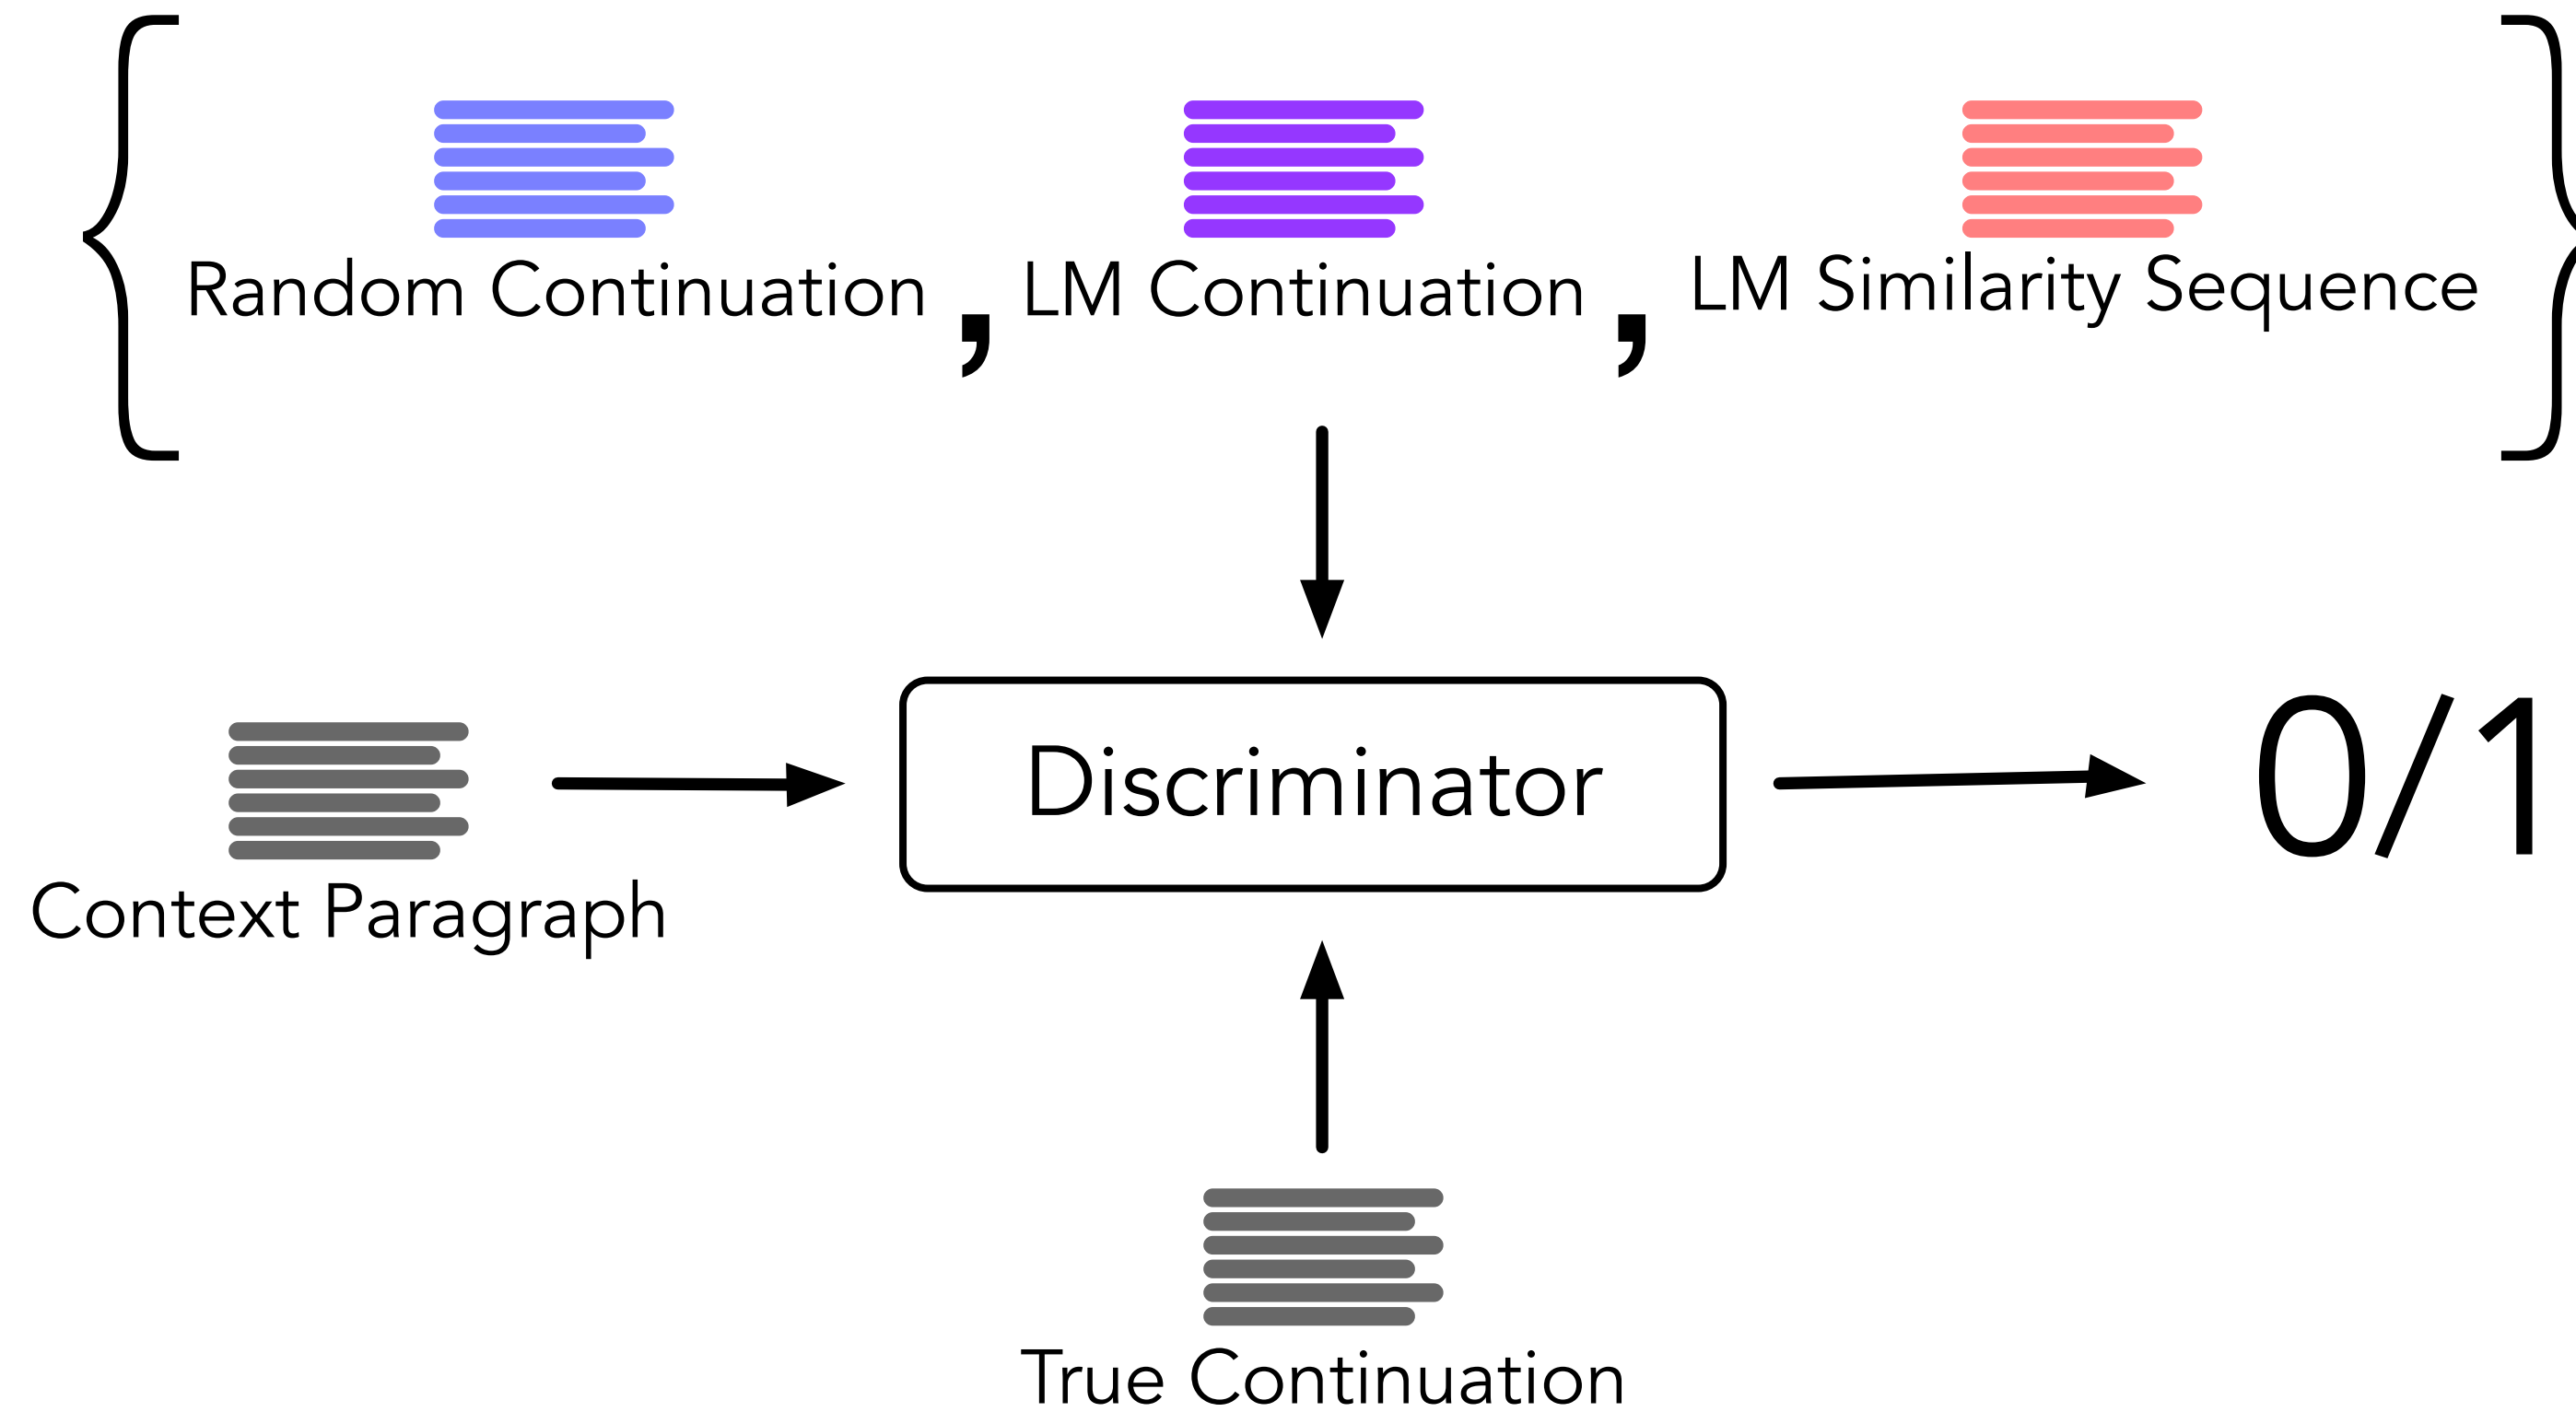


SYNOPSIS

- Generations from RNN-based models tend to be **generic, repetitive, and contradictory**.
- The current loss scheme used for RNNs (cross-entropy & teacher-forcing) is fundamentally **shallow & passive**.
- To capture more abstract **communicative goals** we train discriminators to express elements of **Grice's Maxims** [5]
- These discriminators **learn to cooperate** by **practicing writing** text and then comparing the result to human text

