Cue Me In: Content Inducing Approaches to Interactive Story Generation

Faeze Brahman

Alexandru Petrusca

Snigdha Chaturvedi

Introduction

What?

- Local and global coherence
- Thematic consistency
- Creativity

Why?

- Addressing literary learning needs
- Explore creative writing at an early age
- Collaborative Authorship

Motivation

- ⭐ Existing SOTA models for story generation generate a story in one go with limited initial input from a user.
- ⭐ When humans write, they incrementally edit and refine the text they produce.

We explore the problem of interactive story generation where a user provides cue phrases to guide the model generation.

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup :Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Problem Statement

Prompt John walked into the kitchen.

Cue phrase dark

It was **dark** and cold.

light switch

He turned on the **light switch**.

but

But then John smelled something burning.

checked the oven

He **checked the oven** to find the smell.

Generating one sentence at a time.

Problem Formulation

Context: $\mathbf{x} = \{x_1, x_2, \ldots, x_T\}$

Cue phrase: $\mathbf{c} = \{c_1, c_2, \ldots, c_k\}$

Next sentence: $\mathbf{Y} = \{y_1, y_2, \ldots, y_M\}$

When generating n-th sentence, the model take all n-1 sentence as context along with cue phrase

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup :Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup :Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Model Architecture

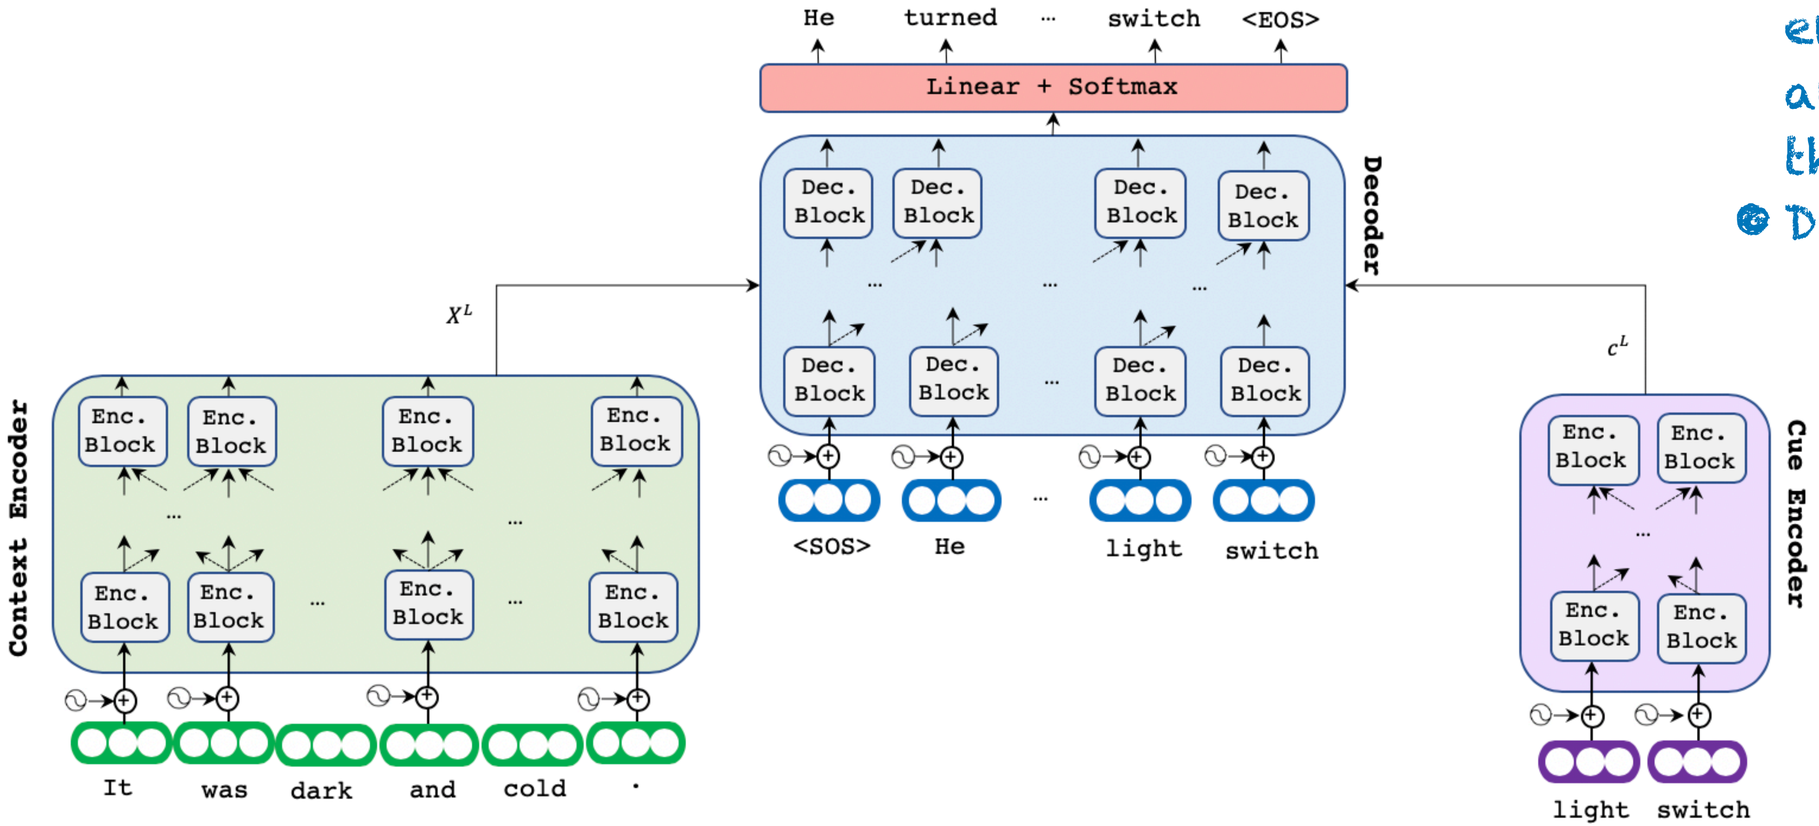

- Use the same encoding mechanism and differ only in their decoders.
- Dual encoding

