

Go Back in Time: Generating Flashbacks in Stories with Event Temporal Prompts

Rujun Han¹ Hong Chen² Yufei Tian³ Nanyun Peng³

¹University of Southern California ²Tokyo University

³University of California, Los Angeles

rujunhan@usc.edu; chen@nlab.ci.i.u-tokyo.ac.jp

yufeit@ucla.edu; violetpeng@cs.ucla.edu

Abstract

Stories or narratives are comprised of a sequence of events. To compose interesting stories, professional writers often leverage a creative writing technique called *flashback* that inserts past events into current storylines as we commonly observe in novels and plays. However, it is challenging for machines to generate *flashbacks* as it requires solid understanding of event **temporal order** (e.g. *feeling hungry* *<before>* *eat*, not vice versa), and the creativity to arrange storylines so that earlier events do not always appear first in **narrative order**. Two major issues in existing systems exacerbate the challenges: 1) temporal bias in pre-training and story datasets that leads to monotonic event temporal orders; 2) lack of explicit guidance that helps machines decide where to insert *flashbacks*. We propose to address these issues using structured storylines to encode events and their pair-wise temporal relations (*<before>*, *<after>* and *<vague>*) as **temporal prompts** that guide how stories should unfold temporally. We leverage a Plan-and-Write framework enhanced by reinforcement learning to generate storylines and stories end-to-end. Evaluation results show that the proposed method can generate more interesting stories with *flashbacks* while maintaining textual diversity, fluency and temporal coherence.¹

1 Introduction

Flashback is a popular creative writing technique that brings the readers from the present moment to the past via inserting earlier events in order to provide background or context of the current narrative (Pavis, 1998; Kenny, 2004; Gebeyehu, 2019). For example, in Figure 1a, the “GHOST” in Shakespeare’s play *Hamlet* instruments a *flashback* by interrupting the main narrative and describing a historical event to the audience that *Hamlet*’s father was killed by the current king rather than a snake.

¹Code, data and trained models are available here: https://github.com/PlusLabNLP/flashback_gen

GHOST: I am glad you’re eager...Now listen, Hamlet.
Everyone is told that *a poisonous snake bit me when I was sleeping in the orchard*. But that’s a lie...You should know, my noble son, *the real snake that stung your father* is now wearing his crown.
-- *Hamlet, Act 1 scene 5*

(a) A *flashback* example from William Shakespeare’s famous play *Hamlet* (in plain English). Red text indicates past events.

Story (1): [e1] *John lost consciousness when ambulance arrived.*
[e2] *He woke up in the hospital in the morning.*
Story (2): [e1] *John woke up in the hospital in the morning.*
[e2] *He lost consciousness when ambulance arrived.*

(b) Two-sentence stories with the same event **temporal order** but different **narrative order**. The second one with a *flashback* is intuitively more interesting than the first one.

Figure 1: (a) *flashback* (b) temporal v.s. narrative order.

Flashback, by manipulating the event temporal orders in narrative structure, can arouse readers’ emotions such as surprise, suspense, and curiosity (Brewer and Lichtenstein, 1981, 1982). These emotions stimulate readers’ interests and eventually contribute to the satisfaction of reading (Tan, 1996; Alwitt, 2002), which improves the **interest level** of a story. The example in Figure 1a injects historical events in the middle of the narrative. This arrangement of events can surprise readers and therefore, makes the story more interesting than a straightforward storyline where the past events are shown in the beginning.

Similarly, consider the pair of two-sentence stories in Figure 1b. Both stories are composed of the same events with the **temporal order** “*lost consciousness*” *<before>* “*woke up in the hospital.*” In Story (1), seeing [e1], readers can make a relatively easy educated guess of [e2], but it is more subtle in Story (2) as there are many different ways to end up in a hospital. By showing the ending event first, the *flashback* in Story (2) creates suspense that makes the following sentences less predictable, and thus arouses readers’ curiosity and makes the reading more interesting.

While human writers are capable of maneuver-

ing event temporal orders to compose coherent and interesting stories, it remains challenging for machines. The challenge is partially attributed to data bias. Ning et al. (2018a) shows that the pattern in Story (1) is dominant in human-written texts, where neighboring events with $\langle \textit{before} \rangle$ temporal relations (i.e., **narrative order** indicates **temporal order**) occur 60 – 70% of the time. This is also manifested in our experiments with vanilla language models amplifying this ratio and producing more than 80% $\langle \textit{before} \rangle$ relations for neighboring events in the generated stories. Furthermore, current state-of-the-art story generation systems that incorporate event representations usually assume event **temporal order** follows **narrative order** (Goldfarb-Tarrant et al., 2020; Lin et al., 2021). There are no explicit prompts in these systems that help determine when *flashback* should be used, leaving models to produce dull stories consisting of event sequences with monotonic $\langle \textit{before} \rangle$ relations.

To facilitate more effective *flashback*, we propose to incorporate **temporal prompts** in an end-to-end story generation framework inspired by the Plan-and-Write paradigm (Yao et al., 2019; Xu et al., 2020; Goldfarb-Tarrant et al., 2020), where machines first learn to plot a storyline, and then generate the story based on the storyline. Specifically, we encode predefined event **temporal prompts** in structured storylines. As the bottom block of Figure 2 shows, a structured storyline contains two components: 1) **event representations** where an event trigger (“grabbed”) and two arguments (“she” and “the dog”) are extracted from the original story sentences; and 2) **temporal prompts**: the temporal order between neighboring events, e.g. event 1: (“she”, “grabbed”, “the dog”) is $\langle \textit{after} \rangle$ event 2: (“white snow”, “blanketed”, “the ground”). By training our storyline generation model with these predefined pair-wise temporal relations, models capture how neighboring events are temporally related to each other; while during storyline decoding, supplying predefined **temporal prompts** can guide models to generate reasonable narratives with desirable event temporal orders.

Prior works (Fan et al., 2019; Goldfarb-Tarrant et al., 2020) build the storyline and story models separately, which creates a discrepancy where gold storylines are used during training, but predicted storylines are used during inference. To mitigate this training-inference discrepancy, we leverage reinforcement learning (RL) to train our systems end-

to-end. It enables the story model to train on generated storylines and updates the storyline model with the feedback from the story model. Our experimental results show that the RL-based models can leverage **temporal prompts** more effectively, resulting in more effective *flashback* generation and more interesting stories.

We summarize the contributions of this paper as follows: 1) To facilitate effective *flashback*, we propose to leverage structured storylines with **temporal prompts** to arrange events in story generation. 2) We integrate reinforcement learning in our story generation pipeline, which can help models better leverage **temporal prompts**. 3) We test our framework on two open-domain story datasets and show more effective *flashbacks* and increased **interest level** while maintaining fluency and temporal commonsense in the generated stories. To our best knowledge, this is a pioneering study on *flashbacks* in *neural story generation*.

2 Background and Task Definitions

In this section, we describe the key components: **events** and **temporal prompts** in our proposed structured storylines, and then define the Plan-and-Write generation task.

Event Representation. Following the definitions of ACE (2005), we define an event as a trigger word and its arguments. In this work, we simplify the representation by leveraging semantic role labeling (SRL) tools (Shi and Lin, 2019) to parse two arguments as shown in Figure 2. We only consider one event per story sentence and denote the k -th event in story i as $e_{i,k}$. We leave more complicated representations for future study.

Temporal Prompts. Let $r_i = \{r_{i,k}\}$ denotes the set of temporal relations between the k -th and the $(k+1)$ -th event in story i . If k indexes the last event, $r_{i,k}$ is not defined. Following the event relation definition of (Ning et al., 2018b), we use events’ *start time* to evaluate temporal order.

Structured Storyline. Figure 2 provides a storyline consisting of five event representations extracted from our data. More formally, let $\mathcal{S}_i = \{e_{i,1}, e_{i,2}, \dots, e_{i,k}, \dots, e_{i,n}\}$ indicates a storyline with n events. Encoding **temporal prompts**, \mathcal{S}_i becomes $\mathcal{S}_i^r = \{e_{i,1}, r_{i,1}, e_{i,2}, r_{i,2}, \dots, e_{i,k}, r_{i,k}, \dots, e_{i,n}\}$. Note that in this work, r_i is provided as **predefined prompts** rather than predicted as $e_{i,k}$.

Story. Our ultimate goal is to generate *flashbacks* in stories. We denote the story associated with the

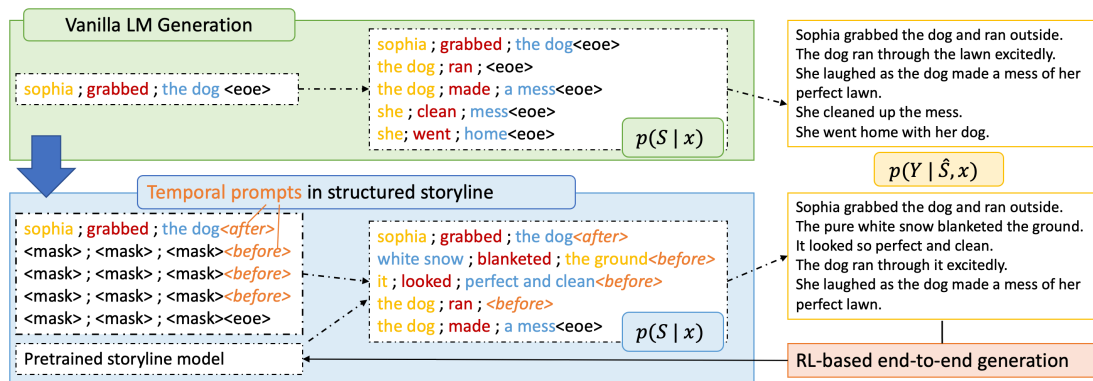


Figure 2: An illustration of our overall model. Here we use the first sentence of the story (and its associated event representation) as input x . The upper block shows the vanilla implementation of the Plan-and-Write workflow. The bottom block is our core novel design by leveraging **temporal prompts** in structured storylines to generate *flashbacks*. For illustration purposes, we re-order the triggers and arguments, and storylines are ground-truths (i.e. not predicted by models). Our final model uses reinforcement learning to implement end-to-end training.

storyline \mathcal{S}_i as \mathcal{Y}_i .

