

Characterizing Variation in Crowd-Sourced Data for Training Neural Language Generators to Produce Stylistically Varied Outputs

Juraj Juraska and Marilyn Walker
Natural Language and Dialogue Systems Lab
University of California, Santa Cruz
{jjuraska,mawalker}@ucsc.edu

Abstract

One of the biggest challenges of end-to-end language generation from meaning representations in dialogue systems is making the outputs more natural and varied. Here we take a large corpus of 50K crowd-sourced utterances in the restaurant domain and develop text analysis methods that systematically characterize types of sentences in the training data. We then automatically label the training data to allow us to conduct two kinds of experiments with a neural generator. First, we test the effect of training the system with different stylistic partitions and quantify the effect of smaller, but more stylistically controlled training data. Second, we propose a method of labeling the style variants during training, and show that we can modify the style of the generated utterances using our stylistic labels. We contrast and compare these methods that can be used with any existing large corpus, showing how they vary in terms of semantic quality and stylistic control.

1 Introduction

Dialogue systems have become one of the key applications in natural language processing, but there are still many ways in which these systems can be improved. One obvious possible improvement is in the system’s language generation to make it more natural and more varied. Both a benefit and a challenge of neural natural language generation (NLG) models is that they are very good at reducing noise in the training data. When they are trained on a sufficiently large dataset, they learn to generalize and become capable of applying the acquired knowledge to unseen inputs. The more

data the models are trained on, the more robust they become, which minimizes the effect of noise in the data on their learning. However, the higher amount of training data can also drown out interesting stylistic features and variations that may not be very frequent in the data. In other words, the model, being statistical, will prefer producing the most common sentence structures, i.e. those which it observed most frequently in the training data and is thus most confident about.

In our work, we consider language generators whose inputs are structured *meaning representations* (MRs) describing a list of key concepts to be conveyed to the human user during the dialogue. Each piece of information is represented by a slot-value pair, where the *slot* identifies the type of information and the *value* is the corresponding content. A language generator must produce a syntactically and semantically correct utterance from a given MR. The *utterance* should express all the information contained in the MR, in a natural and conversational way. Table 1 shows an example MR for a restaurant called “The Waterman” paired with two (out of many) possible output utterances, the first of which might be considered stylistically interesting, since the name of the restaurant follows some aspects of the description and contains a concession, while the second example might be considered as more stylistically conventional.

Recently, the size of training corpora for NLG has become larger, and these same corpora have begun to manifest interesting stylistic variations. Here we start from the recently released E2E dataset (Novikova et al., 2017b) with nearly 50K samples of crowd-sourced utterances in the restaurant domain provided as part of the E2E NLG Challenge.¹ We first develop text analysis methods that systematically characterize types of sen-

¹<http://www.macs.hw.ac.uk/InteractionLab/E2E/>

MR	name [The Waterman], food [English], priceRange [cheap], customer rating [low], area [city centre], familyFriendly [yes]
Utt. #1	There is a cheap, family-friendly restaurant in the city centre , called The Waterman . It serves English food, but received a low rating by customers.
Utt. #2	The Waterman is a family-friendly restaurant in the city centre . It serves English food at a cheap price. It has a low customer rating.

Table 1: Example of a meaning representation and two corresponding utterances of different styles.

tences in the training data. We then automatically label the training data – with the help of a heuristic slot aligner and a handful of domain-independent rules for discourse marker extraction – in order to allow us to conduct two kinds of experiments with a neural language generator: (1) we test the effect of training the system with different stylistic partitions and quantify the effect of smaller, but more stylistically controlled training data; (2) we propose a method of labeling the style variants during training, and show that we can modify the style of the output using our stylistic labels. We contrast these methods, showing how they vary in terms of semantic quality and stylistic control. These methods promise to be usable with any sufficiently large corpus as a simple way of producing stylistic variation.

2 Related Work

The restaurant domain has always been the domain of choice for NLG tasks in dialogue systems (Stent et al., 2004; Gašić et al., 2008; Mairesse et al., 2010; Howcroft et al., 2013), as it offers a good combination of structured information availability, expression complexity, and ease of incorporation into conversation. Hence, even the more recent neural models for NLG continue to be tested primarily on data in this domain (Wen et al., 2015; Dušek and Jurčiček, 2016; Nayak et al., 2017). These tend to focus solely on syntactic and semantic correctness of the generated utterances, nevertheless, there have also been re-

cent efforts to collect training data for NLG with emphasis on stylistic variation (Nayak et al., 2017; Novikova et al., 2017a; Oraby et al., 2017).

