Tracking Perspectives on Event Participants: a Structural Analysis of the Framing of Real-World Events in Co-Referential Corpora

Levi Remijnse, Pia Sommerauer, Antske Fokkens, Piek Vossen
Vrije Universiteit Amsterdam, Computational Linguistics and Text Mining Lab
De Boelelaan 1105, 1081 HV Amsterdam
{l.remijnse, pia.sommerauer, antske.fokkens, p.t.j.m.vossen}@vu.nl

Abstract

In this paper, we present the outcome of a structural linguistic analysis performed on a referentially grounded FrameNet dataset. In this dataset, multiple Dutch events are referenced by multiple co-referential Dutch news texts. Mentions in those documents are annotated with respect to their referential grounding (i.e., links to structured Wikidata), and their conceptual representation (i.e., frames). Provided with each document's temporal reporting distance, we selected documents for two events - the Utrecht shooting and MH17 - and performed an analysis in which we tracked the events' participants over time in both their focalization (number of mentions) and their framing (distribution of frame element labels). This way, we use the carefully collected and annotated data to schematize shifts in focalization and perspectivization of the participants as a result of the constantly developing narrative surrounding the events. This novel type of linguistic research involves reference to the real-world referents and takes into account storytelling in news streams.

Keywords: FrameNet, narratology, referential grounding

1. Introduction

From the moment an event occurs in the world, it generates streams of co-referential news articles.[1] In particular, events that are of interest to society (e.g., mass shootings, music festivals, royal weddings) are reported on by large volumes of documents. Over time, they keep being reported on with regard to their aftermath (e.g., a shooting causes funerals, police investigations, arrests, trials). A narrative develops in which, with every new related topic, different people involved in the event might become the focus. In support of the changing narrative, writers might also change their perspective on those participants. See the examples below, taken from our data, with per sentence the historical distance to the main event, respectively the same day, one day later and the fourth day after:

(1) Vermoedelijk is er daarna **iemand** uit de tram gesprongen.
presumably is there that.after someone out the tram jumped
'Presumably someone jumped out of the tram afterwards.' (day 0) (2019)

(2) De **hoofdverdachte** van de aanslag [...] is de enige verdachte die nog vastzit.
The main.suspect of the attack [...] is the only suspect still that still stuck.sit.
'The main suspect in the attack is the only suspect still in custody.' (day 1) (2019)

(3) **Gökmen Tanis** bekent schietpartij in tram in Utrecht.
Gökmen Tanis confesses shooting.party in tram in Utrecht
'Gökmen Tanis confesses to shooting in tram in Utrecht.' (day 4) (2019)

All three example sentences stem from different documents with different temporal reporting distances, i.e., the temporal distance between the event date and the publication date. The boldfaced mentions co-refer the same entity participating in different events that are part of the same storyline. This entity is perspectivized accordingly. In (1), focus is on his involvement in a mass shooting in a tram; in (2), he is a suspect in custody; and in (3), he is the agent of a confession. Both events in (2) and (3) are related to the event in (1): the mass shooting.

The examples show that, over time, news streams reporting on a single significant event display a continuously developing narrative in which different aspects of the event and its aftermath are topicalized, and the same participant is perspectivized differently.

Suppose we want to perform a structural linguistic analysis in which we get a grip on the way news documents perform storytelling of real-world events. This requires a referentially grounded cor-

[1] We define an event as a specific event instance of a particular event type, e.g., a killing incident happening at a specific location, time, and involving certain participants.

First Workshop on Reference, Framing, and Perspective, pages 1–12

pus with multiple documents reporting on the same event. The documents themselves need to be annotated with information regarding both the referential grounding (which mentions co-refer to which event participant) and conceptual representation (what perspectives do the mention take on in reference to that event participant). Yet, NLP tasks and language resources only cover conceptual representation and reference separately. For instance, the tasks of co-reference resolution (Filatova and Hatzivassiloglou, 2004; Choubey et al., 2018) and entity-linking (Hachey et al., 2013; Getman et al., 2018) contribute to the study of reference. Both Abstract Meaning Representation (Banarescu et al., 2013) as a formal framework and FrameNet (Ruppenhofer et al., 2010; Baker et al., 2003) as a lexicographic paradigm focus on concept description. Yet, for the purpose of a linguistic analysis of referentially grounded data, a dataset should provide information regarding all three components: form, referent and concept.

Therefore, in this paper, we make use of the referentially grounded corpus collected by Remijnse et al. (2022), with documents reporting on real-world events. Focusing on perspectivization as operationalized within frame semantics, this corpus is annotated with FrameNet frames, modeling mentions of participants with semantic types (Remijnse et al., 2022; Postma et al., 2020). This combination of referential grounding and frame semantic information enables us to study the ways in which news streams frame their events over time, as the narrative surrounding those events develops. We take two events commonly known as the Utrecht shooting[2] and MH17[3], focus on the participants of those events, and analyze the ways in which they are framed in our corpus, as a reflection of the developing narrative over time.

We make the following contributions:

- We release a dataset with Dutch reference texts reporting on the Utrecht shooting and MH17. The documents are annotated with links to structured data and FrameNet frames.[4]
- We formulate a model of variation in framing of events that takes into account storytelling and temporal reporting distance.
- Given the dataset and our model, we provide a structural analysis of the linguistic framing of events' participants over time.
- We show patterns in our data of both focalization and perspectivization of participants across events.

This paper is structured as follows. We first discuss related work and background in Section 2. Building on that, we explain our model of computational storytelling in Section 3. We then discuss our selection of data and analysis method in 4. Section 5 provides the results and discussion of the data analysis. We conclude in section 6 and point out limitations in section 7.

