

Boosting Entity Mention Detection for Targetted Twitter Streams with Global Contextual Embeddings

Sataisha Saha Bhownick

Department of Computer Science
Binghamton University
Email: ssahab1@binghamton.edu

Eduard C. Dragut

Department of Computer Science
Temple University
Email: edragut@temple.edu

Weiyi Meng

Department of Computer Science
Binghamton University
Email: meng@binghamton.edu

Abstract—Microblogging sites, like Twitter, have emerged as ubiquitous sources of information. Two important tasks related to the automatic extraction and analysis of information in Microblogs are Entity Mention Detection (EMD) and Entity Detection (ED). The state-of-the-art EMD systems aim to model the non-literary nature of microblog text by training upon offline static datasets. They extract a combination of surface-level features – orthographic, lexical, and semantic – from individual messages for noisy text modeling and entity extraction. But given the constantly evolving nature of microblog streams, detecting all entity mentions from such varying yet limited context of short messages remains a difficult problem. To this end, we propose a framework named EMD Globalizer, better suited for the execution of EMD learners on microblog streams. It deviates from the processing of isolated microblog messages by existing EMD systems, where learned knowledge from the immediate context of a message is used to suggest entities. Instead, it recognizes that messages within a microblog stream are typically related and often repeat entity mentions, thereby leaving the scope for EMD systems to go beyond the localized processing of individual messages. After an initial extraction of entity candidates by an EMD system, the proposed framework leverages occurrence mining to find additional candidate mentions that are missed during this first detection. Aggregating the local contextual representations of these mentions, a global embedding is drawn from the collective context of an entity candidate within a stream. The global embeddings are then utilized to separate entities within the candidates from false positives. All mentions of said entities from the stream are produced in the framework’s final outputs. Our experiments show that EMD Globalizer can enhance the effectiveness of all existing EMD systems that we tested (on average by 25.61%) with a small additional computational overhead.

I. INTRODUCTION

Entity Mention Detection (EMD) is the task of extracting contiguous strings within text that represent entities of interest. These strings (also known as surface forms) are referred to as Entity Mentions (EMs). The benchmarking guidelines set by WNUT [1] identifies EMD and Entity Detection (ED) as two concomitant tasks in this context. ED aims to cover the range of unique entities within text, while EMD compiles the string variations of entities from the text. Together, they form the broader problem of Named Entity Recognition (NER) that has received significant research attention. In this paper, we focus on maximizing effectiveness of state-of-the-art EMD techniques for the microblog streaming environment.

Example 1. Tweets in Figure 1 have entity mentions (in many string variations) from six unique entities: ‘*beshear*’ in T1 and T4, ‘*italy*’ in T2 and T6, ‘*coronavirus*’ in T2, T3 and T5, ‘*trump*’ in T5, ‘*US*’ in T5 and ‘*canada*’ in T6.

Off-the-shelf EMD solutions typically range from statistical Machine Learning models [2], [3], [4], [5], [6] to Deep Neural Networks (DNNs) [7], [8], [9]. However, the commonality among EMD systems focussing on Microblogs resides in their design and offline training process on static datasets. These datasets are curated from randomly sampled messages. As such, they provide a good representation of the non-literary language used in these platforms. Microblog EMD systems primarily study the nuances of its noisy text. They rightfully identify the lack of adherence to proper grammar, syntax or spellings as a key challenge to be addressed for language tasks. To extract contextual information from messages, they use a combination of surface level information – word embeddings, lexical, orthographic and/or semantic information (POS tags, dependency trees), and sometimes even external gazetteers.

For EMD from Microblog streams, existing systems do not take into account any of the streaming environment’s defining traits and simply extend their processing approach for offline datasets. More precisely, these systems run each message in the stream through their EMD pipeline in isolation, one at a time, in the order of its entry into the stream. But given the underspecified context of a single, character-limited message, added to the constantly evolving nature of entity mentions within a microblog stream, detecting every mention of an entity from the stream remains a difficult problem to generalize. The varying textual context where EMs reside in messages is often further complicated by the rarity of many microblog-referenced entities in off-the-shelf lexical resources. This makes it more difficult to consistently extract mentions of novel entities for most EMD tools [3], [5], including even the most effective Deep Learners [8]. To illustrate these problems, we perform EMD on a message stream discussing the most prevalent conversation topic of 2020 – the Coronavirus. We used two existing deep EMD baselines for this task: Aguilar et al. [8] and BERTweet [10] (finetuned for EMD).

A Case Study. The objective of this study is to explore the performance of a state-of-the-art deep EMD tool on a microblog stream and understand its limitations. We run both baselines on a streaming dataset of 2K tweets (D_2 , see Table

Entity Mentions	Aguilar et al. labelled EMs	Entity Mentions	BERTweet labelled EMs
T1: Beshear : Social distancing is not social isolation.		T1: Beshear : Social distancing is not social isolation.	
T2: WE JUST BY-PASS Italy WITH CORONAVIRUS CASES. But @realDonaldTrump wants to relax social distancing.		T2: WE JUST BY-PASS Italy WITH CORONAVIRUS CASES. But @realDonaldTrump wants to relax social distancing.	
T3: Not a bad video to explain how the Coronavirus works as well as the reasoning for social distancing.		T3: Not a bad video to explain how the Coronavirus works as well as the reasoning for social distancing.	
T4: Beshear says he's asking county judges and city mayors to monitor parks and shut them down if people aren't filling		T4: Beshear says he's asking county judges and city mayors to monitor parks and shut them down if people aren't filling	
T5: Trump to rank US counties by Coronavirus risk, may 'relax' social distancing.		T5: Trump to rank US counties by Coronavirus risk, may 'relax' social distancing.	
T6: Canada is rising at a rate similar to the early days in ITALY		T6: Canada is rising at a rate similar to the early days in ITALY	

Fig. 1: EMD on a message stream discussing Coronavirus

1 in Section VI) generated from a Coronavirus tweet stream. We apply their production versions directly in this setting. BERTweet yields a modest performance on this stream subset with an F1 score of 53%. Aguilar et al. fared better at 60% with its reliance on updated Twitter-trained word embeddings and gazetteers, to better cover some rare entities. Apart from missing some entities altogether, both systems show inconsistency in extracting EMs of the same entity throughout the stream, detected in some tweets but missed in others.

Takeaways. A closer look at the EMD results of both systems for tweets in the Coronavirus dataset in Figure 1 shows that they often missed mentions of one of the most important and frequent entities in this stream, i.e. ‘Coronavirus’. Other EMs that frequently came up in the stream but were also frequently missed include ‘Italy’ and ‘US’. For example, the entity ‘Coronavirus’ has three mention variations in the tweets in Figure 1 but only ‘Coronavirus’ was successfully detected while ‘CORONAVIRUS’ and ‘coronavirus’ were not. Note that the problem here is not merely the inability to identify the same entity string across different syntactic formats. It is rather the varying contexts (both syntactic and semantic) in which entities are repeated within a stream that adds to the challenge of generalizing an EMD algorithm that works well across all possible message contexts.

The appearance of entities in their many mention variations is an integral part of the microblog streaming ecosystem that constantly generates messages on multiple contemporaneous topics, i.e. conversation streams, evolving over time [11]. The failure to consistently identify these entity mentions leads to reduced EMD performance. This appears to be a critical weakness of many state-of-the-art EMD techniques. The common solutions to cope with this problem are more training or transfer learning (e.g., fine-tuning) on messages from newer topic streams. But as content in microblog streams evolves or spearheads into different topics, one needs to constantly annotate datasets specific to these emerging topics. This is not a scalable proposition. To this end, we design a framework that boosts the ability of EMD systems to recognize EMs more consistently across contexts.

Local vs Global Context. Majority of EMD systems encode the tokens of an input message with a variety of surface-level information to extract entity mentions. Pre-deep learning sys-

tems directly use this information to perform sequence labeling for EMD and identify entity boundaries within text using some variation of the BIO (beginning-inside-outside) encoding. Deep Neural Networks performing EMD however, use it to generate ‘*contextual word embedding*’ – a representation of a token, in the context that it appears within a sentence. These contextual embeddings are then used for the downstream task of sequence labeling. Irrespective of architectural differences or the resources used to gather contextual information, these systems ultimately follow a traditional computation philosophy of treating individual microblog messages as isolated points of input that are individually processed to generate EMD outputs. We call this the ‘**Local EMD**’ approach. This treats message streams as any other offline dataset – an unchanging collection of messages, not a medium of incremental and topically-related message generation over an extended period of time. Given the various noisy but limited contextual possibilities of microblog messages, it can be untenable to individually analyze them and produce consistent labeling [12]. This provides the motivation to move beyond the localized context of a message and dynamically aggregate token-level local contextual embeddings from the entire stream, which are then used to derive a pooled ‘global’ context for every token encountered within a dataset [13]. The global contextual embeddings are then concatenated to the sentence-level local embeddings for the eventual sequence labeling task. Expanding on this idea we argue that, more so than offline datasets, microblog streams are aptly positioned to collectively view embeddings. Messages within a conversation stream not only repeat a finite set of entities but also the context in which they appear, owing to inter-message topical correlation. Hence global contextual representations of tokens, or rather entity candidates, can be effectively mined and used for EMD in this setting.

