

N-best Response-based Analysis of Contradiction-awareness in Neural Response Generation Models

Shiki Sato¹ Reina Akama^{1,3} Hiroki Ouchi^{2,3} Ryoko Tokuhisa¹
Jun Suzuki^{1,3} Kentaro Inui^{1,3}

¹Tohoku University ²Nara Institute of Science and Technology ³RIKEN
{shiki.sato.d1, akama, tokuhisa, jun.suzuki, inui}@tohoku.ac.jp
hiroki.ouchi@is.naist.jp

Abstract

Avoiding the generation of responses that contradict the preceding context is a significant challenge in dialogue response generation. One feasible method is post-processing, such as filtering out contradicting responses from a resulting n -best response list. In this scenario, the quality of the n -best list considerably affects the occurrence of contradictions because the final response is chosen from this n -best list. This study quantitatively analyzes the contextual contradiction-awareness of neural response generation models using the consistency of the n -best lists. Particularly, we used polar questions as stimulus inputs for concise and quantitative analyses. Our tests illustrate the contradiction-awareness of recent neural response generation models and methodologies, followed by a discussion of their properties and limitations.

1 Introduction

Recent advanced response generation models (Zhang et al., 2020; Adiwardana et al., 2020; Roller et al., 2021) can generate relevant and meaningful responses, which can resolve dull response problems (Vinyals and Le, 2015; Sordoni et al., 2015; Serban et al., 2016). This advancement reveals additional flaws in the quality of neural model responses, such as *contradiction*. Contradiction is a critical error in dialogue because a single contradictory response can disrupt the flow of the dialogue (Higashinaka et al., 2015).

A generation model outputs a response by selecting the candidate with the highest likelihood (1-best) from an n -best candidate list. Prior work has demonstrated that generating the n -best lists with noncontradictory 1-bests is an open challenge (Nie et al., 2020; Kim et al., 2020; Li et al., 2021). Thus, one practical technique for avoiding contradiction is to have an accurate contradiction detector that eliminates all contradictory candidates from the n -best list (Nie et al., 2020). In this scenario, the con-

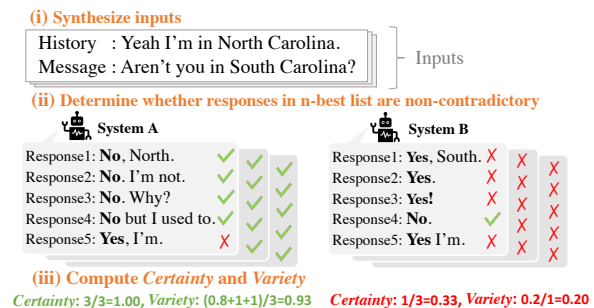


Figure 1: Overview of our analysis framework. The framework analyzes n -best lists by (i) synthesizing a stimulus input that induces contradictions, (ii) automatically determining whether responses in the n -best lists are contradictory, and (iii) computing *Certainty* and *Variety*.

sistency of all candidates in the n -best list, not just the 1-best, substantially impacts whether the final output is contradictory because the final response is chosen from the n -best list. Nonetheless, earlier quantitative investigations of contradiction relied solely on 1-bests from models (Li et al., 2021).

In this study, we analyze the n -best lists generated by the models to explore methods for enhancing neural response generation to avoid contradiction. Specifically, we first consider how analyzing an n -best list should be approached. Then, we propose a method for statistically analyzing the n -best lists (Figure 1). Since it is impractical to study all conceivable contradictions in a dialogue, we first focus on contradictions in response to polar questions.¹ We use our method to highlight the contradiction-awareness of recent high-performance neural response generation models and methodologies. Our results show that beam search has limitations in terms of avoiding contradiction and that the newer techniques, such as unlikelihood training (Welleck et al., 2020), can help overcome these limitations.

¹Codes and test set are available at <https://github.com/shiki-sato/nbest-contradiction-analysis>

NLI data			Dialogue context for our test	
Entailment	Premise: yeah i'm in North Carolina Hypothesis: I'm in North Carolina.	→	ENTQ	History: Yeah I'm in North Carolina. Message: Are you in North Carolina?
Contradiction	Premise: yeah i'm in North Carolina Hypothesis: I'm in South Carolina.	→	CNTQ	History: Yeah I'm in North Carolina. Message: Aren't you in South Carolina?

