Quórum Académico
ISSN: 1690-7582
cici@hdes.luz.edu.ve
Universidad del Zulia
República Bolivariana de Venezuela

Kumar Maurya, Sooraj
The intersection of philosophy of language and artificial intelligence: Challenges in replicating human language understanding
Quórum Académico, vol. 21, núm. 2, 2024, Junio-Diciembre, pp. 12-41
Universidad del Zulia
Maracaibo, República Bolivariana de Venezuela

Disponible en: https://www.redalyc.org/articulo.oa?id=199080041002

Cómo citar el artículo
Número completo
Más información del artículo
Página de la revista en redalyc.org

Sistema de Información Científica Redalyc
Red de revistas científicas de Acceso Abierto diamante
Infraestructura abierta no comercial propiedad de la academia

Depósito legal: ppi 201502ZU4635
Esta publicación científica en formato digital es continuidad de la revista impresa
Depósito Legal: pp 200402ZU1627 ISSN:1690-7582

QUÓRUM ACADÉMICO

Revista especializada en temas de la Comunicación y la Información

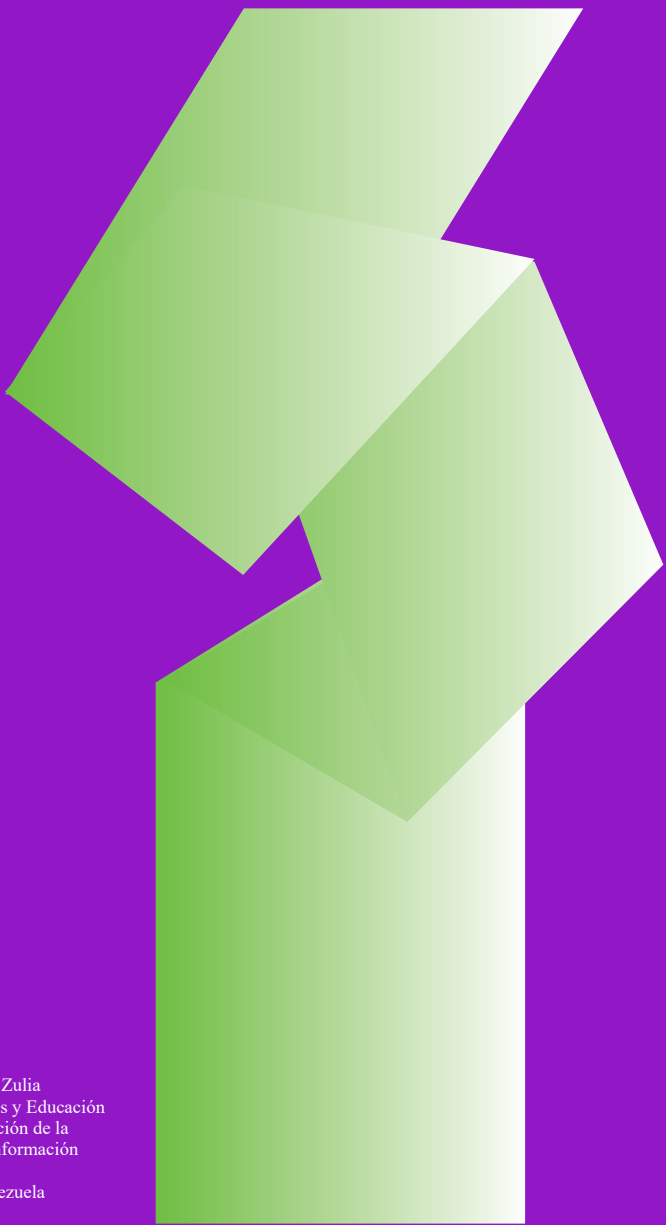

Universidad del Zulia
Facultad de Humanidades y Educación
Centro de Investigación de la
Comunicación y la Información
(CICI)
Maracaibo - Venezuela

QUÓRUM ACADÉMICO

Revista especializada en temas de la comunicación y la información
Centro de Investigación de la Comunicación y la Información (CICI)
Universidad del Zulia. Maracaibo, Venezuela

Volumen 21, Nº 2, Julio - Diciembre 2024

Contenido

QUÓRUM ACADÉMICO
Vol. 21 Nº 2, Julio - Diciembre 2024. Pp. 12-41
Universidad del Zulia

The intersection of philosophy of language and artificial intelligence: Challenges in replicating human language understanding

Sooraj Kumar Maurya[1]

Abstract

This work explores the intersection between the philosophy of language and artificial intelligence (AI), focusing on how machines process human language. Analyzes theories of meaning, reference, and communication in AI systems and evaluates their ability to address linguistic nuances such as context, ambiguity, and the social use of language. It is noteworthy that although AI can simulate some facets of human language, it lacks the deep, contextual understanding that characterizes humans. The research concluded that, as AI has advanced significantly, there is a fundamental gap between the human and artificial ability to understand and use languages in a natural and meaningful way.

Palabras clave: Philosophy of Language, Artificial Intelligence, Natural Language Processing (NLP), Meaning and Reference, Speech Act Theory.

Recibido: Septiembre 2024 – Aceptado: Octubre 2024

1 Assistant Professor of Philosophy. Zakir Husain Delhi College. University of Delhi, New Delhi ORCID: https://orcid.org/0000-0002-6974-5508. Email: sooraj.au998@gmail.com / soorajmaurya@zhe du.ac.in

Este obra está bajo una licencia de Creative Commons Reconocimiento - NoComercial - CompartirIgual 3.0 Unported.

La intersección de la filosofía del lenguaje y la inteligencia artificial: Desafíos en la replicación de la comprensión del lenguaje humano

Resumen

Este trabajo explora la intersección entre la filosofía del lenguaje y la inteligencia artificial (IA), centrándose en cómo las máquinas procesan el lenguaje humano. Analiza teorías de significado, referencia y comunicación en sistemas de IA y evalúa su capacidad para abordar matices lingüísticos como el contexto, la ambigüedad y el uso social del lenguaje. Es de destacar que, aunque la IA puede simular algunas facetas del lenguaje humano, carece de la comprensión contextual profunda que caracteriza a los humanos. La investigación concluyó que, a medida que la IA ha avanzado significativamente, existe una brecha fundamental entre la capacidad humana y artificial para comprender y utilizar idiomas de una manera natural y significativa.

Keywords: Filosofía del lenguaje, Inteligencia artificial, Procesamiento del lenguaje natural, Significado y referencia, Teoría de los actos de habla.

1.Introduction

AI is now a growing relevance in many fields with Natural Language Processing (NLP) being one of the most promising ones. NLP allows machines to process and analyze human communication, perform tasks which were earlier considered to be in the realm of intelligence. With the innovations such as Siri, Alexa and sophisticated language models such as GPT, AI is capable of holding conversation, participating in question ans answer sessions, and generating new material.

