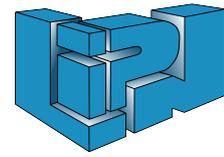




UNIVERSITÉ  
**PARIS  
DESCARTES**



PARIS DESCARTES UNIVERSITY  
DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE

MASTER'S THESIS

---

## Attention-based Encoder-Decoder Networks for Spelling and Grammatical Error Correction

---

*Author:*

Sina AHMADI

*Supervisors:*

Joseph LE ROUX

Nadi TOMEH

*A thesis submitted in partial fulfillment of the requirements for the degree of  
Master of Science in Artificial Intelligence*

September, 2017

*Dedicated to Dad and Mom,  
and my love, Ioanna*

---

## Acknowledgements

---

I would first like to thank my thesis supervisors Dr. Joseph Le Roux and Dr. Nadi Tomeh of RCLN team at the Laboratoire d'Informatique de Paris Nord at the University of Paris 13. During the six months of my internship, they were always around whenever I ran into a trouble spot or had a question about my research or my codes. They consistently allowed this thesis to be my own work, but steered me in the right direction.

I would also like to acknowledge the head of the Machine Learning for Data Science (MLDS) team at Paris Descartes University, Dr. Mohamed Nadif. Furthermore, I am gratefully indebted to Dr. Kyumars S. Esmaili and Jalal Sajadi for their very valuable presence and friendship during my studies.

Finally, I must express my very profound gratitude to my parents, my wife Ioanna and our whole family, for their unconditional love, their forbearance and for providing me with unfailing support and continuous encouragement throughout my years of study and through the process of researching and writing this thesis. This accomplishment would not have been possible without them.

---

## Abstract

---

Automatic spelling and grammatical correction systems are one of the most widely used tools within natural language applications. In this thesis, we assume the task of error correction as a type of monolingual machine translation where the source sentence is potentially erroneous and the target sentence should be the corrected form of the input. Our main focus in this project is building neural network models for the task of error correction. In particular, we investigate sequence-to-sequence and attention-based models which have recently shown a higher performance than the state-of-the-art of many language processing problems. We demonstrate that neural machine translation models can be successfully applied to the task of error correction.

While the experiments of this research are performed on an Arabic corpus, our methods in this thesis can be easily applied to any language.

*Keywords*— natural language error correction, recurrent neural networks, encoder-decoder models, attention mechanism

---

## Contents

---

<b>Acknowledgements</b>	<b>ii</b>
<b>Abstract</b>	<b>iii</b>
<b>List of Figures</b>	<b>viii</b>
<b>List of Tables</b>	<b>ix</b>
<b>List of Abbreviations</b>	<b>x</b>
<b>1 Introduction</b>	<b>11</b>
1.1 Motivations . . . . .	11
1.2 Contributions . . . . .	14
1.3 Thesis Outline . . . . .	15
1.4 Target Audience . . . . .	15
<b>2 Related work</b>	<b>16</b>
2.1 Error detection techniques . . . . .	16
2.1.1 Dictionary lookup . . . . .	16
2.1.2 n-gram analysis . . . . .	17
2.2 Error correction techniques . . . . .	17
2.2.1 Minimum edit distance . . . . .	18
2.2.2 Similarity key technique . . . . .	18
2.2.3 Rule-based techniques . . . . .	18

---

2.2.4	Probabilistic techniques . . . . .	19
<b>3</b>	<b>Background</b>	<b>21</b>
3.1	Neural Networks . . . . .	21
3.1.1	Multilayer Perceptron . . . . .	23
3.1.1.1	Back-propagation . . . . .	24
3.1.2	Recurrent Neural Networks . . . . .	26
3.1.2.1	Backpropagation Through Time . . . . .	27
3.1.2.2	Long short-term memory . . . . .	28
3.1.3	Bidirectional Recurrent Neural Network . . . . .	29
3.1.4	Sequence-to-sequence models . . . . .	31
3.1.5	Attention mechanism . . . . .	32
3.2	Evaluation metrics . . . . .	34
3.2.1	Standard metrics . . . . .	35
3.2.2	MaxMatch ( $M^2$ ) . . . . .	36
3.2.2.1	Method . . . . .	36
3.2.3	I-measure . . . . .	39
3.2.3.1	Method . . . . .	40
3.2.4	BLEU and GLEU . . . . .	41
<b>4</b>	<b>Experimental Setup</b>	<b>43</b>
4.1	Task . . . . .	43
4.2	The QALB corpus . . . . .	44
4.2.1	QALB corpus structure . . . . .	45
4.2.2	Creating data sets . . . . .	47
4.2.3	Refining data sets . . . . .	48
4.3	Technical details . . . . .	48
<b>5</b>	<b>Correction models</b>	<b>50</b>
5.1	Sentence Representation . . . . .	50
5.2	Models . . . . .	51
5.2.1	RNN model . . . . .	51
5.2.2	BRNN model . . . . .	52
5.2.3	Encoder-decoder model . . . . .	52

---

5.2.4	Attention-based encoder-decoder . . . . .	53
5.3	Output generation . . . . .	53
<b>6</b>	<b>Experiments and Results</b>	<b>55</b>
6.1	Data . . . . .	55
6.2	Baseline system . . . . .	56
6.3	Results of the models . . . . .	57
6.3.1	MaxMatch $M^2$ . . . . .	57
6.3.2	I-measure . . . . .	58
6.3.3	BLEU and GLEU . . . . .	59
<b>7</b>	<b>Conclusion and future work</b>	<b>60</b>
7.1	Limitations . . . . .	60
7.2	Future studies . . . . .	61
7.3	Conclusion . . . . .	61
<b>A</b>	<b>Appendix</b>	<b>63</b>
A.1	Adam . . . . .	63
A.2	Annotation example from the QALB corpus . . . . .	64
A.3	Error calculation . . . . .	65
A.4	Qualitative comparison of the correction models . . . . .	67
	<b>Bibliography</b>	<b>67</b>

---

## List of Figures

---

1.1	A recurrent network (in left) vs. a feed-forward network (in right). . . . .	13
1.2	Encoder-Decoder model. Encoder units are specified in cyan and decoder units in gray. Green rectangles refer to the Softmax function, used for a probabilistic interpretation for the output. . . . .	14
3.1	A simple Perceptron (in left) vs. a Multi-Layer Perceptron (in right) . . . . .	22
3.2	Most common activation functions: sigmoid, tanh, ReLU. . . . .	23
3.3	A Recurrent Neural Network with one hidden layer in different time steps. The gradient calculation for time step $t$ is shown in red. . . . .	26
3.4	Activation function tanh and its derivative . . . . .	28
3.5	Architecture of LSTM. The gates are implemented to compute a value between 0 and 1 and are multiplied to partially allow or deny information to flow into or out of the memory (cell state). The final value for a given input is then transferred to the output connection of the layer. . . . .	29
3.6	A bidirectional recurrent neural network model . . . . .	30
3.7	Sequence-to-sequence model for three time steps. The state vector of the last layer of the encoder is passed to the decoder as input. Consequently, each output of the decoder is used as the input of the succeeding layer in the decoder. . . . .	31

---

3.8	Attention-based sequence-to-sequence model for three time steps. The red arrows depict the connection of the intermediate vector $c$ with the state units of the decoder. . . . .	33
3.9	Levenshtein edit distance matrix for the input and the prediction in example 3.2.2.1 . . . . .	37
3.10	Edit lattice of matrix 3.9. Transitive edits are weighted by the sum of the consisting parts. . . . .	38
3.11	An example of conversion of the gold-standard $m2$ format to the annotation scheme of I-measure . . . . .	39
4.1	An example of the correction process based on the structure of the QALB corpus. $S$ and $A$ refer respectively to the beginning of the sentence and the gold annotators' correction . . . . .	45
4.2	Actions used by annotators for correction in the QALB corpus . . . . .	46
4.3	Existing characters in the QALB corpus, each one separated by a space . . . . .	48
4.4	Number of words based on frequency in the QALB corpus . . . . .	49
5.1	A recurrent neural network for error correction. The input is a potentially erroneous sequence. The output of the model is supposed to be the corrected form of the input. Note the fix size of the input and output sequence. . . . .	51
A.1	An annotation example from the QALB corpus. Following line $S$ which refers to the potentially wrong original text, the correction actions are added for each incorrect token in $A$ lines. . . . .	64
A.2	Correction of an input using different models. The size of each sequence is marked in the header of each sequence. . . . .	67

---

## List of Tables

---

3.1	Extended Writer-Annotator-System evaluation system . . . . .	41
6.1	Train, validation and test set of the QALB corpus. . . . .	56
6.2	Evaluation results of the models using MaxMatch $M^2$ metric. Bold numbers indicate the best method per measure. . . . .	57
6.3	Evaluation of the models using I-measure metric . . . . .	58
6.4	Evaluation of the models using BLEU and GLEU metrics . . . . .	59
7.1	Evaluation of models using fixed-size sequences . . . . .	61

---

## List of Abbreviations

---

**ANN** Artificial Neural Networks.

**BPTT** Back-Propagation Through Time.

**BRNN** Bidirectional Recurrent Neural Network.

**CRF** Conditional Random Field.

**GRU** Gated Recurrent Unit.

**HMM** Hidden Markov Models.

**LSTM** Long Short-Term Memory.

**MLP** Multilayer Perceptron.

**NMT** Neural Machine Translation.

**QALB** Qatar Arabic Language Bank.

**RNN** Recurrent Neural Network.

**SGD** Stochastic Gradient Descent.

**SMT** Statistical Machine Translation.

### 1.1 Motivations

Automatic spelling and grammar correction is the task of automatically correcting errors in written text. If you type an incorrect word in a text environment in a phone or on text-editor software, it would be detected as an error, which means that the word is not the best choice for that context. It is then auto-corrected using another word, or a list of possible alternative words is suggested. Nowadays, error correction systems are inseparable components of any text-related application.

A few months ago, a tweet from Donald Trump went viral: “Despite the constant negative press *covfefe*”. Based on our prior knowledge from the context, we can guess that he meant to use the word “coverage” rather than “*covfefe*”. But how can this prior knowledge be represented in a computer system?

We can look at the problem of error correction in multiple ways. In traditional methods, correction models were generally based on the linguistic nature of errors. However, because of the notorious complexity and irregularity of human language, and the great variability in the types of error as well as their syntactic and semantic dependencies on the context, more models that performed better were needed. Example 1.1 demonstrates examples of error types that seem to need more complex methods than those that are language-based.

**Example: Different types of errors**

- We went to the store and bought *new stove*. [*new stove* → a new stove]
- Mary wishes that she *does* not live in a dormitory. [*does* → did]
- Anni said that he wants *too* give her a gift. [*too* → to]
- He was so quiet that *hardly he noticed* her. [*hardly he noticed* → he hardly noticed]
- I drank a double espresso. Now I'm *entirely* awake. [*entirely* → wide]

Given a potentially erroneous input phrase, some approaches use classifiers to generate corrections by modeling their interactions with, for example, an n-gram [1] or a Conditional Random Fields model (CRF) [2]. Statistical Machine Translation (SMT) systems have been used successfully in this context, in particular as a result of the increasing availability of manually annotated corpora. However, their major drawback is the difficulty of modeling corrections in different granularities, e.g., characters and words, which is necessary to reduce the rate of unknown words that are detrimental to their proper functioning. More recently, the use of neural networks has delivered significant gains for mapping tasks between pairs of sequences, including translation and spelling correction, due to their ability to learn a better representation of the data and a better consideration of the context.

In a simple feed-forward neural network, such as Multi-Layer Perceptron (MLP), we assume that all inputs and outputs of the network are independent on each other, which is not the best method for many tasks, particularly for language modeling. By sharing parameters across different parts of a model, we can turn an MLP into a Recurrent Neural Network (RNN). RNNs are a set of neural networks for processing sequential data and modeling long-distance dependencies which is a common phenomenon in human language. The simplest form of a recurrent neural network is an MLP with the previous set of hidden unit activations feeding back into the network together with the inputs, so that the activations can follow a loop. Therefore, unlike feed-forward networks — which are amnesiacs regarding their recent past — the MLP enables the network to do temporal processing and learn sequences, which are important features for language modeling. Figure 1.1 illustrates a feed-forward network and a recurrent network.

In theory, RNNs are capable of handling long-term dependencies, but this is not the case in practice. Different architectures have been proposed to tackle this problem, from which we have used Long Short Term Memory (LSTM) architecture. LSTMs help main-

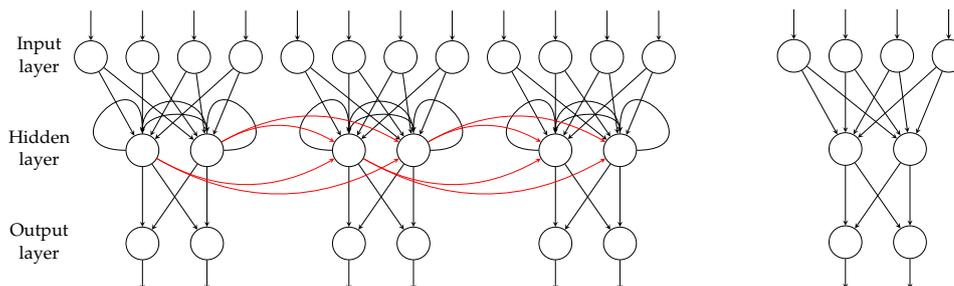


Figure 1.1: A recurrent network (in left) vs. a feed-forward network (in right).

tain a constant error that can be back-propagated through time and layers, so that they allow recurrent networks to continue learning gradually over time. Henceforth, our implementation of an RNN model is based on the LSTM architecture.

In this project, we are particularly interested in Encoder-Decoder models. The basic idea of the model is relatively simple. We use two RNN models: the first encodes the information of the potentially erroneous input text, as a vector of real-value numbers, while the second model decodes this information into the target sentence which is the corrected prediction for the input text. Figure 1.2 demonstrates an encoder-decoder model for error correction.

The fact that the encoder-decoder attempts to store whole sentences of any arbitrary length in a hidden vector of fixed size, leads to a substantial size for the network which hinders the training process for large data sets. On the other hand, even if the network were large, it would not be efficient in terms of memory and time when processing shorter phrases. In addition, the long-distance dependencies would be ignored by time, even if LSTM were supposed to correct it. To remedy this problem, an attention mechanism is used to allow the decoder to *attend* different parts of the source sentence at each step of the output generation.

The ability to correct errors accurately will improve the reliability of the underlying applications and thus numerous commercial and academic implications. It facilitates the construction of software to help foreign language learning, as it reduces noise in the entry of NLP tools, thus improving their performance, especially on unedited texts that can be found on the Web.

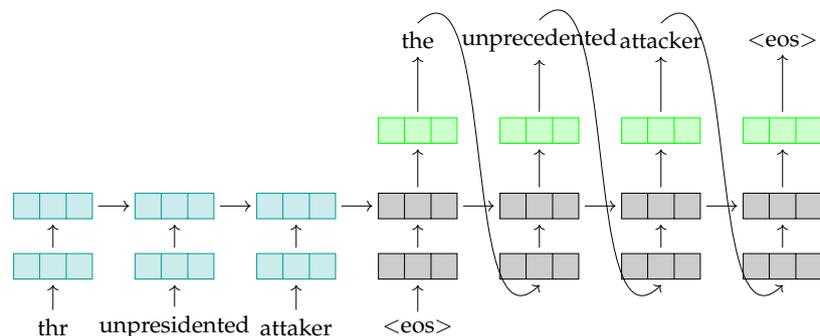


Figure 1.2: Encoder-Decoder model. Encoder units are specified in cyan and decoder units in gray. Green rectangles refer to the Softmax function, used for a probabilistic interpretation for the output.

