

Discourse-Aware Neural Rewards for Coherent Text Generation

Antoine Bosselut^{1,*}, Asli Celikyilmaz², Xiaodong He³,
Jianfeng Gao², Po-Sen Huang² and Yejin Choi^{1,4}

¹Paul G. Allen School of Computer Science & Engineering, University of Washington

²Microsoft Research

³JD AI Research

⁴Allen Institute for Artificial Intelligence

{antoineb, yejin}@cs.washington.edu {xiaodong.he}@jd.com

{aslicel, jfgao, pshuang}@microsoft.com

Abstract

In this paper, we investigate the use of discourse-aware rewards with reinforcement learning to guide a model to generate long, coherent text. In particular, we propose to learn neural rewards to model cross-sentence ordering as a means to approximate desired discourse structure. Empirical results demonstrate that a generator trained with the learned reward produces more coherent and less repetitive text than models trained with cross-entropy or with reinforcement learning with commonly used scores as rewards.

1 Introduction

Defining an ideal loss for training text generation models remains an open research question. Many existing approaches based on variants of recurrent neural networks (Hochreiter and Schmidhuber, 1997; Cho et al., 2014) are trained using cross-entropy loss (Bahdanau et al., 2015; Vinyals et al., 2015; Xu et al., 2015; Rush et al., 2015), often augmented with additional terms for topic coverage or task-specific supervision (Kiddon et al., 2016; Yang et al., 2017).

Training with cross-entropy, however, does not always correlate well with achieving high scores on commonly used evaluation measures such as ROUGE (Lin, 2004), BLEU (Papineni et al., 2002), or CIDEr (Vedantam et al., 2015). Another current line of research therefore explores training generation models that directly optimize the target evaluation measure (Wu et al., 2016; Ranzato et al., 2015; Paulus et al., 2018; Rennie et al., 2017) using reinforcement learning methods such as the REINFORCE algorithm (Williams, 1992).

* Work done while author was at Microsoft Research

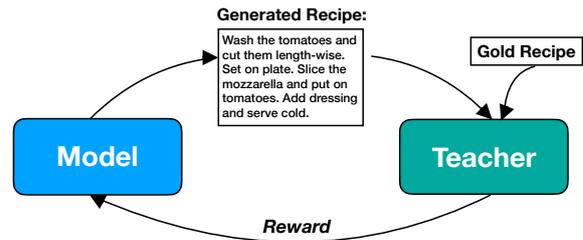


Figure 1: The generator is rewarded for imitating the discourse structure of the gold sequence.

Importantly, most automatic measures are based on local n -gram patterns, providing only a limited and myopic perspective of overall text quality. As a result, while models trained to directly optimize these measures can yield improvements on the same measures, they may not lead to better quality in terms of overall coherence or discourse structure. Indeed, recent studies have reported cases where commonly used measures do not align well with desired aspects of generation quality (Rennie et al., 2017; Li et al., 2016).

The challenge, however, is to define a global score that can measure the complex aspects of text quality beyond local n -gram patterns. In this paper, we investigate *learning* neural rewards and their use in a reinforcement learning regime with a specific focus on learning more discourse-aware and coherent text generation. Our approach shares the spirit of the work of Lowe et al. (2017), where neural scores were learned to approximate human judgments of dialogue quality. The key difference is that our rewards can be fully automatically constructed without requiring human judgments and can be trained in an unsupervised manner.

More specifically, we propose a neural reward learning scheme that is trained to capture cross-sentence ordering structure as a means to approximate the desired discourse structure in documents. The learned *teacher* computes rewards for the

underlying text generator (see Figure 1), which is trained using self-critical reinforcement learning (Rennie et al., 2017). We also present a new method for distributing sentence-level rewards for more accurate credit assignment.

We test our approach on the task of generating cooking recipes, and evaluate using automatic overlap metrics that measure discourse structure. We also provide human judgments that yield comprehensive insights into the model behavior induced by the learned neural rewards. Empirical results demonstrate that a generator trained with the discourse-aware rewards produces text that is more coherent and less repetitive than models trained with cross-entropy or reinforcement learning with other commonly used scores.

2 Neural Teachers

Recent work in image captioning (Rennie et al., 2017), machine translation (Wu et al., 2016), and summarization (Paulus et al., 2018) has investigated using policy gradient methods to fine-tune neural generation models using automatic measures such as CIDEr as the reward. However, because most existing automatic measures focus on local n -gram patterns, fine-tuning on those measures may yield deteriorated text despite increased automatic scores, especially for tasks that require long coherent generation (§6.1).

Since writing out a scoring term that quantifies the quality of discourse coherence is an open research question, we take inspiration from previous research that *learns* the overall ordering structure of a document as an approximation of the discourse structure (Barzilay and Lapata, 2005, 2008; Barzilay and Lee, 2004; Li and Hovy, 2014), and propose two neural teachers that can learn to score an ordered sequence of sentences. The scores from these neural teachers are then used to formulate rewards (§4.2) that guide coherent long text generation systems in a policy gradient reinforcement learning setup. Notably, the neural teachers are trained offline on gold sequences in an unsupervised manner prior to training the generator. They are not trained jointly with the generator and their parameters are fixed during policy learning.

2.1 Notation

We define a document of n sentences as $S = \{s_0, \dots, s_n\}$ where each sentence s_j has L_j words.

