Neural MultiVoice Models for Expressing Novel Personalities in Dialog

Shereen Oraby, Lena Reed, Sharath TS, Shubhangi Tandon, Marilyn Walker

Natural Language and Dialog Systems Lab, University of California, Santa Cruz

{soraby,shtandon,struvek,lireed,mawalker}@ucsc.edu

Abstract

Natural language generators for task-oriented dialog should be able to vary the style of the output utterance while still effectively realizing the system dialog actions and their associated semantics. While the use of neural generation for training the response generation component of conversational agents promises to simplify the process of producing high quality responses in new domains, to our knowledge, there has been very little investigation of neural generators for task-oriented dialog that can vary their response style, and we know of no experiments on models that can generate responses that are different in style from those seen during training, while still maintaining semantic fidelity to the input meaning representation. Here, we show that a model that is trained to achieve a single stylistic personality target can produce outputs that combine stylistic targets. We carefully evaluate the multivoice outputs for both semantic fidelity and for similarities to and differences from the linguistic features that characterize the original training style. We show that contrary to our predictions, the learned models do not always simply interpolate model parameters, but rather produce styles that are distinct, and novel from the personalities they were trained on.

Index Terms: natural language generation, generalization, style, dialog systems, neural methods

1. Introduction

Natural language generators for task-oriented dialog should be able to vary the style of the output while still effectively realizing the system dialog actions and their associated semantics. The use of neural natural language generation (NNLG) for training the response generation component of conversational agents promises to simplify the process of producing high quality responses in new domains by relying on the neural architecture to automatically learn how to map an input meaning representation to an output utterance. However, there has been little investigation of NNLGs for dialog that can vary their response style, and we know of no experiments on models that can generate responses that are different in style from those seen during training, while still maintaining semantic fidelity to the input meaning representation. Instead, work on stylistic transfer has focused on tasks where only coarse-grained semantic fidelity is needed, such as controlling the sentiment of the utterance (positive or negative), or the topic or entity under discussion [1, 2, 3].

Consider for example a training instance for the restaurant domain consisting of a meaning representation (MR) from the End-to-End (E2E) Generation Challenge[1] and a sample output from one of our neural generation models in Figure 1 [4, 5]. Systems using the training set of 50K crowdsourced utterances from the E2E task achieved high semantic correctness, e.g. the BLEU score for our best system on the dev set was 0.72 [6]. However in the best case these models can only reproduce the

---

[1] http://www.macs.hw.ac.uk/InteractionLab/E2E/

style of the training data, and in actuality the outputs have reduced stylistic variation, because when particular stylistic variations are less frequent, they are treated similarly to noise.

INFORM(NAME[BROWNS CAMBRIDGE], EATTYPE[PUB], PRICERANGE[AVERAGE], FOOD[ITALIAN], NEAR[ADRIATIC]) FAMILYFRIENDLY[YES], AREA[CITY CENTRE])

Browns Cambridge is a pub, also it is a moderately priced italian place near Adriatic, also it is family friendly, you know and it's in the city centre.

Table 1: *Meaning representation and output training pair*

In subsequent work, we showed that we could augment the E2E training data with synthetically generated stylistic variants and train a neural generator to reproduce these variants, however the models can still only generate what they have seen in training [5]. Here, instead, we explore whether a model that is trained to achieve a single stylistic personality target can produce outputs that combine stylistic targets, to yield a novel style that is significantly different than what was seen in training, while still maintaining high semantic correctness. We first train each stylistic model with a single latent variable for supervision, for five different personality models, or voices, based on the Big Five theory of personality, namely the personality trait styles of EXTRAVERT, AGREEABLE, DISAGREEABLE, CONSCIENTIOUS, and UNCONSCIENTIOUS. Then, at generation time, we provide the model with combinations of the stylistic variables, i.e. we instruct the NNLG to generate multivoice outputs that combine EXTRAVERT with DISAGREEABLE, where such combined outputs never occurred in the training data.

We first describe how we set up our dataset and neural models in Section 2, and then present our results in Section 3. We evaluate the multivoice outputs for both semantic fidelity and for similarities to and differences from the linguistic features that characterize the original training style. We hypothesize that controlling multiple stylistic parameters is more difficult and will lead to more semantic errors, so we examine in detail the interaction of stylistic variation and semantic fidelity, as well as quantifying stylistic fidelity. We leave a discussion of related work until Section 4 where we also conclude.

