

Conversational Response Re-ranking Based on Event Causality and Role Factored Tensor Event Embedding

Shohei Tanaka¹, Koichiro Yoshino^{1,2}, Katsuhito Sudoh¹, Satoshi Nakamura¹

¹ Nara Institute of Science and Technology

² PRESTO, Japan Science and Technology Agency

{takana.shohei.tj7, koichiro, sudoh, s-nakamura}@is.naist.jp

Abstract

We propose a novel method for selecting coherent and diverse responses for a given dialogue context. The proposed method re-ranks response candidates generated from conversational models by using event causality relations between events in a dialogue history and response candidates (e.g., “be stressed out” precedes “relieve stress”). We use distributed event representation based on the Role Factored Tensor Model for a robust matching of event causality relations due to limited event causality knowledge of the system. Experimental results showed that the proposed method improved coherency and dialogue continuity of system responses.

1 Introduction

While a variety of dialogue models such as the neural conversational model (NCM) (Vinyals and Le, 2015) have been researched widely, such dialogue models often generate simple and dull responses due to the limitation of their ability to take dialogue context into account. It is very difficult for these models to generate coherent responses to a dialogue history. We tackle this problem with a new architecture by incorporating event causality relations between response candidates and a dialogue history. Typical event causality relations are cause-effect relations between two events, such as “be stressed out” precedes “relieve stress.” In this paper, event causality relations are defined that an effect event is likely to happen after a corresponding cause event happens (Shibata and Kurohashi, 2011; Shibata et al., 2014). Event causality relations have been used in why-question answering systems to focus on causalities between questions and answers (Oh et al., 2013, 2016, 2017). It is also reported that a conversational model using event causality relations can generate diverse and coherent responses (Fujita et al., 2011). However,

the relation between dialogue continuity and the coherency of system responses is still an underlying problem.

In this paper, we propose a novel method to select an appropriate response from response candidates generated by NCMs. We define a score for re-ranking to select a response that has an event causality relation to a dialogue history. Re-ranking effectively improves response reliability in language generation tasks such as why-question answering and dialogue systems (Oh et al., 2013; Jansen et al., 2014; Bogdanova and Foster, 2016; Ohmura and Eskenazi, 2018). We used event causality pairs extracted from a large-scale corpus (Shibata and Kurohashi, 2011; Shibata et al., 2014). We also use distributed event representation based on the Role Factored Tensor Model (RFTM) (Weber et al., 2018) to realize a robust matching of event causality relations, even if these causalities are not included in the extracted event causality pairs. In human and automatic evaluations, the proposed method outperformed conventional methods in selecting coherent and diverse responses.

2 Response Re-ranking Using Event Causality Relations

Figure 1 shows an overview of the proposed method. The process consists of four parts. First, N -best response candidates are generated from an NCM given a dialogue history (Figure 1 ①; Section 2.1). Then, events (predicate-argument structures) are extracted by an event parser from both the dialogue history and the response candidates (Figure 1 ②). We used Kurohashi Nagao Parser (KNP)¹ (Kawahara and Kurohashi, 2006; Sasano and Kurohashi, 2011) as the event parser. Next, the extracted events are converted to dis-