Plan-and-Write is a two-stage framework that first generates storyline $\hat{\mathcal{S}}_i$ given some input x (e.g. title, leading sentence), and then generate $\hat{\mathcal{Y}}_i$ based on $\hat{\mathcal{S}}_i$. Again, $r_{i,k}$ are given as *predefined* prompts whereas $e_{i,k}$ are to be predicted as part of the storyline generation shown in Figure 2.

3 Framework for FlashBack Generation

In this section, we first provide an overview of the Plan-and-Write story generation system and introduce a vanilla version of the end-to-end training method. Then we describe the details of our key contribution of leveraging event **temporal prompts** to generate *flashbacks*. After that, we discuss pretraining structured storylines with self-labeled data and incorporating reinforcement learning to jointly train our end-to-end models.

3.1 Plan-and-Write Models

In order to provide better explainability and controllability over the machine generated stories, recent research efforts (Yao et al., 2019; Xu et al., 2020; Goldfarb-Tarrant et al., 2020) explore dividing story generation into two steps: 1) from input or prefix, x , we first produce a storyline, \mathcal{S}_i ; 2) based on the storyline, we generate a story, \mathcal{Y}_i . We describe the details below.

Storyline Model. Let α denote the parameters of the storyline model, per sample training loss can be computed as $\mathcal{L}_\alpha = -\log p(\mathcal{S}_i | x_i, \alpha)$.

Story Model. Let β denote the parameters of the story model, per sample training loss can be computed as $\mathcal{L}_\beta = -\log p(\mathcal{Y}_i | x_i, \mathcal{S}_i, \beta)$.

Inference. Note that \mathcal{S}_i above is the gold storyline extracted from \mathcal{Y}_i . At the inference time, we do not have \mathcal{S}_i , and have to replace it with $\hat{\mathcal{S}}_i$, the predicted storyline. This results in a discrepancy between the training and inference time.

End-to-end Training. Instead of using gold storyline \mathcal{S}_i to train a story model, we can take $\hat{\mathcal{S}}_i$ as its input. Now the per sample training loss for the story model becomes $\mathcal{L}_\theta = -\log p(\mathcal{Y}_i | x_i, \hat{\mathcal{S}}_i, \theta)$, where θ indicates the end-to-end story model parameters. End-to-end training can alleviate the gap between the training and inference time, and potentially lead to more consistent stories².

3.2 Structured Storyline Construction

As Figure 2 shows, for a story sentence, we first use the SRL tool to parse its trigger $t_{i,k}$ and two arguments $a_{i,k}^1$ and $a_{i,k}^2$. We then convert this representation into a textual form: “ $t_{i,k} ; a_{i,k}^1 ; a_{i,k}^2 \langle eoe \rangle$ ”, where “;” separates two event components, and $\langle eoe \rangle$ indicates event ending. For example, the parsed $t_{i,k}$, $a_{i,k}^1$ and $a_{i,k}^2$ in the story sentence “she grabbed the dog and ran outside” are “grabbed,” “she” and “the dog” respectively. They are concatenated into a final textual representation as “grabbed ; she ; the dog $\langle eoe \rangle$.”

Depending on the experimental setup, we may use no or only the leading event as input, x . Inspired by the mask prediction design in Devlin et al. (2019); Liu et al. (2019); Lewis et al. (2020), we represent the remaining missing events in the inputs as “ $\langle \text{mask} \rangle ; \langle \text{mask} \rangle ; \langle \text{mask} \rangle ; \langle eoe \rangle$,” where $\langle \text{mask} \rangle$ indicates either event trigger word or argu-

²End-to-end training may have small coverage of the generated events in stories, which we address in Appendix F.

Algorithm 1 RL-based End-to-end Training

```
1: Randomly initialize  $\alpha$  and  $\theta$ 
2: Pretrain  $\alpha$  ▷ storyline pretraining
3: for  $i \in \mathbf{M}$  do ▷ loop through all data4
4:   Generate  $\hat{S}_i^r$  from storyline model ( $\alpha$ )
5:    $\mathcal{L}_\theta = -\log p(\mathcal{Y}_i | \mathbf{x}_i, \hat{S}_i^r, \theta)$ 
6:    $\nabla J_\alpha = R_i \cdot \nabla \log(p(S_i | \mathbf{x}_i, \mathbf{r}_i, \alpha))$ 
7:    $\alpha = \alpha - \nabla J_\alpha$ 
8:    $\theta = \theta - \nabla \mathcal{L}_\theta$ 
9: end for
```

ments to be predicted by the storyline model.

3.3 Temporal Prompt Encoding

Temporal prompts are used to generate *flashbacks*. As we mentioned in Section 2, we encode a sequence of *predefined* event **temporal prompts** $\mathbf{r}_i = \{r_{i,k}\}$ in storyline for $k \in \{1, n-1\}$ to help models determine whether the next event mention (in narrative order) should start earlier or later than its preceding event mention. We use temporal relation extraction tools to annotate all $r_{i,k}$ in our experimental data. Specifically, we use ECONET (Han et al., 2021b) finetuned on the MATRES dataset (Ning et al., 2018b) to predict the temporal relation between neighboring events.³ The context and the locations of a pair of event trigger words are fed into ECONET to predict their temporal order. The **temporal prompt** set consists of *⟨before⟩*, *⟨after⟩* and *⟨vague⟩* (capturing undetermined temporal order), and are fixed in S_i^r . Note that *⟨vague⟩* indicates undetermined temporal order due to the ambiguity of the context (Cassidy et al., 2014; Ning et al., 2018b) and it does not suggest the context is poor or the relations are wrong. As shown in Figure 2, we replace the end-of-event token *⟨eoe⟩* with **temporal prompts** in storylines, except for the last event which does not have a next event. With the prompt-augmented storylines, S_i^r , we can re-write the storyline loss as $\mathcal{L}_\alpha = -\log p(S_i^r | \mathbf{x}_i, \mathbf{r}_i, \alpha)$, and story loss as $\mathcal{L}_\theta = -\log p(\mathcal{Y}_i | \mathbf{x}_i, \hat{S}_i^r, \theta)$.

3.4 Storyline Pretraining

Using intermediate pretraining to adapt original pretrained language models has been shown to be effective for a variety of downstream tasks such as information extraction (Joshi et al., 2020), question-answering (Khashabi et al., 2020; Garg et al., 2020) and commonsense reasoning (Zhou et al., 2021). To capture more diverse event sequences and facilitate better story generation, we explore pretraining

³The ECONET tool is available here: <https://github.com/PlusLabNLP/ECONET>.

storyline model with SRL extracted storyline from BookCorpus dataset (Zhu et al., 2015), and use learned α to initialize storyline models.

3.5 RL-based End-to-end Model

The end-to-end model described in Sec. 3.1 allows the story model to train with the generated storylines and hence alleviate the gap between training and inference. However, this workflow still lacks a mechanism that enables the storyline model to adjust with the feedback from the story model. The challenges of training storyline and story models jointly originate from decoding storylines as inputs for the story model, which involves non-differentiable token selections. Thus, the final loss L_θ cannot be directly back-propagated into the storyline model.

To overcome this barrier, we adopt reinforcement learning (RL), specifically, the REINFORCE algorithm (Williams and Peng, 1991) in our end-to-end training. Let $R_i = R(\mathbf{x}_i, \mathbf{r}_i)$. The expected reward with respect to the storyline model can be written as $\mathbb{E}_\alpha [R_i] = \mathbb{E} [R_i \cdot \log(p(S_i^r | \mathbf{x}_i, \mathbf{r}_i, \alpha))]$. The gradient to update the storyline model is $\nabla J_\alpha = \mathbb{E} [R_i \cdot \nabla \log(p(S_i^r | \mathbf{x}_i, \mathbf{r}_i, \alpha))]$, which can be approximated with sampling techniques. Motivated by Xu et al. (2018), we use negative loss of the story model to construct rewards, that is, $R = -\mathcal{L}_\theta$.⁵ In other words, smaller loss from the story model is associated with larger reward. Algorithm 1 summarizes the overall method.

4 Experimental Setup

In this section, we start by describing our research objectives, then we describe our data, evaluation metrics, experimental designs and implementation details aiming to achieve these objectives.

The overall research objective is to measure the impact of using **temporal prompts** in structured storylines. Specifically, can *⟨after⟩* successfully induce *flashbacks*? If so, does that contribute to the interest level of the generated stories while maintaining the overall quality of the texts?

4.1 Datasets.

ROCStories (Mostafazadeh et al., 2016a) and WritingPrompts (Fan et al., 2018) are our experimental datasets. We ensured all reported results using the same test data as the baseline systems (Xu et al., 2020) and (Goldfarb-Tarrant et al., 2020). For

⁵We do not use baseline reward as we found this simple reward design works effectively in our experiments.

pretraining data, we use BookCorpus (Zhu et al., 2015). Appendix B shows all details of data splits and pre-processing process.

4.2 Temporal Prompts Constructions

ECONET was finetuned three times with different random seeds, so we take the consensus vote from three models. If there is any disagreement, we label the temporal order as $\langle vague \rangle$. We benchmark ECONET’s annotation performances in Appendix G, which shows it provides highly accurate temporal relations. For human evaluations specifically, we consider two prompt settings in order to gauge different impacts of $\langle after \rangle$. 1) for ROCStories, all structured storylines consist of exactly four *predefined temporal prompts* created following Sec 3.3. We randomly sample stories with one $\langle after \rangle$ prompt from the test data. We will show later in the analysis that vanilla language models would generate more than 80% event pairs with $\langle before \rangle$ relations for ROCStories; $\langle after \rangle$ prompt should bring this ratio down if it is effective. 2) for WritingPrompts, since the number of events is not fixed, we randomly sample test stories generated with $\langle after \rangle$ prompts for evaluation.

