

Can Neural Generators for Dialogue Learn Sentence Planning and Discourse Structuring?

Lena Reed, Shereen Oraby and Marilyn Walker

Natural Language and Dialogue Systems Lab, University of California, Santa Cruz
{lreed, soraby, mawalker}@ucsc.edu

Abstract

Responses in task-oriented dialogue systems often realize multiple propositions whose ultimate form depends on the use of sentence planning and discourse structuring operations. For example a recommendation may consist of an explicitly evaluative utterance e.g. *Chanpen Thai is the best option*, along with content related by the justification discourse relation, e.g. *It has great food and service*, that combines multiple propositions into a single phrase. While neural generation methods integrate sentence planning and surface realization in one end-to-end learning framework, previous work has not shown that neural generators can: (1) perform common sentence planning and discourse structuring operations; (2) make decisions as to whether to realize content in a single sentence or over multiple sentences; (3) generalize sentence planning and discourse relation operations beyond what was seen in training. We systematically create large training corpora that exhibit particular sentence planning operations and then test neural models to see what they learn. We compare models without explicit latent variables for sentence planning with ones that provide explicit supervision during training. We show that only the models with additional supervision can reproduce sentence planning and discourse operations and generalize to situations unseen in training.

1 Introduction

Neural natural language generation (NNLG) promises to simplify the process of producing high

quality responses for conversational agents by relying on the neural architecture to automatically learn how to map an input meaning representation (MR) from the dialogue manager to an output utterance (Gašić et al., 2017; Sutskever et al., 2014). For example, Table 1 shows sample training data for an NNLG with a MR for a restaurant named ZIZZI, along with three reference realizations, that should allow the NNLG to learn to realize the MR as either 1, 3, or 5 sentences.

#	Type	Example
		PRICERANGE[MODERATE], AREA[RIVERSIDE], NAME[ZIZZI], FOOD[ENGLISH], EATTYPE[PUB] NEAR[AVALON], FAMILYFRIENDLY[NO]
1	1 Sent	Zizzi is moderately priced in riverside, also it isn't family friendly, also it's a pub, and it is an English place near Avalon.
2	3 Sents	Moderately priced Zizzi isn't kid friendly, it's in riverside and it is near Avalon. It is a pub. It is an English place.
3	5 Sents	Zizzi is moderately priced near Avalon. It is a pub. It's in riverside. It isn't family friendly. It is an English place.

Table 1: Sentence Scoping: a sentence planning operation that decides what content to place in each sentence of an utterance.

In contrast, earlier models of statistical natural language generation (SNLG) for dialogue were based around the NLG architecture in Figure 1 (Rambow et al., 2001; Stent, 2002; Stent and Molina, 2009).




Figure 1: Statistical NLG Dialogue Architecture

Here the dialogue manager sends one or more dialogue acts and their arguments to the NLG en-

gine, which then makes decisions how to render the utterance using separate modules for content planning and structuring, sentence planning and surface realization (Reiter and Dale, 2000). The sentence planner’s job includes:

- **Sentence Scoping:** deciding how to allocate the content to be expressed across different sentences;
- **Aggregation:** implementing strategies for removing redundancy and constructing compact sentences;
- **Discourse Structuring:** deciding how to express discourse relations that hold between content items, such as causality, contrast, or justification.

Sentence scoping (Table 1) affects the complexity of the sentences that compose an utterance, allowing the NLG to produce simpler sentences when desired that might be easier for particular users to understand. Aggregation reduces redundancy, composing multiple content items into single sentences. Table 2 shows common aggregation operations (Cahill et al., 2001; Shaw, 1998). Discourse structuring is often critical in persuasive settings (Scott and de Souza, 1990; Moore and Paris, 1993), in order to express discourse relations that hold between content items. Table 3 shows how RECOMMEND dialogue acts can be included in the MR, and how content can be related with JUSTIFY and CONTRAST discourse relations (Stent et al., 2002).

Recent work in NNLG explicitly claims that training models end-to-end allows them to do **both** sentence planning and surface realization without the need for intermediate representations (Dusek and Jurcicek, 2016b; Lampouras and Vlachos, 2016; Mei et al., 2016; Wen et al., 2015; Nayak et al., 2017). To date, however, no-one has actually shown that an NNLG can faithfully produce outputs that exhibit the sentence planning and discourse operations in Tables 1, 2 and 3. Instead, NNLG evaluations focus on measuring the semantic correctness of the outputs and their fluency (Novikova et al., 2017; Nayak et al., 2017).

Here, we systematically perform a set of controlled experiments to test whether an NNLG can learn to do sentence planning operations. Section 2 describes our experimental setup and the NNLG architecture that allows us, during training, to vary the amount of supervision provided as to

#	Type	Example
		NAME[THE MILL], EATTYPE[COFFEE SHOP], FOOD[ITALIAN], PRICERANGE[LOW], CUSTOMERRATING[HIGH], NEAR[THE SORRENTO]
4	With, Also	The Mill is a coffee shop with a high rating with a low cost, also The Mill is an Italian place near The Sorrento.
5	With, And	The Mill is a coffee shop with a high rating with a high cost and it is an Italian restaurant near The Sorrento.
6	Distributive	The Mill is a coffee shop with a high rating and cost , also it is an Italian restaurant near The Sorrento.

