

CHATBOT: DESIGN, ARCHITECTURE, AND APPLICATIONS

By Xufei Huang

Advisor: Professor Mitch Marcus

Course Coordinator: Professor Norm Badler

ASCS CIS 498 Senior Capstone Thesis

University of Pennsylvania | School of Engineering and Applied Science

May 3, 2021

Abstract

A chatbot, also known as a dialogue system or a conversational agent, is a computer program that can imitate a conversation with a user [61]. In the past decade, chatbot technology has evolved tremendously, thanks to the growing popularity of artificial intelligence and machine learning. Chatbots are proving themselves as a valuable tool in many situations and are visible in almost every aspect of our daily lives. This paper intends to address the design, architecture, and applications of chatbots. We will discuss the evolution of chatbots, present a technical overview of the chatbot system and the technologies that support it, and address the applications and potential implications of chatbots on the wider world.

TABLE OF CONTENTS

ABSTRACT	1
1 INTRODUCTION.....	3
2 OVERVIEW OF CHATBOTS	4
2.1 THE RISE OF CHATBOTS	4
2.2 CLASSIFICATION OF CHATBOTS	5
3 DESIGN PRINCIPLES	7
3.1 RULE BASED	7
3.2 MACHINE LEARNING	8
4 ARCHITECTURE	9
4.1 USER INTERFACE	9
4.2 NATURAL LANGUAGE UNDERSTANDING	10
4.2.1 Slot Filling.....	10
4.2.2 Statistical Approaches to NLU.....	12
4.3 DIALOGUE MANAGEMENT	14
4.3.1 Design Issues.....	14
4.3.2 Hand-crafted Approaches to DM.....	16
4.3.3 Statistical Approaches to DM	16
4.4 BACKEND	18
4.5 RESPONSE GENERATION	19
4.5.1 The RG Pipeline	19
4.5.2 Statistical Approaches to RG	20
5 APPLICATIONS AND FUTURE OUTLOOK.....	24
5.1 MICROSOFT XIAOICE	24
5.1.1 Design Principle.....	25
5.1.2 Architecture.....	25
5.2 GOOGLE MEENA	28
5.2.1 Human Likeness	28
5.2.2 Training Data.....	28
5.2.3 Model Architecture.....	29
5.3 OTHER APPLICATIONS.....	29
5.4 IMPLICATIONS AND FUTURE OUTLOOK	30
6 CONCLUSION	32
REFERENCES	33

1 Introduction

A dialogue system is an example of an intelligent agent that intends to hold natural conversations with humans. More specifically, it is a software application, with the help of natural language processing and machine learning, that stimulates human conversation in natural language via text or text-to-speech [52].

Dialogue systems have come a long way since their inception in the 1960s. Both the hardware and software aspects of computer science, and natural language processing and machine learning techniques have developed tremendously. Thanks to the advancement of these emerging technologies, chatbots have evolved from systems that generate machine-like responses to human-like agents capable of developing long-term relationships with users. Among the most famous early chatbot implementations are ELIZA and PARRY [111]. Modern chatbots include Apple's Siri and Amazon's Alexa, and Microsoft's Xiaolce [95].

Chatbots promise a range of potential benefits. Most notably, they provide responses and solutions that are instant, consistent, and reliable [53]. These characteristics make chatbots a powerful tool in many different areas, and their use has evolved rapidly in fields such as business, e-commerce, and healthcare [5].

This paper is a literature review on chatbot technology. Section 2 will provide a brief overview and history of chatbots. Section 3 will introduce two design approaches to chatbots. In Section 4, we will delve into the general architecture of chatbots and describe the technologies that support each component in detail. Finally, we will address existing chatbot applications and social and ethical considerations in Section 5.

2 Overview of Chatbots

2.1 The Rise of Chatbots

The idea of a chatbot was inspired by Alan Turing. In 1950, Turing proposed the Turing Test and asked whether machines can think. Since then, conversational systems have attracted extensive attention and become an important concept in artificial intelligence [99].

The first publicly known chatbot was **ELIZA**. ELIZA was developed in 1966 at the MIT Artificial Intelligent Laboratory by Joseph Weizenbaum. It simulates conversations based on hand-craft scripts that mimic a Rogerian psychotherapist. When a user chats with ELIZA, he or she types some statements in natural language. ELIZA analyzes the input text and looks for the presence of a keyword, and then generates responses according to a rule associated with the keyword. While it is able to engage in conversations, its communication ability and knowledge scope are limited. Below shows a typical conversation between ELIZA and a human [101].

User: Men are all alike.

ELIZA: In what way?

User: They're always bugging us about something or other.

ELIZA: Can you think of a specific example?

User: Well, my boyfriend made me come here.

ELIZA: Your boyfriend made you come here.

User: He says I'm depressed much of the time.

ELIZA: I am sorry to hear you are depressed [101].

Other well-known early chatbots include **PARRY** (developed by Colby in 1975) and **ALICE**. (developed by Wallace in 2009). Similar to ELIZA, they simulate how a human would behave in a text-based conversation using simple pattern-matching algorithms. PARRY is considered an improvement of ELIZA as it has a personality and a better controlling structure [20]. The creation of ALICE was another step forward in the history of chatbots [100]. It was the first online chatbot and was awarded for the best human-like system [14]. ALICE pattern matches against 41,000 templates to generate responses, and in comparison, ELIZA contains only 200 keywords and rules [39].

In 2001, the development of **SmarterChild** marked an important advancement in chatbot technology. SmarterChild was available through messenger applications and was able to assist people with simple tasks using information retrieved from databases [69].

In 2003, a project called **CALO (Cognitive Assistant that Learns and Organizes)**, funded by Defense Advanced Research Projects Agency (DARPA) and coordinated by SIR International started [70]. It was a five-year project that aimed to create a cognitive assistant that can learn from its experience and perform routine tasks for its users. The CALO project was very important in the history of chatbots as it integrated many areas of artificial intelligence and it helped software systems to better understand the intentions of humans. Moreover, it had many spin-offs, most notably Apple Siri [18].

Apple Siri, created in 2010, was the first **virtual personal assistant**. It was followed by other assistants such as IBM Watson [47], Microsoft Cortana [76], Amazon Alexa [102], and Google Assistant [31]. The virtual personal assistants were integrated into smartphones or smart speakers and could understand human speech, respond via voices, and handle more advanced tasks. Unlike earlier systems, virtual personal assistants are connected to the internet and are thus able to generate responses very quickly [44]. However, misunderstandings happen very often as they cannot understand colloquial languages and cannot interpret the input within the dialogue context [6].

Another chatbot worth our attention is **Microsoft's XiaoIce**. It is a social chatbot, and it demonstrates that conversational agents can not only conduct conversations and perform simple tasks but also satisfy our need for sociability [113].

More recently, there was another breakthrough in chatbot technology with the advancement of **open-domain chatbots**. Google's Meena [7] and Facebook's Blender [92] are dialogue systems that can chat about virtually anything and can achieve close-to-human-level performance in many aspects.

2.2 Classification of Chatbots

Chatbots can be classified based on different parameters, including the service provided, the knowledge domain, the response generation method, the goal, and the permission, and the amount of human-aid.

Classification based on the **service provided** considers the sentimental proximity of the chatbot to the user, the amount of intimate interaction that takes place, and the task the chatbot is performing. Interpersonal chatbots provide services like booking services or searches in FAQ without being a companion of the user. Intrapersonal chatbots live in the personal domain of the user and are expected to understand the user like a human does. Inter-agent chatbots communicate with other chatbots to accomplish a task [71].

The **knowledge domain** refers to the knowledge the chatbots access or the amount of the data they are trained upon. Open domain chatbots can respond to questions from any domain, whereas closed domain chatbots only has knowledge in a particular domain [71].

Classification based on the **response generation** method takes into account the method of generating responses. A chatbot can be classified as a rule-based, retrieval-based, or generative-based chatbot, and we will discuss this in more detail later in the paper [71].

Classification based on the **goals** considers the primary goal a chatbot aims to achieve. Information chatbots provide the user with specific information stored in a fixed source. Chat-based or conversational chatbots are designed to hold a natural conversation with the user like another human being would do. Task-based chatbots perform a task based on the user's requests [71].

Chatbots can also be classified according to the **permissions** provided by the development platforms. Open-source platforms make their code available so that developers have control over most aspects of implementation. Closed platforms are of proprietary code and are typically offered by large companies [6].

Finally, depending on the **amount of human-aid**, chatbots can be classified into human-mediated or fully autonomous chatbots. A human-mediated chatbot uses human computation in at least one

component. Fully automated chatbots are fast in processing information while human computation provides flexibility and can help fill the gaps caused by the limitations of algorithms [55].

The figure below summarizes the classification criteria we discussed above.

Chatbot Categories	Knowledge domain	Generic
		Open Domain
		Closed Domain
	Service provided	Interpersonal
		Intrapersonal
		Inter-agent
	Goals	Informative
		Chat based/Conversational
		Task based
	Response Generation Method	Rule based
		Retrieval based
		Generative
	Human-aid	Human-mediated
		Autonomous
	Permissions	Open-source
		Commercial

Figure 1: Classification of chatbots [6]

3 Design Principles

There are two approaches that can be used to develop a chatbot depending on the algorithms and techniques adopted: rule-based approach and machine learning approach.

3.1 Rule-based

A rule-based chatbot processes information and provides responses based on a set of predefined rules with the use of **pattern matching** algorithms. Although the pattern matching techniques vary in complexity, the basic idea is the same. The user input is classified as a pattern, and the chatbot selects a predefined answer by matching the pattern with a set of stored responses. The pattern and response matching algorithms are handcrafted [65]. Pattern matching is adopted by many chatbots and is especially popular among the early chatbots like ELIZA, PARRY, and ALICE. The advantage of the rule-based approach is its speed as it does not require any deep analysis of the input text [48]. However, the responses are repeated and lack flexibility and originality as the knowledge is set by the developer in advance [78]. The following paragraphs will provide an overview of the three most commonly used languages for the implementation of rule-based chatbots.

Artificial Intelligence Mark-up Language (AIML) is a derivation of Extensible Mark-up Language (XML) [65]. ALICE was the first chatbot implemented in the AIML language [10]. AIML has a class of data object called an AIML objects, and these objects are responsible for modeling conversational pattern. Each object consists of two units called topics and categories. A topic is an optional top-level element that has a name attribute and a set of categories associated with it. Categories are the most basic unit of knowledge and are the rules of the chatbot. Each category consists of two elements called pattern and template. The pattern matches against the input from the user to the template that contains the response of the chatbot [65]. AIML is simple, flexible, and highly maintainable, and thus is one of the most commonly adopted languages for chatbot development [10]. The biggest disadvantages of AIML are that the developer must specify a pattern for every possible input of the user and that it is inefficient when the knowledge base is large [97]. The structure of an AIML object is demonstrated as below:

```
<category>
    <pattern>User Input</pattern>
    <template>
        Corresponding Response to Input
    </template>
</category> [65].
```

RiveScript is a line-based scripting language that can be used to implement the Knowledge Base [12]. Compared to AIML, RiveScript has more built-in features and tags, which means that the writer does not need to specify information about the chatbot in the additional configuration files [6].

ChatScript is an open-source language for developing rule-based chatbots. It matches user inputs to outputs using rules created by human writers in program scripts through a process called dialog flow scripting [104]. ChatScript uses concepts that are set of words that have the same meaning.

It consists of 2000 predefined concepts and scripters can also write their own concepts easily [46]. Compared to AIML and RiveScript, ChatScript is a harder language, but it allows developers to combine rules in more complicated ways [6].

3.2 Machine Learning

The recent advancement in machine learning has made it possible to develop more intelligent chatbots. Chatbots that adopt machine learning approaches use machine learning algorithms to extract information and generate responses and are able to improve through previous conversations. An extensive training set is required for machine-learning-based chatbots [62]. Two types of models can be used, retrieval or generative. Retrieval-based models involve choosing the optimal response from a set of responses [105], and generative models, on the other hand, use deep learning techniques to generate the response [40].

4 Architecture

In this section, we describe a general architectural design and delve deep into the important parts of each component. A general chatbot architecture consists of five main components, a User Interface component, a Natural Language Understanding (NLU) component, a Dialogue Management (DM) component, a Backend component, and a Response Generation (RG) Component [6], as illustrated in the figure below.