¹<http://nlp.ist.i.kyoto-u.ac.jp/?KNP>

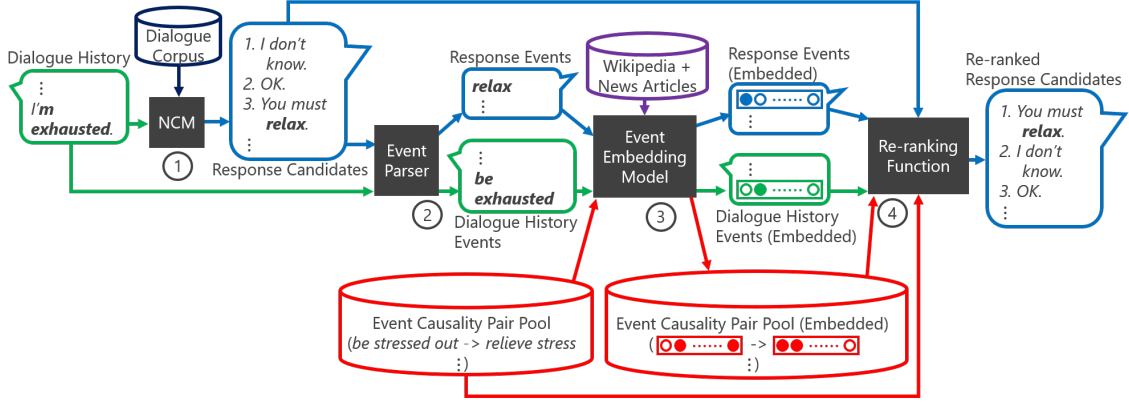


Figure 1: Neural conversational model+re-ranking using event causality; a response that has an event causality relation (“*be exhausted*” → “*relax*”) to the dialogue history is selected by the re-ranking.

predicate 1	argument 1	predicate 2	argument 2	<i>lift</i>
be stressed out	-	relieve	stress	10.02

Table 1: Example of event causality relations included in event causality pairs

tributed event representations by an event embedding model (Figure 1 ③; Section 2.3). Events in event causality pairs are also converted to distributed representations to calculate similarities. The RFTM is used for the embedding. Finally, response candidates are re-ranked (Figure 1 ④; Section 2.2, 2.4). We describe these components in more detail below.

2.1 Neural Conversational Model (NCM)

NCM learns a mapping between input and output word sequences by using recurrent neural networks (RNNs). NCMs can generate N -best response candidates by using beam search or sampling (Macherey et al., 2016).

2.2 Event Causality Pairs

The proposed method uses event causality pairs. Events in a pair, which have cause-effect relations, are extracted from a large-scale corpus on the basis of co-occurring statistics and case frames (Shibata and Kurohashi, 2011; Shibata et al., 2014). 420,000 entries are extracted from 1.6 billion texts: each entry consists of information denoted in Table 1. “predicate 1” and “argument 1” are components of a cause event, and “predicate 2” and “argument 2” are components of an effect event. Each event consists of a predicate and arguments. The predicate is required, and the argument is optional. We used arguments that have the following roles: nominative, accusative, dative, instrumental, and locative cases. *lift* is the mutual

information score between two events, which indicates the strength of the causality relation. Using *lift*, we propose a score for re-ranking as,

$$score = \max_{\langle e_h, e_r \rangle} \frac{\log_2 p}{(\log_2 lift(e_h, e_r))^\lambda}. \quad (1)$$

p is the posterior probability of the response candidate provided by NCM. λ is a hyper parameter to decide the weight of event causality relations. $lift(e_h, e_r)$ is the *lift* score between an event e_h in the dialogue history, and an event e_r in the response candidate, which is equal to 2 if the pair does not appear in the extracted event causality pair pool. Note that $lift(e_h, e_r)$ is log-scaled because it has a wide range of values ($10 < lift(e_h, e_r) < 10,000$). In the case where more than one event causality relations are recognized between the dialogue history and the response candidate, the score of the candidate is determined by the relation with the highest $lift(e_h, e_r)$. We call this model “Re-ranking.”

2.3 Distributed Event Representation Based on Role Factored Tensor Model (RFTM)

It is difficult to determine event causality relations by using only the pairs observed in an actual corpus. Therefore, we introduce a distributed event representation to improve the robustness of matching events in a dialogue with those in the event causality pair pool. Any events are embedded into fixed length vectors to calculate their similarities.




Figure 2: Model architecture of predicate embedding

We define an event with a single predicate or a pair of a predicate and arguments. Argument a of an event is embedded into vector as v_a by using Skip-gram (Mikolov et al., 2013c,a,b). Predicate p of an event is embedded into vector as v_p by using predicate embedding which is based on case-unit Skip-gram. Figure 2 shows the model architecture of predicate embedding. The model learns predicate vector representations which are good at predicting its arguments. To get an event embedding for the pair of v_p and v_a , we propose to use RFTM, which was proposed by Weber et al. (2018). The RFTM embeds a predicate and its arguments into vector e as,

$$e = \sum_a W_a T(v_p, v_a). \quad (2)$$

The relation of a predicate and its arguments is computed using a 3D tensor T and matrices W_a . If the event has no arguments, e is substituted by v_p . The RFTM is trained to predict an event sequence; thus it can represent the meaning of the event in a particular context.

