

# Beyond Facts - a Survey and Conceptualisation of Claims in Online Discourse Analysis

Katarina Boland<sup>a,b,\*</sup>, Pavlos Fafalios<sup>c</sup>, Andon Tchechmedjiev<sup>d</sup>, Stefan Dietze<sup>a,b,\*\*</sup> and Konstantin Todorov<sup>e</sup>

<sup>a</sup> *WTS, GESIS - Leibniz Institute for the Social Sciences, Unter Sachsenhausen 6-8, 50667 Cologne, Germany*  
*E-mail: katarina.boland@gesis.org*

<sup>b</sup> *Data and Knowledge Engineering, University of Düsseldorf, Universitätsstr. 1, 40225 Düsseldorf, Germany*  
*E-mail: stefan.dietze@hhu.de*

<sup>c</sup> *Institute of Computer Science, FORTH-ICS, N. Plastira 100, 70013 Heraklion, Crete, Greece*  
*E-mail: fafalios@ics.forth.gr*

<sup>d</sup> *EuroMov Digital Health in Motion, University of Montpellier, IMT Mines Alès, 34090 Montpellier, France*  
*E-mail: andon.tchechmedjiev@mines-ales.fr*

<sup>e</sup> *LIRMM, CNRS, University of Montpellier, 161 Rue Ada, 34095 Montpellier, France*  
*E-mail: konstantin.todorov@lirmm.fr*

**Editors:** First Editor, University or Company name, Country; Second Editor, University or Company name, Country

**Solicited reviews:** First Solicited Reviewer, University or Company name, Country; Second Solicited Reviewer, University or Company name, Country

**Open reviews:** First Open Reviewer, University or Company name, Country; Second Open Reviewer, University or Company name, Country

**Abstract.** Analyzing statements of facts and claims in online discourse is subject of a multitude of research areas. Methods from natural language processing and computational linguistics help investigate issues such as the spread of biased narratives and falsehoods on the Web. Related tasks include fact-checking, stance detection and argumentation mining. Knowledge-based approaches, in particular works in knowledge base construction and augmentation, are concerned with mining, verifying and representing factual knowledge. While all these fields are concerned with strongly related notions, such as claims, facts and evidence, terminology and conceptualisations used across and within communities vary heavily, making it hard to assess commonalities and relations of related works and how research in one field may contribute to address problems in another. We survey the state-of-the-art from a range of fields in this interdisciplinary area across a range of research tasks. We assess varying definitions and propose a conceptual model—*Open Claims*—for claims and related notions that takes into consideration their inherent complexity, distinguishing between their meaning, linguistic representation and context. We also introduce an implementation of this model by using established vocabularies and discuss applications across various tasks related to online discourse analysis.

**Keywords:** Claims, Facts, Claim Models, Online Discourse Analysis, Fact-checking

## 1. Introduction

The Web has evolved into an ubiquitous platform where many people have the opportunity to be a publisher, express opinions and interact with others. It has

\*Corresponding author. E-mail: katarina.boland@gesis.org,  
Phone: +49 221 47694-513.

\*\*The last two authors share senior authorship.

been widely explored as a source to mine and understand online discourse or to extract knowledge.

On the one hand, understanding and analyzing societal discourse on the Web are becoming increasingly important issues involving computational methods in natural language processing (NLP) or computational linguistics. Related tasks include fact or claim verification, discourse modeling, stance detection or argumentation mining (AM). In this context, a wide range of interdisciplinary research directions have emerged involving a variety of scientific disciplines including investigations into the spreading pattern of false claims on Twitter [1], pipelines for discovering and finding the stance of claim-relevant Web documents [2–4], approaches for classifying sources of news, such as Web pages, pay-level domains (PLDs), users or posts [5], or research into fake news detection [6] and automatic fact-checking [7]. In addition, understanding discourse in scholarly and scientific works has been a long-standing research problem throughout the past years [8–20].

On the other hand, knowledge-based approaches, in particular works in knowledge base construction and augmentation, often are concerned with mining, verifying and representing factual knowledge from the Web. Research in such areas often deploys methods and conceptualisations strongly related to some of the aforementioned computational methods related to claims, e.g., when aiming to verify facts from the Web for augmenting knowledge bases (KBs) [21, 22]. Whereas the focus in knowledge base augmentation (KBA) is on extracting and formally representing trustworthy factual statements as an atomic assertion in the first-order-logic (FOL) sense, research focused on natural language claims tends to put stronger emphasis on understanding the context of a claim, e.g. its source, timing, location or its role as an argument as part of (online) discourse. Capturing the meaning of a claim requires both to preserve the actual claim utterance as natural language text, often carrying a range of statements and sentiments embedded in complex sentences which are easy to process by humans but hard to interpret by machines, as well as structured knowledge about a claim, its context and constituents, which enables machine-interpretation, discoverability and reuse of claims, for instance, to facilitate research in the aforementioned areas.

Despite these differences, methods in various disparate fields, such as claim/fact verification or fact-checking as well as knowledge base augmentation, tend to be based on similar intuitions and heuristics

and are concerned with similar and related notions form different perspectives. Hence, achieving a shared understanding and terminology has become a crucial challenge.

However, both used terminology as well as the underlying conceptual models are still strongly diverging, both in the academic literature as well as the conceptual models deployed by involved applications, for instance, fact-checking sites [23, 24]. This poses challenges for understanding of related works and data by both humans as well as machines and hinders the cross-fertilisation of research across various distinct yet related fields. Thus, our work aims at facilitating a shared understanding of claims and related terminology across diverse communities as well as the representation of semi-structured knowledge about claims and their context, which is a crucial requirement for advancing, replicating, and validating research in the aforementioned fields.

In order to address the aforementioned problems, this paper proposes the following main contributions, building on and extending our previous work on contextualizing claim-related data [25].

- **An extensive survey** (Sect. 3) of related works concerned with defining, understanding and representing online discourse and related notions, most importantly, claims and facts. The survey is the first of its kind, providing a comprehensive overview of definitions, terminology used across various fields and communities.
- **A conceptual model** (Sect. 4), which we call *Open Claims Model*, and corresponding terminology of claims and their constituents and context, that is grounded in both scientific literature in related fields such as argumentation mining or discourse analysis as well as the actual practices of representing and sharing claims on the Web, for instance, as part of fact-checking sites. To this end, we also provide an RDF/S (Ressource Description Framework Schema) implementation of the proposed model that uses state-of-the-art vocabularies, such as *schema.org* and *PROV-O* (Provenance Ontology), in order to facilitate Web-scale sharing, discovery and reuse of claims and their context, for instance through semi-structured Web page markup or as part of dedicated knowledge graphs (KGs) such as *ClaimsKG* [26].
- **An introductory review of related information extraction and knowledge engineering tasks**

(Sect. 5), involved with the extraction, verification and (inter)linking of claim related data. Our aim is to provide an overview of related state-of-the-art works that may be used for populating a knowledge base of claims and their context according to the proposed conceptual model. This also enables us to discover under-researched areas and challenging directions for future work.

This work is meant to facilitate an unambiguous representation of claims across various communities, as is required for inter-disciplinary research. This includes works aimed at detecting and representing the inherent relations of uttered claims among each other or with represented factual knowledge and other resources, such as Web pages or social media posts, e.g., as part of stance detection tasks. Additional use cases involve research into the detection of viewpoints and communities sharing related narratives and viewpoints on the Web, the analysis of quotation pattern involving varied sources and media types or profiling of sources and references used in news media.

## 2. Methodology of the Survey

In this section, we describe the publication selection and review process employed in this survey. An overview of the workflow is given in Figure 1.

### 2.1. Selection of research fields

First, we identified application areas and research fields involved with claims, facts or relevant concepts.

Application domains include, on the one hand, areas related to natural language claims, which are of concern in fact-checking portals, computational journalism or scientific discourse analysis, for instance, as part of scholarly publications, all involving claims of varying complexity. On the other hand, structured knowledge bases such as Wikidata are used in various applications such as Web search and involve factual statements bound to a predefined grammar relying on triples involving a subject (s), predicate (p) and object (o).

It becomes apparent that a more explicit and clear definition of the concepts of *facts* vs. *claims* is needed with both being relevant to this survey. Works focusing on claims made in the context of a discourse can be found in *argumentation mining*, *argumentation theory*, *discourse modeling*, and *pragmatics*. Facts are central

for *knowledge representation / augmentation* works. With claims not only transporting beliefs or knowledge about factual information but also conveying subjective information such as opinions, stances or viewpoints, relevant definitions and concepts can also be found in works targeting *stance detection*, *viewpoint extraction* and *opinion mining / sentiment analysis*. Rumours can be considered specific kinds of claims, thus we include definitions from the *rumour detection* field. Finally, retrieval of claims or respectively facts about specific entities is central to *question answering* and *information retrieval* in general, for instance, in the context of *fact retrieval*, *entity summarisation* or *entity retrieval*. Relevant works from these fields are also taken into account.

### 2.2. Search and review process

Works addressing the aforementioned fields and tasks can be found in a variety of different scientific communities, particularly Natural Language Processing (NLP), Web Mining, Information Retrieval (IR), Knowledge-based Systems and Artificial Intelligence (AI). We screened relevant venues for the aforementioned communities for works related to claims or facts. Table 1 gives an overview of the chosen core journals, conferences, workshops and events. For each of those, we screened the proceedings of the years 2015 - 2019 (incl. 2020 to the extent possible at the time of writing) widened the search across venues using online search engines and databases. Publications cited by relevant publications were also taken into account. For each publication, we extracted formal and informal definitions and descriptions of the concepts of claims and facts which are the basis for the analysis in Sect. 3 and the development of the model introduced in Sect. 4. As part of the modeling process, we defined possible relations between the different classes and mapped the generation of information on classes and relations to knowledge engineering tasks (Sect. 5). We extended our search in the listed venues and beyond to also cover these tasks. The following set of keywords was used for both steps: *fact-checking*, *fact checking*, *fake news*, *fact verification*, *argumentation*, *discourse*, *pragmatics*, *logic*, *knowledge representation*, *knowledge base augmentation*, *knowledge base construction*, *KBA (Knowledge-Base Augmentation)*, *stance*, *viewpoint*, *claim*, *opinion mining*, *sentiment analysis*, *rumour detection*, *rumor detection*, *question answering*, *information extraction*, *relation extraction*, *ontology learning*. This search procedure

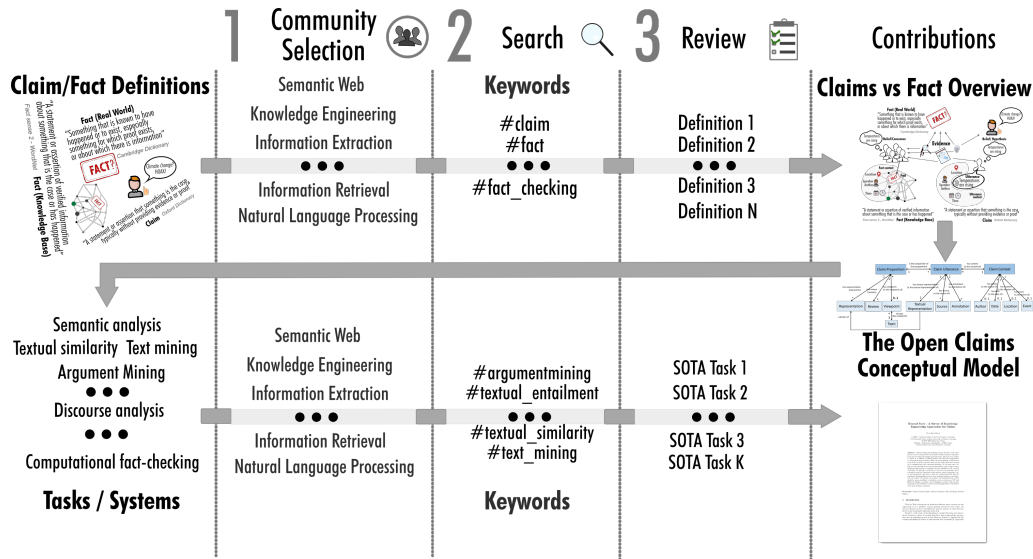


Fig. 1. Publication selection and review workflow

Table 1

Core venues analyzed systematically for the survey of fact and claim definitions and related concepts. Related events and workshops that were also considered: Workshop on Argument Mining, Fake News Challenge (FNC), CLEF Lab: CheckThat!, Fact Extraction and VERification (FEVER) Shared Task.

Community	Journals	Conferences
NLP	Computer Speech and Language	ACL, EMNLP, COLING, NAACL-HLT
Web (Mining)	ACM TWEB	WWW, WSDM
IR	Information Retrieval Journal (Springer)	SIGIR, ECIR
AI		AAAI, IJCAI, ECAI
Knowledge-based Systems & Knowledge Graphs	SWJ, TKDE, JWS, Elsevier KBS	ISWC, ESWC, CIKM

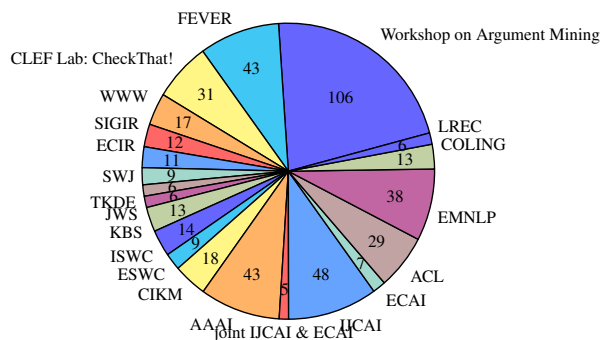


Fig. 2. Analyzed publications and distribution over venues for all venues with at least 5 publications

resulted in a set of 598 publications that we deemed potentially relevant for the topics. Distribution across venues and time periods are displayed in Figures 2 and




Fig. 3. Analyzed publications and distribution over years for all years with at least 5 publications

3. Note that not all of these publications contain relevant definitions or ended up being cited in this survey.

### 2.3. *Related Surveys and Conceptualizations*

While this is, to the best of our knowledge, the first extensive survey on the conceptualization of facts and claims, several works have looked into different aspects of the problem providing overviews of related work in specific areas related to these aspects.

Konstantinovskiy et al. [27] present a novel annotation schema and a benchmark for check-worthy claim detection, providing both an overview of claim definitions from other studies and a new definition of a claim that is constructed as a common denominator of existing ones. The novelty is that the definition is cast in the context of a claim being worthy of fact-checking—an important property of an utterance in view of verifying its veracity. The difficulty of identifying and defining fact-check worthiness of a claim is discussed with regard to the different perspectives that can be given to a single claim according to the human annotator's background.

Daxenberger et al. [23] also take interest in the task of claim identification, but from an argumentation mining (AM) perspective, where this task is defined as recognizing argument components in argumentative discourse. The authors propose a qualitative analysis of claim conceptualization in argumentation mining data sets from six different domains ("different domains" here mean different data distributions). They show that the ways in which claims are conceptualized in each of these data sets are largely diverging and discuss and analyze the presumed harmful impact of these divergences on the task of cross-domain claim identification.

Thorne et al. [28] take a holistic stance on the problem and task of automated fact-checking. They provide an overview of approaches, data sets and methods covering the various steps of the process. This is the first paper of its kind that formulates the ambition to unify the often diverging definitions presented in related works from the fact-checking field by identifying shared concepts, data sets and models. A particularity of the survey is the fact that the authors consider both text-like and structured definitions of claims (e.g. in the form of triples), covering works on knowledge graph building and completion.

Related to fact checking but yet a distinct problem is that of fake news detection. Zhou et al. [29] provide a definition of fake news and present relevant fundamental theories in various disciplines on human cognition and behaviour that are assumed useful for understanding fake news propagation and detection. Their sur-

vey on fake news detection methods is built along four categories of methods: (i) Knowledge-based methods, which verify if the knowledge within the news content matches certified facts; (ii) Style-based methods that look into the form of fake news (e.g., expressing extreme emotions); (iii) Propagation-based methods that are based on online spreading patterns; and (iv) Source-based methods investigating the credibility of sources.

Rumours are often seen as a specific kind of fake news. Zubiaga et al. [30] provide a survey on rumour identification and resolution. Conflicting and diverging definitions of rumours from related works are discussed, but without making parallels to related notions such as fake news or biased discourse. The survey is largely motivated by the assumed impact of social media on rumour generation and spreading. It focuses on datasets for rumour detection, as well as existing tools for accessing, collecting and annotating social media data for the purposes of automated rumour detection. The authors analyse generic rumour detection systems by breaking them down to their different components and subsequently discussing the related approaches to address the challenges related to each of those components. In that, the paper presents rumour tracking systems, rumour stance classification and veracity classification approaches.

Both the lack of and necessity for shared understanding and conceptualization of claims surfaces from all of the above studies, which is underlined as their main motivation. However, the fact that some of these surveys discuss the same notions and refer to the overlapping sets of related work but by using different terminology (like e.g. [30]) comes to show that these works do not fully contribute to closing the terminological and conceptual gap that exists in and across fields.