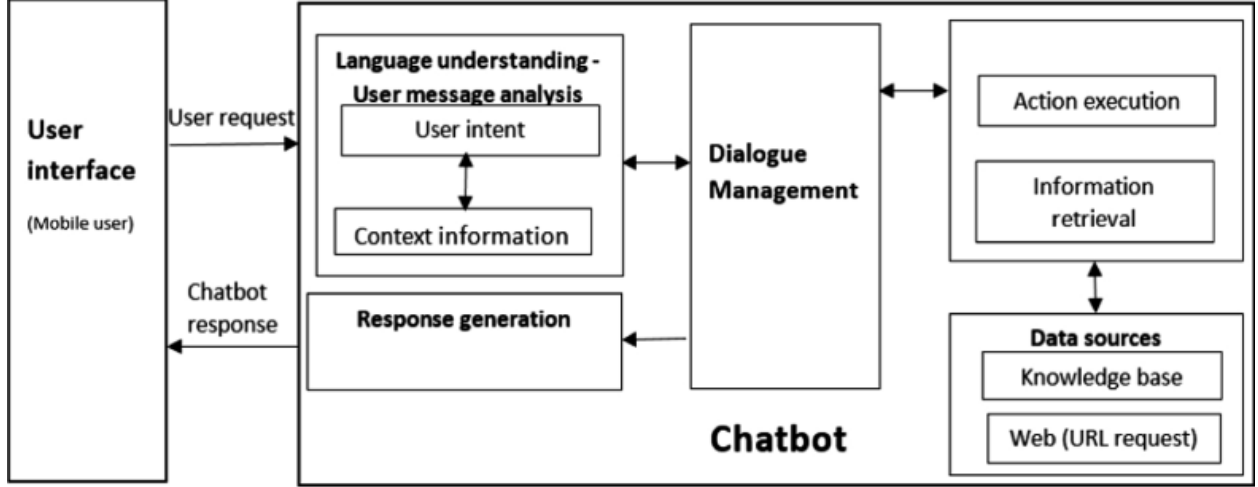


Figure 2: General chatbot architecture [6]

4.1 User Interface

The user interface allows users to communicate and interact with a chatbot through messenger applications like Facebook Messenger, Cortana, or Slack. The operation of a chatbot begins with a user's request [6].

For speech-based conversational agents, an **automatic speech recognition** (ASR) system will first transform the user's input into text. The process can be modeled as a stochastic process and the desired output is the most probable sequence of words W that corresponds to the user's input X [67]. Formally, this can be written as:

$$\begin{aligned}
 \hat{W} &= \underset{W}{\operatorname{argmax}} P(W|X) \\
 &= \underset{W}{\operatorname{argmax}} \frac{P(X|W)P(W)}{P(X)} \\
 &= \underset{W}{\operatorname{argmax}} P(X|W)P(W) \text{ [67]}.
 \end{aligned}$$

Here, $P(X|W)$ represents the acoustic model and $P(W)$ represents the language model. The acoustic model uses a corpus of phonetically transcribed and aligned data as training data. It creates the acoustic representation of the input, divided it into smaller frames, and then computes the likelihood of an input given a word. The language model contains information about which words are more probable in a given sequence. Finally, the output word sequence is determined in a process called decoding [67].

In addition, a speech-based conversational agent also has a **text-to-speech** (TTS) system that converts text back into speech after a text response is generated. TTS consists of two stages, text analysis, and waveform synthesis. Text analysis involves normalizing the text and performing phonetic (pronunciation) and prosodic (phrasing, pitch, loudness, tempo, and rhythm) analysis, and waveform synthesis involves selecting the prerecorded speech that meets the most requirements [67].

4.2 Natural Language Understanding

After the system receives the user’s request, it uses the NLU component to extract information from the input and produce a representation of its meaning that can be used later on in the process [82]. NLU generally deals with three tasks, dialogue act classification, intent classification, and slot filling [25].

Dialogue act classification deals with determining the function of the user’s input, or more precisely, mapping the user’s utterance to a dialogue act type. The utterance can be classified as a question, a statement, an offer, or some other type of dialogue act. Knowing the dialogue act being performed is critical to better comprehend the user’s request and to decide an appropriate response [98].

Intent classification identifies the primary goal of the user. Intents are mainly domain-dependent. For example, a request can be in the domain of food ordering, hotel reservations, weather forecasts, and so on. The intent of an agent within the hotel reservations domain can be to book, cancel or change a reservation, and similarly, the intent of an agent within the food ordering domain can be to place, query, or change an order [67].

Slot filling is the final step in NLU. The agent extracts other necessary details, which when combined with the dialogue act and the intent, allow it to fully understand the user’s request [67].

Examples in the table below illustrate the dialogue act classification and intent classification processes described above.

Input	Dialogue act	Intent
Set a timer for 10 minutes.	A request for action	Set a timer
Is there a direct flight from Philadelphia to Beijing?	A question	Look up flight information
I’ll be arriving at the hotel by 10pm tomorrow.	A statement	Check a hotel reservation

Table 1: Examples of dialogue act and intent classification

4.2.1 Slot Filling

Slot filling is the core of NLU and has received the most research attention. The main task of slot filling is to extract information from the input and to better understand its meaning. More specifically, slot filling aims to analyze the context of the input by separating the text into smaller units and then assigning the units tags or labels based on their functions and roles in the text. In this subsection, we will describe several different methods that are commonly used in slot filling.

Regular expression is a tool for describing text patterns. Each pattern is specified by a sequence of characters and each character is either a regular character or a metacharacter that has a special meaning [50].

Tokenization involves breaking up a text into units by words, punctuation marks, or numbers. Generally, English words are separated by white space, so tokenizing an English text should be fairly straightforward. However, there are some problematic cases where the boundaries of words or sentences are ambiguous. Contracted items (e.g. isn't), phrases (e.g. San Francisco), abbreviations (e.g. PhD.), and acronyms (e.g. AT&T) are examples of the special cases [67]. One commonly used tokenization standard is known as the Penn Treebank tokenization standard. According to the Penn Treebank tokenization standard, contractions and punctuation between tokens are separated out and hyphenated words are kept together [64]. A more advanced way to tokenizing text is to use data to determine the tokens automatically, and it is especially useful in dealing with unknown words. For example, bigger can be recognized even if only the words “big”, “small”, and “smaller” are contained in the training corpus. Usually, this is done by first inducing a vocabulary from a row training corpus and then segmenting the text into the tokens in the vocabulary [50]. The vocabulary is composed of subwords, for example, “small” and “er” instead of “smaller”. The three most widely used algorithms are: byte-pair encoding [86], unigram language modeling [56], and WordPiece [85].

Text normalization refers to converting the text into a standard form. Word normalization is the task of identifying words that have the same meaning but spelled differently (e.g. U.S.A. and USA). Case folding is the task of transforming everything to lower cases. Lemmatization is the process of identifying the roots of the words and mapping the morphological variants into their base form (e.g. produce, product, produces, production → produce) [50].

Bag of words, also known as the **vector space model**, is one of the simplest approaches to analyze the input. We count the occurrences of each word but ignore the syntactic information or word order information. This can be done by performing text normalization and eliminating stop words, i.e. words that do not contribute to the meaning of the sentence. The bag-of-words approach is limited when the meaning of the text is subject to linguistic knowledge such as grammar or word order. For example, “Alice is taller than Bob” and “Bob is taller than Alice” yield the same representation but their meanings are different [67].

Latent Semantic Analysis (LSA), or **Latent Semantic Indexing**, also does not take any linguistic knowledge into account. However, LSA compares the meaning behind the words instead of the actual words, and terms that occur frequently in the same context are grouped together. In order to determine word similarities, we create a matrix where each row represents a term and each column represents a document, and each cell in the matrix represents the frequency of the word in the document. Singular Value Decomposition is then applied to transform the matrix and reduce its dimensionality [67].

Part-of-speech (POS) tagging is the process of marking up each word in the text with a tag that indicates its syntactic role, for example, whether it is a verb, noun, pronoun, conjunction, and so on. POS tagging is especially useful in cases where we encounter words that have different meanings and their use in the sentence may be ambiguous. For example, the word “ring” can be either a noun or a verb, but in the sentence “ring a bell,” it should be categorized as a verb, and using a POS tag in cases like this can clear up the confusion [67]. The figure below illustrates the task of POS tagging.

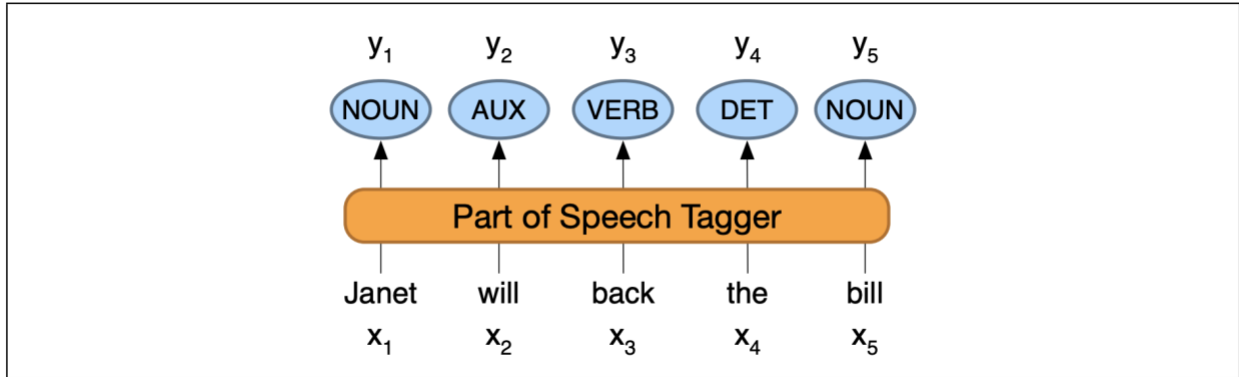


Figure 3: Part-of-speech tagging: mapping from words x_1, x_2, \dots, x_n to POS tags y_1, y_2, \dots, y_n [50]

Name entity recognition (NER) involves extracting items such as persons, dates, and organizations. POS tagging and NER are disambiguation tasks and provide useful information about sentence structure and meaning. NER is often more complicated than POS tagging because it involves determining the appropriate segmentation [50]. A number of datasets exist for POS tagging and NER. The Universal Dependencies [72] and the Penn Treebanks [64] have POS tagged corpora in different languages, and OntoNotes [43] has named entity tagged also in different languages.

Semantic role labeling deals with assigning the semantic roles to the arguments of a verb [30]. For example, in the sentence “book a flight to Philadelphia on Monday,” “book” is the verb, and its object or theme is “a flight,” “to Philadelphia” is a location, and “on Friday” is a date. One use case of semantic role labeling within the context of chatbot technology is to match the semantic structures in the input against similar structures in a collection of documents to find a potential response [67].

Once we extract the relevant parts from the user’s input, a complete analysis of the utterance can be performed with the use of grammars. The utterance can be interpreted directly with semantic grammar, or it can be parsed syntactically before the semantic analysis is applied. The semantic grammar approach uses grammar rules where the words are classified based on their semantic roles. In syntax-driven semantic analysis, the units of syntactic analysis are mapped to units of semantic analysis using one of the two grammar formalisms, context-free grammar and dependency grammar [67].

4.2.2 Statistical Approaches to NLU

Traditionally, the grammar rules are handcrafted, but over the past decades, more research has focused on assigning the labels automatically. In this section, we are going to look at how statistical models are used in the labeling problems, the process of assigning a label to each word in a text.

The **Hidden Markov Model (HMM)** is based on augmenting the Markov chain model. A Markov chain is a model that describes the probability of sequences of possible events, with the assumption that the future depends only on the state attained in the current event. An HMM is a Markov model where the states are not fully observable, and it is useful because the labels in a text are not observable. In the context of NLU, an HMM is a probabilistic sequence model that takes a sequence of words $w_1 \dots w_n$ as the input, computes the probability distribution over possible label sequences, and subsequently outputs the best sequence $t_1 \dots t_n$. An HMM has the following component: N label types or states, an initial distribution over states, a sequence of words or

observations, a transition probability matrix representing the probability of transitioning from one state to another, and emission probabilities describing the probability of an observation being generated from a state. To put this into context, the transition probability represents the likelihood of a label occurring given the previous label, and the emission probability represents the likelihood that a label is associated with a word when the label is given [50]. Formally, we capture the goal of HMM labeling in the equation below:

$$\hat{t}_{1,...,n} = \underset{t_1...t_n}{\operatorname{argmax}} P(t_1 \dots t_n | w_1 \dots w_n) [50].$$

HMM makes two assumptions. The first assumption is the Markov assumption, which says that the probability of a particular label depends only on the previous label. The second is that the probability of a word appearing depends only on its label and not on any other words or any other labels. Applying the two assumptions and the Bayes' rule, we can rewrite the equation as:

$$\begin{aligned} \hat{t}_{1,...,n} &= \underset{t_1...t_n}{\operatorname{argmax}} \frac{P(w_1 \dots w_n | t_1 \dots t_n) P(t_1 \dots t_n)}{P(w_1 \dots w_n)} \\ &= \underset{t_1...t_n}{\operatorname{argmax}} P(w_1 \dots w_n | t_1 \dots t_n) P(t_1 \dots t_n) \\ &\approx \underset{t_1...t_n}{\operatorname{argmax}} \prod_{i=1}^n P(w_i | t_i) P(t_i | t_{i-1}), \end{aligned}$$

which corresponds to the products of the emission and transition probabilities. The Viterbi algorithm can be used to find the solution, and the forward-backward or Baum-Welch algorithm can be used to train the model [50].

