

# Generating Diverse Story Continuations with Controllable Semantics

Lifu Tu<sup>1</sup> Xiaoan Ding<sup>2</sup> Dong Yu<sup>3</sup> Kevin Gimpel<sup>1</sup>

<sup>1</sup>Toyota Technological Institute at Chicago, Chicago, IL, 60637, USA

<sup>2</sup>University of Chicago, Chicago, IL, 60637, USA

<sup>3</sup>Tencent AI Lab, Bellevue, WA, 98004, USA

{lifu, kgimpel}@ttic.edu, xiaoanding@uchicago.edu, dyu@tencent.com

## Abstract

We propose a simple and effective modeling framework for controlled generation of multiple, diverse outputs. We focus on the setting of generating the next sentence of a story given its context. As controllable dimensions, we consider several sentence attributes, including sentiment, length, predicates, frames, and automatically-induced clusters. Our empirical results demonstrate: (1) our framework is accurate in terms of generating outputs that match the target control values; (2) our model yields increased maximum metric scores compared to standard  $n$ -best list generation via beam search; (3) controlling generation with semantic frames leads to a stronger combination of diversity and quality than other control variables as measured by automatic metrics. We also conduct a human evaluation to assess the utility of providing multiple suggestions for creative writing, demonstrating promising results for the potential of controllable, diverse generation in a collaborative writing system.

## 1 Introduction

We consider the problem of automatic story continuation generation, i.e., how to generate story continuations conditioned on the story context. Inspired by recent work in controllable generation (Hu et al., 2017; Ficer and Goldberg, 2017), we propose a simple and effective modeling framework for controlled generation of multiple, diverse outputs based on interpretable control variables. Each control variable corresponds to an attribute of a sentence. Compared to previous work that only seeks to control the values of sentiment (Hu et al., 2017) and length (Kikuchi et al., 2016), we further explore neural text generation with particular verbal predicates, semantic frames, and automatically-induced clusters.

We compare the diversity of story continuations controlled by different sentence attributes and find

<b>Context:</b>	
sandra needed a new phone . her old one had broken . she walked into the store and bought a new one . she was very excited .	
<b>Control Attributes and Generated Continuations:</b>	
<b>Sentiment:</b>	
positive	sandra was happy to have a new phone .
<b>Predicates:</b>	
loved	sandra loved the new one .
gave	sandra 's mom gave sandra a refund .
<b>Frames:</b>	
Calendric_unit	it was the perfect day of her life .
Cardinal_numbers	it was a perfect one .
Activity_ongoing	she kept it on the couch .

Table 1: Story continuations conditioned on various control attributes generated from our framework.

that using frames yields a stronger combination of diversity and quality than other control variables as measured by automatic metrics. Unlike certain other attributes, frames have hundreds of possible values. Some frame values can help to get a natural continuation, while others are not applicable when considering the story context.

We quantitatively evaluate both controllability and diversity. Our empirical results show that: (1) our framework is accurate in terms of generating outputs that match the target attribute values; (2) our framework increases maximum metrics scores compared to  $n$ -best list generation with beam search; (3) controlling generation with frames leads to a stronger combination of diversity and quality than other control variables as measured by automatic metrics.

We also explore methods to rerank continuations to choose attribute values automatically and obtain a small number of high-quality outputs. We consider two frame ranking methods: one reranks the generated continuations using a reverse scoring model (Li et al., 2016) and returns the  $k$ -best generations; the second first predicts the  $k$  most likely frames based on the context, and uses these

frames to generate continuations. One potential use case of controllable, diverse story generation is collaborative writing applications (Clark et al., 2018b). We conduct a human evaluation to assess the utility of providing multiple suggestions from our models in this setting, demonstrating promising results for the potential of controllable generation for collaborative writing.

## 2 Task Description and Definitions

Given a story context and a control attribute value, our goal is to generate a story continuation that: (1) conforms to the given attribute value, (2) is relevant to the story context, and (3) is complementary to continuations with other control attribute values, thereby providing a diverse set of continuations when used with multiple attribute values.

We use  $x = \langle x_1, x_2, \dots, x_{|x|} \rangle$  to denote a story context and  $y = \langle y_1, y_2, \dots, y_{|y|} \rangle$  to denote a story continuation. The last token  $y_{|y|}$  is assumed to be  $\langle eos \rangle$ . We develop a framework to model tuples  $(x, l, y)$ , where  $l$  is a control attribute. The control attribute represents a characteristic of the continuation, such as its length, sentiment, automatically-induced cluster, verbal predicate, or set of semantic frames. Table 1 shows several examples of control attributes and generated continuations corresponding to them from our model.

## 3 Model

Our controllable generation framework is a variation of a sequence-to-sequence (seq2seq) model with attention (Sutskever et al., 2014; Bahdanau et al., 2015; Luong et al., 2015). To represent the control attribute values, we define an attribute embedding function  $R$  that maps a given attribute value  $l$  to a vector  $z$ :  $z = R(l)$ . Here  $l$  can be a single discrete number or a set of discrete numbers, depending on what attributes are being used. The control variable  $z$  contains two parts:  $z = [z_{\text{enc}}; z_{\text{dec}}]$  where semicolon (;) denotes vertical concatenation and  $z_{\text{enc}}$  and  $z_{\text{dec}}$  are additional inputs for the encoder and decoder respectively.

**Encoder.** For our encoder, we use a standard bidirectional recurrent neural network (RNN):

$$\begin{aligned} \vec{s}_i &= f_{e1}([v_i; z_{\text{enc}}], \vec{s}_{i-1}) \\ \overleftarrow{s}_i &= f_{e2}([v_i; z_{\text{enc}}], \overleftarrow{s}_{i+1}) \\ s_i &= [\vec{s}_i; \overleftarrow{s}_i] \end{aligned}$$

where  $v_i$  is the vector representation of word  $x_i$ ,  $s_i \in \mathbb{R}^d$  is the hidden state at time  $i$ , and  $f_{e1}$  and  $f_{e2}$  are the forward and backward RNN functions.