Approach Overview. Here we propose the EMD Globalizer framework. It begins with a traditional EMD system that extracts local contextual information for each individual message and uses them to extract entities from messages. We call this step ‘**Local EMD**’, due to the contextual information and inference drawn being locally confined. However, as evidenced before, local EMD tends to be inconsistent in providing the best message-level entity coverage. Hence its EM outputs are not instantly trusted. In our approach, the EM outputs from local EMD are used to generate a set of seed entity candidates. Additionally, in case of deep EMD systems, the token-level contextual word embeddings generated for every message are also stored. We follow this up by an occurrence mining step that finds additional mentions of seed candidates that are initially missed at Local EMD. Whenever an instance of a seed candidate is found, the local contextual information generated from its mention is aggregated to incrementally construct a global candidate embedding. For deep EMD systems, the contextual embeddings generated during Local EMD for the candidate’s tokens are passed through a ‘*Phrase Embedder*’ that converts token-level embeddings into an embedding for the entire candidate string. Non-deep EMD systems, however, do not provide token-level representations and here we resort

to extracting a syntactic embedding of the candidate mention depending on its immediate context as shown in [14]. Note that in either case, these candidate embeddings still capture only local contexts until this point. For all mentions of a seed candidate, the local candidate embeddings are aggregated to form a pooled global embedding. Global candidate embeddings are then passed through an ‘*Entity Classifier*’ to separate entities in the seed set from false positives that arose during Local EMD. All mentions of such discovered entities are considered valid entity mentions and produced as outputs. The steps following Local EMD up to the identification of true positives by the classifier together constitute what we call ‘**Global EMD**’. By decoupling the local EMD step from the global one, we arrive at a stream-aware EMD framework, that can plugin any existing EMD algorithm without training modification/fine-tuning and still enhance its EMD performance within a stream.

Our experiments show that EMD Globalizer effectively performs EMD, especially on microblog streams. We test it with four different EMD systems, including two state-of-the-art deep EMD networks, for local EMD. In each case, the effectiveness of an EMD system was significantly boosted (on average by 25.61% across all datasets) when plugged into the framework. The framework also surpasses the best EMD performance on existing benchmarking datasets. The uniqueness of the framework is that it can accommodate a variety of (local) EMD systems with no algorithmic modification and still achieve more consistent EM labeling across the stream.

This paper makes the following contributions:

- We propose a novel framework for EMD in microblog streams. It consists of a Local EMD step, followed by a Global EMD step that includes an Entity Classifier. Our framework can accommodate both pre deep learning EMD systems and deep EMD systems and effectively collectivize EMD information in either case for Global EMD. It supports continuous and incremental computation which is in tune with the message generation process of streams.
- The local EMD step is decoupled from the rest of the pipeline. This allows us to test the hypothesis of collectivising local embeddings to generate better performance for different (local) EMD systems. We demonstrate the framework’s impact on several state-of-the-art instantiations of the Local EMD step. The contribution of this framework is therefore larger than catering to any single EMD technology.
- We conduct extensive experiments using multiple Twitter datasets to test the proposed framework. We use both in-house streaming datasets to simulate EMD from Twitter streams and third party datasets curated from microblogs.

The complete implementation of EMD Globalizer and all the experimental datasets are available at <https://github.com/satadisha/collective-EMD-framework>.

II. RELATED WORK

The principal issues surrounding the EMD problem from Microblogs are identified in [1] to be the lack of annotated data from this domain, and congruently, the difficulty in identifying emergent entity forms. The EMD literature features a wide

range of supervised systems with either handcrafted linguistic features in a non-neural system or DNNs with minimal feature engineering. The first category of systems like [3], [15], [16] recreate an information extraction pipeline in the form of POS-taggers or noun phrase chunker to extract and feed semantic features to a CRF based entity boundary segmentation module. In some systems [17], [18], [19], [20], [21], [22], [23], the feature set is further enhanced with word embeddings and gazetteers to better suit the limited yet diverse contextual possibilities of microblogs and the rare tokens that inhabit the medium. With advances made in Deep Learning, many deep neural networks (DNNs) [24], [25], [26], [7], [27], [28], [29] have been adopted for the sequence labeling task of NER. The recent WNUT shared task reports [30], [1] delve into a variety of neural architectures specifically designed for Entity Extraction from Tweets. We have chosen Aguilar et al. [8] – a BiLSTM-CNN-CRF architecture that performed best at the WNUT17 task, and BERTweet [10] – a BERT model trained on a large Twitter corpus that we finetune with the WNUT17 training data for EMD, as two of our Local EMD systems. [31], [32] examine the cross-domain transferability of DNN features learned from the more abundantly labelled well-formatted corpora to overcome the lack of annotated data.

Few other alternatives include TwiNER [33], an unsupervised approach using collocation features prepared from the Microsoft Web N-Gram corpus [34] or joint modeling of NER [35], [36], [37] with another information extraction subtask.

The concept of globalizing EMD computation is encountered in traditional NER pipelines for documents. Unlike a stream of tweets produced by multiple authors, documents are structurally more cohesive and contain well-formatted language. But both repeat entities and tokens through the collective span of their contents. Document-level EMD systems like HIRE-NER [38] utilize this tendency to distill non-local information for each unique token, from the entire scope of the document, using a memory-structure and append them to sentence-level contextual embeddings before an EMD decoder draws final output labels. DocL-NER [39] additionally includes a label refinement network to enforce label consistency across documents and improve EMD results. We compare EMD Globalizer with HIRE-NER [38] to test how effectively global information for EMD is compiled in each system.

For microblogs, TwiCS [14] explores the feature of entity mention repetition to efficiently perform EMD on streams. Using a shallow syntactic heuristic it identifies entity candidates and collectively generates syntactic support from their mentions within the entire stream to distinguish legitimate entities from false positives.

Our proposed EMD Globalizer further extends this idea of looking beyond the modeling of singular sentences in a stream. It combines the potential of collective processing of microblogs from a stream with existing EMD techniques that, despite offering robust EMD generalization, are constrained to processing sentences individually. Unlike TwiCS and other standalone EMD systems, what we propose in this paper is not a standalone system but a general EMD framework that aims

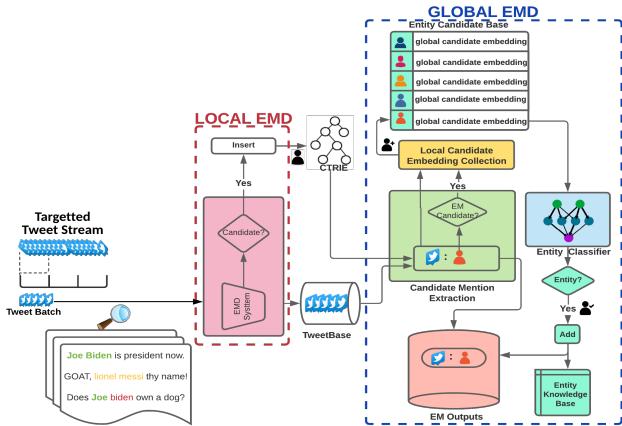


Fig. 2: EMD Globalizer System Architecture

to enhance the effectiveness of existing (local) EMD systems, when applied to microblog streams. Furthermore, in this work, we integrate a diverse set of Local EMD systems into our framework, with a special design focus on Deep EMD learners. TwiCS [14] is not featured as a Local EMD system in our experiments since it does not process sentences individually.

III. SYSTEM ARCHITECTURE

Figure 2 illustrates the overall architecture of EMD Globalizer. Note that, depending on the type of system employed for Local EMD, the components of the rest of the framework are adjusted accordingly. The EMD process in this framework facilitates continuous execution of a tweet stream over multiple iterations. Each iteration consists of a batch of incoming tweets thereby discretizing the evolution of messages within the stream. A single execution cycle through this framework can be divided into the following steps:

- (1) The Streaming module fetches a stream of tweets, on a particular topic, using the Twitter streaming API.
- (2) First we run a batch of tweets through an off-the-shelf (deep or non-deep) EMD system, one sentence at a time, in the **Local EMD** step. Phrases labelled as possible entities are registered as entity candidates. Further, in case of a deep EMD system, the output of the neural network’s final layer before token-level classification are stored, for every tweet in the batch, as ‘entity-aware embeddings’ of sentence tokens.
- (3) Next we initiate **Global EMD**. This includes a few added steps where individual framework components are adjusted according to the type of local EMD system inserted into the framework: (i) First, an additional scan of the tweet batch extracts all mentions of the entity candidates that have been discovered so far. This involves finding candidate mentions that were missed by the local EMD system in the previous step, along with the ones that were already found. (ii) For every mention we find, a candidate embedding is constructed based only on the immediate local information. With a regular (i.e., non-deep) local EMD system, we construct a syntactic embedding for the candidate mention from its immediate context. With a deep local EMD system however, token-level contextual embeddings are also available at the end of the local EMD. Hence, in this case, the token-level embeddings

for the candidate mention phrase are passed through the Phrase Embedder to obtain a unified contextual embedding for the entire phrase. (iii) Local candidate embeddings of every mention of an entity candidate found within a batch of tweets are aggregated to generate the candidate’s pooled global embedding. The global embedding can be incrementally updated by adding local embeddings into the pool as and when new mentions arrive. (iv) The final step is to pass global candidate embeddings through the entity classifier to separate the entities within the seed candidates from non-entities. Mentions of candidates that get labelled as entities are produced as valid entity mentions in the system’s final EMD outputs for the tweet batch.

