

Generate, Filter, and Rank: Grammaticality Classification for Production-Ready NLG Systems

Ashwini Challa* Kartikeya Upasani* Anusha Balakrishnan Rajen Subba

Facebook Conversational AI

{ashwinichalla, kart, anushabala, rasubba}@fb.com

Abstract

Neural approaches to Natural Language Generation (NLG) have been promising for goal-oriented dialogue. One of the challenges of productionizing these approaches, however, is the ability to control response quality, and ensure that generated responses are acceptable. We propose the use of a generate, filter, and rank framework, in which candidate responses are first filtered to eliminate unacceptable responses, and then ranked to select the best response. While acceptability includes grammatical correctness and semantic correctness, we focus only on grammaticality classification in this paper, and show that existing datasets for grammatical error correction don't correctly capture the distribution of errors that data-driven generators are likely to make. We release a grammatical classification and semantic correctness classification dataset for the weather domain that consists of responses generated by 3 data-driven NLG systems. We then explore two supervised learning approaches (CNNs and GBDTs) for classifying grammaticality. Our experiments show that grammaticality classification is very sensitive to the distribution of errors in the data, and that these distributions vary significantly with both the source of the response as well as the domain. We show that it's possible to achieve high precision with reasonable recall on our dataset.

1 Introduction

In recent years, neural network-based approaches have been increasingly promising in the context of goal-oriented Natural Language Generation (NLG). These approaches can effectively learn to generate responses of desired complexity and detail from unaligned data. Additionally, these approaches can be scaled with relatively low effort

to new domains and use cases. However, they are less robust to mistakes and have poor worst case performance. Consistently achieving acceptable response quality in a customer facing product is an immediate blocker to using such models widely.

Controlling quality at generation time in these models is challenging, and there are no guarantees that any of the generated responses are suitable to surface to an end user. Additionally, quality is hard to enforce at data collection time, given the increasingly widespread dependence on large pools of untrained annotators. As a result, classifying acceptability with high precision is extremely desirable. It can be used to establish safe fallbacks to acceptable, but potentially less ideal, responses that are generated by more traditional NLG systems like templates. Such responses are likely to be grammatically and semantically correct, but may sacrifice detail, variety, and naturalness; this trade-off may sometimes be necessary in a consumer-facing product. For example, the system could respond with "Here's your weather forecast", and show a card with relevant weather information, rather than generate an incoherent weather forecast.

Some key aspects of acceptability are **grammaticality** and **semantic correctness**. A grammatical response is one that is well-formed, and a semantically correct response is one that correctly expresses the information that needs to be conveyed. Systems that generate ungrammatical or incorrect responses run the risk of seeming unreliable or unintelligent. Another important facet of acceptability is the **naturalness** (or human likeness) of the response, that can improve the usability of chatbots and other dialogue systems.

In this paper, we first propose the inclusion of a filtering step that performs acceptability classification in the more widely used generate & rank framework (Generate, Filter, and Rank). Then,

*Equal contribution

we narrow our focus to grammaticality classification, and show how this problem calls for datasets of a different nature than typical grammatical error correction (GEC) datasets. We also show that state-of-the-art GEC models trained on general corpora fail to generalize to this problem. Finally, we introduce a dataset of system-generated grammatical errors for the **weather** domain, and demonstrate the performance of some strong baselines for grammatical classification on this data. This dataset can also be used for further research on semantic correctness classification. Our experiments also reinforce the need for the new framework we propose.

2 Generate, Filter, and Rank

In this section, we first review the pros and cons of the traditional generate & rank framework, and then propose a “filter” step that addresses some of its downsides.

The generate & rank framework has been proposed and widely used in several prior works on goal-oriented dialogue (Walker et al. (2001), Langkilde-Geary and Knight (1998)). In NLG systems, the typical use of this framework involves generating multiple candidate responses (often using various different surface realization techniques), and then reranking these using statistical models (most commonly language models). More recent works have also proposed reranking to optimize for certain personality traits or user engagement (Fang et al. (2017)). The input to the generators is usually a structured representation of what the system needs to convey.

