ARTICLE TEMPLATE

Geo-spatial Hypertext in Virtual Reality: Mapping and Navigating Global News Event Spaces

Theresa Berg and Jasper Husetdt and Julian Gagel and Timo Lüttig and Giuseppe Abrami and Alexander Mehler

Goethe University Frankfurt, Texttechnology

ARTICLE HISTORY
Compiled March 2, 2024

ABSTRACT
Everyday, a myriad of events takes place globally, and these occurrences are documented and shared online through news articles from diverse sources. As a result, when users navigate the web, the immense volume of data can lead to information overload, making it difficult to find specific details, for example, about a particular ongoing conflict. We present News in Time and Space (NiTS) to address this issue: NiTS organizes targeted news information within a geo-spatial hypertext system in virtual reality. With NiTS, users can visualize, filter and interact with news reports currently based on GDELT on a virtual globe providing comprehensive document networks to analyse global events and trends. The paper describes NiTS, its event filter semantics and architecture, and discusses several use cases that demonstrate its application potential.

KEYWORDS
virtual reality, geographic information system, virtual hypertext, collaborative interaction, spatio-temporal visualisation, monitoring global events

1. Introduction

Information gathering and processing are of central importance for the information society (Wirfs-Brock and Quehl, 2019). There is a wide range of factors influencing information processing by individuals, including media (e.g. Schneiders et al. (2022)), formats (e.g. Lottridge et al. (2022)), platforms and systems (e.g. Molyneux (2018)), communities (e.g. Bongard et al. (2023)), interfaces (e.g. Chongtay et al. (2018)), types of consumption (e.g. Perez et al. (2020)), its depth and breadth, as well as functions to be fulfilled such as the identification of factuality (Jose et al., 2021; Vosoughi et al., 2018), trustworthiness (Brandtzaeg and Følstad, 2017) or authenticity (Ombelet et al., 2016). Information processing is a challenge, especially given the volume of heterogeneous information produced every day, every hour, even every minute. Several approaches address this challenge. More traditional frameworks address the problem of information overload, such as RSS readers or news aggregators that help curate news based on users' interests (Elyusufi and En-Naimi, 2017). These approaches help to organise and consolidate content from different websites, bringing it together in

Theresa Berg ✉ theresa@wenzler.org
Alexander Mehler ✉ mehler@em.uni-frankfurt.de

one place so that people can stay informed without feeling overwhelmed (Zhang et al., 2022). A more recent approach, or one that is almost on the verge of widespread use, is the use of specialised chatbots based on *Large Language Models* (LLMs) (Cerf, 2023). This approach is characterised by the fact that documents are turned into queryable objects, so to speak, in terms of their content and intertextual relations using the underlying LLM. According to the current possibilities of LLMs, such queries are at an unprecedented quasi-semantic level, so that a broad field of application for information systems based on them is likely to develop quickly. Hardly any information system that addresses the problem of information overload can currently do without such a system. Not even the system presented here. Nevertheless, we believe that there is a lack of access to information that is central to human information processing and that is multimodal information – beyond what is currently processed in terms of images, audio files or videos. That is, there is a need to experiment with new systems that better reflect the characteristics of human information processing in its depth and breadth, that is, multimodal information processing. This is where the information system called News in Time and Space (NiTS – Gagel et al. 2023) comes in: NiTS maps news about events from all around the world and visualizes them on a globe in a Virtual Reality (VR) scenario. This is exemplified by GDELT (*Global Database of Events, Language and Tone*) (Leetaru and Schrodt, 2013) as a dynamic, ever-growing news repository. That is, through the immersive lens of VR and facilitated by the integration of Ubiq (Friston et al., 2021), NiTS introduces a collaborative environment for exploring news about events as disseminated by GDELT in a virtual space where users can interact multimodally in retrieving, filtering and processing these news. Furthermore, by referring to the visual concept of a globe and mapping news onto it in a variable way (i.e., by referring to freely selectable resolutions of visualisation), a geometric map of the relations of documents is created based on the spatially localised events described by them. This means that, unlike in chatbot scenarios (see above), in NiTS the user interacts with large quantities of visualised documents by being presented with a spatial (Schedel and Atzenbeck, 2016) virtual (Norton, 2023) hypertext that models the event-based linking (Wobbrock, 2020) of these documents. Finally, NiTS adds the time layer to this interface, allowing users to filter information according to its spatio-temporal structure.

Creating a collaborative VR-based interface for this sort of information presents several challenges due to the complexity of fusing spatial and temporal data in an immersive 3D environment. The design and development of such a system should therefore follow criteria that facilitate their (re-)use. Such a list of criteria has been considered by Abrami et al. (2023) in the context of VR-based systems. It can be referred to with a focus on the way of interacting with information objects as provided by NiTS. To this end, we consider the following criteria:

(I) **Multiple data sources**: The use of a single source should be avoided; rather, the integration of different sources from different media is required. It is important to integrate not only traditional news sources (e.g. newspapers) but also social media (e.g. Reddit, Blue Sky), audio (e.g. podcasts, interviews) and video sources (e.g. YouTube) in order to provide users with more comprehensive information.

(II) **Flexible explicit views or filters**: In order to visualise and interact with information in a flexible way, it is necessary to implement information views or filters that can be parameterised by users, so that they can selectively access information according to their level of knowledge or focus of interest. This means that filter criteria should be user-centric, flexibly definable, combinable and storable to enable complex queries that can be shared and reused between users. Finally, it should be possible to design filters for specific audiences (e.g. in the context of an educational application as opposed to a political debate).

(III) **Implicit views or filters**: In addition to explicit filters, it should be possible to implicitly filter information based on user behaviour, for which virtual movements in VR systems can be evaluated. It is essential that these movements relate to the interaction with specific information objects, so that not any movement behaviour is evaluated. Implicit filters can lead to unforeseen adaptations of the environment, so it must be possible to switch them off.

(IV) **Virtual (VR), augmented (AR) and mixed reality (XR)**: In order to ensure the visualisation, aggregation and arrangement of information in multi-dimensional, ubiquitous environments (e.g. (Abrami et al., 2020)), as well as its barrier-free consumption and reuse, it is necessary to strive for a solution that supports multiple visualisation and interaction technologies. That is, a system is needed that is platform independent and supports VR, AR and XR devices.

(V) **Integration of *Natural Language Processing* (NLP)**: In order to interact with information at the content level, this information needs to be enriched with annotations (e.g. metadata). This enrichment concerns not only the information objects (e.g. documents) as a whole, but also their parts, so that these can be addressed separately in search processes and linked to other information objects. NLP methods are to be used for this purpose, the application of which should be as simple, standardised and flexible as possible in view of the volume of data to be processed (see criterion I). Such a system also to be integrated with NiTS is introduced by Leonhardt et al. (2023).

(VI) **Integration of large language models (LLM)**: One aspect of NLP is the use of LLMs, as their zero-shot learning capabilities allow for dynamic processing of text at a quasi-semantic level that does not require the predefinition of NLP tasks. In this way, texts captured by explicit filters become queryable objects for users, with the LLM used acting as a kind of oracle that can provide the user with advance information about the document content.

We present NiTS to meet these criteria. That is, we present NiTS as a VR-based system that allows for the flexible filtering of large scale information sources so that they can be processed by single users or groups of interacting users to meet their information needs. An essential part of NiTS is the integration with DUUI. This gives NiTS access to a very wide range of NLP tools that can be used to enrich the underlying documents and integrate the resulting annotations into the creation of explicit filters. The result is NiTS, a system that combines VR-based information access and processing with machine learning-based NLP in the area of large news repositories.

The paper is organized as follows: In Section 2 we present related work. Section 3 describes NiTS in detail, the usage scenario it implements and its technological backbone as well as its entanglement with DUUI. Section 4 presents several specialized use cases of NiTS and Section 5 concludes with an outlook for future work.

No.	Tool	Reference	I	II	III	IV	V	VI
1	VR GIS map	Santos et al. (2018)	not satisfied	not satisfied	not satisfied	satisfied	not satisfied	not satisfied
2	Underground infrastructures	Jurado et al. (2017)	satisfied	unknown	unknown	partially satisfied	not satisfied	not satisfied
3	Blue Link City	Linder et al. (2018)	satisfied	not satisfied	not satisfied	satisfied	not satisfied	not satisfied
4	Tweets in VR Disneyland	Okada et al. (2018)	satisfied	satisfied	not satisfied	satisfied	satisfied	not satisfied
5	NiTS	Gagel et al. (2023)	satisfied	satisfied	not satisfied	satisfied	partially satisfied	not satisfied
satisfied			4	2	0	5	1	0
partially satisfied			0	0	0	0	1	0
not satisfied			1	2	4	0	3	5
unknown			0	1	1	0	0	0

Table 1.: Overview of approaches in relation to criteria I-VI from Section 1. Legend: satisfied (blue), partially satisfied (grey), not satisfied (red), unknown (yellow).