We elaborate on these steps in later sections. The framework, when initialized with a deep neural network for local EMD, consists of a few additional steps. We zoom into it separately from the overall system architecture in Figure 3.

IV. LOCAL EMD

The Local EMD step can be any existing EMD algorithm that processes every single tweet-sentence in a stream, or in a tweet batch, individually and indicates likely entity mentions. A variety of existing systems can be plugged into the framework at the local phase. For majority of these systems, the EMD process is designed as a sequence labeling task where each token is tagged relative to its nearest entity boundary by adopting a variation of the BIO encoding. To facilitate a token’s labeling, local information relative to the token is generated and used. In case of deep learners, this happens to be a token-level contextual embedding obtained at the penultimate layer of the deep neural network, prior to the generation of output labels. For non-deep systems, this can be rich token information like POS-tags or gazetteer features that can aid the labeling process. Therefore, we interpret it as an encoding of the local entity-aware information, extracted from within the context of a single input.

Objectives: For a targeted stream of tweets, Local EMD aims to: (1) identify substrings from individual sentences as mentions of potential entities, (2) encode local entity information for every token in individual sentences, and (3) generate a set of seed entity candidates from tagged mentions.

A. Instantiations

There are different ways to instantiate the Local EMD step. For EMD Globalizer, most off-the-shelf EMD systems that process sentences individually would qualify. We test with four EMD systems, each of which supports a different EMD extraction algorithm, including two state-of-the-art deep learners. We now briefly describe each of these instantiations. Note that Local EMD systems are inserted as blackbox within the framework without any technical alteration during testing.

1. Chunker Based EMD – TweeboParser: This Local EMD system is a dependency parser [40] trained on English Tweets drawn from the POS-tagged tweet corpus of [41]. We use the production version of TweeboParser to generate POS-tags and dependency trees that capture the semantic structure of tweets.

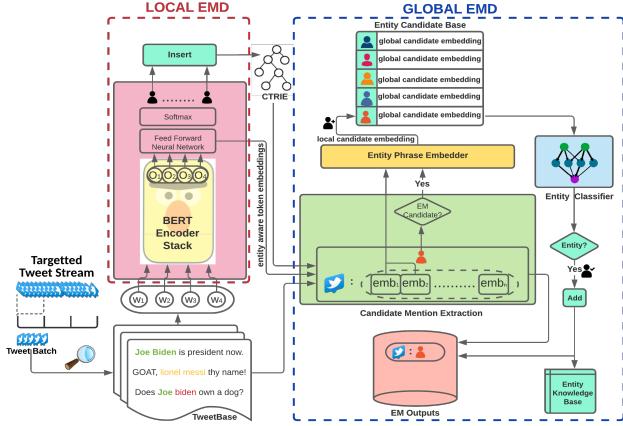


Fig. 3: EMD Globalizer with a Deep EMD System

We implement a procedure (NP Chunker) that extracts noun phrases from the parser-generated dependency trees and pass them as entity candidates to the rest of the framework.

2. CRF Based Tagging – TwitterNLP: TwitterNLP [3] recreates the Information Extraction pipeline for tweets by rebuilding classical NLP tools with microblog-specific considerations. The system begins by building modules for 2 subtasks, that help draw shallow semantic features from the tweet text – 1) Part-of-Speech tagging and 2) Noun Phrase Chunking. These features are later re-used for the actual NER task. The POS tagger (T-POS) and shallow parser (T-CHUNK) are CRF-trained models with additional tweet-sensitive features like @tags, hashtags, and URLs. In addition to shallow parsing, TwitterNLP also constructs a capitalization classifier (T-CAP) to deal with the unreliable capitalization in tweets. The classifier is trained as an SVM and it studies capitalization throughout the entire sentence to predict whether or not it is informative. For EMD, TwitterNLP splits the subtasks of segmenting and classifying Named Entities in two separate modules. T-SEG, a CRF-trained discriminative model, that adopts BIO encoded labels, for entity level segmentation of tweets. The feature set for this task uses orthographic features like capitalization, contextual features gathered from T-POS and T-CHUNK, output of T-CAP as well as dictionary features, including a set of type-lists gathered from Freebase [42], an open-domain ontology, and Brown Clusters. T-SEG uses a manually annotated training set for supervised learning of model parameters. We work with the production version of TwitterNLP available in Github for our experiments.

3. Multi-task Deep Neural Network – Aguilar et al. [8]: The best-performing system for the WNUT 2017 NER task [1] is an effective deep EMD system. It is primarily a BiLSTM-CNN-CRF network that follows the philosophy of multi-task learning. It is multi-task in the sense that it learns higher-order feature representations along three different tasks, each of which includes relevant information for the subsequent token-level entity tag prediction:

a) Character-level representation: character encodings from [43] and an orthographic encoder are fed to a Convolutional Neural Network to learn character-level representations.

b) Token-level representation: here both word level representations and POS representations are concatenated in a unified vector to denote token-level features. Word embeddings from [44] are fed to a BiLSTM to learn word level representations, while a POS encoder is trained using POS tags obtained by parsing the text using TweeboParser [40].

c) Lexical representation: tokens appearing as entities in select gazetteers [23] are encoded into a 6-dimensional vector – one dimension for each gazetteer type. These lexical vectors are fed into a fully-connected layer with ReLU activation function.

The concatenation of these feature vectors is then fed to a common dense layer with a sigmoid activation function. Finally a CRF layer learns dependencies between the neural output nodes and conducts sequential labeling. The token level encoding scheme used for sequence labeling is BIO. We use the production version of the system available online.

4. BERTweet for EMD – Nguyen et al. [10]: Pre-trained language models have become the go-to choice for many NLP tasks and several recent systems have adopted a pre-trained BERT language model that is fine-tuned for the downstream sequence labeling task of EMD. For our Local EMD instantiation we use BERTweet, the first large-scale language model trained on English Tweets. This system has the same architecture as BERT_{base} [45] but uses the RoBERTa [46] pre-training procedure for more robust performance. The pre-training dataset is a collection of 850M tweets, each consisting of at least 10 and at most 64 tokens. fastBPE is applied to segment all tweets into subword units, using a vocabulary of 64K subword types. On average two subword tokens are maintained per Tweet. To fine-tune the language model for EMD, a feed forward neural network (FFNN) layer and a softmax prediction layer are added on top of the output of the last Transformer encoder. The fine-tuning is independent of the underlying BERT model’s training. It is repeated five times with randomly initialized seeds. The reported performance is the average of five test runs. The BERTweet repository is publicly available on GitHub. We use a pre-trained BERTweet model available at the Hugging Face model hub that amasses a large collection of pre-trained language models catering to a variety of downstream NLP tasks.

Every Local EMD system suggests a set of seed entity candidates derived from the EMs that are tagged and discovered by it. These seed entity candidates are stored in a **CandidatePrefixTrie** (CTrie for short). CTrie is a prefix Trie forest for efficient indexing of candidates. It also facilitates subsequent lookups for finding all mentions of discovered entity candidates later during the Global EMD phase. CTrie functions like a regular Trie forest with individual nodes corresponding to a token in a candidate entity string. Entity candidates with overlapping prefixes are part of the same subtree in the forest. Another data structure produced at the end of Local EMD is **TweetBase**. It maintains an individual record for every tweet sentence indexed by a (tweet ID, sentence ID) pair and a list of detected mentions that get updated as the sentences pass through Global EMD.

In addition, deep EMD systems provide token-level contextual embeddings for each tweet-sentence in the input stream. These are also recorded for the computation of local candidate-level embeddings that are then used to generate global candidate embeddings. The token-level embeddings are collected from the final, pre-classification layer of deep EMD. For Aguilar et al. this would be the output of the last fully connected layer, prior to the CRF layer. For BERTweet, this would be the layer prior to the output softmax layer. In either case, these embeddings encode information that demarcates entity boundaries within the sentence tokens. Hence we call them local ‘entity-aware’ embeddings.

V. GLOBAL EMD

At the end of Local EMD an initial entity extraction is accomplished for every tweet-sentence in the TweetBase, sometimes with the generation of token embeddings that are aware of adjacent entity boundary information. The Local EMD outputs suggest a set of seed entity candidates stored in the CTrie. However Local EMD is prone to inconsistent detection of the same entity across the breadth of a stream. Mentions of entity candidates are missed in some sentences while detected in others. Here we introduce the Global EMD module to address some of these inherent limitations. More specifically, the purpose of Global EMD is to shift the focus of its EMD computation, beyond the confines of a single sentence. It views mentions of a candidate collectively, across the entire span of a stream, before determining if it is an entity.

