



# A survey on construction and enhancement methods in service chatbots design

Zhenhui Peng<sup>1</sup> · Xiaojuan Ma<sup>1</sup>

Received: 27 May 2019 / Accepted: 29 August 2019  
© China Computer Federation (CCF) 2019

## Abstract

Chatbots are being widely applied in many service industries to help schedule meetings, online shopping, restaurant reservations, customer care and so on. The key to the success of the service chatbots design is to provide satisfying responses to the given user's requests. This survey aims to provide a comprehensive review of chatbots construction and enhancement methods. We first introduce major techniques for the three core design philosophies, which are rule-based, retrieval-based and generation-based methods, followed by a brief summary of the evaluation metrics. Then we present methods to enhance service chatbot's capabilities with either an ensemble of multiple chatbots, collaborating with human workers or learning from users. Finally, in future directions we discuss the promising response generation models for chatbots using the recent progress in the transformer and contextual embeddings, as well as potential ways to construct a chatbot with personality to achieve a better user experience.

**Keywords** Chatbots · Chatbot design methods · Human–chatbot collaboration · Chatbot survey

## 1 Introduction

Chatbots, also known as dialogue/conversational systems/agents, are computer programs which conduct a natural conversation with users via speech or text (Mauldin 1994; Shawar and Atwell 2007). Initially, chatbots were mainly designed for passing Turing test (Turing 1950) or for fun, such as ELIZA conducting psychology interviews (Weizenbaum 1966), and Microsoft Xiaoice acting as a poet, a singer, a kid audio-books reciter or a journalist (WIKIPEDIA 2018). With the advances in machine learning and the rapid development of social media over the past decade, service chatbots which are designed to provide 24 / 7, easy-to-access services to the users in specific domains, have been a trend recently (Magazine 2018). They can either be task-oriented chatbots which get information from the user to help complete some structured tasks like scheduling meetings (Cranshaw et al. 2017), online shopping (Jain et al. 2018) and restaurant reservations (Luo et al. 2018), or chatbots that accomplish unstructured tasks like customer care in social

media (Xu et al. 2017; Hu et al. 2018) and answering human resource questions (Liao et al. 2018).

Despite the popularity of chatbots, designing service chatbots to provide satisfying responses to the given users' requests still remains a crucial challenge. Given any user's speech or text as an input request, a service chatbot is required to understand the input and give appropriate (e.g., the same topic, make sense), helpful (e.g., contains useful and concrete information) and even tone-aware (e.g., conveys feelings like empathy and passion) responses (Xu et al. 2017; Hu et al. 2018). One popular approach to construct such a chatbot is the frame-based method (Young et al. 2013; Mesnil et al. 2015), which predefines the structure of a dialog state as a set of slots to be filled during a conversation and gives responses based on some hand-crafted rules. For example, in a restaurant booking chatbot, the slots can be the reserved date, the cuisine, or the location of the restaurant, while the rules can be to ask questions until all the slots are filled and to generate responses based on the template. However, such kinds of rule-based chatbots are limited to a narrow domain as manually constructing and updating rules for complex systems are usually expensive.

Different from the rule-based chatbots, service chatbots that are based on data-oriented approaches can handle more types of user's requests (Song et al. 2018). They can either

✉ Zhenhui Peng  
zpengab@connect.ust.hk

<sup>1</sup> Hong Kong University of Science and Technology,  
Clear Water Bay, Hong Kong

retrieve an existing response from a pre-compiled dataset (i.e., retrieval-based) (Yan et al. 2016), or generate a new response word by word based on the input sequence (i.e., generation-based) (Serban et al. 2016). Retrieval-based chatbots can generate literally correct responses to the user, but they are limited by the size of the corpus as they can not generate new responses. While generation-based chatbots could address this problem, they are prone to give responses that are not grammatically correct or contain no useful information (Song et al. 2018).

To build a more powerful chatbot that can handle a broader scope of service requests, chatbot designers should also consider how to enhance a chatbot's capabilities over time. One straightforward way is to make use of multiple chatbots to cover cross-domain requests. The multi-chatbot framework learns to select one response from the responses proposed by each chatbot to address user's requests (Qiu et al. 2017). However, such a framework can not learn new responses or actions outside the capabilities of the chatbot ensemble. Moreover, in complex real-world scenarios, fully automatic chatbots are potentially problematic, as they do not share the same experience as human on how to avoid serious mistakes that might negatively affect users. Therefore, some chatbot designers design chatbots with humans in the loop (Grudin and Jacques 2018). For example, chatbots can be designed to work with human workers collaboratively, and gradually learn to deal with unknown requests (Luo et al. 2018). Chatbots can also be designed to learn from a user's demonstration, where the user teaches the chatbot new skills to provide more personalized services (Li et al. 2017).

### 1.1 Contribution of this work

In this work, we try to systematically present and analyze notable works for chatbots design from both the perspectives of construction and enhancement methods. We aim to inspire more flexible ways to design service chatbots that can meet different user needs. The main contributions of this work are outlined below:

- A comprehensive analysis of different chatbot construction techniques along with their comparison and suitable usage scenarios.
- A brief summary of automatic and human-based metrics for chatbot design evaluations.
- Various chatbot enhancement techniques ranging from multi-chatbots to human-in-the-loop are presented along with their comparison.
- The advanced methods (e.g., transformer) that could help chatbot design is presented.
- Data-driven techniques for chatbot personality design is discussed.

### 1.2 Link and distinction to related surveys

This survey is inspired by several related papers (Ramesh et al. 2017; Song et al. 2018; Chen et al. 2017) that classify the chatbots into rule-based, retrieval-based and generation-based. Apart from the high-level comparison like those papers, our survey uses equations and example models to illustrate the detailed mechanism behind each technique for better understanding of its strengths and drawbacks. Furthermore, to the best of our knowledge, this is the first survey that categorizes the chatbot enhancement methods. In addition, this work provides guidelines for developers to select the appropriate techniques that fit their requirements in terms of performance, usage scenarios, or cost.

### 1.3 Taxonomy and organization of this article

The taxonomy of this survey is showed in Fig. 1.

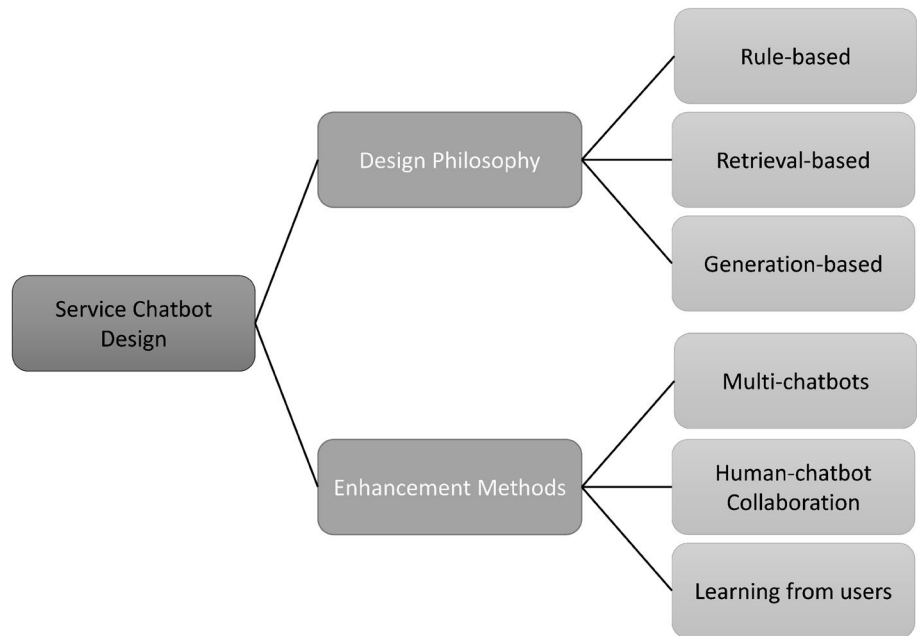
*Design philosophy* The core principle of the chatbot architecture that is designed to give a response based on the request. According to whether it is data-driven and whether it can generate new responses, we classify the chatbot architecture into three categories: (1) Rule-based chatbot, which identifies the characteristic variables of the utterance and gives a pre-defined response based on the variables and hand-crafted rules; (2) Retrieval-based chatbot, which uses the input utterance as the query to search for candidate responses from the dataset by some matching metrics; (3) Generation-based chatbot, which treats the conversation as the input-output mapping problem and learns to generate responses after training on a large amount of data.

*Enhancement methods* The design methods that can enable the chatbots to handle more kinds of requests or extend their functions. According to whether the service chatbots are mainly powered by the ensemble of multiple chatbots, human workers, or users, we classify these methods as: (1) Multi-chatbots methods, which combine rule-based, retrieval-based or generation-based chatbots together into a framework; (2) Human–chatbot collaboration methods, which enable the chatbot to work with and learn from human workers; (3) Learning from users, which improves the chatbot via user's direct feedback or instructions.

The organization is as below:

Section 2 focuses on three kinds of core design philosophies in building a chatbot, i.e., rule-based, retrieval-based, and generation-based methods. We showcase the principle of the representative methods in each design philosophy, as well as compare their advantages and

**Fig. 1** The taxonomy of this survey



disadvantages in terms of complexity and user experience in this section. Then in Sect. 3, we focus on the methods that incorporate different bots, human workers or users into the chatbot design to enhance its capability. Finally in Sect. 4, we discuss the opportunities of using advanced techniques like Transformer in chatbot design and how to incorporate personality into chatbot using data-driven methods.

## 2 Core design philosophy

Given user’s speech or text as the input requests, the chatbots usually pre-process the requests to be a text sequence and fit into the response model to get the responses (Young et al. 2013; Serban et al. 2017). Therefore, in the rest of the survey, we consider both the requests and responses as

textual utterances. There are a number of ways to construct the model that generate responses to the input requests. In this section, we introduce three kinds of core design philosophies (i.e., rule-based, retrieval-based and generation-based) and several popular construction methods under each philosophy (Table 1).

### 2.1 Rule-based methods

Rule-based chatbots are mainly built on manually constructed rules. These manually constructed rules can be the patterns in the input request, the “if-then” logic that triggers the action or response, or the template that is about to be filled in the response. Dating back to 1966, the development of chatbots started from analyzing the input sentence based on decomposition rules which are triggered by key words in a sentence (Weizenbaum 1966). A typical example provided

**Table 1** Chatbot construction techniques overview

Philosophy	Techniques	Representative papers
Rule-based	Pattern matching	Weizenbaum et al. (1966), Colby et al. (1975) and Wallace (2009)
	Modular task-oriented system	Chen et al. SIGKDD (2017)
Retrieval-based	TF-IDF	Lowe et al. SIGDIAL (2015)
	DNN-based	Lu et al. NIPS (2013) and Hu et al. NIPS (2014)
	RNN-based	Lowe et al. SIGDIAL (2015) and Zhou et al. EMNLP (2016)
Generation-based	Statistical Machine Translation	Ritter et al. EMNLP (2011)
	Seq2Seq	Sutskever et al. NIPS (2014) and Xu et al. CHI (2017)
	Seq2Seq + attention mechanism	Shang et al. ACL (2015)
	Seq2Seq + hierarchical structure	Serban et al. AAAI (2016)
	Seq2Seq + memory network	Ghazvininejad et al. AAAI (2018)

by a prominent chatbot “ELIZA” is that the user asks “It seems that you like me”. ELIZA can only recognize the words “you” and “me”, but does not know what “It seems that” and “like” mean. Based on its decomposition rule “O YOU O ME” (where O stands for indefinite number of words) and the reassembly rule “WHAT MAKES YOU THINK I 3 YOU” (where 3 stands for the third component of the the subject decomposition, here is “like”), the ELIZA will reply “What makes you think I like you”. Following the same pattern-response rules as ELIZA, another chatbot “Parry” passed the Turing test (Turing 1950) in 1972 (Colby 1975). Parry adds some affect variables like “fear”, “anger” and “mistrust” to the more complex rules. For example, when a user mentions Parry, Parry decreases fear if mistrust is low and increases anger if mistrust is high. Such rules keep the conversation going and make people feel that they are really chatting to another person.