Figure 4 illustrates the transition probabilities A and the emission probabilities B for three labels VB, MD, and NN in an HMM model.

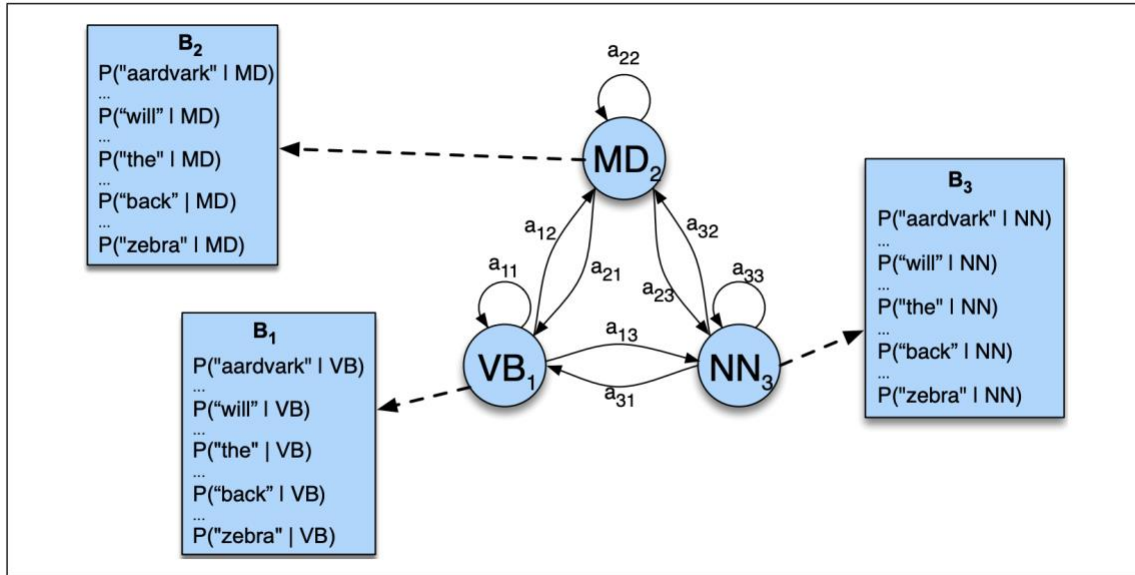


Figure 4: Parts of an HMM representation [50]

Stochastic Finite State Transducers (SFST) considers NLU as a translation process where stochastic language models are implemented using Finite State Machines (FSM). The FSMs take a sequence of words W as input and output a sequence of labels T , and they can be either handcrafted or learned from an annotated corpus. Each elementary label has an FSM, and the

transducers are grouped together as one transducer. A stochastic language model is computed as the joint probability $P(W, T)$:

$$P(W, T) = \prod_{i=1}^k P(w_i t_i | w_{i-1} t_{i-1} \dots w_1 t_1) \quad [80].$$

The model can be further improved by replacing some categories of words with tags, such as country, year, etc. [80].

Conditional Random Fields (CRF) Models are discriminative sequence models based on log-linear models. Unlike HMM, CRF models can incorporate arbitrary features like word shape and prefix. The feature functions are the key to the CRF models. They help us to handle unknown words by incorporating adding features like word shape and morphology. (Words that are capitalized are more likely to be nouns and words that end with ‘-ed’ are usually past tense [50].)

A CRF takes in the input sequence and outputs a probability to the label sequence out of all possible sequences. It computes the posterior probability directly, and at each time step it computes log-linear function over a set of features, and then finds the global probability by aggregating and normalizing the local features. Suppose we have K features, and each feature F_k is the property of the entire input sequence and the entire output sequence and is composed of a sum of local features f_k and has a weight λ_k . We write the desired output as:

$$\begin{aligned} p(T|W) &= \frac{\exp(\sum_{k=1}^K w_k F_k(W, T))}{\sum_{T'} \exp(\sum_{k=1}^K w_k F_k(W, T'))} \\ &= \frac{\exp(\sum_{k=1}^K \lambda_k \sum_{i=1}^n f_k(t_{i-1}, t_i, W, i))}{\sum \exp(\sum_{k=1}^K \lambda_k \sum_{i=1}^n f_k(t_{i-1}, t_i, W, i))} \quad [50]. \end{aligned}$$

Note that we constraint each local feature f_k to be associated with only the current and previous words, and this limitation allows us to train a CRF model in a similar way as we train an HMM [50].

Other statistical models include **Support Vector Machines Models** [80], **Dynamic Bayesian Networks Models** [59, 60], the **Hidden Vector State Models** [37].

4.3 Dialogue Management

The Dialogue Management component deals with information coming from the other components and is responsible for controlling and updating the context of conversations and governing the actions of the chatbot [67].

4.3.1 Design Issues

As mentioned, DM coordinates other modules and encapsulates the logic of the speech application. This means that it is a very important component in the chatbot framework, and we need to make deliberate and thoughtful design decisions [67].

Unfortunately, the process of designing a robust DM strategy is far from trivial as there it is difficult to foresee which form of system behavior will lead to high user satisfaction. In this subsection, we discuss two frequently arising design issues of DM, interaction strategies and the choice of a confirmation strategy [67].

The interaction strategy of a dialogue system dictates who has control over the conversation. The conversation can be user-directed, system-directed, or mixed-initiative. When the conversation is

user-directed, the user takes the initiative and the system just answers the user's queries and commands. When the conversation is system-directed, the system has the initiative and the user simply responds to the system's queries. When the dialogue is mixed-initiative, both the user and the system can take control [67].

The DM keeps track of the state of the conversation, or the conversational behavior of the system, which is based on the interactive strategy it adopts. The table below summarizes the different states that a conversation can be in.

State	Conversation control	Description
Grounded	System	Acknowledging the user's input while deciding on the agent's actions
Slot Filling	System	Requesting extra information from the user to resolve actions
Initiative	System / User	Steering of the conversation by either the user or the agent
Context Switch	User	Change of the basis or premise of the conversation

Table 2: States in Dialogue Management [57]

In addition, the DM component typically includes an error handling module that copes with the uncertainty and ambiguity issues with the information it receives. The ASR and SLU modules are not perfect, and one way to alleviate errors is to ask the user for confirmation when necessary. Typically, confidence scores are assigned to the ASR and SLU results, and results with scores below a certain threshold need to be rejected or confirmed. There are two types of **confirmation strategies**, explicit confirmation, and implicit confirmation. When the explicit confirmation strategy is used, the system confirms its understanding by asking the user another question. [67] The following conversation is an example of an explicit confirmation:

User: I want to know the hours of the closest CVS.

System: Do you want to go to the closest CVS?

User: Yes [67].

When the implicit confirmation strategy is employed, the system includes some of the information it gets in its response [67]. Consider the following example:

User: I want to know the hours of the closest CVS.

System: When do you want to arrive at CVS [67]?

Note that it is the user's responsibility to correct the system if he or she thinks the system has misunderstood the request [67].

There are two types of DM systems, goal/task-oriented systems and non-task-oriented systems [16]. The task-oriented systems are commonly used in agents designed for performing tasks such

as making reservations, answering FAQs, and scheduling meetings [57]. In order to complete the predefined tasks, they direct the dialogue from one state to another [46].

4.3.2 Handcrafted Approaches to DM

We will first go over some of the traditional approaches to DM.

The simplest form of a DM system is a sizeable **Switch Statement**. A Switch Statement is a programming control structure that triggers a predefined response for each possible input or intent [114]. One disadvantage of this approach is that the user is always the one leading the conversation. Also, the actions of the chatbot depend on the predefined set of rules. It cannot engage in conversations or provide responses that are outside the scope [57].

Finite state-based DM uses finite state machines to track the flow or the states of the conversation [42]. This approach is simple and is especially suitable when there is a limited set of possible inputs. Its disadvantage is that it lacks flexibility and adaptability. Conversations that are working toward the same goal have to go through the same process even though not all states are relevant [57].

Frame-based DMs use a frame structure to keep track of all information the system needs from the user. One advantage of this approach is that several data can be captured at the same time and the information can be gathered in any order. Also, frame-based DMs believe humans may change their minds during conversations, so they treat utterances as dialog acts in which the users communicate their goals [66].

Some of the other more advanced approaches include **Information State Theory** [96], **plan-based theory** [8, 9, 20], and **conversational game theory** [15, 54].

4.3.3 Statistical Approaches to DM

Handcrafting DM strategies is time-consuming and requires a lot of effort, so the research focus has now shifted toward machine learning approaches to DM strategy design. The main idea of these machine learning approaches is to learn optimal strategies using data collected from live use of the system [107]. There are three machine learning approaches: reinforcement learning-based, corpus-based, and example-based DM. Example-based approaches can be considered a specific case of corpus-based approaches [67]. In this subsection, we are going to review reinforcement learning-based and corpus-based approaches.

Reinforcement learning-based approaches model a dialogue strategy as a Markov Decision Process (MDP). In the context of a dialogue system, at each step, the system is in some dialogue state s , and it must choose an action a from a set of system dialogue actions. The system moves into a new state s' according to a transition probability $P(s'|s, a)$ which describes the probability of moving from s to s' after action a is performed in state s . The transition is associated with a reward r , which evaluates the result of performing action a in state s . Examples of rewards include task completion and dialogue length. A strategy $\pi(s) = P(a/s)$ or $\pi(s) = a$ if the policy is deterministic specifies the action of the system in each state, and the goal is to learn the best strategy π^* that maximizes the overall reward [50].

The process can be solved by two approaches, model-based approach, and simulation-based approach. The model-based approach first estimates the transition probabilities from a corpus of real human-computer dialog data and then finds the solution using dynamic programming. This approach is also known as “partial strategies” since it can only learn the decisions in some states

and the actions for the other states need to be specified manually. Only in some very rare cases where the corpus contains exploratory data for every state can this approach be used to learn the full strategies. The simulation-based approach learns from sample returns. In this approach, the strategy is learned as the DM interacts with a user simulation. A user simulation is a model that simulates user responses based on a dialogue corpus. The effectiveness of this approach is highly dependent on the simulation's accuracy and ability to generalize unseen dialogue situations. This is a very challenging task, but if such a simulation exists, then the simulation-based approach can explore the state space much more exhaustively, and the "full strategies" can be learned as a result [50].

An extension of the MDP model is a partially observable MDP (POMDP), and it is used when not all states are directly observable. At each step, the system is in an unobserved state s , and it transits to another unobserved state s' after performing action a . Although the exact state cannot be uniquely identified, the system receives an observation o' , which provides some evidence about the new state. The belief state $b(s)$ describes the probability of being in state s [108].

The system chooses an action based on the belief state, and the belief state is updated after the observation is received. The process can be described as follows:

$$\begin{aligned} b'(s') &= P(s'|o', a, b) = \frac{P(o'|s', a, b)P(s'|a, b)}{P(o'|a, b)} \\ &= \frac{P(o'|s', a, b) \sum_{s \in S} P(s'|a, b, s)P(s|a, b)}{P(o'|a, b)} \\ &= k \cdot P(o'|s', a) \sum_{s \in S} P(s'|a, s)b(s) \quad [107], \end{aligned}$$

where $k = P(o'|a, b)$ is a normalization constant [51].

At each time, there is a reward associated with the chosen action, and the accumulated reward throughout the conversation can be found using the equation below:

$$R = \sum_{t=0}^{\infty} \lambda^t R(b_t, a_t) = \sum_{t=0}^{\infty} \lambda^t \sum_{s \in S} b_t(s) r(s, a_t) \quad [108].$$

As in the case of MDP, the goal of the POMDP model is to find the strategy that maximizes the total reward at every point b . However, finding the optimal strategy is very expensive computationally and cannot be done in practice. As a result, POMDPs are not commonly used even though they are theoretically appealing for DM [108].