## 3. *Facts and Claims - A Multidisciplinary Survey of Definitions*

While the analysis of facts and claims plays a crucial role for a number of fields, the definitions of these concepts vary and are often left to the intuition of the reader. Existing definitions vary considerably not only across different fields but also within a single community. At the same time, different communities use the same terminology to refer to different concepts. In this section, we expatiate on different concepts for facts and claims, explain commonalities and differences and

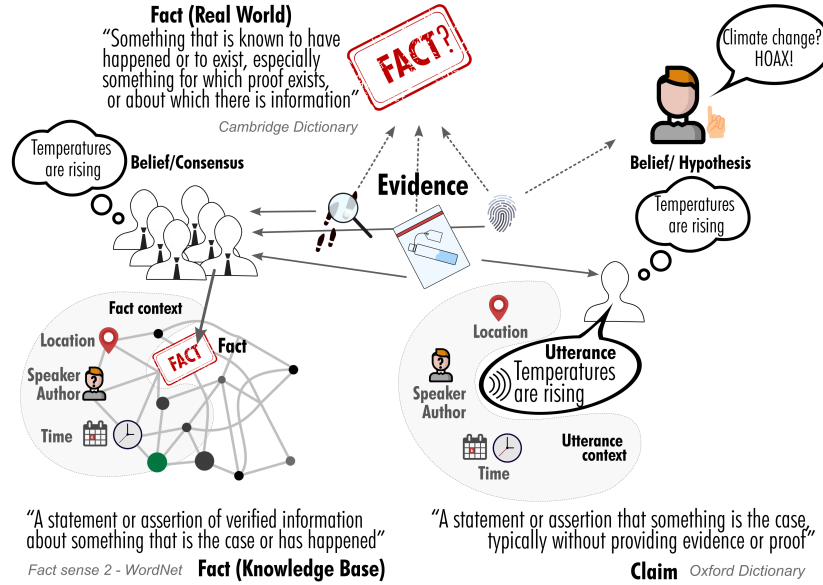


Fig. 4. An overview of definitions and relations between facts and claims

introduce a selected vocabulary to refer to these and related concepts throughout this paper. An overview is given in Figure 4.

### 3.1. Facts

A fact in everyday use of the term (depicted on the top of Figure 4) refers to “A thing that is known or proved to be true”<sup>1</sup>, “something that has actual existence”<sup>2</sup>, “something that is known to have happened or to exist, especially something for which proof exists, or about which there is information”<sup>3</sup>, “something that actually exists; reality; truth”<sup>4</sup>, “an event known to have happened or something known to have existed”<sup>5</sup> or “a concept whose truth can be proved”<sup>6</sup>. Note that facts themselves cannot be observed directly; instead, beliefs about them can be formed by observing *evidence*.

<sup>1</sup>Oxford Dictionary; <https://www.lexico.com/en/definition/fact>

<sup>2</sup>Merriam-Webster Dictionary; <https://www.merriam-webster.com/dictionary/fact>

<sup>3</sup>Cambridge Dictionary; <https://dictionary.cambridge.org/dictionary/english/fact>

<sup>4</sup>Dictionary.com; <https://www.dictionary.com/browse/fact>

<sup>5</sup>WordNet; <http://wordnetweb.princeton.edu/perl/webwn?s=FACT;S3>

<sup>6</sup>WordNet; <http://wordnetweb.princeton.edu/perl/webwn?s=FACT;S4>

#### 3.1.1. Facts in knowledge bases

In the semantic web community and the fields of knowledge representation, knowledge base construction / augmentation, facts are the constituents of KGs or knowledge bases (KBs) [31–36, 36–49]. More precisely, items in KGs or KBs are coined statements of facts or assertions or triples encoding or representing facts [33, 39, 45–47], with the facts being assumed to be true, can be proven to be true or are likely to hold [38, 45, 50]. However, the use of terminology is not consistent: *fact* is often used as synonym for RDF triple [37, 38, 49, 51, 52] or for the representation of a fact, respectively assertion, i.e. there is no distinction made between a fact and the statement of a fact [36, 43, 47, 48]. Having the task of fact-prediction as background, some works coin the relations between entities or the paths in a knowledge base *facts* [34, 38]. The interchangeable use of *statement of fact* and *fact* leads to a widespread terminology of “checking whether facts are true” [53] which is an oxymoron given the definition of a fact being something that is known to be true. As [49] note, facts have a scope, e.g. temporal, that determines the context that has to be taken into account in order to judge their validity.

#### 3.1.2. Types of facts

Several more fine-grained distinctions of different types of facts can be found in the literature. Facts can refer to relations or attributes [52], or can be attributes

of other facts [34]. They can pertain to numerical properties, quotes or other object properties [51]. They can be assessed according to their “*check-worthiness*” [51] or *importance* for the containing KB [34]. Another interesting distinction is made by Tsurel et al. [54] who aim at identifying facts that are suitable to be used as interesting trivia by developing a measure for *trivia-worthiness* that relies on surprise and cohesiveness of the contained information.

### 3.1.3. Facts vs. evidence

Related to the notion of *fact* is the notion of *evidence*. [40, p. 1] define *evidence* “as text, e.g. web-pages and documents, that can be used to prove if news content is or is not true”. More generally, evidence is seen as something to support or contradict a claim [55–57]. As Stahlhu et al. [55] note, the task of evidence detection is similar to premise detection in Argumentation Mining. A premise in Argumentation Mining is, as Stab et al. [58, p. 1] put it, “a reason given by an author for persuading the readers of the claim”. *Evidence* and *premise* directly correspond to each other, both terms are often used interchangeably [59–61].

*Evidence* can be categorized into many different types, such as *expert opinion*, *anecdote*, or *study data* [55], or, with slightly different wording, *study*, *expert* or *anecdotal* [56]. Walker et al. [62] distinguish *lay testimony*, *medical records*, *performance evaluations*, *other service records*, *other expert opinions*, *other records*. Niculae et al. [63] include *references* such as URLs or citations as pointers to evidence. Premises can refer to *logos*, *pathos* or *ethos* [64]. For scientific articles, Mayer et al. [65] distinguish the classes *comparative*, *significance*, *side-effect*, *other*.

While some works refer to knowledge found in texts or other resources as *evidence* [43, 44, 48, 66] for a fact and call it *fact* only after the truthfulness has been determined and that knowledge is entered into a knowledge base, other works assume the truthfulness of the mentions and refer to them or the knowledge they represent as *facts* directly [67, 68]. Very related is the task of *Truth Discovery* (TD). “TD aims to identify facts (true claims) when conflicting claims are made by several sources” [69]. In this domain, the terms *data items* and *truths* are used to refer to invalidated mentions of knowledge and the true values respectively [70, 71].

## 3.2. Claims

A claim is “a statement or assertion that something is the case, typically without providing evidence or proof”<sup>7</sup>.

### 3.2.1. Claims in argumentation

In line with this definition, works in Argumentation Mining and Argumentation Theory focus on claims as the key components of arguments [23], as statements that are made to convince others or express someone’s views, evaluations or interpretations [64, 72–74]. Claims denominate the conclusion of an argument, the assertion the argument aims to prove or the thesis to be justified [58, 61, 73, 75–78]. Claims correspond to propositions in argumentation models and both terms are often used interchangeably, “The claim is a proposition, an idea which is either true or false, put forward by somebody as true” [76]. As Daxenberger et al. [23] point out, the exact definition of a claim even inside the field of Argumentation Mining depends on the domain or task and is somewhat arbitrary. Also, as Torsi et al. [24] show, related annotation categories are often not well defined.

With the use-case of scientific articles in mind, Mayer et al. [79] define a claim as a concluding statement made by the author about the outcome of the study. Focusing on debates, [56, 2], also used by Rinott et al. [57, 1], define a claim as “general, concise statement that directly supports or contests the topic”. A topic here is defined as “a short, usually controversial statement that defines the subject of interest” or “a short phrase that frames the discussion” respectively. Examples for such topics are “Use of performance enhancing drugs (PEDs) in professional sports” with a claim being “PEDs can be harmful to athletes health” [57, p. 2] or “The sale of violent video games to minors should be banned” with a claim being “Violent video games can increase children’s aggression” [56, p. 3]. Note that these definitions diverge from the common definition of a *topic* as the *underlying semantic theme* of a document with a topic being a *probability distribution over terms in a vocabulary* [80] as used in topic modelling and document classification. There, a *topic* may be represented by terms on a coarse-grained level such as *Health* or *Computers & Internet* [81]. Also, the second example of a topic can be seen as a claim or stance itself. Durmus et al. [82] represent topics by tags of pre-defined categories similar to the above de-

<sup>7</sup><https://www.lexico.com/en/definition/claim>

scribed semantic themes plus what they call a *thesis*, corresponding to Aharoni et al. [56]’s claim-like topics, e.g. “free Press is necessary to democracy.”, “All drugs should be legalised.”.

*Topic* will in the following be used to refer to the frame of the discussion, as defined by Rinott et al. [57] while the *underlying semantic theme* will be referred to as the *subject*.

### 3.2.2. Types of claims

According to Lippi et al. [73], there are three different types of claims: 1) *epistemic*, i.e., claims about knowledge or beliefs, 2) *practical*, i.e., claims about actions, alternatives and consequences, and 3) *moral*, i.e., claims about values or preferences.

Similarly, Schiappa et al. [83, 84] differentiate *claims of fact, value and policy*. Claims of fact state that something is true, i.e. they express a belief about a fact. This corresponds to the epistemic claims according to Lippi et al. [73]’s taxonomy with claims of value and policy corresponding to moral and practical claims, respectively. Epistemic claims are also referred to as *factoid claims* [85, 86] or, more commonly, *factual claims*, e.g. [7, 27, 77, 87–91]. However, assessing the *factuality* of a claim may refer to assessing a claim’s veracity [92] rather than assessing whether it is a factual or non-factual claim.

Some works propose a more fine-grained differentiation of claims according to their use-cases, e.g. [8] distinguish *Own Claims* vs. *Background Claims* vs. *references to Data* for argumentation mining of scientific texts. Hassan et al. [93] distinguish between the classes *Non-statistical* (e.g. quotes), *Statistical*, *Media* (e.g., photo or video), and *Other*, [94] between categorical vs. numerical claims. Park et al. [95] categorize claims according to their verifiability and distinguish between *unverifiable*, *verifiable nonexperiential*, *verifiable experiential* claims with *experiential* referring to whether the claim refers to the writer’s personal state or experience or not. Closely related to this is the notion of a *rumour*. In an attempt to unite various definitions found in works addressing the identification and veracity assessment of rumours, [30, p. 1] define rumours as “items of information that are unverified at the time of posting”. The authors further distinguish between different types of rumours, with respect to their currentness (emerging vs. longstanding rumours).

### 3.2.3. Claims vs. stances vs. viewpoints

Habernal et al. [96] explain that the term claim in the context of argumentation theory is a synonym for

standpoint or point of view referring to what is being argued about, i.e. the topic. This is in line with [97] and [56]’s debate-oriented definition and [64, p. 4] defining claims as “proposition that expresses the speaker’s stance on a certain matter”. *Standpoint, point of view and stance* in these definitions do not mean the content of the claim has to be of unverifiable or purely opinionated nature. [98] see *stance* as an attribute of a claim.

Stances are usually defined as text fragments representing opinions, perspectives, points of views or attitudes with respect to a target [4, 90, 99–101]. They can be expressed explicitly or implicitly [102]. Fragments can be messages such as tweets or posts [103, 104], paragraphs [105] or complete articles [4]. Joseph et al. [103] see stances as latent properties of users rather than text fragments. Text fragments can however reveal a user’s stance. As Joseph et al. [103] point out, stance and sentiment are related, but not the same: a negative sentiment of a text can be paired with a positive stance towards a particular target and vice versa. Also the tasks of aspect-based sentiment analysis and stance detection differ, even though both aim at detecting opinions towards a target. For example, a piece of text may express a positive sentiment towards a specific aspect of a person, e.g. their personality, but still argue against this person’s claim.

Stance detection has been used to determine opinions on the veracity of claims [106, 107]. Stances in these works are similar to what is coined *evidence* in fact-checking works, as described above, although they do not necessarily contain factual information that can be used to verify information. Note that this may be the case for *evidence* as well, depending on the precise definition. The fact that a claim is supported by another entity than the source can be seen as evidence for the claim’s truthfulness in itself (cf. *expert-type evidence*).

Stances have been classified into different categories such as *for, against and observing* [106], *pro and con* [108] and *none* [109], *agree, disagree, discuss*, or *unrelated* [110] and, in a hierarchical model, *related: agree, related: disagree, related: discuss, and unrelated* [40]. A further more fine-grained distinction can be found in [64] who distinguish *interpretations, rational evaluations, emotional evaluations, agreement and disagreement*. As Kotonya et al. [101] note, the task of stance classification is closely related to relation-based argumentation mining determining attack and support relations between argumentative units. Another related task is that of viewpoint discovery. Thonet et al. [111] define a viewpoint as “the standpoint of one or sev-



eral authors on a set of topics”. A viewpoint goes beyond a person’s stance on a specific subject and represents their global standpoint or side they are taking. As Thonet et al. [111] explain, a viewpoint in a debate about the building of Israeli communities on disputed lands can for example be summarized as “pro-Palestine” or “pro-Israel”. Consequently, Viewpoint Discovery is considered a sub-task of Opinion Mining [111, 112].

#### 3.2.4. Claims in journalism and fact-checking

Works outside of the area of argumentation focus less on the role of the claim in the context of the discourse and more on the content of the claims.

A very general definition is given by [94, p. 2] for their truth discovery approach: “A claim is defined as a piece of information provided by a source towards an entity”.

From a journalistic fact-checking perspective, dedicated platforms focus on statements supported by (a group of) people or organizations that appear news-worthy, check-worthy, significant and verifiable (cf. definitions from, e.g., politifact.com,<sup>8</sup> truthorfiction.com,<sup>9</sup> or checkyourfact.com<sup>10</sup>). Newsworthiness and significance are not only subjective, both can also vary depending on historical or political context [113].

For other use-cases, different definitions or restrictions of what is considered a claim are employed.

Automatic fact-checking often constrains the problem by limiting the kinds of claims being checked to focusing on simple declarative statements (*short factoid sentences* [114]) or claims about statistical properties [113, 115, 116]. For the Fast & Furious Fact-Check Challenge, four primary types of claims were distinguished and further differentiated into more fine-grained sub-categories:<sup>11</sup> 1) *numerical claims* (involving numerical properties about entities and comparisons among them), 2) *entity and event properties* (such as professional qualifications and event participants), 3) *position statements* (such as whether a political entity supported a certain policy) and 4) *quote verification* (assessing whether the claim states precisely the author of a quote, its content, and the event at which it supposedly occurred). Note that fact-checking portals contain many quoted claims but it is not always clearly marked whether the quote itself is verified (i.e.

did the person indeed make the claim?) or the content of the quoted claim (i.e. is the claim allegedly made by the person correct?).

#### 3.2.5. Claims in databases

From a database-centric perspective, [117, 118] represent a claim as a “parametrized query over a database”. This allows to computationally study the impact of modifying a claim (i.e. its parameters) on the result of the query and to thus identify claim properties, such as claim robustness which may serve as evidence to detect potential misleadingness i.e. due to cherry-picking. A related perspective has been proposed by [119] in the field of Computational Journalism.

#### 3.2.6. Claims in information retrieval and question answering

Similar conceptualizations can be found in the area of Information Retrieval and Question Answering for finding scientific claims and claims in digital libraries. Here, a claim is defined as a statement formulating a problem together with a concrete solution [120] or a sentence in a scientific document that relates two entities given in a query [10, 14].

### 3.3. Summary

In summary, works focusing on the argumentation domain investigate claims in the context of a discourse, i.e. taking their pragmatic role into account. Claims are uttered by the author or speaker to achieve an aim through a speech act [121]. In order to recognize the meaning of an utterance and draw conclusions about the intention of the author, the pragmatic context has to be taken into account. A claim often carries a variety of intentional or unintended meanings, where subtle changes in the wording or context can have significant effects on its validity [113]. Works in other areas, such as Knowledge Bases and Fact-checking, typically focus on the content of epistemic claims, i.e. rather than trying to analyze intended meanings or messages, they try to find and check evidence for assertions and find facts vs. false claims of fact. Works in the area of information retrieval focus more on the surface of claims, trying to retrieve relevant texts without necessarily analyzing their content or contexts. These differences are reflected in the claim definitions found in the respective works.

Note that due to these different foci, there is a difference in what is referred to as *claim* in argumentation mining vs. in the automatic fact-checking com-

<sup>8</sup><https://www.politifact.com/truth-o-meter/article/2018/feb/12/principles-truth-o-meter-politifacts-methodology-i/>

<sup>9</sup><https://www.truthorfiction.com/about/>

<sup>10</sup><https://checkyourfact.com/about-us>

<sup>11</sup><https://www.herox.com/factcheck/5-practise-claims>

munity: what is used as premise or evidence in an argument is often selected as check-worthy *claim* by fact-checking sites, not the evaluative component of the argument that is coined *claim* in Argumentation Mining. Generally, the distinction of argumentative units such as claims and evidence in argumentation mining is based on the statements' usage or its relations in an argument while fact-checking classifies statements into claims, stances and other categories considering features inherent to the statement itself (such as their subjectivity) regardless of their connection to the discourse. Thus, what is identified as *claim* in works of one research field or labelled *claim* in a ground truth corpus may or may not be called *claim* in the other, depending on the specific use-case and context. Likewise, some works focus on identifying claims that belong to a pre-defined topic (called *corpus wide topic-dependent* [59], context-dependent [56] or the *information-seeking* perspective [60]), while others aim at extracting any units that act as claims for any topic (*closed-domain discourse-level* [60] or context/topic-independent). Using topic-dependent annotations as ground truth for topic-independent extraction approaches leads to impaired precision values [122].

Lastly, another difference between statements of fact in knowledge bases and claims is that for the former, a certain level of consensus at least in the respective community can be assumed, while claims may only represent the beliefs of one person or be uttered by them to achieve a certain goal such as spreading disinformation. Thus, it makes sense to model truth values for claims while statements in knowledge bases are assumed to be true. There may be errors in knowledge bases, however. Thus, modeling uncertainty or confidence values is applicable for them.

The task of assessing the correctness of a statement of fact is called *fact validation*. The task of assessing the veracity of a claim is called *fact-checking*. Fact-checking has also been modeled as a specific *stance detection* task where the stance of a source or evidence unit towards an epistemic claim is used to assess the claim's veracity. Finding the true values in case of conflicting evidence is the aim of *truth discovery*.

