



Exploring Descriptions of Movement Through Geovisual Analytics

Scott Pezanowski¹ · Prasenjit Mitra¹ · Alan M. MacEachren^{1,2}

Received: 15 December 2021 / Accepted: 31 January 2022 / Published online: 24 February 2022
© Deutsche Gesellschaft für Kartographie e.V. 2022

Abstract

Sensemaking using automatically extracted information from text is a challenging problem. In this paper, we address a specific type of information extraction, namely extracting information related to descriptions of movement. Aggregating and understanding information related to descriptions of movement and lack of movement specified in text can lead to an improved understanding and sensemaking of movement phenomena of various types, e.g., migration of people and animals, impediments to travel due to COVID-19, etc. We present GeoMovement, a system that is based on combining machine learning and rule-based extraction of movement-related information with state-of-the-art visualization techniques. Along with the depiction of movement, our tool can extract and present a lack of movement. Very little prior work exists on automatically extracting descriptions of movement, especially negation and movement. Apart from addressing these, GeoMovement also provides a novel integrated framework for combining these extraction modules with visualization. We include two systematic case studies of GeoMovement that show how humans can derive meaningful geographic movement information. GeoMovement can complement precise movement data, e.g., obtained using sensors, or be used by itself when precise data is unavailable.

Keywords Geographic movement · Geovisual analytics · Machine learning · Natural language processing · Big data analytics

Erforschung von Bewegungsbeschreibungen durch geovisuelle Analytik

Zusammenfassung

Sensemaking mit automatisch extrahierten Informationen aus Text ist ein anspruchsvolles Problem. In diesem Beitrag befassen wir uns mit einer bestimmten Art der Informationsextraktion, nämlich der Extraktion von Informationen im Zusammenhang mit Bewegungsbeschreibungen. Das Sammeln und Verstehen von Informationen in Bezug auf Beschreibungen von Bewegung und Bewegungsmangel, die in Texten angegeben sind, kann zu einem verbesserten Verständnis und einer Sinngebung von Bewegungsphänomenen verschiedener Art führen, z. B. Migration von Menschen und Tieren, Reisebehinderungen aufgrund von COVID-19 usw. Wir präsentieren GeoMovement, ein System, das auf der Kombination von maschinellem Lernen und regelbasierter Extraktion bewegungsbezogener Informationen mit modernsten Visualisierungstechniken basiert. Neben der Bewegungsdarstellung kann unser Tool Bewegungsmangel extrahieren und darstellen. Es gibt nur sehr wenige Vorarbeiten zum automatischen Extrahieren von Bewegungsbeschreibungen, insbesondere von Negation und Bewegung. Abgesehen davon, dass GeoMovement diese anspricht, bietet GeoMovement auch ein neuartiges integriertes Framework zur Kombination dieser Extraktionsmodule mit Visualisierung. Wir schließen zwei systematische Fallstudien von GeoMovement ein, die zeigen, wie Menschen aussagekräftige geografische Bewegungsinformationen ableiten können. GeoMovement kann präzise Bewegungsdaten ergänzen, die z. B. mithilfe von Sensoren gewonnen werden, oder allein verwendet werden, wenn keine präzisen Daten verfügbar sind.

✉ Scott Pezanowski
scottpez@brightworldlabs.com
Prasenjit Mitra
pum10@psu.edu
Alan M. MacEachren
maceachren@psu.edu

¹ Information Sciences and Technology, The Pennsylvania State University, Westgate Building, University Park, PA 16802, USA

² Department of Geography, The Pennsylvania State University, Walker Building, University Park, PA 16802, USA

1 Introduction

Automated methods proposed by the natural language processing and information retrieval communities often form the basic building blocks in an application. However, in this paper, we argue that such automated tools, even though they have achieved some level of maturity, are not enough for the needs of the end-users, especially for domains that require higher-level information assimilation and cognition like foraging and sensemaking over spatial information. For example, Lai et al. (2022) have recently used natural language processing (NLP) to understand context in the extraction and geocoding of historical floods, storms, and adaptation measures. They extend the state-of-the-art for low-level information extraction, e.g., named entity extraction and geocoding, but do not provide a holistic understanding of the story underlying these events.

We posit that our research community needs to “see the forest for the trees.” Sensemaking is an integral part of information processing, and tighter coupling between the lower levels (information extraction) and the higher levels (information understanding and sensemaking) can improve the state-of-the-art. Specifically, we call for the community to look more at “higher-level tools and systems” that enable end-users to complete tasks. Towards this goal, we study the case of extracting geospatial information from text using visual analytics (VA) (Andrienko et al. 2020; Yuan et al. 2021) to perform tasks over the extracted data.

Since text is unstructured data and the information within the text is often messy, the output from computational techniques includes associated errors. Therefore, it is not sufficient to explore mentions of movement in text without human expertise. VA can address this issue through human-in-the-loop strategies that enable analysts to work iteratively with computational methods that extract knowledge from messy data, cope with uncertainties in computational results, and improve those results over time (Endert et al. 2014; Robinson 2017). VA is especially suitable for big, diverse, messy data that can be interpreted differently (Tapia-McClung and Silván-Cárdenas 2021; Angelini et al. 2018; Ninkov and Sedig 2019; Snyder et al. 2020; MacEachren et al. 2011).

Our research objective was to determine if computational techniques and geovisual analytics can leverage *large volumes of movement statements to enable an end-user to understand the movement described* quickly. If successful, research can then take advantage of the wealth of movement data found in written descriptions about people, wildlife, goods, and other things moving throughout our world. Text statements about movement can be used to understand what is moving, when it is moving, why it is moving, and how it is moving.

For our research, “geographic movement” refers to the movement of people, animals, objects, goods, information,

natural physical processes, and similar things through spaces ranging in size from multiple buildings to the whole earth. We applied computational methods to identify and extract movement statements, and present them in GeoMovement, a human-in-the-loop web-based geovisual analytics system for identification, processing, and exploration of descriptions of movement. GeoMovement involves computational (1) cleaning of the messy text, (2) predicting the statements that describe movement using a machine learning (ML) model, (3) applying Geographic Information Retrieval (GIR) techniques to identify places mentioned, and (4) predicting statements that describe restricted movement or desired movement that is not possible (hereafter referred to as “impaired movement”). While there is substantial research on some of these subtasks, *integrating these techniques with VA* and demonstrating its success in our chosen domain is the main contribution of this paper.

While some progress has happened in processing descriptions of movement in text, there is very little work on detecting and understanding descriptions of impaired movement. For example, the COVID-19 pandemic prompted us to focus on impaired movement due to the importance of disruptions and restrictions in global movement patterns of people and, perhaps equally important, the movement of goods like food and medicines. Therefore, another important contribution of this work is that we show an adaptation of an existing approach developed to detect negation in picture descriptions (van Miltenburg et al. 2016) can successfully uncover impaired geographic movement in text documents.

Specifically, we integrated (and adapted or extended) many existing computational and VA methods to produce a system that supports information foraging related to geographic movement as reflected in text statements. GeoMovement is unique in identifying movement statements and filtering them by place and time.

Figure 1 shows GeoMovement’s user interface. Users can search and filter based on search terms, the statements’ dates, and impaired movement status. The statements originate from three sources ingested into GeoMovement to demonstrate its capabilities and utility for investigating movement:

- 398 thousand News articles from August 2019 to May 2020,
- 328 million Twitter tweets from February 2020 to May 2020,
- 15.6 thousand Scientific articles from August 2019 to November 2020.

Over 520 million total statements contain diverse movement patterns, things moving, geographic coverage, and temporal differences.

Existing geographic movement research has improved analysis methods (Dodge et al. 2012; Dodge 2016a; Dodge

et al. 2016; Dodge 2016b; Graser et al. 2021, 2020, 2020; Soares Junior et al. 2017; Huang 2017) and shown how these methods can derive valuable information about human movement and wildlife movement (Wang et al. 2020a; Dodge et al. 2014; Miller et al. 2019; Zhu et al. 2021; Li et al. 2021). However, partly due to the challenges and because computational techniques to address them are relatively new, this prior research focused on geographic movement in precise movement trajectories from sensors such as GPS and mostly ignored movement described in text documents. The limited existing research to analyze movement described in text has only focused on narrow tasks like mapping route descriptions. This work fills a gap in the research to support the study of geographic scale movement by developing and demonstrating methods for analyzing movement described in text documents at a broad scale. We anticipate that our effort will spur more research to leverage this underutilized geospatial movement information in text documents.

