

A Nested Attention Neural Hybrid Model for Grammatical Error Correction

Jianshu Ji[†], Qinlong Wang[†], Kristina Toutanova^{‡,*},

Yonggen Gong[†], Steven Truong[†], Jianfeng Gao[§]

[†]Microsoft AI & Research [‡]Google Research [§]Microsoft Research, Redmond
^{†§}{jianshuj, qinlwang, yongeg, stevetr, jfgao}@microsoft.com
[‡]kristout@google.com

Abstract

Grammatical error correction (GEC) systems strive to correct both global errors in word order and usage, and local errors in spelling and inflection. Further developing upon recent work on neural machine translation, we propose a new hybrid neural model with nested attention layers for GEC. Experiments show that the new model can effectively correct errors of both types by incorporating word and character-level information, and that the model significantly outperforms previous neural models for GEC as measured on the standard CoNLL-14 benchmark dataset. Further analysis also shows that the superiority of the proposed model can be largely attributed to the use of the nested attention mechanism, which has proven particularly effective in correcting local errors that involve small edits in orthography.

1 Introduction

One of the most successful approaches to grammatical error correction (GEC) is to cast the problem as (monolingual) machine translation (MT), where we translate from possibly ungrammatical English sentences to corrected ones (Brockett et al., 2006; Gao et al., 2010; Junczys-Dowmunt and Grundkiewicz, 2016). Such systems, which are based on phrase-based MT models that are typically trained on large sets of sentence-correction pairs, can correct global errors such as word order and usage and local errors in spelling and inflection. The approach has proven superior to systems based on local classifiers that can only fix focused errors in prepositions, determiners, or inflected forms (Rozovskaya and Roth, 2016).

*This work was conducted while the third author worked at Microsoft Research.

Recently, neural machine translation (NMT) systems have achieved substantial improvements in translation quality over phrase-based MT systems (Sutskever et al., 2014; Bahdanau et al., 2015). Thus, there is growing interest in applying neural systems to GEC (Yuan and Briscoe, 2016; Xie et al., 2016). In this paper, we significantly extend previous work, and explore new neural models to meet the unique challenges of GEC.

The core component of most NMT systems is a sequence-to-sequence (S2S) model which encodes a sequence of source words into a vector and then generates a sequence of target words from the vector. Unlike the phrase-based MT models, the S2S model can capture long-distance, or even global, word dependencies, which are crucial to correcting global grammatical errors and helping users achieve native speaker fluency (Sakaguchi et al., 2016). Thus, the S2S model is expected to perform better on GEC than phrase-based models. However, as we will show in this paper, to achieve the best performance on GEC, we still need to extend the standard S2S model to address several task-specific challenges, which we will describe below.

First, a GEC model needs to deal with an extremely large vocabulary that consists of a large number of words and their (mis)spelling variations. Second, the GEC model needs to capture structure at different levels of granularity in order to correct errors of different types. For example, while correcting spelling and local grammar errors requires only word-level or sub-word level information, e.g., *violets* → *violates* (spelling) or *violate* → *violates* (verb form), correcting errors in word order or usage requires global semantic relationships among phrases and words.

Standard approaches in neural machine translation, also applied to grammatical error correction by Yuan and Briscoe (2016), address the large vocabulary problem by restricting the vocabulary to a limited number of high-frequency words and re-

sorting to standard word translation dictionaries to provide translations for the words that are out of the vocabulary (OOV). However, this approach often fails to take into account the OOVs in context for making correction decisions, and does not generalize well to correcting words that are unseen in the parallel training data. An alternative approach, proposed by Xie et al. (2016), applies a character-level sequence to sequence neural model. Although the model eliminates the OOV issue, it cannot effectively leverage word-level information for GEC, even if it is used together with a separate word-based language model.

Our solution to the challenges mentioned above is a novel, hybrid neural model with nested attention layers that infuse both word-level and character-level information. The architecture of the model is illustrated in Figure 1. The word-level information is used for correcting global grammar and fluency errors while the character-level information is used for correcting local errors in spelling or inflected forms. Contextual information is crucial for GEC. Using the proposed model, by combining embedding vectors and attention at both word and character levels, we model all contextual words, including OOVs, in a unified context vector representation. In particular, as we will discuss in Section 5, the character-level attention layer captures most useful information for correcting local errors that involve small edits in orthography.