This setup allows for the use of multiple generator models, as proposed in Serban et al. (2017) and Pichl et al. (2018), among others. This greatly increases the number of possible responses that can be surfaced, which can improve both diversity and naturalness. The use of statistical rerankers also allows systems under this framework to optimize for naturalness as well as acceptability (primarily grammaticality), since typical statistical models should easily be able to downrank potentially ungrammatical candidates. However, there are a few practical concerns that arise with using this framework in production:

1. Data sparsity: The space of unseen named entities like locations, datetimes, etc., and other sparse token types is potentially very large. This can result in suboptimal language

modeling behaviors, in which language models downrank valid candidates with sparse surface forms.

2. Statistical models that are typically used for reranking cannot capture semantic correctness without conditioning on the goal and arguments. They also run the risk of accidentally biasing towards more likely (but semantically incorrect) responses. This is particularly tricky for ML-based generators, where the generated responses can easily leave out important information. For example, the best models from Nayak et al. (2017) have error rates between 2-5%.
3. There is a significant risk that none of the responses generated by data-driven models is acceptable. For example, in the dataset that we release in this work, there were no grammatical responses generated for around 12% of the scenarios (see Section 4).

The common thread in these issues is that the generate & rank framework conflates acceptability, which is objective, with naturalness and other traits, which are subjective. To address, we propose the addition of a high-precision “filter” step that eliminates any unacceptable responses before the ranking stage, allowing the reranker to focus on optimizing for naturalness and other desirable properties. Since we found grammaticality to be a more serious issue than semantic correctness in our dataset (Table 2), we explore methods to implement a grammaticality “filter” in the following sections.

	CoNLL 2014	Our dataset
# grammatical	53426	18494
# ungrammatical	21638	14511
% scenarios with no grammatical responses	N/A	12%
Avg. length	22.8	17.9
Vocab size	28180	5669
# goals	N/A	2
# semantically correct	N/A	28475
# semantically incorrect	N/A	4530

Table 2: Comparison of weather responses dataset against the NUCLE corpus

Error Category	Examples
Repeated words like “with”, “and”.	In Grand Prairie , it ’s 100 degrees fahrenheit with cloudy skies with snow showers.
Agreement	Friday, September 15 in Branford , it’ll be cloudy with a high of 73 de-grees fahrenheit with an 61 percent chance of snow showers .
Dangling modifiers	In Tongan Qu on Monday, May 22 will be scattered clouds with Patches of Fog , with a high of 18 degrees celsius and a low of 7 degrees .
Incorrect word choice	In Larne on Thursday, November 23 , it’ll be scattered clouds with Fog , with a high of 46 and a low of 35 degrees fahrenheit.
Ungrammatical n-grams	In Funabashi-shi on Monday, March 20 , there will be a low of 31 with a high of 47 degrees fahrenheit with scattered clouds skies and a Light Drizzle
Missing contextual words, like “degrees”	In Caloocan City , expect a temperature of 3 celsius with mostly sunny skies and Fog Patches
Linking words/phrases	Right now in Arrondissement de Besancon , it ’s 2 degrees fahrenheit with sunny and Light Fog

Table 1: Mistakes involving grammatical errors and other cases of unacceptability in model-generated weather responses

Generator	Train		Eval		Test	
	# gr	# ungr	# gr	# ungr	# gr	# ungr
SC-LSTM Lex	4957	2386	1565	882	1712	757
SC-LSTM Delex	1083	2078	365	679	377	657
IR	1530	2513	532	839	493	833
Gen LSTM	3614	1624	1133	600	1247	549

Table 3: Distribution of positive and negative examples in weather responses dataset. # gr and # ungr denote number of grammatical and ungrammatical samples respectively.

3 Mismatched Error Distributions

The CoNLL-2014 shared task on grammatical error correction (Ng et al. (2014)) released the NUCLE corpus for grammatical error correction (GEC), written by students learning English. Ungrammatical sentences in this dataset contain annotations and corrections of each individual error. From a classification perspective, each original ungrammatical utterance in the dataset is a negative example, and the final corrected utterance (obtained by applying all of the corrections to the original ungrammatical utterance) is a positive example. Additionally, sentences without any corrections are positive examples as well.

