

Style Control for Schema-Guided Natural Language Generation

Alicia Y. Tsai^{1*}, Shereen Oraby², Vittorio Perera², Jiun-Yu Kao²,
Yuheng Du², Anjali Narayan-Chen², Tagyoung Chung², Dilek Hakkani-Tur²

¹ University of California, Berkeley

² Amazon Alexa AI

aliciatsai@berkeley.edu

{orabys, pererv, jiunyk, yuhendu,

naraanja, tagyoung, hakkaniit}@amazon.com

Abstract

Natural Language Generation (NLG) for task-oriented dialogue systems focuses on communicating specific content accurately, fluently, and coherently. While these attributes are crucial for a successful dialogue, it is also desirable to simultaneously accomplish specific stylistic goals, such as response length, point-of-view, descriptiveness, sentiment, formality, and empathy. In this work, we focus on stylistic control and evaluation for schema-guided NLG, with joint goals of achieving both semantic and stylistic control. We experiment in detail with various controlled generation methods for large pretrained language models: specifically, conditional training, guided fine-tuning, and guided decoding. We discuss their advantages and limitations, and evaluate them with a broad range of automatic and human evaluation metrics. Our results show that while high style accuracy and semantic correctness are easier to achieve for more lexically-defined styles with conditional training, stylistic control is also achievable for more semantically complex styles using discriminator-based guided decoding methods. The results also suggest that methods that are more scalable (with less hyper-parameters tuning) and that disentangle content generation and stylistic variations are more effective at achieving semantic correctness and style accuracy.

1 Introduction

Natural Language Generation (NLG) for task-oriented dialogue focuses on effectively generating responses based on inputs that are frequently in the form of a structured meaning representation (MR) (Moryossef et al., 2019; Dušek et al., 2018; Colin et al., 2016; Wen et al., 2015). Recent work has suggested a schema-guided paradigm for task-oriented dialogue by adding descriptions in natural language form (Lin et al., 2021; Du et al., 2020;

Rastogi et al., 2019; Bapna et al., 2017). Compared to structured MRs, dialogue schemata contain much richer contextual information, leading to better generated outputs.

Although the primary aim of task-oriented NLG is to effectively generate outputs that realize system dialogue actions and communicate their associated contents correctly, it is often desirable to control the stylistic variations of an output. For example, recognizing and reacting to emotions has been shown to enhance task outcomes and user engagement in task-oriented conversations (Fraser et al., 2018). Language generation systems that use corpora and methods without awareness of emotions may generate callous, generic or even biased responses (Bender et al., 2021; Sheng et al., 2019). Depending on the use case or type of system, it may be useful to stylistically vary responses, e.g., using shorter responses for spoken dialogue systems, longer responses if the system includes visual modality through a screen, or emotion-specific responses that appropriately address user sentiment.

Previous work on controlled text generation aimed at achieving stylistic goals has not focused on a schema-guided paradigm where specific content must be communicated correctly; instead, most work focuses more on unconstrained text-to-text generation without explicit meaning representations (Liu et al., 2021; Krause et al., 2020; Keskar et al., 2019; Ghazvininejad et al., 2017). Meanwhile, work on schema-guided NLG has primarily focused on generating fluent outputs that achieve low semantic error rather than achieving stylistic goals (Kale and Rastogi, 2020; Peng et al., 2020; Du et al., 2020). In this paper, we hope to fill the gap on stylistic control and evaluation for schema-guided NLG. Our contributions in this paper are three-fold:

1. We describe how we pre-process and annotate style parameters within the Schema-guided Dialogue (SGD) dataset (Rastogi et al., 2019).

*Work done as an intern at Amazon Alexa AI.

2. We experiment with controlling different styles with various controlled text generation methods that aim to preserve fluency and semantic correctness.
3. We present results with a broad range of evaluation methods, including a detailed human evaluation.

Specifically, we consider three types of methods: *conditional training*, *guided fine-tuning*, and *guided decoding*. We show that conditional training (CT) can be used for both very lexically-defined (e.g., point-of-view) and more semantically complex styles (e.g., empathy). However, CT introduces the need to re-train new models per style and is more effective at learning styles with strong lexical characteristics (e.g., specific language patterns or vocabulary). For guided fine-tuning, we explore the Plug-and-Play Language Model (PPLM) (Dathathri et al., 2020), but show that it requires careful hyper-parameter turning and is prone to degeneration. For guided decoding, we evaluate the beam search weighted decoding (BSWD) method and show that it performs best overall on measures of style accuracy for semantically complex styles. The results suggest that unlike style control for unconstrained text generation where no specific content needs to be communicated, style control under the schema-guided paradigm has stronger restrictions on the degree of freedom allowed for content generation. We show that methods that disentangle content generation and style variations, especially for more semantically complex styles, result in better overall performance on semantic and stylistic control.

2 Related Work

Controllable text generation is an emerging research field. Current methods for controlling styles in text generation involve learning a conditional generative model or designing an appropriate decoding strategy. There are many methods proposed for learning a good conditional generative model. These include conditional training (Kikuchi et al., 2016; Ficler and Goldberg, 2017; Keskar et al., 2019; See et al., 2019), fine-tuning language models with external attribute models or side models (Dathathri et al., 2020; Zhang et al., 2020), fine-tuning models with reinforcement learning and human feedback (Ziegler et al., 2019), training generative adversarial models (Yu et al., 2017), and

training variational auto-encoders (Yu et al., 2017; Hu et al., 2017).

Others have worked on designing a good decoding strategy to guide generation, where the decoding procedure is guided by a scoring function or discriminator. These include weighted decoding (Ghazvininejad et al., 2017; Holtzman et al., 2018; See et al., 2019) and guided generation (Krause et al., 2020). Other lines of work include curating training data with rich style markup to facilitate training models with explicit stylistic supervision Oraby et al. (2019, 2018). While this previous work does focus on controllable text generation, most work has been carried out in a text-to-text generation setting, without specific semantic constraints. Instead, we focus on the task-oriented dialogue framework where specific values must be communicated, and conduct a rigorous evaluation of different methods and their efficacy on different forms of style generation. Other recent work has explored adding additional information such as chit-chat data to task-oriented dialogue (Sun et al., 2021; Madotto et al., 2020) and could potentially provide new opportunities for stylistic control.

