

# Recurrent Inference in Text Editing

**Ning Shi**

Georgia Institute of Technology  
ning.shi@gatech.edu

**Haotian Zhang**

Learnable, Inc.  
haotian.zhang@learnable.ai

**Ziheng Zeng**

University of Illinois  
zzeng13@illinois.edu

**Yichen Gong**

Horizon Robotics, Inc.  
yichen01.gong@horizon.ai

## Abstract

In neural text editing, prevalent sequence-to-sequence based approaches directly map the unedited text either to the edited text or the editing operations, in which the performance is degraded by the limited source text encoding and long, varying decoding steps. To address this problem, we propose a new inference method, *Recurrence*, that iteratively performs editing actions, significantly narrowing the problem space. In each iteration, encoding the partially edited text, *Recurrence* decodes the latent representation, generates an action of short, fixed-length, and applies the action to complete a single edit. For a comprehensive comparison, we introduce three types of text editing tasks: Arithmetic Operators Restoration (AOR), Arithmetic Equation Simplification (AES), Arithmetic Equation Correction (AEC). Extensive experiments on these tasks with varying difficulties demonstrate that *Recurrence* achieves improvements over conventional inference methods.

## 1 Introduction

For text editing, the sequence-to-sequence (seq2seq) framework has been applied to text simplification (Narayan and Gardent, 2014; Dong et al., 2019), punctuation restoration (Tilk and Alumäe, 2016; Kim, 2019), grammatical error correction (Ge et al., 2018; Lichtarge et al., 2018; Zhao et al., 2019), machine translation post-editing (Libovický et al., 2016; Bérard et al., 2017), and etc. We observe that current inference methods can be roughly grouped into two categories: *End-to-end* (End2end) (Nisioi et al., 2017; See et al., 2017; Tan et al., 2017; Junczys-Dowmunt et al., 2018) and *Tagging* (Filippova et al., 2015; Che et al., 2016; Libovický et al., 2016; Wang et al., 2017; Alva-Manchego et al., 2017; Kim, 2019). For models from both categories, the encoders extract and encode information from the source text sequence. Yet, the goal of the

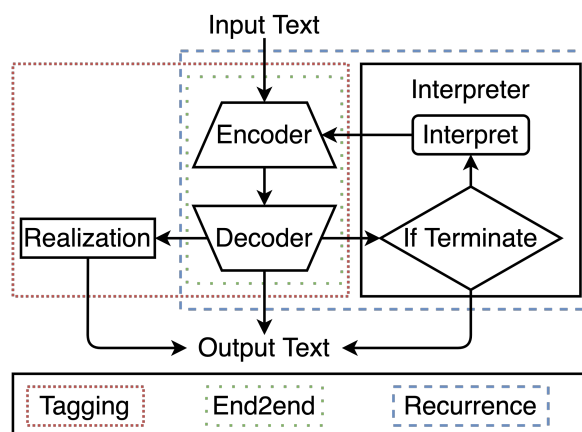


Figure 1: High-level illustration of End2end, Tagging, and Recurrence in text editing.

decoders is different for End2end and Tagging. Upon receiving the encoder’s hidden states that comprise the source text information, the decoder of End2end directly decodes the hidden states and generates the completely edited target text sequence. But, the decoder of Tagging produces a sequence of editing operations, such as deletion and insertion, that is later applied to the source text to yield the edited text via a realization step (Malmi et al., 2019). The mechanisms of End2end and Tagging are illustrated in Figure 1.

However, both End2end and Tagging are problematic because as decoding progresses, the divergence between the partially edited text and the original text grows, rendering the encoder hidden states less and less helpful for decoding the edited text or editing operations toward the end of the editing process; and as the number of decoding steps increases with edited text length, decoding the completely edited text or the full editing operation sequence becomes more and more demanding.

To tackle the aforementioned issues, we propose a recurrent inference method, *Recurrence*, for text-editing with the encoder-decoder framework. *Recurrence* consists of two components as illustrated

in Figure 1: (i) an encoder-decoder model, namely the *programmer*; (ii) an *interpreter*. For a given source sequence, the programmer determines an *editing action* that consists of an editing operation with the tokens it needs and the position in the source sequence to apply the operation. After the interpreter executes the editing action, the partially edited text is again fed to the programmer to determine the next appropriate editing action. This process repeats until the programmer decides that no further editing is needed.

Intuitively, Recurrence is advantageous because (i) as a novel recurrent inference process, it is not constrained by model structures and generally applicable; (ii) the programmer only produces one single editing step, easing the learning difficulty; (iii) the encoder hidden states are updated for each decoding step, providing faithful latent representations; (iv) the decoder outputs an editing action of fixed sequence length, alleviating the problem caused by long decoding steps. Empirically, through three text editing tasks, namely Arithmetic Operators Restoration (AOR), Arithmetic Equation Simplification (AES) and Arithmetic Equation Correction (AEC), we show that Recurrence is data-efficient and more resilient to the text sequence length and the vocabulary size.

Our contributions are the followings: (1) we demonstrate that many text editing tasks can be solved by multiple inference steps recurrently; (2) we propose a novel recurrent inference method, Recurrence, for text editing that tears an editing task down into iterations of editing actions; (3) we design three easily reproducible, proof-of-concept text editing tasks, AOR, AES and AEC; (4) we exhibit that Recurrence outperforms End2end and Tagging in all three text editing tasks and is (i) less sensitive to longer sequences; (ii) less sensitive to larger vocab sizes; (iii) less data-hungry to achieve superior or competitive performances.

The code for three inference methods, text editing tasks, data generation, and experiments in this work is available at: <https://github.com/ShiningLab/Recurrent-Text-Editing>.

## 2 Related Work

**Text Editing** is an Natural Language Processing (NLP) task in that systems change texts by inserting, deleting and rephrasing the words to meet certain needs. According to the length relationship between input and output texts, we summarize text

editing tasks into three types: *short-to-long*, *long-to-short*, and *mixed*.

**End-to-end** is one of the early methods to perform text editing by casting the job as seq2seq (Sutskever et al., 2014) text generation. Without complicated preparation and subsequent processing, End2end has been proven to accomplish text editing well, in all three types (Tilk and Alumäe, 2016; Nisioi et al., 2017; See et al., 2017; Tan et al., 2017; Junczys-Dowmunt et al., 2018; Zhao et al., 2019). Yet, conventional seq2seq-based approaches are well-known for their drawbacks, including dependency on large amounts of data, unexplainable processes, and uncontrollable outcomes (Wiseman et al., 2018). When texts do not need a complete modification, there are more appropriate methods than learning a direct mapping from unedited texts to edited texts.

