

Time to Take Emoji Seriously: They Vastly Improve Casual Conversational Models

Pieter Delobelle and Bettina Berendt

KU Leuven, Department of Computer Science
Celestijnenlaan 200A box 2402, 3000 Leuven, Belgium
{pieter.delobelle, bettina.berendt}@cs.kuleuven.be

Abstract. Graphical emoji are ubiquitous in modern-day online conversations. So is a single thumbs-up emoji able to signify an agreement, without any words. We argue that the current state-of-the-art systems are ill-equipped to correctly interpret these emoji, especially in a conversational context. However, in a casual context, the benefits might be high: a better understanding of users' utterances and more natural, emoji-rich responses.

With this in mind, we modify BERT to fully support emoji, both from the Unicode Standard and custom emoji. This modified BERT is then trained on a corpus of question-answer (QA) tuples with a high number of emoji, where we're able to increase the *1-of-100 accuracy* from 12.7% for the current state-of-the-art to 17.8% for our model with emoji support.

Keywords: Natural language processing · Emoji · Datasets · Conversational modeling · Representation learning

1 Introduction

The prevalent use of emoji—and their text-based precursors—is mostly unaddressed in current natural language processing (NLP) tasks. The support of the Unicode Standard [31] for emoji characters in 2010 ushered in a wide-spread, international adoption of these graphical elements in casual contexts. Interpreting the meaning of these characters has been challenging however, since they take on multiple semantic roles [24].

Whether or not emoji are used depends on the context of a text or conversation, with more formal settings generally being less tolerating. So is the popular aligned corpus Europarl [16] naturally devoid of emoji. Technical limitations, like no Unicode support, also limit its use. This in turn affects commonly used corpora, tokenizers, and pre-trained networks.

Take for example the Ubuntu Dialog Corpus by Lowe et al. [20], a commonly used corpus for multi-turn systems. This dataset was collected from an Internet Relay Chat (IRC) room casually discussing the operating system Ubuntu. IRC nodes usually support the ASCII text encoding, so there's no support for graphical emoji. However, in the

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

7,189,051 utterances, there are only 9946 happy emoticons (i.e. :-)) and the cruelly denosed :) version) and 2125 sad emoticons.

Word embeddings are also handling emoji poorly: Word2vec [22] with the commonly used pre-trained Google News vectors doesn't support the graphical emoji at all and vectors for textual emoticons are inconsistent. As another example with contextualized word embeddings, there are also no emoji or textual emoticons in the vocabulary list of BERT [6] by default and support for emoji is only recently added to the tokenizer. The same is true for GPT-2 [26]. As all downstream systems, ranging from multilingual résumé parsing to fallacy detection [5], rely on the completeness of these embeddings, this lack of emoji support can affect the performance of some of these systems.

Another challenge is that emoji usage isn't static. Think of shifting conventions, different cultures, and newly added emoji to the Unicode list. Several applications also use their own custom emoji, like chat application Slack and streaming service Twitch. This becomes an issue for methods that leverage the Unicode description [7] or that rely on manual annotations [32].

Our contribution with this paper is two-fold: firstly, we argue that the current use—or rather non-existing use—of emoji in the tokenizing, training, and the datasets themselves is insufficient. Secondly, we attempt to quantify the significance of incorporating emoji-based features by presenting a fine-tuned model. We then compare this model to a baseline, but without special attention to emoji.

Section 2 will start with an overview of work on emoji representations, emoji-based models and analysis of emoji usage. A brief introduction in conversational systems will also be given. Section 3 will then look into popular datasets with and without emoji and then introduce the dataset we used.

Our model will then be discussed in Section 4, including the tokenization in Subsection 4.1, training setup in Subsection 4.2 and evaluation in Subsection 4.3. This brings us to the results of our experiment, which is discussed in Section 5 and finally our conclusion and future work are presented in Section 6.

2 Related work

Inspired by the work on word representations, Eisner et al. [7] presented *Emoji2vec*. This system generates a vector representation that's even compatible with the *Word2vec* representations, so they can be used together. This compatibility makes it easy to quickly incorporate *Emoji2vec* in existing systems that use *Word2vec*.