Our model differs substantially from the word-level S2S model of Yuan and Briscoe (2016) and the character-level S2S model of Xie et al. (2016) in the way we infuse information at both the word level and the character level. We extend the word-character hybrid model of Luong and Manning (2016), which was originally developed for machine translation, by introducing a character attention layer. This allows the model to learn substitution patterns at both the character level and the word level in an end-to-end fashion, using sentence-correction pairs.

We validate the effectiveness of our model on the CoNLL-14 benchmark dataset (Ng et al., 2014). Results show that the proposed model outperforms all previous neural models for GEC, including the hybrid model of Luong and Manning (2016), which we apply to GEC for the first time. When integrated with a large word-based n-gram language model, our GEC system achieves an $F_{0.5}$ of 45.15 on CoNLL-14, substantially exceeding the previ-




Figure 1: Architecture of Nested Attention Hybrid Model

ously reported top performance of 40.56 achieved by using a neural model and an external language model (Xie et al., 2016).

2 Related Work

A variety of classifier-based and MT-based techniques have been applied to grammatical error correction. The CoNLL-14 shared task overview paper of Ng et al. (2014) provides a comparative evaluation of approaches. Two notable advances after the shared task have been in the areas of combining classifiers and phrase-based MT (Rozovskaya and Roth, 2016) and adapting phrase-based MT to the GEC task (Junczys-Dowmunt and Grundkiewicz, 2016). The latter work has reported the highest performance to date on the task of 49.5 in $F_{0.5}$ score on the CoNLL-14 test set. This method integrates discriminative training toward the task-specific evaluation function, a rich set of features, and multiple large language models. Neural approaches to the task are less explored. We believe that the advances from Junczys-Dowmunt and Grundkiewicz (2016) are complementary to the ones we propose for neural MT, and could be integrated with neural models to achieve even higher performance.

Two prior works explored sequence to sequence neural models for GEC (Xie et al., 2016; Yuan and Briscoe, 2016), while Chollampatt et al. (2016) integrated neural features in a phrase-based system for the task. Neural models were also applied to the related sub-task of grammatical error identification (Schmaltz et al., 2016). Yuan and Briscoe (2016) demonstrated the promise of neural MT for GEC but did not adapt the basic sequence-to-sequence with attention to its unique challenges, falling back to traditional word-alignment models to address vocabulary coverage with a post-processing heuristic. Xie et al. (2016) built a character-level sequence

to sequence model, which achieves open vocabulary and character-level modeling, but has difficulty with global word-level decisions.

The primary focus of our work is integration of character and word-level reasoning in neural models for GEC, to capture global fluency errors and local errors in spelling and closely related morphological variants, while obtaining open vocabulary coverage. This is achieved with the help of character and word-level encoders and decoders with two nested levels of attention. Our model is inspired by advances in sub-word level modeling in neural machine translation. We build mostly on the hybrid model of [Luong and Manning \(2016\)](#) to expand its capability to correct rare words by fine-grained character-level attention. We directly compare our model to the one of [Luong and Manning \(2016\)](#) on the grammar correction task. Alternative methods for MT include modeling of word pieces to achieve open vocabulary ([Sennrich et al., 2016](#)), and more recently, fully character-level modeling ([Lee et al., 2017](#)). None of these models integrate two nested levels of attention although an empirical evaluation of these approaches for GEC would also be interesting.

3 Nested Attention Hybrid Model

Our model is hybrid, and uses both word-level and character-level representations. It consists of a word-based sequence-to-sequence model as a backbone, and additional character-level encoder, decoder, and attention components, which focus on words that are outside the word-level model’s vocabulary.

3.1 Word-based sequence-to-sequence model as backbone

The word-based backbone closely follows the basic neural sequence-to-sequence architecture with attention as proposed by [Bahdanau et al. \(2015\)](#) and applied to grammatical error correction by [Yuan and Briscoe \(2016\)](#). For completeness, we give a sketch here. It uses recurrent neural networks to encode the input sentence and to decode the output sentence.