We present two case studies (and a third case study in a supplementary video demonstration) related to the global pandemic that demonstrates the potential of our approach and show, despite the messy data and imprecise computational predictions, a human can derive meaningful and

essential information about geographic movement described in text with GeoMovement. The case studies also demonstrate that GeoMovement can support multi-scale information foraging through vast volumes of messy text statements about geographic movement. Furthermore, the space-time concept/attribute filtering methods implemented effectively narrow in on information relevant to an analyst's objectives.

2 Related Work

First, we review related efforts to analyze various geographic movement types described in text documents. Existing research into studying descriptions of geographic movement has primarily focused on reconstructing route descriptions. These routes (driving directions Jaiswal et al. 2010; Drymonas and Pfoser 2010), hiking and general route descriptions (Moncla et al. 2014b, a, 2016; Piotrowski et al. 2010), and historical exploration routes (Bekele et al. 2016; Blank and Henrich 2015) form a constrained subset of movement statements that simplify and thus do not address many of the challenges with a broader set of movement statements. Other research took the opposite approach of textualizing

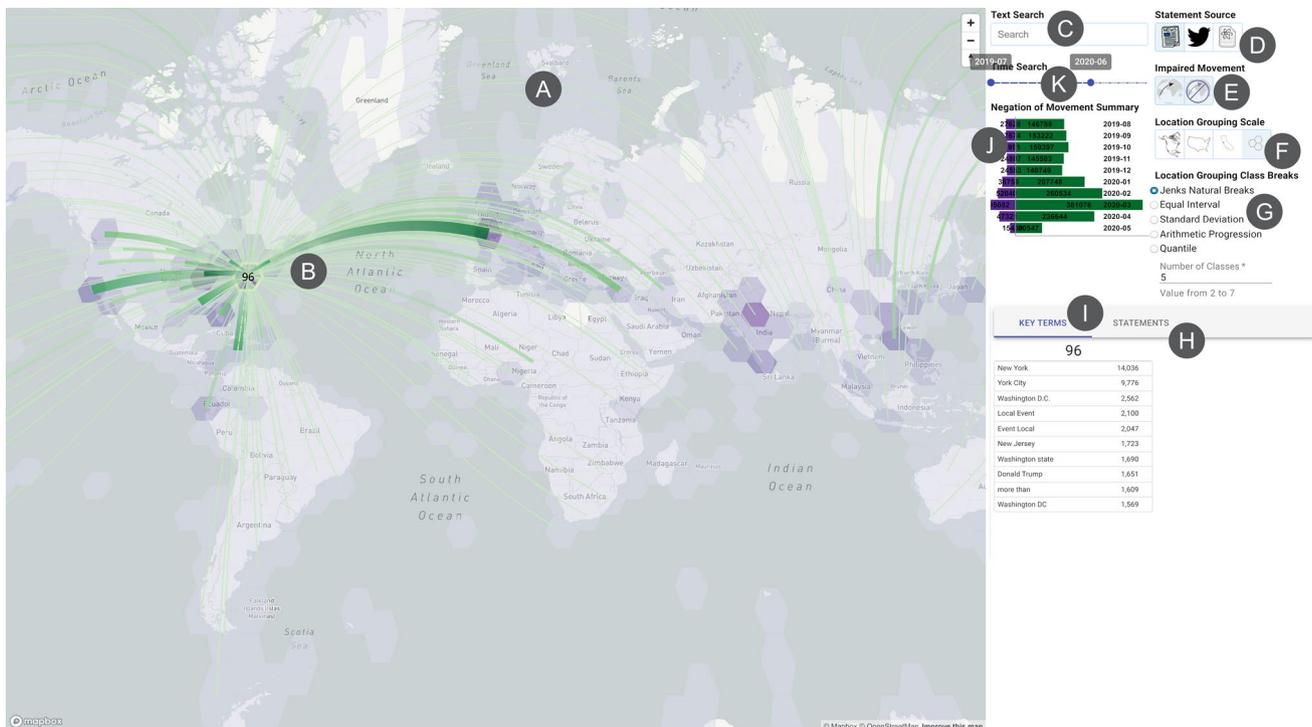


Fig. 1 The GeoMovement user interface: **A** Main mapping interface with hexagons selected for place aggregation. **B** Connection lines between co-occurring place mentions from an area near New York City selected. **C** Free text key term search. **D** Text data source buttons. **E** Normal movement vs. impaired movement buttons. **F** Geo-binning buttons. **G** Classification method for bins and connection

lines colors and the number of classes chooser. **H** Statements that match the search. **I** Most common bi-grams that match the search. **J** Two-sided Temporal Bar Chart that shows the number of statements, by date, with a divided bar depicting normal movement on the right (green) and negated movement on the left (purple) **K** Time range slider

(convert routes to description) to take advantage of the benefits of text Chu et al. 2014; Al-Dohuki et al. 2017). Additionally, Huang et al. (2020) showed how geovisual analytics could improve the retrieval of trajectories in a search. But, their focus on text analytics was only on users' queries of the data (which was precise sensor-based trajectories). Furthermore, in addition to addressing narrow domains, most of this research has restricted the data to a small amount of text.

The GeoCAM project (Jaiswal et al. 2010) created an application that identified, extracted, and generated maps of route directions found on webpages providing textual (often formatted) directions to reach a location. While the GeoCAM project complements our current research, their problem is simpler since route directions are just a small subset of movement descriptions. Furthermore, route directions typically follow a semi-structured pattern that allows for simpler ML models and rules-based approaches. In a precursor to the GeoCAM project, Drymonas and Pfoser (2010) used similar techniques to map route directions. These projects encouraged future work like ours to go beyond route directions to general movement descriptions (Klippel et al. 2008).

Second, we describe related efforts that use geovisual analytics on geographic movement described in text documents. The complicated nature of analyzing place in text documents, especially the need to represent spatial relationships best rendered visually on a map, has prompted other researchers to take a geovisual analytics approach. SensePlace was a system to analyze place mentions in text and pull information from other sources to aid analysis (Tomaszewski et al. 2011). SensePlace2 developed and applied geovisual analytics to methods that focus on the extent to which Twitter users' tweet location compared to the places they discuss in their tweets (MacEachren et al. 2011). Robinson et al. (2013) performed a user study with experts that showed both the advantages, usefulness, and difficulties of such a system for crisis management. SensePlace3 extended this effort by advancing the geovisual analytics techniques and scaling the system to work on millions of tweets per month, thereby improving analysis (Pezanowski et al. 2017).

The SMART system complements the SensePlace versions' focus by providing a visual interface enabling human analysts to explore text's spatial, temporal, and topical components (Snyder et al. 2020; Karimzadeh et al. 2019b). SMART implemented advanced geovisual analytics techniques, including a tweet classifier to filter semantically and a cluster lens to visualize keywords at a large scale. However, like SensePlace3, their system did not focus on analyzing movement. Finally, the NewsStand system (Teitler et al. 2008; Samet et al. 2020) mapped places where news articles are written compared to the places they discuss but focuses less on geovisual analytics for analysis and more on correctly mapping the text.

A few other efforts use geovisual analytics and mapping systems to analyze places mentioned in tweets (Thom et al. 2012; Bosch et al. 2013, 2011; Felmler et al. 2020) and show the potential to take advantage of this geographic data source, albeit mainly focusing on tweets with a geocoded location that makes the challenge different. Mapping and geovisualizations have also been combined with topic analysis and network graph analysis to show similarities between cities (Hu et al. 2017).

Jamonnak et al. (2020) combined location information associated with videos and the narrations of those videos to show their locations on the map and the topics and sentiment discussed at those locations. Xu et al. (2018) created a system to explore Yelp business reviews in areas and their change over time. Ma et al. (2020) showed that geovisual analytics is vital to understanding critical local places that need immediate help in a disaster from 911 call transcripts and clusters of specific crime types from police reports. Although these systems successfully demonstrated geovisual analytics on text, they focused on particular topics and did not consider movement. Therefore, they could not be applied to our goal of analyzing wide-ranging types of movement described in text. Since impaired movement detection is not the primary focus of this research, we describe work related to it in Sect. 3.3.

3 Text Computational Processing Predictions

We first acquired three different sets of documents consisting of 398 thousand news articles, 328 million tweets, and 15.6 thousand scientific articles. Our document sources and the keyword and time parameters used to obtain them are described in detail in Appendix B.1. We cleaned the documents using typical text pre-processing methods and applied computational techniques to identify places in the text, predict the statements that describe movement, and predict statements that describe impaired movement.