2.4 Event Causality Relation Matching Based on Distributed Event Representation

Figure 3 illustrates the process of matching events on the basis of distributed event representation. Given an event pair from a response candidate and a dialogue history, the proposed method finds an event causality pair that has the highest cosine similarity from the pool. $lift$ score, strength of the event causality relation, is extended as,

$$lift_{emb}(e_h, e_r) = lift(e_c, e_e) * mean(sim(e_h, e_c), sim(e_r, e_e)). \quad (3)$$

e_h is an event in the dialogue history, e_r is an event in the response candidate. e_c and e_e are respectively a cause and an effect event of an event




Figure 3: Event causality relation matching; the $lift$ of the event causality relation in which “be exhausted” precedes “relax,” is calculated from the $lift$ of the most similar event causality relation where “be stressed out” precedes “relieve stress.”

	Ave.dist-1	Ave.dist-2
EncDec	0.44	0.56
HRED	0.33	0.42

Table 2: Diversity of N -best Response Candidates

causality pair. We also calculate the score for the case in which the cause and effect events are exchanged to deal with the inverse case. Note that both sim values have a threshold to prevent over-generalization. The threshold was empirically decided as $\sqrt{3}/2$. Replacing $lift(e_h, e_r)$ in Eq. (1) with $lift_{emb}(e_h, e_r)$, the score using distributed event representation is defined as,

$$score = \max_{\langle e_h, e_r \rangle} \frac{\log_2 p}{(\log_2 lift_{emb}(e_h, e_r))^\lambda}. \quad (4)$$

We call this model “Re-ranking (emb).”

3 Experiments

We conducted automatic and human evaluations to compare responses with and without the re-ranking. We evaluated our proposed re-ranking method on a conventional Encoder-Decoder with Attention (EncDec) model (Bahdanau et al., 2015; Luong et al., 2015) and a Hierarchical Recurrent Encoder-Decoder (HRED) model (Sordoni et al., 2015; Serban et al., 2016). While HRED tries to generate more coherent responses to dialogue context than a simple Encoder-Decoder, the diversity of responses is small due to context constraints.

We used the Japanese data from a Wikipedia dump for training Skip-gram and predicate word embeddings of RFTM, and the *Maichichi* newspaper dataset 2017² for training RFTM. We col-

²<http://www.nichigai.co.jp/sales/mainichi/mainichi-data.html>

Method			Evaluation							
NCM	history	re-ranking	re-ranked (%)	BLEU	NIST	extrema	dist-1	dist-2	PMI	length
reference	-	-	-	-	-	-	0.06	0.40	1.86	21.43
EncDec	-	1-best	-	1.12	1.19	0.42	0.06	0.18	1.77	15.55
EncDec	1	Re-ranking	4,016 (7.90)	1.10	1.18	0.42	0.06	0.19	1.78	15.52
EncDec	1	Re-ranking (emb)	29,343 (57.71)	1.02	1.07	0.40	0.06	0.20	1.77	15.64
EncDec	5	Re-ranking	6,469 (12.72)	1.09	1.17	0.42	0.06	0.19	1.78	15.50
EncDec	5	Re-ranking (emb)	35,284 (69.39)	1.00	1.04	0.39	0.07	0.21	1.77	15.66
HRED	-	1-best	-	1.34	2.74	0.42	0.07	0.20	1.84	35.05
HRED	1	Re-ranking	3,671 (7.22)	1.33	2.74	0.42	0.06	0.20	1.84	35.20
HRED	1	Re-ranking (emb)	30,992 (60.95)	1.28	2.74	0.41	0.06	0.20	1.86	34.80
HRED	5	Re-ranking	6,231 (12.25)	1.33	2.73	0.42	0.06	0.20	1.84	35.30
HRED	5	Re-ranking (emb)	36,373(71.53)	1.28	2.74	0.41	0.06	0.20	1.86	34.60

Table 3: Comparison results before and after re-ranking

lected 2,632,114 dialogues from Japanese micro blogs (Twitter) to train and test the dialogue models. The average dialogue turn was 21.99, and the average utterance length was 22.08 words. We removed emoticons from utterances to reduce vocabulary size and accelerate the training. The dialogue corpus was split into 2,509,836, 63,308, and 58,970 dialogues as training, validation, and testing data, respectively.

