Modeling and Contextualizing Claims

Katarina Boland[1], Pavlos Fafalios[2], Andon Tchechmedjiev[3], Konstantin Todorov[4], Stefan Dietze[1,5]

[1] GESIS - Leibniz Institute for the Social Sciences, Germany
[2] Institute of Computer Science, FORTH-ICS, Greece
[3] LGI2P, IMT Mines-Ales, France
[4] LIRMM / University of Montpellier / CNRS, France
[5] Heinrich-Heine-University Düsseldorf, Germany
katarina.boland@gesis.org, fafalios@ics.forth.gr,
andon.tchechmedjiev@mines-ales.fr, konstantin.todorov@lirmm.fr
stefan.dietze@gesis.org

Abstract. Understanding societal debates on the Web and how they are impacted by the spread of biased narratives and falsehoods are becoming increasingly important issues. The notion of a *claim* is central in a number of related studies into fake-news propagation or computational fact-checking. While the understanding of this notion varies from one field to another, there are few studies that have focused on the conceptual modeling of claims and their context. We attempt to contribute to this area by proposing a novel conceptual model for claims and related notions, such as attitudes, reviews and annotations, that aims to take into consideration the claims inherent complexity, distinguishing between their meaning, linguistic representation and context. We provide an example of an implementation of this model by using established vocabularies, such as schema.org, Open Annotation and PROV-O, and discuss the challenges related to this work. [1]

Keywords: Claims; Conceptual Modeling; Claim Context; Societal Debates; Fact-checking

1 Introduction

The spread of biased narratives and falsehoods on the Web and the analyses of online discourse have become increasingly important issues [1, 13] that led to a wide range of interdisciplinary research involving a variety of scientific disciplines. Such works include investigations, for instance, into the spreading pattern of false claims on Twitter [13], or the development of computational methods, such as pipelines for detecting the stance of claim-relevant Web documents [14], classifying sources of news, such as Web pages, PLDs, users or posts [10], or for fake news detection [12] and automatic fact-checking [4].

Whereas techniques for knowledge graph construction and augmentation often deploy methods strongly related to the aforementioned computational methods related to claims, e.g., when aiming to verify facts from the Web for augmenting knowledge bases [2,15], the notion of a *claim* is fundamentally different from the notion of a *fact* as an atomic assertion in the first-order-logic sense. This is due to the inherent complexity of a claim, where its interpretation usually is strongly dependent on its context, such as its source, timing, or location. Moreover, 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 [3]. Ambiguity also arises with respect to claims involving quotations (*"X reported that Y said Z"*), where often fact-checking results remain vague about what part such a nested claim actually has been validated.

In order to facilitate the advancement of tasks such as claim verification or fact checking, it is crucial to capture the complexity of a claim in a way which enables unambiguous interpretation by both humans and machines. However, both the used terminology and the underlying conceptual models are still strongly diverging in academic literature (Sect. 2) as well as in the conceptual models deployed by fact-checking sites.

Therefore, capturing the meaning of a claim requires both the preservation of 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.

This paper makes the following main contributions: i) a *conceptual model* and corresponding terminology of claims and their constituents and context, grounded in both the scientific state-of-the-art in related fields such as argumentation mining as well as the actual practices of representing and sharing claims on the Web, for instance, as part of fact-checking sites; ii) an *RDF implementation* of the proposed conceptual model that uses W3C standards for data sharing, namely RDFS, and is informed by established vocabularies, such as *schema.org*, *Open Annotation*, and the *PROV* data model, 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 such as *ClaimsKG* [11].

2 Background

While the analysis of claims plays a crucial role for a number of fields, the definition of the very concept of a claim is often left to the intuition of the reader. Existing definitions vary considerably across and also within fields.

According to the Oxford English Dictionary, a claim is a statement or assertion that something is the case, typically without providing evidence or proof.[2] Platforms dedicated to journalistic fact-checking refer to claims as statements supported by (a group of) people or organizations that appear newsworthy, significant and verifiable.[3] An RDFS-based model for such fact-checked claims is introduced in [11].

In *argumentation mining*, claims denominate the conclusion of an argument or the assertion the argument aims to prove [6, 7]. A variety of additional definitions can be found for specific tasks in other fields like information retrieval, e.g. a statement formulating a problem together with a concrete solution [8] or a sentence in a scientific document that relates two entities given in a query [9].