**Tagging** solves text editing in two steps instead. It firstly employs a seq2seq framework to produce tag sequences, and secondly, edits input texts according to the tag sequences (the “realization” step) (Malmi et al., 2019). Tagging assigns the tag `KEEP` for words that do not need to be changed so that it does not need to learn a copy mechanism. Some have reported that Tagging is better than End2end in short-to-long (Che et al., 2016; Kim, 2019), long-to-short (Filippova et al., 2015; Alva-Manchego et al., 2017; Wang et al., 2017), and mixed editing (Libovický et al., 2016; Bérard et al., 2017; Malmi et al., 2019). One notable member of the Tagging family is Neural Programmer-Interpreter (NPI), a recurrent and compositional neural network (Reed and de Freitas, 2016). NPI is adopted in text editing to predict tags, such as `KEEP`, `DELETE`, and `INSERT`, and execute operations during decoding simultaneously. NPI-based methods have achieved state-of-the-art results in long-to-short (Dong et al., 2019; Gu et al., 2019), and mixed editing (Vu and Haffari, 2018). Nevertheless, like other Tagging methods, NPI’s encoder hidden states are not updated during editing. Its decoder considers operations and executions from previous time steps to predict the current operation while putting massive pressure on the decoder (Hochreiter, 1998; Bahdanau et al., 2015; Cho et al., 2014). Also, Tagging in general suffers from a performance decline caused by a large vocabulary that combines tags and words or too many decoding steps to assign tags. To resolve the aforementioned problems with Tagging, in Recurrence, we update the encoder

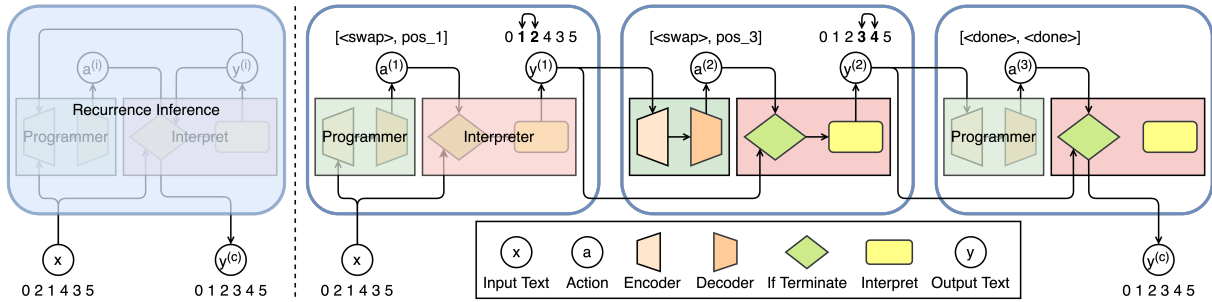


Figure 2: Illustrate Recurrence inference for text editing; the example shows an number ordering task where the number sequence  $[0, 2, 1, 4, 3, 5]$  is edited to  $[0, 1, 2, 3, 4, 5]$  via action  $\mathbf{a}^{(1)}$ ,  $[\text{<swap>, pos}_1]$ , which instructs the interpreter to swap number 1 and 2, and action  $\mathbf{a}^{(2)}$ ,  $[\text{<swap>, pos}_3]$ , which instructs the interpreter to swap number 3 and 4, imitating the bubble sort algorithm; finally, the interpreter halts inference and outputs the completely edited sequence  $\mathbf{y}^{(e)}$  after receiving the termination action  $\mathbf{a}^{(3)} = [\text{<done>, <done>}]$ .

hidden states iteratively and free the interpreter from the decoder to complete text editing in several program-interpret iterations (recurrent inference). NPI belongs to neural program induction (Devlin et al., 2017), but Recurrence is part of neural program synthesis (Ellis et al., 2019). Consequently, Recurrence always follows the latest hidden representation of its input text rather than a static context matrix and only needs to decode an editing action of a fixed length in each iteration.

**Multi-Step Learning** is a manner to solve a problem in several steps. Recent work in text editing prefers multi-step learning, especially for long-to-short (Narayan and Gardent, 2014; Zhang and Lapata, 2017), and mixed editing (Ge et al., 2018; Lichtarge et al., 2018). For example, Tagging can also be regarded as a two-steps learning. However, these studies usually edit texts incrementally through a multi-round seq2seq inference. To the best of our knowledge, our Recurrence is the first inference method that divides a text editing task into multiple independent sub-tasks and completes them recurrently.

### 3 Recurrent Inference

#### 3.1 Method Overview

Recurrence breaks the text editing task down into iterations of editing actions and each editing action is determined on the hidden representation of the partially edited sequence. Conceptually, it is performing a predefined underlying iterative algorithm that is designed to achieve some text editing goals. There are two components in Recurrence: programmer and interpreter. Given a source sequence  $\mathbf{x} = x_1, \dots, x_{|\mathbf{x}|}$ , the programmer determines a single editing action,  $\mathbf{a}^{(1)}$ , to be applied on

$\mathbf{x}$ . Then, the interpreter executes the action  $\mathbf{a}^{(1)}$  on  $\mathbf{x}$  and produces the partially edited sequence with one edit,  $\mathbf{y}^{(1)}$ . Then,  $\mathbf{y}^{(1)}$  is fed to the programmer to determine the next action  $\mathbf{a}^{(2)}$ . This process continues until the programmer determines the text is fully edited and outputs a termination action to stop further editing. The inference also ends if the number of iterations reaches a predefined limit. Finally, the interpreter outputs the completely edited sequence  $\mathbf{y}^{(\text{complete})}$ . This recurrent editing process is illustrated with an example in the number ordering task in Figure 2.

The hypothesis is that it is easier to let a model learn a single editing step than the whole mapping between original and edited sequences. Also, being able to observe the latest text status leads to a more accurate input representation. Furthermore, Recurrence is explainable in the sense that not only we can understand the intention of each editing step done by the model, but we can also actively participate in designing the editing procedure.