Table 2: Aggregation Operation Examples

#	Discourse Rel’n	Example
		NAME[BABBO], RECOMMEND[YES], FOOD[ITALIAN], PRICE[CHEAP], QUAL[EXCELLENT], NEAR[THE SORRENTO], LOCATION[WEST VILLAGE], SERVICE[POOR]
7	JUSTIFY ([RECOMMEND] [FOOD, PRICE, QUAL])	I would suggest Babbo because it serves Italian food with excellent quality and it is inexpensive . The service is poor and it is near the Sorrento in the West Village.
8	CONTRAST [PRICE, SERVICE]	I would suggest Babbo because it serves Italian food with excellent quality and it is inexpensive. However the service is poor . It is near the Sorrento in the West Village.

Table 3: Justify & Contrast Discourse Relations

which sentence planning operations appear in the outputs. To ensure that the training data contains enough examples of particular phenomena, we experiment with supplementing crowdsourced data with automatically generated stylistically-varied data from PERSONAGE (Mairesse and Walker, 2011). To achieve sufficient control for some experiments, we exclusively use Personage training data where we can specify exactly which sentence planning operations will be used and in what frequency. It is not possible to do this with crowdsourced data. While our expectation was that an NNLG can reproduce any sentence planning operation that appears frequently enough in the training data, the results in Sections 3, 4 and 5 show that explicit supervision improves the **semantic accuracy** of the NNLG, provides the capability to **control** variation in the output, and enables **generalizing** to unseen value combinations.

2 Model Architecture and Experimental Overview

Our experiments focus on sentence planning operations for: (1) sentence scoping, as in Table 1, where we experiment with controlling the number

of sentences in the generated output; (2) distributive aggregation, as in Example 6 in Table 2, an aggregation operation that can compactly express a description when two attributes share the same value; and (3) discourse contrast, as in Example 8 in Table 3.

Distributive aggregation requires learning a proxy for the semantic property of **equality** along with the standard mathematical **distributive** operation, while discourse contrast requires learning a proxy for semantic comparison, i.e. that some attribute values are evaluated as positive (*inexpensive*) while others are evaluated negatively (*poor service*), and that a successful contrast can only be produced when two attributes are on opposite poles (in either order), as defined in Figure 2.¹

DISTRIBUTIVE AGGREGATION
if $ATTR_1 := ADJ_i$ and $ATTR_2 := ADJ_j$ and $ADJ_i = ADJ_j$ then $DISTRIB(ATTR_1, ATTR_2)$
DISCOURSE CONTRAST
if $EVAL(ADJ_i(ATTR_1)) = POS$ and $EVAL(ADJ_j(ATTR_2)) = NEG$ then $CONTRAST(ATTR_1, ATTR_2)$

Figure 2: Semantic operations underlying distributive aggregation and contrast

Our goal is to test how well NNLG models can produce realizations of these sentence planning operations with varying levels of supervision, while simultaneously achieving high semantic fidelity. Figure 3 shows the general architecture, implemented in Tensorflow, based on TGen, an open-source sequence-to-sequence (seq2seq) neural generation framework (Abadi and others., 2015; Dusek and Jurcicek, 2016a).² The model uses seq2seq generation with attention (Bahdanau et al., 2014; Sutskever et al., 2014) with a sequence of LSTMs (Hochreiter and Schmidhuber, 1997) for encoding and decoding, along with beam-search and an n-best reranker.

The input to the sequence to sequence model is a sequence of tokens $x_t, t \in \{0, \dots, n\}$ that represent the dialogue act and associated arguments. Each x_i is associated with an embedding vector w_i of some fixed length. Thus for each MR, TGen takes as input the dialogue acts represent-

¹We also note that the evaluation of an attribute may come from the attribute itself, e.g. “kid friendly”, or from its adjective, e.g. “excellent service”.

²<https://github.com/UFAL-DSG/tgen>

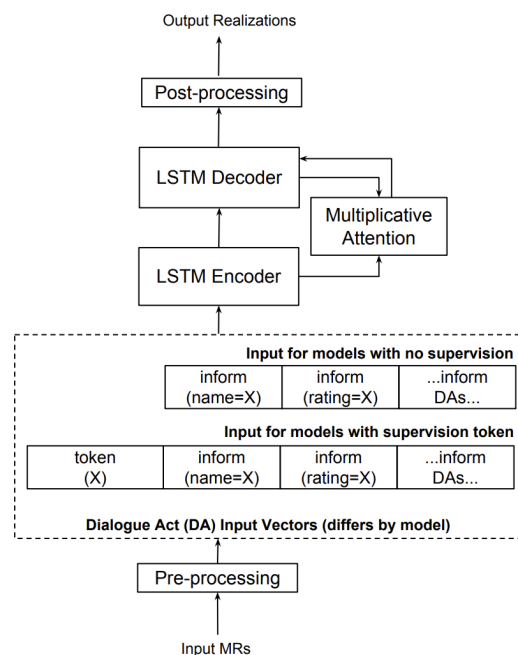


Figure 3: Neural Network Model Architecture, illustrating both the NO SUPERVISION baseline and models that add the TOKEN supervision

ing system actions (*recommend* and *inform* acts) and the attributes and their values (for example, an attribute might be *price range*, and its value might be *moderate*), as shown in Table 1. The MRs (and resultant embeddings) are sorted internally by dialogue act tag and attribute name. For every MR in training, we have a matching reference text, which we delexicalize in pre-processing, then re-lexicalize in the generated outputs. The encoder reads all the input vectors and encodes the sequence into a vector h_n . At each time step t , it computes the hidden layer h_t from the input w_t and hidden vector at the previous time step h_{t-1} , following:

$$h_t = (W_1 \cdot x_t + W_2 \cdot h_{t-1}) + b$$

All experiments use a standard LSTM decoder.