Thus, what is identified as "claim" in a particular work may or may not be called "claim" in another. While it is the belief of a person about a fact that is called "claim" in argumentation mining, it is the fact itself that is coined "claim" in the fact-checking community. Similarly, the belief and opinion about certain consequences are the argumentative "claim", while fact-checking may verify whether the anticipated consequences would indeed follow an action. Statements expressing the position of a person towards a proposition or target are not susceptible to fact-checking (unless the correctness of the quotation is to be verified) but are a prevalent claim type in argumentation mining. Moreover, what is used as premise or evidence in an argument is often selected as check-worthy "claim" by fact-checking sites. Generally, the distinction of argumentative units such as claims and evidence is based on the statements' usage in an argument while fact-checking classifies statements as claims depending primarily on features inherent to the statement itself.

In an effort to reconcile these different understandings of the concept of a "claim", we propose a model considering requirements from various research fields. While in argumentation mining, the meaning of a claim in the context of the current discourse is the significant part, many tasks from the fact-checking community, e.g. those aiming at matching unchecked statements to fact-checked claims [5], focus on the surface form. Thus, going beyond the model introduced in [11], we propose differentiating between the meaning or proposition of a claim and its utterance, representation and context.

3 Conceptual Model

Overview. We distinguish three main components of a claim, represented by three central classes: (1) *claim proposition*, (2) *claim utterance*, and (3) *claim context*. A *claim proposition* is the meaning of a statement or assertion that something is the case. It is usually related to a controversial topic and can be

[2] `https://www.lexico.com/en/definition/claim`

[3] `https://www.truthorfiction.com/about/`
`https://checkyourfact.com/about-us`
`https://www.politifact.com/truth-o-meter/article/2018/feb/12/principles-truth-o-meter-politifacts-methodology-i/`

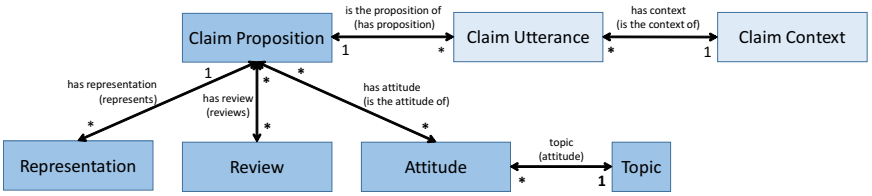

Fig. 1: The main concepts related to a *claim proposition*.

factual or subjective (expressing an opinion). 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 may be expressed in different languages, using different words in the same language, or uttered by different persons and/or in different points in time. On the contrary, a specific claim utterance can be associated to only one proposition, i.e., it has a single meaning. The claim proposition can be represented in different ways, for example, by selecting a representative utterance or through a more formal model. Each claim utterance is related to a specific *claim context*, like the author of the claim or its date. It provides the means to interpret the claim utterance and thus understand its proposition. Below, we provide details and the main properties of each of these three main classes (without repeating the associations among them).

Claim Proposition. A claim proposition reflects the meaning of one or more semantically equivalent claim utterances expressed in different linguistic forms or contexts. A claim proposition is associated with i) one or more preferred *representations*, ii) one or more *reviews*, and iii) one or more *attitudes* (Fig. 1). A *representation* can have the form of free text, e.g., a sentence that best describes the proposition (like the text of one of the corresponding utterances), or be more complex, 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 that not all claims have a review or verdict. For instance, the claim *"the presence of a gun makes a conflict more likely to become violent"* represents a hypothesis and is difficult to be associated with a correctness score (there may be mixed evidence supporting and contradicting it). 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 *pro-Remain* are two different attitudes for the Brexit topic. A claim proposition can be associated with several attitudes for different topics. For example, the claim *"immigrants are taking our jobs"* supports both the *against immigration* attitude (for the Immigration topic) and the *pro-Brexit* attitude (for the Brexit topic).

Claim Utterance. A claim utterance is the act of expressing a claim proposition in a specific natural language and form (like text or speech). Among other things, it may be something said by a politician during an interview, a text within

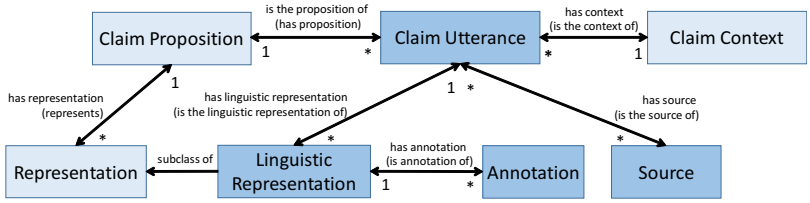

Fig. 2: The main concepts related to a *claim utterance*.