#### 3.2 Programmer

Broadly speaking, the programmer determines the action for a given input in accordance with the underlying algorithm that the programmer is trained to mimic. In the programmer, the encoder extracts relevant information from an input text sequence  $\mathbf{x}$  and then the decoder decides a single step of action that should be applied to  $\mathbf{x}$ . The programmer can be any model that is able to produce editing actions based on textual information. In our experiments, the programmer is a seq2seq model with an encoder-decoder architecture.

	AOR ( $N = 10, L = 6$ )	AES ( $N = 10, L = 4$ )	AEC ( $N = 10, L = 5$ )
Input	6 10 9 5 2 3	$-3 + 10 / 2 == 2$	$4 - 3 / 6 * 4 == 2$
Output	$-6 / 10 + 9 / 5 * 2 == 3$	$-3 + 10 / 2 == 2$	$4 - 3 / 6 * 4 == 2$
Input	2 2 4 8 2 4	$(-2 + 4) / 7 * 7 == 2$	$6 7 * + / 7 + / 7 == 2$
Output	$2 * 2 - 4 + 8 / 2 == 4$	$2 / 7 * 7 == 2$	$-7 * 5 / 7 + 7 == 2$
Input	2 2 4 8 2 4	$2 / 7 * (11 - 4) == (4 - 2)$	$-6 5 + 11 - 2$
Output	$-2 + 2 / 4 * 8 + 2 == 4$	$2 / 7 * 7 == 2$	$-6 + 5 + 11 - 8 == 2$

Table 1: Examples from AOR, AES, and AEC with  $N$  and  $L$ .

### 3.3 Editing Actions

An editing action contains (i) the type of editing operation, (ii) the position the editing occurs, and (iii) a text symbol.

Formally, the set of editing actions is defined by  $\mathcal{A} := \{\mathbf{a} = (e, p, s) | \forall e \in \mathcal{E}, p \in \mathcal{P}, s \in \mathcal{S}\}$ , where  $\mathcal{E}$  is the set of all operations,  $\mathcal{P}$  is the set of all positions, and  $\mathcal{S}$  is the set of symbols. The definition of  $\mathcal{E}$ ,  $\mathcal{P}$  and  $\mathcal{S}$  is determined by the specific text editing task and the underlying text editing algorithm. For example, each  $p \in \mathcal{P}$  would contain a single position or multiple positions (i.e., a tuple of position indices) depending on the operation. Also, if an editing task contains only a single type of operation, then the operation can be omitted. Some operations, such as deletion, do not need a symbol input, so, the symbol component can also be omitted. It is required that  $\text{DONE} \in \mathcal{E}, \mathcal{P}, \mathcal{S}$  to indicate termination.

Editing actions allow the design of the editing order. Given  $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(n)}$ , the position sequence  $p^{(1)} \in \mathbf{a}^{(1)}, \dots, p^{(n)} \in \mathbf{a}^{(n)}$  determines the editing order. This could be beneficial since empirical results have shown that ordering matters for text generation (Ford et al., 2018). For the sake of simplicity, in our experiments, we choose to arrange positions across actions in an increasing order, editing a sequence from left to right.

Due to liberty given by the definition of the action, we believe Recurrence can be applied to a much border field of applications. In the scope of this paper, we only concern about text editing.

### 3.4 Interpreter

The interpreter is a parameter-free function that executes the editing action produced by the programmer. Specifically, the interpreter first checks if the action is the termination action. If so, the interpreter will halt inference and directly output its input sequence as the completely edited text,  $\mathbf{y}^{(\text{complete})}$ . Otherwise, the interpreter carries out

the received action to its input sequence and produces a partially edited sequence. Then, the Recurrence continues by feeding the partially edited sequence into the programmer to determine the next editing action.

It is possible for the programmer to output illegal actions that do not follow the predefined action template (e.g., actions with missing a position component), especially when the programmer is not fully trained. Therefore, the interpreter checks if an action is valid and skips invalid actions by returning the input sequence.

### 3.5 Offline Training

Training text editing models requires pairs of source sequence  $\mathbf{x}$  and target sequence  $\mathbf{y}$ , but different inference methods employ different generation algorithms to produce appropriate target sequences to form suitable training pairs. For the conventional inference methods, End2end map unedited text sequences to target text sequences directly, and Tagging map unedited text sequences to target tag sequences before realizing the target text sequences. Hence, for the training data, the source sequences are the original, while the target sequences are edited text sequences for End2end and editing operation sequences for Tagging. In our experiments, we name the training modes used by the conventional methods *offline training*.

### 3.6 Online Training

To train the programmer, we compute all intermediate actions  $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(n)}$  that are required to edit input  $\mathbf{x}$  to target  $\mathbf{y}^{(\text{complete})}$ . Applying these editing actions, we obtain the partially edited sequences  $\mathbf{y}^{(1)} = \mathbf{x}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(n)} = \mathbf{y}^{(\text{complete})}$ . After that, the training list of pairs for the programmer is  $(\mathbf{y}^{(1)}, \mathbf{a}^{(1)}), (\mathbf{y}^{(2)}, \mathbf{a}^{(2)}), \dots, (\mathbf{y}^{(n)}, \mathbf{a}^{(n)})$ , where  $\mathbf{a}^{(n)}$  is the termination action. We uniformly sample one source-target pair from this list as the training data instance. Due to the fact that selected training pairs for each source sequence  $\mathbf{x}$  varies

during training, we name this training mode *online training*. For the thoroughness of experiments, we examine three inference methods with both training modes. In the training phase, intermediate training instances are exposed to `End2endOnline` and `TaggingOnline`. Only the immediate editing action  $(\mathbf{y}^{(1)}, \mathbf{a}^{(1)})$  are fed to `RecurrenceOffline`.

## 4 Tasks

An increasing number of studies takes synthetic benchmark tasks to examine ideas before extending to open-domain natural language data (Zaremba and Sutskever, 2014; Lake and Baroni, 2018; Nangia and Bowman, 2018; Lample and Charton, 2020). Following the fruitful results of previous work, we aim to evaluate three inference methods in the domain of arithmetic problems (Hosseini et al., 2014; Roy and Roth, 2015; Ling et al., 2017) that can be treated as the test-beds for text editing. We introduce three tasks, namely AOR, AES, and AEC, corresponding to the three types of text editing tasks: short-to-long, long-to-short, and mixed. Being able to control the aspects of the datasets allows us to compare the characteristics of the three inference methods more thoroughly and analyze the appropriate situations to apply each method.