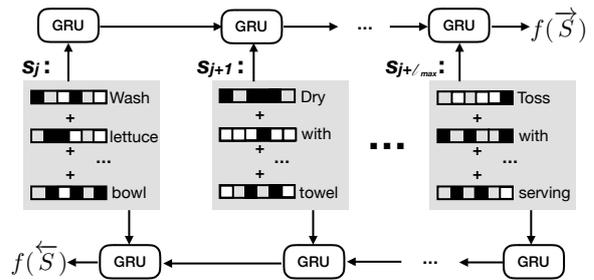


Figure 2: The teacher encodes the sentences of the document in the forward and reverse order.

2.2 Absolute Order Teacher

The first teacher explored is motivated by work on deep semantic similarity models (Huang et al., 2013), which approximated the similarity between queries and documents in information retrieval tasks. We extend this approach to modeling temporal patterns by training a sentence encoder to minimize the similarity between a sequence encoded in its forward order, and the same sequence encoded in the reverse order (see Figure 2).

To focus the teacher on discourse structure, we design the encoder to capture *sentence order*, instead of *word order*. Words in each sentence s_j are encoded using a bag of words:

$$\mathbf{s}_j = \sum_{i=1}^{L_j} x_{ij} \quad (1)$$

where x_{ij} is a word embedding and \mathbf{s}_j is a sentence embedding. Each \mathbf{s}_j is passed to a gated recurrent unit (GRU) and the final output of the hidden unit is used as the representation for the full document:

$$h_j = \text{GRU}(\mathbf{s}_j, h_{j-1}) \quad (2)$$

$$f(S) = h_n \quad (3)$$

where $f(S)$ is the representation of the sentences of the document and h_n is the final output vector of the GRU. To capture properties of temporal coherence among document sentences, the teacher is trained to minimize \mathcal{L}_{abs} , the cosine similarity between the sentence embedding from reading the sentences in the forward order, \vec{S} and from reading the sentences in the reverse order, \overleftarrow{S} :

$$\mathcal{L}_{abs} = \frac{\langle f(\vec{S}), f(\overleftarrow{S}) \rangle}{\|f(\vec{S})\| \|f(\overleftarrow{S})\|} \quad (4)$$

Intuitively, by parametrizing only relations between *sentences* (with the GRU layer) and not those between *words*, the teacher only captures sentence ordering properties. When training the neural generator (§4), we use this learned teacher to generate a reward that judges the generated sequence’s ordering similarity to the gold sequence.

2.3 Relative Order Teacher

While the absolute ordering teacher evaluates the temporal coherence of the entire generation, we may want our teacher to be able to judge finer-grained patterns between sentences. In recipes, for example, where sentences correspond to process steps, the teacher should capture implicit script knowledge (Schank and Abelson, 1975) among groups of sentences. Consequently, the teacher should reward sentences individually for how they fit with surrounding sentences.

In many current approaches for using policy gradient methods to optimize a model with respect to a global score, each sentence receives the same reward. This framework assumes each sentence is equally responsible for the reward gathered by the full sequence, allowing potentially appropriate subsequences to be incorrectly penalized. We design the relative order teacher to address this issue.

The relative order teacher is trained in the same way as the absolute order model. A bag of words embedding is computed for each sentence in the gold sequence. Subsequences of the gold document that have ℓ sentences are selected where $\ell \in (\ell_{min}, \ell_{max})$. For a subsequence beginning at sentence j , the model computes:

$$f(S_{j:j+\ell}) = \text{GRU}(s_{j+\ell}, h_{j+\ell-1}) \quad (5)$$

where $f(S_{j:j+\ell})$ is the encoded representation of sentences $\{s_j, \dots, s_{j+\ell}\}$ and h_{j-1} would be initialized as a vector of zeros. The relative ordering teacher is trained to minimize \mathcal{L}_{rel} , the cosine similarity between gold orders of subsequences:

$$\mathcal{L}_{rel} = \frac{\langle f(\vec{S}_{j:j+\ell}), f(\overleftarrow{S}_{j:j+\ell}) \rangle}{\|f(\vec{S}_{j:j+\ell})\| \|f(\overleftarrow{S}_{j:j+\ell})\|} \quad (6)$$

where the arrow above S signifies the order in which the sentences are processed. The relative ordering teacher learns to identify local sentence patterns among ordered sentences, thereby learning how to reward sequences that are temporally coherent.

3 Generator Architecture

In the task of recipe generation, the model is given a title of a recipe such as “*Cheese Sandwich*” and a list of ingredients (e.g., cheese, bread, etc.) and must generate the full multi-sentence recipe text. Similar to data to document generation tasks, the model must generate a full long-form text from sparse input signal, filling in missing information on its own (Wiseman et al., 2017).

3.1 Notation

Using the same notation as Kiddon et al. (2016), we are given a set of recipe title words $\{g_1, \dots, g_n\}$ (e.g., {“cheese”, “sandwich”}) and a list of ingredients $E = \{\mathbf{i}_1, \dots, \mathbf{i}_{|E|}\}$ where each \mathbf{i} can be a single- or multi-word ingredient phrase (e.g., “onions” or “onions, chopped”). In the following paragraphs, all W variables are projection matrices and all b variables are bias vectors.