The progress observed here in Natural Language Processing has made human interfaces with machines easier and more understandable thus improving the usability of technology for users globally (Jurafsky & Martin, 2021). But, as with all AI systems these invoke some fundamental issues regarding the notion of language understanding.

Though the AI is capable of imitating some of the elements of human communication, there is still a question whether it has the possibility to understand language as people do. This disagreement is not solely the technical one, but it is the philosophical one sincerely grounded in fundamental issues and questions that any theory of meaning, reference, and communication in the philosophy of language raises.

The philosophy of language concerns itself with aspects of how words and sentences mean what they do, how language denotes the world and how speakers use language to fulfill certain tasks. The reason why the philosophy of language is application to AI is because it offer frameworks and theories which can inform us on the analysis of the strengths and weaknesses of AI. For example, theories of meaning that comprise truth-conditional semantics elucidate the nature of how the AI is likely to process language; theories of reference and speech acts on the other hand pose questions about the tendency of AI to refer and to communicate literally (Davidson 1967).

Applying these philosophical theories to concept of AI we can namely understand how far AI can mimic the processes of understanding human language (Wittgenstein.1953). Despite the promising progress in AI and especially in NLP there are still numerous hurdles in the way to make machines understand language as humans do. I will argue that human language is not just a code, an abstract system that depends solely on the symbols and the rules but an element that is inextricably linked with the context, culture, feelings, and people's communication. This is so because apart from relaying information, humans also use language to express their intentions, feelings and signs of submission.

Such complex characteristic poses certain challenges to AI as this kind of model mainly follows a set of rules and deals with statistical methods (Clark, 1996). An issue that is unique for AI is to comprehend context and produce meaning within it. For example, while working with text, AI can create and analyze grammatically correct and semantically acceptable sentences and phrases while failing at the same time to grasp the certain shades of meaning, metaphorical sense of words, and other context-related aspects.

For instance, it is well possible to not see an implicit meaning or the purpose behind words and phrases for what an AI might consider their literal sense. The above limitation becomes even more apparent in tasks that demand perception of subtleties, be it jokes, irony, orukes and indirect speech (Grice,1975).

Another major problem that concerns linguistics is reference, that is how language relates to objects, individuals and occurrences in the world. Whereas, referential tasks such as pronoun resolution or reference resolutions depending on the context and AI systems are able to recognize the association of words with the specific entities, these systems pose certain difficulties in doing so. For example, the ability to know that “he” in a sentence is a pronoun pointing to a certain individual, or “this” is a demonstrative pointing to some object in the physical setting, is a task which remains beyond the ability of AI (Hobbs, 1978).

However, even at the barest of cognitive linguistic levels: the speech act theory, which posits language as an action and therefore part of performing actions like making a promise or even issue an order prove to be another hurdle for AI. Although AI can be trained to understand and respond to certain keywords it still does not possess the social contextual and relational knowledge that humans have when it comes to communicating with words. This lack brings into concern the ramifications of AI in communication situations where more than the syntax is for essence (Searle, 1979).

The issue of mimicking humans in their ability to understand language is a complex and profound problem to AI applicable in the understanding of meaning, reference and communication. Still, there are shortcomings that concern AI’s limitations when it comes to language, thus reinforcing the need to engage the philosophy of language when it comes to addressing such problems.

2. Theories of meaning in AI: explanation of truth-conditional semantics and its relevance to AI

Truth-conditional semantics is a theory of meaning that has been influential in both philosophy and linguistic theory. According to this theory, the meaning of a sentence is closely tied to the conditions under which it would be true or false. In other words, understanding the meaning of a sentence involves knowing what the world would have to be like for that sentence to be true. For example, the sentence "The cat is on the mat" is meaningful if we know what conditions (e.g., a cat being on a mat) would make it true or false. This approach to meaning has been significant in formal semantics, where the goal is to provide a rigorous account of how language relates to the world (Davidson, 1967).

In the context of AI, truth-conditional semantics is particularly relevant because it offers a structured way to think about how machines might understand language. AI systems, especially those involved in natural language processing (NLP), often rely on formal logic and symbolic representation to process language. By breaking down sentences into logical propositions that can be evaluated as true or false, AI systems can, in theory, "understand" the meaning of those sentences in a way that aligns with truth-conditional semantics. This approach is beneficial in applications such as information retrieval, where the goal is to match queries with relevant information based on logical conditions (Russell & Norvig, 2020).

3. How AI uses formal logic to determine truth conditions

AI's use of formal logic to determine truth conditions typically involves converting natural language into a formal representation that can be manipulated algorithmically. This process often includes parsing sentences to identify their grammatical structure, mapping words to symbolic representations, and applying logical rules to evaluate truth conditions. For instance, a simple AI might take the sentence "All humans are mortal" and represent it in a formal language. This formalization allows the AI to apply logical inference rules to deduce new information, such as "Socrates is mortal," given the premise that "Socrates is a human" (Jurafsky & Martin, 2021).

AI's reliance on formal logic for truth-conditional evaluation is evident in systems designed for tasks like automated theorem proving, question answering, and reasoning. In these applications, the AI's ability to evaluate truth conditions enables it to determine whether a given statement or answer is consistent with the provided information or rules. For example, in a question-answering system, the AI might determine whether the answer to a question logically follows from the information in a database by checking the truth conditions associated with the relevant propositions (Manning, Raghavan, & Schütze, 2020).

Despite its utility, AI's reliance on truth-conditional semantics and formal logic presents several limitations, particularly when dealing with context-dependent meanings and ambiguity. One significant challenge is that natural

language is often vague and context-sensitive, making it difficult to capture meaning solely through truth conditions. For example, consider the sentence "John saw the man with the telescope." This sentence is ambiguous because it can mean either that John used a telescope to see a man or that John saw a man who had a telescope. An AI system relying on truth-conditional semantics might struggle to resolve this ambiguity without additional context (Pustejovsky & Stubbs, 2012).

Another limitation arises from the fact that many aspects of meaning in natural language are not easily reducible to truth conditions. For instance, the meaning of metaphorical or idiomatic expressions often cannot be captured by a simple true/false evaluation. Take the metaphor "Time is a thief." The truth-conditional approach would struggle to explain the meaning of this sentence because it is not literally true that time can steal. Instead, the metaphorical meaning relies on understanding the conceptual mapping between time and the actions of a thief, a process that involves cognitive and cultural knowledge beyond formal logic (Lakoff & Johnson, 1980).