3.1 Predicting Movement

In Pezanowski and Mitra (2020)'s work, humans label the statements with a binary class as either describing geographic movement or not. Since they took this initial step and created a corpus to train a model, we can use this model to predict statements that describe movement. The prediction of this ML model is a probability value between zero (no movement) and one (movement). We set an arbitrary cut-off for GeoMovement to only show statements with a probability greater than 0.6 that the statements are about movement. This relatively loose threshold was arrived at by trial and error and is acceptable for all three sources.

3.2 Predicting Geographic Location

We used the GeoTxt system to perform GIR on our text sources (Karimzadeh et al. 2013, 2019a). We chose this because GeoTxt performs comparably or better than other state-of-the-art geoparsers and performs best without case sensitivity, which is common in Twitter data (one of our data sources) (Gritta et al. 2018). A small evaluation corpus had an F1-score of 0.78 for geoparsing place names and an accuracy of 0.91 in resolving those place names correctly. Moreover, it performed even better on higher-order administrative places such as countries and states, which are common in our statements. For enterprise projects, paid commercial sources also exist from software companies such as Esri, Google, and Microsoft. However, we chose not to use these products because they do not allow for customization compared to GeoTxt, which is open-source software and allows adjustments to the software in the future.

3.3 Predicting Impaired Movement

The global pandemic of 2020–2021 brought attention to global movement and how it spreads, and how the pandemic disrupted or prevented regular global movement. Because the pandemic highlighted the importance of analyzing disruptions to movement, we investigated potential strategies for detecting statements about impaired movement.

To detect impaired movement in our statements, we looked to adapt existing methods of negation detection in text. Negation detection strategies can potentially uncover statements about formal restrictions on movement (of the sort imposed by governments), decisions not to move taken by individuals for their safety, and impediments to movement created by limited public/commercial transport such as canceled flights due to the lack of passengers or ill crew.

In this section, first, we review existing related research on negation detection and its everyday use cases. Second, we describe how we adapted an existing approach from the literature that detects negation in picture descriptions (van Miltenburg et al. 2016) to our challenge of detecting impaired geographic movement described in text. Third, we show how we improved upon our initial attempt to detect impaired movement using our geovisual analytics methods to analyze initial mistakes in predictions and then modify the rules specifically to detect impaired movement.

3.3.1 Existing Approaches to Negation Detection

Addressing negation in text has been identified as a problem in several existing works. For example, Fialho et al. (2020) had remarked that “when a negation was involved in a sentence, the classifiers found more difficult to return the appropriate label” in the context of negation in sentences

as identified as part of discourse representation structures. Negation detection in text is vital in challenges like automated summarization of medical reports (Vincze et al. 2008; Slater et al. 2021), summarizing picture descriptions (van Miltenburg et al. 2016), and as a hint in identifying sarcasm (Reyes and Rosso 2014). Hiremath and Patil (2021) show that sarcasm detection depends upon detecting negative sentences in positive situations and positive sentences in negative situations.

Much of the current state-of-the-art research on negation detection was influenced by a Workshop titled *Resolving the Scope and Focus of Negation* (Morante and Blanco 2012), which also produced labeled datasets that continue to be used in training and evaluating the success of new methods. Supervised ML-based solutions such as the LSM Network (Zhao et al. 2021) have learned negative terms while performing sentiment mining automatically from large-scale training data. The current state-of-the-art method, NegBert, is based on ML (Khandelwal and Sawant 2020; Khandelwal and Britto 2020). Although this ML approach is the current state-of-the-art, the challenge in using ML approaches is the need for time-consuming labeling of large amounts of training data. NegBert was trained and tested on datasets designed explicitly for negation detection evaluation and therefore could not be used for our tangential challenge of detecting impaired movement. In the absence of training data, we show that rules-based approaches can still be used.

van Miltenburg et al. (2016) and van Son et al. (2016) have used rules to detect negation. They had humans annotate Flickr picture descriptions for negations and define categories of negations. This annotation exercise produced simple clues for negation, thereby allowing their rules-based approach to be effective on picture descriptions. In general, rules-based approaches have been proven to work in negation detection when the domain is relatively narrow and, like most rule-based systems, typically provides high precision but low recall. Rule-based methods can work fine in our application, where a sample of the negative sentences suffices, but having false positives can result in incorrect conclusions.

3.3.2 Applying Negation Detection to Descriptions of Geographic Movement

Detecting impaired movement is similar to previous negation detection using specific key terms. However, what constitutes a negated word is ill-defined and varies from domain to domain and problem to problem since terms, like *canceled* or *diverted*, would not always be considered negated. But, when applied to movement, they are. This ambiguity complicates the task. We investigated if we could adapt the current methods that are focused on the negation of words (ex. She does *not* have cancer. The alarm clock did *not* have the

feature I wanted.) to detect impaired geographic movement (ex. Our flight to England was *canceled*. The fruit was *stuck* in Brazil because of initial fears early in the pandemic.).

We used van Miltenburg et al. (2016)'s rules-based approach for negation detection. As discussed above, we adapted this rules-based approach, as opposed to an ML approach, to (1) avoid costly labeling of training data, (2) achieve transparency in how the results are obtained, and (3) because our needs in detecting impaired movement are relatively narrow, which the literature suggests (van Miltenburg et al. 2016; van Son et al. 2016; Slater et al. 2021) is a good fit for a rules-based approach.

van Miltenburg et al. (2016)'s method tags part-of-speech in text (such as verbs, nouns, adjectives, etc.) and then searches for a list of negation keywords, prefixes, and suffixes to detect negations. Some negation rules examples are a) exact negated words like “no” and “not,” b) verbs that start with “de,” “mis,” “dis,” or c) adjectives that end with “less.” Overall, all rules are relatively simple and are easily understandable and reproducible. If the input sentence matches a rule, it contains negation.

We first applied van Miltenburg et al. (2016)'s exact rules for negation to our statements previously predicted to describe movement (as described in Section 3.1) to predict impaired movement. After 800,000 movement statements were predicted, we stopped and obtained a summary count. The model predicted about 28% of these movement statements as describing impaired movement.

We selected a stratified random sample of 50 predicted impaired movement statements and 50 predicted normal movement statements. The initial model did not do very well to predict impaired movement correctly. There were 23 true positives where the statement was correctly predicted as impaired movement compared to 27 false positives where the statement was predicted to be impaired movement, but it was normal movement. It did slightly better to correctly predict normal movement with 42 true negatives as opposed to eight false negatives. Table 1 shows the confusion matrix for these predictions. These values equate to a precision of 0.46, a recall of 0.74, an F1-score of 0.57, and an accuracy of 0.65 on the stratified sample.

3.3.3 Improving the Detection of Impaired Movement

To improve our detection of impaired movement, we either added new rules or removed existing rules. Since there were more false positives than false negatives, this hints that rules should mostly be removed so that fewer statements are predicted to have impaired movement.

Not only did we analyze the errors in statement predictions, but we also consulted a list online of verbs common to movement see Appendix A and <https://archiewahwah.wordpress.com/2019/04/16/movement-verbs-list/>). First,

we removed rules specific to verbs that had the prefixes “de,” “mis,” and “dis.” These prefixes could rarely negate the verbs we found related to movement. Second, based on an online list of adjectives related to movement (see Appendix A and https://dSPACE.ut.ee/bitstream/handle/10062/18053/adjectives_and_adverbs_of_movement.html), we removed the adjectives' rules beginning with “a” and “dis” because they are much too broad and did not make sense with those commonly related to movement. Third, we added the specific lemmas “cancel,” “postpone,” “prevent,” and “avoid” because they relate to impaired movement. Lemmas allow us to match many different synonyms of these words that all mean the same.

These rules modifications improved predictions on the existing stratified random sample with a precision of 0.74, a recall of 0.76, an F1-score of 0.75, and an accuracy of 0.79. Therefore, the F1-score improved from 0.57 to 0.75.

Since we used the first set of 100 sampled statements to determine some of our rules, it would not be fair to rely on these improved metrics alone to conclude we improved our method. The rules may over-fit the sample. Therefore, we selected a second unseen test set of 100 statements using the same random stratified technique. On this unseen set, the original rules' prediction F1-score was 0.59, while our new rules were again much better with a 0.65 F1-score. Overall, our minor modifications to the rules produced substantial improvement.