We test three different dialogue act and input vector representations, based on the level of supervision, as shown by the two input vectors in Figure 3: (1) **models with no supervision**, where the input vector simply consists of a set of *inform* or *recommend* tokens each specifying an attribute and value pair, and (2) **models with a supervision token**, where the input vector is supplemented with a new token (either *period* or *distribute* or *contrast*), to represent a latent variable to guide

the NNLG to produce the correct type of sentence planning operation; **(3) models with semantic supervision**, tested only on distributive aggregation, where the input vector is supplemented with specific instructions of which attribute value to distribute over, e.g. *low*, *average* or *high*, in the DISTRIBUTE token. We describe the specific model variations for each experiment below.

Data Sets. One challenge is that NNLG models are highly sensitive to the distribution of phenomena in training data, and our previous work has shown that the outputs of NNLG models exhibit less stylistic variation than their training data (Oraby et al., 2018b). Moreover, even large corpora, such as the 50K E2E Generation Challenge corpus, may not contain particular stylistic variations. For example, out of 50K crowdsourced examples in the E2E corpus, there are 1,956 examples of contrast with the operator “but”. There is only 1 instance of distributive aggregation because attribute values are rarely lexicalized identically in E2E. To ensure that the training data contains enough examples of particular phenomena, our experiments combine crowdsourced E2E data³ with automatically generated data from PERSONAGE (Mairesse and Walker, 2011).⁴ This allows us to systematically create training data that exhibits particular sentence planning operations, or combinations of them. The E2E dataset consists of pairs of reference utterances and their meaning representations (MRs), where each utterance contains up to 8 unique attributes, and each MR has multiple references. We populate PERSONAGE with the syntax/meaning mappings that it needs to produce output for the E2E meaning representations, and then automatically produce a very large (204,955 utterance/MR pairs) systematically varied sentence planning corpus.⁵

Evaluation metrics. It is well known that evaluation metrics used for translation such as BLEU are not well suited to evaluating generation outputs (Belz and Reiter, 2006; Liu et al., 2016; Novikova et al., 2017): they penalize stylistic variation, and don’t account for the fact that different dialogue responses can be equally good, and can vary due to contextual factors (Jordan, 2000;

³<http://www.macs.hw.ac.uk/InteractionLab/E2E/>

⁴Source code for PERSONAGE was provided by François Mairesse.

⁵We make available the sentence planning for NLG corpus at: nlds.soe.ucsc.edu/sentence-planning-NLG.

Krahmer et al., 2002). We also note that previous work on sentence planning has always assumed that sentence planning operations improve the quality of the output (Barzilay and Lapata, 2006; Shaw, 1998), while our primary focus here is to determine whether an NNLG can be trained to perform such operations while maintaining semantic fidelity. Moreover, due to the large size of our controlled training sets, we observe few problems with output quality and fluency.

Thus we leave an evaluation of fluency and naturalness to future work, and focus here on evaluating the multiple targets of semantic accuracy and sentence planning accuracy. Because the MR is clearly defined, we define scripts (information extraction patterns) to measure the occurrence of the MR attributes and their values in the outputs. We then compute Slot Error Rate (SER) using a variant of word error rate:

$$SER = \frac{S + D + I + H}{N}$$

where S is the number of substitutions, D is the number of deletions, I is the number of insertions, H is the number of hallucinations and N is the number of slots in the input MR.

We also define scripts for evaluating the accuracy of the sentence planner’s operations. We check whether: (1) the output has the right number of sentences; (2) attributes with equal values are realized using distributive aggregation, and (3) discourse contrast is used when semantically appropriate. Descriptions of each experiment and the results are in Section 3, Section 4, and Section 5.

3 Sentence Scoping Experiment

To test whether it is possible to control basic sentence scoping with an NNLG, we experiment first with controlling the number of sentences in the generated output, as measured using the period operator. See Table 1. We experiment with two different models:

- **No Supervision:** no additional information in the MR (only attributes and their values)
- **Period Count Supervision:** has an additional supervision token, PERIOD, specifying the number of periods (i.e. the number of sentences) to be used in the output realization.

For sentence scoping, we construct a training set of 64,442 output/MR pairs and a test set of 398 output/MR pairs where the reference utterances

for the outputs are generated from PERSONAGE. Table 4 shows the number of training instances for each MR size for each period count. The right frontier of the table shows that there are low frequencies of training instances where each proposition in the MR is realized in its own sentence (Period = Number of MR attrs -1). The lower left hand side of the table shows that as the MRs get longer, there are lower frequencies of utterances with Period=1.

		Number of Periods						
		1	2	3	4	5	6	7
Attributes	3	3745	167	0	0	0	0	0
	4	5231	8355	333	0	0	0	0
	5	2948	9510	7367	225	0	0	0
	6	821	5002	7591	3448	102	0	0
	7	150	1207	2983	2764	910	15	0
	8	11	115	396	575	388	82	1

Table 4: Distribution of Training Data

We start with the default TGen parameters and monitor the losses on Tensorboard on a subset of 3,000 validation instances from the 64,000 training set. The best settings use a batch size of 20, with a minimum of 5 epochs and a maximum of 20 (with early-stopping based on validation loss). We generate outputs on the test set of 398 MRs.

