

A multi-cascaded model with data augmentation for enhanced paraphrase detection in short texts



Muhammad Haroon Shakeel*, Asim Karim, Imdadullah Khan

Department of Computer Science, Syed Babar Ali School of Science and Engineering, Lahore University of Management Sciences, Lahore, Pakistan

ARTICLE INFO

Keywords:

Paraphrase detection
Deep learning
Data augmentation
Sentence similarity

ABSTRACT

Paraphrase detection is an important task in text analytics with numerous applications such as plagiarism detection, duplicate question identification, and enhanced customer support help-desks. Deep models have been proposed for representing and classifying paraphrases. These models, however, require large quantities of human-labeled data, which is expensive to obtain. In this work, we present a data augmentation strategy and a multi-cascaded model for improved paraphrase detection in short texts. Our data augmentation strategy considers the notions of paraphrases and non-paraphrases as binary relations over the set of texts. Subsequently, it uses graph theoretic concepts to efficiently generate additional paraphrase and non-paraphrase pairs in a sound manner. Our multi-cascaded model employs three supervised feature learners (cascades) based on CNN and LSTM networks with and without soft-attention. The learned features, together with hand-crafted linguistic features, are then forwarded to a discriminator network for final classification. Our model is both wide and deep and provides greater robustness across clean and noisy short texts. We evaluate our approach on three benchmark datasets and show that it produces a comparable or state-of-the-art performance on all three.

1. Introduction

In recent years, short text in the form of posts on microblogs, question answer forums, news headlines, and tweets is being generated in abundance (Cagnina, Errecalde, Ingaramo, & Rosso, 2014). Performing NLP tasks is relatively easier in longer documents (e.g. news articles) than in short texts (e.g. headlines) because, in longer documents, greater context is available for semantic understanding (Rashid, Shah, & Irtaza, 2019). Moreover, in many cases, short texts (e.g. tweets) tend to use informal language (spelling variations, improper grammar, slang) compared to longer documents (e.g. blogs) (Shakeel, Karim, & Khan, 2019). Thus, the techniques tailored for formal and clean text do not perform well on informal one (Dey, Shrivastava, & Kaushik, 2016), which call for a need to develop an approach that can work in both settings (i.e., clean and noisy informal text) (Agarwal, Ramampiaro, Langseth, & Ruocco, 2018).

A paraphrase of a document is another document that can be different in syntax, but that expresses the same meaning in the same language. Automatically detecting paraphrases among a set of documents has many significant applications in natural language processing (NLP) and information retrieval (IR) such as plagiarism detection (Barrón-Cedeño, Vila, Martí, & Rosso, 2013), query ranking (Figuerola & Neumann, 2013), duplicate question detection (Bogdanova, dos Santos, Barbosa, & Zadrozny, 2015; Wang, Hamza, & Florian, 2017), web searching (Wang, Duan, Zhou, & Zhang, 2013), and automatic question answering (Fader, Zettlemoyer, & Etzioni, 2013).

* Corresponding author.

E-mail addresses: m.shakeel@lums.edu.pk (M.H. Shakeel), akarim@lums.edu.pk (A. Karim), imdad.khan@lums.edu.pk (I. Khan).

Paraphrase detection is a binary classification problem in which pairs of texts are labeled as either positive (paraphrase) or negative (non-paraphrase). In this setting, pairs of texts are mapped into a fixed-dimensional feature-space, where a standard classifier is learned. Feature maps based on lexical, syntactic and semantic similarities in conjunction with SVM are proposed in [Dey et al. \(2016\)](#) and [Eyecioglu and Keller \(2015\)](#). More recently, it has been demonstrated that for short text, deep learning-based pairs representations and classification yield better accuracy ([Agarwal et al., 2018](#)).

Many deep learning-based schemes employ one or two Convolutional Neural Network (CNN) or Long Short Term Memory (LSTM) based models to learn features and make predictions on clean texts ([Hu, Lu, Li, & Chen, 2014](#); [Tomar, Duque, Täckström, Uszkoreit, & Das, 2017](#); [Wang et al., 2017](#)), while a recent model also incorporates linguistic features to detect paraphrases in both clean and noisy short texts ([Agarwal et al., 2018](#)). For many NLP tasks involving short texts, it has been shown that developing wider models can yield significant gains ([Reimers & Gurevych, 2017](#)).

While deep models produce richer representations, they require large amounts of training data for a robust paraphrase detection system ([Dey et al., 2016](#)). Thus, for small datasets, such as Microsoft Research Paraphrase (MSRP) corpus and SemEval-2015 Twitter paraphrase dataset (SemEval), handcrafted features and SVM classifier have been widely used ([Dey et al., 2016](#); [Ji & Eisenstein, 2013](#)). Labeling pairs of documents in a human-based computation setting (e.g. crowd-sourcing) is costly ([Xu, Callison-Burch, & Dolan, 2015](#)). Therefore, [Agarwal et al. \(2018\)](#) and [Tomar et al. \(2017\)](#) add to the training set each labeled pair also in the reversed order. However, this simple data augmentation strategy can be extended in a systematic manner by relying upon set and graph theory. For instance, consider four documents: (a) *How can I lose weight quickly?* (b) *How can I lose weight fast?* (c) *What are the ways to lose weight as soon as possible?* (d) *Will Trump win US elections?*. If in the annotated corpus, documents (a) and (b) and documents (b) and (c) are marked as paraphrases, then by transitive extension, documents (a) and (c) can also be considered as paraphrases. Similarly, if documents (a) and (b) are labeled as paraphrase, while documents (b) and (d) are labeled as non-paraphrase, then a new non-paraphrase pair based on document (a) and (d) can be inferred reliably. Such a strategy can be used to generate additional annotations in a sound and cost-effective manner, and potentially enhance the performance of deep learning models for paraphrase detection.

In this paper, we propose a data augmentation strategy for generating additional paraphrase and non-paraphrase annotations reliably from existing annotations. We consider notions of paraphrases and non-paraphrases as binary relations over the set of documents. By representing the binary relation induced by the paraphrase labels as an undirected graph and performing transitive closure on this graph, we include additional paraphrase annotation in the training set. Similarly, by comparing paraphrase and non-paraphrase annotations, we infer additional non-paraphrase annotations for inclusion in the training corpus. Our strategy involves several steps and a parameter through which the data augmentation can be tuned for enhanced paraphrase detection.

We also present a robust multi-cascaded deep learning model for paraphrase detection in short texts. Our model utilizes three independent CNN and LSTM (with and without soft attention) cascades for feature learning in a supervised manner. We also employ a number of additional linguistic features after corpus-specific text preprocessing. All these features are fed into a discriminator network for final classification.

To show effectiveness of our approach we evaluate the data augmentation and deep model on three benchmark short text datasets (MSRP and Quora (clean), and SemEval (noisy)). We also perform extensive comparisons with the state-of-the-art methods. We make the following key contributions in this work:

- We present an efficient strategy for augmenting existing paraphrase and non-paraphrase annotations in a consistent manner. This strategy generates additional annotations and enhances the performance of the data-hungry deep learning models.
- We develop a multi-cascaded learning model for robust paraphrase detection in both clean and noisy texts. This model incorporates multiple learned and linguistic features in a wide and deep network for paraphrase detection.
- We address both clean and noisy texts in our presentation and show that the proposed model matches current best performances on benchmark datasets of both types.
- We analyze the impact of various data augmentation steps and different components of the multi-cascaded model on paraphrase detection performance.

The rest of the paper is organized as follows: We discuss the related work on paraphrase detection and data augmentation in [Section 2](#). We present our data augmentation strategy in [Section 3](#). Our multi-cascaded model for paraphrase detection is presented in [Section 4](#). [Section 5](#) outlines the experimental evaluation setup including discussion of data augmentation. We present and discuss the results of our approach in [Section 6](#). Finally, we present our concluding remarks in [Section 7](#).

2. Related work

Automatic paraphrase detection has been widely studied in the NLP and IR communities. The problem is posed as classification of pairs of text into one of paraphrase and non-paraphrase classes. In this setting, first, the pairs of texts are represented as fixed-length vectors. This representation tends to be sparse due to short length of texts ([Rashid et al., 2019](#)). Thus, this representation must be efficiently computable and should preserve as much contextual information as possible. Standard classifiers, such as SVM, are then learned in this representation scheme for detecting paraphrases. Short text can be clean (i.e., that follows proper grammar and formal diction like news headlines) or noisy (i.e., having informal verbiage and spelling variations like tweets). An abundance of work has been done on clean text paraphrase detection. A weighted term frequency approach for text pair representation along with n -gram overlap features between two texts is proposed in [Ji and Eisenstein \(2013\)](#) to train SVM classifier. In [Mohammad, Jaradat, Mahmoud,](#)

and Jararweh (2017), lexical (Parts of Speech (POS) overlap, minimum edit distance, text alignment) and semantic (Named-entity overlap, topic modeling) features between a pair of input text are used to train support vector regressor for paraphrase identification of news tweets in Arabic language. A probabilistic model that relies upon the similarity between syntactic trees of two input documents is proposed in Das and Smith (2009). The authors in Oliva, Serrano, del Castillo, and Iglesias (2011) devise a method for computing semantic similarity based on the lexical database. The model depends on WordNet meanings and syntactic roles among words in two documents.

Several studies have been carried out that utilize deep learning architectures for paraphrase detection in clean short text. A recursive auto-encoder for reliably understanding the context of texts and performing paraphrase detection is proposed in Socher, Huang, Pennin, Manning, and Ng (2011). This architecture forms a recursive tree and performs dynamic pooling to convert the input into fixed-sized representations. However, making a tree requires parsing hence this approach is less scalable. In Hu et al. (2014), patterns learned on a pair of text through CNN are matched at different levels of abstraction, introducing explicit interaction between the two documents during the learning process. In El-Alfy, Abdel-Aal, Al-Khatib, and Alvi (2015), five lexical metrics are used for reliable semantic similarity detection, where abductive networks are employed to get a composite metric that is used for classification. A method of decomposing text pair to similar and dissimilar components is proposed in Wang, Mi, and Ittycheriah (2016b). A CNN is then trained to convert these components into a fixed dimension vector and classification. In Yin, Schütze, Xiang, and Zhou (2016), three attention schemes are used in CNN to form interdependent document representations. Many neural network models are proposed to match documents from multiple levels of granularity. A multi-perspective matching model (BiMPM) is proposed in Wang et al. (2017), that uses character-based LSTM to learn word representations and a bi-directional LSTM for document representation for each text. After that, it performs four types of matching in “matching-layers” and finally, all these representations are aggregated by an additional bi-directional LSTM for paraphrase detection. One extension of this work is the neural paraphrase detection model based on a self-attended feed-forward network with pretrained embeddings on a huge corpus of another paraphrase data (Tomar et al., 2017). It is shown that this approach outperforms BiMPM model in terms of testing accuracy with fewer parameters.

