

An Emotional Comfort Framework for Improving User Satisfaction in E-Commerce Customer Service Chatbots

Shuangyong Song, Chao Wang, Haiqing Chen, Huan Chen

Alibaba Groups, Hangzhou 311121, China

{shuangyong.ssy, chaowang.wc, haiqing.chenhq, shiwan.ch}
@alibaba-inc.com

Abstract

E-commerce has grown substantially over the last several years, and chatbots for intelligent customer service are concurrently drawing attention. We presented *AliMe Assist*, a Chinese intelligent assistant designed for creating an innovative online shopping experience in E-commerce. Based on question answering (QA), *AliMe Assist* offers assistance service, customer service, and chatting service. According to the survey of user studies and the real online testing, emotional comfort of customers' negative emotions, which make up more than 5% of whole number of customer visits on *AliMe*, is a key point for providing considerate service. In this paper, we propose a framework to obtain proper replies to customers' emotional questions. The framework takes emotion classification model as a core, and the final reply selection is based on topic classification and text matching. Our experiments on real online systems show that the framework is very promising.

1 Introduction

A chatbot is considered as a question answering system in which experts provide knowledge on users' behalf. Meanwhile, chatbots are not just question answering systems, since they can carry out a lot of tasks depending on how you design it (Zhu et al., 2018). As chatbot has become an important solution to rapidly increasing customer service demands in recent years, many companies have recently launched their own intelligent customer service (ICS) chatbots for providing customer service, such as Lenovo (Li et al., 2018), Fujitsu (Okuda and Shoda, 2018), JD.com (Zhu, 2019) and Alibaba (Li et al., 2017).

For customers' emotional questions, proper emotional comfort can help improve the service. This is not only applicable to customer service staffs, but also a key point of ICS chatbots, while demonstrating human-like service is the ultimate goal of

ICS chatbots. Emotional quotient (EQ) has been a core competence of chatbot (Zhou et al., 2020), and about EQ, we can roughly categorize it into two key components: identifying users' emotions and giving users proper emotional responses. Besides, chatbots' EQ is domain-specific, since it is mainly based on emotion analyzing, and emotion-analyzing technologies are mostly domain specific.

In this paper, we introduce an emotional comfort framework for the e-commerce chatbots. E-commerce customers usually complain of slow delivery, poor quality of goods and difficulty of contacting sellers, etc. Traditional question answering based ICS chatbots may just reply customers with some pieces of 'knowledge' such as 'how to speed up the delivery', 'how to report the quality issues of goods' and 'how to contact sellers'. Without responses that are emotionally appropriate, ICS robots are too 'robotic' to users. Human-like empathy and appropriate emotional reply can help the users regain their confidence and move forward with a positive attitude. Besides, in our framework we don't consider emotional response generation models, such as (Huo et al., 2020) and (Zhou et al., 2018), since we should meet the high Queries-per-second (QPS) needs of real online applications.

Figure 1 gives two simple examples for the comparison of traditional ICS chatbots and emotional ICS chatbots, which are without or with emotional comforts. Without emotional comfort, the response appears abruptly.

2 Related Work

Classification model: classification model training is strongly based on extraction of textual semantic features, and textual semantic features can be roughly separated into word- (or character-) level features (Wang et al., 2018; Song et al., 2017; Kusner et al., 2015), n-gram level features (Yin et al., 2016; Wan et al., 2016) and sentence level features (Shen et al., 2018; Arora et al., 2016).

