

Contradiction Detection for Rumorous Claims

Piroska Lendvai

Computational Linguistics
Saarland University
Saarbrücken, Germany
piroska.r@gmail.com

Uwe D. Reichel

Research Institute for Linguistics
Hungarian Academy of Sciences
Budapest, Hungary
uwe.reichel@nytud.mta.hu

Abstract

The utilization of social media material in journalistic workflows is increasing, demanding automated methods for the identification of mis- and disinformation. Since textual contradiction across social media posts can be a signal of rumorosity, we seek to model how claims in Twitter posts are being textually contradicted. We identify two different contexts in which contradiction emerges: its broader form can be observed across independently posted tweets and its more specific form in threaded conversations. We define how the two scenarios differ in terms of central elements of argumentation: claims and conversation structure. We design and evaluate models for the two scenarios uniformly as 3-way Recognizing Textual Entailment tasks in order to represent claims and conversation structure implicitly in a generic inference model, while previous studies used explicit or no representation of these properties. To address noisy text, our classifiers use simple similarity features derived from the string and part-of-speech level. Corpus statistics reveal distribution differences for these features in contradictory as opposed to non-contradictory tweet relations, and the classifiers yield state of the art performance.

1 Introduction and Task Definition

Assigning a veracity judgment to a claim appearing on social media requires complex procedures including reasoning on claims aggregated from multiple microposts, to establish claim veracity status (resolved or not) and veracity value (true or false). Until resolution, a claim circulating on social media platforms is regarded as a rumor (Mendoza et al., 2010). The detection of contradicting and disagreeing microposts supplies important cues to claim veracity processing procedures. These tasks are challenging to automate not only due to the surface noisiness and conciseness of user generated content. One complicating factor is that claim denial or rejection is linguistically often not explicitly expressed, but appears without classical rejection markers or modality and speculation cues (Morante and Sporleder, 2012). Explicit and implicit contradictions furthermore arise in different contexts: in threaded discussions, but also across independently posted messages; both contexts are exemplified in Figure 1 on Twitter data.

Language technology has not yet solved the processing of contradiction-powering phenomena, such as negation (Morante and Blanco, 2012) and stance detection (Mohammad et al., 2016), where stance is defined to express speaker favorability towards an evaluation target, usually an entity or concept. In the veracity computation scenario we can speak of *claim targets* that are above the entity level: targets are entire rumors, such as '11 people died during the Charlie Hebdo attack'. Contradiction and stance detection have so far only marginally been addressed in the veracity context (de Marneffe et al., 2012; Ferreira and Vlachos, 2016; Lukasik et al., 2016).

We propose investigating the advantages of incorporating claim target and conversation context as premises in the Recognizing Textual Entailment (RTE) framework for contradiction detection in rumorous tweets. Our goals are manifold: (a) to offer richer context in contradiction modeling than what would be available on the level of individual tweets, the typical unit of analysis in previous studies; (b) to train and test supervised classifiers for contradiction detection in the RTE inference framework; (c) to address contradiction detection at the level of text similarity only, as opposed to semantic similarity (Xu et al., 2015); (d) to distinguish and focus on two different contradiction relationship types, each involving specific combinations of claim target mention, polarity, and contextual proximity, in particular:

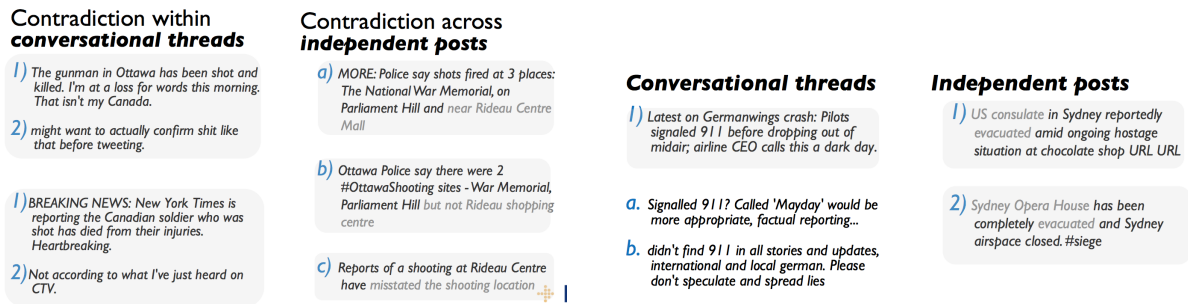


Figure 1: Explicit (far left: in threads, left: in independent posts) vs implicit (right: in threads, far right: in independent posts) contradictions in threaded discussions and in independent posts.

1. **Independent contradictions:** Contradictory relation between independent posts, in which two tweets contain different information about the same claim target that cannot simultaneously hold. The two messages are independently posted, i.e., not occurring within a structured conversation.
2. **Disagreeing replies:** Contradictory relation between a claim-originating tweet and a direct reply to it, whereby the reply expresses disagreement with respect to the claim-introducing tweet.