Extended from ELIZA, Richard Wallace developed the Artificial Linguistic Internet Computer Entity (ALICE) in 1995 (Wallace 2009). ALICE applies the Artificial Intelligence Markup Language (AIML) and is primarily designed for historical and philosophical ruminations on human consciousness. The key part in AIML is the category that forms the unit of knowledge. As the example shown below, a category combines a pattern (e.g., question or stimulus), a template (e.g., answer or response), and optional context (e.g., a previous utterance).

```
< category >
< pattern > YES < /pattern >
< that > DO YOU LIKE MOVIES < /that >
< template > What is your favorite movie? < /template >
< /category >
```

Since the 1990s, a lot of research has been done on designing similar rule-based chatbots to provide services in specific

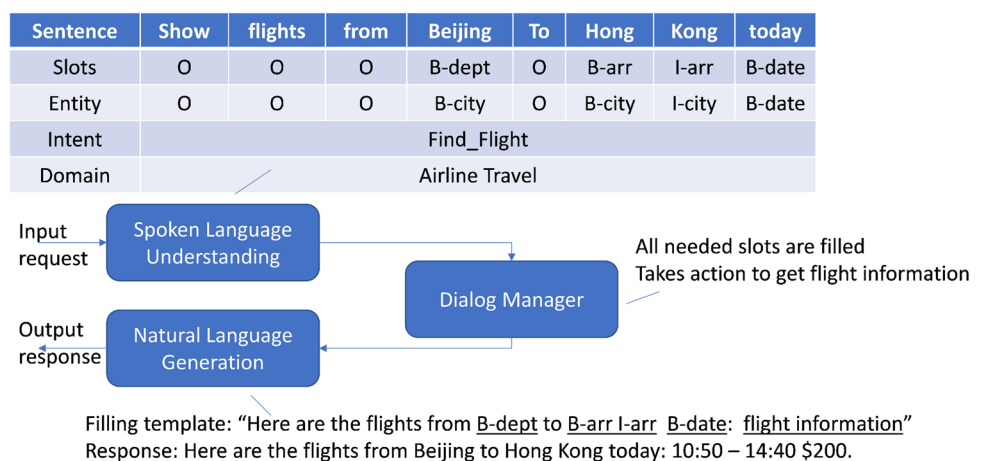
domains (Walker et al. 2001; Chai et al. 2001). These chatbots are also known as modular task-oriented dialog systems, which guide a user to finish some well-structured tasks such as restaurant reservations and movie bookings. As shown in Fig. 2, modular task-oriented dialog system mainly consists of three components (Shum et al. 2018):

1. Spoken Language Understanding (SLU), which turns user utterance into user intention and slot-values, and outputs structured user action;
2. Dialogue Manager (DM), which tracks the dialog state based on past state, action and current user action, and outputs system action based on some policies;
3. Natural Language Generation (NLG), which turns a system action into natural language and outputs it to the user.

There are multiple available platforms such as Microsoft LUIS (2018), IBM Watson Assistant (2018) and Dialogflow (2018) that provide easy-to-use SLU, DM and NLG services to help chatbot designers build service chatbots. For instance in designing an airline travel planning chatbot (Fig. 2), the designer needs to define the intents (i.e., user’s goal or purpose) and entities (i.e., terms or objects that provide clarification or specific context for a particular intent) that might occur in the user’s input. The designer also needs to define the rules to ask for the missing slots in the DM module and define the templates to generate the task-completion information in the NLG module. Inside each of the SLU, DM and NLG services provided by these platforms, there are techniques like Recurrent Neural Network (Schuster and Paliwal 1997), Long Short-Term Memory (Hochreiter and Schmidhuber 1997) and Reinforcement Learning (Kaelbling et al. 1996) to handle each module. A survey of task-oriented chatbots can be found in Chen et al. (2017).

Building a modular task-oriented chatbot could be easy in the domains that have expert knowledge and a well-defined

**Fig. 2** Main components and example of a modular task-oriented chatbot



structure. For example, in an online shoe shopping chatbot (Jain et al. 2018), the shoes have limited features such as price, color, material, style and brand, which can easily be used to filter the shoes interactively. Also in the restaurant reservation, food ordering, and movie ticket booking domains, the users' goals are clear and the designer can easily design the logic to get the needed information (e.g., which restaurant/food/movie, when, and how many people) for task completion. However, as the knowledge space gets larger and users' expectations get higher [e.g., wanting some playful interaction like chitchat (Liao et al. 2018)], it becomes difficult and expensive to anticipate users' intentions and design the rules to handle them. Therefore, researchers seek to use data-driven methods to automatically build up the service chatbots, which are discussed in next two sections.

## 2.2 Retrieval-based methods

The retrieval-based chatbots select the response that best matches the users' requests by searching a pre-constructed conversational repository (Lee et al. 2000; Ji et al. 2014; Yan et al. 2016). The key to retrieval-based chatbots is request-response matching. Given a request  $q$  and a repository of request/response pairs, there are two types of strategies to retrieve the response  $r'$  (Ritter et al. 2011):

1. Request-based strategy  $[r_{\text{argmax}_i \text{sim}(q, q_i)}]$ , which retrieves the response  $r_i$  whose associated request  $q_i$  is most similar to the user's input  $q$ ;
2. Response-based strategy  $[r_{\text{argmax}_i \text{sim}(q, r_i)}]$ , which retrieves the  $r_i$  that is most similar to the user's input  $q$ .

To compute the similarity between the  $q$  and  $r_i$ , it is prescribed to transform them into some numeric or vector representation. As early as in 2000, Isbell et al. developed a retrieval-based chatbot Cobot which interacted with the users in a game LambdaMOO (Lee et al. 2000). Cobot used a request-response matching method based on word occurrence, in which each potential response is weighted by the number of words that match to the words in the request. A more popular occurrence-based method used in chatbots is TF-IDF (Lowe et al. 2015). TF-IDF stands for "term frequency—inverse document frequency" and it measures the importance of a word in a document (i.e., the request in chatbot case) in relation to the whole repository. The "term frequency" is the number of times the word appears in a given request, while the "inverse document frequency" puts a penalty on how often this word appears elsewhere in the repository. The final score is calculated as:

$$\text{tf-idf}(w, q, D) = f(w, q) \times \log \frac{N}{|\{q \in D : w \in q\}|} \quad (1)$$

where  $f(w, q)$  is the number of times word  $w$  appears in the request  $q$ ,  $N$  is the total number of requests or responses in the repository based on a different strategy,  $D$  represents the collection of the requests or responses, and  $|\{q \in D : w \in q\}|$  is the number of requests or responses in which  $w$  appears. Take the response-based strategy as an example, for the request and each candidate responses, their TF-IDF vectors are calculated by concatenating the all tf-idf scores together. The candidate response with the highest cosine similarity to the request vector will be selected as the final response.

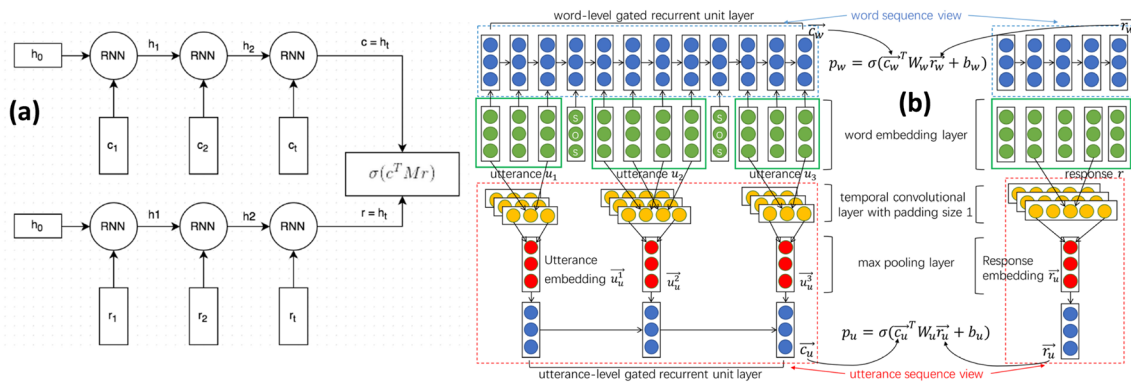
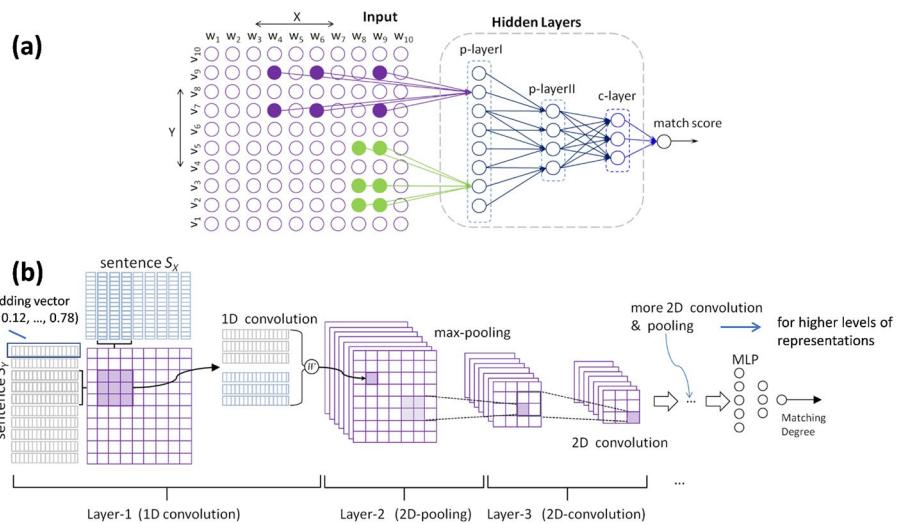
Although TF-IDF is simple to use without the need of training on dataset, it is not accurate enough since it does not consider the location of the word in the sentence and can not efficiently represent the importance of the words. With the success of deep neural networks (DNN) in computer vision, a couple of works have sought to use DNN-based methods to handle request-response matching in natural language processing (Lu and Li 2013; Hu et al. 2014). In 2013, Lu et al. proposed a DNN matching architecture DEEPMATCH to attack the problem of matching short texts (Lu and Li 2013). As shown in Fig. 3a, DEEPMATCH models the similarity between two utterances by first constructing a two-dimensional interaction space of their bag-of-words vectors, then going through two patch-induced layers and one fully connected layer, finally summarizing the decisions with the logistic regression unit. Their experiments on a traveling-related (Question, Answer) pairs dataset and a (Weibo, comment) dataset showed that DEEPMATCH's retrieval performance is better than three inner-product based models.

Later in 2014, Hu et al. proposed a convolutional DNN matching model ARC-II Fig. 3b, which first models all the possible combinations of the word embedding vectors [trained with the Word2Vec (Mikolov et al. 2013)] of two utterances via one-dimensional (1D) convolutions, then performs a 2D max-pooling to shrink the size of the representation by half, and uses more 2D convolution and pooling layers to obtain higher levels of representation. The final level of representation fits into a multi-layer perception (MLP) (Bengio 2009) to get the matching degree of two utterances. Compared with DEEPMATCH, ARC-II can better represent the patterns shared in two utterances and thus perform better in Weibo (tweet, response) retrieval task.