### 4.1 Arithmetic Equation

Our arithmetic equation consists of integer numbers  $\mathcal{N} \in \mathbb{Z}^{\geq 2}$ , an equal sign (“==”), and operators<sup>1</sup>  $\mathcal{O} = \{“+”, “-”, “*”, “/”\}$ . For convenience, we restrict the right-hand side of the equation to a number. The equation holds if the value of the left-hand side equals the number on the right-hand side. Operators  $\mathcal{O}$  are placed between two numbers, where the subtraction operator “-” can also be put to the left of any single number. We consider equations as sequences of mathematical symbols (Saxton et al., 2019) instead of tree structures (Lample and Charton, 2020). We describe an arithmetic equation dataset from three aspects: (1)  $N = |\mathcal{N}|$  defines the number of unique integers; (2)  $L \in \mathbb{Z}^*$  defines the number of integers in an equation; (3)  $D \in \mathbb{Z}^*$  defines the number of unique equations.

Note that since we only consider binary operations, the sequence length of a valid arithmetic expression is always  $2L$  or  $2L - 1$ , depending on if there is a subtraction operator before the first num-

<sup>1</sup>We use these symbols to apply the Python built-in function `eval()`.

ber. Intuitively, it is reasonable to assume that the greater  $N$  and  $L$  become, the harder the task gets. Whereas, the larger  $D$ , the easier the task becomes.

### 4.2 Arithmetic Operators Restoration

The goal of AOR is to convert a sequence of integer numbers into a valid arithmetic equation. For a given source sequence of integer numbers,  $\mathbf{x} \in \mathcal{N}^L$ , a model for AOR inserts appropriate operators from  $\mathcal{O}$  in between the first  $L - 1$  integers in  $\mathbf{x}$  and inserts an equal sign (“==”) before the  $L^{\text{th}}$  element in  $\mathbf{x}$  so that the resulting arithmetic expression sequence (target sequence) is valid. Each integer sequence potentially corresponds to different valid arithmetic equations. Thus, AOR is *one-to-many* learning. To obtain integer sequences for AOR, we first generate valid arithmetic equations and then remove all the operators and equal signs (see Table 1).

### 4.3 Arithmetic Equation Simplification

Here, we involve two more mathematical symbols (“(”, “)”). In an equation, parentheses help to group parts of an expression and indicate the order of precedence. In this task, we aim to simplify equations by calculating the parts in parentheses and removing parentheses from equations. Equation that has no parentheses is already in the simplest form, so there is no need to change. We generate complicated versions of a simplified equation by randomly replacing some integers (including the one on the right-hand side) with their equivalent bracketed expressions. Since these variants share the same simplified form, AES is *many-to-one* learning (see Table 1).

### 4.4 Arithmetic Equation Correction

AEC is a more comprehensive text editing task in that a model needs to detect and correct possible mistakes. To generate mistakes, we inverse a valid equation by deleting, substituting, or inserting random tokens at random positions. We do not touch the right-hand side integer to guarantee that the corrected left-hand side (include “==”) equals the same value to assert equality. We fix the maximum number of errors to three, regardless the values of  $N$ ,  $L$ , and  $D$ . No change is made if there is no error. We generate many wrong equations based on one correct equation. Meanwhile, a wrong equation can be modified into multiple correct equations. Hence, AEC is *many-to-many* learning (see Table 1).

	AOR ( $N = 10, L = 5$ )	AES ( $N = 100, L = 5$ )	AEC ( $N = 10, L = 5$ )
Source	8 2 8 4 2	$- 33 + 25 + 75 - 60 == ( 30 - 23 )$	$7 * 8 / 4 8 2 - == 6$
Target <sub>End2end</sub>	$- 8 * 2 / 8 + 4 == 2$	$- 33 + 25 + 75 - 60 == 7$	$7 * 8 / 4 - 8 == 6$
Target <sub>Tagging</sub>	$\langle \text{insert}_- \rangle \langle \text{keep} \rangle \langle \text{insert}_* \rangle \langle \text{keep} \rangle$ $\langle \text{insert}_/ \rangle \langle \text{keep} \rangle \langle \text{insert}_+ \rangle \langle \text{keep} \rangle$ $\langle \text{insert}_= \rangle \langle \text{keep} \rangle$	$\langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle$ $\langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{sub}_7 \rangle$ $\langle \text{delete} \rangle \langle \text{delete} \rangle \langle \text{delete} \rangle \langle \text{delete} \rangle$	$\langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle \langle \text{keep} \rangle$ $\langle \text{keep} \rangle \langle \text{delete} \rangle \langle \text{sub}_- \rangle \langle \text{sub}_8 \rangle$ $\langle \text{keep} \rangle \langle \text{keep} \rangle$
Target <sub>Recurrence</sub>	$\langle \text{pos}_0 \rangle -$	$\langle \text{pos}_9 \rangle \langle \text{pos}_{13} \rangle 7$	$\langle \text{delete} \rangle \langle \text{pos}_5 \rangle \langle \text{pos}_5 \rangle$

Table 2: Example target sequences given the same source sequence in AOR, AES, and AEC.

## 5 Experiments

We test Recurrence in comparison with End2end and Tagging across AOR, AES, and AEC. We describe the results conditioned on specific  $N$ ,  $L$ , and  $D$ . Later, we analyze the impact of each of them in Section 6.

**Data.** In all tasks, the dataset is divided into three subsets: 70% for training, 15% for validation, and 15% for testing. For AES (many-to-one learning) and AEC (many-to-many learning), we feed the training set to a data generator in every epoch to expose all the variants of targets as input sequences (see Section 4). For the sake of fairness, we examine three methods in both online and offline training modes. To train End2end<sub>online</sub> and Tagging<sub>online</sub>, in each epoch, we keep the targets, but uniformly pick a partially edited  $y^{(i)}$  to alternate the original input  $x$  as the source sequence. The target equations can be used to train End2end directly. By contrast, further pre-processing is necessary for Tagging and Recurrence. Training targets for Tagging are tag sequences, while those for Recurrence are editing actions.