Contradiction between independently posted tweets typically arises in a broad discourse setting, and may feature larger distance in terms of time, space, and source of information. The claim target is mentioned in both posts in the contradiction pair, since these posts are uninformed about each other or assume uninformedness of the reader, and thus do not or can not make coreference to their shared claim target. Due to the same reason, the polarity of both posts with respect to the claim can be identical. Texts paired in this type of contradiction resemble those of the recent Interpretable Semantic Similarity shared task (Agirre et al., 2016) that calls to identify five chunk level semantic relation types (equivalence, opposition, specificity, similarity or relatedness) between two texts that originate from headlines or captions.

Disagreeing replies are more specific instances of contradiction: contextual proximity is small and trivially identifiable by means of e.g. social media platform metadata, for example the property encoding the tweet ID to which the reply was sent, which in our setup is always a thread-initiating tweet. The claim target is by definition assumed to be contained in the thread-initiating tweet (sometimes termed as claim- or rumor-source tweet). It can be the case that the claim target is not contained in the reply, which can be explained by the proximity and thus shared context of the two posts. The polarity values in source and reply must by definition be different; we refer to this scenario as Disagreeing replies. Importantly, replies may not contain a (counter-)claim on their own but some other form to express disagreement and polarity – for example in terms of speculative language use, or the presence of extra-linguistic cues such as a URL pointing to an online article that holds contradictory content. Such cues are difficult to decode for a machine, and their representation for training automatic classifiers is largely unexplored. Note that we do not make assumptions or restrictions about how the claim target is encoded textually in any of the two scenarios.

In this study, we tackle both contradiction types using a single generic approach: we recast them as three-way RTE tasks on pairs of tweets. The findings of our previous study in which semantic inference systems with sophisticated, corpus-based or manually created syntactico-semantic features were applied to contradiction-labeled data indicate the lack of robust syntactic and semantic analysis for short and noisy texts; cf. Chapter 3 in (Lendvai et al., 2016b). This motivates our current simple text similarity metrics in search of alternative methods for the contradiction processing task.

In Section 2 we introduce related work and resources, in Sections 3 and 4 present and motivate the collections and the features used for modeling. After the description of method and scores in Section 5, findings are discussed in Section 6.

2 Related work and resources

Recognizing Textual Entailment (RTE) Processing semantic inference phenomena such as contradiction, entailment and stance between text pairs has been gaining momentum in language technology. Inference has been suggested to be conveniently formalized in the generic framework of RTE¹ (Dagan et al., 2006). As an improvement over the binary Entailment vs Non-entailment scenario, three-way RTE has appeared but is still scarcely investigated (Ferreira and Vlachos, 2016; Lendvai et al., 2016a). The *Entailment* relation between two text snippets holds if the claim present in snippet B can be concluded from snippet A. The *Contradiction* relation applies when the claim in A and the claim in B cannot be simultaneously true. The *Unknown* relation applies if A and B neither entail nor contradict each other.

The RTE-3 benchmark dataset is the first resource that labels paired text snippets in terms of 3-way RTE judgments (De Marneffe et al., 2008), but it is comprised of general newswire texts. Similarly, the new large annotated corpus used for deep models for entailment (Bowman et al., 2015) labeled text pairs as Contradiction are too broadly defined, i.e., expressing generic semantic incoherence rather than semantically motivated polarization and mismatch that we are after, which questions its utility in the rumor verification context.

As far as contradiction processing is concerned, accounting for negation in RTE is the focus of a recent study (Madhumita, 2016), but it is still set according to the binary RTE setup. A standalone contradiction detection system was implemented by (De Marneffe et al., 2008), using complex rule-based features. A specific RTE application, the Excitement Open Platform² (Padó et al., 2015) has been developed to provide a generic platform for applied RTE. It integrates several entailment decision algorithms, while only the Maximum Entropy-based model (Wang and Neumann, 2007) is available for 3-way RTE classification. This model implements state-of-the-art linguistic preprocessing augmented with lexical resources (WordNet, VerbOcean), and uses the output of part-of-speech and dependency parsing in its structure-oriented, overlap-based approach for classification and was tested for both our tasks as explained in (Lendvai et al., 2016b).