**Decoder.** Our decoder uses an RNN with the general global attention scheme from Luong et al. (2015). An additional input  $z_{\text{dec}}$  is fed to the decoder at each time step to reflect the characteristics of the control variable:

$$h_j = f_d([y_{j-1}; z_{\text{dec}}], h_{j-1})$$

where  $h_j$  is the decoder RNN hidden vector at time step  $j$  and  $f_d$  is the decoder RNN function. Then, the conditional probability with controllable generation can be decomposed as follows:

$$\log p_{\Theta}(y | x, l) = \sum_{j=1}^{|y|} \log p_{\Theta}(y_j | y_{<j}, s, l)$$

Here  $s$  represents the hidden states of the source sequence and  $\Theta$  are the parameters of the seq2seq attention model.

**Training.** Our training objective is:

$$\min_{\Theta, R} \sum_{(x, l, y) \in \mathcal{D}} -\log p_{\Theta}(y | x, l) \quad (1)$$

where  $\mathcal{D}$  is the training data, i.e., we assume attribute values  $l$  are provided during training. In practice, these will be predicted automatically using a linguistic analyzer. With certain attributes, we do not update the attribute value embeddings  $R$ , i.e., we fix  $z_{\text{enc}}$  and  $z_{\text{dec}}$  to one-hot vectors.

**Inference.** We can specify the value of the control variable and generate specific outputs. By changing the variable values, we obtain multiple continuations for the same context. Beam search is used for decoding.

## 4 Control Attributes

In this section, we describe the control attributes we explored using our framework. Table 2 shows examples of generated continuations for a single story context with several values for the control attributes described below. Given our simple modeling framework, it would be natural to experiment with combining control attributes via summation or averaging of the attribute representations, but we leave an investigation of this to future work, focusing here on using one attribute at a time.

<b>Context:</b> i bought a pair of earbuds at target . i spent ten dollars . someone told me they were cheaper at the dollar store . they were only a dollar .	
<b>Gold Continuation:</b> i wish i had known that before .	
<b>Sentiment</b>	
<b>negative</b>	<b>i was disappointed .</b>
neutral	i decided to keep them .
positive	i was able to get a new pair .
<b>Length</b>	
4	i was shocked .
5	i was rather disappointed .
6	i bought a new pair .
7	i was able to buy them .
<b>8</b>	<b>i was glad i bought them .</b>
9	i was able to buy a new pair .
10	i was able to buy a new pair .
11	i was able to buy a new pair of shoes .
12	i was able to buy a new pair of new ones .
13	i was able to buy a new pair of rubber flavored items .
14	i was able to buy a new pair and i was very happy .
<b>Verbal Predicates</b>	
<b>wish, known</b>	<b>i wish i will always recognize them .</b>
got	i got a new one .
decided	i never decided on the new restaurants .
went	i went home with a new friend .
is	it is a great price .
get	now i have to get a new one .
felt	i felt very disappointed .
go	i will go to the grocery store .
took	i took them to the store .
left	i left the store with a new tip .
realized	after many years , i realized that .
loved	i loved them .
ran	i ran back to the store .
<b>Frame Semantics</b>	
<b>gold frame set*</b>	<b>i wanted to be a professional photographer .</b>
Arriving	i got the earbuds .
Quantity	it was a lot of fun .
Becoming	it ended up being a target .
Cardinal_numbers	i paid for \$ 50 dollars .
Being_obligated	i guess i had a similar card .
Kinship	my parents were proud .
Statement	i told them i would not be a target .
Causation	they sent me to the seller 's desk .
Opinion	i think i was a millionaire .
Perception_experience	it was a hit .
<b>Clusters</b>	
0	i ended up buying a new one .
1	i bought a new pair of shoes .
<b>2</b>	<b>i was a good price .</b>
3	i bought a new pair of shoes for free .
4	i decided to buy a new pair of headphones .
<b>Oracle BOW</b>	
-	i then wish that i had no time .

Table 2: Generated continuations from our framework with different control attribute values. Boldface indicates attribute values of the human-written continuation. \* = the frame set in the human-written continuation contains the following frames: *Desiring, Experienter, Event, Being\_named, and Entity*.

**Sentiment.** Stories may express sentiment regarding their characters or circumstances. We acquire sentiment labels by running the pretrained

analyzer from Socher et al. (2013) on the continuations in the training data. The analyzer produces three labels: “negative”, “neutral”, or “positive”. During training,  $z_{enc}$  and  $z_{dec}$  are fixed one-hot vectors for each value.

**Length.** Some prior work has generated summaries with a desired length (Kikuchi et al., 2016; Fan et al., 2018a). We similarly use length of the continuation as a control attribute. Instead of using an embedding for each integer length value, we group the lengths into a small number of bins (details are provided below).  $z_{enc}$  and  $z_{dec}$  are fixed one-hot vectors for each bin.

**Verbal Predicates.** Semantic role labeling (SRL) is a form of shallow semantic parsing that annotates predicates and their arguments in sentences. We consider predicates from a semantic role labeling as control attributes. We use the SRL system from AllenNLP (Gardner et al., 2018) to automatically obtain predicates for the continuations in our training set. Then, a predicate vector is obtained by first summing up 100-dimensional GloVe embeddings (Pennington et al., 2014) of the predicted predicates (if there is more than one), then reducing the dimension to 64 using principal component analysis.<sup>1</sup> We wish to clarify that we do not use the argument structure from the SRL system. We restrict our focus to simply the set of verbal predicates in the SRL structure; this would presumably be simpler to use in interactive settings where users would specify attribute values for generating continuations.

**Frame Semantics.** A story is composed of a sequence of meaningful events (Chatman, 1980), often following particular patterns described in various terms such as scripts (Schank and Abelson, 1977) and frames. FrameNet (Baker et al., 1998) is an inventory of semantic frames, which are semantic abstractions describing universal categories of events, concepts, and relationships.

We consider frame semantics as another control attribute in our framework. In order to get a frame semantic representation for a continuation, we use SEMAFOR (Das et al., 2014). SEMAFOR automatically produces a frame-semantic parse for a sentence, which consists of spans that evoke particular frames in FrameNet as well as annotations

<sup>1</sup>The reason we use PCA here is to make all attribute embeddings have comparable embedding size, though we did not systematically evaluate the effect of this choice.

of textual spans that correspond to frame-specific arguments. For our purposes, we drop the arguments and only use the set containing all frames that are evoked in the sentence. A sentence may contain multiple frames. For example, in the sentence “Roa’s advice made Emma a lot happier in her life!”, “a lot” evokes the *Quantity* frame while “Emma a lot happier” evokes the *Effect* frame.