**Models.** After testing Transformer (Vaswani et al., 2017) and a range of modern RNNs (Mikolov et al., 2010; Sutskever et al., 2014; LeCun et al., 2015), we focus on the overall best-performed architecture — bidirectional LSTM (Schuster and Paliwal, 1997; Hochreiter and Schmidhuber, 1997) with an attention mechanism (Luong et al., 2015). Throughout all the experiments, three inference methods share the same model structure with  $d_{\text{model}} = 512$ ,  $d_{\text{embedding}} = 512$ ,  $n_{\text{layers}} = 1$ ,  $r_{\text{learning}} = 10^{-5}$ ,  $r_{\text{teacher forcing}} = 0.5$ , and  $r_{\text{dropout}} = 0.5$  (Srivastava et al., 2014). Parameters are uniformly initialized from  $[-\sqrt{\frac{1}{d}}, \sqrt{\frac{1}{d}}]$ . To prevent uncontrolled interference, we train all models from scratch instead of pre-training. We use Adam optimizer (Kingma and Ba, 2015) with an L2 gradient clipping of 5.0 (Pascanu et al., 2013).

**Evaluation.** We evaluate methods by three metrics: *token accuracy*, *sequence accuracy*, and *equation*

*accuracy*. Token accuracy marks the correct predictions at the token-level divided by the target sequence length and then averaged by the test size. Sequence accuracy stands for the correct predictions at the sequence-level divided by the test size. Equation accuracy is the number of true predicted equations divided by the test size; it emphasizes on whether an equation holds rather than whether an equation is the same as the target. We evaluate the performance via equation accuracy for AOR (one-to-many), sequence accuracy for AES (many-to-one), and both equation accuracy and sequence accuracy for AEC (many-to-many). Sequence accuracy is accompanied by token accuracy for additional reference.

**Training.** We train on a single GeForce RTX Titan with a batch size of 256. The last batch is dropped if it does not contain 256 samples. To ensure convergence, we adopt early stopping (Prechelt, 1998) with a patience of 512 epochs.

### 5.1 Arithmetic Operators Restoration

**Data.** Experiments are performed on a dataset with  $N = 10$ ,  $L = 5$ , and  $D = 10K$ . For Tagging, the tags are KEEP and INSERT<sub>TOKEN<sub>AOR</sub></sub>, where  $\text{TOKEN}_{\text{AOR}} = \mathcal{O} \cup \{“ == ”\}$ .

For Recurrence, the set of editing actions is defined as  $\mathcal{A}_{\text{AOR}} := \{\mathbf{a} = (e, p, s) \mid \forall e \in \mathcal{E}, p \in \mathcal{P}, s \in \mathcal{S}\}$ , where  $\mathcal{E}$  is an empty set since there is only one operation, insertion, and thus omitted;  $\mathcal{P} := \{p \mid p \in \{0, \dots, |x|\}\}$ ; and  $\mathcal{S} = \text{TOKEN}_{\text{AOR}}$ . For a given action  $\mathbf{a} = (p, s)$ , the interpreter inserts  $s$  before  $x_p$  (see Table 2).

**Results.** As shown in Table 3, Recurrence<sub>Online</sub> outperforms End2end<sub>Online</sub> by 29.20% and Tagging<sub>Online</sub> by 7.13%, achieving an equation accuracy of 58.53%. Hence, Recurrence<sub>Online</sub> has the best performance. Note that online training is critical for Recurrence to achieve good performance as Recurrence<sub>Online</sub> outperforms Recurrence<sub>Offline</sub> by 27.40%, whilst online training only helps to improve the performance of Tagging by 0.87% and End2end by 2.86%.

		AOR ( $N = 10, L = 5, D = 10K$ )		AES ( $N = 100, L = 5, D = 10K$ )			AEC ( $N = 10, L = 5, D = 10K$ )			
Method	Training	#Epoch	Equ Acc.%	#Epoch	Token Acc.%	Seq Acc.%	#Epoch	Token Acc.%	Seq Acc.%	Equ Acc.%
End2end	Offline	3352	26.47	5063	75.49	3.27	72144	87.78	54.67	55.13
	Online	2640	29.33	7795	84.60	25.20	112482	<b>88.08*</b>	57.27	57.73
Tagging	Offline	1149	50.53	5223	90.10	43.80	135729	82.29	44.20	44.40
	Online	2245	51.40	4520	87.00	36.67	112968	84.46	46.93	47.33
Recurrence	Offline	1281	31.13	7603	94.92	62.07	203067	81.85	55.87	56.20
	Online	1898	<b>58.53*</b>	7088	<b>98.63*</b>	<b>87.73*</b>	152982	83.64	<b>57.47*</b>	<b>58.27*</b>

Table 3: Evaluation results of three inference methods on AOR, AES, and AEC with specific  $N$ ,  $L$ , and  $D$ .

## 5.2 Arithmetic Equation Simplification




Figure 3: Testing sequence accuracy per epoch in AES with  $N = 10$ ,  $L = 5$ , and  $D = 10K$ ; all methods achieved near-perfect performances, but Recurrence<sub>Online</sub> converges the fastest.

**Data.** We first experiment with  $N = 10$ ,  $L = 5$ , and  $D = 10K$ , but all methods can reach a near-perfect sequence accuracy (see Figure 3). Therefore, we adjust  $N$  from 10 to 100 to make the task more challenging. A target sequence to train Tagging is a sequence of tags consisting of KEEP, DELETE, and SUBSTITUTE<sub>TOKEN<sub>AES</sub></sub>, where  $TOKEN_{AES} \in \mathcal{N}$ . For Recurrence, target editing actions are  $\mathcal{A}_{AES} := \{\mathbf{a} = (e, p, s) \mid \forall e \in \mathcal{E}, p \in \mathcal{P}, s \in \mathcal{S}\}$ , where the default operation is substitution, so  $\mathcal{E}$  is an empty set and omitted;  $\mathcal{P} := \{\mathbf{p} = [p_1, p_2] \mid p_i \in \{0, \dots, |\mathbf{x}|\}, \forall i = 1, 2\}$ ;  $\mathcal{S} = TOKEN_{AES}$ . This editing action instructs the interpreter to replace the part between  $x_{p_1}$  and  $x_{p_2}$  with  $TOKEN_{AES}$  (see Table 2).

