

Grammatical Error Correction with Neural Reinforcement Learning

Keisuke Sakaguchi[†] and Matt Post[‡] and Benjamin Van Durme^{†‡}

[†]Center for Language and Speech Processing, Johns Hopkins University

[‡]Human Language Technology Center of Excellence, Johns Hopkins University

{keisuke, post, vandurme}@cs.jhu.edu

Abstract

We propose a neural encoder-decoder model with reinforcement learning (NRL) for grammatical error correction (GEC). Unlike conventional maximum likelihood estimation (MLE), the model directly optimizes towards an objective that considers a *sentence-level*, task-specific evaluation metric, avoiding the exposure bias issue in MLE. We demonstrate that NRL outperforms MLE both in human and automated evaluation metrics, achieving the state-of-the-art on a fluency-oriented GEC corpus.

1 Introduction

Research in automated Grammatical Error Correction (GEC) has expanded from token-level, closed class corrections (e.g., determiners, prepositions, verb forms) to phrase-level, open class issues that consider fluency (e.g., content word choice, idiomatic collocation, word order, etc.).

The expanded goals of GEC have led to new proposed models deriving from techniques in data-driven machine translation, including phrase-based MT (PBMT) (Felice et al., 2014; Chollampatt et al., 2016; Junczys-Dowmunt and Grundkiewicz, 2016) and neural encoder-decoder models (Yuan and Briscoe, 2016). Napoles et al. (2017) recently showed that a neural encoder-decoder can outperform PBMT on a fluency-oriented GEC data and metric.

We investigate training methodologies in the neural encoder-decoder for GEC. To train the neural encoder-decoder models, maximum likelihood estimation (MLE) has been used, where the objective is to maximize the (log) likelihood of the parameters for a given training data.

As Ranzato et al. (2015) indicates, however, MLE has drawbacks. The MLE objective is based

Algorithm 1: Reinforcement learning for neural encoder-decoder model.

Input: Pairs of source (X) and target (Y)

Output: Model parameter $\hat{\theta}$

```
1 initialize( $\hat{\theta}$ )
2 for  $(x, y) \in (X, Y)$  do
3    $(\hat{y}_1, \dots, \hat{y}_k), (p(\hat{y}_1), \dots, p(\hat{y}_k)) = \text{sample}(x, k, \hat{\theta})$ 
4    $p(\hat{y}) = \text{normalize}(p(\hat{y}))$ 
5    $\bar{r}(\hat{y}) = 0$  // expected reward
6   for  $\hat{y}_i \in \hat{y}$  do
7      $\bar{r}(\hat{y}) += p(\hat{y}_i) \cdot \text{score}(\hat{y}_i, y)$ 
8   backprop( $\hat{\theta}, \bar{r}$ ) // policy gradient  $\frac{\partial}{\partial \hat{\theta}}$ 
9 return  $\hat{\theta}$ 
```

on *word-level* accuracy against the reference, and the model is not exposed to the predicted output during training (exposure bias). This becomes problematic, because once the model fails to predict a correct word, it falls off the right track and does not come back to it easily.

To address the issues, we employ a neural encoder-decoder GEC model with a reinforcement learning approach in which we directly optimize the model toward our final objective (i.e., evaluation metric). The objective of the neural reinforcement learning model (NRL) is to maximize the expected reward on the training data. The model updates the parameters through back-propagation according to the reward from predicted outputs. The high-level description of the training procedure is shown in Algorithm 1, and more details are elaborated in §2. To our knowledge, this is the first attempt to employ reinforcement learning for directly optimizing the encoder-decoder model for GEC task.

We run GEC experiments on a fluency-oriented GEC corpus (§3), demonstrating that NRL outperforms the MLE baseline both in human and automated evaluation metrics.

2 Model and Optimization

We use the *attentional neural encoder-decoder* model (Bahdanau et al., 2014) as a basis for both NRL and MLE. The model takes (possibly ungrammatical) source sentences $x \in X$ as an input, and predicts grammatical and fluent output sentences $y \in Y$ according to the model parameter θ . The model consists of two sub-modules, *encoder* and *decoder*. The encoder transforms x into a sequence of vector representations (hidden states) using a bidirectional gated recurrent neural network (GRU) (Chung et al., 2014). The decoder predicts a word y_t at a time, using previous token y_{t-1} and linear combination of encoder information as attention.