Sentence Scoping Results. Table 5 shows the accuracy of both models in terms of the counts of the output utterances that realize the MR attributes in the specified number of sentences. In the case of NOSUP, we compare the number of sentences in the generated output to those in the corresponding test reference, and for PERIODCOUNT, we compare the number of sentences in the generated output to the number of sentences we *explicitly* encode in the MR. The table shows that the NOSUP setting fails to output the correct number of sentences in most cases (only a 22% accuracy), but the PERIODCOUNT setting makes only 2 mistakes (almost perfect accuracy), demonstrating almost perfect control of the number of output sentences with the single-token supervision. We also show correlation levels with the gold-standard references (all correlations significant at $p \leq 0.01$).

Model	Slot Error	Period Accuracy	Period Correlation
NOSUP	.06	0.216	0.455
PERIOD COUNT	.03	0.995	0.998

Table 5: Sentence Scoping Results

Generalization Test. We carry out an additional experiment to test generalization of the PERIOD-COUNT model, where we randomly select a set of 31 MRs from the test set, then create a set instance for each possible PERIOD count value, from 1 to the N-1, where N is the number of attributes in that MR (i.e. PERIOD=1 means all attributes are realized in the same sentence, and PERIOD=N-1 means that each attribute is realized in its own sentence, except for the restaurant name which is never realized in its own sentence). This yields 196 MR and reference pairs.

This experiment results in an 84% accuracy (with correlation of 0.802 with the test refs, $p \leq 0.01$). When analyzing the mistakes, we observe that all of the scoping mistakes the model makes (31 in total) are the case of PERIOD=N-1. These cases correspond to the right frontier of Table 4 where there were fewer training instances. Thus while the period supervision improves the model, it still fails on cases where there were few instances in training.

Complexity Experiment. We performed an additional sentence scoping experiment where we specified a target sentence complexity instead of a target number of sentences, since this may more intuitively correspond to a notion of reading level or sentence complexity, where the assumption is that longer sentences are more complex (Howcroft et al., 2017; Siddharthan et al., 2004). We used the same training and test data, but labeled each reference as either high, medium or low complexity. The number of attributes in the MR does not include the name attribute, since that is the subject of the review. A reference was labeled high when there are > 2 attributes per sentence, medium when the number of attributes per sentence is > 1.5 and ≤ 2 and low when there are ≤ 1.5 attributes per sentence.

This experiment results in 89% accuracy. Most of the errors occur when the labeled complexity was medium. This is most likely because there is often only one sentence difference between the two complexity labels. This indicates that sentence scoping can be used to create references with either exactly the number of sentences requested or categories of sentence complexity.

4 Distributive Aggregation Experiment

Aggregation describes a set of sentence planning operations that combine multiple attributes into

Operation	Example
PERIOD	<i>X serves Y. It is in Z.</i>
“WITH” CUE	<i>X is in Y, with Z.</i>
CONJUNCTION	<i>X is Y and it is Z. & X is Y, it is Z.</i>
ALL MERGE	<i>X is Y, W and Z & X is Y in Z</i>
“ALSO” CUE	<i>X has Y, also it has Z.</i>
DISTRIB	<i>X has Y and Z.</i>

Table 6: Scoping and Aggregation Operations in PERSONAGE

single sentences or phrases. We focus here on distributive aggregation as defined in Figure 2 and illustrated in Row 6 of Table 2. In an SNLG setting, the generator achieves this type of aggregation by operating on syntactic trees (Shaw, 1998; Scott and de Souza, 1990; Stent et al., 2004; Walker et al., 2002b). In an NNLG setting, we hope the model will induce the syntactic structure and the mathematical operation underlying it, automatically, without explicit training supervision.

To prepare the training data, we limit the values for PRICE and RATING attributes to LOW, AVERAGE, and HIGH. We reserve the combination {PRICE=HIGH, RATING=HIGH} for test, leaving two combinations of values where distribution is possible ({PRICE=LOW, RATING=LOW} and {PRICE=AVERAGE, RATING=AVERAGE}). We then use all three values in MRs where the price and rating are not the same {PRICE=LOW, RATING=HIGH}. This ensures that the model *does* see the value HIGH in training, but never in a setting where distribution is possible. We always distribute when possible, so every MR where the values are the same uses distribution. All other opportunities for aggregation, in the same sentence or in other training sentences, use the other aggregation operations defined in PERSONAGE as specified in Table 6, with equal probability.

Model	Slot Error	Distrib Accuracy	Distrib Accuracy (on HIGH)
NO SUP	.12	0.29	0.00
BINARY	.07	0.99	0.98
SEMANTIC	.25	0.36	0.09

Table 7: Distributive Aggregation Results

The aggregation training set contains 63,690 total instances, with 19,107 instances for each of the two combinations that can distribute, and 4,246 instances for each of the six combinations that can’t distribute. The test set contains 408 MRs, 288 specify distribution over HIGH (which we note is *not* a setting seen in train, and explicitly tests the

models’ ability to generalize), 30 specify distribution over AVERAGE, 30 over LOW, and 60 are examples that do not require distribution (NONE). We test whether the model will learn the equality relation independent of the value (HIGH vs. LOW), and thus realize the aggregation with HIGH. The distributive aggregation experiment is based on three different models:

- **No Supervision:** no additional information in the MR (only attributes and their values)
- **Binary Supervision:** we add a supervision token, DISTRIBUTE, containing a binary 0 or 1 indicating whether or not the corresponding reference text contains an aggregation operation over attributes *price range* and *rating*.
- **Semantic Supervision:** we add a supervision token, DISTRIBUTE, containing a string that is either *none* if there is no aggregation over *price range* and *rating* in the corresponding reference text, or a value of LOW, AVERAGE, or HIGH for aggregation.