As compared to clean short text, less work has been done for noisy short text paraphrase detection. In Xu, Ritter, Callison-Burch, Dolan, and Ji (2014), string-based features (whether the two words, their stemmed forms, and their normalized forms are the same, similar or dissimilar), common POS, and common topic (word’s association with a particular topic) between a pair of text are used as features and a novel multi-instance learning paraphrase model (MultiP) is proposed. Simple lexical features based on word and character n -gram overlaps between two texts are constructed to train SVM classifier in Eyecioglu and Keller (2015). An approach in Zhao and Lan (2015) uses corpus-based (similarity between sum of word vectors of two sentences), syntactic, and sentiment polarity based features and train SVM classifier. While Zarrella, Henderson, Merkhofer, and Strickhart (2015) uses ensemble approach based on seven models using word embeddings. In Dey et al. (2016), a set of lexical (character and word level n -grams), syntactic (words with matching POS tag, same verbs), semantic (adjective overlap), and pragmatic (subjective/objective agreement) features are identified between pair of input text, along with extensive preprocessing (spelling correction, stemming, stopwords removal, synonymous replacement). Although these features perform well on noisy short text, it is shown that they fail to get good predictive performance on clean text corpus of Microsoft Research Paraphrase (MSRP).

All previously reported approaches were focused either on clean text or noisy text. A recent study focuses to develop a single deep learning model that performs well on both clean and noisy short text paraphrase identification (Agarwal et al., 2018). A CNN and LSTM-based approach is adopted to learn sentence representations, while a feedforward network is used for classification. It also utilizes hand-crafted linguistic features, which improves paraphrase detection accuracy. Our approach follows this trend and we focus on a single model for paraphrase identification in both clean and noisy datasets. We also use linguistic features but our set of these features is different than (Agarwal et al., 2018).

Deep learning approaches need a large-scale annotated dataset for developing a robust model. For instance, the AskUbuntu dataset (dos Santos, Barbosa, Bogdanova, & Zadrozny, 2015) contains very few annotations, thus limiting the generalization performance of the model (Tomar et al., 2017). The ability to augment the data with additional sound annotations without requiring human intervention can improve the performance of deep models (Agarwal et al., 2018; Tomar et al., 2017). Such data augmentation has been shown to be fruitful for data analytics when only a piece of limited ordinal information about the pairwise distance between objects is provided (Heikinheimo & Ukkonen, 2013; Kleindessner & von Luxburg, 2017; Kleindessner & Von Luxburg, 2017). Data augmentation has also been shown to be prolific in image classification (Wang & Perez, 2017). Here, standard image processing such as cropping, rotation, and object translation is done to generate additional image samples. To the best of our knowledge, a systematic procedure for augmenting paired data without relying upon the object’s content has not been presented earlier.

3. Data augmentation for paraphrase detection

We start the presentation of our enhanced approach for paraphrase detection by discussing the proposed data augmentation strategy. Paraphrase annotation is costly and time-consuming while deep learning approaches demand a large corpus of paraphrase and non-paraphrase annotations. To address this problem, we develop strategies for generating additional annotations efficiently in a sound manner. We rely upon set theory and graph theory to model the problem and present an algorithm for generating additional data for training.

Let $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$ be the set of documents in the annotated corpus. The corpus contains annotations for paraphrases and non-paraphrases. The triplet $(d_i, d_j, 1)$ indicates that documents $d_i \in \mathcal{D}$ and $d_j \in \mathcal{D}$ are considered as paraphrases, and the triplet $(d_i, d_j, 0)$ denotes that documents $d_i \in \mathcal{D}$ and $d_j \in \mathcal{D}$ are considered as non-paraphrases. Let N_p and N_{np} denote the numbers of paraphrase

and non-paraphrase annotations in the corpus, and $N = N_p + N_{np}$ be the total number of annotations in the corpus. Note that in practice, only a fraction of the pairs of documents in \mathcal{D} will be annotated in the corpus, i.e., $N \ll |\mathcal{D}|^2$.

The information contained in the annotated corpus can also be represented as a graph over the vertex set in \mathcal{D} . Each triplet corresponds to an edge in the graph with its label (1 or 0) indicating whether the two documents are considered paraphrases or not. For example, the triplet $(d_i, d_j, 1)$ is represented by an edge between vertex $d_i \in \mathcal{D}$ and vertex $d_j \in \mathcal{D}$ with an edge-label 1. We assume that each edge can have a single label only, i.e., there are no conflicts in the annotated corpus whereby the same pairs of documents are labeled as both 1 and 0. If such conflicts do exist, they are removed from the corpus.

3.1. Generating additional paraphrase annotations

The notion of pairs of documents in \mathcal{D} being considered paraphrases can be captured by the notion of binary relation in set theory. Let $\mathcal{R}_p \subset (\mathcal{D} \times \mathcal{D})$ define the binary relation over \mathcal{D} such that $\forall i \forall j (d_i, d_j) \in \mathcal{R}_p$ implies that d_i is a paraphrase of d_j . In general, the following two properties hold for \mathcal{R}_p :

1. \mathcal{R}_p is reflexive, i.e., $\forall i, (d_i, d_i) \in \mathcal{R}_p$.
2. \mathcal{R}_p is symmetric, i.e., $\forall (i, j), (d_i, d_j) \in \mathcal{R}_p \Rightarrow (d_j, d_i) \in \mathcal{R}_p$.

The notion of paraphrasing is not defined precisely in linguistics. The boundary between paraphrases and non-paraphrases can lie on the continuum between (strong) paraphrases on one end and (strong) non-paraphrases on the other (Vila, Martí, & Rodríguez, 2014). For clean texts (e.g., news headlines), we may approximate the notion of paraphrase by the notion of semantic duplicate, i.e., $(d_i, d_j, 1)$ implies that documents d_i and d_j are considered duplicates semantically. In set theory, this corresponds to the equivalence relation. The equivalence relation \mathcal{R}_e over \mathcal{D} possesses the following property in addition to properties 1 and 2 listed above:

3. \mathcal{R}_e is transitive, i.e., $\forall (i \neq j \neq k), [(d_i, d_j) \in \mathcal{R}_e \wedge (d_j, d_k) \in \mathcal{R}_e] \Rightarrow (d_i, d_k) \in \mathcal{R}_e$.

Transitivity can be a strong property when applied to the notion of paraphrases, especially for noisy text. Therefore, we consider \mathcal{R}_p to include transitive extensions of the direct relation \mathcal{R} to a pre-selected order $K \geq 1$. That is, $\mathcal{R}_p = \mathcal{R}_p \cup \mathcal{R} \cup \mathcal{R}^1 \cup \dots \cup \mathcal{R}^K$ where \mathcal{R} denotes the relation that two documents are paraphrases due to a direct relationship between them and relation \mathcal{R}^K indicates that two documents are considered paraphrases because there are $K \geq 1$ intermediate documents relating them (\mathcal{R}^1 is the transitive extension of \mathcal{R} , and \mathcal{R}^K is the transitive extension to order K of \mathcal{R}). If $K = *$ (i.e., maximum order extension is done) then we achieve a transitive closure of \mathcal{R} .

Now, consider the graph on the vertex set \mathcal{D} induced by edges labeled with 1, i.e., pairs considered to be paraphrases in the annotated corpus. These pairs do not necessarily induce the relation \mathcal{R}_p on \mathcal{D} . We therefore add more pairs to transform it into the desired relation. More formally:

1. For each $d_i \in \mathcal{D}$, we add $(d_i, d_i, 1)$ in the corpus, i.e., we declare each d_i a paraphrase of itself. We call this step generation by reflexivity.
2. For each $(d_i, d_j, 1)$ in the corpus, we add $(d_j, d_i, 1)$ in the corpus, i.e., we consider d_j and d_i to be paraphrases of each other. We call this step generation by paraphrase symmetry.
3. For every chain of annotations $(d_i, d_{j_1}, 1), (d_{j_1}, d_{j_2}, 1), \dots, (d_{j_{K-1}}, d_{j_K}, 1), (d_{j_K}, d_k, 1)$ starting at d_i and ending at d_k with at most K intermediate documents in the annotated corpus, we add $(d_i, d_k, 1)$ and $(d_k, d_i, 1)$ in the corpus. Thus, we consider d_i and d_k to be paraphrases of each other if these documents are connected by at most K intermediate documents. We call this step generation by paraphrase transitive extension.

In graph terminology, step 1 corresponds to adding self-loops on each vertex with label 1 while in step 2, we ignore the direction on all edges by considering the graph as an undirected graph. Step 3 corresponds to performing a transitive extension on the undirected graph induced by edges labeled with 1. Every vertex needs to be made adjacent to all vertices that are reachable from it in $\leq K$ hops. This can be done with a single BFS (breadth first search) or DFS (depth first search) on the graph.

Note that transitive extension transforms the graph into a collection of cliques (fully connected vertices) where each clique is a paraphrase class representing a unique concept.

3.2. Generating additional non-paraphrase annotations

Let \mathcal{R}_{np} denote the binary relation that two documents in \mathcal{D} are considered non-paraphrases. By definition, this relation is irreflexive, i.e., $\forall i, (d_i, d_i) \notin \mathcal{R}_{np}$. Obviously, a document cannot be a non-paraphrase of itself. The relation \mathcal{R}_{np} is symmetric, i.e., $\forall i, j, (d_i, d_j) \in \mathcal{R}_{np} \Rightarrow (d_j, d_i) \in \mathcal{R}_{np}$. Transitivity does not hold for relation \mathcal{R}_{np} ; if $(d_i, d_j) \in \mathcal{R}_{np}$ and $(d_j, d_k) \in \mathcal{R}_{np}$, then we cannot say for sure that d_i is not a paraphrase of d_k .

Additional non-paraphrase annotations can also be inferred by comparing them with paraphrase annotations. For example, if $(d_i, d_j, 0)$ and $(d_j, d_k, 1)$ (i.e., d_i and d_j are non-paraphrases, while d_j and d_k are paraphrases), then $(d_i, d_k, 0)$ must also be true (i.e., d_i and d_k are non-paraphrases). Of course, if d_i and/or d_k lie within two different paraphrase classes/cliques, then all pairs of documents across the classes/cliques will be considered non-paraphrases. These relations are also included in \mathcal{R}_{np} .

Now, consider the graph on vertex set \mathcal{D} , induced by edges labeled with 0, i.e., non-paraphrases in the annotated corpus. These pairs do not necessarily induce the relation \mathcal{R}_{np} on \mathcal{D} . We, therefore, add more pairs to transform it into the desired relation as follows:

1. For each $(d_i, d_j, 0)$, in the corpus, we add $(d_j, d_i, 0)$ in the corpus, i.e., we consider d_j and d_i to be non-paraphrases as well. We call this step generation by non-paraphrase symmetry.
2. For every $(d_i, d_j, 0)$ in the corpus, let C and C' be the cliques (paraphrase classes) containing d_i and d_j , respectively, we add $(d_m, d_n, 0)$ and $(d_n, d_m, 0)$ for each $d_m \in C$ and $d_n \in C'$ to the corpus. We call this step generation by non-paraphrase transitive extension.

In terms of graph, the second step corresponds to making a complete bipartite graph between the vertex set of C and that of C' .