3.2 Encoder

We use a modification of the baseline encoder of Kiddon et al. (2016). First, the title words are encoded as a bag of embeddings, \mathbf{g} . Second, each ingredient phrase \mathbf{i} is encoded as a bag of embeddings vector, \mathbf{e}_i . The ingredient embeddings are inputs to a bidirectional gated recurrent unit, which yields an output vector \mathbf{e} . The final encoder output is the concatenation of these two representations, $\mathbf{h}^e = [\mathbf{g}, \mathbf{e}]$.

3.3 Decoder

The decoder is a separate gated recurrent unit that receives \mathbf{h}^e from the encoder to initialize its hidden state \mathbf{h}_0^d and must generate a full recipe word by word. At each time step, the model receives an input token embedding, x_t , as well as the output from the encoder \mathbf{h}^e :

$$a_t = \sigma(W_1 h_{t-1}^d + W_2 x_t + b_1) \quad (7)$$

$$\mathbf{z}_t = a_t \mathbf{h}^e \quad (8)$$

$$\tilde{x}_t = [x_t, \mathbf{z}_t] \quad (9)$$

where \tilde{x}_t is the input to the recurrent unit at every time step. The recipe generator is pretrained to minimize the negative loglikelihood of predicting the next token in the recipe:

$$L_{mle} = - \sum_{t=1}^T \log P(x_t | x_0, \dots, x_{t-1}, \mathbf{h}^e) \quad (10)$$

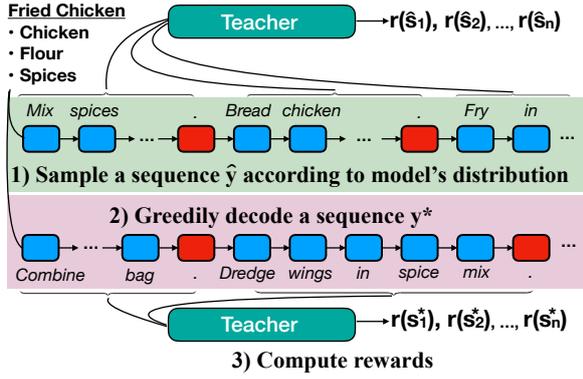


Figure 3: The model generates a recipe by sampling from its output vocabulary distribution and greedily decodes a baseline recipe. The generated sentences are passed to the teacher, which yields a reward for each sentence in each recipe.

where \mathbf{h}^e is the encoded representation of the title and ingredients from Section 3.2 and T is the number of words in the gold recipe.

4 Policy Learning

Training a recipe generation model using maximum likelihood estimation produces generations that are locally coherent, but lack understanding of domain knowledge. By using a teacher that rewards the model for capturing cooking recipe discourse semantics, the model learns a policy that produces generations that better model the underlying recipe process. We learn a policy using the self-critical approach of Rennie et al. (2017).

4.1 Self-critical sequence training

In self-critical sequence training, outlined in Figure 3, the model learns by being rewarded for sampling sequences that receive more reward than a greedily decoded sequence. For each training example, a sequence \hat{y} is generated by sampling from the model’s distribution $P(\hat{y}_t | \hat{y}_0, \dots, \hat{y}_{t-1}, \mathbf{h}^e)$ at each time step t . Once the sequence is generated, the teacher produces a reward $r(\hat{y}_t)$ for each token in the sequence. A second sequence y^* is generated by argmax decoding from $P(y_t^* | y_0^*, \dots, y_{t-1}^*, \mathbf{h}^e)$ at each time step t . The model is trained to minimize:

$$\mathcal{L}_{rl} = - \sum_{t=1}^T (r(\hat{y}_t) - r(y_t^*)) \log P(\hat{y}_t | \hat{y}_0, \dots, \hat{y}_{t-1}, \mathbf{h}^e) \quad (11)$$

where $r(y_t^*)$ is the reward produced by the teacher for tokens of the greedily decoded sequence. Be-

cause $r(y^*)$ can be viewed as a *baseline* reward that sampled sequences should receive more than, the model learns to generate sequences that receive more reward from the teacher than the best sequence that can be greedily decoded from the current policy. This approach allows the model to *explore* sequences that yield higher reward than the current best policy.

4.2 Rewards

As we decode a sequence $y = \{y_0, \dots, y_t\}$, we track a sentence index that is the number of sentence delimiter tokens (e.g., “.”) generated by the model. The model then implicitly decodes a set of generated sentences, $S' = \{s_0, \dots, s_n\}$. These sentences are provided to the teachers defined in Section 2, which compute a score for the generated sequence. We explain the procedure for producing a token reward $r(y_t)$ from these scores below.

Absolute Order Once a sequence has been generated, the absolute order teacher computes a reward for y in the following way:

$$r_{abs}(y) = \frac{\langle f(S'), f(\vec{S}) \rangle}{\|f(S')\| \|f(\vec{S})\|} - \frac{\langle f(S'), f(\overleftarrow{S}) \rangle}{\|f(S')\| \|f(\overleftarrow{S})\|} \quad (12)$$

where \vec{S} is the forward-ordered corresponding gold sequence and \overleftarrow{S} is the reverse-ordered gold sequence. Both terms in the reward computation are variations of the loss function on which the absolute order teacher was trained (Equation (4)). This reward compares the generated sequence to both sentence orders of the gold sequence, and rewards generations that are more similar to the forward order of the gold sequence. Because the cosine similarity terms in Equation (12) are bounded in $[-1, 1]$, the model receives additional reward for generating sequences that are different from the reverse-ordered gold sequence.