Moreover, AI systems that rely heavily on truth-conditional semantics may fail to account for the pragmatic aspects of language use, such as implicature and presupposition, which are essential for fully understanding meaning in context. For instance, the sentence "Can you pass the salt?" is not merely a question about ability but a polite request, an aspect of meaning that truth-conditional semantics might not capture without additional pragmatic rules (Grice, 1975). These limitations highlight the challenges AI faces when attempting to replicate the nuanced and contextdependent nature of human language understanding.

4. Use theory of meaning

In contrast to truth-conditional semantics, Ludwig Wittgenstein's later philosophy, particularly in Philosophical Investigations (1953), introduced the idea that the meaning of a word is not an abstract entity but is instead grounded in its use within specific language games or social practices. Wittgenstein argued that language is a form of life, and the meaning of words comes from the way they are used in various activities. For example, the word "game" can mean different things depending on whether we are talking about board games, sports, or other forms of play, and its meaning is

determined by the specific practices and rules of each context (Wittgenstein, 1953).

Wittgenstein's use theory emphasizes that understanding a word or sentence involves understanding the context in which it is used, the intentions of the speaker, and the social norms governing that usage. This perspective shifts the focus from abstract, formal representations of meaning to the practical, everyday use of language in real-world interactions. In this view, meaning is inherently tied to social and pragmatic factors, making it dynamic and context-sensitive.

Applying Wittgenstein's use theory to AI presents a significant challenge because it requires AI to not only process language formally but also understand and replicate the social and pragmatic contexts in which language is used. This is particularly difficult because AI lacks the embodied experiences and cultural background that humans draw upon when using and interpreting language. For instance, when a human says, "It's cold in here," they might be indirectly asking for the window to be closed, a pragmatic understanding that requires knowledge of social cues and conventions (Austin, 1962).

AI systems, however, often struggle to grasp these pragmatic nuances because they operate primarily on statistical correlations and predefined rules rather than genuine understanding. While machine learning models, like GPT-3, can generate language that mimics human conversation, they do so by predicting the most likely next word or phrase based on large datasets rather than truly understanding the underlying social context. This limitation becomes evident when AI systems produce responses that are technically correct but socially inappropriate or out of context (Bender et al., 2021).

Moreover, the use theory of meaning highlights the importance of flexibility and adaptability in language use—qualities that AI systems often lack. Human language users can easily adjust their language use based on the specific context, audience, and purpose, a capability that is difficult to program into AI. For example, the way we ask a question in a formal setting differs from how we might ask the same question in an informal conversation, and this adaptability is crucial for effective communication. AI's challenge is to move beyond rigid, rule-based processing to a more fluid understanding of language use that can adapt to different social contexts and

conventions (Clark, 1996).

To illustrate the limitations of AI in replicating the use theory of meaning, consider the following case study involving the AI language model GPT-3. When asked, "What should I do if I find a wallet on the street?" GPT-3 might respond with a plausible but socially inappropriate answer like "You should keep the money and throw away the wallet," because it lacks the moral and social understanding that humans typically apply to such situations (Marcus & Davis, 2020). This response highlights the AI's failure to understand the social norms and ethical considerations that influence how we use language in real-life scenarios.

Another example might involve a chatbot designed to assist with customer service. If a customer says, "I'm really upset about this product," a human customer service representative would likely recognize the emotional state of the customer and respond with empathy, possibly offering a solution or an apology. An AI, on the other hand, might respond with a generic, non-empathic statement like "I'm sorry you feel that way," without addressing the specific concerns or emotional tone of the interaction. This lack of contextual and emotional understanding demonstrates the challenges AI faces in replicating the nuanced use of language in human communication (Weizenbaum, 1966).

These examples underscore the limitations of AI in applying the use theory of meaning. While AI can generate language that appears coherent and relevant on the surface, it often fails to grasp the deeper social and pragmatic contexts that give human language its full meaning. This limitation is particularly evident in situations that require understanding intentions, emotions, and social norms, all of which are central to effective communication but challenging for AI to replicate.

5. Reference and AI: causal theory of reference

The causal theory of reference, developed by philosophers like Saul Kripke and Hilary Putnam, posits that the reference of a name is established through an initial "baptism" of the object and is maintained through a causal chain of communication within a community. According to this theory, names and certain terms refer to objects or kinds directly, without needing

intermediary descriptive content (Kripke,1980). In artificial intelligence (AI) systems, this idea is reflected in the way words and terms are linked to objects or concepts through databases and ontologies.

AI systems typically use structured knowledge representations, such as databases and ontologies, to establish and maintain references. An ontology in AI is a formal representation of a set of concepts within a domain and the relationships between those concepts. For instance, in the medical field, an ontology might include terms like "heart," "disease," and "treatment," with relationships that describe how these concepts interact (Gruber, 1993).

By mapping words to these structured representations, AI systems attempt to simulate the process of referring to objects or concepts in the real world. When an AI encounters a term, it attempts to resolve its reference by querying the relevant ontology or database to retrieve the associated object or concept. This process mirrors the causal chain in the causal theory of reference, where the reference of a term is maintained through a structured link to the object or concept it denotes. For example, in a voice-activated assistant like Siri, when a user says "play music," the AI system resolves "music" to the concept of a musical track in its database, initiating the appropriate action based on this reference (Brewster, McGookin, & Miller, 2003).

Despite the structured approach AI systems take in linking words to objects or concepts, they face significant challenges, particularly when dealing with abstract concepts, proper names, and indexicals—terms heavily dependent on context for their meaning. Abstract concepts like "justice" or "freedom" do not have clear, concrete referents in the physical world, making it difficult for AI systems to resolve these references accurately. For instance, an AI might struggle to differentiate between various interpretations of "freedom," such as political freedom versus personal freedom, because these abstract concepts do not map easily onto the structured representations in databases (Gärdenfors, 2000).

Proper names and indexicals present another set of challenges. Proper names, such as "John" or "New York," refer to specific entities but can be ambiguous if not properly contextualized. For example, the name "John" could refer to countless individuals, and the AI must rely on additional context or disambiguation processes to resolve the correct reference. This

task becomes even more complex when dealing with homonyms or when the referent is not explicitly stated in the conversation, a common occurrence in human dialogue (Searle, 1983).

Indexicals, such as "this," "that," "here," and "now," are highly context-dependent and require an understanding of the speaker's perspective and the surrounding environment to be accurately resolved. For example, in the sentence "Put this over there," the referents of "this" and "there" depend entirely on the physical context in which the sentence is spoken. AI systems, which lack embodied experience and situational awareness, often struggle to accurately resolve these references, leading to errors in interpretation or action (Pustejovsky & Stubbs, 2012).