2.1 Maximum Likelihood Estimation

Maximum Likelihood Estimation training (MLE) is a standard optimization method for encoder-decoder models. In MLE, the objective is to maximize the log likelihood of the correct sequence for a given sequence for the entire training data.

$$L(\theta) = \sum_{\langle X, Y \rangle} \sum_{t=1}^T \log p(y_t | x, y_1^{t-1}; \theta) \quad (1)$$

The gradient of $L(\theta)$ is as follows:

$$\frac{\partial L(\theta)}{\partial \theta} = \sum_{\langle X, Y \rangle} \sum_{t=1}^T \frac{\nabla p(y_t | x, y_1^{t-1}; \theta)}{p(y_t | x, y_1^{t-1}; \theta)} \quad (2)$$

One drawback of MLE is the *exposure bias* (Ranzato et al., 2015). The decoder predicts a word conditioned on the correct word sequence (y_1^{t-1}) during training, whereas it does with the predicted word sequence (\hat{y}_1^{t-1}) at test time. Namely, the model is not exposed to the predicted words in training time. This is problematic, because once the model fails to predict a correct word at test time, it falls off the right track and does not come back to it easily. Furthermore, in most sentence generation tasks, the MLE objective does not necessarily correlate with our final evaluation metrics, such as BLEU (Papineni et al., 2002) in machine translation and ROUGE (Lin, 2004) in summarization. This is because MLE optimizes word level predictions at each time step instead of evaluating sentences as a whole.

GEC is no exception. It depends on sentence-level evaluation that considers grammaticality and fluency. For this purpose, it is natural to use GLEU (Napoles et al., 2015), which has been used as a

fluency-oriented GEC metric. We explain more details of this metric in §2.3.

2.2 Neural Reinforcement Learning

To address the issues in MLE, we directly optimize the neural encoder-decoder model toward our final objective for GEC using reinforcement learning. In reinforcement learning, *agents* aim to maximize expected *rewards* by taking *actions* and updating the *policy* under a given *state*. In the neural encoder-decoder model, we treat the encoder-decoder as an agent which predicts a word from a fixed vocabulary at each time step (the action), given the hidden states of the neural encoder-decoder representation. The key difference from MLE is that the reward is not restricted to token-level accuracy. Namely, any arbitrary metric is applicable as the reward.¹

Since we use GLEU as the final evaluation metric, the objective of NRL is to maximize the expected GLEU by learning the model parameter.

$$\begin{aligned} J(\theta) &= \mathbb{E}[r(\hat{y}, y)] \\ &= \sum_{\hat{y} \in S(x)} p(\hat{y} | x; \theta) r(\hat{y}, y) \end{aligned} \quad (3)$$

where $S(x)$ is a sampling function that produces k samples $\hat{y}_1, \dots, \hat{y}_k$, $p(\hat{y} | x; \theta)$ is a probability of the output sentence, and $r(\hat{y}, y)$ is the reward for \hat{y}_k given a reference set y . As described in Algorithm 1, given a pair of source sentence and the reference (x, y) , NRL takes k sample outputs $\hat{y}_1, \dots, \hat{y}_k$ and their probabilities $p(\hat{y}_1), \dots, p(\hat{y}_k)$ (line 3).² Then, the expected reward is computed by multiplying the probability and metric score for each sample \hat{y}_i (line 7).

In the encoder-decoder model, the parameters θ are updated through back-propagation and the number of parameter updates is determined by the partial derivative of $J(\theta)$, called the *policy gradient* (Williams, 1992; Sutton et al., 1999) in reinforcement learning:

$$\frac{\partial J(\theta)}{\partial \theta} = \alpha \mathbb{E}[\nabla \log p(\hat{y}) \{r(\hat{y}, y) - b\}] \quad (4)$$

where α is a learning rate and b is an arbitrary baseline reward to reduce the variance. The sample mean reward is often used for b (Williams, 1992), and we follow it in NRL.