2. Background

In this section, we cover related work that the research in this paper is built on, namely referentially grounded corpora (2.1), work in perspectivization (2.2) and narratology (2.3).

2.1. Referentially Grounded Corpora

If the aim is to perform a linguistic analysis that incorporates referential grounding of events, then this grounding affects the very first step of data collection: a corpus needs to exhibit a large variety of documents referencing the same events. Much work in corpus building follows a **text-to-data** method, i.e., starting from text to derive annotation sentence-by-sentence (e.g., OntoNotes (Pradhan et al., 2007), ECB (Bejan and Harabagiu, 2010) ECB+ (Cybulska and Vossen, 2014) and AEC2005 (Peng et al., 2016)). This work is evaluated as labour intensive and time-consuming, thus resulting in small numbers of annotated texts with low intra-document co-reference (10 mentions on average) and low cross-document co-reference (Vossen et al., 2018). Moreover, concrete links between mentions and structured data are absent.

In order to efficiently aggregate multiple co-referential texts per event, Vossen et al. (2020) reversed the process by developing **data-to-text** based software called MWEP, which takes a Wikidata identifier denoting an event type as input, and returns structured Wikidata concerning all events that are grouped under this event type in the Wikidata knowledge base (Vrandečić and Krötzsch, 2014). Per event, the structured output is accompanied with referential news texts crawled from corresponding Wikipedia pages (Simpson et al., 2010). By starting from structured data in aggregating documents, those documents are by default grouped under their event and annotation merely serves as validation.

For the purpose of our research, since we need both structured data and unstructured data grouped under events that have entries in Wikidata, we consider MWEP to be the most useful corpus compilation software.

[2] https://www.wikidata.org/wiki/Q62090804

[3] https://www.wikidata.org/wiki/Q17374096

[4] Our data is freely available at http://dutchframenet.nl/data-releases/

2.2. Perspectivization

The phenomenon of perspectivization (or "framing") has been analyzed in many different fields, e.g, cognitive linguistics (Horst, 2020; Ziem et al., 2018), political studies (Druckman, 2001; Iyengar, 1994; Entman, 1993), and media studies (Bryant and Finklea, 2022; Cacciatore et al., 2016; Van der Pas, 2014). Framing analysis applied to written text has been the focus of frameworks, such as Critical Discourse Analysis (CDA) (Van Dijk, 2015; Fairclough, 2013), a paradigm that views language as social practises and critically reads discourse as expressions of social power. Also, the above approaches to framing all have to some extent been gaining attention in computational research. For example, Mendelsohn et al. (2021) model political framing in immigration discourse on social media using multiple framing typologies from political communication theory. Walter and Ophir (2021) predict media framing of election candidate campaigns using variables at the level of candidate, state, and electoral race.

FrameNet has a more lexicographic focus, interpreting words in terms of **semantic** frames, i.e., schematized events with participants modeled as highly specified semantic roles, i.e., **frame elements** (Ruppenhofer et al., 2010).

The abovementioned fields of research use different definitions of framing, ranging from fine-grained semantic framing to coarse-grained political framing. Although these are distinct definitions, the types of framing are interrelated (Sullivan, 2023). For example, the fine-grained semantic framing of events and their participants as evoked by constructions is foundational to language, but can be used in combination with communicative means to shape political frames. In this paper, we focus on the more fine-grained semantic framing, and implement certain notions of narratology to see how this framing reflects some of the higher order developments in the narrative surrounding an event over time.

In recent decades, besides the creation of cross-linguistic FrameNets (Torrent et al., 2018; Djemaa et al., 2016; Burchardt et al., 2009; Ohara et al., 2004), the database has been used in many different NLP annotation platforms, like Webanno (Eckart de Castilho et al., 2016) and Salto (Burchardt et al., 2006). More recently, Xia et al. (2021) created LOME, a multilingual end-to-end frame parsing system. With this Large Language Model, texts from any target language can be parsed with both frames and frame elements. Minnema et al. (2022b) implemented LOME in a multilingual tool called SocioFillmore, which performs a large-scale analysis of perspectivization strategies across texts. All those different FrameNet databases and parsing implementations have contributed substantial insight in perspectivization of events. Minnema et al. (2022a) investigate how responsibility is framed by linguistic expressions in news texts reporting on events of gender-based violence. As far as we know, they are the first to use FrameNet to analyze perspectives on referents in corpora. Yet, they did not involve the aspect of narratology and how this affects the framing of a participant over time.

In order to track an event's participant with respect to its conceptual role in the narrative of a corpus, FrameNet can serve as a suitable proxy for modeling semantic framing evoked by the mentions in that corpus. Its definition of linguistic framing depicts frames as events, its database is extensive and cross-domain, and the labels and definitions of each frame's frame elements are highly specific. Given a frame, its frame elements are proxys for the perspectives taken on the event's participants within that frame. For example, in a news text reporting on a shooting, when a predicate evokes the **Killing** frame, the frame elements realize the perspectives: the **Killing**@KILLER perspective or the **Killing**@VICTIM perspective. Yet, on top of information about perspectives, we need referential grounding: information about the referential links between mentions and structured participants given a real-world event. This way, we can get insights about which perspectives are projected on which participants.

In evaluating the aforementioned FrameNet contributions, Remijnse et al. (2022) concluded that referential grounding is still absent from the annotated data. With only the frames, we lack information about who is mentioned and who is framed. The authors built the DFN annotation tool, a resource that combines frame annotations and co-referential annotations in a dual annotation layer. After loading documents grouped under their real-world event and paired with structured data in the interface, co-referential annotation is achieved by linking in-text mentions to that structured data. On top, the same mentions are annotated with semantic frames. In the resulting annotation scheme, per participant of an event, all mentions linked to that participant are schematized with their frame-annotations.