Finally, since we selected the two sets of 100 statements from new articles, as a final assessment, we sampled 100 statements from both the tweets and scientific articles using the same stratified random sampling technique. Predictions using our modified rules for impaired movement on the tweets resulted in an F1-score of 0.61 and an F1-score of 0.60 on scientific articles. Therefore, our predictions for impaired movement performed slightly less accurately for these two sources than news articles, but still respectable. This lower accuracy is likely because both sources contain language more common for that audience (i.e., slang and other informal languages in tweets and technical language in scientific articles). Based on our experience modifying the rules for news articles, we estimate that additional rule changes would also improve predictions for tweets and scientific articles. However, based on our experience and the literature, we surmise that given resources to label a large amount of training data, using an ML approach as

Table 1 The confusion matrix for prediction of negated movement with unmodified rules from general negation

	Actual impaired	Actual normal
Predicted impaired	23 (TP)	27 (FP)
Predicted normal	8 (FN)	42 (TN)

in Khandelwal and Sawant (2020) would likely produce more accurate predictions. In summary, we are using a rules-based approach as a proof-of-concept prediction that is important to show the potential benefits of detecting impaired movement statements and use this as an attribute to analyze geographic movement described in text.

4 Geovisual Analytics to Find Meaning in Descriptions of Movement

GeoMovement is a web-based geovisual analytics system that allows users to explore descriptions of geographic movement. GeoMovement serves as an interface between the human and the data described above and is summarized in Fig. 1. The statements' content, place mentions, date of creation, and impaired movement prediction are all searchable. The map visualizes place mentions in multiple levels of aggregate geo-bins. A geo-bin is many smaller places (e.g., cities) aggregated and displayed as one larger place that they all reside within (e.g., country). Geo-bins allow for a clear summary of location-based data. By choosing particular places of interest on the map, the user can visualize the co-occurrences of places. The individual statements view completes the overview-first + detail approach. This workflow matches the information-seeking mantra of overview-first to gain an awareness of the information and details of interest on-demand (Shneiderman 1996). This section is divided into three subsections that follow the information-seeking mantra that starts with the overviews, then options for the user to search to filter to statements of interest, then detailed views of the filtered statements.

For detailed information on GeoMovement, Appendix B.1 describes the sources and nature of the data sources and statements. In addition, Appendix B.2 includes technical details of the application development that allow for fast user queries on large volumes of text, thereby enabling efficient sensemaking. An important technical component of our approach is our use of Elasticsearch (<https://www.elastic.co/>) as the primary information storage and retrieval software for GeoMovement. Elasticsearch is a search engine that accepts many search parameters like free text and time, and returns matching results ordered by most relevant to the user's search. In addition, it can group results by attributes like place mentions in the text. Most impressively, Elasticsearch does all of this and returns results very quickly, most often in the matter of milliseconds.

4.1 Overviews

There are two primary overview means to explore the movement descriptions. The first overview is the map (Fig. 1A) that displays the number of statements spatially aggregated

by their place mentions. Since we have a large number of statements from multiple sources, GeoTxt extracted over 98,000 unique place mentions from them. Displaying this large number of places on a map using points would likely be very confusing for users. Many places would overlap, and quickly seeing overall patterns and comparing quantities between places would be difficult. This is the primary reason for aggregating the statements place mentions. Polygon bins that contain many place mentions in the statements are colored darker than those with fewer place mentions. Users get a clear overview of place mentions in the statements and a way to visually compare places.

We chose five shapes to spatially aggregate the number of statements by their place mentions. The geo-bins scales include continent, country, administrative 1 (the worldwide equivalent of a state in the US), and two sizes of a hexagonal pattern. Each scale allows users to explore different types of movement data, like long-range bird migrations by continent, traded goods by country or administrative level, and detailed movement through the hexagons that do not adhere to political boundaries like the spread of disease. We chose hexagons as an aggregation shape since they tessellate and will distort values less than squares (Birch et al. 2007; Esri 2021). The user interface provides an option shown in Fig. 1F, where the user can choose the aggregate level.

The bin counts are divided into classes to ease visual comparisons between bins. The user has control of the bin count classification technique and the number of classes, as shown in Fig. 1G. The user has five choices on how to classify the aggregate counts. Jenks Natural Breaks, Equal Interval, Standard Deviation, Arithmetic Progression, and Quantile are options, and they can enter the number of classes between two to seven. These classification techniques are well-accepted statistical methods to make the data more understandable.

We used the geostats JavaScript library (<https://github.com/simogeo/geostats>) to calculate the class breakpoints for the chosen users' classification method. Each classification method is valuable depending on the user search and resulting data. Darker colored bins represent areas that have a larger number of place mentions within the bin. The color scheme is a sequential color scheme chosen from ColorBrewer (<https://colorbrewer2.org/>) (Brewer et al. 2003) to ensure the classes of statement counts are easily distinguishable.

The second overview in GeoMovement is a Two-sided Temporal Bar Chart. Research has shown differences in how people understand geographic movement, such as those who think more spatially than others (Liben and Downs 1993; Ishikawa 2016). Based on this research, GeoMovement provides the user with multiple views of the data. The Two-sided Temporal Bar Chart shows aggregate counts of statements grouped by the month they were

published (Fig. 1J). The oldest month is at the top and the youngest month is at the bottom. The Two-sided Temporal Bar Chart allows for visualization of the amount and changes over time, both in impaired movement on the left side and normal movement on the right side. When hovering the mouse over the chart, the number of statements that match the current user search are shown in bars in the foreground, while the total number of statements in GeoMovement are shown in bars in the background. Figure 6 shows an example of the number of statements that match the user search in the foreground once the user has hovered the cursor over the chart.

4.2 Search

After the overviews give the user an understanding of the data, they can begin filtering it through search options. The five key ways to search the data are (1) free-text search (Fig. 1C), (2) buttons to select from the three sources for statements — news articles, tweets, and scientific articles (Fig. 1D), (3) buttons to select impaired movement statements or normal movement statements (Fig. 1E), (4) a time-range slider (Fig. 1K), (5) and clicking location(s) on the map to filter by location. Multiple features can be chosen by holding the Ctrl-key on their keyboard and selecting the next feature with a mouse-click. After any of the first four searches are performed, the geo-bins and Two-sided Temporal Bar Chart update to show counts of statements that match the searches. After the fifth search, the detailed views appear.

4.3 Detailed Views

The detailed views of statements and their attributes include (a) connecting great circles drawn on the map between place mentions in the statements, and co-occurring place mentions in the same statement, (b) the five most common bi-grams for the set of statements matching the search, (c) and a table showing the actual statements that match the search. All connections between places are aggregated to show places commonly used together. The aggregate count classes adhere to the users' currently chosen map classification method and total class number in a sequential green color scheme that is color-blind friendly, also selected from ColorBrewer (Brewer et al. 2003).

It is important to note that the connection lines are completely accurate in showing movement between the places since they are drawn solely by the places' co-occurrence in a statement. For example, the statement below has three place mentions and, therefore, three possible movement pairs: Sydney–New York, Sydney–London, New York–London. The first two are correct concerning

movement references in the statement, but the third is not correct because there is no direct movement between New York and London. One could argue that the movement could be from New York to Sydney and then next to London from Sydney, but this is not probable. However, these connection lines give the user overall patterns of interest that they can confirm through inspection of the statements.

But this will be the first time a commercial flight is flown from Sydney to New York, and just the second time from Sydney to London, Qantas said (Garber 2019).

Second, the ten most common bi-grams are shown for the set of statements that match the search and have place mentions in the chosen location bin. This provides details about the actual statements behind the overviews by showing the most common words and topics in the selected statements. Since many of the statements contain place mentions and the search location is an essential parameter for the matching statements, two-word place names are often in this bi-grams list. These place mentions may be valuable, but we also found that it was often more important to see bi-grams about topics and not necessarily places when exploring the data. Therefore, we allow users to double-click a bi-gram to remove the bi-gram from the list, and the next most common bi-gram appears, up to the 20th most common bi-gram. To re-populate the bi-grams list, the user can re-run the search. The user can also select multiple bins by holding the Ctrl-key and clicking another map bin. Users can then compare the most common bi-grams for both map bins to see differences between statements with place mentions in each bin.

To complete the overview first + detail approach to analysis, once the user found sets of statements of interest in the overviews and chose a map bin, the user can see the actual statements matching the search in a paged list. The displayed statements match all search parameters (when selecting multiple bins, statements can have place mentions in either bin). The user can scroll through the pages of statements. Each statement's published date is also shown. If there is a particular statement of interest, holding the Ctrl-key while clicking on the statement will open the original document on the Web in a new browser window.

GeoMovement's tight use of modern search engine technologies and Information Retrieval (IR) allows for extremely fast human-in-the-loop sensemaking for the most relevant information on movement. Again, Appendix B provides further details, and our supplemental video showcases a real-time live demonstration of efficient knowledge discovery using GeoMovement.