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* in Fig. 1), and ii) one or more *sources* (Fig. 2). 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 existence. For instance, it can be the URL of an interview video, a news article, or a tweet. A linguistic representation can have one or more *annotations* which provide formal linguistic characteristics, like an entity or date mentioned in the text of the claim utterance, the polarity of this 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* tool for the case of entity annotation in text).

Claim Context. The *claim context* provides background information about the claim utterance (Fig. 3). Together with the linguistic representation of the claim utterance, it 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 said), iv) *where* it was said (location the claim was said), and v) *why* it was said (event or activity in the context of which the claim was said). 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 additional context information about the claim utterance (like the topic of the underlying discourse or the medium used for uttering the claim).

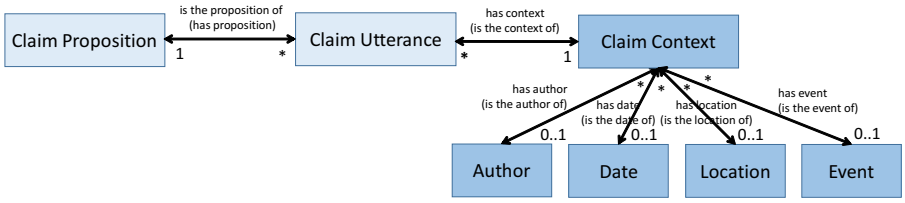

Fig. 3: The main concepts related to a *claim context*.

Instantiation Example. Fig. 4 depicts an instantiation example of the proposed conceptual model. The example shows information for two *claim utterances* (in pink background): i) one said by David Dimbleby during a topical debate in Dover *("We are going to be paying until 2064, apparently")*, and ii) one extracted from a news article of The Independent *("UK will be paying Brexit 'divoce bill' until 2064")*. Both utterances correspond to the same *claim proposition* (in green background) and each one has its own *context* information (in yellow background). 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* (UK). 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* attitude for 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). We can also see the URL of the review article as well as a reference to a PDF file which provides evidence for its correctness. The *context* of each claim utterance provides additional metadata 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*.

4 RDF Implementation

We introduce an RDF/S implementation of the proposed conceptual model using established vocabularies, in particular schema.org,[4] the Open Annotation (OA) Data Model,[5] the Marl Ontology,[6] the NLP Interchange Format (NIF),[7] and the PROV Data Model.[8] The selection of these vocabularies was based on the following three main objectives: i) relying on stable term identifiers and persistent hosting, ii) being supported by a community, iii) being extensible.

Fig. 5 depicts the proposed schema. For representing the main concepts of our conceptual model, we exploit classes and properties of schema.org, a collaborative, community activity with a mission to maintain and promote a common schema for structured data on the Web and beyond. We make use of the class `schema:Claim` (currently under integration in schema.org) to describe a *claim utterance*. According to schema.org, this class represents a specific, factually-oriented claim. For the *claim proposition*, we use the class `schema:Intangible`, a utility class that serves as the umbrella for a number of 'intangible' things. Although this class does not sufficiently reflect the semantics of a claim proposition, it appears to be the most reasonable term for representing a proposition. For the same reason, we use `schema:Intangible` to describe a *claim context*.

[4] `https://schema.org/`
[5] `http://www.openannotation.org/`
[6] `http://www.gsi.dit.upm.es/ontologies/marl/`
[7] `https://persistence.uni-leipzig.org/nlp2rdf/`
[8] `https://www.w3.org/TR/prov-dm/`

