

LESA: Linguistic Encapsulation and Semantic Amalgamation Based Generalised Claim Detection from Online Content

Shreya Gupta^{†*}, Parantak Singh^{‡*}, Megha Sundriyal[†],
Md Shad Akhtar[†], Tanmoy Chakraborty[†]

[†] *IIIT-Delhi, India.* [‡] *Birla Institute of Technology and Science, Pilani, Goa, India.*
{shreyag, meghas, shad.akhtar, tanmoy}@iiitd.ac.in,
f20170109@goa.bits-pilani.ac.in

Abstract

The conceptualization of a claim lies at the core of argument mining. The segregation of claims is complex, owing to the divergence in textual syntax and context across different distributions. Another pressing issue is the unavailability of *labeled* unstructured text for experimentation. In this paper, we propose *LESA*, a framework which aims at advancing headfirst into expunging the former issue by assembling a source-independent generalized model that captures syntactic features through part-of-speech and dependency embeddings, as well as contextual features through a fine-tuned language model. We resolve the latter issue by annotating a Twitter dataset which aims at providing a testing ground on a large unstructured dataset. Experimental results show that *LESA* improves upon the state-of-the-art performance across six benchmark claim datasets by an average of 3 claim-F1 points for in-domain experiments and by 2 claim-F1 points for general-domain experiments. On our dataset too, *LESA* outperforms existing baselines by 1 claim-F1 point on the in-domain experiments and 2 claim-F1 points on the general-domain experiments. We also release comprehensive data annotation guidelines compiled during the annotation phase (which was missing in the current literature).

1 Introduction

The concept of a claim lies at the core of the argument mining task. [Toulmin \(2003\)](#), in his argumentation theory, described the term ‘claim’ as ‘*an assertion that deserves our attention*’; albeit not very precise, it still serves as an initial insight. In recent years, [Govier \(2013\)](#) described a ‘claim’ as

* First two authors have equal contributions. The work was done when Parantak was an intern at LCS2 Lab, IIIT-Delhi.

Text	Claim?
Alcohol cures corona.	Yes
Wearing mask can prevent corona.	Yes
Lord, please protect my family & the Philippines from the corona virus.	No
If this corona scare doesn’t end soon imma have to intervene	No

Table 1: A few examples of claim and non-claim.

‘*a disputed statement that we try to support with reasons.*’

The predicament behind the claim detection task exists given the disparity in conceptualization and lack of a proper definition of a claim. The task of claim detection across different domains has garnered tremendous attention so far owing to an uprise in social media consumption and by extension the existence of fake news, online debates, widely-read blogs, etc. As an elementary example, claim detection can be used as a precursor to fact-checking; wherein segregation of claims aids in restricting the corpus that needs a fact-check. A few examples are shown in Table 1.

Most of the existing works are built upon two fundamental pillars - semantic encapsulation ([Daxenberger et al., 2017](#); [Chakraborty et al., 2019](#)) and syntactic encapsulation ([Levy et al., 2014](#); [Lippi and Torroni, 2015](#)). They mainly focus on adapting to texts from similar distributions or topics or both. Secondly, they often exercise against well-structured and laboriously pre-processed formal texts owing to the lack of a labeled corpus consisting of unstructured texts. As a result, claim detection from unstructured raw data still lies under a relatively less explored umbrella.

Motivation: Claims can be sourced from a variety of sources, e.g., online social media texts, microblogs, Wikipedia articles, etc. It is, however, crucial to pay special attention to claims observed

on online social media (OSM) sites (Matthew A. Baum; who). Twitter, being a major OSM platform, provides the perfect playground for different ideologies and perspectives. Over time, Twitter has emerged as the hub for short, unstructured pieces of text that describe anything from news to personal life. Most individuals view and believe things that align with their compass and prior knowledge, *aka conformity bias* (Whalen and Laland, 2015) – users tend to make bold claims that usually create a clash between users of varied opinions. At times, these claims incite a negative impact on individuals and society. As an example, a tweet that reads “*alcohol cures corona*” can lead to massive retweeting and consequential unrest, especially in times of a pandemic, when people are more vulnerable to suggestions. In such cases, automated promotion of claims for immediate further checks could prove to be of utmost importance. An automated system is pivotal since OSM data is far too voluminous to allow for manual human checks, even if it was an expert.

At the same time deploying separate systems contingent on the source of a text is inefficient and moves away from the goal of attaining human intelligence in natural language processing tasks. An ideal situation would be a framework that can effectively detect claims in the general setting. However, a major bottleneck towards this goal is the unavailability of an annotated dataset from noisy platforms like Twitter. We acknowledge this bottleneck and, in addition to proposing a generalised framework, we develop a qualitative annotated resource and guidelines for claim detection in tweets.

Proposed Method: There exists several claim detection models; however, the downside is that most of them are trained on structured text from a specific domain. Therefore, in this work, we propose **LESA**, a Linguistic Encapsulation and Semantic Amalgamation based generalized claim detection model that is capable of accounting for different text distributions, simultaneously. To formalize this, we divide the text, contingent upon their structure, into three broad categories – noisy text (*tweets*), semi-noisy text (*comments*), and non-noisy text (*news, essays, etc.*). We model each category separately in a joint framework and fuse them together using attention layers.

Since the task of claim detection has a strong association with the structure of the input, as argued by Lippi and Torroni (2015), we leverage two

linguistic properties – part-of-speech (POS) and dependency tree, to capture the linguistic variations of each category. Subsequently, we amalgamate these features with BERT (Devlin et al., 2019) for classification.

We evaluate **LESA** on seven different datasets (including our Twitter dataset) and observe efficient performance in each case. Moreover, we compare **LESA**’s performance against various state-of-the-art systems for all seven datasets in the general and individual settings. The comparative study advocates the superior performance of **LESA**.