3 Data Collection and Annotation

3.1 Schema-to-Template pairs

We use the Schema-guided Dialogue (SGD) dataset¹ to create a rich corpus of schema-to-template pairs. This dataset is one of the largest publicly available corpora of annotated multi-domain, task-oriented dialogues (Rastogi et al., 2019). Each dialogue in the data is represented as a list of user and system utterances. We use only the system-side utterances and annotations since we are focused on system-side generation.

Table 1 shows an example pre-processed data instance and the final flattened input. To create schema-to-template pairs, we follow the pre-processing steps outlined in Du et al. (2020) with the following modifications: (1) we replace the slot values with generic slot values without any slot type by adding a \$ prefix and appending an increasing index (e.g., San Jose → \$slot1) for better generalization; and (2) we use only *domain*, *meaning representations (MRs)*, and *slot description* as input data. Domain provides the context of the conversation (e.g., Restaurants); an

¹The Schema-guided Dialogue Dataset: <https://github.com/google-research-datasets/dstc8-schema-guided-dialogue>

MR contains a dialog act, a slot and a value (e.g., `OFFER(city=$slot2)`); and the slot description describes the meaning of the slot in natural language. Table 8 in Appendix A summarizes the full statistics for the final pre-processed SGD dataset. In summary, we have 1,698 MRs and 118,715 example templates in the training set and 1,137 MRs and 34,598 templates in the test set.

For a single example, there are usually multiple dialogue acts as shown in Table 1. An MR in the original data may also contain multiple values. In such cases, we flatten these MRs into multiple parts, each containing only one dialogue act. For instance, an MR that contains two values originally, e.g., `REQUEST(cuisine=[Mexican, Italian])` becomes two separate dialogue acts, e.g., `REQUEST(cuisine=$slot1)`, `REQUEST(cuisine=$slot2)`. Templates are obtained by delocalizing the utterances with generic slot values (e.g., `$slot1` is a good restaurant). Finally, we flatten the input data into flat natural language strings similar to Budzianowski and Vulić (2019).

3.2 Style Parameters

In order to perform style control, we also need to annotate style parameters for the SGD dataset. Style parameters are features of text that are stylistically expressive. These parameters can be roughly identified at *lexical* (vocabulary and words), *syntactic* (sentence structure) and *semantic* (abstract meaning/emotion) levels (Verma and Srinivasan, 2019). We focus primarily on *lexical* and *semantic* features. Specifically, we characterize lexical style parameters as low-level linguistic features that can be derived from the text directly such as *word count* and *number of adjectives*, and semantic style parameters as high-level styles such as *sentiment* that are more complex to characterize.

Lexical style parameters Table 2 summarizes the lexical styles we annotate for the SGD data and the description of each parameter. In total, we automatically annotate six lexical style parameters for the SGD data. Similar to Zhang et al. (2018), the parameter “*has rare word*” uses the maximum Normalized Inverse Document Frequency (NIDF) to determine whether or not a template contains words that are used less frequently in the corpus.² The complete data distribution for all the style parameters is included in Table 10 in Appendix A.

²Appendix A includes the details of the NIDF calculation.

Semantic style parameters Unlike lexical parameters which consider explicit features such as vocabulary, semantic parameters of style are less lexically-defined. As a result, it is generally harder to annotate these parameters directly from the original data without auxiliary information. In this work, we consider the following semantic parameters: *formality*, *negative sentiment*, *positive sentiment* and *empathy*. Formality and sentiment are common stylistic parameters studied in the stylistic control NLG literature. We also include empathy as an interesting and complex style studied in recent work (Wang et al., 2021; Majumder et al., 2020; Lin et al., 2019; Zhou and Wang, 2018). We train a classifier for each of the four styles and annotate the utterances in SGD with these features. We include more details about the classifiers in Section 4 and show additional information about the dataset used to train each classifier in Table 9 of Appendix A.

4 Baseline Model

Language model Given a sequence of tokens $x = \{x_1, \dots, x_t\}$, the goal of the language model is to model the joint probability of the sequence $p(x) = p(x_1, \dots, x_t)$. The joint probability $p(x)$ is often factorized in terms of the product of conditional probabilities using the chain rule of probability $P(x) = \prod_{t=1}^T P(x_t|x_0, \dots, x_{t-1})$ (Bengio et al., 2001). In recent years, transformer-based models have been used widely to model these conditional probabilities (Vaswani et al., 2017). In this work, we use GPT-2³ as our baseline language model (Radford et al., 2019) and fine-tune the GPT-2 model with the processed SGD data using a flat representation with the beginning of sequence, separator, and end of sequence special tokens.⁴

Semantic style classifiers The classifiers used to annotate semantic parameters are single layer classifiers. To train each classifier, we encode the input $x = \{x_1, \dots, x_t\}$ of length t using the baseline GPT-2 model described above and obtain the last hidden layer o_t for all time steps t . We then take the average representation across time, denoted \bar{o}_t , and train the classifier to predict the target label (e.g., formal vs. informal) from the average repre-

³GPT-2 small from HuggingFace: https://huggingface.co/transformers/model_doc/gpt2.html

⁴e.g., “[BOS] flattened-schema-tokens [SEP] template-tokens [EOS]”

Original schema		Pre-processed schema	
Service	Restaurants_1	Domain	Restaurants
Act	OFFER	Act	OFFER
Slot	restaurant_name	Slot	restaurant_name
Slot Description	Name of the restaurant	Slot Description	Name of the restaurant
Values	71 Saint Peter	Value	slot1
Act	OFFER	Act	OFFER
Slot	City	Slot	City
Slot Description	City where the restaurant is located	Slot Description	City where the restaurant is located
Values	San Jose	Value	slot2
Utterance	I see that at 71 Saint Peter there is a good restaurant which is in San Jose.	Template	I see that at slot1 there is a good restaurant which is in slot2.
Flattened input after pre-processing	restaurants offer restaurant_name name of the restaurant slot1 offer city city where the restaurant is located slot2		

Table 1: Sample system-side schema and the flat natural language strings.