Given a sequence of embedding vectors, corresponding to a sequence of input words \mathbf{x} :

$$\mathbf{x} = (x_1, \dots, x_T), \quad (1)$$

the encoder creates a corresponding context-

specific sequence of hidden state vectors \mathbf{e} :

$$\mathbf{e} = (h_1, \dots, h_T)$$

The hidden state h_t at time t is computed as: $f_t = \text{GRU}_{\text{enc}_f}(f_{t-1}, x_t)$, $b_t = \text{GRU}_{\text{enc}_b}(b_{t+1}, x_t)$, $h_t = [f_t; b_t]$, where $\text{GRU}_{\text{enc}_f}$ and $\text{GRU}_{\text{enc}_b}$ stand for gated recurrent unit functions as described in [Cho et al. \(2014\)](#). We use the symbol GRU with different subscripts to represent GRU functions using different sets of parameters (for example, we used the enc_f and enc_b subscripts to denote the parameters of the forward and backward word-level encoder units.)

The decoder network is also an RNN using GRU units, and defines a sequence of hidden states $\bar{d}_1, \dots, \bar{d}_S$ used to define the probability of an output sequence y_1, \dots, y_S as follows:

The context vector c_s at time step s is computed as follows:

$$c_s = \sum_{j=1}^T \alpha_{sj} h_j \quad (2)$$

where:

$$\alpha_{sk} = \frac{u_{sk}}{\sum_{j=1}^T u_{sj}} \quad (3)$$

$$u_{sk} = \phi_1(d_s)^T \phi_2(h_k) \quad (4)$$

Here ϕ_1 and ϕ_2 denote feedforward linear transformations followed by a tanh nonlinearity. The next hidden state \bar{d}_s is then defined as:

$$d_s = \text{GRU}_{\text{dec}}(\bar{d}_{s-1}, y_{s-1}),$$

$$\bar{d}_s = \text{ReLU}(W[c_s; d_s])$$

where y_{s-1} is the embedding of the output token at time $s-1$. ReLU indicates rectified linear units ([Hahnloser et al., 2000](#)).

The probability of each target word y_s is computed as: $p(y_s|y_{<s}, \mathbf{x}) = \text{softmax}(g(\bar{d}_s))$, where g is a function that maps the decoder state into a vector of size the dimensionality of the target vocabulary.

The model is trained by minimizing the cross-entropy loss, which for a given (\mathbf{x}, \mathbf{y}) pair is:

$$\text{Loss}(\mathbf{x}, \mathbf{y}) = - \sum_{s=1}^S \log p(y_s|y_{<s}, \mathbf{x}) \quad (5)$$

For parallel training data \mathbb{C} , the loss is:

$$\text{Loss} = - \sum_{(\mathbf{x}, \mathbf{y}) \in \mathbb{C}} \sum_{s=1}^S \log p(y_s|y_{<s}, \mathbf{x})$$

3.2 Hybrid encoder and decoder with two nested levels of attention

The word-level backbone models a limited vocabulary of source and target words, and represents out-of-vocabulary tokens with special UNK symbols. In the standard word-level NMT approach, valuable information is lost for source OOV words and target OOV words are predicted using post-processing heuristics.

Hybrid encoder

Our hybrid architecture overcomes the loss of source information in the word-level backbone by building up compositional representations of the source OOV words using a character-level recurrent neural network with GRU units. These representations are used in place of the special source UNK embeddings in the backbone, and contribute to the contextual encoding of all source tokens.

For example, a three word input sentence where the last term is out-of-vocabulary will be represented as the following vector of embeddings in the word-level model: $\mathbf{x} = (x_1, x_2, x_3)$, where x_3 would be the embedding for the UNK symbol.

The hybrid encoder builds up a word embedding for the third word based on its character sequence: x^c_1, \dots, x^c_M . The encoder computes a sequence of hidden states \mathbf{e}_c for this character sequence, by a forward character-level GRU network:

$$\mathbf{e}_c = (h^c_1, \dots, h^c_M), \quad (6)$$

The last state h^c_M is used as an embedding of the unknown word. The sequence of embeddings for our example three-word sequence becomes: $\mathbf{x} = (x_1, x_2, h^c_M)$. We use the same dimensionality for word embedding vectors x_i and composed character sequence vectors h^c_M to ensure the two ways to define embeddings are compatible. Our hybrid source encoder architecture is similar to the one proposed by [Luong and Manning \(2016\)](#).

Nested attention hybrid decoder

In traditional word-based sequence-to-sequence models special target UNK tokens are used to represent outputs that are outside the target vocabulary. A post-processing UNK-replacement method is then used ([Cho et al., 2015](#); [Yuan and Briscoe, 2016](#)) to replace these special tokens with target words. The hybrid model of ([Luong and Manning, 2016](#)) uses a jointly trained character-level decoder

to generate target words corresponding to UNK tokens, and outperforms the traditional approach in the machine translation task.