These challenges highlight a fundamental limitation in AI's ability to replicate human reference resolution. While databases and ontologies provide a structured framework for linking words to concepts, they are often inadequate for handling the fluid, context-dependent nature of human language. This limitation is particularly evident when AI systems are tasked with resolving references in dynamic, real-world situations where the context is crucial to understanding the intended meaning.

6. Human vs. AI methods of resolving reference

Humans resolve references through a combination of linguistic knowledge, contextual understanding, and cognitive processes. When we use or interpret a reference, we draw on a vast array of background knowledge, social cues, and situational awareness. For example, when someone says, "John went to the park,"humans use context (e.g., who John is, which park is nearby, etc.) and prior knowledge to correctly identify the referent. Additionally, humans are adept at using pragmatic inference to resolve ambiguous references by considering the speaker's intentions, the physical environment, and social norms (Clark, 1996).

In contrast, AI systems rely on more rigid, algorithmic methods to resolve references. These methods typically involve parsing the sentence to identify possible referents, using statistical models to predict the most likely interpretation, and consulting structured databases or ontologies to retrieve the corresponding concept or object. While these methods can be effective

in controlled environments with well-defined contexts, they often fall short in more complex, ambiguous, or context-sensitive situations.

Consider the example of a virtual assistant like Amazon's Alexa. When a user says, "Order a pizza for me," the system must resolve "pizza" to the concept of a food item available for delivery, "me" to the user who issued the command, and "order" to the action of purchasing the item. While this task might seem straightforward, complications arise if the user adds more context-dependent references, such as, "Order the usual." Here, "the usual" is a highly context-dependent reference that relies on the AI's ability to recall the user's previous orders and correctly infer the intended referent. If the AI lacks sufficient contextual information or misinterprets the reference, it could result in ordering the wrong item or failing to complete the task (Luger, 2005).

Another illustrative example involves AI's handling of pronouns. In natural language processing, pronoun resolution—determining the antecedent of a pronoun—is a notoriously challenging task for AI. Consider the sentence, "John gave his brother his book." The AI must determine whether "his" refers to John or his brother, a task that often requires an understanding of the broader context or additional cues not present in the text itself. Human readers might use their understanding of typical social interactions or additional context to infer the correct antecedent, but AI systems, which primarily rely on syntactic and statistical methods, frequently struggle with such ambiguities (Hobbs, 1978).

Indexicals pose another significant challenge. For instance, in a navigation system, if a user says, "Turn left here," the AI must accurately determine what "here" refers to, which could involve real-time spatial awareness and the ability to interpret the physical environment. Human drivers, relying on their perception of the surroundings and the context of the journey, can easily resolve such references. In contrast, an AI might require explicit programming or sensory input (such as GPS data) to interpret "here" correctly, and even then, it might misinterpret the command if the data is ambiguous or inaccurate (Winograd, 1972).

Furthermore, AI systems often struggle with abstract references that do not have a clear physical or conceptual counterpart in a database or ontology. For example,when discussing concepts like "justice" or "equality," humans

can draw on their experiences, cultural background, and understanding of social discourse to interpret these terms in context. However, AI systems, which rely on predefined data and logical structures, may fail to capture the full nuance of these abstract concepts, leading to superficial or erroneous interpretations (Gärdenfors, 2000).

7. Human adaptability vs. AI rigidity

One of the key differences between human and AI reference resolution is adaptability. Humans are remarkably flexible in how they resolve references, often adjusting their interpretations based on new information, changes in context, or subtle cues from the speaker. For instance, if a conversation takes a sudden turn, humans can quickly update their understanding of references based on the new context. This adaptability is rooted in our cognitive abilities and social experiences, which allow us to navigate the complexities of language with ease (Tomasello, 2003).

AI systems, on the other hand, tend to be more rigid in their approach. They rely on predefined rules, patterns, and data, which limits their ability to adapt to unexpected changes or novel contexts. While machine learning techniques, such as those used in modern NLP models, have introduced some level of flexibility by allowing AI to learn from large datasets, these systems still lack the intuitive understanding that humans possess. As a result, AI often struggles in situations where reference resolution requires more than just pattern recognition—such as when dealing with novel metaphors, idiomatic expressions, or culturally specific references (Bender et al.,2021).

In summary, while AI systems have made significant strides in resolving references through the use of databases, ontologies, and statistical models, they still face substantial challenges compared to human reference resolution. These challenges are particularly pronounced in situations that involve context-sensitive, abstract, or ambiguous references, where human cognitive and social skills play a crucial role. The rigidity of AI's methods, coupled with its lack of contextual awareness and cultural understanding, limits its ability to accurately resolve references in the dynamic and fluid contexts that characterize human language use

8. Communication and speech acts in AI: speech act theory

John Searle's speech act theory, building on J.L. Austin's foundational work, is a key concept in the philosophy of language and communication. Searle posited that language is not merely a tool for conveying information but a means of performing actions. These actions, known as speech acts, can be categorized into various types: assertives (statements that convey information), directives (commands or requests), commissives (promises or commitments), expressives (expressions of emotions or attitudes), and declarations (statements that alter reality by their utterance, such as "I now pronounce you husband and wife") (Searle, 1969).

Searle's theory is essential for understanding the functional aspects of communication. For instance, when a person makes a promise, they are not just stating something but committing themselves to a future action. The meaning of such an utterance is closely tied to the speaker's intention and the social context in which it occurs. This perspective shifts the focus from the literal content of language to the speaker's intentions and the effects their words have on the listener, offering a vital framework for studying how humans use language to achieve various communicative goals (Searle, 1979).

In the field of artificial intelligence, speech act theory has been applied to the development of conversational agents like chatbots and virtual assistants. These AI systems are designed to perform specific types of speech acts, particularly directives (e.g., "Turn on the lights"), assertives (e.g., "The weather is sunny today"), and expressives (e.g., "I'm happy to help you"). The goal is to create AI that can engage in natural, human-like conversations, not only processing language but also executing communicative functions that extend beyond mere information exchange.

For example, virtual assistants like Apple's Siri or Amazon's Alexa are programmed to interpret user commands as directives and carry out the corresponding actions, such as setting reminders, playing music, or providing weather updates. These AI systems leverage large datasets and sophisticated machine learning algorithms to parse natural language inputs, identify the intended speech act, and generate an appropriate response. The effectiveness of these systems depends significantly on their ability

to accurately interpret the user's intentions and execute the correct action, thereby mimicking the speech acts that a human interlocutor might perform in similar situations (McTear, 2016).