Relative Order Similarly, the relative order reward is generated by the relative order teacher (§2.3), which evaluates subsequences of sentences, rather than the whole sequence. For a sentence s_j , the reward is computed as:

$$r_{rel}(s_j) = \frac{1}{L} \sum_{\ell=\ell_{min}}^{\ell_{max}} \left(\frac{\langle f(S'_{j-\ell:j}), f(\vec{S}_{j-\ell:j}) \rangle}{\|f(S'_{j-\ell:j})\| \|f(\vec{S}_{j-\ell:j})\|} - \frac{\langle f(S'_{j-\ell:j}), f(\overleftarrow{S}_{j-\ell:j}) \rangle}{\|f(S'_{j-\ell:j})\| \|f(\overleftarrow{S}_{j-\ell:j})\|} \right) \quad (13)$$

where ℓ_{min} and ℓ_{max} define the window of sentences to include in the computation of the reward. Similar to the absolute order teacher, the relative order teacher produces scores bounded in $[-1, 1]$, giving the model additional reward for generating sequences that are different from the reverse-ordered gold subsequences.

Credit Assignment When rewarding tokens with the absolute ordering teacher, each generated token receives the same sequence-level reward from the absolute order teacher:

$$r(y_t) = r_{abs}(y) \quad (14)$$

The relative order teacher, meanwhile, computes rewards for sentences based on their imitation of nearby sentences in the gold recipe. Rather than combining all rewards from the teacher to compute a full sequence reward, sentences should only be rewarded for their own quality. Each token in a sentence corresponds to a position in the full sequence. When relative order rewards are computed by the teacher, the correct sentence reward is indexed for each token. Consequently, when training with a relative order teacher, words only receive rewards for the sentences they belong to:

$$r(y_t) = \sum_{j=1}^{|S|} \mathbb{1}(y_t \in \hat{s}_j) r_{rel}(\hat{s}_j) \quad (15)$$

where $|S|$ is the number of sentences in the generated recipe, and $\mathbb{1}$ is an indicator variable identifying word y_t belonging to sentence s_j .

4.3 Mixed Training

As the model learns parameters to optimize the amount of reward it receives from the teacher, it is not explicitly encouraged to produce fluent generations. The model quickly learns to generate simple sequences that *exploit* the teacher for high rewards despite being incoherent recipes (e.g., Figure 4). Consequently, it is possible that generated sequences are no longer readable (Pasunuru and Bansal, 2017; Paulus et al., 2018).

Title: Chili Grits

Ingredients: boiling water, butter, shredded cheddar cheese, jalapenos, eggs, chicken cream of soup, salt

Generated Recipe: Here .

Figure 4: Recipe generated from a self-critical model with no mixed training

To remedy this effect, the model optimizes a mixed objective that balances learning the discourse-focused policy while maintaining the generator’s language model:

$$\mathcal{L}_{mix} = \gamma \mathcal{L}_{rl} + (1 - \gamma) \mathcal{L}_{mle} \quad (16)$$

where L_{mle} is the objective from Equation (10), L_{rl} is the objective from either Equation (11), and γ is a hyperparameter in $[0, 1]$.

5 Experimental Setup

5.1 Datasets

We use the Now You’re Cooking dataset with the same training/test/development splits from Kiddon et al. (2016). For training, we use 109567 recipes with 1000 recipes set aside for both development and test.

5.2 Training

Teacher Models The teachers are trained before the recipe generator and their parameters are fixed during generation. We tune hyperparameters on the development set. To train the relative order teacher, we sample 20 subsequences from each recipe of $\ell_{min} = 3$ to $\ell_{max} = 6$ sentences. Additional details are provided in Appendix A.2.

Recipe Generator We pretrain a recipe generator using a variant of the encoder-decoder baseline from Kiddon et al. (2016). Comprehensive hyperparameter details can be found in Appendix A.3.

Policy Learning We train a different model for three different teacher-provided rewards: absolute ordering (AO), relative ordering (RO) and a joint reward of relative ordering and BLEU-4 (RO + B4), where the full-sequence BLEU-4 reward and the sentence-level relative ordering reward are summed at each time step. The best model for the absolute and relative ordering rewards are the ones that receive the highest average reward on the development set. The best model for the mixed reward was chosen as the one that achieved the highest average geometric mean of BLEU-4 reward and average relative ordering reward for each generated sequence y in the development set:

$$\bar{r} = \frac{r_{b4}(y)}{T} \sum_{t=1}^T r_{RO}(y_t) \quad (17)$$

where r_{b4} is the BLEU-4 score of the whole generated sequence, and r_{RO} is computed using Equa-