Besides CNN, the recurrent neural network (RNN) (Schuster and Paliwal 1997) is also widely used in the request-response matching deep architectures (Lowe et al. 2015; Zhou et al. 2016, 2018). In their released Ubuntu dialogue corpus which is about technical support for Ubuntu-related problems (Lowe et al. 2015), Lowe et al. proposed an RNN model (Fig. 4a) to capture the representation for the request ( $c$  in their case) and response ( $r$ ), and then use the dot product to measure the similarity. The RNN can model the time-dependent sequence data like the sentence, and can be formally represented by:



**Fig. 3** Two DNN-based matching models: **a** DEEPMATCH [from Lu and Li (2013), with permission], which represents the interaction of two utterances via an overlap in their bag-of-words vectors; **b** ARC-II [adapted from Hu et al. (2014)], which represents the interactions via a convolution of their word embedding vectors



**Fig. 4** Two RNN-based matching models: **a** The RNN model that only considers single-turn information [from Lowe et al. (2015), with permission]; **b** the multi-view model that considers both word-level

and utterance-level dependencies in multi-turn dialogues [adapted from Zhou et al. (2016)]

$$h_t = f(h_{t-1}, x_t) = f(W_h h_{t-1} + W_x x_t) \tag{2}$$

where  $h_t$  is the hidden state at time step  $t$  (e.g., index of word),  $x_t$  is the observed variable (e.g., word) at the same time step, and  $W_h, W_x$  are the weights learned from the training data. However, a conventional RNN can easily fail to capture the long-term dependence due to the gradient vanishing (or exploding) problem. To relieve this problem, Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber 1997) or Gated Recurrent Unit (GRU) (Chung et al. 2014) is commonly used to replace the hidden units in the RNN models (Lowe et al. 2015; Zhou et al. 2016). In Lowe’s work 2015, the LSTM model outperforms the RNN model and the TF-IDF method in terms of the Recall@k metric in the response selection task.

Nevertheless, the models above only consider single-turn information to retrieve a response from the repository, while in the conversation the chatbot’s response at the current turn

usually needs the information from previous turns. To this end, Zhou et al. proposed a multi-view response selection model (Fig. 4b) which uses GRU to learn the word-level semantics and dependencies in the connected utterances (i.e., context,  $u_1, u_2, u_3$  in this case), as well as the utterance-level semantic and discourse information in a convolutional manner (Zhou et al. 2016). In the same response selection task, this model performs better than the LSTM model in Lowe’s work, and better than the GRU model that only learns the word-level or utterance-level information. Recently, Zhou et al. proposed a deep attention matching (DAM) network which extends the attention mechanism of Transformer (Vaswani et al. 2017) in the self-attention (for intra word-level dependencies) and in the cross-attention (for dependence between a latently matched pair) (Zhou et al. 2018). The DAM achieves the best result in response selection task in the Ubuntu corpus (Lowe et al. 2015) compared to RNN models.

Although not all of these request-response matching methods are originally designed for conversation, they are effective in the conversational studies (Ritter et al. 2011; Xu et al. 2017). The advantages of retrieval-based chatbots are that they need much fewer hand-crafted features than the rule-based chatbots and promise that the response is grammatically correct and diverse. However, they can easily provide an inappropriate response which is not really in the context of the request. It is also hard to persist with the same tone since the repository is diverse. Therefore, retrieval-based methods are more commonly used in chit-chat chatbots like Xiaoice WIKIPEDIA (2018) in open domain and question-answering chatbots which answer questions in specific domains such as travel (Lu and Li 2013) or give technical support (Xu et al. 2017; Lowe et al. 2015).

### 2.3 Generation-based methods

The generation-based chatbots synthesize a new sentence word by word as the response to the user's requests (Sutskever et al. 2014). Originally, research on generation-based chatbots was mainly in chit-chat-style, open domains (Ritter et al. 2011; Serban et al. 2016; Shang et al. 2015). However, they have huge potential to be deployed directly for the tasks [e.g., customer care service (Xu et al. 2017; Hu et al. 2018)] which do not have a measurable goal if the task-related data is available. Therefore, in this section, we summarize the generation-based construction methods of general chatbots (either domain-specific or open domain).

The development of generation-based chatbots is inspired by the work in machine translation. In 2011, Ritter et al. used

a phrase-based statistical machine translation (SMT) method in the response generation for the first time, and showed that their SMT techniques are better-suited than retrieval-based methods on the Twitter dataset (Ritter et al. (2011)). More specifically, the phrase-based SMT model considers the strong structural relation between many request-response pairs (e.g., “the soup smells delicious”—“I will bet it looks gorgeous too”), and extracts phrases like “smell-look” and “delicious-gorgeous” from the dataset to translate the request to the response. However, this model could work badly since the responses are often not semantically matched to the requests as in a translation. For example, it is likely that for a request the responses “having my fruit salad now”, “but it is 2 a.m. now” and “which restaurant” are both appropriate.

Sequence-to-sequence (seq2seq) Sutskever et al. (2014), another technique derived from machine translation, is shown to work better in response generation tasks. The seq2seq model contains an encoder which encodes a request word by word and represents it as a vector, and a decoder which decodes the vector and generates the response word by word (Fig. 5a). To better capture the dependencies in the utterance, LSTM (Hochreiter and Schmidhuber 1997), GRU (Chung et al. 2014), and techniques like bidirection and reverse order are commonly used to design the seq2seq models (Serban et al. 2016; Xu et al. 2017). Formally, a standard seq2seq model uses RNN to encode the request into a vector  $c$  (i.e., the hidden state of the last word) using the Eq. 2 (shown in the previous section). Then at the time step  $t$  in the decoder, the probability distribution  $p_t$  of candidate words for the response is:

$$s_t = f(y_{t-1}, s_{t-1}, c) \tag{3}$$

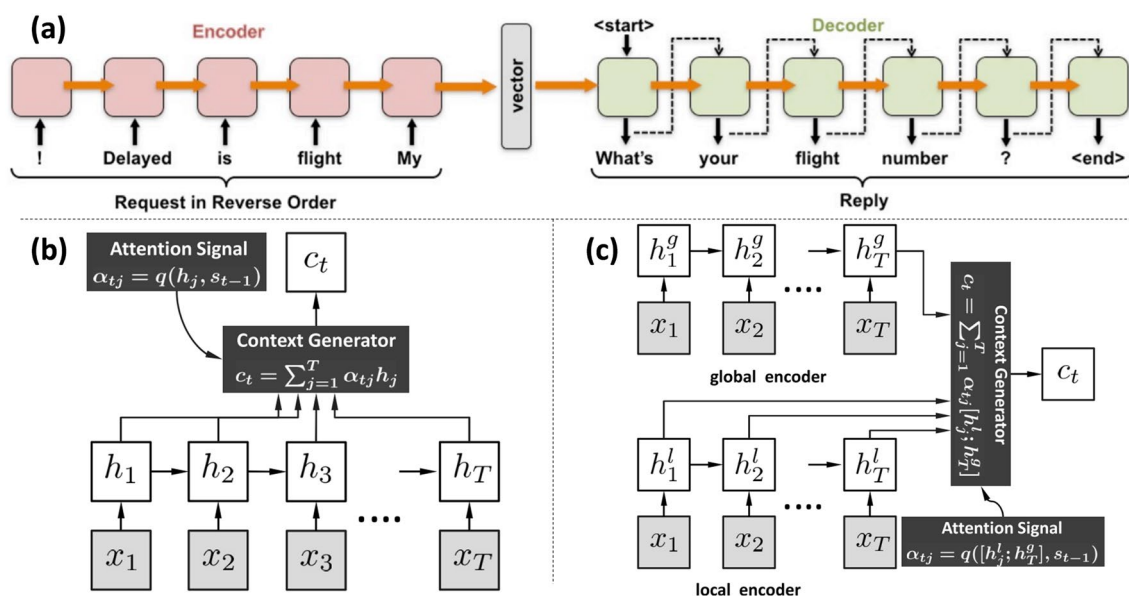


Fig. 5 The seq2seq models: **a** With LSTM neural networks, used in customer care services [from Xu et al. (2017), with permission]; **b** with attention mechanism (Shang et al. 2015); **c** combine (b) and (a) for a new attention signal (Shang et al. 2015)

$$p_t = \text{softmax}(s_t, y_{t-1}) \tag{4}$$

where  $s_t$  is the hidden state of the decoder RNN and  $y_t$  is the selected word for the response sequence at time step  $t$ . The objective function is:

$$p(y_1, \dots, y_{T'} | x_1, \dots, x_T) = p(y_1 | c) \prod_{t=2}^{T'} p(y_t | c, y_1, \dots, y_{t-1}) \tag{5}$$

where  $T$  and  $T'$  are the length of the request and response sequence respectively. Vinyals et al. showed that such a seq2seq model can generate more relatively preferable responses than the rule-based CleverBot, in both a domain-specific IT Helpdesk Troubleshooting dataset and a open-domain OpenSubtitles dataset (Vinyals and Le 2015). Xu et al. further suggested that it can generate more appropriate, more helpful and more empathetic responses than the retrieval-based method in branded customer care service (Xu et al. 2017). However, it could be problematic to simply use the hidden state of the last word as the context vector  $c$ , since each word in the response may strongly relate to different parts of the words in the request. The attention mechanism (Bahdanau et al. 2014) is then introduced to address this problem. As shown in Fig. 5b, for the word  $y_t$  at time step  $t$  in the decoder, there is a corresponding context vector  $c_t$  specifically contributing to it, where  $c_t$  is calculated by the sum of all the hidden states in the request sequence weighted by the attention signal  $\alpha$ . The attention mechanism can be formulated as the following two equations:

$$c_t = \sum_{j=1}^T \alpha_{tj} h_j \tag{6}$$

$$\alpha_{tj} = q(h_j, s_{t-1}) \tag{7}$$

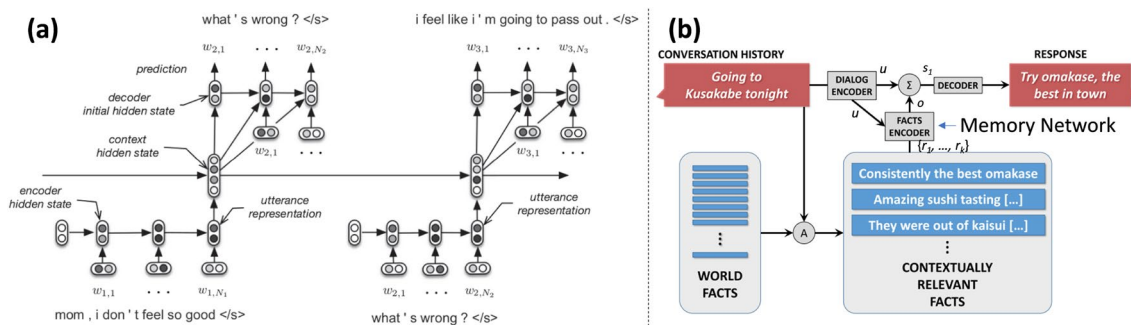
where  $q$  is a score function (e.g.,  $h_j^T W s_{t-1}$ ) that calculates how much the hidden state  $h_j$  relates to the previous hidden state  $s_{t-1}$  in the decoder at time step  $t$ . Shang et al. proposed

a neural responding machine that combines both the global encoder that only considers the last hidden state and the local encoder that considers all the hidden states in the request by concatenation (Fig. 5c) (Shang et al. 2015). The responses generated by this hybrid model in the Twitter-like Weibo dataset are perceived more suitable than those generated by the models that only consider global/local encoders, or by the models that are retrieval-based or SMT-based.