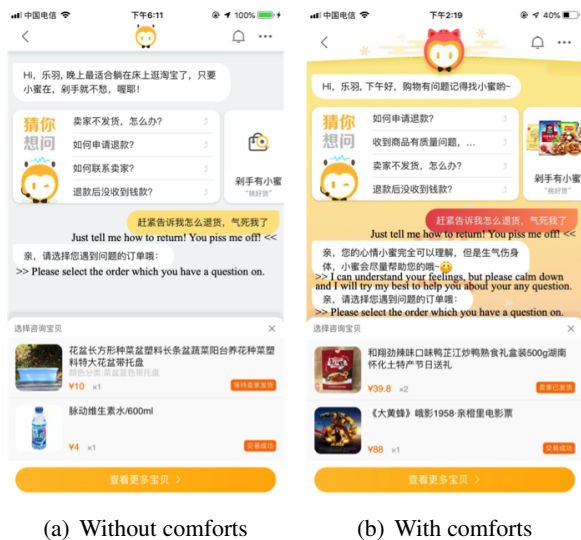


Figure 1: Comparison of conversations with or without emotional comforts.

1) Word-level features: Kusner et al. (Kusner et al., 2015) proposed word mover’s distance (WMD), a distance function between two documents, which measures the minimum traveling distance from the embedded words of one document to another one. WMD achieved good performance in the document classification task (Ma et al., 2018). Referring to WMD, Song et al. (Song et al., 2017) proposed Word Similarity Maximization (WSM), which is a faster method for calculating similarity between two short texts with word embeddings, and WSM can achieve even better results than WMD on short text classification task. Wang et al. (Wang et al., 2018) proposed a novel classification model that considers correlation between embeddings of category labels and word embeddings (LEAM), which has further enriched the word-level features of text classification.

2) N-gram level features: Yin et al. (Yin et al., 2016) proposed Attention based CNN (ABCNN) model to extract n-gram features of each of two texts, and then combine those features as input of Logistic regression model to obtain semantic similarity between two texts. Wan et al. proposed a MV-LSTM model, which utilize Bi-LSTM model to obtain multiple positional sentence representations as a kind of ‘dynamic’ n-gram features.

3) Sentence level features: Arora et al. (Arora et al., 2016) represent a sentence with a weighted average of word embeddings, with their projection onto the first principal component across all sentences in the corpus removed. Shen et al. (Shen

et al., 2018) thoroughly analyzed the effect of pooling mechanisms on representing sentences with simple word embeddings. With those sentence-level features, classification task, text sequence matching task and some other feature based tasks can all achieve good performance.

In our sentiment classification model and topic classification model, we combine those multiple-level features, and prove that our model can achieve significantly improved results.

Emotional chatbot: the most famous emotional chatbot is Xiaoice (Zhou et al., 2020), which was designed about 6 years ago. Understanding and responding to users’ emotions are two dimensions of the ability of emotional chatbots. For realizing a human-like customer service chatbot, we try to understand users’ emotions with an emotion classification model, and detect topics in user questions with a topic classification model. Then for responding users’ emotions, we design an emotional comfort framework including matching based comfort, comfort with considering both emotion and topic, and a base comfort with just considering emotion.

Text matching: text matching needs to capture the rich interaction structures in the matching process, and this process can be conducted between abstract features of two texts (Yin et al., 2016; Hu et al., 2014; Qiu and Huang, 2015) or between word embedding of two texts (Pang et al., 2016; Hu et al., 2014; Lu and Li, 2013). In papers (Yin et al., 2016; Hu et al., 2014; Qiu and Huang, 2015) (the ARC-I model in (Hu et al., 2014)), they all extract features from each of those two texts and then combine those features as the input of final Logistic regression model. In papers (Pang et al., 2016; Hu et al., 2014; Lu and Li, 2013) (the ARC-II model in (Hu et al., 2014)), they all take the interaction matrix of two texts as input of their models, and extract features from the given interaction matrix to evaluate similarity between two texts. In our matching-based emotional comfort part, we combine a BCNN model (Yin et al., 2016), which is with a text interaction on abstract feature level, and a MatchPyramid model (Pang et al., 2016), which is with a text interaction on word embedding level, to obtain an eligible performance for online service.

3 Framework Description

Our proposed framework consists of two parts (Figure 2), offline part and online part, and each of them consists of three components. With the offline part,

we want to realize the ability to understand users' emotions as detailed as possible, and with the online part, we sequentially run increasingly general comfort strategies for responding users' emotions on a larger scale.