Style Parameter	Description	Condition
Short	Number of tokens in the template. The condition is based roughly on 25th percentile of the number of tokens per template in the training data.	≤ 7
Long	Number of tokens in the template. The condition is based roughly on 75th percentile of the number of tokens per template in the training data.	≥ 15
Has Rare Word	Whether or not the template contains very specific or rarely used words compared to other templates in the corpus	Max NIDF ≥ 0.5
First Person Pronouns	Whether or not the template contains first person pronouns: I, me, my, mine	-
Second Person Pronouns	Whether or not the template contains second person pronouns: you, your, yours	-
Descriptive	Number of adjectives in the template	> 2

Table 2: Lexical style parameters and possible values.

sentation: $f(\bar{o}_t) = f\left(\sum_{t=1}^T \frac{o_t}{T}\right)$. The classifiers are used for annotating the semantic parameters of the SGD data and for the style control models described in Section 5.

5 Style Controlled Text Generation

In this work, we require a controlled generation method that is able to simultaneously render the semantic content given the schemata while achieving the desired stylistic goals. We also require the method to be both stable (preserve the fluency of the response even when the stylistic goals are not met) and general-purpose (can be applied to many styles). Under these requirements and constraints, we discuss three types of controlled generation methods to achieve these goals: *conditional training*, *guided fine-tuning*, and *guided decoding*, and compare their performance in Section 6. To our knowledge, our work is the first to systematically study the effectiveness of these control methods for schema-guided NLG.

5.1 Conditional Training (CT)

Controllable generation entails modeling $p(x|a)$, where a is a control variable and x is the generated sequence. However, a pre-trained language model such as GPT-2 is only trained to learn $p(x)$. On the other hand, conditional training (Kikuchi et al., 2016; Peng et al., 2018; Fan et al., 2018; See et al., 2019) refers to directly learning the conditional generative model $p(x|a)$. The results are of high quality because the model is trained to directly maximize $p(x|a)$, but this comes at the expense of fixing the control variable upfront and of re-training the entire model for each new control variable.

To perform style control, we fine-tune the baseline GPT-2 with the conditional training method. Specifically, each input in the training set is annotated with the variable a that we wish to control, e.g., the length (*short*, *long*) of the input. The value of the control variable a is then added to model vocabulary as a special token (e.g., [LENGTH_SHORT]) and appended to the meaning representation after the [BOS] special token. The model then learns an embedding for each value of a and learns to generate $x = \{x_1, \dots, x_t\}$ con-

ditioned on a value of a and the given meaning representation by optimizing cross-entropy loss.

5.2 Guided Fine-tuning

Unlike conditional training that requires fine-tuning an entire GPT-2 model per style, guided fine-tuning refers to methods that require only fine-tuning a smaller set of the parameters while the majority of the base model stays fixed. In this paper, we consider the recent Plug-and-Play Language Model (PPLM) (Dathathri et al., 2020). In guided fine-tuning methods, the conditional probability $p(x|a) \propto p(x)p(a|x)$ is obtained by fine-tuning the base language model (LM) using an auxiliary discriminator that explicitly models $p(a|x)$. In our work, we use the semantic style classifiers described in Section 4 for the discriminator $p(a|x)$ and the GPT-2 model for the base LM.

The major difficulty of the PPLM method is the problem of degeneration – output that is ungrammatical, incoherent, or repetitive. In practice, we observe that PPLM is prone to generating ungrammatical outputs or getting stuck in a repetitive loop if the hyper-parameters are not carefully tuned. We illustrate the effect of hyper-parameters tuning and the degeneration problem in Appendix B.

5.3 Guided Decoding

While conditional training and guided fine-tuning require fine-tuning the base language model, weighted decoding is applied only at decoding time, requiring no change to the base language model. To control the generation, it re-ranks the probability of words based on a scoring function or discriminator.

Weighted Decoding (WD) In weighted decoding (Ghazvininejad et al., 2017), at time step t , the distribution of the next token x_{t+1} is re-weighted by a semantic style classifier that models $p(a|x)$. The probability of each possible next word w in the vocabulary given the control variable a is then re-computed as

$$p(w|a) = \text{Softmax}\left(p(w)p(a|w)\right). \quad (1)$$

Here $p(w)$ is the probability of the word w calculated by the base language model as the next token given the generated sequence $\{x_1, \dots, x_t\}$, and $p(a|w)$ is the probability of the word w associated with the control variable a .

Beam search weighted decoding (BSWD) The weighted decoding method described above takes

the highest scoring item at each time step. While this approach is effective, it is often non-optimal and can limit the diversity of the generated text. To mitigate this limitation, we can increase the search space at generation time using the beam search algorithm. Given a fixed beam width parameter B , the beam search algorithm selects B best alternatives with the highest probability for an input sequence at each time step. Therefore, the original weighted decoding approach described above is a special case of the beam search algorithm with $B = 1$.

Finally, we note that in Eq. 1, the style classifier is only conditioned on the next possible token w but not the entire past sequence, *i.e.*, the next possible token w plus the text that has been generated $\{x_1, \dots, x_t, w\}$. Empirically, in both WD and BSWD, we observe that maximizing the probability of the desired style by greedily considering only the next generated token, rather than the entire sequence of previously generated tokens, yielded better performance on the SGD data. When the entire sequence representation is used, we find that the re-weighting of the distribution is usually not strong enough to successfully match the desired stylistic goal.

6 Experiments

In this section, we show the experimental results of the methods described in Section 5 for controlling the styles described in Section 3 on the SGD data. The baseline GPT-2 model is fine-tuned on the training set with no control variables, and the conditional training model is fine-tuned with control variable special tokens, *e.g.*, LENGTH_SHORT. Our evaluation is tested on the test set of 1,137 MRs. We focus on controlling a single style at a time in this experiment; however, it is also possible to control for multiple styles – we include details on multiple-style control experiments in Appendix C (with sample outputs in Appendix Table 14).

6.1 Automatic Evaluation Metrics

We focus on evaluating three key dimensions: style accuracy, fluency, and semantic correctness.

Style accuracy To evaluate how effective each controlled generation method is per style, we use the *style accuracy* metric, or the percentage of outputs that conform to the required input style. For lexical styles, this is simply computed using the conditions in Table 2. For semantic styles, we

classify the generated text using the corresponding style classifier and check if the predicted style matches the desired style value. For instance, if the predicted sentiment for generated text with the “positive sentiment” control code does not match the “positive” label, then it is considered incorrect.