### 3.4. Naming conventions

To arrive at a more precise usage of terminology, we will, throughout this paper, refer to items in knowledge bases as *statements of facts* and other mentions or assertions of knowledge as *claims about a fact* that can

act as *evidence* about some information being true and its content being a fact. An index of all naming conventions followed in this work is given in table 2.

## 4. Conceptual Modeling

In this section, we propose a conceptual model for representing claims and related data as well as an example of an implementation of this model in RDF using established vocabularies. We start by giving an overview of the key terminology.

### 4.1. Key Terminology - From Pragmatics to Fact-Checking

For our conceptual model, we follow notions from pragmatics to allow modeling not only a claim in isolation, but also its meaning in a given discourse and its role in communication.

As Green (1996) puts it, "(...) communication is not accomplished by the exchange of symbolic expressions. Communication is, rather, the successful interpretation by an addressee of a speaker's intent in performing a linguistic act." [123, p. 1] "Minimally the context required for the interpretation (...) includes the time, place, speaker, and topic of the utterance." [123, p. 2] While this quote refers to the interpretation of indexical expressions (i.e. words like "here" and "now"), the same holds true for the interpretation of the meaning of an utterance in general.

A linguistic act, or speech act following [121], includes an utterance, a proposition, an illocution and a perlocution. An *utterance* is a grammatically and syntactically meaningful statement. A *proposition* is the semantic content, i.e. meaning. An *illocution* is the intended effect, e.g. persuading the addressee or requesting a service, while a *perlocution* is the achieved effect.

For example, referring to the topic "Brexit", i) British journalist David Dimbleby said during a topical debate in Dover "*We are going to be paying until 2064, apparently*"<sup>12</sup>, and ii) a news article of The Independent on the same topic wrote ("*UK will be paying Brexit 'divorce bill' until 2064*"<sup>13</sup>). While the surface forms of these utterances differ, they express the

<sup>12</sup><https://fullfact.org/europe/brexit-divorce-bill-2064/>

<sup>13</sup><https://www.independent.co.uk/news/uk/politics/uk-brexit-divorce-bill-taxpayers-deadline-treasury-obr-office-budget-responsibility-a82>

html

Table 2  
Index of main notions and definitions as discussed in this paper

Term	Definition
Claim	A statement or assertion that something is the case
Fact	A thing that is known or proved to be true
Statement of fact	A statement in a knowledge base
Evidence	Information that can be used to assess the truthfulness of a claim or claimed value or relation
Topic	Phrase describing the frame of the discussion
Subject	Keyword describing the semantic theme
Stance	Opinion or perspective expressed by a user or text fragment with respect to a given target
Viewpoint	Standpoint of one or several authors on a set of topics

same proposition. At the same time, utterances with equivalent surface forms may be used to express different and even contradicting propositions or viewpoints when embedded in different contexts. Consider the two claims: (i) “*The unemployment rate among Poles in Britain is lower than the unemployment rate among Brits*”, uttered by British public policy analyst and former Labour Party politician David Miliband<sup>14</sup> to argue that immigrants are not a drain on the British welfare system and thus not bad for the British society; (ii) “*EU migrants are MORE likely to have a job in the UK than British citizens*”, written by MailOnline journalist Matt Dathan<sup>15</sup> to make a point that immigrants are taking away British citizens’ jobs and thus are bad for society. The propositions are semantically similar and both utterances aim at persuading the audience (illocutionary act) but the expressed viewpoints are different.

This can only be recognized when taking the context into account which is why we argue that the context should be modeled along with the claim utterance. The importance of contextual information has also been recognized for the task of fact-checking: “Who makes a claim, when they say it, where they say it, and who they say it to, can all affect the conclusion a fact-checker could reach. Whether it’s true to say unemployment is what country or which part of a country a speaker is referring to, and when the speaker makes the claim. An open format for recording public debate should support metadata, including at least the time,

the place, the venue or publication, and the speaker.” [124].

As outlined in the previous section, we see a *fact* as a conceptual object which represents the current consensual knowledge in a given community about something or someone. While this knowledge is relatively stable, a change of its truth value is possible for example when flaws in scientific studies are discovered and findings have to be corrected [125].

Any verified information about a claim, like who uttered it, when and where, can be considered a *fact*. In addition, explicitly uttered facts and factual claims can be modeled as utterances. In the case of a fact extracted from a knowledge base, the speaker equals the knowledge base reporting the fact. In contrast to a claim, it is not necessarily embedded in a discourse. Thus, certain attributes of the context, like the topic of the discourse, remain unknown. Likewise, non-factual claims do not have truth values. Apart from that, we argue that facts, factual claims and non-factual claims can be represented using the same model.

#### 4.2. The Open Claims Conceptual Model

In line with the rationale outlined above, we introduce the *Open Claims* conceptual model, which distinguishes three main components of a claim represented by three central classes: (1) *claim proposition*, (2) *claim utterance*, and (3) *claim context* (Fig. 5). A *claim proposition* is the meaning of a statement or assertion. In the context of fact-checking and argumentation mining, it is usually related to a controversial topic and is supported by one person or a group of people. A *claim proposition* can be expressed in many different ways and in different contexts, thus it has one or more *claim utterances*. For example, it might be expressed in different languages, using different words

<sup>14</sup><https://www.politifact.com/factchecks/2016/jun/26/david-miliband/david-miliband-link-between-jobs-immigration-and-b/>

<sup>15</sup><https://www.dailymail.co.uk/news/article-3628840/The-true-cost-open-borders-revealed-EU-migrants-likely-job-UK-British-citizens.html>

in the same language, or uttered by different persons and/or in different points in time.

In contrast, a specific claim utterance is typically associated with only one claim proposition, i.e., it has a single meaning. However, the claim proposition can be represented in different ways, for example, by selecting a representative utterance with its context, or through a more formal model. Each claim utterance is related to a specific *claim context*, which includes the person who uttered the claim, the time point at which the claim was uttered, the location or the event of the utterance and the topic of the enclosing discourse. The claim context provides information to interpret the claim utterance and thus understand its proposition.

Information on the perlocution (achieved effect) of utterances is usually unavailable. So is explicit information about illocutions (intended effect). Since both are not central to the use-cases motivating this survey, we do not consider them in this model.

Below, we provide details and the main properties of each of the three main classes (Claim Proposition, Claim Utterance, Claim Context).

#### 4.2.1. Claim Proposition.

A claim proposition is the meaning of one or more claim utterances in their respective contexts. A claim proposition is associated with i) one or more *representations*, ii) one or more *reviews*, and iii) one or more *attitudes*. A *representation* can have the form of free text, e.g., a sentence that best describes the proposition, or be more formal, e.g., a first-order logic model. A *review* is a resource (e.g., a document) that analyzes one or more check-worthy claim propositions and provides a verdict about their veracity or trustworthiness. An example of such a review is an article published by a fact-checking organization. Note here that not all factual claims have a clear verdict. For instance, the claim “the presence of a gun makes a conflict more likely to become violent” represents a hypothesis which can be linked to both supporting and contradicting evidence and is thus difficult to be associated with a single overall correctness score. If a claim is associated to a review which gives a true verdict about its veracity, then the claim can be considered a fact (it represents current knowledge about something). Non-factual claims are not linked to any reviews and have no verdicts. An *attitude* is an opinion on a given topic (e.g., a viewpoint), which often underlies a set of specific values, beliefs or principles. For instance, *pro-Brexit* and *anti-Brexit* are two different viewpoints for the Brexit topic. A claim proposition can be associated with several attitudes for

different topics. For example, the proposition linked to the claim “immigrants are taking our jobs” can support both the *against immigration* (for the Immigration topic) and the *pro-Brexit* attitude (for the Brexit topic).

#### 4.2.2. Claim Utterance.

A claim utterance is the expression of a claim in a specific natural language and form, like text or speech. Among other things, it can be something said by a politician during an interview, a text within a news article written by a journalist, or a tweet posted by a celebrity about a controversial topic. It is associated with i) one or more *linguistic representations* (subclass of *representation*), and ii) one or more *sources*. A *linguistic representation* can be, for example, a text in a specific language that best imprints the claim as it was said/appeared, or a sound excerpt from someone’s speech. A *source* provides evidence of the claim’s existence. For instance, it can be the URL (Uniform Resource Locator) of an interview video, a news article, or a tweet, i.e. source here means the medium reporting the utterance, not the originating speaker or author which is part of the context. For this distinction, see also [126]. A linguistic representation can have one or more *annotations* which provide formal linguistic characteristics. For instance, it can be an entity or date mentioned in the text of the claim utterance, the sentiment of the text (e.g., positive, negative, neutral), or the linguistic tone of a speech (like irony). The annotation can enable advanced exploration of the claims (e.g., based on mentioned entities) and can be manually provided by a domain expert or automatically produced using a NLP or speech processing tool (like an *entity linking* [127] tool for the case of entity annotation in text).

#### 4.2.3. Claim Context.

The *claim context* provides background information about the claim utterance. It is associated with meta-data information about the claim utterance and, together with the linguistic representation of the claim utterance, can provide an answer to the *Five W’s*: i) *what* was said (linguistic representation of claim utterance), ii) *who* said it (author of the claim), iii) *when* it was said (date the claim was uttered), iv) *where* it was said (location where claim was uttered), and v) *why* it was said (event or activity in the context of which the claim was uttered). The claim context provides the necessary information for interpreting the claim utterance (and thus understanding its proposition), and can be extended with more concepts that allow describing

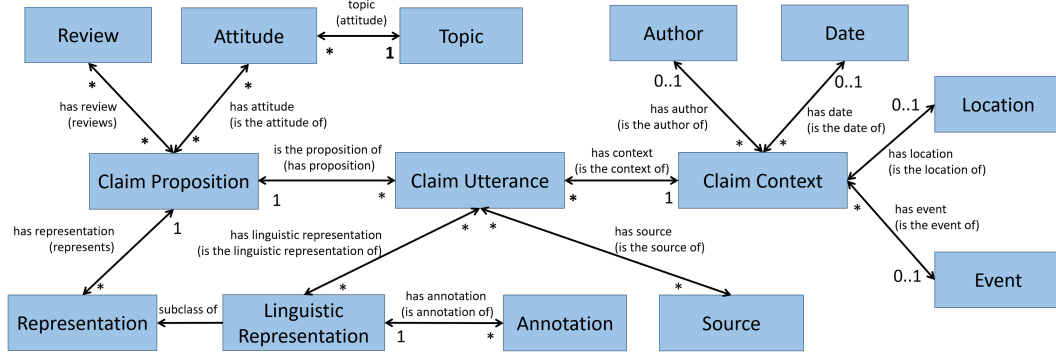


Fig. 5. The Open Claims Conceptual Model.

additional context information (like the topic of the underlying discourse).

#### 4.2.4. Links between the Components of Different Utterances

The role of a claim in a discourse and the relations between different claims can be explicated by creating links between the components. For example, claim utterances can be connected by relations such as *used as evidence for* or *used as evidence against* to model premises, evidence, conclusions and other components and relations in argumentation. Likewise, *supports* and *attacks* relations may hold between utterances to connect stances and their targets, or *conflicting* and *concordant* claims can be modeled by creating typed links between propositions. With this, we follow [128] and the discussion in section 3 with the notion that whether a statement is of type *evidence* or another type depends on its usage in the context of a discourse, e.g. its relations, rather than being an inherent property of the statement in isolation.

#### 4.2.5. Instantiation Example.

Fig. 6 depicts an instantiation example of the proposed conceptual model. The example shows information for two *claim utterances* (in pink background in Fig. 6): i) the one by David Dimbleby (“We are going to be paying until 2064, apparently”), and ii) the one by The Independent (“UK will be paying Brexit ‘divorce bill’ until 2064”). Both utterances correspond to the same claim proposition (in green background in Fig. 6) and each one has its own context information (in yellow background in Fig. 6). The linguistic representation of the first claim utterance has been annotated with one *date annotation* (2064) and that of the second claim utterance with one *entity annotation* (United Kingdom).

The claim proposition has two representations, a textual one (“*Britain will be paying its Brexit bill for 45 years after leaving the EU*”) and a formal one (“*cost = {of= ‘Brexit’, for= ‘UK’ amount= ?, until=2064}*”), and supports the *against-Brexit* viewpoint of the *Brexit* topic. In addition, there is a review of this claim proposition with verdict “true”, published by Full Fact (UK’s independent fact-checking organisation). Moreover, we can see the URL of the review article as well as a reference to a document file which provides evidence for its correctness.

The context of each claim utterance provides additional metadata information about the claim. For example, we see that the first utterance was said by *David Dimbleby* on *15.03.2018*, in the context of a *debate* about *Brexit* which took place in *Dover*. For the second claim utterance, the example only represents its author (*UK Office of Budget Responsibility*) and date (*13.03.2018*).

#### 4.3. RDF Implementation

In order to facilitate the use and operationalisation of our *Open Claims* Conceptual Model, we provide an RDF/S implementation using established vocabularies, depicted in Fig. 7. Vocabulary selection followed three directives: i) relying on stable term identifiers and persistent hosting, ii) being supported by a community, iii) being extensible.

As our base schema, we propose to use *schema.org*.<sup>16</sup> For capturing provenance information of all generated annotations, we employ *PROV*.<sup>17</sup> This includes information about employed tools and confidence val-

<sup>16</sup><https://schema.org>

<sup>17</sup><https://www.w3.org/TR/prov-overview/>

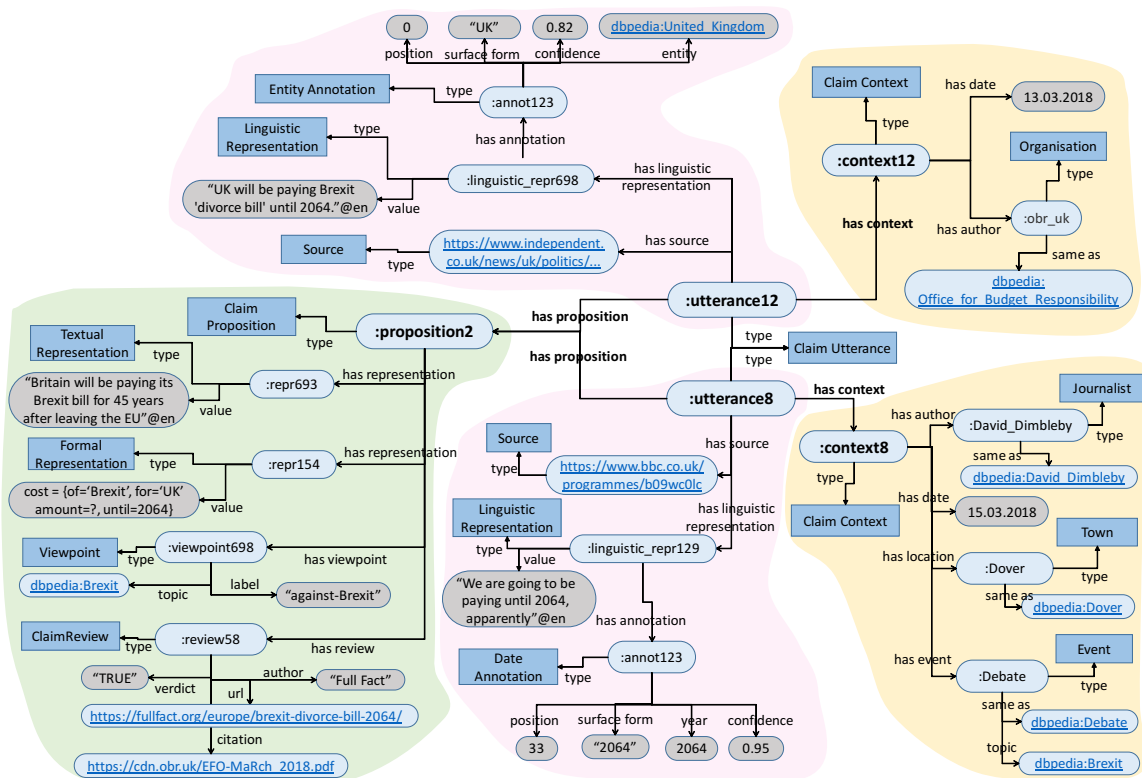


Fig. 6. Instantiation example of the conceptual model.

ues. Claim review verdicts are part of the schema.org *ClaimReview* entity. Viewpoints are represented using MarL Ontology<sup>18</sup>, which is designed to annotate and describe subjective opinions. Linguistic annotations are represented in the NLP Interchange Format (NIF)<sup>19</sup>. To cover other modalities such as video and images, we include the Web Annotation Vocabulary (OA).<sup>20</sup>

More details about the RDF/S implementation of the proposed conceptual model can be found in our previous work [25].

## 5. Related Knowledge Engineering Tasks

In this section, we review different knowledge engineering and information extraction tasks that are involved with claim related data, like utterances, claim verification scores, claim context information (e.g. who uttered the claim, when and where) and other

claim metadata that is described in our *Open Claims* model. We identify three main (sometimes overlapping) categories of tasks: *extraction*, *verification*, *interlinking* and position them within the context of our conceptual model. Note that we do not aim to provide an exhaustive overview of those tasks, but rather introduce examples of works of different relevant areas and show how they are positioned with respect to extracting or generating the information and relations suggested by our model. *Extraction* pertains to detecting statements, utterances and other components and attributes in a corpus of (mainly) textual modality. *Verification* pertains to the assignment of truth ratings or credibility scores to claims or other related components such as information sources. *Interlinking*, finally, includes a range of different tasks aiming at detecting various relations between claims or related components thereof, such as same-as relations, stances or topic-relatedness.