Figure 2: Framework of emotional comfort in ICS chatbots.

Offline Part: 1) Emotion classification model is trained with considering word-level features, n-gram level features and sentence level features. We consider seven different emotions as fear, abuse, disappointed, aggrieved, anxious, anger and grateful. 2) Topic classification model is trained with a same way as the emotion classification model, and we choose 35 high frequency service classes, such as ‘complaints about the quality of service’ and ‘complaints of slow Delivery’, etc. 3) Knowledge construction is for collecting some user questions with very specific content that needs to response emotional comforts. Those specific questions are with high frequency, but they are hard to be classified into a topic or cannot get well treated with just topic-level comforts. For each question, our service experts will design a professional reply, and for each ‘question-reply’ pair we call it as a piece of ‘knowledge’.

Online Part: 1) Knowledge-based comfort is for users with specific questions, and we use a text-matching model to match a user’s question and the high-frequent questions in collected pieces of knowledge. If we can get a prepared question, which has the biggest similarity with the given user’s question and also the similarity value is bigger than a particular threshold, the corresponding reply will be taken as the emotional comfort result to this user. 2) Emotion & topic comfort means the comfort based on both users’ emotions and the topics of users’ questions. 3) Emotion-level comfort is a backup component to the emotion & topic comfort, since we cannot list all topics. So for other emotional queries without listed topics, we use this component to reply a general emotional response.

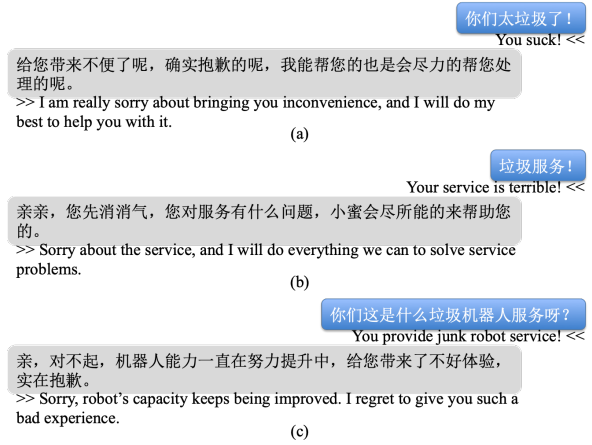


Figure 3: Examples of comforts: (a) emotion-level; (b) emotion & topic level; (c) knowledge-based level

Figure 3 gives examples of online emotional comforts. (a) shows an emotion-level comfort example. This user just complains, without any topic or any reason, so we can just give this user a very general comfort. (b) shows a comfort considering both emotions and topics. This user complains about service, so we can pointedly give a comfort about service. (c) shows a user’s complain about bad robot service, and for this kind of questions with very specific content, we utilize knowledge-based matching models to give proper responses.

4 Offline Part

4.1 Emotion Classification

Emotion classification is the base and core of whole emotional comfort framework. We propose an ensemble classification model MLC (Multi-Level feature based Classification), which combines sentence level features, n-gram level features and word-level features. Figure 4 gives the description of this model, and from left to right, sentence level features, n-gram level features and word-level features are respectively obtained. Given the word embedding of which the dimension is set as M , we also define a series of embedding of labels (emotions) of which the dimension is also set as M . Below we discuss the feature extraction steps:

1) **Sentence level features:** Simple Word-Embedding based Models (SWEM) (Shen et al., 2018), which employs simple pooling strategies operated over word embeddings, shows close performance to some classic CNN- or RNN-based text matching models or classification models. In our work we use those simple pooling strategies to obtain sentence-level features of users’ ques-

