International Journal of Modern Developments in Engineering and Science
Volume 5, Issue 2, February 2026
www.ijmdes.com | E-ISSN: 2583-3138 | RESAIM Publishing | www.resaim.com

Understanding Chatbots and Conversational AI

Abhisha Ganesh Whaval*
Department of Artificial Intelligence and Data Science, Datta Meghe College of Engineering, Mumbai, India

***Abstract*: Conversational AI, including chatbots and virtual assistants, has rapidly advanced with the rise of large language models (LLMs). This survey examines the historical evolution of chatbots from rule-based systems to retrieval-based and generative AI models, while dissecting key components such as NLU, dialogue management, and NLG. It also covers major applications in customer service, healthcare, and education, as well as the ethical and societal implications of AI-driven dialogue systems. Challenges such as bias, hallucination, privacy concerns, and job displacement are discussed along with future directions in this evolving field.**

***Keywords*: Attention Mechanism, BERT, Chatbots, Conversational AI, Deep Learning, Dialogue Management, Generative Models, Machine Learning, Natural Language Processing (NLP), Transformers.**

1. Introduction

Conversational Artificial Intelligence (AI) has become a key enabler of human–computer interaction, powering chatbots, virtual assistants, and dialogue systems across domains such as customer service, healthcare, and education. By allowing users to communicate in natural language, these systems provide intuitive and scalable alternatives to traditional interfaces.

The progression of chatbots reflects major advances in natural language processing and machine learning. From rule-based scripts to today's transformer-based large language models, conversational AI has grown more context-aware and human-like. At the same time, issues of bias, transparency, and privacy highlight the need for responsible development. This paper surveys the evolution, architecture, applications, and challenges of conversational AI, and complements the discussion with a data-driven comparative analysis to illustrate adoption trends, framework usage, and ethical considerations.

2. Literature Survey

Early chatbot research began with rule-based systems such as ELIZA (Weizenbaum, 1966), which demonstrated scripted interactions but lacked contextual depth. Retrieval-based approaches later improved relevance by matching queries to predefined responses, while intent classification and slot filling models advanced task-oriented dialogue.

The shift to neural models marked a breakthrough: Seq2Seq with attention (Sutskever et al., 2014; Bahdanau et al., 2015) enabled more coherent responses, while transformer architectures (Vaswani et al., 2017) powered models like BERT (Devlin et al., 2018) and GPT (Radford et al., 2019; OpenAI, 2023). These large language models set new benchmarks in fluency and contextual understanding. Complementary research, such as knowledge graph embeddings (Wang et al., 2021), has enhanced reasoning, and open-source frameworks like Rasa (Bocklisch et al., 2017) have made conversational AI widely accessible.

3. Evolution of Chatbots

The journey of chatbots across several decades reveals a fascinating progression, primarily driven by advancements in computational linguistics and artificial intelligence. This evolution can be broadly categorized into three primary phases, each overcoming limitations of its predecessor through innovative technical paradigms.

1) Rule-Based Chatbots: The Dawn of Dialogue Simulation

The earliest conversational systems, known as rule-based chatbots, operated on a foundation of meticulously crafted scripts and decision trees. Their functionality was entirely dictated by predefined if-then logic, which required exact keyword matching to trigger a response. This inherent rigidity meant they were highly inflexible to linguistic variations, synonyms, or nuanced user expressions, often leading to conversational breakdowns if input deviated from their programmed pathways. A seminal example is ELIZA, developed by Joseph Weizenbaum in the 1960s, which simulated a Rogerian psychotherapist by rephrasing user statements into questions, creating an illusion of understanding without proper comprehension. While pioneering, these systems were limited by the sheer effort of manually encoding every possible interaction and their inability to learn or adapt.

2) Retrieval-Based Chatbots: Embracing Machine Learning for Better Matching

The limitations of purely rule-based systems paved the way for retrieval-based chatbots, which marked a significant leap forward by incorporating data-driven approaches. Instead of generating responses from scratch, these bots leverage a vast database of pre-written replies. Their sophistication stems from the application of machine learning techniques, particularly in developing sophisticated similarity metrics to match user input to the most appropriate canned response from their repos- itory. Techniques such as TF-IDF (Term Frequency-Inverse Document Frequency) or word embeddings (like Word2Vec) are widely used. The pervasive presence of conversational AI systems across diverse sectors highlights their transformative

*Corresponding author: abhishawhaval@gmail.com

potential, streamlining operations, enhancing user experiences, and opening up entirely new modes of interaction.