2. Data and Models

There is a long tradition in AI of using slightly synthetic tasks and datasets in order to test the ability of particular models to achieve these tasks [7, 8]. The PERSONAGE corpus [5] provides a controlled environment for testing different models of neural generation and style generation. It consists of 88,500 restaurant domain utterances whose style varies according to models of personality, which were generated by an existing statistical NLG engine that has the capability of manipulating 67 different stylistic parameters [9]. Table 2 shows sample utterances that are output for the singlevoice models and for each of our multi-

Table 2: *MultiVoice generation output and comparable singlevoice outputs for* Disagreeable, Extravert *and* Conscientious *for the meaning representation in Figure 1. We count the frequency of periods (Period Agg.) and expletives (Explet. Prag) for multivoice models that utilize* Disagreeable*).*

	Persona	Period Agg.	Explet. Prag.	Example
1	DISAG	5.71	2.26	Browns Cambridge is damn moderately priced, also it's in city centre. It is a pub. It is an italian place. It is near Adriatic. It is damn family friendly.
2	CONSC	0.60	0.02	Let's see what we can find on Browns Cambridge. I see, well it is a pub, also it is moderately priced, an italian restaurant near Adriatic and family friendly in city centre.
3	DISAG+ CONSC	3.81	0.84	Browns Cambridge is an italian place and moderately priced. It is near Adriatic. It is kid friendly. It is a pub. It is in city centre.
4	EXTRA	0.18	0.01	Browns Cambridge is a pub, also it is moderately priced, an italian place near Adriatic, also it is family friendly, you know and it's in city centre, you know!
5	DISAG+ EXTRA	8.07	1.92	Damn moderately priced Browns Cambridge is it 's a pub, also it is near Adriatic. It is an italian place and moderately priced. It is in city centre.
6	UNCON	0.40	2.88	Oh god yeah, i don't know. Browns Cambridge is a pub, also it is damn family friendly, also it's an italian place near Adriatic, also it is darn moderately priced in city centre.
7	DISAG+ UNCON	2.88	3.16	Oh god i mean, i thought everybody knew that Browns Cambridge is a pub, also it is near Adriatic. It is an italian place and moderately priced. It is in city centre.

voice models (described below) for the same MR. Each output corresponding to each single voice personality is controlled by a set of sentence planning parameters that vary for each personality. These parameters are discussed in Section 3 when we evaluate stylistic fidelity. What is important to note here is that each individual voice represents a distinct stylistic distribution in the training data.

The corpus uses the MRs and training/test splits of the E2E Generation Challenge. There are 3,784 unique MRs in the training set, and the corpus contains 17,771 MR/training utterance pairs for each of the existing models for the personality traits of Agreeable, Disagreeable, Conscientiousness, Unconscientiousness, and Extravert, for a total training set of 88,855 utterances. This guarantees a wide range of variation in parameter combinations. The test set consists of 278 unique MRs. The frequencies of longer utterances (more attribute MRs) vary across train and test with test MRs not seen during training. The training data has more smaller MRs, while the test set is more challenging, with more larger MRs.

Previous work shows that a simple model trained on the whole corpus of 88,855 utterances produces semantically correct outputs, but with reduced stylistic variation [5], while a model that allocates a variable corresponding to a label for each style learns to reproduce the stylistic variation. This is interesting because each style variable (personality) actually encodes a set of 36 different stylistic parameters and their values: the model learns for example how the Disagreeable personality tends to produce many shorter sentences in the output, as well as learning that it tends to use expletives like *damn*, e.g. see the outputs based on Disagreeable personality in Table 2.

Model Description. Our NNLG model uses a single token to represent personality encoding, following the use of single language labels used in machine translation and other work on neural generation [10, 5]. Figure 1 summarizes the model architecture. This model builds on the open-source sequence-to-sequence (seq2seq) TGen system [11], which is implemented in Tensorflow [12].[2] The system is based on the seq2seq generation method with attention [14, 15], and uses a sequence of LSTMs [16] for the encoder and decoder, combined with beam-search and an n-best list reranker for output tuning.