These positive and negative samples can then be directly used to train the grammaticality filter described in previous sections. In the runtime of the goal-oriented NLG system, this filter would be used to filter out ungrammatical responses that are generated by **models** - even though the filter was trained on human-written responses. This signals

the possibility of a data mismatch.

To better understand the nature of this difference, we collected a corpus of system-generated responses for the **weather** domain (see Section 4) and manually inspected 200 of these responses to identify common categories of model mistakes (see Table 1). Interestingly, we found that the most common mistakes made by our models, like repeated words and missing contextual words, don’t match any of the error categories in NUCLE (see Table 1 from Ng et al. (2014)). There are also qualitative differences stemming from the domains in these datasets. Our corpus has a large number of mentions of sparse entities (particularly locations), dates, and weather-specific constructs like temperatures, while the NUCLE corpus is open-ended and spans a variety of topics. In order to quantify this difference, we measure the performance of open-domain GEC models on our corpus by evaluating a model that achieves state-of-the-art performance on the CoNLL-2014 test set (Chollampatt and Ng, 2018). We found

that this model failed to generalize well to our dataset (see section 6), and missed several classes of errors. For example, the model failed to catch any of the errors in Table 1 (see Appendix A for more examples).

Intuitively, this suggests that training models for response filtering demands datasets very different in distribution from publicly available datasets that only reflect human mistakes. We show this empirically through experiments in section 6, and describe the process for collecting our dataset in the next section.

4 Dataset

We first collected a dataset of human-generated responses for the weather domain, using a process similar to the one used in Novikova et al. (2017). Each of the collected responses is conditioned on a **scenario**, consisting of a goal (the intent to be expressed) and arguments (information to be expressed). In collecting the dataset, we restricted ourselves to the goals `inform_current_condition` and `inform_forecast`.

An example scenario is

“requested_location”: “London”,

“temp”: “32”,

“temp_scale”: “fahrenheit”,

“precip_summary”: “Heavy Blowing Snow”

A possible response for this scenario is In London, it’s currently 32 degrees Fahrenheit with heavy snow..

We then trained some standard NLG models on this corpus. Two of these (`sc-LSTM Lex` and `sc-LSTM Delex`) are semantically conditioned LSTMs as described in Wen et al. (2015); the `genLSTM` model is a vanilla LSTM decoder; and `IR` is a simple retrieval-based generator. The details of these are described in Appendix A.1. We generated $n = 3$ responses from each of these models for each scenario in a held out data set, and deduped generated candidates that differed by a single character (often punctuation). We then asked crowdworkers to judge the grammaticality of these responses. Our final dataset¹ consists of 33K model-generated responses with grammaticality and semantic correctness judgments. Table 3 shows a detailed breakdown of grammatical and ungrammatical responses per model.

¹github.com/facebookresearch/momi

5 Approach

Preprocessing We made the assumption that the specific values of arguments such as locations, dates, and numbers do not affect sentence framing. We therefore replaced locations and dates with placeholder tokens. Numbers are replaced with either `_num_`, `_num_vowel_` if the number begins with a vowel sound (example, 80), or `_num_one_` if the number is 1. Hence the sentence “There is an 85 percent chance of rain in New York on Wednesday, August 25” would become “There is an `_num_vowel_` percent chance of rain in `_location_` on `_date_`”. In case of `sc-LSTM delex`, all remaining arguments (such as weather conditions) are also delexicalized.

To maintain class balance in the train set, for each response source, the class with fewer samples is upsampled to match the number of samples of the other class. When training on samples from multiple generators, the samples of each generator in train set are upsampled to match those of generator with highest number of samples. Upsampling is not done for validation or test sets.

Features from Language Model	
Geometric mean: $(\prod_{i=1}^m p_i)^{(1/m)}$	Arithmetic mean: $\sum_{i=1}^m p_i / m$
$\min P_x$	$\max P_x$
Median: \tilde{p}	Std Dev: σ_{P_x}
$C_{P_x}(0, 0.1), C_{P_x}(0.1, 0.2), \dots, C_{P_x}(0.9, 1.0)$	

Table 4: Features derived from Language Model. $P_x = p_1, p_2, \dots, p_m$ is the set of all n-gram probabilities from an n-gram LM for a sentence x . $C_{P_x}(a, b) \in [0, 1]$ is the ratio of n-gram probabilities $p_i \in P_x$ for which $a \leq p_i < b$.