As above, we start with the default TGen parameters and monitor the losses on Tensorboard on subset of 3,000 validation instances from the 63,000 training set. The best settings use a batch size of 20, with a minimum of 5 epochs and a maximum of 20 epochs with early-stopping.

Distributive Aggregation Results. Table 7 shows the accuracy of each model overall on all 4 values, as well as the accuracy specifically on HIGH, the only distribution value unseen in train. Model NO-SUP has a low overall accuracy, and is completely unable to generalize to HIGH, which is unseen in training. It is frequently able to use the HIGH value, but is not able to distribute (generating output like *high cost and cost*). Model BINARY is by far the best performing model, with an almost perfect accuracy (it is able to distribute over LOW and AVERAGE perfectly), but makes some mistakes when trying to distribute over HIGH; specifically, while it is always able to distribute, it may use an incorrect value (LOW or AVERAGE). Whenever BINARY correctly distributes over HIGH, it interestingly always selects attribute RATING before COST, realizing the output as *high rating and price*. Also, BINARY is consistent even when it incorrectly uses the value LOW instead of HIGH: it always selects the attribute *price* before *rating*. To our surprise, Model SEMANTIC does poorly, with 36% overall accuracy, and only 9% accuracy

Source	MR	Realization
NYC	name[xname], recommend[no], cuisine[xcuisine], decor[bad], qual[acceptable], location[xlocation], price[affordable], service[bad]	I imagine xname isn't great because xname is affordable, but it provides bad ambiance and rude service. It is in xlocation. It's a xcuisine restaurant with acceptable food.
E2E	name[xname], cuisine[xcuisine], location[xlocation], familyFriendly[no]	It might be okay for lunch, but it's not a place for a family outing.
E2E	name[xname], eatType[coffee shop], cuisine[xcuisine], price[more than \$30], customerRating[low], location[xlocation], familyFriendly[yes]	Xname is a low customer rated coffee shop offering xcuisine food in the xlocation. Yes, it is child friendly, but the price range is more than \$30.

Table 8: Training examples of E2E and NYC Contrast sentences

Training Sets	NYC #N	E2E #N
3K	N/A	3,540 contrast
7K	3,500 contrast	3,540 contrast
11K	3,500 contrast	3,540 contrast + 4K random
21K	3,500 contrast	3,540 contrast + 14K random
21K CONTRAST	3,500 contrast	3,540 contrast + 14K random

Table 9: Overview of the training sets for contrast experiments

on HIGH, where most of the mistakes on HIGH include repeating the attribute *high rating and rating*, including examples where it does not distribute at all, e.g. *high rating and high rating*. We plan to explore alternative semantic encodings in future work.

5 Discourse Contrast Experiment

Persuasive settings such as recommending restaurants, hotels or travel options often have a critical discourse structure (Scott and de Souza, 1990; Moore and Paris, 1993; Nakatsu, 2008). For example a recommendation may consist of an explicitly evaluative utterance e.g. *Chanpen Thai is the best option*, along with content related by the justify discourse relation, e.g. *It has great food and service*, as in Table 3.

Our experiments focus on DISCOURSE-CONTRAST. We developed a script to find contrastive sentences in the 40K E2E training set by searching for any instance of a contrast cue word, such as *but*, *although*, and *even if*. This identified 3,540 instances. While this data size is comparable to the 3-4K instances used in prior work (Wen et al., 2015; Nayak et al., 2017), we anticipated that it might not be enough data to properly test whether an NNLG can learn to produce discourse contrast. We were also interested in testing whether synthetic data would improve the ability of the NNLG to produce contrastive utterances while maintaining semantic fidelity. Thus we used PERSONAGE with its native database of New York City restaurants (NYC) to generate an additional 3,500 examples of one form

of contrast using only the discourse marker *but*, which are most similar to the examples in the E2E data. Table 8 illustrates both PERSONAGE and E2E contrast examples. While PERSONAGE also contains JUSTIFICATIONS, which could possibly confuse the NNLG, it offers many more attributes that can be contrasted and thus more unique instances of contrast. We create 4 training datasets with contrast data in order to systematically test the effect of the combined training set. Table 9 provides an overview of the training sets, with their rationales below.

3K Training Set. This dataset consists of all instances of contrast in the E2E training data, i.e. 3,540 E2E references.

7K Training Set. We created a training set of 7k references by supplementing the E2E contrastive references with an equal number of PERSONAGE references.

11K Training Set. Since 7K is smaller than desirable for training an NNLG, we created several additional training sets with the aim of helping the model learn to correctly realize domain semantics while still being able to produce contrastive utterances. We thus added an additional 4K crowd-sourced E2E data that was not contrastive to our training data, for a total of 11,065. See Table 9.

21K Training Set. We created an additional larger training set by adding more E2E data, again to test the effect of increasing the size of the training set on realization of domain semantics, without a significant decrease in our ability to produce contrastive utterances. We added an additional 14K E2E references, for a total of 21,065. See Table 9.

We perform two experiments with the 21K training set. First we trained on the MR and reference exactly as we had done for the 7K and 11K training sets. The second experiment added a contrast token during training time with values of either 1 (contrast) or 0 (no contrast) to test if that would achieve better control of contrast.

Contrast Test Sets. To have a potential for contrast there must be an attribute with a positive value and another attribute with a negative value in the same MR. We constructed 3 different test sets, two for E2E and one for NYC. We created a delexicalized version of the test set used in the E2E generation challenge. This resulted in a test of 82 MRs of which only 25 could support contrast (E2E TEST). In order to allow for a better test of contrast, we constructed an additional test set of 500 E2E MRs all of which could support contrast (E2E CONTRAST TEST). For the NYC test, which provides many opportunities for contrast, we created a dataset of 785 MRs that were different than those seen in training (NYC TEST). At test time, in the 21K contrast token experiment, we utilize the contrast token as we did in training.