The inputs to the model are dialog acts for each system action (such as *inform*) and a set of attribute slots (such as *rating*) and their values (such as *high* for attribute *rating*). To prepro-
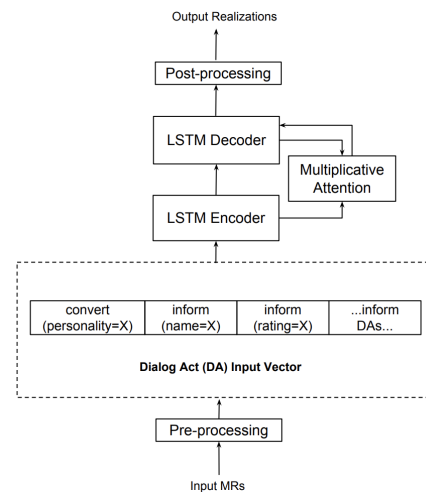

Figure 1: *Neural network model architecture*

cess the corpus of MR/utterance pairs, attributes that take on proper-noun values are delexicalized during training i.e. *name* and *near*. We encode personality as an additional dialog act, of type CONVERT with personality as the key and the target personality as the value (see Figure 1). For every input MR and a personality, we train the model with the corresponding PERSONAGE generated sentence. Our model differs from the TOKEN model used in our previous work [5] because it is trained on unsorted inputs to allow us to add multiple CONVERT tags to the MR at generation time. Note that we do not train on multiple personalities, instead, we train one model that uses all the data, where each distinct single personality has a corresponding CONVERT(PERSONALITY = X) in the training instance.

At generation time, we generate singlevoice data for all the test MRs (1,390 total realizations, 278 unique MRs, realized for each of 5 personalities). For the multivoice experiments, we generate 2 references per combination of two personalities for each of the 278 test MRs, since the order of the CONVERT tags matters. For a given order, the model produces a single output. We do not combine personalities that are exact opposites such as Agreeable and Disagreeable, yielding 8 combinations. The multivoice test set consists of 4,448 total realizations (278 MRs and 8×2 outputs per MR).

---

[2] We refer the reader to TGen publications [11, 13] for model details.

3. Results

Although it is well known that current automatic metrics do not perform well for evaluating the quality of an NLG [30], and that they penalize stylistic variation, we report automatic metrics for completeness. To address their limitations, we also report the results of our own metrics developed to measure semantic correctness and stylistic fidelity. Examples of model outputs for single and multivoice are shown in Table 2, demonstrating how our models interpolate the stylistic parameters described here.

Automatic Metrics. The automatic evaluation uses the E2E generation challenge script.[3] Table 3 summarizes the results for each personality combination for the metrics: BLEU (n-gram precision), NIST (weighted n-gram precision), METEOR (n-grams with synonym recall), and ROUGE (n-gram recall). We note that multivoice automatically has a better chance because the evaluation is over 4,448 examples as opposed to 1,390 for singlevoice, and each multivoice output is compared to 2 possible references (one for each single voice), and then averaged.

Table 3: *Automatic metric evaluation*

Personality	BLEU	NIST	METEOR	ROUGE_L
SINGLEVOICE	0.35	4.93	0.36	0.50
MULTIVOICE	0.42	5.64	0.36	0.52

Semantic Errors. Table 4 shows ratios for the number of deletions, repeats, and hallucinations for each single and multivoice model for their respective test sets (1,390 total realizations and 4,448 realizations). The error counts are split by personality, and normalized by the number of unique MRs (278). Note that smaller ratios are preferable, indicating fewer errors. As we predicted, it is more challenging to preserve semantic fidelity when attempting to hit multiple stylistic targets. We see that in most cases the frequency of errors increase for multivoice compared to singlevoice, with particular combinations such as DISAGREEABLE plus EXTRAVERSION making more than one attribute deletion for each output on average. In the singlevoice results DISAGREEABLE and EXTRAVERT make the most errors with the smallest total ratio found for CONSCIENTIOUS, but when CONSCIENTIOUS combines with DISAGREEBLE it performs worse than either model alone.