It is reasonable to compare NRL to minimum risk training (MRT) (Shen et al., 2016). In fact,

¹The reward is given at the end of the decoder output (i.e., delayed reward).

²We sampled sentences from softmax distribution.

Corpus	# sents.	mean chars per sent.	# sents. edited
NUCLE	57k	115	38%
FCE	34k	74	62%
Lang-8	1M	56	35%

Table 1: Statistics of training corpora

NRL with a *negative expected reward* can be regarded as MRT. The gradient of MRT objective is a special case of *policy gradient* in NRL. We show mathematical details about the relevance between NRL and MRT in the supplemental material (Appendix A).

2.3 Reward in Grammatical Error Correction

To capture fluency as well as grammaticality in evaluation on such references, we use GLEU as the reward. We have shown GLEU to be more strongly preferred than other GEC metrics by native speakers (Sakaguchi et al., 2016). Similar to BLEU in machine translation, GLEU computes n -gram precision between the system hypothesis (H) and the reference (R). In GLEU, however, n -grams in source (S) are also considered. The precision is penalized when the n -gram in H overlaps with the source and not with the reference.

$$\text{GLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^4 \frac{1}{n} \log p'_n\right)$$

$$p'_n = \frac{N(H, R) - [N(H, S) - N(H, S, R)]}{N(H)}$$

$$\text{BP} = \begin{cases} 1 & \text{if } h > r \\ \exp(1 - r/h) & \text{if } h \leq r \end{cases}$$

where $N(A, B, C, \dots)$ is the number of overlapped n -grams among the sets, and BP *brevity penalty* is compute based on token length in the system hypothesis (h) and the reference (r).

3 Experiments

Data For training the models (MLE and NRL), we use the following corpora: the NUS Corpus of Learner English (NUCLE) (Dahlmeier et al., 2013), the Cambridge Learner Corpus First Certificate English (FCE) (Yannakoudakis et al., 2011), and the Lang-8 Corpus of learner English (Tajiri et al., 2012). The basic statistics are shown in Table 1.³ We exclude some unreasonable edits (comments by editors, incomplete sentences such

³All the datasets are publicly available, for purposes of reproducibility. For more details about each dataset, refer to Sakaguchi et al. (2017).

Models	Methods	# sents. (corpora)
CAMB14	Hybrid (rule + PBMT)	155k (NUCLE, FCE, in-house)
AMU	PBMT + GEC-feat.	2.3M (NUCLE, Lang8)
NUS	PBMT + Neural feat.	2.1M (NUCLE, Lang8)
CAMB16	enc-dec (MLE) + unk alignment	1.96M (non-public CLC)
MLE/NRL	enc-dec (MLE/NRL)	720k (NUCLE, Lang8, FCE)

Table 2: Summary of baselines, MLE and NRL models.

as URLs, etc.) using regular expressions and setting a maximum token edit distance within 50% of the original length. We also ignore sentences that are longer than 50 tokens or sentences where more than 5% of tokens are out-of-vocabulary (the vocabulary size is 35k). In total, we use 720k pairs of sentences for training (21k from NUCLE, 32k from FCE, and 667k from Lang-8). Spelling errors are corrected in preprocessing with the Enchant open-source spell checking library.⁴

Hyperparameters For both MLE and NRL, we set the vocabulary size to be 35k for both source and target. Words are represented by a vector with 512 dimensions. Maximum output token length is 50. The size of hidden layer units is 1,000. Gradients are clipped at 1, and beam size during decoding is 5. We regularize the GRU layer with a dropout probability of 0.2.

For MLE we use mini-batches of size 40, and the ADAM optimizer with a learning rate of 10^{-4} . We train the encoder-decoder with MLE for 900k updates, selecting the best model according to the development set evaluation.