4.3 Automatic Evaluation Metrics

We use automatic metrics to evaluate the textual quality of stories. We report **Ref. PPL**: reference stories’ perplexity in our models and **Gen. PPL**: generated stories’ perplexity scored by GPT-2 (Radford et al., 2019). For diversity, we report **Distinct Ratio (%)**: overall vocabulary:token number ratio. We also report standard **BLEU-3** and **ROUGE_L**.

4.4 Human Evaluation Metrics

We rely on human annotators to analyze the effectiveness of *flashback* generations. We request 18 MTurkers who succeeded in our previous annotation tasks (Han et al., 2021a) to evaluate stories produced by our compared models. We host a small qualification round followed by a large annotation task. Only 10 workers are qualified, and we only consider their annotations. Eventually, we collect 106 and 77 sets of valid annotations for ROCStories and WritingPrompts.

Temporal diversity. The dominance of $\langle before \rangle$ relation in our data can make models biased toward generating stories with more $\langle before \rangle$ relations. Therefore, we are interested to see how inserting an $\langle after \rangle$ prompt can help **increase** the percentage of non- $\langle before \rangle$ event relations in the generated stories. Let \hat{R}_r indicate the percentage of a

particular relation annotated by MTurkers. We calculate the entropy of the set $\{\hat{R}_r\}, \forall r \in \{\langle before \rangle, \langle after \rangle, \langle vague \rangle\}$ to measure temporal diversity.

Accuracy measures the percentage of $\langle after \rangle$ being correctly incorporated in the generated stories labeled by human annotators. We used a relaxed version by counting annotated $\langle vague \rangle$ as correct too, as $\langle vague \rangle$ can potentially be $\langle after \rangle$. Both **accuracy** and **temporal diversity** can show the effectiveness of generating *flashbacks* using $\langle after \rangle$.

Temporal coherence indicates if the event sequence in a generated story aligns with an annotator’s **temporal commonsense**.⁶ 1 and 0 correspond to yes and no, respectively.

Interest level. Precisely defining interest level is difficult as it is a broad concept. So we focus on the *unexpectedness* component of cognitive interest. As pointed out by Behrooz (2019), *unexpectedness* can be further explained as how predictive an event is, which is closely related to *flashback* generation. Therefore, we define an *interesting* event as 1) being unexpected or surprising and 2) being logical according to the context and general commonsense.⁷

For the compared models, we ask annotators to provide ranks between 1 to K for the generated stories, with K indicating the most interesting story and 1 indicating the least interesting one. We encourage workers to provide different scores for all compared stories, but equal scores are allowed. The max score K depends on the number of compared models, 5 for ROCStories and 4 for WritingPrompts. We provide detailed instructions in the interface shown in the appendix. Crucially, **interest level** is separately annotated from other metrics and we ensure annotators do not see the same set of stories in both tasks.

4.5 Compared Models

Baselines. Xu et al. (2020), denoted as MEGATRON, is chosen as the baseline as it outperforms previous systems such as Guan et al. (2020) on ROCStories. We also compare with TEMPORAL-BART (Lin et al., 2021) as it is pretrained with temporal ordering and event infilling tasks. For WritingPrompts, we compare with CONTENTPLANNING (Goldfarb-Tarrant et al., 2020) as it also

⁶interface with detailed instructions and examples can be found in Figure 7 of the appendix.

⁷the second definition is not identical to the temporal coherence above. Events contradicting general commonsense can still be temporally coherent (see Figure 7) for examples.

Models	Automatic Evaluation					Human Evaluation			
	Ref. PPL(↓)	Gen. PPL(↓)	Distinct Ratio(↑)	BLEU (↑)	ROUGE _L (↑)	Temporal Diversity(↑)	Accuracy (↑)	Temporal Coherence (↑)	Interest Level (↑)
TEMPORAL-BART	24.65	19.47	4.10	5.01	19.12	-	-	-	-
MEGATRON	-	34.14	4.57	2.57	15.23	1.21*	-	0.78	2.69
VANILLA-GEN	27.30	19.29	3.99	5.13	19.29	0.88	-	0.88	2.95
+ STRUCTURED-PROMPT	22.85	19.94	4.09	5.07	19.39	1.09	55.75	0.82	3.03
+ PRETRAINED	21.16	19.25	4.01	5.06	19.44	1.07	52.21	0.84	2.96
+ RL	15.45	19.42	4.17	5.20	19.49	1.14	56.64	0.86	3.06

Table 1: Evaluation results for ROCStories. All values in the bottom block are averaged over three runs. MEGATRON’s **Ref. PPL** is missing as it is not reported in the paper and the model is not publicly available. MEGATRON and VANILLA-GEN do not use any **temporal prompts**, so **accuracy** is not applied. MEGATRON’s **temporal diversity*** is high due to large amount of *<vague>* annotations in its generated stories.

Input prompt	[s1] my favorite class was english literature. <i><before></i> <i><after></i> <i><before></i> <i><before></i>
baseline	[s2] i was in english class for the first time. [s3] i had to take a class with my local author this semester. [s4] i decided to ask him to speak to me at a local book festival. [s5] he said that he was really good and i liked him.
VANILLA	[s2] one day, i decided to go to the library. [s3] i found a book that i really liked. [s4] i read the whole book in one day. [s5] i was so happy that i went back to school.
RL	[s2] i would go to the library to read all the books. [s3] i had never been to a library before. [s4] i was so excited to go. [s5] i ended up spending a lot of time in the library.
Input prompt	[s1] tom was excited to go on a date. <i><after></i> <i><before></i> <i><before></i> <i><before></i>
baseline	[s2] he went to his favorite restaurant to pick up his date. [s3] the date went well and then tom got drunk. [s4] he told his date they were going to a comedy show. [s5] the date went on and they had a great time at the movie.
VANILLA	[s2] he was going to take his girlfriend to the movies. [s3] they went to the movie theater. [s4] they had a great time. [s5] tom couldn’t wait to go back again.
RL	[s2] he had been looking forward to it for months. [s3] he went to the mall and bought a new dress. [s4] he wore it to his date’s house. [s5] his date loved it and they had a great time.

Table 2: Generated stories by our compared models. baseline: MEGATRON, VANILLA-GEN and RL-based. end-to-end. Input is the first sentence. Highlighted sentences contain an event that started earlier than the preceding event.

adopts the Plan-and-Write workflow as well as structured event representations. Appendix D describes more details of baseline systems.

We describe our own model variants below,

1. VANILLA-GEN uses the parameters of a pre-trained language model (LM), specifically BART-base (Lewis et al., 2020), to initialize both the storyline and story models. Its workflow is illustrated in the upper block of Figure 2. Since no information other than the prefix (first sentence, prompt, etc.) is used to generate the story, we denote this model as vanilla LM generation or VANILLA-GEN.

2. STRUCTURED-PROMPT enhances VANILLA-GEN by using a structured storyline of events to encode **temporal prompts**, which is associated with the workflow of the bottom block of Figure 2.

3. PRETRAINED. For ROCStories data only,

we initialize the storyline model of STRUCTURED-PROMPT with the pretrained parameters.

4. RL uses the same inputs as STRUCTURED-PROMPT. The difference is that reinforcement learning is used to train storyline and story models jointly. As Algorithm 1 shows, RL-based model is trained following the same forward workflow as STRUCTURED-PROMPT, but during backpropagation, the storyline models’ parameters are updated.

5 Results and Analysis

The main results for ROCStories and WritingPrompts are shown in Table 1 and Table 3 respectively. Examples of generated stories can be found in Table 2 and Table 6 for ROCStories and Table 11 in the appendix for WritingPrompts. We organize our discussions and analysis in the following sections by answering the four research questions. **Q1)** Can our proposed models (with **temporal prompts**) produce stories with good **textual quality**? **Q2)** Are our proposed models effective at generating *flashbacks*? **Q3)** Can our proposed models maintain event **temporal coherence** in stories? **Q4)** How do our proposed models contribute to stories’ **interest levels**?

5.1 Textual Quality

We measure the textual quality of stories using a wide range of automatic evaluation metrics.

Perplexity. For ROCStories, all three model variants can improve **Ref. PPL** against VANILLA-GEN and TEMPORAL-BART while maintaining good **Gen. PPL**. The weak **Gen. PPL** of MEGATRON may be attributed to its sentence-by-sentence generation pipeline, whereas our models generate an entire story in an integrated step. For WritingPrompts, both model variants improve **Gen. PPL** over VANILLA-GEN and CONTENTPLANNING while maintaining good **Ref. PPL**.

Token Diversity. For ROCStories, RL-based model improves the VANILLA-GEN by 0.18 per

Models	Ref. PPL (↓)	Gen. PPL (↓)	Distinct Ratio (↑)	BLEU (↑)	ROUGE _L (↑)	Tok. Len. (↑)	Pearson Corr. (↑)	Temporal Coherence (↑)	Interest Level (↑)
CONTENTPLANNING	-	25.52	1.80	3.46	14.40	252.3	0.04	0.57	2.20
VANILLA-GEN	31.04	11.17	3.50	0.67	9.43	160.2	0.09	0.52	2.49
+ STRUCTURED-PROMPT	30.77	9.30	2.86	1.44	10.95	208.6	0.56	0.49	2.49
+ RL	30.98	9.50	2.83	1.39	10.78	203.8	0.57	0.55	2.62

Table 3: Evaluation results for WritingPrompts. Pearson correlation approximates the effectiveness of prompts.

Distinct Ratio. MEGATRON achieves the highest token diversity as it incorporates external knowledge-bases that make the generated stories contain novel tokens. For WritingPrompts, we observe longer stories are associated with poorer scores. However, the large increases in **Distinct Ratio** suggest that the token usages in our proposed models are diverse.

BLEU and ROUGE_L. For ROCStories, the proposed models perform on-par with VANILLA-GEN and TEMPORAL-BART while outperforming MEGATRON, which generates the shortest stories among all compared models. For WritingPrompts, CONTENTPLANNING performs the best partially due to its usage of BART-large models.

The overall performances across these three types of automatic metrics suggest that using **temporal prompts** in the Plan-and-Write framework can produce stories with high textual quality.