For the purpose of this paper, since we need our corpus to be annotated with information regarding both framing and referential grounding, we make use of the DFN annotation tool to annotate our corpus data.

2.3. Narratology

When analyzing the linguistic framing of participants of a real-world event in a referentially grounded corpus, we need to take into account that on a higher order, their fine-grained conceptual representations (i.e., the frames and frame elements) reflects their role in a continuously changing narrative. Describing events by means of creating narratives is an ability inherent to human nature (Boyd, 2009; Gottschall, 2012). In analyzing narratology as a discourse phenomenon, a text displays a sequence of causally related events involving participants, which constitutes a storyline (Mani, 2014; Bal and Van Boheemen, 2009; Forster, 1956). Vossen et al. (2021); Bal and Van Boheemen (2009) point out the following requirements that a storyline needs to meet in order to qualify as a narrative:

- The ordered events lead to a **climax**, which serves the document's topic.
- It follows a **focalizer**, i.e., one of the storyline's participants.
- The focalizer takes on a certain **perspective**, a certain role in the story.

Vossen et al. (2021) further break down the storyline's event sequence in a formal model that classifies pre-climax events and post-climax events, and derive a novel annotation scheme applied to news texts.

NLP tasks that model and extract narratological information from corpus data have been scarce. First attempts resulted in entailment recognition tasks (Dzikovska et al., 2013; Bowman et al., 2015), end-of-story prediction tasks (Mostafazadeh et al., 2016, 2017) and narrative chains (Chambers and Jurafsky, 2008, 2009). Although these NLP systems pose a relevant first step in getting structural insight into storytelling, they still have been evaluated as "limited and in their infancy" (Caselli et al., 2021, 2). Moreover, they still do not combine conceptual representation and referential grounding.

In the next section, we introduce a theoretical model of storytelling that involves referential grounding and framing.

3. A Model for Variation in Framing and Storytelling

In this section, we describe our theoretical model for variation in framing of real-world events with the incorporation of narratology. We start with a description of real-world events, referential grounding, and framing. Given an event instance in the world, this event instance involves structured data involving participants, location and time. Generally, we assume that people describe and report on an event instance at a granularity that fits their daily interest, and consisting of sequences of more fine-grained events. Besides structured data, this event instance generates a stream of co-referential texts with varying temporal reporting distances. The mentions in those texts can be linked to the structured data. On top of this referential relation, the mention also evokes a semantic frame or expresses a frame element. The set of mentions across documents that co-refer the same entity, can exhibit various frame elements. This way, we can model variation in framing.

We can apply the notions of narratology introduced in Section 2.3 as follows. The storyline is conveyed by the set of causally ordered mentions in a reference text. One of the entities in the structured data is selected by the writer as the document's focalizer. When mentioning this focalizer, it is perspectivized by the frame element expressed by the mention.

Instead of topicalizing the event instance, we expect documents to sometimes topicalize distinct yet related events. The motivation for reporting on those distinct events is their relevance to the more salient event instance. In other words, the salient event instance is always referenced, but not always as the document's climax. The climax can also be one of the events that affect the salient event instance or is caused by it in the aftermath. We label the salient event instance as the **anchor incident**. See an example of an anchor incident and both its reference texts and related events on a timeline in Figure 1.

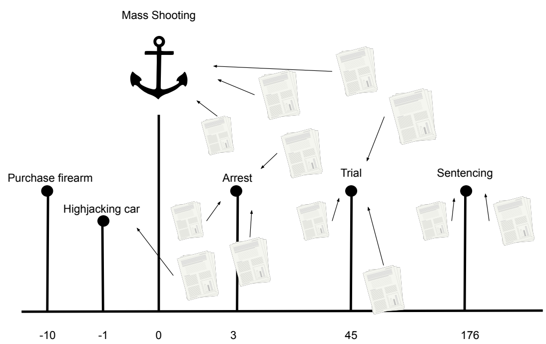

Figure 1: A timeline with a mass shooting as the anchor incident (indicated by an anchor) and different related events occurring before and after (indicated by black dots), together with reference texts (indicated by the newspapers). X-axis = number of days from the onset of the Anchor incident. The arrows from the reference texts to the events indicate the topic of writing, and thus a shift in narrative in the overall news stream over time.

Figure 1 displays a mass shooting as the anchor incident. On the timeline (x-axis), different yet related events occur before and after the shooting. Any of those events can form the climax of a reference text. Yet, in order for all those reference texts to show relevance of reporting their climax, the writers have to ground it in the related Anchor incident. We call this process **anchoring**: using at least minimal reference in your document to ground your climax incident in the anchor incident.[5]

With respect to the anchor incident's participants, we can make the following distinction:

- **DpA**: Directly related participant of the Anchor incident. This participant is present at the scene of the incident itself (e.g., shooters and victims of a mass shooting).
- **IpA**: Indirectly related participant of the Anchor incident. This participant is only indirectly involved in the Anchor incident (e.g., relatives and criminal investigators of a mass shooting).

The structured data exhibits both DpAs and IpAs. We expect IpA's to be referenced occasionally. Assuming that every reference text needs anchoring, we expect that the DpAs will always be referenced, if only briefly. Yet, when the narrative evolves around them, they might also undergo a change in their perspectivization, i.e., how they are framed.