Summary of the Contributions: We summarize our major contributions below:

- **Twitter claim detection dataset and comprehensive annotation guidelines.** To mitigate the unavailability of an annotated dataset for claim detection in Twitter, we develop a large COVID-19 Twitter dataset, the first of its kind, with $\sim 10,000$ labeled tweets, following a comprehensive set of claim annotation guidelines.
- **LESA, a generalized claim detection system.** We propose a generalized claim detection model, **LESA**, that identifies the presence of claims in *any* online text, without prior knowledge of the source and independent of the domain. To the best of our knowledge, this is the first attempt to define a model that handles claim detection from both structured and unstructured data in conjunction.
- **Exhaustive evaluation and superior results.** We evaluate **LESA** against multiple state-of-the-art models on six benchmark claim detection datasets and our own Twitter dataset. Comparison suggests **LESA**’s superior performance across datasets and the significance of each model component.

Reproducibility: Code and dataset is publicly available at <https://github.com/LCS2-IIITD/LESA-EACL-2021>. Appendix comprises of detailed dataset description, annotation guidelines, hyperparameters, and additional results.

2 Related Work

In the past decade, the task of claim detection has become a popular research area in text processing with an initial pioneering attempt by Rosenthal and McKeown (2012). They worked on mining claims from discussion forums and employed a

supervised approach with features based on sentiment and word-grams. [Levy et al. \(2014\)](#) proposed a context dependent claim detection (CDCD) approach. They described CDC as ‘*a general, concise statement that directly supports or contests the given topic.*’ Their approach was evaluated over Wikipedia articles; it detected sentences that include CDCs using context-based and context-free features. This was followed by ranking and detecting CDCs using logistic regression. [Lippi and Torroni \(2015\)](#) proposed context-independent claim detection (CICD) using linguistic reasoning, and encapsulated structural information to detect claims. They used constituency parsed trees to extract structural information and predicted parts of the sentence holding a claim using SVM. Although their approach achieved promising results, they also used a Wikipedia dataset which was highly engineered and domain dependent.

[Daxenberger et al. \(2017\)](#) used six disparate datasets and contrasted the performance of several supervised models. They performed two sets of experiments – in-domain CD (trained and tested on the same dataset) and cross-domain CD (trained on one and tested on another unseen dataset). They learned divergent conceptualisations of claims over cross-domain datasets. [Levy et al. \(2017\)](#) proposed the first unsupervised approach for claim detection. They hypothesised a “claim sentence query” as an ordered triplet: $\langle \text{that} \rightarrow \text{MC} \rightarrow \text{CL} \rangle$. According to the authors, a claim begins with the word ‘that’ and is followed by the main concept (MC) or topic name which is further followed by words from a pre-defined claim lexicon (CL). This approach would not fit well for text stemming from social media platforms owing to a lack of structure and the use of ‘that’ as an offset for claim.

In recent years transformer-based language models have been employed for claim detection. [Chakrabarty et al. \(2019\)](#) used over 5 million self-labeled Reddit comments to fine-tune their model. However, they made no attempt to explicitly encapsulate the structure of a sentence. Recently, the CLEF-2020 shared task ([Barrón-Cedeño et al., 2020](#)) attracted multiple models which are tweaked specifically for claim detection. [Williams et al. \(2020\)](#) bagged the first position in the task using a fine-tuned RoBERTa ([Liu et al., 2019](#)) model with mean pooling and dropout. First runner up of the challenge, [Nikolov et al. \(2020\)](#) used logistic regression on various meta-data tweet features

and a RoBERTa-based prediction. [Cheema et al. \(2020\)](#), the second runner up, incorporated pre-trained BERT embeddings along with POS and dependency tags as features trained using SVM.

Traditional approaches focused primarily on the syntactic representations of claims and textual feature generation, while recent neural methods leverage transformer models. With LESA, we attempt to learn from the past while building for the future – we propose encapsulating syntactic representations in the form of POS tags and dependency sequences along with the semantics of the input text using transformer-based BERT ([Devlin et al., 2019](#)). Another key observation has been the use of highly structured and domain-engineered datasets for training the existing models in claim detection. In the current age of alarming disinformation, we recognise the augmented need for a system that can detect claims in online text independent of its origin, context or domain. Therefore, in addition to considering texts from different online mediums, we incorporate, for the first time, a self-annotated large Twitter dataset to the relatively structured datasets that exist in this field.

3 Proposed Methodology

Traditionally, the narrative on claim detection is built around either syntactic ([Levy et al., 2017](#); [Lippi and Torroni, 2015](#)) or semantic ([Daxenberger et al., 2017](#); [Chakrabarty et al., 2019](#)) properties of the text. However, given our purview on the integration of both, we propose a combined model, LESA, that incorporates exclusively linguistic features leveraged from part-of-speech (POS) tags and dependency tree (DEP) as well as semantic features leveraged from transformer-based model, BERT ([Devlin et al., 2019](#)).