Gradient Boosted Decision Tree Using LM Features (LM-GBDT) Language models (Brown et al. (1992)) can effectively capture n-gram patterns that occur in grammatical data, making features derived from them good candidates to distinguish between grammatical and ungrammatical responses. We train a 7-gram LM² on human-written weather responses described in Section 4. The trained LM is then used to extract features listed in Table 4 for each model-generated response. Finally, we feed these features into a

²We found that 7 gram LM performed slightly better than other lower n-gram LMs. LM with larger n-grams may be better at catching model mistakes that require looking at long-range dependencies.

Model	Training Data	Test Data	R@P98	R@P
Chollampatt and Ng (2018)	NUCLE	Weather	-	75 @ 64
CNN	NUCLE	NUCLE	62.4	-
		Weather	-	80 @ 56.8
	NUCLE + Weather	Weather	52.5	-
CNN	Weather	Weather	71.9	-
CNN + source	Weather	Weather	72.8	-
LM-GBDT	Weather	Weather	63.8	-

Table 5: Training on NUCLE and weather data

gradient boosted decision tree (GBDT) (Friedman (2001)) to classify the model-generated response as grammatical or not.

CNN-based Classification Model We used a convolutional neural network (CNN) for sentence classification in an approach similar to Kim (2014). After pooling convolutional features along the time dimension, the result can be optionally concatenated with additional features. A one-hot vector of length 4 encoding the source of the response (IR, GenLSTM, sc-LSTM delex, sc-LSTM lex) is passed as an additional feature when training on responses from multiple sources.

6 Experiments

We try different combinations of NUCLE corpus and our dataset as train and test sets to learn a grammaticality classifier for model-generated weather responses. Table 5 and 7 lists the results of these experiments described above. As discussed before, since the goal is to build a classifier for use in production systems, we report the recall of models for grammatical class when the precision is very high (98%). In cases where the model does not achieve this precision, we report recall at the highest precision achieved by the model.

CNN + source represents the case when the source of response is passed as an additional feature to the CNN. We used filters with widths 2, 3, 4, 5 for the CNN. Performance did not change with different number and sizes of filters.

	Ungrammatical	Semantically incorrect
Ranker	29.4%	8.2%
Filter + Ranker	2.4%	0.75%

Table 6: Comparison of number of times the top ranked response is unacceptable with and without filtering.

6.1 Ranker vs Filter + Ranker

In order to validate the Generate, Filter, and Rank framework, we used our trained n-gram language model³ (from Section 4) to rank all the responses for each scenario in our dataset. We then measured the % of times the top ranked candidate is ungrammatical, to understand how many times the final response would be ungrammatical in a traditional generate & rank framework. We repeat the experiment with our proposed framework, by filtering ungrammatical responses using a CNN-based filter with 98% precision before the ranking step. The results are shown in Table 6.

The filtering step increases the overall response quality, but comes at the cost of losing genuine grammatical candidates because of slightly lower recall, 72.8%, (the best recall we achieved on the weather data set). This is a fundamental trade-off of our proposed framework; we sacrifice recall for the sake of precision, in order to ensure that users of the system very rarely see an unacceptable response. The semantic correctness also improves, but this doesn't indicate that grammatical filter is enough to solve both grammaticality and correctness problems.

6.2 Performance of filters on NUCLE and weather data

Table 5 compares performance of CNN, LM-GBDT, and the GEC model used by Chollampatt and Ng (2018). The GEC model is adopted for binary classification by checking whether the model makes a correction for an ungrammatical sentence, and doesn't make any corrections for a grammati-

³an n-gram based language model is a simple baseline. It is possible to use more sophisticated rankers (such as RNN-LMs) to achieve better results. However, ranking approaches will still fail over filters when there are no grammatical candidates at all.