The main drawback is that the emoji representations are trained on the Unicode descriptions. As a consequence, the representations only capture a limited meaning and do not account for shifting or incorrect use of emoji in the real world. For example, a peach emoji could be considered a double entendre, due to the resemblance to a woman's posterior. This is of course mentioned nowhere in the Unicode description. Which shows that the meaning of an emoji is how users interpret it, so also accidental incorrect use can cause issues [23].

In spirit, Felbo et al. [8] is similar to our work. Their system, *DeepMoji*, illustrates the importance of emoji for sentiment, emotion, and sarcasm classification. For these tasks, they used a dataset of 1246 million tweets containing at least one emoji. However,

Table 1: Overview of the support for both graphical emoji and textual emoticons for common corpora. No support is indicated with an em bar (—). Otherwise, we provide the fraction of utterances with at least one emoji or emoticon.

Dataset	Size	Language	Emoji	Emoticons
Ubuntu [20]	7,189,051	EN	—	0.17%
Amazon QA [33]	1,396,896	EN	—	0.14%
OpenSubtitles [18]	316,891,717	EN	—	—
ConvAI2 Persona-Chat [36]	47,234	EN	—	1.25%
Twitter customer support [12]	2,000	EN	8.75%	7.50%

the authors use the emoji in those tweets not for the DeepMoji model input, but as a target label. With a slightly better agreement score than humans on the sentiment task, this supports our hypothesis that emoji carry the overall meaning of an utterance.

Barbieri et al. [2] focus on a predicting one emoji based on the textual content. Interestingly, they looked into both English and Spanish tweets and compared a range of systems for a shared task at SemEval 2018: Multilingual Emoji Prediction. This shared task shows that emoji are getting more attention, but how their task is set up also highlights the current lack of high quality datasets with emoji.

The same shared task was tackled by [4] and a year later by Huang et al. [13], which made use of a pre-processor and tokenizer from Baziotis et al. [3]. This tokenizer replaces some emoji and emoticons by tokens related to their meaning. So is `\o/` replaced with `<happy>`. Naturally, this approach suffers from the same issues as described before. And even though it’s really useful to have some basic, out-of-the-box support for emoticons thanks to this work, we think that this strategy is too reducing to capture subtle nuances.

An analysis on the use of emoji on a global scale is done by Ljubešić and Fišer [19]. For this, the authors used geo-tagged tweets, which also allowed them to correlate the popularity of certain emoji with development indicators. This shows that the information encoded by emoji—and of course the accompanying tweet—is not limited to sentiment or emotion. Also Ai et al. [1] analyze the uses of emoji on social networks. Their approach consists of finding information networks between emoji and English words with LINE [30].

An interesting aspect of emoji usage is analyzed by Robertson et al. [28]. In this work, the correlation between the use of Fitzpatrick skin tone [9] modifiers and the perceived skin tone of the user. This research shows that users are inclined to use representing emoji for themselves. Robertson et al. [28] reported that no negative sentiment was associated with specific skin tone modifiers.

Conversational AI systems The research on conversational AI has been focussing on various aspects, including building high-quality datasets [11, 17, 18, 20, 27, 36], adding customizable personalities [21, 36, 37] or conjoining the efforts with regard to different datasets, models and evaluation practices [11]. With these combined efforts, businesses




Fig. 1: Illustration of the use of BERT for a response selection task.

and the general public quickly began developing ambitious use-cases, like customer support agents on social networks.

The proposed models in this field are diverse and largely depending on how the problem is formulated. When considering free-form responses, generative models like GPT [25], GPT-2 [26] or seq2seq [29] are appropriate. When the conversational task is modeled as a response selection task to pick the correct response out of N candidates [10, 11, 35], this can be a language model like BERT [6] with a dedicated head.

3 Emoji-rich datasets are hard to find

Emoji are commonly used in casual settings, like on social media or in casual conversations. In conversations—as opposed to relatively context-free social media posts—an emoji alone can be an utterance by itself. And with a direct impact for some applications, like customer support, we focus on conversational datasets. We hope the conversational community has the most direct benefit from these emoji-enabled models. Of course, the conclusions we’ll draw don’t have to be limited to this field.

Table 1 gives an overview of frequently used and interesting conversational datasets. The lacuna of emoji-rich reference datasets was already mentioned in Section 1 and is in our opinion one of the factors that emoji remain fairly underutilized.