Objectives: The Global EMD step addresses:

1. Removal of False Negatives: False Negatives happen when Local EMD fails to tag true EMs. For example, in Figure 1, *coronavirus* in T4 is a false negative.

2. Removal of False Positives: False Positives happen when Local EMD extracts non-entity phrases as entity candidates.

3. Correction of Partial Extraction: Partial extractions happen usually due to mislabeling of multi-token entities, where a part of an entity string is left outside of an entity window under consideration. Correcting such partial extractions improves both recall and precision.

Execution: The execution with Global EMD is broken down into three separate components. First, an additional scan of the tweet-sentences alongside a lookup through the CTrie, reveals all existing mentions of entity candidates, including ones that were previously missed by Local EMD. For every candidate mention encountered here, a local candidate embedding is extracted and recorded to its entry in a data structure called the ‘CandidateBase’. Depending on the local EMD system, the process of collecting local embedding varies. Non-deep systems do not generate token-level contextual embeddings along with their EM suggestions for a sentence. In this case, we generate a syntactic embedding for the candidate mention found from its immediate local context in a sentence. For deep EMD systems, token embeddings are collected from the TweetBase entry of a sentence recorded at the end of Local EMD. Next, a candidate’s token embeddings are together fed

to an Entity Phrase Embedder to generate a unified local contextual embedding for the entire phrase, for this mention of the candidate. Then, a pooling operation on mention level contextual embeddings gives us the ‘global candidate embedding’. Note that we update the global embedding of a candidate incrementally as we find new mentions in the stream. Finally, an Entity Classifier receives the ‘global candidate embeddings’ to label every candidate as an ‘entity’ or a ‘non-entity’. Candidates recognized as entity find their mentions produced as valid EMs in the final output for the stream. We now describe the components of Global EMD in more detail.

A. Candidate Mention Extraction

In theory the purpose of the Candidate Mention Extraction step is to detect EM boundaries within text. Most EMD systems [5], [16] treat this as a sequence labeling problem. However, empowered by the seed candidates from Local EMD registered in the CTrie, the problem of segmenting tweets into candidate (EM) boundaries here is simplified to that of a lookup in the CTrie. The module analyzes every token in a tweet sentence, in conjunction with a CTrie traversal. With a case-insensitive comparison of tokens with CTrie nodes, this results in two possibilities:

(i) A token that matches a candidate node on the current CTrie path, when cases are ignored.

(ii) A token matching no node in current path.

The problem is to check if a token forms a candidate mention alone or together with up to k tokens following it.

The extraction process scans a tweet-sentence and identifies the set of longest possible subsequences matching with candidates in the CTrie, while case is ignored (e.g., “coronavirus” is a match for “Coronavirus”). As a consequence, candidate mentions extracted during Local EMD are verified, and sometimes corrected. For example, if the Local EMD system finds only a partial excerpt ‘Andy’ of the EM ‘Andy Beshear’ in a tweet, but nonetheless recognized the entire string in other tweets, the candidate (‘andy beshear’) will be registered in the CTrie. This partial extraction can now be rectified and corrected to the complete mention. The resulting process is syntax agnostic. It initiates a window that incrementally scans through a sequence of consecutive tokens. In each step it checks:

a) whether the subsequence within the current scan window corresponds to an existing path in the CTrie. If true, it implies that the search can continue along the same path, by including the token to the right within the window in the next iteration.

b) whether the node on this path, that corresponds to the last token of the subsequence, refers to a valid candidate. If true, it implies that the subsequence can be recorded as the current longest match, before next iteration begins.

In case of a mismatch, i.e. scenario (ii), the module stores the last matched subsequence within the current window, and then skips ahead, by initializing a new window from the position of the token, next to it. The search for a new matching subsequence path is initiated from the root of the CTrie. However, if the last search had failed to yield a match with any of the existing candidates in CTrie, the new window is

initialized from a position that is to the immediate right of the first token in the previous window. The process is repeated until all tokens are consumed. In the end we obtain a collection of mention variations for each entity candidate.

B. Local Candidate Embedding Collection

The process of collecting candidate embeddings from the local context of a candidate’s mention depends on the type of Local EMD system. Here we broadly categorize them into two groups: 1) pre-deep learning systems– systems that do not use (deep) neural networks EMD systems and 2) deep EMD systems– systems that use deep neural networks. For the former, EMD execution is limited to generating BIO labels of sentence tokens that suggest entities. Hence, for these systems, we provide a workaround to study the immediate context of mentions and generate a local candidate embedding. Deep EMD systems, however, provide token-level contextual embeddings that encode latent relations between tokens in a sentence. Hence for deep learners we fully utilize them when preparing candidate embeddings.

1) **Syntactic Distribution for Non-deep Local EMD:** For non-deep systems we follow [14] and extract an embedding that encodes the local syntactic context of a candidate mention into an embedding of 6 dimensions. It enumerates all the syntactic possibilities in which a candidate can be presented. (1) **Proper Capitalization:** This corresponds to the first character of every candidate token capitalized.

(2) **Start-of-sentence capitalization:** A unigram candidate capitalized at the start of sentence.

(3) **Substring capitalization:** Only a proper substring of a multi-gram candidate is capitalized.

(4) **Full capitalization:** Abbreviations like ‘UN’ or ‘UK’ where the entire string is capitalized.

(5) **No capitalization:** The entire string is in lower case.

(6) **Non-discriminative:** A sentence is entirely in upper or lowercase, or has first character of every word capitalized, so candidate mentions found are not syntactically informative.

In the end, the local syntactic embeddings are aggregated and pooled to derive a global candidate embedding.

2) **Entity Phrase Embedder for Deep Local EMD:** The ‘entity aware embeddings’ generated by a Local Deep EMD system are for individual tokens. However, the Entity Classifier verifies candidates based on their global contextual representation generated by aggregating local contextual representations of their mentions. So we need semantically meaningful representations of candidate mentions before they can be aggregated. Given that entity candidates have variable number of tokens, we need to combine the token-level embeddings into a unified, fixed-size embedding of the candidate phrase. This is the role of the Entity Phrase Embedder.

To generate phrase embeddings, we refer to the work on sentence embedding for Semantic Textual Similarity (STS) tasks. Since the contextual embeddings provided by language models are token-level, the intuitive solution for a sentence embedding sourced from multiple token embeddings is to add an average (or max) pooling operation and arrive at an

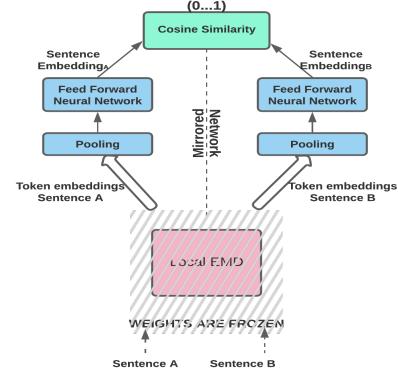


Fig. 4: Entity Phrase Embedder in Modified Siamese Network

average embedding to represent the sentence. Alternatively, one can add a CLS (classification) token at the end of sentences and train them for a Natural Language Inference task. The embedding for the CLS token would be considered representative of the entire sentence. We however follow the approach in Sentence-BERT or SBERT [47]. SBERT argues that using the aforementioned approaches do not yield good sentence embeddings, and can often be worse than averaging Glove embeddings [48] of tokens.

SBERT uses a ‘siamese network structure’ and trains it for different STS tasks, including sentence similarity. The input set in this case consists of pairs of sentences whose similarity is calculated such that sentences conveying similar content have a higher score than those that do not exhibit any content similarity. A siamese network consists of two identical sub-networks that have the same architecture and parameters to generate representations of pair-wise inputs that are then compared to generate a similarity score which is the network’s output. SBERT, in its implementation, uses the same BERT model as sub-networks in its siamese structure. It also adds an average pooling layer that generates a mean sentence-level embedding from the token-level embeddings of the BERT. Finally, the Cosine Similarity function is used to generate a similarity score upon which the loss function is calculated. Mean squared error loss is used as the regression objective function to train SBERT for this task. The updation of weights during back-propagation is mirrored across both sub-networks.