## 1.2 Contributions

Inspired by the prevailing successes of neural networks, we decided to explore in detail certain models for the task of spelling and grammar correction. This thesis focuses on how to create such systems using sequence-to-sequence models together with attention mechanism. The first of two contributions is the architecture and training of such a model at character-level. The ultimate aim is to explore whether NMT methods can deliver competitive results.

Although our research does not aim at a specific language, since models are tested on an Arabic corpus, this thesis also explores whether this project does contribute better error correcting than previous works for Arabic. To summarize, the following tasks have been performed in chronological order:

- Collection and preprocessing data.
- Implementation of our target models using DyNet. Although similar models were already implemented in other toolkits, e.g., Theano <sup>1</sup>, TensorFlow <sup>2</sup>, we had to code the models entirely in DyNet <sup>3</sup>.
- Experiments and comparison of our results with the existing results
- Documentation of the research project

<sup>1</sup><http://www.deeplearning.net/software/theano/>

<sup>2</sup><https://www.tensorflow.org/>

<sup>3</sup><https://github.com/clab/dynet>

## 1.3 Thesis Outline

This introductory chapter is followed by a review in chapter 2 of previous research on the same theme. Chapter 3 describes the background methods needed for this research. In chapter 4, the experimental setup of the work is described. In this section the corpus is preprocessed and data sets are created. Chapter 5 describes further details about the models that are used for error correction. In chapter 6, the results of our correction system are presented. This work is concluded in chapter 7. We will also propose several new ideas for future works in this domain. The final sections of this thesis are devoted to the appendices in A and the bibliography in A.4.

## 1.4 Target Audience

This thesis endeavors to provide clear and coherent information for two audiences: developers and computational linguists. Since this work does not address the linguistic aspects of spelling and grammatical error correction, it may not be useful for linguistic researchers.

This chapter reviews previous research on automatic spelling and grammar correction. Section 2.1 presents a general survey on the methods often used for error detection tasks. Following this section, section 2.2 explains the most common error correction techniques. Since our principal task is predominantly related to error correction using neural networks, the emphasis in chapter 3 is on introducing the basic notions of neural networks, particularly Recurrent Neural Networks.

### 2.1 Error detection techniques

In most systems, before any correction is carried out, a detection process is conducted on the input sentence to extract the potentially incorrect words. Two main techniques [3] are used to detect non-word spelling errors that are explained in the following sections. A non-word error refers to a potentially incorrect word that does not exist in a given dictionary.

#### 2.1.1 Dictionary lookup

Dictionary lookup is one of the basic techniques employed to compare input strings with the entries of a language resource, e.g., lexicon or corpus. Such a language resource must contain all inflected forms of the words and it should be updated regularly. If a

given word does not exist in the language resource, it will be marked as a potentially incorrect word. Reducing the size of resources and ameliorating the search performance, through morphological analysis and pattern-matching algorithms (e.g., hashing, search trees), presents a challenge in this method.

### 2.1.2 n-gram analysis

Statistical models were designed to assign a probability to a sequence of symbols, e.g. characters and words. *n-gram* is one of the popular models that represents an attempt to count a structure by counting its substructure. Given a sentence  $C = c_1, c_2, \dots, c_l = c_1^l$ , we can define the probability over a character  $c_i$  as follows:

$$P(C = c_1, c_2, \dots, c_l = c_1^l) = \prod_{i=1}^l P(c_i | c_1, c_2, \dots, c_{i-1}) \quad (2.1)$$

Assuming different  $n$ -values in the  $n$ -gram model, we can create different probabilistic models. In a unigram model ( $n = 1$ ), the probability of each character is calculated independently from previous characters, i.e.,  $P(c_i | c_1, c_2, \dots, c_{i-1}) \approx P(c_i)$ . Although the unigram model ignores the context by taking only the current character  $c_i$  into account, the models where  $n > 1$  can represent a more accurate probability of the characters, since  $P(c_i | c_1, c_2, \dots, c_{i-1}) \approx P(c_i | c_{i-n}, \dots, c_{i-1})$ . A simple way to estimate probabilities for each character is known as maximum likelihood estimation (MLE). Equation 2.2 calculates the  $n$ -gram probability by dividing the number of times a particular string is observed by the frequency of the context.

$$P(c_i | c_{i-n+1}^{i-1}) = \frac{\text{count}(c_{i-n+1}^i)}{\text{count}(c_{i-n+1}^{i-1})} \quad (2.2)$$

For the task of error detection,  $n$ -gram analysis estimates the likelihood that a given input will be correctly spelled. To accomplish this, a pre-calculated  $n$ -gram statistic from the language resource is provided for comparative tasks. The choice of  $n$  depends on the size of the training corpus.

## 2.2 Error correction techniques

The task of machine correction is defined as correcting a  $N$ -character source sentence  $S = s_1, \dots, s_N = s_1^N$  into a  $M$  - character target sentence  $T = t_1, \dots, t_M = t_1^M$ . Thus, any

type of correction system can be defined as a function  $MC$ :

$$\hat{T} = MC(S) \quad (2.3)$$

which returns a correction hypothesis  $\hat{T}$  given an input word  $S$ . In the case of word-level models, inputs can be considered as the source words and the output are the predicted words based on the same function.

### 2.2.1 Minimum edit distance

Minimum edit distance is one of the most studied techniques for error correction. It is based on counting editing operations, which are defined in most systems as insertion, deletion, substitution and transposition, in order to transform an incorrect input into the most probable word, i.e., the one with least edit distance. Hamming [4], Jaro–Winkler [5], Wagner–Fischer [6], Damerau-Levenshtein [7] and Levenshtein [8] are among the most famous edit distance algorithms. We will Levenshtein distance later in one of the evaluation metrics in section 3.2.2.

### 2.2.2 Similarity key technique

Similarity key technique is another technique for error correction. It is based on classifying characters into the groups with similar keys. When a word is detected to be potentially incorrect, its characters are mapped into the pre-defined keys so that all other words in the language resource that have the same key are suggested as correction candidates. The candidates are then ranked according to the minimum edit distance.

Various algorithms propose different approaches to design key groups for characters of a language. In general, the similarity is measured based on the position and the order of the characters in the words. Anagrams and the phonology of language are suitable factors to construct corresponding keys for each character. Soundex [9], Metaphone [10] and SPEEDCOP [11] and Caverphone [12] are among the most common methods using similarity key technique.

### 2.2.3 Rule-based techniques

By analyzing the most common spelling errors, some researchers have attempted to create a knowledge base of errors for the task of correction [13, 14, 15] using rule-based

models that encode grammatical knowledge. These rules are generally based on the morphological characteristics of the language.

#### 2.2.4 Probabilistic techniques

Approaches based on stochastic models that learn a suitable representation from data have been also used for error correction tasks. Statistical machine correction systems, like the similar systems in machine translation [16]. They perform their task by creating a probabilistic model on sequences of symbols, e.g., characters and words, also known as *language model*, in such a way that the desired output has the highest probability giving specific parameters of the model  $\theta$ . Therefore, we have:

$$\hat{T} = \underset{T}{\operatorname{argmax}} P(T|S; \theta) \quad (2.4)$$

which is called the Fundamental Equation of Machine Translation [17].

Various approaches are used to model the probability. n-gram-based language models, as mentioned in section 2.1.2, often also function as a probabilistic technique for error correction. Another common method is the log-linear language model [18] which calculates the probability by creating a feature vector that describes the context using different features and then calculating a score vector that corresponds to the likelihood of each symbol.

More highly complex probabilistic network representations have been also introduced [16, 19]. Hidden Markov Models (HMM), for example, have demonstrated a strong ability to model human language [20], but since HMMs assume conditional independence of the previous words except the last one, they are not a practicable choice for modeling the long-distance dependencies. On the other hand, the recurrent neural networks (RNN) have been demonstrated to be more capable of modeling such dependencies [21].

Neural networks are also based on the probability distribution of language and they have shown recent success in different tasks related to natural language processing. This study focuses on attention-based encoder decoder neural networks for error correction. Such systems consist of two components: an encoder that computes a representation for an input sentence, and a decoder that generates one target word at a time and decomposes the conditional probability as

$$\log(p(T|S)) = \sum_{j=1}^m \log p(t_j; t_{<j}, S) \quad (2.5)$$

where  $t_j$  is the current word of sentence  $T$ ,  $t_{<j}$  refers to the previous words and  $S$  is the source sentence.

These neural machine translation (NMT) systems have been used in recent related works with different architectures of the decoder and different representations in the encoder. Kalchbrenner and Blunsom [22] used an RNN for the decoder and a convolutional neural network for encoding the source sentence. Sutskever et al. [23] and Luong et al. [24], on the other hand, have used multiple layers of an RNN with a Long Short-Term Memory (LSTM) hidden unit for both encoder and decoder. Cho et al. [25], Bahdanau et al. [26] and Jean et al. [27] have all used another architecture for hidden units, the Gated Recurrent Unit (GRU), for both encoding and decoding components.

The work of Luong et al. [24] is the most relevant to the present study. While our aim is to apply the same principals to error correction task in character-level models. We can also say that this work is the continuation of the generalized character-level spelling error correction presented in [28] which is a model that maps input characters into output characters using supervised learning. Although the model is applied to correct errors in Arabic, it does not consider morphological or linguistic features. On the other hand, the model is context-sensitive, and it explores beyond the context of the given word.

This chapter is divided into two sections and provides an overview of the background needed for constructing our error correction system: section 3.1 explains the most important characteristics of neural networks and special methods that can be used for modeling long-distance dependencies. Section 3.2 then addresses the most common evaluation metrics for error correction systems. The evaluation results of our designed models in chapter 5 are analyzed in chapter 6.

### 3.1 Neural Networks

Artificial neural networks (ANN) were originally developed as mathematical models inspired by the biological neural networks. A biological neural network is an interconnected network of neurons transmitting elaborate patterns of electrical signals. Dendrites receive input signals and, based on those inputs, trigger an output signal via an axon. An ANN follows the same principle using mathematical computations to process inputs.

In 1943, McCulloch and Pitts modeled mathematically the biological neural networks [29]. This network could solve binary problems but without learning. Later, in 1957, Rosenblatt presented the first ANN that was able to vary its own weights, in other words, the network could learn and ameliorate itself (Figure 3.1).

The perceptron computes a single output from multiple real-valued inputs by form-

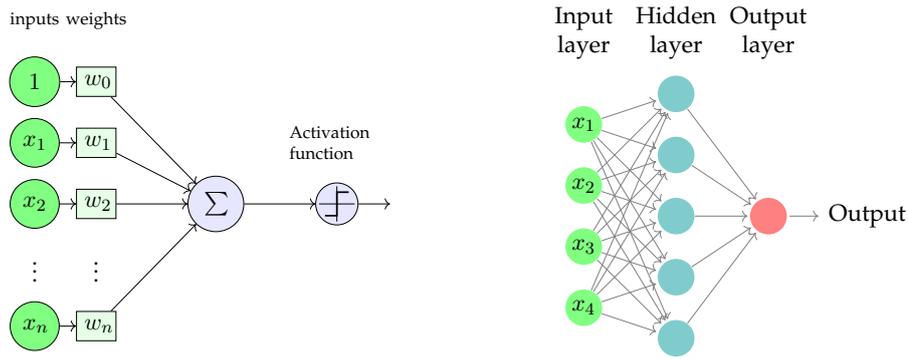


Figure 3.1: A simple Perceptron (in left) vs. a Multi-Layer Perceptron (in right)

ing a linear combination according to its input weights and then possibly putting the output through some nonlinear activation function. Technically speaking, let  $x_1, x_2, \dots, x_n \in \mathbb{R}$  be  $n$  scalar inputs of the network. Output  $y$ , as function of a sum of weighted inputs is computed by:

$$z = \sum_{i=1}^n w_i x_i + b = w^T x + b \quad (3.1)$$

where  $w = (w_1, w_2, \dots, w_n) \in \mathbb{R}^n$  is corresponding weights,  $b \in \mathbb{R}$  is bias and  $\varphi$  is an activation function. The weight is often referred to as *pre-activation* and the bias is replaced by an equivalent input term  $x_0 = 1$  weighted by  $w_0 = b$ . Note that this is simply the dot product of the weight vector  $w$  and the input vector  $x$ . The result of this affine transformation  $z$  is then passed through a step function that determines the binary output of the Perceptron.

$$y = \begin{cases} 1 & \text{if } z \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

Given a training pair  $(x, y)$ , the parameters of the network, i.e., weights and bias, are learning using the following rule:

$$w_i \leftarrow w_i - \eta \cdot (\hat{y} - y) \cdot x_i \quad (3.3)$$

where  $\hat{y}$  is the output of the Perceptron,  $y$  is the desired output and  $\eta$  is the learning rate to adjust the updating magnitude. Note that the  $\leftarrow$  refers to the assignment of the new value of  $w_i$ .

### 3.1.1 Multilayer Perceptron

Although conceptually remarkable, Perceptron is significantly limited in modeling non-linearly separable problems. A famous example is that of the XOR logic operation which cannot be solved by a linear transformation. An evolution of Perceptrons introduces an intermediate between the input and the output, which is known as the hidden layer. Hidden layers allow the pre-activation to be followed by a nonlinearity, which produces a nonlinear transformation that can map the input into a linearly separable space, such as the case of XOR. Figure 3.1 demonstrates an MLP.

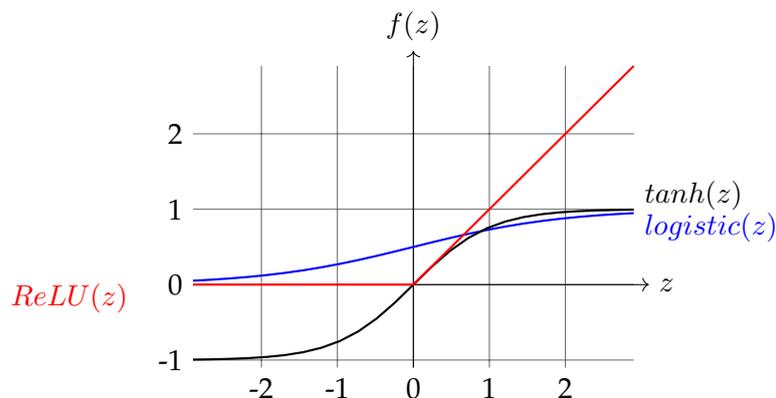


Figure 3.2: Most common activation functions: sigmoid, tanh, ReLU.