Model	BLEU-1	BLEU-4	R-L	AB1	AB4	AR-L	SCB1	SCB4	SCR-L
Cross-entropy (MLE)	26.86	4.74	28.86	31.23	4.83	28.51	51.92	26.35	50.21
BLEU-4 (Rennie et al., 2017)	7.75	1.38	13.93	5.69	0.84	10.37	10.76	5.05	20.87
CIDEr (Rennie et al., 2017)	12.67	1.90	21.20	14.61	1.79	21.70	26.07	12.30	41.65
ROUGE-L (Paulus et al., 2018)	29.00	4.86	29.10	33.49	4.73	28.11	56.86	27.83	51.26
BLEU-1 ($\gamma = 0.97$)	31.16	5.60	29.53	32.28	5.09	29.34	52.63	25.43	51.58
BLEU-4 ($\gamma = 0.99$)	30.56	5.42	29.16	32.53	4.99	28.99	53.48	26.35	51.02
CIDEr ($\gamma = 0.97$)	29.60	5.10	28.79	33.93	4.81	28.41	57.00	27.55	50.57
ROUGE-L ($\gamma = 0.97$)	26.88	4.66	29.49	31.85	5.01	29.25	53.84	26.77	51.88
Absolute Ordering (AO)	23.70	4.25	28.43	28.22	4.44	27.88	47.93	24.47	50.15
Relative Ordering (RO)	27.75	4.88	29.60	34.37	5.60	29.36	58.31	29.14	53.08
Relative Ordering + BLEU-4	29.58	5.26	29.78	35.13	5.55	29.33	59.13	29.19	52.46

Table 1: Evaluation results for generated sequences by models and baselines. We **bold** the top performing result. The second to fourth columns list word-level scores. Columns AB1, AB4, and AR-L list action-level scores (§6.1). Columns SCB1, SCB4, and SCR-L list state change level scores (§6.1).

tion (15). Our best models use $\gamma = 0.97$ when training with the mixed objective from Equation (16).

5.3 Baselines

As baselines, we report results for a model trained only with cross-entropy loss (MLE) and for re-implemented versions of models from Rennie et al. (2017) and Paulus et al. (2018). These baselines achieved state of the art results in image captioning and document summarization tasks. We found, however, that their high γ (1 and 0.9984, respectively) led to low fluency, resulting in reduced performance on word-level scores. To control for this effect, we trained additional versions of each baseline with different values for γ and report the best performing configurations (see Table 1).

6 Results

6.1 Overlap Metrics

Scores We compute the example-level BLEU-1, BLEU-4, and ROUGE-L (R-L) scores for all recipes in the test set. A generated recipe, however, must be coherent at both the *word-level*, linking words and phrases sensibly, and the *world-level*, describing events that are grounded in real-world actions. Because n -gram scores do not evaluate if a generated recipe models this latent process, we also report these scores on the *action* and *state change* sequence described in the recipe. These words depict a *simulated* world where actions are taken and state changes are induced. A generated recipe should follow the sequence of actions taken in the gold recipe, and induce the same state changes as those in the gold recipe.

We use the state change lexicon from Bosselut et al. (2018) to map recipe words to ordered sequences of actions and state changes. Each entry in the lexicon contains an action in the cooking domain as well as the state changes that result from that action in the set of {LOCATION, COMPOSITION, COOKEDNESS, TEMPERATURE, SHAPE, CLEANLINESS}.

Action sequences are formed by mapping lemmas of words in generated sequences to entries in the lexicon. We compare these event sequences to the gold event sequences using the same scores as for words – BLEU-1, BLEU-4, and ROUGE-L. Intuitively, these scores can be seen as evaluating the following: whether the generated recipe depicts the same actions (AB1), subsequences of consecutive actions (AB4), and full action sequence (AR-L) as the gold recipe.

State change sequences are more coarse-grained than action sequences, and are formed by mapping actions to their state changes in the lexicon from Bosselut et al. (2018). These scores evaluate whether the generated recipe implies the same induced state changes (SCB1), subsequences of consecutive state changes (SCB4), and global state change order (SCR-L) as the gold recipe.

Results Our results in Table 1 show that models optimized on word overlap metrics achieve the greatest improvements for those scores. Optimizing scores such as BLEU-1 encourages the model to output words and phrases that overlap often with reference sequences, but that may not describe main events in the recipe process.

When examining models trained using a neural teacher, we see that the model optimized with

	MLE	RO + B4	Tie
Fluency	0.330	0.447	0.223
Ingredient Use	0.350	0.440	0.210
Title Completion	0.347	0.430	0.223
Action Order	0.377	0.453	0.170
	BLEU-1	RO + B4	Tie
Fluency	0.387	0.373	0.240
Ingredient Use	0.327	0.363	0.310
Title Completion	0.353	0.377	0.270
Action Order	0.410	0.403	0.187

Table 2: Human evaluation measuring proportion of winners. Upper table compares MLE baseline with RO + B4 model. Lower table compares BLEU-1 baseline with RO + B4 model.

the absolute ordering reward performs worse than most baselines for every word-level score. The relative ordering model, however, raises every word-level score above the cross-entropy baseline, indicating the importance of fine-grained credit assignment at the sentence-level. The model trained with mixed rewards from the teacher and BLEU-4 achieves even higher scores, showing the benefits of training with diverse rewards.

When evaluating these metrics for the action and state change sequence, the models trained with feedback from the relative ordering teacher show large improvement over the baselines, indicating that the models exhibit more understanding of the latent process underlying the task. While optimizing word-level scores teaches the generator to output common sequences of words, the relative ordering reward teaches the model to focus on learning co-occurrences between recipe events.

6.2 Human Evaluation

We perform a human evaluation on 100 recipes sampled from the test set to evaluate our model on four aspects of recipe quality: fluency, ingredient use, title completion, and action ordering. For each example, three judges from Amazon Mechanical Turk are shown a pair of recipes, each generated by a different model and asked to select the recipe that is better according to the criteria above. For ingredient use, judges select the recipe that uses more of the ingredients correctly. For title completion, we ask judges to select the recipe that best completes the dish described in the recipe title. Finally, for action ordering, judges choose the recipe that better links subtasks in the recipes.