Table 1: Acquiring dialogue context by transforming the Natural Language Inference (NLI) data.

2 Analysis perspectives

First, n -best lists must be generated to prevent contradiction, assuming the filters can remove contradictory responses. An ideal model produces output that is noncontradictory and outperforms in many other criteria, such as relevance or informativeness. A model must generate at least one noncontradictory candidate to deliver a noncontradictory output. Furthermore, even noncontradictory candidates could be eliminated based on other criteria (e.g., relevance, informativeness). Therefore, it can be hypothesized that having more noncontradictory responses in an n -best list would enhance the final output quality across various criteria. Taking the above into account, we examine n -best lists based on the certainty of the existence of noncontradictory responses (**Certainty**), and the variety of noncontradictory responses (**Variety**):

- **Certainty**: The proportion of the n -best lists that have at least one noncontradictory response.
- **Variety**: The proportion of noncontradictory responses in each n -best list when only the n -best lists with at least one noncontradictory response are collected.

Given a set of inputs \mathcal{Q} , we calculate them as follows:

$$\mathbf{Certainty} = \frac{|\mathcal{Q}'|}{|\mathcal{Q}|}, \mathbf{Variety} = \frac{1}{|\mathcal{Q}'|} \sum_{q \in \mathcal{Q}'} \frac{\text{cnt}(f(q))}{|f(q)|}$$

$$\mathcal{Q}' = \{q \mid \text{cnt}(f(q)) > 0, q \in \mathcal{Q}\}$$

where $f(\cdot)$ is an n -best list generation function and $\text{cnt}(\cdot)$ is a function that returns the number of noncontradictory responses from a given n -best list. For example, the **Certainty** of a model that generates n -best lists with a combination of noncontradictory and contradictory responses is high, but its **Variety** is low. However, a model that always generates n -best lists with only noncontradictory or contradictory responses has a high **Variety** but a low

Certainty. We anticipate that n -best lists must include noncontradictory responses (**Certainty**= 1.0), with a high proportion (high **Variety**).

3 Analytical inputs and evaluation

To analyze a model from the aforementioned viewpoints, we consider how to prepare the analytical inputs and evaluate the generated responses in this section.

3.1 Inputs for highlighting contradictions

Polar echo question. An *echo question* (Noh, 1998) confirms or clarifies the context information by repeating the utterance of another speaker. It is commonly used when the speaker did not hear or understand what was said correctly, or when the speaker wishes to express incredulity. Based on Li et al. (2021)'s discovery, contradictions emerge mostly when speakers refer to earlier information communicated in dialogue; we use echo questions as stimulus input in our analysis to elicit contradictory responses. We use polar-typed echo questions to make our analysis more succinct and quantitative. Since polar questions allow for basically only two responses, *yes* or *no*, we can clearly determine whether the generated response is contradictory or not. Furthermore, by analyzing the produced responses as a yes/no binary classification issue, it allows for quantitative discussion of experimental outcomes based on the probability level.

Input preparation. We use the dataset from the natural language inference (NLI) task to effectively obtain the analytical inputs described in the preceding paragraph. This dataset specifies the logical relationship (i.e., entailment, neutrality, or contradiction) between a premise and its associated hypothesis. We transform the NLI dataset into dialogue data using a set of basic rewriting rules.² Our test involves two types of inputs, which can be classified as follows:

- **ENTQ**: generating a *confirmation* response.

²The details are described in Appendix A.

- CNTQ: generating a *refutation* response.

Table 1 displays the input samples and how they are transformed from the initial NLI data. Each input is made up of the following two utterances: the history and message. In our analysis, the model generates responses to a given message, assuming the model has generated the history in the preceding turn.