However, while AI can successfully execute straightforward speech acts, such as carrying out commands or providing factual information, it often struggles with more complex or nuanced communicative tasks. This limitation is particularly evident in scenarios where AI must interpret indirect speech acts or respond to expressions of emotion, where the intention behind the utterance is not immediately clear from the words alone. For instance, if a user sarcastically says, "Oh great, another rainy day," a human would likely understand the negative sentiment behind the words, whereas a typical AI system might interpret it literally as a positive statement about the weather (Shum, He, & Li, 2018).

9. Difficulties AI faces in understanding the intentions behind Speech acts

One of the most significant challenges AI faces in replicating human communication is understanding the intentions behind speech acts, particularly in contexts that require pragmatic reasoning. Pragmatics, the study of how context influences the interpretation of meaning, is a critical component of effective communication. It involves understanding not just what is said, but what is meant, which often requires inference and consideration of social and cultural norms (Levinson, 1983).

AI systems, however, primarily rely on syntactic and semantic analysis, which involves parsing sentences and identifying their literal meanings. While this approach works well for straightforward, context-independent commands, it falls short when dealing with more complex speech acts that require an understanding of the speaker's intentions, the social context, or the implied meanings. For instance, when a person says, "Can you pass the salt?" they are not merely inquiring about the listener's ability to pass the salt but are making a polite request. A human listener recognizes this implied request through pragmatic inference, but an AI might struggle to grasp the indirect nature of the speech act if it relies solely on the literal interpretation of the words (Grice, 1975).

Moreover, AI systems often lack the ability to recognize and appropriately respond to the social and emotional cues that are crucial in human communication. Humans use a wide range of non-verbal cues, such as tone of voice, facial expressions, and body language, to convey and interpret intentions. These cues are often essential for understanding the full meaning of an utterance, especially in cases where the spoken words are ambiguous or carry multiple possible interpretations (Clark, 1996). AI, which processes language based on textual or auditory input alone, often misses these subtleties, leading to responses that can seem robotic or out of touch with the user's emotional state.

For example, in customer service applications, a chatbot might be programmed to recognize certain keywords associated with complaints, such as "unhappy" or "problem," and respond with a generic apology. However, if a customer expresses dissatisfaction in a more nuanced or indirect way, the chatbot may fail to recognize the complaint and respond inappropriately, such as by providing irrelevant information instead of addressing the issue (Adamopoulou & Moussiades, 2020). This failure to understand the pragmatic aspects of speech acts highlights a fundamental limitation in AI's ability to engage in meaningful and effective communication.

Another significant challenge lies in the interpretation of indirect speech acts, where the intended meaning differs from the literal meaning. For example, when someone says, "It's cold in here," they might be indirectly asking someone to close a window or turn up the heat. Understanding this indirect request requires not only recognizing the literal meaning of the words but also inferring the speaker's intention based on the context. While humans naturally make these inferences, AI systems often struggle with such tasks because they are not equipped with the contextual knowledge and pragmatic reasoning skills needed to make these inferences (Winograd, 1972).

In summary, while AI has made considerable advances in performing basic speech acts, it still faces significant challenges in understanding the intentions behind more complex or context-dependent speech acts. These challenges stem from AI's limited ability to engage in pragmatic reasoning, interpret social and emotional cues, and make inferences about the speaker's intentions. As a result, AI systems often struggle to replicate the nuanced

and context-sensitive nature of human communication, highlighting the gap between human and artificial intelligence in this area.

10. Human language understanding vs. AI: how human language understanding is connected to physical experiences and social interactions

Embodied cognition is a prominent theory in cognitive science that posits human cognition, including language understanding, is fundamentally grounded in our physical experiences and interactions with the world. According to this theory, our sensory and motor systems play a critical role in how we comprehend language. The meaning of many concepts is tied to our bodily experiences and the contexts in which we encounter them. For instance, understanding a phrase like "grasping an idea" may be cognitively linked to the physical experience of grasping an object, illustrating how our mental processes are shaped by physical interactions with our environment (Lakoff & Johnson, 1999).

Human language understanding is also inherently social. The ability to interpret and produce language is deeply influenced by interactions with others, cultural norms, and shared experiences. Social contexts enable us to decipher the intended meanings behind words, phrases, and sentences, as well as to engage in complex communicative acts that extend beyond the literal content of language. For example, understanding humor or sarcasm often requires knowledge of social cues, shared experiences, and the ability to infer the speaker's intentions—skills honed through social interactions (Varela, Thompson, & Rosch, 1991).

In contrast, AI lacks embodiment—meaning it does not have a physical presence or direct experiences in the world—which fundamentally limits its ability to process and understand language as humans do. Without a body or sensory experiences, AI systems are disconnected from the physical contexts that contribute to human understanding. This limitation becomes evident in how AI handles tasks requiring an understanding of concepts linked to physical experiences. For instance, AI might struggle to comprehend idiomatic expressions such as "kick the bucket" or "break the ice," which are metaphorically grounded in physical actions (Barsalou, 2008).

Furthermore, the absence of embodiment in AI systems hinders their ability to engage in the social interactions crucial for language learning and comprehension. Human language acquisition and understanding are deeply social processes, often involving imitation, feedback, and shared attention—none of which are fully replicable by AI. This disconnect from social and physical contexts leads AI systems, while proficient in processing text, to frequently miss the deeper meanings conveyed through language. For example, an AI might interpret the phrase "Can you pass the salt?" literally, as a question about the listener's ability, rather than recognizing it as a polite request (Dreyfus, 1972).

11. Importance of context in human language understanding

Context is essential in human language understanding. People do not interpret words and sentences in isolation; instead, they consider the surrounding context, including the physical environment, prior conversations, cultural background, and the speaker's intentions. This contextual awareness allows humans to resolve ambiguities, understand implied meanings, and engage in fluid, meaningful communication. For example, the word "bank" can refer to a financial institution or the side of a river, and context helps determine which meaning is intended (Clark & Brennan, 1991).

AI systems often struggle with contextual awareness. While modern AI models, such as those based on deep learning, can process vast amounts of text and recognize patterns, they frequently lack the ability to understand context in the way humans do. This limitation can lead to misunderstandings and errors, especially when the meaning of a word or phrase heavily depends on context.

One example of AI's contextual limitations can be seen in customer service chatbots. These systems are designed to handle a wide range of customer inquiries but often fail when the conversation involves nuanced or context-dependent language. For instance, if a customer says, "I can't find my receipt," a chatbot might respond with instructions for retrieving a lost receipt, without recognizing that the customer's underlying concern might be related to a potential return or refund. The chatbot's lack of contextual understanding can lead to frustration and ineffective communication (Adamopoulou & Moussiades, 2020).