By the virtue of digital media, we generally deal with texts from three kind of environments: (a) a controlled platform where content is pre-reviewed (e.g., news, essays, etc.); (b) a free platform where authors have the freedom to express themselves without any restrictions on the length (e.g., online comments, Wikipedia talk pages); and (c) a free platform with restrictions on the text length (e.g., tweets). The texts in the first category is usually free of any grammatical and typographical mistakes, and thus belong to the non-noisy category. On the other hand, in the third case, texts exhibit a significant amount of noise, in terms of spelling variations, hashtags, emojis, emoticons, abbrevia-

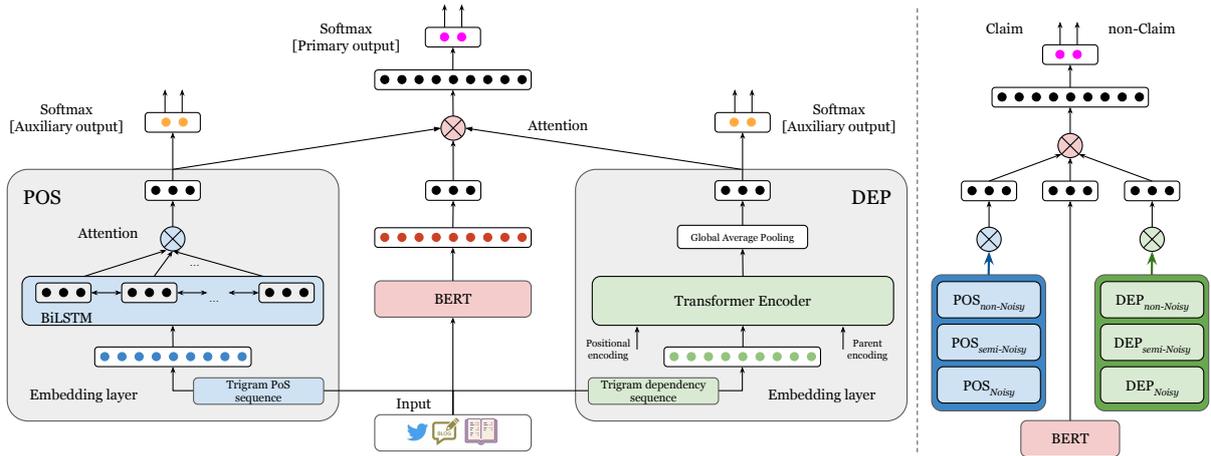


Figure 1: Schematic diagram of our proposed LESA model. The structure on the right is a high level schematic diagram. Structure on the left shows POS and DEP for one viewpoint.

tions, etc., to express the desired information within the permissible limit, thus it belongs to the noisy class. The second case is a mixture of the two extreme cases and hence, constitutes the semi-noisy category. We employ three pre-trained models representing noisy, semi-noisy, and non-noisy data for both POS and dependency-based features. The intuition is to leverage the structure-specific linguistic features in a joint framework.

Domain adaptation from a structured environment to an unstructured one is non-trivial and requires specific processing. Therefore, to ensure generalization, we choose to process each input text from three different viewpoints (structure-based segregation), and intelligently select the contributing features among them through an attention mechanism. We use it to extract the POS and DEP-based linguistic features. Subsequently, we fuse the linguistic and semantic features using another attention layer before feeding it to a multilayer perceptron (MLP) based classifier. The idea is to amalgamate diverse set of features from different perspectives and leverage them for the final classification. A high-level architectural diagram is depicted in Figure 1. We design parallel pillars for each viewpoint (right side of Figure 1) such that the noisy pillar contains pre-trained information from the noisy source and so on. When the common data is passed through the three pillars we hypothesize each pillar’s contribution depending on the type of input data. For example, if the data source is from a noisy platform, we hypothesize that the noisy pillars will have more significance than the other two viewpoints. We demonstrate this effect in Table 5.

A. Part-of-speech (POS) Module

The POS module consists of an embedding layer followed by a BiLSTM and an attention layer to extract the syntactic formation of the input text. We pre-train the POS module for each viewpoint, and later fine-tune them while training the integrated model.

At first, each sequence of tokens $\{x_1, x_2, \dots, x_n\}$ is converted to a sequence of corresponding POS tags resulting into the set $\{p_1, p_2, \dots, p_n\}$. However, the foremost limitation of this modeling strategy is the limited and small vocabulary size of 19 owing to a specific number of POS tags. To tackle this, we resort to using k -grams of the sequence. The sequence of POS tags (with $k = 3$) now becomes $\{(p_0, p_1, p_2), (p_1, p_2, p_3), (p_2, p_3, p_4), \dots, (p_{n-2}, p_{n-1}, p_n), (p_{n-1}, p_n, p_{n+1})\}$, where p_0 and p_{n+1} are dummy tags. Subsequently, a skip-gram model (Mikolov et al., 2013) is trained on the POS-transformed corpus of each dataset, which thereby translates to a POS embedding, E_P .

B. Dependency Tree (DEP) Module

Dependency parsing is the function of abstracting the grammatical assembly of a sequence of tokens $\{x_1, x_2, \dots, x_n\}$ such that there exists a directed relation (dependency), $d(x_i, x_j)$, between any two tokens x_i and x_j , where x_i is the headword and x_j is modified by the headword. We obtain these dependency relations through spaCy¹ which uses the clearNLP guidelines. Initially, each sequence is rendered into a combination of the *dependency-tag* arrangement $\{d_1, d_2, \dots, d_n\}$ and a *parent-position*

¹www.spacy.io

arrangement $\{pp_1, pp_2, \dots, pp_n\}$. Here, each d_j represents a dependency tag, where x_j is modified by x_i , and pp_j is the index of the modifier (headword) x_i .

We then leverage the transformer encoder (Vaswani et al., 2017), where traditionally, a position-based signal is added to each token’s embedding to help encode the placement of tokens. In our modified version, the token sequence is the *dependency-tag* sequence $d_e = \{d_1, d_2, \dots, d_n\}$, wherein a *parent-position* based signal is additionally added to encode the position of the modifier words.

$$d'_e = d_e + [(E_{pp_1}, E_{pp_1}), \dots, (E_{pp_n}, E_{pp_n})] \quad (1)$$

where $d'_e \in \mathbf{R}^{d \times n}$ is the modified dependency embedding of a sequence of length n , E_{p_i} and E_{pp_i} are the encodings for the *token-position* and the *parent-position* (position of token’s modifier), and $(,)$ represent tuple brackets.