While there is previous work on stylistic variation in NLG (Paiva and Evans, 2004; Mairesse and Walker, 2007), this work did not use crowd-sourced utterances for training. More recent work in neural NLG that explores stylistic control has not needed to control semantic correctness, or examined the interaction between semantic correctness and stylistic variation (Sennrich et al., 2016; Ficler and Goldberg, 2017). Also related is the work of Niu and Carpuat (2017) that analyzes how dense word embeddings capture style variations, Kabbara and Cheung (2016) who explore the ability of neural NLG systems to transfer style without the need for parallel corpora, which are difficult to collect (Rao and Tetreault, 2018), while Li et al. (2018) use a simple delete-and-retrieve method also without alignment to outperform adversarial methods in style transfer. Finally, Oraby et al. (2018) propose two different methods that give neural generators control over the language style, corresponding to the Big Five personalities, while maintaining semantic fidelity of the generated utterances.

To our knowledge, there is no previous work exploring the use of and utility of stylistic selection for controlling stylistic variation in NLG from structured MRs. This may be either because there have not been sufficiently large corpora in a particular domain, or because it is surprising, as we show, that relatively small corpora (2000 samples) whose style is controlled can be used to train a neural generator to achieve high semantic correctness while producing stylistic variation.

3 Dataset

We perform the stylistic selection on the E2E dataset (Novikova et al., 2017b). It is by far the largest dataset available for task-oriented language generation in the restaurant domain. It offers almost 10 times more data than the San Francisco restaurant dataset (Wen et al., 2015), which had frequently been used for NLG benchmarks. This significant increase in size allows successful training of neural models on smaller subsets of the dataset. Careful selection of the training subset can be used to influence the style of the utterances produced by the model, as we show in this paper.

A portion of the human reference utterances

	Samples	Unique MRs
Training	42,061	4,862
Validation	4,672	547
Test	630	630
Total	47,363	6,039

Table 2: Number of samples vs. unique meaning representations in the training, validation and test set of the E2E dataset.

Slots	3	4	5	6	7	8
Sentences	1.09	1.23	1.41	1.65	1.84	1.92
Proportion	5%	18%	32%	28%	14%	3%

Table 3: Average number of sentences in the reference utterance for a given number of slots in the corresponding MR, along with the proportion of MRs with specific slot counts.

was collected using pictures as the source of information, which was shown to inspire more natural utterances compared to textual MRs (Novikova et al., 2016). The reference utterances in the E2E dataset exhibit superior lexical richness and syntactic variation, including more complex discourse phenomena. It aims to provide higher-quality training data for end-to-end NLG systems to learn to produce better phrased and more naturally sounding utterances.

Although the E2E dataset contains a large number of samples, each MR is associated on average with more than 8 different reference utterances, effectively supplying almost 5K unique MRs in the training set (Table 2). It thus offers multiple alternative ways of expressing the same information in an utterance, which the model can learn. We take advantage of this aspect of the dataset when selecting the subset of samples for training with a particular purpose of stylistic variation.

The dataset contains 8 different slot types, which are fairly equally distributed in the dataset. Each MR comprises 3 to 8 slots, whereas the majority of MRs consist of 5 and 6 slots. Even though most of the MRs contain many slots, the majority of the corresponding human utterances consist of one or two sentences only (Table 3), suggesting a reasonably high level of sentence complexity in the references.

Domain	Utterance
TV	You might like the Dionysus 44 television that has an a+ eco rating and 720p resolution, <i>while</i> only using 32 watts in power consumption. (Wen et al., 2016)
Laptop	<i>For the price of 449 dollars</i> , you could purchase the Satellite Hypnos 38 laptop. (Wen et al., 2016)
People	<i>Born in the London Borough of Havering</i> , Alex Day started performing in 2006. (Gardent et al., 2017)
Food	Sago is the main ingredient in binignit, <i>but</i> sweet potatoes are also used in it. (Gardent et al., 2017)

Table 4: Examples of utterances in different datasets/domains, also exhibiting interesting discourse phenomena.