For our models, we’ll use a customer support dataset with a relatively high usage of emoji. The dataset contains 2000 tuples collected by Hori and Hori [12] that are sourced from Twitter. They provide conversations, which consist of at least one question and one free-form answer. Some conversations are longer, in this case we ignored the previous context and only looked at the last tuple. This dataset illustrates that even when contacting companies, Twitter users keep using emoji relatively often, 8.75% of all utterances.

The tweets were filtered on hyper links and personal identifiers, but Unicode emoji characters were preserved. As emoji are frequently used on Twitter, this resulted in a dataset with 170 of the 2000 tuples containing at least one emoji character.

4 Fine-tuning BERT with emoji support

We continue training of a multilingual BERT model [6] with new tokens for emoji and fine-tune this model and a baseline on the dataset discussed in Section 4. This approach

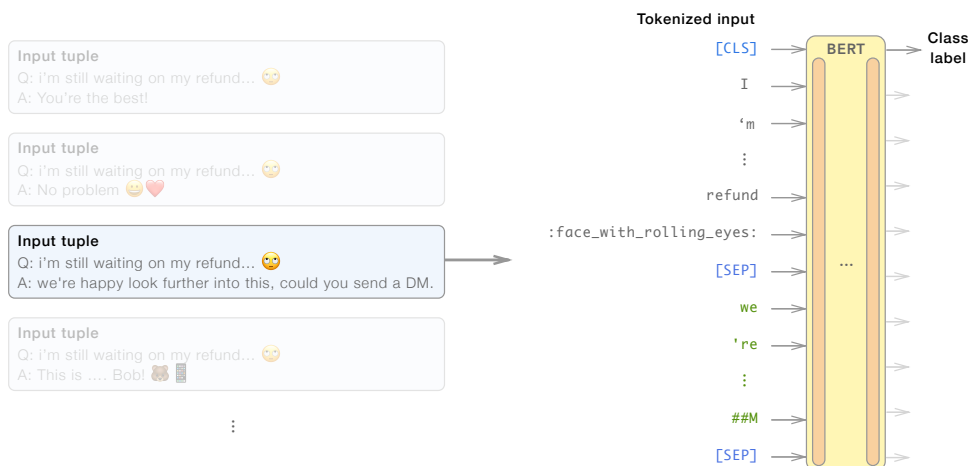


Fig. 2: Illustration the tokenization process for BERT. The illustrated input is an actual sample from the test set. Notice the emoji is replaced by a descriptive token.

is explained in Subsection 4.1 and the training itself is discussed in Subsection 4.2. At last, the evaluation is then discussed in Subsection 4.3.

4.1 Tokenizing emoji

We add new tokens to the BERT tokenizer for 2740 emoji from the Unicode Full Emoji List [31, v12.0], as well as some aliases (in the form of `:happy:` as is a common notation for emoji). In total, 3627 emoji tokens are added to the vocabulary.

We converted all UTF-8 encoded emoji to a textual alias for two reasons. First, this mitigates potential issues with text encodings that could drop the emoji. Second, this is also a common notation format for custom emoji, so we have one uniform token format. Aside from this attention to emoji, we use WordPiece embeddings [34] in the same manner as Devlin et al. [6].

4.2 Training and fine-tuning

We start from 12-headed multilingual BERT (`bert-base-multilingual-cased`), which has 110M parameters. For the model with emoji support, the number of tokens is increased, so new vectors are appended at the end of the embeddings matrix. We then continue training on the language modeling task. We use the default configuration as is also used by Devlin et al. [6] where randomly selected tokens are replaced by:

- **a mask token:** 80% chance,
- **another random word:** 10% chance,
- **the original word:** 10% chance.

This model is trained for 100 epochs with the Adam [15] optimizer. The learning rate is set to the commonly used $lr = 5 \cdot 10^{-5}$ and $\epsilon = 10^{-8}$. No hyper-parameter tuning was done, as the results are acceptable on their own and are sufficient to allow conclusions for this paper. The loss is cross entropy [14].

We then fine-tune both models, with and without emoji tokenization, on the sentence prediction task with a training set of 70%. We again use the Adam optimizer with the same settings and with binary cross entropy. In this case, the training was limited to 10 epochs. To mitigate the need for weighting and other class imbalance issues, we trained with pairs of positive and negative candidates. This is in contrast to the evaluation, where 99 negative candidates are used. However, since each candidate is considered on its own merit during evaluation, this discrepancy won't affect the performance.

