consciousness' in Remarkable Claim." *The New York Times*, July 23, 2022. <https://www.nytimes.com/2022/07/23/technology/google-engineer-artificial-intelligence.html>

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Translated by G.E.M. Anscombe (2001). Blackwell.