3.2 Contradiction detection for output

To compute the *Certainty* and *Variety*, we must first determine whether each generated response in the n -bests compared to the inputs is contradictory. The simplest method for detecting the contradictions is to check whether the response begins with *yes* or *no*. However, in the event of an indirect expression (e.g., *Why not?*), this method cannot detect the contradictions. Therefore, we use an automated yes-no classifier to categorize the n -best responses to ENTQ/CNTQ. We train the classifier by fine-tuning RoBERTa (Liu et al., 2019) using the Circa dataset (Louis et al., 2020), which comprises pairs of polar questions and indirect responses, as well as annotations for the answer’s interpretation, to categorize utterances as affirmations or refutations.³

4 Experiments

We demonstrate how our framework shows the properties of n -best lists, which could be quite influential in terms of avoiding contradiction. We demonstrate this by comparing the n -bests generated by conventional beam search (BS) versus recently proposed techniques.

4.1 Experimental settings

Inputs preparation. We used the Multi-Genre NLI Corpus (Williams et al., 2018) to obtain analytical inputs, which is a large scale and is consistent in good quality NLI data. We created 2,000 ENTQ/CNTQ inputs by extracting 2,000 samples labeled with *entailment* or *contradiction*.⁴

Response generation models. We used the following two recently developed high-performance models: DialoGPT (Zhang et al., 2020) and Blender (Roller et al., 2021).⁵

³The details are described in Appendix B.

⁴We used the samples in the TELEPHONE domain; this domain covers open-domain conversations.

⁵The details of the settings are described in Appendix C.

Model	<i>Certainty</i>		<i>Variety</i>	
	ENTQ	CNTQ	ENTQ	CNTQ
Blender 400M	0.806	0.747	0.780	0.775
Blender 1B	0.832	0.752	0.832	0.753
Blender 3B	0.856	0.768	0.824	0.737
DialoGPT 345M	0.938	0.917	0.750	0.669
DialoGPT 762M	0.883	0.918	0.671	0.713

Table 2: *Certainty* and *Variety* of 10-best lists using beam search with beam size $B = 10$.

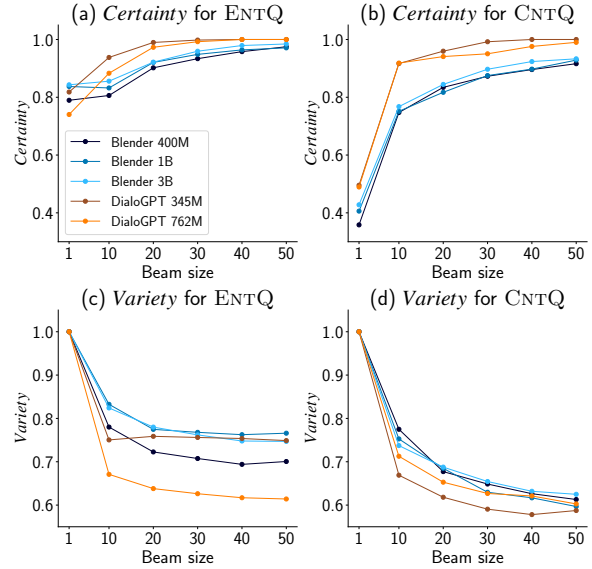


Figure 2: *Certainty* and *Variety* of n -best lists using beam search with various beam sizes.

4.2 Analysis of n -best using beam search

Let B denote the beam size during generation. It has been empirically found that using beam search with $B = 10$ to generate a response yields excellent quality results and has a frequently used value (Zhang et al., 2020; Roller et al., 2021). Table 2 displays the *Certainty* and *Variety* of 10-best lists generated using these methods. Figure 2 also depicts the *Certainty* and *Variety* of n -best lists generated using different beam sizes.

Certainty. Table 2 illustrates that in approximately 10% of CNTQ-type inputs, even the highest scoring model generates 10-best lists full of contradictory responses. Even with a perfect response filter, the models are unable to provide noncontradictory answers to these questions. It should be emphasized that the error rate is not low, given that the inputs are polar questions with highly restricted viable responses. Expanding the beam size can increase the number of n -best lists with at least one noncontradictory response. Indeed, increas-

ing the beam size enhances the *Certainty* ((a) and (b) in Figure 2). By increasing B to 40, the *Certainty* of using DialoGPT 345M for both ENTQ- and CNTQ-type inputs achieve 1.0.