An MLP is a network of Perceptrons grouped in interconnected layers. Each connection from the  $i$ -th neuron of layer  $l-1$  to the  $j$ -th neuron of layer  $l$  is associated to a weight  $w_{ij}^{(l)}$ , and all the weights are stored in a matrix  $W^{(l)}$ . It is a feed-forward neural network with one or more hidden layers between the input and output layer. Feed-forward means that data flows in one direction from input to output layer. Similar to the Perceptron, each layer computes an affine transformation as follows:

$$z^{(l)} = W^{(l)} \cdot a^{(l-1)} \quad (3.4)$$

Generally, each layer of an ANN computes some activation based on its input and a nonlinear activation function. The choice of which activation function to use in each layer can have a significant impact on the performance of the model. An activation function is a non-linear function that maps any real-valued number to a reduced range. Hyperbolic tangent (tanh), logistic function and rectified linear unit (ReLU) function are among the most common activation function. They do the aforementioned mapping in  $[-1, 1]$ ,  $[0, 1]$  and  $[0, x]$  respectively. These functions are illustrated in figure 3.2.

$$\tanh(x) = \frac{e^{2x} - 1}{e^{2x} + 1} \quad (3.5)$$

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (3.6)$$

$$\text{ReLU}(x) = \max(0, x) \quad (3.7)$$

Both the number of units in the output layer and the choice of output activation function depend on the task. Binary classification can be performed with a one-neuron output layer. An output can be interpreted as the probability  $Pr(\text{True}|x)$  of assigning a predefined class to an input  $x$ . In the case that a logistic sigmoid is used as the activation function, the model is called *logistic regression*, or a *logit* model. In a more general way, we can solve the same problematic for an N-class classification ( $N > 2$ ) task, by using N output neurons normalizing by a softmax function [30]. The softmax is a squashing function that maps its input to a categorical distribution. It is defined as follows:

$$\text{softmax}_i(z) = \frac{\exp(z_i)}{\sum_{n=0}^N \exp(z_n)} \quad (3.8)$$

where  $z$  refers to the affine transformation mentioned in equation 3.1.

### 3.1.1.1 Back-propagation

Since the solution of the network exists in the values of the parameters (weights and bias)  $\theta$  between neurons, an optimization process is needed to acquire parameters that fit the training corpus. One of the most common algorithms is termed back-propagation. The basic idea is to modify each weight of the network by a factor proportional to the error, i.e., the difference between the desired output and the output of the network. Different functions have been introduced as objective function to calculate the error from which we use the Cross-Entropy cost as follows:

$$H(\hat{y} - y) = \sum_i y_i \log \frac{1}{\hat{y}_i} = - \sum_i y_i \log \hat{y}_i \quad (3.9)$$

where  $\hat{y}_i = MLP(x_i; \theta)$  refers to the output of the network and  $y$  is the desired output. The summation computes the overall loss of the network on a single pair. An objective function is either a loss function or its negative form. Here, we assume that equation 3.9 is our objective function. In fact, we calculate the error in order to reduce it for each pair

by searching for the best parameters  $\theta$ :

$$\theta^* = \underset{\theta}{\operatorname{argmin}} H(\hat{y} - y) \quad (3.10)$$

Since MLPs are differentiable operators, except for the case where ReLU function is used, they can be trained to minimize any differentiable cost function using *gradient descent*. The gradient descent aims to determine the derivative of the objective function with respect to each of the network weights, then adjust the weights in the direction of the negative slope. Having said that since the ReLU function is not differentiable, Leaky ReLU is generally used for back-propagation where:

$$\operatorname{LeakyReLU}(x) = \begin{cases} x & \\ \text{if } x > 0.01x & \text{otherwise} \end{cases} \quad (3.11)$$

**Stochastic Gradient Descent** Stochastic Gradient Descent (SGD) is an iterative process where we pick an observation uniformly at random, say  $i$  and attempt to improve the likelihood with respect to the observation. Each iteration, also called epoch,  $\theta$  is updated as follows:

$$\theta \leftarrow \theta - \eta \nabla_{x_i} H(\hat{y} - y) \quad (3.12)$$

where  $\eta$  is the learning rate that determines the size of the steps needed to reach a local minimum. In other words, the gradient follows the slope direction of the surface created by the objective function until it reaches a minimum.

While SGD performs one update at a time, two variants of SGD perform this update differently. In the Vanilla Gradient Descent method, the gradient of the loss function is computed for the entire training set. Thus, it can be extremely slow and intractable for large training sets. On the other hand, the Mini-batch Gradient Descent — instead of calculating the entire training set — divides it into small subsets and computes the gradient with respect to a subset of the training set. In other words, this model uses the efficiency of the Vanilla Gradient Descent in finding the optimal convergence and the performance of the stochastic variant.

In appendix [A.1](#) we will introduce Adaptive Moment Estimation (Adam) optimization method as well.

### 3.1.2 Recurrent Neural Networks

Recurrent neural network (RNN) is another type of neural network that is particularly used for language processing tasks. The main characteristic of an RNN, in addition to those of any ANN (e.g., units and parameters), is the cyclic connections that each unit can have. This provides a memory-like functionality for the RNN which could be an appropriate model for language processing. Unlike a simple MLP that maps inputs into the outputs directly, an RNN can make use of, theoretically, all the previous inputs for each output, thus providing an internal state other than the network's parameters.

In order to create an RNN, we first need to compress a sequence of input symbols  $X = (x_1, x_2, \dots, x_n)$  into a fixed-dimensional vector by using recursion. Using this vector enables us to deal with variable-length inputs and outputs. Then, assume at step  $t$  that we have a vector  $h_{t-1}$  which is the history of all previous symbols before the current one. The RNN will compute the new vector, or its internal state,  $h_t$  which compresses all the previous symbols  $(x_1, x_2, \dots, x_{t-1})$  as well as the new symbol  $x_t$  by

$$h_t = \tanh(Wx_t + Uh_{t-1} + b) \quad (3.13)$$

$$\hat{y}_t = \text{softmax}(Vh_t) \quad (3.14)$$

where  $W$  is the input weight matrix,  $U$  is the recurrent weight matrix,  $V$  is the weight of the hidden layer and  $b$  is the bias vector. Although this is not the only way to model an RNN [31], we have used this formulation in our implementation (which is known as Elman network [32]).

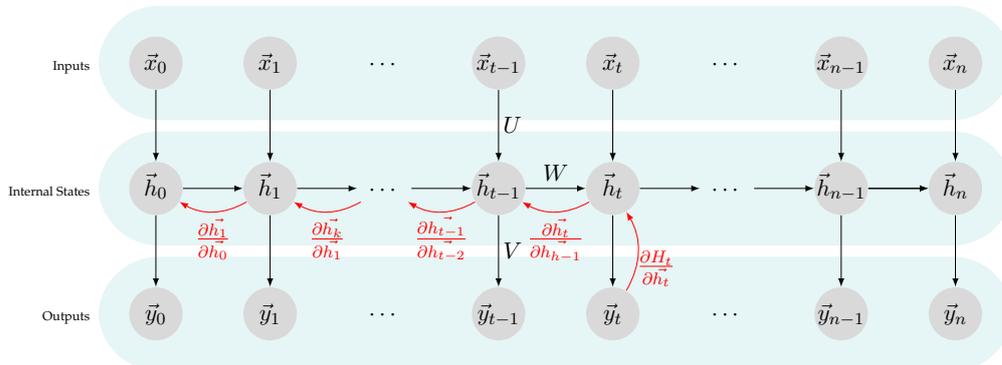


Figure 3.3: A Recurrent Neural Network with one hidden layer in different time steps. The gradient calculation for time step  $t$  is shown in red.

Training an RNN is similar to training a traditional Neural Network. We also use

the back-propagation algorithm, but with some modifications. Since the parameters are shared by all time steps in the network, the gradient at each output depends not only on the calculations of the current time step, but also the previous time steps. For instance, in order to calculate the gradient at  $t = 4$ , we need to back-propagate three steps and sum up the gradients. This method is called *Back-propagation Through Time* (BPTT).

### 3.1.2.1 Backpropagation Through Time

In the previous section, we explained the basic mechanism of an RNN. In calculating the back-propagation, we mentioned that the gradients of the error should be calculated with respect to the parameters of the network, i.e.,  $U$ ,  $V$  and  $W$ , for the current and previous time steps. To do so, we take use of the BPTT algorithm to do so.

Let us assume that we want to calculate the gradients for time step  $t$ . To calculate the gradients, using the chain rule of differentiation, we have:

$$\frac{\partial H_t}{\partial V} = \frac{\partial H_t}{\partial \hat{y}_t} \frac{\partial \hat{y}_t}{\partial V} = \frac{\partial H_t}{\partial \hat{y}_t} \frac{\partial \hat{y}_t}{\partial z_t} \frac{\partial z_t}{\partial V} = (\hat{y}_t - y_t) \otimes h_t \quad (3.15)$$

where  $z_t = Vh_t$  and  $\otimes$  refers to the outer product of two vectors. As we see,  $\frac{\partial H_t}{\partial V}$  depends only on the values at the current time step  $y_t$ ,  $\hat{y}_t$  and  $h_t$ . Although the gradient with respect to  $W$  and  $U$  is a little different:

$$\frac{\partial H_t}{\partial W} = \frac{\partial H_t}{\partial \hat{y}_t} \frac{\partial \hat{y}_t}{\partial h_t} \frac{\partial h_t}{\partial W} \quad (3.16)$$

where  $h_t$  is the activation function 3.14 which depends on the previous internal state  $h_{t-1}$  which depends on the previous states  $h_{t-2}, h_{t-3}, \dots, h_0$  and  $W$  (shown in figure 3.3 in red). So we need to apply the chain rule on all the previous states to obtain the gradient  $\frac{\partial H_t}{\partial W}$  as follows:

$$\frac{\partial H_t}{\partial W} = \sum_{k=0}^t \frac{\partial H_t}{\partial \hat{y}_t} \frac{\partial \hat{y}_t}{\partial h_t} \left( \prod_{j=k+1}^t \frac{\partial h_j}{\partial h_{j-1}} \right) \frac{\partial h_k}{\partial W} \quad (3.17)$$

where the inner product refers to the chain rule until the lower limit of  $k$ . The only difference between this calculation and that of the standard back-propagation is that we sum up the gradients for  $W$  at each time step. In an MLP, for example, since we don't share parameters across layers, so we don't need to add anything. The calculation above is the same when we calculate the gradient with respect to  $U$  as well.

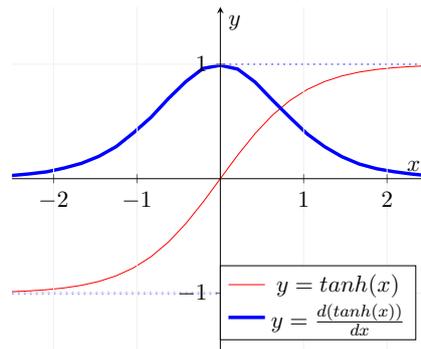


Figure 3.4: Activation function tanh and its derivative

Equation 3.17 results in a square matrix including first-order partial derivatives of the vector of error function, also known as Jacobian matrix. One of the difficulties in the training process is the tendency of the derivatives of the activation function towards 0 as both extremes (the blue curve in figure 3.4) with a flat curve. Consequently, when applying the chain rule, the gradient of a layer tends towards 0 and vanishes completely after a few time steps. Vanishing gradients is a common problem, particularly in the networks that are intended to model a language. Long-distance dependencies are one of the characteristics of human language.

A more recent popular solution to confront the problem of the vanishing gradient is using Gated Recurrent Unit (GRU) or Long Short-Term Memory (LSTM) architectures. In our project, we used the latter, which is explained in the next section.

### 3.1.2.2 Long short-term memory

In the previous section, we stated that the vanishing gradient problem prevents standard RNNs from learning long-term dependencies. Long Short-Term Memory (LSTM) network is an architecture for the RNN, capable of learning long-term dependencies. They were initially introduced by Hochreiter and Schmidhuber [33] and were popularized by many other works in the following years (particularly in [34, 35, 36, 37, 38]). It is among the most widely used models in Deep Learning for NLP today.

An LSTM network is composed of memory cells and gate units to calculate the internal states. Three gates control the behavior of the memory cells: ignore gate  $i$ , forget gate  $f$  and output gate  $o$ . These gates intuitively determine precisely how much of a given vector is to be taken into account at a specific time step  $t$ . For instance, the input gate defines how much of current input should go through, the forget input defines much of

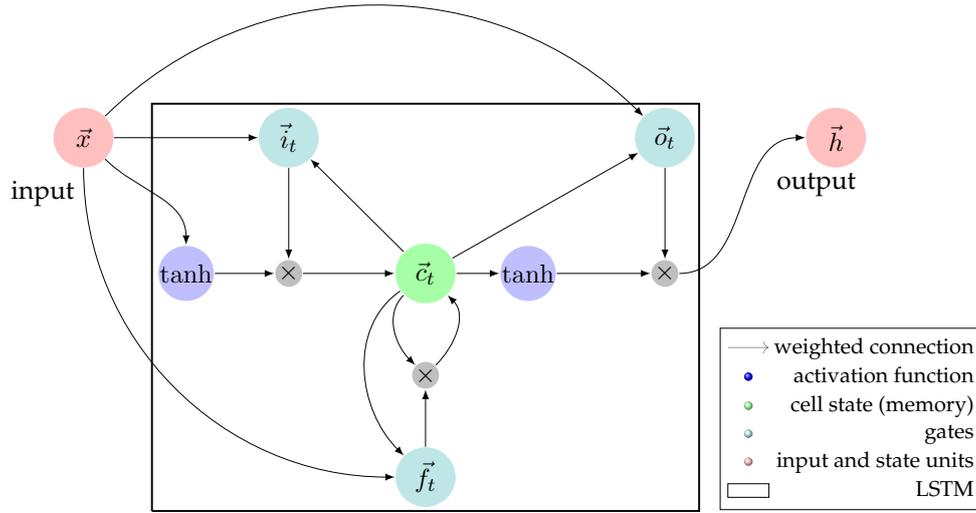


Figure 3.5: Architecture of LSTM. The gates are implemented to compute a value between 0 and 1 and are multiplied to partially allow or deny information to flow into or out of the memory (cell state). The final value for a given input is then transferred to the output connection of the layer.

the previous state pass and the output gate indicates how much of the internal state be influenced in the next time steps and next layers. The full architecture of the LSTM is illustrated in figure 3.5. Formally, an LSTM is defined as:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{ci}c_{t-1} + b_i) \quad (3.18)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{cf}c_{t-1} + b_f) \quad (3.19)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{co}c_t + b_o) \quad (3.20)$$

$$c_t = f_t \cdot c_{t-1} + i_t \cdot \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c) \quad (3.21)$$

$$h_t = o_t \tanh(c_t) \quad (3.22)$$

where  $i_t$ ,  $f_t$  and  $o_t$  respectively refer to the input, forget and output gates at time step  $t$ .  $c_t$  is the memory cell vector at time step  $t$  and  $h_t$  is the output layer. Note that the gates have the same equations, but with different parameter matrices. Various variations of LSTM can be found in [39].

### 3.1.3 Bidirectional Recurrent Neural Network

So far, we have focused on the RNNs that consider the past observation in order to predict the next symbol in a sequence. For many sequence-labeling tasks, e.g., predicting the

correct character for a given input, in addition to the past context, we would like to have access to the future context. To do so, we can use two RNN models; one that reads through the input sequence forwards (left-to-right propagation) and the other backwards (right-to-left propagation), both with two different hidden units but connected to the same output (figure 3.6). This model is called Bidirectional Recurrent Neural Network (BRNN) [40].