2. Related Work

There is relatively little research on mapping information in VR, especially in a globe format, as this is a relatively new area of research. Ongoing research is focusing on several sub-problems that are relevant to NiTS. This includes GIS (Geographic Information Systems) (Wei and Yao, 2022), mapping of spatio-temporal information (e.g. Nikolaou et al. 2015), data visualization in VR, studies of conflict data based on GDELT and other sources, and the use of VR in collaborative research, especially in the context of journalism or international development cooperation.

Santos et al. (2018) present a VR map for geographic information systems, the so-called VR GIS map, which allows users to interact with maps in an immersive way. Other research focuses on the handling of and interaction with large maps in VR (Giannopoulos et al., 2017), however, without addressing the integration of geo-located data such as news reports and the analysis of real-time data in VR. Due to the lack of external data in VR GIS map criteria (I), (III), (IV), (V), (VI) remain unfulfilled. However, the implementation in Unity ensures platform independence and the integration of body- and device-based interaction in VR GIS map contributes to fulfilling criterion (IV).

Research underscores the significance of geo-spatial (Ma et al., 2008) VR, offering a novel horizon for interactive and compelling learning experiences (Cho and Chun, 2019). Beyond education, geographic information in VR also serves as a valuable source for visualizing and interacting with virtualized infrastructures (Underground infrastructure through VR; Jurado et al. 2017) or public transportation systems (León-Paredes et al., 2020) due to its realistic adaptation to the real world. Underground infrastructure through VR is a multi modal approach (I) using different types of data, pictures and CAD maps to reconstruct underground structures in a web environment for public access in VR so that criterion (IV) is partially met. Whether implicit(II) or explicit (III) filters are used is unclear. However, NLP (V) and especially LLMs (VI) are not used in this system.

The integration of social media data with the goal of providing a more comprehensive perspective within VR is an emerging field (Moran et al., 2015). Examples include the construction of a fictional landscape based on social media (Blue Link City; Linder et al. 2018) or the geographic localization of tweets in a specific area (Tweets in VR Disneyland; Okada et al. 2018). Blue Link City is a cross-platform VR application (IV) implemented with Unity, and is multimodal (I) in that it incorporates text resources from Reddit and allows the display of linked videos and images without having to open an external browser. However, there are no explicit or implicit filters in the sense that a user can modify the cityscape itself and customize the information it should contain (II). Furthermore, Blue Link City does not include NLP (V) or LLM (VI). Tweets in VR Disneyland can be considered multimodal (I), as it incorporates both textual information and images, enhancing the overall location information of Tweets within Disneyland. With the implementation in Unity and the possibility to fly around Disneyland in two different display modes (WorldView and MiniMap) and to select a time range for this, criteria (II) and (IV) are fulfilled. In addition, topic modeling based on Latent Dirichlet Allocation is used to improve the selection of tweets based on their content, which satisfies criterion (V). However, implicit filters (III) and LLMs (VI) are not integrated.

Although there are a number of projects that use geographic information to enrich VR applications, there is no such reference for conveying the content of very large volumes of documents. Using GDELT data, the goal of NiTS is to provide the user with a form of information processing of news in which event networks are made accessible through the articles describing them by selecting alternative contextual filters. The advantage of NiTS is that its information is spatially located via an interactive, zoomable globe. This gives the user a dual, spatial and symbolic access to the underlying information and distinguishes NiTS from the systems presented here.

3. News in Time and Space

3.1. The Usage Scenario of NiTS*: Locating Event Networks on an Interactive Globe*

NiTS maps news from GDELT and visualizes them on a virtualized globe. Based on the Unity engine Juliani et al. (2020), NiTS runs on both VR headsets and desktop platforms; using Ubiq (Friston et al., 2021), it enables collaborative working so that users can simultaneously exchange ideas and coordinate their work effectively. NiTS is open source and available on GitHub[1]. GDELT is a collection of news articles, online posts, and other textual sources from around the world. GDELT is one of the largest global event databases, capturing news from almost every region of the world in multiple languages and providing researchers and analysts with a comprehensive dataset. It includes over 1 billion entries dating back to 1979 and is updated every 15 minutes (Leetaru and Schrodt, 2013) by means of a CSV file, which NiTS pulls into its MongoDB (see subsection 3.2). Each CSV file contains a list of names of events taking place anywhere in the world and documented in news articles and other media sources. They range from political protests and conflicts to diplomatic activities, speeches and various other social and political interactions.

NiTS can be seen as a system for the spatio-temporal visualisation of this GDELT data in VR, in which users refer to events by means of explicit filters (see section 1), so that the news relating to these events are displayed on its globe as an interactive document network, possibly in conjunction with a timeline. *Globe*, *event*, *filter*, *timeline* and *document network* are thus the essential components of NiTS data model. The usage scenario addressed by NiTS concerns browsing and interacting with docu-

[1] `https://github.com/texttechnologylab/NewsInTimeAndSpace/`

Figure 1.: Extract from NiTS's globe showing the distribution of news articles in West Africa (left), marked by green dots which indicate the number of corresponding articles. By zooming in, we reach the second aggregation level, which shows Nigeria in the centre and the distribution of events across its regions. If we zoom in further, a detailed network unfolds, allowing us to observe the distribution of articles aggregated by city in Nigeria.

ment networks, subject to their event-related intertextual structure. This structure is made accessible by NiTS by means of so-called invariants, i.e. constituents that occur in related events (persons, places, artefacts, etc. – see subsection 3.3). Depending on their choice of filters (manifesting their point of interest; see below), users get representations of the corresponding document networks, where each document is mapped onto NiTS's globe based on the location information provided by GDELT.[2] We now describe the components of NiTS step by step, starting with the globe functionality and the timeline, before explaining the filter mechanism of NiTS in detail.

The most visible part of NiTS is the globe it uses to organise the presentation of news (see Figure 3a): GDELT events are marked on the globe with green dots, the size of which varies based on the number of events occurring at a specific location (see Figure 1 and (see Figure 3c)). Calculating locations and mapping coordinates onto the surface of NiTS's globe is realised by utilizing the Robinson Projection (Ipbuker, 2004). To visualize locations, NiTS implements a dynamic zoom feature based on three levels of aggregation: the country, the region and the city level. Figure 1 shows the zooming function, using the geospatial distribution of events in Nigeria as an example. We adapted the globe and its shaders from Lague et al. (2023) game "Geographical Adventures". The use of a globe in conjunction with its interactive use presented unforeseen challenges. This concerns the globe's shader, where the shading is highly dependent on the orientation of the globe. To overcome this problem, we decided to move the player virtually instead of the globe. This gives the player the illusion that the globe is moving, when in fact it is the player's avatar that is moved. Zooming and rotating remain unchanged from this operation. Enabling an intuitive VR experience, we draw inspiration from a method commonly used in smartphone map apps to adapt to what users are accustomed to. That is, scaling the globe up and down is achieved by moving hands closer together or farther apart (see Figure 3b). In response to any such gesture, NiTS aggregates the number of corresponding events (see Figure 1 for an example). In any event, navigating on the globe is similar to using a tabletop globe, i.e. the virtual globe can be rotated by hand to reach the desired destination.

GDELT organises news in a timely manner and also provides real-time updates. In this way, time points and time intervals are accessible as explicit information filters. To

[2] In conjunction with the use of an LLM, NiTS would enable the combination of chat-based and browsing-based document processing.

map this, NiTS implements a timescale (the so-called Time Filter) to make events that occur on a specific date or in a specific time interval selectable for display on the globe. Figure 2 provides an overview of all filters implemented in NiTS.

Each event is assigned a unique Global Event ID, the GDELT ID, which can be used to reference information about the actors involved, the location, date, time and type of event. GDELT uses a non-public algorithm to identify and categorise events based on the information contained in the source text, and assigns events to specific event types derived from a hierarchical classification system. To do so, GDELT employs the CAMEO (*Conflict and Mediation Event Observations*) taxonomy to categorize information; this taxonomy serves as the foundational structure for implementing our frontend filters (Gerner et al., 2002). CAMEO offers a standardised method to systematically categorize events in international relations such as political conflict or diplomatic engagements (Keertipati et al., 2014). CAMEO compromises 20 main categories known as *Event Root Codes* (ERC) that are classified by two digits 01-20 Schrodt (2012). These represent high-level classifications such as *conflict*, *cooperation* or *political interactions*. Each ERC has several sub-categories, which are more specific types of events that fall under the main category indicated by the third digit of the CAMEO code. Subcategories can be expanded into detailed descriptions represented by the fourth digit. For example, in the CAMEO code "0213", "02" represents the ERC "Appeal", the third digit "1" specifies the subcategory *Appeal for material cooperation* and the last digit "3" describes the subcategory to the type of cooperation made "Appeal for judicial cooperation". Here is an example of a sentence that would be coded with CAMEO "0213": "*Turkey renewed an appeal to Belgium to extradite a far-left militant wanted for murder, Justice Minister Cemil Cicek said Thursday, slamming what he called lax international cooperation against terrorism.*" (Gerner et al., 2002). In addition to the Time Filter, NiTS offers four additional filters as illustrated in Figure 2. We now explain these filters, which are all implemented on the basis of CAMEO, that is, the Event Type Filter, the Actor Filter, the Tone Filter and the Goldstein Scale Filter. To facilitate navigation through news about events characterized by CAMEO codes, we provide the Event Type Filter, located on the left hand side of the globe. For the sake of clarity, the user interface contains a sorted list of the 20 ERCs, labelled with their names and sorted according to the frequency of the events currently displayed on the globe. Users can click on these codes to explore the subcategories associated with the events. These subcategories are also sorted by frequency and named according to the CAMEO taxonomy. Each subcategory is represented by a filter sphere containing a 3D object, indicating the respective ERC to which the subcategory belongs (see Figure 3e). For instance, all subcategories that fall under the ERC Protest have a blue fist as their 3D object inside the sphere, indicating that the subcategory is a specification of that ERC (see Figure 3e).