3.3. Conflicts and errors in annotations

Using the principled strategy outlined earlier, conflicts and errors in annotations can be identified and potentially fixed. A conflict occurs when a particular pair of documents is found to be paraphrase and non-paraphrase either in the original annotated corpus or during augmentation. Based on our data augmentation strategy described above, the following conflicts can arise: (1) In the original annotated corpus, a pair of documents is labeled as both paraphrase and non-paraphrase. (2) Erroneous annotations can be generated during our data augmentation strategy in the following two cases: (a) when generating additional paraphrases by a transitive extension (b) when generating additional non-paraphrases by non-paraphrase transitive extension. A detailed analysis of conflicts and errors and their resolution is beyond the scope of this paper in which we focus on data augmentation and its impact on paraphrase detection. Nonetheless, we believe that this is a fruitful area for future research.

In this work, we resolve the first type of conflict by removing the conflicting annotations and the second type of conflict by removing the conflicting non-paraphrase annotation and retaining the generated paraphrase annotation. We control the number of conflicts and errors by varying K and selecting/dropping specific generation steps and we evaluate these variations by their performances on paraphrase detection.

3.4. Algorithm

Algorithm 1 outlines our proposed strategy for augmenting data for enhanced paraphrase detection. The algorithm takes as input the set of documents \mathcal{D} , the (original) annotated corpus or dataset \mathcal{A} , and the parameter K (extension order) and it outputs the augmented annotated corpus or dataset $\tilde{\mathcal{A}}$. After removing conflicts in the dataset (lines 4 – 5), the algorithm proceeds with generating additional paraphrase annotations (lines 6 – 15) followed by generating additional non-paraphrase annotations (lines 16 – 19). Generating paraphrase annotations involve 3 steps i.e., P1 (lines 6 – 7), P2 (lines 8 – 9), and P3 (lines 10 – 15), while generating non-paraphrase annotations consists of two steps i.e., NP1 (lines 16 – 17) and NP2 (lines 18 – 19). It is worth noting that step P1 can be performed at any sequence while the other steps must follow the given sequence. In our experiments, we perform P1 after P2 and P3. Note that step P1 means additionally generated paraphrase pairs would be equal to the number of unique texts in the training data (minus number of pairs that were already part of the original annotations).

The worst-case computational complexity of the algorithm is defined by step P3. If Z is the size of the largest paraphrase class (max. clique size or max. node degree in graph terminology), then the computational complexity of the algorithm is $O(|\mathcal{D}|ZK)$. Note that in practice both Z and K will be much less than $|\mathcal{D}|$.

4. Multi-cascaded deep model for paraphrase detection

In this section, we present our multi-cascaded deep learning model for enhanced paraphrase detection. While deep models have been popularly used in recent years for paraphrase detection, they are typically tailored to either clean text (e.g., news headlines) or noisy text (e.g., tweets). Similarly, most employ a single model for feature learning and discrimination, and some do not utilize linguistic features in their models. In order to benefit from previous insights and to produce robust paraphrase detection for both clean and noisy texts, we propose three independent feature learners and a discriminator model that can consider both learned and linguistic features for paraphrase detection.

Fig. 1 shows the architecture of our multi-cascaded model for paraphrase detection. We employ three feature learners that are trained to distinguish between paraphrases and non-paraphrases independently (in parallel). The features from these models (one layer before the output layer) are subsequently fed into a discriminator network together with any additional linguistic features to make the final prediction.

Our model takes as input a pair of documents d_i and d_j and outputs the label 1 if the documents are considered paraphrases and the label 0 if the documents are considered non-paraphrases. Let $d_i = \langle w_1^i, w_2^i, \dots, w_T^i \rangle$ be the sequence of words in a document. We represent each word in the sequence by an e dimensional fixed-length vector (word embedding). For a given paraphrase detection problem, we empirically decide length T of each document to use such that longer documents are truncated and shorter ones are padded with zero vectors. This decision is made to ensure compatible length vectors for all document pairs across different cascades. As discussed in [Section 6](#), we select an appropriate length T for each dataset based on the distribution of lengths of documents in the dataset. We experiment with several linguistic features (syntactic and lexical) for both clean and noisy text paraphrase detection. The details of these features are given in [Section 5.3](#).

- 1: **Input:** \mathcal{D} (documents), \mathcal{A} (original corpus), K (extension order)
- 2: **Output:** $\bar{\mathcal{A}}$ (augmented corpus)
 - Remove conflicting annotations
- 3: **for all** $[(d_i, d_j, 0) \in \mathcal{A} \vee (d_j, d_i, 0) \in \mathcal{A}] \wedge [(d_i, d_j, 1) \in \mathcal{A} \vee (d_j, d_i, 1) \in \mathcal{A}]$ **do**
- 4: $\mathcal{A} \leftarrow \mathcal{A} \setminus \{(d_i, d_j, 0), (d_j, d_i, 0), (d_i, d_j, 1), (d_j, d_i, 1)\}$
- 5: $\bar{\mathcal{A}} \leftarrow \mathcal{A}$
 - Paraphrase augmentation
 - Step P1: by reflexivity
- 6: **for all** $d_i \in \mathcal{D}$ **do**
- 7: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \cup (d_i, d_i, 1)$
 - Step P2: by paraphrase symmetry
- 8: **for all** $(d_i, d_j, 1) \in \bar{\mathcal{A}}$
- 9: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \cup (d_j, d_i, 1)$
 - Step P3: by paraphrase transitive extension
- 10: $\mathcal{A} \leftarrow \bar{\mathcal{A}}$
- 11: **for** $(n \leftarrow 1 \rightarrow K)$ **do**
- 12: **for all** $(d_i, d_{j_1}, 1) \in \mathcal{A} \wedge \dots \wedge (d_{j_n}, d_k, 1) \in \mathcal{A}$ **do**
- 13: **if** $(d_i, d_k, 0) \in \bar{\mathcal{A}} \vee (d_k, d_i, 0) \in \bar{\mathcal{A}}$ **then**
- 14: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \setminus \{(d_i, d_k, 0), (d_k, d_i, 0)\}$
- 15: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \cup \{(d_i, d_k, 1), (d_k, d_i, 1)\}$
 - Non-paraphrase augmentation
 - Step NP1: by non-paraphrase symmetry
- 16: **for all** $(d_i, d_j, 0) \in \bar{\mathcal{A}}$ **do**
- 17: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \cup (d_j, d_i, 0)$
 - Step NP2: by non-paraphrase transitive extension
- 18: **for all** $(d_i, d_j, 0) \in \bar{\mathcal{A}} \wedge (d_j, d_k, 1) \in \bar{\mathcal{A}}$ **do**
- 19: $\bar{\mathcal{A}} \leftarrow \bar{\mathcal{A}} \cup \{(d_i, d_k, 0), (d_k, d_i, 0)\}$

Algorithm 1: Data augmentation strategy

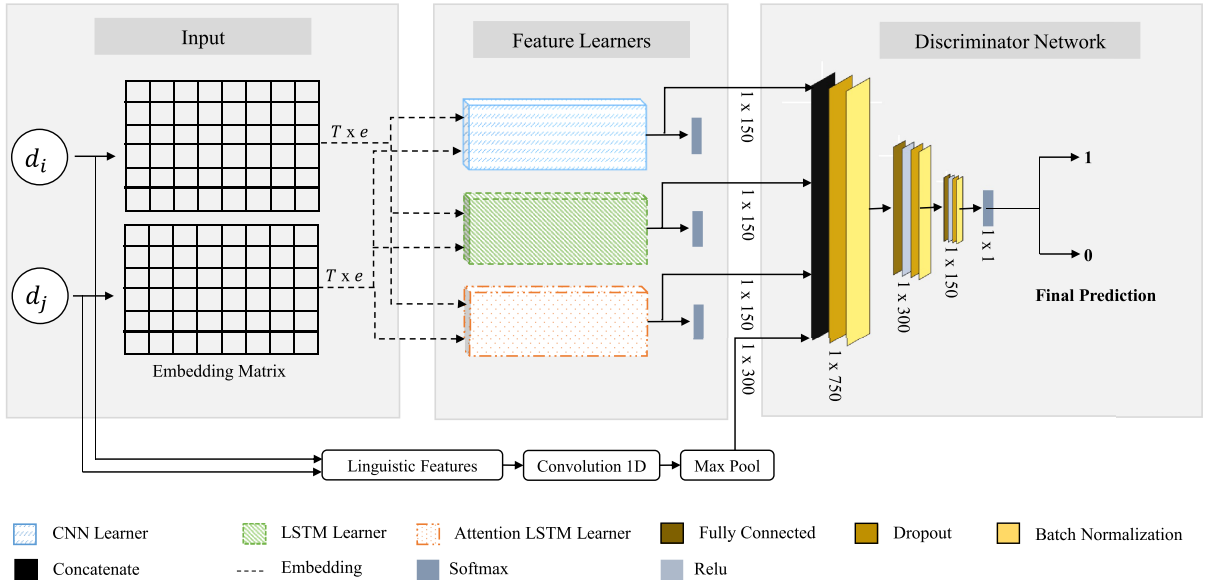


Fig. 1. Multi-cascaded model for enhanced paraphrase detection. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

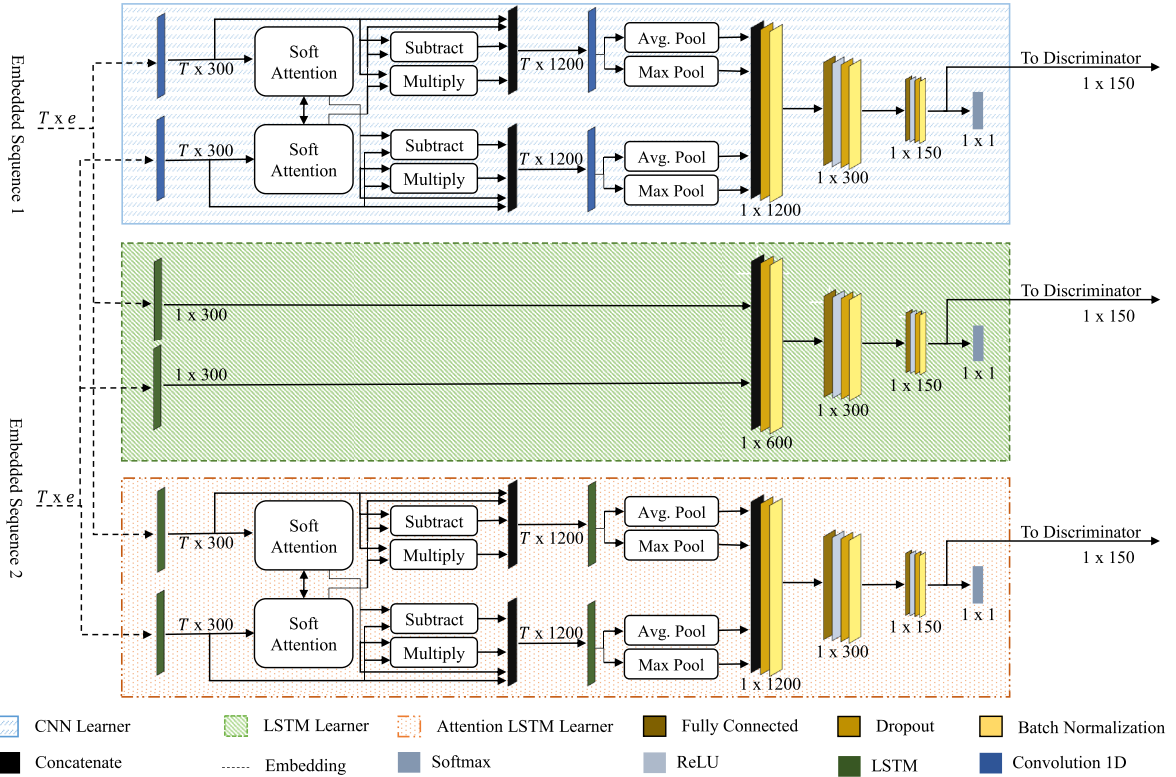


Fig. 2. Architecture and dimensions of output for each cascade in our multi-cascaded model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

The details of the feature learners and the discriminator network are given in the following subsections.