	MLE	RO + B4	Tie
Fluency	0.317	0.425	0.258
Ingredient Use	0.342	0.458	0.200
Title Completion	0.358	0.450	0.192
Action Order	0.367	0.483	0.150
	BLEU-1	RO + B4	Tie
Fluency	0.391	0.383	0.225
Ingredient Use	0.267	0.392	0.342
Title Completion	0.325	0.418	0.258
Action Order	0.433	0.442	0.125

Table 3: Proportion of winners for **long** generated recipes. Upper table compares MLE baseline with RO + B4 model. Lower table compares BLEU-1 baseline with mixed RO + B4 model.

Models We use the Relative Ordering + BLEU-4 model (RO + B4) and compared to two baselines, the cross-entropy model (MLE), and the BLEU-1 model, which achieved the best scores on several word-level metrics (§6.1).

Results We report results in Table 2. Our model outperforms the cross-entropy baseline, consistently being preferred on aggregate for every question. Workers preferred the BLEU-1 baseline for the fluency and action order questions, while preferring recipes generated by the teacher-trained model for the ingredient use and title ordering questions. Upon further analysis, we see that the strength of the BLEU-1 model depends on the length of the original reference sequence. In Table 3, we show evaluation scores for recipes where the gold recipe was longer than 100 words. Our model’s performance rises compared to the BLEU-1 model for every question, showing that modeling discourse structure as learned reward improves global coherence in long text.

6.3 Insights

Qualitative Analysis In Table 4, we see the effect that the neural teacher has on the recipe generator. The teacher rewards behavior that more closely imitates the actions in the gold recipe. In the first example, the generator learns to complete the actions of placing the mixture into the a *greased casserole* and then *baking* it, which the MLE model misses. The teacher also discourages repetitive phrases, as they provide no increase in reward during training. One weakness of our teacher models, however, is that they encourage common temporal patterns, such as in the third

Title	Royale Casserole
Ings	Broccoli, chopped; butter; flour; salt; black pepper; nutmeg; milk; cheese shredded swiss; lemon juice; cooked cubed turkey; mushrooms; grated parmesan cheese; biscuits, refrigerated
MLE	Cook broccoli according to package directions. Drain. Melt butter in skillet. Add flour, stirring constantly. Cook, stirring constantly, until thickened. Add milk, salt, pepper, and nutmeg. Cook over low heat, stirring constantly, until thickened. Add cheese, salt, pepper and pimento. Cook and stir till thickened. Add cheese, stirring until melted. Add broccoli and mushrooms . Cook over low heat for 5 minutes. Add cheese. Cook and stir until cheese melts. Add cheese. Cook and stir until cheese melts. Add cheese. Cook and stir until cheese melts. Add cheese and stir until melted. Serve over hot noodles.
RO	Melt butter in saucepan. Add flour and cook, stirring constantly, until thickened. Add milk, salt, pepper, and nutmeg. Cook, stirring constantly, until thickened. Add cheese, salt and pepper. Cook over low heat, stirring constantly, until mixture is thickened and smooth. Remove from heat. Stir in cheese. Spoon into greased casserole . Top with cheese. Bake in 350 f oven for 30 minutes. Serves 6.
Gold	Preheat oven to 375. Melt butter in saucepan. Blend in flour, salt, pepper, and nutmeg; cook 1-2 minutes. Gradually add milk; cook, stirring, until slightly thickened. Stir in frozen vegetables. Remove from heat; stir in cheese until melted. Add lemon juice, turkey, mushrooms and broccoli. Pour mixture into a lightly greased baking dish; sprinkle with parmesan and top with biscuits. Bake 20 minutes, or until biscuits are golden brown.
Title	Wine Sauerkraut
Ings	butter; dry white wine; 16 oz drained sauerkraut
MLE	In a saucepan, combine soup , wine, and butter. Heat, uncovered, in microwave oven 2 minutes or until mixture is heated through.
RO	Melt butter in skillet. Add sauerkraut and wine; heat to boiling. Reduce heat; cover and simmer 15 minutes. Add wine and heat to boiling; reduce heat. Cover and cook 10 minutes or until mixture is tender. Serve on rice.
Gold	Melt butter in 10-inch skillet over low heat; add sauerkraut. Cover and cook, stirring occasionally, 30 minutes; add wine. Heat to boiling; reduce heat. Cover and simmer until liquid is absorbed, about 45 minutes. Follow directions except simmer until liquid is absorbed, about 1 hour.
Title	Strawberry Glazed Pie
Ings	fresh strawberries; granulated sugar; salt; cornstarch; lemon juice; baked, cooled pie shell
MLE	Combine sugar, cornstarch and salt in top of a double boiler. Cook and stir over medium heat until mixture thickens and boils. Boil and stir 1 minute. Remove from heat and stir in lemon juice. Cool. Add strawberries and lemon juice . Pour into baked pie shell. Top with whipped cream. Serve immediately.
RO	Combine sugar, cornstarch, and salt in a saucepan; cook over medium heat, stirring constantly, until mixture boils and thickens. Remove from heat; stir in lemon juice. Cool. Add strawberries; mix well. Pour into baked pie shell. Bake in preheated 325-degree oven for 10 minutes . Cool. Garnish with whipped cream.
Gold	Wash, drain thoroughly, and hull strawberries. Arrange about 3 cups of whole berries over bottom of baked pastry shell. Crush remaining berries in a saucepan. In a bowl, mix sugar, salt and cornstarch; stir into crushed berries. Heat slowly, stirring constantly, until mixture comes to a boil and thickens. Remove from heat and stir in lemon juice. Cool, then spoon over berries in pie shell chill until glaze is set. Garnish with whipped cream.