Without having to consult the CAMEO Event Code documentation (Schrodt, 2012), users can interact directly with the Event Type Filter: The 3D figures and objects within the sphere are thematically aligned with the 20 ERCs to ensure an intuitive understanding of the filters. As soon as a user has decided what they want to search for, they can activate the filter by picking up the sphere with the desired event code object and placing it in the filter pool below the sphere. The filter pool provides an overview of all activated filters, where only events with the corresponding CAMEO code are displayed on the globe. Deactivation is just as easy as activation by moving the sphere out of the frame of the filter pool. Unlike classic drop-down menus, the filter pool is an innovative approach to filtering in VR, allowing the user to keep

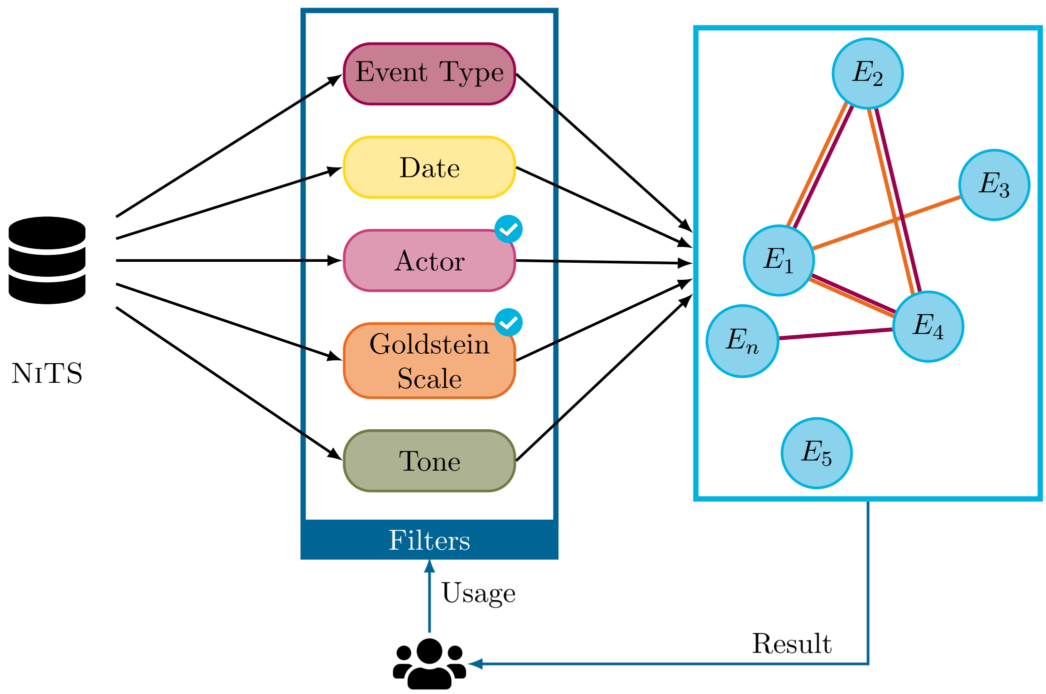

Figure 2.: NiTS implements five explicit filters to select the data to be displayed according to the user's preferences (see figures 3d–3f). Filtering ensures that only articles about events that match either all or one of the selected filters (✓) are displayed (see also Figure 3g).

track of all active filters (Xiao et al., 2020). As soon as the filter pool is used and the parameters are changed, the front-end triggers a request to the back-end. The back end, implemented in Java using *Java Spark* (Wendel and Löfdahl, 2023) as a RESTful web service, queries the data stored in a MongoDB and returns the relevant information. In addition, an interface with *OpenApi3* (Miller et al., 2021) was specified to facilitate data preparation and make the collected data available for visualisation.

In addition to the Event Type Filter, users have the option to filter by actor. In GDELT, actors refer to the entities or parties involved in events, that is, organizations, countries, political groups, businesses, or individuals mentioned in news articles or other media sources relating to an event. An event can have up to two actors: a source actor and a target actor, representing the entities involved in the event. The source actor is the initiating party or the origin of the event while the target actor is the party or entity the event pertains to. For instance, the source actor can be a group leading a protest. The target actor would be the government or specific individuals the protest is directed against. By analysing such actors, it is possible to study event structures at the level of social networks. For this purpose, NiTS provides the Actor Filter (see Figure 3f), which, like the Event Type Filter, is located on the right side of the globe. Mirroring the latter, the Actor Filter comprises spheres with relevant items corresponding to actors (such as national flags representing countries). Dragging a sphere into the filter pool will activate the filter, which will only show events if at least one of their actors matches the sphere placed in the pool. This operation maps a logical *or* (see Figure 3g). Since events can involve two actors,

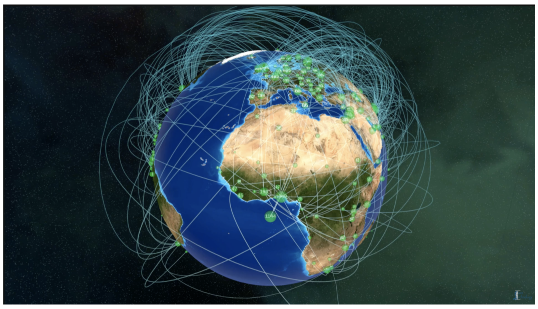

(a) NiTS globe with events.

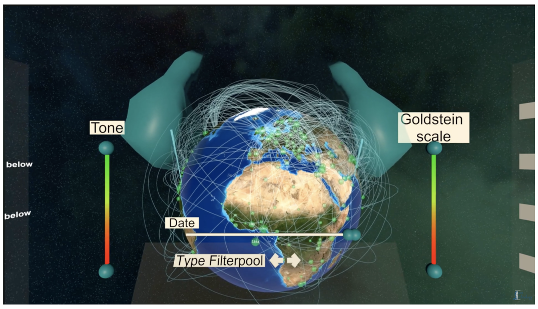

(b) Scaling the globe up and down by moving hands together and apart.

(c) Further zoom in to explore events.

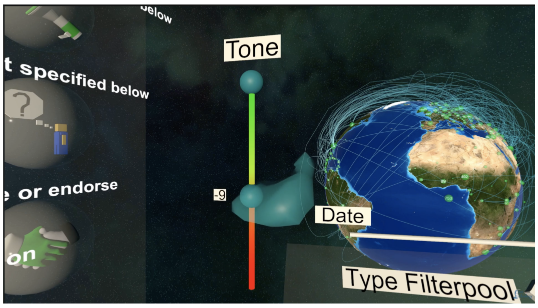

(d) Adjusting the tone slider to see how the event in perceived in the media.

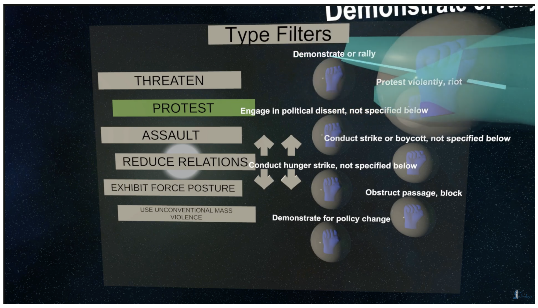

(e) Browse through the event type filters.

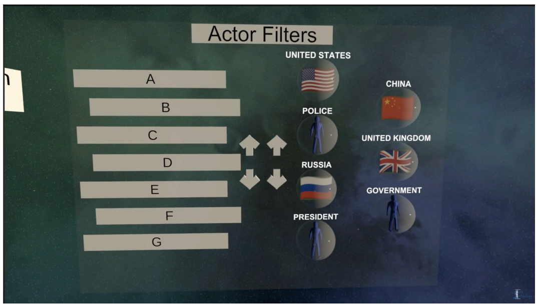

(f) Select the actors you want to be displayed such as countries.

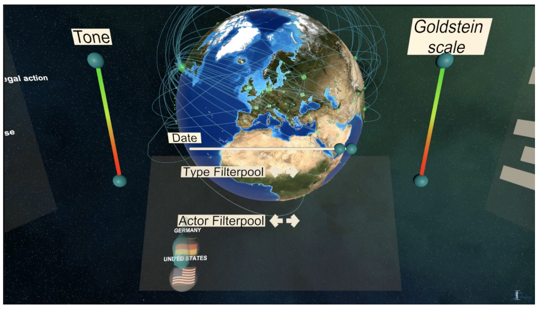

(g) Here you have selected Germany or the United States.