However, unlike machine translation, models for grammar correction conduct “translation” in the same language, and often need to apply a small number of local edits to the character sequence of a source word corresponding to the target UNK word. For example, rare but correct words such as entity names need to be copied as is, and local spelling errors or errors in inflection need to be corrected. The architecture of [Luong and Manning \(2016\)](#) does not have direct access to a source character sequence, but only uses a single fixed-dimensionality embedding of source unknown words aggregated with additional contextual information from the source.

To address the needs of the grammatical error correction task, we propose a novel hybrid decoder with two nested levels of attention: word level and character-level. The character-level attention serves to provide the decoder with direct access to the relevant source character sequence.

More specifically, the probability of each target word is defined as follows: For words in the target vocabulary, the probability is defined by the word-level backbone. For words outside the vocabulary, the probability of each token is the probability of UNK according to the backbone, multiplied by the probability of the word’s character sequence.

The probability of the target character sequence corresponding to an UNK token at position s in the target is defined using a character-level decoder. As in [Luong and Manning \(2016\)](#), the “separate path” architecture is used to capture the relevant context and define the initial state for the character-level decoder:

$$\hat{d}_s = \text{ReLU}(\hat{W}[c_s; d_s])$$

where \hat{W} are parameters different from W , and \hat{d}_s is not used by the word-level model in predicting the subsequent tokens, but is only used to initialize the character-level decoder.

To be able to attend to the relevant source character sequence when generating the target character sequence, we use the concept of hard attention ([Xu et al., 2015](#)), but use an arg-max approximation for inference instead of sampling. A similar approach to represent discrete hidden structure in a variety of architectures is used in [Kong et al. \(2017\)](#).

The source index z_s corresponding to the target

position s is defined according to the word-level attention model:

$$z_s = \arg \max_{k \in 0 \dots T-1} \alpha_{sk}$$

where α_{sk} are the intermediate outputs of the word-level attention model we described in Eq.(3).

The character-level decoder generates a character sequence $\mathbf{y}^c = (y^c_1, \dots, y^c_N)$, conditioned on the initial vector \hat{d}_s and the source index z_s . The characters are generated using a hidden state vector d^c_n at each time step, via a $\text{softmax}(gc(d^c_n))$, where gc maps the state to the target character vocabulary space.

If the source word x_{z_s} is in the source vocabulary, the model is analogous to the one of [Luong and Manning \(2016\)](#) and does not use character-level attention: the source context is available only in aggregated form to initialize the state of the decoder. The state d^c_n for step n in the character-level decoder is defined as follows, where $\text{GRU}^c_{\text{dec}}$ are parameters for the gated recurrent cell of this decoder:

$$d^c_n = \begin{cases} \text{GRU}^c_{\text{dec}}(\hat{d}_s, y^c_{n-1}) & n = 0 \\ \text{GRU}^c_{\text{dec}}(d^c_{n-1}, y^c_{n-1}) & n > 0 \end{cases}$$

In contrast, if the corresponding token in the source x_{z_s} is also an out-of-vocabulary word, we define a second nested level of character attention and use it in the character-level decoder. The character-level attention focuses on individual characters from the source word x_{z_s} . If \mathbf{e}_c are the source character hidden vectors computed as in Eq.(6), the recurrence equations for the character-level decoder with nested attention are:

$$\bar{d}^c_n = \text{ReLU}(W_c[c^c_n; d^c_n])$$

$$d^c_n = \begin{cases} \text{GRU}^c_{\text{decNested}}(\hat{d}_s, y^c_{n-1}) & n = 0 \\ \text{GRU}^c_{\text{decNested}}(\bar{d}^c_{n-1}, y^c_{n-1}) & n > 0 \end{cases}$$

where c^c_n is the context vector obtained using character-level attention on the sequence \mathbf{e}^c and the last state of the character-level decoder d^c_n , computed following equations 2, 3 and 4, but using a different set of parameters.

These equations show that the character-level decoder with nested attention can use both the word-level state \hat{d}_s , and the character-level context c^c_n

and hidden state d^c_n to perform global and local editing operations.

Since we introduced two architectures for the character-level decoder depending on whether the source word x_{z_s} is OOV, the combined loss function is defined as follows for end-to-end training:

$$Loss_{total} = Loss_w + \alpha Loss_{c1} + \beta Loss_{c2}$$

Here $Loss_w$ is the standard word-level loss in Eq.(5); character level losses $Loss_{c1}$ and $Loss_{c2}$ are losses for target OOV words corresponding to source known and unknown tokens, respectively. α and β are hyper-parameters to balance the loss terms.