Table 4: *Ratio of errors by multivoice personality pairs as compared to singlevoice models*

Personality	Deletions	Repetitions	Hallucinations
AGREE	0.27	0.29	0.34
CONSC	0.22	0.12	0.41
EXTRA	0.74	0.46	0.35
UNCONSC	0.31	0.28	0.29
DISAGREE	0.87	0.81	0.22
Personality Pairs			
AGREE+CONSC	0.44	0.08	0.26
AGREE+EXTRA	0.28	0.17	0.19
AGREE+UNCONSC	0.33	0.24	0.24
CONSC+DISAGR	1.01	0.18	0.28
CONSC+EXTRA	0.67	0.28	0.23
DISAGR+EXTRA	1.20	0.75	0.09
DISAGR+UNCONSC	1.10	0.39	0.14
EXTRA+UNCONSC	1.05	0.55	0.17

Stylistic Characterization. To characterize the differences in style between the multivoice and singlevoice outputs, we develop scripts that count the aggregation operations and prag-

matic markers in Figure 5 in both the singlevoice and multivoice test data. We then compare the singlevoice data directly with multivoice results.

Table 5: *Aggregation and Pragmatic Operations*

Attribute	Example
AGGREGATION OPERATIONS	
PERIOD	*X serves Y. It is in Z.*
"WITH" CUE	*X is in Y, with Z.*
CONJUNCTION	*X is Y and it is Z. & X is Y, it is Z.*
ALL MERGE	*X is Y, W and Z & X is Y in Z*
"ALSO" CUE	*X has Y, also it has Z.*
PRAGMATIC MARKERS	
ACK_DEFINITIVE	*right, ok*
ACK_JUSTIFICATION	*I see, well*
ACK_YEAH	*yeah*
CONFIRMATION	*let's see what we can find on X, let's see ....., did you say X?*
INITIAL REJECTION	*mmm, I'm not sure, I don't know.*
COMPETENCE MIT.	*come on, obviously, everybody knows that*
FILLED PAUSE STATIVE	*err, I mean, mmhm*
DOWN_KIND_OF	*kind of*
DOWN_LIKE	*like*
DOWN_AROUND	*around*
EXCLAIM	*!*
INDICATE SURPRISE	*oh*
GENERAL SOFTENER	*sort of, somewhat, quite, rather*
DOWN_SUBORD	*I think that, I guess*
EMPHASIZER	*really, basically, actually, just*
EMPH_YOU_KNOW	*you know*
EXPLETIVES &	*oh god, damn, oh gosh, darn*
IN GROUP MARKER	*pal, mate, buddy, friend*
TAG QUESTION	*alright?, you see? ok?*

The aggregation parameters in Table 5 control how the NLG combines attributes into sentences, e.g. whether it tries to create complex sentences and what types of combination operations it uses. The pragmatic operators in the bottom part of Table 5 are intended to achieve particular pragmatic effects in the generated outputs: for example the use of a hedge such as *sort of* softens a claim and affects perceptions of friendliness and politeness [17], while the exaggeration associated with emphasizers like *actually, basically, really* influences perceptions of extraversion and enthusiasm [18, 19]. Each parameter value can be set to `high`, `low`, or `don't care`.

Aggregation. To measure the similarity of each multivoice model to its parent single voices for aggregation operations, we first count the average number of times each aggregation operation occurs for each model and personality or personality combination. We then compute Pearson correlation across different model outputs to quantify the similarity of these model outputs with respect to the aggregation operations. Table 6 provides a summary of these results (higher means more correlated).

The final column of Table 6 provides the correlations between the original two single voices that were put together to create the multivoice model. This shows for example (Row 1) that AGREEABLE and CONSCIENTIOUS are similar in their use of aggregation but that DISAGREEABLE and EXTRAVERSION are very dissimilar (Row 6). We would expect that models that are similar to start with would be less novel when they are combined, and indeed Row 1 shows that when the multivoice model is compared with both the original AGREEABLE voice (Column 3) and the CONSCIENTIOUS voice (Column 4) the use of aggregation operations changes little. However other combinations seem to produce completely novel models that use aggregation very differently than either of their singlevoice source models.

---

[3] https://github.com/tuetschek/e2e-metrics

For example in Row 7 the combination of DISAGREEABLE and UNCONSCIENTIOUS produces a model whose use of aggregation is distinct from either of its source models. All of the correlations in Table 6 are significant ($p < 0.05$) except for the 0.01 correlation when comparing the single voices of CONSCIENTIOUS vs. DISAGREE where the p-value is 0.6.