Train	E2E Test (N = 82)		
	SLOT ERRORS	CONTRAST ATTEMPTS	CONTRAST CORRECT
3K	.38	13	.15
7K	.56	61	.41
11K	.31	24	.33
21K	.28	2	.50
21K CONTRAST	.24	25	.84

Table 10: Slot Error Rates and Contrast for E2E

Train	E2E Contrast Test (N=500)		
	SLOT ERRORS	CONTRAST ATTEMPTS	CONTRAST CORRECT
3K	.70	213	.19
7K	.45	325	.22
11K	.23	227	.70
21K	.17	13	.62
21K CONTRAST	.16	422	.75

Table 11: Slot Error Rates and Contrast for E2E, Contrast Only

Train	NYC Test (N = 785)		
	SLOT ERRORS	CONTRAST ATTEMPTS	CONTRAST CORRECT
3K	N/A	N/A	N/A
7K	.29	784	.65
11K	.26	696	.71
21K	.25	659	.82
21K CONTRAST	.24	566	.85

Table 12: Slot Error Rates and Contrast for NYC

Contrast Results. We present the results for both slot error rates and contrast for the E2E test set in Table 10, E2E Contrast in Table 11, and NYC test set in Table 12.

Table 10 shows the results for testing on the original E2E test set, where we only have 25 instances with the possibility for contrast. Overall, the table shows large performance improvements with the CONTRAST token supervision for 21K for both slot errors and correct contrast. On the E2E test set, the the 3K E2E training set gives a slot error rate of .38 and only 15% correct contrast. The 7K training set, supplemented with additional generated contrast examples gets a correct contrast of .41 but a much higher slot error rate. Interestingly, the 11K dataset is much better than the 3K for contrast correct, suggesting a positive effect for the automatically generated contrast examples along with more E2E training data. The 21K set without the contrast token does not attempt contrast since the frequency of contrast data is low, but with the CONTRAST token, it attempts contrast every time it is possible (25/25 instances).

In Table 11 with only contrast data, we see similar trends, with the lowest slot error rate (.16) and highest correct contrast (.75) ratios for the experiment with token supervision on 21K. Again, we see much better performance from the 11K set than the 3K and 7K in terms of slot error and correct contrast, indicating that more training data (even if that data does not contain contrast) helps the model. As before, we see very low contrast attempts with 21K without supervision, with a huge increase in the number of contrast attempts when using token supervision (422/500).

Table 12 also shows large performance improvements from the use of the CONTRAST token supervision for the NYC test set, again with improvements for the 21K CONTRAST in both slot error rate and in correct contrast. Interestingly, while we get the highest correct contrast ratio of .85 with 21K CONTRAST, we actually see *fewer* contrast attempts, showing that the most explicitly supervised model is becoming more selective when deciding when to do contrast. When training on the 7K dataset, the neural model **always** produces a contrastive utterance for the NYC MRs (all the NYC data is contrastive). Although it never sees any NYC non-contrastive MRs, the additional E2E training data allows it to improve its ability to decide when to contrast (Row 21K CONTRAST) as well as improving the slot error rate in the final experiment.

6 Related Work

Much of the previous work focused on sentence planning was done in the framework of statistical NLG, where each module was assumed to require training data that matched its representational requirements. Methods focused on training individual modules for content selection and linearization (Marcu, 1997; Lapata, 2003; Barzilay and Lapata, 2005), and trainable sentence planning for discourse structure and aggregation operations (Stent and Molina, 2009; Walker et al., 2007; Paiva and Evans, 2004; Sauper and Barzilay, 2009; H. Cheng and Mellish, 2001). Previous work also explored statistical and hybrid methods for surface realization (Langkilde and Knight, 1998; Bangalore and Rambow, 2000; Oh and Rudnicky, 2002). and text-to-speech realizations (Hitzeman et al., 1998; Bulyko and Ostendorf, 2001; Hirschberg, 1993).

Other work on NNLG also uses token supervision and modifications of the architecture in order to control stylistic aspects of the output in the context of text-to-text or paraphrase generation. Some types of stylistic variation correspond to sentence planning operations, e.g. to express a particular personality type (Oraby et al., 2018b; Mairesse and Walker, 2011; Oraby et al., 2018a), or to control sentiment and sentence theme (Ficler and Goldberg, 2017). Herzig et al. (2017) automatically label the personality of customer care agents and then control the personality during generation. Rao and Tetreault (2018) train a model to paraphrase from formal to informal style and Niu and Bansal (2018) use a high precision classifier and a blended language model to control utterance politeness.

Previous work on contrast has explored how the user model determines which values should be contrasted, since people may have differing opinions about whether an attribute value is positive or negative (e.g. *family friendly*) (Carenini and Moore, 1993; Walker et al., 2002a; White et al., 2010). To our knowledge, no-one has yet trained an NNLG to use a model of user preferences for content selection. Here, values are treated as inherently good or bad, e.g. service is ranked from great to terrible.

7 Discussion and Conclusion

This paper presents detailed, systematic experiments to test the ability of NNLG models to produce complex sentence planning operations for re-

sponse generation. We create new training and test sets designed specifically for testing sentence planning operations for sentence scoping, aggregation and discourse contrast, and train novel models with increasing levels of supervision to examine how much information is required to control neural sentence planning. The results show that the models benefit from extra latent variable supervision, which improves the **semantic accuracy** of the NNLG, provides the capability to **control** variation in the output, and enables **generalizing** to unseen value combinations.