**Results.** Our Recurrence<sub>Online</sub> obtains a sequence accuracy of 87.73%, outperforming End2end<sub>Online</sub> by 62.53% and Tagging<sub>Offline</sub> by 43.93%. We also find that the performance of Recurrence is impaired significantly without online training. Besides, online training saves 515 epochs and achieves a better

outcome. Both facts demonstrate the necessity of intermediate steps for training Recurrence.

## 5.3 Arithmetic Equation Correction

**Data.** We use a dataset with  $N = 10$ ,  $L = 5$ , and  $D = 10K$ . A tag sequence is made of tags including KEEP, DELETE, SUBSTITUTE<sub>TOKEN<sub>AEC</sub></sub>, and INSERT<sub>TOKEN<sub>AEC</sub></sub>, where  $TOKEN_{AEC} \in \mathcal{N} \cup \mathcal{O} \cup \{“ == ”\}$ . For Recurrence, we define target editing actions as  $\mathcal{A}_{AEC} := \{\mathbf{a} = (e, p, s) \mid \forall e \in \mathcal{E}, p \in \mathcal{P}, s \in \mathcal{S}\}$ , where  $\mathcal{E} := \{e \mid e \in \{DELETE, SUBSTITUTE, INSERT\}\}$ ;  $\mathcal{P} := \{p \mid p \in \{0, \dots, |\mathbf{x}|\}\}$ ;  $\mathcal{S} := TOKEN_{AEC}$ . To freeze the sequence length of  $\mathbf{a}$ , we repeat  $p$  at  $\mathbf{a}_3$  to replace  $s$  when  $e = DELETE$ . During interpreting,  $e = DELETE$  directs to remove  $x_p$ ;  $e = SUBSTITUTE$  guides to replace  $x_p$  with  $s$ ;  $e = INSERT$  means to insert  $s$  before  $x_p$  (see Table 2).

**Results.** Recurrence<sub>Online</sub> attains higher scores over the other two methods, resulting in a sequence accuracy of 57.47% and an equation accuracy of 58.27%. The performance edge of Recurrence is not obvious due to the task setting. In section 6, we adjust the task to distinguish the performance of each method more easily. When applying online training, we observe improvements in all three methods. Particularly, Recurrence<sub>Online</sub> takes around 50K epochs less than Recurrence<sub>Offline</sub> and attains a better performance.

## 6 Analysis

As shown in section 5, Recurrence outperforms End2end and Tagging in all three tasks in our experiment settings. In this section, we explore the limits of Recurrence by running experiments with varying values of  $N$ ,  $L$  and  $D$ , so as to determine in what scenario Recurrence performs well (see Figure 4).

**The Impact of  $N$ .** We conduct experiments with  $L = 5$ ,  $D = 50K$ , and  $N$  increasing from 10 to 50 with an interval of 10 for AOR;  $L = 5$ ,  $D = 10K$ , and  $N$  increasing from 100 to 300 with