The Entity Phrase Embedder used in our framework is shown in Figure 4. It follows a modified design of SBERT and is also trained on the sentence similarity task. We use the deep neural network used for Local EMD to generate token-level embeddings as the principal component of the mirrored subnetwork in our siamese structure. In addition we add an average pooling layer to combine token-level representations into an average embedding that is then passed on to a dense layer. The Cosine similarity score between the dense layer outputs of the two subnetworks gives the final output upon which the regression loss function is calculated and backpropagated. Unlike SBERT however the gradient computation is not backpropagated all the way back to the deep neural network (the BERT engine in case of SBERT). In other words, the DNN’s weights remain frozen in our siamese network and

Newly Detected EMs	Aguilar et al. labelled EMs	Newly Detected EMs	BERTweet labelled EMs
T1: Beshear : Social distancing is not social isolation.		T1: Beshear : Social distancing is not social isolation.	
T2: WE JUST BY-PASS Italy WITH CORONAVIRUS CASES. But @realDonaldTrump wants to relax social distancing.		T2: WE JUST BY-PASS Italy WITH CORONAVIRUS CASES. But @realDonaldTrump wants to relax social distancing.	
T3: Not a bad video to explain how the Coronavirus works as well as the reasoning for social distancing.		T3: Not a bad video to explain how the Coronavirus works as well as the reasoning for social distancing.	
T4: Beshear says he's asking county judges and city mayors to monitor parks and shut them down if people aren't filling		T4: Beshear says he's asking county judges and city mayors to monitor parks and shut them down if people aren't filling	
T5: Trump to rank US counties by coronavirus risk, may 'relax' social distancing.		T5: Trump to rank US counties by coronavirus risk, may 'relax' social distancing.	
T6: Canada is rising at a rate similar to the early days in ITALY		T6: Canada is rising at a rate similar to the early days in ITALY	

Fig. 5: EMD Globalizer on a message stream on Coronavirus only the weights of the layers following it are updated. This is because the deep neural network’s role in our framework is to produce (Local) EMD results for which it had already been optimized. The rest of the subnetwork however work on the ‘entity-aware’ token embeddings to produce an optimal sentence/phrase embedding that performs well on the semantic similarity task. The dense layer is also useful to customize the phrase embedding size, separate from the underlying token embedding size generated by the DNN. For example, when training the Entity Phrase Embedder with BERTweet, where token embeddings are of 768 dimensions, we bring down the candidate embeddings from the Entity Phrase Embedder to 300 dimensions. Further details on the training process of the Entity Phrase Embedder is provided in Section VI.

A candidate’s local embedding ($local_emb \in \mathbb{R}^d$) from its token-level embeddings can be computed using one of the Entity Phrase Embedder sub-networks as

$$pooled_emb = \frac{1}{|T|} \sum_{j=1}^{|T|} token_emb_{T_j} \quad (1)$$

$$local_emb = W_{ff}pooled_emb + b_{ff} \quad (2)$$

where T denotes the set of tokens in the candidate phrase, $token_emb_{T_j} \in \mathbb{R}^d$ is the contextual embedding of the j -th token in T . The weight matrix $W_{ff} \in \mathbb{R}^{d \times d}$ and bias $b_{ff} \in \mathbb{R}^d$ are trainable parameters from the mirrored sub-networks.

C. Entity Classifier

The information encoded in the local embeddings of individual mentions of a candidate is limited to the context of the sentence containing it. We add these local embeddings to the candidate’s record in a data structure called the CandidateBase, which maintains an entry for every entity candidate discovered for a stream during Local EMD. It is incrementally updated with the local embeddings of a candidate’s mentions. Next, a pooling operation conducted over all the local contextual embeddings of an entity candidate gives the ‘**global candidate embedding**’. It is global in the sense that it aggregates all contextual possibilities in which a candidate appears in the stream, to generate a consensus representation.

The global candidate embeddings are then fed to a multi-layer network of feed-forward layers with ReLU activation function followed by a sigmoid output layer. We call this module the Entity Classifier. It is trained to determine the

likelihood of a candidate being an entity. The sigmoid output gives the probability of a candidate being a true entity and is divided into three ranges, that we empirically determined from variation in the Classifier’s entity detection performance over different values. Each range corresponds to a class label:

- (1) $\alpha: \geq 0.55$, candidate is confidently labelled as an **entity**.
- (2) $\beta: \leq 0.4$, candidate is confidently labelled as a **non-entity**.
- (3) $\gamma: \epsilon(0.4, 0.55)$, deemed **ambiguous**; requires more evidence downstream for confident labeling.

Note that a candidate’s global embedding over mention variations is more reliable when its frequency of occurrence is high and is computed over more than just a few mentions (avoiding randomness). Consequently, the classifier performs better in distinguishing entities among more frequent candidates.

The classifier is supervised with the training performed using labelled global embedding records of entity candidates extracted from D_5 (see Table 1). Further details on the classifier training will be provided in Section VI.

Follow-up on case study on dataset D_2 : In Figure 5, it can be seen that all the entity mentions that were missed by Aguilar et al. [8] and BERTweet [10] are discovered by the end of the EMD Globalizer pipeline.

TABLE I: Twitter Datasets

Dataset	Size	#Topics	#Hashtags	#Entities
D_1	1K	1	1	283
D_2	2K	1	1	461
D_3	3K	3	6	906
D_4	6K	5	5	674
D_5	38K	1	1	≈ 7000
WNUT17	1287	-	-	-
BTC	9553	-	-	-

VI. EXPERIMENTS

We conducted extensive experiments to test the effectiveness of EMD Globalizer for entity mention detection in tweets. We used four existing EMD systems, including two deep EMD systems, for Local EMD. In each case, we evaluate the performance gain, when plugged into the framework. We implemented the framework in Python 3.8 and executed it on a NVIDIA Tesla T4 GPU on Google Colaboratory. Our datasets and code used for experiments are available at Github.

Datasets: We use a combination of third-party datasets, along with the streaming datasets used in [14] that include crawled message streams from Twitter, for the purpose of evaluating the effectiveness of EMD Globalizer. The datasets are listed in Table I. D_1 - D_4 are **streaming datasets** that contain subsets of Twitter streams. The topics covered here are Politics, Sports, Entertainment, Science and Health, with (D_2) curated from a Covid-19 tweet stream. Having datasets directly covering a variety of tweet streams helps preserve their naturally occurring topic-specificity that often repeats a finite set of entities. This is exploited in Global EMD for collectively processing candidates, without making the analysis biased towards a particular topic. In real-world deployment, a topic classifier [49] could precede an EMD tool launched for streams.

Other than the four streaming datasets D_1 - D_4 , two datasets popular for EMD benchmarking, WNUT17 [1] and BTC [50], are also included in our evaluation. These are **non-streaming datasets** curated to accommodate a random sampling of tweets. Although they do not characterize the application setting for EMD Globalizer, we use them to gauge the framework’s effectiveness against pre-established benchmarks.

Also like [14], we use dataset D_5 , a collection of 38K tweets from a single tweet stream to generate entity candidates. The candidates are labelled as ‘entity’/‘non-entity’ and used to train the Entity classifier to learn optimal global embeddings and generate correct candidate labels for Global EMD.

Performance Metrics: We use Precision (P), Recall (R) and F1-score to evaluate EMD effectiveness. EMD requires detection of all occurrences of entities in their various string forms within a dataset. It is captured in WNUT17 shared task [1] as **F1 (Surface)**. Here we simply call it F1. Our framework does not involve entity typing. So the evaluation here only includes EMD and not their type classification. We also record execution times in seconds to check the run-time overhead for executing Local EMD systems within the framework.

Local EMD Instantiations: We run our EMD framework with four different Local EMD instantiations (see Section IV-A): 1) **NP Chunker** – a Chunking based Tagger that uses noun phrase chunking on Twitter dependency parser [40] to project entity candidates; 2) **Twitter NLP** [3] – a CRF based Machine Learning model; 3) **Aguilar et al.** [8] – a Deep Learning architecture that won the WNUT 2017 [1] NER challenge; and finally, 4) **BERTweet** [10] – a BERT language model trained on a large twitter corpus that we fine-tune using the WNUT2017 training data for the downstream EMD task.

Baseline for testing Global EMD: We use the production version of the Document EMD system HIRE-NER [38] as a baseline for testing Global EMD. We compare the performance of this system with EMD Globalizer on our Twitter datasets. HIRE-NER treats messages in a stream as composite content, much like a document.

Training Entity Phrase Embedder: When using Deep EMD systems for Local EMD, we employ an Entity Phrase Embedder in the Global EMD step to combine the contextual embeddings of a candidate’s tokens provided by the deep EMD system into a unified local embedding for the entire candidate phrase. As mentioned earlier, the Entity Phrase Embedder is trained using the STS Benchmark (STS-b) data [51]. This dataset contains 5749 sentence pairs in the training set and 1500 sentence pairs in the validation set. Every sentence pair is given a score between 0-5, indicating the semantic similarity between the sentences in the pair. To evaluate the Entity Phrase Embedder, we divide the integer scores by 5 to normalize them into a range of [0, 1], and then compare this with the Cosine similarities between the embeddings of the sentence pairs generated by the Entity Phrase Embedder. Here we use mean squared loss as the regression objective function to optimize training and estimate performance on the validation set.

We use Adam optimizer [52] with a fixed learning rate of 0.001 and batch size of 32. We compute performance on the validation set after each training epoch, and save the best model checkpoint to execute test sets. Here, we also enforce early stopping when validation performance does not improve for 25 continuous epochs. Note that S-BERT [47] tests the quality of sentence embeddings by employing them in downstream tasks like Paraphrase Mining and Natural Language Inference. Since we simply use the Entity Phrase Embedder to generate embeddings for candidate phrases, such detailed evaluation is not carried out here.