5 GeoMovement Assessment

First, we show the challenge of interpreting movement described in statements without GeoMovement by discussing summary statistics of the data. Second, we present two case studies that show how our approach can retrieve information about movement from vast quantities of statements. A third case study is included in a supplemental video. It is recorded in real-time to show that a user can quickly extract meaningful information about movement despite the challenges posed by the large quantity of messy text, ambiguity in text descriptions, and imprecise computational predictions. Third, we show how these case studies also generated future GeoMovement needs. Finally, in Appendix C, we discuss the skill level and hardware and software requirements for GeoMovement users.

5.1 Illustrative Data Summary Statistics

To show the value of GeoMovement, we created summary statistics of the data to clearly illustrate how it is unreasonable to think a human can analyze and understand large quantities of movement statements without such a system. We chose three keywords related to our case studies and three prominent place names: one being a country, one a state, and one a city. The number of our statements that match these keywords and place names is shown in Table 2. To relate these statistics to our first case study in Sect. 5.2, we filter GeoMovement's 520 million statements (36 thousand of those contain the term *smuggling*, 275 thousand of them include the term *gold*, and 201 thousand mention *India*) and efficiently identify important gold smuggling patterns around and in India. To produce these statistics, we used our GIR extracted place names to determine the number of statements that contain each selected place name. And, to find the number of statements containing each of the keywords, we searched the statements in a Postgresql database using a full-text search (<https://www.postgresql.org/docs/14/textsearch.html>) so that different variations of the same word will be matched (e.g., *smuggling*, *smuggled*, *smuggle*).

This exercise to generate summary statistics is meant to show that without such a geovisual analytics system like ours, it would be extremely difficult, if not unfeasible, to perform geographic, temporal, and attribute sensemaking of the described movement. Although GeoMovement's 520 million statements are substantial, it is still a fraction of the

accessible text available that could be included and analyzed in GeoMovement, given more development and computational resources.

5.2 Case Studies

To confirm our claim that the geovisual analytics interface helps users understand and make sense from the statements and multiple computational predictions, we provide two detailed case studies below from different types of (prototypical, fictitious) potential users. The case studies presented provide evidence of usability. A third case study, given only in the video supplement (due to space limitations in the text), adds additional evidence about the flexibility and utility of GeoMovement to explore the mix of text data sources from different perspectives.

5.2.1 Understanding International Crime Affecting India

Jennifer Lang is a college student who wants to write a class report about different types of international crime affecting India. She opens her web browser to GeoMovement and types *smuggling* to begin her search. She notices that many statements involving smuggling are in October 2019, despite that month having fewer statements overall. She adjusts the time range slider to filter statements to that month and sees a hotspot of activity in England. She clicks the hexagon bin in England and is reminded of a significant human smuggling event in that month where many people lost their lives after being trapped in a truck that was smuggling them (Fig. 2). Although this is a significant smuggling event from a British perspective and also highlighted in the U.S. news, she decides to look for other ways to focus on India.

As her next step, she chooses to aggregate place mentions by the country level and select India. Figure 3 shows the results, and she sees that most results related to smuggling are affecting India from the neighboring countries of Pakistan and Bangladesh, and a few more countries. In the bigrams list, she sees many references to drugs and gold being smuggled into India.

Next, after reading some individual statements, she finds out more about multiple cases of gold smuggling from both outside and within India (Fig. 4). To get a more detailed analysis of the movement, she types *gold* in the search box and chooses to aggregate place mentions by the state level. By clicking the neighboring state of Sindh, Pakistan, which has many place mentions in the statements, connections with

Table 2 Summary statistics of the number of statements that contain selected key terms and place mentions plus the total number of statements

Smuggling	Gold	Sports	London	California	India	Total Statements
36 K	275 K	336 K	589 K	165 K	201 K	520 M

many states in India are highlighted, including a state in southeast India.

She clicks this state in India and discovers that the state is Tamil Nadu, where the large city of Chennai is located. From reading a few statements, she knows that gold smuggling in Chennai is arriving at the airport and through the Chennai Express train from Mumbai. The connection with Mumbai through the train is confirmed on the map by the strong connection with Mumbai's state, Maharashtra, on the west coast, closer to Pakistan (Fig. 5). She can now clearly visualize and report on some of the prominent drugs and gold smuggling sources into and throughout India and read more detailed descriptions about individual incidents.

5.2.2 Examining the Impact of the Pandemic on Travel for Tourism

Arti Reddy is a travel agent in India. She uses GeoMovement to understand the pandemic's impact on global travel and travel related to India. As of May 2020, like other countries in the World, India was dealing with a global pandemic. Since Arti had previously planned to advertise to potential customers traveling for sporting events, she chose this topic to investigate. She loads GeoMovement and enters the term *sports* in the text search box. With the geospatial hexagon

geo-bins as a layer, she quickly sees a hotspot of discussion in Japan. She selects to filter statements for impaired movement. While mousing over the Two-sided Temporal Bar Chart, she sees that the impaired movement statements are more prevalent in recent months (except May, where there is less data), as seen in the Two-sided Temporal Bar Chart in Fig. 6.

Arti switches to the country geo-binning, chooses to view both types of movement again, and moves the timeline from July 2019 to December 2019, and Japan is still a popular location. After selecting Japan on the map, she sees in the bi-grams list that the upcoming Olympics in Japan are prominent, and on the map, she sees that there are connections between Japan with many other places in the World (Fig. 7). There is much discussion about the potential impact of the pandemic on the Olympics. She is intrigued to see that there are connections between Japan and her home country that implies her business will be affected.

Since she is most familiar with India's geography, she changes the geo-binning to the state level for more detail and selects two states in India with much discussion (Fig. 8). A quick look through the statements shows that the Olympics are in jeopardy, and a closer-to-home event of the under-17 women's soccer World Cup scheduled to be in India was unfortunately postponed. Teams would have come from around the World for this event. This sad news prompts

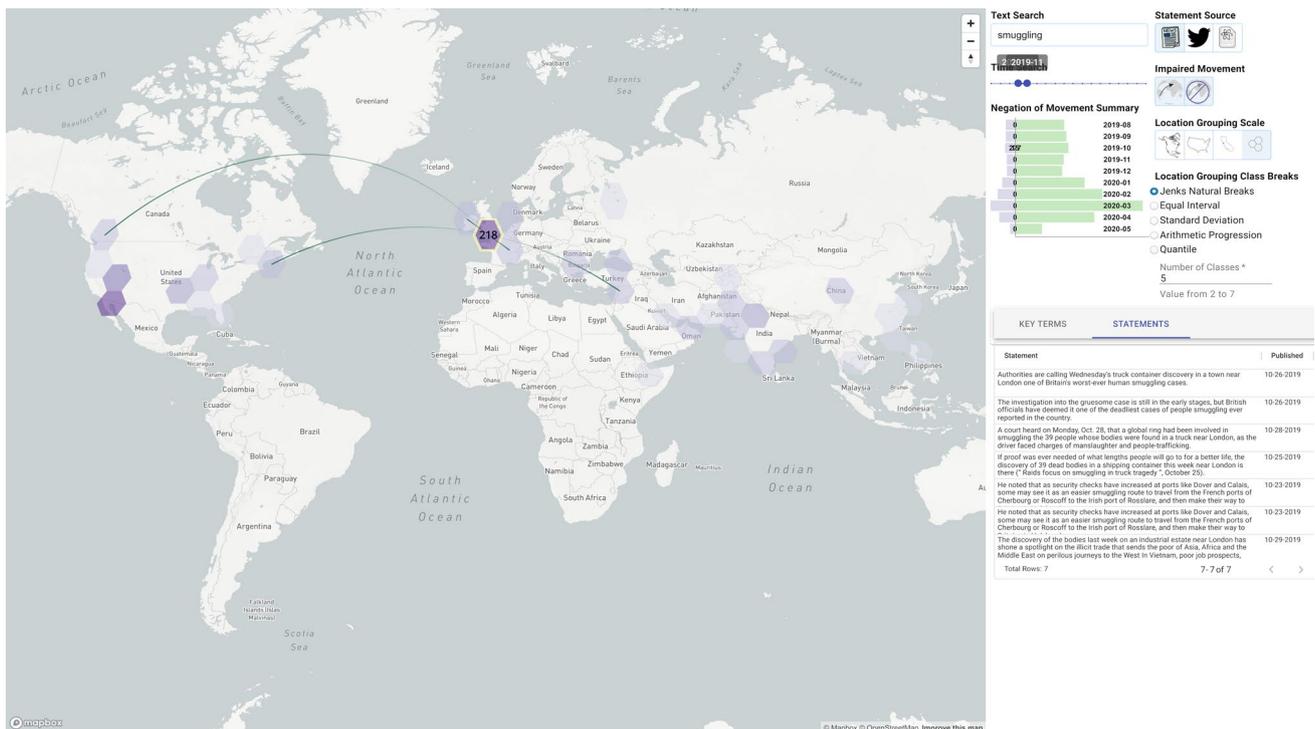


Fig. 2 By filtering the search results for *smuggling* to only show statements from October 2019, she sees many statements from a significant human trafficking event in England where many people lost their lives