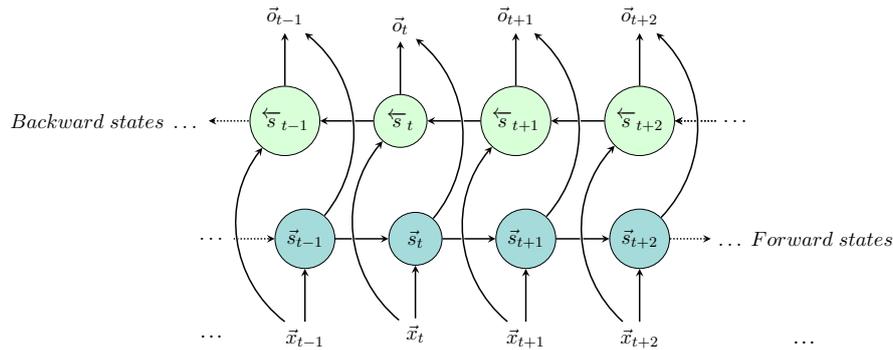


Figure 3.6: A bidirectional recurrent neural network model

The right-to-left propagation for the BRNN hidden layers is the same as for a simple RNN, except that the input sequence is presented in the opposite direction to the two hidden layers. Processing the hidden layers of the two RNN models is necessary to update the output layer. The following equation shows the mathematical formulation behind setting up the bidirectional RNN hidden layer. Note that calculating the hidden layers follows the same formulation mentioned in 3.14:

$$\vec{h}_t = \tanh(\vec{W}x_t + \vec{U}\vec{h}_{t-1} + \vec{b}) \quad (3.23)$$

$$\overleftarrow{h}_t = \tanh(\overleftarrow{W}x_t + \overleftarrow{U}\overleftarrow{h}_{t+1} + \overleftarrow{b}) \quad (3.24)$$

The corresponding calculation for each direction is shown by using vector on each formulation. These two calculations are concatenated as  $h_t$  in equation 3.14. In order to train such a model, we will need two twice as much memory for the parameters.

BRNNs have previously given improved results in various domains in speech processing [41, 42]. More recently, deep BRNN have been used where each lower layer feeds the next layer [43, 44].

### 3.1.4 Sequence-to-sequence models

Apart from long dependency, another challenge in language models using neural networks is the variable-length output spaces, e.g., words and sentences. Sequence-to-sequence neural models, or Encoder-Decoder network, have demonstrated the ability to address this challenge of the variable length of input and output. It is a generative neural network model that, given a string of inputs, produces a string of outputs, both of arbitrary lengths.

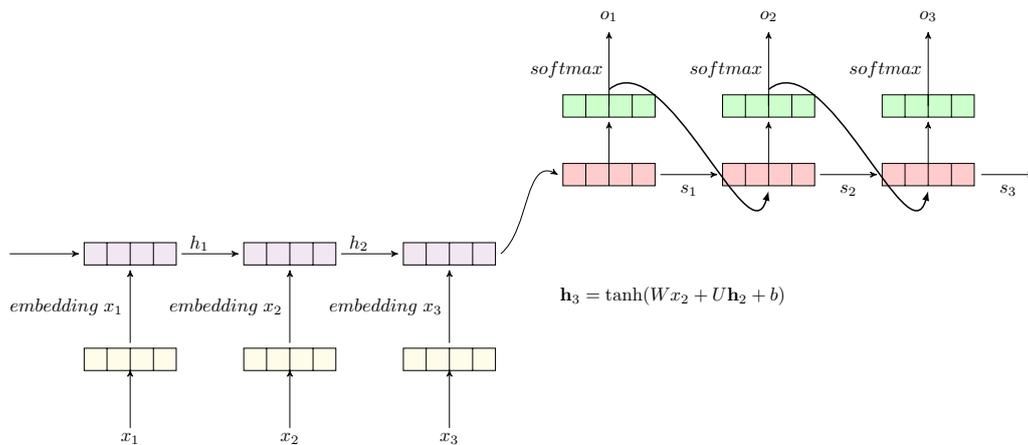


Figure 3.7: Sequence-to-sequence model for three time steps. The state vector of the last layer of the encoder is passed to the decoder as input. Consequently, each output of the decoder is used as the input of the succeeding layer in the decoder.

The sequence-to-sequence model is composed of two processes : *encoding* and *decoding*. In the encoding process, the input sequence  $x = (x_1, x_2, \dots, x_T)$  is fed into the encoder, which is basically an RNN model. Then, unlike a simple RNN where the output of each state unit is taken into account, only the hidden unit of the last layer  $h_t$  is kept; this vector, often called *sentence embedding* or *context vector*  $c$ , is intended to contain a representation of the input sentence:

$$h_t = LSTM(x_t, h_{t-1}) \quad (3.25)$$

$$c = \tanh(h_T) \quad (3.26)$$

where  $h_t$  is a hidden state at time  $t$ , and  $c$  is the context vector of the hidden layers of the encoder.

On the other hand, by passing the context vector  $c$  and all the previously predicted words  $\{y_1, y_2, \dots, y_{t-1}\}$  to the decoder, the decoding process predicts the next word  $y_t$ .

In other words, the decoder defines a probability over the output  $y$  by decomposing the joint probability as follows:

$$p(y) = \prod_{t=1}^T p(y_t | \{y_1, y_2, \dots, y_{t-1}\}, c) \quad (3.27)$$

$$p(y_t | \{y_1, y_2, \dots, y_{t-1}\}, c) = \tanh(y_{t-1}, h_t, c) \quad (3.28)$$

where  $y = (y_1, y_2, \dots, y_T)$  and  $h_t$  is the hidden unit.

Figure 3.7 depicts an encoding and a decoding process the input sequence of  $x_1, x_2, x_3$ . Softmax function is used in order to normalize the probability distribution for the output. Later, the Softmax output would be used for calculating the error loss as well.

In some of the first studies that used encoder-decoder RNN models, a fixed-length vector is used to represent the context of a source sentence [25, 23]. The fact that a fixed-length, regardless of the size of the input, is used, means a limitation on the model to represent variable-size input strings. More recently, Bahdanau et al. [26] explored the attention mechanism in NMT which is discussed in the following section.

### 3.1.5 Attention mechanism

Attention mechanism was introduced to address the limitation of modeling long dependencies and the efficient usage of memory for computation. The attention mechanism intervenes as an intermediate layer between the encoder and the decoder, having the objective of capturing the information from the sequence of tokens that are relevant to the contents of the sentence [26].

In an attention-based model, a set of attention weights is first calculated. These are multiplied by the encoder output vectors to create a weighted combination. The result should contain information about that specific part of the input sequence, and thus help the decoder select the correct output symbol. Therefore, the decoder network can use different portions of the encoder sequence as context, while it is processing the decoder sequence. This justifies the variable-size representation of the input sequence. This mechanism is shown in figure 3.8.

Unlike the encoder-decoder model that uses the same context vector for every hidden state of the decoder, the attention mechanism calculates a new vector  $c_t$  for the output word  $y_t$  at the decoding step  $t$ . It can be defined as:

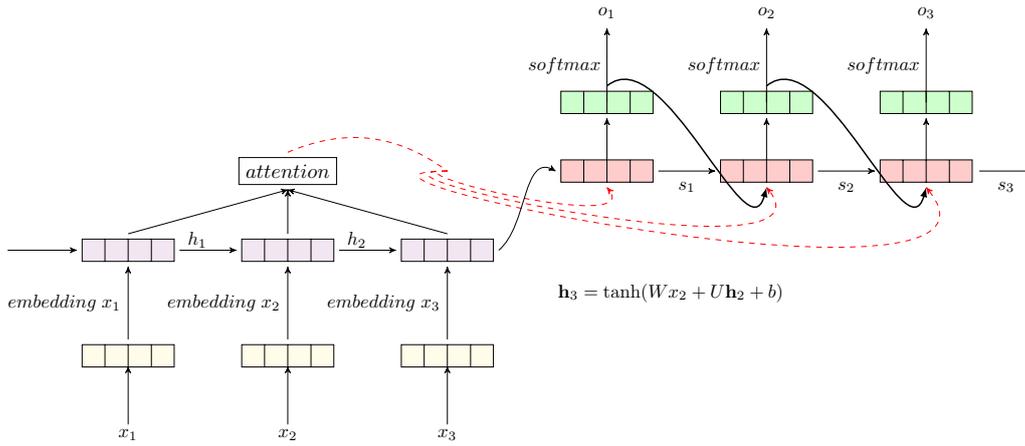


Figure 3.8: Attention-based sequence-to-sequence model for three time steps. The red arrows depict the connection of the intermediate vector  $c$  with the state units of the decoder.

$$c_t = \sum_{j=1}^T a_{tj} h_j \quad (3.29)$$

where  $h_j$  is the hidden state of the word  $x_j$ , and  $a_{tj}$  is the weight of  $h_j$  for predicting  $y_t$ . This vector is also called *attention vector* which is generally calculated with a softmax function:

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^T \exp(e_{ik})} \quad (3.30)$$

$$e_{ij} = \text{attentionScore}(s_{i-1}, h_j) \quad (3.31)$$

where the `attentionScore`, as an arbitrary function in the calculation, which scores the quality of alignment between input  $j$  and output  $i$ . It is parameterized as a feed-forward neural network in the original paper [26]. This vector is then normalized into the actual attention vector by using a softmax function over the scores:

$$\alpha_t = \text{softmax}(\text{attentionScore})_t \quad (3.32)$$

This attention vector is then used to weight the encoded representation of the hidden state  $h_j$ . Having said that in addition to this formulation, some others have been proposed for attention, particularly in [24].

## 3.2 Evaluation metrics

This section presents evaluation methods that would be used to evaluate the models. Although various methods can be used to evaluate a machine translation system, in the case of machine correction, evaluation may be limited to a binary classification of predictions where the matching elements are considered as true predictions and the others, incorrect predictions. However, a good evaluation must include more details about this comparison.

### Example: Comparison of two error correction systems

---

**Input:** The burgening cherry trees are a signe that spring ishere.

**Gold-standard:** {burgening → burgeoning, signe → sign, ishere → is here}

---

Hypothesis A: {burgening → burning, a signe → assigned}

---

Hypothesis B: {burgening → burgeoning, signe → sign, ishere. → is here., . → ∅ }

---

Example 3.2 illustrates a comparison between the results of two different correction systems. For the given input phrase, two hypotheses are proposed by the two correction systems: Hypothesis A yields, *"The burning cherry trees are assigned that spring ishere."*, while hypothesis B produces the same expected phrase of the gold-standard annotation.

How would one evaluate the performance of these correction systems? The output of system A and B includes respectively 9 and 12 tokens. How can we compare the elements of the predicted and gold-standard phrases pairwise, while they have different lengths? A further challenge arises when we have different edit actions that yield the same output which is the case of correction system B. Finally, how can we distinguish between two incorrect predictions and determine their quality? For instance, correction system A predicts *"burning"* as a correction for *"burgening"*, which is not true, but its prediction includes similar letters that indicate a better performance in comparison to a completely non-related correction.

Over the years, a number of metrics have been proposed for error correction systems evaluation, each motivated by weaknesses of previous metrics. There is no single best evaluation metric and the performance of a metric depends on the research goals and application [45]. Thus, we have evaluated our system based on the most popular metrics to date.

### 3.2.1 Standard metrics

In classification tasks, *accuracy* is one of the most used performance measures. This metric corresponds to the ratio of correctly classified inputs to the total number of inputs. More formally, accuracy can be defined as:

$$acc = \frac{TP + TN}{TP + TN + FP + FN} \quad (3.33)$$

where TN (true negative) corresponds to the number of false inputs that were correctly classified, TP (true positive) corresponds to the number of positive inputs that were classified correctly, FN (false negative) corresponds to the number of positive samples incorrectly classified as positive, and finally, FP (false positive) corresponds to the number of negative inputs incorrectly classified as positive.

One drawback of this metric is that correction actions are completely ignored. Consider correction system B in example 3.2 which yields the expected output but with different correction actions. Even if  $accuracy = 1$  in this case, the performance of actions are not reflected.

In order to take the correction actions into account, we define two additional performance metrics, *precision* and *recall*. Given a set of  $N$  sentences, where  $G_i$  is the set of gold-standard edits for sentence  $i$ , and  $P_i$  is the set of predicted edits for sentence  $i$ , we define precision, recall and F-measure as follows:

$$precision = \frac{TP}{TP + FP} = \frac{\sum_{i=1}^N |G_i \cap P_i|}{\sum_{i=1}^N |P_i|} \quad (3.34)$$

$$recall = \frac{TP}{TP + FN} = \frac{\sum_{i=1}^N |G_i \cap P_i|}{\sum_{i=1}^N |G_i|} \quad (3.35)$$

where the intersection between  $G_i$  and  $P_i$  indicates the matches between the two sets. A measure that combines *precision* and *recall* is the harmonic mean of them, also known as  $F_\beta$  - *measure*. It is defined as:

$$F_\beta = (1 + \beta^2) \cdot \frac{precision \times recall}{\beta^2 \times precision + recall} \quad (3.36)$$

Depending on the usage, different values are attributed to  $\beta$ . The most common case is in  $F_{\beta=1}$ , also known as *F-score*, which weights recall and precision evenly. As the evaluation metric in our research, we have used  $F_{0.5}$ , since it places twice as much emphasis

on precision than recall, while  $F_1$  weighs precision and recall equally. This metric has also been used in CoNLL-2014 shared task [46]. Our  $F_{0.5}$  is defined as follows:

$$F_{0.5} = (1 + 0.5^2) \cdot \frac{\textit{precision} \times \textit{recall}}{0.5^2 \times \textit{precision} + \textit{recall}} \quad (3.37)$$

To illustrate, consider correction system B in example 3.2. Based on our defined metrics, the performance of this correction is:  $P = 2/4$ ,  $R = 2/3$  and  $F_{0.5} = (1 + 0.5^2) \times P \times R / (0.5^2 \times P + R) = 10/19$ . Although this correction system yields the expected phrase of the gold-standard annotation in the output, these metrics still do not provide sufficient information. Our correction system produces its predicted text for the given source phrase, and we should provide correction actions to be compared with the gold-standard annotations.

In the case of evaluation of machine correction systems,  $P$ ,  $R$  and  $F_{0.5}$  do not provide an indicator of improvement on the original text. Given this, we opt to use more ameliorated metrics such as *MaxMatch* ( $M^2$ ), *I-measure*, *BLEU* and *GLEU*, thus gaining better insight into system performance. Each metric is explained in the following sections.

### 3.2.2 MaxMatch ( $M^2$ )

Dahlmeier et al. [47] proposed a more accurate evaluation method, which is known as MaxMatch( $M^2$ ), for efficiently computing the sequence of phrase-level edits between a source text and the predicted text. The main idea is to determine the set of edits that most often matches the gold-standard. Thus, the algorithm yields the set of phrase-level edits with the maximum overlap with the gold-standard. The edits are, then, scored using  $F_\beta$  measure.

#### 3.2.2.1 Method

Given a set of source sentences  $S = \{s_1, \dots, s_n\}$  with a set of gold-standard annotation  $G = \{g_1, \dots, g_n\}$ , our error correction system yields a set of predictions  $P = \{p_1, \dots, p_n\}$ . Each annotation  $g_i$  contains a set of edits  $g_i = \{g_i^1, \dots, g_i^m\}$ .  $M^2$  method, in the first step, extracts a set of system edits  $e_i$  for each pair of  $(s_i, p_i)$ . Thus, an edit lattice structure is constructed from a source-prediction pair. Then, the optimal sequence of edits is determined by solving a shortest path search through the lattice.