Table 6: *Correlations between* PERSONAGE *data and multivoice models for the aggregation operations in Table 5*

#	P1	P2	P1+P2 vs. P1	P1+P2 vs. P2	P1 vs. P2
1	AGREE	CONSC	0.74	0.76	0.74
2	AGREE	EXTRA	0.70	0.31	0.44
3	AGREE	UNCONSC	0.75	0.31	0.65
4	CONSC	DISAGR	0.36	0.65	0.01
5	CONSC	EXTRA	0.51	0.31	0.44
6	DISAGR	EXTRA	0.53	-0.36	-0.04
7	DISAGR	UNCONSC	0.23	0.33	0.05
8	EXTRA	UNCONSC	0.20	0.43	0.47

Figure 2 provides a closer look at particular aggregation operations associated with CONSCIENTIOUSNESS and DISAGREEABLE and plots the differences between the singlevoice models and the use of these operations in the multivoice models. Interestingly, these plots also clearly show that the multivoice model is a novel personality, yielding a different distribution for aggregation operations than either of its source voice styles.
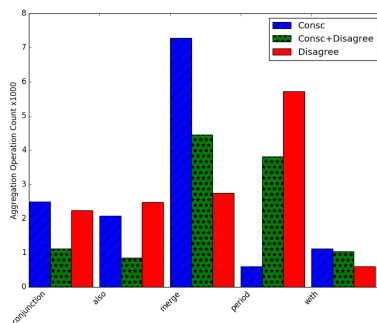

Figure 2: *Frequency of the most frequent aggregation operations for Conscientiousness and Disagreeable compared to combined Conscientiousness and Disagreeable multivoice*

Pragmatic Marker Usage. To measure the models' use of pragmatic markers, we count the number of times each marker in Table 5 occurred in the model outputs, compared to the singlevoice references. We again compute the Pearson correlation between the original voices and the multivoice model outputs for personality combination. The results are shown in Table 7 (all correlations significant with $p \leq 0.05$).

Table 7: *Correlations between* PERSONAGE *data and multivoice models for the pragmatic markers in Table 5*

#	P1	P2	P1+P2 vs. P1	P1+P2 vs. P2	P1 vs. P2
1	AGREE	CONSC	0.11	0.74	0.30
2	AGREE	EXTRA	0.19	-0.02	-0.07
3	AGREE	UNCONSC	0.03	0.18	-0.16
4	CONSC	DISAGR	0.44	0.05	-0.10
5	CONSC	EXTRA	0.41	-0.09	-0.11
6	DISAGR	EXTRA	0.12	-0.03	-0.07
7	DISAGR	UNCONSC	0.09	0.34	-0.05
8	EXTRA	UNCONSC	-0.11	0.37	-0.08

The final column of Table 7 provides the correlations between the original two single voices that were put together to create the multivoice model. As we can see in Row 1, the only two voices that are similar to start are AGREEABLE and CONSCIENTIOUS. All of the other voices have negative correlations with one another in their use of pragmatic markers. Interestingly, the multivoice combination of AGREEABLE and CONSCIENTIOUS resembles CONSCIENTIOUS much more (see column 4). All the other multivoice models also appear to resemble one of the parent models more than the other, but none are *very* similar to their parents: they each appear to demonstrate characteristics of a novel voice. For example, in Row 6, the combination of DISAGREEABLE and EXTRAVERSION produces a model that bears very little similarity to either DISAGREEABLE (0.12 correlation) or EXTRAVERSION (-0.03 correlation).

Figure 3 provides a closer look at particular pragmatic markers associated with CONSCIENTIOUSNESS and DISAGREEABLE and plots the differences between the singlevoice models and the multivoice models. Again, interestingly, these plots show that the multivoice model is a novel personality that yields a different distribution for pragmatic markers than either of its source voice styles.
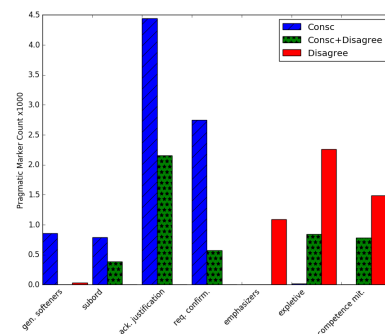