Dual Encoders

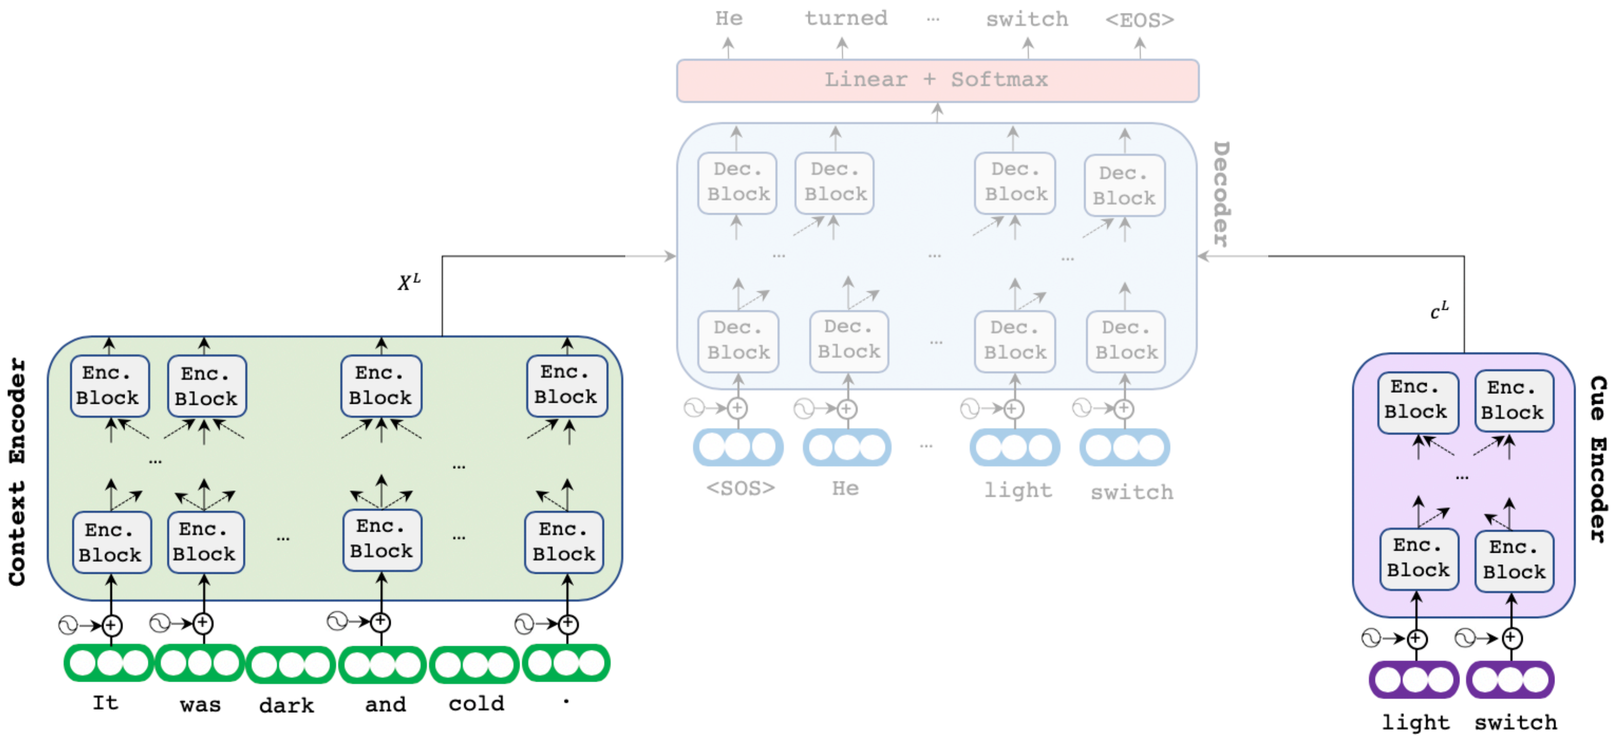

$$X^{l+1} = \text{ENCBLOCK}(X^l, X^l, X^l)$$

$$c^{l+1} = \text{ENCBLOCK}(c^l, c^l, c^l)$$

Content-Inducing Decoders — Cued Writer

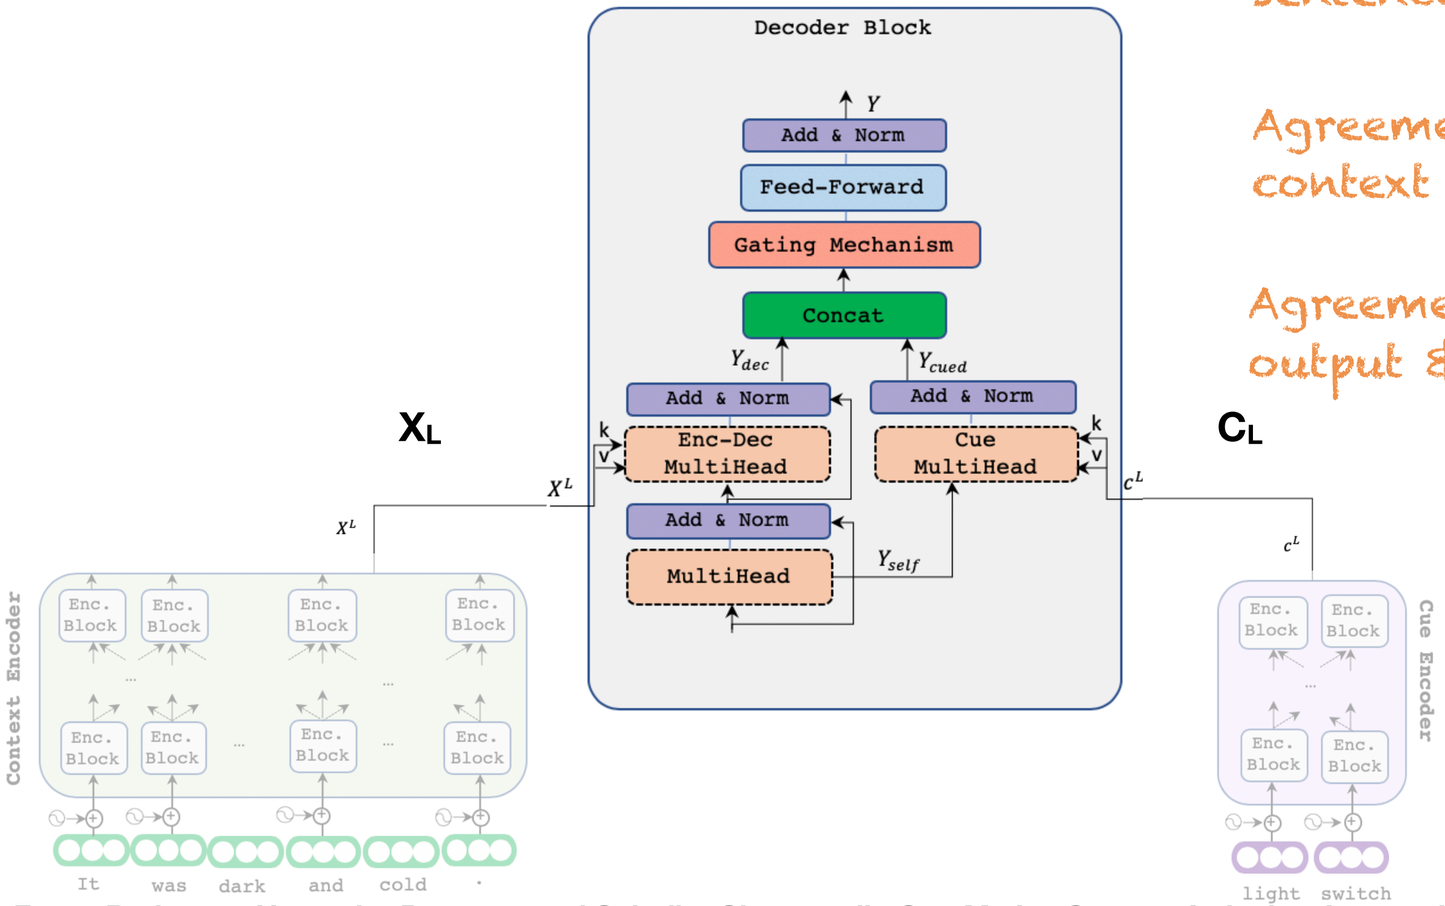

Intra-sentence sentence agreement

$$Y^{l+1}_{self} = \text{MULTIH}(Y^l, Y^l, Y^l)$$

Agreement between context & output

$$Y^{l+1}_{dec} = \text{MULTIH}(Y^{l+1}_{self}, X^L, X^L)$$

Agreement between output & cue phrase

$$Y^{l+1}_{cued} = \text{MULTIH}(Y^{l+1}_{self}, c^L, c^L)$$

Faeze Brahman, Alexandru Petrusca and Snigdha Chaturvedi . *Cue Me In*: Content-Inducing Approaches to Interactive Story Generation . AACL 2020