3.1 Model Settings

The hidden unit size of Skip-gram (Mikolov et al., 2013c,a,b), predicate embedding, and RFTM (Weber et al., 2018) was 100. We used gated recurrent units (GRUs) (Cho et al., 2014; Chung et al., 2014) whose number of layers was 2 and hidden unit size was 256, for the encoder and decoder of the NCMs. The batch size was 100, the dropout probability was 0.1, and the teacher forcing rate was 1.0. We used Adam (Kingma and Ba, 2015) as the optimizer. The gradient clipping was 50, the learning rate for the encoder and the context RNN of HRED was $1e^{-4}$, and the learning rate for the decoder was $5e^{-4}$. The loss function was inverse token frequency (ITF) loss (Nakamura et al., 2019). We used sentencepiece (Kudo and Richardson, 2018) as the tokenizer, and the vocabulary size was 32,000. These settings were the same in all models.

Repetitive suppression (Nakamura et al., 2019) and length normalization (Macherey et al., 2016) were used at the decoding step. Finally, λ of Eq. (1) and Eq. (4) was set to 1.0.

3.2 Diversity of Beam Search

We investigated internal diversity of N -best response candidates generated from each dialogue model. It is expected that the higher diversity is, the more effective re-ranking is. Hence, we evaluated diversity on the test data by dist-1, 2 (Li et al., 2016). Beam width was set to 20; it is same in the

following experiments.

The result is shown in Table 2: Ave.dists are averages of dist computed internal N -best response candidates. The diversity of EncDec is higher than that of HRED.

3.3 Comparison in Automatic Metrics

Table 3 shows the results of our evaluation using automatic metrics. We compared the results by referring to the ratio of responses different from the without re-ranking method (“re-ranked”), bilingual evaluation understudy (BLEU) (Papineni et al., 2002), NIST (Doddington, 2002), and vector extrema (Gabriel et al., 2014) (“extrema”) score. NIST is based on BLEU, but heavily weights less frequent N-grams to focus on content words. Vector extrema computes cosine similarity between sentence vectors of a reference and a generated response from a model. Each sentence vector e_s is computed by taking extrema of Skip-gram word vectors e_w in each dimension d as,

$$e_{sd} = \begin{cases} \max_{w \in s} e_{wd} & \text{if } e_{wd} > |\min_{w' \in s} e_{w'd}| \\ \min_{w \in s} e_{wd} & \text{otherwise} \end{cases} \quad (5)$$

e_{sd} and e_{wd} are the d th dimensions of e_s and e_w respectively. Additionally, we evaluated dist (Li et al., 2016), Pointwise Mutual Information (PMI) (Newman et al., 2010), and average response length (“length”). Dist and PMI are used to evaluate diversity and coherency respectively. PMI between a response and a dialogue history is defined as,

$$\text{PMI} = \frac{1}{|\text{response}|} \sum_{wr}^{|\text{response}|} \max_{wh} \text{PMI}(wr, wh). \quad (6)$$

wr and wh are words in the response and the dialogue history respectively. Each method used a

	word coherency	dialogue continuity
1-best	28.62	40.84
Re-ranking	33.91	38.53
neither	37.47	20.62

Table 4: 1-best v.s. Re-ranking; # dialogues is 100.

	word coherency	dialogue continuity
1-best	30.10	35.50
Re-ranking (emb)	25.40	38.20
neither	44.50	26.30

Table 5: 1-best v.s. Re-ranking (emb); # dialogues is 100.

specific NCM, a range of dialogue history used for re-ranking, and re-ranking method. Methods with “1-best” used neither re-ranking and event embedding. Those with “Re-ranking” used re-ranking but did not use event embedding. Those with “Re-ranking (emb)” used both the re-ranking and the proposed event embedding method.