However, the works discussed above only consider generating a response based on one previous request, while the huge amount of information derived from previous turns of the dialogue are ignored. To incorporate dialogue history into response generation, Serban et al. adopted the hierarchical recurrent encoder-decoder (HRED) and showed that their model is competitive with other models in the *MovieTriples* in terms of word perplexity (Serban et al. 2016). The HRED uses RNN to not only encode the tokens appearing within the utterance, but also encode the temporal structure so far in the dialogue. It treats the dialogue as a sequence of utterances  $\{U_1, \dots, U_M\}$  and each utterance  $U_m$  with  $N_m$  tokens is represented as  $U_m = \{w_{m,1}, \dots, w_{m,N_m}\}$ , where  $w_{m,1}$  is a random variable taking values in the vocabulary  $V$  and represents the token at position  $n$  in utterance  $m$ . The probability distribution  $P$  over the set of all possible dialogues is:

$$P_\theta(U_1, \dots, U_M) = \prod_{m=1}^M P_\theta(U_m | U_{<m}) = \prod_{m=1}^M \prod_{n=1}^{N_m} P_\theta(w_{m,n} | w_{m,<n}, U_{<m}) \tag{8}$$

where  $\theta$  represents all parameters,  $U_{<m} = \{U_1, \dots, U_{m-1}\}$  is the input utterances so far,  $w_{m,<n} = \{w_{m,1}, \dots, w_{m,n-1}\}$  is the tokens preceding  $n$  in the utterance  $U_m$ . The response generation can be performed as in standard language modeling: sampling one word at a time from  $P_\theta(w_{m,n} | w_{m,<n}, U_{<m})$  conditioned on previous sampled words. As shown in Fig. 6a, the context hidden state  $c$  takes in the information from both the current utterance representation and the historical utterance information represented in the previous context hidden state.



**Fig. 6** The modified seq2seq models: **a** Hierarchical Recurrent Encoder–Decoder that adds historical utterances in [from Serban et al. (2016), with permission]; **b** Knowledge-Grounded Neural Con-

versation Model that adds external knowledge in [adapted from Ghazvininejad et al. (2018)]



To provide more appropriate and more informative responses to the requests, it usually needs to make use of external knowledge outside the dialogue history. For this purpose, Ghazvininejad et al. took advantage of the Memory Network (Sukhbaatar et al. 2015) and proposed to condition responses on both the dialogue history and external “facts” (Ghazvininejad et al. 2018). As shown in Fig. 6b, they used the request (e.g., conversation history) to retrieve all contextually relevant facts  $F = \{f_1, \dots, f_k\}$ , where each  $f_i$  is represented by bag of words as  $r_i$  to input to the facts encoder (i.e. Memory Network). Together with the summary of the request  $u$  (i.e.,  $h_T$ ) from the encoder RNN, the memory network calculates the initialized hidden state  $s_1$  of the decoder as:

$$m_i = Ar_i; c_i = Cr_i; p_i = \text{softmax}(u^T m_i); o = \sum_{i=1}^k p_i c_i; s_1 = o + u \tag{9}$$

where  $A, C$  are the parameters of the memory network. Specifically, they used multi-task learning (Luong et al. 2015) to train their model on the tasks with/without the fact encoder, and the task with the fact encoder but replace the response with each of the facts in the training data. They demonstrated that their model is informative and applicable in the Twitter dataset grounded by Foursquare tips (e.g., comments about restaurant and other commercial establishments).

Despite the progress in seq2seq style models, the generation-based chatbots are still prone to generate universal sentences (e.g. “I’m OK”) that lack helpful information. Also, a good generation-based service chatbot needs a well-collected dataset that is full of domain-specific conversation and less biased (Schlesinger et al. 2018). Nevertheless, the generation-based method is still the research hot spot on building customized and context-aware chatbots.

## 2.4 Evaluation methods

In this section, we look into the common methods which evaluate the quality of chatbot’s response given the request, rather than the methods which evaluate the chatbot’s objective performance like the task-success rate in spoken language understanding tasks and recall rate in response retrieval tasks. We summarize these methods as automatic metrics and human-based metrics.

### 1. Automatic metrics

Although it is still an open question to have a well-established method for automatic evaluation of the response quality, there are two common automatic metrics for reference. The first one is word perplexity originally for probabilistic language models (Bengio et al. 2003) and later adapted to evaluate some generative dialogue models (Serban et al. 2016; Ghazvininejad et al. 2018). For example, for a model

with parameters  $\theta$ , dataset with  $N$  triples  $\{U_1^n, U_2^n, U_3^n\}_{n=1}^N$ , the word perplexity can be defined as (Serban et al. 2016):

$$\exp\left(-\frac{1}{N_W} \sum_{n=1}^N \log P_{\theta}(U_1^n, U_2^n, U_3^n)\right) \tag{10}$$

where  $N_W$  is the number of tokens in the entire dataset and  $P_{\theta}(U_1^n, U_2^n, U_3^n)$  is the probability of regenerating the exact  $(U_1^n, U_2^n, U_3^n)$ . A better model will have lower perplexity. The word perplexity is appropriate because it measures the model’s ability to account for the syntactic structure of the dialogue (e.g., turn taking) and the syntactic structure of each utterance (e.g., punctuation marks).

The second commonly used automatic metric is the BLEU (bilingual evaluation understudy) Papineni et al. (2002) originally for machine translation and later adapted to the evaluation of some dialogue models (Xu et al. 2017; Ghazvininejad et al. 2018). BLEU grades an output response according to n-gram matches to the reference, i.e., the real response from a human worker. Basically, it can be defined as:

$$\text{BLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^N w_n \log p_n\right) \tag{11}$$

where BP is the brevity penalty on the length of the utterance,  $p_n$  is probability that the n-grams in a generated response occur in the real response,  $N$  is the max number of gram (normally 4) and  $w_n$  is the weight for each n-gram (normally  $\frac{1}{4}$ ). The higher BLEU score is indicative of a better model as the generated response is closer to the real one.

### 2. Human-based metrics

The automatic metrics can save time and money compared to human judgments, but can not really examine whether the generated response is appropriate, helpful and tone-aware. Liu et al. (2016) showed that BLEU could correlate poorly with human judgment. Currently human evaluation is still the more convincing method for judging the response quality and is widely applied in chatbot evaluation. One common human-based metric is conducting pair-wise comparison to let humans choose which of the two responses is more suitable, more appropriate, and more helpful, etc (Ritter et al. 2011; Shang et al. 2015). Another way is to ask people to rate the appropriateness, helpfulness, passion, etc, of each response given the same request on a 5-point or 7-point Likert Scale (Xu et al. 2017; Hu et al. 2018). The comparison among their proposed method and other methods are conducted on the average score of randomly sampled responses. Besides that, case studies with response example are often needed to analyze the quality of the response in depth (Ghazvininejad et al. 2018; Hu et al. 2018).

## 2.5 Summary

Based on the discussion above, we summarize the rule-based, retrieval-based and generation based chatbot construction methods in terms of major techniques, pros, cons and suitable application scenarios in Table 2. However, service chatbots built merely on these techniques are usually not enough to meet user's various needs. To handle more types of requests (e.g., cross-domain) and expand its functionality, chatbots need to be designed to enhance themselves over time. In the next section, we discuss some of these enhancement methods (Table 3).

## 3 Enhancement of chatbot

### 3.1 Multi-chatbot design

Most available service chatbots are designed to perform tasks in highly specific domains like flight and hotel bookings. To satisfy the user's multiple-domain needs, there is a need to combine available chatbots that have different expertise (Ferrucci et al. 2010). Besides, even in the same domain, chatbots with different design philosophies have their own strengths and drawbacks (Table 2). Some researchers then seek to assemble chatbots built on different methods (e.g., retrieval-based and generation-based) to absorb their merits (Song et al. 2018).

In a typical multi-chatbot framework, the responses from the different chatbots form the candidate responses pool for the same given request, and the key is to design a proper re-ranker (or response selection policy) which scores all candidates to pick the highest-scored response (Song et al. 2018; Qiu et al. 2017; Serban et al. 2017). In 2016 Amazon Alexa Prize competition, Serban et al. proposed MILABOT, which combines 22 response models including neural retrieval-based models, neural network based generative models and template-based systems (Serban et al. 2017). The MILABOT was trained on crowdsourced data (from Amazon Mechanical Turk) and real-world user interactions via reinforcement learning, to select an appropriate response from the models in its ensemble. In the reinforcement learning framework (Sutton and Barto 1998), they consider selecting the appropriate response as a sequential decision making problem. They then treat the re-ranker as the agent which observes the dialogue history  $h_t$  at each time step  $t = 1, 2, \dots, T$  and selects one of the  $K$  actions (i.e. candidate responses  $a_t^1, \dots, a_t^K$ ). The goal of the agent is to maximize the accumulated reward  $R$ :

$$R = \sum_{t=1}^T \gamma^t r_t \quad (12)$$

where  $\gamma \in (0, 1]$  is a discount factor, and  $r_t$  is the reward after taking the action at time step  $t$ , which is the labeled

**Table 2** Summary of three core design philosophies for chatbot construction

Philosophy	Techniques	Pros	Cons	Scenarios
Rule-based	Pattern matching; modular task-oriented system	Easy to start; robust, safe in narrow domains; context-aware	Hard to extend; needs structured domain knowledge; needs a lot of hand-crafted features	Restaurant reservations; movie bookings; ordering food; online shopping
Retrieval-based	TF-IDF; DNN-based; RNN-based	Can handle more requests; literal human utterances; various expression	Easily out of context; inconsistent personality; limited by size of repository	Domain-specific (e.g., travel) question answering; technical support
Generation-based	Phrase-based SMT; Seq2Seq-based: + attention; + hierarchical structure; + memory network	Can generate new responses; can add in external knowledge; highly coherent	Prono to generate universal sentences; easily not informative; needs a huge training dataset	Online customer care; technical support; entertainment

**Table 3** Chatbot enhancement techniques

Enhancement methods	Techniques	Representative papers
Multi-chatbots design	Reinforcement learning for re-ranker policy	Serban et al. CoRR (2017)
	Data-driven re-ranker models	Qiu et al. ACL (2017) and Song et al., IJCAI (2018)
Human–chatbot collaboration	CoChat: external memory + HRNN	Luo et al., AAAI (2018)
	Evorus: crowd-powered, automates itself over time	Huang et al., CHI (2018)
Learning from users	Programming by demonstration	Li et al., CHI (2017)
	Verbal instruction	Azaria et al., AAAI (2016)

1-5 points of appropriateness in the MILABOT case. The MILABOT trained by the off-policy REINFORCE approach or the Q-learning approach achieved a great performance in terms of user average score, dialogue length and positive utterances.