As seen, our proposed nested attention hybrid model uses character-level attention only when both a predicted target word and its corresponding source input word are OOV. While the model can be naturally generalized to integrate character-level attention for known words, the original hybrid model proposed by [Luong and Manning \(2016\)](#) does not use any character-level information for known words. Thus for a controlled evaluation of the impact of the addition of character-level attention only, in this paper we limit character-level attention to OOV words, which already use characters as a basis for building their embedding vectors. A thorough investigation of the impact of character-level information in the encoder, attention, and decoder for known words as well is an interesting topic for future research.

Decoding for word-level and hybrid models

Beam-search is used to decode hypotheses according to the word-level backbone model. For the hybrid model architecture, word-level beam search is conducted first; for each target UNK token, character-level beam-search is used to generate a corresponding target word.

4 Experiments

4.1 Dataset and Evaluation

We use standard publicly available datasets for training and evaluation. One data source is the NUS Corpus of Learner English (NUCLE) ([Dahlmeier et al., 2013](#)), which is provided as a training set for the CoNLL-13 and CoNLL-14 shared tasks. From the original corpus of size about 60K parallel sentences, we randomly selected close to 5K sentence pairs for use as a validation set, and 45K parallel sentences for use in training. A second data source

| | Training | Validation | Development | Test |
|-------------|-----------|------------|-------------|-------|
| #Sent pairs | 2,608,679 | 4,771 | 1,381 | 1,312 |

Table 1: Overview of the datasets used.

| Source | #Sent pairs |
|--------|-------------|
| NUCLE | 45,422 |
| CLC | 1,517,174 |
| lang-8 | 1,046,083 |
| Total | 2,608,679 |

Table 2: Training data by source.

is the Cambridge Learner Corpus (CLC) (Nicholls, 2003), from which we extracted a substantially larger set of parallel sentences. Finally, we used additional training examples from the Lang-8 Corpus of Learner English v1.0 (Tajiri et al., 2012). As Lang-8 data is crowd-sourced, we used heuristics to filter out noisy examples: we removed sentences longer than 100 words and sentence pairs where the correction was substantially shorter than the input text. Table 2 shows the number of sentence pairs from each source used for training.

We evaluate the performance of the models on the standard sets from the CoNLL-14 shared task (Ng et al., 2014). We report final performance on the CoNLL-14 test set without alternatives, and analyze model performance on the CoNLL-13 development set (Dahlmeier et al., 2013). We use the development and validation sets for model selection. The sizes of all datasets in number of sentences are shown in Table 1. We report performance in $F_{0.5}$ -measure, as calculated by the `m2scorer`—the official implementation of the scoring metric in the shared task.¹ Given system outputs and gold-standard edits, `m2scorer` computes the $F_{0.5}$ measure of a set of system edits against a set of gold-standard edits.

4.2 Baseline

We evaluate our model in comparison to the strong baseline of a word-based neural sequence-to-sequence model with attention, with post-processing for handling out-of-vocabulary words (Yuan and Briscoe, 2016); we refer to this model as word NMT+UNK replacement. Like Yuan and Briscoe (2016), we use a traditional word-alignment model (GIZA++) to derive a word-correction lexicon from the parallel training set. However, in decoding, we don’t use GIZA++ to find the corresponding source word for each tar-

¹<http://www.comp.nus.edu.sg/~nlp/sw/m2scorer.tar.gz>

get OOV, but follow Cho et al. (2015), Section 3.3 to use the NMT system’s attention weights instead. The target OOV is then replaced by the most likely correction of the source word from the word-correction lexicon, or by the source word itself if there are no available corrections.

4.3 Training Details and Results

The embedding size for all word and character-level encoders and decoders is set to 1000, and the hidden unit size is also 1000. To reproduce the model of Yuan and Briscoe (2016), we selected the word vocabulary for the baseline by choosing the 30K most frequent words in the source and target respectively to form the source and target vocabularies. In preliminary experiments for the hybrid models, we found that selecting the same vocabulary of 30K words for the source and target based on combined frequency was better (.003 in $F_{0.5}$) and use that method for vocabulary selection instead. However, there was no gain observed by using such a vocabulary selection method in the baseline. Although the source and target vocabularies in the hybrid models are the same, like in the word-level model, the embedding parameters for source and target words are not shared.