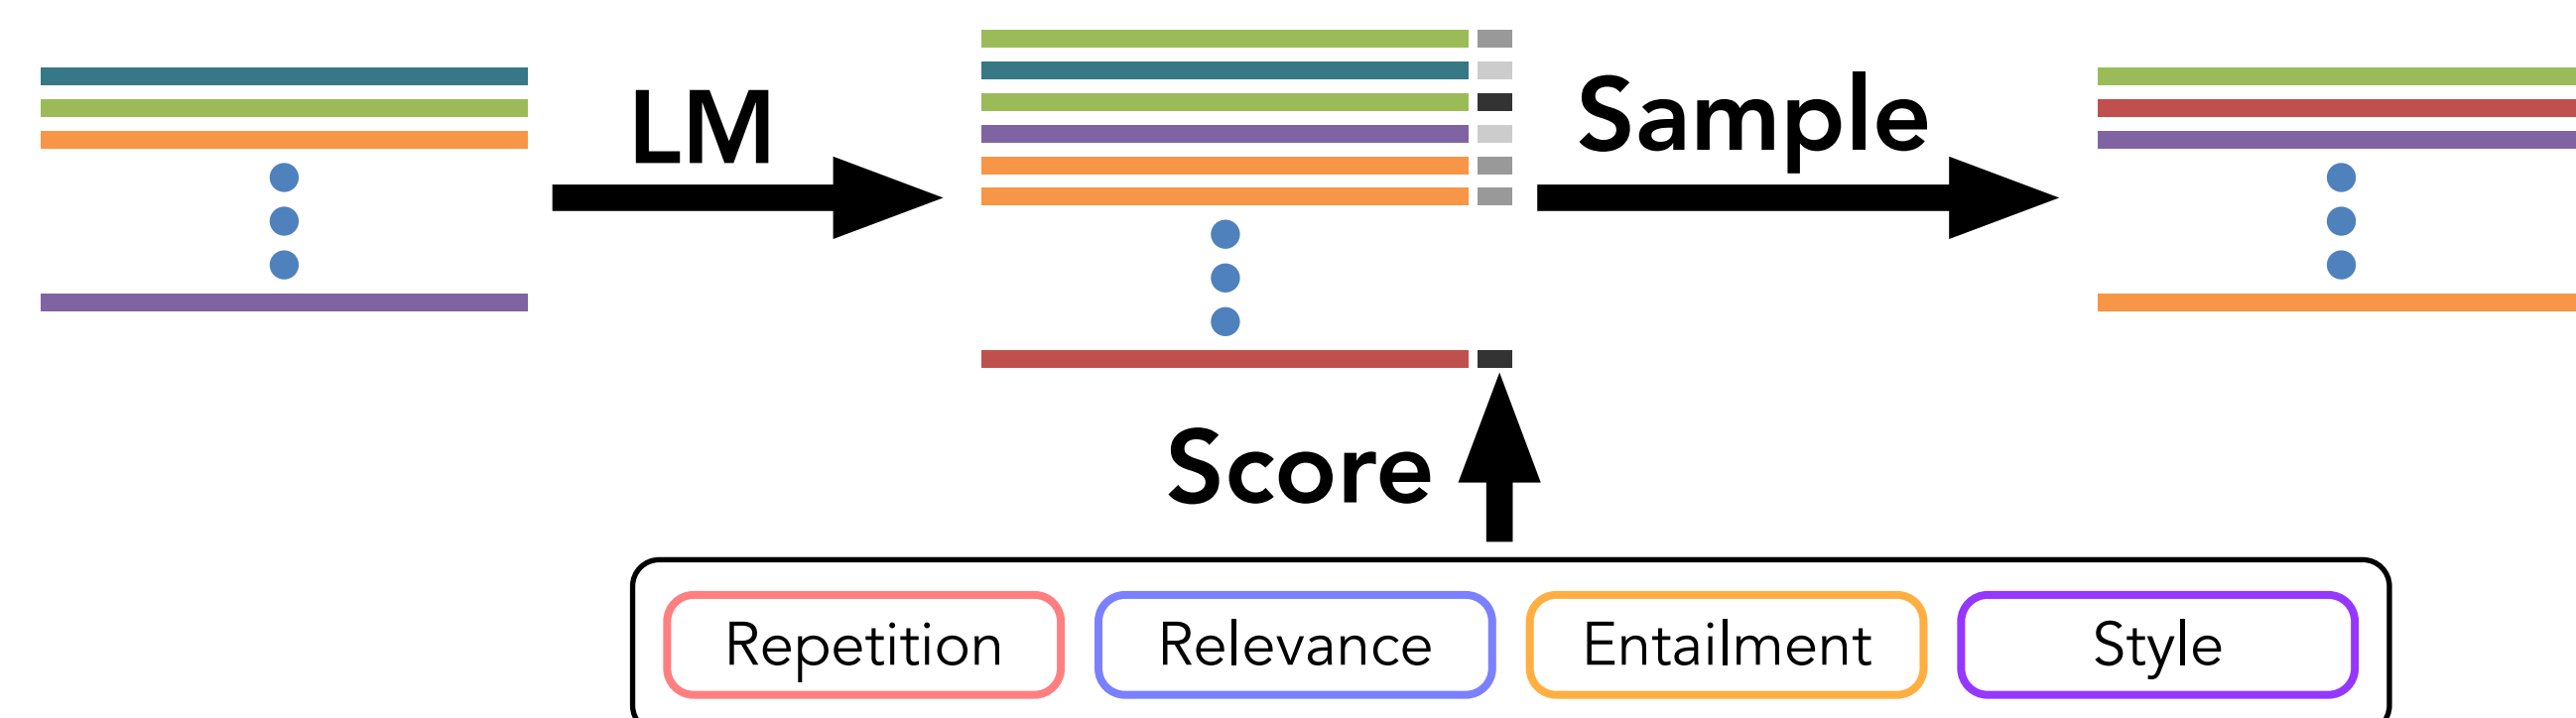
Communicative Goals

Discriminators are trained to identify abstract properties of text from adversarially extracted data:



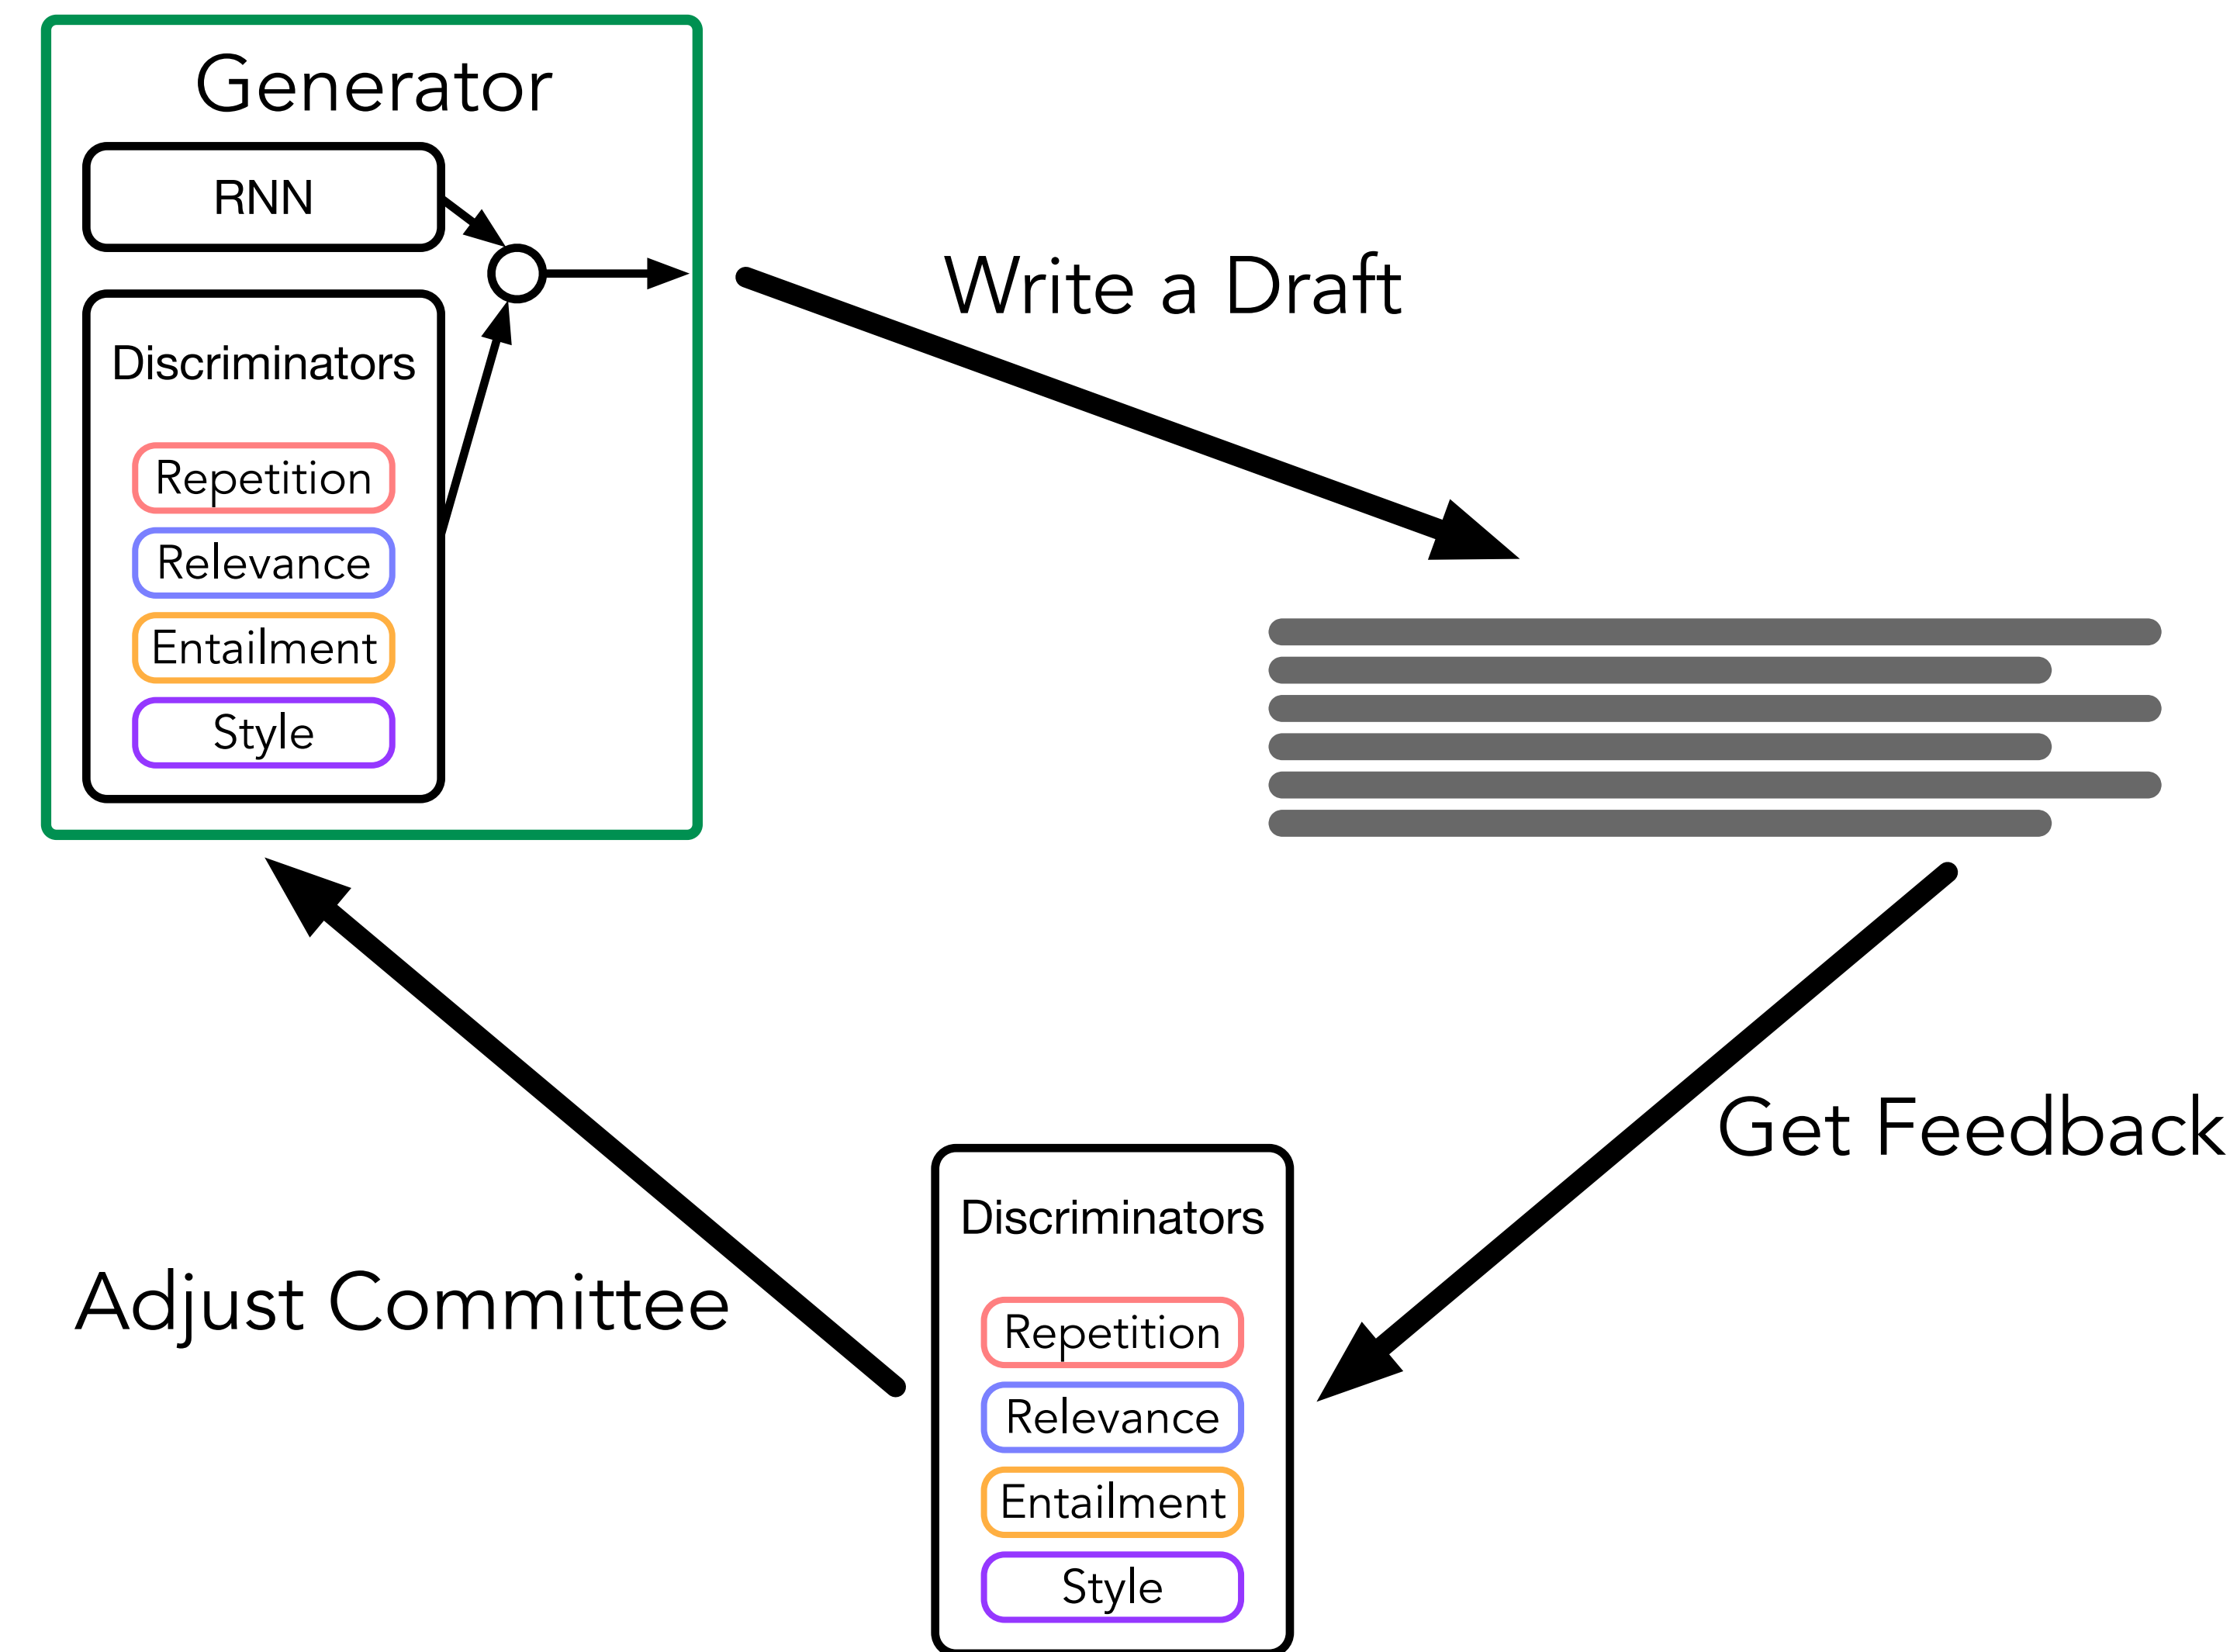
Cooperative Discriminators

A generalization of **Product of Experts** [4] is used to combine signals from different discriminators



$$f_{\lambda}(\mathbf{x}, \mathbf{y}) = \log(P_{LM}(\mathbf{y}|\mathbf{x})) + \sum_k \lambda_k s_k((\mathbf{x}), \mathbf{y})$$

Learning by Practice



Discussion

GAN vs L2W	
Modeling space is constrained by RNN limitations	Modeling space is enriched by cooperative discriminators
Word level only loss	Word & paragraph loss

RNN vs L2W	
Teacher-forcing means RNNs only learn to read passively	Learning by practice means generators learn to write
Cross-entropy only measures word prediction error	L2W loss enhances generator with communicative goals

Works Cited

- [1] Grave, Edouard, et al. "Efficient softmax approximation for GPUs." *arXiv preprint*. 2016.