3) Generative Chatbots: The Transformative Era of Large Language Models (LLMs)

The current paradigm in chatbot evolution is dominated by generative models, representing a profound shift in capability. These systems harness the power of deep neural networks, especially the transformer architecture, to synthesize entirely novel and contextually relevant responses. Unlike retrieval-based systems, generative chatbots do not simply select from a database; they ”create” human-like text on the fly. Ground-breaking models, such as OpenAI’s ChatGPT and Google’s Bard, exemplify this category. Their ability to adapt to diverse conversational contexts and produce fluid, coherent prose has revolutionized human-machine interaction. This capacity stems from their training on enormous datasets, allowing them to learn intricate patterns of language, grammar, and even world knowledge. However, this impressive generative power comes with its own set of challenges, including the propensity for ”hallucinations”—the generation of factually incorrect or nonsensical information—and the potential to perpetuate biases present in their vast training data. Addressing these limitations remains a critical area of ongoing research.

4. Core Components of Conversational AI

Effective human-machine interaction through natural language is underpinned by three interdependent core components within conversational AI systems: Natural Language Understanding (NLU), Dialogue Management, and Natural Language Generation (NLG). Each plays a distinct yet crucial role in processing user input, maintaining conversational flow, and formulating appropriate responses.

A. Natural Language Understanding (NLU): Deciphering Human Intent

Natural Language Understanding (NLU) is the foundational process through which a machine interprets and derives meaning from human linguistic input. This goes beyond mere keyword matching, aiming to grasp the user’s underlying intent and extract relevant information. Key tasks within NLU include:

1. *Intent Recognition*: Identifying the primary goal or purpose behind a user’s utterance (e.g., ”book a flight,” ”check account balance,” ”play music”). This often in- volves classification models trained on labelled datasets, utilizing techniques such as Support Vector Machines (SVMs), Convolutional Neural Networks (CNNs), or more recently, transformer-based models for superior performance.
2. *Entity Extraction (Named Entity Recognition - NER)*: Identifying and classifying key pieces of information within the text, such as dates, times, locations, product names, or person names (e.g., extracting ”next Monday” as a date or ”New York” as a location). This typically involves sequence labeling models, such as Conditional Random Fields (CRFs) or Recurrent Neural Networks (RNNs), with attention mechanisms.
3. *Sentiment Analysis*: Determining the emotional tone or attitude expressed in the user’s input (e.g., positive, negative, neutral). This provides critical context, especially in customer service applications, allowing the system to adapt its response or escalate complicated interactions. Accurate and robust NLU is paramount, as any misinterpretation at this stage propagates errors throughout the entire conversational pipeline, severely compromising the relevance and effectiveness of the chatbots subsequent responses.

B. Dialogue Management: Orchestrating the Conversation Flow

Dialogue management is the central intelligence that governs the progression and coherence of a conversation. It maintains the conversational state, tracks context across multiple turns, and decides the following appropriate action or response. This component ensures that interactions are not merely a series of isolated exchanges, but a logical and contextually aware dialogue. Common approaches include:

1. *State-based Dialogue Managers*: These systems rely on predefined finite state machines, where each state represents a point in the conversation, and transitions between states are triggered by recognized intents and extracted entities. They are predictable but can be less flexible for complex, open-ended dialogues.
2. *Frame-based Dialogue Managers*: Here, the system attempts to fill ”slots” within a predefined ”frame” for a specific task (e.g., for booking a flight, slots might include ’destination’, ’date’, ’number of passengers’). The manager prompts the user for missing information until all slots are filled.
3. *Reinforcement Learning (RL) based Dialogue Managers*: More advanced systems, particularly for complex and adaptive dialogues, use reinforcement learning. The dialogue manager learns optimal policies for interacting with users by receiving rewards for successful task completion and penalties for errors. This enables the system to develop more flexible and robust conversational strategies over time, often by leveraging techniques such as Deep Q-Networks (DQNs) or Policy Gradients. Rasa, an open- source framework, integrates sophisticated dialogue management capabilities that can learn from conversational data.