Stance detection Stance classification and stance-labeled corpora are relevant for contradiction detection, because the relationship of two texts expressing opposite stance (positive and negative) can in some contexts be judged to be contradictory: this is exactly what our Disagreeing reply scenario covers. Stance classification for rumors was introduced by (Qazvinian et al., 2011) where the goal was to generate a binary (for or against) stance judgment. Stance is typically classified on the level of individual tweets: reported approaches predominantly utilize statistical models, involving supervised machine learning (de Marneffe et al., 2012) and RTE (Ferreira and Vlachos, 2016). Another relevant aspect of stance detection for our current study is the presence of the stance target in the text to be stance-labeled. A recent shared task on social media data defined separate challenges depending on whether target-specific training data is included in the task or not (Mohammad et al., 2016); the latter requires additional effort to encode information about the stance target, cf. e.g. (Augenstein et al., 2016). The PHEME project released a new stance-labeled social media dataset (Zubiaga et al., 2015) that we also utilize as described next.

3 Data

The two datasets corresponding to our two tasks are drawn from a freely available, annotated social media corpus³ that was collected from the Twitter platform⁴ via filtering on event-related keywords and hashtags in the Twitter Streaming API. We worked with English tweets related to four events: the Ottawa shooting⁵, the Sydney Siege⁶, the Germanwings crash⁷, and the Charlie Hebdo shooting⁸. Each event in

¹http://www.aclweb.org/aclwiki/index.php?title=Recognizing_Textual_Entailment

²<http://hltfbk.github.io/Excitement-Open-Platform>

³https://figshare.com/articles/PHEME_rumour_scheme_dataset_journalism_use_case/2068650

⁴twitter.com

⁵https://en.wikipedia.org/wiki/2014_shootings_at_Parliament_Hill,_Ottawa

⁶https://en.wikipedia.org/wiki/2014_Sydney_hostage_crisis

⁷https://en.wikipedia.org/wiki/Germanwings_Flight_9525

⁸https://en.wikipedia.org/wiki/Charlie_Hebdo_shooting

event	ENT	CON	UNK	#uniq clms	#uniq tws	ENT	CON	UNK	#uniq clms	#uniq tws
chebdo	143	34	486	36	736	647	427	866	27	199
gwings	39	6	107	13	176	461	257	447	4	29
ottawa	79	37	292	28	465	555	377	168	18	125
ssiege	112	59	456	37	697	332	317	565	21	143
	373	136	1341	114	2074	1995	1378	2046	70	496

Table 1: *Threads* (left) and *iPosts* (right) RTE datasets compiled from 4 crisis events: amount of pairs per entailment type (*ENT*, *CON*, *UNK*), amount of unique rumorous claims (*#uniq clms*) used for creating the pairs, amount of unique tweets discussing these claims (*#uniq tws*).

the corpus was pre-annotated as explained in (Zubiaga et al., 2015) for several rumorous claims⁹ – officially not yet confirmed statements lexicalized by a concise proposition, e.g. "Four cartoonists were killed in the Charlie Hebdo attack" and "French media outlets to be placed under police protection". The corpus collection method was based on a retweet threshold, therefore most tweets originate from authoritative sources using relatively well-formed language, whereas replying tweets often feature non-standard language use.

Tweets are organized into threaded conversations in the corpus and are marked up with respect to stance, certainty, evidentiality, and other veracity-related properties; for full details on released data we refer to (Zubiaga et al., 2015). The dataset on which we run disagreeing reply detection (henceforth: *Threads*) was converted by us to RTE format based on the threaded conversations labeled in this corpus. We created the *Threads* RTE dataset drawing on manually pre-assigned Response Type labels by (Zubiaga et al., 2015) that were meant to characterize source tweet – replying tweet relations in terms of four categories. We mapped these four categories onto three RTE labels: a reply pre-labeled as *Agreed* with respect to its source tweet was mapped to *Entailment*, a reply pre-labeled as *Disagreed* was mapped to *Contradiction*, while replies pre-labeled as *AppealforMoreInfo* and *Comment* were mapped to *Unknown*. Only direct replies to source tweets relating to the same four events as in the independent posts RTE dataset were kept. There are 1,850 tweet pairs in this set; the proportion of contradiction instances amounts to 7%. The *Threads* dataset holds *CON*, *ENT* and *UNK* pairs as exemplified below. Conform the RTE format, pair elements are termed *text* and *hypothesis* – note that directionality between *t* and *h* is assumed as symmetric in our current context so *t* and *h* are assigned based on token-level length.

- **CON** <t>We understand there are two gunmen and up to a dozen hostages inside the cafe under siege at Sydney.. ISIS flags remain on display 7News</t> <h>not ISIS flags</h>
- **ENT** <t>Report: Co-Pilot Locked Out Of Cockpit Before Fatal Plane Crash URL Germanwings URL</t> <h>This sounds like pilot suicide.</h>
- **UNK** <t>BREAKING NEWS: At least 3 shots fired at Ottawa War Memorial. One soldier confirmed shot - URL URL</t> <h>All our domestic military should be armed, now.</h>.