4.1. Feature learners

We employ three independent cascades to learn contextual features for paraphrase detection. Each cascade focuses on a different sequential learning model to extract features from different perspectives. Each cascade takes as input the pair of documents to be classified as paraphrase or non-paraphrase, and each is trained independently (in parallel) on the annotated corpus. The details of all three cascades are discussed in the following paragraphs.

The first cascade is based on CNN with soft-attention (Fig. 2). The first layer has 300 CNN filters with kernel size of 1. Subsequently, soft-attention is applied to highlight words in documents d_i and d_j that are more important to achieve correct prediction. The result of soft-attention for document t_i is then subtracted and multiplied with that of document t_j to learn semantic contrariety. The next layer concatenates the output of CNN layer, soft-attention output, subtracted output, and multiplied output for both documents (one layer concatenates these 4 outputs for d_i and other for d_j independent of each other). This output is forwarded to another CNN layer containing 300 filters with kernel size of 2. The purpose of this CNN layer is to learn bi-gram feature representation. After this representation, global max-pooling and global average-pooling is performed to obtain most important bi-gram and average of all bi-grams respectively for both documents. These two representations are then concatenated to make a single vector of both documents (up till now, both documents were being treated separately so two streams of same functions were being produced). This outputs a vector of length 1,200. After this concatenation, a dropout and batch normalization layer is deployed to avoid any feature co-adaptation. Then this representation is forwarded to a fully connected layer of 300 units with ReLU activation followed by drop out and batch normalization. Finally, this 300 dimensional representation is squashed into a 150-dimensional vector by another fully connected layer. This cascade is learned to detect paraphrases and non-paraphrases, and the learned 150-dimensional representation is forwarded to discriminator network for final classification.

The second cascade utilizes LSTM to encode long-term dependencies in the documents (Wang, Jiang, & Luo, 2016a) (Fig. 2). This cascade learns contextual representations without any attention or semantic contrariety. As such, it comprises of a single LSTM layer with 300 units followed by fully connected hidden layers with 300 and 150 units each and an output layer. Again, the 150-dimensional representation serves as an input to the discriminator network for final classification.

The third cascade is based on LSTM with soft-attention (Fig. 2). This cascade is similar to the first cascade but uses LSTM units instead of CNN filters. As such, it does not learn bi-gram feature representation but instead learns long-term dependencies with

important sequences highlighted by attention mechanism. The 150-dimensional vector obtained before the output layer is used as an input to the discriminator network for final classification.

We empirically decide to use 300 CNN filters, 300 LSTM units, 300 units for first fully-connected layer, and 150 units for second fully-connected layer in all cascades. In all feature learning cascades, weights of CNN filters as well as weights of LSTM layers are shared among d_i and d_j . Weight sharing reduces the number of parameters required and both documents are converted to deep representations in same embedding space. We use ReLU activation for every layer except last prediction layer, which uses the softmax function. We train the model using categorical cross entropy loss. Dropout and batch normalization is used after every fully-connected layer.

4.2. Discriminator network

Fig. 1 shows the discriminator network utilized to make the final prediction. For a given pair of documents d_i and d_j , this network takes as input the features from all three cascades plus any linguistic features with dropout and batch normalization. Two fully-connected hidden layers are then used with 300 and 150 units, respectively, where each is followed by dropout and batch normalization layer. The activation function for both fully-connected layer is set to ReLU. On final layer softmax activation function is used with categorical cross-entropy as loss. This network outputs final prediction for documents d_i and d_j as either paraphrase or non-paraphrase. We also do early-stopping if validation accuracy of discriminator is not improved for 10 epochs. A checkpoint of the model is created after the training epoch at which validation accuracy is improved.

5. Experimental setup

In this section, we describe the settings for the experimental evaluation of our enhanced paraphrase detection model. We discuss the key characteristics of the three datasets used in our evaluations before and after augmentation. We also present the parameter settings of the model, the text preprocessing performed, and the linguistic features used in our experiments.

5.1. Datasets and their augmentation

We use three real-world datasets in our experimental evaluations. The Quora questions pairs dataset (Quora) (Wang et al., 2017) contains pairs of questions in English with their annotations (paraphrase or non-paraphrase).¹ This dataset is collected from questions posted on the Quora question-answering website. This is a clean text dataset. The Microsoft Research Paraphrase Corpus (MSRP) (Dolan, Quirk, & Brockett, 2004) contains pairs of sentences in English together with their paraphrase/non-paraphrase annotation.² The sentences are extracted from news articles on the web. Thus, this is another example of a clean text dataset. The SemEval-2015 Twitter paraphrase dataset (SemEval) (Xu et al., 2014) contains pairs of tweets in English with their paraphrase/non-paraphrase annotation.³ This is a noisy text dataset. The Quora and SemEval datasets have pre-defined training, development, and test sets available. The MSRP dataset has predefined training and test sets only. The statistics for each split of all the datasets are presented in Table 1. We describe the characteristics of the *training* set of each dataset before and after augmentation in the following subsections.

5.1.1. Quora dataset

The Quora dataset contains annotations for 517,968 unique questions. We notice that there are 30 incorrect non-paraphrase annotations in the dataset. The texts in these 30 annotations are identical (reflexive) but they are marked as non-paraphrases. We remove these annotations from the dataset.

An analysis of this dataset reveals that questions have an average length of about 13 words with a standard deviation of 6.7 words and a maximum length of 272 words. Based on this analysis, we select $T = 40$ for this dataset.

Table 2 shows the numbers of paraphrases and non-paraphrases in the Quora dataset before and after each step of augmentation. For this clean text dataset, we perform step P3 (paraphrase generation by transitive extension) using $K = *$, i.e., we generate by transitive closure. Performing transitive closure (step P3) generates 168,766 additional paraphrase annotations. Applying step P2 (paraphrase generation by symmetry) almost doubles the number of paraphrase annotations. Subsequently, generating paraphrase annotations by reflexivity (step P1) produces an additional 517,131 paraphrase pairs, bringing the total of paraphrase annotations to 964,509. As an example of generated paraphrase pairs through transitive closure, consider following three questions. (a) *What causes deja vu ?* (b) *Read below, what causes deja vu ?* (c) *What is the cause of someone frequently experiencing deja vu ?*. The questions (a) and (b) and questions (b) and (c) are marked as paraphrases in the original annotations. After transitive closure, a new pair based on questions (a) and (c) is correctly inferred.

For augmentation of non-paraphrase annotations, step NP1 (non-paraphrase generation by symmetry) brings the total number of non-paraphrase pairs to 490,015. Subsequently, performing step NP2 (generation by non-paraphrase transitive extension) an additional 165,204 non-paraphrase pairs are produced, bringing the total number of non-paraphrase annotations to 655,219. Similar to the example of paraphrase annotation generation, the proposed data augmentation strategy successfully generates non-paraphrase

¹ <https://github.com/zhiguowang/BiMPPM>.

² <https://www.microsoft.com/en-us/download/details.aspx?id=52398>.

³ <https://github.com/cocoxu/SemEval-PIT2015>.

Table 1
Statistics for all datasets.

Dataset	Split	Total pairs	Paraphrase pairs	Non-paraphrase pairs	Debatable pairs
Quora	Train	384,348	139,306	245,042	–
	Dev	10,000	5,000	5,000	–
	Test	10,000	5,000	5,000	–
MSRP	Train	4,076	2,753	1,323	–
	Test	1,725	1,147	578	–
SemEval	Train	13,063	3,996	7,534	1,533
	Dev	4,727	1,470	2,672	585
	Test	972	175	663	134

Table 2
Numbers of paraphrase and non-paraphrase annotations in the datasets before and after augmentation.

Dataset	Paraphrases				Non-Paraphrases		
	Original	P2	P3	P1	Original	NP1	NP2
Quora	139,306	278,612	447,378	964,509	245,042	490,015	655,219
MSRP	2,753	5,506	5,874	13,688	1,323	2,646	2,647
SemEval	3,996	7,992	94,722	107,953	7,534	15,068	23,205

annotations also. Consider following four questions. (a) *Why does Quora censor opinions and answers ?* (b) *Is Quora censored ?* (c) *Why does Quora want to censor/collapse the truth ?* (d) *Why does Quora have a character limit for question titles and details ?*. In original annotations, questions (a), (b), and (c) are recorded as paraphrases while questions (b) and (d) are recorded as non-paraphrases. Using non-paraphrase transitive extension, additional non-paraphrase pairs (a), (d) and (c), (d) are correctly inferred.

After performing step P3 (generation by transitive closure), it is observed that there are 57,119 unique paraphrase classes (cliques) corresponding to distinct concepts about which questions are being asked. Note that transitive closure converts the graph on paraphrase annotations into a disjoint collection of cliques. Therefore, the node degree after transitive closure is one less than the number of questions (nodes) of the clique in which that node lies.

As discussed in Section 3.3, our data augmentation strategy can highlight conflicts and errors in the original annotations. For example, a paraphrase annotation generated by transitive closure can be in conflict with an existing non-paraphrase annotation. Usually, a generated annotation is based on strong evidence of related annotations, hence, such conflicts may indicate an error in the existing annotations. We find 214 such conflicts in the training set of the Quora dataset. The number of conflicts can also be considered as a measure of annotation quality. Upon subjective evaluation of these conflicts, it is observed that annotations generated by the proposed data augmentation strategy are correct while the original annotations are erroneous. Table 3 shows some examples of such conflicts in the Quora dataset. If we consider the first example, the only difference between both question is the spelling of the word “color”, yet it is marked as non-paraphrase in the original label while the paraphrase transitive extension marks it as paraphrase, which is correct. Similarly, in second example, same question is asked with different wordings but the original annotations mark it as non-paraphrase while the proposed data augmentation strategy correctly marks it as paraphrase. These results confirm that our data augmentation strategy works well to detect and distinguish unique concepts and generate new pairs along with their associated labels from existing annotations reliably.

5.1.2. MSRP dataset

The MSRP dataset is a much smaller dataset containing annotations for 7,814 unique sentences. There are no conflicting annotations found in this dataset. We select $T = 25$ based on exploratory analysis of the dataset which reveals that the average length of sentences is about 19 words with a standard deviation of 5.1 word and a maximum length of 31 words.

Table 2 gives the numbers of paraphrases and non-paraphrases before and after augmentation in this dataset. There are only 4,076 annotations in the original train split of the dataset out of which 2,753 are for paraphrases and 1,323 are for non-paraphrases. The

Table 3

Examples of errors detected during paraphrase augmentation using transitive extension on Quora dataset (1 = paraphrase, 0 = non-paraphrase).