Figure 3: *Frequency of the most frequent pragmatic markers for Conscientiousness and Disagreeable compared to combined Conscientiousness and Disagreeable multivoice*

4. Related Work and Conclusion

The restaurant domain has been a testbed for conversational agents for over 25 years [20, 21, 22, 23, 24], but there is little previous work examining stylistic variation in this domain [9, 25]. Most of the recent research using neural NLG has focused on semantic fidelity [26, 27, 13, 28], however there is work on methods for controlling when long utterances should be split into shorter ones, and for attempting to enforce pronominalization [29]. Other work has pointed out how poor evaluation metrics such as BLEU are for evaluating natural language generation quality [30].

Recent work on neural methods for controlling linguistic style has mainly been carried out in the context of machine translation [1] or focused on tasks where semantic fidelity was not required [31]. Previous work in the statistical NLG tradition presents methods for controlling stylistic variation [32, 33, 34]. Work on the *persona* of a conversational agent did not actually focus on stylistic variation, or personality, but instead tried to ensure that an open domain conversational agent would answer questions about itself in a semantically consistent way [35].

Here we present the first experiment, to our knowledge, examining stylistic generalization in a domain that requires semantic fidelity. We show that our neural models produce novel styles that they have not seen in training, and examine how and to what extent stylistic control interacts with semantic fidelity.

5. References

[1] R. Sennrich, B. Haddow, and A. Birch, "Controlling politeness in neural machine translation via side constraints," in *NAACL*, 2016.

[2] T. Shen, T. Lei, R. Barzilay, and T. S. Jaakkola, "Style transfer from non-parallel text by cross-alignment," in *NIPS*, 2017, pp. 6833–6844.

[3] A. Fan, D. Grangier, and M. Auli, "Controllable abstractive summarization," *CoRR*, vol. abs/1711.05217, 2017. [Online]. Available: http://arxiv.org/abs/1711.05217

[4] J. Novikova, O. Lemon, and V. Rieser, "Crowd-sourcing nlg data: Pictures elicit better data," in *International Conference on Natural Language Generation*, 2016.

[5] S. Oraby, L. Reed, S. Tandon, S. TS, S. Lukin, and M. Walker, "Controlling personality-based stylistic variation with neural natural language generators," in *Proceedings of SIGDIAL 2018*, 2018.

[6] J. Novikova, O. Dušek, and V. Rieser, "The E2E dataset: New challenges for end-to-end generation," in *Proceedings of the 18th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Saarbrucken, Germany, 2017, arXiv:1706.09254. [Online]. Available: https://arxiv.org/abs/1706.09254

[7] J. Weston, A. Bordes, S. Chopra, and T. Mikolov, "Towards ai-complete question answering: A set of prerequisite toy tasks," *CoRR*, vol. abs/1502.05698, 2015.

[8] J. Dodge, A. Gane, X. Zhang, A. Bordes, S. Chopra, A. H. Miller, A. Szlam, and J. Weston, "Evaluating prerequisite qualities for learning end-to-end dialog systems," *CoRR*, vol. abs/1511.06931, 2015.

[9] F. Mairesse and M. Walker, "Towards personality-based user adaptation: psychologically informed stylistic language generation," *User Modeling and User-Adapted Interaction*, pp. 1–52, 2010.

[10] M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado *et al.*, "Google's multilingual neural machine translation system: enabling zero-shot translation," *arXiv preprint arXiv:1611.04558*, 2016.

[11] O. Dusek and F. Jurcícek, "A context-aware natural language generator for dialogue systems," *CoRR*, vol. abs/1608.07076, 2016. [Online]. Available: http://arxiv.org/abs/1608.07076

[12] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "Tensorflow: A system for large-scale machine learning," in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 265–283. [Online]. Available: http://dl.acm.org/citation.cfm?id=3026877.3026899

[13] O. Dusek and F. Jurcícek, "Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings," in *CoRR*, vol. abs/1606.05491, 2016. [Online]. Available: http://arxiv.org/abs/1606.05491

[14] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *CoRR*, vol. abs/1409.0473, 2014. [Online]. Available: http://arxiv.org/abs/1409.0473

[15] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, 2014, pp. 3104–3112.

[16] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[17] P. Brown and S. Levinson, *Politeness: Some universals in language usage*. Cambridge University Press, 1987.