This helps us create a flat representation for a dependency graph. The transformer-architecture that we employ comprises of 5 attention heads with an embedding size of 20. Given that there are only a handful of dependency relations, this, still poses the problem of a limited vocabulary size of 37. Having accounted for the parent positions already, we decide to again employ tri-gram sequences $\{(d_0, d_1, d_2), (d_1, d_2, d_3), (d_2, d_3, d_4), \dots, (d_{n-2}, d_{n-1}, d_n), (d_{n-1}, d_n, d_{n+1})\}$ in place of uni-grams.

4 Datasets

A number of datasets exist for the task of claim detection in online text (Peldszus and Stede, 2015; Stab and Gurevych, 2017); however, most of them are formal and structured texts. As we discussed earlier, OSM platforms are overwhelmed with various claim-ridden posts. Despite the abundance of tweets, literature does not suggest any significant effort for detecting claims in Twitter; Arguably, the prime reason is the lack of a large-scale dataset. Recently, a workshop on claim detection and verification in Twitter was organized under CLEF-2020 (Barrón-Cedeño et al., 2020). It had two subtasks related to claim identification with separate datasets. The first dataset consists of 1, 060 COVID-19 tweets for claim detection; whereas, the second one comprises of another 1, 000 tweets for claim retrieval. In total, there were 2, 069 annotated tweets of which 1, 704 had claims and 365 were

CLEF-2020	Our Annotation	
	Non-claim	Claim
Non-claim	301	47
Claim	64	550

Table 2: Confusion matrix highlighting the differences and similarities between Alam et al. (2020) and our annotation guidelines for CLEF-2020 claim dataset.

non-claims. Another recent in-progress dataset on claim detection, which currently has only 305 claim and 199 non-claim tweets, was released by Alam et al. (2020).

Unfortunately, the aforementioned limited instances are insufficient to develop an efficient model. Therefore, we attempt to develop a new and relatively larger dataset for claim detection in OSM platforms. We collected $\sim 40, 000$ tweets from various sources (Carlson, 2020; Smith, 2020; Celin, 2020; Chen et al., 2020; Qazi et al., 2020) and manually annotated them. We additionally included claim detection datasets of Alam et al. (2020) and CLEF-2020 (Barrón-Cedeño et al., 2020) and re-annotated them in accordance with our guidelines. During the cleaning process, we filtered a majority of tweets due to their irrelevancy and duplicacy. To ensure removal of duplicates, we performed manual checking and exhaustive preprocessing.

Data Annotation: To annotate the tweets, we extend and adapt the claim annotation guidelines of Alam et al. (2020). The authors targeted and annotated only a subset of claims, i.e., factually verifiable claims. They did not consider personal opinions, sarcastic comments, implicit claims, or claims existing in a sub-sentence or sub-clause level. Subsequently, we propose our definition of claims and extrapolate the existing guidelines to be more inclusive, nuanced and applicable to a diverse set of claims. Our official definition for claims, adopted from Oxford dictionary², is to *state or assert that something is the case, with or without providing evidence or proof*.

We present the details of annotation guidelines in Gupta et al. (2021). Following the guidelines, we annotated the collected tweets, and to ensure coherence and conformity, we re-annotated the tweets of Alam et al. (2020) and CLEF-2020 (Barrón-Cedeño et al., 2020). It is intriguing to see the differences and similarities of the two guidelines; therefore, we compile a confusion matrix for CLEF-

²<https://www.lexico.com/definition/claim>

Dataset		Text
Noisy	TWR	@realDonaldTrump Does ingesting bleach and shining a bright light in the rectal area really cure #COVID19? Have you tried it? Is that what killed Kim Jong Un? #TrumpIsALaughingStock #TrumpIsALoser
Semi-noisy	OC	*smacks blonde wig on Axel* I think as far as DiZ is concerned, he is very smart but also in certain areas very dumb - - witness the fact that he didn't notice his apprentices were going to turn on him, when some of them (cough Vexen cough) aren't exactly subtle by nature.
	WTP	<u>Not to mention one</u> without any anonymous users TALKING IN CAPITAL LETTERS !!!!!!!
Non-noisy	MT	Tax data that are not made available for free should not be acquired by the state.
	PE	I believe that education is the single most important factor in the development of a country.
	VG	When's the last time you slipped on the concept of truth?
	WD	The public schools are a bad place to send a kid for a good education anymore.

Table 3: One example from each dataset. Underlined text highlights noisy and semi-noisy phrases.

Dataset		Noisy	Semi-noisy			Non-noisy			
		TWR	OC	WTP	MT	PE	VG	WD	
Tr	Cl	7354	623	1030	100	1885	495	190	
	N-cl	1055	7387	7174	301	4499	2012	3332	
Ts	Cl	1296	64	105	12	223	57	14	
	N-cl	189	730	759	36	509	221	221	
Tot	Cl	8650	687	1,135	112	2,108	552	204	
	N-cl	1244	8117	7933	337	5008	2233	3553	

Table 4: Statistics of the datasets (Abbreviations: Cl: Claim, N-Cl: Non-claim, Tr: Train set, Ts: Test set; Tot: Total).

2020 claim dataset, as presented in Table 2. Each tweet in our corpus of 9,894 tweets has been annotated by at least two annotators, with an average Cohen’s kappa inter-annotator agreement (Cohen, 1960) score of 0.62. In case of a disagreement, the third annotator was considered and a majority vote was used for the final label. All annotators were linguists.