The hyper-parameters for the losses in our models are selected based on the development set and set as follows: $\alpha = \beta = 0.5$. All models are trained with mini-batch size of 128 (batches are shuffled), initial learning rate of 0.0003 and a 0.95 decay ratio if the cost increases in two consecutive 100 iterations. The gradient is rescaled whenever its norm exceeds 10, and dropout is used with a probability of 0.15. Parameters are uniformly initialized in $[-\frac{\sqrt{3}}{\sqrt{1000}}, \frac{\sqrt{3}}{\sqrt{1000}}]$.

We perform inference on the validation set every 5000 iterations to log word-level cost and character-level costs; we save parameter values for the model every 10000 iterations as well as the end of each epoch. The stopping point for training is selected based on development set $F_{0.5}$ among the top 20 parameter sets with best validation set value of the loss function. Training of the nested attention hybrid model takes approximately five days on a Tesla k40m GPU. The basic hybrid model trains in around four days and the word-level backbone trains in approximately three days.

Table 3 shows the performance of the baseline and our nested attention hybrid model on the development and test sets. In addition to the word-level

| Model | Performance | |
|-------------------------------|--------------|--------------|
| | Dev | Test |
| Word NMT + UNK replacement | 26.17 | 38.77 |
| Hybrid model | 28.49 | 40.44 |
| Nested Attention Hybrid Model | 28.61 | 41.53 |

Table 3: $F_{0.5}$ results on the CoNLL-13 and CoNLL-14 test sets of main model architectures.

baseline, we include the performance of a hybrid model with a single level of attention, which follows the work of [Luong and Manning \(2016\)](#) for machine translation, and is the first application of a hybrid word/character-level model to grammatical error correction. Based on hyper-parameter selection, the character-level component weight of the loss is $\alpha = 1$ for the basic hybrid model.

As shown in Table 3, our implementation of the word NMT+UNK replacement baseline approaches the performance of the one reported in [Yuan and Briscoe \(2016\)](#) (38.77 versus 39.9). We attribute the difference to differences in the training set and the word-alignment methods used. Our reimplementation serves to provide a controlled experimental evaluation of the impact of hybrid models and nested attention on the GEC task. As seen, our nested attention hybrid model substantially improves upon the baseline, achieving a gain of close to 3 points on the test set. The hybrid word/character model with a single level of attention brings a large improvement as well, showing the importance of character-level information for this task. We delve deeper into the impact of nested attention for the hybrid model in Section 5.

4.4 Integrating a Web-scale Language Model

The value of large language models for grammatical error correction is well known, and such models have been used in classifier and MT-based systems. To establish the potential of such models in word-based neural sequence-to-sequence systems, we integrate a web-scale count-based language model. In particular, we use the modified Kneser-Ney 5-gram language model trained from Common Crawl ([Buck et al., 2014](#)), made available for download by [Junczys-Dowmunt and Grundkiewicz \(2016\)](#).

Candidates generated by neural models are re-ranked using the following linear interpolation of log probabilities: $s_{y|x} = \log P_{NN}(y|x) + \lambda \log P_{LM}(y)$. Here λ is a hyper-parameter that balances the weights of the neural network model and the language model. We tuned λ separately

| Model | Performance | |
|---|--------------|--------------|
| | Dev | Test |
| Character-based NMT + LM (Xie et al., 2016) | | 40.56 |
| Word NMT + UNK replacement + LM | 31.73 | 42.82 |
| Hybrid model + LM | 33.21 | 44.99 |
| Nested Attention Hybrid Model + LM | 33.47 | 45.15 |

Table 4: $F_{0.5}$ results on the CoNLL-13 and CoNLL-14 test sets of main model architectures, when combined with a large language model.

for each neural model variant, by exploring values in the range $[0.0, 2.0]$ with step size 0.1, and selecting according to development set $F_{0.5}$. The selected values of λ are: 1.6 for word NMT + UNK replacement and 1.0 for the nested attention model.

Table 4 shows the impact of the LM when combined with the neural models implemented in this work. The table also lists the results reported by [Xie et al. \(2016\)](#), for their character-level neural model combined with a large word-level language model. Our best results exceed the ones reported in the prior work by more than 4 points, although we should note that [Xie et al. \(2016\)](#) used a smaller parallel data set for training.