To conclude, based on our theoretical model, we formulate the following hypotheses about the focalization and perspectivization of the participants of the anchor incidents in our data:

- On a referential level, we expect shifts in focalization between participants as a result of different topics that push the narrative surrounding the Anchor incident forward. We expect that the documents keep referencing DpAs over time, while IpAs are introduced occasionally.
- On a conceptual level, we expect that the focalized participants also show variation in framing: they show different perspectives as a result of their role in a new related topic. However, DpAs can be frequently referenced while not showing variation in framing. This is then a result of anchoring: they have to be mentioned to anchor the document's climax, but if they do not play an active role in the storyline, there is no need to change their framing.

In the next section, we describe our methodology with respect to the analysis of two anchor incidents.

4. Methodology

In this section, we describe our method. This includes corpus compilation (4.1), the annotation process (4.2), document clustering (4.3), participant selection (4.4) and data analysis (4.5).

[5]One could argue that anchoring is a product of following the Gricean maxim of Relevance, i.e., make your contribution, this news report, relevant to the reader. (see Grice (1991) and Grice (1975)).

4.1. Corpus Data

We make use of the corpus data collected by Remijnse et al. (2022). They used MWEP to query Wikidata with preset event type identifiers. For each event type, MWEP returned both structured and unstructured data for two anchor incidents. We selected two anchor incidents. The first incident is the 2019 Utrecht shooting (Q62090804), which is an instance of mass shooting (Q21480300). For this incident, MWEP returned 42 Dutch reference texts. The second incident is Malaysia Airlines Flight 17 (a.k.a MH17, Q17374096), which is an instance of aircraft shootdown (Q6539177). For this incident, MWEP returned 117 Dutch reference texts.

4.2. Annotation Process

Provided with both a reference text and the anchor incident's structured data in the annotation tool's interface, four annotators were trained to annotate different texts. Per text, they first linked in-text mentions to the structured data (entity-links). Then, they performed frame-annotations on the same text. Whenever the annotators could not find an antecedent for a frame element within its predicate's sentence boundaries, they were instructed to look for an antecedent across sentences. The main motivation for this instruction is that we assume that participants are sometimes implicitly involved in descriptions elsewhere in discourse, contributing to the storyline. As a result, mentions of participants get *n* annotations of frame elements belonging to frames evoked in different sentences.

For the Utrecht Shooting subcorpus, this process resulted in 1,459 links to 13 different entities, 1,830 frame-annotations and 5,807 assignments of frame elements. For the MH17 subcorpus, the annotation process resulted in 3,390 links to 37 different entities, 3.436 frame-annotations and 10,978 frame element assignments.

4.3. Temporal Distance Clustering

As discussed in Section 3, we expect that given an anchor incident, whenever a distinct yet related event occurs, this triggers a stream of reference texts reporting on this event as their climax. Those reports push the narrative of the anchor incident forward, possibly changing the framing of its participants. In order to model and vizualize this shift in narrative and analyze participant's framing, we first visualized the distribution of the reference texts per anchor incident, see Figure 2. We observed that in certain time periods, the documents appear to cluster. We take these clusters as a proxy for finding news streams reporting on a novel event that is topicalized as a new climax.

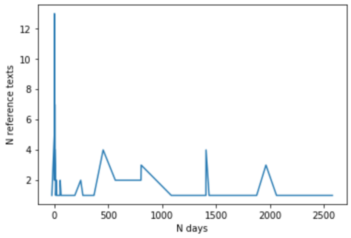

Figure 2: The distribution of reference texts reporting on MH17 over time, from the onset of the anchor incident.

For both anchor incidents, we selected the four periods on the timeline in which most reference texts were published.[6] Right after the incidents, those periods are days with few days of silence (i.e., no publications). Here, we started to widen the scope of a high frequent publication period, but already set borders when encountering a day of silence. The peaks of reference texts later on in the timeline show a larger spread of documents. Thus, we started to widen the scope of the cluster, but still included documents that were published after few days of silence. Here, the silence period had to be longer. We attempted to put the borders of the temporal classes where the reference texts would be equally balanced, as a means of normalization. Documents that fell between the borders of the classes, were removed from the final dataset. This results in the temporal distance classes (TDC) shown in Table 1.

Utrecht shooting	
Temporal Distance Class	N docs
1. Day 0	13
2. Day 1	11
3. Day 4-12	9
4. Day 37-703	9
MH17	
Temporal Distance Class	N docs
1. Day 0-1	18
2. Day 6-22	20
3. Day 333-1212	18
4. Day 1407-2581	18

Table 1: Per Anchor incident, the temporal distance classes with N of documents.

We checked the titles of the documents within each class to see if they would be reporting on the same events. Overall, this turned out to be the case, e.g., most reference texts on Day 0 of the Utrecht shooting cover a manhunt, while reference texts in Day 37-703 largely cover a trial.

[6] Considerable related work contributes to the process of clustering documents of various publication dates with respect to their topic (Wang et al., 2014; Wang and McCallum, 2006; Blei and Lafferty, 2006). Yet, the proposed models involve linguistic information from those documents in their techniques. Since the linguistic information is the object of our analysis, we had to find a different way of setting cluster boundaries, in order to avoid circularity.

4.4. Participant Selection

Both anchor incidents contain a large set of structured entities. For clear vizualization purposes, we selected all DpAs and the top three IpAs showing the highest number of entity-links. Table 2 shows the statistics in terms of number of entity-links and frame elements per participant (after temporal distance clustering). We will analyse these across TDCs in Section 4.5.

Utrecht shooting			
Participant	status	N entity-links	N FEs
Gökmen Tanis	DpA	329	983
victims	DpA	197	525
police officers	IpA	74	137
other suspects	IpA	60	89
Utrecht citizens	IpA	46	62
MH17			
Participant	status	N entity-links	N FEs
victims	DpA	239	462
suspects	DpA	126	356
relatives	IpA	123	169
Russia	IpA	166	283
Dutch Safety Board	IpA	69	96

Table 2: Selected participants per anchor incident with involvement status, number of entity-links and number of frame element annotations.