We separately train the Entity Phrase Embedder for the two different variants of our framework with Aguilar et al. [8] and BERTweet [10] as Local EMD systems. For Aguilar et al. the size of the candidate embeddings generated by the Entity Phrase Embedder is of 100 dimensions, the same as the system’s output vectors. When trained with token embeddings from Aguilar et al., the best validation loss obtained is 0.185. For BERTweet, we tested EMD Globalizer with candidate embeddings of size 768 – the size of the BERT encoder’s output layer – and 300. In our experiments, we obtained slightly better effectiveness with candidate embeddings of size 300 and hence we present those results in our evaluation in Table III. Nonetheless, these hyperparameters are easily customizable. When trained with token embeddings from BERTweet, the best validation loss is 0.167.

TABLE II: Validation Performance of Entity Classifier

Local EMD	Local EMD System Type	Entity Embedding Size	Validation F1
NP Chunker	CRF Chunker	6+1	0.936
TwitterNLP	CRF EMD Tagger	6+1	0.936
Aguilar et al.	BiLSTM-CNN-CRF	100+1	0.908
BERTweet	BERT-FFNN	300+1	0.941

Training Entity Classifier: We train the Entity Classifier everytime we initialize a variant of the framework with a different Local EMD instantiation. The ‘+1’ in the column for Entity Embedding Size (see Table II) indicates that we also add length of the candidate string as an additional feature, along with the global entity embedding. We use a 80-20 training to validation split and train over 1000 epochs. We use Adam optimizer [52] with a fixed learning rate of 0.0015 and batch size of 128. We compute the task performance after each training epoch on the validation set, and select the best model checkpoint to compute the performance score on the test set. Here, we also apply early stopping when no improvement is observed on the validation loss after 20 continuous epochs. These validation performance obtained when training the Entity Classifier for different variants of our framework are compiled in Table II.

A. Evaluating EMD Globalizer

We first evaluate EMD Globalizer on its primary objective of boosting EMD for the Local EMD systems. To this end, we test EMD Globalizer with four different Local EMD

TABLE III: Effectiveness and Execution Time (in seconds) with EMD Globalizer

Dataset	Local EMD					Global EMD				F1 Gain	Time Overhead
	System Name	P	R	F1	Execution Time	P	R	F1	Execution Time		
D1	NP Chunker	0.3	0.58	0.4	100.4	0.81	0.63	0.71	101.6	77.5%	1.2
	TwitterNLP	0.65	0.47	0.55	7.07	0.8	0.66	0.72	8.03	36.4%	0.96
	Aguilar et al.	0.76	0.55	0.64	124.8	0.87	0.66	0.75	126.07	17.3%	1.27
	BERTweet	0.66	0.49	0.56	33.16	0.84	0.66	0.74	34.32	32.1%	1.16
D2	NP Chunker	0.40	0.47	0.43	123.62	0.59	0.62	0.60	125.71	39.5%	2.09
	TwitterNLP	0.33	0.52	0.41	18.91	0.71	0.55	0.62	20.57	51.2%	1.66
	Aguilar et al.	0.63	0.57	0.60	296	0.69	0.67	0.68	297.7	13.3%	1.7
	BERTweet	0.56	0.51	0.53	40.23	0.65	0.64	0.64	42.58	20.8%	2.35
D3	NP Chunker	0.59	0.54	0.56	175.3	0.71	0.66	0.68	177.9	21.4%	2.6
	TwitterNLP	0.75	0.64	0.69	15.1	0.88	0.71	0.78	18	13.04%	2.9
	Aguilar et al.	0.77	0.64	0.70	298.2	0.82	0.77	0.794	301.34	13.6%	3.14
	BERTweet	0.77	0.63	0.69	58.6	0.83	0.82	0.83	62.18	20.3%	3.58
D4	NP Chunker	0.47	0.59	0.52	551.3	0.83	0.73	0.77	556.7	48.1%	5.4
	TwitterNLP	0.67	0.41	0.52	35.24	0.89	0.64	0.74	41.06	42.3%	5.82
	Aguilar et al.	0.82	0.61	0.69	588.24	0.88	0.75	0.81	594.22	17.4%	5.98
	BERTweet	0.69	0.58	0.62	230.75	0.81	0.76	0.78	237.53	26.1%	6.78
WNUT17	NP Chunker	0.42	0.35	0.39	121.22	0.63	0.35	0.44	123.56	12.8%	2.34
	TwitterNLP	0.35	0.42	0.39	14.25	0.65	0.52	0.58	16.72	48.7%	2.47
	Aguilar et al.	0.68	0.47	0.56	229.32	0.72	0.5	0.59	231.04	5.4%	1.72
	BERTweet	0.61	0.43	0.51	24.40	0.73	0.48	0.58	26.15	13.7%	1.75
BTC	NP Chunker	0.46	0.51	0.48	627.98	0.66	0.52	0.58	642.02	20.8%	14.04
	TwitterNLP	0.69	0.43	0.53	77.15	0.74	0.45	0.56	87.8	5.7%	10.65
	Aguilar et al.	0.75	0.56	0.64	685.36	0.77	0.59	0.67	695.56	4.7%	10.2
	BERTweet	0.63	0.50	0.56	193.8	0.69	0.58	0.63	204.49	12.5%	10.69

TABLE IV: Effectiveness of Global EMD systems

Dataset	Global EMD System	P	R	F1
D1	EMD Globalizer	0.87	0.66	0.75
	HIRE-NER	0.65	0.62	0.63
D2	EMD Globalizer	0.69	0.67	0.68
	HIRE-NER	0.46	0.56	0.51
D3	EMD Globalizer	0.82	0.77	0.79
	HIRE-NER	0.75	0.73	0.74
D4	EMD Globalizer	0.88	0.75	0.81
	HIRE-NER	0.58	0.68	0.61
WNUT	EMD Globalizer	0.72	0.5	0.59
	HIRE-NER	0.5	0.49	0.5
BTC	EMD Globalizer	0.77	0.59	0.67
	HIRE-NER	0.6	0.49	0.54

instantiations. In each case we check the improvement that is achieved from the Local EMD’s initial F1 score by executing it with the rest of the framework. Table III summarizes these results along with the execution time overhead brought in for each Local EMD system post plugin. We also compare the best performing EMD Globalizer variant to a state-of-the-art Document EMD system to understand how effectively global information is mined in each case and utilized for EMD.

Local EMD Performance: The columns under ‘Local EMD’ in Table III show the EMD performances of each of the four Local EMD systems, along with their computation times.

Performance improvement with EMD Globalizer: The columns in Table III under ‘Global EMD’ show the EMD performance once the Global EMD components have been run on top of a Local EMD system and the total run-time at the end of its execution. Comparing the F1 of a Local EMD system with that of its Global EMD counterpart gives the improvement achieved by the framework. The difference in execution times at the end of the two stages gives the additional time required by Global EMD.

Global EMD is able to make considerable improvement in

performance with only minor execution time overhead across all datasets. For each Local EMD system, we compute the percentage gain in F1 score across all datasets. For time overhead, we calculate the absolute increment in the execution time of a system (in seconds) when passed through EMD Globalizer. This provides better context. As is the case of every Local EMD system, the absolute overhead incurred by injecting it into the framework is only a few additional seconds. For computationally expensive EMD systems that already have higher execution times, the time overhead is nominal, thereby making the performance gain obtained all that more significant. In summary, the average performance gain across all datasets and all local EMD systems is 25.61%. The average individual performance gains for the four Local EMD systems are: a) 36.69% for NP Chunker, b) 31.06% for TwitterNLP, c) 11.91% for Aguilar et al., and, d) 20.66% for BERTweet. Note that two types of datasets are used in our evaluation, one is the streaming datasets – the application setting for which EMD Globalizer was originally designed, and the other is non-streaming datasets – popularly used for EMD benchmarking. EMD Globalizer yields different improvements over these dataset types as discussed below.

Improvement on Streaming Datasets: For datasets D_1-D_4 that retain the inherent properties of Twitter streams, EMD Globalizer yields an average F1 gain of 30.29% across all Local EMD systems. For individual Local EMD instantiations, the average F1 gains are: a) 46.63% for NP Chunker, b) 34.36% for TwitterNLP, c) 15.36% for Aguilar et al., and, d) 24.82% for BERTweet.

Improvement on Non-Streaming Datasets: For datasets WNUT17 and BTC, there is no adherence to specific Twitter streams but rather a random sampling off the Twittersphere, avoiding the latter’s tendency to repeat entity mentions within

streams. However, EMD Globalizer is still able to improve effectiveness for its Local EMD instantiations, albeit to a less significant degree than streaming datasets. In this case, the average F1 gain across all Local EMD systems is 15.53%. For individual Local EMD instantiations, the average F1 gains are: a) 16.82% for NP Chunker, b) 24.47% for TwitterNLP, c) 5.04% for Aguilar et al., and, d) 12.22% for BERTweet.

Comparison with Document EMD method on Global EMD: We compare the performance of Aguilar et al. instantiated EMD Globalizer and HIRE-NER [38] (both BiLSTM architectures) on all the annotated datasets in Table IV. Here we test how effectively global information is captured in each system when performing EMD. As evident from Table IV, EMD Globalizer consistently outperforms HIRE-NER across all datasets, especially attaining higher precision. HIRE-NER simultaneously updates global features in the memory structure and appends them with local embeddings to infer final output labels of tokens in a sentence. Adding non-local contextual information inevitably introduces noise which can interfere with the decoder’s inference of output labels. Distinct from this, we limit the curation of global contextual representations only for entity candidates. First the local context suggests the entity candidates situated within a sentence. The local contextual embeddings of the various candidate mentions are then aggregated in the memory structure to generate global candidate embeddings. Using this the entity classifier is able to better separate true entities from noisy candidates.