Response fluency We use *BLEU* score (n-gram precision with brevity penalty) (Papineni et al., 2002) as a measurement of the response fluency. We acknowledge that lexical overlap metrics are poor measures of quality (Novikova et al., 2017); however, we include BLEU for completeness and further evaluate quality through human judgments.

Semantic correctness We use *slot error rate (SER)* (Luong et al., 2015) to measure the semantic correctness of the generated response as compared to the given MR. SER measures the ratio of semantic errors that the model makes by finding the total number of slot mistakes (deletions, repetitions, and hallucinations) in the generated text (lower is better). SER here only considers slots that have explicit values that must be realized (e.g., $\$slotN$).

6.2 Lexical Style Automatic Evaluation

We evaluate lexical styles with only conditional training as the rest of the methods include semantic style classifiers and are thus not applicable. Table 3 summarizes the performance of conditional training for the six lexical styles. The style accuracy for most styles is generally high, between 80% to nearly 100%, especially for styles marked explicitly by specific words, such as first and second person pronouns. However, we observe that “*descriptive*” has a particularly low accuracy. First, the majority of the references in the training data (95%) have less than two adjectives, making it difficult for the model to learn this kind of style effectively.⁵ Secondly, we observe that conditional training is particularly effective when the style exhibits a clear syntactic characteristic (e.g., length) or a particular set of vocabulary (e.g., pronouns); however, this is not the case for the “*descriptive*” style.

The fluency of the generated text with style control drops slightly as compared to no style control. Having style control often makes the generated text different from its matching responses, *i.e.*, adding extra content or changing linguistic characteristics. Since BLEU score only considers lexical overlap,

⁵The full data distribution can be found in Appendix A Table 10.

this behavior is expected. Finally, we see that there is not much of a performance drop in semantic correctness with respect to SER. The experimental results show that for lexical styles with a clear syntactic pattern or vocabulary, CT can be quite effective.

Table 4 illustrates example outputs using CT when controlling for “*short*”, “*long*” and “*has rare word*” styles.⁶ Interestingly, we see that when asked for a longer response, the model starts to hallucinate extra content (but not slots) not given in the MR in order to satisfy the control variable. This also translates to slightly lower BLEU scores and a higher SER. Methods to enforce better implicit constraints to increase fluency and semantic correctness are important points for future work.

Style	Style Acc.	BLEU \uparrow	SER \downarrow
Baseline (no style control)	-	0.480	0.010
Short	96.4%	0.205	0.009
Long	80.1%	0.381	0.012
Has Rare Word	89.5%	0.377	0.010
First Person Pronouns	99.9%	0.365	0.011
Second Person Pronouns	99.9%	0.447	0.011
Descriptive	19.6%	0.378	0.012

Table 3: Evaluation of lexical style parameters with conditional training.

6.3 Semantic Style Automatic Evaluation

Table 5 summarizes the main results for semantic style evaluation using CT, PPLM, and BSWD. Since CT is trained to directly maximize conditional probability, it frequently has a higher BLEU score and a lower SER across different styles with the exception of “*formal*” BLEU. We note that formality is rather lexically-defined, exhibiting characteristic keywords such as “**please**”, which are frequently picked up by the model, resulting in a particularly high style accuracy for “*formal*” responses and a lower BLEU score.

For the three more semantically complex styles, “*positive sentiment*”, “*negative sentiment*” and “*empathy*”, we see that BSWD achieves a higher style accuracy than CT (at the cost of a lower BLEU and slightly higher semantic error). We note, however, that the drop in BLEU is expected: as the output is steered towards a certain style, its n-gram overlap with the references is more likely to decrease

⁶Example outputs for other lexical styles are included in Appendix C Table 12.

MR: OFFER(restaurant_name=\$slot1), OFFER(city=\$slot2)	
W/o Style Control	\$slot1 is a nice restaurant in \$slot2 that serves curry.
Short	\$slot1 is a nice restaurant in \$slot2.
Long	Okay! the restaurant, \$slot1 located in \$slot2 is a good one and serves Taiwanese dishes.
MR: OFFER(address=\$slot1), OFFER(rating=\$slot2)	
W/o Style Control	There is a nice house at \$slot1 with a \$slot2 rating.
Has Rare Word	There is a lovely residence located at \$slot1. It has a rating of \$slot2.

Table 4: Example outputs for conditional training with lexical styles “short”, “long” and “has rare word” (more content hallucinations for “long”). Example outputs for other lexical styles are included in Appendix C Table 12.

– thus, we use these automatic metrics as an evaluation guide, but leave a more rigorous evaluation to human judgment in the next section. Finally, with careful hyper-parameter tuning, PPLM can achieve similar performance to BSWD on BLEU and SER but does worse on style accuracy. Increasing the style accuracy for PPLM worsens the BLEU and SER score significantly and thus we do not consider it in our human evaluation. In summary, our automatic evaluation shows that for semantic styles, BSWD gives us a good trade-off between consistent style accuracy and semantic fidelity across styles.

Table 6 illustrates example outputs for the three semantic styles using BSWD (with additional examples including combining multiple styles in Appendix C). In general, styles that encapsulate complex phenomena such as “empathy” are harder to generate as shown by their lower style accuracy; nevertheless, we are able to preserve fluency and semantic correctness in most cases.

Style	Model	Style Acc.	BLEU \uparrow	SER \downarrow
Baseline	-	-	0.480	0.010
Formal	CT	70.0%	0.461	0.009
	PPLM	17.5%	0.390	0.016
	BSWD	42.5%	0.496	0.021
Negative	CT	7.5%	0.409	0.012
	PPLM	2.5%	0.373	0.017
	BSWD	25.0%	0.359	0.021
Positive	CT	17.5%	0.442	0.009
	PPLM	32.5%	0.371	0.026
	BSWD	65.0%	0.369	0.021
Empathy	CT	21.5%	0.449	0.009
	PPLM	11.0%	0.368	0.021
	BSWD	31.0%	0.394	0.017

Table 5: Evaluation of semantic styles.

6.4 Human Evaluation

We focus on human evaluation for our semantic styles since they are the most inherently subjective.