### 5.1. Extraction

Given the complexity and varying definitions of what is or what constitutes a claim, a number of differ-

<sup>18</sup><http://www.gsi.dit.upm.es/ontologies/marl>

<sup>19</sup><https://persistence.uni-leipzig.org/nlp2rdf/index.html>

<sup>20</sup><https://www.w3.org/TR/annotation-vocab/>

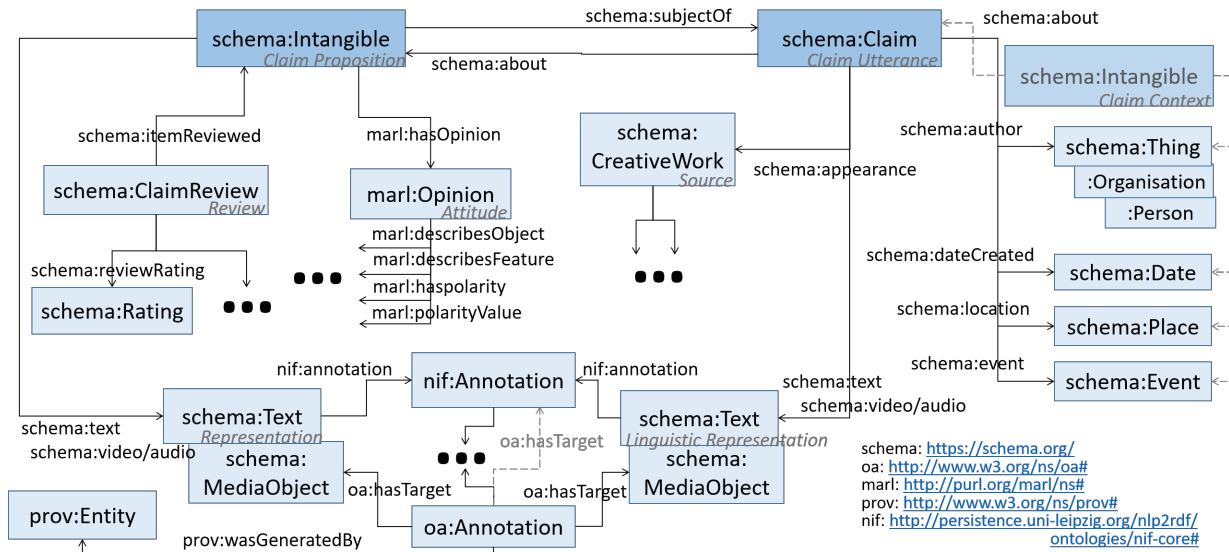


Fig. 7. RDF implementation of the Open Claims Conceptual Model

ent knowledge extraction approaches can be associated to the tasks in each of the three groups outlined above. We will follow the definition of a claim and its components as given by our model (Sect. 4) in order to review the existing techniques for knowledge extraction pertaining to each of these components and attributes. In parallel, we identify challenging problems that are underrepresented in the literature.

#### 5.1.1. Extracting claim propositions

The task of extracting a claim proposition can be reformulated as assigning an identifier to a group of statements that are assumed to be semantically equivalent. Our model suggests that the meaning of a claim can be captured both by the means of natural language as well as formal knowledge representation framework, e.g. description logics.

Extracting formally represented claim propositions at different levels of formality is of main interest in the field of knowledge extraction from unstructured (web pages, social networks) or semi-structured (Wikipedia) sources. Populate and build knowledge bases and thus provide structured knowledge on the web has been of central interest in the NLP, the web data mining and the semantic web communities over the past decades, focusing on a variety of tasks such as named entity recognition, entity linking, relation extraction or word sense disambiguation. The extensive research in this field has lead to a very broad range of works. A comparison of generic information extraction tools and systems is provided in [129], while [36] and [130] fo-

cus on semantic web approaches (aiming at the provision of structured knowledge for populating ontologies, linked data and knowledge graphs). In that relation, we point to the book on NLP methods for building the semantic web [131] as well as a recent survey on fact extraction from the web [132]. Relation extraction and ontology learning from text are overviewed in [133] and [134], respectively, while [135] dedicate their survey on the particular task of extracting event-related knowledge. Finally, [136] consider methods that take the inverse approach of annotating documents with entities or statements of facts based on existing knowledge bases.

If we look at textual representations of a claim, the task can be approached by first extracting textual utterances (see below), then grouping them together according to their meaning by the help of textual similarity methods (some of them described in 5.3.2) and then identifying in a cluster of semantically equivalent utterances one that will serve as an identifier for the meaning of the claim. A formal approach to the assignment of textual identifiers to a set of equivalent claims has not been discussed in the literature, to the best of our knowledge, but the task relates closely to the text summarization task, which is surveyed in [137].

**Detecting claim topics.** Detecting what claims are about is a challenging issue. If available, context such as the source articles the claim was extracted from, a claim review article or the discourse the utterance was embedded in, e.g. the given subject in a debating por-



tal can be considered for claim topic detection, and standard NLP methods of topic extraction, modeling or detection from text [36] can be employed. However, detecting the topic when only the textual content of a claim utterance can be considered or the textual context is also sparse is challenging. Approaches developed for extracting topics from short text (like tweets and micro-blogs) can be adapted for claim topic modeling [138]. However, the complex structure and positioning in a context of elements (such as sources, authors and other entities) has to be taken into consideration when predicting topics of claims. Topics can be seen as groups of equivalent claims (e.g. all claims pertaining to “US immigration policies”) situated in a network of contextual entities (e.g. a knowledge graph such as the one given in our model implementation example in Fig. 7). Therefore, link prediction methods on knowledge graphs may be used, where a recent work [139] studies the effectiveness of neural graph embedding features for claim topic prediction as well as their complementarity with text embeddings. The authors show, however, that state-of-the-art link prediction models fail to capture equivalence structures and transfer poorly to downstream tasks such as claim topic prediction, which may, however, also be connected to the lack of sufficiently large and reliable ground truth data (topic-labeled claims) that would allow to train neural embedding models. This calls for the development of novel methods that surpass the state-of-the-art graph embedding model’s reliance on a local link prediction objective, which likely limits the ability of these models to capture more complex relationships (e.g. equivalence cliques between claims, keywords and topic concepts) and the generation of suitable ground truth data.

*Extracting viewpoints and stances.* Existing computational models [140] describe viewpoints via a summarization framework, able to find phrases that best reflect them. In [111, 141], unsupervised topic models are proposed to jointly discover viewpoints, aspects and opinions in text and social media. An unsupervised model for viewpoint detection in online debate forums, proposed in [142], favors ‘heterophily’ over ‘homophily’ when encoding the nature of the authors’ interactions in online debates. With respect to viewpoint detection in social media, the model in [143] groups Twitter users along a common ideological dimension based on who they follow. A graph partitioning method that exploits social interactions for the discovery of different user groups (representing distinct

viewpoints) discussing about a controversial topic in a social network is proposed in [112], also providing a method to explain the discovered viewpoints by detecting descriptive terms that characterize them.

Our model suggests, in line with the current research, that viewpoints with respect to topics take the form of polarized opinions. Given a controversial topic, for example an issue like *climate change*, viewpoint discovery aims to find the general viewpoint expressed in a piece of text or supported by a user. This task can indeed be considered a sub-task of opinion mining, which aims to analyze opinionated documents and to infer properties such as subjectivity or polarity. The survey in [144] provides a general review of the opinion mining and sentiment analysis tasks. However, for some topics, there may be more than two viewpoints. As of yet, there is limited research that studies these cases.

Viewpoint extraction is closely connected to the *stance detection* problem, a supervised classification problem in natural language processing where the stance of a piece of text towards a particular *target* is explored. Stance detection has been applied in different contexts, including *social media* (stance of a tweet towards an entity or topic) [145–151], *online debates* (stance of a user post or argument/claim towards a controversial topic or statement) [152–155], and *news media* (stance of an article towards a claim) [2–4, 40, 156]. In contrast to the works on viewpoint extraction described previously, works in stance detection focus more on supervised models and textual features (like the sentiment expressed in the text, or the use of polarised words), and less on the structure of the underlying network of users or documents which can be exploited by unsupervised approaches. For two recent surveys of stance detection works, we refer to [157, 158].

In a recent work, [159] compare untargeted and targeted opinion mining methods (sentiment analysis, aspect-based sentiment analysis, stance detection) to infer approval of political actors in tweets. They can show that the compared targeted approaches have low generalizability on unseen and unfamiliar targets and that indirectly expressed stances are hard to detect, identifying the need for further research in this area.

### 5.1.2. Extracting claim utterances

*Textual utterance extraction.* In this survey, we focus on methods for extracting information from language rather than other modalities such as speech or video. The methods discussed in the literature, with



few exceptions, are tailored towards a particular context, topic or type of targeted utterances, usually referred to as *claims* in these works. Identifying and extracting argumentative components such as claims (also called propositions in these works) or evidence units (also called premises) is a central task in the *argumentation mining* (AM) field [23, 160]. The first survey on the topic [161] assumes the availability of an argumentative text and focuses on the problem of analyzing the underlying structure of the presented argument from two perspectives: (1) argument annotation schemes drawing from works in the classical AI field of argumentation and (2) automatic argumentation mining, discussing the first approaches that enhance the historical field with data-centered machine learning approaches. A more recent survey [162] provides a structured view on the existing models, methods, and applications in AM attempting to draw a single unifying view over a plethora of related sub-tasks and dispersed efforts. The authors define the AM problem as a pipeline consisting of the detection of argument components in raw text and predicting the structure (or relations) between these components, where the former is of particular interest to the task that we consider in this section. Building on and completing these surveys, [163] adopt a data-driven perspective of the existing work in AM with a focus on applications, algorithms, features, and resources for evaluation of state of the art systems. Taking also a data-driven perspective, the difficulty of devising cross-domain claim identification approaches has been discussed and analyzed in [23] by using multiple domain-specific data sets. In that, the authors address the analysis of the generalization properties of systems and features across heterogeneous domains and study their robustness across the underlying fields.

The extraction of a claim is the first step in a *computational fact-checking* pipeline, where it is common to see fact verification as a three-step process: (i) detecting/extracting a check-worthy claim, (ii) reviewing the claim with respect to its veracity and (iii) publishing the reviewed claim [28, 164].<sup>21</sup> In [164], the authors propose a first version of the Claimbuster tool with a particular focus on the extraction of check-worthy claims. The claim-spotting problem is defined as a two step task, comprising (1) classification of pieces of text as check-worthy or not and (2) their ranking with re-

spect to their check-worthiness. An end-to-end fact-checking platform, including both steps (1) and (2), is presented in a follow up work [87]. To overpass the limitations of using hand-crafted features for claim detection, [165] propose a neural check-worthiness ranking model that represents a claim as a set of features, where each word is accounted for by its embedding (capturing its semantics) and its syntactic dependencies (capturing its relation to other terms in the sentence). The extraction of simple claims about statistical properties to be subjected to verification is addressed in [115]. The authors apply a distantly supervised claim identification approach that relies on approximate matching between numerical values in text and a knowledge base (Freebase). A relevant line of work has been followed in the field of subjectivity analysis of text, proposing approaches which aim at classifying sentences into objective and subjective categories, e.g., [166–168]. It has been shown in [87] that subjectivity identifiers are limited in discerning factual claims as compared to the method presented by Claimsbuster.

*Annotating utterances.* In our model, we discuss an annotation of utterances based on (1) entities (such as names, dates, locations, etc.) and (2) lower-level linguistic features extracted from the text that can be useful for a number of tasks, such as bias detection or fake-news analysis, as discussed in [169]. For (1), one can turn to the literature surrounding (end-to-end) Entity Linking<sup>22</sup>, particularly the exhaustive survey in [170]. The features in (2) include characteristics of the discourse, such as shades of irony or the overall polarity score of the expression, as well as linguistic or syntactic cues (part-of-speech (POS) tags, syntax, dependencies, semantic parsing, punctuation or capitalization) that can be indicative of a certain intention. For the identification of such cues, one could turn to NLP annotation pipelines (with standardized annotation type taxonomy). The industrial standard is UIMA (Unstructured Information Management Applications) [171], a comprehensive meta-framework for inter-operable linguistic annotation. Recent developments in deep approaches to NLP have led to the development of ad-hoc annotation models such as SpaCy.<sup>23</sup>

<sup>21</sup>[https://fullfact.org/media/uploads/full\\_fact-the\\_state\\_of\\_automated\\_factchecking\\_aug\\_2016.pdf](https://fullfact.org/media/uploads/full_fact-the_state_of_automated_factchecking_aug_2016.pdf)

<sup>22</sup>[http://nlpprogress.com/english/entity\\_linking.html](http://nlpprogress.com/english/entity_linking.html)

<sup>23</sup><https://spacy.io/>

*Claim utterance source extraction.* Sources are identified as the media that publishes a claim. Their extraction can be straightforward in many cases (e.g. when the utterance itself is extracted directly from its source). In certain cases identifying the original source may be more challenging and would require tracking down the claim to its original publication by, e.g. following cascades of retweets or identifying and analysing quotations [1, 172–174].

### 5.1.3. Extracting claim context.

This group of approaches deals with annotating a claim with contextual information that helps reply to the questions who uttered the claim when and where. In order to extract a date or a location one can rely on Entity Linking (EL) or Named Entity Recognition (NER) techniques outlined in the previous section. We focus in more detail on the tasks of event detection and author identification and attribution.

*Event detection.* The event in which a claim was uttered is an important component from the context that defines a claim. An event can be seen as a complex entity defined by a set of attributes, such as a date, persons involved and a location. Following this definition, one can apply the methods described in the previous paragraph in order to extract independently these components to populate an event. However, recent approaches consider an event as an atomic entity that can be detected from web corpora (often social networks) [175, 176].

*Author identification and extraction.* Identifying the author of an utterance is not trivial [124] yet authorship is crucial for interpreting its meaning. Moreover, claims are often quoted by distant sources, e.g. in news articles or other media. The attribution of content to an author<sup>24</sup> is consequently gaining increased attention in the context of the analysis of news articles, e.g. by [173, 177] who build structured databases of claims with extracted quotes and author information. Approaches for quotation extraction and attribution from newspaper articles for both direct and indirect speech usually comprise three different component identification steps: (1) cue phrases signalling the presence of a quotation (e.g. “say” or “criticize”) are identified using manually curated word lists [178] or classifiers trained

on labelled data [173, 179, 180]. On this basis, (2) quotation content spans are identified using manually defined syntactic rules [178], conditional random fields (CRFs) [173, 179] or semi-markov models [180]. Finally, (3) author entities are identified, typically using sequence models such as CRFs [173, 179, 181]. In that, [173] extend [181]’s sequence-based quote attribution to a two-stage approach using maximum entropy classifiers for connecting cue and content spans and cue and author spans, respectively, allowing multiple content and cue spans to take part in an attribution relation. A different approach is followed by [182] who employ pattern-based bootstrapping to extract quotation-speaker pairs. A recent paper [183] extracts structured information from fact-checking articles, including the “claimant”. This corresponds to either the source or the author of the claim, depending on which of those is mentioned in the fact-checks where usually, this distinction is not made.

## 5.2. Claim verification

A number of terms, such as fact-checking, truth discovery, claim or fact verification pertain to a large degree to the process of the automatic assignment of a veracity score to a statement uttered by a particular person or a group of people [28]. Note that the analysis of false or mis-information spread, or fake-news detection,<sup>25</sup> defined and surveyed in [184], often deal with entire news articles or outlets and are, therefore, broader problems where claim verification can be seen as one of their ingredients.

Claim truthfulness verification is reviewed in [28, 185], where [28] in particular propose to unify diverging definitions of the task and its components from various disciplines, such as NLP, machine learning, knowledge representation, databases, and journalism. Indeed, most of the existing techniques rely on background knowledge sources (e.g. encyclopedic knowledge graphs, such as DBpedia or Freebase) that provide a “truthfulness context” [116, 164] and a combination of various computational methods in order to infer the veracity scores of a claim either from those background knowledge sources or, more rarely, in a self-dependent manner. In addition, versatile features pertaining to all three main components of our model (meaning, utterance and context) are often considered in a combined manner, making it difficult to

<sup>24</sup>Coined “source” in the respective works; in order to not confuse different terminologies, we are referring to these entities as “authors” in the following text although this diverges from the naming used in the literature in this field.

<sup>25</sup>“False and often sensational information disseminated under the guise of news reporting”, according to Collins English Dictionary.

break down claim verification approaches along each of these three axes independently.

In certain cases, claims are given a structured form (e.g. triples or database queries), which allows for the verification of entity-centric kind of information by calling on machine learning techniques [21]. In that, fact verification can be seen as a particular kind of a link prediction or knowledge base augmentation task [45, 46]. In contrast, certain methods apply symbolic inference approaches on KGs in order to infer the truth value of a statement [69], or to identify potential errors [186]. A multitude of features, machine learning models and inference techniques are combined together in the KB construction approach presented in [22].

In other cases, statements are taken in their textual form [5, 187], while again largely Machine Learning (ML) techniques are applied in order to assess their veracity. Training data in the form of examples of true and false claims come either from archives of fact-checked statements [187–189] or from manually labelled (often crowdsourced) collections of claims [190–192]. Statistical (topic) models as well as standard NLP filters are used in order to construct a feature space. Note that the majority of the ML-based approaches rely primarily on highly contextualized features on document/text level, such as words, n-grams, salient entities and topics [7]. Additional context- and aspect-related features such as provenance, time and sources are considered in [5, 188]. An analysis of news corpora is provided by [169] in an attempt to identify linguistic and stylistic cues that help discriminate facts from false information. In addition, certain approaches, like e.g. [193], look at how a claim spreads through a crowd or how sources and claims are connected, exploiting social/community-related features.

### 5.3. Interlinking

There exist a variety of types of relations between claims and in particular between their components as introduced in our conceptual model. We consider that the problem of claim relatedness depends on the particular perspective and application context—for example, two claims can be considered contextually similar because they have been uttered at the same event by the same person, but still differ in their meaning and textual expression. Following the main building components of our model, we identify a number of dimensions on which this problem can be studied. One could be interested in relating instances of propositions, utterances or contexts within each of these three groups.