5.2 Effectiveness on Flashback Generation

The second research question probes the effectiveness of using **temporal prompts** on generating *flashbacks*. For ROCStories, all models can generate stories with the same number of events/sentences as the gold stories. This allows annotators to judge pairwise event relations in the generated stories and help us check whether the generated events have relations truthfully reflecting the **temporal prompts** used. **Accuracy** is the perfect metric for this. As Table 1 shows, the final RL-based model achieves the highest score, which indicates the strongest effectiveness of generating *flashbacks*.

However, **temporal prompts** are not used in the baselines and VANILLA-GEN. So we compute an approximate measures of effectiveness, **temporal diversity**, which indicates how many non-*<before>* relations *<after>* prompt can induce. Table 1 shows that STRUCTURED-PROMPT, PRETRAINED and RL-based models can help improve VANILLA-GEN with more than 80% generated *<before>* relations. MEGATRON achieves the highest score due to the largest amount (29%) of *<vague>* relations (complex or undetermined) annotated by MTurkers shown in Figure 3, which is associated with its lowest **temporal coherence** score.

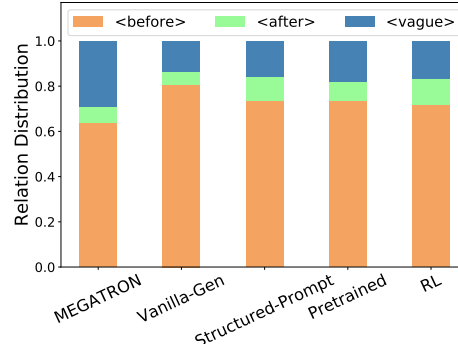


Figure 3: Relation annotation distribution by MTurkers for compared models trained on ROCStories.

For WritingPrompts, stories are long and can contain dialogues or short phrases without events at all. These make the sentence or event alignments between the gold and generated stories worse than ROCStories, i.e. $e_{i,k}, r_{i,k}$ may not correspond to the k -th sentence in \mathcal{Y}_i . Therefore, **accuracy** cannot be computed. To obtain an approximate metric, we use the tool described in Sec. 3.3 to annotate neighboring event temporal relations in the *generated* test stories for all the compared models. Slightly different from **temporal diversity**, we calculate the total number of machine annotated *<after>* relations, denoted as $\hat{N}_{i,A}$ in each \hat{S}_i . Let $\{N_{i,A}\}$ denote the number of *<after>* **temporal prompts** extracted in gold stories. We compute the **Pearson correlation** (Benesty et al., 2009) between the sets $\{\hat{N}_{i,A}\}$ and $\{N_{i,A}\}$ as the measure.

As Table 3 shows, for CONTENTPLANNING and VANILLA-GEN without **temporal prompts**, the correlations are weak; whereas when **temporal prompts** are used in both STRUCTURED-PROMPT and RL-based models, the correlations are strong. Although using models’ temporal annotations for the generated stories is not as precise as human annotations, the large differences in correlation provide another piece of evidence that our proposed methods are effective at generating *flashbacks*.

5.3 Temporal Coherence

Generating *flashbacks* requires a system to disrupt the monotonic *<before>* sequence, which is the dominant temporal pattern generated by VANILLA-GEN (see Figure 3). In other words, *flashbacks* with at least one *<after>* are minor patterns that can be

hard to learn in our data, which may result in event sequences violating our temporal commonsense. Thus, we need to check that stories generated with *flashbacks* maintain good **temporal coherence**. As shown in Table 1 and 3, our proposed models with **temporal prompts** can achieve on-par or slightly lower scores, suggesting little trade-off of **temporal coherence** in generating *flashbacks*. We will discuss this more in the error analysis (Sec. 6).

5.4 Contributions to the Interest Level

As we can observe in Table 1, the impact of **temporal diversity** and **coherence** on the **interest level** appears to be complex. To better understand the dynamics among these metrics, we run ordinary least square regressions (OLS) (Kenney and Keep-ing, 1965) by setting **interest level** as the target variable and **temporal diversity**, **coherence** and the number of *<after>*s as predictors. Since all of these metrics apply to each of the compared stories, the total instances are 530 and 308 for ROCStories and WritingPrompts, respectively.

As Table 4 shows, for ROCStories, holding other metrics constant, adding 1 unit to **temporal coherence** and **diversity** leads to a 0.609 increase and 0.532 decrease of the **interest level**. The former result implies that a story lacking event temporal coherence tends to be less interesting. The latter result suggests that increasing **temporal diversity** may lead to less interesting stories, which we hypothesize could be attributed to two factors: 1) *<before>* is dominant in ROCStories, and by using *<after>* as prompt, we force models to generate relations less seen in data. 2) Figure 3 shows **temporal diversity** can increase with more *<vague>* relations. Since *<vague>* is an undetermined temporal relation even for our annotators, it could make the storyline confusing and thus lead to less interesting stories. The coefficient for the number of *<after>* indicators is positive with strong statistical significance. It suggests that holding the other two metrics constant, adding the number of *<after>* indicators by 1 contributes to 0.387 increases of the **interest level**.

For WritingPrompts, although we are not able to conclude that the estimates are statistically significant, the coefficients have the same signs as ROCStories. Also, we observe that the p-value of the number of *<after>* indicators is much lower than the other two variables, which implies a relatively stronger (positive) impact.

Since **temporal prompts** in human evaluations all contain at least one *<after>*, these results show

	ROCStories		WritingPrompts	
	Coef.	p-value	Coef.	p-value
Temporal Coherence	0.609	0.000*	0.006	0.963
Temporal Diversity	-0.532	0.004*	-0.279	0.410
# <i><after></i> prompt	0.387	0.000*	0.034	0.238

Table 4: OLS regression results on temporal coherence, diversity and number of *<after>* indicators. The coefficients for the intercept are omitted. * means statistically significant with 99% confidence.

that when *<after>* prompt successfully produces event pairs with *<after>* relation in the final stories, it makes stories more interesting. Now, we can answer the final research question: improving **temporal diversity** can help **interest level** when *<after>* prompts are effective at generating *<after>* relations in stories; that is, when *flashbacks* truly work, stories become more interesting.

6 Error Analysis

In Table 1, we observe that our final models cannot outperform VANILLA-GEN for **Temporal Coherence** and fall behind MEGATRON for **Temporal Diversity**. We show examples below to provide feasible explanations.

Temporal Coherence. Table 5 shows three comparisons between our final RL-based model and the VANILLA-GEN baseline. In all three pairs, human evaluators mark VANILLA-GEN’s generated stories as temporally coherent but not for RL’s stories. As we can see in these examples, the red highlighted events temporally contradict the preceding context. In Input 1, the context shows “I have listened to the new album” so “can’t wait to listen to the new album” shouldn’t follow the former event. For Input 2, if “Anna had gotten her hair permed,” it does make sense that the event “went to apply it (the perm)” follows. In Input 3, if the officer “gave [a] ticket” already, it does quite make sense to give “a warning” afterward as it is a weaker penalty than a “ticket.” However, we observe that in all three cases, *<after>* relations are successfully reflected (*flashback*). Since stories with *<after>* relations are minority cases in the data, we hypothesize that they have not been perfectly learned by our proposed models. We leave more rigorous investigation for future research efforts.

Temporal Diversity. Table 10 in the appendix shows three comparisons between our final RL-based model and the MEGATRON baseline. In all three pairs, MEGATRON’s generated stories are more temporally diverse based on the **predicted relations** provided by human evaluators. However, MEGATRON’s stories are either contradic-

Input 1 prompt	[s1] the new album from my favorite pop star came out today. <after> <before> <before> <before>
VANILLA	[s2] i was so excited to listen to it. [s3] i listened to it all day long. [s4] i couldn't believe how good it was. [s5] red i can't wait to listen again.
RL	[s2] i had never heard the album before. [s3] i went to the store and bought the album. [s4] when i got home, i listened to the whole thing. [s5] i can't wait to listen to the new album.
Input 2 prompt	[s1] anna had just gotten her hair permed. <before> <before> <after> <before>
VANILLA	[s2] she was excited to see how it looked. [s3] but then she saw that her hair had grown out. [s4] anna was horrified! [s5] she had to go to the salon to get it fixed.
RL	[s2] she was excited to see how it looked. [s3] but when she went to apply it, it was terrible. [s4] her hair had grown way too long! [s5] anna had to go to the salon to get her hair fixed.
Input 3 prompt	[s1] last week we saw a police car in front of a parked car. <after> <before> <before> <before>
VANILLA	[s2] the police officer pulled over the car. [s3] he asked the driver if he had insurance. [s4] the driver said he did not. [s5] the officer gave him a ticket.
RL	[s2] the car was parked in the parking lot. [s3] the police officer pulled over the car. [s4] he gave the driver a ticket. [s5] the officer gave him a warning.

Table 5: Examples from our RL-based model that are not as temporally coherent as examples generated by the VANILLA-GEN model. Highlighted sentences contain an event that started earlier than the preceding event. Red text indicates temporal incoherence.

tory (Input 1), incoherent (Input 2) or repetitive (Input 3), resulting in higher ambiguous event relations, i.e. <vague> annotations (consistent with Figure 3). Therefore, despite lower temporal diversity, our proposed models can still produce stories with higher quality, which is demonstrated via other metrics such as **Temporal Coherence** and **Interest Level**.

7 Related Work

Generating flashbacks has been studied in a few prior works. [Bae and Young \(2008\)](#) is one of the early efforts proposing a planning-based approach to generate flashbacks to evoke surprise in the readers. Follow-up works proposed a cognitive-based model that finds the best location in the original stories to insert a past event ([Wu et al., 2016](#)). Our work differs from this line of research by using temporal prompts with pretrained language models to generate integrated *flashback* in stories. [Hoek et al. \(2014\)](#) studies *flashback* in game narrative generation, which is remotely related to our work.

Plan-and-Write framework has been shown to be an effective method to enhance the explainability and controllability of story generation. [Yao et al. \(2019\)](#) enables machines to produce a sequence of

keywords prior to generating stories. Follow-up works leverage commonsense or external knowledge to enhance the quality of stories ([Guan et al., 2020](#); [Xu et al., 2020](#); [Tan et al., 2021](#)). [Goldfarb-Tarrant et al. \(2020\)](#) is one of our compared works that incorporates SRL extracted event representations in storylines and train models with several event-related decoding objectives. Our work differs from it by explicitly encoding **temporal prompts** in event plots that facilitates *flashback*.