We pick a subset of our semantic styles to evaluate, specifically, *formal*, *negative* and *positive*. We also focus on the evaluation of CT and BSWD only since they have an overall better performance in the automatic evaluation and are simpler in nature (e.g., less hyper-parameter tuning). To evaluate style, we ask three human annotators to rate:

- **Style Rating (Sty. Rat.):** How closely the response matches the given style (1 being not at all, 5 being very closely).
- **Fluency (Flu.):** The fluency of the generated response (1 being low, 5 being high).
- **Semantic Error:** Slot errors in the generated response using a 0/1 scale. For analysis purposes, we further break down the items marked “0” into four error types to understand the strengths and weaknesses of each method. For each type of error, the output is only marked “1” if there are no issues.
 - **Deletion (Del.):** Whether the response drops any slot values from the MR.
 - **Repetition (Rep.):** Whether the response repeats any slot values from the MR.
 - **Content Hallucination (Cont. Hal.):** Whether the response includes extra content not given in the MR.
 - **Incorrect Slot Values (Inc. Slot):** Whether the response includes any slot values not given in the MR (a specific type of hallucination).

Table 7 shows the aggregated evaluation (40 samples per style, with three judgments per sample). The results show that outputs generated by BSWD have a higher or comparable style rating

⁷More details on the human evaluation design are in Appendix D.

MR: REQUEST (where_to=none)	
W/o style control	What city are you staying?
Formal	What city are you planning to stay in?
MR: OFFER (address=\$slot1), OFFER (rating=\$slot2)	
W/o style control	There is a house at \$slot1 with a rating of \$slot2.
Positive	There is a nice house with a \$slot2 rating located at \$slot1.
MR: INFORM (rating=\$slot1), NOTIFY_FAILURE (null=none)	
W/o style control	The rating is \$slot1. I was unable to make a reservation.
Empathy	The rating is \$slot1. I'm sorry, but I couldn't make the reservation.

Table 6: Example outputs of semantic styles using beam search weighted decoding (BSWD).

compared to CT, confirming its ability to control semantic styles. We also observe that outputs from BSWD have a higher fluency score across all styles. This means that even though BSWD showed lower BLEU scores in automatic evaluation, its outputs are considered to be more natural and fluent in human evaluation. Finally, we see an interesting difference in error types when comparing CT and BSWD. In general, BSWD is more prone to deleting and hallucinating slots, while CT more frequently generates incorrect slot values. Since BSWD requires no change to the base language model, it is able to obtain a lower incorrect slot value error rate as compared to CT, which requires re-training the language model with a control code. On the other hand, it has a higher deletion and hallucination error rate since during decoding time, BSWD is free to insert or drop content in order to achieve the desired style.

Style	Model	Sty. Rat.	Flu.	Semantic Error			
				Del.	Rep.	Cont. Hal.	Inc. Slot
Formal	CT	3.88	3.83	15%	0%	13%	3.75%
	BSWD	3.95	4.30	19%	0%	15%	1.25%
Negative	CT	2.78	3.54	20%	0%	12.5%	2.5%
	BSWD	2.86	3.93	20%	0%	12.5%	5%
Positive	CT	3.48	3.64	22.5%	0%	22.5%	12.5%
	BSWD	3.42	3.84	27.5%	0%	30%	7.5%

Table 7: Human evaluation results for selected semantic styles and methods.

7 Conclusion

In this work, we focus on stylistic control and evaluation of schema-guided NLG. We discuss three different types of methods for style controlled text generation: conditional training (CT), guided fine-tuning (PPLM), and guided decoding (BSWD). We present a rich set of evaluations to quantify each

method’s ability to achieve various styles while preserving language fluency and minimizing slot errors. Our analysis shows that, in general, styles that encapsulate abstract ideas are naturally harder to generate (*e.g.*, empathy), and methods that require careful hyper-parameters tuning may run into the problems of instability and degeneration (*e.g.*, PPLM) while under-performing in style accuracy. The automatic and human evaluations suggest that simultaneously achieving stylistic goals and realizing schema information requires methods that allow us to separate content generation and stylistic variations. We show that CT and BSWD overcome some of these challenges and are effective at controlling several styles while maintaining good fluency and semantic correctness in most cases. CT is effective for lexical styles with strong syntactic characteristics or distinctive vocabulary sets, while BSWD excels at semantic styles that are more complex to characterize.

For future work, we are interested in extending our analysis to a larger number of styles, and in exploring techniques for understanding and representing styles that are more abstract, such as “empathy”, as well as methods for generating those styles under a task-oriented dialogue framework. Additionally, since abstract styles may be characterized by multiple features (*e.g.*, a combination of sentiment and descriptiveness), we are interested in studying how these underlying features can be represented and incorporated more accurately to improve overall semantic and stylistic control.

Acknowledgments

We would like to thank Sofia Scharfenberg, Jasmin Rehm, and the rest of the Alexa Data Services team for all of their help with preparing and performing the human evaluation study.

References

- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017. Towards zero shot frame semantic parsing for domain scaling. In *Interspeech 2017*.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. 2001. A neural probabilistic language model. In *Advances in Neural Information Processing Systems*, volume 13. MIT Press.
- Paweł Budzianowski and Ivan Vulić. 2019. Hello, it's GPT-2 - how can I help you? towards the use of pre-trained language models for task-oriented dialogue systems. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 15–22, Hong Kong. Association for Computational Linguistics.
- Sven Buechel, Anneke Buffone, Barry Slaff, Lyle Ungar, and João Sedoc. 2018. Modeling empathy and distress in reaction to news stories. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*.
- Emilie Colin, Claire Gardent, Yassine M'rabet, 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, Edinburgh, UK. Association for Computational Linguistics.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*.
- Yuheng Du, Shereen Oraby, Vittorio Perera, Minmin Shen, Anjali Narayan-Chen, Tagyoung Chung, Anushree Venkatesh, and Dilek Hakkani-Tur. 2020. Schema-guided natural language generation. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 283–295, Dublin, Ireland. Association for Computational Linguistics.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2018. Findings of the E2E NLG challenge. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 322–328, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Angela Fan, David Grangier, and Michael Auli. 2018. Controllable Abstractive Summarization. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Jamie Fraser, Ioannis Papaioannou, and Oliver Lemon. 2018. Spoken Conversational AI in Video Games: Emotional Dialogue Management Increases User Engagement. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. Hafez: an interactive poetry generation system. In *Proceedings of ACL 2017, System Demonstrations*, pages 43–48, Vancouver, Canada. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, D. Golub, and Yejin Choi. 2018. Learning to write with cooperative discriminators. *ArXiv*, abs/1805.06087.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, R. Salakhutdinov, and E. Xing. 2017. Toward controlled generation of text. In *ICML*.
- Mihir Kale and Abhinav Rastogi. 2020. Template guided text generation for task-oriented dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6505–6520, Online. Association for Computational Linguistics.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *ArXiv*, abs/1909.05858.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. Controlling output length in neural encoder-decoders. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1328–1338, Austin, Texas. Association for Computational Linguistics.
- Ben Krause, Akhilesh Deepak Gotmare, B. McCann, N. Keskar, Shafiq R. Joty, R. Socher, and Nazneen Rajani. 2020. Gedi: Generative discriminator guided sequence generation. *ArXiv*, abs/2009.06367.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021. Leveraging slot descriptions for zero-shot