Re-ranking lowered scores of the similarity to reference: BLEU, NIST, and extrema, because normal NCM models were trained to generate similar responses to the references, generated top 1 response before re-ranking should have the highest scores in those similarity metrics. Dist-2 and PMI were improved by re-ranking. This indicates that words in re-ranked responses are diverse and coherent to dialogue histories. However, ratios of re-ranked responses were around 10%; hence, the effect of re-ranking was limited. By introducing the proposed event embedding method, the ratios of re-ranked responses improved drastically (Re-ranking vs. Re-ranking (emb)). Moreover, the re-ranking models with event embedding have highest dist-1, dist-2, and PMI. As the HRED models had higher BLEU, NIST, and PMI values than those of EncDec models in all re-ranking methods, we conducted a human evaluation by comparing HRED model-based systems.

3.4 Human Evaluation

It is difficult to evaluate system performances only with automatic metrics (Liu et al., 2016). Hence, we compared a baseline model and our models in a human evaluation to confirm coherency and dialogue continuity of responses selected by our proposed methods. We compared baseline HRED model with our proposed models, re-ranked without embedding and with embedding using the last

	word coherency	dialogue continuity
Re-ranking	23.70	35.53
Re-ranking (emb)	22.91	35.65
neither	55.39	28.83

Table 6: Re-ranking v.s. Re-ranking (emb); # dialogues is 100.

five histories. To reduce evaluators’ workload, we used test data whose the number of user utterances is less than three, and removed dialogues which need external knowledge to evaluate. We used crowdsourcing for the human evaluation. Ten crowd-workers compared responses selected by two of three models in the following two subjective criteria. The first one is “which words in a response are more related to a dialogue history” (word coherency), which indicates system response coherency to dialogue histories. The second criterion is “which response is easier to respond to” (dialogue continuity), which indicates how much dialogue continuity system responses have. We were inspired to make these criteria by those of the Alexa Prize (Ram et al., 2018).

The results are shown in Table 4, 5, and 6. Word coherency was improved by our model without embedding, but lowered by the model with embedding. This is because workers acknowledged causality relations included in the event causality pair pool, but did not acknowledge generalized causalities with event embedding. However, dialogue continuity was improved by the proposed re-ranking model with embedding, it is probably because the proposed model reduced the number of dull responses. We need to investigate the better threshold in the event embedding to balance out the coherency and the continuity as the future work.

4 Discussion

We analyzed an adequacy of re-ranking using event causality relations. Here are system response examples of our proposed method. “()” indicates original Japanese sentences, “[]” indicates event causality relations used for re-ranking, and “<>” indicates responses before re-ranking. All examples are translated from Japanese to English.

Conversation 1:

User 1: Because of my fears, I have **been stressed out.**

(*Mou fuan-na koto ga oosugite sutoresu ga tamatteku.*)

User 2 (System): Are you OK? Don't work too hard.

β (*Daijobu desuka muri shinaide kudasaine*)

[work too hard → be stressed out
(*muri wo suru → sutoresu ga tamaru*)]
<Are you OK? (*Daijobu desuka*)>

Conversation 2:

User 1: Someone, let's go to eat lunch around *Yayoi-ga-oka* tomorrow. Everyone is welcome.

(*Dareka, ashita, yayoi-ga-oka shuhen de hiru-gohan tabe ni iko daredemo iiken onegai*)

User 2 : Let's go. (*Ikimasuka*)

User 1: Where will we go?

(*Doko kui iku*)

User 2 (System): Anywhere near *Shinjuku* station

(*Shinjuku-eki no chikaku*)

[be near → go (*chiaku → ikeru*)]
<Where are you? (*doko inno*)>

In above examples, appropriate event causality relations are used to select logical, coherent, and diverse responses. However, we found that these cases are not majority. Our method used inadequate event causality relations even if coherent responses were selected as a result (“Conversation 3”). Responses selected by our method are sometimes unnatural and incoherent than those before re-ranking as “Conversation 4”, “Conversation 5”, or “Conversation 6”.