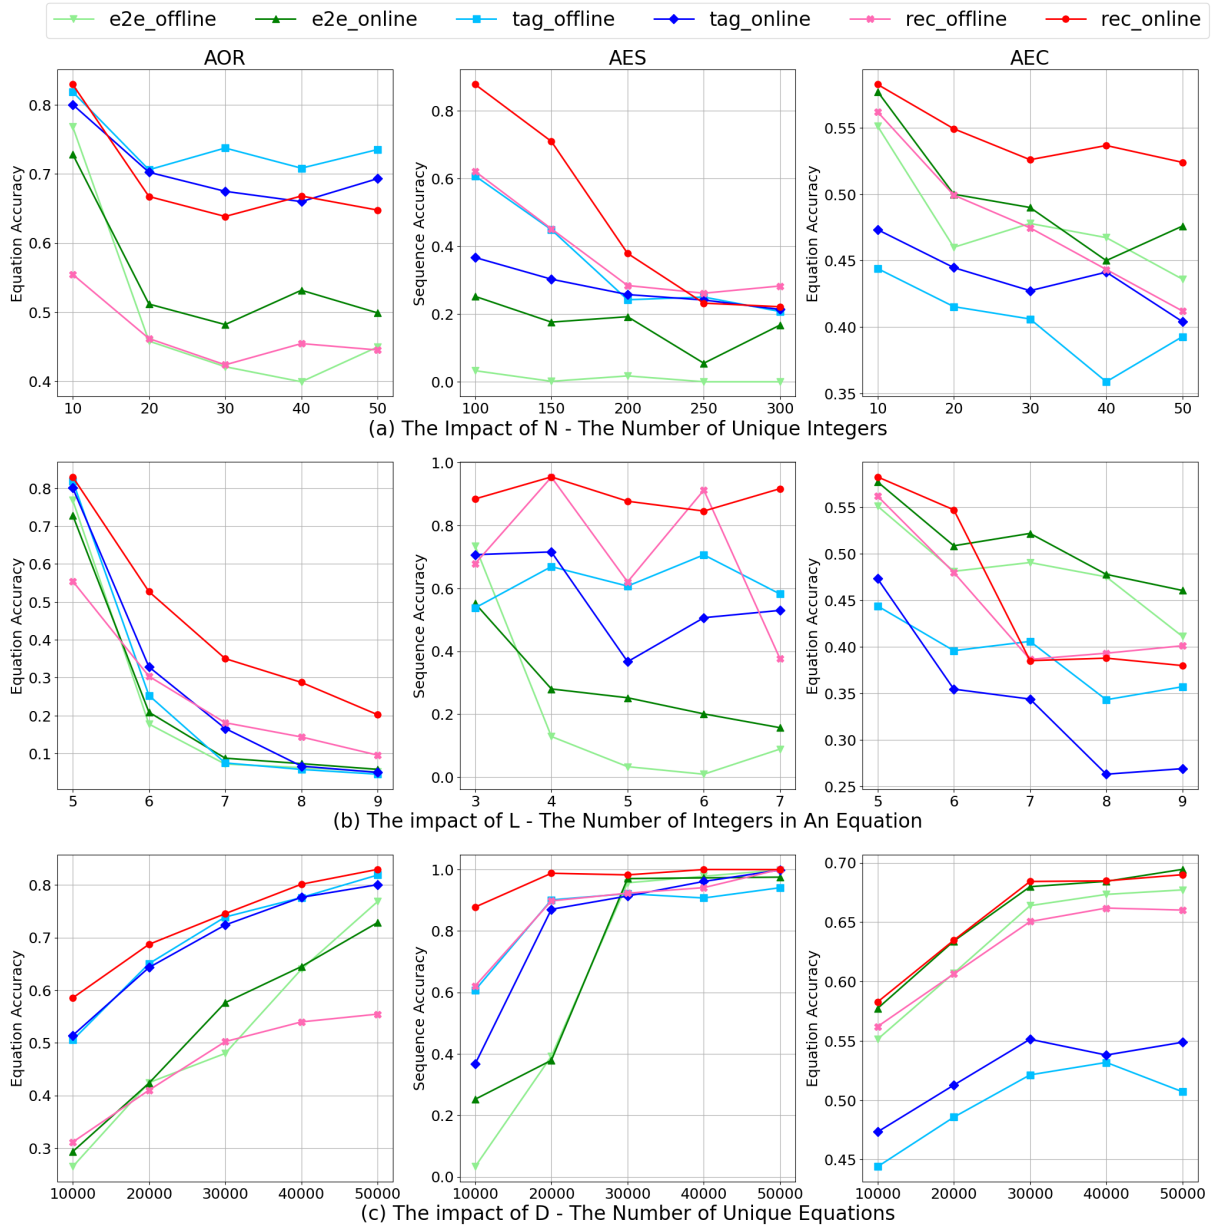


Figure 4: Evaluation results of three inference methods on AOR, AES, and AEC under the control of (a)  $N$ , (b)  $L$ , and (c)  $D$ , respectively.

an interval of 50 for AES; and  $L = 5$ ,  $D = 10K$ , and  $N$  increasing from 10 to 50 with an interval of 10 for AEC. For AOR,  $\text{Recurrence}_{\text{Online}}$  and  $\text{Tagging}$  show similar resilience, however,  $\text{Tagging}_{\text{Offline}}$  performs better when  $N \geq 20$ . For AES,  $\text{Recurrence}_{\text{Online}}$  performs much better than  $\text{Tagging}$  and  $\text{End2end}$  (by at least 20%) when  $N \leq 150$ . Note that  $\text{End2end}_{\text{Offline}}$  learns hardly anything. We also observe that  $\text{End2end}_{\text{Offline}}$  can achieve a near-perfect performance when  $N = 10$ . These results indicate that the  $\text{End2end}_{\text{Offline}}$ 's performance declines rapidly as  $N$  increases and re-

quires a much larger  $D$ -to- $N$  ratio to perform well. Finally, for AEC,  $\text{Recurrence}_{\text{Online}}$  displays the most resilience and performs the best.

**The Impact of  $L$ .** We conduct experiments with  $N = 10$ ,  $D = 50K$ , and  $L$  increasing from 5 to 9 with an interval of 1 for AOR;  $N = 10$ ,  $D = 50K$ , and  $L$  increasing from 3 to 7 with an interval of 1 for AES; and settings identical to AOR for AEC. For AOR and AES, both  $\text{Recurrence}_{\text{Online}}$  and  $\text{Tagging}$  show similar trend, however,  $\text{Recurrence}_{\text{Online}}$  performs the best. For AEC, while  $\text{Recurrence}_{\text{Online}}$  still outperforms  $\text{Tagging}$ ,  $\text{End2end}_{\text{Offline}}$  performs the best for  $L \geq$



7 and shows more resilience. We think when  $N = 10$ , the AEC task is too easy for End2end with 50K training data. Thus, we increase  $N$  from 10 to 100 and find that End2end cannot gain any performance within 512 epochs (i.e., accuracy is 0%). We want to stress that when the amount of data cannot counter the increase of  $L$ , which is the case for AOR and AES, End2end’s performance declines faster than Recurrence and Tagging.

**The Impact of  $D$ .** We conduct experiments with  $N = 10$ ,  $L = 5$ , and  $D$  increasing from 10K to 50K with an interval of 10K for AOR;  $N = 100$ ,  $L = 5$ , and  $D$  increasing from 10K to 50K with an interval of 10K for AES; and settings identical to AOR for AEC. All models benefit from the increasing of  $D$  as expected. However, it is clear that Recurrence<sub>Online</sub> is the best performing model when  $D$  is small. The only exception is that for AEC, End2end has similar performance trend as Recurrence. As discussed before, this is likely because End2end performs well with small  $N$ .

**The Impact of Online Training.** When comparing the performance between online and offline training, the online training, as expected, generally has better performances than offline training for End2end and Tagging with only a few exceptions. Note that online training is not part of the standard training procedure for End2end and Tagging, however, we use online training with End2end and Tagging for the sake of a fair comparison. Therefore, for End2end and Tagging, the online training acts like a data augmentation technique, providing more data points for training. Surprisingly, offline training also allows Recurrence to gain some editing ability, at times better than End2end and Tagging. We believe for text editing tasks with very localized editing actions, such as AES, showing the immediate editing actions are enough for the model to generalize proper editing actions. In other words, when the editing actions are less sequentially dependent, even offline training enables Recurrence to achieve performance better than End2end and Tagging. This supports our intuition that letting the programmer produce one single editing step reduces the learning difficulty.

**The Impact of Ordering.** In early experiments, We find that the programmer cannot converge if the data guide it to edit a sequence in a random order (a mixture of both left-to-right and right-to-left). Hence, we think ordering matters for not only text generation (Ford et al., 2018) but also Recurrence

in text editing. One of our assumptions is that random ordering may assign various actions to the same text state, and thus causes confusion in the list of actions used to edit the input text  $\mathbf{x}$  to the output text  $\mathbf{y}$ . When there are conflicting sample pairs in the training data set, the model cannot easily converge. We leave this problem for future work.

To summarize our findings, under settings with moderate or large  $N$  and  $L$ , End2end performs much worse than Tagging and Recurrence with limited data. Tagging performs slightly better than Recurrence when  $N$  gets larger with fixed  $D$  and  $L$  in AOR (short-to-long). However, Tagging performs worse than Recurrence in all other cases. Therefore, we conclude that Recurrence is more data-efficient and overall better performs than End2end and Tagging in most situations, especially in AES (long-to-short).