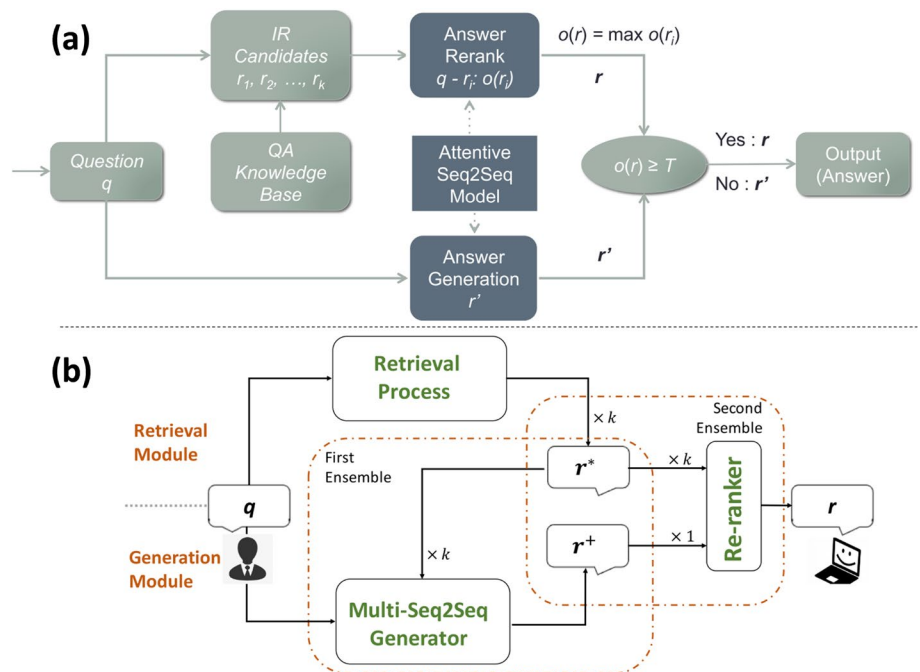
However, the reinforcement learning method for the re-ranker design in MILABOT needs a lot of labeled data. For a multi-chatbot framework that only incorporates the retrieval-based and generation-based chatbots and only targets the one-round response generation to the request, there are some data-driven models for the re-ranker design (Song et al. 2018; Qiu et al. 2017). For example, in *AliMe Chat*, Qiu et al. used an attentive Seq2Seq model to score the candidate answers (i.e. responses) with an input question (Qiu et al. 2017) (Fig. 7a). For each of the retrieved  $k$  candidates, the score function  $s^{\text{Mean-Prob}}$  is defined as:

$$s^{\text{Mean-Prob}} = \frac{1}{n} \sum_{i=1}^n p(y_i = w_i | \theta_i) \tag{13}$$

where  $w_1, \dots, w_n$  are the word sequences of the candidate answer, and  $\theta_i$  is the parameters for generating  $w_i$  at time step  $i$ . If the maximum candidate score is not less than a threshold  $T$  (determined by cross validation on top-1 accuracy), *AliMe Chat* selects that candidate as the final answer. Otherwise, it selects the answer generated by the same attentive Seq2Seq model. With such a re-ranker, *AliMe Chat* performs better than merely a retrieval-based chatbot in the online A/B test.

To further make use of the retrieved real responses, Song et al. proposed to feed the retrieved candidates to the generative model in addition to the original request, and re-rank both the retrieved candidates and generated candidate (Song et al. 2018) (Fig. 7b). They deployed a Gradient Boosting Decision Tree (Ye et al. 2009) in the re-ranker, and utilized several high-level features for training and scoring, such as term similarity, entity similarity, topic similarity, statistical machine translation, length and fluency. Trained on 1,606,741 query-reply pairs from Baidu Tieba, their chatbots

**Fig. 7** Two multi-chatbot frameworks: **a** *AliMe Chat* that uses an attentive Seq2Seq model as the re-ranker (Qiu et al. 2017); **b** an ensemble of retrieval-based and generation-based chatbots that use a decision tree as the re-ranker [from Song et al. (2018), with permission]



ensemble framework outperformed other methods in terms of BLEU and human score.

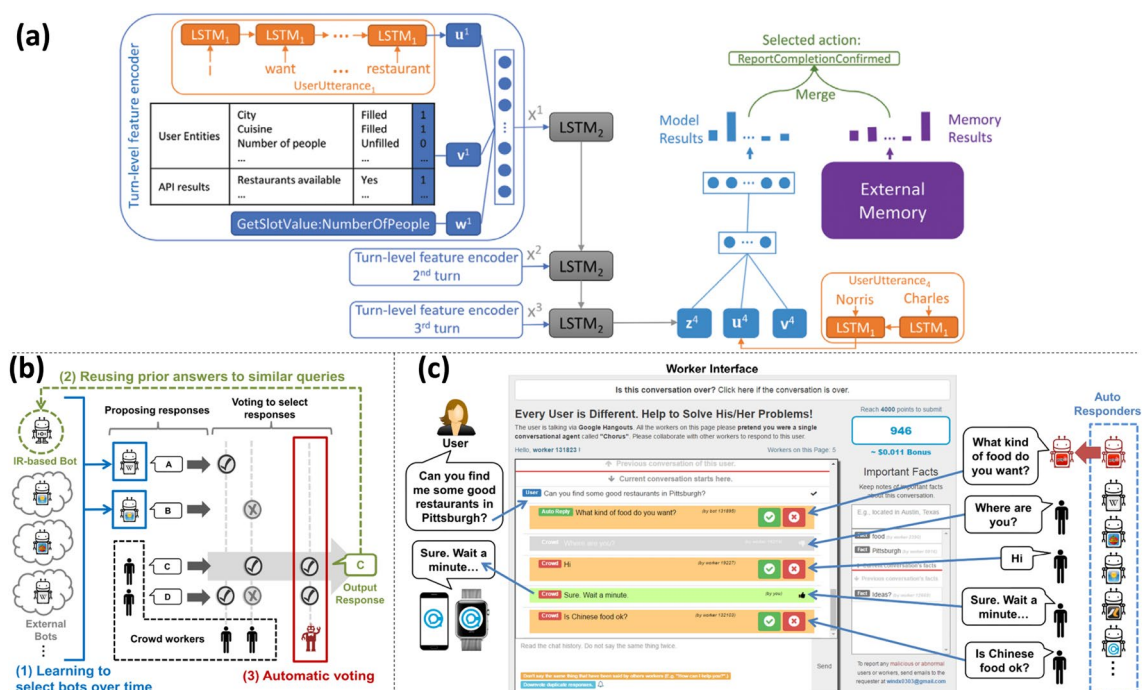
Although a hybrid of different chatbots can help to handle cross-domain requests and generate more appropriate responses, it can not learn new skills outside the capability of the chatbot ensemble and will probably fail when facing out-of-domain requests. To deal with such circumstances, designers may need to add human intervention or user customization mechanisms into the chatbot, which will be introduced in the following two sections.

### 3.2 Human–chatbot collaboration

As the fully automatic service chatbots are still potentially problematic by generating the wrong responses or taking unexpected actions, it is necessary to involve human workers in some failure-sensitive real-world domains. One of the main challenges in such human–chatbot collaboration scenarios is to enable the chatbots to learn a new action from a human worker for handling similar cases later. To handle this challenge, Luo et al. introduced CoChat, a dialogue manager framework that utilizes the MemHRNN (i.e., an external memory + hierarchical recurrent neural network) to enable human intervention and chatbot improvement at any time (Luo et al. 2018) (Fig. 8a). Similar to the HRED introduced in Sect. 2.3, MemHRNN concatenates the utterance  $u^m$ , entity form  $v^m$  and taken action  $w^m$  features into a vector  $x^m$  in each user's turn, and then feeds this vector

into a higher level recurrent network for dialogue encoding  $z^{m+1}$  in the next turn. The latest dialogue encoding  $z^{m+1}$ , together with the current  $u^{m+1}$  and  $v^{m+1}$ , is then used to select action for the current turn through a two-layer fully-connected network. The external memory is introduced to address the one-shot learning challenges of the new actions by increasing their possibilities if their recorded dialogue states are similar to the current dialogue state. For example, when a human worker refuses all existing candidate actions and inputs a new action, the new action will be added into the action pool, and the corresponding dialogue state (i.e., entity form + context vector) will be stored in the memory. When a similar user utterance comes later, the possibility of selecting the newly added action from the HRNN is still low because of the limited samples, but inside the external memory the possibility is high and it is merged with the model results to increase its chances.

However, the MemHRNN above is only suitable in narrow domains with structured knowledge. To incorporate human workers with service chatbots in relatively unstructured domains, Huang et al. introduced Evorus, a crowd-powered chatbot built to automate itself over time (Huang et al. 2018) (Fig. 8b, c). Evorus selects the response to the user's request via a voting mechanism, in which each human worker votes for the candidate responses proposed by different chatbots and other human workers. With these upvotes and downvotes labels, the re-ranker of the multiple chatbots is trained to select candidate responses to the given request.



**Fig. 8** Two human–chatbot collaboration frameworks: **a** CoChat Luo et al. (2018) that utilizes external memory to handle one-shot learning challenges for new action; **b, c** Huang et al. (2018) that learns to automate itself from crowd workers over time



Also, these labels are used to train a LibLinear (Fan et al. 2008) classifier for automatic voting, which aims to reduce human workload in the future. Another way to save human effort in Evorus is that it uses an information-retrieval-based (IR-based) method to find proper answers to similar queries in prior conversations. With such techniques, Evorus is able to improve itself over time. However, the cost paid for the crowd is expensive, and the answers will not be consistent.

The drawback of the human–chatbot collaboration method is that it always consumes human workers effort. In the cases that human intervention is the must, the future research should focus on providing a better experience for the human workers and improving the chatbots more quickly.

### 3.3 Learning from users

As another way to enhance service chatbots’ abilities, enabling the chatbots to directly learn from users’ instructions can also help in handling more type of requests. For example, when chatbots serving as a smart phone personal agent to handle a request “order a cup of greentea”, it is possible

that the chatbot can not fulfill the user’s request due to the undefined intent or lack of corresponding built-in APIs. In this case, a general response like “here is what I found from google” could be frustrating, while a mechanism for the user to teach it how to handle similar requests could be more helpful.

To this end, Li et al. designed a programming-by-demonstration system SUGLITE which enables the user to teach the chatbot how to execute unknown commands (Li et al. 2017) (Fig. 9a). Take the unknown “green tea ordering” command as an example, SUGILITE will answer “Sorry but I don’t understand. Would you like to teach me?” The user can then say “yes” and start to demonstrate the procedure of ordering a cup of green tea via the drink ordering app (e.g., Starbucks) step by step. SUGLITE will guide and confirm with the user to finish the demonstration, and after that it parameterizes the script based on the command and the UI hierarchy of that app. The next time the user says “order a DRINK please”, SUGLITE is able to understand and order that DRINK directly using that app if available. Alternatively, the user can view/edit the recording and script

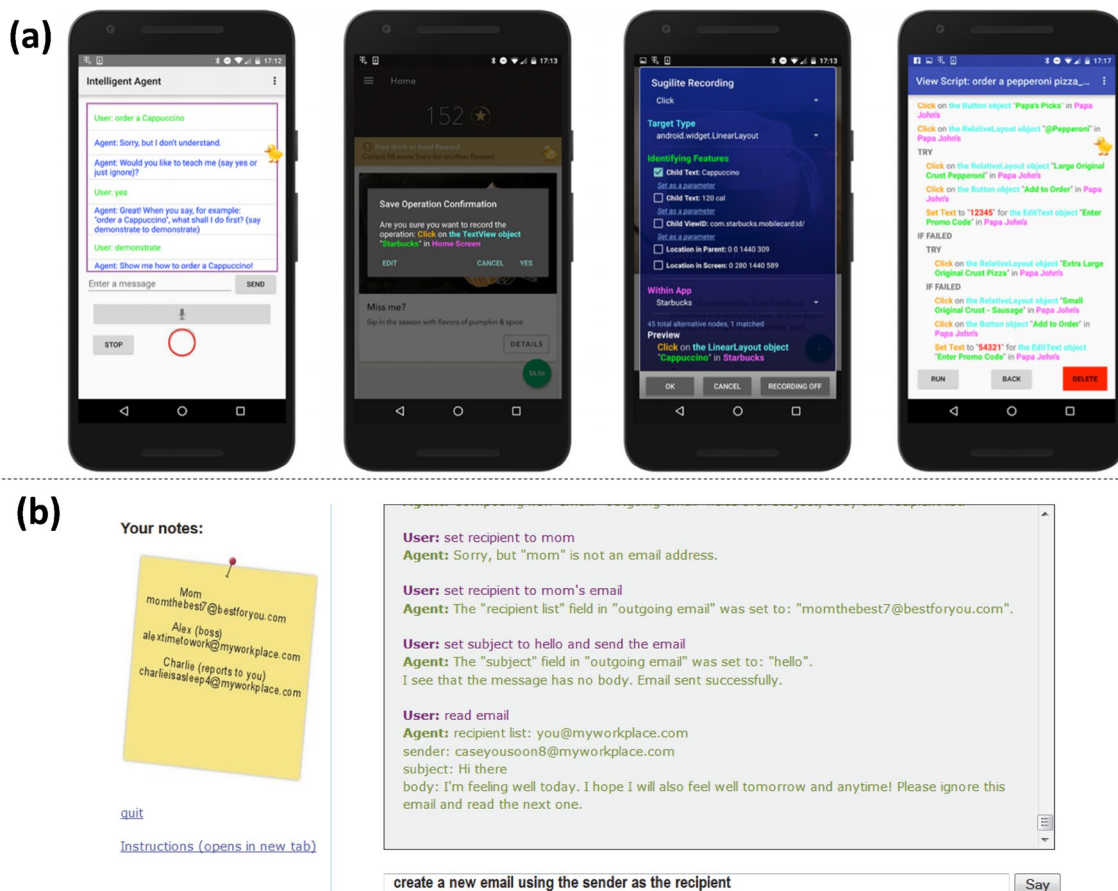


Fig. 9 Two chatbots learning from users examples: a SUGILITE (Li et al. 2017), which learns how to execute new commands from user’s demonstration; b LIA [from Azaria et al. (2016), with permission],

which learns new concepts or new procedural knowledge from user’s verbal instructions



to deal with ambiguities and failure issues. The authors also showed that they can successfully teach SUGLITE in many other scenarios like “send [AMOUNT] dollars to [NAME]” and “show me the CHI paper by [NAME]”. However, SUGLITE is limited to the structured tasks through structured apps. SUGLITE can not understand the semantics in the demonstration and be applied in web-based applications.