5 Analysis

We analyze the impact of sub-word level information and the two nested levels of attention in more detail by looking at the performance of the models on different segments of the data. In particular, we analyze the performance of the models on sentences containing OOV source words versus ones without OOV words, and corrections to orthographically similar versus dissimilar word forms.

5.1 Performance by Segment: OOV versus Non-OOV

We present a comparative performance analysis of models on the CoNLL-13 development set. First, we divide the set into two segments: OOV and NonOOV, based on whether there is at least one OOV word in the given source input. Table 5 shows that both hybrid architectures substantially outperform the word-level model in both segments of the data. The additional nested character-level attention of our hybrid model brings a sizable improvement over the basic hybrid model in the OOV segment and a small degradation in the non-OOV segment. We should note that in future work character-level attention can be added for non-OOV source words in the nested attention model, which could improve performance on this segment as well.

| Model | NonOOV | OOV | Overall |
|-------------------------------|--------------|--------------|--------------|
| Word NMT + UNK replacement | 27.61 | 21.57 | 26.17 |
| Hybrid model | 29.36 | 25.92 | 28.49 |
| Nested Attention Hybrid Model | 29.00 | 27.39 | 28.61 |

Table 5: $F_{0.5}$ results on the CoNLL-13 set of main model architectures, on different segments of the set according to whether the input contains OOVs.

| | |
|-------------------------------|---|
| source | This greatly violets the rights of people . |
| gold | This greatly violates the rights of people . |
| word NMT + UNK replacement | This greatly violets the rights of people . |
| Nested Attention Hybrid Model | This greatly violates the rights of people . |

Table 6: An example sentence from the OOV segment where the nested attention hybrid model improves performance.

Table 6 shows an example where the nested attention hybrid model successfully corrects a misspelling resulting in an OOV word on the source, whereas the baseline word-level system simply copies the source word without fixing the error (since this particular error is not observed in the parallel training set).

5.2 Impact of Nested Attention on Different Error Types

To analyze more precisely the impact of the additional character-level attention introduced by our design, we continue to investigate the OOV segment in more detail.

The concept of *edit*, which is also used by the official M2 score metric, is defined as a minimal pair of corresponding sub-strings in a source sentence and a correction. For example, in the sentence fragment pair: “Even though there is a risk of causing **harms** to someone, people still **are prefers** to keep their pets without a leash.” → “Even though there is a risk of causing **harm** to someone, people still **prefer** to keep their pets without a leash.”, the minimal edits are “harms → harm” and “are prefers → prefer”. The $F_{0.5}$ score is computed using weighted precision and recall of the set of a system’s edits against one or more sets of reference edits.

For our in-depth analysis, we classify edits in the OOV segment into two types: *small changes* and *large changes*, based on whether the source and target phrase of the edit are orthographically similar or not. More specifically, we say that the target and

| Model | Performance | | |
|-------------------------------|-------------|-------|-----------|
| | P | R | $F_{0.5}$ |
| Small Changes Portion | | | |
| Hybrid model | 43.86 | 16.29 | 32.77 |
| Nested Attention Hybrid Model | 48.25 | 17.92 | 36.04 |
| Large Changes Portion | | | |
| Hybrid model | 32.52 | 8.32 | 20.56 |
| Nested Attention Hybrid Model | 33.05 | 8.11 | 20.46 |

Table 7: Precision, Recall and $F_{0.5}$ results on CoNLL-13, on the “small changes” and “large changes” portions of the OOV segment.

source phrases are orthographically similar, iff: the character edit distance is at most 2 and the source or target is at most 8 characters long, or $edit_ratio < 0.25$, where $edit_ratio = \frac{character_edit_distance}{\min(len(src), len(tar)) + 0.1}$, $len(*)$ denotes number of characters in $*$, and src and tgt denote the pairs in the edit. There are 307 gold edits in the “small changes” portion of the CoNLL-13 OOV segment, and 481 gold edits in the “large changes” portion.

Our hypothesis is that the additional character-level attention layer is particularly useful to model edits among orthographically similar words. Table 7 contrasts the impact of character-level attention on the two portions of the data. We can see that the gains in the “small changes” portion are indeed quite large, indicating that the fine-grained character-level attention empowers the model to more accurately correct confusions among phrases with high character-level similarity. The impact in the “large changes” portion is slightly positive in precision and slightly negative in recall. Thus most of the benefit of the additional character-level attention stems from improvements in the “small changes” portion.