Content-Inducing Decoders — Cued Writer

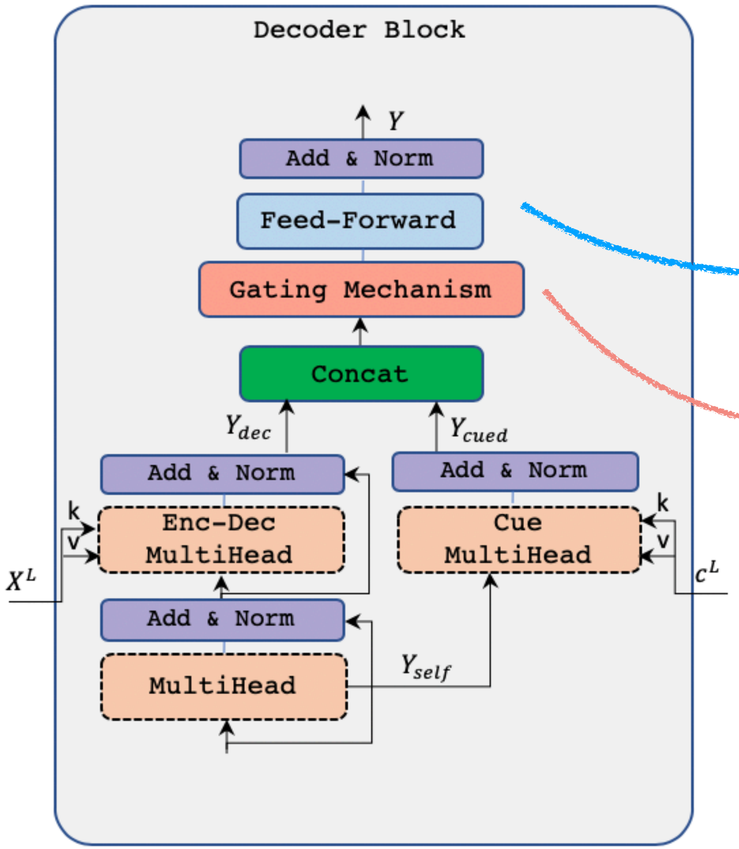

Feed-forward with ReLU

$$Y^{l+1} = \mathrm{FFN}(Y_{int}^{l+1})$$

Integrate semantic representations

$$g^{l+1} = \sigma(W_1[Y_{dec}^{l+1}; Y_{cued}^{l+1}])$$

$$Y_{int}^{l+1} = W_2(g^{l+1} \circ [Y_{dec}^{l+1}; Y_{cued}^{l+1}])$$

Faeze Brahman, Alexandru Petrusca and Snigdha Chaturvedi . *Cue Me In*: Content-Inducing Approaches to Interactive Story Generation . AACL 2020