In future work we plan to test these methods in different domains, e.g. the WebNLG challenge or WikiBio dataset (Wiseman et al., 2018; Colin et al., 2016). We also plan to experiment with more complex sentence planning operations and test whether an NNLG system can be endowed with fine-tuned control, e.g. controlling multiple aggregation operations. Another possibility is that hierarchical input representations representing the sentence plan might improve performance or allow finer-grained control (Moore et al., 2004; Su and Chen, 2018; Bangalore and Rambow, 2000). It may be desirable to control which attributes are aggregated together, distributed or contrasted, and to allow more than two values to be contrasted.

Here, our main goal was to test the ability of different neural architectures to learn particular sentence planning operations that have been used in previous work in SNLG. Because we don't make claims about fluency or naturalness, we did not evaluate these with human judgements. Instead, we focused our evaluation on automatic assessment of semantic fidelity, and the extent to which the neural architecture could reproduce the desired sentence planning operations. In future work, we hope to quantify the extent to which human subjects prefer the outputs where the sentence planning operations have been applied.

8 Acknowledgments

This work was supported by NSF Cyberlearning EAGER grant IIS 1748056 and NSF Robust Intelligence IIS 1302668-002 as well as an Amazon Alexa Prize Gift 2017 and Grant 2018 awarded to the Natural Language and Dialogue Systems Lab at UCSC.

References

- Martín Abadi and others. 2015. [TensorFlow: Large-scale machine learning on heterogeneous systems](#). Software available from tensorflow.org.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *CoRR*, abs/1409.0473.
- S. Bangalore and O. Rambow. 2000. Exploiting a probabilistic hierarchical model for generation. In *Proc. of the 18th Conference on Computational Linguistics*, pages 42–48.
- Regina Barzilay and Mirella Lapata. 2005. Collective content selection for concept-to-text generation. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 331–338. Association for Computational Linguistics.
- Regina Barzilay and Mirella Lapata. 2006. Aggregation via set partitioning for natural language generation. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 359–366. Association for Computational Linguistics.
- Anja Belz and Ehud Reiter. 2006. Comparing automatic and human evaluation of nlg systems. In *EACL*.
- I. Bulyko and M. Ostendorf. 2001. Joint Prosody Prediction and Unit Selection for Concatenative Speech Synthesis. In *ICASSP*, volume II, pages 781–784.
- Lynne Cahill, John Carroll, Roger Evans, Daniel Paiva, Richard Power, Donia Scott, and Kees van Deemter. 2001. From rags to riches: exploiting the potential of a flexible generation architecture. In *Meeting of the Association for Computational Linguistics*.
- Giuseppe Carenini and Johanna Moore. 1993. Generating explanation in context. In *Proc. of the International Workshop on Intelligent User Interfaces*.
- Emilie Colin, Claire Gardent, Yassine Mrabet, Shashi Narayan, and Laura Perez-Beltrachini. 2016. The webnlg challenge: Generating text from dbpedia data. In *Proceedings of the 9th International Natural Language Generation conference*, pages 163–167.
- Ondrej Dusek and Filip Jurčicek. 2016a. [A context-aware natural language generator for dialogue systems](#). volume abs/1608.07076.
- Ondrej Dusek and Filip Jurčicek. 2016b. [Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings](#). volume abs/1606.05491.
- Jessica Fidler and Yoav Goldberg. 2017. Controlling linguistic style aspects in neural language generation. In *Proceedings of the EMNLP Workshop on Stylistic Variation*, pages 94–104.
- Milica Gašić, Dilek Hakkani-Tür, and Asli Celikyilmaz. 2017. Spoken language understanding and interaction: machine learning for human-like conversational systems. *Computer Speech and Language* 46, pages 249 – 251.
- Renate Henschel H. Cheng, Massimo Poesio and Chris Mellish. 2001. Corpus-based np modifier generation. In *Proc. of the NAACL*.
- Jonathan Herzig, Michal Shmueli-Scheuer, Tommy Sandbank, and David Konopnicki. 2017. Neural response generation for customer service based on personality traits. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 252–256.
- Julia B. Hirschberg. 1993. Pitch accent in context: predicting intonational prominence from text. *Artificial Intelligence Journal*, 63:305–340.
- Janet Hitzeman, Alan W. Black, Paul Taylor, Chris Mellish, and Jon Oberlander. 1998. On the use of automatically generated discourse-level information in a concept-to-speech synthesis system. In *Proc. of the International Conference on Spoken Language Processing, ICSLP98*, pages 2763–2766.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- David M. Howcroft, Dietrich Klakow, and Vera Demberg. 2017. [The extended sparky restaurant corpus: Designing a corpus with variable information density](#). In *Proc. Interspeech 2017*, pages 3757–3761.
- Pamela W. Jordan. 2000. Influences on attribute selection in redescription: A corpus study. In *Proc. of CogSci2000*.
- Emiel Kraemer, André Verleg, and Sebastiaan van Erk. 2002. Graph-based generation of referring expressions. *Computational Linguistics*, page to appear.
- Gerasimos Lampouras and Andreas Vlachos. 2016. Imitation learning for language generation from unaligned data. In *COLING*, pages 1101–1112. ACL.
- I. Langkilde and K. Knight. 1998. Generation that exploits corpus-based statistical knowledge. In *Proc. of COLING-ACL*.
- M. Lapata. 2003. Probabilistic text structuring: Experiments with sentence ordering. In *Proc. of the ACL*.
- Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proc. of Empirical Methods in Natural Language Processing (EMNLP)*.