Question pair	Original label	Generated label
What is the colour of the Sun? What is the color of the sun?	0	1
Is pro wrestling fake? Wwe is real fight?	0	1
How can I get free iTunes gift cards online? What 's the best way to legally get free iTunes gift cards?	0	1

Table 4

Numbers of paraphrase and non-paraphrase annotations after step P3 and step NP2, respectively, in SemEval dataset with different orders of transitive extension.

Tr. extension	Paraphrases		Non-Paraphrases
	P3	P1	NP2
$K = 1$	30,738	43,969	18,539
$K = 2$	59,538	72,769	23,490
$K = 3$	90,042	103,273	24,175
$K = *$	94,722	107,952	23,205

number of paraphrases is doubled after performing step P2 (generation by paraphrase symmetry), an additional 368 paraphrase annotations are generated in step P3 (generation by transitive closure), and step P1 generates an additional 7,814 paraphrase annotations. Step NP1 doubles the number of non-paraphrase annotations while step NP2 generates just one more non-paraphrase annotation. As seen from these results, transitive closure does not add many annotations implying that there are generally small paraphrase classes in this dataset.

In this dataset, there were no conflicts detected while employing the proposed augmentation strategy.

5.1.3. SemEval dataset

The SemEval dataset has paraphrase and non-paraphrase annotations for 13,231 unique tweets. The original train split of the dataset has 3,996 paraphrase, 7,534 non-paraphrase, and 1,533 *debatable* annotations (Xu et al., 2015). All existing studies on this dataset ignore the debatable annotations while reporting their results; therefore, we also choose to ignore these annotations in our experiments. No conflicting annotations are found in the dataset.

For this dataset, we choose $T = 19$. This number corresponds to the maximum length of tweets in the dataset with the average length being around 8 words.

Table 2 presents the numbers of paraphrases and non-paraphrases in this dataset before and after each step of our data augmentation strategy. These numbers are obtained when transitive closure is performed during step P3, i.e., $K = *$. It is seen that the number of paraphrases jumps from 3,996 to 107,953 while the number of non-paraphrases increases from 7,534 to 23,205. A major increase in paraphrases occurs during step P3. This can be attributed to the following two reasons: (1) The SemEval is based on a dataset developed for the task of semantic similarity estimation. As such, the notions of paraphrase and non-paraphrase are not well separated which is then blurred by the process of transitive closure. (2) The SemEval dataset contains short and noisy text that makes annotation difficult and prone to errors. For such datasets, transitive closure can be a “blunt instrument” during the data augmentation strategy. Therefore, we utilize parameter K to control the transitive extension levels.

Table 4 shows the numbers of paraphrases and non-paraphrases after steps P3, P1, and NP2 when transitive extension of orders $K = 1$, $K = 2$, $K = 3$, and $K = *$ (transitive closure) is performed. The numbers for step P2 and NP1 are not shown in this table as they are identical to those given in Table 2. It is clear from this table that order K controls the number of paraphrases that are generated during step P3. For example, when $K = 1$ the number of paraphrases after step P3 is 30,738 which is significantly lower than 94,722 when $K = *$. Note that the number of non-paraphrases actually increases slightly as the order of transitive extension is reduced. This is due to the fact that more conflicts arise as K is increased that are resolved by retaining the generated paraphrase annotation and discarding the conflicting non-paraphrase annotation, thus fewer non-paraphrase annotations will produce fewer additional non-paraphrase annotations during non-paraphrase extension (step NP2). The number of conflicts during step P3 are 4, 14, 22, and 23 for $K = 1$, $K = 2$, $K = 3$, and $K = *$, respectively.

5.2. Preprocessing

Text preprocessing is an essential component of many NLP applications. However, in case of short text, common text preprocessing steps such as removing punctuations and stopwords can result in loss of information critical to the application (Eyecioglu & Keller, 2015). Therefore, we keep preprocessing to a minimum in our experiments. For SemEval dataset, which represents noisy texts, we perform lemmatization and correct commonly misspelled words such as *dnt* to *do not*. We use a predefined dictionary to map misspelled words to their standard forms. Preprocessed tweets are used while training our multi-cascaded model as well as to extract linguistic features. For the other two datasets, which represent clean texts, we perform preprocessing (stopword removal and lemmatization) only for computing linguistic features while the raw text is used in our multi-cascaded model. The details of the linguistic features are given in the next subsection.

5.3. Linguistic features

We employ a set of NLP/linguistic features in our experiments as it has been shown that including linguistic features for paraphrase identification in short text can improve the performance of deep learning models (Agarwal et al., 2018). We identify the following linguistic and statistical features to be used alongside learned features in our multi-cascaded model.

1. 2 features based on cosine similarity between TF-IDF vectors of documents d_i and d_j , before and after removing stopwords and doing lemmatization.
2. 4 n-gram overlapping ratio features based on unigrams and bigrams that are common to a given document pair, divided by total number of n-grams in d_i and d_j respectively.
3. 2 features based on cosine similarity between ELMo (Peters et al., 2018) embeddings vectors of d_i with d_j , before and after removing stopwords and doing lemmatization.
4. 2 features based on cosine similarity between Universal Encoder⁴ vectors of d_i and d_j , before and after removing stopwords and doing lemmatization.
5. 1 feature based on count of unigrams that has same POS tag in d_i and d_j .
6. 6 features based on length of intersection of character bigrams, trigrams and quadgrams of d_i in d_j , before and after removing stopwords and doing lemmatization.
7. 2 features based on longest substring match in d_i and d_j , before and after removing stopwords and doing lemmatization.
8. 2 features based on longest subsequence match in d_i and d_j , before and after removing stopwords and doing lemmatization.
9. 3 syntactic features based on number of Verbs, Nouns and Adjectives common in d_i and d_j .
10. 2 Named-entity Recognition (NER) features in d_i and d_j based on number of same NER tags and numbers of same word-NER tuple.

We use all linguistic features for clean text datasets and linguistic features 1 – 4 only for the noisy text dataset. Linguistic features are passed through a single CNN layer with 300 dimensions and then provided as input to discriminator network.

5.4. Hyper-parameters tuning

A number of hyper-parameters are required in our multi-cascaded model for paraphrase detection. We tune and select these hyper-parameters on the development sets of each dataset using a grid search. As MSRP dataset has train and test splits only, we hold-out 10% of the training set as the development set for the purpose of hyper-parameters tuning. We decide the type of word embeddings among GloVe⁵ (Pennington, Socher, & Manning, 2014), Word2Vec (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013), and ELMo. For selecting an optimizer, we decide among nAdam, Adam, Adadelta, and SGD. We consider dropout rates 0.1, 0.2, 0.3, 0.4 and 0.5 in our model.

5.5. Performance evaluation and comparison

Since paraphrase detection is a binary classification problem, standard measures of performance can be used for evaluation. We report performances as percent accuracy, precision, recall, and F1-value on the test sets after training over the respective training set of the datasets. Each dataset has a fixed training and test sets. Therefore, results can be compared with previously reported results on the same datasets.

5.6. Implementation

We use networkX library in Python for graph analysis and data augmentation.⁶ For implementing and evaluating our multi-cascaded model we use Keras⁷ as the front-end with TensorFlow⁸ on the backside. All model parameters or weights are initialized randomly, and to ensure reproducibility the random seed is fixed. The code and implementation setting used in our experimental evaluation is available from the website.⁹

6. Results and discussion

In this section, we present and discuss the evaluation of our multi-cascaded model and data augmentations strategy in terms of paraphrase detection predictive performance. We first present results on each dataset with and without data augmentation and linguistic features. Subsequently, we discuss the performance of different components of our multi-cascaded model. Finally, we present the key takeaways from our experimental study.

6.1. Quora dataset

For this dataset, the hyper-parameters of our model are tuned on the provided dev set. GloVe embeddings are selected as the best choice while among optimizers, nAdam with learning rate of 0.002 is found to be optimal. Similarly, the dropout rate of 0.1 is found

⁴ <https://tfhub.dev/google/universal-sentence-encoder/2>.

⁵ <https://nlp.stanford.edu/projects/glove/>.

⁶ <https://networkx.github.io/>.

⁷ <https://keras.io/>.

⁸ <https://www.tensorflow.org/>.

⁹ <https://github.com/haroonshakeel/multi-cascaded-deep-network-for-paraphrase-identification>.

Table 5
Paraphrase detection performance on Quora dataset.

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
None	89.4	90.3	88.4	89.3	89.8	90.0	89.6	89.8
P2, NP1	89.6	90.6	88.4	89.5	74.9	67.6	95.5	79.2
P2, P3	89.9	90.0	89.7	89.9	90.1	91.2	88.7	89.9
P2, P3, P1	90.2	90.3	90.1	90.2	90.0	90.2	89.6	89.9
P2, P3, NP1	90.0	90.5	89.5	90.0	90.0	89.9	90.1	90.0
P2, P3, P1, NP1	90.3	90.9	89.5	90.2	90.1	89.7	90.5	90.1
P2, P3, P1, NP1, NP2	89.9	90.7	89.0	89.8	89.9	90.2	89.5	89.9

to be optimal.

Table 5 shows the predictive performance of our enhanced paraphrase detection model on the Quora dataset. Performances (accuracy, precision, recall, and F1-score) on test sets are given for different data augmentation steps with learned features and learned plus linguistic features. We first discuss results obtained by using learned features only. Using the original data without any augmentation, our model achieves an accuracy of 89.4%. Augmenting the data with step P2 and NP1 (paraphrase and non-paraphrase generation by symmetry) increases the accuracy slightly. Note that this augmentation step has also been performed in earlier works (Agarwal et al., 2018; Tomar et al., 2017). However, noticeable increase in accuracy is observed when the data is augmented with additional paraphrases using step P2, step P3 (generation by transitive closure), and step P1 (generation by reflexivity), jumping the accuracy to 90.2%. We obtain the best performance of 90.3% when in addition to augmenting paraphrases via steps P3, P2, and P1 additional non-paraphrases are generated via step NP1. This is also the current state-of-the-art performance on this dataset.

Note that by including additional non-paraphrase annotations using step NP2 (generation by non-paraphrase transitive extension) decreases the accuracy to 89.9% from the high of 90.3% obtained when steps P2, P3, P1, and NP1 are executed. The reason behind this decrease can be determined by analyzing paraphrase concepts (clique in paraphrase graph). Recall from Section 3.2 that even a single edge with label 0 (a non-paraphrase annotation) between two cliques will generate a complete set of edges between the nodes of the two cliques with label 0. Thus, any error in such non-paraphrase annotations gets magnified during the NP2 non-paraphrase augmentation step and degrades the quality of the dataset for paraphrase detection. For example, the incorrect annotation of questions *Is there a way to hack Facebook account ?* and *How can I hack Facebook ?* as non-paraphrase generates numerous erroneous non-paraphrase annotations between paraphrases of the first and second question. Therefore, step NP2 has the potential to degrade paraphrase detection performance when errors exist in existing non-paraphrase annotations that link large paraphrase concepts.

When we perform experiments by including linguistic features with learned features, slightly lower performances are obtained. This highlights that when sufficiently large dataset is available, deep learning models can effectively capture the semantics and contexts of short texts for improved paraphrase detection. For such datasets, the extra effort of including linguistic features is not beneficial.