These are the kind of relations that will be discussed in this section. Else, one can look into cross-class relations (e.g. establishing the association between an utterance and its author or viewpoint). Such relations result from knowledge extraction processes already discussed in section 5.1. Although most of these problems can be considered as challenging with little existing work that approach them directly, we will outline below relevant works.

#### 5.3.1. Relating Propositions

According to our model, the proposition, or meaning, of a claim is materialized via a particular representation (e.g. a natural language or a logical expression) and is further described by its topics to which we associate viewpoints. As discussed in section 5.1, different extraction methods can be applied in order to derive those representations. Independent on the particular representation type, we outline three general types of relations that we can establish between proposition instances: *equivalence (same-as)*, *similarity* and *relatedness*.

*Same-as.* The equivalence or identity relation bounds together claims that have the exact same meaning. In the case of textual expression of the meaning of a claim, when two propositions are expressed differently although they convey the same message (have the same meaning), we talk of a relation of paraphrase. Paraphrasing detection allows to discover equivalent text fragments that differ (to a given extent), where neural language models are currently largely applied to the task [194, 195]. In the case of a symbolic or formal expression of a claim (or a fact), we outline works on relation alignment, such as [196].

*Similar.* Two propositions can be similar to a given degree on a scale between “identical” (represented by the *same-as* relation) and “dissimilar”. This notion relates to that of semantic similarity discussed, for example, in [197] and tackled in the Semantic Textual Similarity task [198, 199]. A first systematic study on finding similar claims is proposed by [200].

*Related.* Relatedness, as opposed to similarity, covers “any kind of lexical or functional association” [197] and is, hence, a more general concept than semantic similarity. Relatedness encloses various relationships, such as meronymy (a relation of composition (part-of) that is such that the meaning of a complex expression relates and can be expressed by the meanings of the parts from which it is constructed), antonymy (opposite meanings, including conflicting / contradicting claims), logical or textual entailment

[201], same topic, or any kind of functional relationships or relationships or frequent association. A survey of semantic relatedness methods, evaluation and datasets is given by [202, 203]. As opposed to logical entailment, textual entailment is understood as a relationship between pairs of text fragments where one entails the other if a human reading the former would be able to infer that the latter holds.

### 5.3.2. Relating utterances

Several works address finding equivalent claims in the context of claim verification [87, 164], where a claim matcher is a component of a fact-checking system matching new claims to claims that have already been checked. Concerning the matching of text fragments more generally, recent advances in neural NLP and the advent of deep contextualized language models for language understanding, have allowed a renewal state-of-the-art techniques for matching text fragments through the pooling or aggregation of classical [204] and contextualized word-embeddings [205, 206] into phrase, sentence or document embeddings [207, 208] and the computation of distance metrics to find the closest matching utterances.

In the context of the *Open Claims* Model, relations between utterances can further be derived from the relations between their constituents. For example, an utterance is a repetition of another utterance when all constituents are equal except for at least one attribute of the context such as the author or the date. An utterance is a paraphrase of another utterance when the propositions are equal but the (linguistic) representations differ. Deriving relations when some of the constituents are similar or related, rather than equal, remains a question for further research.

Other types of relations comprise support/attack relations or pro/con stances. Many works treat this as an extraction and classification task, e.g. classifying an argumentative unit as *evidence* (see section 5.1, while others treat this as an argumentative relation extraction task, e.g. relating two units with a *supports* or *attacks* relation [128, 209, 210].

### 5.3.3. Relating contexts

A context is broken down to its constituents: events, entities, dates, etc. Establishing links among contexts comes down to linking their respective components. For that purpose, one may call upon state-of-the-art approaches to data linking, where, following years of research and practice, a wealth of methodological approaches and tools are currently out there [211]. Among those, property-centric approaches (e.g. [212,

213]) can be of particular interest in order to establish relations (like identity or overlap) between different contexts, comparing their elements individually by the help of well-suited similarity measures (e.g. measures of similarity between proper names or dates).

## 6. Conclusion

This paper bridges the gap between various disciplines involved with online discourse analysis from a range of perspectives by (a) surveying definitions of claims, facts and related concepts across different research areas and communities, (b) establishing a shared conceptualisation and vocabulary in this context and (c) discussing a range of tasks involved with such notions, for instance, for extracting or interlinking related concepts through NLP techniques. We contribute to a shared understanding of a wide range of disparate yet strongly related research areas, facilitating a deeper understanding of shared methods, approaches and concepts and the potential for reuse and cross-fertilisation across communities. Below, we highlight under-researched areas and potential future directions.

Currently, a framework for claim relatedness and similarity is missing. Several works from different fields appear to deal with the problem from different perspectives, but an approach that takes into consideration the various aspects of a claim, as well as its various representations, as defined in our *Open Claims* model in order to discover claim relatedness or similarity of different kinds is yet to be proposed. While there are works addressing the extraction of structured information from claims [183], allowing for example the detection of nested claims, current fact-checking methods and sites largely ignore such issues. For instance, a complex claim can have different (or no) truth rating as compared to its constituents. For instance, “Colin Kaepernick says Winston Churchill said, “A lie gets halfway around the world before the truth has a chance to get its pants on.””<sup>26</sup> The claim that Kaepernick uttered this is true, while the claim within the claim is false since Churchill never said that. Using the proposed model, such cases can be modeled precisely and unambiguously.

Given the subtle differences between claims, where meaning often derives from subtext and context, dis-

<sup>26</sup><https://www.politifact.com/factchecks/2017/oct/09/colin-kaepernick/nfls-colin-kaepernick-incorrectly-credits-winston/>

ambiguating claim representations, e.g. when mapping novel claims to knowledge bases of fact-checked claims, appears challenging. Even for humans, deciding on the type of relationship of two claims is a non-trivial task. For example, the claim “Interest on the debt will exceed defense spending by 2022” provides an exact date, while “Interest on debt will exceed security spending” does not provide a date. Can these two claims be considered the same, and, if not, what is their relation? Using the proposed model, such subtle differences can be made explicit. In addition, automated fact-checking has the potential to elevate the problem given its lack of maturity at different steps, where for instance, the classification of half-correct or poorly disambiguated claims as correct may introduce further false claims into the wild.

Similarly, the process of stance detection is challenging as it has been shown to not work well for the minority class, i.e. documents disagreeing with the claim [4, 214], and for unseen targets [159]. Little research in viewpoint discovery deals with extracting viewpoints for more than two polarized positions, a topic that could be worthwhile researching for the analysis of debates.

Detecting claim topics and linking those to a specific commonly shared vocabulary or thesaurus of topics (like, e.g., the TheSoZ [215] or the Unesco<sup>27</sup> thesauri) appears to be a difficult and under-researched topic that promises to enhance claim retrieval, improve search and interoperability across sources, and facilitate access to currently existing or yet to be constructed structured resources of claims [26].

Generally speaking, considering the wide variety of methods and datasets involving claims and related notions, adopting a shared and well-defined vocabulary has the potential to significantly increase impact and reuse of research methods and data.

## References

- [1] S. Vosoughi, D. Roy, and S. Aral, “The spread of true and false news online,” *Science*, vol. 359, no. 6380, pp. 1146–1151, Mar. 2018, publisher: American Association for the Advancement of Science Section: Report. [Online]. Available: <https://science.sciencemag.org/content/359/6380/1146>
- [2] X. Wang, C. Yu, S. Baumgartner, and F. Korn, “Relevant Document Discovery for Fact-Checking Articles,” in *Companion Proceedings of the The Web Conference 2018*, ser. WWW ’18. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, Apr. 2018, pp. 525–533. [Online]. Available: <https://doi.org/10.1145/3184558.3188723>
- [3] G. Bhatt, A. Sharma, S. Sharma, A. Nagpal, B. Raman, and A. Mittal, “Combining Neural, Statistical and External Features for Fake News Stance Identification,” in *Companion Proceedings of the The Web Conference 2018*, ser. WWW ’18. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, Apr. 2018, pp. 1353–1357. [Online]. Available: <https://doi.org/10.1145/3184558.3191577>
- [4] A. Hanselowski, A. Pvs, B. Schiller, F. Caspelherr, D. Chaudhuri, C. M. Meyer, and I. Gurevych, “A Retrospective Analysis of the Fake News Challenge Stance-Detection Task,” in *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, 2018, p. 16.
- [5] K. Popat, S. Mukherjee, J. Strögen, and G. Weikum, “Where the Truth Lies: Explaining the Credibility of Emerging Claims on the Web and Social Media,” in *Proceedings of the 26th International Conference on World Wide Web Companion - WWW ’17 Companion*. Perth, Australia: ACM Press, 2017, pp. 1003–1012. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3041021.3055133>
- [6] S. Tschatschek, A. Singla, M. Gomez Rodriguez, A. Merchant, and A. Krause, “Fake News Detection in Social Networks via Crowd Signals,” in *Companion Proceedings of the The Web Conference 2018*, ser. WWW ’18. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, Apr. 2018, pp. 517–524. [Online]. Available: <https://doi.org/10.1145/3184558.3188722>
- [7] N. Hassan, B. Adair, J. T. Hamilton, C. Li, M. Tremayne, J. Yang, and C. Yu, “The Quest to Automate Fact-Checking,” *Proceedings of the 2015 Computation+ Journalism Symposium*, 2015.
- [8] A. Lauscher, G. Glavaš, and S. P. Ponzetto, “An Argument-Annotated Corpus of Scientific Publications,” in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 40–46. [Online]. Available: <http://aclweb.org/anthology/W18-5206>
- [9] A. Lauscher, G. Glavaš, and K. Eckert, “ArguminSci: A Tool for Analyzing Argumentation and Rhetorical Aspects in Scientific Writing,” in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 22–28. [Online]. Available: <http://aclweb.org/anthology/W18-5203>
- [10] J. M. Gonzalez Pinto, J. Wawrzinek, and W.-T. Balke, “What Drives Research Efforts? Find Scientific Claims that Count!” in *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. Champaign, IL, USA: IEEE, Jun. 2019, pp. 217–226. [Online]. Available: <https://ieeexplore.ieee.org/document/8791177/>
- [11] P. Accuosto and H. Saggion, “Transferring Knowledge from Discourse to Arguments: A Case Study with Scientific Abstracts,” in *Proceedings of the 6th Workshop on Argument Mining*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 41–51. [Online]. Available: <https://www.aclweb.org/anthology/W19-4505>

<sup>27</sup><http://vocabularies.unesco.org/thesaurus>

- [12] N. Green, "Towards Creation of a Corpus for Argumentation Mining the Biomedical Genetics Research Literature," in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 11–18. [Online]. Available: <http://aclweb.org/anthology/W14-2102>
- [13] H. Graves, R. Graves, R. Mercer, and M. Akter, "Titles That Announce Argumentative Claims in Biomedical Research Articles," in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 98–99. [Online]. Available: <http://aclweb.org/anthology/W14-2113>
- [14] J. M. González Pinto and W.-T. Balke, "Scientific Claims Characterization for Claim-Based Analysis in Digital Libraries," in *Digital Libraries for Open Knowledge*, E. Méndez, F. Crestani, C. Ribeiro, G. David, and J. C. Lopes, Eds. Cham: Springer International Publishing, 2018, vol. 11057, pp. 257–269, series Title: Lecture Notes in Computer Science. [Online]. Available: [http://link.springer.com/10.1007/978-3-030-00066-0\\_22](http://link.springer.com/10.1007/978-3-030-00066-0_22)
- [15] C. Kirschner, J. Eckle-Kohler, and I. Gurevych, "Linking the Thoughts: Analysis of Argumentation Structures in Scientific Publications," in *Proceedings of the 2nd Workshop on Argumentation Mining*. Denver, CO: Association for Computational Linguistics, 2015, pp. 1–11. [Online]. Available: <http://aclweb.org/anthology/W15-0501>
- [16] K. Hyland, *Hedging in Scientific Research Articles*. John Benjamins Publishing, Mar. 1998, google-Books-ID: UUp-CAAAQBAJ.
- [17] M. Neves, D. Butzke, and B. Grune, "Evaluation of Scientific Elements for Text Similarity in Biomedical Publications," in *Proceedings of the 6th Workshop on Argument Mining*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 124–135. [Online]. Available: <https://www.aclweb.org/anthology/W19-4515>
- [18] N. Green, "Identifying Argumentation Schemes in Genetics Research Articles," in *Proceedings of the 2nd Workshop on Argumentation Mining*. Denver, CO: Association for Computational Linguistics, 2015, pp. 12–21. [Online]. Available: <http://aclweb.org/anthology/W15-0502>
- [19] —, "Proposed Method for Annotation of Scientific Arguments in Terms of Semantic Relations and Argument Schemes," in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 105–110. [Online]. Available: <http://aclweb.org/anthology/W18-5213>
- [20] M. Hernández A. and J. M. Gómez, "Survey in sentiment, polarity and function analysis of citation," in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 102–103. [Online]. Available: <http://aclweb.org/anthology/W14-2115>
- [21] R. Yu, U. Gadiraju, B. Fetahu, O. Lehmborg, D. Ritze, and S. Dietze, "KnowMore – knowledge base augmentation with structured web markup," *Semantic Web*, vol. 10, no. 1, pp. 159–180, Dec. 2018. [Online]. Available: <https://www.medra.org/serve/aliasResolver?alias=iospress&doi=10.3233/SW-180304>
- [22] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, and W. Zhang, "Knowledge vault: a web-scale approach to probabilistic knowledge fusion," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14*. New York, New York, USA: ACM Press, 2014, pp. 601–610. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2623330.2623623>
- [23] J. Daxenberger, S. Eger, I. Habernal, C. Stab, and I. Gurevych, "What is the Essence of a Claim? Cross-Domain Claim Identification," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 2055–2066. [Online]. Available: <http://aclweb.org/anthology/D17-1218>
- [24] B. Torsi and R. Morante, "Annotating Claims in the Vaccination Debate," in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 47–56. [Online]. Available: <http://aclweb.org/anthology/W18-5207>
- [25] K. Boland, P. Fafalios, A. Tchechmedjiev, K. Todorov, and S. Dietze, "Modeling and Contextualizing Claims," in *Second International Workshop on Contextualized Knowledge Graphs (CKG2019) @ ISWC*, 2019.
- [26] A. Tchechmedjiev, P. Fafalios, K. Boland, M. Gasquet, M. Zloch, B. Zapilko, S. Dietze, and K. Todorov, "Claim-sKG: A Knowledge Graph of Fact-Checked Claims," in *The Semantic Web – ISWC 2019*, ser. Lecture Notes in Computer Science, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, and F. Gandon, Eds. Cham: Springer International Publishing, 2019, pp. 309–324.
- [27] L. Konstantinovskiy, O. Price, M. Babakar, and A. Zubiaga, "Towards Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection," *arXiv:1809.08193 [cs]*, Sep. 2018, arXiv: 1809.08193. [Online]. Available: <http://arxiv.org/abs/1809.08193>
- [28] J. Thorne and A. Vlachos, "Automated Fact Checking: Task Formulations, Methods and Future Directions," in *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, Aug. 2018, pp. 3346–3359. [Online]. Available: <https://www.aclweb.org/anthology/C18-1283>
- [29] X. Zhou, R. Zafarani, K. Shu, and H. Liu, "Fake News: Fundamental Theories, Detection Strategies and Challenges," in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, ser. WSDM '19. New York, NY, USA: Association for Computing Machinery, Jan. 2019, pp. 836–837. [Online]. Available: <https://doi.org/10.1145/3289600.3291382>
- [30] A. Zubiaga, A. Aker, K. Bontcheva, M. Liakata, and R. Procter, "Detection and Resolution of Rumours in Social Media: A Survey," *ACM Computing Surveys*, vol. 51, no. 2, pp. 32:1–32:36, Feb. 2018. [Online]. Available: <https://doi.org/10.1145/3161603>
- [31] I. Balazevic, C. Allen, and T. Hospedales, "TuckER: Tensor Factorization for Knowledge Graph Completion," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computa-