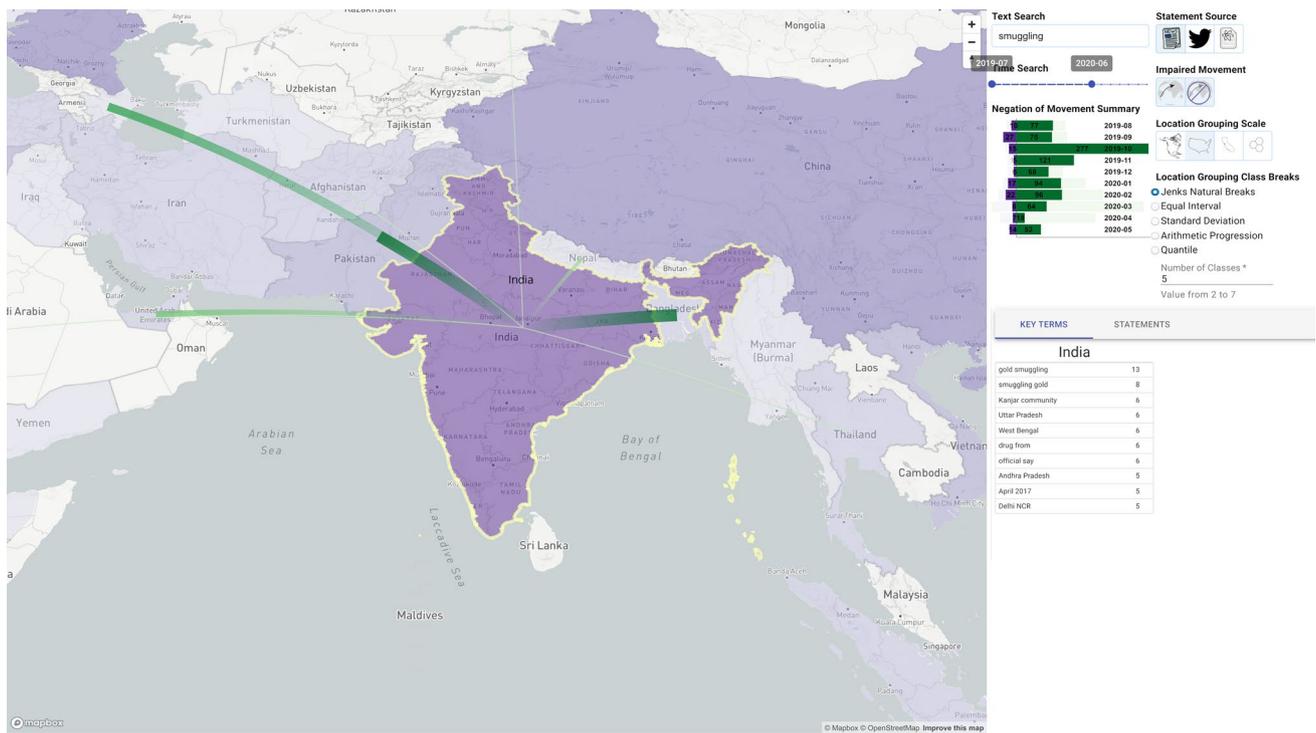


Fig. 3 A search for *smuggling* shows important location sources for India and a spike of activity in October 2009

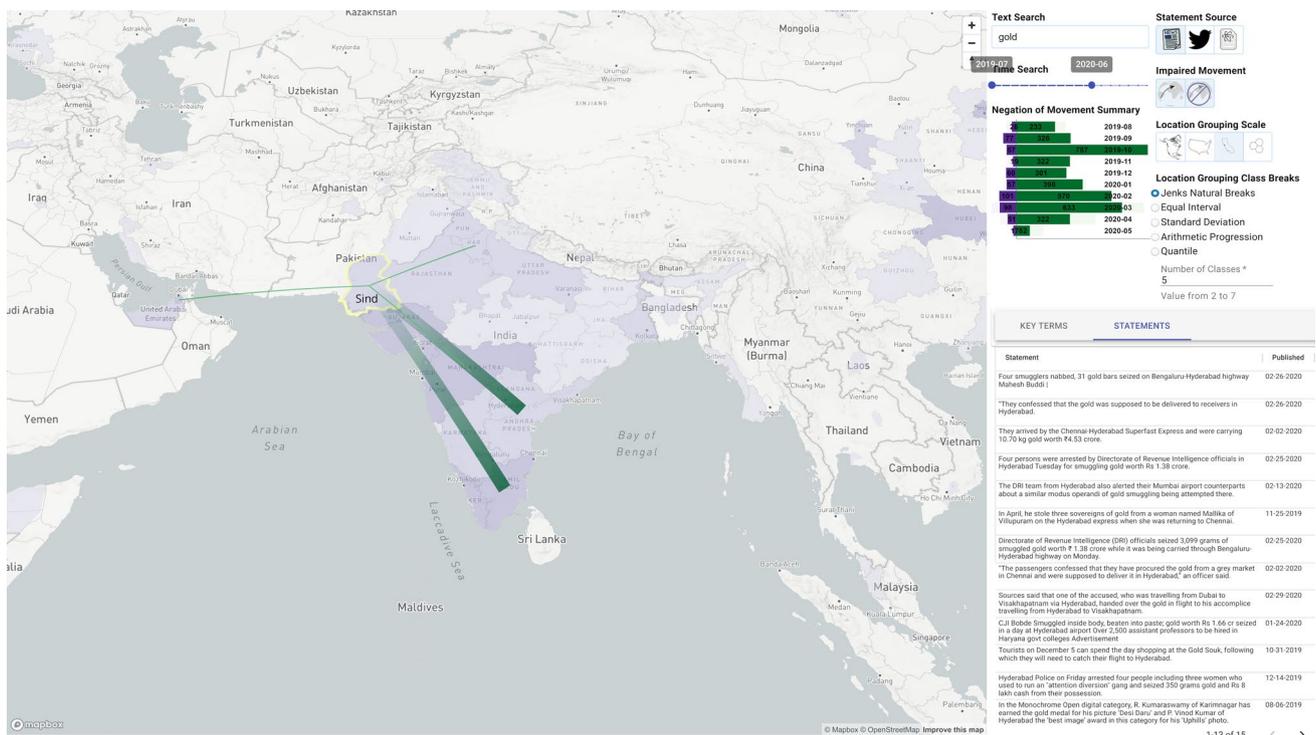


Fig. 4 A search for *gold* produces many statements about gold smuggling and possible routes