Table 6 presents the performance of previously published work on this dataset. Accuracy values are given in this table because previous works report accuracies only. In Lan and Xu (2018), 7 different models are re-implemented on several tasks involving sentence pairs. Quora dataset is used to get results for paraphrase detection task. We only include results of best performing model among all 7. They find that Shortcut-Stacked Sentence Encoder Model (SSE) (Nie & Bansal, 2017) performs the best, giving testing accuracy of 87.8%. The previous best accuracy is 88.4% (Tomar et al., 2017). Our multi-cascaded model beats this result without any data augmentation with an accuracy of 89.4%. As seen from the table, our enhanced model outperforms all previous results. Ensemble BiMPM model achieves an accuracy of 88.2% accuracy while pt-DECATT_{char} shows a slightly better performance with an accuracy of 88.4%. In comparison, our model achieves an accuracy of 90.3%, which is almost 2% improvement over the previous best performance. In contrast to Wang et al. (2017), we avoid using ensemble model approach (which is computationally costly).

Table 6
Comparison of our model's performance with previously published performances on Quora dataset.

Model	Accuracy
(Saimese-CNN) Wang et al. (2017)	79.6
(Multi-Perspective-CNN) Wang et al. (2017)	81.4
(Saimese-LSTM) Wang et al. (2017)	82.58
(Multi-Perspective-LSTM) Wang et al. (2017)	83.2
(L.D.C) Wang et al. (2017)	85.6
(BiMPM) Wang et al. (2017)	88.2
(pt-DECATT _{word}) Tomar et al. (2017)	87.5
(pt-DECATT _{char}) Tomar et al. (2017)	88.4
(SSE) Lan and Xu (2018)	87.8
Our model	90.3

Table 7
Paraphrase detection performance on MSRP dataset.

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
None	74.1	75.3	90.9	82.4	74.7	76.2	90.2	82.6
P2, NP1	74.8	76.9	88.9	82.5	75.2	77.1	89.2	82.7
P2, P3	74.4	76.0	89.9	82.3	75.0	76.8	89.5	82.7
P2, P3, P1	76.8	77.1	92.7	84.2	77.4	77.0	91.7	84.2
P2, P3, NP1	74.4	75.6	90.7	82.5	74.1	74.7	92.5	82.6
P2, P3, P1, NP1	77.0	77.3	92.7	84.3	78.3	79.3	91.0	84.8
P2, P3, P1, NP1, NP2	77.0	77.3	92.7	84.3	78.3	79.3	91.0	84.8

Similarly, contrary to results in Tomar et al. (2017), our results are based on word features only and do not use computationally expensive character-based features.

6.2. MSRP dataset

We use the pre-defined split provided in MSRP dataset for training and testing our model. No dev set is provided with this dataset; hence, we hold-out 10% of the training split randomly as dev set. By using a grid search, we find optimal hyper-parameters on this dev set. ELMo embeddings are found to be better for this dataset, while Adam optimizer is selected as optimal one with learning rate of 0.002. When optimizing dropout rate on this dataset, 0.5 is found to yield the best results.

Table 7 shows the predictive performance of our enhanced paraphrase detection model on MSRP dataset. Performances are given for configurations with and without linguistic features after applying various data augmentation steps. It is observed that without data augmentation, we achieve an F1-score of 82.4%. Doubling the data (step P2 and NP1) increases performance slightly in terms of F1-score but the significant gain is obtained when augmenting the data using symmetry, transitivity, and reflexivity (steps P2, P3, P1). This configuration of augmentation gives F1-score of 84.2%. The highest F1-score in experiments without any linguistic features of 84.3% is achieved when P2, P3, P1, and NP1 augmentation steps are performed. Please note that adding a pair generated by NP2 augmentation did not affect the performance of the model as it only produces 1 additional pair (recall Section 5.1.2).

Using linguistic features along with learned features boosts the performance of the model. In comparison with experiments without linguistic features, this gain is consistent with the exception of P2, P3 and P1 augmentation scheme, where the F1-score for experiments without and with linguistic features remains the same. Without any augmentation, the F1-score is increased from 82.4% to 82.6% by introducing linguistic features, while with P2 and NP1 augmentation scheme, it is improved from 82.5 to 82.7. We achieve maximum performance using data augmentation schemes of P2, P3, P1, and NP1 while using linguistic features. This configuration yields an accuracy of 78.3% and an F1-Score of 84.8%, which is the highest among all of our experiments. Note that this augmentation scheme also has the highest F1-score when no linguistic features were used. Adding one additional pair generated by NP2 augmentation did not affect the model's performance.

These results prove that the usefulness of linguistic features is dataset domain and size-specific, and is not generalizable. In particular, the impact of linguistic features is limited when the dataset is large (e.g., Quora dataset) while it is more significant for small datasets (e.g., MSRP dataset).

We compare our best performing variation of experiments with existing state of the art approaches on MSRP dataset in Table 8.

Table 8
Comparison of our model's performance with previously published performances on MSRP dataset.

Model	Accuracy	F1-score
Socher et al. (2011)	76.8	83.6
Madnani, Tetreault, and Chodorow (2012)	77.4	84.1
Ji and Eisenstein (2013)	77.8	84.3
Hu et al. (2014)	69.6	80.3
Hu et al. (2014)	69.9	80.9
El-Alfy et al. (2015)	73.9	81.2
Kenter and De Rijke (2015)	76.6	83.9
Eyecioglu and Keller (2015)	74.4	82.2
He et al. (2015)	78.6	84.7
Dey et al. (2016)	–	82.5
Wang et al. (2016b)	78.4	84.7
Yin et al. (2016)	78.9	84.8
Pagliardini, Gupta, and Jaggi (2018)	76.4	83.4
Ferreira et al. (2018)	74.08	83.1
Agarwal et al. (2018)	77.7	84.5
Arora and Kansal (2019)	79.0	–
Our model	78.3	84.8

Table 9
Paraphrase detection performance on SemEval dataset.

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
None	51.1	27.1	78.9	40.2	89.0	70.8	80.6	75.4
P2, NP1	78.8	48.5	21.1	29.4	88.9	82.5	59.4	69.1
P2, P3	82.5	57.7	60.0	58.9	84.6	63.7	61.1	62.4
P2, P3, P1	82.2	56.1	68.6	61.7	87.6	77.5	57.1	65.8
P2, P3, NP1	84.4	62.6	62.3	62.5	84.0	60.5	67.4	63.8
P2, P3, P1, NP1	84.1	60.7	68.0	64.2	84.4	60.0	76.0	67.0
P2, P3, P1, NP1, NP2	82.7	57.1	68.6	62.3	84.7	62.4	67.4	64.8

The current state of the art on MSRP is reported to be in Ji and Eisenstein (2013). They report their best results as 80.4% accuracy and an F1-score of 85.9%. However, they assume that they have access to testing data at the time of training a model and call it a form of transductive learning. On the other hand, our model is based on inductive learning where training is done in total isolation from the test split. Therefore, it is only fair to compare both models in inductive setup. The training of the model by Ji and Eisenstein (2013) in inductive setup yields 77.8% accuracy and F1-score of 84.3% (Agarwal et al., 2018). These results are far less than what were originally reported in Ji and Eisenstein (2013) using transductive learning. Thus, the current state-of-the-art results in inductive setup are reported by He, Gimpel, and Lin (2015) and Wang et al. (2016b), with an F1-score of 84.7%. Our best model yields an F1-score of 84.8%, which is slightly higher than previous state-of-the-art.

In terms of accuracy, the highest performance is reported in Arora and Kansal (2019) which is 79.0%. However, the F1-score was not reported by the authors. As the class label distribution is highly skewed in this dataset (recall Table 1), the accuracy is not a good measure of performance. In such cases, it is plausible to use F1-score to measure the predictive performance of the models (Sokolova & Lapalme, 2009). Therefore, despite higher accuracy reported by the authors, it cannot be concluded that their model yields higher performance than other reported results.

6.3. SemEval dataset

In SemEval dataset, the provided dev split is used to fine-tune hyper-parameters using grid search. We find that ELMo embeddings yield the best results when Adam is used as optimizer with learning rate 0.002. The dropout rate is found to be 0.2.

Table 9 presents the predictive performance of our model on this dataset. In this table, data augmentation is done with full transitive closure, i.e., $K = *$ for step P3. Without data augmentation and linguistic features, our model achieves an F1-score of 40.2%. However, when we apply data augmentation, the F1-score tends to increase. The maximum F1-score without linguistic features is 64.2% which is achieved with P2, P3, P1, and NP1 augmentation. This is consistent with the results of other two datasets used in the study, which also yield maximum performance on this particular augmentation combination. In noisy text, linguistic features affect performance by a larger margin.

Our enhanced model outperforms existing approaches on SemEval dataset in terms of F1-score when no augmentation is done but linguistic features are provided along with learned features. This particular experiment yields Precision and Recall of 70.8% and 80.6%, respectively. However, it is noticed that precision is lower than recall. When we augment the data using P2 and P3, the precision and recall values tend to become closer, highlighting that model is more robust in distinguishing the two classes. But, the F1-score drops significantly.

The reason behind this degradation in performance can be found by investigating pairs generated by transitive closure. It is observed that as we move along the path in graph to find transitive pairs, the meaning of text tends to change, and it becomes more probable that the two documents are no longer paraphrases. This phenomenon is more likely in a noisy short text such as the SemEval dataset. Therefore, instead of applying transitive closure with $K = *$, transitive extension to an order K is plausible. To investigate the effect of K in transitive extension, we perform same experiments with $K = 1$, $K = 2$, and $K = 3$. The results of these experiments in terms of F1-score is given in Fig. 3 (plots (a) and (b) are for without and with linguistic features, respectively). Complete tables of results are included in Appendix A.

These results show that using linguistic features improves predictive performance as compared to when only learned features are used, and this improvement is consistent for all K and all data augmentation configurations. For order K of transitive extension, we note that moving beyond $K = 1$ yield minor improvements in performance. Without any linguistic features, $K = 1$ produces the highest F1-score overall, with exception of just one configuration, i.e., P2, P3, P1, NP1, NP2. Similarly, maximum F1-score with linguistic features is also achieved with $K = 1$. These observations prove that in noisy text, full transitive closure can produce lower performances and the order K of transitive extension needs to be investigated to determine the optimal data augmentation strategy. Without linguistic features, as opposed to $K = *$, augmenting data with NP2 using $K = 1, 2$ or 3 does not drop the predictive performance of the model drastically but rather shows a noticeable improvement. However, with inclusion of linguistic features, NP2 augmentation has variable effect on the performance for each order of K .