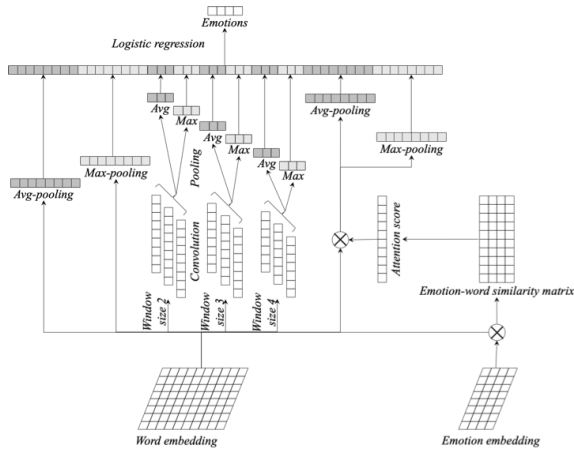


Figure 4: Emotion classification model.

tions for the emotion classification task. For combining the features obtained from average-pooling strategy and max-pooling strategy, two different methods are proposed as concatenating method and hierarchical method. Under the design idea of whole emotion classification model, we choose the SWEM-concat method to combine SWEM-max features and SWEM-avg features.

2) n -gram level features: Traditional CNN is used to obtain n -gram level features, and n is a variate denoting the convolution window size. In this paper, we set n as 2, 3 and 4 respectively, and for each window size, 16 convolution kernels are used to extract plentiful information from the original word embedding matrix. Pooling steps are similar as that in extraction of sentence level features.

3) Word-level features: We use the Label-Embedding Attentive Model (LEAM) proposed in (Wang et al., 2018) to extract word-level features. LEAM embeds the words and labels in the same joint space for text classification. It utilizes label descriptions for increasing the interaction between labels and words, which can obtains deeper consideration of semantic information of words. In our model, each ‘label’ means a kind of emotion, such as ‘anger’ or ‘disappointment’, etc. In our online service, 6 negative emotions and a ‘grateful’ emotion are considered.

Finally, features of different levels are put together for the output layer trained with logistic regression model.

4.2 Topic Classification

We summarize high frequent service topics with referring the experience of service experts, and then use the same model design with the emotion

classification step to realize topic classification.

4.3 Knowledge Construction

Besides ICS chatbots, we also have human customer services. For extracting users’ high frequent questions and also the high-quality replies, we can all refer to the chat log data of human customer services. We combine the chat log of chatbots and human customer services together, and utilize a self-adapting clustering method proposed in (Song et al.) to cluster similar user questions. With the arrangement of professional service experts, we finally choose 649 high-frequent user questions as basis of constructing ‘question-reply’ pairs. For each high-frequent user question, we collect referenceable replies from log of human customer services. Then with those referenceable replies, professional service experts can reorganize them to obtain final 649 ‘question-reply’ pairs as our ‘knowledge base’.

5 Online Part

5.1 Knowledge-based Comfort

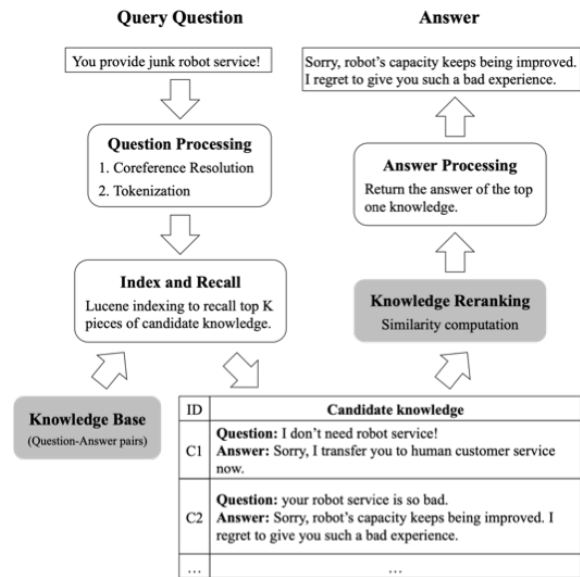


Figure 5: The workflow of retrieval-based QA systems.