The frame set variable  $z$  is computed by summing embeddings for the frames in the set:

$$z = R(l) = \sum_{j \in l} R_j \quad (2)$$

where  $l$  is the frame set and  $R_j$  is the representation of frame  $j$ . The frame embeddings are learned during training.<sup>2</sup> For modeling purposes, we restrict our attention to the 100 most frequent frames in the training data. The rest of the frames are pooled together to form a single additional “catch-all” frame.

**Automatically-Induced Clusters.** We also experiment with running  $k$ -means clustering on the bag-of-words sentence representations of the continuations in the training set. We treat these automatically-induced cluster labels as control attribute values. Below we describe experiments with different cluster labels and analyze the characteristics of the generated outputs.

**Oracle Bag-of-Words Sentence Representations.** We also consider the use of a bag-of-words (BOW) sentence representation as a control attribute. Naturally, the sentence representation of the continuation is not available before generating the continuation in practice. However, we can use this attribute to verify the capability of our model to reconstruct the continuation from its bag-of-words representation.

## 5 Experimental Setup

### 5.1 Datasets

We experiment with the publicly available ROC story corpus developed by Mostafazadeh et al. (2016). It consists of approximately 100K five-sentence stories of everyday events. We sample 2000 stories as a development set and 2000 as our test set. The remaining stories form our training set. Our goal is to generate the fifth sentence (the

<sup>2</sup>The dimension of the frame vector is 64 in our experiments.

“continuation”) given the previous four sentences. We use the 10k most frequent words in the training set as our vocabulary. A special token  $\langle unk \rangle$  is introduced for unknown words.

### 5.2 Evaluation

Previous work evaluates generation tasks with automatic metrics, such as perplexity (PPL), BLEU (Papineni et al., 2002),<sup>3</sup> and ROUGE (Lin, 2004). We adopt these in our evaluation and add three more metrics using the pretrained story scorer from Sagarkar et al. (2018). The scorer rates a generated continuation given its context along three dimensions: **relevance (R)**, **interest- ingness (I)**, and **overall quality (O)**. The story scorer does not use a gold standard continuation.

In addition, to evaluate the diversity of the generation, we use **Max-BLEU**<sup>4</sup> and **Max-ROUGE**. First, we compute BLEU and ROUGE scores over a set of outputs  $(y_1, y_2, \dots, y_n)$  with different attribute values given the same story context, then we compute the max scores:

$$\text{Max-BLEU} = \max_i \text{BLEU}(y_i, r)$$

$$\text{Max-ROUGE} = \max_i \text{ROUGE}(y_i, r)$$

where  $r$  is the gold standard continuation.

We also use **Self-BLEU** (Zhu et al., 2018) to evaluate the diversity of a set of outputs. It is calculated by averaging the BLEU scores computed between all pairs of generated continuations for a given context, then averaging this quantity over all contexts. The smaller the Self-BLEU score is, the more diverse are the generated outputs.

### 5.3 Training Details

Our seq2seq model has a 2-layer biLSTM (Hochreiter and Schmidhuber, 1997) encoder and a 1-layer LSTM decoder. The hidden dimension of all layers is 512. The word embedding dimension is also 512. For optimization, we use Adam (Kingma and Ba, 2014) with learning rate 0.001. We use early stopping based on perplexity on the development set.

<sup>3</sup>In this paper, all BLEU scores are BLEU-2 scores (i.e., using unigrams and bigrams).

<sup>4</sup>While max metric scores are not solely measuring diversity, they do provide a sense of the potential of the list. If all entries on the list are the same, the max metric scores would equal the average metric scores. The difference between the max and average metric scores therefore can be viewed as providing a bound on the diversity.



Target Sentiment	Generated Continuations		
	negative	neutral	positive
negative	<b>68.5</b>	19.6	12.0
neutral	7.0	<b>73.9</b>	19.2
positive	0.8	3.1	<b>96.2</b>

Table 3: Sentiment match percentages of generated continuations and target sentiment values.

	dif = 0	dif ≤ 1	dif ≤ 2	dif ≤ 3
3 bins	95.8	100	-	-
30 bins	70.35	94.8	99.25	99.9

Table 4: Frequency (%) of the generated continuations in the range of  $\text{dif} = |l - l_p|$  where  $l$  is the continuation length and  $l_p$  is the desired length.

## 6 Results

We now present our experimental results. Section 6.1 includes results related to how well our generated output matches the desired attribute values. Section 6.2 presents results when generating continuations with oracle attribute values. In Section 6.3 we use our set-level metrics to evaluate sets of outputs with various attribute values. In Section 6.4 we report results when attempting to automatically infer attribute values to generate a small set of high-quality outputs.

### 6.1 Controllability Evaluation

In this section, we evaluate the controllability accuracy of our framework by automatically measuring the match between the attribute values of the generated continuations and the desired values. For certain control variables, like sentiment and frames, this automatic evaluation is prone to errors in the associated analyzers. That is, the metrics that rely on automatic analyzers could become artificially high if our generation models learn to produce outputs that match the biases of the analyzers. We could instead consider manual evaluation of control accuracy. However, we were more interested in devoting our manual evaluation to the question of whether users would find the system outputs useful for a particular goal.

**Sentiment.** We generate three continuations for each story in the development set, one for each sentiment label. Using the same sentiment analyzer from Socher et al. (2013) as above, we obtain predicted sentiment labels for the continuations. Table 3 shows the sentiment distribution for each label. We see that the vast majority of the time, the continuations match the desired val-

Predicate	M%	Predicate	M%
was	100	got	100
had	100	decided	94.4
went	99.9	is	100
made	99.25	were	100
found	100	get	99.95
felt	99.55	go	100
took	99.2	ended	98.25
be	99.95	told	99.9
gave	99.95	left	99.85
said	100	bought	100

Table 5: Match percentages (M%) showing fraction of stories for which generated continuations contain the desired predicate. The 20 most frequent predicates are shown; additional results are in the Appendix.

ues. Matching positive sentiment is easiest for our model, followed by neutral.

**Length.** We quantize the generation lengths into bins, each representing a size range. Below are the two settings we consider:

- 3 bins: We use three bins with the following length ranges: [1,7], [8,13], and [14, ∞).
- 30 bins: We use a bin for each length. No sentence is longer than 30.

During training, we do not update the representations of the length control variable. After training, we treat the length of the continuation in the development set as the target control variable and generate continuations for each length. The results are shown in Table 4 and demonstrate that our model can generate continuations with the desired length with only small differences.