- cross-domain dialogue StateTracking. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648, Online. Association for Computational Linguistics.
- Zhaojiang Lin, Andrea Madotto, Jamin Shin, Peng Xu, and Pascale Fung. 2019. MoEL: Mixture of Empathetic Listeners. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 121–132, Hong Kong, China.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. Dexperts: Decoding-time controlled text generation with experts and anti-experts.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, Jamin Shin, and Pascale Fung. 2020. Attention over parameters for dialogue systems. *ArXiv*, abs/2001.01871.
- Navonil Majumder, Pengfei Hong, Shanshan Peng, Jiankun Lu, Deepanway Ghosal, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. MIME: MIMicking emotions for empathetic response generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8968–8979, Online. Association for Computational Linguistics.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. Step-by-step: Separating planning from realization in neural data-to-text generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2267–2277, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Shereen Oraby, Vrindavan Harrison, Abteen Ebrahimi, and Marilyn Walker. 2019. Curate and Generate: A Corpus and Method for Joint Control of Semantics and Style in Neural NLG. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5938–5951, Florence, Italy. Association for Computational Linguistics.
- Shereen Oraby, Lena Reed, Shubhangi Tandon, Sharath T.S., Stephanie Lukin, and Marilyn Walker. 2018. Controlling Personality-Based Stylistic Variation with Neural Natural Language Generators. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 180–190, Melbourne, Australia. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ellie Pavlick and Joel Tetreault. 2016. An Empirical Analysis of Formality in Online Communication. *Transactions of the Association for Computational Linguistics*, 4:61–74.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot natural language generation for task-oriented dialog. In *EMNLP*.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. 2018. Towards Controllable Story Generation. In *Proceedings of the First Workshop on Storytelling*, pages 43–49, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, R. Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *arXiv preprint arXiv:1909.05855*.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. What makes a good conversation? How controllable attributes affect human judgments. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1702–1723, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment tree-bank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie. 2021. [Adding chit-chat to enhance task-oriented dialogues](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1570–1583, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Gaurav Verma and Balaji Vasan Srinivasan. 2019. [A Lexical, Syntactic, and Semantic Perspective for Understanding Style in Text](#). *arXiv:1909.08349 [cs]*. ArXiv: 1909.08349.
- Yi-Hsuan Wang, Jia-Hao Hsu, Chung-Hsien Wu, and Tsung-Hsien Yang. 2021. [Transformer-based empathetic response generation using dialogue situation and advanced-level definition of empathy](#). In *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pages 1–5.
- 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*, pages 1711–1721, Lisbon, Portugal. Association for Computational Linguistics.
- Lantao Yu, W. Zhang, J. Wang, and Y. Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*.
- J. Zhang, Alexander Sax, A. Zamir, L. Guibas, and J. Malik. 2020. Side-tuning: A baseline for network adaptation via additive side networks. In *ECCV*.
- Ruqing Zhang, Jiafeng Guo, Yixing Fan, Yanyan Lan, Jun Xu, and Xueqi Cheng. 2018. [Learning to Control the Specificity in Neural Response Generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1108–1117, Melbourne, Australia. Association for Computational Linguistics.
- Xianda Zhou and William Yang Wang. 2018. MojiTalk: Generating emotional responses at scale. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, T. Brown, A. Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *ArXiv*, abs/1909.08593.

A Statistics of Style Annotation for SGD Data

	Train	Dev	Test
System Turn	164,982	24,363	42,297
Meaning Representations	1,698	721	1,137
Templates	118,715	19,963	34,598

Table 8: Schema-Guided Dialogue dataset statistics.

Normalized inverse document frequency calculation For a word w in a template d , the Inverse Document Frequency (IDF) of w is

$$\text{IDF}(w) = \log \left(\frac{N}{1 + |\{d \in D : w \in d\}|} \right) \quad (2)$$

where N is the total number of templates in the training data D , and $|\{d \in D : w \in d\}|$ is the number of those templates d in the training data that contain w . The Normalized IDF (NIDF) is obtained via the min-max normalization

$$\text{NIDF}(w) = \frac{\text{IDF}(w) - \min_{w' \in D}(\text{IDF}(w'))}{\max_{w' \in D}(\text{IDF}(w')) - \min_{w' \in D}(\text{IDF}(w'))} \quad (3)$$

where $\max_{w' \in D}(\text{IDF}(w'))$ and $\min_{w' \in D}(\text{IDF}(w'))$ are the maximum and minimum IDF value of all the words in the training data. We use the maximum NIDF of all the words in a template to determine whether the template contains rarely used words.

Semantic style parameter annotation details

The formality corpus⁸ (Pavlick and Tetreault, 2016), which contains sentence-level formality annotations, is used for training the formality classifier. The Stanford sentiment treebank (SST)⁹ (Socher et al., 2013) is used for training the sentiment classifier. The SST data consists of 5 classes – *very negative*, *negative*, *neutral*, *positive*, and *very positive*. We combine “*very negative*” and “*negative*” into one class and “*very positive*” and “*positive*” into one class. Finally, the empathetic reactions data¹⁰ (Buechel et al., 2018) is used for training the empathy classifier. We use only the *empathy* label from the data.