(h) You can browse through different events getting detailed information and the corresponding news article. Only events involving Germany or the United States are displayed.

Figure 3.: Virtual walkthrough of NiTS.

users need to be able to filter events that involve both. To achieve this, the user can pick up two actor spheres, one in each hand, and then connect both hands . This action merges the two spheres into one, which can then be placed in the filter pool. In this way, we implement an *and* operation. To visualise event networks or corresponding document networks, NiTS integrates a line renderer component that connects events if they share actors. This makes it easy to see which events are most closely connected and which countries share which events, or to observe areas with many connections as well as regions with a relatively low density of connections. On this basis, users can carry out analyses, e.g. focusing on events involving particular actors and event types. The visual representation of interconnected events on NiTS's globe enhances their data-driven analysis by facilitating the identification of trends or patterns in particular areas or regions, especially when these are manifested by large amounts of news. This type of document processing goes beyond the current capabilities of LLMs and is additionally based on the extended interface possibilities of VR. In addition to the Actor Filter and the Event Type Filter, we introduced two more filters to enrich event exploration: the Tone Filter and the Goldstein Scale. Tone refers to the sentiment associated with a specific event and helps to understand the connotation of an event described in news articles or other media sources. Tone can vary widely, including positive, negative, neutral or mixed sentiments – for the notion of sentiment as part of linguistic meaning see Bender and Lascarides (2020). To set the Tone Filter, a slider on the left-hand side of the globe can be used, ranging from green, indicating positive sentiment, through yellow, for neutral or mixed sentiment, to red, representing negative sentiment (see Figure 3d). For example, a news article about a peaceful diplomatic agreement might have a positive tone, suggesting optimism and cooperation. An article about a violent conflict, on the other hand, could have a negative tone, reflecting tension and hostility. By looking at tone, users can assess events based on the sentiment of the texts describing them and gain insight into how the event is perceived or portrayed in the media. In addition to the Tone Filter, we introduced a slider to adjust the *Goldstein Scale*, positioned on the right-hand side of the globe, operating in a similar manner. The Goldstein scale, widely utilized in political science and conflict studies, assigns numerical values to events, indicating the potential effect they could have on the stability of a nation. Events with higher *Goldstein scale* values are considered to have a more significant impact on a country's political stability, since the scale is designed to measure potential consequences. For instance, diplomatic events or agreements might have a positive *Goldstein scale* value, indicating a stabilizing influence: For instance, the statement "Promise Material support" with a *Goldstein scale* value of 5.2 suggests that the underlying event has a moderate or high impact on stability. Negative scale values suggest potential destabilization. For example, the statement "Military attack; clash; assault" with a Goldstein value of -10 indicates a highly significant and profoundly destabilising event. Events with little deviation from zero do not have a significant impact on stabilization or destabilization and fall within the yellow spectrum of the slider. For instance, the statement "Deny an accusation" having a Goldstein Value of -0.9 suggests that this event is perceived as having relatively minor impact on political stability (Goldstein, 1992). The Goldstein Scale Filter allows users to categorise events according to their expected impact on political processes. In this way, NiTS opens up another semantic criterion for displaying and filtering document networks. Beyond filtering, NiTS provides additional functionality for interacting with document content. This means that users can not only observe connected event networks, but also immerse themselves in individual events. By clicking on an event, a window opens with detailed information on this event. For this purpose, NiTS explores data from the Global Knowledge Graph (gkg) provided by GDELT. For example, topic markers, links to articles or to relevant images are displayed as additional information (see Figure 3h for an example).

Using the Actor Filter, the Event Type Filter, the Goldstein Scale Filter, the Time Filter and the Tone Filter, users can create very flexible filters to extract suitable document networks as hypertexts that can be traversed in VR, depending on their varying interests. The resulting networks, displayed on NiTS's globe, provide a visually mediated, diagrammatic perspective on the involvement of actors in events and their interrelationships. Users can utilise these representations to explore regions in relation to events, their actors, and event structures to conduct research on, for example, geopolitical dynamics or international relations. In this way, NiTS meets the criteria I, II and IV of section 1. Criterion III (non-intentional implicit filters as a result of user behaviour) is not yet realised. How NiTS addresses Criterion V (Natural Language Processing) is addressed in the next section.

3.2. The Architecture Behind **NiTS**

NiTS consists of two components: the backend and the visualization component. This section describes the architecture of the backend and the interfaces that are available to the client application for visualisation, modification and interaction, whereby communication is generally handled via REST. As shown in Figure 4, NiTS consists of several components that are responsible for collecting, transforming, analysing, mapping, providing and aggregating event data, as well as visualising and interacting with it. NiTS is implemented as a multi-user interface which, with the help of Ubiq (Friston et al., 2021), allows users to collaborate and communicate with each other. This includes collaborative interaction with NiTS's information units, which are retrieved from its backend via REST. By encapsulating the database connection, defining and providing the REST endpoint and implementing the periodic retrieval, mapping and importing of the required data sources, NiTS's backend forms the central component of the entire application. In the backend, implemented in Java, a strong focus is placed on lightweight and component-based implementation, which means as a result that Java Spark (Wendel and Löfdahl, 2023) is used to serve as a REST endpoint; it enables a straightforward reuse via the OpenAPI standard (Miller et al., 2021), even independently of the frontend. A snapshot of the latest available data sources is periodically imported by performing an interface-based mapping in NiTS, which enables an effortless extension to integrate additional data sources. The database management system chosen is MongoDB, a distributed, replicable and scalable database solution due to its document-based and schema-free approach.

The imported data is mapped to the database schema of NiTS so that it can be queried via the REST endpoint. Currently, GDELT-annotated data is used for filtering, selecting and visualizing events. Post-processing routines are required to mobilize additional data. This concerns NLP methods for named entity recognition (e.g. Honnibal et al., 2020), text similarity measurement (e.g. Viji and Revathy, 2022), topic modeling (e.g. Klamm et al., 2022) or sentiment analysis (e.g. Camacho-collados et al., 2022). This also includes methods for analyzing large amounts of text, which are often difficult to reuse, especially when implemented in different programming languages. This can be done with the help of Docker Unified UIMA Interface (DUUI – Leonhardt et al. 2023), which combines NLP processes in pipelines and executes them

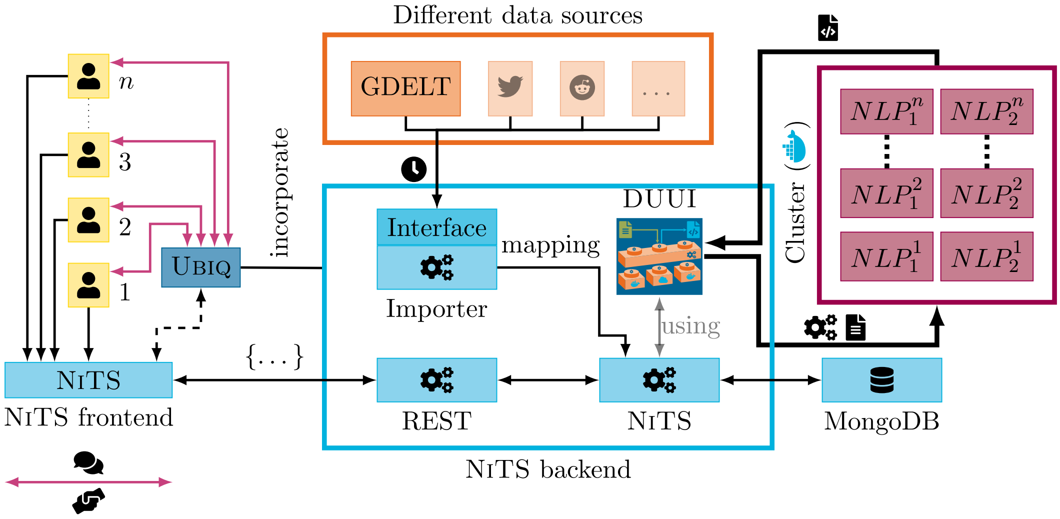

Figure 4.: The architecture of NiTS has a multi-level structure. Here a distinction is made between the frontend and the backend, in which operational and interactive processes occur as well as being mapped. This involves communication and data exchange between the two components via REST. For users, the frontend is the primary interface which allows them to talk to each other and interact with objects and each other collaboratively using VR glasses supplemented by an Ubiq extension Mehler et al. (2023). However, the heart of NiTS is the backend, which provides all REST functionalities and provides an interface-based importer to map and render various data sources for NiTS. Various data sources can therefore be integrated into NiTS, provided an implementation is available. Processed data are stored in a MongoDB database. In addition, an automated analysis of the existing texts based on linguistic features could be supported with the help of Docker Unified UIMA Interface (DUUI- Leonhardt et al. (2023) however, this has not yet been integrated.

in scalable clusters. DUUI uses microservices to instantiate processes as independent Docker images, which can be used independently of programming languages, with *UIMA* Ferrucci et al. (2009) serving as the annotation framework. In this way, text corpora can be processed in a reproducible manner where all annotations are stored in the underlying database for reuse (cf. Abrami and Mehler, 2018).