C. Natural Language Generation (NLG): Crafting Human-Like Responses

Natural Language Generation (NLG) is the crucial final step in the conversational AI pipeline, responsible for transforming structured data, system intents, or abstract representations into coherent, grammatically correct, and human-readable text. The quality of NLG has a direct impact on the user experience and the perceived intelligence of the chatbot. Early NLG systems often relied on templates or rule-based concatenation of pre-written phrases. However, modern generative chatbots, powered by transformer architectures like those in GPT-4, utilize sophisticated neural networks to generate novel text.

5. Applications of Conversational AI

The pervasive presence of conversational AI systems across diverse sectors highlights their transformative potential, streamlining operations, enhancing user experiences, and opening up entirely new modes of interaction.

1) Customer Service: The Frontline of Automation

Customer service is the most common use of conversational AI, handling routine queries, bookings, and troubleshooting while reducing costs and improving satisfaction through 24/7 support. The main challenge is smoothly escalating complex or emotional cases to human agents without frustrating users.

2) Healthcare: Enhancing Accessibility and Support

In healthcare, conversational AI improves patient engagement, scheduling, reminders, and health guidance, with tools like Woebot (mental health support) and Ada (symptom assessment). While it expands access to care, especially in underserved areas, challenges around accuracy, privacy, and regulatory compliance remain critical.

3) Education: Personalizing Learning Journeys

Educational chatbots are revolutionizing learning by providing personalized tutoring, offering assistance with quizzes, and supporting language acquisition. AI-powered platforms are uniquely capable of adapting to individual learner levels, identifying knowledge gaps, and tailoring content delivery to optimize comprehension. This fosters a more accessible and individualized educational experience, moving beyond the traditional one-size-fits-all model. The challenge lies in designing systems that offer a profound conceptual understanding rather than just factual recall, and in ensuring equitable access to these advanced tools, thereby bridging the digital divide.

4) Entertainment: Immersive and Interactive Experiences

The entertainment sector is leveraging chatbots to create novel and engaging user experiences. These systems can simulate conversations with fictional characters, enhancing storytelling in interactive narratives, or act as helpful guides and companions within gaming environments. The ability to generate dynamic, unique dialogues adds a layer of immersion and creativity, opening up possibilities for personalized plotlines and character interactions. As these systems become more sophisticated, they blur the lines between passive consumption and active participation, offering users unprecedented ways to engage with digital content.

5) Bias and Fairness: Confronting Algorithmic Inequality

AI models risk perpetuating societal biases embedded in training data, leading to unfair or discriminatory outputs against certain groups. For example, systems trained on biased job application data may discourage specific genders or ethnicities from career paths. Addressing this requires diverse datasets, careful model tuning, adversarial training, and continuous auditing to ensure algorithmic fairness.

6) Hallucinations: The Peril of Fabricated Information

A major challenge in generative models is hallucination, where AI produces confident but incorrect or fabricated information. This undermines trust, especially in critical domains like healthcare or finance, as models prioritize plausible text over factual accuracy. Research focuses on grounding responses in external knowledge and improving uncertainty handling to reduce this issue.

7) Privacy Concerns: Safeguarding Sensitive User Data

Conversational AI often handles sensitive data, raising concerns about privacy, security, and consent. Risks include breaches or misuse of personal information, prompting regulations like GDPR.

8) Job Displacement: Navigating Economic and Social Shifts

The rise of conversational AI raises concerns about job displacement in roles like customer support and content creation, even as it boosts efficiency. Yet, technological shifts also create opportunities in areas such as AI training, monitoring, and ethical development. Preparing the workforce through reskilling and upskilling is vital to ensure a fair transition.

Table 1
NLU techniques comparison

Technique	Model Type	Use Case	Advantages
TF-IDF + Logistic Regression	Statistical	Simple intent detection	Fast, interpretable
SVM	Supervised ML	Intent sorting	Accurate on small datasets
CRF	Probabilistic	Entity extraction	Handles sequence labeling well
LSTM	Deep learning	Context aware NLU	Remembers sequence order
BERT	Transformer Based	Contextual NLU	State of the art accuracy

Table 2
Evolution metrics

Metrics	Measures	Used for	Strengths
Accuracy	Correct prediction/ total	Intent classification	Easy to interpret
Precision/ Recall	True +ve false -ve	NER, intent detection	Balances false positive/ negatives
F1-Score	Harmonic mean of P and R	NER, classification task	Useful with imbalanced classes
BLEU	Overlap reference responses	NLG evaluations	hallucinate, resource intensive
Human Evaluation	Coherence, engagement, empathy	NLG, overall conversation	Most reliable