Model	Training Data	Test Data	R@P98	R@P
CNN + source	Weather	IR	9.8	-
	IR	IR	23.2	-
	Weather	GenLSTM	95.5	-
	GenLSTM	GenLSTM	92.2	-
	Weather	SC-LSTM Delex	25.2	-
	SC-LSTM Delex	SC-LSTM Delex	-	45.9@80
	Weather	SC-LSTM Lex	96.8	-
	SC-LSTM Lex	SC-LSTM Lex	94.6	-
LM-GBDT	Weather	IR	-	8@95.5
	IR	IR	18	-
	Weather	GenLSTM	83.4	-
	GenLSTM	GenLSTM	76	-
	Weather	SC-LSTM Delex	2	-
	SC-LSTM Delex	SC-LSTM Delex	-	65.5@70.5
	Weather	SC-LSTM Lex	90.6	-
	SC-LSTM Lex	SC-LSTM Lex	88.4	-

Table 7: Performance of filter for individual generators

cal sentence⁴. This model achieves poor precision and recall on our dataset, and we found that it fails to generalize adequately to the novel error types in our data.

We also train the CNN on NUCLE data and find that it similarly achieves poor recall when classifying weather responses. This is attributed to the fact that the domain and category of errors in both datasets are different. Comparing Table 1 in Ng et al. (2014) and Table 1 of this work further supports this observation.

The CNN and LM-GBDT are trained and tested on our weather dataset. We report the performance of these models on the **complete** weather test set, not just on individual generators, since this is closest to the setting in which such models would be used in a production system. The CNN consistently has better recall than LM-GBDT at the same precision. CNN + source performs better than the CNN, indicating that information regarding source helps in classifying responses from multiple generators.

Augmenting the weather responses with NUCLE corpus while training the CNN did not help performance.

⁴We assume that the GEC model has classified the response as ungrammatical if an edit is made. This does not account for cases in which the edited response is still ungrammatical. As a result, the precision of this model in the true setting would be lower than that reported in this setting.

6.3 Performance of filter for individual generators

Table 7 presents results on test sets of each generator for classifiers trained together on all generators and trained on individual generators. Models trained individually on IR and SC-LSTM Delex responses perform poorly compared to GenLSTM and SC-LSTM Lex as the training set size is much smaller for former. The recall for individual generators is higher when training is done on data from all generators, indicating that the approach generalizes across sources. An exception to this is IR which does better when trained just on IR responses. This may be due to errors of retrieval based approach being different in nature compared to LSTM-based approach.

Tables 8 and 9 in Appendix shows the errors in responses from different generators. Some errors occur more frequently with one generator than another, for example, the problem of repeating words (like *with* and *and*) is dominant in responses generated by the LSTMs, but very rarely seen in IR since it is a retrieval based approach.

6.4 Comparison of LM-GBDT and CNN

The recall of CNN is slightly better than LM-GBDT consistently across experiments. Both approaches do well in catching types of errors listed in Table 1. One difference between the two is the ability of CNN-based models to successfully catch errors such as “1 degrees”, while the LM-

GBDT fails to do so. On further inspection, we noticed that the human generated weather responses, which were used as training data for the language model, contained several instances of “1 degrees”. The LM-GBDT has a heavy dependency on the quality of features generated by LM (which in turn depends on the quality of the LM training corpus), and this is a disadvantage compared to the CNNs.

7 Related Work & Conclusion

Several previous works have established the need for a generate & rank framework in a goal-oriented NLG system (Walker et al. (2001), Langkilde-Geary and Knight (1998)). Recent work on the Alexa prize (Ram et al. (2018)) has demonstrated that this architecture is beneficial for systems that bridge the gap between task-oriented and open-ended dialogue (Serban et al. (2017), Fang et al. (2017), Pichl et al. (2018)). In such systems, the ranker needs to choose between a much more diverse set of candidates, and potentially optimize for other objectives like personality or user satisfaction. To make such systems practical for production-scale usage, our work proposes the inclusion of a high precision filter step that precedes ranking and can mark responses as acceptable. Our experiments show that this filter with sufficient fallbacks guarantees response quality with high precision, while simply reranking does not (Section 6.1).