3.3. From Events to News about these Events and their Interactive Contextualization with **NiTS**

NiTS is a system for integrating and interacting with news data (van Dijk, 1988) in VR. As such, it serves as a bridge between data about (real) events, their description in (press) texts and the information needs of users. NiTS can be seen as a walk-in hypertext visualisation machine that enables its users to ground information read in selectable texts in a virtually reconstructed world view. The grounding concerns representations of points in time, spaces and actors. This grounding can be collaborative in the sense that several users interact simultaneously with the same virtual hypertext in order to be informed about the same events, people, places and their relations, and to start a joint discussion based on the information thus obtained. From a cognitive perspective, NiTS influences the construction of a potentially shared (Pickering and Garrod, 2004) situation model (Zwaan, 2016) based on a multiple text (List and Alexander, 2017) presented as a spatial hypertext (Conklin, 1987; Francisco-Revilla and Figueira, 2012; Marshall and Shipman III, 1993).

To systematize the information structure on which this approach is based, we utilize situation theory (Barwise and Perry, 1983; Devlin, 1991; Lalmas and van Rijsbergen, 1992; Seligman and Moss, 1997). That is, we start from basic types, such as the type of temporal locations, spatial locations, actors involved in actions, and the topics that are dealt with in order to form so-called infons $\langle\!\langle R, a_1, \ldots, a_n, w \rangle\!\rangle$ as elementary building blocks of information according to which a certain relation R exists ($w = 1$) or does not exist ($w = 0$) between the instances $a_1, \ldots, a_n$ of the latter types at certain time-space locations. Infons can be elementary (i.e., defined around the attribution of a single relation) or compound (i.e., defined by the conjunctive or disjunctive combination of several elementary infons). In parameterized infons, at least one argument of the underlying relation(s) is a typed variable. Parameterized infons are *supported* by *situation types*, while non-parameterized infons are *supported* by *situations*. A situation s is said to support an infon $i = \langle\!\langle R, a_1, \ldots, a_n, w \rangle\!\rangle$ if the objects $a_1, \ldots, a_n$ stand ($w = 1$) in the relation R in s or not ($w = 0$). Since *situation* is an abstract term, we need to define how situations can be accessed in the application context of NiTS. To do this, we take advantage of the fact that situations are described by texts or other media (images, audio or video) and use this as our access point. Thus, a text x is said to support an infon i if the relation instantiated in i is described in x in relation to the argument-forming entities of i. In the same way we can refer to images or videos as supporting infons. Since infons can be parameterised, we arrive at the notion of situation *types* which, in the context of NiTS, serve as abstract filters for the creation of hypertextual views on news (or, more generally, on textually manifested information). Given a (possibly parameterized) infon i, a *text type* is defined by means of a situation parameter $\dot{s}$: $[\dot{s} \mid \dot{s} \vDash i]$ (Devlin, 1991), that is, the type of texts describing i.[3] In the same way, we may speak about image, audio or video types. A *filter in the sense of* NiTS used to generate a view on a text database (e.g. GDELT) is then defined as an infon (parameterized or not). While such a filter selects a set of texts that support the underlying infon, their networking occurs according to constraints, which in the case of NiTS are detected by means of invariants across situations. More precisely, two texts are considered to be connected if there is an instance of one of the above types occurring in both of them. In this way, we can start from a filter to select a set of texts that support the corresponding infon(s), and then link these texts according to shared actors. The same can, in principle, be achieved by referring to shared elements of other types (e.g. topic, sentiment, or any other kind of information that can be attributed to or observed in texts).

Take a step back and consider what we have gained from this sketch of the information structure covered by NiTS: Suppose we have two (parameterized or non-parameterized) filters with which we select two different sets of text, each of which supports one of these filters as an access point to corresponding situations or situation types. For each set of texts, we then obtain a separate text network by recourse to the concept of an invariant: two texts of any set are linked if they describe the same invariant of the type now selected to filter links. So far, NiTS has implemented this at the level of persons. However, we can also think of the two sets A, B of texts as modes of a bipartite graph and link them according to the same principle. For example, if A

[3] Note that we do not define situation types in any other way. To think of something concrete, imagine a situation type as an abstract representation of a set of texts describing the content specified by a particular parameterized infon.

is a set of texts that deal with a certain topic from a certain perspective x (e.g. pro), and B is a set of texts that deal with the same topic from another perspective (e.g. contra), we arrive at a contrastive view of the same superordinate topic, and see how A and B are or are not networked in space and time according to shared invariants (e.g. opinion leaders, disseminators, etc.). The opportunities opened up by the further development of NiTS in this respect are fourfold:

(1) First and foremost, we can imagine more than two perspectives to span polypartite graphs that show the diversity with which a given topic is treated.
(2) The second extensions then concerns the variability of the concept of invariant, which can include all the types differentiated above, i.e., in addition to persons, variables such as places and times, but also text attributes such as sentiment and topic.
(3) The third extension concerns the fact that this set of variables can be extended to any text attribute that is the subject of natural language processing (NLP) methods. This concerns attributes such as claim (Salek Faramarzi et al., 2023), misinformation (Ferrara, 2022), stance (Küçük and Can, 2020) of the respective author, trustworthiness (Castillo et al., 2011), bias, toxicity (Arora et al., 2023), or any other text attribute that computational linguistics aims to detect automatically.
(4) The fourth extension concerns the aspect that the condition of co-occurrence in texts to be linked can be abandoned in favour of a similarity-oriented approach, which again makes use of machine learning. This means that two texts can be linked, for example, even if they do not necessarily contain the same but similar topics (Mehler et al., 2020), if they deal with related claims or if (parts of) one of the texts entails (parts of) the other (Dagan et al., 2013).

In this way, NiTS becomes a system for visualizing intertextual relations in the context of the spatio-temporal localization of the underlying texts. This provides the user with a spatial orientation for grounding (inter)textual content – provided that it is correctly annotated in the underlying text base. This grounding is subject to the virtualization of interaction spaces and can therefore be connected to a wide range of application scenarios, which will be discussed in the next section.

4. Use Cases

4.1. Critical online reasoning based on simulation-based learning with VR-based multiple documents

In subsection 3.3 we argued that NiTS can be seen as a tool that supports the collaborative interpretation of multiple documents through shared situation models. In this sense, NiTS qualifies as a means of improving simulation-based learning in VR. Simulation-based learning is a method in which participants practise dealing with problem situations, especially when practising in real environments is difficult (Baker et al., 2009; Hollan et al., 1984; Ziv et al., 2000). VR is a suitable environment for this kind of learning as it enables the simulation of application scenarios. It allows the flexible, parameter-controlled change of (e.g. spatial, temporal, social or semiotic) conditions to test their effect on learning. Consequently, there have been approaches to simulation-based learning in the context of educational science to investigate the processing of multiple documents (List and Alexander, 2017). An example is Va.Si.Li-

Lab (Mehler et al., 2023), which has been developed to study learning processes in social and educational environments. Multiple documents (Britt et al., 2012; Perfetti et al., 1999) are the result of non-linear reading processes, especially in the Web, where segments of different texts are combined to produce sequences for which coherence relations cannot necessarily be assumed. Since reading is practically shifting to the Web, this type of reading can be considered the default case, even in the field of higher education (Yuan et al., 2022). What makes the VR-based study of multiple documents interesting is the possibility of interacting by means of information systems that largely capture the multimodal data generated by the interactions involved. In this way, learning becomes the subject of multimodal machine learning (Zhang et al., 2023). This is where NiTS comes into play, as it goes beyond the dissemination of news in the Web by explicitly allowing the selection of reading criteria by collaborating readers – see the notion of a filter in subsection 3.3. It should be noted that the selection process itself can be collaborative in NiTS, as can the subsequent process of traversing the generated text networks from which the multiple documents are ultimately drawn. Thus, by embedding tools such as NiTS in Va.Si.Li-Lab, simulation-based learning is given a considerable degree of freedom with regard to the study of reading processes based on multiple documents. In this way, the concept of hypertext undergoes a predicted but so far hardly realised extension towards the combination of spatial hypertexts (Marshall and Shipman III, 1993), VR-based interactive learning and multimodal openness (Mehler et al., 2023), which together provide access to investigating collaborative information processing.

4.2. Political Discourse, Peace and Conflict Studies

Organizations around the world, including NGOs, development cooperatives and journalistic research departments, often produce press clippings from local media. This involves time-consuming tasks, such as copying links from news reports. Exploring and collecting a wide range of news can help people stay up to date. The further development of NiTS may allow the extraction of such press clippings. This can help save time and increase the efficiency and accuracy of message monitoring processes.