Other Datasets: Since we attempt to create a generalized model that is able to detect the presence of a claim in any online text, we accumulate, in addition to the Twitter dataset, six publicly available benchmark datasets: (i) Online Comments (OC) containing Blog threads of LiveJournal (Biran and Rambow, 2011), (ii) Wiki Talk Pages (WTP) (Biran and Rambow, 2011), (iii) German Microtext (MT) (Peldszus and Stede, 2015), (iv) Persuasive Student Essay (PE) (Stab and Gurevych, 2017), (v) Various Genres (VG) containing newspaper editorials, parliamentary records and judicial summaries, and (vi) Web Discourse (WD) containing blog posts or user comments (Habernal and Gurevych, 2015). All datasets utilised in this paper contain English texts only. For German Microtexts (MT), we used the publicly available English translated version published by MT’s original authors (Peldszus and Stede, 2015)). The same was utilized by Chakrabarty et al. (2019).

The datasets are formed by considering text at

the sentence level. For example, in Persuasive Essays (PE) dataset, each essay is broken into sentences and each sentence is individually annotated for a claim. Considering the structure of the input texts in these datasets, we group them into three categories as follows: Noisy (Twitter), Semi-noisy (OC, WTP), Non-noisy (MT, PE, VG, WD). We list one example from each dataset in Table 3. We also highlight the noisy and semi-noisy phrases in Twitter, and OC and WTP datasets respectively. Moreover, we present detailed statistics of all seven datasets in Table 4.

5 Experimental Setup

For all datasets besides twitter, we use the train, validation, and test splits as provided by UKP Lab³. A mutually exhaustive 70:15:15 split was maintained for Twitter dataset. We compute POS embeddings by learning word2vec skip-gram model (Mikolov et al., 2013) on the tri-gram⁴ POS sequence. For the skip-gram model, we set *context window* = 6, *embedding dimension* = 20, and discard the POS sequence with *frequency* ≤ 2. Subsequently, we compute dependency embeddings with *dimension* = 20 using Transformer (Vaswani et al., 2017) encoder with 5 *attention heads*. Please note that the choice of using Bi-LSTM, as oppose to Transformers, for extracting the POS features is empirical⁵.

The outputs of the POS and dependency embedding layers are subsequently fed to a BiLSTM and *GlobalAveragePooling* layers, respectively. Their respective outputs are projected to a 32-dimensional representation for the fusion. We employ HuggingFace’s BERT implementation for computing the tweet representation. The

³<https://tinyurl.com/yyckv29p>

⁴Choice of $n = 3$ is empirical. We report supporting experimental results in Gupta et al. (2021).

⁵Gupta et al. (2021) accompanies the supporting results.

Models	Noisy		Semi-Noisy				Non-Noisy								Wt Avg	
	Twitter		OC		WTP		MT		PE		VG		WD			
	<i>m-F1</i>	<i>c-F1</i>														
BERT	0.60	0.83	0.52	0.24	0.53	0.32	0.70	0.63	0.69	0.64	0.58	0.43	0.48	0.22	0.58	0.73
BERT + POS	0.61	0.84	0.53	0.24	0.54	0.31	0.75	0.69	0.72	0.64	0.59	0.43	0.51	0.24	0.60	0.74
BERT + Dependency	0.59	0.82	0.51	0.23	0.52	0.30	0.79	0.73	0.69	0.62	0.56	0.41	0.48	0.22	0.57	0.72
POS + Dependency	0.45	0.70	0.48	0.19	0.47	0.25	0.57	0.46	0.50	0.45	0.56	0.41	0.44	0.17	0.48	0.61
LESA (Combined-view)	0.61	0.85	0.51	0.23	0.53	0.31	0.77	0.71	0.71	0.64	0.57	0.40	0.48	0.22	0.59	0.75
LESA(768dim)	0.58	0.80	0.52	0.24	0.57	0.29	0.77	0.71	0.73	0.65	0.60	0.43	0.52	0.25	0.59	0.71
LESA(32dim)	0.62	0.85	0.53	0.24	0.55	0.32	0.77	0.69	0.74	0.66	0.68	0.41	0.52	0.25	0.61	0.75

Table 5: Macro F1 (*m-F1*) and Claim-F1 (*c-F1*) for ablation studies.

768-dimensional embeddings is projected to a 32-dimensional representation using linear layers. We progress with the 32-dimensional representation of BERT as we observe no elevation in results on using the 768-dimensional representation, as can be seen in Table 5. Besides, the latter results in ~ 2 million trainable parameters, whereas the former requires ~ 1.2 million trainable parameters. We employ sparse categorical cross-entropy loss with Adam optimizer and use softmax for the final classification.⁶ For evaluation, we adopt **macro-F1** (*m-F1*) and **claim-F1** (*c-F1*) scores used by the existing methods (Daxenberger et al., 2017; Chakrabarty et al., 2019).

We perform our experiments in two setups. In the **first in-domain setup**, we train, validate and test on the same dataset and repeat it for all seven datasets independently. In the **second general-domain setup**, we combine all datasets and train a unified generic model. Subsequently, we evaluate the trained model on all seven datasets individually.

Furthermore, for each experiment, we ensure a balanced training set by down-sampling the dominant class at 1 : 1 ratio. However, we use the original test set for a fair comparison against the existing baselines and state-of-the-art models.

6 Experimental Results

Table 5 shows *m-F1* and *c-F1* for different variants of LESA. We begin with a fine-tuned BERT model and observe the performance on test sets of all seven datasets. On the Twitter dataset, the BERT architecture yields *m-F1* score of 0.60 and *c-F1* score of 0.83. We also report the weighted-average score as 0.58 *m-F1* and 0.73 *c-F1*, in the last two columns of Table 5. Since we hypothesize that claim detection has a strong association with the structure of the text, we amalgamate POS and dependency (DEP) information with the BERT architecture in a step-wise manner. The BERT+POS model reports an increase of 1% *m-F1* and *c-F1*

⁶Gupta et al. (2021) accompanies other hyperparameters.

scores on the Twitter dataset. We observe similar trends in other datasets and the overall weighted-average score as well. We also perform experiments on other permutations, and their results are listed in Table 5. Finally, we combine both POS and DEP modules with the BERT architecture (*aka.* LESA). It obtains improved results for most of the cases, as shown in the last row of Table 5. The best result on average stands at 0.61 *m-F1* and 0.75 *c-F1* for the proposed LESA model. This serves as a testament to our hypothesis, validating our assumption that combining syntactic and semantic representations leads to better detection of claims.