For the formulation of the fine-tuning task, we use the same approach as Devlin et al. [6]. The first input sentence is joined with the second sentence, separated by a special [SEP] token, as can be seen in Figure 2. The model, with a specialized head for *next sentence prediction*, then outputs a correlation score.

4.3 evaluation metrics

Finally, our model is compared against the pre-trained version of BERT without special emoji tokens. We evaluate both this baseline and our model as a response selection task. In this case, the system has to select the most appropriate response out $N = 100$ candidates. This is a more restricted problem, where the *1-of-100 accuracy* [11] is a popular evaluation metric.

Note that 1-in-100 accuracy gives a summary of the model performance for a particular dataset. Since not all 99 negative responses are necessarily bad choices, the resulting score is in part dependent on the prior distribution of a dataset. For example, Henderson et al. [11] compares models for three datasets, where the best performing model has a score of 30.6 for OpenSubtitles [18] and 84.2 for AmazonQA [33].

Aside from the 1-of-100 accuracy, we also present the mean rank of the correct response. Since the Twitter dataset is focussed on customer service, the correct response is sometimes similar to others. The mean rank, also out of $N = 100$, can differentiate whether or not the model is still selecting good responses. For each input sentence, a rank of 1 means the positive response is ranked highest and is thus correctly selected and a rank of N signifies the positive response was—incorrectly—considered the worst-matching candidate.

5 Emoji provide additional context to response selection models

After training of the language model with additional tokens for all Unicode emoji, we achieved a final perplexity of 2.0281. For comparison, the BERT model with 16 heads achieved a perplexity of 3.23 [6], but this is on a general dataset.

For the sentence prediction task, Table 2 shows the results of the baseline and our model with additional emoji tokens. For each of the 600 utterance pairs of the held-out test set, we added 99 randomly selected negative candidates, as described in Subsection 4.3. The *1-out-of-100 accuracy* measures how often the true candidate was cor-

Table 2: 1-of-100 accuracy and mean ranking out of 100 candidates on the test set of 600 utterance tuples.

Model	1-of-100 ACC	Mean rank
bert-base-multilingual-cased	12.7%	33
bert-base-multilingual-cased with emoji tokens	17.8%	31

rectly selected and the *mean rank* gives an indication of how the model performs if it fails to correctly select the positive candidate.

The baseline correctly picks 12.7% of all candidate responses, out of 100. Given that the dataset is focussed on support questions and multiple responses are likely to be relevant, this baseline already performs admirably. For reference, a BERT model on the OpenSubtitles dataset [18] achieves a *1-of-100 accuracy* between 12.2% and 17.5%, depending on the model size [11].

Our model improves on this baseline with a *1-of-100 accuracy* of 17.8%. The mean rank remains almost the same. This indicates that the emoji tokens do help with picking the correct response, but don't really aid when selecting alternative suitable candidates. One possible explanation is that when emoji are used (this is the case for 8.75% of all utterances), including those tokens helps matching those based on those emoji and their meaning. When there are no emoji present, our model might be just as clueless as the baseline.

6 Conclusion and future work

In this paper we discussed the current state of emoji usage for conversational systems, which mainly lacks large baseline datasets. When looking at public datasets, conversational AI makers have to choose between dataset size and emoji support, with some datasets at least containing a few textual emoticons. We argued that this duality results in systems that fail to capture some information encoded in those emoji and in turn fail to respond adequately.

Based on this premise, we investigated how a response selection system based on BERT can be modified to support emoji. We proposed a format and tokenization method that's indifferent to current Unicode specifications, and thus also works for datasets containing custom emoji.

Evaluation of this emoji-aware system increased the *1-of-100 accuracy* from 12.7% for the baseline to 17.8%. Thus showing that supporting emoji correctly can help increasing performance for more casual systems, without having to rely on labeling or external descriptions for those emoji.

However, the lack of high-quality, general datasets with emoji limits our conversational model. Working towards larger casual conversational datasets would help both for our model, and for the conversational NLP community in general.

We investigated the impact of emoji for conversational models and one could argue that these conclusions—or even the BERT model—can be generalized. We didn't in-

investigate whether other tasks also benefited from our fine-tuned BERT model with the additional emoji tokens.