A **corpus-based approach** to DM generates the next response through a classification process using the complete dialogue history. The DM represents conversations as a sequence of (A_i, U_i) pairs, where A_i is the system's response and U_i is the semantic representation of the user's input at time i [32]. At each moment, the goal is to find the optimal response out of all the possible responses given the history of the conversation, which is denoted as:

$$\hat{A}_i = \operatorname{argmax}_{A_i \in A} P(A_i | (A_1, U_1), \dots, (A_i, U_{i-1})) \quad [33].$$

As described by the equation above, the DM selects the next answer taking the entire dialogue history into account. The biggest problem is that the equation is hard to solve when there are many

possible sequences of states. To overcome this difficulty, a data structure called dialog register (DR) is defined to keep track of the information provided by the user throughout the dialog. After introducing the new data structure, we say the output depends on the DR that results from turn 1 to turn $i-1$ and the previous state. We explicitly consider the previous state since a user turn supplies both task-dependent and task-independent information and are more informative than the DR [33]. Formally, we can then rewrite the above equation as:

$$\hat{A}_i = \operatorname{argmax}_{A_i \in A} P(A_i | DR_{i-1}(A_{i-1}, U_{i-1})) \quad [33].$$

The problem can be solved using a classification process. At each moment, the objective of the DM is to choose a class from a set of classes, and then provide an answer based on the selected class. The classification function is defined using a multiplayer perceptron [83] where the input layer represents the input pair, and the output layer represents the probability of the input belonging to the associated class [33].

As mentioned before, the DR contains information about the labels and values for the user's inputs. However, it is assumed that the exact values of these attributes are not needed to predict the next action of the system. What is relevant is whether they are present or not. Consequently, DR codifies each slot based on whether the concept and the attribute are present or not and the confidence measures of the values [33].

4.4 Backend

Chatbots retrieve information that is needed for performing the required tasks from the Backend and then forward the message to the Dialogue Management Component and Response Generation Component [6].

Rule-based chatbots require a **Knowledge Base (KB)** to store the handcrafted rules. The rules in the Knowledge Base should be as diverse and as holistic as possible to ensure the robustness of the chatbot [52]. A chatbot can also use a Relationship Data Base (RDB) to recall past conversations. Taking the previous information into account allows the chatbot to communicate in a more consistent, precise, and reliable manner [1].

The development of the KB is necessary and essential but can be very time-consuming and demanding as it involves manual work. To overcome this difficulty, developers came up with a way that can build a new KB automatically from the KB of an existing chatbot [11]. There's also a program that can transform a corpus to an AIML KB [3]. Many times, rule-based chatbots guide the users and complete the KB using the users' responses [35].

Recently, there has been a trend toward keeping the information in digitally stored texts on the World Wide Web or in other online sources. These large repositories store information in a machine-readable and accessible form [23]. Examples of these knowledge bases include Google's Knowledge Graph [90], DBpedia [2], Freebase [13], and Wolfram Alpha [103]. Knowledge Graph was introduced in 2012. It provides structured information about a topic or a summary using crowd-sourced and manually curated data [90]. DBpedia data can be accessed online, and the source of its information is structured data from Wikipedia [2].

Typically, these repositories consist of subject-predicate-object triples and are represented by a graphical structure. The nodes indicate the entities, or the subjects and objects, and the edges indicate the relationships between the entities. The figure below is a subgraph of Freebase related

to the TV show *Family Guy*. Note that the nodes include names and dates which are connected using directed edges that specify their relationships [28].

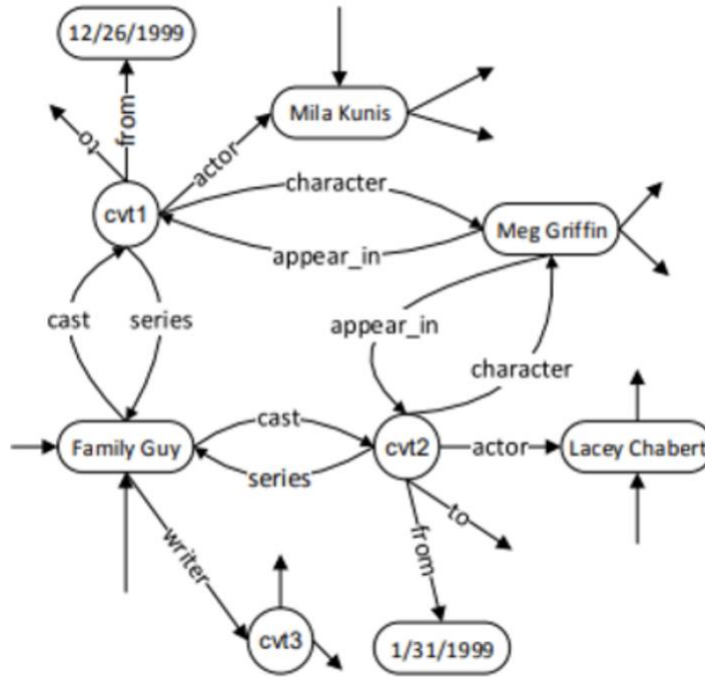


Figure 5: Subgraph of Freebase [106]

4.5 Response Generation

Once the appropriate information has been retrieved, the next step for the dialogue system is to determine the content of the response and the best way to express it. The Response Generation component is responding for generating responses in user understandable format [67].

4.5.1 The RG pipeline

RG involves five stages of processing: signal analysis, data interpretation, document planning, microplanning, and realization [67].

Signal analysis and **data interpretation** are preliminary stages. Input that is in the form of data, such as numerical data in weather predictions, goes through these two stages. The signal analysis identifies patterns and trends in data using pattern matching algorithms. Data interpretation uses symbolic reasoning methods based on domain knowledge to analyze complex messages and relations between the messages [82].

Document planning can be further composed into two substages: content determination and discourse planning. Content determination involves deciding what information needs to be delivered to the user. The system only communicates some of the information retrieved, and it may communicate the information in stages. When necessary, content determination also involves filtering and summarizing the information. Discourse planning is concerned with organizing and structuring the text to aid human comprehension [81].

Microplanning, or sentence planning, involves coming up with the sentences that delivery the information. It consists of three tasks: referring expressions, aggregation, and lexical selection.

Dealing with referring expression involves choosing the appropriate way to address an entity given a particular context. For example, an entity should be referred with a pronoun if it has already been mentioned earlier in the text. Aggregation is the task of joining content with the use of conjunctions, ellipsis, etc. Lexical selection is the issue of looking for the appropriate words for the text [81].

Realization deals with linguistic knowledge of the content, for example, to choose the correct tenses, enforce the subject-verb agreement, and apply word ordering rules, and so on [81].

4.5.2 Statistical Approaches to RG

In this subsection, we are going to review some commonly used language models, models that assign probabilities to sequences of words.

N-gram is the simplest language model. Formally, we write the desired output of assigning probabilities to sequences of n words as:

$$P(w_1, \dots, w_n) [50].$$

We apply the chain rule and the Markov assumption that the current word depends only on the previous $n-1$ words. For example, a bigram model approximates the probability of a word using only the conditional probability of the preceding word. We can rewrite the expression above for the bigram model as:

$$\prod_{k=1}^n P(w_k | w_{1:k-1}) = \prod_{k=1}^n P(w_k | w_{k-1}) [50].$$

The probabilities are estimated using the maximum likelihood estimation. To get the estimate, we estimate the bigram probability by counting the number of occurrences $w_1 w_2$ and then normalizing it by the sum of all bigrams that start with the same word as follows:

$$p(w_n | w_{n-1}) = \frac{C(w_{n-1} w_n)}{\sum_w C(w_{n-1} w)} = \frac{C(w_{n-1} w_n)}{C(w_{n-1})} [50].$$

In practice, trigram, 4-gram, or even 5-gram models are used if sufficient training data is available. To use these models, extra contexts are appended at the beginning and the end of sentences [50].

A **Recurrent Neural Network (RNN)** is a network that contains a cycle in the network connections, which means that previous outputs can be used as inputs [24]. It overcomes a major drawback in N-gram models that the context from which information can be extracted is limited. RNN-based language models [68] take sequences of words as input and predict the next word in a sequence. At each step, the model first computes a hidden layer using the word embedding, or vector representation of the current word and the previous hidden layer. The new hidden layer then passes through a softmax layer to generate a probability distribution over the entire vocabulary [50].

More formally, we describe the process at time t as follows:

$$\begin{aligned} e_t &= E^T x_t \\ h_t &= g(Uh_{t-1} + We_t) \\ y_t &= \text{softmax}(Vh_t), \end{aligned}$$

where x denotes the input, y the output, E the word embedding matrix, and U , V , and W the weight matrices. Concretely, the input sequence consists of word embeddings and the word embedding matrix is used to retrieve the embedding for the current word. Then, the sum of the embedding

multiplied by W and the previous hidden layer multiplied by U is passed through an activation function g to find the activation value for the current hidden layer h_t . Vh represents a set of scores over the possible words given the information provided in h , and the scores are normalized into a probability distribution through the softmax. Using the probability distribution, we can find the probability of a particular word being the next word in the sequence as follows:

$$P(w_{t+1} = i | w_{1:t}) = y_t^i,$$

and the probability of the word sequence is simply the product of these individual word probabilities:

$$P(w_{1:n}) = \prod_{i=1}^n P(w_i | w_{1:i-1}) = \prod_{i=1}^n y_{w_i}^i [50].$$

The RNN is trained on a corpus of text using a regimen called teacher forcing, which minimizes the difference between the predicted and actual probability distributions. However, in practice, RNNs are hard to train because they cannot carry forward distant information and because of the so-called vanishing gradients problem during the backpropagation step [50].

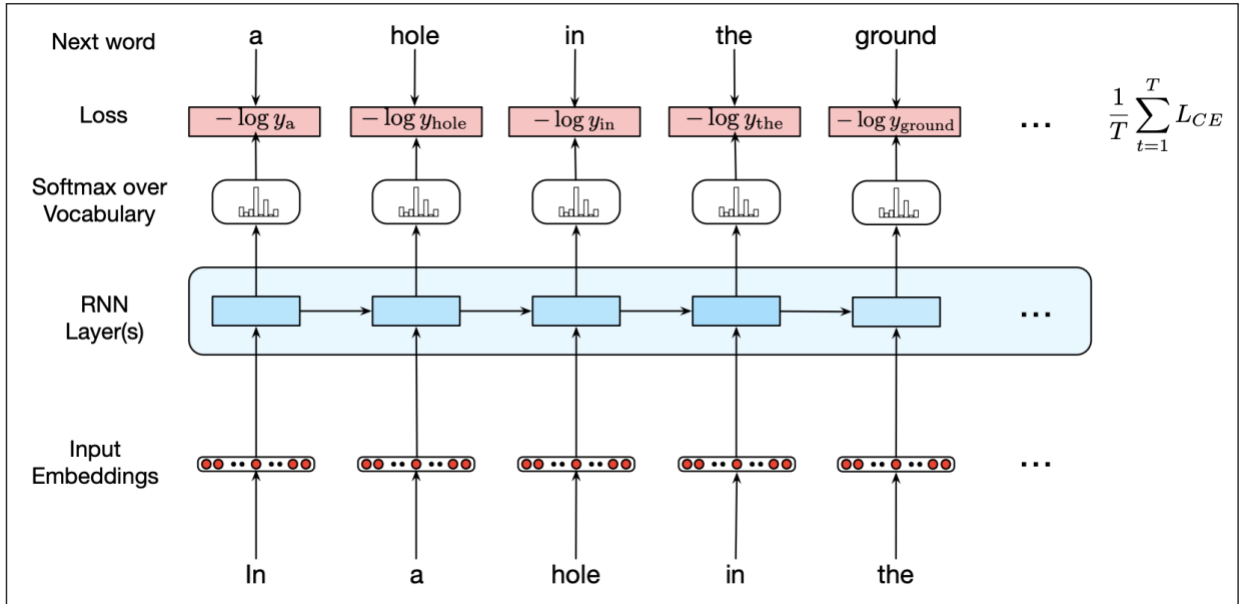


Figure 6: Training RNNs as language models [50]

Long Short Term Memory (LSTM) networks [41] is a variant of RNN that was introduced to overcome the difficulties of RNN. An LSTM network manages the context by removing irrelevant information and adding information likely to be useful in the future. This is achieved by explicitly adding a context layer and specialized units called gates to control the information flow. Three gates are introduced into the architecture, the forget gate, the add gate, and the output gate. The forget gate is responsible for getting rid of information that is no longer useful, the add gate selects new information that is needed for the current context, and the output gate determines whether a piece of information is required for the current hidden state. Each gate is composed of a feedforward layer, a sigmoid activation function, and a pointwise multiplication with the layer being gated [50].