NiTS offers comprehensive topic overviews and the ability to collaborate in VR with partner organizations. Moreover, GDELT's extension Translingual translates news from over 65 languages. Since the vast majority of news comes from non-English language sources, this extension is beneficial for international organizations trying to understand contexts within countries through the lens of local media that would normally need to be translated. For example, conflicts in the global South (such as the war in Yemen) are not necessarily adequately covered by news agencies in the global North (Bachman and Ruiz, 2023). Thematic preferences arise here, as a result of which conflicts are underreported or even overlooked (Brenner and Han, 2021). It is necessary to have access to a wide range of local media in order not to overlook conflicts anywhere in the world , whereby gaining access to different perspectives beyond one's own perspective is an added value.

Scientists and policy makers could use an advanced version of NiTS to monitor and analyze conflicts (including their spatial distribution and development patterns). Decision makers can benefit from bringing together information from different sources (Keertipati et al., 2014). Analysts could use NiTS to monitor events and stability indicators across countries, which could help to assess political risks and prepare decisions. The combination of different filters by means of NiTS creates representations of past events that allow comparisons with related events.

Understanding political processes is often difficult for individuals. There are many reasons for this: political processes are complex and often have legal frameworks that must be viewed from multiple perspectives and interpreted according to the context. Political decisions are often influenced by divergent interests and agendas, which may be political, economic, social, or religious. Through the consolidation of such information, the integration of NLP methods and the possibility of 3D topic-related mapping with NiTS, political texts (e.g. from parliamentary debates – see GerParCor (Abrami et al., 2022) or Hansard (Rayson et al., 2015)) can be made interactively accessible and more transparent. However, as with any tool, there is a risk of adopting a closed view if the underlying data is not questioned or its origin is not sufficiently checked.

5. Outlook and Future Work

The current focus of NiTS on GDELT can be overcome by integrating other types of sources. Databases with a similar objective include ACLED (*Armed Conflict Location and Event Data Project*), ICEWS (*Integrated Crisis Early Warning System*), Cline Center Phoenix Historical Data and UCDP (*Uppsala Conflict Data Program*). These event-oriented databases with different focuses would be a worthwhile extension of NiTS. The challenge would be to combine different conflict data (Raleigh and Kishi, 2019) and to integrate different coding schemes that deviate from the CAMEO taxonomy (Raleigh et al., 2010). The inclusion of sources beyond these databases could also be beneficial. For example, this could refer to geo-located social media data (e.g., tweets from micro-blogging services).

Based on the information from GDELT, events can be assigned to three types of locations: countries, regions, and cities. Currently, cities represent the most detailed level of aggregation, as we lack more precise coordinates of the locations where the events took place. Refining the spatial coordinates could help to draw a more detailed picture of the respective event network. This can be achieved using NLP methods that analyze news articles to determine location data. Matching this information with geographic databases such as Geonames or Open Street Map can then provide the desired coordinates, which in turn would allow a more accurate drawing of the event location on NiTS's globe.

NiTS is already a highly interactive tool, but some features that are crucial for an effective immersive VR application are not yet implemented (see Section 1). This concerns, among other things, the lack of integration of language models that control, for example, avatars that act as interaction partners for users. In order to combine the advantages of both systems, NiTS will be integrated into Va.Si.Li-Lab: in this way, it will eventually fulfill criterion IV (extension to mixed reality), criterion V (regarding the usability of the full range of NLP), and especially criterion VI (use of LLMs).

References

Giuseppe Abrami, Mevlüt Bagci, Leon Hammerla, and Alexander Mehler. 2022. German Parliamentary Corpus (GerParCor). In *Proceedings of the Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 1900–1906.

Giuseppe Abrami and Alexander Mehler. 2018. A UIMA Database Interface for Managing

NLP-related Text Annotations. In *Proceedings of the 11th edition of the Language Resources and Evaluation Conference, May 7 - 12 (LREC 2018)*. European Language Resources Association (ELRA), Miyazaki, Japan, 6 pages. https://aclanthology.org/L18-1212/

Giuseppe Abrami, Alexander Mehler, Mevlüt Bagci, Patrick Schrottenbacher, Alexander Henlein, Christian Spiekermann, Juliane Engel, and Jakob Schreiber. 2023. Va.Si.Li-Lab as a Collaborative Multi-User Annotation Tool in Virtual Reality and Its Potential Fields of Application. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media* (Rome, Italy) *(HT '23)*. Association for Computing Machinery, New York, NY, USA, Article 22, 9 pages. https://doi.org/10.1145/3603163.3609076

Giuseppe Abrami, Alexander Mehler, Christian Spiekermann, Attila Kett, Simon Lööck, and Lukas Schwarz. 2020. *Educational Technologies in the area of ubiquitous historical computing in virtual reality.* Taylor & Francis, Chapter 16, 240–261. https://www.routledge.com/New-Perspectives-on-Virtual-and-Augmented-Reality-Finding-New-Ways-to-Teach/Daniela/p/book/9780367432119

Arnav Arora, Preslav Nakov, Momchil Hardalov, Sheikh Muhammad Sarwar, Vibha Nayak, Yoan Dinkov, Dimitrina Zlatkova, Kyle Dent, Ameya Bhatawdekar, Guillaume Bouchard, and Isabelle Augenstein. 2023. Detecting Harmful Content on Online Platforms: What Platforms Need vs. Where Research Efforts Go. *ACM Comput. Surv.* 56, 3, Article 72 (oct 2023), 17 pages. https://doi.org/10.1145/3603399

Jeffrey Bachman and Esther Ruiz. 2023. The geopolitics of human suffering: a comparative study of media coverage of the conflicts in Yemen and Ukraine. *Third World Quarterly* 0, 0 (07 2023), 1–19. https://doi.org/10.1080/01436597.2023.2228715

Suzanne C Baker, Ryan K Wentz, and Madison M Woods. 2009. Using virtual worlds in education: Second Life® as an educational tool. *Teaching of Psychology* 36, 1 (2009), 59–64.

Jon Barwise and John Perry. 1983. *Situations and Attitudes.* MIT Press, Cambridge.

Emily M. Bender and Alex Lascarides (Eds.). 2020. *Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Semantics and Pragmatics.* Morgan & Claypool Publishers, San Rafael.

Jan Bongard, Tobias Elßner, Xuke Hu, and Jens Kersten. 2023. The Geography of Conflicts: A Comparative Assessment of Web Data. In *Proceedings of the 8th ACM SIGSPATIAL International Workshop on Security Response Using GIS* (Hamburg, Germany) *(EM-GIS '23)*. Association for Computing Machinery, New York, NY, USA, 18–23. https://doi.org/10.1145/3615884.3629424

Petter Bae Brandtzaeg and Asbjørn Følstad. 2017. Trust and Distrust in Online Fact-Checking Services. *Commun. ACM* 60, 9 (aug 2017), 65–71. https://doi.org/10.1145/3122803

David Brenner and Enze Han. 2021. Forgotten Conflicts: Producing Knowledge and Ignorance in Security Studies. *Journal of Global Security Studies* 7 (09 2021). https://doi.org/10.1093/jogss/ogab022

M Anne Britt, Jean-François Rouet, and Jason LG Braasch. 2012. Documents as entities: Extending the situation model theory of comprehension. In *Reading-from words to multiple texts*, M Anne Britt, Susan R Goldmann, and Jean-François Rouet (Eds.). Routledge, 161–179.

Jose Camacho-collados, Kiamehr Rezaee, Talayeh Riahi, Asahi Ushio, Daniel Loureiro, Dimosthenis Antypas, Joanne Boisson, Luis Espinosa Anke, Fangyu Liu, and Eugenio Martínez Cámara. 2022. TweetNLP: Cutting-Edge Natural Language Processing for Social Media. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations.* Association for Computational Linguistics, Abu Dhabi, UAE, 38–49. https://aclanthology.org/2022.emnlp-demos.5

Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information Credibility on Twitter. In *Proceedings of the 20th International Conference on World Wide Web* (Hyderabad, India) *(WWW '11)*. Association for Computing Machinery, New York, NY, USA, 675–684. https://doi.org/10.1145/1963405.1963500

Vinton G. Cerf. 2023. Large Language Models. *Commun. ACM* 66, 8 (jul 2023), 7. https:

//doi.org/10.1145/3606337

Daeheon Cho and Bo Ae Chun. 2019. Virtual Reality as a New Opportunity in Geography Education: From the Teachers' Perspectives in Korea. In *Proceedings of the 2019 5th International Conference on Education and Training Technologies* (Seoul, Republic of Korea) *(ICETT '19)*. Association for Computing Machinery, New York, NY, USA, 140–145. https://doi.org/10.1145/3337682.3337701

Rocio Chongtay, Mark Last, and Bettina Berendt. 2018. Responsive News Summarization for Ubiquitous Consumption on Multiple Mobile Devices. In *23rd International Conference on Intelligent User Interfaces* (Tokyo, Japan) *(IUI '18)*. Association for Computing Machinery, New York, NY, USA, 433–437. https://doi.org/10.1145/3172944.3172992

J. Conklin. 1987. Hypertext: An Introduction and Survey. *Computer* 20, 9 (1987), 17–41.

Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto. 2013. *Recognizing textual entailment: Models and applications.* Morgan & Claypool Publishers, San Rafael.