In all aforementioned experiments, we use our pre-defined concept of three viewpoints, i.e., noisy, semi-noisy and non-noisy. Therefore, for completeness, we also construct a combined viewpoint which does not contain any structure-specific pillar in POS or DEP branches. The results from this ablation experiment are reported in LESA (Combined-view) row. We observe that the combined-view results are inferior to the variant with separate viewpoints for each component (c.f. second last and last row of Table 5 respectively). Thus, providing attention to datasets based on the noise in their content is demonstrated by a significant increase of $\sim 2\%$ *m-F1* from combined viewpoint to separate viewpoints experiment.

A. Baselines and Comparative Analysis

We employ the following baselines (some of them being state-of-the-art systems for claim detection and text classification): \triangleright **XLNet** (Yang et al., 2019): It is similar to the BERT model, where we fine-tune XLNet for the claim detection; \triangleright **Accenture** (Williams et al., 2020): A RoBERTa-based system that ranked first in the CLEF-2020 claim detection task (Barrón-Cedeño et al., 2020); \triangleright **Team Alex** (Nikolov et al., 2020): The second-ranked system at CLEF-2020 task that fused tweet meta-data into RoBERTa for the final prediction; \triangleright **Check-Square** (Cheema et al., 2020): An SVM-based system designed on top of pre-trained BERT em-

Models	Noisy		Semi-Noisy				Non-Noisy								Wt Avg	
	Twitter		OC		WTP		MT		PE		VG		WD			
	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>		
BERT	0.50	0.67	0.50	0.24	0.36	0.27	0.75	0.69	0.73	0.67	0.61	0.48	0.48	0.23	0.52	0.62
XLNet	0.52	0.70	0.45	0.24	0.55	0.30	0.49	0.43	0.71	0.64	0.53	0.43	0.51	0.12	0.54	0.64
Accenture	0.48	0.15	0.44	0.16	0.50	0.23	0.48	0.28	0.45	0.11	0.39	0.27	0.34	0.11	0.46	0.15
Team Alex	0.70	0.88	0.46	0.23	0.52	0.21	0.75	0.69	0.64	0.32	0.38	0.60	0.34	0.59	0.76	0.76
Check Square	0.12	0.02	0.49	0.25	0.39	0.26	0.57	0.32	0.47	0.11	0.32	0.37	0.76	0.56	0.35	0.07
CrossDomain	0.67	0.84	0.61*	0.26*	0.59*	0.29*	0.79*	0.67*	0.74*	0.61*	0.66*	0.45*	0.63*	0.29*	0.65	0.74
CrossDomain†	0.67	0.84	0.50	0.24	0.52	0.27	0.85	0.79	0.71	0.63	0.60	0.46	0.59	0.31	0.61	0.74
LESA	0.67	0.89	0.51	0.26	0.57	0.33	0.80	0.71	0.73	0.67	0.68	0.52	0.61	0.35	0.63	0.79

Table 6: Macro F1 (*m-F1*) and F1 for claims (*c-F1*) in the in-domain setup. For CrossDomain, the asterisk (*) indicates results taken from Daxenberger et al. (2017) and the dagger (†) represents the reproduced results.

Model	Noisy		Semi-Noisy				Non-Noisy								Wt Avg	
	Twitter		OC		WTP		MT		PE		VG		WD			
	<i>m-F1</i>	<i>c-F1</i>														
BERT	0.60	0.83	0.52	0.24	0.53	0.32	0.70	0.63	0.69	0.64	0.58	0.43	0.48	0.22	0.58	0.73
XLNet	0.59	0.81	0.56	0.28	0.57	0.29	0.68	0.69	0.71	0.64	0.61	0.44	0.52	0.25	0.59	0.72
Accenture	0.49	0.43	0.31	0.12	0.40	0.18	0.36	0.13	0.51	0.36	0.38	0.17	0.37	0.04	0.43	0.38
Team Alex	0.54	0.75	0.54	0.25	0.54	0.30	0.71	0.65	0.71	0.63	0.61	0.43	0.48	0.19	0.57	0.67
Check Square	0.58	0.82	0.51	0.23	0.48	0.28	0.56	0.53	0.68	0.59	0.56	0.38	0.47	0.21	0.56	0.72
CrossDomain	0.65	0.82	0.57	0.27	0.53	0.28	0.71	0.63	0.66	0.57	0.61	0.43	0.52	0.25	0.60	0.71
LESA	0.62	0.85	0.53	0.24	0.55	0.32	0.77	0.69	0.74	0.66	0.68	0.41	0.52	0.25	0.61	0.75

Table 7: Macro F1 (*m-F1*) and Claim-F1 (*c-F1*) in the general-domain setup.