In this work, we focus specifically on filtering ungrammatical responses. Previous work in this space has focused on classifying (and sometimes correcting) errors made by humans (Ng et al. (2014)) or synthetically induced errors (Foster (2007)). We found, however, that the domain and error distribution in such datasets is significantly different from that of typical data-driven generation techniques. To address this gap, we release grammatical and semantic correctness classification data generated by these models, and present a reasonable baseline for grammatical classification. The approaches we present are similar to work on grammatical classification using features from generative models of language, like language models (Wagner et al. (2009)). One future direction is to explore modeling semantic correctness classification with the dataset we release.

We compare the performance of two approaches for classifying grammaticality: CNNs, and GBDTs with language model features. Both are stan-

dard classifiers that are easy to deploy in production systems with low latency. An interesting future direction would be to explore model architectures that scale better to new domains and generation approaches. This could include models that take advantage of existing GEC data consisting of human responses, as well as datasets similar to ours for other domains. Models that successfully make use of these datasets may have a more holistic understanding of grammar and thus be domain- and generator-agnostic.

A drawback of the generate-filter-rank framework is the increased reliance on a fallback response in case no candidate clears the filtering stage. This is an acceptable trade-off when the goal is to serve responses in production systems where the standards of acceptability are high. One way to alleviate this is to do grammatical error correction instead of simply removing unacceptable candidates from the pipeline. Correcting errors instead of rejecting candidates can be of value for trivial mistakes such as missing articles or punctuation. However, doing this with high precision and correcting semantic errors remains a challenge.

References

- Peter F Brown, Peter V Desouza, Robert L Mercer, Vincent J Della Pietra, and Jenifer C Lai. 1992. Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–479.
- Shamil Chollampatt and Hwee Tou Ng. 2018. A multilayer convolutional encoder-decoder neural network for grammatical error correction. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hao Fang, Hao Cheng, Elizabeth Clark, Ariel Holtzman, Maarten Sap, Mari Ostendorf, Yejin Choi, and Noah A Smith. 2017. Sounding board–university of washingtons alexa prize submission. *Alexa Prize Proceedings*.
- Jennifer Foster. 2007. Treebanks gone bad. *International Journal of Document Analysis and Recognition (IJ DAR)*, 10(3-4):129–145.
- Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751.

Irene Langkilde-Geary and Kevin Knight. 1998. Generation that exploits corpus-based statistical knowledge. In *COLING-ACL*.

Neha Nayak, Dilek Hakkani-Tur, Marilyn Walker, and Larry Heck. 2017. To plan or not to plan? discourse planning in slot-value informed sequence to sequence models for language generation. In *INTERSPEECH*.

Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. The conll-2014 shared task on grammatical error correction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*.

Jekaterina Novikova, Ondej Due, and Verena Rieser. 2017. The e2e dataset: New challenges for end-to-end generation. In *Proceedings of the SIGDIAL 2017 Conference*, pages 201–206.

Jan Pichl, Petr Marek, Jakub Konrád, Martin Matulík, Hoang Long Nguyen, and Jan Sedivý. 2018. [Alquist: The alexa prize socialbot](#). *CoRR*, abs/1804.06705.

Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, et al. 2018. Conversational ai: The science behind the alexa prize. *arXiv preprint arXiv:1801.03604*.

Iulian Vlad Serban, Chinnadhurai Sankar, Mathieu Germain, Saizheng Zhang, Zhouhan Lin, Sandeep Subramanian, Taesup Kim, Michael Pieper, Sarath Chandar, Nan Rosemary Ke, Sai Mudumba, Alexandre de Brébisson, Jose Sotelo, Dendi Suhubdy, Vincent Michalski, Alexandre Nguyen, Joelle Pineau, and Yoshua Bengio. 2017. [A deep reinforcement learning chatbot](#). *CoRR*, abs/1709.02349.

Joachim Wagner, Jennifer Foster, and Josef van Genabith. 2009. Judging grammaticality: Experiments in sentence classification. In *CALICO Journal*, pages 474–490.

Marilyn A. Walker, Owen Rambow, and Monica Rogati. 2001. Spot: A trainable sentence planner. In *NAACL*.

Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *EMNLP*.

A Appendices

A.1 NLG Models for Generating Weather Responses

The dataset we present in this paper consists of responses generated by 4 model types:

1. `sc-LSTM delex`: An `sc-LSTM` trained on fully delexicalized human responses, where

delexicalization refers to the process of replacing spans corresponding to specific arguments by placeholder strings.

2. `sc-LSTM lex`: An `sc-LSTM` trained on partly delexicalized human responses. For this model, we only delexicalize locations, dates, and temperatures, thus allowing the model to freely choose surface forms for any other arguments.
3. `GenLSTM`: A vanilla LSTM-based decoder model, where the decoder hidden state is initialized using embeddings of the goal and arguments. This model is also trained on fully delexicalized responses.
4. `IR`: A simple retrieval approach in which n random candidates that satisfy the given goal and arguments are retrieved. The retrieved candidates are delexicalized, and any candidates that contain the right arguments (regardless of argument value) are considered valid.

For all models, the final response is obtained by replacing argument placeholders by the canonical values of those arguments in the scenario.

Since our goal was just to get responses from a diverse set of data-driven generators with a reasonable distribution of errors, we did not experiment too much with improving `IR` and `genLSTM`, which are much weaker than the `sc-LSTM` models.

A.2 Model-Generated Responses: Error Analysis

Tables 8 and 9 show errors made by different generators. While there is an overlap in the category of grammatical errors made by different generators, the frequency of the errors is largely different. There are also a few generator-specific errors. For example, the problem of repeating words (like *with* and *and*) is dominant in responses generated by `GenLSTM`, `sc-LSTM delex`, `sc-LSTM lex`, but very rarely seen in `IR`. This is because human responses themselves are unlikely to have repeating words, however the LSTM-based generators tend to make these mistakes while trying to fit all information into the response. Ungrammatical n-grams like *scattered clouds skies* are very infrequent in `sc-LSTM lex` responses while more commonly seen with

other generators. This is because the `sc-LSTM` `lex` generator directly produces surface forms of weather conditions. LSTM models doesn't tend to generate responses with out of vocabulary words, but it is something common with IR responses usually because of spelling mistakes in templates.

A.3 General GEC Model Performance

Table 10 shows examples of ungrammatical responses that the general GEC model ((Chollampatt and Ng, 2018)) failed to correct, and Table 11 shows examples of ungrammatical responses that the model correctly edited. The model corrects mistakes that are much more likely to occur in the GEC data (like verb agreement), but fails to catch model-specific error types like stuttering and other ungrammatical n-grams.

Generator	Error Category	Examples
SC-LSTM Lex	Repeating words like “with”, “and”.	In Grand Prairie , it ’s 100 degrees fahrenheit with cloudy skies with snow showers.
	Poor choice of words to connect 2 phrases	Right now in Medford , with a temperature of -10 degrees celsius .
	Wrong Plurals/singulars	In Yushu , it’s 1 degrees celsius and cloudy .
	Missing words that forms incomplete sentences	In Tongan Qu on Monday, May 22 will be scattered clouds with Patches of Fog , with a high of 18 degrees celsius and a low of 7 degrees May 22 there will be scattered clouds ... Right now in East Liverpool it is -3 fahrenheit with Heavy Rain. *missing word degrees*
	Wrong articles	Friday, September 15 in Branford , it’ll be cloudy with a high of 73 degrees fahrenheit with an 61 percent chance of snow showers .
	Wrong pronouns	In Larnie on Thursday, November 23 , it’ll be scattered clouds with Fog , with a high of 46 and a low of 35 degrees fahrenheit. *there’ll be scattered clouds*
SC-LSTM Delex	Repeating words like “with”, “and”	In Chengtangcun on Wednesday, April 12 expect a high of 2 degrees and a low of -10 degrees celsius with cloudy skies with Snow Showers .
	Wrong word choices	In Funabashi-shi on Monday, March 20 , there will be a low of 31 with a high of 47 degrees fahrenheit with scattered clouds skies and a Light Drizzle
	Wrong articles	In Newbury on Tuesday, February 07 , there will be an 46 percent chance of Heavy Rain Showers with a high of 5 degrees celsius with overcast skies .
	Wrong Plurals/Singulars	In Shiselweni District on Tuesday, March 21 , it will be overcast with a high of 8 degrees celsius and a low of 1 degrees .
Missing contextual words like “degrees”	In Caloocan City , expect a temperature of 3 celsius with mostly sunny skies and Fog Patches	

Table 8: Some more examples of grammatical errors made by different generation models in our dataset.