**Constructing edit lattice** Assuming that we have tokenized sentence  $i$  in the source of sentences  $S$  as  $s_i = \{s_i^1, \dots, s_i^r\}$  for which its relevant prediction is  $p_i = \{p_i^1, \dots, p_i^r\}$ . We start by calculating the Levenshtein distance in pair  $(s_i, h_i)$ . In the resulting edit distance matrix, the last element represents the edit distance value for the two sets. By determining the shortest paths from the first element of the matrix([0,0]) to the last element, we can construct a lattice structure in which each vertex corresponds to a cell in the matrix and each edge to an atomic edit operations. Atomic edit operation in the Levenshtein algorithm involves deleting a token, inserting a token, substituting a token or leaving a token unchanged.

**Example:**

**Input:** the greater the Levenshtein distances , more different strings are .

**Prediction:** the greater the Levenshtein distances , the more different strings are .

$G = \{\text{distances} \rightarrow \text{distance}, \text{more} \rightarrow \text{the more}\}$

$P = \{\text{more} \rightarrow \text{the more}\}$

Example 3.2.2.1 presents a phrase in input and the prediction of the correction system. Figure 3.9 indicates the Levenshtein distance matrix between the input and the prediction. The shortest path that leads to the minimum edit distance in the bottom-right corner of the matrix is identified by circles. Therefore, by applying a breadth-first search on the matrix, we construct the corresponding lattice.

		the	greater	the	Levenshtein	distances	,	the	more	different	strings	are	.
	⓪	1	2	3	4	5	6	7	8	9	10	11	12
the	1	⓪	1	2	3	4	5	6	7	8	9	10	11
greater	2	1	⓪	1	2	3	4	5	6	7	8	9	10
the	3	2	1	⓪	1	2	3	4	5	6	7	8	9
Levenshtein	4	3	2	1	⓪	1	2	3	4	5	6	7	8
distances	5	4	3	2	1	⓪	1	2	3	4	5	6	7
,	6	5	4	3	2	1	⓪	⓪	2	3	4	5	6
more	7	6	5	4	3	2	1	1	⓪	2	3	4	5
different	8	7	6	5	4	3	2	2	2	⓪	2	3	4
strings	9	8	7	6	5	4	3	3	3	2	⓪	2	3
are	10	9	8	7	6	5	4	4	4	3	2	⓪	2
.	11	10	9	8	7	6	5	5	5	4	3	2	⓪

Figure 3.9: Levenshtein edit distance matrix for the input and the prediction in example 3.2.2.1

Figure 3.10 illustrates the corresponding edit lattice of matrix 3.9. Each edge is weighted by the sum of the costs of its parts. Each unique part has a unit cost. The goal is to determine the sequence of edits in the lattice structure that has the maximum overlap with

the gold-standard annotation. To accomplish this, we change the cost of those edges that have an equivalent match to the gold-standard to  $-(u + 1) \times |E|$ , e.g. the edge between 6,6 and 7,8. Parameter  $u$  is defined as a threshold to avoid a substantially large number of unchanged edits. Transitive edits are allowed in the range of this threshold which is set to 2 in this lattice.

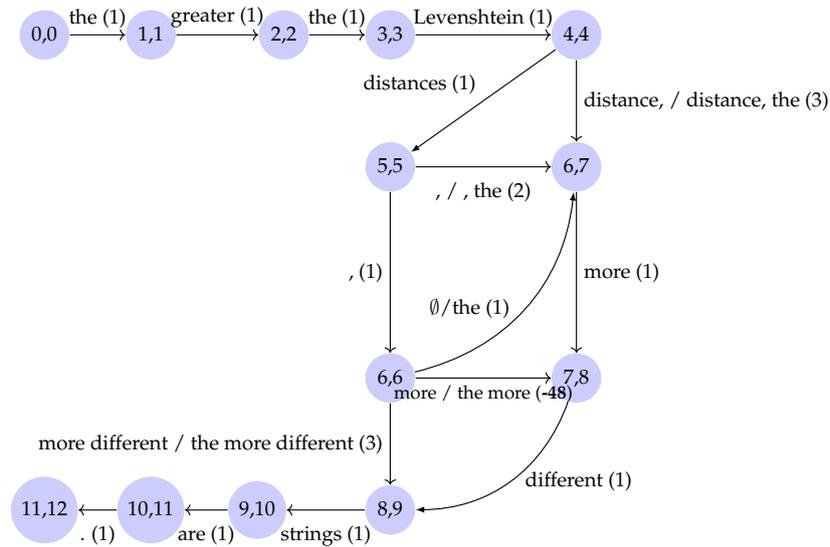


Figure 3.10: Edit lattice of matrix 3.9. Transitive edits are weighted by the sum of the consisting parts.

By performing a shortest path algorithm on the lattice, e.g. using Bellman-Ford algorithm, we obtain the shortest path in the lattice (a single-source problem with negative edges). Dahlmeier et al. [47] proved that the edits with the maximum overlap with the gold-standard edits are those of the shortest-path.

Back to our example, the  $M^2$  Scorer evaluates the correction system with  $Precision = 0.5000$ ,  $Recall = 0.5000$  and  $F_{0.5} = 0.5000$  which seems to be justifiable since the correction system has predicted only one correction that exists also in the gold-standard corrections.

Although  $M^2$  Scorer is currently a standard metric in evaluating error correction systems, having been used to rank error correction systems in the 2013 and 2014 CoNLL [48, 46] and EMNLP 2014 [49] shared tasks, it also has some weak points. In experimenting with the baseline system, since there is theoretically no correction to be done, the results of  $M^2$  Scorer appear to be non-interpretable ( $P = 100\%$ ,  $R = 0$ ,  $F_{0.5} = 0$ ). On the other hand, partial edits are also ignored. Although the system prediction  $P = \{\text{more} \rightarrow \text{the more}\}$  includes a part of the gold-standard annotation, if assumed as the only pre-

dicted edit by the correction system, it is evaluated with no difference with a completely incorrect prediction ( $P = 0, R = 0, F_{0.5} = 0$ ). These limitations have been the motivation to present the standard evaluation method for error correction systems that we discuss in the following section.

### 3.2.3 I-measure

I-measure is a novel method in response to the limitations of the previous methods, particularly the  $M^2$  Scorer [50]. Unlike  $M^2$  Scorer that uses phrase-level edits, I-measure uses tokens as the unit of evaluation. It also provides scores for both detection and correction tasks and it is sensitive to different types of edit corrections. I-measure designs a new evaluation system by presenting a new format for the gold-standard annotations. Then, by determining an optimal alignment between an input, a prediction and a gold-standard annotation, it computes the metrics that evaluates the performance of the correction system.

```
S And we keep track of all family members
health conditions .
A 2 2|||Vm|||can|||REQUIRED|||-NONE-|||0
A 8 8|||Mec|||'| |||REQUIRED|||-NONE-|||0
A 9 10|||Rloc-|||'| |||REQUIRED|||-NONE-|||0
```

```
1 <?xml version='1.0' encoding='UTF-8' ?>
2 <scripts>
3   <script id="1">
4     <sentence id="1" numann="1">
5       <text>
6         And we keep track of all family
members health conditions .
7       </text>
8       <error-list>
9         <error id="1" req="yes" type="Vm">
10          <alt ann="0">
11            <c end="2" start="2">can</c>
12          </alt>
13        </error>
14        <error id="2" req="yes" type="Mec">
15          <alt ann="0">
16            <c end="8" start="8">'</c>
17          </alt>
18        </error>
19        <error id="3" req="yes" type="Rloc->
20          <alt ann="0">
21            <c end="10" start="9" />
22          </alt>
23        </error>
24      </error-list>
25    </sentence>
26  </script>
27 </scripts>
```

Figure 3.11: An example of conversion of the gold-standard  $m2$  format to the annotation scheme of I-measure

### 3.2.3.1 Method

In the first step, I-measure defines a new scheme for the gold-standard annotations where each sentence is annotated with a set of errors and their possible corrections. Figure 3.2.3 illustrates the equivalent scheme of a gold-standard annotation example from the CoNLL 2014 Shared Task corpus. Each error is detected based on the location. Since all the correction alternatives are mutually exclusive, different mixture of corrections from different annotators lead to a better estimation of the evaluation.

Using dynamic programming implementation of the Sum of Pairs alignment, a globally optimal alignment between the source sentence, the predicted output and the gold-standard annotation is created. Sum of Pairs method is a common method in Bioinformatics that scores a multiple sequence alignment by summing the scores of each pairwise alignment. These scores are based on of three constrains in error correction:

- $c_{match}$ : cost of matches
- $c_{gap} > c_{match}$ : cost of insertions and deletions
- $c_{mismatch} > c_{gap}$ : cost of substations
- $2c_{gap} > c_{mismatch}$ : to ensure that the desired alignment has a lower cost than a initial alignment

Once the optimal alignment found, the evaluation is calculated using an extended version of the Writer-Annotator-System (WAS) introduced in [45] and is shown in Table 3.1. In this table, a, b and c are used to indicate the agreements, e.g., if the source, the hypothesis and the gold-standard annotation are token a, a is a true-negative detection and correction. The cases where  $Source \neq Hypothesis \neq Gold - standard$  belong to false-positive and false negative classes. These cases are specific by FPN class.

Traditionally, NLP-related systems are evaluated using a  $2 \times 2$  contingency table that provides a comparison between system outputs and the gold-standard annotations. Considering the input text in the evaluation, the WAS table is an alternative to represent the  $2 \times 2 \times 2$  possible combinations in correction system evaluation.

Finally, evaluation metrics are calculated using *precision*, *recall* and *F - measure*, defined in section 3.2.1. In order to reward correction more than preservation (i.e., "doing nothing"), a weighted version of accuracy is also defined:

$$WAcc = \frac{w.TP + TN}{w.(TF + FP) + TN + FN - (w + 1) \cdot \frac{FPN}{2}} \quad (3.38)$$

Tokens			Classification	
Source	Hypothesis	Gold-standard	Detection	Correction
a	a	a	TN	TN
a	a	b	FN	FN
a	a	-	FN	FN
a	b	a	FP	FP
a	b	b	TP	TP
a	b	c	TP	FP, FN, FPN
a	b	-	TP	FP, FN, FPN
a	-	a	FP	FP
a	-	b	TP	FP, FN, FPN
a	-	-	TP	TP
-	a	a	TP	TP
-	a	b	TP	FP, FN, FPN
-	a	-	FP	FP
-	-	a	FN	FN

Table 3.1: Extended Writer-Annotator-System evaluation system

where  $w > 1$ .

### 3.2.4 BLEU and GLEU

One of the most widely used automatic evaluation metrics is BLEU score [51]. It is computed as the geometric mean of the modified n-gram *precision*, multiplied by a brevity penalty,  $\rho$ , to control for recall by penalizing short translations:

$$BLEU = \rho \left( \prod_{i=1}^N precision_i \right)^{\frac{1}{N}} \quad (3.39)$$

where  $n$  is the n-gram order that is most often 4, and  $precision_i$  and  $\rho$  are calculated as follow:

$$Precision_i = \frac{\sum_{t_i} \min\{C_h(t_i), \max_j C_{h_j}(t_i)\}}{H(i)} \quad (3.40) \quad \rho = \exp\left\{\min\left(0, \frac{n-L}{n}\right)\right\} \quad (3.41)$$

where  $H(i)$  is the number of  $i$ -gram tuples in the hypothesis,  $t_i$  is an  $i$ -gram tuple in hypothesis  $h$ ,  $C_h(t_i)$  is the number of times  $t_i$  occurs in the hypothesis and  $C_{h_j}(t_i)$  is the number of times  $t_i$  occurs in the gold-standard annotation  $j$  of this hypothesis. In calculating brevity penalty  $\rho$ ,  $n$  refers to the length of the hypothesis and  $L$  refers to the length of the gold-standard annotation.

Recently, Napoles et al. ameliorated BLEU metric for evaluation of grammatical error correction systems and proposed the *Generalized Language Evaluation Understanding* (GLEU) [52, 53]. Similar to the I-measure metric, the precision calculation is modified to assign extra weight to n-grams that are present in the gold-standard annotation and

the system prediction, but not those of the input (set of n-grams  $R \setminus S$ ), and it also penalizes the grams in the candidates that are present in the source but not the gold-standard annotation ( $S \setminus R$ ). This metric is defined as:

$$p_n' = \frac{\sum_{n\text{-gram} \in C} \text{count}_{R \setminus S}(n\text{-gram}) - \lambda(\text{count}_{S \setminus R}(n\text{-gram})) + \text{count}_R(n\text{-gram})}{\sum_{n\text{-gram} \in C} \text{count}_S(n\text{-gram}) + \sum_{n\text{-gram} \in R \setminus S} \text{count}_{R \setminus S}(n\text{-gram})} \quad (3.42)$$

where  $\lambda$  refers to the rate of penalty for incorrectly changed  $n\text{-grams}$ , and function  $\text{count}$  is calculated for a set of n-grams  $B$  as:

$$\text{count}_B(n\text{-gram}) = \sum_{n\text{-gram}' \in B} d(n\text{-gram}, n\text{-gram}') \quad (3.43)$$

where  $d(n\text{-gram}, n\text{-gram}')$  has a unit value in the case that  $n\text{-gram}$  and  $n\text{-gram}'$  are equal. Using these parameters, GLEU( $C, R, S$ ) score for a candidate  $C$  for input  $S$  and reference  $R$  is defined as:

$$BP = \begin{cases} 1 & c > r \\ e^{(1-c)/r} & c \leq r \end{cases} \quad (3.44)$$

$$GLEU(C, R, S) = BP \cdot \exp\left(\sum_{n=1}^N w_n \log p_n'\right) \quad (3.45)$$

Since the task of evaluation in machine correction systems is limited to a comparative method between the prediction results and the gold-standard annotations, some of the evaluation measures in machine translation systems, e.g., manual evaluation for adequacy and fluency, does not appear to be pertinent in the case of machine correction systems. Results are reported in section 6 using these metrics.

---

## Experimental Setup

---

This chapter discusses the experimental procedure used to create our error correction system. First, in section 4.1, we present our task and the pipeline the we followed in developing the project. Section 4.2 introduces the QALB corpus used for our experiments. In this section, the corpus is preprocessed and different details are analyzed. Finally, section 4.3 contains the technical details for the implementation of the project.

### 4.1 Task

Conceptually, the task of creating our error correction system can be classified in three modules:

1. **Modeling:** Given a parallel corpus of potentially incorrect words versus their corrected forms, we define a model by defining the representation of data and the approach of the correction.
2. **Training:** In order to estimate parameters  $\theta$  that most most suite to the corpus, we need to train a method on the training data.
3. **Searching:** In this last step, we solve the selection mechanism  $argmax$  by searching for the hypothesis with the highest probability.

Although these three models are practically addressed in chapter 5, this chapter covers a partial description regarding the QALB corpus that is used in the training and evaluating of the models.

## 4.2 The QALB corpus

Qatar Arabic Language Bank (QALB)<sup>1</sup> is created as a part of a collaborative project between Columbia University and the Carnegie Mellon University Qatar, funded by the Qatar National Research Fund. In our research, we used the release of QALB at ANLP-ACL 2015 which includes data sets of native and non-native (L2 data sets) Arabic speakers. [54, 55].

QALB corpus is provided in three subsets: training, development and test, for which (*\*.train.\**), (*\*.dev.\**) and (*\*.test.\**) are respectively used as the file extension. This is not the case for the non-native corpus(L2) where the training and development sets are only provided. To evaluate different models on the development data, QALB corpus comes with some evaluation scripts that may be useful to evaluate the performance of different models. However, we have not used them in this project.