Table 4: Example recipe generations from our model and comparative baselines. **Boxed** spans indicate recipe events missed by another model’s generation. **Red** spans indicate superfluous events. The **Ings** row lists the ingredients (separated by semicolons) provided to make the dish in the title.

example in Table 4, where the generator mentions *baking the pie*. The model recognizes pies are generally supposed to be baked, even if it is not appropriate for that particular recipe.

Teacher Feedback Frequency We design the reward functions in Eq. 12 and Eq. 13 to require two passes through the teacher, one comparing the generated sequence to the forward gold sequence, and one comparing it to the reverse gold sequence. With no teacher comparison to the reverse-ordered sequence, the generator learns to *exploit* the teacher for reward with very simple sequences such as “Serve.” and “Here’s direction.” When comparing with both orders, however, this effect is dampened, hinting at the importance of

ensembling feedback from multiple sources for robust reward production. Another solution to this effect was mixing policy learning and maximum likelihood learning (Eq. 16) as the underlying language model of the generator did not deteriorate.

Impact of ℓ_{max} and γ Two hyperparameters to tune when training with teacher models are the mixed loss coefficient γ , which balances MLE learning with policy learning, and $[\ell_{min}, \ell_{max}]$, the number of sentences to consider when computing the relative order reward. We fix $\ell_{min} = 3$, and vary $\ell_{max} \in [3, 6]$ and $\gamma \in \{0.95, 0.97, 0.98\}$. Figure 5 shows the importance of tuning γ . A low γ will not allow the teacher to guide the model’s learning, while a high γ causes the lan-

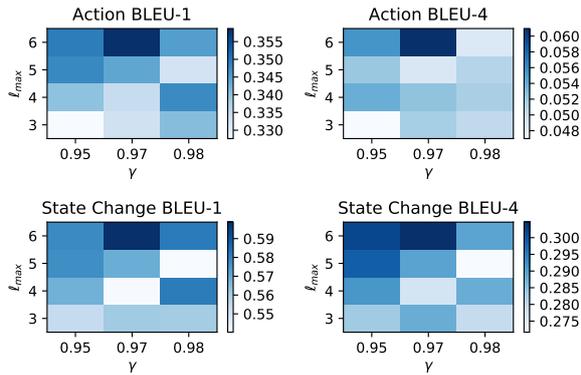


Figure 5: Action and State Change BLEU Metrics for different initializations of ℓ_{max} and γ

guage model to deteriorate. Interestingly, a higher ℓ_{max} leads to better performance on global coherence scores, implying that relative order rewards conditioned on more sentences allow the model to learn longer-range context co-occurrences.

7 Related Work

The field of neural text generation has received considerable attention in tasks such as image captioning (Vinyals et al., 2015; Xu et al., 2015), summarization (Rush et al., 2015; See et al., 2017), machine translation (Bahdanau et al., 2015), and recipe generation (Kiddon et al., 2016). While these works have focused on developing new neural architectures that introduce structural biases for easier *learning*, our work uses a simple architecture and focuses on improving the optimization of the learner (i.e., better *teaching*).

The importance of better teaching for RNN generators was outlined in Bengio et al. (2015), which showed that *exposure* bias from a misaligned train and test setup limited the capabilities of sequence-to-sequence models. This limitation had been addressed in previous work by augmenting training data with examples generated by pretrained models to make models robust to their own errors (Daumé III et al., 2009; Ross et al., 2011).

More recent work on training RNNs for generation has used sequence scores such as ROUGE (Paulus et al., 2018), CIDEr (Rennie et al., 2017; Pasunuru and Bansal, 2017), BLEU (Ranzato et al., 2015) and mixtures of them (Liu et al., 2017) as a global reward to train a policy with the REINFORCE algorithm (Williams, 1992). In contrast, our work uses a neural teacher to reward a model for capturing discourse semantics.

Most similar to our work is work on using neural and embedding rewards to improve dialogue (Li et al., 2016), image captioning (Ren et al., 2017), simplification (Zhang and Lapata, 2017), and paraphrase generation (Li et al., 2017). While these works use single-sentence similarity rewards for short generation tasks, our work designs teachers to reward long-range ordering patterns.

Finally, our teachers can be seen as rewarding generators that approximate script patterns in recipes. Previous work in learning script knowledge (Schank and Abelson, 1975) has focused on extracting scripts from long texts (Chambers and Jurafsky, 2009; Pichotta and Mooney, 2016), with some of that work focusing on recipes (Kiddon et al., 2015; Mori et al., 2014, 2012). Our teachers implicitly learn this script knowledge and reward recipe generators for exhibiting it.

8 Conclusion

We introduce the *absolute ordering* and *relative ordering* teachers, two neural networks that score a sequence’s adherence to discourse structure in long text. The teachers are used to compute rewards for a self-critical reinforcement learning framework, allowing a recipe generator to be rewarded for capturing temporal semantics of the cooking domain. Empirical results demonstrate that our teacher-trained generator better models the latent event sequences of cooking recipes, and a human evaluation shows that this improvement is mainly due to maintaining semantic coherence in longer recipes.