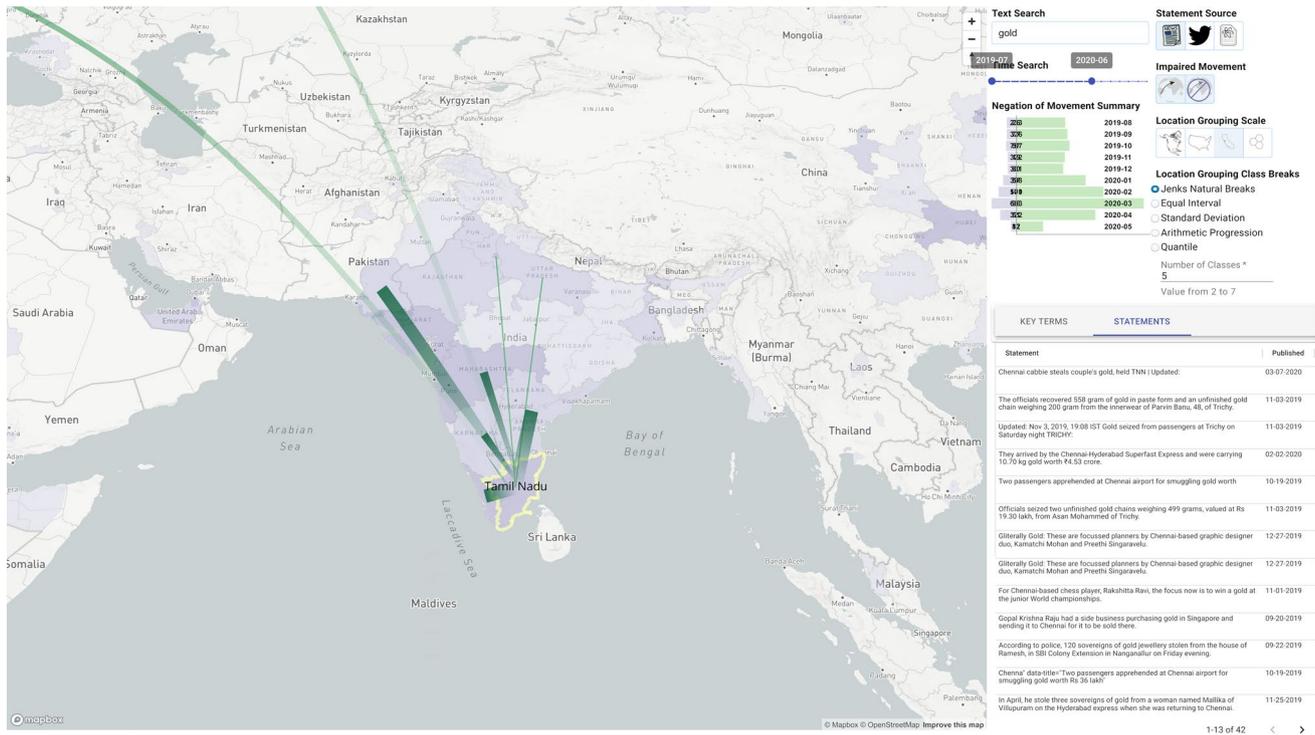


Fig. 5 While looking for more detail about gold smuggling, statements suggest a smuggling route from Mumbai to Chennai along the Chennai Express train

her to follow up on the story by clicking the statements to view the original articles in her web browser to see if it will be rescheduled and thus if there will be a future need for travel. From reading the statements, she sees that an auto show, the AP World Indoor Sporting Championships, and the Australian Grand Prix auto racing's China leg are three events postponed in China. Cricket in Australia was also severely affected, with teams planning to come from India, England, and many other countries to compete. Major sporting leagues in the US, like the NBA basketball league, were also interrupted. The disruptions to nearby events like cricket are particularly concerning given the sport's popularity with Indians and their potential spectator travel related to her business.

Lastly, to see how widespread the pandemic's impact is on global sporting events where Indians may travel, Arti explores other countries and sees that the pandemic has significantly impacted professional soccer in Spain and other European countries. Madrid, Milan, and cities in Germany where prominent soccer clubs play all show up clearly on the map as having their games affected (Fig. 9).

These use cases highlight the potential for information foraging and acquisition from statements describing movement. Geovisual analytics that combines multiple views of the data and overview + detail capabilities allow users to quickly identify important locations, connections between

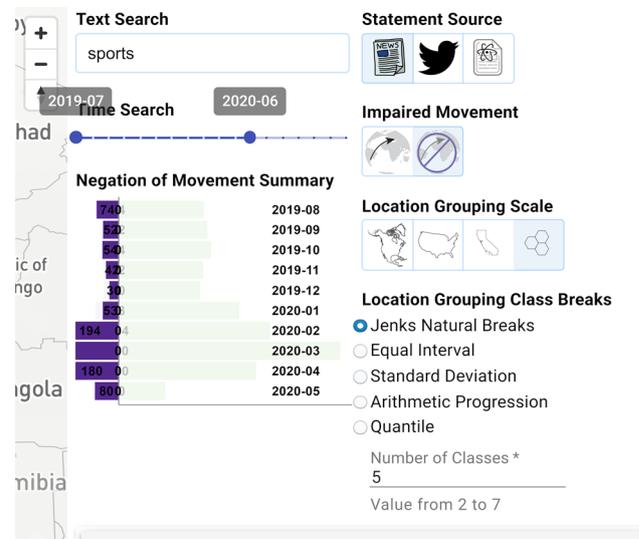


Fig. 6 Mousing over the Two-sided Temporal Bar Chart shows the number of filtered statements for impaired movement (purple bars to the left of the axis) involving sports has been increasing since the pandemic began (bars are aggregate counts of statements by month, with the oldest month at the top). Filtered statements by the search are shown in the foreground, while all statements are shown in the background on the mouse hover

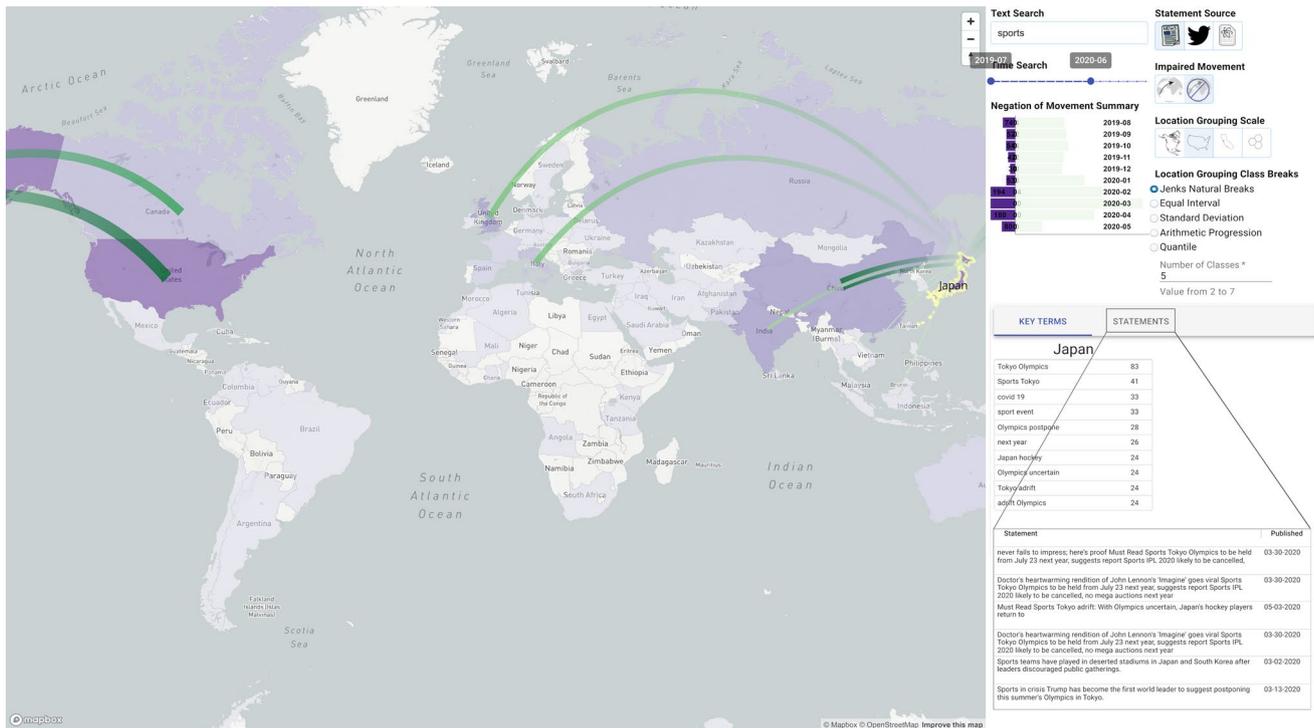


Fig. 7 Sports have been disrupted during the pandemic, including the Olympics, where athletes travel from around the World. Mousing over the Two-sided Temporal Bar Chart shows that the filtered statements for impaired movement have increased since the pandemic began



Fig. 8 Sports in India have also been affected, including competitions where local teams compete against other teams from around the World

locations, time periods, and topics of interest. We showed how the map view could show critical hotspots, and the Two-sided Temporal Bar Chart can show essential time periods, prompting the user to investigate details.

In this section, our use cases show GeoMovement's current sensemaking capabilities. Moreover, the examples of computational prediction errors and suggested solutions highlight how geovisual analytics is well suited for allowing user interactions to either explicitly or implicitly improve GeoMovement's computational predictions. Improved prediction accuracy can make the sensemaking process more efficient by reducing the human effort to sift through incorrect intermediate results.

5.3 Case Studies Needs Assessment

In addition to our case studies demonstrating the effectiveness of GeoMovement in extracting information about movement, we used them as a needs assessment to propose additional functionality. In the first case study, Jennifer identified statements from an event that she was not interested in. Although she quickly found what she was looking for, other problems may require much further analysis, and future results should not include statements deemed not relevant. Therefore, a future addition to GeoMovement can consist of a mechanism to either mark statements as completely "not relevant" and not show them or as "less important."