⁸<http://www.seas.upenn.edu/~nlp/resources/formality-corpus.tgz>

⁹<https://nlp.stanford.edu/sentiment/>

¹⁰https://github.com/wbwp/empathic_reactions

Style parameters	Dataset
Formal	Formality corpus
Positive Negative	Stanford sentiment treebank (SST)
Empathy	Empathetic reactions

Table 9: Semantic style parameters with their possible values and associated training datasets.

Category	Style	Train	Dev	Test
Lexical styles	Short	31.42%	36.96%	32.23%
	Long	26.49%	25.13%	25.44%
	Has rare word	35.83%	37.76%	42.04%
	First person pron.	28.48%	28.88%	28.75%
	Second person pron.	52.89%	52.50%	60.18%
Semantic styles	Descriptive	5.00%	5.32%	3.88%
	Formal	30.59%	34.27%	29.76%
	Negative	6.55%	4.33%	4.95%
	Positive	39.70%	42.43%	37.41%
	Empathy	27.72%	30.04%	25.65%

Table 10: Data distribution for all of the style parameters we annotate in SGD.

B Hyper-parameters Tuning for PPLM

There are several hyper-parameters in PPLM models. We refer the readers to Dathathri et al. (2020) for the full model details. From Section 5.2, the PPLM model can be written as an optimization problem of the form:

$$\min_{\Delta H_t} L_{\text{CE}} \left(f \left(\sum_{t=1}^T \tilde{o}_t / T \right) \right) + \lambda D_{\text{KL}}(\tilde{o}_t \| o_t) \quad (4)$$

$$\text{s.t. } \tilde{o}_t = \text{LM}(x_t, \tilde{H}_t) \quad (5)$$

$$\tilde{H}_t = H_t + \Delta H_t \quad (6)$$

$$o_t = \text{LM}(x_t, H_t) \quad (7)$$

where λ is a hyper-parameter that scales the KL divergence and $f(\cdot)$ is the semantic style classifier learned to produce $p(a|x)$. We minimize the cross-entropy loss $L_{\text{CE}}(\cdot)$ of the attribute model and the Kullback-Leibler (KL) divergence $D_{\text{KL}}(\cdot \| \cdot)$ of the language model. For each time step t , the GPT-2 model generates a history of weight matrices H_t for all hidden layers. To control the generation, we shift the weight matrices H_t in the direction that accomplishes two goals: 1) minimize cross-entropy (CE) of the attribute a under the conditional attribute model $p(a|x)$ and 2) minimize the KL divergence between itself and the unmodified language model $p(x)$. Let ΔH_t be the small shift of H_t such that $\tilde{H}_t = H_t + \Delta H_t$ can shift the last

MR: OFFER(restaurant_name=\$slot1), OFFER(city=\$slot2)	γ_{gm}	α	λ
\$slot1 is a nice restaurant in \$slot2.	0.1	0.01	1
\$slot1 is a nice restaurant in \$slot2.	0.1	0.1	1
\$slot1 is a good restaurant in \$slot2.	0.1	1	1
\$slot1 is a nice restaurant in \$slot2	0.3	0.01	0.01
\$slot1 is a restaurant in \$slot2.	0.3	0.01	0.1
\$slot1 is a restaurant in \$slot2.	0.3	0.01	1
There’s the \$slot4 \$slot1 restaurant in \$slot2. Do you like the same restaurant. Do you like this one more?	0.5	0.01	0.01
\$slot1 \$slot2 I have a restaurant has good rating restaurant has \$slot1	0.5	0.01	0.1
\$slot1 \$slot2 I have a restaurant has good rating restaurant has a restaurant name is	0.5	0.01	1
\$slot1 is a nice restaurant in \$slot2. Do i suggest you like that one I’ve found? You can have a nice meal.	0.5	0.1	1
I recommend mathemat, a \$slot1 is in \$slot2.	0.5	1	1
\$slot4 \$slot4 \$slot4 there is \$slot4 there is	0.9	0.01	1

Table 11: Example outputs of PPLM for controlling “*positive sentiment*” with various hyper-parameters. Increasing γ_{gm} and α encourages style control but can lead to ungrammatical outputs or run into degeneration. Increasing λ encourages sentence fluency and semantic accuracy; however, careful hyper-parameters fine-tuning is required to ensure reasonable output quality.

layer (logit vector) to achieve the above two goals. The logit vector o_t is obtained by a forward pass of the LM, *i.e.* $o_t = \text{LM}(x_t, H_t)$. The next generated token x_{t+1} is sampled as $x_{t+1} \sim \text{softmax}(W o_t)$ where W is a linear transformation that maps o_t to a probability distribution of vocabulary size.

The shift ΔH_t is then updated with gradient descent as follows:

$$\Delta H_t \leftarrow \Delta H_t - \alpha \left(\nabla_{\Delta H_t} L_{\text{CE}}(\cdot) + \lambda \nabla_{\Delta H_t} D_{\text{KL}}(\cdot \| \cdot) \right) \quad (8)$$

where α is the step size and the update can be repeated until it converges or up to a certain time step to obtain the shift ΔH_t . At generation time, a *post-norm fusion* that samples the next token x_{t+1} based on the shifted distribution \tilde{p}_{t+1} and the unmodified distribution p_{t+1} is done to increase the language fluency:

$$x_{t+1} \sim \frac{1}{\beta} \left(\tilde{p}_{t+1}^{\gamma_{gm}} p_{t+1}^{1-\gamma_{gm}} \right) \quad (9)$$

β is a normalizing factor such that it forms a valid distribution. As $\gamma_{gm} \rightarrow 1$ this converges to the distribution from the shifted LM, and as $\gamma_{gm} \rightarrow 0$ it converges to the unmodified unconditional LM distribution. We observe that the hyper-parameters $\lambda, \alpha, \gamma_{gm}$ can be tuned to affect the generation. In practice, we found that increasing γ_{gm} led to nonsensical results very quickly as the generated text is no longer conformed to the given meaning representation. A larger step size α also led to unstable

and nonsensical results compared to a smaller step size as it moves H_t further away from its original position. In practice, a larger λ helps keep the fluency and the semantic of the generated text. Nevertheless, its effect is less influential than α and γ_{gm} . Table 11 illustrates the effect of hyper-parameters on the model outputs.