Models	Noisy		Semi-Noisy		Non-Noisy	
	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>	<i>m-F1</i>	<i>c-F1</i>
BERT	0.60	0.83	0.52	0.29	0.63	0.58
XLNet	0.59	0.81	0.57	0.29	0.65	0.59
Accenture	0.49	0.43	0.36	0.16	0.45	0.30
Team Alex	0.54	0.75	0.54	0.28	0.65	0.57
CheckSquare	0.58	0.82	0.49	0.26	0.61	0.53
CrossDomain	0.65	0.82	0.55	0.28	0.63	0.53
LESA	0.62	0.85	0.54	0.29	0.69	0.60

Table 8: Category-wise weighted-average F1 scores.

beddings in addition to incorporating POS and dependency tags as external features. \triangleright **CrossDomain** (Daxenberger et al., 2017): Among several variations reported in the paper, their best model incorporates CNN (random initialization) for the detection. We reproduce the top submissions from CLEF-2020 challenge using the best performing models mentioned in the referenced papers. Code for CheckSquare was provided online. For Accenture and Team Alex we reproduce their methods using the hyper-parameters mentioned in the paper. We evaluate all baselines using the same train and test set as for LESA.

We report our comparative analysis for the in-domain setup in Table 6. We observe that LESA obtains best *c-F1* scores for six out of seven datasets. Additionally, it achieves a weighted average *c-F1* of 0.79 which is 3.95% improvement over the best performing baseline. In terms of *m-F1* values, our weighted average ranks second next to CrossDomain. We reproduced CrossDomain baseline using their GitHub code (UKPLab). If the reproduced values are considered, our model outperforms all

other models in *m-F1* value as well.

Similarly, we compile the results for the general-domain setup in Table 7. In the non-noisy category, LESA obtains better *m-F1* scores than three of the four state-of-the-art systems, i.e., it reports 0.77, 0.74, and 0.68 *m-F1* scores compared to 0.71, 0.71, and 0.61 *m-F1* scores of the comparative systems on MT, PE, and VG test sets, respectively. On WD, we observe similar *m-F1* and *c-F1* scores for both the best baseline and LESA. On the datasets in other categories, we observe comparative *m-F1* scores; however, none of the baselines are consistent across all dataset – e.g., CrossDomain (Daxenberger et al., 2017) reports the best *m-F1* scores on Twitter and OC, but yields (joint) fourth-best performance on WTP. Moreover, LESA yields the best *m-F1* score across the seven datasets on average with $\geq 1\%$ improvements. On the other hand, we obtain best *c-F1* scores for five out of seven datasets. In addition, LESA reports overall *c-F1* of 0.75 with a significant improvement of $\geq 3\%$. Using a paired T-test, LESA showed significant statistical improvement compared against BERT in *m-F1* and *c-F1* for the noisy dataset with p-values .00017 and $<.00001$ respectively. Results were also significant for *m-F1* and *c-F1* for PE and *m-F1* for WD. The small sample size in some datasets like MT and VG does not allow a reliable calculation of test statistics.

Since our work intends to developing a model that is able to detect claims irrespective of the source and origin of text, we also analyse the

		Example	Gold	Prediction	
				LESA	CrossDomain
TWR	x_1	28 coronaoutbreak cases thus far in india italian tourists 16 their driver 1 kerala 3 cureddischarged agra 6 delhi 1 noida school dad telangana 1 coronavirusindia	1	0	0
	x_2	can we just call this a cure now	0	0	1
MT	x_3	Besides it should be in the interest of the health insurers to recognize alternative medicine as treatment, since there is a chance of recovery.	0	1	1
PE	x_4	On the other hand, fossil fuels are abundant and inexpensive in many areas	0	1	1
	x_5	Daily exercise will help also to develop children’s brain function.	1	1	0
OC	x_6	Skinny Puppy is headlining Festival Kinetik !	0	1	1
	x_7	I guess I’m not desensitized enough to just forget about people being murdered in my neighborhood.	1	1	0
WD	x_8	No wonder 50 million babies have been aborted since 1973 .	0	1	1

Table 9: Error analysis of the outputs. Red texts highlight errors.

weighted-average scores for each category in Table 8. We observe that LESA obtains the best c -F1 scores in each category, in addition to the best m -F1 score in non-noisy category as well. For the other two categories, LESA yields comparative performances. The results are better for noisy data than for non-noisy owing to the small size and skewness against claims in the latter’s test set. Therefore, misclassification of a single claim causes severe penalization to c -F1.

B. Error Analysis

It is apparent from the results that all systems (including LESA) committed some errors in claim detection. Thus, in this section, we explore where our system misclassified the inputs by analysing some examples. Table 9 presents a few instances along with the gold labels and the predictions of the best-performing baseline, CrossDomain (Daxenberger et al., 2017), for comparison. In some cases, both LESA and CrossDomain failed to classify the instances correctly, whereas in others, LESA classifies the instances correctly but CrossDomain could not. We also report intuitions for the misclassification by LESA in some cases. The presence of numbers and statistics could be the reason behind the misclassifications in examples x_1 and x_8 . Example x_3 contains two weak phrases (‘*alternative medicine as treatment*’ and ‘*there is a chance of recovery*’) which are most likely the cause of misclassification. The former might have been interpreted as suggestion backed up by some evidence, while in the latter phrase, LESA might have misinterpreted the optimism with claim. Furthermore, the phrase ‘*fossil fuels are abundant*’ in example x_4 reflects world knowledge instead of a claim, as interpreted by LESA.

7 Conclusion

In this paper, we addressed the task of claim detection from online posts. To do this, we proposed a generic and novel deep neural framework, LESA, that leverages the pre-trained language model and two linguistic features, corresponding to the syntactic properties of input texts, for the final classification. Additionally, we tackled the texts from distinct sources for the claim detection task in a novel way. In particular, we categorized the input text as noisy, non-noisy, and semi-noisy based on the source, and modeled them separately. Subsequently, we fused them together through an attention module as the combined representation.