4.5. Data Analysis

We analyze the participants' referential grounding as well as their perspectivization separetely per anchor incident. With respect to referential grounding, we take frequency distribution of the entity-links as a proxy for focalization: we expect participants that are referenced most are focalized in storytelling. We distributed the number of entity-links of the participants per TDC to observe any shift in this focalization.

Regarding the perspectivization of the participants, we take frequency distribution of frame elements as a proxy for variation in framing: within a participant, a strong shift in frame element frequency shows a change in perspectivization. This new perspective that the participant is assigned with, is part of a change in narrative. Given that the frequency distribution of frame elements given a participant shows a long tail, we limited the vizualization of the data to the proportionally most salient frame elements. Per participant and per TDC, we sliced the top frequent 35% of frame elements. After experimenting with different percentages, this proportional number gave for every participant a sufficient number of different frame element types to interpret. Then, per participant, we took the union of those slices across TDCs. The resulting set contains the frame elements we assume convey most information about how the participant in general has been framed over time. Next, per participant, we plotted per frame element type its proportional frequency distribution across time buckets.

In the next section, we present and discuss the results of our data analysis.

5. Results and Discussion

In this section, we present and discuss the results of both the referential part and the framing part of our analysis. We first turn to the data of the Utrecht shooting. See Figure 3a for the referential grounding of the participants of the Utrecht shooting across TDCs. Overall, we find that the narrative evolves around the DpAs as they show the highest number of references. The focus shifts over time, from Gökmen Tanis (the perpetrator of the shooting) on Day 1 to the victims on Day 4-12, then back to Tanis on Day 37-703.

See Figure 3b-d for the proportional distribution of frame elements that were used to frame the participants across TDCs. We plotted this distribution for a selection of participants. A first observation is that all three participants show variation in framing over time. Their perspectives are not fixed, but subject to change. This is only possible if the narrative changes and writers choose to give the participants new perspectives.

Second, we notice a correlation between each participant's framing and its referential grounding. When a participant shows a significant change in number of references, its framing changes accordingly. For example, whereas Gökmen Tanis is focalized in Day 37-703 in 3a, in that same TDC in Figure 3b, frame elements such as **Hit_target**@AGENT decrease, while many frame elements, such as **Trial**@DEFENDANT, start to increase. Similarly, the victims are focalized in Day 4-12 in 3a, whereas their framing changes significantly in that same TDC in 3c. Finally, we see that the police officers show a decrease of references on Day 1 in Figure 3a, and similarly in Figure 3d, most frame elements drop to zero in that same TDC, while **Intentionally_act**@AGENT is introduced to henceforth frame this group.

Turning to the data of MH17, see Figure 4a for the referential grounding of the participants of MH17 across TDCs. This figure shows a shift in top frequent references between a variety of participants. Here, the IpAs play a larger role in the narrative as compared to the Utrecht shooting. Whereas the Dutch Safety Board (an investigation board) peaks in number of references at Day 333-1212, that number drops to zero in subsequent TDC, when the suspects are introduced. Together with the Russian government, they take over the narrative, completely backgrounding the Dutch Safety Board.

See Figure 4b-d for the proportional distribution of frame elements that were used to frame the participants across TDCs. We plotted this distribution for a selection of the participants. In Figure 4c-d, again we observe strong variation in framing over time. In fact, each TDC in Figure 4d shows different top frequent frame elements. Figure 4c and 4a show both a clear change in mention frequency and a change in frame element types in Day 333-1212.

Figure 4b shows no variation in framing over time. In fact, **Catastrophe**@PATIENT increases in the final TDC, whereas the participant simultaneously decreases in number of references in Figure 4a. This means that the victims do not play an active part in the narrative anymore and their perspective freezes. As a DpA, this group is still mentioned across documents in order to anchor the document's climax to MH17, but with a fixed set of frames.[7]

To conclude, we interpret the shift in frequency distribution of references in both Figures 3a and 4a as a shift in focalization between participants that play an active role in the development of the narrative. Furthermore, we notice that DpAs and IpAs behave differently between anchor incidents in their mention frequencies. In the Utrecht shooting, overall focus is on the DpAs, whereas in MH17, the IpAs are focalized to a stronger degree over time, and the DpA suspects is only introduced in the final TDC. It seems that the narrative surrounding each anchor incident is unique and affects different patterns of referential grounding of participants over time. The current analysis captures this development.

With respect to framing, we observe an overall correlation between a participant's focalization (sudden steep change in number of references) and variation in framing (shift in dominant frame type). The victims of MH17 are consistently framed with the same frame elements. Even when their number of references decreases over time, they are still necessarily mentioned to anchor the climax, but their part in the narrative has not changed, i.e., they do not require new perspectives.

[7] Remijnse et al. (2021) describe a similar process. They derive a fixed set of what they call **Anchor** frames that writers consistently evoke over time when anchoring a documents climax. This set of anchor frames is dependent of the anchor incident's event type.