We compare our results with existing work in terms of precision, recall and F1-score on test split in Table 10. On this dataset,

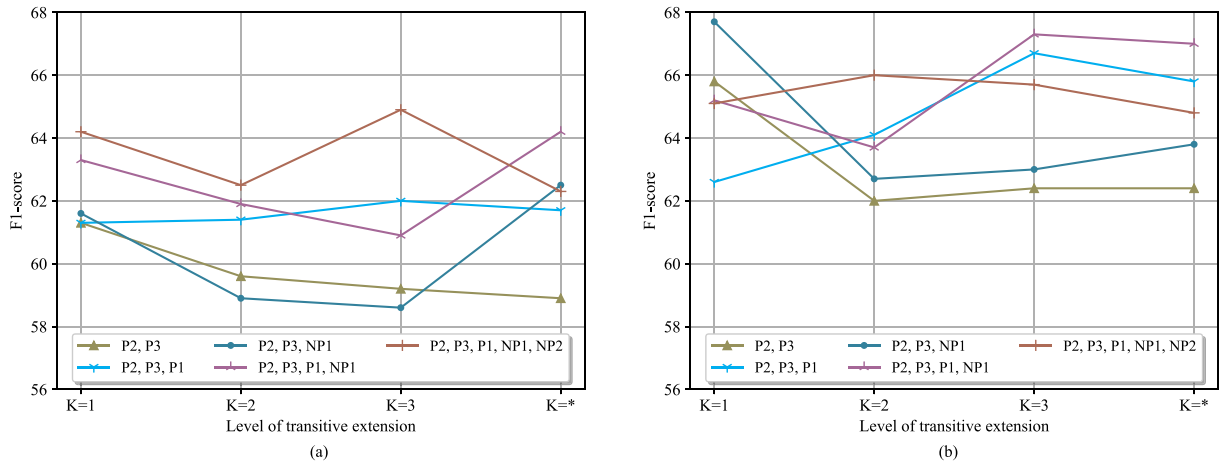


Fig. 3. Performance of our model with paraphrase transitive extensions of order $K = 1$, $K = 2$, $K = 3$, and $K = *$ on SemEval dataset; (a) learned features, (b) learned + linguistic features.

Table 10

Comparison of our model’s performance with previously published performances on SemEval dataset.

Model	Precision	Recall	F1-score
Das and Smith (2009)	62.9	63.2	63.0
Guo and Diab (2012)	58.3	52.5	65.5
Ji and Eisenstein (2013)	66.4	62.8	64.5
Xu et al. (2014)	72.2	72.6	72.4
Eyecioglu and Keller (2015)	68.0	66.9	67.4
Zarrella et al. (2015)	56.9	80.6	66.7
Zhao and Lan (2015)	76.7	58.3	66.2
Vo, Magnolini, and Popescu (2015)	68.5	63.4	65.9
Karan et al. (2015)	64.5	67.4	65.9
Dey et al. (2016)	75.6	72.6	74.1
Huang, Yao, Lyu, and Ji (2017)	64.3	65.7	65.0
Lan and Xu (2018)	-	-	65.6
Agarwal et al. (2018)	76.0	74.2	75.1
Our model	70.8	80.6	75.4

Lan and Xu (2018) in their comparative study of re-implementing 7 different models, found that Pairwise Word Interaction Model (PWIM) (He & Lin, 2016) performs the best and achieves F1-score of 65.6%, which we report in the table. State of the art performance on the SemEval dataset is reported by Agarwal et al. (2018) with F1-score of 75.1%. Our best performing model outperforms state of the art by achieving an F1-score of 75.4%.




Fig. 4. Accuracy versus training epochs for each feature learning cascades and discriminator network on (a) Quora, (b) MSRP, and (c) SemEval datasets.

6.4. Impact of multiple cascades

As all cascades are supervised, we can record predictive performances (on test set) for every cascade after each training epoch and compare them with that produced by discriminator network (complete multi-cascaded model). Fig. 4 shows the performance of every cascade and the complete multi-cascaded model on Quora (plot (a)), MSRP (plot (b)), and SemEval (plot (c)) datasets.

It can be observed that the LSTM-based feature learner yields least performance and remains at the bottom throughout the training, while attention-based convolution feature learner is the next best performer. The attention-based LSTM learner is closely following the discriminator's performance but discriminator remains at the top during the training process after every epoch. This confirms that training a discriminator based on features learned from multiple perspectives is more fruitful as compared to relying on features learned by only one type of deep learning model. The same trend is followed in all the datasets. These results are presented only for the configuration which yields best results on each dataset.

6.5. Summary

The extensive experiments evaluated our enhanced paraphrase detection model along several dimensions. These dimensions include noisy versus clean datasets, large versus small datasets, data augmentation steps and their variations, learned and linguistic features, and cascades in the multi-cascaded model. The key findings of our experiments are summarized below.

1. Data augmentation improves paraphrase detection predictive performance on all datasets (noisy, clean, large, and small). These easy steps can generate additional annotations that translate into the higher predictive performance from deep learning models.
2. Each step for generating additional paraphrase annotations produces an improvement in the prediction performance. On the other hand, only step NP1 for generating additional non-paraphrase annotations consistently improve performance; step NP2 can sometimes cause a decrease in predictive performance especially when the annotation is error-prone (e.g., the notion of paraphrase is not well defined, noisy text).
3. Linguistic features are important if a dataset is relatively small and noisy in nature. For such datasets, including linguistic features can produce significant boost in the predictive performance of the model.
4. When a dataset is sufficiently large, using linguistic features does not have any effect on the predictive performance.
5. For clean text, augmentation scheme of P2, P3, P1, and NP1 gives maximum performance in terms of F1-score on both datasets while for user-generated noisy text, this scheme yields maximum performance with learned features only. When linguistic features are used, maximum performance is achieved without any augmentation.
6. It is not recommended to use full transitive closure ($K = *$) for user-generated noisy datasets as no noticeable and consistent improvement in performance is observed. For such datasets, data augmentation with the transitive extension of order K should be investigated.

7. Conclusion

We present a data augmentation strategy and a multi-cascaded model for enhanced paraphrase detection. The data augmentation strategy generates additional paraphrase and non-paraphrase annotations based on the graph analysis of the existing annotations. This strategy is easy to implement and yields significant improvement in the performance of deep learning models for paraphrase detection. The multi-cascaded model employs multiple feature learners to encode and classify short text pairs. As such, it exploits multiple semantic cues to distinguish between paraphrases and non-paraphrases. The proposed multi-cascaded model is both deep and wide in architecture, and it embodies previous best practices in deep models for paraphrase detection.

We evaluate our enhanced model on three benchmark datasets representing noisy and clean text. Our model produces a higher predictive performance on all three datasets beating all previously published results on them. We also study the impact of different steps in data augmentation, the use of linguistic features in conjunction with learned features, and different deep models. The results show that data augmentation is generally beneficial and linguistic features only help for small and noisy text datasets. Furthermore, it is seen that multiple models can boost predictive performance beyond that achievable from any single model.

This work provides a comprehensive treatment of paraphrase detection that includes small and large datasets, clean and noisy texts, CNN and LSTM-based models, learned features, hand-crafted linguistic features, and a new data augmentation strategy. In the future, it would be beneficial to carry out a deeper analysis of conflicts arising from data augmentation and investigate strategies for resolving them in an effort to improve annotation quality in noisy data.

CRedit authorship contribution statement

Muhammad Haroon Shakeel: Conceptualization, Methodology, Software, Investigation, Visualization, Writing - original draft. **Asim Karim:** Conceptualization, Supervision, Project administration, Writing - review & editing. **Imdadullah Khan:** Conceptualization, Methodology, Supervision, Writing - review & editing.

Declaration of Competing Interest

The authors declare that they do not have any financial or nonfinancial conflict of interests.

Appendix A. Effect of paraphrase transitive extension levels on SemEval dataset

Tables A.11 –A.13.

Table A1

Paraphrase detection performance on SemEval dataset with $K = 1$.

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
P2, P3	82.8	57.9	65.1	61.3	86.8	71.3	61.1	65.8
P2, P3, P1	82.1	55.9	68.0	61.3	85.3	66.9	58.9	62.6
P2, P3, NP1	81.9	55.2	69.7	61.6	87.4	72.5	63.4	67.7
P2, P3, P1, NP1	82.8	57.1	70.9	63.3	86.8	72.2	59.4	65.2
P2, P3, P1, NP1, NP2	85.2	64.9	63.4	64.2	84.6	61.9	68.6	65.1

Table A2

Paraphrase detection performance on SemEval dataset with $K = 2$.

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
P2, P3	82.2	56.7	62.9	59.6	85.0	65.6	58.9	62.0
P2, P3, P1	81.1	53.6	72.0	61.4	83.3	58.1	71.4	64.1
P2, P3, NP1	82.0	56.3	61.7	58.9	85.1	65.6	60.0	62.7
P2, P3, P1, NP1	83.5	59.9	64.0	61.9	85.4	66.5	61.1	63.7
P2, P3, P1, NP1, NP2	84.2	62.1	62.9	62.5	83.2	57.1	78.3	66.0

Table A3

Paraphrase detection performance on SemEval dataset with $K = 3$

Augmentation	Learned features				Learned + Linguistic features			
	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score
P2, P3	82.6	57.9	60.6	59.2	86.0	63.7	61.1	62.4
P2, P3, P1	81.7	54.8	71.4	62.0	86.3	67.6	65.7	66.7
P2, P3, NP1	85.0	69.0	50.9	58.6	81.5	54.1	75.4	63.0
P2, P3, P1, NP1	86.0	73.4	52.0	60.9	86.8	69.5	65.1	67.3
P3, P2, P1, NP1, NP2	85.6	65.9	64.0	64.9	85.6	65.2	66.3	65.7

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.ipm.2020.102204](https://doi.org/10.1016/j.ipm.2020.102204).

References

- Agarwal, B., Ramampiaro, H., Langseth, H., & Ruocco, M. (2018). A deep network model for paraphrase detection in short text messages. *Information Processing & Management, (IPM)*, 54(6), 922–937. <https://doi.org/10.1016/j.ipm.2018.06.005>.
- Arora, M., & Kansal, V. (2019). Character level embedding with deep convolutional neural network for text normalization of unstructured data for twitter sentiment analysis. *Social Network Analysis and Mining, (SNAM)*, 9(1), 1–14. <https://doi.org/10.1007/s13278-019-0557-y>.
- Barrón-Cedeño, A., Vila, M., Martí, M. A., & Rosso, P. (2013). Plagiarism meets paraphrasing: Insights for the next generation in automatic plagiarism detection. *Computational Linguistics, (CL)*, 39(4), 917–947. https://doi.org/10.1162/COLI_a_00153.
- Bogdanova, D., dos Santos, C., Barbosa, L., & Zadrozny, B. (2015). Detecting semantically equivalent questions in online user forums. *Proceedings of the conference on computational natural language learning, (CNL)*123–131. <https://doi.org/10.18653/v1/K15-1013>.
- Cagnina, L., Errecalde, M., Ingaramo, D., & Rosso, P. (2014). An efficient particle swarm optimization approach to cluster short texts. *Information Sciences*, 265, 36–49. <https://doi.org/10.1016/j.ins.2014.05.015>.