One of the major bottlenecks of claim detection in online social media platforms is the lack of qualitative annotation guidelines and a sufficiently large annotated dataset. Therefore, we developed a large Twitter dataset of $\sim 10,000$ manually annotated tweets for claim detection. In addition to our twitter dataset, we employed six benchmark datasets (representing either semi-noisy or non-noisy input channels) for evaluation of the proposed model. We compared the performance of LESA against four state-of-the-art systems and two pre-trained language models. Comparison showed the superiority of the proposed model with $\geq 3\%$ claim-F1 and $\geq 1\%$ macro-F1 improvements compared to the best performing baselines on average. As a by-product of the study, we released a comprehensive guideline for claim annotation.

Acknowledgement

We would like to thank Rituparna and LCS2 members for helping in data annotation. The work was partially supported by Accenture Research Grant, Ramanujan Fellowship, and CAI, IIIT-Delhi.

References

Immunizing the public against misinformation.

- Alam, Shaden Shaar, Fahim Dalvi, Hassan Sajjad, Alex Nikolov, Hamdy Mubarak, Giovanni Da San Martino, Ahmed Abdelali, Nadir Durrani, Kareem Darwish, and Preslav Nakov. 2020. [Fighting the covid-19 infodemic: Modeling the perspective of journalists, fact-checkers, social media platforms, policy makers, and the society.](#) *arXiv preprint arXiv:2005.00033*.
- Alberto Barrón-Cedeño, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain, Reem Suwaileh, and Fatima Haouari. 2020. Check-that! at clef 2020: Enabling the automatic identification and verification of claims in social media. In *Advances in Information Retrieval*, pages 499–507, Cham. Springer International Publishing.
- O. Biran and O. Rambow. 2011. Identifying justifications in written dialogs. In *2011 IEEE Fifth International Conference on Semantic Computing*, pages 162–168.
- Or Biran and Owen Rambow. 2011. [Identifying justifications in written dialogs by classifying text as argumentative.](#) *International Journal of Semantic Computing*, 05:363–381.
- Carlson. 2020. [Coronavirus tweets.](#)
- Sven Celin. 2020. [Covid-19 tweets afternoon 31.03.2020.](#)
- Tuhin Chakrabarty, Christopher Hidey, and Kathleen McKeown. 2019. Imho fine-tuning improves claim detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 558–563.
- Gullal S. Cheema, Sherzod Hakimov, and Ralph Ewerth. 2020. [Check_square at checkthat! 2020: Claim detection in social media via fusion of transformer and syntactic features.](#) *arXiv: 2007.10534*.
- Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. [Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set.](#) *JMIR Public Health Surveill*, 6(2):e19273.
- J. Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37.
- Johannes Daxenberger, Steffen Eger, Ivan Habernal, Christian Stab, and Iryna Gurevych. 2017. What is the essence of a claim? cross-domain claim identification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2055–2066.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Trudy Govier. 2013. *A practical study of argument.* Cengage Learning.
- Shreya Gupta, Parantak Singh, Megha Sundriyal, Md Shad Akhtar, and Tanmoy Chakraborty. 2021. [LESA: Linguistic Encapsulation and Semantic Amalgamation Based Generalised Claim Detection from Online Content \(Supplementary\).](#)
- Ivan Habernal and Iryna Gurevych. 2015. [Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse.](#) In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2127–2137, Lisbon, Portugal. Association for Computational Linguistics.
- Ran Levy, Yonatan Bilu, Daniel Hershcovich, Ehud Aharoni, and Noam Slonim. 2014. Context dependent claim detection. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1489–1500.
- Ran Levy, Shai Gretz, Benjamin Sznajder, Shay Hummel, Ranit Aharonov, and Noam Slonim. 2017. Un-supervised corpus-wide claim detection. In *Proceedings of the 4th Workshop on Argument Mining*, pages 79–84.
- Marco Lippi and Paolo Torroni. 2015. Context-independent claim detection for argument mining. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 185–191.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Hanyu Chwe Alexi Quintana Roy H. Perlis David Lazer James Druckman Mauricio Santillana Jennifer Lin John Della Volpe Matthew Simonson Jon Green Matthew A. Baum, Katherine Ognyanova. [The state of the nation: A 50-state covid-19 survey report 14: Misinformation and vaccine acceptance.](#)
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space.](#) *arXiv: 1301.3781*.
- Alex Nikolov, Giovanni Da San Martino, Ivan Koychev, and Preslav Nakov. 2020. [Team alex at clef checkthat! 2020: Identifying check-worthy tweets with transformer models.](#) *arXiv:2009.02931*.

- Andreas Peldszus and Manfred Stede. 2015. [Joint prediction in MST-style discourse parsing for argumentation mining](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 938–948, Lisbon, Portugal. Association for Computational Linguistics.
- Umair Qazi, Muhammad Imran, and Ferda Ofli. 2020. Geocov19: a dataset of hundreds of millions of multilingual covid-19 tweets with location information. *SIGSPATIAL Special*, 12(1):6–15.
- Sara Rosenthal and Kathleen McKeown. 2012. Detecting opinionated claims in online discussions. In *2012 IEEE sixth international conference on semantic computing*, pages 30–37. IEEE.
- Shane Smith. 2020. [Coronavirus \(covid19\) tweets - early april](#).
- Christian Stab and Iryna Gurevych. 2017. [Parsing argumentation structures in persuasive essays](#). *Computational Linguistics*, 43(3):619–659.
- Stephen E Toulmin. 2003. *The uses of argument*. Cambridge university press.
- UKPLab. [Ukplab/emnlp2017-claim-identification](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Andrew Whalen and Kevin Laland. 2015. Conformity biased transmission in social networks. *Journal of Theoretical Biology*, 380:542–549.
- Evan Williams, Paul Rodrigues, and Valerie Novak. 2020. [Accenture at checkthat! 2020: If you say so: Post-hoc fact-checking of claims using transformer-based models](#). *arXiv: 2009.02431*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5753–5763.