Structured representation such as discourse structure ([Guan et al., 2021](#)), story keywords ([Peng et al., 2018](#); [Goldfarb-Tarrant et al., 2019](#)) and event/plot graph ([Ammanabrolu et al., 2019, 2021](#)) have been widely used in story generation to enable models to output diverse stories, but they are remotely related to our *flashback* generation task.

Reinforcement learning has also been explored in two-stage story generation such as [Xu et al. \(2018\)](#) and [Tambwekar et al. \(2019\)](#). Our motivation of using RL-based generation is to enhance the effectiveness of **temporal prompts**.

Event temporal reasoning helps the construction of the **temporal prompts**. It has been studied in story understanding ([Han et al., 2019b](#)), information extraction ([Ning et al., 2017](#); [Han et al., 2019c,a, 2020](#); [Ma et al., 2021](#)), QA ([Ning et al., 2020](#); [Zhou et al., 2019](#)), and event generation ([Lin et al., 2021](#); [Li et al., 2021](#)). Our proposed method is the pioneer work to introduce event **temporal prompts** in *flashback* generation.

8 Conclusions

We propose to generate *flashbacks* in stories by encoding temporal prompts in structured storylines. Experimental and evaluation results show our proposed systems can produce fluent, temporally coherent, and more interesting stories. Future work can focus on improving temporal prompts so that they handle minority cases (i.e. prompts with one or more <after> indicators) more effectively. Designing more powerful rewards in reinforcement learning is another promising research direction.

Acknowledgments

We thank the PlusLab members and the anonymous reviewers for their constructive feedback. This work is supported in part by IARPA, via Contract No. 2019-19051600007, the DARPA Machine Common Sense (MCS) program under Cooperative Agreement N66001-19-2-4032, and a CISCO research award.

References

- ACE. 2005. The ace 2005 (ace 05) evaluation plan evaluation of the detection and recognition of ace entities , values , temporal expressions , relations , and events 1.
- L. Alwitt. 2002. Maintaining attention to a narrative event. *Shohov, S.P. (eds.) Advances in Psychology Research*, 18:99–114.
- Prithviraj Ammanabrolu, Wesley Cheung, William Broniec, and Mark O Riedl. 2021. Automated storytelling via causal, commonsense plot ordering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5859–5867.
- Prithviraj Ammanabrolu, Ethan Tien, Wesley Cheung, Zhaochen Luo, William Ma, Lara J. Martin, and Mark O. Riedl. 2019. **Story realization: Expanding plot events into sentences**. In *Proceedings of AAAI 2020*.
- Byung-Chull Bae and R. Michael Young. 2008. A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In *Interactive Storytelling*, pages 156–167, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Morteza Behrooz. 2019. *Curating Interest in Open Story Generation*. Ph.D. thesis. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2021-05-25.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 37–40. Springer.
- W.F. Brewer and E. H. Lichtenstein. 1981. Event schemas, story schemas, and story grammars. *Long, J. and Baddeley, A. (eds.) Attention and Performance*, 9:263–279.
- W.F. Brewer and E. H. Lichtenstein. 1982. Stories are to entertain: A structural-affect theory of stories. *Journal of Pragmatics*, 6:473–483.
- Taylor Cassidy, Bill McDowell, Nathanael Chambers, and Steven Bethard. 2014. **An annotation framework for dense event ordering**. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 501–506, Baltimore, Maryland. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. **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.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2020. **Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7780–7788.
- Tsegaye Abie Gebeyehu. 2019. **Analysis of literary techniques employed in the novel the revelation: Flashback in focus**. *Journal of Literature, Languages and Linguistics*, 63.
- Seraphina Goldfarb-Tarrant, Tuhin Chakrabarty, Ralph Weischedel, and Nanyun Peng. 2020. **Content planning for neural story generation with aristotelian rescoring**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4319–4338, Online. Association for Computational Linguistics.
- Seraphina Goldfarb-Tarrant, Haining Feng, and Nanyun Peng. 2019. **Plan, write, and revise: an interactive system for open-domain story generation**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 89–97, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. 2020. **A knowledge-enhanced pre-training model for commonsense story generation**. *Transactions of the Association for Computational Linguistics*, 8:93–108.
- Jian Guan, Xiaoxi Mao, Changjie Fan, Zitao Liu, Wenbiao Ding, and Minlie Huang. 2021. **Long text generation by modeling sentence-level and discourse-level coherence**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6379–6393, Online. Association for Computational Linguistics.
- Rujun Han, I-Hung Hsu, Jiao Sun, Julia Baylon, Qiang Ning, Dan Roth, and Nanyun Peng. 2021a. **ESTER: A machine reading comprehension dataset for reasoning about event semantic relations**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7543–7559, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Rujun Han, I-Hung Hsu, Mu Yang, Aram Galstyan, Ralph Weischedel, and Nanyun Peng. 2019a. [Deep structured neural network for event temporal relation extraction](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 666–106, Hong Kong, China. Association for Computational Linguistics.
- Rujun Han, Mengyue Liang, Bashar Alhafni, and Nanyun Peng. 2019b. Contextualized word embeddings enhanced event temporal relation extraction for story understanding. In *2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT 2019), Workshop on Narrative Understanding*.
- Rujun Han, Qiang Ning, and Nanyun Peng. 2019c. [Joint event and temporal relation extraction with shared representations and structured prediction](#). 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 434–444, Hong Kong, China. Association for Computational Linguistics.
- Rujun Han, Xiang Ren, and Nanyun Peng. 2021b. [ECONET: Effective continual pretraining of language models for event temporal reasoning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5367–5380, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rujun Han, Yichao Zhou, and Nanyun Peng. 2020. [Domain knowledge empowered structured neural net for end-to-end event temporal relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5717–5729, Online. Association for Computational Linguistics.
- Marissa Hoek, Mariët Theune, and Jeroen Linssen. 2014. [Generating game narratives with focalization and flashbacks](#). volume 10.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- J. F. Kenney and E. S. Keeping. 1965. Ols regression via normal equations. In *Mathematics of Statistics*, 3 edition, chapter 15, pages 252–285. Princeton, NJ.
- Kenny. 2004. *Teaching Tv Production in a Digital World: Integrating Media Literacy*. Libraries Unltd Incorporated.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Manling Li, Sha Li, Zhenhailong Wang, Lifu Huang, Kyunghyun Cho, Heng Ji, Jiawei Han, and Clare R. Voss. 2021. [Future is not one-dimensional: Graph modeling based complex event schema induction for event prediction](#). *arXiv preprint arXiv:2104.06344*.
- Shih-Ting Lin, Nathanael Chambers, and Greg Durrett. 2021. [Conditional generation of temporally-ordered event sequences](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7142–7157, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint*, arXiv:1907.11692.
- Mingyu Derek Ma, Jiao Sun, Mu Yang, Kung-Hsiang Huang, Nuan Wen, Shikhar Singh, Rujun Han, and Nanyun Peng. 2021. Eventplus: A temporal event understanding pipeline. In *2021 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), Demonstrations Track*.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016a. [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.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016b. [CaTeRS: Causal and temporal relation scheme for semantic annotation of event structures](#). In *Proceedings of the Fourth Workshop on Events*, pages 51–61, San Diego, California. Association for Computational Linguistics.
- Qiang Ning, Zhili Feng, and Dan Roth. 2017. [A structured learning approach to temporal relation extraction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1027–1037, Copenhagen, Denmark. Association for Computational Linguistics.

- Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. [TORQUE: A reading comprehension dataset of temporal ordering questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, Haoruo Peng, and Dan Roth. 2018a. [Improving temporal relation extraction with a globally acquired statistical resource](#). 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 841–851, New Orleans, Louisiana. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, and Dan Roth. 2018b. [A multi-axis annotation scheme for event temporal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1318–1328, Melbourne, Australia. Association for Computational Linguistics.
- Shantz Pavis. 1998. *Dictionary of the Theatre: Terms, Concepts, and Analysis*. University of Toronto Press.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. 2018. Towards controllable story generation. In *NAACL Workshop*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Peng Shi and Jimmy J. Lin. 2019. Simple bert models for relation extraction and semantic role labeling. *ArXiv*, abs/1904.05255.
- Pradyumna Tambwekar, Murtaza Dhuliawala, Lara J. Martin, Animesh Mehta, Brent Harrison, and Mark O. Riedl. 2019. [Controllable neural story plot generation via reward shaping](#). In *Proceedings of the 2019 International Joint Conference on Artificial Intelligence*.
- Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric P Xing, and Zhiting Hu. 2021. Progressive generation of long text with pretrained language models.
- Ed S. Tan. 1996. *Emotion and the structure of narrative film: Film as an emotion machine*. Mahwah, NJ. Erlbaum Associates.
- Ronald J. Williams and Jing Peng. 1991. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3:241–268.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Hui-Yin Wu, Michael Young, and Marc Christie. 2016. [A Cognitive-Based Model of Flashbacks for Computational Narratives](#). In *AIIDE 2016 - 12th Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 239–245, San Francisco, United States.
- Jingjing Xu, Xuancheng Ren, Yi Zhang, Qi Zeng, Xiaoyan Cai, and Xu Sun. 2018. [A skeleton-based model for promoting coherence among sentences in narrative story generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4306–4315, Brussels, Belgium. Association for Computational Linguistics.
- Peng Xu, Mostofa Patwary, Mohammad Shoeybi, Raul Puri, Pascale Fung, Anima Anandkumar, and Bryan Catanzaro. 2020. [MEGATRON-CNTRL: Controllable story generation with external knowledge using large-scale language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2831–2845, Online. Association for Computational Linguistics.
- Lili Yao, Nanyun Peng, Weischedel Ralph, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343, Florence, Italy. Association for Computational Linguistics.
- Ben Zhou, Daniel Khashabi, Qiang Ning, and Dan Roth. 2019. [“going on a vacation” takes longer than “going for a walk”: A study of temporal commonsense understanding](#). 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 3363–3369, Hong Kong, China. Association for Computational Linguistics.
- Wangchunshu Zhou, Dong-Ho Lee, Ravi Kiran Selvam, Seyeon Lee, Bill Yuchen Lin, and Xiang Ren. 2021. Pre-training text-to-text transformers for concept-centric common sense.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards

story-like visual explanations by watching movies and reading books. In *The IEEE International Conference on Computer Vision (ICCV)*.