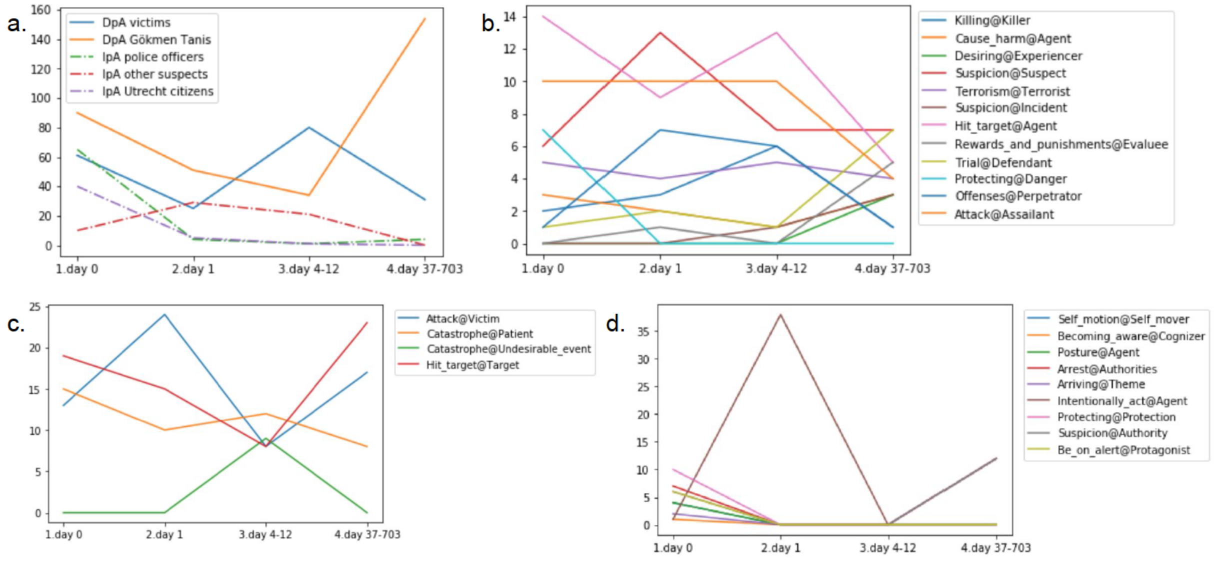

Figure 3: a. The frequency distribution of mentions of participants of the Utrecht shooting across TDCs; b-d. The proportional distribution of the top frequent frame elements across TDCs in reference to participants of the Utrecht shooting. b. Gökmen Tanis; c. victims; d. police officers. The frame element notations in the index can be read as frame@frame_element.

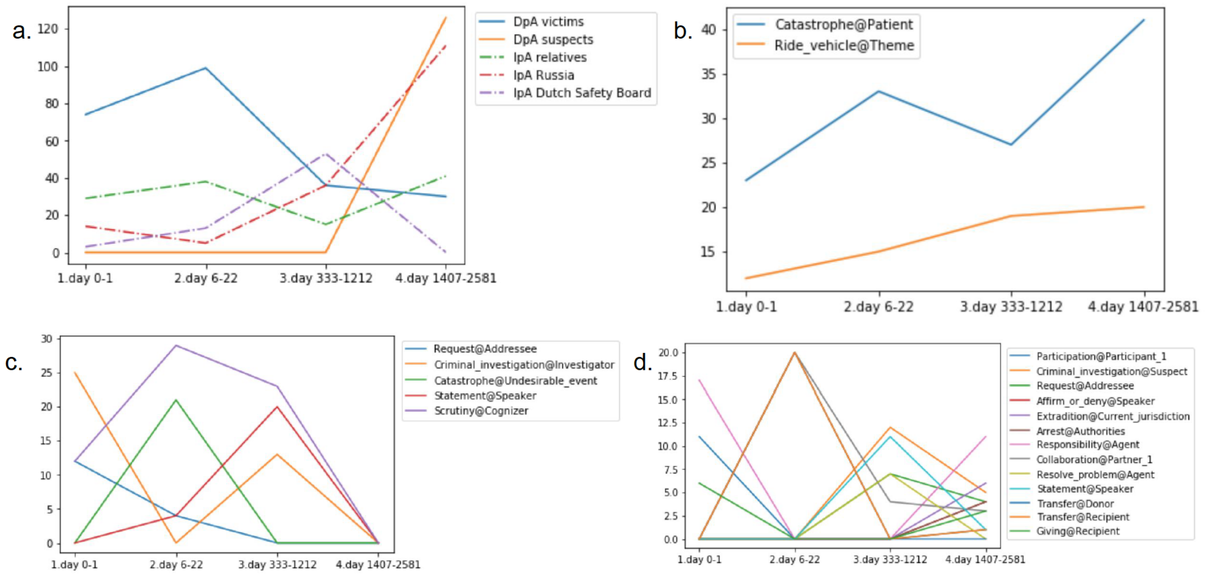

Figure 4: a. The frequency distribution of mentions of participants of MH17 across TDCs; b-d. The proportional distribution of the top frequent frame elements across TDCs in reference to participants of the Utrecht shooting. b. victims; c. Dutch Safety Board; d. Russian government. The frame element notations in the index can be read as frame@frame_element.

6. Conclusion

In this paper, we performed a structural linguistic analysis of variation in framing of participants of real-world events over time. In order to perform such an analysis, we met multiple requirements: collect a referentially grounded corpus accompanied with structured data; annotate the data with both entity-links and frames; and describe a theoretical model of variation in framing and narratology.

For two anchor incidents, we first analyzed the frequency distribution of the entity-links for the participants on a timeline, to observe a strong shift in focalization between participants, an indication of a change in narrative.

We then analyzed the frequency distribution of the participants' frame elements over time in order to measure the extent of variation in framing. Overall, we observe a correlation between shift in focalization between participants, and variation in framing within a participant: when a different participant is getting the focus of the narrative, this participant also gets a new perspective.