B. Ablation Study on Framework Components

While it is evident that the proposed EMD framework is capable of enhancing performance for its various Local EMD instantiations, we wanted to take a closer look at how the individual framework components contribute towards the EMD overall performance. To this end, we execute the framework with Aguilar et al. [8] as the Local EMD system in it as this instantiation is the best performer among the local EMD systems used in this paper. Here, we use the entire collection of annotated streaming datasets (D_1 - D_4) as the test set. Figure 6 shows the improvement in performance as individual system components are added. From bottom to top, the first curve (with only Local EMD) reports the weakest performance proving the limitations of the standalone system in capturing all the entity mention variations within the stream. The middle curve is the EMD performance we get just by following up the entity (candidate) extraction by Local EMD with the mention extraction process that simply adds missed mentions of candidates detected as likely entities by the local EMD system. The topmost curve is the performance yielded at the end of run of the entire framework with Aguilar et al. [8]. Aguilar et al. [8] is a very competitive EMD system – the best among the four (local) systems we evaluated in this paper. Even for such a system, EMD Globalizer is still able to significantly improve on its EMD effectiveness over the streaming datasets. Following up its execution with just the candidate mention extraction process gives a modest improvement of 5.06%. This simply recovers missed mentions

of candidates identified as entities elsewhere in the stream. The focus here is mainly on improving the recall by yielding more consistent mention detection across tweets. With EMD Globalizer however, the average overall improvement across all streaming datasets is 15.36%. This is because with all components of Global EMD in place, candidates suggested by Aguilar et al. are further verified and false positives are removed. Table III shows that both precision and recall are improved with the induction of Global EMD.

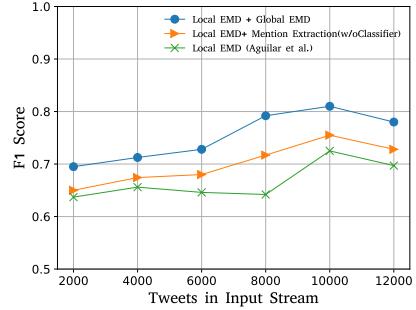


Fig. 6: Impact of Components on Performance

C. Error Analysis

Though EMD Globalizer improves over Local EMDs, it is not perfect. We give here an analysis of its errors.

1) If Local EMD misses every mention of an entity, the entity will not be added as a candidate to the CTrie and will also go undetected in the global phase. As a result, all mentions of the entity will not be detected by the proposed framework. Of the 11412 mentions in our streaming datasets from 2306 unique entities, the BERTweet instantiated EMD Globalizer failed to find 3008 (26.35%) mentions of 1018 entities that are entirely missed by the BERTweet system.

2) If Global EMD mislabels a candidate that happens to be an entity, then all of its mentions will be left out of the final output. This would include the mentions that the Local EMD did correctly find at first. But more importantly, a false negative from Entity Classifier would hinder EMD Globalizer’s objective of recovering mentions of the entity that the Local EMD missed. However, in our experience, it is rare that an entity found by Local EMD is missed at the global step. Of the 11412 entity mentions in our streaming datasets, BERTweet instantiated EMD Globalizer missed only 469 mentions (4.1%) due to the mislabeling of 81 entities as false negatives by the Entity Classifier.

3) Handling of long-tailed entities: To better understand the false negatives yielded by Global EMD, we take a look at how the Entity Classifier’s performance changes as more mentions of an entity are found in a stream. Figure 7 shows that it is consistently able to detect high-frequency entities from the streaming datasets. We group entities of different mention frequency in bins of width 5 and track the classifier’s recall in detecting them. For infrequent entities, the recall is modest – around 56% for entities with 5 or less mentions. But it increases quickly with mention frequency and most frequent entities are correctly labelled. This ensures that their mentions

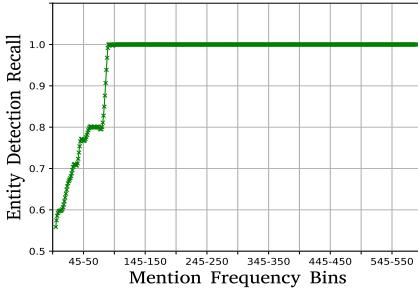


Fig. 7: Impact of Frequency on Detecting Entities

are also included in the final EMD Globalizer output. This confirms our initial intuition that as more mentions of an entity are found, better global contextual embeddings can be learned, leading to correct classification. Although long-tailed entities are a common issue in EMD, EMD Globalizer can still rectify the mislabeling of many such entities by collecting more instances further downstream.

EMD Globalizer aims to correct the often irregular detection of mentions of the same entity and we expect an improvement in EMD Recall from the Local EMD step. As seen in Table III, the EMD Globalizer is able to improve Aguilar et al.’s recall on the streaming datasets by 20.2% and BERTweet’s recall by 30.3% on average. Moreover, EMD Globalizer also does a good job of filtering noise when aggregating non-local information. As observed in the Global EMD baseline, this issue can affect the EMD precision. EMD Globalizer not only improves the precision yielded by Local EMD, it also yields fewer false positives than HIRE-NER (in Table IV) on all datasets. From Table III, EMD Globalizer is able to improve Aguilar et al.’s precision by 10.1% on average and BERTweet’s by 17.1% for the streaming datasets.

D. Discussions

We make some additional interesting observations from Table III. Here we summarize their implications for EMD:

- **Streaming vs Non-streaming Datasets:** Even though EMD Globalizer was conceptualized with the streaming setting in mind, all Local EMD instantiations, including Aguilar et al. – the erstwhile topper of WNUT17 NER challenge [1] – experience performance gain within the framework, on the two widely-used non-streaming datasets. This further validates the power of this framework in maximizing performance across different dataset types. Nonetheless, the framework is designed mainly to improve EMD performance on message streams. The idea of generating global contextual embeddings guided by a candidate mention extraction process specifically relies on the recurrence of entities across messages – a phenomenon more typical of social media message streams. Hence for streaming datasets, higher EMD performance is achieved with EMD Globalizer.
- **Design Flexibility:** We deliberately decouple the Local EMD step – which can be any existing EMD tool – from the rest of the EMD Globalizer framework. The advantage here is that the Local EMD tool can be inserted as is without any algorithmic modification. Also, depending on the type

of Local EMD tool, the individual components of Global EMD are separately customizable.

- **Improvement for low-performing EMD systems:** Local EMD variants like the NP Chunker that initially produced sporadic and relatively lower effectiveness, also yield competitive performance when aided by the framework.
- **New state-of-the-art record on existing benchmarks:** Aguilar et al. outperformed its original performance of the WNUT17 challenge when executed within the framework, and, along with BERTweet and TwitterNLP, all record much improved results on WNUT17. BERTweet also improves upon its performance on WNUT17 as reported in [10] when executed within the framework.
- **Time Overhead:** The time overhead brought by inserting a Local EMD system into EMD Globalizer is a fraction of its standalone execution time and depends on the input/stream size. The absolute overhead is a few additional seconds.

VII. CONCLUSIONS

In this paper, we presented EMD Globalizer – a novel two-phase EMD framework designed to address the limitations of existing EMD systems when executed on microblog streams and improve their effectiveness. Although EMD Globalizer is itself not a standalone EMD system, it is capable of significantly improving the EMD effectiveness of existing EMD systems that perform EMD on microblogs individually. In this paper, we tested EMD Globalizer with four existing EMD systems of various types (two deep EMD systems and two non-deep EMD systems) on both streaming and non-streaming datasets. Remarkable improvement on effectiveness (based on F1-measure) is achieved for each of these systems, averaging over 25% across all the four EMD systems. The improvement is even more remarkable for streaming datasets, averaging over 30% across the four systems. EMD Globalizer is specifically designed for streaming datasets. These improvements were achieved with only a small execution time overhead.

A big reason that EMD Globalizer can achieve these improvements is its ability to aggregate local candidate embeddings (contexts of local entity candidates) into a global candidate embedding. It is global in the sense that it considers all the contexts in which a candidate appears within a stream and derives a consensus representation. The global candidate embedding is then used to determine the likelihood of a candidate being an entity. Candidates labelled as entities find their mentions in the final output.

An encouraging result from our experiments was that EMD Globalizer also achieved good improvement for standard non-streaming datasets, averaging over 15%. Based on this, we believe that EMD Globalizer is a powerful tool that could be applied for different EMD application settings, not just microblog streams. In future work, we aim to expand the idea of collective processing for the entire NER pipeline.

Acknowledgements: This work was supported in part by U.S. National Science Foundation BIGDATA 1546480 and 1546441 grants, and a gift from NVIDIA Corporation.