Input	[s1] nina needed blood tests done.
prompt	<i><before></i> <i><after></i> <i><before></i> <i><before></i>
baseline	[s2] she was very nervous. [s3] she went to the local clinic. [s4] they told her about her bad breath. [s5] she was relieved her nerves were n't too much to handle. [s2] she was terrified of needles.
VANILLA	[s3] she went to the doctor. [s4] the doctor gave her a needle. [s5] nina was relieved.
RL	[s2] she was afraid to go to the hospital. [s3] she had never been to a hospital before. [s4] but when she got there, she was relieved. [s5] she felt so relieved!
Input	[s1] mickey was always a fan of webcomics.
prompt	<i><after></i> <i><before></i> <i><before></i> <i><before></i>
baseline	[s2] he loved writing his own comics one night. [s3] then he had a great idea. [s4] he created a story called the superman story. [s5] the story was real as a kids cartoon. [s2] one day, he decided to make a webcomic himself.
VANILLA	[s3] he created a website and sent it to a publisher. [s4] the publisher loved it and contacted mickey. [s5] mickey was happy to receive a publishing deal.
RL	[s2] he had been writing for years. [s3] one day, mickey saw an ad for a webcomic. [s4] he signed up for it. [s5] mickey was able to make a lot of money from his work.
Input	[s1] anna had just gotten her hair permed.
prompt	<i><before></i> <i><before></i> <i><after></i> <i><before></i>
baseline	[s2] she was nervous about it. [s3] she was too nervous to do it. [s4] she tried to cut hair to look better. [s5] when she was done it looked great ! [s2] she was excited to see how it looked.
VANILLA	[s3] but then she saw that her hair had grown out. [s4] anna was horrified! [s5] she had to go to the salon to get it fixed.
RL	[s2] she was excited to see how it looked. [s3] but when she went to apply it, it was terrible. [s4] her hair had grown way too long! [s5] anna had to go to the salon to get her hair fixed.

Table 6: Additional generated stories for ROCStories.

A Additional Generated Stories

Please see Table 6 and Table 11 for more examples on ROCStories and WritingPrompt respectively.

B Data

ROCStories (Mostafazadeh et al., 2016a) contains 5-sentence stories. Following (Xu et al., 2020), we split data into 88,344/4,908/4,909 for train/validation/test sets.

WritingPrompt (Fan et al., 2018) contains 30,335 pairs of prompts and stories. With an average of more than 700 words per story, Writing Prompts are much longer than ROCStories. These stories are also much less structured as some dialogues and short phrases may be included. To speed up our experiments, we select stories with a maximum of 500 words, resulting in a total number

of 96,488 training and 5,784 validation prompt-story pairs, respectively. For the test set, we use the 1,000 prompt-story pairs provided by the baseline paper (Goldfarb-Tarrant et al., 2020) for reporting automatic evaluation results.

Pretraining Data. As we mention in Section 3.4, we pretrain storyline models for ROCStories. To be consistent with ROCStories inputs, we divide BookCorpus data (Zhu et al., 2015) into 5 consecutive sentences and filter out those with noisy tokens. We randomly select 1 million such 5-sentence text spans and extract their storylines following Section 3.2.

C More Details for Evaluation Metrics

Automatic evaluation metrics are used to measure textual quality of stories. We report 1) **Ref. PPL**: reference stories' perplexity in a model; 2) **Gen. PPL**: generated stories' perplexity scored by GPT-2 (Radford et al., 2019), i.e. we feed the generated stories into GPT-2 to compute perplexity scores. For diversity scores, we found our models implemented by Huggingface (Wolf et al., 2020) can achieve nearly 0 **Repeat-3** and 100% **Distinct-3** scores, so we follow Goldfarb-Tarrant et al. (2020) to compute the overall vocabulary:token number ratio, which we denote as 3) **Distinct Ratio (%)**. We also report standard 4) **BLEU-3** and 5) **ROUGE_L** scores.

D More Details for Baseline Models

MEGATRON-CNTRL Xu et al. (2020), denoted as MEGATRON for brevity, is chosen as the baseline as it outperforms previous systems such as Guan et al. (2020) on ROCStories. We do not perform delexicalization that replaces names and entities with *[MALE]*, *[FEMALE]* and *[NEUTRAL]* tokens, as we found our models work well by recognizing names and entities. When conducting evaluations, we try our best to **map these special tokens back to their original texts** by using the given first sentence. For rare undetermined cases, we manually examine the generated stories and swap in names or entities that make the most sense in the context. To be fair, we compare with the 124M-parameter version.

ContentPlanning (Goldfarb-Tarrant et al., 2020) is chosen as the baseline for WritingPrompt, as it also adopts the Plan-and-Write workflow as well as structured event representations. However, their models are based on BART-large and do not train

with an end-to-end framework. They use 65% of the original training data and also filter out samples with non-[WP] prompts. Our final training data is about 2/3 of theirs.

TemporalBART (Lin et al., 2021) is designed for two event temporal relation related tasks: temporal ordering and event infilling. Although TemporalBART does not tackle story generation directly, it encodes event temporal information via pretraining tasks. So we consider TemporalBART as another baseline model by initialing the storyline model with their parameters and training the STRUCTURED-PROMPT workflow on ROCStories.⁸

E Reproduction Check List

We finetune BART-base. For ROCStories, hyper-parameters are learning rate: $5e^{-5}$; batch size: 10. We use 3 random seeds: (5, 9998, 20016) and report the average performances for all end-to-end models. For Writing Prompts, hyper-parameters are learning rate: $1e^{-4}$; batch size: 64; gradient accumulation: 8.

For ROCStories, we were able to finetune on a single Nvidia GTX2020 GPU with 11G memory, and training time is 3-4 hours per epoch. For WritingPrompt, we have to use a much larger Nvidia A100 GPU with 40G memory, and the training time is 20 hours per epoch. We train all models for 10 epochs and save the model with the best evaluation perplexity. All reproduction details can be found in the separately submitted code.

F Perplexity and Event Coverage Trade-off

One caveat of using end-to-end training is that there is no guarantee that the generated events will appear in the final stories; whereas in two-stage models, the story model learns a mapping from reference storylines to stories, which leads to a higher coverage rate of the generated events. To provide a potential solution, we experiment with the mixture-training method proposed by Zhang et al. (2019),

$$p = \frac{\mu}{\mu + \exp(e/\mu)}$$

where p controls the ratio of reference storylines used in training and e is the training step. Here the larger the hyper-parameter μ , the slower p decays to 0 as training proceeds.

⁸More implementation details such as hyper-parameters and software can be found in the appendix.

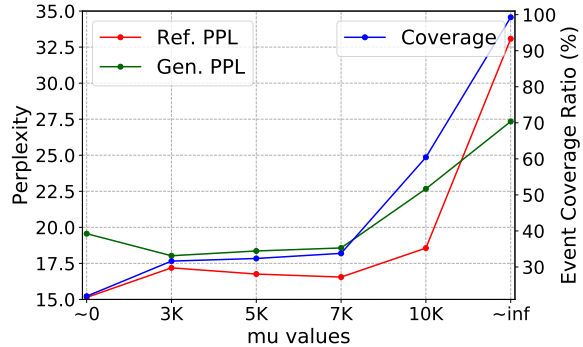


Figure 4: Trade-off between perplexity scores and generated event (trigger) coverage.

In Figure 4, we show the trade-off between perplexity and μ values by training our final RL-based models. When μ is nearly zero, it corresponds to always using predicted storylines, and hence a smaller predicted event coverage rate in the final stories. When μ gets larger, eventually making p close to 1, the training corresponds to always using gold storylines, which leads to a very high event coverage rate, but relatively poor perplexity (still much stronger than two-stage results). We leave the search for optimal μ for future research.

CaTeRS	MATRES
A <i><Before></i> B	A <i><before></i> B
A <i><Identity></i> B	A <i><vague></i> B
A <i><Contains></i> B	A <i><before></i> B
A <i><Overlaps></i> B	A <i><before></i> , <i><after></i> , <i><vague></i> B

Table 7: Label mapping from CaTeRS to MATRES. *<after>* is ignored in CaTeRS by flipping event physical order.

G Benchmark Event Temporal Relation Annotations

The experimental results in the main text demonstrate the effectiveness of using temporal prompts. Here, we further show that the tool to produce temporal prompts, i.e. ECONET, provides reliable event temporal relation annotations. We benchmark ECONET’s performances using CaTeRS (Mostafazadeh et al., 2016b), which annotates 4 types of temporal relations for event pairs in a small amount of ROCStories. However, CaTeRS’s annotations are based on event time interval rather than event start time as used in MATRES, which ECONET is finetuned on.

In Tabel 7, we provide a mapping from CaTeRS’s temporal relations to MATRES labels.

The only non-unique mapping is $\langle \textit{Overlap} \rangle$. In other words, when ECONET predicts $\langle \textit{before} \rangle$ for a CaTeRS sample $\langle \textit{Overlaps} \rangle$, we have to manually examine whether it is correct or not. We found that when ECONET predicts $\langle \textit{before} \rangle$ for CeTeRS data, the precision rate is 65.53% due to a large amount of $\langle \textit{Overlaps} \rangle$ event pairs being predicted as $\langle \textit{before} \rangle$. But we emphasize here that this low number is **caused by label mismatch** as shown in Table 7, which does not truthfully reflect the ECONET’s accuracy.