We utilize a retrieval-based QA system (Yu et al., 2018) to realize knowledge-based comfort, of which the workflow is shown in figure 5. Collected knowledge base is indexed by Lucene, and for each emotional user question, we recall top K pieces of candidate knowledge from Lucene index, and then rerank those candidates to get a final reply. Similarity computation in ‘Knowledge Reranking’

module is the key component, and with different situations we have designed different models.

An unsupervised text similarity computation model: For making our framework applicable to some domains with no domain-sensitive labeled data, we use an unsupervised text matching model to rank candidates and decide which is most similar with the given user question. We use Word Similarity Maximization (WSM) (Song et al., 2017), which is an optimization of Word Mover’s Distance (WMD) proposed in (Kusner et al., 2015), to realize this unsupervised text matching step. Compared to WMD, WSM can get a normalized similarity value restricted to [0,1] instead of the distance value of WMD of which is not normalized, and computational complexity of WSM can be greatly decreased compared to WMD.

A supervised deep text similarity computation model: With the discussion of ‘text matching’ in related work section, we choose two well-performing models, MatchPyramid (Pang et al., 2016) and BCNN (Yin et al., 2016), as baselines, and we realize a combined model PBmatch, with considering features in both MatchPyramid and BCNN. Feature extraction steps of MatchPyramid and BCNN are separated and then on the Logistic regressions step, features extracted from both models are combined together, and the whole framework makes a joint training of both models.

5.2 Emotion & Topic Comfort

Emotion classification and topic classification are all run on a given user question, and for each possible ‘emotion+topic’ combination, our service experts have set different comfortable replies for realizing diversified emotional comfort. These ‘emotion+topic’ sensitive replies are randomly responded when needed.

5.3 Emotion-level Comfort

Similar with the description in above subsection, with user questions without obvious topical content, we just consider the emotional information contained in questions. For each emotion, our service experts have also set different emotion-level comfortable replies for realizing diversified emotional comfort. Compared with comfortable replies considering both emotion and topic, emotion-level comfortable replies are more general, which are like the example in figure 3(a).

6 Experiments and Evaluations

6.1 Dataset and Evaluation Metric

Dataset: 1) Emotion classification dataset: Since we annotate that just about 5% of user questions are with emotion, a manual labeling on all user questions for emotion classification is a waste. We first extract some suspicious emotional questions with an emotional dictionary, which is empirically collected, and then we published crowdsourcing tasks with checking and revising those dictionary-based labels. Each question was labeled by 3 annotators, with one of the given emotions or ‘no emotion’. If 3 annotators give 3 different labels, we delete this question, otherwise we label this question as the emotion labeled by at least 2 annotators. Finally, we got a totally 46,000 labeled questions with 8 different classes: 6 negative emotions, 1 grateful emotion and a class ‘other’.

2) Topic classification dataset: Similar with the creation of the emotion classification dataset, we also firstly extract some suspicious topical questions with an empirically collected topical dictionary, which contains 35 topics such as ‘poor service attitude’, ‘recharge slow’ and ‘urging a refund’, and similar crowdsourcing tasks were also published. Finally, we got totally 98,000 labeled questions.

3) Text matching dataset: For creating enough dataset for training the text matching model, we implement following strategies: we randomly select 10,000 user questions from chatbot log, and top 15 candidates for each of them can be obtained with Lucene index. Then 8 service experts labeled those candidates with right/wrong, and some examples are shown in Table 1. Serious data unbalance shows in above labeled data, since just 14.3% candidates are labeled as right ones (positive samples). For balancing the data, we randomly extract about 20% candidates, which are labeled as wrong, of whole dataset as negative samples.

User questions	Candidate knowledge titles	Labels
The seller does not refund, how should I do?	After my application, the seller still won't refund, how should I do?	right
	Seller does not refund shipping charge, how should I do?	wrong
	Buyer does not finish payment after a successful auction.	wrong
Fill a fault phone number	Can I change the phone number if I have filled a fault one?	right
	I filled a fault phone number, how should I do?	right
	I filled a fault phone number. Does it impact my ticket service?	wrong

Table 1: Examples of Labeled Training Dataset (Translated into English).