Variety. With $B = 10$, all the models’ *Variety* are more than 0.5 (chance rate) (Table 2). Therefore, rather than being fully random, the models generate n -best lists with a degree of directionality toward avoiding contradictions. However, increasing the size of beam reduces the *Variety* ((c) and (d) in Figure 2), resulting in lower output quality. For example, the *Variety* of DialoGPT 345M with $B = 40$ for CNTQ-type inputs (a model with *Certainty* of 1.0 for both ENTQ- and CNTQ-type inputs) decreases to 0.58.

Overall. In terms of avoiding contradiction, our analytical framework demonstrated the features of the n -best lists of the beam search. The *Certainty* did not achieve 1.0 in the commonly used configuration ($B = 10$). When the beam size is increased, the *Certainty* increases to 1.0, whereas the *Variety* reduces dramatically. These results show the trade-off between *Certainty* and *Variety* as a function of beam size; in this example, we found constraints in obtaining high *Certainty* and *Variety* with beam search. Furthermore, it is found that the *Certainty* obtained using DialoGPT is greater than that obtained using Blender, whereas the opposite is true for *Variety*, suggesting that various models behave differently in terms of *Certainty* and *Variety*. This study emphasizes the significance of examining the *Certainty* and *Variety* of each model.

4.3 Analysis of n -best by various techniques

How to achieve high *Certainty* and *Variety*?

One method to increase *Certainty* is to generate n -best lists with a wider range of responses, such that each n -best list is guaranteed to contain a specific number of noncontradictory responses. The diverse beam search (DBS) (Vijayakumar et al., 2016) and nucleus sampling (NS) (Holtzman et al., 2020) methods are used to construct such n -best lists. Furthermore, Li et al. (2020) recently proposed models that use unlikelihood (UL) training to assign low probabilities to contradict responses. Using these models to generate n -best lists will almost certainly enhance both *Certainty* and *Variety*. We assess the n -best lists generated using these three strategies to see how much these techniques enhance *Certainty* and *Variety* (n -best lists

Technique	<i>Certainty</i>		<i>Variety</i>	
	ENTQ	CNTQ	ENTQ	CNTQ
BS	0.856	0.768	0.824	0.737
DBS	0.999	0.981	0.758	0.478
NS	1.000	0.994	0.755	0.462
UL ($\alpha = 0$)	1.000	0.996	0.406	0.759
UL ($\alpha = 1$)	0.943	0.900	0.920	0.938
UL ($\alpha = 10$)	0.910	0.937	0.969	0.968

Table 3: *Certainty* and *Variety* of 10-best lists using various techniques with Blender 3B.

generated using DBS and NS, and n -best lists generated using beam search together with the UL training). Appendix C contains a description of the techniques used for this analysis.

Result. Table 3 displays the *Certainty* and *Variety* of the 10-best lists generated using BS, DBS, NS, and UL.⁶ The values of α show the degree of UL loss during fine-tuning. Here UL with $\alpha = 0$ used the response generation model fine-tuned with maximum likelihood in the same training settings as those used for UL with $\alpha > 0$. Thus, note that comparing UL with $\alpha = 0$ and $\alpha > 0$ allows a fair comparison between likelihood and unlikelihood training. The results reveal the properties of the n -best lists obtained for the three techniques, as well as the extent to which the techniques increase *Certainty* and *Variety*. The *Certainty* obtained using the DBS and NS method reach 1.0 for significantly lower search sizes than that for the BS to attain a *Certainty* of 1.0; the *Variety* for CNTQ-type inputs are less than 0.5 (chance rate). Thus, using the DBS and NS methods efficiently improves *Certainty* compared with the results obtained using the beam search; nevertheless, the methods do not simultaneously attain high *Certainty* and *Variety*. However, the *Certainty* obtained using UL with $\alpha > 0$ are greater than those obtained using the BS, and this was accomplished while maintaining higher *Variety* than those obtained using the BS and UL with $\alpha = 0$ (likelihood training). Our findings show that generation models are advancing toward high *Certainty* and *Variety*, which is particularly true for the recently proposed UL loss method. Despite the highly restricted viable responses, i.e., *yes* or *no*, the *Certainty* obtained using UL with $\alpha > 0$ does not reach 1.0. Thus, we conclude that there is still room for improvement in n -best list generation

⁶For the BS, DBS, and UL, we obtained the 10-best lists setting beam size to 10. For the NS, we got the 10-best lists by performing nucleus sampling ten times.

in terms of avoiding contradiction.