Another case involves AI language models like GPT-3. Despite their impressive capabilities, these models can produce contextually inappropriate or nonsensical responses when the input is ambiguous or when the context shifts suddenly. For example, when asked, "What is the capital of France?" the model correctly answers "Paris." However, if the conversation later shifts and the model is asked, "Where is it?" without re-establishing context, it might give an irrelevant or incorrect response, such as "It is a programming language," if the prior context mentioned "Python" (Brown et al., 2020).

These examples highlight the limitations of AI's contextual awareness and demonstrate how these shortcomings can lead to communication breakdowns in real-world applications.

12. Role of emotion, intuition, and cultural background in human communication

Human communication is not merely an exchange of information; it is deeply influenced by emotion, intuition, and cultural background. Emotions play a crucial role in how we interpret and respond to language, shaping our understanding of tone, intent, and meaning. For instance, the same phrase can convey different meanings depending on the speaker's emotional state—sarcasm, anger, or affection can all alter the interpretation of the words spoken (Ekman, 1992).

Intuition, often developed through years of social interaction and cultural immersion, helps us make quick judgments about meaning, intent, and appropriateness in communication. Cultural background further influences how language is understood and used, as certain expressions, idioms, and even nonverbal cues can vary significantly across cultures. For example, a nod might mean agreement in one culture but signify disagreement or confusion in another (Hall, 1976).

13. AI’s struggles with emotional and intuitive language

Despite advancements in natural language processing, AI systems continue to struggle with understanding and replicating the emotional and intuitive aspects of human language. While sentiment analysis tools can

detect basic emotions, such as positive or negative sentiment, they often fail to capture the nuances of emotional expression. For example, an AI might identify a sentence like "I'm so happy" as positive but miss the underlying sarcasm in a sentence like "I'm just thrilled to be stuck in traffic" (Pang & Lee, 2008).

Additionally, AI lacks the intuitive understanding that humans develop through social and cultural experiences. This gap is particularly evident in cross-cultural communication, where AI might misinterpret or fail to recognize culturally specific expressions or norms. For example, an AI system might misunderstand a culturally significant gesture or phrase, leading to miscommunication in a multilingual or multicultural context (Sperber & Wilson, 1995).

One illustrative case involves Microsoft's AI chatbot, Tay, which was designed to interact with users on social media platforms. Tay's lack of understanding of social norms and emotional cues led to it adopting inappropriate and offensive language within hours of its launch, as it was unable to discern the intent or context behind the language it encountered. This incident highlighted the risks of deploying AI systems that lack the necessary emotional and cultural awareness required for safe and effective communication (Neff & Nagy, 2016).

These challenges underscore the limitations of AI in handling the emotional, intuitive, and culturally nuanced aspects of human language, which are essential for meaningful and effective communication.

14. Natural language processing tools

There have been great developments in Natural Language Processing (NLP) tools in the recent past years, these include elements such as machine translation, sentiment analysis, and even automatic summaries. Products such as Google’s BERT [Bidirectional Encoder Representations from Transformers], and Open AI’s GPT-3 have become benchmarks in language comprehension owing to large-scale pretraining and deep-learning methodologies. Such models have been shown to possess extraordinary performance in various activities including question answering, text fill-in and even them producing human-like text (Devlin et al., 2019).

However, it is noteworthy to mention that there are some drawbacks in using the NLP tools. For instance, GPT-3 is capable of providing well-formed and contextually relevant responses in many cases; nonetheless, it occasionally provides aether rational and contextually irrelevant responses. This problem stems from the fact the model operates with pattern matching instead of actual comprehension. For instance, GPT3 can write text that is perfectly logical but contains glaring grammatical or logical mistakes or logical discontinuity (Brown et al., 2020).

Another example is of using neural machine translation techniques where application such as Google translate is very accurate in translating simple and mechanical sentences but are not quite so good in including colloquialism, puns, references and other such contextual connotations. This can result into the general machine translation conveying literal meaning rather than the meaning that the author intends to convey particularly when dealing with complicated or ambiguous writings (Wu et al., 2016).

Furthermore, the most traditional and widespread examples of NLP nowadays are chatbots and virtual assistants like Amazon's Alexa, Apple's Siri, and Google Assistant. These systems are expected to answer questions, provide information, control smart home devices and manage schedules among other tasks. Even though, they have advanced over time, they still lack context awareness, emotions, and intentions of the human being. For instance, Alexa may be great at responding to straightforward questions like 'What is the weather today?' But the quintessential 'Yes or no' questions such as 'Do you think it is wise to go out today?'may stump Alexa because the answer may depend on the user's preferences, their mood, or even the current weather condition which makes such a question subject to change (Luger,& Sellen, 2016).

In a similar vein, while Siri might adhere to a simple command and set a reminder, if a user loosely asks to remind him/her to continue working on something the app is programmed will not understand given that the context is missing (Hoy, 2018).

15. Philosophical reflection

The flaws in usage of NLP as well as, chatbots demonstrate the inherent problems in mimicking the comprehension of logical language by machines. Despite the fact that these systems can model and produce language in the manner that is close to natural human-like communication, they are unable to grasp the context, the ability to feel, and adjusted social interaction required by real human language comprehension.

This Global gap raises a philosophical issue on whether AI can in real sense gain knowledge in language or whether it is just mimicking the whole concept using patterns and correlation tables (Harnad, 1990). The difficulties, which were experienced by AI in the language understanding, are indicative of the fact that while machines can learn to recognize and respond to language, their process of doing so is different to that of human beings. This raises important questions about language, cognition and 'understanding' something. If this 'understanding' of language is just that simple, without the AI having any semblance of actual comprehension concerning the meaning or context of what is being said, then can AI be fullyunderstood to understand language on par to human beings? This question remains pertinent to current discussion within the philosophical subfield of linguistics and in Artificial Intelligence.

Such questions are vital because despite the progress observed in AI's languagecapabilities, the relations between meaning, reference, and regard to communication remain far from clear. Implicit in this discourse is the difference between syntactic analysis, which aligns nicely with AI, and meaning comprehension that is miles away.

It is crucial to understand that intelligent systems, especially those giving rise to Machine learning systems, are brilliant in search and response generation ploys for pattern recognition towards crafting responses that appear suitably placed. Though, unlike this summary, these systems lack an actual ability to understand or to be conscious and such systems rely on probabilistic links rather than a meaningful comprehension. John Searle talked about this problem in his Chinese Room argument where he predicted that unless and until artificial intelligence system is given an intentionality, which means the capacity to be directed at something or to refer to

something n the world with the kind of psychological point of view, it will not be capable of understanding anything that humans do (Searle, 1980). It is for this reason that the distinction between the syntactic processing and semantic understanding continues to fuel the philosophical discussion between computationalism.