During evaluation, we also observed utterances with only emoji characters. Even with our model that supports emoji, it could still be difficult to extract information like the subject of a conversation. Some of these utterances—but not all—were part of a larger conversation, so an interesting question could be how additional information affects the model.

Acknowledgements

This work was supported by the Research Foundation - Flanders under EOS No. 30992574.

Bibliography

- [1] Ai, W., Lu, X., Liu, X., Wang, N., Huang, G., Mei, Q.: Untangling emoji popularity through semantic embeddings. In: Eleventh International AAAI Conference on Web and Social Media (2017)
- [2] Barbieri, F., Ronzano, F., Saggion, H.: What does this Emoji Mean? A Vector Space Skip-Gram Model for Twitter Emojis. In: LREC (2016)
- [3] Baziotis, C., Pelekis, N., Doulkeridis, C.: DataStories at SemEval-2017 Task 4: Deep LSTM with Attention for Message-level and Topic-based Sentiment Analysis. In: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), pp. 747–754, Association for Computational Linguistics, Vancouver, Canada (Aug 2017), <https://doi.org/10.18653/v1/S17-2126>
- [4] Chen, J., Yang, D., Li, X., Chen, W., Wang, T.: Peperomia at SemEval-2018 Task 2: Vector Similarity Based Approach for Emoji Prediction. In: Proceedings of The 12th International Workshop on Semantic Evaluation, pp. 428–432, Association for Computational Linguistics, New Orleans, Louisiana (Jun 2018), <https://doi.org/10.18653/v1/S18-1067>
- [5] Delobelle, P., Cunha, M., Massip Cano, E., Peperkamp, J., Berendt, B.: Computational Ad Hominem Detection. In: Proceedings of the 57th Conference of the Association for Computational Linguistics: Student Research Workshop, pp. 203–209, Association for Computational Linguistics (Jul 2019)
- [6] Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 4171–4186 (Jun 2019), <https://doi.org/10.18653/v1/N19-1423>
- [7] Eisner, B., Rocktäschel, T., Augenstein, I., Bošnjak, M., Riedel, S.: Emoji2vec: Learning Emoji Representations from their Description. In: Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media, pp. 48–54, Association for Computational Linguistics, Austin, TX, USA (Nov 2016), <https://doi.org/10.18653/v1/W16-6208>
- [8] Felbo, B., Mislove, A., Sjøgaard, A., Rahwan, I., Lehmann, S.: Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In: Proceedings of the 2017 Con-

- ference on Empirical Methods in Natural Language Processing, pp. 1615–1625, Association for Computational Linguistics, Copenhagen, Denmark (2017), <https://doi.org/10.18653/v1/D17-1169>
- [9] Fitzpatrick, T.B.: The Validity and Practicality of Sun-Reactive Skin Types I Through VI. *Archives of Dermatology* **124**(6), 869–871 (Jun 1988), ISSN 0003-987X, <https://doi.org/10.1001/archderm.1988.01670060015008>
- [10] Henderson, M., Al-Rfou, R., Strope, B., Sung, Y.h., Lukacs, L., Guo, R., Kumar, S., Miklos, B., Kurzweil, R.: Efficient Natural Language Response Suggestion for Smart Reply. arXiv:1705.00652 [cs] (May 2017)
- [11] Henderson, M., Budzianowski, P., Casanueva, I., Coope, S., Gerz, D., Kumar, G., Mrkšić, N., Spithourakis, G., Su, P.H., Vulić, I., Wen, T.H.: A Repository of Conversational Datasets. arXiv:1904.06472 [cs] (Apr 2019)
- [12] Hori, C., Hori, T.: End-to-end Conversation Modeling Track in DSTC6. arXiv:1706.07440 [cs] (Jun 2017)
- [13] Huang, C., Trabelsi, A., Zaiane, O.R.: ANA at SemEval-2019 Task 3: Contextual Emotion detection in Conversations through hierarchical LSTMs and BERT. arXiv:1904.00132 [cs] (Mar 2019)
- [14] Janocha, K., Czarnecki, W.M.: On Loss Functions for Deep Neural Networks in Classification. arXiv:1702.05659 [cs] (Feb 2017)
- [15] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs] (Dec 2014)
- [16] Koehn, P.: Europarl: A Parallel Corpus for Statistical Machine Translation p. 8 (2012)
- [17] Kummerfeld, J.K., Gouravajhala, S.R., Peper, J.J., Athreya, V., Gunasekara, C., Ganhotra, J., Patel, S.S., Polymenakos, L.C., Lasecki, W.: A Large-Scale Corpus for Conversation Disentanglement. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 3846–3856, Association for Computational Linguistics, Florence, Italy (Jul 2019)
- [18] Lison, P., Tiedemann, J.: OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles p. 7 (2016)
- [19] Ljubešić, N., Fišer, D.: A Global Analysis of Emoji Usage. In: Proceedings of the 10th Web as Corpus Workshop, pp. 82–89, Association for Computational Linguistics, Berlin (Aug 2016), <https://doi.org/10.18653/v1/W16-2610>
- [20] Lowe, R., Pow, N., Serban, I., Pineau, J.: The Ubuntu Dialogue Corpus: A Large Dataset for Research in Unstructured Multi-Turn Dialogue Systems. arXiv:1506.08909 [cs] (Jun 2015)
- [21] Mazaré, P.E., Humeau, S., Raison, M., Bordes, A.: Training Millions of Personalized Dialogue Agents. arXiv:1809.01984 [cs] (Sep 2018)
- [22] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
- [23] Miller, H., Kluver, D., Thebault-Spieker, J., Terveen, L., Hecht, B.: Understanding Emoji Ambiguity in Context: The Role of Text in Emoji-Related Miscommunication. Proceedings of the Eleventh International AAAI Conference on Web and Social Media p. 10 (2017)
- [24] Na’aman, N., Provenza, H., Montoya, O.: Varying Linguistic Purposes of Emoji in (Twitter) Context. In: Proceedings of ACL 2017, Student Research Workshop,

