Controllable Neural Story Plot Generation via Reinforcement Learning

Pradyumna Tambwekar, Murtaza Dhuliawala,
Lara J. Martin, Animesh Mehta, Brent Harrison,
and Mark O. Riedl

Neural *story plot generation* can be **controlled** via guiding a language model through intermediate plot points toward a **desired goal**

Motivation

Automated Plot Generation is the problem of creating a sequence of main plot points that create a story.

Existing story and plot generators lack **controllability**—the ability to receive guidance to achieve a particular goal.

Simply sampling from a probability distribution creates aimless stories that lack plot **coherence.**

Reward Function

Reward a seq2seq model when it moves the story progressively toward a target

1. Compute reward for all verbs using:
   - **Frequency** before **target verb**

$$r_2(v) = \log \frac{k_{v,g}}{N_v}$$

   - **Distance** from **target verb**

$$r_1(v) = \log \sum_{s \in S_{v,g}} l_s - d_s(v, g)$$

2. Create **clusters** of verbs based on their **reward values**

$$R(v) = \alpha \times r_1(v) \times r_2(v)$$

using Jenks Natural Breaks

3. **Constrain output verb** selection to the **next cluster.**

➡ Use the **REINFORCE** method to backpropagate the reward

https://arxiv.org/abs/1809.10736

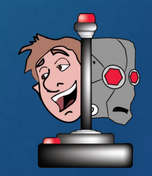

entertainment intelligence lab

Goal

Find a **coherent** sequence of **events** that results in a world state where a **desired goal** holds.

RL allows to address the **control** and **coherence** in plot generation

Contributions

We present a **policy gradients** approach to plot generation.

Our novel **reward shaping** technique facilitates **intermediate** rewards, enabling the system to progress towards a **specific goal**.

Results

Goal	Model	Goal Achievement Rate	Average Perplexity	Average Story Length
admire	Seq2Seq	35.52%	48.06	7.11
	DRL	15.82%	**5.73**	7.32
	DRL + Clustering	**94.29%**	7.61	4.90
marry	Seq2Seq	39.92%	48.06	6.94
	DRL	24.05%	9.78	7.38
	DRL + Clustering	**93.35%**	**7.05**	5.76