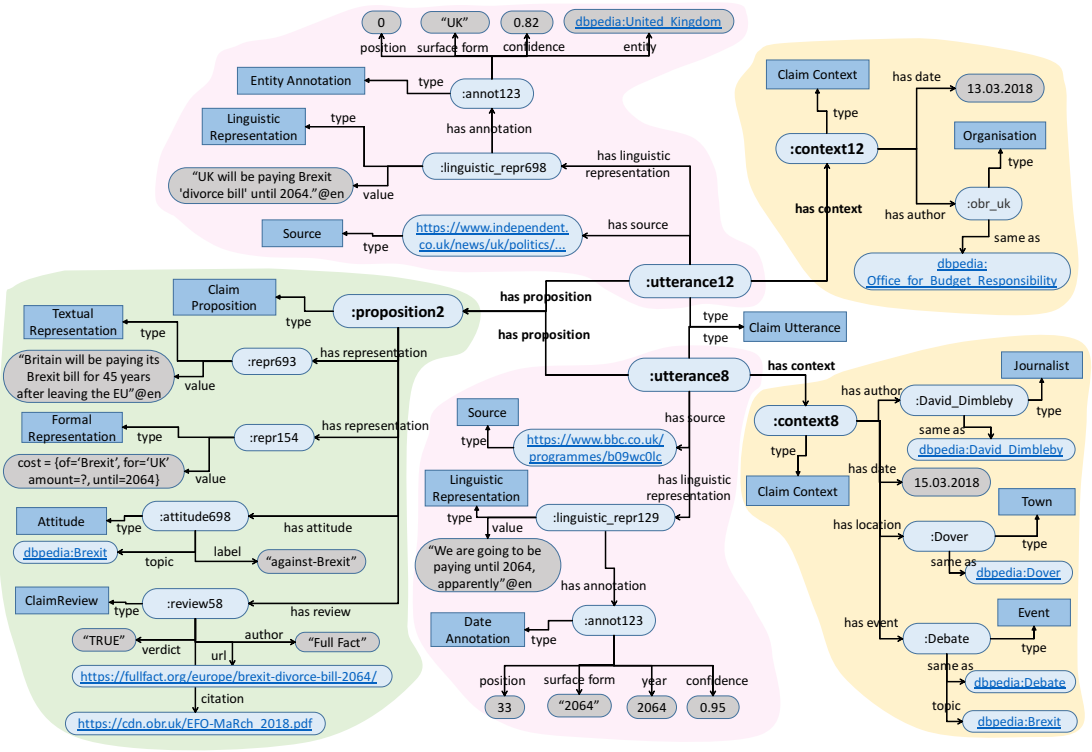

Fig. 4: Instantiation example of the conceptual model.

An alternative solution is to bypass the *claim context* class and directly link an instance of `schema:Claim` to instances of the four classes connected to the claim context (*author*, *date*, *location*, *event*). These four classes are described through corresponding schema.org classes: `schema:Thing` (e.g., a person, an organization, a blog, etc.), `schema:Date`, `schema:Place`, `schema:Event`. For connecting a `schema:Claim` to a `schema:Intangible`, we can use the property `schema:about` or its inverse `schema:subjectOf`.

For representing a *source*, we use the class `schema:CreativeWork` (or one of its sub-classes). Thereby, we take advantage of its properties and can describe additional information about the source, such as headline, language, keywords, publisher, etc. The *linguistic representation* of a claim utterance, as well as the (preferred) *representation* of a claim proposition, can be described through the class `schema:Text` (for textual representations) or `schema:MediaObject` (for image, audio or video representations). For describing *annotations*, we make use of the widely-used OA and NIF data models, while provenance information is represented though the PROV data model. NIF allows us to include detailed information about the outcome of an NLP process on textual representations (like begin/end indexes and confidence scores). The *review* of a claim proposition

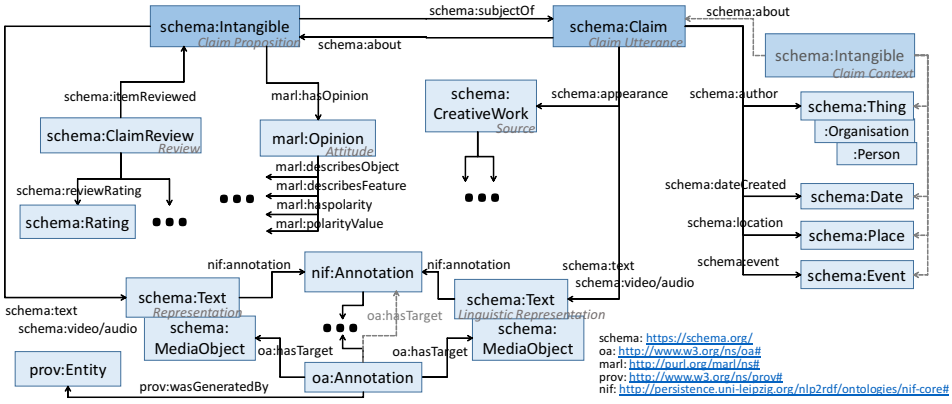

Fig. 5: An RDF implementation of the conceptual model.

is described through the class `schema:ClaimReview`, which in turn is connected to a `schema:Rating` for assigning a rating score about the veracity of the claim proposition. Finally, we exploit the Marl ontology to represent *attitudes*. Marl is a data schema designed to annotate and describe subjective opinions, and provides the attributes that enable to connect opinions with contextual information.