Evaluation Metric: User Satisfaction.

Same as other kind chatbots, accuracy rating of

single-turn response can also be taken to measure the performance of an ICS chatbot. However, ‘User Satisfaction’ is a much more important metric for ICS domain and we also take it as a mirror of the performance of our proposed framework. In practice, about 1.5K conversation sessions per day are labeled by users with a satisfaction degree of 1, 2 and 3, which respectively mean ‘very satisfied’, ‘so-so’ and ‘unsatisfied’. We take the percentage of the label ‘1’ as final ‘User Satisfaction’.

We choose the final period of data for ‘User Satisfaction’ evaluation as from Oct. 15, 2020 to Nov. 15, 2020, which consist of almost 20,000 labeled data by user research experts. Besides, our emotional comfort framework was deployed in the online system on Oct. 31, 2020.

6.2 Results and Discussions

	CNN	SWEM	LEAM	MLC
Fear	0.680	0.688	0.652	0.701
Abuse	0.940	0.925	0.889	0.945
Disappointed	0.902	0.921	0.905	0.920
Aggrieved	0.840	0.821	0.812	0.847
Anxious	0.921	0.949	0.911	0.953
Anger	0.930	0.948	0.932	0.955
Grateful	0.955	0.987	0.952	0.997
Total	0.881	0.891	0.865	0.903

Table 2: Comparison of emotion classification models

First, we check the performance of the emotion classification model. Table 2 gives an emotion-level performance comparison of different models, which are CNN, SWEM, LEAM and our model. With more diversified features, our model can get better results than all the baseline models. And a total precision of 0.903 has reached the standard of online service when we set an optimum threshold of the classification probability as 0.625. Besides, topic classification is with a same model design of emotion classification. Since the topics are too many to show up all of them, we just give a total precision result comparison in table 3.

	CNN	SWEM	LEAM	MLC
Total	0.801	0.809	0.793	0.817

Table 3: Comparison of topic classification models

Table 4 gives the comparison of different models’ performance on text matching, and we can see the PBmatch model can get a higher F-value than either BCNN or MatchPyramid models, with setting an optimum threshold. Besides, the two unsupervised models can also get passable experimental

Models	Threshold	Precision	Recall	F-value
WMD	0.73	0.823	0.782	0.802
WSM	0.75	0.845	0.823	0.834
BCNN	0.87	0.876	0.858	0.862
MatchPyramid	0.93	0.873	0.866	0.869
PBmatch	0.85	0.901	0.878	0.889

Table 4: Comparison of Text Matching Models.

results. For the Lucene recalling before the text matching step, we set the maximum number of recalled candidates as 20, considering the high ‘query per second’ (QPS) demand of our online system.

Comfort strategies Percentages	Knowledge-	Emotion- & topic-	Emotion-
	21.64%	26.69%	51.67%

Table 5: Percentages of Different Comfort Strategies.

Table 5 gives the coverages of different comfort strategies on emotional user questions. We can see the emotion-level comfort strategy is with the largest percentage, since most of the user questions are usually very short and the emotional expression of users are without specific content or specific topics.

	Without our framework	With our framework
User Satisfaction	0.214	0.301

Table 6: User Satisfaction with or without Our Framework on Negative Emotions.

Table 6 shows the comparison results of user satisfaction with or without our framework on 6 negative emotions. We can see that those chat sessions with users’ negative emotions have a very low user satisfaction, and our emotional comfort framework can help slightly raise the user satisfaction with 8.7 percent. Table 7 shows the comparison results of user satisfaction with or without our framework on the grateful emotion. With our framework, users may feel more comfortable and satisfied with the responses to their grateful emotion. So, more human-like service can get more customers’ satisfaction.

	Without our framework	With our framework
User Satisfaction	0.589	0.723

Table 7: User Satisfaction with or without Our Framework on the Grateful Emotion.