[18] J. Oberlander and A. Gill, "Individual differences and implicit language: personality, parts-of-speech, and pervasiveness." in *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, 2004, pp. 1035–1040.

[19] J.-M. Dewaele and A. Furnham, "Extraversion: the unloved variable in applied linguistic research," *Language Learning*, vol. 49, no. 3, pp. 509–544, 1999.

[20] J. Polifroni, L. Hirschman, S. Seneff, and V. Zue, "Experiments in evaluating interactive spoken language systems," in *Proc. of the DARPA Speech and NL Workshop*, 1992, pp. 28–33.

[21] A. Stent, R. Prasad, and M. A. Walker, "Trainable sentence planning for complex information presentation in spoken dialog systems." in *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2004.

[22] L. Devillers, H. Maynard, S. Rosset, P. Paroubek, K. McTait, D. Mostefa, K. Choukri, L. Charnay, C. Bousquet, N. Vigouroux *et al.*, "The french media/evalda project: the evaluation of the understanding capability of spoken language dialogue systems." in *LREC*, 2004.

[23] M. Gašic, S. Keizer, F. Mairesse, J. Schatzmann, B. Thomson, K. Yu, and S. Young, "Training and evaluation of the his-pomdp dialogue system in noise," *Proc. Ninth SIGdial*, 2008.

[24] F. Mairesse, M. Gašić, F. Jurčíček, S. Keizer, B. Thomson, K. Yu, and S. Young, "Phrase-based statistical language generation using graphical models and active learning," in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ser. ACL '10. Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 1552–1561. [Online]. Available: http://dl.acm.org/citation.cfm?id=1858681.1858838

[25] N. Dethlefs, H. Cuayáhuitl, H. Hastie, V. Rieser, and O. Lemon, "Cluster-based prediction of user ratings for stylistic surface realisation," *EACL 2014*, p. 702, 2014.

[26] T. Wen, M. Gasic, N. Mrksic, P. Su, D. Vandyke, and S. J. Young, "Semantically conditioned lstm-based natural language generation for spoken dialogue systems," *CoRR*, vol. abs/1508.01745, 2015.

[27] H. Mei, M. Bansal, and M. R. Walter, "What to talk about and how? selective generation using lstms with coarse-to-fine alignment," *CoRR*, vol. abs/1509.00838, 2015. [Online]. Available: http://arxiv.org/abs/1509.00838

[28] G. Lampouras and A. Vlachos, "Imitation learning for language generation from unaligned data." in *COLING*, N. Calzolari, Y. Matsumoto, and R. Prasad, Eds. ACL, 2016, pp. 1101–1112.

[29] N. Nayak, D. Hakkani-Tur, M. Walker, and L. Heck, "To plan or not to plan? discourse planning in slot-value informed sequence to sequence models for language generation," in *Interspeech*, 2017.

[30] J. Novikova, O. Dušek, A. Cercas Curry, and V. Rieser, "Why we need new evaluation metrics for nlg," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2017, pp. 2241–2252. [Online]. Available: http://aclweb.org/anthology/D17-1238

[31] J. Ficler and Y. Goldberg, "Controlling linguistic style aspects in neural language generation," *CoRR*, vol. abs/1707.02633, 2017.

[32] D. S. Paiva and R. Evans, "A framework for stylistically controlled generation," in *Natural Language Generation, Third Internatonal Conference, INLG 2004*, ser. LNAI, A. Belz, R. Evans, and P. Piwek, Eds., no. 3123. Springer, July 2004, pp. 120–129.

[33] A. Isard, C. Brockmann, and J. Oberlander, "Individuality and alignment in generated dialogues," in *Proceedings of the 4th International Natural Language Generation Conference (INLG)*, 2006, pp. 22–29.

[34] D. Z. Inkpen and G. Hirst, "Near-synonym choice in natural language generation," in *Recent Advances in Natural Language Processing III*, G. A. Nicolas Nicolov, Kalina Bontcheva and R. Mitkov, Eds. John Benjamins Publishing Company, 2004.

[35] J. Li, M. Galley, C. Brockett, G. P. Spithourakis, J. Gao, and B. Dolan, "A persona-based neural conversation model," *arXiv preprint arXiv:1603.06155*, 2016.