- org/10.1016/j.ins.2013.12.010.
- Das, D., & Smith, N. A. (2009). Paraphrase identification as probabilistic quasi-synchronous recognition. *Proceedings of the joint conference of the annual meeting of the ACL and the international joint conference on natural language processing of the AFNLP, (ACL-IJCNLP) 1*. *Proceedings of the joint conference of the annual meeting of the ACL and the international joint conference on natural language processing of the AFNLP, (ACL-IJCNLP)* 468–476.
- Dey, K., Shrivastava, R., & Kaushik, S. (2016). A paraphrase and semantic similarity detection system for user generated short-text content on microblogs. *Proceedings of the international conference on computational linguistics, (COLING)*2880–2890.
- Dolan, B., Quirk, C., & Brockett, C. (2004). Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. *Proceedings of the international conference on computational linguistics, (COLING)*350–357.
- El-Alfy, E.-S. M., Abdel-Aal, R. E., Al-Khatib, W. G., & Alvi, F. (2015). Boosting paraphrase detection through textual similarity metrics with abductive networks. *Applied Soft Computing, (ASC)*, 26, 444–453. <https://doi.org/10.1016/j.asoc.2014.10.021>.
- Eyecioglu, A., & Keller, B. (2015). Twitter paraphrase identification with simple overlap features and SVMs. *Proceedings of the international workshop on semantic evaluation, (SemEval)* 64–69. <https://doi.org/10.18653/v1/S15-2011>.
- Fader, A., Zettlemoyer, L., & Etzioni, O. (2013). Paraphrase-driven learning for open question answering. *Proceedings of the annual meeting of the association for computational linguistics, (ACL) (long papers)1*. *Proceedings of the annual meeting of the association for computational linguistics, (ACL) (long papers)* 1608–1618.
- Ferreira, R., Cavalcanti, G. D., Freitas, F., Lins, R. D., Simske, S. J., & Riss, M. (2018). Combining sentence similarities measures to identify paraphrases. *Computer Speech & Language, (CSL)*, 47, 59–73. <https://doi.org/10.1016/j.csl.2017.07.002>.
- Figueroa, A., & Neumann, G. (2013). Learning to rank effective paraphrases from query logs for community question answering. *Proceedings of the AAAI conference on artificial intelligence, (AAAI)*1099–1105.
- Guo, W., & Diab, M. (2012). Modeling sentences in the latent space. *Proceedings of the annual meeting of the association for computational linguistics, (ACL): Long papers*Vol. 1. *Proceedings of the annual meeting of the association for computational linguistics, (ACL): Long papers* 864–872.
- He, H., Gimpel, K., & Lin, J. (2015). Multi-perspective sentence similarity modeling with convolutional neural networks. *Proceedings of the conference on empirical methods in natural language processing, (EMNLP)*1576–1586. <https://doi.org/10.18653/v1/D15-1181>.
- He, H., & Lin, J. (2016). Pairwise word interaction modeling with deep neural networks for semantic similarity measurement. *Proceedings of the conference of the north American chapter of the association for computational linguistics, (NAACL): Human language technologies*937–948. <https://doi.org/10.18653/v1/N16-1108>.
- Heikinheimo, H., & Ukkonen, A. (2013). The crowd-median algorithm. *Association for the advancement of artificial intelligence conference on human computation and crowdsourcing, (HCOMP)*69–77.
- Hu, B., Lu, Z., Li, H., & Chen, Q. (2014). Convolutional neural network architectures for matching natural language sentences. *Proceedings of the international conference on neural information processing systems, (NIPS)*2042–2050.
- Huang, J., Yao, S., Lyu, C., & Ji, D. (2017). Multi-granularity neural sentence model for measuring short text similarity. *International conference on database systems for advanced applications, (DASFAA)*439–455. https://doi.org/10.1007/978-3-319-55753-3_28.
- Ji, Y., & Eisenstein, J. (2013). Discriminative improvements to distributional sentence similarity. *Proceedings of the conference on empirical methods in natural language processing, (EMNLP)*891–896.
- Karan, M., Glavaš, G., Šnajder, J., Bašić, B. D., Vulić, I., & Moens, M.-F. (2015). Tklblir: Detecting twitter paraphrases with tweetingjay. *Proceedings of the international workshop on semantic evaluation, (SemEval)*70–74. <https://doi.org/10.18653/v1/S15-2012>.
- Kenter, T., & De Rijke, M. (2015). Short text similarity with word embeddings. *Proceedings of the international conference on information and knowledge management, (CIKM)*1411–1420.
- Kleindessner, M., & von Luxburg, U. (2017). Kernel functions based on triplet comparisons. *Proceedings of the international conference on neural information processing systems, (NIPS)* 6807–6817.
- Kleindessner, M., & Von Luxburg, U. (2017). Lens depth function and k-relative neighborhood graph: Versatile tool for ordinal data analysis. *Journal of Machine Learning Research, (JMLR)*, 18(1), 1889–1940.
- Lan, W., & Xu, W. (2018). Neural network models for paraphrase identification, semantic textual similarity, natural language inference, and question answering. *Proceedings of the international conference on computational linguistics, (COLING)*3890–3902.
- Madnani, N., Tetreault, J., & Chodorow, M. (2012). Re-examining machine translation metrics for paraphrase identification. *Proceedings of the conference of the north American chapter of the association for computational linguistics, (NAACL): Human language technologies*182–190.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Proceedings of the international conference on neural information processing systems, (NIPS)*3111–3119.
- Mohammad, A.-S., Jaradat, Z., Mahmoud, A.-A., & Jararweh, Y. (2017). Paraphrase identification and semantic text similarity analysis in arabic news tweets using lexical, syntactic, and semantic features. *Information Processing & Management, (IPM)*, 53(3), 640–652. <https://doi.org/10.1016/j.ipm.2017.01.002>.
- Nie, Y., & Bansal, M. (2017). Shortcut-stacked sentence encoders for multi-domain inference. *Proceedings of the workshop on evaluating vector space representations for NLP, (repeval@emnlp)*41–45. <https://doi.org/10.18653/v1/W17-5308>.
- Oliva, J., Serrano, J. I., del Castillo, M. D., & Iglesias, A. (2011). Syms: A syntax-based measure for short-text semantic similarity. *Data and Knowledge Engineering, (DKE)*, 70(4), 390–405. <https://doi.org/10.1016/j.datak.2011.01.002>.
- Pagliardini, M., Gupta, P., & Jaggi, M. (2018). Unsupervised learning of sentence embeddings using compositional n-gram features. *Proceedings of the conference of the north American chapter of the association for computational linguistics, (NAACL): Human language technologies*528–540. <https://doi.org/10.18653/v1/N18-1049>.
- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. *Proceedings of the conference on empirical methods in natural language processing, (EMNLP)*1532–1543. <https://doi.org/10.3115/v1/D14-1162>.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *Proceedings of the conference of the north American chapter of the association for computational linguistics, (NAACL): Human language technologies*2227–2237. <https://doi.org/10.18653/v1/N18-1202>.
- Rashid, J., Shah, S. M. A., & Irtaza, A. (2019). Fuzzy topic modeling approach for text mining over short text. *Information Processing & Management, (IPM)*, 56(6), 102060. <https://doi.org/10.1016/j.ipm.2019.102060>.
- Reimers, N., & Gurevych, I. (2017). Reporting score distributions makes a difference: Performance study of LSTM-networks for sequence tagging. *Proceedings of the conference on empirical methods in natural language processing, (EMNLP)*338–348. <https://doi.org/10.18653/v1/D17-1035>.
- dos Santos, C., Barbosa, L., Bogdanova, D., & Zadrozny, B. (2015). Learning hybrid representations to retrieve semantically equivalent questions. *Proceedings of the annual meeting of the association for computational linguistics and the 7th international joint conference on natural language processing, (ACL-IJCNLP)*694–699. <https://doi.org/10.3115/v1/P15-2114>.
- Shakeel, M. H., Karim, A., & Khan, I. (2019). A Multi-cascaded Deep Model for Bilingual SMS Classification. *International Conference on Neural Information Processing, (ICONIP)*. Springer287–298.
- Socher, R., Huang, E. H., Pennin, J., Manning, C. D., & Ng, A. Y. (2011). Dynamic pooling and unfolding recursive autoencoders for paraphrase detection. *Proceedings of the international conference on neural information processing systems, (NIPS)*801–809.
- Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management, (IPM)*, 45(4), 427–437. <https://doi.org/10.1016/j.ipm.2009.03.002>.
- Tomar, G. S., Duque, T., Täckström, O., Uszkoreit, J., & Das, D. (2017). Neural paraphrase identification of questions with noisy pretraining. *Proceedings of the workshop on subword and character level models in NLP, (SCLeM)*142–147. <https://doi.org/10.18653/v1/W17-4121>.
- Vila, M., Martí, M. A., & Rodríguez, H. (2014). Is this a paraphrase? What kind? Paraphrase boundaries and typology. *Open Journal of Modern Linguistics, (OJML)*, 4(01), 205–218. <https://doi.org/10.4236/ojml.2014.41016>.
- Vo, N. P. A., Magnolini, S., & Popescu, O. (2015). Paraphrase identification and semantic similarity in twitter with simple features. *Proceedings of the international workshop on natural language processing for social media, (Social NLP)*10–19. <https://doi.org/10.3115/v1/W15-1702>.
- Wang, C., Duan, N., Zhou, M., & Zhang, M. (2013). Paraphrasing adaptation for web search ranking. *Proceedings of the annual meeting of the association for computational linguistics, (ACL) (short papers)2*. *Proceedings of the annual meeting of the association for computational linguistics, (ACL) (short papers)* 41–46.
- Wang, J., & Perez, L. (2017). The effectiveness of data augmentation in image classification using deep learning. *Convolutional Neural Networks for Visual Recognition, (CNNVR)*, 1–8.

- Wang, X., Jiang, W., & Luo, Z. (Jiang, Luo, 2016a). *Combination of convolutional and recurrent neural network for sentiment analysis of short texts. Proceedings of the international conference on computational linguistics, (COLING): Technical papers*2428–2437.
- Wang, Z., Hamza, W., & Florian, R. (2017). *Bilateral multi-perspective matching for natural language sentences. Proceedings of the international joint conference on artificial intelligence, (IJCAI)*4144–4150. <https://doi.org/10.24963/ijcai.2017/579>.
- Wang, Z., Mi, H., & Ittycheriah, A. (Mi, Ittycheriah, 2016b). *Sentence similarity learning by lexical decomposition and composition. Proceedings of the international conference on computational linguistics, (COLING)*1340–1349.
- Xu, W., Callison-Burch, C., & Dolan, B. (2015). *SemEval-2015 task 1: Paraphrase and semantic similarity in twitter (PIT). Proceedings of the international workshop on semantic evaluation, (SemEval)*1–11. <https://doi.org/10.18653/v1/S15-2001>.
- Xu, W., Ritter, A., Callison-Burch, C., Dolan, W. B., & Ji, Y. (2014). *Extracting lexically divergent paraphrases from twitter. Transactions of the Association for Computational Linguistics, (TACL)*, 2, 435–448. https://doi.org/10.1162/tacl_a_00194.
- Yin, W., Schütze, H., Xiang, B., & Zhou, B. (2016). *Abcnn: Attention-based convolutional neural network for modeling sentence pairs. Transactions of the Association for Computational Linguistics, (TACL)*, 4, 259–272. https://doi.org/10.1162/tacl_a_00097.
- Zarella, G., Henderson, J., Merkhofer, E. M., & Strickhart, L. (2015). *Mitre: Seven systems for semantic similarity in tweets. Proceedings of the international workshop on semantic evaluation, (SemEval)*12–17. <https://doi.org/10.18653/v1/S15-2002>.
- Zhao, J., & Lan, M. (2015). *Ecnv: Leveraging word embeddings to boost performance for paraphrase in twitter. Proceedings of the international workshop on semantic evaluation, (SemEval)* 34–39. <https://doi.org/10.18653/v1/S15-2006>.