- Francois Mairesse and Marilyn A. Walker. 2011. Controlling user perceptions of linguistic style: Trainable generation of personality traits. *Computational Linguistics*.
- Daniel Marcu. 1997. From local to global coherence: A bottom-up approach to text planning. In *Proc. of the 14th National Conference on Artificial Intelligence and 9th Innovative Applications of Artificial Intelligence Conference (AAAI-97/IAAI-97)*, pages 629–636, Menlo Park. AAAI Press.
- Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. 2016. What to talk about and how? selective generation using lstms with coarse-to-fine alignment. In *Proceedings of NAACL-HLT*, pages 720–730.
- J. D. Moore and C. L. Paris. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4).
- Johanna Moore, Mary Ellen Foster, Oliver Lemon, and Michael White. 2004. Generating tailored, comparative descriptions in spoken dialogue. In *Proc. FLAIRS-04*.
- Crystal Nakatsu. 2008. [Learning contrastive connectives in sentence realization ranking](#). In *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, pages 76–79, Columbus, Ohio. Association for Computational Linguistics.
- 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 *Proc. of Interspeech 2017*.
- Tong Niu and Mohit Bansal. 2018. Polite dialogue generation without parallel data. *Transactions of the Association for Computational Linguistics*, 6:273–289.
- Jekaterina Novikova, Ondrej Dúek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for nlg. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252.
- Alice H. Oh and Alexander I. Rudnicky. 2002. Stochastic natural language generation for spoken dialog systems. *Computer Speech and Language: Special Issue on Spoken Language Generation*, 16(3-4):387–407.
- Shereen Oraby, Lena Reed, TS Sharath, Shubhangi Tandon, and Marilyn Walker. 2018a. Neural multivoice models for expressing novel personalities in dialog. *Proc. Interspeech 2018*, pages 3057–3061.
- Shereen Oraby, Lena Reed, Shubhangi Tandon, TS Sharath, Stephanie Lukin, and Marilyn Walker. 2018b. Controlling personality-based stylistic variation with neural natural language generators. In *SIGDIAL*.
- Daniel S. Paiva and Roger Evans. 2004. A framework for stylistically controlled generation. In *Natural Language Generation, Third International Conference, INLG 2004*, number 3123 in LNAI, pages 120–129. Springer.
- O. Rambow, M. Rogati, and M. Walker. 2001. Evaluating a trainable sentence planner for a spoken dialogue travel system. In *Proc. of the Meeting of the Association for Computational Linguistics, ACL 2001*.
- 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 *North American Association of Computational Linguistics Conference, NAACL-18*.
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*. Cambridge University Press.
- Christina Sauper and Regina Barzilay. 2009. Automatically generating wikipedia articles: A structure-aware approach. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*, pages 208–216. Association for Computational Linguistics.
- Donia R. Scott and Clarisse Sieckenius de Souza. 1990. Getting the message across in RST-based text generation. In Robert Dale, Chris Mellish, and Michael Zock, editors, *Current Research in Natural Language Generation*. Academic Press, London.
- James Shaw. 1998. Clause aggregation using linguistic knowledge. In *Proc. of the 8th International Workshop on Natural Language Generation*, Niagara-on-the-Lake, Ontario.
- A. Siddharthan, A. Nenkova, and K. McKeown. 2004. Syntactic simplification for improving content selection in multi-document summarization. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING 2004)*.
- Amanda Stent. 2002. A conversation acts model for generating spoken dialogue contributions. *Computer Speech and Language: Special Issue on Spoken Language Generation*.
- Amanda Stent and Martin Molina. 2009. Evaluating automatic extraction of rules for sentence plan construction. In *Proc. of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 290–297.
- Amanda Stent, Rashmi Prasad, and Marilyn Walker. 2004. Trainable sentence planning for complex information presentation in spoken dialogue systems. In *Meeting of the Association for Computational Linguistics*.

- Amanda Stent, Marilyn Walker, Steve Whittaker, and Preetam Maloor. 2002. User-tailored generation for spoken dialogue: An experiment. In *ICSLP*.
- Shang-Yu Su and Yun-Nung Chen. 2018. Investigating linguistic pattern ordering in hierarchical natural language generation. In *7th IEEE Workshop on Spoken Language Technology (SLT 2018)*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- M. A. Walker, S. J. Whittaker, A. Stent, P. Maloor, J. D. Moore, M. Johnston, and G. Vasireddy. 2002a. Speech-Plans: Generating evaluative responses in spoken dialogue. In *In Proc. of INLG-02*.
- Marilyn Walker, Owen Rambow, and Monica Rogati. 2002b. Training a sentence planner for spoken dialogue using boosting. *Computer Speech and Language: Special Issue on Spoken Language Generation*, 16(3-4):409–433.
- Marilyn A. Walker, Amanda Stent, François Mairesse, and Rashmi Prasad. 2007. Individual and domain adaptation in sentence planning for dialogue. *Journal of Artificial Intelligence Research (JAIR)*, 30:413–456.
- 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 *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Michael White, Robert A. J. Clark, and Johanna D. Moore. 2010. Generating tailored, comparative descriptions with contextually appropriate intonation. *Computational Linguistics*, 36(2):159–201.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2018. Learning neural templates for text generation. *CoRR*, abs/1808.10122.