Generator	Error Category	Examples
Gen LSTM	Repeating words like “with”, “and”	Right now in Wojewodztwo Maopolskie , it ’s sunny with Light Thunderstorms with Hail and a temperature of 13 degrees fahrenheit .
	Poor word choices	Right now in Franklin Square , it ’s 96 degrees fahrenheit with scattered clouds skies .
	Wrong articles	On Friday, November 17 in San-Pedro , expect a low of 44 and a high of 68 degrees fahrenheit with an 41 percent chance of Flurries .
	Wrong Plurals/Singulars	Right now in Minnetonka Mills , it ’s 1 degrees celsius with sunny skies .
	Wrong pronouns	On Monday, July 03 in Himeji Shi , it’ll be scattered clouds with a high of 48 degrees fahrenheit. *there’ll be scattered clouds*
	Wrong connecting words	Right now in Arrondissement de Besancon , it ’s 2 degrees fahrenheit with sunny and Light Fog *... and sunny with light fog ... would make it grammatical*
IR	Wrong articles	In Shiraki Marz on Thursday, November 09, there will be an 51 percent chance of Heavy Blowing Snow and a high of 39 degrees fahrenheit
	Wrong ordinal indicators	On Friday, June 02th in Selma there will be a low of 82 degrees fahrenheit with Light Thunderstorms with Hail
	Wrong Plurals/Singulars	On Tuesday, June 13, in Wilayat-e Paktiya, there will be Heavy Snow Showers and the high will be 1 degrees celsius.
	Wrong helping verbs (Plural versus singular)	N/A
	Wrong Pronoun	On Wednesday, October 18, in Reus, it’ll be scattered clouds and 3 degrees celsius. *... there’ll be scattered clouds ...*
	Poor templates like one with repeating words, spelling mistakes, missing words like degrees	In Rudraprayag on Tuesday, November Tuesday, June 13 temp is -8 to 0 celsius with Low Drifting Snow Snow Showers and overcast cloud
	Out of vocabulary words	It’s currently -15 degrees fahrenheit t and mostly clear with gentle breeze in Dammam

Table 9: Some more examples of grammatical errors made by different generation models in our dataset.

Response	Error type
On Friday, February 17 in Changwat Samut Songkhram, expect a likely of heavy rain showers and a high of 15 degrees celsius.	Agreement
Currently in Maastricht there is fog and is -3 Fahrenheit.	Missing “degrees”
Right now in Westminster it is 1 degrees Fahrenheit with partly cloudy skies.	Numerical agreement
In Ayacucho on Wednesday, February 22, with a high of 83 degrees fahrenheit with a 98 percent chance of light snow.	Repeated “with”, and missing linking words
In kajiado on Thursday, January 12, expect a high of 82 degrees and a low of 61 degrees Fahrenheit with mostly sunny .	Incomplete response

Table 10: Mistakes involving grammatical errors and other cases of unacceptability in model-generated weather responses

Original Response	Corrected Response
The weather for Wednesday, December 27 in Oak Hill will includes a high of 14 Celsius and a 37 percent chance of heavy freezing rain.	The weather for Wednesday, December 27 in Oak Hill will include a high of 14 Celsius and a 37 percent chance of heavy freezing rain.
In Ocean County, it is 34 degrees Fahrenheit with sunny.	In Ocean County, it is 34 degrees Fahrenheit with sunny weather .
In Bim Son, it is 1 degrees fahrenheit with funnel cloud.	In Bim Son, it is 1 degrees fahrenheit with funnel clouds .

Table 11: Mistakes involving grammatical errors and other cases of unacceptability in model-generated weather responses