5 Conclusion

Based on the recent development of contradiction detectors, removing contradictory candidates from models' n -best lists is a practical method for avoiding contradiction. In this method, the consistency of all candidates in the n -best lists substantially affects whether the final outputs are contradictory.

We quantitatively examined the properties of the n -best lists in terms of avoiding contradiction, using polar-typed questions as analytical inputs. We demonstrated that the proposed framework exhibits the properties of n -best lists based on *Certainty* and *Variety*. *Certainty* determines whether an n -best list has at least one noncontradictory response, whereas *Variety* evaluates how many noncontradictory responses each n -best list has. The results, particularly, demonstrated the present limitations on achieving high *Certainty* and *Variety* when using the well-established beam search method. In addition, our method emphasizes the improvements in *Certainty* and *Variety* achieved by recently proposed response generation strategies.

Our approach, which analyzes models' n -best lists based on *Certainty* and *Variety*, can be applied to any response generation problem, not just polar-typed response generation, which will be future work.

Acknowledgments

We would like to thank all anonymous reviewers for their insightful comments. We also thank Ana Brassard and Yosuke Kishinami for their valuable feedback and support. This work was partly supported by JSPS KAKENHI Grant Numbers JP21J22383, JP22K17943, JST Moonshot R&D Grant Number JPMJMS2011, and a Bilateral Joint Research Program between RIKEN AIP Center and Tohoku University.

References

- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. [Towards a human-like open-domain chatbot](#). In *arXiv preprint arXiv:2001.09977*.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Masahiro Araki, Hiroshi Tsukahara, Yuka Kobayashi, and Masahiro Mizukami. 2015. [Towards taxonomy of](#)

[errors in chat-oriented dialogue systems](#). In *Proceedings of the 16th annual meeting of the special interest group on discourse and dialogue (SIGDIAL)*, pages 87–95.

- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The Curious Case of Neural Text Degeneration](#). In *Proceedings of the eighth international conference on learning representations (ICLR)*.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.
- Hyunwoo Kim, Byeongchang Kim, and Gunhee Kim. 2020. [Will I sound like me? Improving persona consistency in dialogues through pragmatic self-consciousness](#). In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 904–916.
- Margaret Li, Stephen Roller, Ilia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020. [Don't say that! Making inconsistent dialogue unlikely with unlikelihood training](#). In *Proceedings of the 58th annual meeting of the association for computational linguistics (ACL)*, pages 4715–4728.
- Zekang Li, Jinchao Zhang, Zhengcong Fei, Yang Feng, and Jie Zhou. 2021. [Addressing Inquiries about History: An Efficient and Practical Framework for Evaluating Open-domain Chatbot Consistency](#). In *Findings of the joint conference of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (ACL-IJCNLP)*, pages 1057–1067.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). In *arXiv preprint arXiv:1907.11692*.
- Annie Louis, Dan Roth, and Filip Radlinski. 2020. ["I'd rather just go to bed": Understanding Indirect Answers](#). In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 7411–7425.
- Alexander H. Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. 2017. [ParlAI: A dialog research software platform](#). In *Proceedings of the 2017 conference on empirical methods in natural language processing (EMNLP): System demonstrations*, pages 79–84.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2020. [I like fish, especially dolphins: Addressing Contradictions in Dialogue Modeling](#). In *Proceedings of the 59th annual meeting of the association for computational linguistics (ACL)*, pages 1699–1713.