For NRL we set the sample size to be 20. We use the SGD optimizer with a learning rate of 10^{-4} . For the *baseline reward*, we use average of sampled reward following Williams (1992). The sentence GLEU score is used as the reward $r(\hat{y}, y)$. Following a similar (but not the same) strategy of the Mixed Incremental Cross-Entropy Reinforce (MIXER) algorithm (Ranzato et al., 2015), we initialize the model by MLE for 600k updates, followed by another 600k updates using NRL, and select the best model according to the development set evaluation. Our NRL is implemented by extending the Nematus toolkit (Sennrich et al., 2017).⁵

⁴<https://github.com/AbiWord/enchant>

⁵NRL code is available at <https://github.com/keisks/nematus/tree/nrl-gleu>

Models	dev set		test set	
	Human	GLEU	Human	GLEU
Original	-1.072	38.21	-0.760	40.54
AMU	-0.405	41.74	-0.168	44.85
CAMB14	-0.160	42.81	-0.225	46.04
NUS	-0.131	46.27	-0.249	50.13
CAMB16	-0.117	47.20	-0.164	52.05
MLE	-0.052	48.24	-0.110	52.75
NRL	0.169	49.82	0.111	53.98
Reference	1.769	55.26	1.565	62.37

Table 3: Human (TrueSkill) and GLEU evaluation of system outputs on the development and test set.

Baselines In addition to our MLE baseline, we compare four leading GEC systems. All the systems are based on SMT, but they take different approaches. The first model, proposed by Felice et al. (2014), uses a combination of a rule-based system and PBMT with language model reranking (referring as CAMB14). Junczys-Dowmunt and Grundkiewicz (2016) proposed a PBMT model that incorporates linguistic and GEC-oriented sparse features (AMU). Another PBMT model, proposed by Chollampatt et al. (2016), is integrated with neural contextual features (NUS). Finally, Yuan and Briscoe (2016) proposed a neural encoder-decoder model with MLE training (CAMB16). This model is similar to our MLE model, but CAMB16 additionally trains an unsupervised alignment model to handle spelling errors as well as unknown words, and it uses 1.96M sentence pairs extracted from the non-public Cambridge Learner Corpus (CLC). The summary of baselines is shown in Table 2.⁶

Evaluation For evaluation, we use the JFLEG corpus (Heilman et al., 2014; Napoles et al., 2017), which consists of 1501 sentences (754: dev, 747: test) with four fluency-oriented references.

In addition to the automated metric (GLEU), we run a human evaluation using Amazon Mechanical Turk (MTurk). We randomly select 200 sentences each from the dev and test set. For each sentence, two turkers are repeatedly asked to rank five systems randomly selected from all eight: the four baseline models, MLE, NRL, one randomly selected human correction, and the original sentence. We infer the evaluation scores by comparing pairwise rankings with the TrueSkill algorithm (Herbrich et al., 2006; Sakaguchi et al., 2014).

⁶The four baselines are not tuned toward the same dev set as MLE and NRL. Also, they use different training set (Table 2). We compare them just for reference.

Models	Precision	Recall	M2 (F _{0.5})
AMU	69.95	18.81	45.32
CAMB14	65.09	22.84	47.51
NUS	69.59	29.19	54.50
CAMB16	64.35	32.26	53.67
MLE	66.00	34.62	55.87
NRL	65.93	37.28	57.15

Table 4: M2 (F_{0.5}) scores on the dev set.

Models	Precision	Recall	M2 (F _{0.5})
AMU	69.39	20.79	47.29
CAMB14	63.52	23.44	47.33
NUS	68.08	32.30	55.73
CAMB16	65.66	35.93	56.34
MLE	65.19	37.66	56.88
NRL	65.80	40.96	58.68

Table 5: M2 (F_{0.5}) scores on the test set.

	NRL > MLE	NRL = MLE	NRL < MLE
Dev	33%	45%	22%
Test	30%	57%	13%

Table 6: Ratio of pairwise (preference) judgments between NRL and MLE. NRL >MLE: NRL correction is preferred over MLE. NRL <MLE: MLE is preferred over NRL. NRL =MLE: NRL and MLE are tied.

Results Table 3 shows the human evaluation by TrueSkill and automated metric (GLEU). In both dev and test set, NRL outperforms MLE and other baselines in both the human and automatic evaluations. Human evaluation and GLEU scores correlate highly, corroborating the reliability of GLEU. With respect to inter-annotator agreement, Spearman’s rank correlation between Turkers is 55.6 for the dev set and 49.2 for the test set. The correlations are sufficiently high to show the agreement between Turkers, considering the low chance level (i.e., ranking five randomly selected systems consistently between two Turkers).