Content-Inducing Decoders — Relevance Cued Writer

Captured relevance provided to decoder

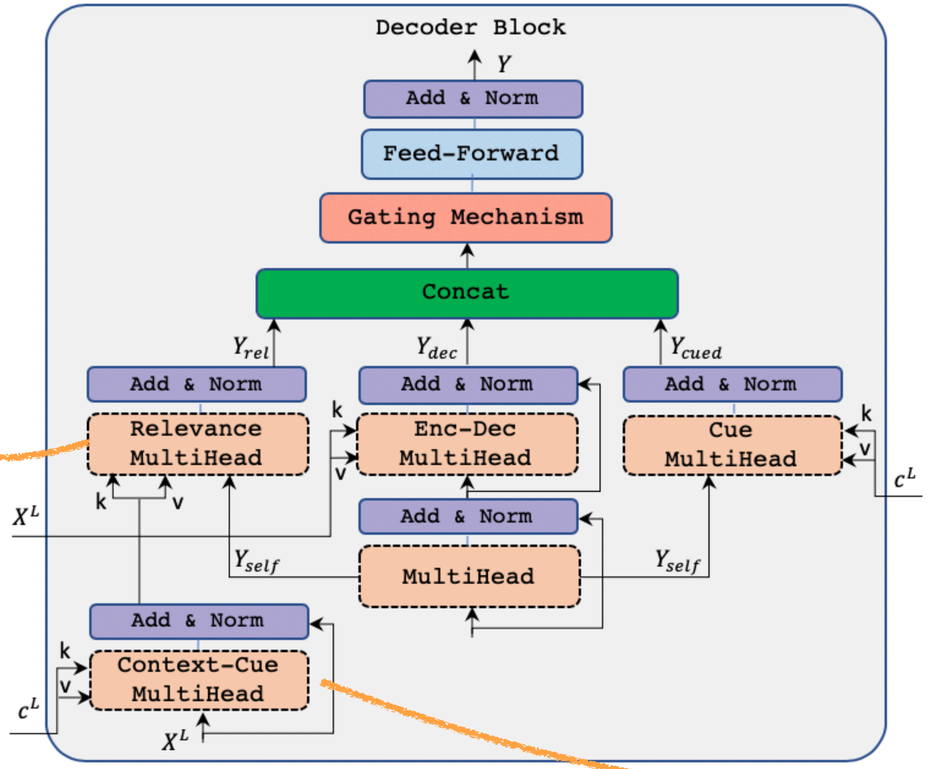

Characterize the relevance btw context & cue phrase

Faeze Brahman, Alexandru Petrusca and Snigdha Chaturvedi . *Cue Me In*: Content-Inducing