- tional Linguistics, 2019, pp. 5184–5193. [Online]. Available: <https://www.aclweb.org/anthology/D19-1522>
- [32] M. H. Gad-Elrab, D. Stepanova, J. Urbani, and G. Weikum, “ExFaKT: A Framework for Explaining Facts over Knowledge Graphs and Text,” in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining - WSDM '19*. Melbourne VIC, Australia: ACM Press, 2019, pp. 87–95. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3289600.3290996>
- [33] N. Veira, B. Keng, K. Padmanabhan, and A. Veneris, “Unsupervised Embedding Enhancements of Knowledge Graphs using Textual Associations,” in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. Macao, China: International Joint Conferences on Artificial Intelligence Organization, Aug. 2019, pp. 5218–5225. [Online]. Available: <https://www.ijcai.org/proceedings/2019/725>
- [34] N. Voskarides, E. Meij, R. Reinanda, A. Khaitan, M. Osborne, G. Stefanoni, P. Kambadur, and M. de Rijke, “Weakly-supervised Contextualization of Knowledge Graph Facts,” *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval - SIGIR '18*, pp. 765–774, 2018, arXiv: 1805.02393. [Online]. Available: <http://arxiv.org/abs/1805.02393>
- [35] M. Al-Bakri, M. Atencia, S. Lalande, and M.-C. Rousset, “Inferring Same-as Facts from Linked Data: An Iterative Import-by-Query Approach,” in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, ser. AAAI'15, 2015, pp. 9–15.
- [36] J. L. Martinez-Rodriguez, A. Hogan, and I. Lopez-Arevalo, “Information extraction meets the Semantic Web: A survey,” *Semantic Web*, vol. 11, no. 2, pp. 255–335, Feb. 2020. [Online]. Available: <https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/SW-180333>
- [37] H. Zhu, R. Xie, Z. Liu, and M. Sun, “Iterative Entity Alignment via Joint Knowledge Embeddings,” in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization, Aug. 2017, pp. 4258–4264. [Online]. Available: <https://www.ijcai.org/proceedings/2017/595>
- [38] A. Padiá, K. Kalpakis, F. Ferraro, and T. Finin, “Knowledge Graph Fact Prediction via Knowledge-Enriched Tensor Factorization,” *Journal of Web Semantics*, 2019, arXiv: 1902.03077. [Online]. Available: <http://arxiv.org/abs/1902.03077>
- [39] M. Chen, W. Zhang, W. Zhang, Q. Chen, and H. Chen, “Meta Relational Learning for Few-Shot Link Prediction in Knowledge Graphs,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, 2019, pp. 4216–4225. [Online]. Available: <https://www.aclweb.org/anthology/D19-1431>
- [40] Zhan, Q. Liang, S. Lipani, A. Ren, Z. and Yilmaz, E., “From Stances’ Imbalance to Their Hierarchical Representation and Detection,” in *Proceedings of WWW '19: The World Wide Web Conference*, ser. WWW '19, San Francisco, CA, USA, 2019, pp. 2323–2332. [Online]. Available: [https://www.researchgate.net/publication/331304021\\_From\\_Stances%27\\_Imbalance\\_to\\_Their\\_Hierarchical\\_Representation\\_and\\_Detection](https://www.researchgate.net/publication/331304021_From_Stances%27_Imbalance_to_Their_Hierarchical_Representation_and_Detection)
- [41] B. D. Trisedya, G. Weikum, J. Qi, and R. Zhang, “Neural Relation Extraction for Knowledge Base Enrichment,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 229–240. [Online]. Available: <https://www.aclweb.org/anthology/P19-1023>
- [42] C. Wang, M. Yan, C. Yi, and Y. Sha, “Capturing Semantic and Syntactic Information for Link Prediction in Knowledge Graphs,” in *The Semantic Web – ISWC 2019*, ser. Lecture Notes in Computer Science, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, and F. Gandon, Eds. Cham: Springer International Publishing, 2019, pp. 664–679.
- [43] M. Rospocher, M. van Erp, P. Vossen, A. Fokkens, I. Aldabe, A. Soroa, T. Ploeger, and T. Bogaard, “Building Event-Centric Knowledge Graphs from News,” *Journal of Web Semantics*, vol. 37–38, pp. 132–151, 2016.
- [44] I. Augenstein, D. Maynard, and F. Ciravegna, “Distantly supervised Web relation extraction for knowledge base population,” *Semantic Web*, vol. 7, no. 4, pp. 335–349, May 2016. [Online]. Available: <https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/SW-150180>
- [45] G. L. Ciampaglia, P. Shiralkar, L. M. Rocha, J. Bollen, F. Menczer, and A. Flammini, “Computational Fact Checking from Knowledge Networks,” *PLOS ONE*, vol. 10, no. 6, p. e0128193, Jun. 2015. [Online]. Available: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0128193>
- [46] B. Shi and T. Weninger, “Discriminative Predicate Path Mining for Fact Checking in Knowledge Graphs,” *arXiv:1510.05911 [cs]*, vol. 104, pp. 123–133, 2016, arXiv: 1510.05911. [Online]. Available: <http://arxiv.org/abs/1510.05911>
- [47] Pasquale Minervini, Volker Tresp, Claudia d’Amato, and Nicola Fanizzi, “Adaptive Knowledge Propagation in Web Ontologies,” 2, vol. 12, no. 1, p. 28, 2017. [Online]. Available: [https://www.researchgate.net/publication/319282592\\_Adaptive\\_Knowledge\\_Propagation\\_in\\_Web\\_Ontologies](https://www.researchgate.net/publication/319282592_Adaptive_Knowledge_Propagation_in_Web_Ontologies)
- [48] V. Fionda and G. Pirrò, “Fact Checking via Evidence Patterns,” in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. Stockholm, Sweden: International Joint Conferences on Artificial Intelligence Organization, Jul. 2018, pp. 3755–3761. [Online]. Available: <https://www.ijcai.org/proceedings/2018/522>
- [49] D. Gerber, D. Esteves, J. Lehmann, L. Bühmann, R. Usbeck, A.-C. N. Ngomo, and R. Speck, “DeFacto - Temporal and Multilingual Deep Fact Validation,” *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 35, no. P2, 2015.
- [50] K. Popat, “Credibility Analysis of Textual Claims with Explainable Evidence,” Doctoral Thesis, Saarbrücken, 2019.
- [51] V.-P. Huynh and P. Papotti, “Towards a Benchmark for Fact Checking with Knowledge Bases,” in *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*. Lyon, France: ACM Press, 2018, pp. 1595–1598. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3184558.3191616>
- [52] Q. Zhang, Z. Sun, W. Hu, M. Chen, L. Guo, and Y. Qu, “Multi-view Knowledge Graph Embedding for Entity Alignment,” in *Proceedings of the Twenty-Eighth*

- International Joint Conference on Artificial Intelligence. Macao, China: International Joint Conferences on Artificial Intelligence Organization, Aug. 2019, pp. 5429–5435. [Online]. Available: <https://www.ijcai.org/proceedings/2019/754>
- [53] Z. H. Syed, N. Srivastava, M. Roder, and A.-C. Ngonga, “CO-PAAL – An Interface for Explaining Facts using Corroborative Paths,” in *ISWC Satellites*, 2019, pp. 201–204.
- [54] D. Tsurel, D. Pelleg, I. Guy, and D. Shahaf, “Fun Facts: Automatic Trivia Fact Extraction from Wikipedia,” in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, ser. WSDM ’17, 2017, pp. 345–354, arXiv: 1612.03896. [Online]. Available: <http://arxiv.org/abs/1612.03896>
- [55] C. Stahlhut, “Interactive Evidence Detection: train state-of-the-art model out-of-domain or simple model interactively?” in *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*. Hong Kong, China: Association for Computational Linguistics, 2019, pp. 79–89. [Online]. Available: <https://www.aclweb.org/anthology/D19-6613>
- [56] E. Aharoni, A. Polnarov, T. Lavee, D. Hershcovich, R. Levy, R. Rinott, D. Gutfreund, and N. Slonim, “A Benchmark Dataset for Automatic Detection of Claims and Evidence in the Context of Controversial Topics,” in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 64–68. [Online]. Available: <http://aclweb.org/anthology/W14-2109>
- [57] R. Rinott, L. Dankin, C. Alzate Perez, M. M. Khapra, E. Aharoni, and N. Slonim, “Show Me Your Evidence - an Automatic Method for Context Dependent Evidence Detection,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, 2015, pp. 440–450. [Online]. Available: <http://aclweb.org/anthology/D15-1050>
- [58] C. Stab and I. Gurevych, “Identifying Argumentative Discourse Structures in Persuasive Essays,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, 2014, pp. 46–56. [Online]. Available: <http://aclweb.org/anthology/D14-1006>
- [59] E. Shnarch, C. Alzate, L. Dankin, M. Gleize, Y. Hou, L. Choshen, R. Aharonov, and N. Slonim, “Will it Blend? Blending Weak and Strong Labeled Data in a Neural Network for Argumentation Mining,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Melbourne, Australia: Association for Computational Linguistics, 2018, pp. 599–605. [Online]. Available: <http://aclweb.org/anthology/P18-2095>
- [60] D. Trautmann, J. Daxenberger, C. Stab, H. Schütze, and I. Gurevych, “Fine-Grained Argument Unit Recognition and Classification,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, ser. AAAI’ 20, 2020, arXiv: 1904.09688. [Online]. Available: <http://arxiv.org/abs/1904.09688>
- [61] M. Lippi and P. Torroni, “Context-independent claim detection for argument mining,” in *Proceedings of the 24th International Conference on Artificial Intelligence*, ser. IJCAI’15. Buenos Aires, Argentina: AAAI Press, Jul. 2015, pp. 185–191.
- [62] V. R. Walker, D. Foerster, J. M. Ponce, and M. Rosen, “Evidence Types, Credibility Factors, and Patterns or Soft Rules for Weighing Conflicting Evidence: Argument Mining in the Context of Legal Rules Governing Evidence Assessment,” in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 68–78. [Online]. Available: <http://aclweb.org/anthology/W18-5209>
- [63] V. Niculae, J. Park, and C. Cardie, “Argument Mining with Structured SVMs and RNNs,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada: Association for Computational Linguistics, 2017, pp. 985–995, arXiv: 1704.06869. [Online]. Available: <https://www.aclweb.org/anthology/P17-1091>
- [64] C. Hidey, E. Musi, A. Hwang, S. Muresan, and K. McKeown, “Analyzing the Semantic Types of Claims and Premises in an Online Persuasive Forum,” in *Proceedings of the 4th Workshop on Argument Mining*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 11–21. [Online]. Available: <http://aclweb.org/anthology/W17-5102>
- [65] T. Mayer, E. Cabrio, and S. Villata, “Evidence Type Classification in Randomized Controlled Trials,” in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 29–34. [Online]. Available: <http://aclweb.org/anthology/W18-5204>
- [66] Y. Nie, H. Chen, and M. Bansal, “Combining Fact Extraction and Verification with Neural Semantic Matching Networks,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 6859–6866, 2019, number: 01. [Online]. Available: <https://aaai.org/ojs/index.php/AAAI/article/view/4662>
- [67] R. Clancy, I. F. Ilyas, and J. Lin, “Scalable Knowledge Graph Construction from Text Collections,” in *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*. Hong Kong, China: Association for Computational Linguistics, 2019, pp. 39–46. [Online]. Available: <https://www.aclweb.org/anthology/D19-6607>
- [68] A. Hanselowski, H. Zhang, Z. Li, D. Sorokin, B. Schiller, C. Schulz, and I. Gurevych, “UKP-Athene: Multi-Sentence Textual Entailment for Claim Verification,” in *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 103–108. [Online]. Available: <http://aclweb.org/anthology/W18-5516>
- [69] V. Beretta, S. Harispe, S. Ranwez, and I. Mougenot, “Combining Truth Discovery and RDF Knowledge Bases to Their Mutual Advantage,” in *The Semantic Web – ISWC 2018*, ser. Lecture Notes in Computer Science, D. Vrandečić, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee, and E. Simperl, Eds. Cham: Springer International Publishing, 2018, pp. 652–668.
- [70] X. Wang, Q. Z. Sheng, L. Yao, X. Li, X. S. Fang, X. Xu, and B. Benatallah, “Empowering Truth Discovery with Multi-Truth Prediction,” in *Proceedings of the 25th ACM International Conference on Information and Knowledge Management - CIKM ’16*. Indianapolis, Indiana, USA: ACM Press, 2016, pp. 881–890. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2983323.2983767>



- [71] H. Xiao, J. Gao, Q. Li, F. Ma, L. Su, Y. Feng, and A. Zhang, "Towards Confidence Interval Estimation in Truth Discovery," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 3, pp. 575–588, Mar. 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8359426/>
- [72] S. Rosenthal and K. McKeown, "Detecting Opinionated Claims in Online Discussions," in *2012 IEEE Sixth International Conference on Semantic Computing*. Palermo, Italy: IEEE, Sep. 2012, pp. 30–37. [Online]. Available: <http://ieeexplore.ieee.org/document/6337079/>
- [73] M. Lippi and P. Torroni, "Argument mining from speech: detecting claims in political debates," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, ser. AAAI'16. Phoenix, Arizona: AAAI Press, Feb. 2016, pp. 2979–2985. [Online]. Available: [https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwiWvdK116bqAhUJmIsKHY\\_-DKUQFjABegQIBhAB&url=https%3A%2F%2Fwww.aaai.org%2Focs%2Findex.php%2FAAAI%2FAAAI16%2Fpaper%2Fdownload%2F12164%2F12050&usg=AOvVaw3ZP2OEAfjCKAFnTrVUSiKm](https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwiWvdK116bqAhUJmIsKHY_-DKUQFjABegQIBhAB&url=https%3A%2F%2Fwww.aaai.org%2Focs%2Findex.php%2FAAAI%2FAAAI16%2Fpaper%2Fdownload%2F12164%2F12050&usg=AOvVaw3ZP2OEAfjCKAFnTrVUSiKm)
- [74] L. Lugini and D. Litman, "Argument Component Classification for Classroom Discussions," in *Proceedings of the 5th Workshop on Argument Mining*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 57–67. [Online]. Available: <http://aclweb.org/anthology/W18-5208>
- [75] P. Besnard and A. Hunter, *Elements of Argumentation*. The MIT Press, 2008.
- [76] R. M. Palau and M.-F. Moens, "Argumentation mining: the detection, classification and structure of arguments in text," in *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, ser. ICAIL '09. New York, NY, USA: Association for Computing Machinery, Jun. 2009, pp. 98–107. [Online]. Available: <https://doi.org/10.1145/1568234.1568246>
- [77] R. Levy, B. Bogin, S. Gretz, R. Aharonov, and N. Slonim, "Towards an argumentative content search engine using weak supervision," in *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, Aug. 2018, pp. 2066–2081. [Online]. Available: <https://www.aclweb.org/anthology/C18-1176>
- [78] M. Lippi, M. Mamei, S. Mariani, and F. Zambonelli, "An Argumentation-Based Perspective Over the Social IoT," *IEEE Internet of Things Journal*, vol. 5, no. 4, pp. 2537–2547, Aug. 2018, conference Name: IEEE Internet of Things Journal.
- [79] T. Mayer, E. Cabrio, M. Lippi, P. Torroni, and S. Villata, "Argument Mining on Clinical Trials," in *COMMA 2018 - 7th International Conference on Computational Models of Argument Proceedings*, ser. Frontiers in Artificial Intelligence and Applications, vol. 305, Warsaw, Poland, Sep. 2018, pp. 137 – 148. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01876462>
- [80] D. M. Blei and J. D. McAuliffe, "Supervised topic models," in *Proceedings of the 20th International Conference on Neural Information Processing Systems*, ser. NIPS'07. Red Hook, NY, USA: Curran Associates Inc., Dec. 2007, pp. 121–128.
- [81] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy, "Hierarchical attention networks for document classification," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, Jun. 2016, pp. 1480–1489. [Online]. Available: <https://www.aclweb.org/anthology/N16-1174>
- [82] E. Durmus, F. Ladhak, and C. Cardie, "Determining Relative Argument Specificity and Stance for Complex Argumentative Structures," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 4630–4641. [Online]. Available: <https://www.aclweb.org/anthology/P19-1456>
- [83] E. Schiappa and J. P. Nordin, *Argumentation: Keeping Faith with Reason Keeping Faith With Reason*. Pearson Education, Jul. 2013, google-Books-ID: ZOn8nQEACAAJ.
- [84] C. Fierro, C. Fuentes, J. Pérez, and M. Quezada, "200K+ Crowdsourced Political Arguments for a New Chilean Constitution," in *Proceedings of the 4th Workshop on Argument Mining*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 1–10. [Online]. Available: <http://aclweb.org/anthology/W17-5101>
- [85] J. Thorne, A. Vlachos, O. Cocarascu, C. Christodoulopoulos, and A. Mittal, "The Fact Extraction and VERification (FEVER) Shared Task," in *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 1–9. [Online]. Available: <http://aclweb.org/anthology/W18-5501>
- [86] —, "The FEVER2.0 Shared Task," in *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*. Hong Kong, China: Association for Computational Linguistics, 2019, pp. 1–6. [Online]. Available: <https://www.aclweb.org/anthology/D19-6601>
- [87] N. Hassan, F. Arslan, C. Li, and M. Tremayne, "Toward Automated Fact-Checking: Detecting Check-worthy Factual Claims by ClaimBuster," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '17. New York, NY, USA: Association for Computing Machinery, Aug. 2017, pp. 1803–1812. [Online]. Available: <https://doi.org/10.1145/3097983.3098131>
- [88] S. Mohtaj, T. Himmelsbach, V. Woloszyn, and S. Moller, "Using External Knowledge Bases and Coreference Resolution for Detecting Check-Worthy Statements," in *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*, ser. CLEF 2019 Working Notes, Lugano, Switzerland, 2019, p. 8.
- [89] B. Ghanem, G. Glavačs, A. Giachanou, S. Paolo, P. Rosso, and F. Rangel, "UPV-UMA at CheckThat! Lab: Verifying Arabic Claims using a Cross Lingual Approach," in *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*, ser. CLEF 2019 Working Notes, Lugano, Switzerland, 2019, p. 10.
- [90] B. Ghanem, P. Rosso, and F. Rangel, "Stance Detection in Fake News A Combined Feature Representation," in *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 66–71. [Online]. Available: <http://aclweb.org/anthology/W18-5510>