Three versions of the data are included in the corpus: *\*.m2* is the format accepted by the scorer for evaluation including the gold annotations, *\*.sent* includes plain documents with identifiers, and finally *\*.column* provides the feature files in column format generated using [56].

The QALB corpus is composed of machine translation texts, human user comments, weblog contents, student essays and non-native essays. Their goal in the annotation is to correct all the errors by maintaining the minimum edit actions. Corrected errors in the QALB corpus can be classified in the following six classes:

1. Spelling errors: common typographical errors
2. Punctuation errors
3. Lexical errors: inadequate lexicon
4. Morphological errors: incorrect derivation or inflection
5. Syntactic errors: grammatical errors, e.g., agreement of gender, number, definiteness

---

<sup>1</sup><http://nlp.qatar.cmu.edu/qalb/>

6. Dialectal errors: specific words are detected as incorrect if not present in the Almaany reference dictionary <sup>2</sup>.

### 4.2.1 QALB corpus structure

Using the annotation style of the CoNLL-2013 shared task in [48], the QALB corpus is consists of blocks of annotated phrases. In each block, the first line starting by 'S' is a document token that may encompass a single sentence or a paragraph of different sentences. Following the sentence line, the corrected form of each token appears in lines starting by 'A'. In each line containing a correction, three different pieces of information are important to create our data set:

1. location of each token represented by *beginning* – *ending* identifiers (indexed from 0). This location refers to the spaces between each token
2. correction type needed to correct a given token
3. the corrected form of the token

Each block is separated by an empty line.

```

S_0Tih1s2text3is4meant5tu6be7in8gAra9bic10but11is12it13not14:15
A 0 1|||Edit|||This|||REQUIRED|||-NONE-|||0
A 2 3|||Split|||is meant|||REQUIRED|||-NONE-|||0
A 3 4|||Edit|||to|||REQUIRED|||-NONE-|||0
A 6 8|||Merge|||Arabic|||REQUIRED|||-NONE-|||0
A 9 9|||Add_before|||,|||REQUIRED|||-NONE-|||0
A 9 11|||Move|||it is|||REQUIRED|||-NONE-|||0
A 12 12|||Add_after|||.|||REQUIRED|||-NONE-|||0
A 12 13|||Delete||| |||REQUIRED|||-NONE-|||0

```

Figure 4.1: An example of the correction process based on the structure of the QALB corpus. *S* and *A* refer respectively to the beginning of the sentence and the gold annotators' correction

Figure 4.1 demonstrates a correction annotation based on the QALB corpus for a dummy writer's text. In this example, the phrase is started by *S* and then the correction annotations are followed in *A* lines. The first correction concerning the token that starts from location 0 to location 1, i.e. *Tih<sub>1</sub>s*, is to be corrected to *This* by an *Edit* action. The token in location 2 to 3 needs a *Split* action to be corrected as *is meant* instead of *ismean<sub>2</sub>t*. Following all the *A* lines, these annotations yield the correct phrase "*This text*

<sup>2</sup><http://www.almaany.com/>

is meant to be in Arabic but, it is not.”. Locations are detectable based on the space characters in the phrases, which are enumerated in red in this example. The last index  $A$  lines, i.e.,  $|||0$ , indicate the number of annotators. This may influence the evaluation results depending on the metric. However, for the current version of the QALB corpus, we could not find more than one annotator. An example of the QALB corpus can be found in A.2.

If we consider each action as a function  $\forall x_i, x_{i+n} \text{Correct}(x_i, x_{i+n})$ , we can define the correction actions in the QALB corpus as follows:

- *Add\_before*: insert a token in front of another token ( $n = 0$ ).
- *Merge*: merge multiple tokens ( $n \geq 2$ ). The maximum value that we observed for this action was 4.
- *Split*: split a token into multiple tokens ( $n = 1$ ).
- *Delete*: delete a token (*action* = ”” and  $n = 1$ ).
- *Edit*: replace a token with a different token ( $n = 1$ ).
- *Move*: move a token to a different location in the sentence ( $n \geq 2$ ). The maximum value that we observed for this action was 7.
- *Add\_after*: insert a token after another token. This action is always accompanied by a *Delete* action ( $n = 0$ ).
- *Other*: a complex action that may involve multiple tokens ( $n \geq 1$ ). The maximum value that we observed for this action was 6.

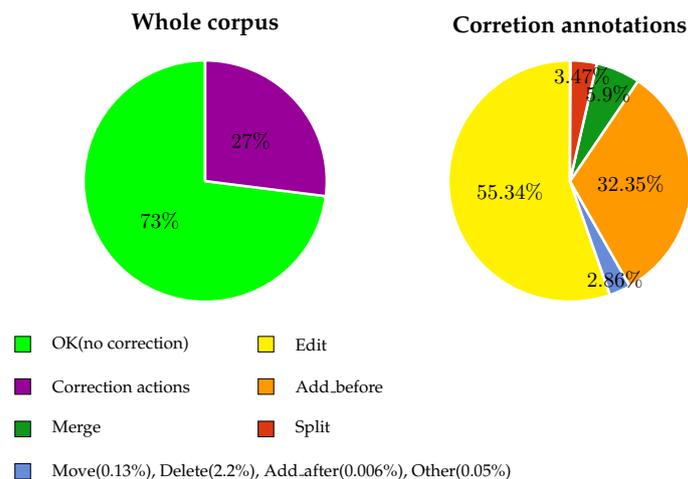


Figure 4.2: Actions used by annotators for correction in the QALB corpus

Figure 4.2 shows the pie charts of the correction actions proportionally. The 73% of the tokens in the QALB corpus that are not annotated are considered correct, i.e., *action* =

OK. The proportion of each action is as follows: *Edit*: 169 769, *Move*: 427, *Add\_before*: 99258, *Merge*: 18267, *Delete*: 6778, *Split*: 10675, *Other*: 1563 and *Add\_after*: 20, which are illustrated in the pie chart on the right.

### 4.2.2 Creating data sets

Based on the level of modeling, we need to manipulate the corpus in order to extract the needed data. For instance, we can create our data set based on three pieces of information: **original tokens** (*S* line), **corrected tokens** and the **correction actions** (both in *A* line). With reference to the example in figure 4.2, the structure of the data set of this block will be:

```

1 [
2   ["Tih's", "This", "Edit"],
3   ["text", "text", "OK"],
4   ["ismeant", "is meant", "Split"],
5   ["tu", "to", "Edit"],
6   ["in", "in", "OK"],
7   ["Ara bic", "Arabic", "Merge"],
8   ["but", "but", "OK"],
9   [",", ",", "Add_before"],
10  ["is it", "it is", "Move"],
11  [".", ".", "Add_after"],
12  [":", ":", "Delete"]
13 ]

```

Listing 4.1: Processing the annotation block of figure 4.2

One of the challenges was the case of having more than one action for a given token. This amounts to 701 tokens requiring 306 757 correction actions. Thus, for juxtaposing actions to the tokens, in the case of a 1 → 1 juxtaposition between an annotation and a token, each action would be assigned as the correction action of the token. However, if more than one correction action is available for a given token, we concatenate all the actions into one action. Since we have not found two different actions for a single token in the current version of QALB, we simply concatenate the corrections keeping the same action.

After analyzing different aspects of the corpus, we created different data sets to train our models. The corpus has a density of 15.80 errors/block and it has 115 characters (including space) from all 1 022 126 tokens in the corpus. We have also considered <



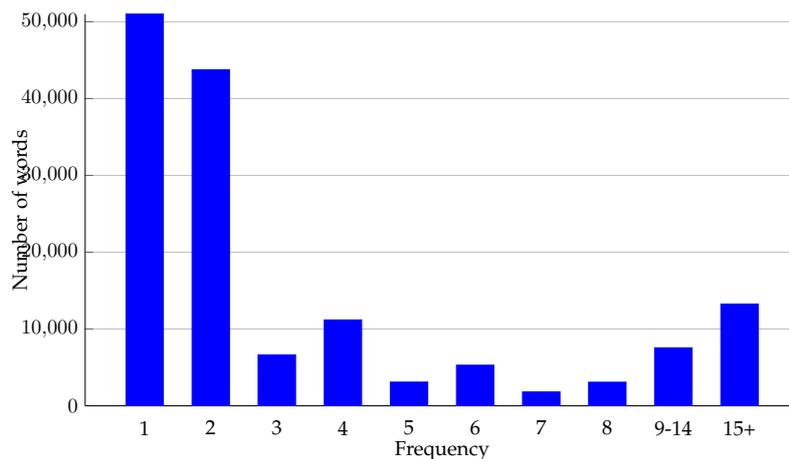


Figure 4.4: Number of words based on frequency in the QALB corpus

works that have dynamic structures. DyNet supports both static and dynamic declaration strategies used in neural networks computations. In the dynamic declaration, each network is built by using a directed and acyclic computation graph that is composed of expressions and parameters that define the model.

Working efficiently on CPU or GPU, DyNet has powered a number of NLP research papers and projects recently.

This chapter will describe how experiments were implemented within this research and what parameters were used in order to carry them out. We also present the training details of our models for the task of error correction. Results of these models are evaluated in chapter 6.

## 5.1 Sentence Representation

The first step in creating a model, is to choose a representation for the input and output. For instance, a sentence can be represented in word-level, which is transferring the data with indices that refer to the words. However, in this work, we anticipate the need to handle an exponentially large input space of tokens by choosing a character-level sentence representation. Thus, model inputs are individual characters which are a sequence of indices corresponding characters. We have defined a specific character  $\langle EOS \rangle$  in order to specify the end of each sentence. This character will be later used in the training process as a criteria. Note that in the character-level representation, we have also considered space as a unique character, i.e. 116 unique characters in all 4.3.

Once a sentence is mapped into indices, the characters are embedded [58], i.e. each character is represented by a one-hot vector and is multiplied by a trainable matrix with the size of the input  $\times$  embeddings size. The embedding allows the model to

group together characters that are similar for the task. These embeddings are fed into the models.

## 5.2 Models

Previously in chapter 3, we reviewed a background of the models that we're using in this research. In this section, we present the technical points of the models that we have used. Having said that we explored different hyper-parameters, here we present those for which the experiments are done.

### 5.2.1 RNN model

The first model that we examined is an RNN as the simplest architecture. By passing the embedding of the input sentence as the input of the network, the corrected form of the input is provided as the output of each state of the network. We use softmax layer to produce probability distribution over output alphabet at each time step. Figure 5.1 illustrates this process.

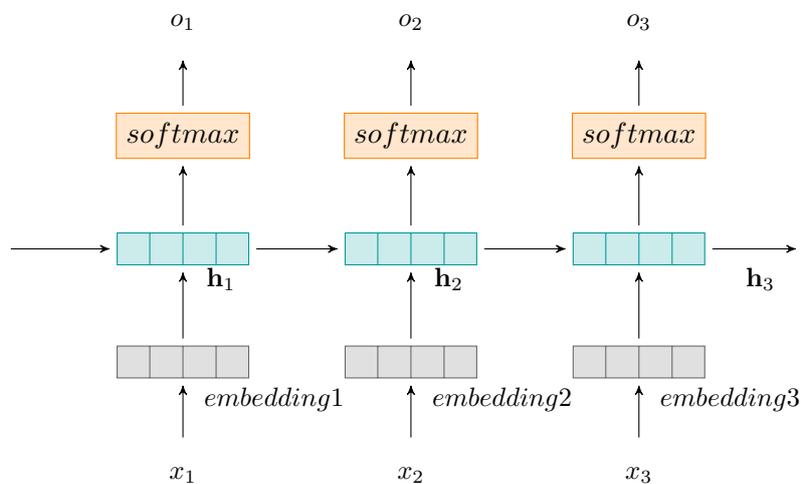


Figure 5.1: A recurrent neural network for error correction. The input is a potentially erroneous sequence. The output of the model is supposed to be the corrected form of the input. Note the fix size of the input and output sequence.

The hyper-parameters of the network are:

- Architecture: LSTM
- Number of layers: 2
- Cells in each layer: 80

- Number of embeddings: 50
- Parameters initialization: [-0.1, 0.1]
- Number of epochs: 20
- Learning rate: Default value of `SimpleSGDTrainer` in `DyNet`, i.e.,  $\eta_t = \frac{\eta_0}{1+\eta_{decay}t}$  at epoch  $t$ .

We train the model by minimizing the negative log likelihood of the training data using stochastic gradient. The codes of the error calculation for each model is available in the appendix [A.3](#).

### 5.2.2 BRNN model

The structure of our BRNN model is exactly the same with the RNN model. The only difference is that a bidirectional RNN consists of two unidirectional RNNs, one reading the inputs in forward direction and the other in backward direction. So, a sum of forward and backward RNN outputs is the output of bidirectional RNN. The hyper-parameters of our model is the following:

- Architecture: LSTM
- Number of layers: 4
- Cells in each layer: 90
- Number of embeddings: 50
- Parameters initialization: [-0.1, 0.1]
- Number of epochs: 20
- Learning rate: Default value of `SimpleSGDTrainer` in `DyNet`, i.e.,  $\eta_t = \frac{\eta_0}{1+\eta_{decay}t}$  at epoch  $t$ .

BRNN in addition to the previous observations, can also take next observations into account. This enables it to be more efficient in dealing with all kinds of errors.

### 5.2.3 Encoder-decoder model

Using two RNN models, we created an encoder-decoder model with the following hyper-parameters:

- Architecture: LSTM
- Number of encoder layers: 4

- Number of decoder layers: 4
- Cells in each layer of the encoder: 100
- Cells in each layer of the decoder: 100
- Number of embeddings: 50
- Parameters initialization: [-0.1, 0.1]
- Number of epochs: 10
- Learning rate: Default value of `SimpleSGDTrainer` in `DyNet`, i.e.,  $\eta_t = \frac{\eta_0}{1+\eta_{decay}t}$  at epoch  $t$ .

#### 5.2.4 Attention-based encoder-decoder

We added an attention mechanism between the encoder and the decoder model. This mechanism helps the decoder to pick only the encoded inputs that are important for each step of the decoding process. Once we calculate the importance of each encoded vector, we normalize the vectors using softmax and multiply each encoded vector by its weight to obtain a "time dependent" input encoding which is fed to each step of the decoder RNN. The only difference of this model with the encoder-decoder model in terms of hyper-parameters, is the number of epochs. We trained the model in 7 epochs.

### 5.3 Output generation

For an error correction system, sentences with spelling or grammatical errors should have lower probability than their corrected versions, since they are less present in the whole corpus. Based on this assumption, as we explained in the description of the models, an output is the most probable element in the softmax probability distribution (demonstrated in algorithm 1). Although, this is not the only approach to generate an output. Dahlmeier and Ng [59] developed a *beam-search* decoder to iteratively generate sentence-level candidates and ranking them.