Table 3
Use Cases V/S Chatbots Types

Application Area	Chatbot Type	Why?
Customer Support	Retrieval/Hybrid Based	Fast, reliable canned responses
Healthcare	Generative + Rule Based	Needs empathy + accuracy fallback
Education	Generative Based	Personalized, contextual responses

6. Data-Driven Comparative Analysis Using Power BI

A. Chatbot Adoption Over Time

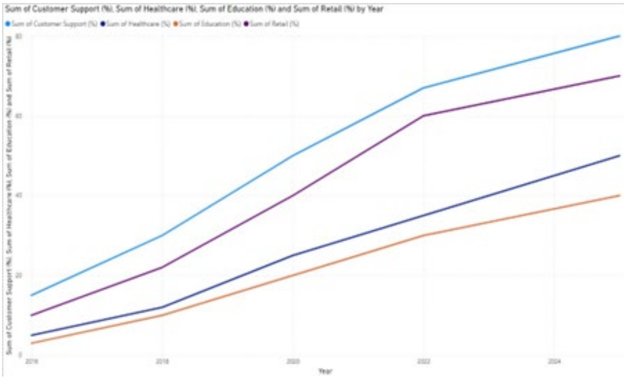

Fig. 1. Chatbot adoption trends across industries (2016–2025)

The line chart illustrates the growth in chatbot adoption across four key industries: Customer Support, Retail, Healthcare, and Education, from 2016 to 2025. Customer Support leads the trend, showing rapid growth from 15 percent to 80 percent, driven by automation of FAQs and live chat support. Retail follows closely, fueled by order tracking and product recommendation bots. Healthcare and Education, though slower in adoption, show steady increases as use cases like symptom checkers and virtual tutoring become mainstream. The consistent upward trend across all sectors underscores the expanding role of chatbots in digital transformation across domains.

B. Chatbot Use Cases by Industry

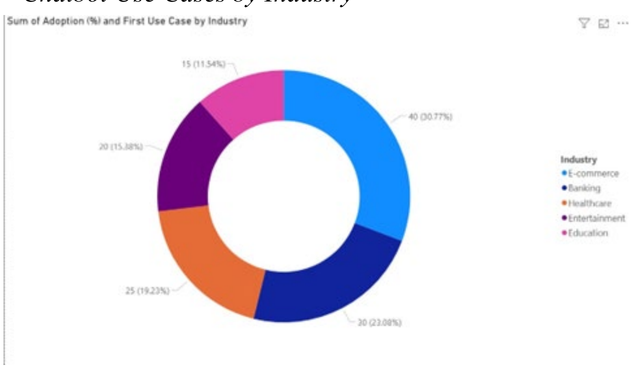

Fig. 2. Chatbot use case distribution across industries

The donut chart highlights the distribution of chatbot adoption across various industries based on primary use cases. E- commerce leads with 30.77 percent, driven by use cases like order tracking, product recommendation, and customer interaction. Banking follows at 23.08 percent, utilizing chatbots for balance inquiries, fraud alerts, and customer queries. Healthcare holds 19.23 percent, focusing on symptom checkers and appointment scheduling, while Entertainment (15.38 percent) uses bots for content suggestions and fan engagement. Education (11.54 percent), though smaller in share, is growing with virtual tutors and learning assistants. This chart emphasizes how sector-specific needs shape chatbot deployment.