4 Stylistic Selection

We note that the E2E dataset is significantly larger than what is needed for a neural model to learn to produce correct utterances in this domain. Thus, we seek a way to help the model learn more than just to be correct. We strive to achieve higher stylistic diversity of the utterances generated by the model through stylistic selection of the training samples. We start by characterizing variation in the crowd-sourced dataset and detect what opportunities it offers for the model to learn more advanced sentence structures. Table 5 illustrates some of the stylistic variation that we observe, which we describe in more detail below. We then judge the level of desirability of specific discourse phenomena in our context, and devise rules based on the parse tree to extract the samples that manifest those stylistic phenomena. This gives us the ability to create subsets of the samples with an arbitrary combination of stylistic features that we are interested in. We then explore the extent to which we can make the model’s output demonstrate these stylistic features.

4.1 Stylistic Variation in the Dataset

This section gives an overview of different discourse phenomena in the E2E dataset that we consider relevant in the context of a task-oriented dialogue in the restaurant domain. The majority of

Category	Utterance
Aggregation	Located in the city centre is a family-friendly coffee shop called Fitzbillies. It is both inexpensive and highly rated .
Contrast	The Rice Boat is a Chinese restaurant in the riverside area. It has a customer rating of 5 out of 5 but is not family friendly.
Fronting	With a 1 out of 5 rating Midsummer House serves Italian cuisine in the high price range, found not far from All Bar One.
Subordination	Wildwood pub is serving 5 star food while keeping their prices low .
Exist. clause	In the city center, there is an average priced, non-family-friendly, Japanese restaurant called Alimentum.
Imperative/modal	In Riverside, you’ll find Fitzbillies. It is a passable, affordable coffee shop which interestingly serves Chinese food. Don’t bring your family though.

Table 5: Examples of the categories of discourse phenomena extracted from the utterances in the E2E dataset.

these would, however, generalize to other domains too, since they appear not only in summaries of restaurants, but, for example, in those of TVs, laptops (Wen et al., 2016), people and food (Gardent et al., 2017) too (see examples in Table 4). The extraction rules we have implemented can thus be widely used in task-oriented data-to-text language generators. We split the sentence features in the following six categories. An example of each is given in Table 5:

- **Aggregation:** Discourse phenomena grouping information together in a more concise way. This includes specifiers such as “both” or “also”, as well as apposition and gerunds. Another type of aggregation uses the same quantitative adjective for characterizing multiple different qualities (such as “It has a *low customer rating and price range*.”).

Note that some of the following categories contain other markers that also represent aggregation.

- **Contrast:** Connectors and adverbs expressing concession or contrast between two or more qualities, such as “but”, “despite”, “however”, or “yet”.
- **Fronting:** Fronted adjective, verb and prepositional phrases, typically highlighting qualities of the eatery before its name is given.

In this category we also include specificational copular constructions, which are for-

mulations with inverted predication around a copula, bringing a particular quality of the eatery in the front (e.g. “*A family friendly option is The Rice Boat*.”).

- **Subordination:** Clauses introduced by a subordinating conjunction (such as “if” or “while”), or by a relative pronoun (such as “whose” or “that”).
- **Existential clause:** Sentences formulated using the expletive “there”.
- **Imperative and modal verb:** Sentences involving a verb in the imperative form or a modal verb, making the utterance sound more personal and interactive.

4.2 Discourse Marker Weighting

Many human-produced utterances, naturally, contain multiple of the discourse phenomena described in Section 4.1. Such utterances are preferred to those only containing a single discourse phenomenon of interest, especially if it is a common one, such as the existential clause. We therefore devise a weighting schema for different groups of discourse markers, whose purpose is to represent the markers’ general desirability in the output utterances, as well as to counteract the sparsity of some of the markers compared to others. In other words, the weighting is supposed to ensure all the most desirable utterances are picked from the training set during the selection, but some that only contain less interesting (and typically more

prevalent) discourse phenomena would be omitted in favor of the more complex ones. Our reasoning behind that is that the greater the proportion of the most desirable discourse phenomena in the stylistically selected training set, the more confidently the model is expected to generate utterances in which they are present.