The independently posted tweets dataset (henceforth: *iPosts*) that we used for contradiction detection between independently emerging claim-initiating tweets is described in (Lendvai et al., 2016a). This collection is holds 5.4k RTE pairs generated from about 500 English tweets using semi-automatic 3-way RTE labeling, based on semantic or numeric mismatches between the rumorous claims annotated in the data. The proportion of contradictory pairs (*CON*) amounts to 25%. The two collections are quantified in Table 1. *iPosts* dataset examples are given below.

- **CON** <t>12 people now known to have died after gunmen stormed the Paris HQ of magazine CharlieHebdo URL URL</t> <h>Awful. 11 shot dead in an assault on a Paris magazine. URL CharlieHebdo URL</h>
- **ENT** <t>SYDNEY ATTACK - Hostages at Sydney cafe - Up to 20 hostages - Up to 2 gunmen - Hostages seen holding ISIS flag DEVELOPING..</t> <h>Up to 20 held hostage in Sydney Lindt Cafe siege URL URL</h>
- **UNK** <t>BREAKING: NSW police have confirmed the siege in Sydney's CBD is now over, a police officer is reportedly among the several injured.</t> <h>Update: Airspace over Sydney has been shut down. Live coverage: URL sydneyseige</h>.

⁹*Rumor*, *rumorous claim* and *claim* are used interchangeably throughout the paper to refer to the same concept.

4 Text similarity features

Data preprocessing on both datasets included screen name and hashtag sign removal and URL masking. Then, for each tweet pair we extracted vocabulary overlap and local text alignment features. The tweets were part-of-speech-tagged using the Balloon toolkit (Reichel, 2012) (PENN tagset, (Marcus et al., 1999)), normalized to lowercase and stemmed using an adapted version of the Porter stemmer (Porter, 1980). Content words were defined to belong to the set of nouns, verbs, adjectives, adverbs, and numbers, and were identified by their part of speech labels. All punctuation was removed.

4.1 Vocabulary overlap

Vocabulary overlap was calculated for content word stem types in terms of the Cosine similarity and the F1 score. The Cosine similarity of two tweets is defined as $C(X, Y) = \frac{|X \cap Y|}{\sqrt{|X| \cdot |Y|}}$, where X and Y denote the sets of content word stems in the tweet pair.

The F1 score is defined as the harmonic mean of precision and recall. Precision and recall here refer to covering the vocabulary X of one tweet by the vocabulary Y of another tweet (or vice versa). It is given by $F1 = 2 \cdot \frac{\frac{|X \cap Y|}{|X|} \cdot \frac{|X \cap Y|}{|Y|}}{\frac{|X \cap Y|}{|X|} + \frac{|X \cap Y|}{|Y|}}$. Again the vocabularies X and Y consist of stemmed content words. Just like the Cosine index, the F1 score is a symmetric similarity metric.

These two metrics are additionally applied to the content word POS label inventories within the tweet pair, which gives the four features *cosine*, *cosine_pos*, *f_score*, and *f_score_pos*, respectively.

4.2 Local alignment

The amount of stemmed word token overlap was measured by applying local alignment of the token sequences using the Smith-Waterman algorithm (Smith and Waterman, 1981). We chose a score function rewarding zero substitutions by +1, and punishing insertions, deletions, and substitutions each by 0-reset. Having filled in the score matrix H , alignment was iteratively applied the following way:

while $\max(H) \geq t$

- trace back from the cell containing this maximum the path leading to it until a zero-cell is reached
- add the substring collected on this way to the set of aligned substrings
- set all traversed cells to 0.

The threshold t defines the required minimum length of aligned substrings. It is set to 1 in this study, thus it supports a complete alignment of any pair of permutations of x . The traversed cells are set to 0 after each iteration step to prevent that one substring would be related to more than one alignment pair. This approach would allow for two restrictions: to prevent cross alignment not just the traversed cells $[i, j]$ but for each of these cells its entire row i and column j needs to be set to 0. Second, if only the longest common substring is of interest, then the iteration is trivially to be stopped after the first step. Since we did not make use of these restrictions, in our case the alignment supports cross-dependencies and can be regarded as an iterative application of a longest common substring match.

From the substring pairs in tweets x and y aligned this way, we extracted two text similarity measures:

- *laProp*: the proportion of locally aligned tokens over both tweets $\frac{m(x)+m(y)}{n(x)+n(y)}$
- *laPropS*: the proportion of aligned tokens in the shorter tweet $\frac{m(\hat{z})}{n(\hat{z})}$, $\hat{z} = \arg \min_{z \in \{x, y\}} [n(z)]$, where $n(z)$ denotes the number of all tokens and $m(z)$ the number of aligned tokens in tweet z .

4.3 Corpus statistics

Figures 2 and 3 show the distribution of the features introduced above each for a selected event in both datasets. Each figure half represents a dataset; each subplot shows the distribution of a feature in dependence of the three RTE classes for the selected event in that dataset.