7. Limitations

The first limitation of our research is that it is limited to two incidents and one incident type. In order to find stronger patterns of focalization and perspectivization, we need to scale this to many more event types, incidents per event type and sources of reference text. We hope to do this in future work using automated techniques for frame annotation, coreference resolution and entity-linking. The MWEP tool can be used to collect large corpora of referentially grounded texts for this.

Another limitation is that we have now manually annotated the texts but need to rely on automatic techniques to scale our research. Automatic techniques will be less accurate and biased to assign more dominant frames and frame elements which may cause a bias in the analysis. Furthermore, current tools are not trained to annotate frame and frame element relation at the discourse level.

Finally, we hand-picked the TDCs for our analysis on the collected reference texts. This should be automated as well to apply this on a larger scale. However, what is a pause or not in the publication also depends on the ability to find all sources at any point of time that report on each incident. Furthermore, there could be multiple new events in the same TDC and related to the same anchor incident that are being reported.

8. Bibliographical References

Mieke Bal and Christine Van Boheemen. 2009. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.

David M Blei and John D Lafferty. 2006. Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120.

Ilona Bos, Amber Dujardin, Hans Marijnissen, Gidi Pols, and Petra Vissers. 2019. https://www.trouw.nl/home/om-kwart-voor-elf-staat-de-tram-stil-daarna-heel-utrecht a1d5dacf/.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. pages 632–642.

Brian Boyd. 2009. *On the origin of stories: Evolution, cognition, and fiction*. Harvard University Press.

Jennings Bryant and Bruce W Finklea. 2022. *Fundamentals of media effects*. Waveland Press.

Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski, and Sebastian Pado. 2006. SALTO - a versatile multi-level annotation tool. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).

Michael A Cacciatore, Dietram A Scheufele, and Shanto Iyengar. 2016. The end of framing as we know it... and the future of media effects. *Mass communication and society*, 19(1):7–23.

Tommaso Caselli, Eduard Hovy, Martha Palmer, and Piek Vossen. 2021. *Computational Analysis of Storylines: Making Sense of Events*. Cambridge University Press.

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.

Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2 - Volume 2*, ACL '09, page 602–610, USA. Association for Computational Linguistics.

Prafulla Kumar Choubey, Kaushik Raju, and Ruihong Huang. 2018. Identifying the most dominant event in a news article by mining event coreference relations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 340–345, New Orleans, Louisiana. Association for Computational Linguistics.

James N Druckman. 2001. The implications of framing effects for citizen competence. *Political behavior*, 23:225–256.

Myroslava Dzikovska, Rodney Nielsen, Chris Brew, Claudia Leacock, Danilo Giampiccolo, Luisa Bentivogli, Peter Clark, Ido Dagan, and Hoa Trang Dang. 2013. SemEval-2013 task 7: The joint student response analysis and 8th recognizing textual entailment challenge. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 263–274, Atlanta, Georgia, USA. Association for Computational Linguistics.

Richard Eckart de Castilho, Éva Mújdricza-Maydt, Seid Muhie Yimam, Silvana Hartmann, Iryna Gurevych, Anette Frank, and Chris Biemann. 2016. A web-based tool for the integrated annotation of semantic and syntactic structures. In *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH)*, pages 76–84, Osaka, Japan. The COLING 2016 Organizing Committee.

Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43(4):51–58.

Norman Fairclough. 2013. *Critical discourse analysis: The critical study of language*. Routledge.

Elena Filatova and Vasileios Hatzivassiloglou. 2004. Event-based extractive summarization. In *Text Summarization Branches Out*, pages 104–111, Barcelona, Spain. Association for Computational Linguistics.

Edward Morgan Forster. 1956. *Aspects of the Novel*. Harcourt, Brace.

Roeland Franck. 2019. https://www.ad.nl/binnenland/gokmen-tanis-bekent-schietpartij-in-tram-in-utrecht a239a24a/.

Jeremy Getman, Joe Ellis, Stephanie Strassel, Zhiyi Song, and Jennifer Tracey. 2018. Laying the groundwork for knowledge base population: Nine years of linguistic resources for tac kbp. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.

Jonathan Gottschall. 2012. *The storytelling animal: How stories make us human*. Houghton Mifflin Harcourt.

H Paul Grice. 1975. Logic and conversation. *Cole, P., and Morgan, J.(Eds.)*, 3.

Paul Grice. 1991. *Studies in the Way of Words*. Harvard University Press.

Ben Hachey, Will Radford, Joel Nothman, Matthew Honnibal, and James R. Curran. 2013. Evaluating entity linking with wikipedia. *Artificial Intelligence*, 194:130–150. Artificial Intelligence, Wikipedia and Semi-Structured Resources.

Kasper Hermans. 2019. https://www.metronieuws.nl/in-het-nieuws/2019/03/van-zanen-twee-verdachten-vrij-politie-ontkent/.

Dorothea Horst. 2020. Patterns 'we' think by? critical cognitive linguistics between language system and language use. *Yearbook of the German Cognitive Linguistics Association*, 8(1):67–82.

Shanto Iyengar. 1994. *Is anyone responsible?: How television frames political issues*. University of Chicago Press.

Inderjeet Mani. 2014. Computational narratology. *Handbook of narratology*, pages 84–92.

Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. Modeling framing in immigration discourse on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.

Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022a. Dead or murdered? predicting responsibility perception in femicide news reports. pages 1078–1090.

Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022b. SocioFillmore: A tool for discovering perspectives. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 240–250, Dublin, Ireland. Association for Computational Linguistics.

Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.

Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. LSDSem 2017 shared task: The story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pages 46–51, Valencia, Spain. Association for Computational Linguistics.

Levi Remijnse, Marten Postma, and Piek Vossen. 2021. Variation in framing as a function of temporal reporting distance. In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 228–238, Groningen, The Netherlands (online). Association for Computational Linguistics.