In the case of encoder-decoder and attention-based encoder-decoder, since the output size can be variant, we have used two different conditions. Algorithm 2 shows the generation algorithm in an encoder-decoder model or an attention-based encoder-decoder model:

---

**Algorithm 1** Correction of an input sequence using a character-level model(RNN or BRNN)

---

**Input:** Input sequence  $C \leftarrow c_1, c_2, \dots, c_N$   
**Output:** Corrected sequence  $O \leftarrow o_1, o_2, \dots, o_N$

- 1: **procedure** GENERATEOUTPUT( $C$ )
- 2:    $embeddedInput \leftarrow embedSequence(C)$
- 3:    $modelOutputs \leftarrow runModel(embeddedInput)$
- 4:    $p_0, p_1, \dots, p_N \leftarrow getProbabilities(modelOutputs)$
- 5:    $index \leftarrow 0$
- 6:   **while**  $i < N$  **do**
- 7:      $predictedIndex \leftarrow \operatorname{argmax}(p_i)$
- 8:      $o_i \leftarrow indexToCharacter(predictedIndex)$
- 9:      $i \leftarrow i + 1$
- 10:   **return**  $O$

---



---

**Algorithm 2** Correction of an input sequence using a character-level model(encoder-decoder or attention-based encoder-decoder)

---

**Input:** Input sequence  $C \leftarrow c_1, c_2, \dots, c_N$   
**Output:** Corrected sequence  $O \leftarrow o_1, o_2, \dots, o_M$

- 1: **procedure** GENERATEOUTPUT( $C$ )
- 2:    $embeddedInput \leftarrow embedSequence(C)$
- 3:    $modelOutputs \leftarrow runModel(embeddedInput)$
- 4:    $p_0, p_1, \dots, p_N \leftarrow getProbabilities(modelOutputs)$
- 5:    $index \leftarrow 0$
- 6:    $MaxSize \leftarrow N \times 2$
- 7:   **while**  $i < MaxSize$  **do**
- 8:      $predictedIndex \leftarrow \operatorname{argmax}(p_i)$
- 9:      $o_i \leftarrow indexToCharacter(predictedIndex)$
- 10:      $i \leftarrow i + 1$
- 11:     **if**  $o_i = "$  < EOS > " **then**
- 12:       **break**
- 13:   **return**  $O$

---

---

## Experiments and Results

---

In the previous chapter, we presented four models for the task of error correction. The details of the models were explained. In this chapter, the results obtained from the models are discussed. First we analyze the data set and the error rate of the gold-standard corpus. The evaluation metrics are then applied on the results of the models. As we explained in section 6, using different metrics can reflect a better interpretation of the evaluation of the results.

### 6.1 Data

A sequence in input or output of a error correction system can be represented in different levels. In character-level, a sequence is processed character by character. When searching for errors, humans often consider a bigger sequence of characters at word-level. Clause-level, phrase-level, sentence-level and text-level are other common representations for modeling sequences.

In our research, we worked at character-level where each character in an input sequence is mapped to an real-valued number and then it is embedded. In order to model linguistic dependencies in each sequence, we took every character into account, including space. This enables us to deal with different kinds of errors and a larger range of characters in each sequence. On the other hand, the output of the models are also at

character level. We will later see that the fact that we train the methods using the whole sentence size reduces the accuracy of the models remarkably.

	#sentences	#words	#characters	Corrected tokens
train set	19411	1041537	10231305	29.45%
validation set	1017	54880	537182	30.35%
test set	968	52290	510582	31.32%

Table 6.1: Train, validation and test set of the QALB corpus.

Three data sets are provided in the QALB corpus:

- Train set is used to build up our prediction model. Our algorithm tries to tune itself to find optimal parameters with the back-propagation. This set is usually used to create multiple algorithms in order to compare their performances during the validation phase.
- Validation set is used to compare the performances of the prediction parameters created based on the training set. We select the parameters that has the best performance.
- Test set is a set of examples used only to assess the performance of a fully-trained classifier. In the RNN case, we would use the test set to estimate the error rate after we have chosen the final model.

Table 6.1 summarizes each data set in details. We used these data sets with the given sizes.

## 6.2 Baseline system

We define the pair of source sentences and the gold-standard annotations as the baseline of the models. In this baseline, we assume that non of our implemented models intervene in the task of correction and only references are considered as correction. Simply saying, Baseline system is the system that makes no corrections on the input text. The baseline enables us to interpret the performance of each model in comparison to the default results.

## 6.3 Results of the models

This section presents evaluation results of the models using the metrics that we introduced in section 3.2. Since each of these metrics reflect a different aspect of evaluation of each model, comparing them together enables us to obtain a better description of the performance of the models.

### 6.3.1 MaxMatch $M^2$

As explained in detail in section 3.2.2, the MaxMatch  $M_2$  metric computes the sequence of phrase-level edits between the source text and the predicted text in a way that it maximizes their overlap. The edits are then scored using precision  $P$ , recall  $R$  and F-score  $F_{0.5}$ .

Model	$M^2$ scorer		
	P	R	$F_{0.5}$
Baseline	1.0000	0.0000	0.0000
RNN	0.5397	0.2487	0.4373
BiRNN	0.5544	0.2943	0.4711
Encoder-decoder	<b>0.5835</b>	<b>0.3249</b>	<b>0.5034</b>
Attention	0.5132	0.2132	0.4155

Table 6.2: Evaluation results of the models using MaxMatch  $M^2$  metric. Bold numbers indicate the best method per measure.

Table 6.2 demonstrates the results of the models using MaxMatch  $M^2$  metric. In the baseline system, the metric detects complete precision of the system. This is because in the baseline system the gold-standard annotations are directly applied as the hypothesis of a model. So, we can justify the zero value of the recall and the  $F_{0.5}$  consequently. However, these scores are not as high as the baseline system for other models.

We set the value of the transitive edits in the range of 2 for calculating the corresponding lattice of each edit. We set also the value of  $\beta = 2$  in calculating the F-score.

In addition to the `m2` format of the gold-standard annotations which is specifically used for MaxMatch  $M^2$  metric, the reference correction of the QALB corpus is provided also in raw text with `.sent` format. Thus, we could evaluate the models without taking any possible preprocessing noises into account. For instance, evaluating the baseline model, we could have the following results  $P = 0.0179$ ,  $R = 0.0004$ ,  $F_{0.5} = 0.0017$ . Even if these values are not remarkably different from those of the reference correction, it indicate

probable mismatching during the preprocessing of the .m2 files of the corpus.

We have used the original implementation of the MaxMatch  $M^2$  introduced in [47].<sup>1</sup>

### 6.3.2 I-measure

I-measure metric is a method in response to the limitations of the MaxMatch  $M^2$  metric. As we explained in section 3.2.3, the I-measure determines an optimal alignment between an input, a hypothesis and a gold-standard text. In addition to the standard metrics introduced in section 6, in this metric an Improvement (I) score is computed by comparing system performance with that of the baseline which leaves the original text uncorrected. Table 6.3 demonstrates the evaluation results of the trained models using I-measure metric.

	Baseline		RNN		BRNN		Encoder-decoder		Attention	
	Detection	Correction	Detection	Correction	Detection	Correction	Detection	Correction	Detection	Correction
TP	0	0	11286	208	11800	223	12213	247	11557	201
TN	39589	39589	20881	20881	20378	20378	22204	22204	21070	21070
FP	0	0	18730	29808	19220	30797	17410	29376	18533	29889
FN	98318	98318	87031	98109	86518	98095	86111	98077	86758	98114
FPN	0	0	0	11078	0	11577	0	11966	0	11356
P	100.00	100.00	37.60	0.69	38.04	0.72	<b>41.23</b>	<b>0.83</b>	38.41	0.67
R	0.00	0.00	11.48	0.21	12.00	0.23	12.42	0.25	11.76	0.20
$F_{0.5}$	0.00	0.00	25.84	0.48	26.53	0.50	28.16	<b>0.57</b>	26.43	0.46
Acc	28.71	28.71	23.32	15.29	23.33	14.94	24.95	16.28	23.66	15.42
$Acc_b$	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71
WAcc	28.71	28.71	25.87	13.11	26.03	12.76	27.83	<b>14.05</b>	26.30	13.23
$WAcc_b$	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71	28.71
I	0.00	0.00	-9.87	-54.32	-9.32	-55.54	-3.06	<b>-51.07</b>	-8.39	-53.92

Table 6.3: Evaluation of the models using I-measure metric

The I-measure evaluates a prediction in terms detection and correction. However, we are interested only in the correction results, the detection results can also be informative about the quality of correction of a model. Similar to the evaluation results of the MaxMatch  $M^2$  metric in table 6.2, the baseline systems is completely precise in terms of correction, i.e.,  $F_{0.5} = 0.00$ ,  $P = 100.00$ . On the other hand, other models could not obtain competitive results. This is mainly because of the poor alignment between the source text, the hypothesis and the reference. In comparison to other models, the encoder-decoder demonstrates better performance with an Improvement score of  $I = -51.07$  and a weighted accuracy of  $WAcc = 14.05$ . The weighted version of accuracy  $WAcc$  rewards correction more than preservation.

<sup>1</sup>Available at <http://nlp.comp.nus.edu.sg/software/>

We set the values of  $W Acc$  and  $F-\beta$  for this metric to 2 and 0.5 respectively. We used the original implementation of the I-measure introduced in [50].<sup>2</sup>

### 6.3.3 BLEU and GLEU

BLEU was one of the first automatic metrics used in measuring translation accuracy and has become one of most common metrics in machine translation systems evaluation [51]. In addition to this metric, we also use GLEU which is a simple variant of BLEU showing a better correlation with human judgments in the evaluation task [52]. The following table shows the evaluation results of our correction models using BLEU and GLEU metrics.

Model	BLEU score	GLEU score
Baseline	1.000000	1.000000
RNN	0.349212	0.313035
BiRNN	0.305852	0.269368
Encoder-decoder	0.327858	0.292956
Attention	0.312947	0.287664

Table 6.4: Evaluation of the models using BLEU and GLEU metrics

As we expect, since the baseline system contains the gold-standard correction, the BLEU and the GLEU scores for the baseline system have the maximum value 1.00. Using metric, the RNN model shows higher scores in comparison to other models, i.e.,  $BLEU = 0.3492$  and  $GLEU = 0.3130$ . Note that in the GLEU metric, the precision is modified to assign extra weight to the n-grams that are present in the reference and the hypothesis, but not those of the input.

We have used the last update of the original implementation of the GLEU [52] introduced in [53].<sup>3</sup>

<sup>2</sup>Available at <https://github.com/mfelice/imeasure/>

<sup>3</sup>Available at <https://github.com/cnap/gec-ranking>

---

## Conclusion and future work

---

Thus far, the results of each method are presented in depth. In this last chapter, these results will be discussed in more detail based on the research aims. Furthermore, several ideas about future works in this domain are proposed.

### 7.1 Limitations

A comparative example of a source text, its gold-standard correction and output of trained model is illustrated in appendix A.2. We tagged manually four kinds of errors in different colors: the green tags refer to the incorrect tokens in the input which are predicted correctly by a trained model, the yellow tags are the incorrect input words, the orange tags are the incorrectly predicted tokens of an incorrect input word and finally the cyan tags show the incorrect prediction of a correct input token.

The purpose of using these tags is to demonstrate different kinds of correction and the quality of correction by each model for the same symbol. For instance, the incorrect form of "لا زال" is corrected by none of the models as "لا زال", while "الشبيحه" is predicted correctly by the RNN, BRNN and encoder-decoder models as "الشبيحة". Looking carefully at the distribution of incorrect prediction of correct input words (colored in cyan), one can deduce that the models perform less sensibly when the size the sequence become gradually bigger. To prove our observation, we tried to evaluate the models by limiting

the sequences to a fixed size.

	30		50		70		100	
	BLEU	GLEU	BLEU	GLEU	BLEU	GLEU	BLEU	GLEU
Baseline	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
RNN	0.575715	0.476466	<b>0.579784</b>	0.545291	0.556254	0.543040	0.519782	0.514380
BRNN	0.481971	0.397735	0.501957	0.469293	0.49238	0.480605	0.461176	0.455698
Encoder-decoder	0.471226	0.421491	0.488849	0.457512	0.479811	0.468159	0.456483	0.451533
Attention	0.506216	0.389947	0.507188	0.477017	0.490448	0.478636	0.457459	0.452876

Table 7.1: Evaluation of models using fixed-size sequences

Table 7.1 demonstrates the BLEU and GLEU scores for a limited size of sequences of 30, 50, 70 and 100 characters. In comparison to the results of the models in table 6.4, this limitation shows considerably higher scores.

We could not do the same experiment using MaxMatch nor I-measure methods. MaxMatch  $M^2$  and I-measure need the gold-standard correction annotations for the evaluation, while BLEU and GLEU evaluate the models using the gold-standard raw text, without any specific information about the annotation.

## 7.2 Future studies

For the future studies, we have the following suggestions:

- Models to be explored in more levels, e.g., action-level, word-level and sentence-level.
- Limiting the length of the sequences in training models.
- Using deeper networks with larger embedding size.
- Preventing over-learning of models by not training them over correct input tokens (action = "OK").

## 7.3 Conclusion

The basic problem that the attention mechanism solves is that instead of forcing the network to encode all parameters into one fixed-length vector, it allows the network to take use of the input sequence. In other words, the attention mechanism gives the network

access to the hidden state of the encoder which acts like an internal memory. Thus, the network is able to retrieve from the memory the related parameters for the current input.

There is a tight relation between the attention mechanism and the memory mechanisms. For instance, the hidden state of a simple RNN by itself is a type of internal memory. However, the RNNs suffer from the vanishing gradient problem that hinders modeling long distance dependencies, which is a common phenomenon in human language. On the other hand, LSTM uses a gating mechanism that provides an explicit memory for deletions and updates.

Recently more researchers tend towards more complex memory structures. For instance, End-to-End Memory Networks [60] allow the network to read multiple times an input sequence in order to make an output and also to update the memory contents at each step. These models are used in question answering [61] and to language modeling successfully, since it is able to make multiple computational steps, also known as hop, over an input story.

Future researches may reveal a clearer distinction between attention and memory mechanisms. In the recent studies, reinforcement learning techniques have been also used for error correction task [62, 63].

## A.1 Adam

Adaptive Moment Estimation (Adam) [64] is a stochastic optimization method that computes learning rates for each parameter. In addition to storing an exponentially decaying average of past squared gradients  $v_t$ , Adam also maintains an exponentially decaying average of past gradients  $m_t$ . The step taken at each update is proportional to the ratio of these two moments for which  $\beta_1$  and  $\beta_2$  are respectively used as the exponential decay rates. We can define each parameter as follow:

$$g_t = \nabla H_t(\theta_t - 1) \tag{A.1}$$

$$m_t = \frac{\beta_1 m_{t-1} + (1 - \beta_1) g_t}{1 - \beta_1^2} \tag{A.2}$$

$$v_t = \frac{\beta_2 v_{t-1} + (1 - \beta_2) g_t^2}{1 - \beta_2^2} \tag{A.3}$$

$$\theta_t = \theta_{t_1} - \alpha \cdot \frac{m_t}{\sqrt{v_t} + \epsilon} \tag{A.4}$$

where  $\epsilon$  is a very small number to prevent division by zero.