For an illustration, let us assume there are eight different reference utterances for an MR. All of them will be scored based on the discourse markers they contain, but only those that score above a certain threshold will be selected, while the rest will be ignored. The purpose of that is to encourage the model to learn to use, say, a contrastive phrase if there is an opportunity for it in the MR, and not be distracted by other possible realizations of the same MR, which are not as elegant (such as the example utterance #1 vs. #2 in Table 1). Thus, we can set the weighting schema in such a way that sentences containing only, for example, “which” or an existential clause, will not be picked. However, if there is no high scoring utterance for an MR, the utterance with the highest score is picked so that the model would not miss an opportunity to learn from any MR samples.

Our final weighting schema is specified in Table 6. When there are discourse markers from multiple subsets present in the utterance, the weights are accumulated. It is then the total weight that is used to determine whether the utterance satisfies the stylistic threshold or should be eliminated.

The weights have been determined through a combination of the discourse markers’ frequency in the dataset, their intra-category variation, as well as their general desirability in the particular domain of our task. The weights can be easily adjusted for any new domain according to the above, or any other factors. As an example, another such factor could be the length of the utterance. We have experimented with a length penalty, i.e. giving an utterance that contains a verb in gerund form as the only advanced construct, but that is composed of three sentences, a lower score than a short one-sentence utterance with a gerund verb. However, we did not find the use of this extra coefficient helpful in our domain, as it resulted in eliminating a significant proportion of desirable utterances too.

5 Data Annotation

5.1 Contrastive Relation

One of the discourse phenomena whose actualization could benefit from explicit indication of when it should be applied, is the contrastive relation between two (or more) slot realizations in the utterance. There are several reasons why such a comparison of specific slots would be desired in the restaurant domain. One of them is to provide emphasis that one attribute is positive, whereas the other is negative. Another natural reason in dialogue systems could be to indicate that the closest match to the user’s query that was found is a restaurant that does not satisfy one of the requested criteria. A third instance is when the value of one attribute creates the expectation of a particular value of another attribute, but the latter has in reality the opposite value.

Some of the above could presumably be learned by the model if sufficient training data was available. However, they involve fairly complex sentence constructs with various potentially confusing rules for the neural network. The slightly more than 2K samples with a contrasting relation can be drowned among the thousands of other samples in the E2E dataset, meaning that it is difficult for the learned model to produce them.

Hence, we augment the input given to the model with the information about which slots should be put into a contrastive relation. We hypothesize that this explicit indication will help the model to learn to apply contrasting significantly more easily despite the small proportion of training samples exhibiting the property.

In order to extract the information as exactly as possible from the training utterance, we use a heuristic slot aligner (Juraska et al., 2018) to identify two slots that are in a contrastive relation. For the relation we only consider the two scalar slots (*price range* and *customer rating*), plus the boolean slot *family friendly*. Whenever a contrastive relation appears to the aligner to involve a slot other than the above three, we discard it as an undesirable utterance formulation. Depending on the values of the two identified slots, we assign the sample either of the following labels:

- **Contrast:** If the slots have different values on the 3-level positivity scale that they can be mapped to (the *family friendly* slot is only mapped to levels {1, 3}). An example would

Category	Subset of markers	Proportion	Weight
Aggregation	“also, both, neither,...”, quantitative adjectives	1.8%	3
	apposition	4.6%	2
	gerund	11.2%	2
Contrast	“but, however, despite, although,...”	5.4%	3
Fronting	fronted adjective/prepositional/verb clause	14.5%	2
Subordination	subordinating conj.	2.9%	2
	relative pronouns	19.3%	1
Existential clause	expletive “there”	10.0%	1
Imperative/modal	imperative	1.0%	2
	modal verb	4.1%	2

Table 6: The weighting schema for different discourse markers for each introduced category of discourse phenomena. For each set of markers we indicate the heuristically determined proportion of reference utterances in the training set they appear in.

be *customer rating* being “low” (\rightarrow 1) and *family friendly* having value “yes” (\rightarrow 3).

- **Concession:** If the slots have an equivalent value. For instance, *customer rating* being “5 out of 5” (\rightarrow 3) and *price range* having value “cheap” (\rightarrow 3).

The label is added in the form of a new auxiliary slot in the MR, containing the names of the two corresponding slots as its value, such as `<contrast> [priceRange customerRating]`.