## 7 Conclusions and Future Work

We propose a recurrent inference method, Recurrence, that edits a given text sequence iteratively such that in each iteration the programmer determines a single step of editing action and the interpreter executes the action. Our method outperforms the other two inference methods, End2end and Tagging, in three arithmetic equation editing tasks we introduced. For future work, we plan to apply Recurrence to open-domain natural language data and investigate on how to relax its need for intermediate editing steps as extra supervision signals. We also wish to experiment with applying pointer attention (Vinyals et al., 2015) to replace the position component in actions.

## Acknowledgments

We give thanks to Keran Zhao, Yun He, Che Wang (Watcher) for sharing their pearls of wisdom. We also would like to thank EMNLP anonymous reviewers for their helpful insights, comments, and corrections. This research was supported by the Shining Lab.

## References

Fernando Alva-Manchego, Joachim Bingel, Gustavo Paetzold, Carolina Scarton, and Lucia Specia. 2017. Learning how to simplify from explicit labeling of complex-simplified text pairs. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 295–305, Taipei, Taiwan. Asian Federation of Natural Language Processing.

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Alexandre Bérard, Laurent Besacier, and Olivier Pietquin. 2017. LIG-CRISAL submission for the WMT 2017 automatic post-editing task. In *Proceedings of the Second Conference on Machine Translation*, pages 623–629, Copenhagen, Denmark. Association for Computational Linguistics.
- Xiaoyin Che, Cheng Wang, Haojin Yang, and Christoph Meinel. 2016. Punctuation prediction for unsegmented transcript based on word vector. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 654–658.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.
- Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. 2017. Robustfill: Neural program learning under noisy i/o. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 990–998. JMLR.org.
- Yue Dong, Zichao Li, Mehdi Rezagholizadeh, and Jackie Chi Kit Cheung. 2019. EditNTS: An neural programmer-interpreter model for sentence simplification through explicit editing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3393–3402, Florence, Italy. Association for Computational Linguistics.
- Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. 2019. Write, execute, assess: Program synthesis with a repl. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9169–9178. Curran Associates, Inc.
- Katja Filippova, Enrique Alfonseca, Carlos A Colmenares, Łukasz Kaiser, and Oriol Vinyals. 2015. Sentence compression by deletion with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 360–368.
- Nicolas Ford, Daniel Duckworth, Mohammad Norouzi, and George E. Dahl. 2018. The importance of generation order in language modeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2942–2946. Association for Computational Linguistics.
- Tao Ge, Furu Wei, and Ming Zhou. 2018. Fluency boost learning and inference for neural grammatical error correction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1055–1065, Melbourne, Australia. Association for Computational Linguistics.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. Levenshtein transformer. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 11181–11191. Curran Associates, Inc.
- Sepp Hochreiter. 1998. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. Learning to solve arithmetic word problems with verb categorization. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–533, Doha, Qatar. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Shubha Guha, and Kenneth Heafield. 2018. Approaching neural grammatical error correction as a low-resource machine translation task. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 595–606, New Orleans, Louisiana. Association for Computational Linguistics.
- Seokhwan Kim. 2019. Deep recurrent neural networks with layer-wise multi-head attentions for punctuation restoration. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7280–7284. IEEE.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *35th International Conference on Machine Learning, ICML 2018, 35th International Conference on Machine Learning, ICML 2018*, pages 4487–4499. International Machine Learning Society (IMLS). 35th International Conference on Machine Learning, ICML 2018 ; Conference date: 10-07-2018 Through 15-07-2018.

- Guillaume Lample and François Charton. 2020. [Deep learning for symbolic mathematics](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature*, 521(7553):436–444.
- Jindřich Libovický, Jindřich Helcl, Marek Tlustý, Ondřej Bojar, and Pavel Pecina. 2016. CUNI system for WMT16 automatic post-editing and multimodal translation tasks. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 646–654, Berlin, Germany. Association for Computational Linguistics.
- Jared Lichtarge, Christopher Alberti, Shankar Kumar, Noam Shazeer, and Niki Parmar. 2018. Weakly supervised grammatical error correction using iterative decoding. *CoRR*, abs/1811.01710.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167, Vancouver, Canada. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. Encode, tag, realize: High-precision text editing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5054–5065, Hong Kong, China. Association for Computational Linguistics.
- Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*.
- Nikita Nangia and Samuel Bowman. 2018. [ListOps: A diagnostic dataset for latent tree learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 92–99, New Orleans, Louisiana, USA. Association for Computational Linguistics.
- Shashi Narayan and Claire Gardent. 2014. Hybrid simplification using deep semantics and machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 435–445, Baltimore, Maryland. Association for Computational Linguistics.
- Sergiu Nisioi, Sanja Štajner, Simone Paolo Ponzetto, and Liviu P Dinu. 2017. Exploring neural text simplification models. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 85–91.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, page III–1310–III–1318. JMLR.org.
- Lutz Prechelt. 1998. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer.
- Scott E. Reed and Nando de Freitas. 2016. [Neural programmer-interpreters](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Subhro Roy and Dan Roth. 2015. Solving general arithmetic word problems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal. Association for Computational Linguistics.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. [Analysing mathematical reasoning abilities of neural models](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mike Schuster and Kuldeep K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D.

- Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. Abstractive document summarization with a graph-based attentional neural model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1171–1181.
- Ottokar Tilk and Tanel Alumäe. 2016. Bidirectional recurrent neural network with attention mechanism for punctuation restoration. In *Interspeech*, pages 3047–3051.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.
- Thuy-Trang Vu and Gholamreza Haffari. 2018. Automatic post-editing of machine translation: A neural programmer-interpreter approach. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium. Association for Computational Linguistics.
- Liangguo Wang, Jing Jiang, Hai Leong Chieu, Chen Hui Ong, Dandan Song, and Lejian Liao. 2017. Can syntax help? improving an LSTM-based sentence compression model for new domains. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vancouver, Canada. Association for Computational Linguistics.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2018. Learning neural templates for text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3174–3187, Brussels, Belgium. Association for Computational Linguistics.
- Wojciech Zaremba and Ilya Sutskever. 2014. [Learning to execute](#). *CoRR*, abs/1410.4615.
- 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*, pages 584–594, Copenhagen, Denmark. Association for Computational Linguistics.
- Wei Zhao, Liang Wang, Kewei Shen, Ruoyu Jia, and Jingming Liu. 2019. Improving grammatical error correction via pre-training a copy-augmented architecture with unlabeled data. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 156–165, Minneapolis, Minnesota. Association for Computational Linguistics.