7 Conclusion

In this paper, we focus on an emotional comfort framework in e-commerce chatbots, and the experiments show such a framework can effectively

improve user satisfaction. About the future work, we will consider more emotions in this framework. Besides, we will automatically evaluate users' satisfaction with technologies on emotion analysis and sequence labeling.

References

- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2016. A simple but tough-to-beat baseline for sentence embeddings.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. *Advances in neural information processing systems*, 27:2042–2050.
- Pei Huo, Yan Yang, Jie Zhou, Chengcai Chen, and Liang He. 2020. Terg: Topic-aware emotional response generation for chatbot. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966.
- Feng-Lin Li, Minghui Qiu, Haiqing Chen, Xiongwei Wang, Xing Gao, Jun Huang, Juwei Ren, Zhongzhou Zhao, Weipeng Zhao, Lei Wang, et al. 2017. Alime assist: An intelligent assistant for creating an innovative e-commerce experience. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 2495–2498.
- Yang Li, Qingliang Miao, Ji Geng, Christoph Alt, Robert Schwarzenberg, Leonhard Hennig, Changjian Hu, and Feiyu Xu. 2018. Question answering for technical customer support. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 3–15. Springer.
- Zhengdong Lu and Hang Li. 2013. A deep architecture for matching short texts. *Advances in neural information processing systems*, 26:1367–1375.
- Yinglong Ma, Peng Zhang, and Jiangang Ma. 2018. An ontology driven knowledge block summarization approach for chinese judgment document classification. *IEEE Access*, 6:71327–71338.
- Takuma Okuda and Sanae Shoda. 2018. Ai-based chatbot service for financial industry. *Fujitsu Scientific and Technical Journal*, 54(2):4–8.
- Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Xipeng Qiu and Xuanjing Huang. 2015. Convolutional neural tensor network architecture for community-based question answering. In *Twenty-Fourth international joint conference on artificial intelligence*.
- Dinghan Shen, Guoyin Wang, Wenlin Wang, Martin Renqiang Min, Qinliang Su, Yizhe Zhang, Chunyuan Li, Ricardo Henao, and Lawrence Carin. 2018. Baseline needs more love: On simple word-embedding-based models and associated pooling mechanisms. *arXiv preprint arXiv:1805.09843*.
- Shuangyong Song, Haiqing Chen, and Zhiwei Shi. 2017. Intention classification of user queries in intelligent customer service system. In *2017 International Conference on Asian Language Processing (IALP)*, pages 83–86. IEEE.
- Shuangyong Song, Yao Meng, and Zhongguang Zheng. Summarizing microblogging users with existing well-defined hashtags. *International Journal of Asian Language Processing*, 23(2):111–125.
- Shengxian Wan, Yanyan Lan, Jiafeng Guo, Jun Xu, Liang Pang, and Xueqi Cheng. 2016. A deep architecture for semantic matching with multiple positional sentence representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. 2018. Joint embedding of words and labels for text classification. *arXiv preprint arXiv:1805.04174*.
- Wenpeng Yin, Hinrich Schütze, Bing Xiang, and Bowen Zhou. 2016. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics*, 4:259–272.
- Jianfei Yu, Minghui Qiu, Jing Jiang, Jun Huang, Shuangyong Song, Wei Chu, and Haiqing Chen. 2018. Modelling domain relationships for transfer learning on retrieval-based question answering systems in e-commerce. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 682–690.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020. The design and implementation of xiaoice, an empathetic social chatbot. *Computational Linguistics*, 46(1):53–93.
- Hongyuan Zhu, Qi Liu, Nicholas Jing Yuan, Chuan Qin, Jiawei Li, Kun Zhang, Guang Zhou, Furu Wei, Yuanchun Xu, and Enhong Chen. 2018. Xiaoice

band: A melody and arrangement generation framework for pop music. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2837–2846.

Xiaoming Zhu. 2019. Case ii (part a): Jimi’s growth path: Artificial intelligence has redefined the customer service of jd. com. In *Emerging Champions in the Digital Economy*, pages 91–103. Springer.