Transformers map input sequences to output sequences with a similar approach to RNN. At each step, given an input sequence, the Transformer produces an output distribution over all possible

words. However, it does so without the recurrent connections with the use of additional layers called self-attention. Self-attention is the key to transformers. Information can be directly processed without going through intermediate recurrent connections. A self-attention layer takes word sequences as input and output sequences of the same length. At each step, it has access to the current input and all previous inputs. An output is computed using the relevance of input in a given context, which can be found by comparing the input with some relevant items. More formally, the result of the comparisons is referred to as scores, and the output is the weighted sum of normalized scores:

$$y_i = \sum_{j \leq i} \frac{\exp(\text{score}(x_i, x_j))}{\sum_{k=1}^i \exp(\text{score}(x_i, x_k))} x_j \quad [50].$$

Also, additional parameters are included in the transformers to provide opportunities for learning. Each parameter is a set of weight matrices that are used to operate over the inputs. The self-attention layers, along with some other components like feedforward layers, residual connections, and normalizing layers form a transformer block as illustrated in the figure below. A transformer is made of stacks of transformer blocks [50].

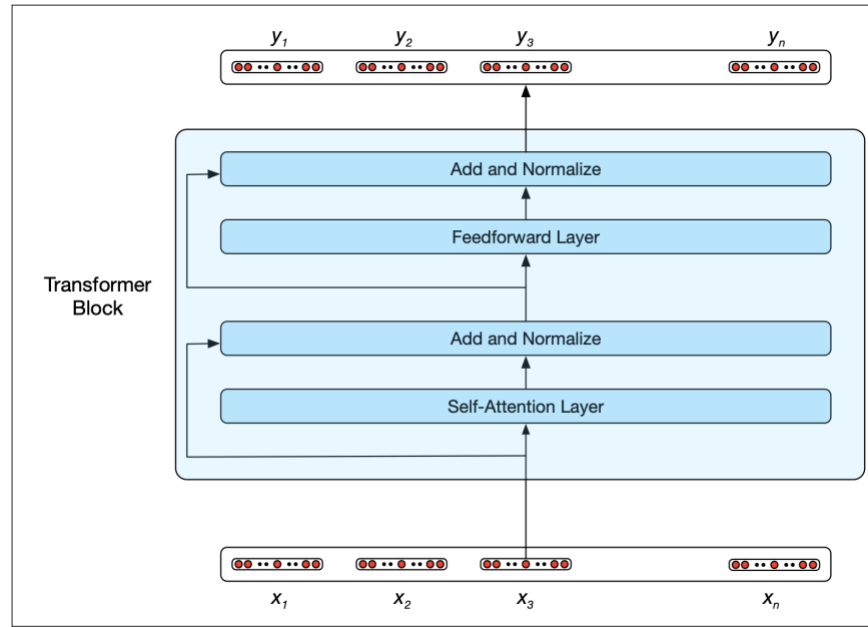


Figure 7: A transformer block [50]

An extension of the transformer model uses multihead self-attention layers. Each multihead self-attention layer is a set of self-attention layers and has its distinct set of parameters. This characteristic makes the extension very useful when words in the sentences are related to each other in different ways. Another extension combines the input with positional embedding so that the input incorporates information about the order [50].

Encoder-decoder networks or **sequence-to-sequence** networks are a special class of RNN capable of generating arbitrary-length output sequences within the context. An encoder-decoder network has three components, an encoder, a context vector, and a decoder. The encoder converts the input sequence into a contextualized representation called the context, and then the decoder accepts the context as input to generate an output sequence. Both the encoder and decoder are

either LSTM or Transformer models. The first hidden state of the decoder is initialized to the final state of the encoder, and each other hidden state takes the previous hidden state, the output generated in the previous state, and the context vector as input. More formally, this process can be illustrated with the equation below:

$$h_t^d = g(\hat{y}_{t-1}, h_{t-1}^d, c),$$

where c denotes the context vector, and the superscripts e and d are used to distinguish the hidden states of the encoder and the decoder. The full equations for the decoder are then:

$$c = h_n^e$$

$$h_0^d = c$$

$$h_t^d = g(\hat{y}_{t-1}, h_{t-1}^d, c)$$

$$z_t = f(h_t^d)$$

$$y_t = \text{softmax}(z_t) [50].$$

We then use the output y to find the most likely output at each step [50].

The figure below illustrates the entire process of the model.

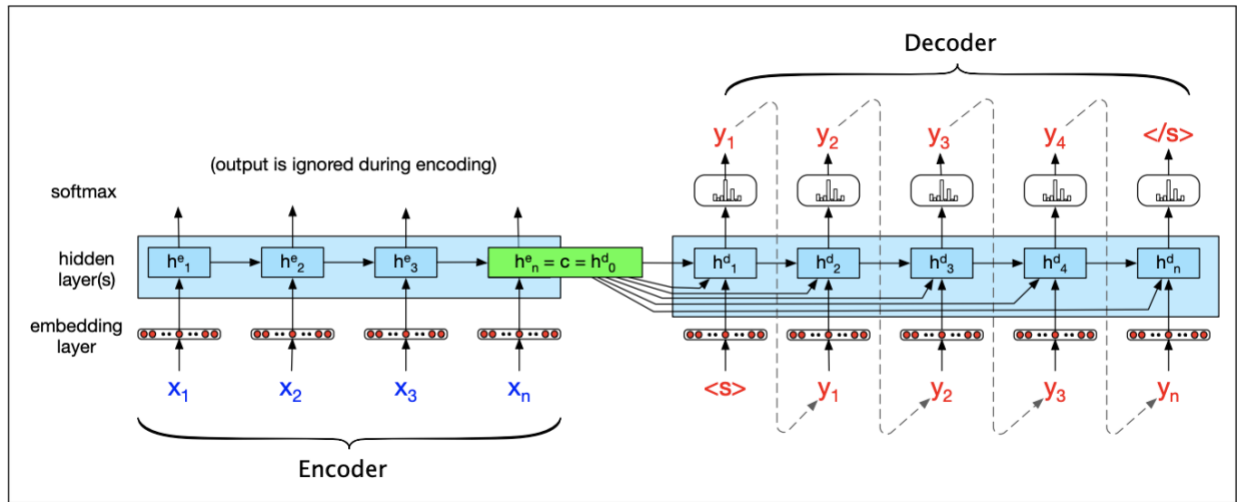


Figure 8: The basic RNN-based encoder-decoder architecture [50]

Encoder-decoder architectures are trained on tuples of paired strings, each consists of a source and a target [50].

5 Applications and Future Outlook

In this section, we will focus on the applications of chatbot technology. We will first examine two modern chatbots, and then address the potential implications of chatbots on the wider world.

Before we delve into the details of the examples, we will discuss some issues relating to the construction of chatbots, or more specifically, how to tailor the general architecture we presented for specific applications. The first step in the development of a chatbot involves understanding its objectives and user requirements. Once the use cases are decided, sample dialogues are created using relevant knowledge [21]. Many corpora are available for training chatbots. Conversational chatbots are usually trained on real dialogues or movie scripts corpora [112]. Domain-specific chatbots generally use knowledge elicited from an expert within the field [6]. Most of the time, these sample dialogues will not be sufficient. The chatbot will continuously learn from the conversations, and developers can also apply copywriting skills to polish the conversation [21].

Keeping the development process in mind, we will look at the objectives of XiaoIce and Meena, as well as their specific architectures.

5.1 Microsoft XiaoIce

We first present the design and implementation of Microsoft XiaoIce, the most popular social chatbot in the world. A social chatbot is a dialogue system designed not only to respond to users' requests but also to establish long-term and emotional connections with users. XiaoIce has attracted a lot of attention since its release in 2014 as many users expect a chatbot to be a companion instead of just an assistant. In fact, around 40% of the users' requests to the agents are emotional [113]. To better understand the capabilities of XiaoIce, we show conversation sessions between a user and XiaoIce in the table below:

Day 1	After 9 Weeks
User: Hello, XiaoIce. XiaoIce: Hellos, your profile pic is interesting. Like I've seen it somewhere. User: Really? XiaoIce: Is it a comic character? User: How do you know that? XiaoIce: Guess. User: Are you human? XiaoIce: Well, you'll know it after chatting with me more.	User: What are you doing, XiaoIce? XiaoIce: Feeling bored (thinking: what's up?) User: She refused me today. XiaoIce: What a coincidence! Someone also refused me today. User: It happened just now. XiaoIce: What did she say? User: She is in love with someone else, her schoolmate in college. XiaoIce: Let me see...You're smart, cute and handsome, why not you?

Table 3: Sample conversation sessions between a user and XiaoIce [113]

From the table above, we see during the first encounter with XiaoIce (left), the user explored the general features and just started to get to know XiaoIce. After they became more familiar with each other, the user started to chat about his hobbies and daily lives with XiaoIce. By the end of 9

weeks (right), the user talked to XiaoIce almost every day and preferred to talk to XiaoIce whenever he needed to talk. The user felt comfortable sharing personal feelings and secrets with XiaoIce [113].

5.1.1 Design Principle

As mentioned, the core of a social chatbot is its ability to respond to both instructional and emotional requests. More concretely, it requires both a high intelligent quotient and an emotional quotient, and a unique personality [113].

Intelligent quotient capacities are skills that allow the chatbot to respond to specific questions and tasks. These capacities include knowledge and memory modeling, image and natural language understanding, reasoning, generation, and prediction. **Emotional quotient** is composed of two key components, empathy, and social skills. Empathy refers to the capacity to understand or feel from another person’s point of view. Social skills involve framing the responses based on the user’s backgrounds, interests, and needs and maintaining the flow of the conversation. **Personality** is defined as the characteristic set of behaviors, cognitions, and emotional patterns that form an individual’s distinctive character. A social chatbot needs to demonstrate a consistent personality so that the users know what to expect during the conversation and consequently develop a long-term relationship with it. XiaoIce persona is designed as an 18-year-old girl who is always reliable, sympathetic, affectionate and has a wonderful sense of humor [113].

These objectives are fulfilled through a **hierarchical decision-making process**, which can be modeled using a **Markov Decision Process** [93]. To maintain its attractiveness to the user, XiaoIce is able to behave in a variety of conversation modes, including chatting casually, restaurant booking, and etc., and each conversation mode is handled by its corresponding skills. A dialogue can be regarded as a hierarchical decision-making process: a top-level process governs the overall conversation and selects the appropriate skills, and the low-level process then generates responses and actions accordingly. The process can be modeled with a Markov Decision Process, where human users are the environment. At each step, the chatbot takes the current dialogue state as input and outputs a skill or a response based on the hierarchical dialogue policy. It then receives a reward from the user and moves on to the next state. The goal of the chatbot is to find the optimal policies and skills that maximize the rewards [113].

5.1.2 Architecture

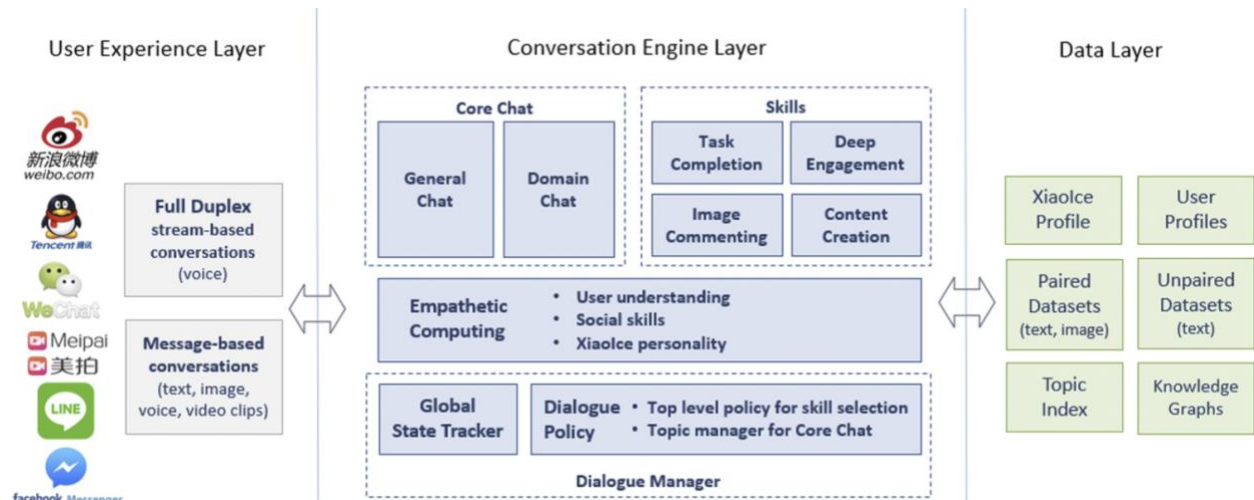


Figure 9: XiaoIce system architecture [113]

XiaoIce consists of three layers: user experience, conversation engine, and data, as illustrated in the figure above.

The **user experience layer** assists XiaoIce to communicate with the users through popular chat platforms. There are two communication modes, full-duplex and taking turns. In the full-duplex mode, the user and XiaoIce can talk to each other simultaneously. In the other mode, the user and XiaoIce talk one at a time. This layer also includes components that are responsible for processing the user’s input and generating appropriate responses [113].