Table 4 and 5 show the M2 (F_{0.5}) scores (Dahlmeier and Ng, 2012), which compute phrase-level edits between the system hypothesis and source and compare them with the oracle edits. Although this metric has several drawbacks such as underestimation of system performance and indiscrimination between “no change” and “wrong edits” (Felice et al., 2014), we see that the correlation between the M2 scores and human evaluation is still high in the result.

Finally, Table 6 shows the percentages of preference in the pairwise comparisons between NRL and MLE. In both the dev and test sets, around 30% of NRL corrections are preferred over MLE and approximately 50% are tied.

Orig.	but found that successful people use the people money and use there idea for a way to success .
Ref.	But it was found that successful people use other people 's money and use their ideas as a way to success .
MLE	But found that successful people use the people money and use it for a way to success .
NRL	But found that successful people use the people 's money and use their idea for a way to success .
Orig.	Fish firming uses the lots of special products such as fish meal .
Ref.	Fish firming uses a lot of special products such as fish meal .
MLE	Fish contains a lot of special products such as fish meals .
NRL	Fish shops use the lots of special products such as fish meal .

Table 7: Example outputs by MLE and NRL

Analysis Table 7 presents example outputs from MLE and NRL. In the first example, both MLE and NRL successfully corrected the homophone error (*there vs. their*), but MLE changed the meaning of the original sentence by replacing *their idea* to *it*. Meanwhile, NRL made the sentence more grammatical by adding a possessive 's. The second example demonstrates challenging issues for future work in GEC. The correction by MLE looks fairly fluent as well as grammatical, but it is semantically nonsense. The correction by NRL is also fairly fluent and makes sense, but the meaning has been changed too much. For further improvement, better GEC models that are aware of the context or possess world knowledge are needed.

4 Conclusions

We have presented a neural encoder-decoder model with reinforcement learning for GEC. To alleviate the MLE issues (exposure bias and token-level optimization), NRL learns the policy (model parameters) by directly optimizing toward the final objective by treating the final objective as the reward for the encoder-decoder agent. Using a GEC-specific metric, GLEU, we have demonstrated that NRL outperforms the MLE baseline on the fluency-oriented GEC corpus both in human and automated evaluation metrics. As a supplement, we have explained the relevance between minimum risk training (MRT) and NRL, claiming that MRT is a special case of NRL.

Acknowledgments

This work was supported in part by the JHU Human Language Technology Center of Excellence (HLTCOE), and DARPA LORELEI. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes. The views and conclusions contained in this publication are those of the authors and should not be interpreted as representing official policies or endorsements of DARPA or the U.S. Government.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv:1409.0473*.
- Shamil Chollampatt, Duc Tam Hoang, and Hwee Tou Ng. 2016. Adapting grammatical error correction based on the native language of writers with neural network joint models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1901–1911, Austin, Texas. Association for Computational Linguistics.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv:1412.3555*.
- Daniel Dahlmeier and Hwee Tou Ng. 2012. Better evaluation for grammatical error correction. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572, Montréal, Canada. Association for Computational Linguistics.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. Building a large annotated corpus of learner english: The NUS Corpus of Learner English. In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 22–31, Atlanta, Georgia. Association for Computational Linguistics.
- Mariano Felice, Zheng Yuan, Øistein E. Andersen, Helen Yannakoudakis, and Ekaterina Kochmar. 2014. Grammatical error correction using hybrid systems and type filtering. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 15–24, Baltimore, Maryland. Association for Computational Linguistics.
- Michael Heilman, Aoife Cahill, Nitin Madnani, Melissa Lopez, Matthew Mulholland, and Joel Tetreault. 2014. Predicting grammaticality on an ordinal scale. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 174–180, Baltimore, Maryland. Association for Computational Linguistics.
- Ralf Herbrich, Tom Minka, and Thore Graepel. 2006. TrueSkillTM: A Bayesian skill rating system. In

- Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems*, pages 569–576, Vancouver, British Columbia, Canada. MIT Press.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. Phrase-based machine translation is state-of-the-art for automatic grammatical error correction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1546–1556, Austin, Texas. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. Ground truth for grammatical error correction metrics. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 588–593, Beijing, China. Association for Computational Linguistics.
- Courtney Napoles, Keisuke Sakaguchi, and Joel Tetreault. 2017. JFLEG: A fluency corpus and benchmark for grammatical error correction. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 229–234, Valencia, Spain. Association for Computational Linguistics.
- Franz Josef Och. 2003. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 160–167, Sapporo, Japan. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv:1511.06732*.
- Keisuke Sakaguchi, Courtney Napoles, Matt Post, and Joel Tetreault. 2016. Reassessing the goals of grammatical error correction: Fluency instead of grammaticality. *Transactions of the Association for Computational Linguistics*, 4:169–182.
- Keisuke Sakaguchi, Courtney Napoles, and Joel Tetreault. 2017. GEC into the future: Where are we going and how do we get there? In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 180–187, Copenhagen, Denmark. Association for Computational Linguistics.
- Keisuke Sakaguchi, Matt Post, and Benjamin Van Durme. 2014. Efficient elicitation of annotations for human evaluation of machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 1–11, Baltimore, Maryland. Association for Computational Linguistics.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch, Barry Haddow, Julian Hitschler, Marcin Junczys-Dowmunt, Samuel Läubli, Antonio Valerio Miceli Barone, Jozef Mokry, and Maria Nadejde. 2017. Nematus: a toolkit for neural machine translation. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain. Association for Computational Linguistics.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics.
- Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. 1999. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, volume 99, pages 1057–1063.
- Toshikazu Tajiri, Mamoru Komachi, and Yuji Matsumoto. 2012. Tense and aspect error correction for ESL learners using global context. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 198–202, Jeju Island, Korea. Association for Computational Linguistics.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading ESOL texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 180–189, Portland, Oregon. Association for Computational Linguistics.
- Zheng Yuan and Ted Briscoe. 2016. Grammatical error correction using neural machine translation. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 380–386, San Diego, California. Association for Computational Linguistics.

A Minimum Risk Training and Policy Gradient in Reinforcement Learning

We explain the relevance between minimum risk training (MRT) (Shen et al., 2016) and neural reinforcement learning (NRL) for training neural encoder-decoder models. We describe the detailed derivation of gradient in MRT, and show that MRT is a special case of NRL.

As introduced in §2, the model takes ungrammatical source sentences $x \in X$ as an input, and predicts grammatical and fluent output sentences $y \in Y$. The objective function in NRL and MRT are written as follows.

$$J(\theta) = \mathbb{E}[r(\hat{y}, y)] \quad (5)$$

$$R(\theta) = \sum_{(x, Y)} \mathbb{E}[\Delta(\hat{y}, y)] \quad (6)$$

where $r(\hat{y}, y)$ is the *reward* and $\Delta(\hat{y}, y)$ is the *risk* for an output (\hat{y}).

For the sake of simplicity, we consider expected loss in MRT for a single training pair:

$$\begin{aligned} \tilde{R}(\theta) &= \mathbb{E}[\Delta(\hat{y}, y)] \\ &= \sum_{\hat{y} \in S(x)} q(\hat{y}|x; \theta, \alpha) \Delta(\hat{y}, y) \end{aligned} \quad (7)$$

where

$$q(\hat{y}|x; \theta, \alpha) = \frac{p(\hat{y}|x; \theta)^\alpha}{\sum_{\hat{y}' \in S(x)} p(\hat{y}'|x; \theta)^\alpha} \quad (8)$$

$S(x)$ is a sampling function that produces k samples $\hat{y}_1, \dots, \hat{y}_k$, and α is a smoothing parameter for the samples (Och, 2003). Although the direction to optimize (i.e., minimizing or maximizing) is different, we see the similarity between $J(\theta)$ and $\tilde{R}(\theta)$ in the sense that they both optimize models directly towards evaluation metrics.