To have a better understanding, we randomly selected 20 such pairs and manually examine their temporal relations in the context and found that 90% of such pairs are indeed correctly predicted by ECONET. Adjusting for this factor, the precision rate for the annotated $\langle \textit{before} \rangle$ relation would be 92.07%, indicating highly accurate predictions. We do not claim the final accuracy is 92.07%, but simply argue that the annotations provided by ECONET are helpful as our main experimental results demonstrate.

H Two-stage Model Results

Two-stage Model. As we mentioned in Section 3, another way to implement Plan-and-Write framework is to train storyline and story models separately with gold input and outputs, and replace story models’ inputs with the storyline model’s predictions during inference. We found this variant’s performances fall far behind other compared models. So we do not use them in human evaluations and simply show their automatic evaluation results in Table 8.

Models	Ref. PPL(↓)	Gen. PPL(↓)	Distinct Ratio(↑)	BLEU (↑)	ROUGE _L (↑)
ROCStories	197.97	27.15	4.26	2.98	16.39
WritingPrompt	76.69	39.43	8.12	0.00	6.24

Table 8: Automatic Evaluation Results for the Two-stage models.

I Storyline Model Results

Our primary goal is to improve temporal control in final stories. However, as aforementioned, the motivation for using reinforcement learning is to allow the storyline model to adapt together with the story model. In Table 9, we compare storyline prediction performances between the vanilla and the RL-based end-to-end models. These results suggest that our final framework combining structured storylines, temporal prompts, pretraining and

Models	Ref. PPL(↓)	BLEU (↑)	Distinct Ratio (↑)
Vanilla-Gen + RL	13.97	3.21	3.68
	12.71	3.49	4.09

Table 9: Improvements of RL-based model against vanilla end-to-end model on ROCStories data.

RL also helps storyline models fit reference storylines better, resulting in lower perplexity, higher overlapped scores and better diversity.

J Error Analysis

Table 10 shows examples for the temporal diversity in Section 6.

Input 1	[s1] justin is a very charismatic man who loves women.
baseline	[s2] he always said that he wanted to be a man. [s3] he thought that men looked bad. [s4] justin decided to become a boy scout. [s5] he is now a successful realtor. predicted relations: $\langle \textit{vague} \rangle$ $\langle \textit{vague} \rangle$ $\langle \textit{vague} \rangle$ $\langle \textit{before} \rangle$
RL	[s2] one day he met a beautiful woman and they fell in love. [s3] justin decided he wanted to marry her. [s4] he went to the mall and bought a ring. [s5] he proposed to her and she said yes. predicted relations: $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$
Input 2	[s1] the smoke was beginning to get to me.
baseline	[s2] mary friend said he had to come to my house. [s3] i wanted to have a smoke fight. [s4] he said i had to light a cigarette. [s5] i smoked for a few hours before i gave up. predicted relations: $\langle \textit{after} \rangle$ $\langle \textit{vague} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$
RL	[s2] i had just gotten out of my car. [s3] i ran to the front door. [s4] i opened the door. [s5] the fire department came and put out the fire. predicted relations: $\langle \textit{vague} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$
Input 3	[s1] tom was in court for a parking ticket.
baseline	[s2] he did not have to pay his court date. [s3] the judge gave him a parking ticket. [s4] tom paid his court date. [s5] the judge gave tom a parking ticket. predicted relations: $\langle \textit{before} \rangle$ $\langle \textit{vague} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{after} \rangle$
RL	[s2] he had been speeding for a while. [s3] the judge gave him a warning. [s4] tom didn’t listen to the warning. [s5] he got a ticket anyway. predicted relations: $\langle \textit{after} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{before} \rangle$ $\langle \textit{after} \rangle$

Table 10: Examples from our RL-based model that are not as temporally diverse as examples generated by the baseline MEGATRON model. **Predicted relations** are provided by human evaluators.

K Potential Risks

Since our models deal with open-domain generation, it is conceivable that the generated stories could contain biases, malicious languages, and hallucinations. We refer readers to the work in fairness and fact checking to address these issues.

Input	artic outpost 9
prompt	<i>(after)</i> <i>(vague)</i> <i>(before)</i> <i>(after)</i> <i>(before)</i> <i>(vague)</i> <i>(vague)</i>
baseline	she stood alone in the dimly lit corridor , staring at the light that flooded into the room . her eyes scanned the room , finally settling on the small , round table . ꞑꞑꞑ ꞑꞑꞑ “ hello , everyone . ” ꞑꞑꞑ the woman spoke in a voice that echoed the room in a series of hushed tones , as though it was a voice only audible to those who heard it . “ do you know why we are here ? do you know where we come from ? ” our leader , the man with the black suit and black hat , stood at the end of the hall . a man in a black coat and black trousers , a long coat , and a black hat stood behind him , a black cloak and a dark hat , a pair of gloves . the black suit had a black mask on , and she was wearing her best , black dress shirt and a pair of pants , and black shoes . on her right wrist , the small of one shoulder , was an amulet , and on the other , a small device hidden in the pocket of the man ’s coat , a single small device . there was nothing else . just this man . . . and our leader
VANILLA	i’ve never seen anything like it before. it’s a strange place. i don’t know why i chose it, but i’ve always wanted to see it. it was the first time i saw it. it was a small building with a single window. i couldn’t believe it. i had never seen it before, but it was here. i’m not sure what it is, but there’s something about it that makes me want to run. i can’t remember the last time i ’d seen it. but i do know that it was there. the first thing i noticed was the door opening. i wasn’t sure if it was my own, or my own. i didn’t want to be here. i just wanted to get out. i wanted to go back to my home. as i approached the door, i saw a man standing there. he was wearing a trench coat. he had a gun pointed at his head. he pointed it at my head. “ what are you doing here? ” i asked. he looked at me. he looked like he was about to say something. “ i need you to come with me. ” he said, and i followed him. we were in the middle of nowhere when the man shot me in the head. i looked at him, and he looked back at me, then at the gun. i felt a sharp pain in my chest. i started to cry. i tried to speak, but my throat was dry. i heard the man say something, and then he shot me. i fell to the ground, bleeding. i woke up in a hospital bed. there was a note on my bedside table. i opened it and read it aloud. *i’m sorry, but this isn’t what i was expecting.
RL	what’s that? we’ve been waiting for you for a long time, i said, i’m not sure how long you’ll be waiting for me. you don’t have much time to think about it, do you? i mean, i can see it, but i can feel it in your eyes, and i can hear it in the way you can hear me. i can smell it, and it can feel your heart beating faster than you can think. i think you know what i am talking about. i know what you are talking about, but you can not understand what i mean. i want you to know that i love you. i need you to understand that you are my love, and that i will always be there for you. that you will always love me, that i am the love of my life, and you are the love that i have always wanted to be, and i will never let you down. i will not let you go. i promise you that. with that, i walked out of the building, into the street, and into the night.
Input	100000 year high
prompt	<i>(before)</i> <i>(after)</i> <i>(before)</i> <i>(before)</i>
baseline	for a long time it seemed the world ’s people all were living in the same way . the same . they were all the same . the same , they had the same families and their children and they all had the same friends . but these families were all different . they may have been the last of them , and the world was a better place without it , but those people knew about it , they just did n’t know the truth . some of them lived for years , some of them a century . and a new age started , one that would make this world a better places , a world without war or war . it seemed the first people could see the world and see the future and people living it . they were the first to do it . ” “ the first people on earth who could live a normal life ! the first to have a normal family and a normal home ! the first ones who we could live on ! who could have imagined that our world was as good as ours was and that the next life would be better than what i ’ ve been , that there ’s a whole new life ahead of us , that this world is a living one ! ” and that ’s when it happened . that was the day . on march 20th , 2036 . march 20, 2036 , the first time that day
VANILLA	it was a beautiful day. the sun was shining and the sky was blue. the clouds were bright and the trees were green and the birds were singing. it was a wonderful day. i was walking down the street when i saw it. the sky above me was blue and it was beautiful. i had never seen anything like it before. i looked down at the ground and saw nothing. i didn’t know what to do. i took a step back and looked up at the sky. i couldn’t see anything at all. i walked towards the sky and looked at the stars. they were beautiful. they seemed like they were floating. i reached out and touched them. i felt a warmth in my chest. i touched them again and again. it was amazing. i wasn’t sure if i was dreaming or not, but i knew i was.
RL	it’s been a long, long time since i’ve seen the light of the sun. i’m not sure how long it has been, but i think it’s a long time. i ’d like to say that the last time i saw the light, it was a thousand years ago, but it wasn’t. it was only a few years ago that i first saw the sun, when it was just a few seconds away from reaching the horizon. i don’t know how long i have been seeing it, but the light is blinding, and i can’t help but wonder what it could have been like. maybe i ’ll see it again, but for now, i just want to see the light again.
Input	paint a picture of loneliness , without using the word , or any of its forms .
prompt	<i>(before)</i> <i>(vague)</i> <i>(vague)</i> <i>(vague)</i> <i>(before)</i> <i>(before)</i> <i>(before)</i> <i>(after)</i> <i>(after)</i> <i>(before)</i>
baseline	the glass is cold , but the sun shines brightly through . on the glass , in a small room , a woman stands . it stands before her , its hands clutching the glass in front of it like a shield . it 2019 s not the sunlight , but it hits the sky like a storm . the light makes the clouds roll in from all directions . a soft breeze , the wind blowing through the trees , the water rustling in the water . there are people in the room , but most do n’t look at their feet , they ’re too busy staring out the window . some are silent . most look at the woman . one of the people looks back at the woman who seems to be alone in this room , who is alone with a glass of champagne . the wind and water make the trees sway . you hear the birds chirping from outside . a few laugh , some are sad , some angry , but most are happy , and some are happy . there ’s a small child sitting on her lap . in the corner . on the ground , on a bench . and it ’s just this woman and the glass .
VANILLA	i’ve always been lonely. i’m not sure why. maybe it’s the fact that i don’t have any friends. maybe i’ve been alone for a long time. i haven’t seen anyone. i can’t even remember where i am. i have no friends. there’s no one to talk to. no one who can help me. i just want to be alone. sometimes i think that i ’ll just be alone forever. but then i think back to the first time i saw someone else. it was a girl. she’s beautiful, but she looks like she doesn’t belong here. she has no friends, no family, no one. then i thought of her. i think of the way she looked at me. the way i looked at her, the way her eyes looked at the people around her. that’s when i realized that i was alone.
RL	i don’t know why i’m doing this, but i’ve been doing it for a long time now. it’s been a long, long time since i ’d seen a single person. i remember the first time i saw her. she was beautiful, and i loved her. but i couldn’t tell you how much i missed her. i didn’t want to know how much she meant to me. and i can’t even remember the last time i talked to her, because i was too busy trying to figure out what to say, she was always there for me, and she always made me feel so alone. sometimes i wish i could go back, but sometimes i think i ’ll just stay here forever. maybe someday, when she’s gone, she ’ll come back.