- [91] A. Padia, F. Ferraro, and T. Finin, "Team UMBC-FEVER : Claim verification using Semantic Lexical Resources," in *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 161–165. [Online]. Available: <http://aclweb.org/anthology/W18-5527>
- [92] M. Hasanain, R. Suwaileh, T. Elsayed, A. Barron-Cedeno, and P. Nakov, "Overview of the CLEF-2019 CheckThat! Lab: Automatic Identification and Verification of Claims. Task 2: Evidence and Factuality," in *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*, ser. CLEF 2019 Working Notes, 2019, p. 15.
- [93] N. Hassan, M. Yousuf, M. Mahfuzul Haque, J. A. Suarez Rivas, and M. Khadimul Islam, "Examining the Roles of Automation, Crowds and Professionals Towards Sustainable Fact-checking," in *Companion Proceedings of The 2019 World Wide Web Conference*, ser. WWW '19. San Francisco, USA: Association for Computing Machinery, May 2019, pp. 1001–1006. [Online]. Available: <https://doi.org/10.1145/3308560.3316734>
- [94] H. Zhang, Q. Li, F. Ma, H. Xiao, Y. Li, J. Gao, and L. Su, "Influence-Aware Truth Discovery," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management - CIKM '16*. Indianapolis, Indiana, USA: ACM Press, 2016, pp. 851–860. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2983323.2983785>
- [95] J. Park and C. Cardie, "Identifying Appropriate Support for Propositions in Online User Comments," in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 29–38. [Online]. Available: <http://aclweb.org/anthology/W14-2105>
- [96] I. Habernal and I. Gurevych, "Argumentation Mining in User-Generated Web Discourse," *Computational Linguistics*, vol. 43, no. 1, pp. 125–179, Apr. 2017. [Online]. Available: <https://www.aclweb.org/anthology/J17-1004>
- [97] M. Liebeck, K. Esau, and S. Conrad, "What to Do with an Airport? Mining Arguments in the German Online Participation Project Tempelhofer Feld," in *Proceedings of the Third Workshop on Argument Mining (ArgMining2016)*. Berlin, Germany: Association for Computational Linguistics, 2016, pp. 144–153. [Online]. Available: <http://aclweb.org/anthology/W16-2817>
- [98] C. Stab and I. Gurevych, "Parsing Argumentation Structures in Persuasive Essays," *Computational Linguistics*, vol. 43, no. 3, pp. 619–659, Sep. 2017. [Online]. Available: [https://www.mitpressjournals.org/doi/abs/10.1162/COLI\\_a\\_00295](https://www.mitpressjournals.org/doi/abs/10.1162/COLI_a_00295)
- [99] S. Zhi, Y. Sun, J. Liu, C. Zhang, and J. Han, "ClaimVerif: A Real-time Claim Verification System Using the Web and Fact Databases," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management - CIKM '17*. Singapore, Singapore: ACM Press, 2017, pp. 2555–2558. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3132847.3133182>
- [100] D. Ghosh, S. Muresan, N. Wacholder, M. Aakhus, and M. Mitsui, "Analyzing Argumentative Discourse Units in Online Interactions," in *Proceedings of the First Workshop on Argumentation Mining*. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 39–48. [Online]. Available: <http://aclweb.org/anthology/W14-2106>
- [101] N. Kotonya and F. Toni, "Gradual Argumentation Evaluation for Stance Aggregation in Automated Fake News Detection," in *Proceedings of the 6th Workshop on Argument Mining*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 156–166. [Online]. Available: <https://www.aclweb.org/anthology/W19-4518>
- [102] P. Rajendran, D. Bollegala, and S. Parsons, "Contextual stance classification of opinions: A step towards enthymeme reconstruction in online reviews," in *Proceedings of the Third Workshop on Argument Mining (ArgMining2016)*. Berlin, Germany: Association for Computational Linguistics, 2016, pp. 31–39. [Online]. Available: <http://aclweb.org/anthology/W16-2804>
- [103] K. Joseph, L. Friedland, W. Hobbs, D. Lazer, and O. Tsur, "ConStance: Modeling Annotation Contexts to Improve Stance Classification," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 1115–1124. [Online]. Available: <http://aclweb.org/anthology/D17-1116>
- [104] G. Giasemidis, N. Kaplis, I. Agraftiotis, and J. R. C. Nurse, "A semi-supervised approach to message stance classification," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 1, pp. 1–11, Jan. 2020, arXiv: 1902.03097. [Online]. Available: <http://arxiv.org/abs/1902.03097>
- [105] P. Potash, A. Ferguson, and T. J. Hazen, "Ranking Passages for Argument Convincingness," in *Proceedings of the 6th Workshop on Argument Mining*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 146–155. [Online]. Available: <https://www.aclweb.org/anthology/W19-4517>
- [106] W. Ferreira and A. Vlachos, "Emergent: a novel data-set for stance classification," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, 2016, pp. 1163–1168. [Online]. Available: <http://aclweb.org/anthology/N16-1138>
- [107] J. Ma, W. Gao, and K.-F. Wong, "Detect Rumor and Stance Jointly by Neural Multi-task Learning," in *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*. Lyon, France: ACM Press, 2018, pp. 585–593. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3184558.3188729>
- [108] R. Bar-Haim, L. Edelstein, C. Jochim, and N. Slonim, "Improving Claim Stance Classification with Lexical Knowledge Expansion and Context Utilization," in *Proceedings of the 4th Workshop on Argument Mining*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 32–38. [Online]. Available: <http://aclweb.org/anthology/W17-5104>
- [109] O. Toledo-Ronen, R. Bar-Haim, and N. Slonim, "Expert Stance Graphs for Computational Argumentation," in *Proceedings of the Third Workshop on Argument Mining (ArgMining2016)*. Berlin, Germany: Association for Computational Linguistics, 2016, pp. 119–123. [Online]. Available: <http://aclweb.org/anthology/W16-2814>
- [110] M. Nadeem, W. Fang, B. Xu, M. Mohtarami, and J. Glass, "FAKTA: An Automatic End-to-End Fact Checking System," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*. Minneapolis, Minnesota: Association for

- Computational Linguistics, Jun. 2019, pp. 78–83. [Online]. Available: <https://www.aclweb.org/anthology/N19-4014>
- [111] T. Thonet, G. Cabanac, M. Boughanem, and K. Pinel-Sauvagnat, “VODUM: A Topic Model Unifying Viewpoint, Topic and Opinion Discovery,” in *Advances in Information Retrieval*, ser. Lecture Notes in Computer Science, N. Ferro, F. Crestani, M.-F. Moens, J. Mothe, F. Silvestri, G. M. Di Nunzio, C. Hauff, and G. Silvello, Eds. Cham: Springer International Publishing, 2016, pp. 533–545.
- [112] M. Quraishi, P. Fafalios, and E. Herder, “Viewpoint Discovery and Understanding in Social Networks,” in *Proceedings of the 10th ACM Conference on Web Science*, 2018, pp. 47–56, arXiv: 1810.11047. [Online]. Available: <http://arxiv.org/abs/1810.11047>
- [113] L. Graves, “Understanding the Promise and Limits of Automated Fact-Checking,” Reuters Institute, University of Oxford, Factsheet, Feb. 2018.
- [114] James Thorne and Andreas Vlachos, “Adversarial attacks against Fact Extraction and VERification,” *CoRR*, vol. abs/1903.05543, 2019. [Online]. Available: <http://arxiv.org/abs/1903.05543>
- [115] A. Vlachos and S. Riedel, “Identification and Verification of Simple Claims about Statistical Properties,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 2596–2601. [Online]. Available: <https://www.aclweb.org/anthology/D15-1312>
- [116] J. Thorne and A. Vlachos, “An Extensible Framework for Verification of Numerical Claims,” in *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*. Valencia, Spain: Association for Computational Linguistics, Apr. 2017, pp. 37–40. [Online]. Available: <https://www.aclweb.org/anthology/E17-3010>
- [117] Y. Wu, P. K. Agarwal, C. Li, J. Yang, and C. Yu, “Toward computational fact-checking,” *Proceedings of the VLDB Endowment*, vol. 7, no. 7, pp. 589–600, Mar. 2014. [Online]. Available: <https://dl.acm.org/doi/10.14778/2732286.2732295>
- [118] —, “Computational Fact Checking through Query Perturbations,” *ACM Transactions on Database Systems*, vol. 42, no. 1, pp. 1–41, Mar. 2017. [Online]. Available: <https://dl.acm.org/doi/10.1145/2996453>
- [119] S. Cohen, C. Li, J. Yang, and C. Yu, “Computational Journalism: A Call to Arms to Database Researchers,” in *5th Biennial Conference on Innovative Data Systems Research (CIDR '11)*, Asilomar, California, USA, 2011, pp. 148–151.
- [120] J. M. González Pinto and W.-T. Balke, “Offering Answers for Claim-Based Queries: A New Challenge for Digital Libraries,” in *Digital Libraries: Data, Information, and Knowledge for Digital Lives*, ser. Lecture Notes in Computer Science, S. Choemprayong, F. Crestani, and S. J. Cunningham, Eds. Cham: Springer International Publishing, 2017, pp. 3–13.
- [121] J. R. Searle, *Speech acts. An essay in the philosophy of language*. Cambridge: Cambridge University Press, 1969.
- [122] M. Lippi and P. Torroni, “MARGOT: A web server for argumentation mining,” *Expert Systems with Applications*, vol. 65, pp. 292–303, Dec. 2016. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0957417416304493>
- [123] Georgia M. Green, *Pragmatics and Natural Language Understanding*, 2nd ed. Hilldale, NJ: Lawrence Erlbaum Associates, 1996.
- [124] Mevan Babakar and Will Moy, “The State of Automated Factchecking - How to make factchecking dramatically more effective with technology we have now,” Full Fact, Tech. Rep., Aug. 2016. [Online]. Available: [https://fullfact.org/media/uploads/full\\_fact-the\\_state\\_of\\_automated\\_factchecking\\_aug\\_2016.pdf](https://fullfact.org/media/uploads/full_fact-the_state_of_automated_factchecking_aug_2016.pdf)
- [125] O. B. Rekdal, “Academic urban legends,” *Social Studies of Science*, vol. 44, no. 4, pp. 638–654, Aug. 2014. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0306312714535679>
- [126] E. Newell, D. Margolin, and D. Ruths, “An Attribution Relations Corpus for Political News,” in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA), 2018, p. 8.
- [127] W. Shen, J. Wang, and J. Han, “Entity Linking with a Knowledge Base: Issues, Techniques, and Solutions,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 2, pp. 443–460, Feb. 2015, conference Name: IEEE Transactions on Knowledge and Data Engineering.
- [128] L. Carstens and F. Toni, “Towards relation based Argumentation Mining,” in *Proceedings of the 2nd Workshop on Argumentation Mining*. Denver, CO: Association for Computational Linguistics, 2015, pp. 29–34. [Online]. Available: <http://aclweb.org/anthology/W15-0504>
- [129] A. Gangemi, “A Comparison of Knowledge Extraction Tools for the Semantic Web,” in *The Semantic Web: Semantics and Big Data*, ser. Lecture Notes in Computer Science, P. Cimini, O. Corcho, V. Presutti, L. Hollink, and S. Rudolph, Eds. Berlin, Heidelberg: Springer, 2013, pp. 351–366.
- [130] P. Ristoski and H. Paulheim, “Semantic Web in data mining and knowledge discovery: A comprehensive survey,” *Journal of Web Semantics*, vol. 36, pp. 1–22, Jan. 2016. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1570826816000020>
- [131] D. Maynard, K. Bontcheva, and I. Augenstein, *Natural Language Processing for the Semantic Web*. Morgan & Claypool Publishers, 2016.
- [132] G. Weikum, J. Hoffart, and F. Suchanek, “Knowledge Harvesting: Achievements and Challenges,” in *Computing and Software Science: State of the Art and Perspectives*, ser. Lecture Notes in Computer Science, B. Steffen and G. Woeginger, Eds. Cham: Springer International Publishing, 2019, pp. 217–235. [Online]. Available: [https://doi.org/10.1007/978-3-319-91908-9\\_13](https://doi.org/10.1007/978-3-319-91908-9_13)
- [133] S. Kumar, “A Survey of Deep Learning Methods for Relation Extraction,” *arXiv:1705.03645 [cs]*, May 2017, arXiv: 1705.03645. [Online]. Available: <http://arxiv.org/abs/1705.03645>
- [134] W. Wong, W. Liu, and M. Bennamoun, “Ontology learning from text: A look back and into the future,” *ACM Computing Surveys*, vol. 44, no. 4, pp. 20:1–20:36, Sep. 2012. [Online]. Available: <https://doi.org/10.1145/2333112.2333115>
- [135] F. Atefeh and W. Khreich, “A Survey of Techniques for Event Detection in Twitter,” *Computational Intelligence*, vol. 31, no. 1, pp. 132–164, Feb. 2015. [Online]. Available: <https://doi.org/10.1111/coin.12017>

- [136] V. Uren, P. Cimiano, J. Iria, S. Handschuh, M. Vargas-Vera, E. Motta, and F. Ciravegna, "Semantic annotation for knowledge management: Requirements and a survey of the state of the art," *Journal of Web Semantics*, vol. 4, no. 1, pp. 14–28, Jan. 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1570826805000338>
- [137] H. Lin and V. Ng, "Abstractive Summarization: A Survey of the State of the Art," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 9815–9822, Jul. 2019, number: 01. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/5056>
- [138] B. Sriram, D. Fuhry, E. Demir, H. Ferhatosmanoglu, and M. Demirbas, "Short text classification in twitter to improve information filtering," in *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, ser. SIGIR '10. New York, NY, USA: Association for Computing Machinery, Jul. 2010, pp. 841–842. [Online]. Available: <https://doi.org/10.1145/1835449.1835643>
- [139] V. Beretta, S. Harispe, K. Boland, L. Lo Seen, K. Todorov, and A. Tchechmedjiev, "Can Knowledge Graph Embeddings Tell Us What Fact-checked Claims Are About?" in *Proceedings of the First Workshop on Insights from Negative Results in NLP*. Online: Association for Computational Linguistics, Nov. 2020, pp. 71–75. [Online]. Available: <https://www.aclweb.org/anthology/2020.insights-1.11>
- [140] M. J. Paul, C. Zhai, and R. Girju, "Summarizing contrastive viewpoints in opinionated text," in *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, ser. EMNLP '10. USA: Association for Computational Linguistics, Oct. 2010, pp. 66–76.
- [141] T. Thonet, G. Cabanac, M. Boughanem, and K. Pinel-Sauvagnat, "Users Are Known by the Company They Keep: Topic Models for Viewpoint Discovery in Social Networks," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, ser. CIKM '17. New York, NY, USA: Association for Computing Machinery, Nov. 2017, pp. 87–96. [Online]. Available: <https://doi.org/10.1145/3132847.3132897>
- [142] A. Trabelsi and O. R. Zaiane, "Unsupervised Model for Topic Viewpoint Discovery in Online Debates Leveraging Author Interactions," in *Proceedings of the Twelfth International AAAI Conference on Web and Social Media (ICWSM 2018)*. Association for the Advancement of Artificial Intelligence, 2018, pp. 425–433.
- [143] P. Barberá, "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data," *Political Analysis*, vol. 23, no. 1, pp. 76–91, 2015, publisher: Cambridge University Press.
- [144] B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," *Foundations and Trends in Information Retrieval*, vol. 2, no. 1-2, pp. 1–135, Jan. 2008. [Online]. Available: <https://doi.org/10.1561/15000000011>
- [145] S. Mohammad, S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry, "SemEval-2016 Task 6: Detecting Stance in Tweets," in *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. San Diego, California: Association for Computational Linguistics, Jun. 2016, pp. 31–41. [Online]. Available: <https://www.aclweb.org/anthology/S16-1003>
- [146] J. Du, R. Xu, Y. He, and L. Gui, "Stance classification with target-specific neural attention networks," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, ser. IJCAI'17. Melbourne, Australia: AAAI Press, Aug. 2017, pp. 3988–3994.
- [147] I. Augenstein, T. Rocktäschel, A. Vlachos, and K. Bontcheva, "Stance Detection with Bidirectional Conditional Encoding," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 876–885. [Online]. Available: <https://www.aclweb.org/anthology/D16-1084>
- [148] M. Lai, D. I. Hernández Farías, V. Patti, and P. Rosso, "Friends and Enemies of Clinton and Trump: Using Context for Detecting Stance in Political Tweets," in *Advances in Computational Intelligence*, ser. Lecture Notes in Computer Science, G. Sidorov and O. Herrera-Alcántara, Eds. Cham: Springer International Publishing, 2017, pp. 155–168.
- [149] Q. Sun, Z. Wang, Q. Zhu, and G. Zhou, "Stance Detection with Hierarchical Attention Network," in *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, 2018, pp. 2399–2409. [Online]. Available: <https://www.aclweb.org/anthology/C18-1203>
- [150] J. Ebrahimi, D. Dou, and D. Lowd, "Weakly Supervised Tweet Stance Classification by Relational Bootstrapping," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 1012–1017. [Online]. Available: <https://www.aclweb.org/anthology/D16-1105>
- [151] C. Xu, C. Paris, S. Nepal, and R. Sparks, "Cross-Target Stance Classification with Self-Attention Networks," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 778–783. [Online]. Available: <https://www.aclweb.org/anthology/P18-2123>
- [152] M. A. Walker, P. Anand, R. Abbott, and R. Grant, "Stance classification using dialogic properties of persuasion," in *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, ser. NAACL HLT '12. USA: Association for Computational Linguistics, Jun. 2012, pp. 592–596.
- [153] D. Sridhar, J. Foulds, B. Huang, L. Getoor, and M. Walker, "Joint Models of Disagreement and Stance in Online Debate," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Beijing, China: Association for Computational Linguistics, 2015, pp. 116–125. [Online]. Available: <http://aclweb.org/anthology/P15-1012>
- [154] R. Bar-Haim, I. Bhattacharya, F. Dinuzzo, A. Saha, and N. Slonim, "Stance Classification of Context-Dependent Claims," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. Valencia, Spain: Association for Computational Linguistics, Apr. 2017, pp. 251–261. [Online]. Available: <https://www.aclweb.org/anthology/E17-1024>