Levi Remijnse, Piek Vossen, Antske Fokkens, and Sam Titarsolej. 2022. Introducing frege to Fillmore: A FrameNet dataset that captures both sense and reference. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 39–50, Marseille, France. European Language Resources Association.

Heather Simpson, Stephanie Strassel, Robert Parker, and Paul McNamee. 2010. Wikipedia and the web of confusable entities: Experience from entity linking query creation for TAC 2009 knowledge base population. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).

Karen Sullivan. 2023. Three levels of framing. *Wiley Interdisciplinary Reviews: Cognitive Science*, 14(5):e1651.

Daphne Van der Pas. 2014. Making hay while the sun shines: Do parties only respond to media attention when the framing is right? *The International Journal of Press/Politics*, 19(1):42–65.

Teun A Van Dijk. 2015. Critical discourse analysis. *The handbook of discourse analysis*, pages 466–485.

Piek Vossen, Tommaso Caselli, and Roxane Segers. 2021. A narratology-based framework for storyline extraction. *Computational Analysis of Storylines: Making Sense of Events*, 125.

Piek Vossen, Filip Ilievski, Marten Postma, Antske Fokkens, Gosse Minnema, and Levi Remijnse. 2020. Large-scale cross-lingual language resources for referencing and framing. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3162–3171, Marseille, France. European Language Resources Association.

Piek Vossen, Filip Ilievski, Marten Postma, and Roxane Segers. 2018. Don't annotate, but validate: a data-to-text method for capturing event data. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.

Dror Walter and Yotam Ophir. 2021. Strategy framing in news coverage and electoral success: An analysis of topic model networks approach. *Political Communication*, 38(6):707–730.

Jian Wang, Xianhui Liu, Junli Wang, and Weidong Zhao. 2014. News topic evolution tracking by incorporating temporal information. In *Natural Language Processing and Chinese Computing: Third CCF Conference, NLPCC 2014, Shenzhen, China, December 5-9, 2014. Proceedings 3*, pages 465–472. Springer.

Xuerui Wang and Andrew McCallum. 2006. Topics over time: a non-markov continuous-time model of topical trends. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 424–433.

Patrick Xia, Guanghui Qin, Siddharth Vashishtha, Yunmo Chen, Tongfei Chen, Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, and Benjamin Van Durme. 2021. LOME: Large ontology multilingual extraction. pages 149–159.

Alexander Ziem, Christian Pentzold, and Claudia Fraas. 2018. Medien-frames als semantische frames: Aspekte ihrer methodischen und analytischen verschränkung am beispiel der ‚snowden-affäre '. *Frames interdisziplinär: Modelle, Anwendungsfelder, Methoden*, pages 155–182.

9. Language Resource References

Collin F. Baker, Charles J. Fillmore, and Beau Cronin. 2003. The Structure of the FrameNet Database. *International Journal of Lexicography*, 16(3):281–296.

Cosmin Bejan and Sanda Harabagiu. 2010. Unsupervised event coreference resolution with rich linguistic features. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1412–1422, Uppsala, Sweden. Association for Computational Linguistics.

Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski, Sebastian Padó, and Manfred Pinkal. 2009. 8. using framenet for the semantic analysis of german: Annotation, representation, and automation. In *Multilingual FrameNets in computational lexicography*, pages 209–244. De Gruyter Mouton.

Agata Cybulska and Piek Vossen. 2014. Using a sledgehammer to crack a nut? lexical diversity and event coreference resolution. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4545–4552, Reykjavik, Iceland. European Language Resources Association (ELRA).

Marianne Djemaa, Marie Candito, Philippe Muller, and Laure Vieu. 2016. Corpus annotation within the French FrameNet: a domain-by-domain methodology. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3794–3801, Portorož, Slovenia. European Language Resources Association (ELRA).

Kyoko Hirose Ohara, Seiko Fujii, Toshio Ohori, Ryoko Suzuki, Hiroaki Saito, and Shun Ishizaki. 2004. The japanese framenet project: An introduction. In *Proceedings of LREC-04 Satellite Workshop "Building Lexical Resources from Semantically Annotated Corpora"(LREC 2004)*, pages 9–11.

Haoruo Peng, Yangqiu Song, and Dan Roth. 2016. Event detection and co-reference with minimal supervision. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 392–402, Austin, Texas. Association for Computational Linguistics.

Marten Postma, Levi Remijnse, Filip Ilievski, Antske Fokkens, Sam Titarsolej, and Piek Vossen. 2020. Combining conceptual and referential annotation to study variation in framing. In *Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet*, pages 31–40, Marseille, France. European Language Resources Association.

Sameer Pradhan, Edward Loper, Dmitriy Dligach, and Martha Palmer. 2007. SemEval-2007 task-17: English lexical sample, SRL and all words. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 87–92, Prague, Czech Republic. Association for Computational Linguistics.

Levi Remijnse, Piek Vossen, Antske Fokkens, and Sam Titarsolej. 2022. Introducing frege to Fillmore: A FrameNet dataset that captures both sense and reference. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 39–50, Marseille, France. European Language Resources Association.

Josef Ruppenhofer, Michael Ellsworth, Miriam R L Petruck, Christopher R Johnson, and Jan Schefczyk. 2010. Framenet ii: Extended theory and practice.

Tiago Timponi Torrent, E Matos, Ludmila Lage, Adrieli Laviola, Tatiane Tavares, VG Almeida, and Natália Sigiliano. 2018. Towards continuity between the lexicon and the constructicon in framenet brasil. *Constructicography: Constructicon development across languages*, 22:107.