Apart from the graphical programming-by-demonstration techniques, using verbal instruction to teach the chatbots opens another opportunity to generate new commands. For example, in an email environment, Azaria et al. proposed a Learning by Instruction Agent (LIA), which allows users to teach it some primitive commands such as sending emails via solely natural language interaction (Azaria et al. 2016) (Fig. 9b). LIA uses a semantic parser to assign executable semantics to each natural language command (i.e., text command paired with logical form). During the verbal teaching, the user can teach LIA new concepts along with their fields and concepts (e.g., “a contact has an email address” and “bob is a contact”). The user can also teach procedural knowledge such as how to forward an email by updating the semantic parser using known natural language instruction. The authors showed that people with little or no programming knowledge could easily teach LIA with very little training. Yet the rules behind each task are complicated and it is thus inconvenient to generalize LIA in other domains.

The chatbot learning from users through demonstration or verbal instruction is a promising way to enhance and customize the service chatbots. However, existing methods can only achieve it in narrow domains and need a lot of hand-crafted features. It will be an interesting research direction to leverage big data and machine learning to automatically learn from users and provide a better user experience.

### 3.4 Summary

According to the example and discussion above, we summarize the advantages and disadvantages of current techniques with three kinds of enhancement methods in Table 4. These techniques show a promising future for building a stronger

and a more user-friendly service chatbot. The construction of these chatbots not only needs the steady development of natural language processing, but also the human–computer interaction design considerations of human workers as well as ended users. In the next section, we discuss two potential future directions on service chatbot design.

## 4 Future directions

Previous construction methods of service chatbots normally suffer from out-of-context, universal response generation as well as inconsistent personality. In this section, we discuss how the recent Transformer and contextual embedding techniques can potentially help to generate more context-related responses. We also discuss the need for including personality into service chatbots and building a tone-aware chatbot.

### 4.1 Response generation with transformer and contextual embedding

Over the past decade, the success of many natural language processing (NLP) tasks (including the response generation given a request in this survey) is mainly built on the improvement of deep learning models and language models. Up to 2017, most state-of-the-art deep learning models on response generation were limited to the ensemble of sequential RNN and other components like LSTM, GRU, attention mechanism and memory network. These sequential models can capture the time information of the input sentences (e.g., how next word/sentence is conditioned on previous words/sentences), but largely ignore the global information of the whole sentence (e.g., a word could have different relations with all the words in the sentence). This could be the main reasons for the out-of-context and universal responses generated by these response generation models. Moreover, these models usually use pre-trained language models like word2vec Mikolov et al. (2013) and Glove Pennington et al. (2014) to represent the input, while these language models are context-free, i.e. the word (e.g., blue) is always represented by the same vector

**Table 4** Summary of three kinds of chatbot enhancement methods

Enhancement methods	Techniques	Pro	Cons
Multi-chatbots design	Reinforcement learning for re-ranker policy Data-driven re-ranker models	Can handle cross-domain requests; make use of existing chatbots	Unable to learn new skills outside the capability of ensemble; fails in out-of-domain requests
Human–chatbot collaboration	CoChat: external memory + HRNN Evorus: crowd-powered, automates itself over time	Learn new skills from human workers; more robust and able to handle complex requests	Expensive; not consistent; long delays
Learning from users	Programming by demonstration Verbal instruction	Learn unknown commands from users; friendly to novice	Limited to narrow domains; need to design hand-crafted rules

regardless of different contexts (e.g., color or mood). These context-free language models may not properly represent the words and sentences in a certain context, which also limits the performance of the response generation models in previous service chatbots. Therefore, to generate a context-related and informative response, the service chatbots should better capture the global information of the requests. The recent developments in Transformer (Vaswani et al. 2017) and contextual embedding (Peters et al. 2018; Devlin et al. 2018) point out a potentially better way for this purpose.

Unlike RNN-based models which condition the next word on previous words, the Transformer enables words to connect with each other in the sentence. It uses Scaled Dot-Product Attention (Fig. 10a), which can be defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{14}$$

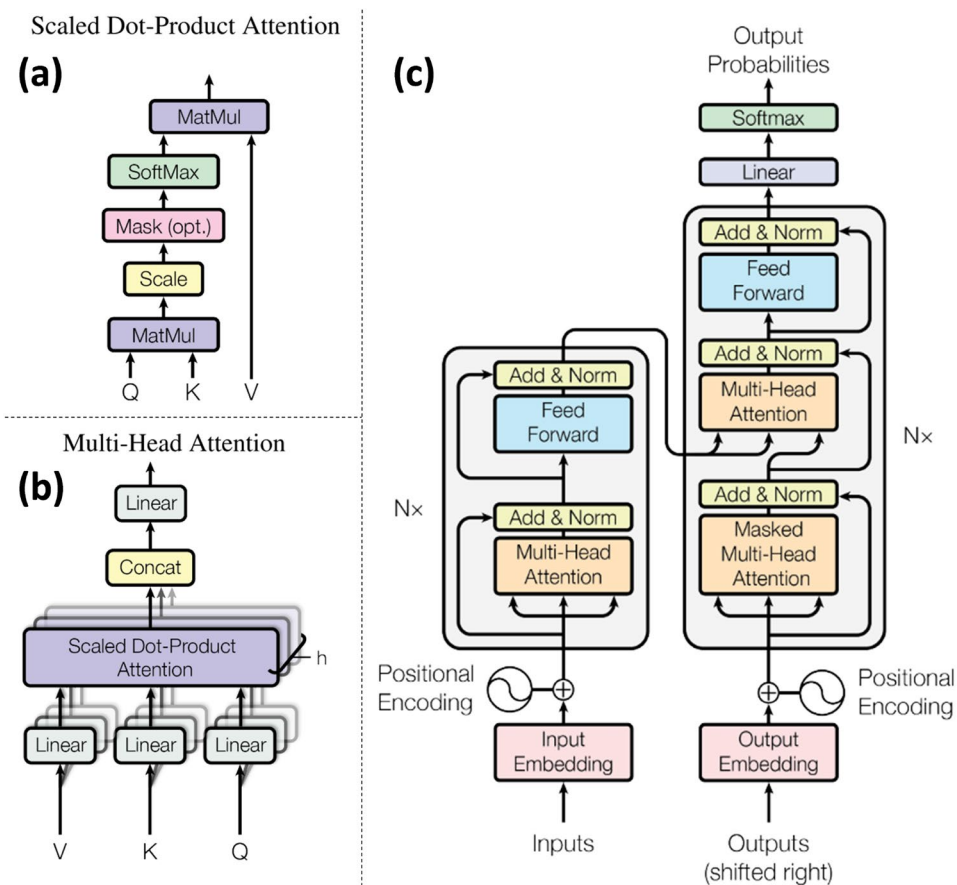
where  $Q$  and  $K$  are a set of queries and keys of dimension  $d_k$ , and  $V$  is the values of dimension  $d_v$  paired with the keys. In the cases of machine translation or response generation, the  $Q, K, V$  can be set to the same set of words  $X$  in the sentence, i.e., self-attention. Then through the dot-production

operation, each word (i.e.,  $q_t = x_t$ ) builds connections with other words (i.e.,  $K = V = X$ ) in the sentence. Furthermore the Transformer uses Multi-Head Attention (Fig. 10b) to jointly attend to information from different representation subspaces at different positions, which is somehow similar to the idea of convolutional neural network. Multi-head attention can be formulated as:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \tag{15}$$

where  $\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V)$ , and  $W_i^Q, W_i^K, W_i^V, W^O$  are parameter matrices. The final transformer model (Fig. 10c) also takes positional encoding as another input to make use of the order of the sequence. The transformer model has an encoder-decoder structure, where the encoder maps an input sequence  $(x_1, \dots, x_n)$  to a sequence of continuous representations  $z = (z_1, \dots, z_n)$ , and the decoder generates an output sequence  $(y_1, \dots, y_m)$  of symbols one element at a time given  $z$  and previously generated symbols. The Transformer is shown to outperform the best previously reported models on the WMT English-to-German translation task in terms of BLEU score. Since the task of response generation in service chatbots is similar to machine translation (MT) in some aspects, and previously it has successfully

**Fig. 10** The transformer model [adapted from Vaswani et al. (2017)]: **a** scaled dot-product attention; **b** multi-head attention; **c** model architecture



used MT models (e.g. SMT, Seq2Seq) for response generation, it will be a promising future direction to apply the Transformer to the service chatbots design (especially in retrieval-based and generation-based methods).

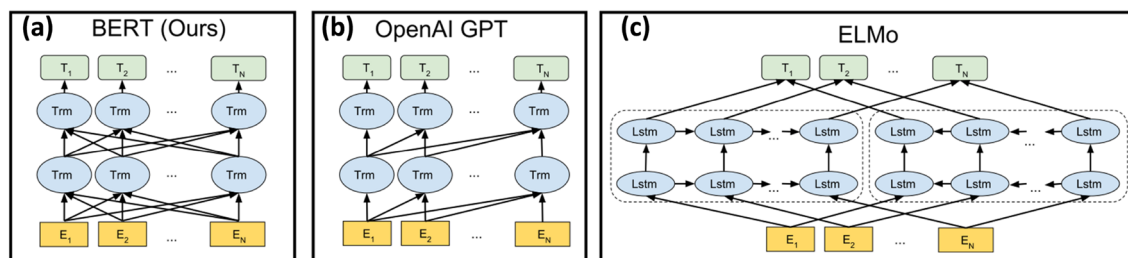
Besides the transformer, the rapid development of contextual word embedding in 2018 is also an exciting direction for better response generation. Unlike context-free word embedding like word2vec and Glove, the contextual word embedding represents each word in the context of the sentences through some pre-training models like ELMo (Peters et al. 2018), OpenAI GPT (Alec and Sutskever 2018) and BERT (Devlin et al. 2018). As shown in Fig. 11, these models take some context-free word embedding as the input, generate the features of the words and sentences inside the model, and output the contextual word embedding. ELMo (Fig. 11c) uses the concatenation of independently trained left-to-right and right-to-left LSTM to generate features, and it improved the state-of-the-art results on some question answering and language inference tasks at that time. Later, the OpenAI GPT (Fig. 11b) uses a left-to-right transformer to replace LSTM, which could potentially improve the model but is still unidirectional. The latest BERT (i.e., Bidirectional Encoder Representation from Transformers) model proposed by Devlin et al. (2018) instead jointly conditions the representations on both the left and right context in all layers (Fig. 11a). It takes the sum of the token embeddings, the segmentation embeddings and the positional embeddings as the input embeddings, and is trained on a masked LM (MLM) task and a next sentence prediction task. In the MLM task, 15% of the tokens are randomly masked and the model is trained to predict these masked tokens. While the next sentence prediction task aims to capture the relationship between two text sentences. The upper models built on the pre-trained BERT representations achieve the current state-of-the-art results on eleven NLP tasks, including the sentence pair classification tasks and SQuAD v1.1 question answering. It is possible that applying these models to construct contextual word embeddings of the original request can also improve the model performance in response generation tasks for chatbots.