We observed instances in the dataset that, semantically, can be classified neither as contrast nor as concession, but using our above rules, they would be considered a concession. An example of such a reference utterance is: “Strada is a low price restaurant located near Rainbow Vegetarian Café serving English food with a *low customer rating but not family-friendly*.” Notice that the emphasized part of the utterance contains a questionable use of the word “but”, as both of the attributes of the restaurant (customer rating and family-friendliness) are negative. Such utterances were, however, scarce, and thus we considered them as an acceptable noise.

5.2 Emphasis

Another utterance property that might in practice be desired to be indicated explicitly and, in that way, enforced in the output utterance, is emphasis. Through fronting discourse phenomena, such as specificational copular constructions or fronted

User query	Is there a family-friendly Indian restaurant nearby?
Response with no emphasis	<i>The Rice Boat</i> in city centre near Express by Holiday Inn is serving Indian food at a high price. It is family-friendly and received a customer rating of 1 out of 5.
Response with emphasis	A family-friendly option is <i>The Rice Boat</i> . This Indian cuisine is priced on the higher end and has a rating of 1 out of 5. They are located near Express by Holiday Inn in the city centre.

Table 7: Example of emphasizing the information about family-friendliness in an utterance conveying the same content.

prepositional phrases, certain information about the subject can be emphasized at the beginning of the utterance.

This could be used to make the dialogue system’s responses sound more context-aware and thus natural. Consider the following example in the restaurant domain. Assume the user asks the agent for a recommendation of a family-friendly Indian restaurant (see Table 7). Considering they have explicitly specified the “family-friendly” requirement in the query, it is arguably more natural for the response utterance to be in the form of the second response example in the table rather than

MR	name [Wildwood], <i>eatType</i> [coffee shop], <i>food</i> [English], priceRange [moderate], customer rating [1 out of 5], <i>near</i> [Ranch]
Reference	A low rated English style coffee shop around Ranch called Wildwood has moderately priced food.
No emph.	Wildwood is a coffee shop providing English food in the moderate price range. It is located near Ranch.
With emph.	There is an English coffee shop near Ranch called Wildwood. It has a moderate price range and a customer rating of 1 out of 5.

Table 8: Examples of generated utterances with or without an explicit emphasis annotation.

the first.

We argue that the order of the information given in the response matters and should not be entirely random. That motivated us to identify instances in the training set where some information about the restaurant is provided in the utterance before its name. In order to do so, and to extract the information about which slot(s) the segment of the utterance represents, we employ the heuristic slot aligner once again. Subsequently, we augment the corresponding input to the model with additional `<emph>` tokens before the slots that should be emphasized in the output utterance. This additional indication will give the model an incentive to learn to realize such slots at the beginning of the utterance when desired. From the perspective of the dialogue manager in a dialogue system, it simply needs to indicate slots to emphasize along with the generated MR whenever applicable.

6 Evaluation

6.1 Experimental Setup

For our sequence-to-sequence NLG model we use the standard encoder-decoder (Cho et al., 2014) architecture equipped with an attention mechanism as defined in Bahdanau et al. (2015). The samples are delexicalized before being fed into the model as input, so as to enhance the ability of the model to generalize the learned concepts to unseen MRs. We only delexicalize categorical slots whose values always propagate verbatim from the MR to the utterance. The corresponding values in the input MR get thus replaced with placeholder tokens for which the values from the original MR are eventually substituted in the output utterance as a part of post-processing.

We use a 4-layer bidirectional LSTM (Hochreiter and Schmidhuber, 1997) encoder and a 4-layer

LSTM decoder, both with 512 cells per layer. During inference time, we use beam search with the beam width of 10 and length normalization of the beams as defined in Wu et al. (2016). The length penalty that we determined was providing the best results on the E2E dataset was 0.6. The beam search candidates are reranked using a heuristic slot aligner as described in Juraska et al. (2018), and the top candidate is returned as the final utterance.

6.2 Style Subsets

In the initial experiments, we trained the model on the reduced training set, which only contains the utterances filtered out based on the weighting schema defined in Table 6. Setting the threshold to 2, we obtained a training set of 17.5K samples, which is approximately 40% of the original training set. Although this reduced training set had a higher concentration of more desirable reference utterances, the dataset turned out to be still too general with most of the rare discourse phenomena drowned out. However, many of them, including contrast, apposition and fronting, appeared multiple times in the generated utterances in the test set, which was not the case for a model trained on the full training set.