REFERENCES

- [1] L. Derczynski, E. Nichols, M. van Erp, and N. Limsopatham, “Results of the wnut2017 shared task on novel and emerging entity recognition,” in *WNUT*, 2017, pp. 140–147.
- [2] L. Cen, E. C. Dragut, L. Si, and M. Ouzzani, “Author disambiguation by hierarchical agglomerative clustering with adaptive stopping criterion,” in *SIGIR*, 2013, pp. 741–744.
- [3] A. Ritter, S. Clark, Mausam, and O. Etzioni, “Named entity recognition in tweets: An experimental study,” in *EMNLP*, 2011, pp. 1524–1534.
- [4] M. B. Habib and M. van Keulen, “Need4tweet: a twitterbot for tweets named entity extraction and disambiguation,” 2015.
- [5] J. R. Finkel, T. Grenager, and C. Manning, “Incorporating non-local information into information extraction systems by gibbs sampling,” in *ACL*, 2005, pp. 363–370.
- [6] A. T. Schneider, A. Mukherjee, and E. C. Dragut, “Leveraging social media signals for record linkage,” in *WWW*, P. Champin, F. Gandon, M. Lalmas, and P. G. Ipeirotis, Eds., 2018, pp. 1195–1204.
- [7] F. Dernoncourt, J. Y. Lee, and P. Szolovits, “NeuroNER: an easy-to-use program for named-entity recognition based on neural networks,” 2017.
- [8] G. Aguilar, S. Maharjan, A. P. López-Monroy, and T. Solorio, “A multi-task approach for named entity recognition in social media data,” in *W-NUT*. ACL, 2017.
- [9] M. N. Gerguis, C. Salama, and M. W. El-Kharashi, “Asu: An experimental study on applying deep learning in twitter named entity recognition,” *WNUT*, 2016.
- [10] D. Q. Nguyen, T. Vu, and A. T. Nguyen, “BERTweet: A pre-trained language model for English Tweets,” in *EMNLP*, 2020, pp. 9–14.
- [11] A. Aljebreen, W. Meng, and E. C. Dragut, “Segmentation of tweets with urls and its applications to sentiment analysis,” in *AAAI*. AAAI Press, 2021, pp. 12480–12488.
- [12] L. He, C. Han, A. Mukherjee, Z. Obradovic, and E. C. Dragut, “On the dynamics of user engagement in news comment media,” *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.*, vol. 10, no. 1, 2020.
- [13] A. Alkbik, T. Bergmann, and R. Vollgraf, “Pooled contextualized embeddings for named entity recognition,” in *NAACL-HLT*, 2019, pp. 724–728.
- [14] S. Saha Bhowmick, E. C. Dragut, and W. Meng, “Twics: Lightweight entity mention detection in targeted twitter streams,” *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [15] K. Bontcheva, L. Derczynski, A. Funk, M. A. Greenwood, D. Maynard, and N. Aswani, “TwitIE: An open-source information extraction pipeline for microblog text,” in *RANLP*, 2013.
- [16] X. Liu, S. Zhang, F. Wei, and M. Zhou, “Recognizing named entities in tweets,” in *HLT*, 2011.
- [17] C. Cherry and H. Guo, “The unreasonable effectiveness of word representations for twitter named entity recognition,” in *HLT 2015*.
- [18] E.-S. Yang and Y.-S. Kim, “Hallym: Named entity recognition on twitter with induced word representation,” in *ACL-IJCNLP 2015*.
- [19] Z. Toh, B. Chen, and J. Su, “Improving twitter named entity recognition using word representations,” in *WNUT 2015*.
- [20] E.-S. Yang, Y.-B. Kim, R. Sarikaya, and Y.-S. Kim, “Drop-out conditional random fields for twitter with huge mined gazetteer,” in *HLT-NAACL*, 2016, pp. 282–288.
- [21] M. S. Akhtar, U. K. Sikdar, and A. Ekbal, “Iitp: Multiobjective differential evolution based twitter named entity recognition,” in *W-NUT*, 2015, pp. 61–67.
- [22] T. Tian, M. Dinarelli, and I. Tellier, “Lattice: Data adaptation for named entity recognition on tweets with features-rich crf,” in *ACL 2015 workshop WNUT*, 2015.
- [23] S. Mishra and J. Diesner, “Semi-supervised named entity recognition in noisy-text,” in *WNUT*, 2016, pp. 203–212.
- [24] P. von Däniken and M. Cieliebak, “Transfer learning and sentence level features for named entity recognition on tweets,” in *WNUT 2017*.
- [25] F. Dugas and E. Nichols, “Deepnnner: Applying blstm-cnns and extended lexicons to named entity recognition in tweets,” in *WNUT*, 2016.
- [26] Q. Zhang, J. Fu, X. Liu, and X. Huang, “Adaptive co-attention network for named entity recognition in tweets,” in *AAAI 2018*.
- [27] Y. Shen, H. Yun, Z. C. Lipton, Y. Kronrod, and A. Anandkumar, “Deep active learning for named entity recognition,” *arXiv preprint arXiv:1707.05928*, 2017.
- [28] S. Zhang, L. He, E. C. Dragut, and S. Vucetic, “How to invest my time: Lessons from human-in-the-loop entity extraction,” in *SIGKDD*, 2019, pp. 2305–2313.
- [29] S. Zhang, L. He, S. Vucetic, and E. C. Dragut, “Regular expression guided entity mention mining from noisy web data,” in *EMNLP*, pp. 1991–2000.
- [30] B. Strauss, B. Toma, A. Ritter, M.-C. de Marneffe, and W. Xu, “Results of the wnut16 named entity recognition shared task,” in *WNUT 2016*.
- [31] B. Y. Lin and W. Lu, “Neural adaptation layers for cross-domain named entity recognition,” *arXiv preprint arXiv:1810.06368*, 2018.
- [32] N. Peng and M. Dredze, “Multi-task domain adaptation for sequence tagging,” *arXiv preprint arXiv:1608.02689*, 2016.
- [33] C. Li, J. Weng, Q. He, Y. Yao, A. Datta, A. Sun, and B.-S. Lee, “Twiner: Named entity recognition in targeted twitter stream,” in *SIGIR*, 2012, pp. 721–730.
- [34] K. Wang, C. Thrasher, E. Viegas, X. Li, and B.-j. P. Hsu, “An overview of microsoft web n-gram corpus and applications,” in *NAACL HLT 2010 DEMO*, 2010, pp. 45–48.
- [35] X. Liu, M. Zhou, F. Wei, Z. Fu, and X. Zhou, “Joint inference of named entity recognition and normalization for tweets,” in *ACL*, 2012.
- [36] I. Yamada, H. Takeda, and Y. Takefuji, “Enhancing named entity recognition in twitter messages using entity linking,” in *WNUT 2015*.
- [37] A. Gattani, D. S. Lamba, N. Garera, M. Tiwari, X. Chai, S. Das, S. Subramaniam, A. Rajaraman, V. Harinarayan, and A. Doan, “Entity extraction, linking, classification, and tagging for social media: a wikipedia-based approach,” *PVLDB*, vol. 6, no. 11, pp. 1126–1137, 2013.
- [38] Y. Luo, F. Xiao, and H. Zhao, “Hierarchical contextualized representation for named entity recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8441–8448.
- [39] T. Gui, J. Ye, Q. Zhang, Y. Zhou, Y. Gong, and X. Huang, “Leveraging document-level label consistency for named entity recognition,” in *IJCAI*, 2020, pp. 3976–3982.
- [40] L. Kong, N. Schneider, S. Swayamdipta, A. Bhatia, C. Dyer, and N. A. Smith, “A dependency parser for tweets,” 2014.
- [41] O. Owoputi, B. O’Connor, C. Dyer, K. Gimpel, N. Schneider, and N. A. Smith, “Improved part-of-speech tagging for online conversational text with word clusters,” in *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies*, 2013, pp. 380–390.
- [42] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor, “Freebase: a collaboratively created graph database for structuring human knowledge,” in *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. ACM, 2008, pp. 1247–1250.
- [43] X. Ma and E. Hovy, “End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF,” in *ACL*, 2016, pp. 1064–1074.
- [44] F. Godin, B. Vandersmissen, W. De Neve, and R. Van de Walle, “Multimedia lab@ acl wnut ner shared task: Named entity recognition for twitter microposts using distributed word representations,” in *WNUT*, 2015, pp. 146–153.
- [45] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *NAACL-HLT*. Minneapolis, Minnesota: Association for Computational Linguistics, 2019, pp. 4171–4186.
- [46] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [47] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in *EMNLP-IJCNLP*. Association for Computational Linguistics, 2019, pp. 3982–3992.
- [48] J. Pennington, R. Socher, and C. Manning, “GloVe: Global vectors for word representation,” in *EMNLP*. Association for Computational Linguistics, 2014, pp. 1532–1543.
- [49] K. Lee, D. Palsetia, R. Narayanan, M. M. A. Patwary, A. Agrawal, and A. Choudhary, “Twitter trending topic classification,” in *ICDM*, 2011.
- [50] L. Derczynski, K. Bontcheva, and I. Roberts, “Broad Twitter corpus: A diverse named entity recognition resource,” in *COLING*, 2016, pp. 1169–1179.
- [51] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia, “SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation,” *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 2017.
- [52] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.