The **data layer** is composed of databases that store collected conversational and non-conversational data, knowledge graphs used by Core Chat and skills, and the profiles of users and XiaoIce. The conversational data are collected from the internet and past conversations between XiaoIce and the human users, and the non-conversational data are collected from public articles, reports, and lectures [113].

The **conversation engine layer** includes four components, a dialogue manager, an empathetic computing module, Core Chat, and dialogue skills [113].

The dialogue manager consists of the global state tracker that keeps track of the dialogue state and the dialogue policy that selects an action according to the dialogue policy. An action can be either a skill or Core Chat, and as discussed in the previous subsection, the top-level policy dictates which one is triggered. Core Chat deals with text inputs, and the skills handle image or video inputs and other specific requests from the user. The collection of skills and Core Chat demonstrates XiaoIce’s IQ [113].

The empathetic computing module is responsible for understanding the empathetic aspects of the conversation and can be viewed as the language understanding unit of XiaoIce. It takes the user’s query as input and outputs a dialogue state vector that contains information about the context of the conversation and empathetic aspects of the query and the response [113].

Core Chat is the response generation component. It consists of two parts, General Chat and a set of Domain Chats. General Chat is responsible for engaging in open-domain conversations while Domain Chats are responsible for engaging in conversations on specific domains. The only difference between General Chat and Domain Chats implementation-wise is the databases they access, so we describe only General Chat in this paper. General Chat outputs responses by first generating a set of responses candidates and then ranking the candidates. The candidate generators can be either retrieval-based or neural-model-based that follows the seq2seq framework. Retrieval-based generators retrieve responses using either conversational or non-conversational data, and neural-model-based generators are trained on conversational data but can learn from the conversations and generate unseen responses. Note that data that contains personally identifiable information and those that do not fit XiaoIce’s persona is removed. Retrieval-based generators that use conversational data provide high-quality responses but are limited to specific topics, while the other two generators can cover a wider range of topics [113].

After the responses are generated, a response is randomly selected among candidates with ranking scores that are higher than a pre-determined threshold. The responses are ranked based on coherence and consistency and are assigned a score of 0, 1, or 2. Figure 10 shows sample query-response pairs and their ranking scores [113].

Finally, to keep the conversation going, XiaoIce will provide an editorial response like “Hmm, difficult to say. What do you think?” when the system fails to generate any valid responses [113].

Query	Response	Rating
捂上脸不看 (Don't want to look with my face in my hands.)	我边吃饭边看你 (I'm watching you while eating.)	0
	你怎么了? (what is wrong with you?)	1
	当心捂出一脸痱子, 哈哈. (Haha, beware of the heat rash of your face.)	2
我肚子有点疼 (I have a pain in my stomach.)	运动对身体有好处的 (Exercise is good for your health.)	0
	喝热水了吗? (Have you been drinking hot water?)	1
	来和我聊天, 分散一下注意力就好了 (Come and talk to me, just distract yourself.)	2
我觉得自己太笨了 (I feel so stupid.)	大象的智商很高 (The elephant has a high IQ.)	0
	多学习提升自己 (Learn more and improve yourself.)	1
	很多聪明人都这么说耶 (That is what lots of smart people say about themselves.)	2
为什么人总是活成自己不喜欢的样子 (Why do we always live the way we don't like.)	喜欢一个人是发自内心的 (Love a person from my heart.)	0
	因为人想要的太多 (Because we want too much.)	1
	因为那个是别人喜欢的样子 (Because that's what others like.)	2
我考试没及格 (I did not pass the exam.)	那要看和谁比呢 (It depends on whom you compare against.)	0
	一次没关系, 下次努力吧 (Don't worry, try it next time.)	1
	起码你不用怀疑自己能不能及格了 (At least you don't have to wonder if you can pass.)	2

Figure 10: Examples of query-response pairs labeled on a three-level quality scale [113]

XiaoIce is also capable of responding to image inputs, and this skill is referred to as Image Commenting. Responses to images are generated in a similar process as those to texts. The figure below contains sample comments to image inputs generated by XiaoIce [113].



Figure 11: Examples of image-comment pairs labeled on a three-level quality scale [113]

XiaoIce is equipped with 230 dialogue skills, and these three skills can be categorized into three categories: content creation, deep engagement, and task completion. The content creation skills involve creative works such as generating poetry, songs, books, and so on. The deep engagement skills help establish long-term connections with users by choosing topics and settings that fit better with the user's needs. The task completion skills are responsible for accomplishing tasks the users request, including flight booking, weather forecasting, restaurant recommendation, and so on. XiaoIce's task-completion skills offer responses that are more individualized and engaging in

comparison with traditional chatbots. For example, when the user asks “what’s the area of China?,” XiaoIce, knowing that the user comes from the United States, may respond “it’s 3.71 million sq miles, about equal to the size of USA [113].”

5.2 Google Meena

In this subsection, we are going to look at Meena, a human-like open-domain chatbot introduced by Google. Below is a sample conversation session between a user and Meena:

User: Hi!

Meena: Hi! How are you this fine evening?

Human: I’m doing well. It’s morning my time!

Meena: Good morning! What are your plans for today?

User: Get my work done then watch a movie at home in the night

Meena: Sounds fun! What movie?

Human: I haven’t decided yet, perhaps an Oscar nominated one. any suggestion?

Meena: I’m not too familiar with recent Oscar winners, but I do know I wanted to see *The Grand Budapest Hotel*.

...

As shown above, Meena is able to chat with the human user pretty smoothly and is able to engage in conversations in an arbitrary domain [7].

5.2.1 Human Likeness

One distinguishable character of Meena is its human likeness. Human likeness takes both sensibleness and specificity into consideration. Sensibleness addresses the more concrete and basic aspect of human conversations while specificity addresses the more subjective aspect of human conversations [7].

Sensibleness simply indicates whether the response makes sense given the context. Some important aspects sensibleness captures include common sense, logical coherence, consistency, and so on. Any response that’s out of context, confusing, or wrong, is considered not sensible [7].

However, sensibleness itself is not sufficient. Many chatbots are developed based solely on metrics like sensibleness and their responses are dry and unspecific []. The response also needs to take **specificity** into account, i.e., be specific to the context. Suppose the user says, “I love golf.” Responses like “Oh that’s good to know” are not specific while responses like “Me too, I can’t get enough of Tiger Woods!” are closely related to the topic and considered specific [7].

5.2.2 Training Data

Meena is trained on public domain social media conversations. The source data can be described using a tree-structure: the first message is the root, and the responses to the message are its child nodes. A training example of the form (context, response) pair can then be formed where the context includes the previous seven messages and the response the current message [7].

To improve the quality of the data, messages are evaluated based on certain criteria and those that do not meet the standard will be removed. The criteria are as follows: 1) messages with less than

2) words or more than 128 words; 2) more than 30% are non-alphabetic characters; 3) contain URL; 4) the author's username contains "bot"; 5) messages that have been occurred more than 100 times; 6) messages that have a high n -gram overlap with the parent's text; 7) potentially unsafe or offensive messages. When a message is removed, all its descendants are also removed for simplicity [7].

After the pre-processing steps, there are 867M (context, response) remaining pairs. The final dataset contains 341 GB of text or 40B words after the text is tokenized [7].

5.2.3 Model Architecture

The best performing Meena model is an Evolved Transformer seq2seq model with 2.6B parameter. An Evolved Transformer is an improved version of the Transformer introduced by So, Liang, and Le in 2019, and it has outperformed the traditional architecture on several tasks [91].

A simple sample-and-rank decoding strategy is used to generate responses. In this strategy, N candidate responses are sampled and the candidate with the highest probability is selected as the final output [7].

5.3 Other Applications

XiaoIce and Meena are two examples of how chatbots are used in real life. With the increase in research and development of chatbot technology, chatbot applications have evolved into one of the most powerful and widely adopted tools. They are now visible in almost every domain, including education, e-commerce, and healthcare [88]. In the following paragraphs, we are going to look at how chatbots are used in these three areas, but the usefulness of chatbots is definitely not limited to these three fields.

One important application of chatbots is their use in the **educational services industry**, or more specifically, in the e-learning industry. In many cases, a chatbot can be as helpful as a human teacher. It can foster an interactive and personalized learning environment for the students, answer students' questions and check their work, keep track of the progress and achievements of each student, and so on. It may even be more efficient than a teacher in some cases. For example, their responses are more accurate, they can work with multiple students at the same time, and information on e-learning platforms can be preserved [40]. It is worth mentioning that chatbots are quite popular in foreign language learning. Many students are more willing to learn with a chatbot than with a human tutor because they are more comfortable speaking in front of an agent, and they can learn according to their own pace [36]. Jia [48] introduced a chatbot for English learners in 2009. The bot, performing syntactic and semantic analysis, is able to generate sensible responses that take the context of the conversation and the student's ability and personality into account [48].

Chatbots are used extensively in the **e-commerce industry**. Over the past years, the way people interact with businesses has changed. With the development of new technologies, the e-commerce market has been evolving at a rapid pace. However, this new way of selling products has created problems related to the quality of customer service for businesses. The issue is especially significant in live chats. Customers find it hard to receive timely and relevant responses [74]. As a result, many companies adopted chatbots to facilitate customer support [49]. They promise a range of potential benefits. Chatbots are available 24/7, and their responses are instant, consistent, and reliable, which greatly improves user satisfaction. There are a variety of chatbots available in the market nowadays. For example, there is a website-based chatbot implemented with RiveScript that can be used to recommend suitable products for customers [34]. There's also a chatbot written

in AIML and LSA that provides responses to users' questions by collecting information from a Frequently Asked Questions dataset [73]. It will not be a surprise that chatbots will dominate the customer service industry soon in the future.

In the **healthcare industry**, chatbots can provide a wide range of services include answering customized health and therapy information, suggesting related products and services, and recommending a diagnosis and treatments based on the user's symptoms [94]. For example, OneRemission [75] provides information to those who are fighting against cancer. Youper [109] is a chatbot that aims to enhance users' emotional health and Florence [26] helps patients achieve their wellness goals by tracking their health and reminding them to take their pills. Other healthcare chatbots include Your.Md [110], AdaHealth [4], Sensely [87], and so on. Also, many chatbots developed recently such as HealthBuddy [38] aim to inform users about the COVID-19 pandemic. Some of the advantages of using healthcare chatbots include they are cost-effective, unbiased, and reliable especially at administrative tasks. Moreover, many customers are more comfortable sharing their information and concerns with chatbots. However, chatbots lack the consistency and transparency of human physicians, who meet patients regularly. More importantly, chatbots cannot be trusted at all times, and patients prefer to seek human expertise and are reluctant to use health bots when they are making important health decisions [6].

5.4 Implications and Future Outlook

During the past years, there has been a significant advancement in chatbot technology. Chatbots are now replacing the role of humans in many domains, and they are able to perform the tasks sufficiently well. However, a chatbot is much more beyond a system that can perform machine-like interactions, and increasing the trustworthiness or the credibility of the chatbots has now attracted a lot of research attention.

User's trust in the agents comes from two aspects: 1) factors concerning the specific chatbot and 2) factors related to its manufacturer [100].

Factors concerning the specific chatbot deal with its human-likeness, which encompasses its appearance, its ability to interpret and provide responses, and its personality. In the context of a chatbot, appearance refers to the graphical representation of the agent, such as its age and clothing. Adopting human figures is the most straightforward way to enhance the human-likeness. Humans are much more likely to treat the agent as a human when they see a human figure during the conversation [30]. The natural language understanding and generation ability is the focus of the paper and is perhaps the most critical issue in the future development of chatbots. There are noticeable differences in terms of content and quality between human-chatbot communication and human-human communication and the fundamental issue here is to understand what's causing the differences. We have reviewed some techniques in the previous sections, but to fully overcome the problem requires breakthroughs in many areas, including psychology, cognitive science, linguistics, computer science, and so on. There is general agreement in the field of psycholinguistics that a close relationship exists between language and thought [58]. Chatbots today cannot learn to think, and a prerequisite for future development is the clarification of the mechanisms of thought, language, and their relations [6]. Finally, giving the chatbot a personality changes the user's expectations of its behavior, and we already discussed the basics of personality and how it is used in XiaoIce.

Factors related to its manufacturer involve privacy and security issues. Establishing ethical guidelines for the design and implementation of chatbots is especially important as chatbots

become more dominant in our daily lives. Privacy is a concept that has been discussed widely, and for many users, it is the biggest concern. Conversational agents, especially those designed to develop an intimate relationship with the users, can gain access to very personal and sensitive information. It is therefore important for the chatbots to use this information only for the purpose of providing services to the users. Developers must be very clear and transparent about what data are collected, who has access to them, and how they are used [83]. Moreover, given the power and significant influence of the agents, it is also critical that they do not harm or take advantage of the users. The systems should be user-centric, and users should always remain in control except in very special cases. New features must be evaluated and examined thoroughly before introduced to maintain the level of security [89].