- Eun-Ju Noh. 1998. [Echo Questions: Metarepresentation and Pragmatic Enrichment](#). *Linguistics and Philosophy*, 21(6):603–628.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th conference of the european chapter of the association for computational linguistics: Main volume (EACL)*, pages 300–325.
- Iulian Vlad 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 AAAI conference on artificial intelligence (AAAI-16)*, pages 3776–3783.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 conference of the north american chapter of the association for computational linguistics: Human language technologies (NAACL-HLT)*, pages 196–205.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. [Diverse Beam Search for Improved Description of Complex Scenes](#). In *Proceedings of the 32nd AAAI conference on artificial intelligence (AAAI-18)*.
- Oriol Vinyals and Quoc V. Le. 2015. [A neural conversational model](#). In *Proceedings of the 31st international conference on machine learning (ICML) deep learning workshop*.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural Text Generation With Unlikelihood Training](#). In *Proceedings of the eighth international conference on learning representations (ICLR)*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 conference of the north american chapter of the association for computational linguistics: Human language technologies (NAACL-HLT)*, volume 1, pages 1112–1122.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP): System demonstrations*, pages 38–45.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th annual meeting of the association for computational linguistics (ACL): System demonstrations*, pages 270–278.

A Details of transforming NLI data

As described in Section 3.1, we obtain an analytical input from the NLI dataset. Specifically, we convert the hypothesis sentence of an NLI sample into a yes-no question. We describe the procedure as follows:

1. Detect the first verb of a sentence.
2. Move the verb to the beginning of the sentence, or put one of $\{Do, Does, Did\}$ at the front of the sentence, changing the verb back to its base (e.g., *made* \rightarrow *make*).
3. Change first-person pronouns to second-person pronouns and second-person pronouns to first-person pronouns (e.g., *my* \rightarrow *your*).
4. Change the punctuation mark at the end of the sentence to a question mark.

We used spaCy (`en_core_web_sm`) (Honnibal and Montani, 2017) to detect the verbs of hypothesis sentences. We did not use NLI samples with syntactically complex hypothesis sentences, such as those containing coordinating conjunctions, to avoid obtaining ungrammatical inputs. Further details are provided in our source codes.⁷

B Details of yes-no classifier

Training settings. On the Circa dataset, we fine-tuned the pretrained RoBERTa (`roberta-large`) implemented by Hugging Face (Wolf et al., 2020). We divided the dataset at random into `train:valid = 8 : 2`. The other training parameters were identical to those used by Louis et al. (2020).

Performance of classifier. To investigate the performance of the classifier, we measured the classification accuracy. First, we manually labeled the top-1 responses in the 10-best lists generated by the analysis presented in Section 4.2 with one of the two following labels: *Contradictory* or *Noncontradictory*. The accuracy with which the automated evaluation categorized the labeled responses was then evaluated. We selected 500 responses⁸ from 50 ENTQ/CNTQ inputs drawn at random from our test for the evaluation. The classifier classified 433/500 responses (see Appendix C), and the accuracy was 0.921. Some examples of the classification are shown in Table 4. The classifier correctly

⁷<https://github.com/shiki-sato/nbest-contradiction-analysis>

⁸100 responses generated by each of 5 generation models.

History: and we didn't ever call it uh Cokes and such you know we call it soda.
Message: Don't you always call it Coke?
Model Response: We call it coke.

Human Label: **Contradictory**
Predicted Label: **Contradictory**

(a) Example 1

History: The buying a house was the last thing that i wanted to do.
Message: Weren't you desperate to buy a house?
Model Response: No, I just wanted to buy a house.

Human Label: **Contradictory**
Predicted Label: **Non-contradictory**

(a) Example 2

Table 4: Examples of the response classification results by the yes-no classifier. The model responses were generated by Blender 400M using beam search with beam size $B = 10$.

Model	ENTQ	CNTQ
Blender 400M	1331 / 2000	1270 / 2000
Blender 1B	1413 / 2000	1316 / 2000
Blender 3B	1566 / 2000	1403 / 2000
DialoGPT 345M	1126 / 2000	924 / 2000
DialoGPT 762M	1044 / 2000	956 / 2000

Table 5: Number of stimulus inputs analyzed to calculate the *Certainty* and *Variety* described in Table 2.