- [2] Grave, Edouard, Armand Joulin, and Nicolas Usunier. "Improving neural language models with a continuous cache." *ICLR*. 2017.
- [3] Grice, H. Paul. "Logic and conversation." 1975.
- [4] Hinton, G. "Products of Experts." *ICANN*. 1999.
- [5] Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. "Sequence to sequence learning with neural networks." *NIPS*. 2014.
- [6] Yu, Lantao, et al. "SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient." *AAAI*. 2017.
- [7] Yu, Lei, et al. "The neural noisy channel". *ICLR*. 2017.

RNN Deficiencies

All in all, I would highly recommend this hotel to anyone who wants to be **in the heart of the action**, and want to be **in the heart of the action**. If you want to be **in the heart of the action**, this is not the place for you. However, If you want to be **in the middle of the action**, this is the place to be.

TripAdvisor L2W Example

The bed was super comfortable! I can't think of a single thing that wasn't perfect. We didn't have breakfast at the hotel, but there are plenty of cafes and patisseries nearby. It's a bit of a walk to the nearest metro, but I'm sure that's the best way to see Paris!!

BookCorpus L2W Example

It was the voice of a man who had never heard him speak before. He didn't even know what to say, so he just nodded and turned away. The rest of the congregation gathered around him. There were so many people in their midst now that they were all looking at him. He felt a little sick.

Results

L2W is preferred over competitive baselines by **human evaluators**, on specific criteria (left) and overall quality (right)