It has become a trend to apply the Transformer and contextual word embeddings to many NLP tasks. Similarly in the future, we can use the Transformer to replace the RNN, LSTM and GRU in previous retrieval-based and generation-based models, and use the contextual word embeddings as the new input features. With these techniques, the service chatbots are supposed to provide more appropriate and informative responses given the request.

## 4.2 Building chatbot with personality

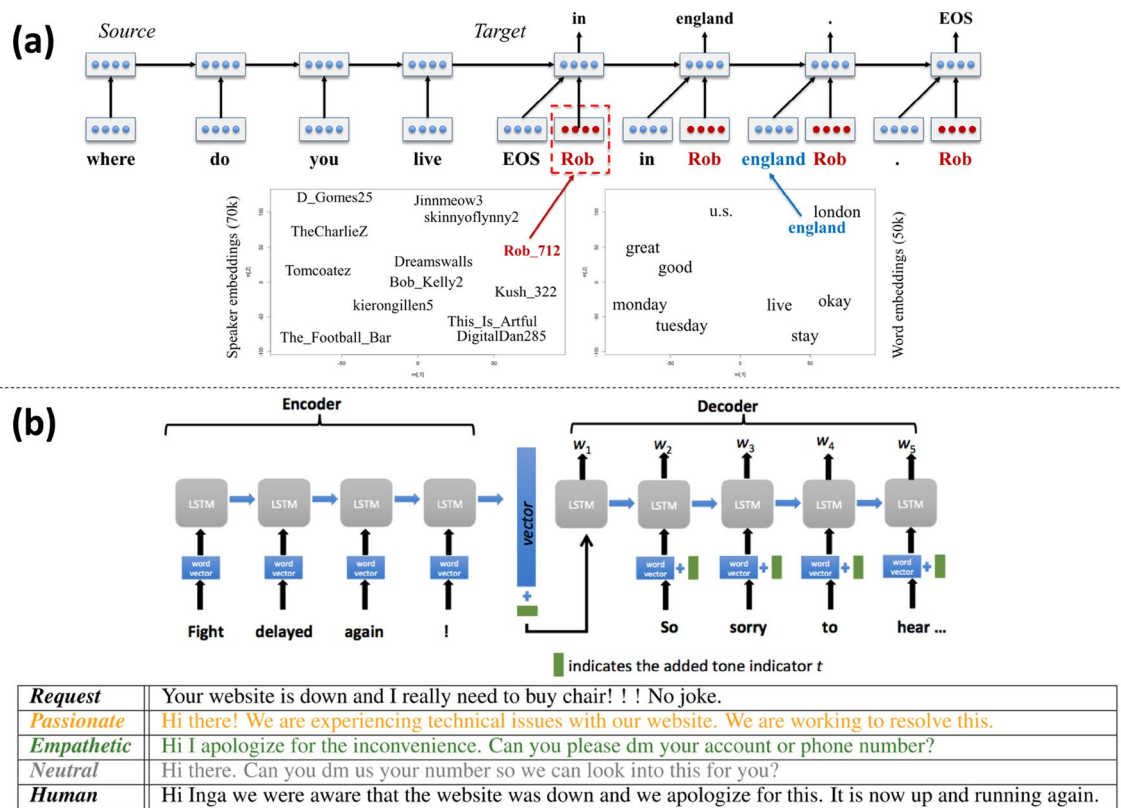
Although recent progress in machine learning has enabled chatbots to generate more grammatically correct responses, it is still hard to design a chatbot that has consistent personality, which will have a significant impact on user experience (Hu et al. 2018). A chatbot with a consistent personality should have a corresponding persona [e.g., background facts or profile (Li et al. 2016)] and have a certain speaking style [e.g., toned responses (Hu et al. 2018)]. The designers can easily design the personality for simple rule-based chatbots, but it becomes impossible when chatbots get more and more complicated. For retrieval-based and generation-based chatbots, a straightforward thought of designing a consistent personality is to build a highly consistent dataset, but it would be unrealistic as the dataset is usually diverse. Merging personality information into the response generation models, instead, is a promising way to design chatbots with a consistent personality in the wild.

In Li et al.'s work (2016), additional speaker embeddings are concatenated with the word embeddings in the decoder of the standard Seq2Seq model (Fig. 12a). The speaker model cluster users along some of the traits (e.g., age, country of residence) based on the responses alone, and the embeddings are jointly learned with word embeddings via backpropagation. Specifically during the training, each source-target pair has already been clustered into one corresponding speaker, and the randomly initiated speaker embeddings will be trained with all the pairs inside the same speaker cluster. Then given the same request as the input,



**Fig. 11** The different pre-training models of contextual word embeddings [from Devlin et al. (2018), with permission]: **a** BERT Devlin et al. (2018) that uses a bidirectional transformer; **b** OpenAI GPT

Alec and Sutskever (2018) that uses a left-to-right transformer; **c** ELMo Peters et al. (2018) that uses the concatenation of independently trained left-to-right and right-to-left LSTM



**Fig. 12** Two example of building chatbots with personality: **a** Persona-based neural conversation model that uses speaker embeddings [adapted from Li et al. (2016)]; **b** tone-aware seq2seq model and examples of toned responses [from Hu et al. (2018), with permission]

the model can give a different answer based on the speakers' personality.

Inspired by this model in the open-domain, Hu et al. explored how to generate toned responses for customer care chatbots on social medial (Hu et al. 2018). They first conducted a formative study to annotate the initial 53 tones and summarize them into eight major tones, i.e., *anxious*, *frustrated*, *impolite*, *passionate*, *polite*, *sad*, *satisfied* and *empathetic*. Then they established eight linear regression analyses for the eight major tones, and found that *empathetic* and *passionate* tones are beneficial for customer care. A keywords extraction was then conducted on these two selected tones. Using these keywords as indicators of different tones, their tone-aware seq2seq model concatenates the indicator and the word embedding in the decoder to train itself on the pre-processed conversations data with tone information (Fig. 12b). Their user studied shows that the toned responses are perceived to be significantly more appropriate and helpful than the neural responses.

However, the two works above do not evaluate how different chatbot personalities could affect the end user engagement and how to automatically adjust the speaking style. In the future, we can investigate the user satisfaction relationship between the tones in the requests and the tones in the

responses via the annotated data. We can then train a service chatbot that not only has a consistent personality, but also senses a user's emotional status and adjust its speaking style.

## 5 Conclusion

Using chatbots to assist or replace human workers is a trend in service industries, owing to chatbots being cheaper, more easy-to-access and more objective. However, it is still a huge challenge to design service chatbots which provide as satisfying responses to users' requests as human workers can. This survey presents the common construction methods of service chatbots, which can be classified into rule-based, retrieval-based and generation-based methods. We further present some enhancement methods, which can power the chatbots either by an ensemble of multiple chatbots, collaboration with human workers, or by instruction from the users. With the progress seen in both the Transformer and contextual embedding over the past two years, we discuss promising future directions in the design of chatbots which could better understand users' requests and generate more helpful responses. Moreover, we show a potential way to



build service chatbots with the persona and speaking style of a real human, which can provide a better user experience.

## References

- Azaria, A., Krishnamurthy, J., Mitchell, T.M.: Instructable intelligent personal agent. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI'16, pp. 2681–2689. AAAI Press (2016). <http://dl.acm.org/citation.cfm?id=3016100.3016277>
- Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. CoRR [arXiv:abs/1409.0473](https://arxiv.org/abs/1409.0473) (2014)
- Bengio, Y.: Learning deep architectures for AI. Found. Trends Mach. Learn. **2**(1), 1–127 (2009). <https://doi.org/10.1561/2200000006>
- Bengio, Y., Ducharme, R., Vincent, P., Janvin, C.: A neural probabilistic language model. J. Mach. Learn. Res. **3**, 1137–1155 (2003)
- Chai, J., Lin, J., Zadrozny, W., Ye, Y., Stys-Budzikowska, M., Horvath, V., Kambhatla, N., Wolf, C.: The role of a natural language conversational interface in online sales: a case study. Int. J. Speech Technol. **4**(3), 285–295 (2001). <https://doi.org/10.1023/A:1011316909641>
- Chen, H., Liu, X., Yin, D., Tang, J.: A survey on dialogue systems: recent advances and new frontiers. SIGKDD Explor. Newsl. **19**(2), 25–35 (2017). <https://doi.org/10.1145/3166054.3166058>
- Chung, J., Çağlar Gülçehre, Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. CoRR [arXiv:abs/1412.3555](https://arxiv.org/abs/1412.3555) (2014)
- Colby, K.M.: Artificial Paranoia: A Computer Simulation of Paranoid Processes. Elsevier Science Inc., New York (1975)
- Cranshaw, J., Elwany, E., Newman, T., Kocielnik, R., Yu, B., Soni, S., Teevan, J., Monroy-Hernández, A.: Calendar.help: designing a workflow-based scheduling agent with humans in the loop. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17, pp. 2382–2393. ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3025453.3025780>
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. CoRR [arXiv:abs/1810.04805](https://arxiv.org/abs/1810.04805) (2018)
- Dialog flow: Retrieved: 2018-12-10. (2018). <https://dialogflow.com/>. Accessed 10 Dec 2018
- Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: Liblinear: a library for large linear classification. J. Mach. Learn. Res. **9**, 1871–1874 (2008)
- Ferrucci, D.A., Brown, E.W., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J.M., Schlaefer, N., Welty, C.A.: Building watson: an overview of the deepqa project. AI Mag. **31**(3), 59–79 (2010)
- Ghazvininejad, M., Brockett, C., Chang, M.W., Dolan, W.B., Gao, J., tau Yih, W., Galley, M.: A knowledge-grounded neural conversation model. In: AAAI (2018)
- Grudin, J., Jacques, R.: Chatbots, humbots, and the quest for artificial general intelligence. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19, pp. 209:1–209:11. ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3290605.3300439>
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
- Hu, B., Lu, Z., Li, H., Chen, Q.: Convolutional neural network architectures for matching natural language sentences. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 27, pp. 2042–2050. Curran Associates Inc, New York (2014)
- Hu, T., Xu, A., Liu, Z., You, Q., Guo, Y., Sinha, V., Luo, J., Akkiraju, R.: Touch your heart: a tone-aware chatbot for customer care on social media. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, pp. 415:1–415:12. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3173989>
- Huang, T.H.K., Chang, J.C., Bigham, J.P.: Evorus: A crowd-powered conversational assistant built to automate itself over time. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, pp. 295:1–295:13. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3173869>
- Ibm watson assistant: Retrieved: 2018-12-10. (2018). <https://assistant-us-south.watsonplatform.net/us-south/b3a5bd9b-9ea9-4be8-9ec7-145f04f69453/home>. Accessed 10 Dec 2018
- Isbell Jr, C.L., Kearns, M., Kormann, D., Singh, S., Stone, P.: Cobot in lambdamoo: A social statistics agent. pp. 36–41 (2000)
- Jain, M., Kota, R., Kumar, P., Patel, S.N.: Convey: Exploring the use of a context view for chatbots. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, pp. 468:1–468:6. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3174042>
- Ji, Z., Lu, Z., Li, H.: An information retrieval approach to short text conversation. CoRR [arXiv:abs/1408.6988](https://arxiv.org/abs/1408.6988) (2014)
- Kaelbling, L.P., Littman, M.L., Moore, A.W.: Reinforcement learning: a survey. J. Artif. Int. Res. **4**(1), 237–285 (1996)
- Li, J., Galley, M., Brockett, C., Spithourakis, G., Gao, J., Dolan, B.: A persona-based neural conversation model. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 994–1003. Association for Computational Linguistics (2016). <https://doi.org/10.18653/v1/P16-1094>
- Li, T.J.J., Azaria, A., Myers, B.A.: Sugilite: creating multimodal smartphone automation by demonstration. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17, pp. 6038–6049. ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3025453.3025483>
- Li, Y., Luo, X., Zheng, Y., Xu, P., Fu, H.: Sweepcanvas: sketch-based 3d prototyping on an rgb-d image. In: Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology, pp. 387–399. ACM (2017)
- Liao, Q.V., Mas-ud Hussain, M., Chandar, P., Davis, M., Khazaeni, Y., Crasso, M.P., Wang, D., Muller, M., Shami, N.S., Geyer, W.: All work and no play? In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, pp. 3:1–3:13. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3173577>
- Liu, C.W., Lowe, R., Serban, I., Noseworthy, M., Charlin, L., Pineau, J.: How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pp. 2122–2132. Association for Computational Linguistics (2016). <https://doi.org/10.18653/v1/D16-1230>
- Lowe, R., Pow, N., Serban, I., Pineau, J.: The ubuntu dialogue corpus: a large dataset for research in unstructured multi-turn dialogue systems. In: SIGDIAL Conference (2015)
- Lu, Z., Li, H.: A deep architecture for matching short texts. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems, vol. 26, pp. 1367–1375. Curran Associates Inc, New York (2013)
- Luo, X., Lin, Z., Wang, Y., Nie, Z.: Cochat: Enabling bot and human collaboration for task completion. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (2018)
- Luong, M.T., V. Le, Q., Sutskever, I., Vinyals, O., Kaiser, L.: Multi-task sequence to sequence learning. In: Proceedings of ICLR, San Juan, Puerto Rico (2015)