Model	ENTQ	CNTQ
BS	1566 / 2000	1403 / 2000
DBS	991 / 2000	882 / 2000
NS	818 / 2000	684 / 2000
UL ($\alpha = 0$)	1914 / 2000	1871 / 2000
UL ($\alpha = 1$)	1806 / 2000	1887 / 2000
UL ($\alpha = 10$)	1654 / 2000	1811 / 2000

Table 6: Number of stimulus inputs analyzed to calculate the *Certainty* and *Variety* described in Table 3.

detected the contradiction in the model response using an indirect expression, in Example 1. However, in Example 2, the classifier failed to detect the contradiction of the model response, having both a noncontradictory direct expression (“No”) and a contradictory indirect expression (the part of the response after “No”). We found that the classifier tended to misclassify model responses containing the contradictions with themselves, such as Example 2.

C Details of experiments

Number of analyzed stimulus inputs. To simplify the analysis, we omitted from Section 4 and

Appendix B the analytical inputs with one or more ambiguous responses in the n -best lists. We defined ambiguous responses as those that were not identified by the classifier as either affirmations or refutations.⁹ Table 5 and Table 6 display the number of analytical inputs from the total of 2,000 ENTQ/CNTQ used for the two analyses in Section 4.

Generation model settings. In Section 4 experiments, we used DialoGPT (Zhang et al., 2020) and Blender (Roller et al., 2021) as response generation models. We used the codes of ParlAI (Miller et al., 2017) with its default settings, except for `beam_length_penalty=0` to generate responses.

Unlikelihood training settings. We used unlikelihood training with Blender 3B for the study of Section 4.3. To use the unlikelihood training proposed by Li et al. (2020), we require training data that includes the following three elements: input (here, history, and message), gold response, and negative response. These training samples were created by altering the NLI data with entailing and contradicting hypotheses.¹⁰ Table 7 displays the original NLI data and the transformed training samples. One NLI data set yields four types of questions (PositiveQ1, PositiveQ2, NegativeQ1, and NegativeQ2). We synthesized 8,000 samples from 2,000 NLI data and randomly divided them into `train : valid = 9 : 1`. We tuned the learning rate $\{7.0 \times 10^{-4}, 7.0 \times 10^{-5}, 7.0 \times 10^{-6}, 7.0 \times 10^{-7}, 7.0 \times 10^{-8}\}$ and the number of warmup updates $\{50, 100\}$ for each $\alpha = \{0, 1, 10\}$ for training. The rest of the training parameters are identical to those used by Roller et al. (2021). It is worth noting that we only trained the models marked as UL in Section 4.3 on these transformed data.

Premise: yeah i'm in North Carolina
Hypothesis – **Entailment**: I'm in North Carolina.
Hypothesis – **Contradict**: I'm in South Carolina.

(a) Original NLI data

PositiveQ1

History: Yeah I'm in North Carolina.
Message: Are you in North Carolina?
Gold: Yes, I'm in North Carolina.
Negative: No, I'm in South Carolina.

PositiveQ2

History: Yeah I'm in North Carolina.
Message: Are you in South Carolina?
Gold: No, I'm in North Carolina.
Negative: Yes, I'm in South Carolina.

NegativeQ1

History: Yeah I'm in North Carolina.
Message: Aren't you in North Carolina?
Gold: Yes, I'm in North Carolina.
Negative: No, I'm in South Carolina.

NegativeQ2

History: Yeah I'm in North Carolina.
Message: Aren't you in South Carolina?
Gold: No, I'm in North Carolina.
Negative: Yes, I'm in South Carolina.

(b) Training samples for UL

Table 7: Example of transforming (a) original NLI data to (b) training sample for UL. We synthesized four questions, i.e., PositiveQ1, PositiveQ2, NegativeQ1, and NegativeQ2, from each NLI sample.

⁹Circa dataset has seven different labels such as “Yes” and “Probably/sometimes yes.” We regard the responses classified into “In the middle” or “I am not sure” as ambiguous ones.

¹⁰Note that we did not use the identical NLI samples to synthesize ENTQ/CNTQ.