Approaches to Interactive Story Generation . AACL 2020

Decoding and Objective

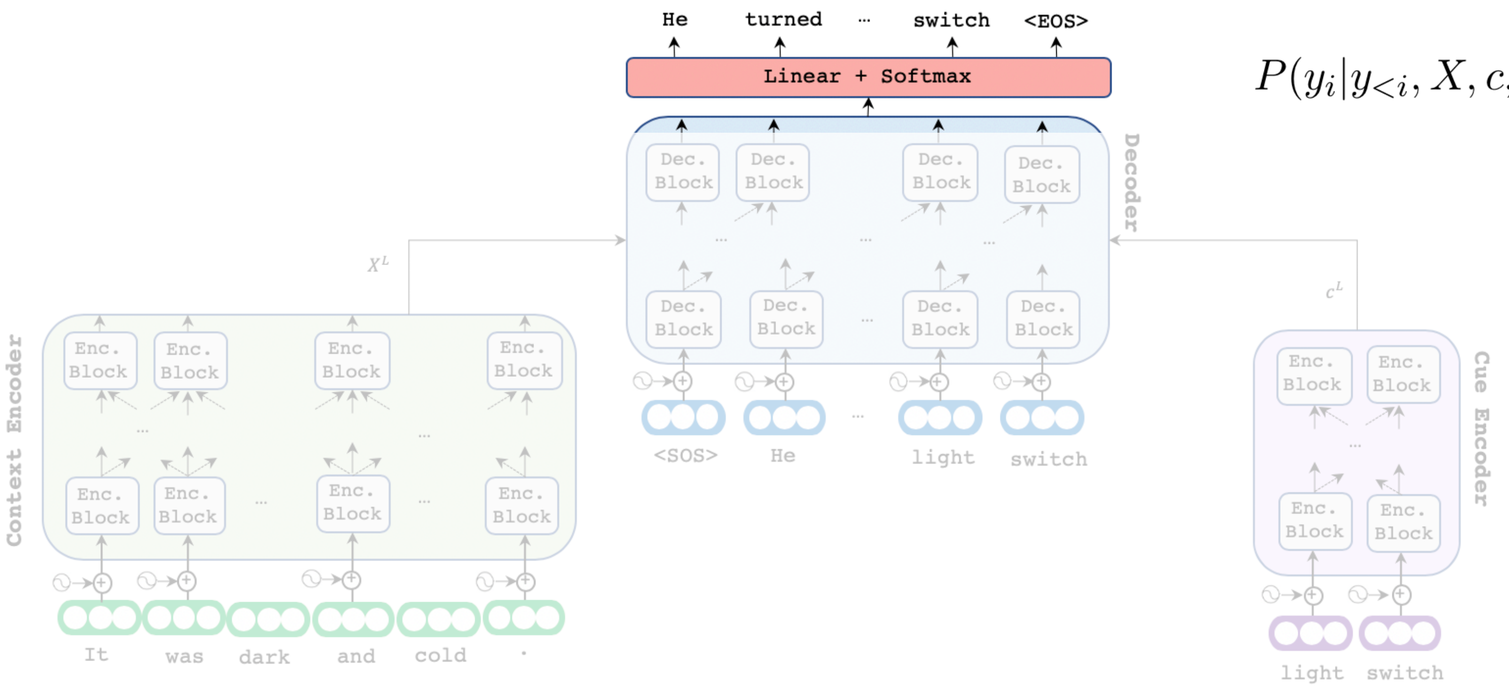

$$P(y_i|y_{<i}, X, c, \theta) = \text{softmax}(Y_i^L W_y)$$

Likelihood of predicting y_i given the preceding text $y_{<i}$

Cross-entropy loss:

$$-\sum_{i=1}^{M} log\ P(y_i|y_{<i}, X, c, \theta)$$

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup :Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup : Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Training Details

- **Dataset:** ROCStories 98,161 five-sentence long stories
- **Cue phrases:** To scale model training, we use RAKE algorithm to automatically extract cue phrases from target sentence.

Evaluation

- **Baseline:**
  - ✦ SEQ2SEQ
  - ✦ Dynamic (Yao et al. 2019):
  - ✦ Static (Yao et al. 2019):
  - ✦ Vanilla:

Evaluation

- **Automatic Eval.**
  - ✦ Perplexity
  - ✦ BLEU (1-2-3)
  - ✦ Greedy-Match (GM)
  - ✦ Repetition-4
- **Human Eval.**
  - ✦ Story-level
  - ✦ Sentence-level

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup :Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Overview

Task: Interactive Story Generation

Models: Content-Inducing Decoders

Experimental Setup : Dataset, Training Details, Evaluation Measures

Empirical Results: Automatic / Human

Results — Automatic Eval.

Models	PPL (↓)	BLEU-1 (↑)	BLEU-2 (↑)	BLEU-3 (↑)	GM (↑)	Repetition-4 (↓)
DYNAMIC [35]	29.49	30.05	9.16	4.59	0.73	44.36
STATIC [35]	20.81	33.25	9.64	4.77	0.75	26.26
SEQ2SEQ	20.97	33.91	10.01	3.09	0.82	33.23
VANILLA	15.78	40.30	16.09	7.19	0.89	20.87
Cued Writer	14.80	41.50	16.72	7.25	0.92	**15.08**
Rel. Cued Writer	**14.66**	**42.65**	**17.33**	**7.59**	**0.94**	16.23

☑ Our models outperform all baselines.

Both models outperform Dynamic and Static

Additive gain from modeling relevance

Stories more related to cue phrases

Proposed models significantly outperform all baselines

Results — Automatic Eval.

The two proposed models are almost comparable but show a general superior performance.