Therefore, our next step was to verify whether our model is capable of learning all the concepts of the discourse phenomena individually and apply them in generated utterances. To that end, we repeatedly trained the model on subsets of the E2E dataset, each containing only samples with a specific group of discourse markers, as listed in the second column of Table 6.² We then evaluated the outputs on the correspondingly reduced

²The samples did not necessarily contain the respective discourse marker exclusively, and many exhibited additional markers.

test set, using the same method we used for identifying samples with specific discourse markers, as described in Section 4.1. In other words, we identified what proportion of the generated utterances did exhibit the desired discourse phenomenon.

The results show that the model is indeed able to learn how to produce various advanced sentence structures that are, moreover, syntactically correct despite being trained on a rather small training set (in certain cases less than 2K samples). In all of the experiments, 97–100% of the generated utterances conformed to the style the model was trained to produce. Any occasional incoherence that we observed (e.g. “It has a high customer rating, but *are* not kid friendly.”) was actually picked up from poor reference utterances in the training set. The only exception in the syntactic correctness was the *Imperative/modal* category. Since this is one of the least represented categories among the six, and due to the particularly high complexity and diversity of the utterances, the model trained exclusively on the samples in this category generated a significant proportion of slightly incoherent utterances.

6.3 Data Annotation

The first set of experiments we performed with the data annotation involved explicit indication of emphasis in the input (see Section 5.2). As the results in Table 9 show, the model trained on data with emphasis annotation reached an almost 98% success rate of generating an utterance with the desired slots emphasized.³ In order to get a better idea of the impact of the annotation, notice that the same model trained on non-annotated data does not produce a single utterance with emphasis. The latter model defaults to producing utterances in a rigid style, which always starts with the name of the restaurant (see Table 8).

We notice that the error rate of the slot realization rises (from 3.45% to 5.82%) when the annotation is introduced. Nevertheless, it is still lower than the error rate among the reference utterances in the test set, in which over 8% of slots have missing mentions. Thus we find it acceptable considering the desired stylistic improvement of the output utterances.

The experiments with contrastive relation annotation also show a significant impact of the added

³There were 3,309 slots across all the test MRs that were labeled as to-be-emphasized.

	Emph. realiz.	Slot error rate
Reference	100.00%	8.48%
No emph.	0.00%	3.45%
With emph.	97.85%	5.82%

Table 9: Comparison of the emphasis realization success rate (precision) and the slot realization error rate in the generated outputs using data annotation against the reference utterances, as well as the outputs of the same model trained on non-annotated data.

labels on the style of the output utterances produced by our model. However, the success rate of the realization of a contrast/concession formulation was only 49.12%, and the slot realization error rate jumped up to 8.34%. The contrast and concession discourse phenomena being syntactically more complex, and at the same time being less prevalent among the training utterances, it is understandable that it was more difficult for the model to learn how to use them properly.

6.4 Aggregation

One of the aggregation discourse markers that we identified in Section 4.1 as contributing to the stylistic variation in an interesting way is, unfortunately, very sparsely represented in the E2E dataset. It is the last aggregation type described in the category overview in Section 4.1. Its scarcity in the training set would not make it feasible to train a successful neural model on the subset of the corresponding samples only.

Nevertheless, we analyze the potential for this aggregation in the training set. Since there are only two scalar slots in this dataset – *price range* and *customer rating* – we obtain the frequencies of their value combinations. Both of these take on values on a scale of 3, however, the values are different for each of the slots. Moreover, there are two sets of values for both slots throughout the dataset. We have observed, however, that the values between the two sets are used somewhat interchangeably in the utterances, e.g. “low” seems to be a valid expression of the “less than £20” value of the *price range* slot, and vice versa.

As can be seen in Table 10, the potential for the aggregation is rather limited. Although the 6,604 samples in which a feasible value combination can be found corresponds to over 15% of the training set, due to the values not matching ex-

Price range	Customer rating	Frequency
less than £20	low	2,153
£20-25	3 out of 5	919
moderate	3 out of 5	1,282
more than £30	high	1,329
more than £30	5 out of 5	921

Table 10: Combinations of the slot values for which aggregation would be feasible. Note that only the combinations with a non-zero frequency are listed.

actively between the two slots, aggregation was not elicited in the utterances. Moreover, a high value in the *customer rating* means it is a positive attribute, while a high value in the *price range* slot indicates a negative attribute. We conjecture this might have also deterred the crowd-source workers who produced the utterances from aggregating the values together.

