

Discourse-Aware Neural Rewards for Coherent Text Generation

Yahui Liu

Tencent AI Lab

yahui.cvr@gmail.com

June 25, 2018

Introduction

Task

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

Challenge

BLEU, ROUGE, CIDEr, etc

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.

Contribution

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*.

Introduction

Main Idea

- teacher should capture sentence order instead of word order
- teacher should reward sentences individually for how they fit with surrounding sentences

Architecture

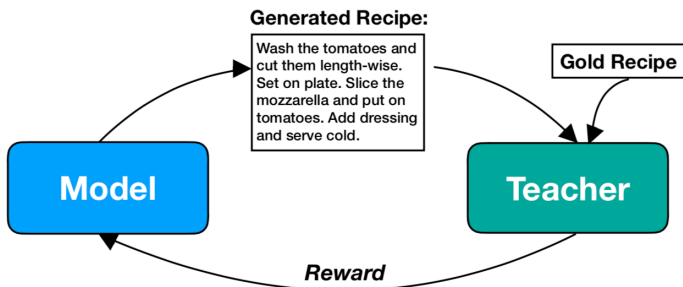
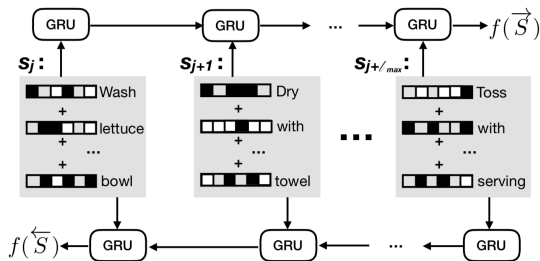


Figure 1: The generator is rewarded for imitating the discourse structure of the gold sequence.

distributing sentence-level rewards

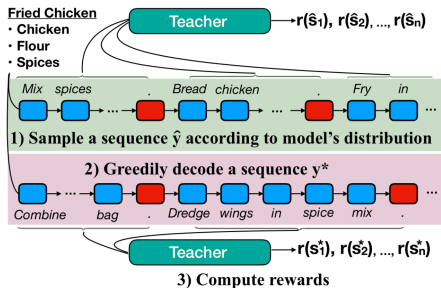
Absolute Order Teacher



$$\mathcal{L}_{abs} = \frac{\langle f(\vec{S}_j), f(\overleftarrow{S}_j) \rangle}{\|f(\vec{S}_j)\| \|f(\overleftarrow{S}_j)\|} \quad (1)$$

$$r_{AO}(y) = \frac{\langle f(S'), f(\vec{S}_j) \rangle}{\|f(S')\| \|f(\vec{S}_j)\|} - \frac{\langle f(S'), f(\overleftarrow{S}_j) \rangle}{\|f(S')\| \|f(\overleftarrow{S}_j)\|} \quad (2)$$

Relative Order Teacher



$$\mathcal{L}_{abs} = \frac{\langle f(\vec{S}_{j:j+l}), f(\overleftarrow{S}_{j:j+l}) \rangle}{\|f(\vec{S}_{j:j+l})\| \|f(\overleftarrow{S}_{j:j+l})\|} \quad (3)$$

$$r_{RO} = \frac{1}{L} \sum_{l=\ell_{min}}^{\ell_{max}} \left(\frac{\langle f(S'_{j-l:j}), f(\vec{S}_{j-l:j}) \rangle}{\|f(S'_{j-l:j})\| \|f(\vec{S}_{j-l:j})\|} - \frac{\langle f(S'_{j-l:j}), f(\overleftarrow{S}_{j-l:j}) \rangle}{\|f(S'_{j-l:j})\| \|f(\overleftarrow{S}_{j-l:j})\|} \right) \quad (4)$$

Experiments

| 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 |

Relative Ordering performs better, Absolute Ordering provides no improvements.

Experiments

| | |
|--------------|---|
| 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. |

Boxed spans indicate recipe events missed by another model's generation.
Red spans indicate superfluous events.

Thanks!