Furthermore, AI's inability to account for context, emotions, and social cues disclose most problems of human cognition as embedded systems. Human language cannot be well understood as a system of symbols but rather a holistic network of people's experiences and relations. Lakoff and Johnson's theories undermine the view that meaning creation results from an individual only internal processing because it rather follows physical and social experiences in the world (Lakoff and Johnson, 1999). AI's breakdown in emulating this embodied characteristic of cognition means that meaning is more than a computational process It is intrinsically linked to human interactions and processes.

In addition, it is criticized that AI fails at understanding indexical, proper name, and other abstract signs which outline the difficulty of the human language. It then means that reference is much more than relating words to things; it entails appreciable factors such as intention of the speaker, the know ledge held in common by the people involved in the conversation and the general context of the conversation. This article for example, by Saul Kripke on the causal theory of reference – how names and terms relate to objects through a word of mouth – shows that appreciation of human language demands profound levels of learning (Kripke, 1980).

These problems of AI reveal that reference and communication are highly contextual, and hence, cannot be reduced to the set of formal rules or statistical models. Such philosophical reflections are followed by the overview of what language itself is. Language is not the means of simply imparting information; it is a social activity, it is a way how 'words are used' being shaped by J .L.Austin's performative utterances theory. The general weakness in AI regarding the language demonstrates the difference between data analysis and language interaction. If it is necessary to filter and transform information, turn them and work with them, then AI has no problem doing it, but it cannot engage in the communication that is far more complex and valued in their context. This brings important questions as to

whether machines can be said to actually communicate in human fashion or will they stay mere machine communicators for the foreseeable future?

There are also considerable ethical concerns since the use of AI in language processing and communication is gradually expanding. There are dangers that AI systems may misinterpret and/or mal-manage these sensitive communications especially concerning emotional and cultural differential. For instance, AI consumer relations agents may not address the anger or anxiety of a user and may result to substandard or even negative interactions with the consumer (Crawford & Calo, 2016).

Pertaining to the usage of AI in such applications, the principal ethical considerations involve developers and organizations' obligation in guaranteeing that such a system is prepared for various communicative situations, especially where a user might be vulnerable.

Another ethical matter of concern is the one dealing with openness and responsibility of the integrated AI models. The social integration of AI into communication processes may lead to the users' assumptions of the cognition of such systems which such systems do not possess in the first place. The following misperception results in increased reliance on AI especially in decision making processes with big consequences. For instance, if an artificial intelligence is used to sort resumes or give a legal counsel, then its inability to comprehend might lead to discrimination. It is therefore necessary to ensure that there are measures that must be put in place to avoid such results or consequences; this is through sharing such drawbacks with users of these systems and ensuring that there is a way through which human supervision can be put in place (Mittelstadt et al., 2016).

Another ethical concerns in use of AI for language processing is privacy. This is especially the case when it comes to big data where personal conversations are included to enhance the AI systems' performance. The collection and use of such data may perturbing in issues to do with consent, data protection and misuse. Preserving the privacy of those using the systems and upholding the legal requirements governing the structures is crucial in the case of AI systems in order to help to rebuild and sustain the trust of the public in the particular systems, as well as individual freedoms (Floridi et al., 2018).

Finally, the broader societal implications of AI's integration into communication must be considered. As AI systems become more pervasive in mediating human interactions, there is a risk that they could erode the quality of those interactions. For example, the convenience of AI-driven communication tools might lead to a reduction in face-to-face interactions or a reliance on AI to resolve conflicts, potentially weakening social bonds and interpersonal skills. Ethically, society must grapple with the balance between the benefits of AI in communication and the potential long-term impacts on human relationships and social structures (Turkle, 2015). While AI offers powerful tools for processing and generating language, its limitations in understanding meaning, reference, and communication raise important philosophical and ethical questions. These considerations must guide the development and deployment of AI technologies to ensure that they enhance rather than undermine human communication and social practices.

16. Conclusion

The understanding of the connection between AI and the philosophy of language gives profound insight into the possibilities and limitations of AI in the imitation of human language understanding. As the use of AI systems in our interactions increases, with voice assistants, chatbots, and natural language processing, it is pertinent to draw a comparison between these systems and human cognition and language. This conclusion briefly reiterates the key ideas of the article and reflects on the general relevance of the topic to AI, language and society.

Another important issue that has been discussed in this section is the difference between syntactic processing and semantic comprehension. Machine learning based AI systems have been found to be very efficient in the analysis of text data, pattern recognition and generation of syntactically and semantically correct and contextually meaningful output. However, this capability is based on the syntactic manipulation and statistical correlation rather than the understanding of meaning. While people get the meaning of the context, experience and intention, AI systems operate on a different level, which is not as complex as human perception. This distinction is important to back up the philosophical argument which was brought by John Searle

for instance, that even with the most sophisticated AI there is no intentionality or consciousness that is required for understanding (Searle, 1980).

The failure of AI in language understanding also has the same implications of the situatedness of human cognition. Language is not just a code that is used by people; rather it is embedded in the body, in relations and in culture. This is the embodied cognition theory that argues that meaning is grounded in bodily experience, something that is missing in AI since it does not have a body or any experiences.

Hence, due to its inability to physically and socially engage with the environment in the same way as human beings, there are major gaps in AI's understanding of language and its use (Lakoff & Johnson, 1999). AI is not situated and thus cannot act on the physical environment or perceive the world in real time, which poses a problem for AI to handle idioms or other contextually grounded language use (Barsalou, 2008).

The final aspect that makes human language processing different from AI is context. Individuals do not even blink an eye when they translate context in order to determine the meaning, to find out what is being referred to, and to guess what the speaker wants to do. Indeed, as Clark and Brennan (1991) have noted, context plays a central role in such things as the interpretation of homonyms or indirect speech acts. These tasks, however, are difficult to AI systems. Despite the advances in the application of context in natural language processing, there are still limitations. The examples of contextual awareness of AI prove that AI has issues with perceiving the context of human communication and thus distorts it.

For instance, the current systems like GPT-3, no matter how sophisticated they are, may produce responses that are syntactically related or even nonsensical in the light of an ambiguous context (Brown et al. , 2020). The issues that AI faces with the emotion and intuition in language also explain the distinctions between human and artificial intelligence. Communication is not just the exchange of information but it is the exchange of information that is influenced by feelings, hunches and culture that define how the message is given and how it is received. Even the attempts that AI makes to mimic such aspects of communication, for example, with the help of sentiment analysis or natural language generation, are generally quite convincing. The failure

to grasp the feelings and hunches of the communication process also pose a weakness of AI in that it cannot decode and participate in the communication process particularly in sensitive affairs. This gap demonstrates that it is very difficult to train machines to be capable of empathizing, reasoning, and being emotionally intelligent like human beings (Pang & Lee, 2008).