The plots indicate a general trend over all events and datasets: the similarity features reach highest values for the ENT class, followed by CON and UNK. Kruskal-Wallis tests applied separately for all combinations of features, events and datasets confirmed these trends, revealing significant differences for all boxplot triplets ($p < 0.001$ after correction for type 1 errors in this high amount of comparisons using

the false discovery rate method of (Benjamini and Yekutieli, 2001)). Dunnett post hoc tests however clarified that for 16 out of 72 comparisons (all POS similarity measures) only UNK but not ENT and CON differ significantly ($\alpha = 0.05$). Both datasets contain the same amount of non-significant cases. Nevertheless, these trends are encouraging to test whether an RTE task can be addressed by string and POS-level similarity features alone, without syntactic or semantic level tweet comparison.

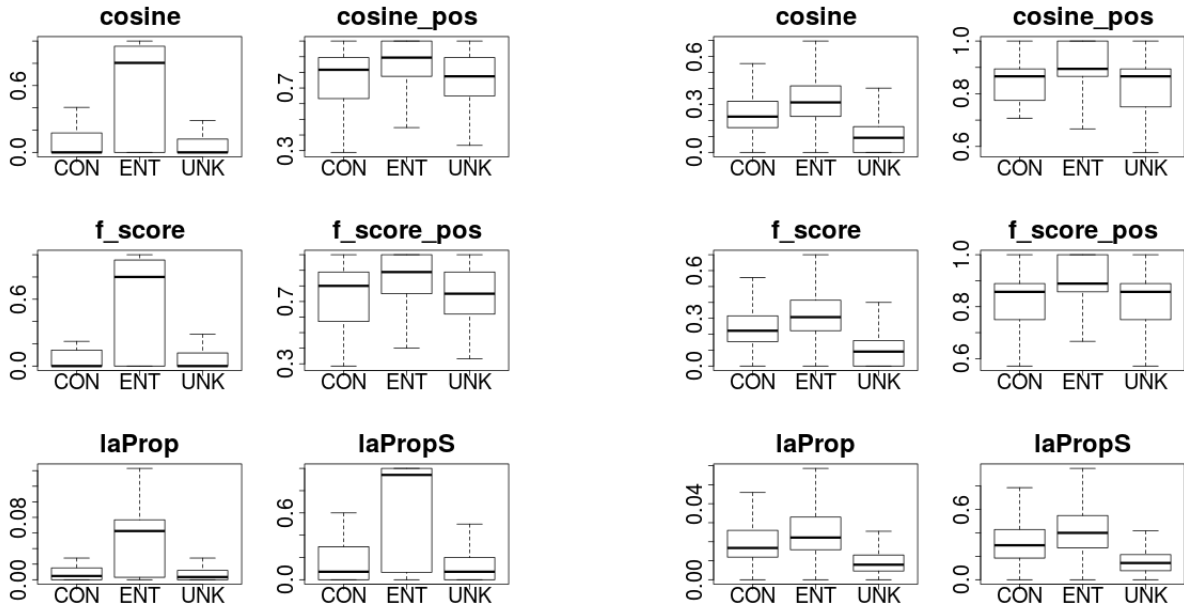


Figure 2: Distributions of the similarity metrics by tweet pair class for the event *chebdo* in the *Threads* (left) and the *iPosts* dataset (right).