Conversation 3:

User 1: I can't stop coughing:_(

(*Hen na seki ga tomaranai...namida*)

User 2 (System): Don't catch a cold.

(*kaze hikanai dene*)

[catch cold → stop coughing
(*kaze wo hiku → seki ga tomaru*)]
<Are you OK? (*daijobu-ka*)>

Conversation 4:

User 1: Our president who cannot drink alcohol gives me beer...

(*Osake wo nomenai shacho ga biru wo kureru...*)

User 2 (System): I will drink alcohol.

(*Osake nomu.*)

[drink alcohol → can drink alcohol
(*osake wo nomu → osake wo nomeru*)]
<I see. (*wakaru*)>

Conversation 5:

User 1: Despite of morning, ambulances and police cars were noisy... I was wondered what happened?

(*Asa nanoni, kyu-kyu-sha to patoka no oto ga sugoi...doshitandaro*)

User 2 (System): Good morning.

(*Ohayogozaimasu*)

[morning → good morning

(*asa → ohayo*)

<Are you OK? (*Daijobu desuka*)>

Conversation 6:

User 1: I, I can't get up... I'm gonna take a day off... I have a fever...

(*O, okiagarenai... kaisha yasumu... netsu...*)

User 2 (System): You must be tired.

(*Otsukaresama desu*)

[be tired → take a day off

(*tsukare → yasumu*)

<Take care of yourself. (*Odaiji-ni*)>

Considering the result of human evaluation and above examples, we hypothesized that our method have two problems in selecting appropriate event causality relations. The first problem is that the event embedding over-generalized events (“Conversation 4”). The causality in Conversation 4 (“drink alcohol” precedes “can drink alcohol”) is obtained by generalizing a causality that “enter restaurant” precedes “order beer”, which is included in the event causality pair pool. It is necessary to prevent over-generalization by improving the embedding architecture. The second problem is that our method focuses on only word coherency, not response naturalness (“Conversation 5” and “Conversation 6”). To solve the problem, our method has to maintain response naturalness while improving coherency of word choices.

5 Conclusion

We proposed a selection of response candidates generated from a neural conversational model (NCM) utilizing event causality relations. The method had a robust matching of event causality relations attributed to distributed event representation. Experimental results showed that the proposed method selects a coherent and diverse response. The proposed method can be applied to any languages that have a semantic parser, because it uses predicate-argument structure based event expressions. However, unnatural responses were sometimes selected due to inadequate event causality relations. Future work will focus on solving the problem by preventing over-generalization of events, and maintaining response naturalness.

Acknowledgments

We would like to thank Sadao Kurohashi, Ph.D. and Tomohide Shibata, Ph.D. of Kurohashi Laboratory in Kyoto University who provided us the event causality pairs.