Acknowledgments

This research was supported in part by NSF (IIS-1524371), DARPA under the CwC program through the ARO (W911NF-15-1-0543) and Samsung Research.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference for Learning Representations*.
- Regina Barzilay and Mirella Lapata. 2005. Modeling local coherence: An entity-based approach. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*.

- Regina Barzilay and Mirella Lapata. 2008. Modeling local coherence: An entity-based approach. *Computational Linguistics* 34(1).
- Regina Barzilay and Lillian Lee. 2004. Catching the drift: Probabilistic content models, with applications to generation and summarization. In *HLT-NAACL*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*.
- Antoine Bosselut, Omer Levy, Ari Holtzman, Corin Ennis, Dieter Fox, and Yejin Choi. 2018. Simulating action dynamics with neural process networks. *Proceedings of the 6th International Conference for Learning Representations*.
- Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Hal Daumé III, John Langford, and Daniel Marcu. 2009. Search-based structured prediction.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9(8).
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*. ACM.
- Chloé Kiddon, Ganesa Thandavam Ponnuraj, Luke Zettlemoyer, and Yejin Choi. 2015. Mise en place: Unsupervised interpretation of instructional recipes. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Chloé Kiddon, Luke Zettlemoyer, and Yejin Choi. 2016. Globally coherent text generation with neural checklist models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Jiwei Li and Eduard H Hovy. 2014. A model of coherence based on distributed sentence representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2017. Paraphrase generation with deep reinforcement learning. *arXiv preprint arXiv:1711.00279*.
- Chin-Yew Lin. 2004. ROUGE: a package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*. Barcelona, Spain, volume 8.
- Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. 2017. Improved image captioning via policy gradient optimization of spider. *Proceedings of the 2017 IEEE International Conference on Computer Vision*.
- Ryan Lowe, Michael Noseworthy, Iulian Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Shinsuke Mori, Hirokuni Maeta, Yoko Yamakata, and Tetsuro Sasada. 2014. Flow graph corpus from recipe texts. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation*.
- Shinsuke Mori, Tetsuro Sasada, Yoko Yamakata, and Koichiro Yoshino. 2012. A machine learning approach to recipe text processing. In *Proceedings of the 1st Cooking with Computer Workshop*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Ramakanth Pasunuru and Mohit Bansal. 2017. Reinforced video captioning with entailment rewards. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *Proceedings of the 6th International Conference for Learning Representations*.
- Karl Pichotta and Raymond J. Mooney. 2016. Using sentence-level lstm language models for script inference. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. In *Proceedings of the 4th International Conference for Learning Representations*.

- Zhou Ren, Xiaoyu Wang, Ning Zhang, Xutao Lv, and Li-Jia Li. 2017. Deep reinforcement learning-based image captioning with embedding reward. *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition* .
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition* .
- Stéphane Ross, Geoffrey J Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Roger C Schank and Robert P Abelson. 1975. *Scripts, plans, and knowledge*. Yale University.
- Abigale See, Peter J. Liu, and Christopher Manning. 2017. Gettothepoint: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition*.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4).
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. Challenges in data-to-document generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* .
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of The 32nd International Conference on Machine Learning*.
- Zichao Yang, Phil Blunsom, Chris Dyer, and Wang Ling. 2017. Reference-aware language models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Xingxing Zhang and Mirella Lapata. 2017. Sentence simplification with deep reinforcement learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.

A Hyperparameters

A.1 Data

Each recipe is batched based on the number of tokens and number of ingredients it has. We use a minibatch size of 32.

A.2 Teachers

The hidden size of the reward generator is 100, the word embeddings have dimensionality 100. We use dropout with a rate of 0.3 between the bag of words layers and the recurrent layers.

A.3 Pretrained Recipe Generator

We use a hidden size of 256 for the encoder and 256 for the decoder. We initialize three different sets of embeddings for the recipe titles, ingredient lists, and text, each of size 256. All models are trained with a dropout rate of 0.3 and are single-layer. We use a temperature coefficient of $\beta = 2$ to make the output word distribution more peaky (Kiddon et al., 2016), allowing for more controlled exploration during self-critical learning. We use scheduled sampling with a linear decay schedule of 5% every 5 epochs up to a max of 50%. We use a learning rate of $\eta = 0.0003$ and train with the Adam optimizer.

A.4 Policy Learning

We use the same model hyperparameters as during pretraining, but re-initialize the Adam optimizer, use $\eta = 3 \times 10^{-5}$ as the learning rate, and do not train with scheduled sampling.

B Baseline Selection

For each baseline we trained, we report the score of the γ setting that achieved the highest score for the metric on which it was trained. For example, for baselines trained with ROUGE-L reward, we report the results for the model trained with the value of γ that scored the highest ROUGE-L score on the development set. For the models trained with the CIDEr reward, we select the model with value of γ that achieved the highest CIDEr score on the development set. We do the same for models trained with BLEU-1 and BLEU-4 rewards. The values of γ yielding the best performance on the development set were 0.97 for the BLEU-1, ROUGE-L, and CIDEr-trained models, and 0.99 for the BLEU-4 trained baseline. For each baseline, the best model is chosen by selecting the checkpoint that achieves the highest reward

(or lowest loss for the MLE model) for the metric it was trained on.