**Verbal Predicates.** We select the top 100 most frequent verbal predicates in the training data. Then for all the stories in the development set, we generate a continuation for each of the 100 predicates. We check whether the predicate appears in the generated continuations. As the results in Table 5 show, the framework can nearly always generate outputs with the desired predicates.

**Frame Semantics.** In order to check how frequently the generated output matches the desired frames, we generate continuations for the top 100 frames (one frame for each continuation) for all stories in the development set. We check whether the frame appears in the specific continuation using SEMAFOR. The results are shown in Table 6. Most frames have very high match accuracies, but there are a few frames with much

Frame	M%	Frame	M%
Calendric_unit	99.5	Locative_relation	87.4
Arriving	98.9	Quantity	92.5
Temporal_collocation	99.9	Becoming	99.6
Cardinal_numbers	89.1	Being_obligated	97.7
Kinship	94.4	Intentionally_act	99.7
Statement	98.0	Causation	98.6
Emotion_directed	98.7	Buildings	93.0
Personal_relationship	92.7	Food	79.2
Self_motion	86.4	Capability	99.9
Desirability	98.1	Observable_body_parts	74.2

Table 6: Match percentages (M%) showing fraction of stories for which generated continuations contain the desired frame. Additional results are in the Appendix.

Target Value	Generated Continuations				
	0	1	2	3	4
0	<b>79.9</b>	2.8	3.1	0.9	13.4
1	5.1	<b>63.1</b>	26.4	1.4	4.2
2	2.6	2.0	<b>90.6</b>	0.3	4.7
3	20.9	20.1	24.6	<b>31.0</b>	3.5
4	0.9	0.3	0.5	0.1	<b>98.3</b>

Table 7: Cluster match percentages (%) for each value of the cluster control variable.

lower accuracy, such as “Food” and “Observable\_body\_parts”. These are more concrete frames that may be difficult to reasonably incorporate in certain story contexts.

**Automatically-Induced Clusters.** Given the cluster, the model generates a continuation. Then, we represent the continuation as a bag-of-words sentence embedding (using the same method as when performing the initial clustering) and find the cluster that has the nearest cluster embedding. Then we check whether the two clusters match.

In analyzing the clusters, we observed that cluster 0 corresponds to simple but reasonable continuations. Cluster 2 corresponds to continuations with positive sentiment. Cluster 4 contains continuations with more actions. Some of the generated outputs are shown in Table 2. From the results in Table 7, we still see controllability for most clusters; however, for target cluster 3, which is rather generic based on our observations, the generated output seems flat.

## 6.2 Evaluation with Oracle Attributes

Table 8 shows automatic metric scores with oracle attribute values, i.e., using the attribute values of the gold standard continuations. Unsurprisingly, compared with the seq2seq baseline, the perplexity decreases and the ROUGE and BLEU scores increase with oracle attributes. We also find that

the scores from the story scorer, which does not use the gold standard while scoring, also show improvements over the baseline. We note that frame semantics form one of the most useful control attributes, aside from those that use words directly.

The oracle BOW representation of the gold standard continuation yields the lowest perplexity and highest ROUGE and BLEU scores. It is not surprising that using this attribute would be useful according to metrics that favor matching the gold standard. However, these results do show that our simple modeling framework can make use of the information in the control variable with a high degree of effectiveness. In addition, while the scores from the story scorer are generally higher than for other control attributes, they are roughly on par with those when using predicates and frames.

## 6.3 Evaluating Sets of Continuations

We now evaluate *sets* of continuations using our set-level metrics. Standard methods to generate sets of outputs include beam search (BS) and temperature-based sampling (TS), which we use as baselines. TS with temperature  $\tau$  corresponds to transforming probabilities  $p_i$  as follows:  $\hat{p}_i \propto p_i^{\frac{1}{\tau}}$ . A high temperature  $\tau$  leads to high variety in generated samples, but also more noise, while lower temperatures lead to samples with less noise but also less diversity.

For each attribute, we generate continuations for each of its values, and compare to BS and TS systems with the same number of outputs. For example, for sentiment, we generate continuations for each of the 3 sentiment values and compare to BS and TS with 3 continuations.

Results are shown in Table 9. BS shows the least diversity (as evidenced by its high self-BLEU scores). However, it generally yields high average ROUGE and BLEU scores. TS does very well in terms of diversity, and this diversity enables it to produce higher max scores than BS, but it has lower averages when using small numbers of continuations (3 or 5).

Our sentiment- and cluster-controlled systems outperform TS in max metric scores and BS in diversity (self-BLEU). They also have the highest average BLEU scores, though the differences are small. With 30 continuations, TS with  $\tau = 0.5$  performs best across all metrics; this number of continuations appears to be well-suited for temperature 0.5. As we move to 100 continuations, we

	PPL ( $\downarrow$ )	ROUGE		BLEU ( $\uparrow$ )	Story Scorer		
		ROUGE-1 ( $\uparrow$ )	ROUGE-L ( $\uparrow$ )		O ( $\uparrow$ )	R ( $\uparrow$ )	I ( $\uparrow$ )
seq2seq	25.8	27.0	23.5	17.7	5.5	5.5	4.8
sentiment	25.0	26.7	23.5	17.7	5.5	5.6	4.8
length	23.1	27.3	24.6	20.3	5.7	5.8	5.0
predicates	17.1	42.9	35.1	26.4	6.0	6.2	5.2
frames	15.0	41.1	35.0	27.2	5.9	6.1	5.2
clusters	24.3	28.6	25.0	18.4	5.5	6.1	5.1
BOW	5.7	64.5	54.5	45.4	6.2	6.2	5.2
gold standard	-	-	-	-	6.5	6.7	5.7

Table 8: Automatic metrics for baseline system and when using oracle values for control attributes. For the gold standard continuation, we report only the story scorer results because they do not require a gold standard (unlike the other metrics).