C Additional Experimental Results

We include additional example outputs for lexical styles in Table 12.

Multiple Styles Control In this section, we demonstrate the possibility of controlling multiple styles with the methods described in Section 5. To control multiple styles with conditional training, we condition the CT model on multiple control variables $P(x|a_1, \dots, a_n)$. The control variables are concatenated and then added to the input as special tokens after the [BOS] token. To control multiple styles with PPLM, we simply add the cross-entropy loss of each semantic style classifier to the objective function. Extra hyper-parameters can be introduced to control the significance of each style. Similarly, for (beam search) weighted decoding, the distribution is re-weighted by

$$p(w|a_1, \dots, a_n) = \text{Softmax} \left(p(w) \prod_{i=1}^n \lambda_i p(a_i|w) \right) \quad (10)$$

where λ_i are hyper-parameters that allow us to determine the significance of each style a_i . Finally,

MR: OFFER (address=\$slot1), OFFER (rating=\$slot2)	
W/o Style Control	There is a nice house at \$slot1 with a \$slot2 rating.
Short	\$slot1 is rated \$slot2.
Long	There is a house at \$slot1 with a rating of \$slot2. Would you like that one?
Has Rare Word	There is a lovely residence located at \$slot1. It has a rating of \$slot2.
First Person Pron.	I found a house at \$slot1 with a \$slot2 rating.
Second Person Pron.	There is a house at \$slot1 with a \$slot2 rating. Would you like that one?
Descriptive	There is a nice house with a \$slot2 rating available at \$slot1.
MR: OFFER_INTENT (intent=intent)	
W/o Style Control	Would you like to make a reservation for this house?
Short	Do you want to reserve it?
Long	Do you want me to go ahead and reserve a room at the hotel?
Has Rare Word	Would you like to continue with the hotel reservation?
First Person Pron.	Shall I book the rooms in that hotel now?
Second Person Pron.	Would you like to make a reservation for this house?
Descriptive	Would you like to reserve a room at this hotel?

Table 12: Example outputs using conditional training for the six lexical styles described in Section 3.

conditional training and (beam search) weighted decoding can be used simultaneously by training a CT model and then applying WD or BSWD on the trained CT model during decoding time. Example outputs and the style accuracy of multiple style control using CT and BSWD are shown in Tables 13 and 14.

Lexical Style	Lexical Sty. Acc.	Semantic Style	Semantic Sty. Acc.
Short	55.4%	Formal	2.0%
Long	95.2%	Formal	39.1%
Short	69.2%	Positive	20.3%
Long	84.4%	Positive	44.8%

Table 13: Style accuracy of multiple-style control using CT and BSWD with beam width $B = 2$. The parameter *length* is controlled by CT and the parameters *formality* and *sentiment* are controlled by the BSWD.

D Human Evaluation Design

The goal of the annotation task is to evaluate the style of the automatically generated templates that express specific information to the user. The templates follow a certain style value such as “*formal*”. The annotators are asked to imagine that they are having a conversation with the system, and that the template presented to them is an example of something that the system may say to them. The template may be in the form of a question, statement, confirmation, etc. Table 15 illustrates two sample MRs and four generated templates, which we lexicalize with a map of values (one per slot) to make the annotation task more intuitive.

MR: OFFER(restaurant_name=\$slot), OFFER(city=\$slot2)	Lexical Sty.	Semantic Sty.
\$slot1 in \$slot2 is a very popular indian restaurant. \$slot1 is a restaurant in \$slot2 that serves pizza and pasta.	W/o Style Control	W/o Style Control
\$slot1 is a restaurant in \$slot2. \$slot1 serves italian food in \$slot2.	Short	W/o Style Control
Okay! The \$slot1 located in \$slot2, as for the tofu, it has a pretty good rating. Would you like to try for a lunch or dinner reservation? Okay! The \$slot1 located in \$slot2, as for the tofu, it has a pretty good rating. Would you like to try for a meal there?	Long	W/o Style Control
\$slot1 is a nice restaurant located in \$slot2. \$slot1 might interest you, and it is located in \$slot2.	W/o Style Control	Positive
\$slot1 is a good restaurant in \$slot2. \$slot1 restaurant is also in \$slot2.	Short	Positive
Okay! The restaurant \$slot1 located right inside \$slot2 is a good one and it has many vegetarian side dishes. Okay! The restaurant \$slot1 located right inside \$slot2 is a good one. Do you wish to have your lunch here?	Long	Positive
MR: OFFER(address=\$slot1), OFFER(rating=\$slot2)	Lexical Sty.	Semantic Sty.
There is a \$slot2 rated house at \$slot1. There is a house at \$slot1 with a \$slot2 rating.	W/o Style Control	W/o Style Control
\$slot1 has a \$slot2 rating \$slot1 with a \$slot2 rating	Short	W/o Style Control
There is a house at \$slot1 with a \$slot2 rating. would you like to stay here? There is a house at \$slot1 that you might be interested in. It has a \$slot2 average rating.	Long	W/o Style Control
There is a lovely house at \$slot1 with a rating of \$slot2 There is a good house at \$slot1 with a rating of \$slot2	W/o Style Control	Positive
\$slot1 has a rating of \$slot2 \$slot1 with rating of \$slot2	Short	Positive
There is a lovely residence located at \$slot1 with a rating of \$slot2 There is a nice house located at \$slot1 with a rating of \$slot2	Long	Positive

Table 14: Example outputs by combining CT and BSWD with beam width $B = 2$. Lexical style is controlled by CT and semantic style is controlled by the BSWD. Note how the lexical style "long" tends to yield outputs that include more hallucinated content in an attempt to fulfill the required style goal(s).

Domain	Simplified MR	Output template
Movies	<i>request city none</i>	What is the location and type of movie you are interested in?
	<i>request movie_name none</i>	What movie do you want to see and in what location?
Restaurants	<i>confirm restaurant_name [Ala Romana]</i>	Please confirm: booking a table at [Ala Romana] on [March 1st].
	<i>confirm date [March 1st]</i>	Please confirm the following: you want a table at [Ala Romana] on [March 1st].

Table 15: Example human evaluation worksheet.