5 Concluding Remarks

We propose a conceptual model and an implementation based on existing RDF vocabularies to represent and contextualize claims and related entities. Our work is meant to advance a shared understanding of claims and related terminology across communities, and the semi-structured representation of claims and their contexts to foster transparency, reproducibility and a joint advancement of research in related fields.

An open challenge is the detection and representation of the inherent relations between claims as well as their relations to other entities or resources. In particular, the semantic relatedness of claims is reflected by the relations between their proposition components. Establishing relations across the model classes, e.g., relating an utterance to a proposition, allows the uncovering of paraphrased claims with identical meaning, while topical relations between claims are of crucial importance to enable retrieval and search of claims.

Other challenges concern information extraction techniques geared towards the extraction of utterances from text or audio together with the attitude towards a particular topic, as well as additional contextual information, such as authors or sources. From a knowledge representation perspective, we emphasize the need for a formal representation of propositions (e.g. by applying dynamic predicate logics), or for extending the *Marl* model in order to represent specifically viewpoints instead of general opinions on objects. This reveals the need for the development of a dedicated ontology for claims representation.

References

1. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. Journal of Economic Perspectives **31**(2), 211–36 (2017)
2. Dong, X., Gabrilovich, E., Heitz, G., Horn, W., Lao, N., Murphy, K., Strohmann, T., Sun, S., Zhang, W.: 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. pp. 601–610. ACM (2014)
3. Graves, D.: Understanding the promise and limits of automated fact-checking (2018)
4. Hassan, N., Adair, B., Hamilton, J.T., Li, C., Tremayne, M., Yang, J., Yu, C.: The quest to automate fact-checking. In: Proceedings of the 2015 Computation+Journalism Symposium (2015)
5. Hassan, N., Zhang, G., Arslan, F., Caraballo, J., Jimenez, D., Gawsane, S., Hasan, S., Joseph, M., Kulkarni, A., Nayak, A.K., et al.: Claimbuster: The first-ever end-to-end fact-checking system. VLDB Endowment **10**(12), 1945–1948 (2017)
6. Levy, R., Bogin, B., Gretz, S., Aharonov, R., Slonim, N.: Towards an argumentative content search engine using weak supervision. In: 27th International Conference on Computational Linguistics. pp. 2066–2081. ACL (Aug 2018)
7. Lippi, M., Torroni, P.: Argument mining from speech: Detecting claims in political debates. In: Thirtieth AAAI Conference on Artificial Intelligence (2016)
8. Pinto, J.M.G., Balke, W.T.: Offering answers for claim-based queries: a new challenge for digital libraries. In: International Conference on Asian Digital Libraries. pp. 3–13. Springer (2017)
9. Pinto, J.M.G., Wawrzinek, J., Balke, W.T.: What drives research efforts? find scientific claims that count! In: 2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 217–226. IEEE (2019)
10. Popat, K., Mukherjee, S., Strötgen, J., Weikum, G.: 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. pp. 1003–1012. International World Wide Web Conferences Steering Committee (2017)
11. Tchechmedjiev, A., Fafalios, P., Boland, K., Gasquet, M., Zloch, M., Zapilko, B., Dietze, S., Todorov, K.: ClaimsKG: A knowledge graph of fact-checked claims. In: International Semantic Web Conference. pp. 309–324. Springer (2019)
12. Tschiatschek, S., Singla, A., Gomez Rodriguez, M., Merchant, A., Krause, A.: Fake news detection in social networks via crowd signals. In: Companion Proceedings of the The Web Conference 2018. pp. 517–524. International World Wide Web Conferences Steering Committee (2018)
13. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. Science **359**(6380), 1146–1151 (2018)
14. Wang, X., Yu, C., Baumgartner, S., Korn, F.: Relevant document discovery for fact-checking articles. In: Companion Proceedings of the The Web Conference 2018. pp. 525–533. International World Wide Web Conferences Steering Committee (2018)
15. Yu, R., Gadiraju, U., Fetahu, B., Lehmberg, O., Ritze, D., Dietze, S.: KnowMore–knowledge base augmentation with structured web markup. Semantic Web pp. 1–22 (2019)