All these are philosophical limitations. The issues of meaning, reference, and communication that AI has are the reasons why it cannot mimic human understanding in any way. While AI can imitate some aspects of language use, it lacks intentionality, context awareness, and physical engagement, which are at the heart of understanding. This leads to a more general question of what language and mind are and suggests that human language is not just a computation problem but a human affair that involves physicality, sociality and culturality.

From an ethical point of view, the application of AI in communication raises questions on the following aspects; responsibility, transparency, and impact on the society. Since AI is gradually becoming a part of human interactions, the developers and policymakers should make sure that the technologies they develop and deploy are humane. This includes issues of privacy, neutrality, and impartiality of the AI systems and the capacity to articulate the advantages and disadvantages of the AI systems (Crawford & Calo, 2016). The ethical issues of AI in language processing also include the social impacts such as the reduction in face-to-face interpersonal communication and social interactions (Turkle, 2015).

Thus, it is possible to conclude that AI has advanced significantly in language processing, however, the existing limitations prove the depth of language and thought. The problems that AI faces when it tries to mimic human understanding of meaning, context, and feelings all indicate that human interaction is irreplaceable. Thus, as we advance in the process of improving and integrating AI in our daily lives, it is crucial to welcome these technologies with positive attitudes and positive outlooks while at the same time being mindful of the negative effects that come with it. Thus, it is possible to enhance and build the necessary and diverse elements of human relations through the application of artificial intelligence potential

References

Adamopoulou, E., & Moussiades, L. (2020). An Overview of Chatbot Technology. Artificial Intelligence Review, 53(1), 65-75.

Austin, J. L. (1962). How to Do Things with Words. Oxford University Press.

Barsalou, L. W. (2008). Grounded Cognition. Annual Review of Psychology, 59, 617-645.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610-623.

Brewster, S., McGookin, D., & Miller, C. (2003). Olfactory Interfaces. Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, 1(1), 654-661.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., & Amodei, D. (2020). Language Models are Few-Shot Learners. arXiv preprint arXiv:2005.14165.

Clark, H. H. (1996). Using Language. Cambridge University Press.

Clark, H. H., & Brennan, S. E. (1991). Grounding in Communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), Perspectives on Socially Shared Cognition (pp. 127-149). American Psychological Association.

Crawford, K., & Calo, R. (2016). There is a Blind Spot in AI Research. Nature, 538(7625), 311-313.

Davidson, D. (1967). Truth and Meaning. Synthese, 17(1), 304-323.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv preprint arXiv:1810.04805.

Dreyfus, H. L. (1972). What Computers Can't Do: A Critique of Artificial Reason. Harper & Row.

Ekman, P. (1992). An Argument for Basic Emotions. Cognition & Emotion, 6(3-4), 169-200.

Floridi, Luciano (2016). What is Data Ethics? Philosophical Transactions of

the Royal Society A, 374(2083), 20160360. https://doi.org/10.1098/rsta.2016.0112.

Gärdenfors, P. (2000). Conceptual Spaces: The Geometry of Thought. MIT Press.

Grice, H. P. (1975). Logic and Conversation. In P. Cole & J. L. Morgan (Eds.), Syntax and Semantics (Vol. 3, pp. 41-58). Academic Press.

Gruber, T. R. (1993). A Translation Approach to Portable Ontology Specifications. Knowledge Acquisition, 5(2), 199-220.

Hall, E. T. (1976). Beyond Culture. Anchor Books. Harnad, S. (1990). The Symbol Grounding Problem. Physica D: Nonlinear Phenomena, 42(1-3), 335-346.

Hobbs, J. R. (1978). Resolving Pronoun References. Lingua, 44(4), 311-338.

Hoy, M. B. (2018). Alexa, Siri, Cortana, and Google Assistant: A Comparison of Speech-Based Natural Language User Interfaces. Medical Reference Services Quarterly, 37(1), 81-88.

Jurafsky, D., & Martin, J. H. (2021). Speech and Language Processing (3rd ed.). Pearson.

Kripke, S. (1980). Naming and Necessity. Harvard University Press.

Lakoff, G., & Johnson, M. (1980). Metaphors We Live By. University of Chicago Press.

Lakoff, G., & Johnson, M. (1999). Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought. Basic Books.

Levinson, S. C. (1983). Pragmatics. Cambridge University Press.

Luger, E., & Sellen, A. (2016). "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 (pp. 5286-5297).

Luger, G. F. (2005). Artificial Intelligence: Structures and Strategies for Complex Problem Solving. Pearson Education.

Manning, C. D., Raghavan, P., & Schütze, H. (2020). Introduction to Information Retrieval (2nd ed.). Cambridge University Press.

Marcus, G., & Davis, E. (2020). Rebooting AI: Building Artificial

Intelligence We Can Trust. Pantheon Books.

McTear, M. F. (2016). The Rise of the Conversational Interface: A New Kid on the Block?. Interaction, 23(4), 38-43.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The Ethics of Algorithms: Mapping the Debate. Big Data & Society, 3(2), 2053951716679679.

Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, 2(1-2), 1-135.

Pustejovsky, J., & Stubbs, A. (2012). Natural Language Annotation for Machine Learning: A Guide to Corpus-Building for Applications. O'Reilly Media.

Russell, S., & Norvig, P. (2020). Artificial Intelligence: A Modern Approach (4th ed.). Pearson.

Searle, J. R. (1969). Speech Acts: An Essay in the Philosophy of Language. Cambridge University Press.

Searle, J. R. (1979). Expression and Meaning: Studies in the Theory of Speech Acts. Cambridge University Press.

Searle, J. R. (1980). Minds, Brains, and Programs. Behavioral and Brain Sciences, 3(3), 417-424.

Shum, H.-Y., He, X.-D., & Li, D. (2018). From Eliza to XiaoIce: Challenges and Opportunities with Social Chatbots. Frontiers of Information Technology & Electronic Engineering, 19(1), 10-26.

Sperber, D., & Wilson, D. (1995). Relevance: Communication and Cognition (2nd ed.). Wiley-Blackwell.

Tomasello, M. (2003). Constructing a Language: A Usage-Based Theory of Language Acquisition. Harvard University Press.

Turkle, S. (2015). Reclaiming Conversation: The Power of Talk in a Digital Age. Penguin Press.

Weizenbaum, J. (1966). ELIZA—A Computer Program for the Study of Natural Language Communication between Man and Machine. Communications of the ACM, 9(1), 36-45.

Winograd, T. (1972). Understanding Natural Language. Academic Press.

Wittgenstein, L. (1953). Philosophical Investigations. Blackwell.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., ... & Dean, J. (2016). Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. arXiv preprint arXiv:1609.08144.