The partial derivative of $\tilde{R}(\theta)$ with respect to the model parameter θ is derived as follows.

$$\begin{aligned} \frac{\partial \tilde{R}(\theta)}{\partial \theta} &= \frac{\partial}{\partial \theta} \sum_{\hat{y} \in S(x)} q(\hat{y}|x; \theta, \alpha) \Delta(\hat{y}, y) \\ &= \sum_{\hat{y} \in S(x)} \Delta(\hat{y}, y) \frac{\partial}{\partial \theta} q(\hat{y}|x; \theta, \alpha) \end{aligned} \quad (9)$$

We need $\frac{\partial}{\partial \theta} q(\hat{y}|x; \theta, \alpha)$ in (9). For space efficiency, we use $q(\hat{y})$ as $q(\hat{y}|x; \theta, \alpha)$ and $p(\hat{y})$ as $p(\hat{y}|x; \theta)$ below.

$$\begin{aligned} \frac{\partial}{\partial \theta} q(\hat{y}) &= \frac{\partial q(\hat{y})}{\partial p(\hat{y})} \frac{\partial p(\hat{y})}{\partial \theta} \quad (\because \text{chain rule}) \\ &= \frac{\partial q(\hat{y})}{\partial p(\hat{y})} \nabla p(\hat{y}) \end{aligned} \quad (10)$$

For $\frac{\partial q(\hat{y})}{\partial p(\hat{y})}$, by applying the quotient rule to (8),

$$\begin{aligned} \frac{\partial q(\hat{y})}{\partial p(\hat{y})} &= \frac{\{\sum_{\hat{y}'} p(\hat{y}')^\alpha\} \frac{\partial}{\partial p(\hat{y})} p(\hat{y})^\alpha - p(\hat{y})^\alpha \frac{\partial}{\partial p(\hat{y})} \sum_{\hat{y}'} p(\hat{y}')^\alpha}{\{\sum_{\hat{y}'} p(\hat{y}')^\alpha\}^2} \\ &= \frac{\alpha p(\hat{y})^{\alpha-1}}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} - \frac{\alpha p(\hat{y})^\alpha p(\hat{y})^{\alpha-1}}{\{\sum_{\hat{y}'} p(\hat{y}')^\alpha\}^2} \\ &= \alpha \frac{p(\hat{y})^{\alpha-1}}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \left\{ 1 - \frac{p(\hat{y})^\alpha}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \right\} \\ &= \alpha \frac{p(\hat{y})^\alpha}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \frac{1}{p(\hat{y})} \left\{ 1 - \frac{p(\hat{y})^\alpha}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \right\} \end{aligned} \quad (11)$$

Thus, from (10) and (11), (9) is

$$\begin{aligned} \frac{\partial \tilde{R}(\theta)}{\partial \theta} &= \sum_{\hat{y} \in S(x)} \Delta(\hat{y}, y) \nabla p(\hat{y}) \\ &= \left[\alpha \frac{p(\hat{y})^\alpha}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \frac{1}{p(\hat{y})} \left\{ 1 - \frac{p(\hat{y})^\alpha}{\sum_{\hat{y}'} p(\hat{y}')^\alpha} \right\} \right] \\ &= \alpha \mathbb{E} \left[\nabla p(\hat{y}) \cdot \frac{1}{p(\hat{y})} \{ \Delta(\hat{y}, y) - \mathbb{E}[\Delta(\hat{y}, y)] \} \right] \\ &= \alpha \mathbb{E} [\nabla \log p(\hat{y}) \{ \Delta(\hat{y}, y) - \mathbb{E}[\Delta(\hat{y}, y)] \}] \end{aligned} \quad (12)$$

According to the policy gradient theorem for REINFORCE (Williams, 1992; Sutton et al., 1999), the partial derivative of (5) is given as follows:

$$\frac{\partial J(\theta)}{\partial \theta} = \tilde{\alpha} \mathbb{E} [\nabla \log p(\hat{y}) \{ r(\hat{y}, y) - b \}] \quad (13)$$

where $\tilde{\alpha}$ is a *learning rate*⁷ and b is arbitrary *baseline reward* to reduce the variance of gradients. Finally, we see that the gradient of MRT (12) is a special case of policy gradient in REINFORCE (13) with $b = \mathbb{E}[\Delta(\hat{y}, y)]$. It is also interesting to see that the smoothing parameter α works as a part of learning rate ($\tilde{\alpha}$) in NRL.

⁷In this appendix, we use $\tilde{\alpha}$ to distinguish it from smoothing parameter α in MRT.