## A.2 Annotation example from the QALB corpus

S	كل ماء يجري من تخريب وتفجير وقتل وسلب ونهب في هذه الايام تم الاعداد له من سابق من قبل القاعده الاحمريه القليله لكى تثير الفتن . وهذا مخطط على وحاشيته من اجل البقاء في استعمار الجنوب . واما انصار الشريعة هي تلك القوات التي درية من قبل امرياء وهم من يحول البلاد الى امارات اسلاميه . والمستفيد من ذلك رئيس العربية اليمانية الشقيقه . ثورة ثورة يا جنوب
A 1 2	Split ما يجري REQUIRED NONE 0
A 4 5	Split وتفجير وقتل REQUIRED NONE 0
A 9 10	Edit الايام REQUIRED NONE 0
A 11 12	Split الإعداد له REQUIRED NONE 0
A 13 14	Edit السابق REQUIRED NONE 0
A 16 17	Edit القاعدة REQUIRED NONE 0
A 17 18	Edit الأحمريه REQUIRED NONE 0
A 18 19	Edit القليله REQUIRED NONE 0
A 20 21	Split تثير الفتن REQUIRED NONE 0
A 22 23	Split وهذا مخطط REQUIRED NONE 0
A 24 25	Edit وحاشيته REQUIRED NONE 0
A 26 27	Edit أجل REQUIRED NONE 0
A 27 28	Split البقاء في REQUIRED NONE 0
A 28 29	Split استعمار الجنوب REQUIRED NONE 0
A 30 31	Other أوأما أنصار الشريعة REQUIRED NONE 0
A 34 35	Edit التي REQUIRED NONE 0
A 35 36	Edit دريتها REQUIRED NONE 0
A 36 39	Move أمريكا من قبل REQUIRED NONE 0
A 38 39	Move وهم REQUIRED NONE 0
A 41 42	Split البلاد إلى REQUIRED NONE 0
A 42 43	Edit إمارات REQUIRED NONE 0
A 43 44	Edit إسلامية REQUIRED NONE 0
A 44 45	Split والمستفيد من REQUIRED NONE 0
A 45 46	Edit ذلك REQUIRED NONE 0
A 46 47	Edit رئيس REQUIRED NONE 0
A 49 50	Edit الشقيقة REQUIRED NONE 0
A 53 54	Split يا جنوب REQUIRED NONE 0

Figure A.1: An annotation example from the QALB corpus. Following line *S* which refers to the potentially wrong original text, the correction actions are added for each incorrect token in *A* lines.

### A.3 Error calculation

```

1 def get_error(input_sequence, output_sequence):
2     error = list()
3     dy.renew_cg()
4     embedded_sequence = embed_string(input_sequence)
5
6     rnn_state = network.initial_state()
7     rnn_outputs = rnn_run(rnn_state, embedded_sequence)
8
9     for rnn_output, output_char in zip(rnn_outputs, output_sequence):
10        w_out = dy.parameter(network.w_out)
11        b_out = dy.parameter(network.b_out)
12        probabilities = dy.softmax(output_w * rnn_output + output_b)
13        error.append(-dy.log(dy.pick(probabilities, output_char)))
14    error = dy.esum(error)
15    return error

```

Listing A.1: Calculating error in training the RNN model

```

1 def get_error(input_sequence, output_sequence):
2     error = list()
3     dy.renew_cg()
4     embedded_sequence = embed_string(input_sequence)
5
6     rnn_bwd_state = bwdRNN.initial_state()
7     rnn_bwd_outputs = run_rnn(rnn_bwd_state, embedded_sequence[::-1])[:, -1]
8     rnn_fwd_state = fwdRNN.initial_state()
9     rnn_fwd_outputs = run_rnn(rnn_fwd_state, embedded_sequence)
10    rnn_outputs = [dy.concatenate([fwd_out, bwd_out]) for fwd_out, bwd_out in
11                  zip(rnn_fwd_outputs, rnn_bwd_outputs)]
12
13    for rnn_output, output_char in zip(rnn_outputs, output_sequence):
14        w_out = dy.parameter(network.w_out)
15        b_out = dy.parameter(network.b_out)
16        probabilities = dy.softmax(output_w * rnn_output + output_b)
17        error.append(-dy.log(dy.pick(probabilities, output_char)))
18    error = dy.esum(error)
19    return error

```

Listing A.2: Calculating error in training the BRNN model

```

1 def get_error(input_sequence, output_sequence):
2     error = list()
3     dy.renew_cg()
4     embedded_sequence = embed_string(input_sequence)
5     encoded_string = encode_string(embedded_sequence)[-1]
6     rnn_state = decoder.initial_state()
7
8     for output_char in output_sequence:
9         rnn_state = rnn_state.add_input(encoded_string)
10        w_out = dy.parameter(network.w_out)
11        b_out = dy.parameter(network.b_out)
12        probabilities = dy.softmax(output_w * rnn_state.output() + output_b)
13        error.append(-dy.log(dy.pick(probabilities, output_char)))
14    error = dy.esum(error)
15    return error

```

Listing A.3: Calculating error in training the encoder-decoder model

```

1 def get_error(self, input_sequence, output_sequence):
2     error = list()
3     dy.renew_cg()
4     embedded_sequence = embed_string(input_sequence)
5     encoded_string = encode_string(embedded_sequence)
6     rnn_state = decoder.initial_state().add_input(dy.vecInput(self.
7     encoder_state_size))
8
9     for output_char in output_sequence:
10        attended_encoding = self._attend(encoded_string, rnn_state)
11        rnn_state = rnn_state.add_input(attended_encoding)
12        w_out = dy.parameter(network.w_out)
13        b_out = dy.parameter(network.b_out)
14        probabilities = dy.softmax(output_w * rnn_state.output() + output_b)
15        error.append(-dy.log(dy.pick(probabilities, output_char)))
16    error = dy.esum(error)
17    return error

```

Listing A.4: Calculating error in training the attention-based encoder-decoder model

## A.4 Qualitative comparison of the correction models

Source sentence	306 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفه من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوه ، أننا يا أحبائي ندفع ثمن أكثر من أربعين عام ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات</p>	
Gold standard reference	309 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفة من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوة . إننا يا أحبائي ندفع ثمن أكثر من أربعين عاما ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات .</p>	
Recurrent neural network model prediction	307 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفه من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوه ، أننا يا أحبائي ندفع ثمن أكثر من أربعين عام ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات</p>	
Bidirectional recurrent neural network prediction	307 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفه من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوه ، أننا يا أحبائي ندفع ثمن أكثر من أربعين عام ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات</p>	
Sequence-to-sequence model prediction	306 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفه من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوه ، أننا يا أحبائي ندفع ثمن أكثر من أربعين عام ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات</p>	
Attention-based sequence-to-sequence model prediction	309 characters
<p>لا زال كبير الشبيحة يظن أن أرواح وآلام الناس أقل كلفه من تخليه عن منصبه ، فلذلك إذا كان السوريون لا يرتضون بهذه المعادلة المهينة ، فعليهم أن يهبوا هبة قوية واحدة ويأخذوا حقوقهم من هذه العصابة عنوه ، أننا يا أحبائي ندفع ثمن أكثر من أربعين عام ، ومن الخنوع ، والذل والشن سيكون غالبا ولكنه يستأهل هذه التضحيات</p>	

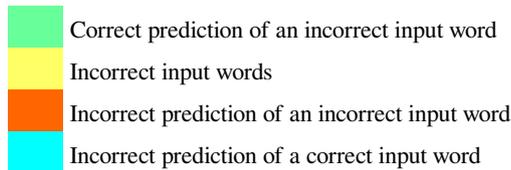


Figure A.2: Correction of an input using different models. The size of each sequence is marked in the header of each sequence.

---

## Bibliography

---

- [1] Julian R. Ullmann. A binary n-gram technique for automatic correction of substitution, deletion, insertion and reversal errors in words. *The Computer Journal*, 20(2):141–147, 1977.
- [2] Andrew McCallum, Kedar Bellare, and Fernando Pereira. A conditional random field for discriminatively-trained finite-state string edit distance. *arXiv preprint arXiv:1207.1406*, 2012.
- [3] Karen Kukich. Techniques for automatically correcting words in text. *ACM Computing Surveys (CSUR)*, 24(4):377–439, 1992.
- [4] Richard W Hamming. Error detecting and error correcting codes. *Bell Labs Technical Journal*, 29(2):147–160, 1950.
- [5] William E Winkler. String comparator metrics and enhanced decision rules in the fellegi-sunter model of record linkage. 1990.
- [6] Robert A Wagner and Michael J Fischer. The string-to-string correction problem. *Journal of the ACM (JACM)*, 21(1):168–173, 1974.
- [7] Fred J. Damerau. A technique for computer detection and correction of spelling errors. *Commun. ACM*, 7(3):171–176, March 1964.
- [8] Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals.

- [9] Leon Davidson. Retrieval of misspelled names in an airlines passenger record system. *Communications of the ACM*, 5(3):169–171, 1962.
- [10] Lawrence Philips. Hanging on the metaphone. *Computer Language Magazine*, 7(12):39–44, December 1990. Accessible at <http://www.cuj.com/documents/s=8038/cuj0006philips/>.
- [11] Joseph J Pollock and Antonio Zamora. Collection and characterization of spelling errors in scientific and scholarly text. *Journal of the Association for Information Science and Technology*, 34(1):51–58, 1983.
- [12] Clifton Phua, Vincent Lee, and K Smith-Miles. The personal name problem and a recommended data mining solution. *Encyclopedia of Data Warehousing and Mining*, 2006.
- [13] Emmanuel J Yannakoudakis and David Fawthrop. The rules of spelling errors. *Information Processing & Management*, 19(2):87–99, 1983.
- [14] Linda G Means. Cn yur cmputr raed ths? In *Proceedings of the second conference on Applied natural language processing*, pages 93–100. Association for Computational Linguistics, 1988.
- [15] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2 edition, 2003.
- [16] Christopher D. Manning and Hinrich Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, USA, 1999.
- [17] Peter F Brown, Vincent J Della Pietra, Stephen A Della Pietra, and Robert L Mercer. The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):263–311, 1993.
- [18] Ronald Rosenfeld. A maximum entropy approach to adaptive statistical language modeling. In *Computer, Speech and Language*, 1996.
- [19] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.

- [20] Claudio Lottaz, Christian Iseli, C Victor Jongeneel, and Philipp Bucher. Modeling sequencing errors by combining hidden markov models. *Bioinformatics*, 19(suppl.2):ii103–ii112, 2003.
- [21] Andrej Karpathy, Justin Johnson, and Fei-Fei Li. Visualizing and understanding recurrent networks. *CoRR*, abs/1506.02078, 2015.
- [22] Nal Kalchbrenner and Phil Blunsom. Recurrent convolutional neural networks for discourse compositionality. *CoRR*, abs/1306.3584, 2013.
- [23] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [24] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- [25] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [26] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [27] Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. On using very large target vocabulary for neural machine translation. *arXiv preprint arXiv:1412.2007*, 2014.
- [28] Noura Farra, Nadi Tomeh, Alla Rozovskaya, and Nizar Habash. Generalized character-level spelling error correction. In *ACL*, 2014.
- [29] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of mathematical biology*, 52(1):99–115, 1990.
- [30] Mario Costa. Probabilistic interpretation of feedforward network outputs, with relationships to statistical prediction of ordinal quantities. *International journal of neural systems*, 7(05):627–637, 1996.

- [31] Razvan Pascanu, Çağlar Gülçehre, Kyunghyun Cho, and Yoshua Bengio. How to construct deep recurrent neural networks. *CoRR*, abs/1312.6026, 2013.
- [32] Eric H Lenneberg. The biological foundations of language. *Hospital Practice*, 2(12):59–67, 1967.
- [33] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997.
- [34] Alex Graves. Generating sequences with recurrent neural networks. *CoRR*, abs/1308.0850, 2013.
- [35] Alex Graves, Santiago Fernández, and Jürgen Schmidhuber. Multi-dimensional recurrent neural networks. *CoRR*, abs/0705.2011, 2007.
- [36] Justin Bayer, Daan Wierstra, Julian Togelius, and Jürgen Schmidhuber. Evolving memory cell structures for sequence learning. *Artificial Neural Networks–ICANN 2009*, pages 755–764, 2009.
- [37] Alex Graves. *Supervised Sequence Labelling with Recurrent Neural Networks*. PhD thesis, Citeseer, 2008.
- [38] Alex Graves and Jürgen Schmidhuber. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5):602–610, 2005.
- [39] Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 2016.
- [40] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [41] Michael Schuster. On supervised learning from sequential data with applications for speech recognition. 04 1999.
- [42] Toshiaki Fukada, Mike Schuster, and Yoshinori Sagisaka. Phoneme boundary estimation using bidirectional recurrent neural networks and its applications. *Systems and Computers in Japan*, 30(4):20–30, 1999.

- [43] Abdel-rahman Mohamed, Frank Seide, Dong Yu, Jasha Droppo, Andreas Stoicke, Geoffrey Zweig, and Gerald Penn. Deep bi-directional recurrent networks over spectral windows. In *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*, pages 78–83. IEEE, 2015.
- [44] Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. Hybrid speech recognition with deep bidirectional lstm. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, pages 273–278. IEEE, 2013.
- [45] Martin Chodorow, Markus Dickinson, Ross Israel, and Joel R Tetreault. Problems in evaluating grammatical error detection systems. In *COLING*, pages 611–628, 2012.
- [46] Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. The conll-2014 shared task on grammatical error correction. 2014.
- [47] Daniel Dahlmeier and Hwee Tou Ng. Better evaluation for grammatical error correction. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572. Association for Computational Linguistics, 2012.
- [48] Ting-Hui Kao, Yu-Wei Chang, Hsun-Wen Chiu, Tzu-Hsi Yen, Joanne Boisson, Jian-Cheng Wu, and Jason S Chang. Conll-2013 shared task: Grammatical error correction nthu system description. In *CoNLL Shared Task*, pages 20–25, 2013.
- [49] Behrang Mohit, Alla Rozovskaya, Nizar Habash, Wajdi Zaghouni, and Ossama Obeid. The first qalb shared task on automatic text correction for arabic. 2014.
- [50] Mariano Felice and Ted Briscoe. Towards a standard evaluation method for grammatical error detection and correction. In *HLT-NAACL*, pages 578–587, 2015.
- [51] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics, 2002.
- [52] Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. Ground truth for grammatical error correction metrics. In *Proceedings of the 53rd Annual Meeting of*

- the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, volume 2, pages 588–593, 2015.
- [53] Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. GLEU without tuning. *eprint arXiv:1605.02592 [cs.CL]*, 2016.
- [54] Wajdi Zaghouani, Behrang Mohit, Nizar Habash, Ossama Obeid, Nadi Tomeh, Alla Rozovskaya, Noura Farra, Sarah Alkuhlani, and Kemal Oflazer. Large scale arabic error annotation: Guidelines and framework.
- [55] Ossama Obeid, Wajdi Zaghouani, Behrang Mohit, Nizar Habash, Kemal Oflazer, and Nadi Tomeh. A web-based annotation framework for large-scale text correction. In *IJCNLP*, pages 1–4, 2013.
- [56] Arfath Pasha, Mohamed Al-Badrashiny, Mona T Diab, Ahmed El Kholy, Ramy Eskander, Nizar Habash, Manoj Pooleery, Owen Rambow, and Ryan Roth. Madamira: A fast, comprehensive tool for morphological analysis and disambiguation of arabic. In *LREC*, volume 14, pages 1094–1101, 2014.
- [57] Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*, 2017.
- [58] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- [59] Daniel Dahlmeier and Hwee Tou Ng. A beam-search decoder for grammatical error correction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 568–578. Association for Computational Linguistics, 2012.
- [60] Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. Weakly supervised memory networks. *CoRR*, abs/1503.08895, 2015.
- [61] Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*, 2015.

- 
- [62] Keisuke Sakaguchi, Matt Post, and Benjamin Van Durme. Grammatical error correction with neural reinforcement learning. *CoRR*, abs/1707.00299, 2017.
- [63] Wojciech Zaremba and Ilya Sutskever. Reinforcement learning neural turing machines. *CoRR*, abs/1505.00521, 2015.
- [64] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.