Keith Devlin. 1991. *Logic and information.* Cambridge University Press, Cambridge.

Yasyn Elyusufi and El Mokhtar En-Naimi. 2017. User Profiling and RSS Feeds Recommendation Based on the Ontological Approach. In *Proceedings of the 2nd International Conference on Computing and Wireless Communication Systems* (Larache, Morocco) *(ICCWCS'17)*. Association for Computing Machinery, New York, NY, USA, Article 10, 8 pages. https://doi.org/10.1145/3167486.3167496

Emilio Ferrara. 2022. Twitter spam and false accounts prevalence, detection, and characterization: A survey. *First Monday* (2022). https://doi.org/10.5210/fm.v27i12.12872

David Ferrucci, Adam Lally, Karin Verspoor, and Eric Nyberg. 2009. Unstructured Information Management Architecture (UIMA) Version 1.0. OASIS Standard. https://docs.oasis-open.org/uima/v1.0/uima-v1.0.html

Luis Francisco-Revilla and Alvaro Figueira. 2012. Adaptive Spatial Hypermedia in Computational Journalism. In *Proceedings of the 23rd ACM Conference on Hypertext and Social Media* (Milwaukee, Wisconsin, USA) *(HT '12)*. ACM, New York, NY, USA, 313–314. https://doi.org/10.1145/2309996.2310050

Sebastian J Friston, Ben J Congdon, David Swapp, Lisa Izzouzi, Klara Brandstätter, Daniel Archer, Otto Olkkonen, Felix Johannes Thiel, and Anthony Steed. 2021. Ubiq: A System to Build Flexible Social Virtual Reality Experiences. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology* (Osaka, Japan) *(VRST '21)*, Yuichi Itoh, Kazuki Takashima, Parinya Punpongsanon, Misha Sra, Kazuyuki Fujita, Shigeo Yoshida, and Piumsomboon. Tham (Eds.). Association for Computing Machinery, New York, NY, USA, Article 6, 11 pages. https://doi.org/10.1145/3489849.3489871

Julian Gagel, Jasper Hustedt, Timo Lüttig, Theresa Berg, Giuseppe Abrami, and Alexander Mehler. 2023. News in Time and Space: Global Event Exploration in Virtual Reality. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media* (Rome, Italy) *(HT '23)*. Association for Computing Machinery, New York, NY, USA, Article 7, 3 pages. https://doi.org/10.1145/3603163.3609080

Deborah Gerner, Rajaa Jabr, and Philip Schrodt. 2002. Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for the Analysis of Foreign Policy Interactions. *International Studies Association, New Orleans* (2002), 35 pages.

Ioannis Giannopoulos, Andreas Komninos, and John Garofalakis. 2017. Natural Interaction with Large Map Interfaces in VR. In *Proceedings of the 21st Pan-Hellenic Conference on Informatics* (Larissa, Greece) *(PCI '17)*. Association for Computing Machinery, New York, NY, USA, Article 56, 6 pages. https://doi.org/10.1145/3139367.3139424

Joshua S. Goldstein. 1992. A Conflict-Cooperation Scale for WEIS Events Data. *The Journal of Conflict Resolution* 36, 2 (1992), 369–385. http://www.jstor.org/stable/174480

James D Hollan, Edwin L Hutchins, and Louis Weitzman. 1984. STEAMER: An interactive inspectable simulation-based training system. *AI magazine* 5, 2 (1984), 15–15.

Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python. https://doi.org/"10.5281/zenodo.1212303"

Cengizhan Ipbuker. 2004. Numerical Evaluation of the Robinson Projection. *Cartography and Geographic Information Science* 31 (2004), 79 – 88. https://api.semanticscholar.org/CorpusID:120568092

Xavier Jose, S.D Madhu Kumar, and Priya Chandran. 2021. Characterization, Classification and Detection of Fake News in Online Social Media Networks. In *2021 IEEE Mysore Sub Section International Conference (MysuruCon)*. IEEE, Virtual, 759–765. https://doi.org/10.1109/MysuruCon52639.2021.9641517

Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. 2020. Unity: A General Platform for Intelligent Agents. arXiv:1809.02627 [cs.LG]

J. M. Jurado, A. Graciano, L. Ortega, and F. R. Feito. 2017. Web-Based GIS Application for Real-Time Interaction of Underground Infrastructure through Virtual Reality. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (Redondo Beach, CA, USA) *(SIGSPATIAL '17)*. Association for Computing Machinery, New York, NY, USA, Article 97, 4 pages. https://doi.org/10.1145/3139958.3140004

Swetha Keertipati, Bastin Tony Roy Savarimuthu, Maryam Purvis, and Martin Purvis. 2014. Multi-Level Analysis of Peace and Conflict Data in GDELT. In *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis* (Gold Coast, Australia QLD, Australia) *(MLSDA'14)*. Association for Computing Machinery, New York, NY, USA, 33–40. https://doi.org/10.1145/2689746.2689750

Christopher Klamm, Ines Rehbein, and Simone Paolo Ponzetto. 2022. FrameASt: A Framework for Second-level Agenda Setting in Parliamentary Debates through the Lense of Comparative Agenda Topics. In *Proceedings of the Workshop ParlaCLARIN III within the 13th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 92–100. https://aclanthology.org/2022.parlaclarin-1.13

Dilek Küçük and Fazli Can. 2020. Stance detection: A survey. *ACM Computing Surveys (CSUR)* 53, 1 (2020), 1–37.

Sebastian Lague, Farès Sioni, Rodrigo Appelt, Davide Marcoli, EvilsSouls, Malte Jürgens, Matic Jeseničnik, MillersMan, Pattrigue, Racooder, Stefano Serioli, and tposejank. 2023. SebLague/Geographical-Adventures. https://github.com/SebLague/Geographical-Adventures

Mounia Lalmas and Cornelis J. van Rijsbergen. 1992. Situation Theory and Dempster-Shafer's Theory of Evidence for Information Retrieval. In *Proceedings of the Workshop on Incompleteness and Uncertainty in Information Retrieval. Concordia University, Montreal.* Springer London, London, 102—-116.

Kalev Leetaru and Philip A Schrodt. 2013. GDELT: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, Vol. 2. Citeseer, International Studies Association, San Francisco, California, 1–49.

Alexander Leonhardt, Giuseppe Abrami, Daniel Baumartz, and Alexander Mehler. 2023. Unlocking the Heterogeneous Landscape of Big Data NLP with DUUI. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Association for Computational Linguistics, Singapore, 15 pages. accepted.

Gabriel A. León-Paredes, Omar G. Bravo-Quezada, Erwin J. Sacoto-Cabrera, Oscar F. Pizarro-Gordillo, Paúl E. Vintimilla-Tapia, Jack F. Bravo-Torres, and Wilson P. Cabrera-Chica. 2020. Virtual reality and data analysis based platform for urban mobility awareness as a tool for road education. In *2020 IEEE ANDESCON*. IEEE, Quito, Ecuador, 1–6. https://doi.org/10.1109/ANDESCON50619.2020.9272084

Rhema Linder, Alexandria M. Stacy, Nic Lupfer, Andruid Kerne, and Eric D. Ragan. 2018. Pop the Feed Filter Bubble: Making Reddit Social Media a VR Cityscape. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, Reutlingen, Germany, 619–620. https://doi.org/10.1109/VR.2018.8446271

Alexandra List and Patricia A Alexander. 2017. Analyzing and integrating models of multiple text comprehension. *Educational Psychologist* 52, 3 (2017), 143–147.

Danielle Lottridge, Katie Quehl, Frank Bentley, Max Silverman, Melissa Ong, Michael Dickard, Brooke White, and Ethan Plaut. 2022. Ubiquitous News Experienced Alone: Interviews with Americans and Their Devices. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1, Article 128 (apr 2022), 29 pages. https://doi.org/10.1145/3512975

Chunyong Ma, Yongyang Qi, Yong Chen, Yong Han, and Ge Chen. 2008. VR-GIS: An Integrated Platform of VR Navigation and GIS Analysis for City/Region Simulation. In *Proceedings of The 7th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry* (Singapore) *(VRCAI '08)*. Association for Computing Machinery, New York, NY, USA, Article 16, 6 pages. https://doi.org/10.1145/1477862.1477883

Catherine C. Marshall and Frank M. Shipman III. 1993. Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext. In *Proceedings of the Fifth ACM Conference on Hypertext*. ACM, Seattle Washington, USA, 217–230.