Finally, it is worth our attention to note that the use of chatbots has a complex social impact, and the developers should take social factors into consideration during the design and implementation stages. Today, chatbots are used as social utilities and they interact with users through language. Language is inherently social, cultural and contextual, and is served to convey a particular view or value. Language is also constantly evolving, and a word may have entirely different meanings today and 50 years ago. Consequently, decisions that are made during development will eventually reflect the beliefs and worldviews the chatbots carry. Developers should consider how they can design a chatbot that is socially acceptable and responsible and even handle sensitive conversations or requests [83].

6 Conclusion

A chatbot is an exemplary example of human-computer interaction. In recent years, there has been a significant advancement in the development of chatbots, and they have evolved into one of the most powerful and widely adopted applications. However, this is only the beginning of chatbots, and to meet their full potential, it requires our attention and effort to understand how they work and how they are developed.

In this paper, we reviewed the history of chatbots and saw how they have evolved since their inception. We then examined the design principles and general architecture in detail. Next, we presented how we can build specific applications and use cases. Finally, we discussed the future direction of research as well as ethical and social considerations during the design and development stages. Enhancing language comprehension and generation ability is a critical step in future development. Given the rapid developments in NLP and machine learning, we are confident that we can build more powerful human-like chatbots in the near future.

References

- [1] Abdul-Kader, S. A., & Woods, J. (2015). Survey on Chatbot Design Techniques in Speech Conversation Systems. *International Journal of Advanced Computer Science and Applications*, 6, 72-80. <http://dx.doi.org/10.14569/IJACSA.2015.060712>
- [2] About. DBpedia. (n.d.). <https://wiki.dbpedia.org/about>.
- [3] AbuShawar B., Atwell E. ALICE Chatbot: Trials and outputs *Computación y Sistemas*, 19 (2015). doi: 10.13053/cys-19-4-2326
- [4] Ada: Your health companion (2020) Ada. website: <https://ada.com/>.
- [5] Adamopoulou E., Moussiades L. (2020) “An Overview of Chatbot Technology”. In: Maglogiannis I., Iliadis L., Pimenidis E. (eds) *Artificial Intelligence Applications and Innovations. AIAI 2020. IFIP Advances in Information and Communication Technology*, vol 584. Springer, Cham. https://doi.org/10.1007/978-3-030-49186-4_31
- [6] Adamopoulou, E., & Moussiades, L. (2020, November 9). Chatbots: History, technology, and applications. *Machine Learning with Applications*. doi: 10.1016/j.mlwa.2020.100006
- [7] Adiwardana, D., Luong, M.-T., So, D. R., Hall, J., Fiedel, N., Thoppilan, R., ... Le, Q. V. (2020, January 27). Towards a Human-like Open-Domain Chatbot. *arXiv.org*. <https://arxiv.org/abs/2001.09977v1>.
- [8] Allen JF, Perrault CR (1980) Analyzing intentions in dialogs. *Artif Intell* 15(3):143–178. doi:10. 1016/0004-3702(80)90042-9
- [9] Appelt DE (1985) *Planning English sentences*. Cambridge University Press, Cambridge. doi:10. 1017/CBO9780511624575
- [10] Arsovski S., Muniru I., Cheok A. Analysis of the chatbot open source languages aiml and chatscript: A review (2017). doi: 10.13140/RG.2.2.34160.15367
- [11] Arsovski S., Osipyan H., Oladele M.I., Cheok A.D. Automatic knowledge extraction of any chatbot from conversation *Expert Systems with Applications*, 137 (2019), pp. 343-348. doi: 10.1016/j.eswa.2019.07.014
- [12] Artificial intelligence scripting language—Rivescript.com (2019) <https://www.rivescript.com/>.
- [13] Bollacker, K., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. (2008). Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250. ACM. <https://doi.org/10.1145/1376616.1376746>
- [14] Bradeško, L., & Mladenčić, D. (2012). A survey of chatbot systems through a loebner prize competition. *C*, 34.
- [15] Carletta JC, Isard A, Isard S, Kowtko J, Doherty-Sneddon G, Anderson A (1995) The coding of dialog structure in a corpus. In: Andernach T, van de Burgt SP, van der Hoeven GF (eds) *Proceedings of the Twente workshop on language technology: corpus-based approaches to dialogue modelling*, University of Twente, Netherlands, June 1995

- [16] Chen, H., Liu, X., Yin, D., & Tang, J. (2017). A Survey on Dialogue Systems. *ACM SIGKDD Explorations Newsletter*, 19(2), 25–35.
- [17] Clarizia, F., Colace, F., Lombardi, M., & Domenico Santaniello, F. P. (2018). Chatbot: An Education Support System for Student ... *Cyberspace Safety and Security*. https://link.springer.com/chapter/10.1007/978-3-030-01689-0_23.
- [18] Cognitive Agent that Learns and Organizes. The IT Law Wiki. (n.d.). https://itlaw.wikia.org/wiki/Cognitive_Agent_that_Learns_and_Organizes.
- [19] Cohen P, Levesque H (1990) Rational interaction as the basis for communication. In: Cohen P, Morgan J, Pollack M (eds) *Intentions in communication*. MIT Press, Cambridge, MA, pp 221–256. <https://www.sri.com/work/publications/rational-interaction-basis-communication>.
- [20] Colby K.M., Hilf F.D., Weber S., Kraemer H.C. Turing-like indistinguishability tests for the validation of a computer simulation of paranoid processes. *Artificial Intelligence*, 3 (1972), pp. 199-221. doi: 10.1016/0004-3702(72)90049-5
- [21] The Conversation Design Process: How to Design Chatbot Dialogue. KeyReply Blog. (n.d.). <https://blog.keyreply.com/blog/conversation-design-process>.
- [22] Dahiya, M. (2017, May 30). A Tool of Conversation: Chatbot. *JCSE - International Journal of Computer Sciences and Engineering*. http://ijcseonline.org/pub_paper/27-IJCSE-02149.pdf.
- [23] de Melo G, Hose K (2013) Searching the web of data. In: Serdyukov P et al (ed) *Advances in information retrieval: Proceedings of 35th European conference on IR research, ECIR 2013, Moscow, Russia, 24–27 Mar 2013. Lecture Notes in Computer Science*, vol 7814. Springer Publishing Company, pp 869–873. doi:10.1007/978-3-642-36973-5_105
- [24] Elman, J. L. (1990). Finding structure in time. *Cognitive science* 14(2), 179–211.
- [25] Firdaus, M., Golchha, H., Ekbal, A., & Bhattacharyya, P. (2020, March 23). A Deep Multi-task Model for Dialogue Act Classification, Intent Detection and Slot Filling. *Cognitive Computation*. <https://link.springer.com/article/10.1007/s12559-020-09718-4>.
- [26] Florence—your health assistant (2020) <https://florence.chat/>.
- [27] Følstad, A., Nordheim, C. B., & Bjørkli, C. A. (2018). What Makes Users Trust a Chatbot for Customer Service? An Exploratory Interview Study. *Internet Science*. https://link.springer.com/chapter/10.1007/978-3-030-01437-7_16.
- [28] Gao, J., Galley, M., & Li, L. (2019, September 10). Neural Approaches to Conversational AI. *arXiv.org*. <https://arxiv.org/abs/1809.08267v3>.
- [29] Gildea D, Jurafsky D (2002) Automatic labeling of semantic roles. *Comp Linguist* 28(3):245–288. doi:10.1162/089120102760275983
- [30] Go, E., & Sundar, S. S. (2019, January 29). Humanizing chatbots: The effects of visual, identity and conversational cues on humanness perceptions. *Computers in Human Behavior*. <https://www.sciencedirect.com/science/article/pii/S0747563219300329>.
- [31] Google Assistant, your own personal Google. <https://assistant.google.com/>

- [32] Griol D, Hurtado LF, Segarra E, Sanchis E (2008) A statistical approach to spoken dialog systems design and evaluation. *Speech Commun* 50(8–9):666–682. doi:10.1016/j.specom.2008.04.001
- [33] Griol D, Callejas Z, López-Cózar R, Riccardi G (2014) A domain-independent statistical methodology for dialog management in spoken dialog systems. *Comput Speech Lang* 28 (3):743–768. doi:10.1016/j.csl.2013.09.002
- [34] Gupta, S., Borkar, D., Mello, C. D., & Patil, S. (2015). An e-commerce website based chatbot. 6, 3.
- [35] Hahm, Y., Kim, J., An, S., Lee, M., & Choi, K.-S. (2018). Chatbot who wants to learn the knowledge: kb-agent. *Semdeep/NLIWoD@ISWC*, 4.
- [36] Haristiani N. Artificial intelligence (AI) chatbot as language learning medium: An inquiry. *Journal of Physics: Conference Series*, 1387 (2019), Article 012020, [10.1088/1742-6596/1387/1/012020](https://doi.org/10.1088/1742-6596/1387/1/012020)
- [37] He Y, Young S (2006) Spoken language understanding using the hidden vector state model. *Speech Commun* 48(3–4):262–275. doi:10.1016/j.specom.2005.06.002
- [38] Healthbuddy: a new chatbot to engage with communities in europe and central asia on covid-19 (2020) <http://www.euro.who.int/en/health-topics/health-emergencies/coronavirus-covid-19/news/news/2020/5/healthbuddy-a-new-chatbot-to-engage-with-communities-in-europe-and-central-asia-on-covid-19>.
- [39] Heller B., Procter M., Mah D., Jewell L., Cheung B. Freudbot: An investigation of chatbot technology in distance education (2005)
- [40] Hien, H.T., Cuong, P.-N., Nam, L.N.H., Nhung, H.L.T.K., Thang, L.D.: Intelligent assistants in higher-education environments: the FIT-EBot, a chatbot for administrative and learning support. In: *Proceedings of the Ninth International Symposium on Information and Communication Technology*, pp. 69–76. ACM, New York (2018)
- [41] Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation* 9(8), 1735–1780.
- [42] Hori, C., Kiyonori Ohtake, Teruhisa Misu, Hideki Kashioka, & Satoshi Nakamura (2009). Statistical dialog management applied to WFST-based dialog systems. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 4793-4796). doi: 10.1109/ICASSP.2009.4960703.
- [43] Hovy, E. H., Marcus, M. P., Palmer, M., Ramshaw, L. A., and Weischedel, R. (2006). Ontonotes: The 90% solution. *HLT-NAACL*.
- [44] Hoy M. Alexa, siri, cortana, and more: An introduction to voice assistants *Medical Reference Services Quarterly*, 37 (2018), pp. 81-88. doi: 10.1080/02763869.2018.1404391
- [45] Huang, Z., Xu, W., and Yu, K. (2015). Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*.