C. Chatbot Framework Usage

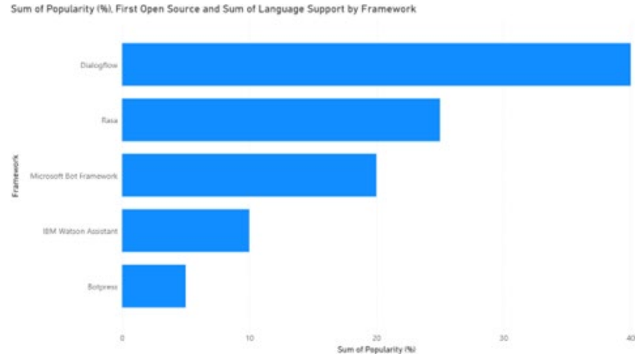

Fig. 3. Popularity of chatbot development frameworks

The bar chart compares the popularity of major chatbot frameworks based on industry usage. Dialogflow ranks highest with around 40 percent, largely due to its seamless integration with Google Cloud and ease of use. Rasa follows with strong adoption (approx 25 percent), favored for its open- source flexibility and full control over deployment. Microsoft Bot Framework maintains moderate popularity (approx 20 percent), particularly among enterprises using Azure and .NET environments. IBM Watson Assistant and Botpress trail with lower adoption, despite offering strong analytics and visual flow design, respectively. This distribution reflects developers' priorities—balancing ease, integration, and customization.

D. Chatbot Model Performance

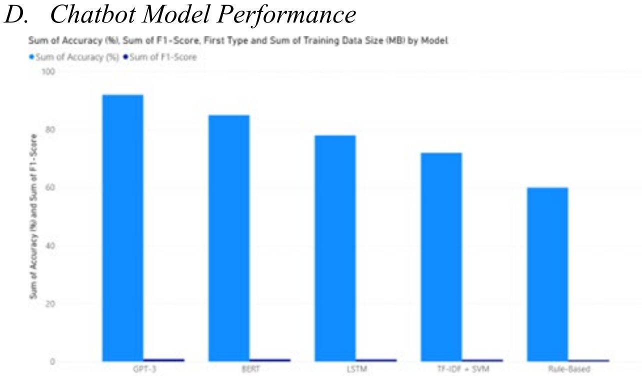

Fig. 4. Performance comparison of chatbot models (Accuracy and F1-Score)

The column chart presents a comparative analysis of popular chatbot models based on accuracy and F1-score. GPT-3 leads with the highest performance, achieving around 95 percent accuracy, followed closely by BERT, a transformer-based model. LSTM models perform moderately with approx. 78 percent accuracy, benefiting from sequential learning capabilities. Traditional models like TF-IDF + SVM and Rule-Based systems show lower effectiveness, highlighting the limitations of non-contextual approaches. This comparison emphasizes the significant performance leap achieved through deep learning and transformer architectures in modern conversational AI systems.

E. Ethical Concerns by Model

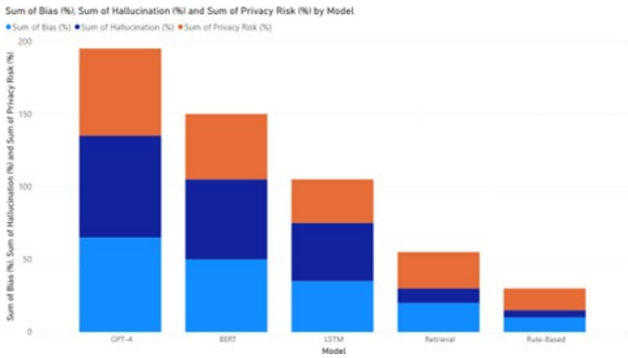

Fig. 5. Ethical risk comparison of chatbot models (Bias, Hallucination, and Privacy risk)

This stacked bar chart evaluates chatbot models based on three critical ethical dimensions: bias, hallucination, and privacy risk. GPT-4 shows the highest cumulative risk, with significant levels of hallucination and bias due to its generative nature. BERT and LSTM also exhibit moderate levels, though BERT faces notable hallucination issues. The graph emphasizes the trade-off between advanced model capabilities and the ethical challenges they introduce in real-world deployments.

7. Future Directions and Research Opportunities

The field of conversational AI is rapidly evolving, and future research is expected to overcome current limitations and enable more natural, accessible, and capable dialogue systems.

A. Enhancing Long-Term Contextual Understanding and Memory

Current conversational agents often struggle to maintain coherence and relevance across extended dialogues, frequently" forgetting" earlier turns or previously provided information. Future research will focus on developing advanced mem- ory mechanisms and contextual awareness models that can track and recall information over many conversational turns, enabling more natural, sustained, and coherent interactions. This involves exploring novel neural architectures and more efficient methods for representing and retrieving dialogue history.

B. Deeper Personalization and Emotional Intelligence

Moving beyond generic responses, a key frontier involves tailoring conversational experiences to individual user preferences, communication styles, and even emotional states. This necessitates research into more sophisticated user profiling adaptive response generation, and the ability to detect and respond appropriately to human emotions (e.g., frustration, confusion, joy). Such emotionally intelligent AI could lead to more empathetic and supportive interactions, particularly in sensitive domains like mental health support or education.