Table 11: Additional examples for WritingPrompt. Possible *(after)* relations are highlighted.

Here is an overview of your task,

1. Judge temporal relations between each pair of neighboring events.
2. Check if the sequence of events in a story contradicts to your commonsense, particularly temporal commonsense.

Always review the instruction page if you are uncertain about a task!

Story 1

[Sentence 1]: \${sentence1_1}

[Sentence 2]: \${sentence1_2}

[Sentence 3]: \${sentence1_3}

[Sentence 4]: \${sentence1_4}

[Sentence 5]: \${sentence1_5}

Q1: Which event started/starts earlier between [Sentence 1] and [Sentence 2]?

- [Sentence 1]: \${sentence1_1}
- [Sentence 2]: \${sentence1_2}
- Not sure

Q2: Which event started/starts earlier between [Sentence 2] and [Sentence 3]?

- [Sentence 2]: \${sentence1_2}
- [Sentence 3]: \${sentence1_3}
- Not sure

Q3: Which event started/starts earlier between [Sentence 3] and [Sentence 4]?

- [Sentence 3]: \${sentence1_3}
- [Sentence 4]: \${sentence1_4}
- Not sure

Q4: Which event started/starts earlier between [Sentence 4] and [Sentence 5]?

- [Sentence 4]: \${sentence1_4}
- [Sentence 5]: \${sentence1_5}
- Not sure

Q5: Is the event sequence in this story consistent with your general commonsense?

- Yes. All events are consistent.
- No.

Q6: Is the event sequence consistent to your temporal commonsense?

- Yes. All events are temporally consistent
- No.

Story 2

[Sentence 1]: \${sentence2_1}

[Sentence 2]: \${sentence2_2}

Figure 5: Annotation interface for pairwise event temporal relations and temporality.

[Task 1] Judge temporal relations between each pair of neighboring events.

Instructions: In this task, we want you to focus on the **temporal order of events**. Occasionally, you may encounter event sequences that do not make sense. For example, "I went to class this afternoon. I found my puppy was giving a speech in class." Although this sequence of events contradict to our commonsense, we want you to **imagine the situation where if these events were to happen**, which event started earlier? In this case, "I went to class this afternoon" started earlier. In [Task 2], we will ask you to check events' general commonsense and temporal commonsense. Again, please try to **evaluate events' temporal order by assuming all events in a story indeed happen** in **Q1 - Q4** below.

Q1: Which event **started/starts earlier** between **[Sentence 1]** and **[Sentence 2]**?

- [Sentence 1]:** tom was about to have a little brother.
- [Sentence 2]:** he was nervous because he didn't know what to name him. [✓]
- Not sure [✓]

Explanation: The event in **[Sentence 1]** is "was about to have," while the events in **[Sentence 2]** are "was nervous" and "didn't know what to name." From the context, it is relatively clear that "was about to have" started earlier than "was nervous." However, the second event "didn't know" is subject to personal interpretation. One way to reason about "didn't know" is that **Tom probably always lacks knowledge or creativity, so "didn't know" is a long-lasting status for him, in which case "didn't know" started earlier than "was about to have,"** and thus, pick the second option. Another way to interpret this is that **"didn't know" started at the time when "he was about to have a little brother," in which case we are not sure which event started earlier.** So pick the last option. When you are working on the real tasks and are uncertain about exactly which event started earlier, "not sure" is a safe option. **But please do carefully consider option 1 and 2 before you pick option 3.**

Q2: Which event **started/starts earlier** between **[Sentence 2]** and **[Sentence 3]**?

- [Sentence 2]:** he was nervous because he didn't know what to name him. [✓]
- [Sentence 3]:** he went to the doctor.
- Not sure

Explanation: The events in **[Sentence 2]** are "was nervous" and "didn't know," while the event in **[Sentence 3]** is "went." **It is clear from the context that both events in [Sentence 2] started earlier. So you should pick the first option here.**

Q3: Which event **started/starts earlier** between **[Sentence 3]** and **[Sentence 4]**?

- [Sentence 3]:** he went to the doctor. [✓]
- [Sentence 4]:** the doctor gave him a name.
- Not sure

Explanation: The events in **[Sentence 3]** is "went" while the event in **[Sentence 4]** is "gave." **Based on the context, it is also clear that "went" started earlier than "gave," so pick option 1 here.**

Q4: Which event **started/starts earlier** between **[Sentence 4]** and **[Sentence 5]**?

- [Sentence 4]:** the doctor gave him a name. [✓]
- [Sentence 5]:** tom named him james.
- Not sure

Explanation: **It is clear from the context that "gave him a name" started earlier than "named him james." So also pick the first option in this case.**

Figure 6: Instruction for pairwise event temporal relation annotation.

[Task 2] Check Event Temporal Commonsense

Instructions: Upon finishing [Task 1], one should have a general picture about the temporal orders of events in the above story (**Story 1**), which helps us understand how the story unfolds. The majority of event sequences you will see in this research project are coherent with our commonsense. However, a few of them do not, and we need you to double check for us.

Definitions: Temporal commonsense belongs to broader human commonsense including physical, scientific, societal and temporal commonsense. It's possible that event sequences may contradict to some human commonsense, but not the temporal commonsense. The example we show above "I went to class this afternoon. I found my puppy was giving a speech in class." contradicts to our scientific commonsense, but if we assume both events indeed happen and evaluate from a pure temporal perspective, this sequence still makes sense ("giving a speech" happens following "went to class").

An example that contradicts to our temporal commonsense would be "I put the sandwich in the microwave. I forgot to turn the microwave on. I burned my sandwich." Here, all individual events can happen in our lives, but "burned sandwich" shouldn't happen following "forgot to turn microwave on."

Here are more examples on temporal commonsense,

- "I had a big meal. Then I felt full." The event order would be temporally wrong if we say "I felt full. Then I had a big meal."
- "I was poor. So I get a well-paid job." The event order would be temporally wrong if we say "I get a well-paid job. So I was poor."

For the next two questions, based on **Story 1** presented above,

- **Q5** asks you to detect any event or event sequence that doesn't make sense to you. We do not need you to be precise on which commonsense an event or event sequence contradicts to. Simply follow your intuition and answer "yes" or "no."
- **Q6** is our main focus, where we request you to be thoughtful regarding temporal commonsense. Again, assuming each event can individually happen, does their temporal order make sense?

Q5: Is the event sequence in this story consistent with your general commonsense?

Yes. All events are consistent. No. [✓]

Explanation: In **Story 1**, it says "Tom went to the doctor to get a name for his little brother." This likely contradicts to our societal commonsense as we usually go to doctor for medical reasons. So you can select "No" here. But if you feel it's reasonable to see doctor if one is nervous; then option 1 is also fine. Again, we don't need you to be very precise for this question.

Q6: Is the event sequence consistent to your **temporal** commonsense?

Yes. All events are temporally consistent. [✓] No.

Explanation: **Q5** above should help you think whether the contradiction comes from temporal commonsense or not. In this case, it is clearly not. So pick the first option.

Figure 7: Instruction for temporality annotation.

We first describe how to determine the **unexpected level** of a story. We will explain it with the examples below. You are given stories consisting of 5 sentences.

1. When you see [Sentence 1], pretend you don't know the rest of the story and predict what are some possible events in the next sentence, or [Sentence 2]. Then evaluate if the events in [Sentence 2] meet your expectation.
2. **Definition:** A valid **unexpected event** occurs when it **surprises** you, but it **makes sense for this event to follow** [Sentence 1].
3. Repeat the same exercise for Sentence 2, 3, 4 and 5 by imagining you don't see the rest of the story.

When you are done with the entire story, reflect on how many **unexpected events** you saw and to what degree they surprised you. The more unexpected events, the higher the **unexpected level**.

Here is an overview of your task. There are 5 stories (each with 5 sentences) provided in this assignment.

1. Read each story carefully and provide a score from 1 - 5, with 5 indicating the most unexpected story.
2. Try to provide ***unique*** scores for different stories by comparing all 5 stories (the most unexpected gets 5; the least unexpected gets 1), but identical scores are allowed if you feel two stories are equally unexpected.

Example 1

[Sentence 1]: tom was about to take the final exam.
[Sentence 2]: he was nervous.
[Sentence 3]: so he felt sick.
[Sentence 4]: he went to the doctor.
[Sentence 5]: the doctor gave him some pills.

Example 2

[Sentence 1]: tom was about to take the final exam.
[Sentence 2]: he was nervous.
[Sentence 3]: so he felt sick.
[Sentence 4]: he went to the doctor.
[Sentence 5]: tom was arrested.

In Sentence 5 of Example 2, the last event "tom was arrested" is surprising. However, this contradicts to our definition above that a valid unexpected event must be both **surprising** and **logical**. From the context of Example 2, we don't see any reason why tom could get arrested. So based on our definition, Example 2 isn't more unexpected than Example 1, given Sentence 1-4 are identical in these two examples. You should also give lower score to Example 2 because it is unreasonable.

The stories can be more complicated than our examples. Please use your best judgement, but always keep in mind of our **unexpected event** definition: "being surprising but logical."

Story 1

[Sentence 1]: \${sentence1_1}
[Sentence 2]: \${sentence1_2}
[Sentence 3]: \${sentence1_3}
[Sentence 4]: \${sentence1_4}
[Sentence 5]: \${sentence1_5}

Figure 8: Annotation interface for interest level (unexpectedness).