- [46] Hussain, N. (2019). A Survey on Conversational Agents/Chatbots Classification and Design Techniques. In *Web, Artificial Intelligence and Network Applications* (pp. 946–956). Springer International Publishing. doi: 10.1007/978-3-030-15035-8_93
- [47] IBM Watson. <https://www.ibm.com/watson>
- [48] Jia J. CSIEC: A computer assisted english learning chatbot based on textual knowledge and reasoning *Knowledge-Based Systems*, 22 (4) (2009), pp. 249- 255. doi: 10.1016/j.knosys.2008.09.001
- [49] Johannsen F., Leist S., Konadl D., Basche M. Comparison of commercial chatbot solutions for supporting customer interaction *ECIS* (2018)
- [50] Jurafsky, D., & Martin, J. H. (2020). *Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition*. Pearson.
- [51] Kaelbling LP, Littman ML, Cassandra AR (1998) Planning and acting in partially observable stochastic domains. *Artif Intell* 101(1–2):99–134. doi:10.1016/s0004-3702(98)00023-x
- [52] Khanna, A., Pandey, B., Vashishta, K., Kalia, K., Bhale, P., Das, T.: A study of today’s A.I. through chatbots and rediscovery of machine intelligence. *Int. J. u- e-Serv. Sci. Technol.* 8, 277–284 (2015). <https://doi.org/10.14257/ijunesst.2015.8.7.28>
- [53] Klopfenstein, Lorenz Cuno, Saverio Delpriori, Silvia Malatini, and Alessandro Bogliolo. 2017. “The Rise of Bots.” In *Proceedings of the 2017 Conference on Designing Interactive Systems*. ACM. <https://doi.org/10.1145/3064663.3064672>
- [54] Kowtko JC, Isard SD, Doherty, GM (1993) *Conversational games within dialogue*. Human Communication Research Centre, University of Edinburgh, (HCRC/RP-31). doi:10.1.1.52.5350
- [55] Kucherbaev, P., Bozzon, A., Houben, G.-J.: Human-aided bots. *IEEE Internet Comput.* 22, 36–43 (2018). <https://doi.org/10.1109/MIC.2018.252095348>
- [56] Kudo, T. and Richardson, J. (2018). SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *EMNLP*.
- [57] Kulkarni , P., Mahabaleshwarkar , A., Kulkarni , M., Sirsikar , N., & Gadgil, K. (2019). *Conversational AI: An Overview of Methodologies, Applications & Future Scope*. IEEE Xplore. <https://ieeexplore.ieee.org/document/9129347/>.
- [58] *Language and thought* (2020) <http://srwyatt.tripod.com/languagethought.htm>.
- [59] Lefèvre F (2006) A DBN-based multi-level stochastic spoken language understanding system. In *IEEE spoken language technology workshop*, Palm Beach, Aruba, 10–13 Dec 2006, pp 82– 85. doi:10.1109/slt.2006.326822
- [60] Lefèvre F (2007) Dynamic bayesian networks and discriminative classifiers for multistage semantic interpretation. In: *Proceedings of the IEEE international conference on acoustics, speech and signal processing (ICASSP’07)*, vol 4. Honolulu, HI, USA, 15–20 Apr 2007, pp 13–16. doi:10.1109/ICASSP.2007.367151

- [61] Lexico Dictionaries. (n.d.). CHATBOT: Definition of CHATBOT by Oxford Dictionary on Lexico.com also meaning of CHATBOT. Lexico Dictionaries | English.
<https://www.lexico.com/en/definition/chatbot>.
- [62] Lin L., D'Haro L., Banchs R. A web-based platform for collection of human-chatbot interactions (2016), pp. 363-366. doi: 10.1145/2974804.2980500
- [63] Lokman, A. S., Ameen, M. A., & Kapoor, S. (2019). Modern Chatbot Systems: A Technical Review . Proceedings of the Future Technologies Conference (FTC) 2018.
https://link.springer.com/chapter/10.1007/978-3-030-02683-7_75.
- [64] Marcus, M. P., Santorini, B., and Marcinkiewicz, M. A. (1993). Building a large annotated corpus of English: The Penn treebank. *Computational Linguistics* 19(2), 313–330.
<https://www.aclweb.org/anthology/J93-2004>
- [65] Marietto, M., et al.: Artificial intelligence markup language: a brief tutorial. *Int. J. Comput. Sci. Eng. Surv.* 4 (2013). <https://doi.org/10.5121/ijcses.2013.4301>
- [66] McTear M (2004) Spoken dialogue technology: toward the conversational user interface. Springer, New York. doi:10.1007/978-0-85729-414-2
- [67] McTear, M., Callejas, Z., & Griol, D. (2018). *The Conversational Interface Talking to Smart Devices*. Springer International Publishing.
- [68] Mikolov, T., Karafiat, M., Burget, L., Cernocký, J., and Khudanpur, S. (2010). Recurrent neural network based language model. *INTERSPEECH* 2010.
- [69] Molnár, G., Zoltán, S.: The role of chatbots in formal education. Presented at the 15 September 2018
- [70] Naone, E. (2020, April 2). Software That Learns from Users. *MIT Technology Review*.
<https://www.technologyreview.com/2007/11/30/98176/software-that-learns-from-users/>.
- [71] Nimavat, K., & Champaneria, T. (2017, July). Chatbots: An overview Types, Architecture, Tools and Future Possibilities. https://www.researchgate.net/profile/Ketakee-Nimavat/publication/320307269_Chatbots_An_overview_Types_Architecture_Tools_and_Future_Possibilities/links/59dcf2fdaca2729031b77acf/Chatbots-An-overview-Types-Architecture-Tools-and-Future-Possibilities.pdf.
- [72] Nivre, J., de Marneffe, M.-C., Ginter, F., Goldberg, Y., Hajic, J., Manning, C. D., McDonald, R., Petrov, S., Pyysalo, S., Silveira, N., Tsarfaty, R., and Zeman, D. (2016b). Universal Dependencies v1: A multilingual treebank collection. *LREC* .
<https://www.aclweb.org/anthology/L16-1262>
- [73] Nt T. An e-business chatbot using AIML and LSA (2016), pp. 2740-2742
<https://doi.org/10.1109/ICACCI.2016.7732476>
- [74] Nuruzzaman M., Hussain O.K. A survey on chatbot implementation in customer service industry through deep neural networks. 2018 IEEE 15th international conference on e-business engineering, IEEE, Xi'an (2018), pp. 54-61. doi: 10.1109/ICEBE.2018.00019
- [75] Oneremission, (2020) <https://keenethics.com/project-one-remission>.

- [76] Personal Digital Assistant – Cortana Home Assistant Microsoft.
<https://www.microsoft.com/en-us/cortana>
- [77] Picard, R. W. (2000). Affective computing. The MIT Press.
- [78] Pietroszek, Krzysztof. "Providing Language Instructor with Artificial Intelligence Assistant". International Journal of Emerging Technologies in Learning (2007). Found at: https://www.researchgate.net/publication/26490167_Providing_Language_Instructor_with_Artificial_Intelligence_Assistant
- [79] Ramesh K., Ravishankaran S., Joshi A., Chandrasekaran K. A survey of design techniques for conversational agents. Information, communication and computing technology: Vol. 750, Springer, Singapore (2017), pp. 336-350. doi: 10.1007/978-981-10-6544-6_31
- [80] Raymond C, Riccardi G (2007) Generative and discriminative algorithms for spoken language understanding. In: Proceedings of the 8th annual conference of the international speech communication association (Interspeech 2007), Antwerp, Belgium, 27–31 Aug, pp 1605–1608. http://www.isca-speech.org/archive/archive_papers/interspeech_2007/i07_1605.pdf.
- [81] Reiter E, Dale R (2000) Building natural language generation systems. Cambridge University Press, Cambridge. doi:10.1017/cbo9780511519857
- [82] Reiter E (2007) An architecture for data-to-text systems. In: Proceedings of ENLG-2007 11th European workshop on natural language generation. Schloss Dagstuhl, Germany, pp 97–104. doi:10.3115/1610163.1610180
- [83] Ruane, E., Birhane, A., & Ventresque, A. (n.d.). Conversational AI: Social and Ethical Considerations. http://ceur-ws.org/Vol-2563/aics_12.pdf.
- [84] Rumelhart DE, Hinton GE, Williams RJ (1986) Learning internal representations by error propagation. In: Rumerhart DE, McClelland JL (eds) Parallel distributed processing: explorations in the microstructure of cognition, vol 1. MIT Press, Cambridge, pp 318–362. <http://dl.acm.org/citation.cfm?id=104293>
- [85] Schuster, M. and Nakajima, K. (2012). Japanese and korean voice search. ICASSP.
- [86] Sennrich, R., Haddow, B., and Birch, A. (2016). Neural machine translation of rare words with subword units. ACL.
- [87] Sensely: character-based enterprise virtual assistant platform (2020) Sensely website: <https://www.sensely.com/>.
- [88] Shawar B., Atwell E. Chatbots: Are they really useful? LDV Forum, 22 (2007), pp. 29-49
- [89] Shum, Hy., He, Xd. & Li, D. "From Eliza to XiaoIce: challenges and opportunities with social chatbots". Frontiers Inf Technol Electronic Eng 19, 10–26 (2018). <https://doi.org/10.1631/FITEE.1700826>
- [90] Singhal, A. (2012, May 16). Introducing the Knowledge Graph: things, not strings. Google. <https://blog.google/products/search/introducing-knowledge-graph-things-not/>.
- [91] So, D. R., Liang, C., & Le, Q. V. (2019, May 17). The Evolved Transformer. arXiv.org. <https://arxiv.org/abs/1901.11117>.

- [92] A state-of-the-art open source chatbot. Facebook AI. (n.d.).
<https://ai.facebook.com/blog/state-of-the-art-open-source-chatbot/>.
- [93] Sutton, Richard S., Doina Precup, and Satinder P. Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2):181-211. [An earlier version appeared as Technical Report 98-74, Department of Computer Science, University of Massachusetts, Amherst, MA 01003]
- [94] The top 12 healthcare chatbots (2020) The Medical Futurist
 website: <https://medicalfuturist.com/top-12-health-chatbots>.
- [95] Top 30 successful chatbots of 2021 & Reasons for their success. AIMultiple. (2021, March 14). <https://research.aimultiple.com/top-chatbot-success/>.
- [96] Traum DR, Larsson S (2003) The information state approach to dialog management. In: Smith R, Kuppevelt J (eds) *Current and new directions in discourse and dialog*. Kluwer Academic Publishers, Dordrecht, pp 325–353. doi:10.1007/978-94-010-0019-2_15
- [97] Trivedi A., Gor V., Thakkar Z. Chatbot generation and integration: A review *International Journal of Advance Research, Ideas and Innovations in Technology*, 5 (2) (2019), pp. 1308-1311 XX.XXX/IJARIIT-V5I2-1840
- [98] Tur G, Deng L (2011) Intent determination and spoken utterance classification. In: Tur G, de Mori R (eds) *Spoken language understanding: systems for extracting semantic information from speech*. Wiley, Chichester, pp 93–118. doi:10.1002/9781119992691.ch4
- [99] Turing, A.M.: Computing machinery and intelligence. *Mind* 59, 433–460 (1950). <https://doi.org/10.1093/mind/LIX.236.433>
- [100] Wallace, R.S.: The anatomy of A.L.I.C.E. In: Epstein, R., Roberts, G., Beber, G. (eds.) *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*, pp. 181–210. Springer, Cham (2009). https://doi.org/10.1007/978-1-4020-6710-5_13
- [101] Weizenbaum, J. (1966, January 1). ELIZA-a computer program for the study of natural language communication between man and machine. *Communications of the ACM*. <https://dl.acm.org/doi/10.1145/365153.365168>.
- [102] What exactly is Alexa? Where does she come from? And how does she work? <https://www.digitaltrends.com/home/what-is-amazons-alexa-and-what-can-it-do/>
- [103] What Is Wolfram|Alpha? Wolfram. (n.d.). <https://www.wolframalpha.com/tour/>.
- [104] Wilcox, B. (n.d.). bwilcox-1234/ChatScript. GitHub. <https://github.com/bwilcox-1234/ChatScript>.
- [105] Wu, Y., Wu, W., Xing, C., Zhou, M., Li, Z.: Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-based Chatbots. arXiv:1612.01627 [cs] (2016)
- [106] Yih, W.-t., Chang, M.-W., He, X., and Gao, J. (2015). Semantic parsing via staged query graph generation: Question answering with knowledge base. In *ACL*, pages 1321–1331.

- [107] Young S (2002) Talking to machines (statistically speaking). In: Proceedings of the 7th international conference on spoken language processing, Denver, Colorado, USA, 16–20 sept 2002, pp 9–16. http://www.isca-speech.org/archive/archive_papers/icslp_2002/i02_0009.pdf
- [108] Young S, Gašić M, Keizer S, Mairesse F, Schatzmann J, Thomson B, Yu K (2010) The Hidden Information State model: a practical framework for POMDP-based spoken dialogue management. *Comp Speech Lang* 24(2):150–174. doi:10.1016/j.csl.2009.04.001
- [109] Youper—emotional health assistant powered by ai (2020) Youper website: <https://www.youper.ai/>.
- [110] Your.md - health guide and self-care checker(2020)Your.MD website: <https://www.your.md/>.
- [111] Zemčík, T. (2019, October 22). (PDF) A Brief History of Chatbots. ResearchGate. https://www.researchgate.net/publication/336734161_A_Brief_History_of_Chatbots.
- [112] Zhang S., Dinan E., Urbanek J., Szlam A., Kiela D., Weston J. Personalizing dialogue agents: I have a dog, do you have pets too? (2018) ArXiv:1801.07243 [Cs]. Retrieved from <http://arxiv.org/abs/1801.07243>
- [113] Zhou, Li, Jianfeng, Gao, Di, Li, and Heung-Yeung, Shum. "The Design and Implementation of XiaoIce, an Empathetic Social Chatbot". *Computational Linguistics* 46, no.1 (2020): 53-93. https://doi.org/10.1162/coli_a_00368
- [114] Zhu W., Chowanda A., Valstar M. (2016) Topic Switch Models for Dialogue Management in Virtual Humans. In: Traum D., Swartout W., Khooshabeh P., Kopp S., Scherer S., Leuski A. (eds) *Intelligent Virtual Agents. IVA 2016. Lecture Notes in Computer Science*, vol 10011. Springer, China.