- [155] C. Guggilla, T. Miller, and I. Gurevych, "CNN- and LSTM-based Claim Classification in Online User Comments," in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. Osaka, Japan: The COLING 2016 Organizing Committee, Dec. 2016, pp. 2740–2751. [Online]. Available: <https://www.aclweb.org/anthology/C16-1258>
- [156] D. Pomerleau and D. Rao, "Fake News Challenge Stage 1 (FNC-I): Stance Detection," 2017. [Online]. Available: <http://www.fakenewschallenge.org/>
- [157] D. Küçük and F. Can, "Stance Detection: A Survey," *ACM Computing Surveys*, vol. 53, no. 1, pp. 12:1–12:37, Feb. 2020. [Online]. Available: <https://doi.org/10.1145/3369026>
- [158] S. Ghosh, P. Singhanian, S. Singh, K. Rudra, and S. Ghosh, "Stance Detection in Web and Social Media: A Comparative Study," *arXiv:2007.05976 [cs]*, vol. 11696, pp. 75–87, 2019, arXiv: 2007.05976. [Online]. Available: <http://arxiv.org/abs/2007.05976>
- [159] I. Sen, F. Flöck, and C. Wagner, "On the Reliability and Validity of Detecting Approval of Political Actors in Tweets," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 1413–1426. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-main.110>
- [160] J. Lawrence and C. Reed, "Argument Mining: A Survey," *Computational Linguistics*, vol. 45, no. 4, pp. 765–818, Dec. 2019. [Online]. Available: <https://www.aclweb.org/anthology/J19-4006>
- [161] A. Peldszus and M. Stede, "From Argument Diagrams to Argumentation Mining in Texts: A Survey," *International Journal of Cognitive Informatics and Natural Intelligence*, vol. 7, no. 1, pp. 1–31, Jan. 2013. [Online]. Available: <https://doi.org/10.4018/jcini.2013010101>
- [162] M. Lippi and P. Torroni, "Argumentation Mining: State of the Art and Emerging Trends," *ACM Transactions on Internet Technology*, vol. 16, no. 2, pp. 10:1–10:25, Mar. 2016. [Online]. Available: <https://doi.org/10.1145/2850417>
- [163] E. Cabrio and S. Villata, "Five years of argument mining: a data-driven analysis," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, ser. IJ-CAI'18. Stockholm, Sweden: AAAI Press, Jul. 2018, pp. 5427–5433.
- [164] N. Hassan, G. Zhang, F. Arslan, J. Caraballo, D. Jimenez, S. Gawsane, S. Hasan, M. Joseph, A. Kulkarni, A. K. Nayak, V. Sable, C. Li, and M. Tremayne, "ClaimBuster: the first-ever end-to-end fact-checking system," *Proceedings of the VLDB Endowment*, vol. 10, no. 12, pp. 1945–1948, Aug. 2017. [Online]. Available: <https://doi.org/10.14778/3137765.3137815>
- [165] C. Hansen, C. Hansen, S. Alstrup, J. Grue Simonsen, and C. Lioma, "Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking," in *Companion Proceedings of The 2019 World Wide Web Conference*, ser. WWW '19. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 994–1000. [Online]. Available: <https://doi.org/10.1145/3308560.3316736>
- [166] P. Biyani, S. Bhatia, C. Caragea, and P. Mitra, "Using non-lexical features for identifying factual and opinionative threads in online forums," *Knowledge-Based Systems*, vol. 69, pp. 170–178, Oct. 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0950705114001786>
- [167] J. Wiebe and E. Riloff, "Creating subjective and objective sentence classifiers from unannotated texts," in *Proceedings of the 6th international conference on Computational Linguistics and Intelligent Text Processing*, ser. CICLing'05. Berlin, Heidelberg: Springer-Verlag, Feb. 2005, pp. 486–497. [Online]. Available: [https://doi.org/10.1007/978-3-540-30586-6\\_53](https://doi.org/10.1007/978-3-540-30586-6_53)
- [168] H. Yu and V. Hatzivassiloglou, "Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences," in *Proceedings of the 2003 conference on Empirical methods in natural language processing*, 2003, pp. 129–136.
- [169] H. Rashkin, E. Choi, J. Y. Jang, S. Volkova, and Y. Choi, "Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics, 2017, pp. 2931–2937. [Online]. Available: <http://aclweb.org/anthology/D17-1317>
- [170] O. Sevgili, A. Shelmanov, M. Arkhipov, A. Panchenko, and C. Biemann, "Neural Entity Linking: A Survey of Models based on Deep Learning," *arXiv:2006.00575 [cs]*, May 2020, arXiv: 2006.00575. [Online]. Available: <http://arxiv.org/abs/2006.00575>
- [171] D. Ferrucci and A. Lally, "UIMA: an architectural approach to unstructured information processing in the corporate research environment," *Natural Language Engineering*, vol. 10, no. 3-4, pp. 327–348, Sep. 2004. [Online]. Available: <https://doi.org/10.1017/S1351324904003523>
- [172] K. Starbird, A. Arif, T. Wilson, K. V. Koevering, K. Yefimova, and D. Scarnecchia, "Ecosystem or Echo-System? Exploring Content Sharing across Alternative Media Domains," in *Proceedings of the International AAAI Conference on Web and Social Media*. Association for the Advancement of Artificial Intelligence, 2018.
- [173] C. Newell, T. Cowlshaw, and D. Man, "Quote Extraction and Analysis for News," in *Data Science, Journalism & Media @KDD 2018 International Conference on Knowledge Discovery and Data Mining*, 2018.
- [174] V. Niculae, C. Suen, J. Zhang, C. Danescu-Niculescu-Mizil, and J. Leskovec, "QUOTUS: The Structure of Political Media Coverage as Revealed by Quoting Patterns," in *Proceedings of the 24th International Conference on World Wide Web*, ser. WWW '15. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, May 2015, pp. 798–808. [Online]. Available: <https://doi.org/10.1145/2736277.2741688>
- [175] M. Hasan, M. A. Orgun, and R. Schwitter, "A survey on real-time event detection from the Twitter data stream," *Journal of Information Science*, vol. 44, no. 4, pp. 443–463, Aug. 2018. [Online]. Available: <https://doi.org/10.1177/0165551517698564>
- [176] X. Chen, S. Wang, Y. Tang, and T. Hao, "A bibliometric analysis of event detection in social media," *Online Information Review*, vol. 43, no. 1, pp. 29–52, Jan. 2019, publisher: Emerald Publishing Limited. [Online]. Available: <https://doi.org/10.1108/OIR-03-2018-0068>

- [177] A. Salway, P. Meurer, K. Hofland, and Ø. Reigem, "Quote Extraction and Attribution from Norwegian Newspapers," in *Proceedings of the 21st Nordic Conference on Computational Linguistics*. Gothenburg, Sweden: Association for Computational Linguistics, May 2017, pp. 293–297. [Online]. Available: <https://www.aclweb.org/anthology/W17-0241>
- [178] R. Krestel, S. Bergler, and R. Witte, "Minding the Source: Automatic Tagging of Reported Speech in Newspaper Articles," in *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*. Marrakech, Morocco: European Language Resources Association (ELRA), May 2008. [Online]. Available: [http://www.lrec-conf.org/proceedings/lrec2008/pdf/718\\_paper.pdf](http://www.lrec-conf.org/proceedings/lrec2008/pdf/718_paper.pdf)
- [179] S. Pareti, "Attribution: A Computational Approach," PhD Thesis, Edinburgh, 2015.
- [180] C. Scheible, R. Klinger, and S. Padó, "Model Architectures for Quotation Detection," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, 2016, pp. 1736–1745. [Online]. Available: <http://aclweb.org/anthology/P16-1164>
- [181] T. O'Keefe, S. Pareti, J. R. Curran, I. Koprinksa, and M. Honnibal, "A Sequence Labelling Approach to Quote Attribution," in *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Jeju Island, Korea: Association for Computational Linguistics, Jul. 2012, pp. 790–799. [Online]. Available: <https://www.aclweb.org/anthology/D12-1072>
- [182] D. Pavlo, T. Piccardi, and R. West, "Quootstrap: Scalable Unsupervised Extraction of Quotation-Speaker Pairs from Large News Corpora via Bootstrapping," in *Proceedings of the Twelfth International AAAI Conference on Web and Social Media (ICWSM 2018)*, ser. ICWSM 2018, 2018, pp. 231–240, arXiv: 1804.02525. [Online]. Available: <http://arxiv.org/abs/1804.02525>
- [183] S. Jiang, S. Baumgartner, A. Ittycheriah, and C. Yu, "Factoring Fact-Checks: Structured Information Extraction from Fact-Checking Articles," in *Proceedings of The Web Conference 2020*. Taipei Taiwan: ACM, Apr. 2020, pp. 1592–1603. [Online]. Available: <https://dl.acm.org/doi/10.1145/3366423.3380231>
- [184] K. Sharma, F. Qian, H. Jiang, N. Ruchansky, M. Zhang, and Y. Liu, "Combating Fake News: A Survey on Identification and Mitigation Techniques," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 3, pp. 21:1–21:42, Apr. 2019. [Online]. Available: <https://doi.org/10.1145/3305260>
- [185] S. Cazalens, P. Lamarre, J. Leblay, I. Manolescu, and X. Tannier, "A Content Management Perspective on Fact-Checking," in *Companion Proceedings of the The Web Conference 2018*, ser. WWW '18. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, Apr. 2018, pp. 565–574. [Online]. Available: <https://doi.org/10.1145/3184558.3188727>
- [186] L. A. Galárraga, C. Teflioudi, K. Hose, and F. Suchanek, "AMIE: association rule mining under incomplete evidence in ontological knowledge bases," in *Proceedings of the 22nd international conference on World Wide Web*, ser. WWW '13. New York, NY, USA: Association for Computing Machinery, May 2013, pp. 413–422. [Online]. Available: <https://doi.org/10.1145/2488388.2488425>
- [187] W. Y. Wang, "Liar, liar pants on fire: A new benchmark dataset for fake news detection," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Short Papers)*. Vancouver, Canada: Association for Computational Linguistics, 2017, pp. 422–426. [Online]. Available: <https://www.aclweb.org/anthology/P17-2067.pdf>
- [188] A. Vlachos and S. Riedel, "Fact Checking: Task definition and dataset construction," in *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*. Baltimore, MD, USA: Association for Computational Linguistics, Jun. 2014, pp. 18–22. [Online]. Available: <https://www.aclweb.org/anthology/W14-2508>
- [189] A. Barrón-Cedeño, T. Elsayed, R. Suwaileh, L. Márquez, P. Atanasova, W. Zaghouani, and S. Kyuchukov, "Overview of the CLEF-2018 CheckThat! Lab on Automatic Identification and Verification of Political Claims. Task 2: Factuality," in *Working Notes of CLEF 2018 - Conference and Labs of the Evaluation Forum*, ser. CLEF 2018 Working Notes, 2018, p. 13.
- [190] G. Gorrell, E. Kochkina, M. Liakata, A. Aker, A. Zubiaga, K. Bontcheva, and L. Derczynski, "RumourEval 2019: Determining Rumour Veracity and Support for Rumours," in *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2019)*. Minneapolis, Minnesota, USA: Association for Computational Linguistics, 2019, pp. 845–854.
- [191] T. Mihaylova, P. Nakov, L. Marquez, A. Barron-Cedeno, M. Mohtarami, G. Karadzhov, and J. Glass, "Fact Checking in Community Forums," in *arXiv:1803.03178 [cs]*, ser. AAAI 2018, 2018, pp. 5309–5316, arXiv: 1803.03178. [Online]. Available: <http://arxiv.org/abs/1803.03178>
- [192] J. Thorne, A. Vlachos, C. Christodoulopoulos, and A. Mittal, "FEVER: a Large-scale Dataset for Fact Extraction and VERification," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, Jun. 2018, pp. 809–819. [Online]. Available: <https://www.aclweb.org/anthology/N18-1074>
- [193] X. Yin, J. Han, and P. S. Yu, "Truth Discovery with Multiple Conflicting Information Providers on the Web," *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 6, pp. 796–808, Jun. 2008. [Online]. Available: <https://doi.org/10.1109/TKDE.2007.190745>
- [194] X. Liu, P. He, W. Chen, and J. Gao, "Multi-Task Deep Neural Networks for Natural Language Understanding," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 4487–4496. [Online]. Available: <https://www.aclweb.org/anthology/P19-1441>
- [195] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, "ERNIE: Enhanced Language Representation with Informative Entities," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 1441–1451. [Online]. Available: <https://www.aclweb.org/anthology/P19-1139>
- [196] B. Pereira Nunes, A. Mera, M. A. Casanova, B. Fetahu, L. A. P. Paes Leme, and S. Dietze, "Complex Matching of RDF Datatype Properties," in *Database and Expert Systems Appli-*

- cations, ser. Lecture Notes in Computer Science, H. Decker, L. Lhotská, S. Link, J. Basl, and A. M. Tjoa, Eds. Berlin, Heidelberg: Springer, 2013, pp. 195–208.
- [197] J. Gracia and E. Mena, “Web-Based Measure of Semantic Relatedness,” in *Proceedings of the 9th international conference on Web Information Systems Engineering*, ser. WISE ’08. Berlin, Heidelberg: Springer-Verlag, Sep. 2008, pp. 136–150. [Online]. Available: [https://doi.org/10.1007/978-3-540-85481-4\\_12](https://doi.org/10.1007/978-3-540-85481-4_12)
- [198] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia, “SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Vancouver, Canada: Association for Computational Linguistics, Aug. 2017, pp. 1–14. [Online]. Available: <https://www.aclweb.org/anthology/S17-2001>
- [199] E. Agirre, C. Banea, D. Cer, M. Diab, A. Gonzalez-Agirre, R. Mihalcea, G. Rigau, and J. Wiebe, “SemEval-2016 Task 1: Semantic Textual Similarity, Monolingual and Cross-Lingual Evaluation,” in *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. San Diego, California: Association for Computational Linguistics, Jun. 2016, pp. 497–511. [Online]. Available: <https://www.aclweb.org/anthology/S16-1081>
- [200] L. Dumani and R. Schenkel, “A Systematic Comparison of Methods for Finding Good Premises for Claims,” in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR’19. New York, NY, USA: Association for Computing Machinery, Jul. 2019, pp. 957–960. [Online]. Available: <https://doi.org/10.1145/3331184.3331282>
- [201] I. Dagan, O. Glickman, and B. Magnini, “The PASCAL Recognising Textual Entailment Challenge,” in *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Textual Entailment*, ser. Lecture Notes in Computer Science, J. Quiñero-Candela, I. Dagan, B. Magnini, and F. d’Alché Buc, Eds. Berlin, Heidelberg: Springer, 2006, pp. 177–190.
- [202] Z. Zhang, A. L. Gentile, and F. Ciravegna, “Recent advances in methods of lexical semantic relatedness – a survey,” *Natural Language Engineering*, vol. 19, no. 4, pp. 411–479, Oct. 2013, publisher: Cambridge University Press.
- [203] M. A. Hadj Taieb, T. Zesch, and M. Ben Aouicha, “A survey of semantic relatedness evaluation datasets and procedures,” *Artificial Intelligence Review*, vol. 53, no. 6, pp. 4407–4448, Aug. 2020. [Online]. Available: <https://doi.org/10.1007/s10462-019-09796-3>
- [204] M. Pagliardini, P. Gupta, and M. Jaggi, “Unsupervised Learning of Sentence Embeddings Using Compositional n-Gram Features,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, Jun. 2018, pp. 528–540. [Online]. Available: <https://www.aclweb.org/anthology/N18-1049>
- [205] L. Liu, X. Ren, J. Shang, X. Gu, J. Peng, and J. Han, “Efficient Contextualized Representation: Language Model Pruning for Sequence Labeling,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, Oct. 2018, pp. 1215–1225. [Online]. Available: <https://www.aclweb.org/anthology/D18-1153>
- [206] J. Howard and S. Ruder, “Universal Language Model Fine-tuning for Text Classification,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 328–339. [Online]. Available: <https://www.aclweb.org/anthology/P18-1031>
- [207] A. Akbik, T. Bergmann, and R. Vollgraf, “Pooled Contextualized Embeddings for Named Entity Recognition,” 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)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 724–728. [Online]. Available: <https://www.aclweb.org/anthology/N19-1078>
- [208] A. Akbik, D. Blythe, and R. Vollgraf, “Contextual String Embeddings for Sequence Labeling,” in *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA: Association for Computational Linguistics, Aug. 2018, pp. 1638–1649. [Online]. Available: <https://www.aclweb.org/anthology/C18-1139>
- [209] J. Opitz and A. Frank, “Dissecting Content and Context in Argumentative Relation Analysis,” in *Proceedings of the 6th Workshop on Argument Mining*. Florence, Italy: Association for Computational Linguistics, 2019, pp. 25–34. [Online]. Available: <https://www.aclweb.org/anthology/W19-4503>
- [210] H. Nguyen and D. Litman, “Context-aware Argumentative Relation Mining,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, 2016, pp. 1127–1137. [Online]. Available: <http://aclweb.org/anthology/P16-1107>
- [211] M. Nentwig, M. Hartung, A.-C. Ngonga Ngomo, and E. Rahm, “A survey of current Link Discovery frameworks,” *Semantic Web*, vol. 8, no. 3, pp. 419–436, Dec. 2016. [Online]. Available: <https://www.medra.org/serve/aliasResolver?alias=iospress&doi=10.3233/SW-150210>
- [212] A. Jentzsch, R. Isele, and C. Bizer, “Silk – Generating RDF Links while publishing or consuming Linked Data,” in *9th International Semantic Web Conference (ISWC’10)*, 2010.
- [213] A.-C. N. Ngomo and S. Auer, “LIMES: a time-efficient approach for large-scale link discovery on the web of data,” in *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence - Volume Volume Three*, ser. IJCAI’11. Barcelona, Catalonia, Spain: AAAI Press, Jul. 2011, pp. 2312–2317.
- [214] A. Roy, P. Fafalios, A. Ekbal, X. Zhu, and S. Dietze, “Exploiting stance hierarchies for cost-sensitive stance detection of Web documents,” *arXiv:2007.15121 [cs]*, Jul. 2020, arXiv: 2007.15121. [Online]. Available: <http://arxiv.org/abs/2007.15121>
- [215] B. Zapolko, J. Schaible, P. Mayr, and B. Mathiak, “TheSoz: A SKOS representation of the thesaurus for the social sciences,” *Semantic Web*, vol. 4, no. 3, pp. 257–263, Jul. 2013.