		ROUGE				BLEU ( $\uparrow$ )		Self-BLEU ( $\downarrow$ )
		ROUGE-1 ( $\uparrow$ )		ROUGE-L ( $\uparrow$ )		Max	(Avg)	
		Max	(Avg)	Max	(Avg)			
3 continuations:	BS, beam = 3	31.8	<b>(26.7)</b>	27.3	<b>(22.6)</b>	19.7	(17.0)	50.5
	TS, $\tau = 0.5$	32.5	(25.5)	27.8	(21.6)	20.3	(16.3)	27.0
	TS, $\tau = 0.6$	30.5	(22.2)	25.9	(18.8)	19.0	(14.8)	<b>23.8</b>
	sentiment	<b>32.8</b>	(25.6)	<b>28.9</b>	(22.5)	<b>21.1</b>	<b>(17.1)</b>	30.7
	predicted frames	30.8	(22.5)	27.0	(19.9)	19.7	(15.7)	30.3
	frames + reranking	30.7	(21.9)	26.3	(18.9)	18.8	(14.7)	25.8
5 continuations:	BS, beam = 5	33.9	<b>(26.3)</b>	29.4	<b>(22.3)</b>	21.3	(16.2)	68.1
	TS, $\tau = 0.5$	35.0	(25.3)	30.0	(21.5)	21.7	(16.2)	<b>40.8</b>
	clusters	<b>36.1</b>	(24.5)	<b>31.7</b>	(21.4)	<b>22.9</b>	<b>(16.4)</b>	43.8
30 continuations:	BS, beam = 30	40.0	<b>(25.4)</b>	34.3	(20.8)	25.6	(16.0)	92.5
	TS, $\tau = 0.5$	<b>43.0</b>	<b>(25.4)</b>	<b>37.5</b>	<b>(21.6)</b>	<b>28.1</b>	<b>(16.3)</b>	<b>74.0</b>
	length	42.1	(24.7)	35.9	(20.0)	26.2	(14.8)	82.2
100 continuations:	BS, beam = 100	44.4	(25.0)	38.6	(20.6)	29.2	(15.9)	96.2
	TS, $\tau = 0.5$	47.8	<b>(25.4)</b>	42.3	<b>(21.6)</b>	<b>32.1</b>	<b>(16.3)</b>	85.6
	frames (individual)	47.0	(24.0)	41.2	(20.8)	29.8	(15.5)	<b>72.1</b>
	frames (sets)	<b>48.3</b>	(23.1)	<b>42.7</b>	(20.1)	31.2	(15.1)	75.5

Table 9: Metric scores to evaluate the potential of a list of continuations. We report the maximum and average metric scores over the continuations in each list to evaluate the quality of the lists, and self-BLEU to evaluate diversity. Best results for each metric and each number of outputs are in bold.

find that using our frame control variable leads to better diversity than TS, suggesting that the move to 100 samples has introduced some amount of repetition. By contrast, the 100 distinct frames and frame sets yield better diversity.

#### 6.4 Automatically Choosing Attribute Values

Using our framework, we can generate continuations with any attribute values. However, if we are interested in generating a single continuation, we do not know the ideal attribute values to use. So, we propose two methods to automatically select a small set of values for the frame attribute.

**Frames + Reranking:** Following Li et al. (2016), we rerank the outputs from the 100 most frequent frame sets by linearly combining the forward score  $p(y | x)$  and the “reverse score”  $\lambda p(x | y)$ , where the latter comes from a separately-trained seq2seq model. The forward score  $p(y | x)$  is adjusted by dividing by the length of  $y$  in order

to not favor shorter outputs.

**Predicted Frames:** We also build a model to automatically predict the frames in the continuation. Given the frames in a sentence  $x$ , we compute a binary frame vector  $f_x$  where entry  $j$  is 1 if frame  $j$  appears in  $x$ . We train a model that predicts the frame vector of the continuation given the frame vectors of the previous 4 sentences. The model is an LSTM followed by averaging of hidden states. Mean squared error is minimized during training. After training, the  $k$  continuations are selected based on the  $k$  frames with the highest predicted score under this frame prediction model.

We use these two methods to produce 3 continuations for each story and report results in Table 9. They both achieve a similar balance of quality and diversity as TS with  $\tau = 0.6$ , with reranking leading to greater diversity than frame prediction and the latter showing higher ROUGE/BLEU scores.

## 7 Human Evaluation

Our previous results demonstrate that our frame control system has strong controllability and diversity in generation. In this section, we conduct a human evaluation to assess the utility of providing multiple suggestions from our models in a creative writing setting. We consider four different systems: **BS** with beam size 3; **TS** with 3 continuations using  $\tau = 0.6$ , which we found to produce outputs with more diversity than 0.5; **reranking** the 100 most frequent frame sets and using the top 3; and using continuations from the top-3 **predicted** frames under our frame prediction model.<sup>5</sup>

To assess which set of generations from these four systems are most helpful in a collaborative writing setting, we collect annotations using Amazon Mechanical Turk. We randomly select 100 stories. For each story, we generate three outputs as a set of suggestions for each system, so there are 600 comparison pairs in total. We show workers two sets of outputs from different systems and ask them to select which suggestion is more helpful for writing the next line in the story. We also provide a choice of “neither one is helpful at all”. We ask them explicitly to imagine they are writing the next line of the given story (see the appendix for more details).

Table 10 shows the results.<sup>6</sup> We observe that workers prefer the BS baseline over TS, although TS yields higher diversity. This could be because the continuations from BS are shorter, simpler, and more fluent. In addition, we observe that workers prefer the outputs from the reranking system over BS more often than not. Although the predicted frame system yields more diverse outputs, workers still prefer BS, likely due to the difficulty in predicting frames. The reranking and predicted frame systems are both preferred to TS, though the gap is smaller with the predicted system. We also see that generating helpful suggestions is a difficult task: in many cases workers thought neither system was helpful, especially when given the outputs from BS/TS or TS/predicted.

One may ask why workers do not show a stronger preference for the more diverse sets of

<sup>5</sup>The BS and TS baselines do not use control variables.

<sup>6</sup>We remove results from 10-question sets where more than half of the questions were answered with the “neither” option, as we were concerned that these annotators did not fully understand the task or did not spend enough time studying the continuations. This occurred in roughly one third of question sets.

system 1	system 2	Human Preference		
		1	2	neither
BS	TS	43	16	29
BS	reranking	18	30	15
BS	predicted	38	29	19
TS	reranking	18	38	16
TS	predicted	18	27	34
reranking	predicted	27	24	21

Table 10: Human preferences when given three continuations from each pair of systems.

outputs. From our own preliminary annotations, we believe this is because diverse outputs tend to be longer and harder to understand, and also because greater diversity increases the chance of producing disfluent or nonsensical outputs. The BS outputs, by comparison, are sensible and mostly on-topic. Even if the suggestions are not creative, they may still help a worker to think about a new direction for the story to take. Nonsensical or disfluent suggestions, however, are rarely helpful.