Alexander Mehler, Mevlüt Bagci, Alexander Henlein, Giuseppe Abrami, Christian Spiekermann, Patrick Schrottenbacher, Maxim Konca, Andy Lücking, Juliane Engel, Marc Quintino, Jakob Schreiber, Kevin Saukel, and Olga Zlatkin-Troitschanskaia. 2023. A Multimodal Data Model for Simulation-Based Learning with Va.Si.Li-Lab. In *Digital Human Modeling and Applications in Health, Safety, Ergonomics and Risk Management*, Vincent G. Duffy (Ed.). Springer Nature Switzerland, Cham, 539–565.

Alexander Mehler, Rüdiger Gleim, Regina Gaitsch, Tolga Uslu, and Wahed Hemati. 2020. From Topic Networks to Distributed Cognitive Maps: Zipfian Topic Universes in the Area of Volunteered Geographic Information. *Complexity* Special Issue on Cognitive Network Science: A New Frontier, 4 (2020), 1–47. https://doi.org/10.1155/2020/4607025

Darrel Miller, Jeremy Whitlock, Marsh Gardiner, Mike Ralphson, Ron Ratovsky, and Uri Sarid. 2021. OpenAPI Specification v3.1.0. https://spec.openapis.org/oas/v3.1.0. Accessed: 2023-06-01.

Logan Molyneux. 2018. Mobile news consumption: A habit of snacking. *Digital journalism* 6, 5 (2018), 634–650. https://doi.org/10.1080/21670811.2017.1334567

Andrew Moran, Vijay Gadepally, Matthew Hubbell, and Jeremy Kepner. 2015. Improving Big Data visual analytics with interactive virtual reality. In *IEEE Conference on High Performance Extreme Computing (HPEC)*. IEEE, Waltham, MA USA, 1–6. https://doi.org/10.1109/HPEC.2015.7322473

Charalampos Nikolaou, Kallirroi Dogani, Konstantina Bereta, George Garbis, Manos Karpathiotakis, Kostis Kyzirakos, and Manolis Koubarakis. 2015. Sextant: Visualizing time-evolving linked geospatial data. *Journal of Web Semantics* 35 (2015), 35–52. https://doi.org/10.1016/j.websem.2015.09.004 Geospatial Semantics.

Emily Norton. 2023. Beyond Hypertexting the Hypertext: Annotated and GIS Adaptations of Joyce's Ulysses as Case Studies for User Experience and Engagement. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media* (Rome, Italy) *(HT '23)*. Association for Computing Machinery, New York, NY, USA, Article 15, 6 pages. https://doi.org/10.1145/3603163.3609051

Kaya Okada, Mitsuo Yoshida, Takayuki Itoh, Tobias Czauderna, and Kingsley Stephens. 2018. Spatio-Temporal Visualization of Tweet Data around Tokyo Disneyland Using VR. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces Companion*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3180308.3180356

Pieter-Jan Ombelet, Aleksandra Kuczerawy, and Peggy Valcke. 2016. Employing Robot Journalists: Legal Implications, Considerations and Recommendations. In *Proceedings of the 25th International Conference Companion on World Wide Web* (Montréal, Québec, Canada) *(WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 731–736. https://doi.org/10.1145/2872518.2890093

Eric Balagtas Perez, James King, Yugo H. Watanabe, and Xiang 'Anthony' Chen. 2020. Counterweight: Diversifying News Consumption. In *Adjunct Proceedings of the 33rd An-*

nual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) *(UIST '20 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 132–134. https://doi.org/10.1145/3379350.3416154

Charles A Perfetti, Jean-François Rouet, and M Anne Britt. 1999. Toward a theory of documents representation. In *The construction of mental representations during reading*, H. van Oostendorp and S. R. Goldman (Eds.). Erlbaum, Mahwah, NJ, 99–122.

Martin J. Pickering and Simon Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27 (2004), 169–226.

Clionadh Raleigh and Roudabeh Kishi. 2019. Comparing Conflict Data: Similarities and differences across conflict datasets. In *ACLED*. ACLED.

Clionadh Raleigh, Andrew Linke, Håvard Hegre, and Joakim Karlsen. 2010. Introducing ACLED: An Armed Conflict Location and Event Dataset: Special Data Feature. *Journal of Peace Research* 47 (09 2010), 651–660. https://doi.org/10.1177/0022343310378914

Paul Rayson, Alistair Baron, Scott Piao, and Steve Wattam. 2015. Large-scale time-sensitive semantic analysis of historical corpora. In *Proceedings of the 36th Meeting of ICAME, Trier, Germany*. Trier, Germany.

Noushin Salek Faramarzi, Fateme Hashemi Chaleshtori, Hossein Shirazi, Indrakshi Ray, and Ritwik Banerjee. 2023. Claim Extraction and Dynamic Stance Detection in COVID-19 Tweets. In *Companion Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) *(WWW '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 1059–1068. https://doi.org/10.1145/3543873.3587643

Andrés Santos, Telmo Zarraonandia, Paloma Díaz, and Ignacio Aedo. 2018. A Virtual Reality Map Interface for Geographical Information Systems. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces*. Association for Computing Machinery, New York, NY, USA, 3 pages. https://doi.org/10.1145/3206505.3206580

Thomas Schedel and Claus Atzenbeck. 2016. Spatio-Temporal Parsing in Spatial Hypermedia. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media* (Halifax, Nova Scotia, Canada) *(HT '16)*. Association for Computing Machinery, New York, NY, USA, 149–157. https://doi.org/10.1145/2914586.2914596

Pascal Schneiders, Edgar L. Kist, and Birgit Stark. 2022. *Media Use*. Springer Fachmedien Wiesbaden, Wiesbaden, 1–28. https://doi.org/10.1007/978-3-658-34048-3_57-2

Philip A Schrodt. 2012. CAMEO: Conflict and Mediation Event Observations Event and Actor Codebook. http://data.gdeltproject.org/documentation/CAMEO.Manual.1.1b3.pdf

Jerry Seligman and Lawrence S. Moss. 1997. Situation Theory. In *Handbook of Logic and Language*, John van Benthem and Alice ter Meulen (Eds.). Elsevier, Amsterdam, 239–309.

Teun A. van Dijk. 1988. *News as Discourse*. Erlbaum, Hillsdale.

D. Viji and S. Revathy. 2022. A hybrid approach of Weighted Fine-Tuned BERT extraction with deep Siamese Bi – LSTM model for semantic text similarity identification. *Multimedia Tools and Applications* 81, 5 (01 Feb 2022), 6131–6157. https://doi.org/10.1007/s11042-021-11771-6

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.

Xuebin Wei and Xiaobai Yao. 2022. A GIS representation framework for location-based social media activities. *Transactions in GIS* 26 (05 2022). https://doi.org/10.1111/tgis.12929

Per Wendel and Love Löfdahl. 2023. Spark - A micro framework for creating web applications in Kotlin and Java 8 with minimal effort. http://sparkjava.com. Accessed: 2023-06-01.

Jordan Wirfs-Brock and Katie Quehl. 2019. News From the Background to the Foreground: How People Use Technology To Manage Media Transitions: A Study of Technology-Mediated News Behaviors in a Hyper-Connected World. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3, Article 110 (sep 2019), 19 pages. https://doi.org/10.1145/3351268

Jacob O. Wobbrock. 2020. Hypertext, Social Media, and Civic Engagement: How Hypertext is Ruining the World, and Might Just Save It. In *Proceedings of the 3rd Workshop on Human Factors in Hypertext* (Virtual Event, USA) *(HUMAN '20)*. Association for Computing Machinery, New York, NY, USA, Article 2, 1 pages. https://doi.org/10.1145/3406853.3432660

Wei-Ye Xiao, Jun-Zou Duan, and Cheng-Qi Xue. 2020. Research on Visual Programming System of Gesture Interaction Interface in VR. In *Proceedings of the 2020 4th International Conference on Big Data and Internet of Things* (Singapore, Singapore) *(BDIOT '20)*. Association for Computing Machinery, New York, NY, USA, 38–42. https://doi.org/10.1145/3421537.3421554

Tangming Yuan, Siyuan Ji, and Guoxiang Zhong. 2022. The Exploration of the Future Teaching Mode in Post-Pandemic Higher Education. In *Proceedings of the 7th International Conference on Distance Education and Learning* (Beijing, China) *(ICDEL '22)*. Association for Computing Machinery, New York, NY, USA, 222–227. https://doi.org/10.1145/3543321.3543358

Xiao Zhang, Shamim Akhter, Abdelmohsen A. Nassani, and Mohamed Haffar. 2022. Impact of News Overload on Social Media News Curation: Mediating Role of News Avoidance. *Frontiers in Psychology* 13 (2022), 12 pages. https://doi.org/10.3389/fpsyg.2022.865246

Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. 2023. Meta-Transformer: A Unified Framework for Multimodal Learning. arXiv:2307.10802 [cs.CV]

Amitai Ziv, Stephen D. Small, and Paul Root Wolpe. 2000. Patient safety and simulation-based medical education. *Medical teacher* 22, 5 (2000), 489–495.

Rolf A Zwaan. 2016. Situation models, mental simulations, and abstract concepts in discourse comprehension. *Psychonomic bulletin & review* 23, 4 (2016), 1028–1034.