- Magazine, C.: Chatbot report 2018: Global trends and analysis (2018). Retrieved: 2018-12-10. <https://chatbotmagazine.com/chatbot-report-2018-global-trends-and-analysis-4d8bbe4d924b>
- Mauldin, M.L.: Chatterbots, tinymuds, and the turing test: entering the loebner prize competition. In: Proceedings of the Twelfth National Conference on Artificial Intelligence (Vol. 1), AAAI '94, pp. 16–21. American Association for Artificial Intelligence, Menlo Park, CA, USA (1994). <http://dl.acm.org/citation.cfm?id=199288.199285>
- Mesnil, G., Dauphin, Y., Yao, K., Bengio, Y., Deng, L., Hakkani-Tur, D., He, X., Heck, L., Tur, G., Yu, D., Zweig, G.: Using recurrent neural networks for slot filling in spoken language understanding. *Trans. Audio Speech Lang. Proc.* **23**(3), 530–539 (2015)
- Microsoft: Cognitive services: language understanding (luis). Retrieved: 2018-12-10. <https://www.luis.ai/home> (2018)
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proceedings of the 26th International Conference on Neural Information Processing Systems, Vol. 2, NIPS'13, pp. 3111–3119. Curran Associates Inc., USA (2013). <http://dl.acm.org/citation.cfm?id=2999792.2999959>
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (2002). <http://aclweb.org/anthology/P02-1040>
- Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP, vol. 14, pp. 1532–1543 (2014)
- Peters, M.E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep contextualized word representations. In: Proc. of NAACL (2018)
- Qiu, M., Li, F.L., Wang, S., Gao, X., Chen, Y., Zhao, W., Chen, H., Huang, J., Chu, W.: Alime chat: A sequence to sequence and rerank based chatbot engine. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp. 498–503. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/P17-2079>
- Qiu, M., Li, F.L., Wang, S., Gao, X., Chen, Y., Zhao, W., Chen, H., Huang, J., Chu, W.: Alime chat: a sequence to sequence and rerank based chatbot engine. In: ACL (2017)
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding with unsupervised learning. In: Technical report, OpenAI (2018)
- Ramesh, K., Ravishankaran, S., Joshi, A., Chandrasekaran, K.: A survey of design techniques for conversational agents. In: Kaushik, S., Gupta, D., Kharb, L., Chahal, D. (eds.) *Information, Communication and Computing Technology*, pp. 336–350. Springer Singapore, Singapore (2017)
- Ritter, A., Cherry, C., Dolan, W.B.: Data-driven response generation in social media. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '11, pp. 583–593. Association for Computational Linguistics, Stroudsburg, PA, USA (2011). <http://dl.acm.org/citation.cfm?id=2145432.2145500>
- Schlesinger, A., O'Hara, K.P., Taylor, A.S.: Let's talk about race: Identity, chatbots, and ai. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, pp. 315:1–315:14. ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3173889>
- Schuster, M., Paliwal, K.: Bidirectional recurrent neural networks. *Trans. Signal Process.* **45**(11), 2673–2681 (1997). <https://doi.org/10.1109/78.650093>
- Serban, I., Sankar, C., Germain, M., Zhang, S., Lin, Z., Subramanian, S., Kim, T., Pieper, M., Chandar, A.P.S., Ke, N.R., Mudumba, S., de Brébisson, A., Sotelo, J., Suhubdy, D., Michalski, V., Nguyen, A., Pineau, J., Bengio, Y.: A deep reinforcement learning chatbot. *CoRR arXiv:abs/1709.02349* (2017)
- Serban, I.V., Sordoni, A., Bengio, Y., Courville, A., Pineau, J.: Building end-to-end dialogue systems using generative hierarchical neural network models. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI'16, pp. 3776–3783. AAAI Press (2016). <http://dl.acm.org/citation.cfm?id=3016387.3016435>
- Shang, L., Lu, Z., Li, H.: Neural responding machine for short-text conversation. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 1577–1586. Association for Computational Linguistics (2015). <https://doi.org/10.3115/v1/P15-1152>
- Shawar, B.A., Atwell, E.: Chatbots: are they really useful? *LDV Forum* **22**, 29–49 (2007)
- Shum, Hy, He, Xd, Li, D.: From eliza to xiaoice: challenges and opportunities with social chatbots. *Front. Inf. Technol. Electron. Eng.* **19**(1), 10–26 (2018). <https://doi.org/10.1631/FITEE.1700826>
- Song, Y., Li, C.T., Nie, J.Y., Zhang, M., Zhao, D., Yan, R.: An ensemble of retrieval-based and generation-based human–computer conversation systems. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18, pp. 4382–4388. International Joint Conferences on Artificial Intelligence Organization (2018). <https://doi.org/10.24963/ijcai.2018/609>
- Sukhbaatar, S., Szlam, A., Weston, J., Fergus, R.: End-to-end memory networks. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 28, pp. 2440–2448. Curran Associates Inc, New York (2015)
- Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14, pp. 3104–3112. MIT Press, Cambridge, MA, USA (2014). <http://dl.acm.org/citation.cfm?id=2969033.2969173>
- Sutton, R.S., Barto, A.G.: *Introduction to Reinforcement Learning*, 1st edn. MIT Press, Cambridge (1998)
- Turing, A.M.: Computing machinery and intelligence. *Mind* **LIX**(236), 433–460 (1950). <https://doi.org/10.1093/mind/LIX.236.433>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds.) *Advances in Neural Information Processing Systems* 30, pp. 5998–6008. Curran Associates, Inc., New York (2017)
- Vinyals, O., Le, Q.V.: A neural conversational model. *CoRR arXiv:abs/1506.05869* (2015)
- WIKIPEDIA: Xiaoice. Retrieved: 2018-12-10. <https://en.wikipedia.org/wiki/Xiaoice> (2018)
- Walker, M., Aberdeen, J., Boland, J., Bratt, E., Garofolo, J., Hirschman, L., Le, A., Lee, S., Narayanan, S., Papineni, K., Pellom, B., Polifroni, J., Potamianos, A., Prabhu, P., Rudnicki, A., Sanders, G., Seneff, S., Stallard, D., Whittaker, S.: Darpa communicator dialog travel planning systems: the June 2000 data collection. pp. 1371–1374 (2001)
- Wallace, R.S.: *The Anatomy of A.L.I.C.E.*, pp. 181–210. Springer Netherlands, Dordrecht (2009). [https://doi.org/10.1007/978-1-4020-6710-5\\_13](https://doi.org/10.1007/978-1-4020-6710-5_13)
- Weizenbaum, J.: Eliza—a computer program for the study of natural language communication between man and machine. *Commun. ACM* **9**(1), 36–45 (1966). <https://doi.org/10.1145/365153.365168>
- Xu, A., Liu, Z., Guo, Y., Sinha, V., Akkiraju, R.: A new chatbot for customer service on social media. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17, pp. 3506–3510. ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3025453.3025496>

- Yan, R., Song, Y., Wu, H.: Learning to respond with deep neural networks for retrieval-based human–computer conversation system. In: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16, pp. 55–64. ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2911451.2911542>
- Ye, J., Chow, J.H., Chen, J., Zheng, Z.: Stochastic gradient boosted distributed decision trees. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09, pp. 2061–2064. ACM, New York, NY, USA (2009). <https://doi.org/10.1145/1645953.1646301>
- Young, S., Gašić, M., Thomson, B., Williams, J.D.: Pomdp-based statistical spoken dialog systems: a review. *Proc. IEEE* **101**(5), 1160–1179 (2013). <https://doi.org/10.1109/JPROC.2012.2225812>
- Zhou, X., Dong, D., Wu, H., Zhao, S., Yu, D., Tian, H., Liu, X., Yan, R.: Multi-view response selection for human–computer conversation. In: EMNLP (2016)
- Zhou, X., Li, L., Dong, D., Liu, Y., Chen, Y., Zhao, W.X., Yu, D., Wu, H.: Multi-turn response selection for chatbots with deep attention matching network. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1118–1127. Association for Computational Linguistics (2018). <http://aclweb.org/anthology/P18-1103>



**Zhenhui Peng** is a PhD student (Since Fall 2017) in the Department of Computer Science and Engineering (CSE) in Hong Kong University of Science and Technology (HKUST). He majors in Human-Computer Interaction (HCI) and is supervised by Dr. Xiaojuan Ma. He got my Bachelor degree in ESE from Nanjing University in 2017. His current research interests are in Human-Robot Interaction, chatbot design and Reinforcement Learning.



**Xiaojuan Ma** is an assistant professor of Human-Computer Interaction (HCI) at the Department of Computer Science and Engineering (CSE), Hong Kong University of Science and Technology (HKUST). She received the Ph.D. degree in Computer Science at Princeton University. She was a post-doctoral researcher at the Human-Computer Interaction Institute (HCII) of Carnegie Mellon University (CMU), and before that a research fellow in the National University of Singapore (NUS) in the Information Systems department. Before joining HKUST, she was a researcher of Human-Computer Interaction at Noah's Ark Lab, Huawei Tech. Investment Co., Ltd. in Hong Kong. Her background is in Human-Computer Interaction. She is particularly interested in data-driven human-engaged AI and Human-Robot Interaction in the domain of education, health, and design.