## 8 Related Work

Automatic story generation has a long history, with early work based primarily on hand-written rules (Klein et al., 1973; Meehan, 1977; Dehn, 1981; Turner, 1993). Subsequent methods were based on planning from artificial intelligence (Theune et al., 2003; Oinonen et al., 2006; Riedl and Young, 2010) and, more recently, data-driven methods have been developed (McIntyre and Lapata, 2010; Elson, 2012; Daza et al., 2016; Roemmele and Gordon, 2015; Clark et al., 2018a; Martin et al., 2018; Fan et al., 2018b; Yao et al., 2019; Fan et al., 2019). In concurrent work, Gupta et al. (2019) also propose methods to generate more diverse and interesting story endings, albeit without control variables.

In stronger relevance to our work, Clark et al. (2018b) explore a creative writing setting with a machine in the loop, albeit with mixed results in terms of the quality of system suggestions. Predicting and controlling with frame values suggests a new way of interacting with collaborative writing systems, as long as frames can be communicated to users in ways they can easily understand. Recently, Clark et al. (2018a) proposed a neural text generation method that explicitly represents and tracks entities. In addition, event sequences (Chaturvedi et al., 2017; Liu et al., 2018) are important elements in narrative texts but under-explored for story generation.



These and related characteristics of creative writing could be incorporated into our framework as control attributes in future work.

The broader neural text generation community has also recently been interested in controllable text generation, i.e., generating text with specified characteristics reflected by control variables. In some previous work, the variables to be controlled are embedded into vectors which are then fed into models to reflect the characteristics of the variables. Kikuchi et al. (2016) and Fan et al. (2018a) developed methods for controllable summarization, for example permitting users to control the length of the generated summary. Related work has controlled style, topic, and sentiment polarity (Hu et al., 2017; Wang et al., 2017; Shen et al., 2017; Yang et al., 2018).

Despite the widespread usage of beam search for neural text generation, it has long been observed that its outputs are lacking in diversity. Several efforts have been made to provide diverse outputs for generation tasks, such as dialogue (Li et al., 2016) and machine translation (Devlin and Matsoukas, 2012; Gimpel et al., 2013; Li and Jurafsky, 2016). Diverse beam search (Vijayakumar et al., 2018) produces a list of outputs with a diversity-augmented objective. Ippolito et al. (2019) compare several methods for producing a set of diverse outputs from conditional language models. We leave a careful comparison to such algorithms to future work.

## 9 Conclusion and Future Work

We proposed a controllable framework that generates the next sentence of a story given its context. We experimented with a broad range of control attributes and demonstrated that our framework can accurately generate outputs that match the target values. Sets of outputs from our method show high diversity and high oracle metric scores. The human evaluation shows that the multiple suggestions from our model show promise for integration in a collaborative writing system. Future work could explore other control attributes as well as a compositional framework to control multiple attributes jointly.

## Acknowledgments

We would like to thank TTIC SL@TTIC researchers who participated in a preliminary annotation study.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of International Conference on Learning Representations*.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *COLING-ACL '98: Proceedings of the Conference*, pages 86–90, Montreal, Canada.
- Seymour Benjamin Chatman. 1980. *Story and discourse: Narrative structure in fiction and film*. Cornell University Press.
- Snigdha Chaturvedi, Haoruo Peng, and Dan Roth. 2017. [Story comprehension for predicting what happens next](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1603–1614. Association for Computational Linguistics.
- Elizabeth Clark, Yangfeng Ji, and Noah A. Smith. 2018a. [Neural text generation in stories using entity representations as context](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2250–2260. Association for Computational Linguistics.
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A. Smith. 2018b. Creative writing with a machine in the loop: Case studies on slogans and stories. In *IUI*, pages 329–340. ACM.
- Dipanjan Das, Desai Chen, Andr F. T. Martins, Nathan Schneider, and Noah A. Smith. 2014. Frame-semantic parsing. *Computational Linguistics*, 40:1:9–56.
- Angel Daza, Hiram Calvo, and Jesús Figueroa-Nazuno. 2016. Automatic text generation by learning from literary structures. In *Proceedings of the Fifth Workshop on Computational Linguistics for Literature*, pages 9–19.
- Natalie Dehn. 1981. Story generation after TALE-SPIN. In *IJCAI*, pages 16–18.
- Jacob Devlin and Spyros Matsoukas. 2012. [Trait-based hypothesis selection for machine translation](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 528–532. Association for Computational Linguistics.
- David K. Elson. 2012. *Modeling Narrative Discourse*. Ph.D. thesis, Columbia University.
- Angela Fan, David Grangier, and Michael Auli. 2018a. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine*

- Translation and Generation*, pages 45–54. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018b. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2019. [Strategies for structuring story generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2650–2660, Florence, Italy. 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.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [Allennlp: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6. Association for Computational Linguistics.
- Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. 2013. [A systematic exploration of diversity in machine translation](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111, Seattle, Washington, USA. Association for Computational Linguistics.
- Prakhar Gupta, Vinayshekhar Bannihatti Kumar, Mukul Bhutani, and Alan W Black. 2019. [Writer-Forcing: Generating more interesting story endings](#). In *Proceedings of the Second Workshop on Storytelling*, Florence, Italy. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proc. of ICML*.
- Daphne Ippolito, Reno Kriz, Joao Sedoc, Maria Kustikova, and Chris Callison-Burch. 2019. [Comparison of diverse decoding methods from conditional language models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3752–3762, Florence, Italy. Association for Computational Linguistics.
- 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.
- Diederik Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Sheldon Klein, John F. Aeschlimann, David F. Balsiger, Steven L. Converse, Claudine Court, Mark Foster, Robin Lao, John D. Oakley, and Joel Smith. 1973. [Automatic novel writing: A status report](#). Technical Report 186, University of Wisconsin-Madison.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li and Daniel Jurafsky. 2016. [Mutual information and diverse decoding improve neural machine translation](#). *CoRR*, abs/1601.00372.
- Chin-Yew Lin. 2004. [Rouge: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Fei Liu, Trevor Cohn, and Timothy Baldwin. 2018. [Narrative modeling with memory chains and semantic supervision](#). *CoRR*, abs/1805.06122.
- 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.
- Lara J Martin, Prithviraj Ammanabrolu, Xinyu Wang, William Hancock, Shruti Singh, Brent Harrison, and Mark O Riedl. 2018. [Event representations for automated story generation with deep neural nets](#). In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Neil McIntyre and Mirella Lapata. 2010. [Plot induction and evolutionary search for story generation](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1562–1572.
- James R. Meehan. 1977. [TALE-SPIN, an interactive program that writes stories](#). In *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–98.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende,

- Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- K.M. Oinonen, Mariet Theune, Antinus Nijholt, and J.R.R. Uijlings. 2006. *Designing a story database for use in automatic story generation*, Lecture Notes in Computer Science, pages 298–301. Springer Verlag.
- 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.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Mark O. Riedl and Robert Michael Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268.
- Melissa Roemmele and Andrew S. Gordon. 2015. Creative help: a story writing assistant. In *International Conference on Interactive Digital Storytelling*, pages 81–92. Springer.
- Manasvi Sagarkar, John Wieting, Lifu Tu, and Kevin Gimpel. 2018. Quality signals in generated stories. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics (\*SEM 2018)*, New Orleans, Louisiana.
- Roger C. Schank and Robert P. Abelson. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. L. Erlbaum, Hillsdale, NJ.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in Neural Information Processing Systems 30*, pages 6833–6844. Curran Associates, Inc.
- 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 treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Mariët Theune, Sander Faas, Anton Nijholt, and Dirk Heylen. 2003. The virtual storyteller: story creation by intelligent agents. In *Proceedings of the 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, pages 204–215. Springer.
- Scott R. Turner. 1993. *Minstrel: a computer model of creativity and storytelling*. Ph.D. thesis, University of California at Los Angeles.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David J. Crandall, and Dhruv Batra. 2018. Diverse beam search: Decoding diverse solutions from neural sequence models. In *Proc. of AAAI*.
- Di Wang, Nebojsa Jojic, Chris Brockett, and Eric Nyberg. 2017. [Steering output style and topic in neural response generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2140–2150, Copenhagen, Denmark. Association for Computational Linguistics.
- Zichao Yang, Zhiting Hu, Chris Dyer, Eric P Xing, and Taylor Berg-Kirkpatrick. 2018. [Unsupervised text style transfer using language models as discriminators](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7287–7298. Curran Associates, Inc.
- Lili Yao, Nanyun Peng, Weischedel Ralph, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling. In *Thirty-Third AAAI Conference on Artificial Intelligence*.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Taxygen: A benchmarking platform for text generation models. *SIGIR*.

## A Appendix

**Model.** Figure 1 shows our modeling framework.

**Additional Experimental Results.** Figure 2 shows the length distribution of the continuations in the training data. Figure 3, Table 11, and Table 12 show the frequency (%) of matching the target control variable for length, predicates, and frames, respectively.

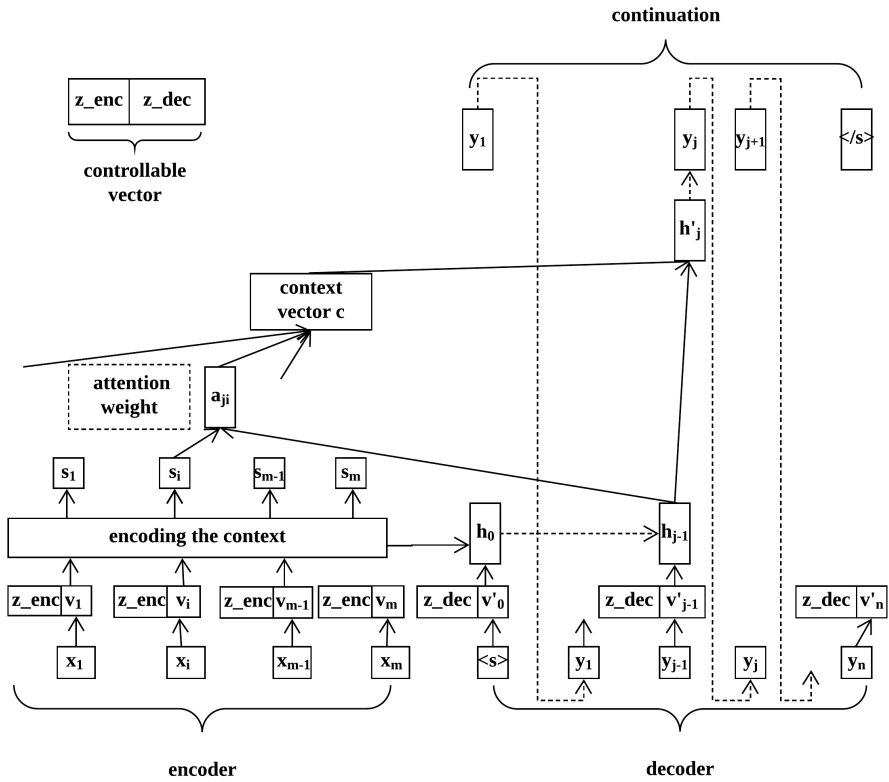


Figure 1: Our modeling framework for controllable generation.

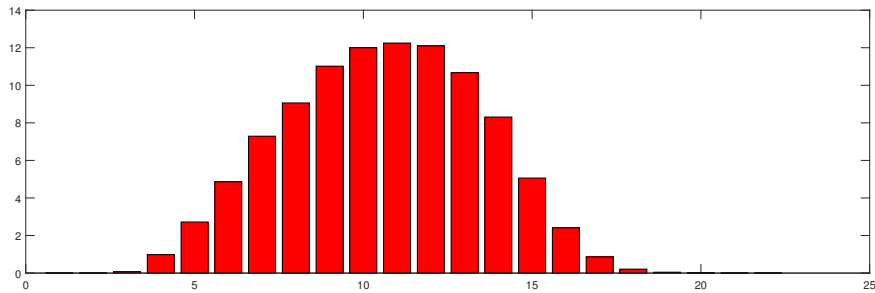


Figure 2: The length distribution of the continuations in the training data. The horizontal axis shows continuation lengths and the vertical axis shows the frequency (%) of the length.