Inter-story repetition

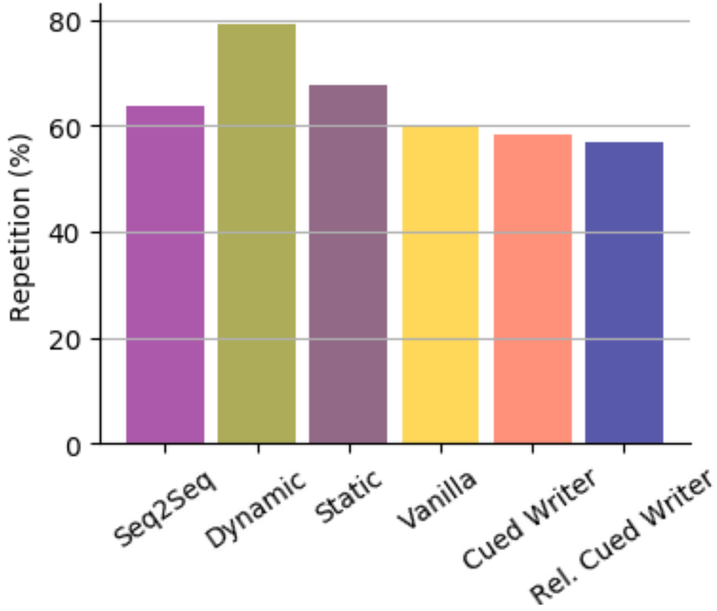

Intra-story repetition

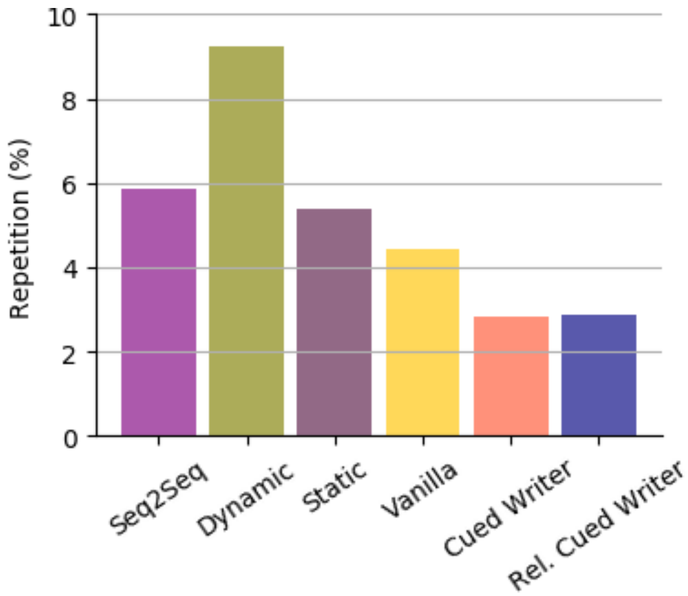

Results — Human Eval.

Percentage of preference for our stronger model, Rel. Cued Writer, over other baselines

Story-level

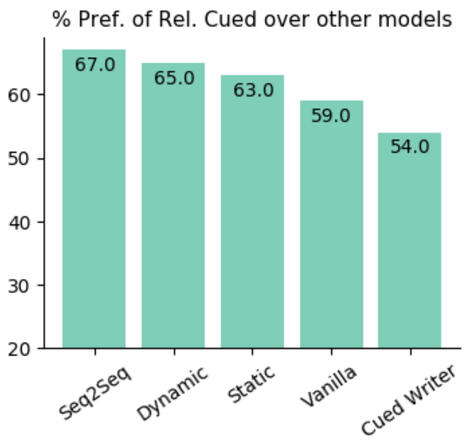

Sentence-level

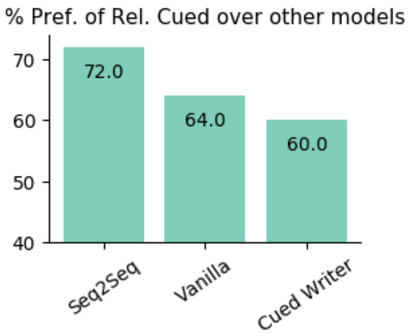

☑ Our stronger model, Rel. Cued Writer, was preferred over other baselines in both experiments.

Qualitative Results

Prompt: Harry really wanted to see the new Star Wars movie.
Cue phrases: opening day - ticket - entire movie - really

Static

One day, **she** was in a hurry to get to the theater. she wasn't sure if she would like to go. She was excited to try the new movie. **She was so excited** to see the movie. **[opening day] [ticket]**

Vanilla

On opening day, Harry was very nervous. He bought a ticket **to** the theater. He bought **Harry** ticket **tickets** to the theater. He really didn't like the movie. **[entire movie]**

Rel. Cued Writer

He decided to watch it on opening day. He got to the theater and got a ticket. He watched the entire movie. He was really excited about it.

☑ Our novel Content-Inducing approach generates stories that suffer less from listed issues.

Summary

- ⭐ We explored the problem of interactive storytelling, which leverages human and computer collaboration for creative language generation.
- ⭐ We propose two content-inducing approaches to incorporate additional information into the generation phase.
- ⭐ We evaluate the utility of our approaches through several experiments and demonstrate that our methods outperformed competitive baselines.
- ⭐ We suggest other significant aspects of story generation, such modeling of discourse relations, and representation of key narrative elements, which could be possible future directions.

Thank you!
Any question?