This work is supported by JST PRESTO (JP-MJPR165B).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Dasha Bogdanova and Jennifer Foster. 2016. This is how we do it: Answer Reranking for Open-Domain How Questions with Paragraph Vectors and Minimal Feature Engineering. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1290–1295.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. In *Proceedings of the 28th Conference Neural Information Processing Systems, Deep Learning and Representation Learning Workshop (NIPS)*.
- George Doddington. 2002. Automatic Evaluation of Machine Translation Quality Using N-gram Co-occurrence Statistics. In *Proceedings of the 2nd International Conference on Human Language Technology Research (HLT)*, pages 138–145.
- Motoyasu Fujita, Rafal Rzepka, and Kenji Araki. 2011. Evaluation of Utterances Based on Causal Knowledge Retrieved from Blogs. In *Proceedings of the 14th IASTED International Conference Artificial Intelligence and Soft Computing (ASC)*, pages 294–299.
- Forgues Gabriel, Joelle Pineau, Jean-Marie Larchevêque, and Réal Tremblay. 2014. Bootstrapping Dialog Systems with Word Embeddings.
- Peter Jansen, Mihai Surdeanu, and Peter Clark. 2014. Discourse Complements Lexical Semantics for Non-factoid Answer Reranking. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 977–986.
- Daisuke Kawahara and Sadao Kurohashi. 2006. A Fully-Lexicalized Probabilistic Model for Japanese Syntactic and Case Structure Analysis. In *Proceedings of Human Language Technology Conference - North American Chapter of the Association for Computational Linguistics Annual Meeting (HLT-NAACL)*, pages 176–183.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Taku Kudo and John Richardson. 2018. Sentence-Piece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 110–119.
- Chia-Wei Liu, Ryan Lowe, Iulian V. Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. 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 (EMNLP)*.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-Based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s Neural Machine Translation System: Bridging the Gap Between Human and Machine Translation. In *arXiv:1609.08144*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Deany. 2013a. Efficient Estimation of Word Representations in Vector Space. In *Proceedings of the 1st International Conference on Learning Representations (ICLR)*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems (NIPS)*, volume 2, pages 3111–3119.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013c. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings of the 12th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 746–751.
- Ryo Nakamura, Katsuhito Sudoh, Koichiro Yoshino, and Satoshi Nakamura. 2019. Another Diversity-Promoting Objective Function for Neural Dialogue Generation. In *Proceedings of the 33rd Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence, Workshop on Reasoning and Learning for Human-Machine Dialogues (DEEP-DIAL 2019) (AAAI)*.
- David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. 2010. Automatic Evaluation of Topic Coherence. In *Proceedings of the 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 100–108.
- Jong-Hoon Oh, Kentaro Torisawa, Chikara Hashimoto, Ryu Iida, Masahiro Tanaka, and Julien Kloetzer. 2016. A Semi-supervised Learning Approach to Why-Question Answering. In *Proceedings of the 30th Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence (AAAI)*, pages 3022–3029.
- Jong-Hoon Oh, Kentaro Torisawa, Chikara Hashimoto, Motoki Sano, Stijn De Saeger, and Kiyonori Ohtake. 2013. Why-Question Answering Using Intra- and Inter-Sentential Causal Relations. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1733–1743.
- Jong-Hoon Oh, Kentaro Torisawa, Canasai Kruegkrai, Ryu Iida, and Julien Kloetzer. 2017. Multi-Column Convolutional Neural Networks with Causality-Attention for Why-Question Answering. In *Proceedings of the 10th Association for Computing Machinery International Conference on Web Search and Data Mining (WSDM)*, pages 415–424.
- Junki Ohmura and Maxine Eskenazi. 2018. Context-Aware Dialog Re-ranking for Task-Oriented Dialog Systems. In *Proceedings of IEEE Spoken Language Technology Workshop (SLT)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2018. Conversational AI: The Science Behind the Alexa Prize. In *arXiv:1801.03604*.
- Ryohei Sasano and Sadao Kurohashi. 2011. A Discriminative Approach to Japanese Zero Anaphora Resolution with Large-Scale Lexicalized Case Frames. In *Proceedings of the 5th International Joint Conference on Natural Language Processing (IJCNLP)*, pages 758–766.
- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. In *Proceedings of the 30th Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence (AAAI)*.
- Tomohide Shibata, Shotaro Kohama, and Sadao Kurohashi. 2014. A Large Scale Database of Strongly-Related Events in Japanese. In *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC)*.
- Tomohide Shibata and Sadao Kurohashi. 2011. Acquiring Strongly-Related Events Using Predicate-Argument Co-occurring Statistics and Case Frames. In *Proceedings of the 5th International Joint Conference on Natural Language Processing (IJCNLP)*, pages 1028–1036.
- Alessandro Sordoni, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob G. Simonsen, and Jian-Yun Nie. 2015. A Hierarchical Recurrent Encoder-Decoder For Generative Context-Aware Query Suggestion. In *Proceedings of the 24th Association for Computing Machinery International Conference on Information Knowledge and Management (ACM)*.
- Oriol Vinyals and Quoc V. Le. 2015. A Neural Conversational Model. In *Proceedings of the 32nd International Conference on Machine Learning, Deep Learning Workshop (ICML)*.

Noah Weber, Niranjan Balasubramanian, and Nathanael Chambers. 2018. Event Representations with Tensor-Based Compositions. In *Proceedings of the 32nd Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence (AAAI)*.