7 Conclusion

In this paper we have presented two different methods of giving a neural language generation system greater stylistic control. Our results indicate that the data annotation method has a significant impact on the model being able to learn how to use a specific style and sentence structures, without an unreasonable impact on the error rate. As our future work, we plan to utilize transfer learning in the style-subset method to improve the model’s ability to apply various different styles at the same time, wherein we would also make further use of the weighting schema. Finally, these methods are a convenient way for achieving the goal of stylistic control when training a neural model with an arbitrary existing large corpus.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *ICLR*.

Merriënboer Cho, Bougares Gulcehre, and Bengio Schwenk. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *EMNLP*.

Ondřej Dušek and Filip Jurčiček. 2016. Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings. *ACL*.

Jessica Fidler and Yoav Goldberg. 2017. Controlling linguistic style aspects in neural language generation. In *Proceedings of the Workshop on Stylistic Variation*, pages 94–104.

Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. Creating training corpora for micro-planners. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vancouver, Canada. Association for Computational Linguistics.

Milica Gašić, Simon Keizer, Francois Mairesse, Jost Schatzmann, Blaise Thomson, Kai Yu, and Steve Young. 2008. Training and evaluation of the HIS POMDP dialogue system in noise. In *Proceedings of the 9th SIGDIAL Workshop on Discourse and Dialogue*, pages 112–119. Association for Computational Linguistics.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

David Howcroft, Crystal Nakatsu, and Michael White. 2013. Enhancing the expression of contrast in the SPaRKY restaurant corpus. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 30–39.

Juraj Juraska, Panagiotis Karagiannis, Kevin K. Bowden, and Marilyn A. Walker. 2018. A deep ensemble model with slot alignment for sequence-to-sequence natural language generation. *NAACL*.

Jad Kabbara and Jackie Chi Kit Cheung. 2016. Stylistic transfer in natural language generation systems using recurrent neural networks. In *Proceedings of the Workshop on Uphill Battles in Language Processing: Scaling Early Achievements to Robust Methods*, pages 43–47.

Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, retrieve, generate: A simple approach to sentiment and style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

François Mairesse, Milica Gašić, Filip Jurčiček, Simon Keizer, Blaise Thomson, Kai Yu, and Steve Young. 2010. Phrase-based statistical language generation using graphical models and active learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1552–1561. Association for Computational Linguistics.

François Mairesse and Marilyn Walker. 2007. Personage: Personality generation for dialogue. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 496–503.

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*.
- Xing Niu and Marine Carpuat. 2017. Discovering stylistic variations in distributional vector space models via lexical paraphrases. In *Proceedings of the Workshop on Stylistic Variation*, pages 20–27.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017a. The E2E dataset: New challenges for end-to-end generation. *arXiv preprint arXiv:1706.09254*.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017b. The E2E NLG shared task. In *SIGDIAL: Conference of the Special Interest Group on Discourse and Dialogue*.
- Jekaterina Novikova, Oliver Lemon, and Verena Rieser. 2016. Crowd-sourcing NLG data: Pictures elicit better data. In *International Conference on Natural Language Generation*.
- Shereen Oraby, Sheideh Homayon, and Marilyn Walker. 2017. Harvesting creative templates for generating stylistically varied restaurant reviews. In *Proceedings of the Workshop on Stylistic Variation*, pages 28–36. Association for Computational Linguistics.
- Shereen Oraby, Lena Reed, Shubhangi Tandon, TS Sharath, Stephanie Lukin, and Marilyn Walker. 2018. Controlling personality-based stylistic variation with neural natural language generators. *CoRR*.
- Daniel S Paiva and Roger Evans. 2004. A framework for stylistically controlled generation. In *Natural Language Generation*, pages 120–129. Springer.
- Sudha Rao and Joel Tetreault. 2018. Dear Sir or Madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 129–140.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40.
- Amanda Stent, Rashmi Prasad, and Marilyn Walker. 2004. Trainable sentence planning for complex information presentation in spoken dialog systems. In *Proceedings of the 42nd annual meeting on association for computational linguistics*, page 79. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve Young. 2016. Multi-domain neural network language generation for spoken dialogue systems. In *Proceedings of the 2016 Conference on North American Chapter of the Association for Computational Linguistics (NAACL)*. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*.