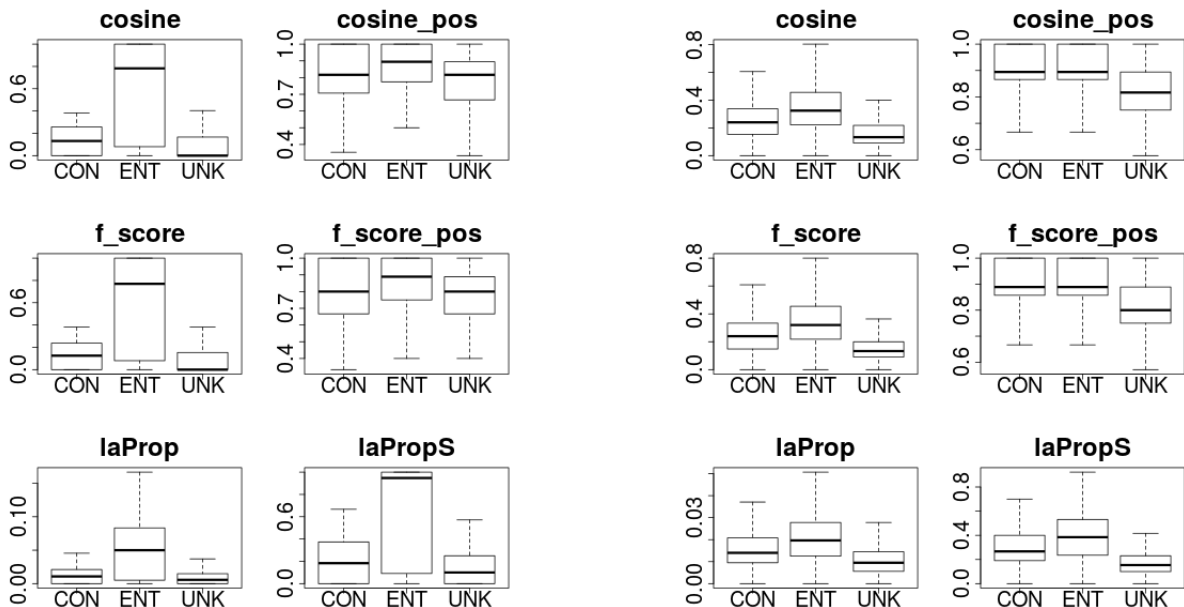


Figure 3: Distributions of the similarity metrics by tweet pair class for the event *ssiege* in the *Threads* (left) and the *iPosts* dataset (right).

5 RTE classification experiments for Contradiction and Disagreeing Reply detection

In order to predict the RTE classes based on the features introduced above, we trained two classifiers: Nearest (shrunken) centroids (NC) (Tibshirani et al., 2003) and Random forest (RF) (Breiman, 2001; Liaw and Wiener, 2002), using the R wrapper package *Caret* (Kuhn, 2016) with the methods *pam* and *rf*, respectively. To derive the same number of instances for all classes, we applied separately for both datasets resampling without replacement, so that the total data amounts about 4,550 feature vectors equally distributed over the three classes, the majority of 4,130 belonging to the iPosts data set. Further, we centered and scaled the feature matrix. Within the Caret framework we optimized the tunable parameters of both classifiers by maximizing the F1 score. This way the NC shrinkage delta was set to 0, which means that the class reference centroids are not modified. For RF the number of variables randomly sampled as candidates at each split was set to 2. The remaining parameters were kept default.

The classifiers were tested on both datasets in a 4-fold event-based held-out setting, training on three events and testing on the remaining one (4-fold cross-validation, CV), quantifying how performance generalizes to new events with unseen claims and unseen targets. The CV scores are summarized in Tables 2 and 3. It turns out generally that classifying CON is more difficult than classifying ENT or UNK. We observe a dependency of the classifier performances on the two contradiction scenarios: for detecting CON, RF achieved higher classification values on Threads, whereas NC performed better on iPosts. General performance across all three classes was better in independent posts than in conversational threads.

Definitions of contradiction, the genre of texts and the features used are dependent on end applications, making performance comparison nontrivial (Lendvai et al., 2016b). On a different subset of the Threads data in terms of events, size of evidence, 4 stance classes and no resampling, (Lukasik et al., 2016) report .40 overall F-score using Gaussian processes, cosine similarity on text vector representation and temporal metadata. Our previous experiments were done using the Excitement Open Platform incorporating syntactico-semantic processing and 4-fold CV. For the non-resampled Threads data we reported .11 F1 on CON via training on iPosts (Lendvai et al., 2016b). On the non-resampled iPosts data we obtained .51 overall F1 score (Lendvai et al., 2016a), F1 on CON being .25 (Lendvai et al., 2016b).

	CON	ENT	UNK
F1 (RF/NC)	0.33/ 0.35	0.55/0.59	0.51/0.57
precision	0.35/0.40	0.54/0.61	0.54/0.57
recall	0.32/0.34	0.58/0.59	0.56/0.67
accuracy		0.47/0.51	
wgt F1		0.48/ 0.51	
wgt prec.		0.51/0.55	
wgt rec.		0.47/0.51	

Table 2: *iPosts* dataset. Mean and weighted (wgt) mean results on held-out data after event held-out cross validation for the Random Forest (RF) and Nearest Centroid (NC) classifiers.

	CON	ENT	UNK
F1 (RF/NC)	0.37 /0.11	0.45/0.50	0.40/0.36
precision	0.42/0.07	0.52/0.56	0.34/0.31
recall	0.35/0.20	0.41/0.47	0.50/0.61
accuracy		0.42/0.39	
wgt F1		0.43 /0.32	
wgt prec.		0.47/0.33	
wgt rec.		0.42/0.39	

Table 3: *Threads* dataset. Mean and weighted (wgt) mean results on held-out data after event held-out cross validation for the Random forest and Nearest Centroid classifiers (RF/NC).

We proposed to model two types of contradictions: in the first both tweets encode the claim target (iPosts), in the second typically only one of them (Threads). The Nearest Centroid algorithm performs poorly on the CON class in Threads where textual overlap is typically small especially for the CON and UNK classes, in part due to the absence of the claim target in replies. However, the Random Forest algorithm’s performance is not affected by this factor. The advantage of RF on the Threads data can be explained by its property of training several weak classifiers on parts of the feature vectors only. By this boosting strategy a usually undesirable combination of relatively long feature vectors but few training observations can be tackled, holding for the Threads data that due to its extreme skewedness (cf. Table 1) shrunk down to only 420 datapoints after our class balancing technique of resampling without replacement. Results indicate the benefit of RF classifiers in such sparse data cases.