**Human Evaluation Setup.** The form used on Amazon Mechanical Turk is shown in Figure 4.



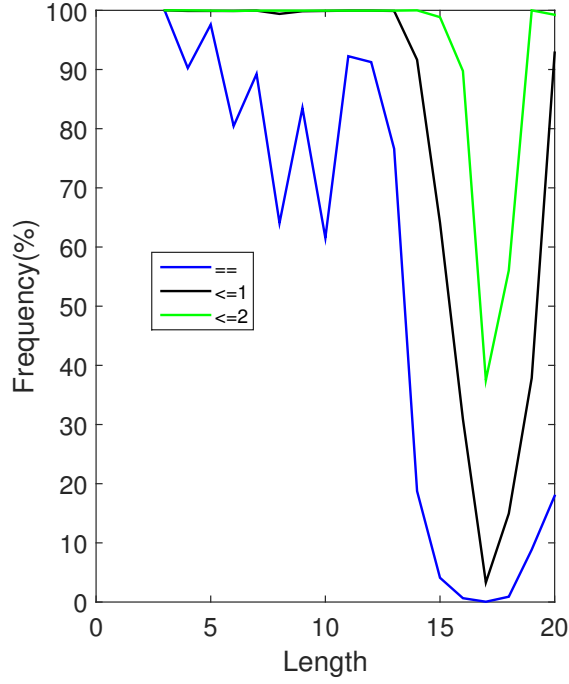


Figure 3: Frequency (%) of the generated continuations for different ranges of  $dif = |l - l_p|$ . The horizontal axis is the desired length. The vertical axis is the percentage of generated outputs with the given values of  $dif$ . Here  $l$  is the story continuation length and  $l_p$  is generated sentence length as the target control variable. “==” refers to  $l = l_p$ , “<=  $i$ ” refers to  $|l - l_p| \leq i$ .

predicate	acc	predicate	acc	predicate	acc	predicate	acc	predicate	acc
was	100	got	100	had	100	decided	94.4	went	99.9
is	100	made	99.25	were	100	found	100	get	99.95
felt	99.55	go	100	took	99.2	ended	98.25	be	99.95
told	99.9	gave	99.95	left	99.85	said	100	bought	100
came	98.7	realized	79.75	loved	99.8	have	99.1	won	100
became	99.75	put	83.6	ran	100	see	99.9	saw	98.75
started	99.85	buy	100	make	97.3	knew	99.6	looked	98.65
lost	99.05	enjoyed	99.95	called	98.5	learned	99.8	ate	95.35
turned	93.05	do	100	wait	99.85	take	99.8	fell	99.35
wanted	90.25	find	100	thought	80.05	stopped	98.65	play	100
has	98.15	eat	100	let	100	walked	99.9	did	96.55
being	99.05	done	94.05	returned	99.6	laughed	86.6	tried	99.6
's	99.25	asked	99.9	began	93.75	helped	97.65	caught	99.6
broke	99.9	getting	99.95	keep	97.15	paid	97.1	been	97
finished	99.9	threw	99.4	spent	99.85	passed	99.95	help	98.85
are	99.65	drove	99.25	work	81.25	going	35.9	leave	99.85
brought	98.4	feel	99.75	picked	73.65	agreed	98.35	needed	24.2
worded	100	used	98.75	pay	100	kept	98.8	liked	98.65
arrived	99.2	played	99.95	come	98.7	thanked	98.3	relieved	97.35
stop	99.15	playing	98.75	cried	100	died	97.75	know	89.05

Table 11: Frequency (%) of the generated continuations containing the desired predicate.

<b>frame</b>	<b>acc</b>	<b>frame</b>	<b>acc</b>	<b>frame</b>	<b>acc</b>	<b>frame</b>	<b>acc</b>
Calendric_unit	99.5	Locative_relation	87.4	Arriving	98.85	Quantity	92.5
Temporal_collocation	99.9	Becoming	99.6	Cardinal_numbers	89.05	Being_obligated	97.65
Kinship	94.35	Intentionally_act	99.65	Statement	98	Causation	98.55
Emotion_directed	98.65	Buildings	93	Personal_relationship	92.7	Food	79.2
Self_motion	86.4	Capability	99.95	Desirability	98.1	Observable_body_parts	74.15
Experiencer_focus	91.75	Time_vector	99	Request	99.35	Deciding	99.6
Relational_quantity	95.7	Measure_duration	96.7	Frequency	99.75	People	90.8
Ingestion	86.25	Vehicle	86.45	Age	98.55	Increment	98.6
Leadership	80.6	Desiring	90.7	Stimulus_focus	93.4	Activity_start	95.95
Education_teaching	96.5	Degree	100	Grasp	92.7	Commerce_buy	98.7
Scrutiny	93.5	Locale_by_use	65.15	Conquering	95.3	Giving	99.3
Clothing	49.05	Attempt	98.95	Becoming_aware	90.25	Building_subparts	73.7
Motion	96.55	Placing	73.15	Natural_features	26.7	Getting	85.05
Locating	76.55	Sufficiency	97.75	Feeling	98.7	Awareness	95.7
Experiencer_obj	87.4	Working_on	67.65	Aggregate	70.3	Performers_and_roles	87.95
Change_position_on_a_scale	72.45	Roadways	71.9	Containers	44.1	Coming_to_believe	92.4
Assistance	99.4	Ordinal_numbers	96.5	Relative_time	92.6	Choosing	96.55
Existence	88.15	Dimension	69.65	Cause_harm	89.75	Perception_active	87.1
Text	63.6	Cause_motion	81.2	Possession	76.85	Type	78.2
Body_movement	66	Opinion	87.7	Removing	74	Money	97.5
Have_as_requirement	25.4	Using	93.1	Storing	71.4	People_by_age	57.25
Contacting	73.35	Make_noise	97.9	Substance	50.4	Required_event	91.7
Political_locales	85.3	Difficulty	89.8	Activity_ongoing	94	Direction	98.6
Perception_experience	35	Impact	95.6	Locale	87.25	Waiting	96.65
Concessive	95.15	Partitive	63.65	Operate_vehicle	61.55	Size	98.3

Table 12: Frequency (%) of the generated continuations containing the desired frame.

Imagine you are writing the next line in the story below and are provided with suggestions from two **collaborative writing systems**. Which set of suggestions is more helpful?

Q298. tom was trying to get some water to drink. he didn't grab the glass properly. it slipped and fell to the flour. water went everywhere.

- 1: tom had to go to the hospital.
- 2: tom had to go to the hospital to get a new glass.
- 3: he had to go to the hospital to get it out.

- 1: tom was drenched in the water.
- 2: tom cleaned up the mess.
- 3: tom quickly walked away.

Neither one is helpful at all



Figure 4: The form for human evaluation of the generation systems for their potential in a collaborative writing setting.