- pp. 136–141, Association for Computational Linguistics, Vancouver, Canada (Jul 2017)
- [25] Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving Language Understanding by Generative Pre-Training p. 12 (Jun 2018)
 - [26] Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. *OpenAI Blog* **1**(8) (2019)
 - [27] Reddy, S., Chen, D., Manning, C.D.: CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics* **7**, 249–266 (Mar 2019), <https://doi.org/10.1162/tacl.a00266>
 - [28] Robertson, A., Magdy, W., Goldwater, S.: Self-representation on twitter using emoji skin color modifiers. In: *Twelfth International AAAI Conference on Web and Social Media* (2018)
 - [29] Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to Sequence Learning with Neural Networks. *arXiv:1409.3215 [cs]* (Sep 2014)
 - [30] Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: Line: Large-scale information network embedding. In: *Proceedings of the 24th International Conference on World Wide Web*, pp. 1067–1077, International World Wide Web Conferences Steering Committee (2015)
 - [31] The Unicode Consortium: *The Unicode Standard* (2010)
 - [32] Tigwell, G.W., Flatla, D.R.: Oh that’s what you meant!: Reducing emoji misunderstanding. In: *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’16*, pp. 859–866, ACM Press, Florence, Italy (2016), ISBN 978-1-4503-4413-5, <https://doi.org/10.1145/2957265.2961844>
 - [33] Wan, M., McAuley, J.: Modeling Ambiguity, Subjectivity, and Diverging Viewpoints in Opinion Question Answering Systems. In: *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pp. 489–498, IEEE, Barcelona, Spain (Dec 2016), ISBN 978-1-5090-5473-2, <https://doi.org/10.1109/ICDM.2016.0060>
 - [34] Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, Ł., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J.: Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *arXiv:1609.08144 [cs]* (Sep 2016)
 - [35] Yu, Z., Xu, Z., Black, A.W., Rudnicky, A.: Strategy and policy learning for non-task-oriented conversational systems. In: *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pp. 404–412, Association for Computational Linguistics, Los Angeles (Sep 2016), <https://doi.org/10.18653/v1/W16-3649>
 - [36] Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., Weston, J.: Personalizing Dialogue Agents: I have a dog, do you have pets too? *arXiv:1801.07243 [cs]* (Jan 2018)
 - [37] Zhang, Z., Li, J., Zhu, P., Zhao, H., Liu, G.: Modeling Multi-turn Conversation with Deep Utterance Aggregation. *arXiv:1806.09102 [cs]* (Jun 2018)