Table 8 shows an example input which illustrates the precision gain of the nested attention hybrid model. The input sentence has a source OOV word which is correct. The hybrid model introduces an error in this word, because it uses only a single source context vector, aggregating the character-level embedding of the source OOV word together with other source words. The additional character-level attention layer in the nested hybrid model enables the correct copying of this long source OOV word, without employing the heuristic mechanism of the word-level NMT system.

| | |
|-------------------------------|---|
| source | Population ageing : A more and more attention-getting topic |
| gold | Population ageing : A more and more attention-getting topic |
| Word NMT + UNK replacement | Population ageing : A more and more attention-getting topic |
| Hybrid Model | Population ageing : A more and more attention-teghting topic |
| Nested Attention Hybrid Model | Population ageing : A more and more attention-getting topic |

Table 8: An example where the nested attention hybrid model outperforms the non-nested model.

6 Conclusions

We have introduced a novel hybrid neural model with two nested levels of attention: word-level and character-level. The model addresses the unique challenges of the grammatical error correction task and achieves the best reported results on the CoNLL-14 benchmark among fully neural systems. Our nested attention hybrid model deeply combines the strengths of word and character level information in all components of an end-to-end neural model: the encoder, the attention layers, and the decoder. This enables it to correct both global word-level and local character-level errors in a unified way. The new architecture contributes substantial improvement in correction of confusions among rare or orthographically similar words compared to word-level sequence-to-sequence and non-nested hybrid models.

Acknowledgements

We would like to thank the ACL reviewers for their insightful suggestions, Victoria Zayats for her help with reproducing the baseline word-level NMT system and Yu Shi, Daxin Jiang and Michael Zeng for the helpful discussions.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.

Chris Brockett, William B Dolan, and Michael Gamon. 2006. Correcting ESL errors using phrasal SMT techniques. In *Proceedings of the 21st International*

Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics. pages 249–256.

Christian Buck, Kenneth Heafield, and Bas Van Ooyen. 2014. N-gram counts and language models from the Common Crawl. In *LREC*.

Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, 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 (EMNLP)*.

Sébastien Jean Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proceedings of ACL*.

Shamil Chollampatt, Kaveh Taghipour, and Hwee Tou Ng. 2016. Neural network translation models for grammatical error correction. In *Proceedings of IJCAI*.

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.

Jianfeng Gao, Xiaolong(Shiao-Long) Li, Daniel Micol, Chris Quirk, and Xu Sun. 2010. A large scale ranker-based system for search query spelling correction. In *The 23rd International Conference on Computational Linguistics*.

Richard HR Hahnloser, Rahul Sarpeshkar, Misha A Mahowald, Rodney J Douglas, and H Sebastian Seung. 2000. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature* 405(6789):947–951.

Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. Phrase-based machine translation is state-of-the-art for automatic grammatical error correction. In *EMNLP*.

Lingpeng Kong, Chris Alberti, Daniel Andor, Ivan Bogatyy, and David Weiss. 2017. Dragnn: A transition-based framework for dynamically connected neural networks.

Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. Fully character-level neural machine translation without explicit segmentation. *TACL* 5.

Minh-Thang Luong and Christopher D. Manning. 2016. Achieving open vocabulary neural machine translation with hybrid word-character models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. The CoNLL-2014 shared task on grammatical error correction. In *CoNLL Shared Task*. pages 1–14.
- Diane Nicholls. 2003. The Cambridge Learner Corpus: Error coding and analysis for lexicography and ELT. In *Proceedings of the Corpus Linguistics 2003 conference*. volume 16, pages 572–581.
- Alla Rozovskaya and Dan Roth. 2016. Grammatical error correction: Machine translation and classifiers. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pages 2205–2215.
- 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.
- Allen Schmaltz, Yoon Kim, Alexander M Rush, and Stuart M Shieber. 2016. Sentence-level grammatical error identification as sequence-to-sequence correction. *arXiv preprint arXiv:1604.04677*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. pages 3104–3112.
- 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: Short Papers-Volume 2*. pages 198–202.
- Ziang Xie, Anand Avati, Naveen Arivazhagan, Dan Jurafsky, and Andrew Y. Ng. 2016. [Neural language correction with character-based attention](http://arxiv.org/abs/1603.09727). *CoRR* abs/1603.09727. <http://arxiv.org/abs/1603.09727>.
- 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 *ICML*. volume 14, pages 77–81.
- 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*.