C. Grounding Responses in Factual Knowledge and Minimizing Hallucinations

To combat the issue of factual inaccuracies or 'hallucinations,' a significant research thrust is dedicated to grounding generative models in verifiable, external knowledge bases. This involves integrating Large Language Models (LLMs) with structured knowledge graphs or real-time access to reliable data sources, ensuring that generated responses are not only fluent but also factually accurate. Techniques like retrieval- augmented generation (RAG) are gaining prominence, where the model first retrieves relevant information before generating a response based on that retrieved context.

D. Multimodal Conversational AI

The future of human-machine interaction is inherently multimodal. Research is increasingly exploring systems that can seamlessly combine and interpret various forms of input—text, speech, images, and video—and generate responses in equally rich modalities. Imagine a chatbot that can not only understand your spoken question but also analyze a screenshot you've provided, or respond with a relevant image or video clip. This convergence will lead to far richer, more intuitive, and human- like interactions.

E. Bridging the Linguistic Divide: Support for Low-Resource Languages and Dialects

Ensuring equitable access to AI technology globally requires significant efforts to enhance support for low-resource languages and diverse dialects. Many of current state- of-the-art models are heavily biased towards high-resource languages, such as English, due to data availability. Future research will concentrate on developing robust techniques for training and fine-tuning models with limited data, leveraging cross-lingual transfer learning, and creating more inclusive language technologies.

F. Advancements in Privacy-Preserving AI

With the increasing amount of sensitive data processed by conversational systems, research into privacy-preserving techniques is paramount. Approaches like federated learning, which allows models to be trained on decentralized data without sharing the raw information, and edge computing, which processes data closer to its source, are expected to play a larger role in protecting user privacy while still enabling robust AI development.

G. Real-time Learning, Reasoning, and Proactive Adaptation

The ultimate goal for many advanced AI systems is to move beyond mere reactivity. Future conversational agents are envisioned to possess capabilities for real-time learning from new interactions, complex reasoning, and proactive behavior.

This means not just responding to explicit commands, but also anticipating user needs, offering relevant suggestions, and adapting their conversational strategy dynamically in diverse and unpredictable real-world scenarios, thereby transforming them from tools into genuine collaborative partners.

8. Conclusion

This paper surveyed chatbots and conversational AI, covering their evolution, architectures, applications, and challenges. From rule-based systems to transformer-based models, the field has advanced toward more context-aware and human-like interactions. Our comparative analysis, supported by Power BI visualizations, highlighted adoption trends, framework usage, model performance, and ethical concerns.

While models like BERT and GPT set new performance standards, they also raise issues of bias, hallucination, and privacy. Future progress will depend on addressing these risks while exploring multimodal systems, personalization, and privacy-preserving AI. With balanced innovation and ethical safeguards, conversational AI can develop into reliable, human-centered technology.

References

[1] J. Weizenbaum, "ELIZA—A computer program for the study of natural language communication between man and machine," *Communications of the ACM*, vol. 9, no. 1, pp. 36–45, Jan. 1966.

[2] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems (NeurIPS 2014)*, 2014, pp. 3104–3112.

[3] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.

[4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS 2017)*, 2017, pp. 5998–6008.

[5] A. Ram, R. Prasad, C. Khatri, A. Venkatesh, R. Gabriel, and J. Nunn, "Conversational AI: The science behind the Alexa Prize," Amazon.com, Inc., Technical Report, 2018.

[6] T. Bocklisch, J. Faulkner, N. Pawlowski, and A. Nichol, "Rasa: Open source language understanding and dialogue management," in *Proc. 16th Conf. European Chapter of the Association for Computational Linguistics (EACL), System Demonstrations*, 2017, pp. 50–54.

[7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2019, pp. 4171–4186.

[8] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac *et al.*, "Transformers: State-of-the-art natural language processing," in *Proc. 2020 Conf. on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP)*, 2020, pp. 38–45.

[9] OpenAI, "GPT-4 technical report," Mar. 2023.

[10] B. Liu and I. Lane, "Attention-based recurrent neural network models for joint intent detection and slot filling," in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 685–689.