The good performance of NC on the much larger amount of data in iPosts is in line with the corpus statistics reported in section 4.3, implying a reasonably small amount of class overlap. The classes are thus relatively well represented by their centroids, which is exploited by the NC classifier. However, as illustrated in Figures 2 and 3, the majority of feature distributions are generally better separated for ENT and UNK, while CON in its mid position shows more overlap to both other classes and is thus overall a less distinct category.

6 Conclusions and Future Work

The detection of contradiction and disagreement in microposts supplies important cues to factuality and veracity assessment, and is a central task in computational journalism. We developed classifiers in a uniform, general inference framework that differentiates two tasks based on contextual proximity of the two posts to be assessed, and if the claim target may or may not be omitted in their content. We utilized simple text similarity metrics that proved to be a good basis for contradiction classification.

Text similarity was measured in terms of vocabulary and token sequence overlap. To derive the latter, local alignment turned out to be a valuable tool: as opposed to standard global alignment (Wagner and Fischer, 1974), it can account for crossing dependencies and thus for varying sequential order of information structure in entailing text pairs, e.g. in ”the cat chased the mouse” and ”the mouse was chased by the cat”, which are differently structured into topic and comment (Halliday, 1967). We expect contradictory content to exhibit similar trends in variation with respect to content unit order – especially in the Threads scenario, where entailment inferred from a reply can become the topic of a subsequent replying tweet. Since local alignment can resolve such word order differences, it is able to preserve text similarity of entailing tweet pairs, which is reflected in the relative *laProp* boxplot heights in Figures 2 and 3.

We have run leave-one-event-out evaluation separately on the independent posts data and on the conversational threads data, which allowed us to compare performances on collections originating from the same genre and platform, but on content where claim targets in the test data are different from the targets in the training data. Our obtained generalization performance over unseen events turns out to be in line with previous reports. Via downsampling, we achieved a balanced performance on both tasks across the three RTE classes; however, in line with previous work, even in this setup the overall performance on contradiction is the lowest, whereas detecting the lack of contradiction can be achieved with much better performance in both contradiction scenarios.

Possible extensions to our approach include incorporating more informed text similarity metrics (Bär et al., 2012), formatting phenomena (Tolosi et al., 2016), and distributed contextual representations (Le and Mikolov, 2014), the utilization of knowledge-intensive resources (Padó et al., 2015), representation of alignment on various content levels (Noh et al., 2015), and formalization of contradiction scenarios in terms of additional layers of perspective (van Son et al., 2016).

7 Acknowledgments

P. Lendvai was supported by the PHEME FP7 project (grant nr. 611233), U. D. Reichel by an Alexander von Humboldt Society grant. We thank anonymous reviewers for their input.

References

- Eneko Agirre, Aitor Gonzalez-Agirre, Inigo Lopez-Gazpio, Montse Maritxalar, German Rigau, and Larraitz Uria. 2016. Semeval-2016 task 2: Interpretable semantic textual similarity. *Proceedings of SemEval*, pages 512–524.
- Isabelle Augenstein, Andreas Vlachos, and Kalina Bontcheva. 2016. USFD: Any-Target Stance Detection on Twitter with Autoencoders. In Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry, editors, *Proceedings of the International Workshop on Semantic Evaluation, SemEval '16*, San Diego, California.
- Daniel Bär, Chris Biemann, Iryna Gurevych, and Torsten Zesch. 2012. UKP: Computing semantic textual similarity by combining multiple content similarity measures. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, pages 435–440. Association for Computational Linguistics.
- Yoav Benjamini and Daniel Yekutieli. 2001. The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 29:1165–1188.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal, September. Association for Computational Linguistics.
- Leo Breiman. 2001. Random forests. *Machine Learning*, 45(1):5–32.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*, pages 177–190. Springer.
- Marie-Catherine De Marneffe, Anna N Rafferty, and Christopher D Manning. 2008. Finding contradictions in text. In *Proc. of ACL*, volume 8, pages 1039–1047.
- Marie-Catherine de Marneffe, Christopher D Manning, and Christopher Potts. 2012. Did it happen? the pragmatic complexity of veridicality assessment. *Computational Linguistics*, 38(2):301–333.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of NAACL*.
- Michael Alexander Kirkwood Halliday. 1967. Notes on transitivity and theme in English, part II. *Journal of Linguistics*, 3(2):199–244.
- Max Kuhn, 2016. *caret: Classification and Regression Training*. R package version 6.0-71.
- Quoc V Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *ICML*, volume 14, pages 1188–1196.
- Piroska Lendvai, Isabelle Augenstein, Kalina Bontcheva, and Thierry Declerck. 2016a. Monolingual social media datasets for detecting contradiction and entailment. In *Proc. of LREC-2016*.
- Piroska Lendvai, Isabelle Augenstein, Dominic Rout, Kalina Bontcheva, and Thierry Declerck. 2016b. Algorithms for Detecting Disputed Information. Deliverable D4.2.2 for FP7-ICT Collaborative Project ICT-2013-611233 PHEME. https://www.pheme.eu/wp-content/uploads/2016/06/D422_final.pdf.
- Andy Liaw and Matthew Wiener. 2002. Classification and regression by randomForest. *R News*, 2(3):18–22.
- Michal Lukasik, P.K. Srijith, Duy Vu, Kalina Bontcheva, Arkaitz Zubiaga, and Trevor Cohn. 2016. Hawkes Processes for Continuous Time Sequence Classification: An Application to Rumour Stance Classification in Twitter. In *Proceedings of ACL-16*.
- Madhumita. 2016. Recognizing textual entailment. Master’s thesis, Saarland University, Saarbrücken, Germany.
- Mitchell P. Marcus, Ann Taylor, Robert MacIntyre, Ann Bies, Constance Cooper, Mark Ferguson, and Alison Littman. 1999. The Penn Treebank Project. <http://www.cis.upenn.edu/~treebank/home.html>. visited on Sep 29th 2016.
- Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. 2010. Twitter Under Crisis: Can We Trust What We RT? In *Proceedings of the First Workshop on Social Media Analytics (SOMA'2010)*, pages 71–79, New York, NY, USA. ACM.

- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting stance in tweets. In *Proceedings of the International Workshop on Semantic Evaluation, SemEval '16*, San Diego, California.
- Roser Morante and Eduardo Blanco. 2012. *SEM 2012 shared task: Resolving the scope and focus of negation. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*.
- Roser Morante and Caroline Sporleder, editors. 2012. *ExProm '12: Proceedings of the ACL-2012 Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*. Association for Computational Linguistics.
- Tae-Gil Noh, Sebastian Padó, Vered Shwartz, Ido Dagan, Vivi Nastase, Kathrin Eichler, Lili Kotlerman, and Meni Adler. 2015. Multi-level alignments as an extensible representation basis for textual entailment algorithms. *Lexical and Computational Semantics (*SEM 2015)*, page 193.
- Sebastian Padó, Tae-Gil Noh, Asher Stern, Rui Wang, and Roberto Zanolli. 2015. Design and Realization of a Modular Architecture for Textual Entailment. *Natural Language Engineering*, 21(02):167–200.
- Martin F. Porter. 1980. An algorithm for suffix stripping. *Program*, 14(3):130–137.
- Vahed Qazvinian, Emily Rosengren, Dragomir R. Radev, and Qiaozhu Mei. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '11*, pages 1589–1599.
- Uwe D. Reichel. 2012. Perma and Balloon: Tools for string alignment and text processing. In *Proc. Interspeech*, page paper no. 346, Portland, Oregon, USA.
- Temple F. Smith and Michael S. Waterman. 1981. Identification of common molecular subsequences. *Journal of Molecular Biology*, 147:195–197.
- Robert Tibshirani, Trevor Hastie, Balasubramanian Narasimhan, and Gilbert Chu. 2003. Class prediction by nearest shrunken centroids, with applications to DNA microarrays. *Statistical Science*, 18(1):104–117.
- Laura Tolosi, Andrey Tagarev, and Georgi Georgiev. 2016. An analysis of event-agnostic features for rumour classification in twitter. In *Proc. of Social Media in the Newsroom Workshop*.
- Chantal van Son, Tommaso Caselli, Antske Fokkens, Isa Maks, Roser Morante, Lora Aroyo, and Piek Vossen. 2016. GRaSP: A Multilayered Annotation Scheme for Perspectives. In *Proceedings of the 10th Edition of the Language Resources and Evaluation Conference (LREC)*.
- Robert A. Wagner and Michael J. Fischer. 1974. The string-to-string correction problem. *Journal of the Association for Computing Machinery*, 21(1):168–173.
- Rui Wang and Günter Neumann. 2007. Recognizing textual entailment using a subsequence kernel method. In *AAAI*, volume 7, pages 937–945.
- Wei Xu, Chris Callison-Burch, and William B Dolan. 2015. SemEval-2015 Task 1: Paraphrase and semantic similarity in Twitter (PIT). *Proceedings of SemEval*.
- Arkaitz Zubiaga, Maria Liakata, Rob Procter, Kalina Bontcheva, and Peter Tolmie. 2015. Towards Detecting Rumours in Social Media. *CoRR*, abs/1504.04712.