If user accounts were added to GeoMovement, these preferences in statements could be stored. If certain types of statements will never be relevant to that user, their marked statements can be used to affect their future search results. Such user input can be used as feedback to the system, and in the future, statements that are very similar to those that are marked not of interest will also be filtered out; i.e., the system learns from the feedback. Elasticsearch allows for on-the-fly criteria provided in searches that promote or demote statements or exclude them, allowing for these criteria to be personalized if user accounts were added and changed for each search.

In the second case study, Arti found that impaired movement became more prevalent in recent months. It would be a fair assumption that this is because of travel restrictions from the pandemic. However, Arti may ask for more detail about "How is impaired movement changing over time?" To do this, she should easily access statements by month. Currently, she would have to change the time slider filter to each month and select the impaired movement button. A straightforward way to answer this question would be to choose any bar in the Two-sided Temporal Bar Chart to update the time range and impaired movement filters. For example, clicking on the impaired movement purple bar for January 2020 would update all views to show these statements. Similarly, when Jennifer found a spike in smuggling activity in India

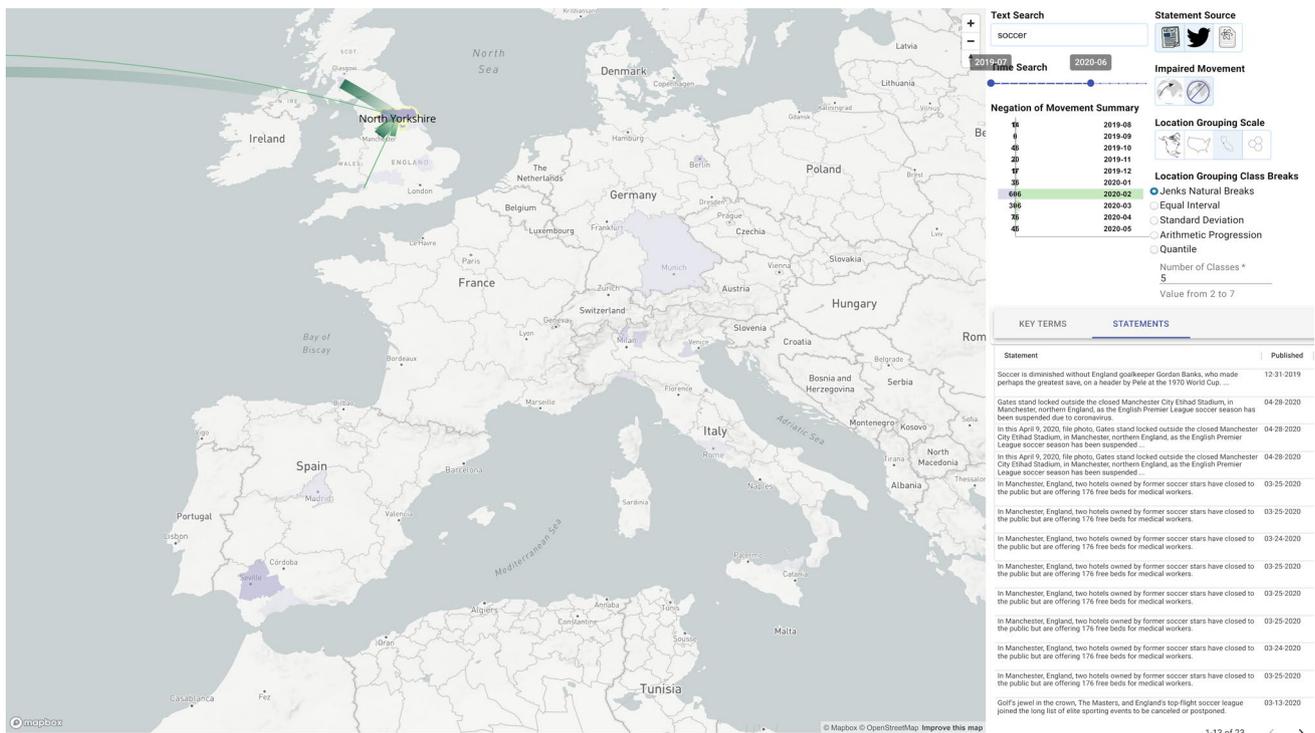


Fig. 9 Europe is another example where the pandemic postponed many of their beloved soccer matches

in October 2009, she would likely ask herself why this is and want a quick way to filter to those statements.

Finally, both Jennifer and Arti were interested in India. A question they both may have is: “What other locations have similar problems as India?” Currently, GeoMovement allows selection on the map of multiple locations to show bigrams lists for both locations’ statements next to each other. However, there are automated ways to give users hints on locations with similar statements. For example, Elasticsearch provides a “percolate” query where, after the user already selects a location, they could choose a user interface control and click a second location. The statements from the first location can be used in the percolate query. The result would be statements like the first set, and therefore other places like that place will show.

Our evaluation of the sensemaking capabilities of GeoMovement includes a statistical summary illustrating the challenge in extracting meaningful information without it, case study demonstrations, a needs assessment, and a list of user requirements. We show that despite the large volume of messy data that is at times ambiguous, GeoMovement can quickly extract meaningful information about geographic movement.

6 Results and Implications

Analyzing a large volume of text describing geographic movement requires imprecise computational processing and predictions on already messy data that different people can perceive differently, resulting in some errors. We present the data in a geovisual analytics web application that follows the overview-first + detail mantra (Shneiderman 1996). Users can find exciting patterns in the overviews and areas that need further investigation. They can also search the data and eventually drill down to the actual statements and articles. Results that appear to be erroneous or merely uninteresting can be hidden so that it is easier for the analyst to focus on more helpful information.

Our approach tightly integrates geovisual analytics with behind-the-scenes computational predictions for statements describing movement from most text that does not, GIR methods to extract geography for mapping and analysis, and a novel prediction for impaired movement. We show the value of GeoMovement through case studies where humans quickly learn about and visualize geographic movements. Furthermore, we demonstrate that a multi-scale system that applies a space-time-concept/attribute filtering process can effectively make sense from large volumes of messy movement statements.

Because of the challenges of utilizing messy big text data for understanding geographic movement, previous work has either focused only on movement data from

precise sensors or addressed narrow domains of movement described in text. However, the need for improved methods to utilize big text data about movement is illustrated in the work described in Sect. 2, and Hultquist and Cervone (2020), Hultquist and Cervone (2019), and Janowicz et al. (2019), where geography in text complements sensor data or stands alone to solve real-world problems. MacEachren (2017) also encouraged such advances in his positional paper. GeoMovement (1) tightly integrates computational predictions and geovisual analytics, and (2) is built with modern web computing technologies that return results from user queries in milliseconds. These features allow users to make sense of and uncover movement patterns quickly. We show how this efficient and broad movement exploration is unfeasible without computational and visual means to assist a human. Moreover, this completes our research objective to show that modern advances in NLP and IR can leverage large volumes of movement statements to understand the movement described quickly.

Although we stopped collecting text for GeoMovement, our techniques can be adapted in a straightforward technical way to accommodate a continually updating source given modest computational and storage requirements. In addition, all of our code is available with an open-source license to prompt future advances (Pezanowski 2022).

We anticipate GeoMovement or a production version to be beneficial for scientists, journalists, or even the general public to find information about the movement of humans, wildlife, goods, disease, and much more. Advances to analyzing movement statements can provide similar knowledge gains as movement recorded by sensors with precise geospatial and temporal data.

7 Conclusions and Future Work

Our research shows that modern computational techniques can be combined with a human-in-the-loop geovisual analytics system to overcome significant challenges and identify, process, and explore large volumes of movement statements to quickly obtain an overview of movement patterns and forage for detailed information of interest. Future research should take advantage of the wealth of context information found alongside geographic movement in text documents about why, when, and how the movement is occurring that is not often present in precise GPS data. Our research methods can likely be adapted to analyze statements involving other attributes like time and more complex spatial analysis like correlation. In addition, future research should explore the integration of precise geospatial movement data with movement in text. One initial way is to link them spatially and temporally. For example, the following steps can involve connecting

entities in the text to other information using the Semantic Web (Berners-Lee et al. 2001) and linked geographic data (Stadler et al. 2012; Janowicz et al. 2012). Future research needs to identify what makes movement statements different from other statements. Pezanowski et al. (2022) took initial steps towards this goal by identifying vital characteristics of movement statements that humans use to differentiate the movement described.

Also, since we chose to illustrate the effectiveness of GeoMovement through the case studies and other assessments in Sect. 5, future work should include a more thorough evaluation of user needs, especially focusing on the additions suggested in this Section.

Moreover, prior research shows that a geovisual analytics system combined with ML predictions can continually improve the accuracy of the ML predictions by having users correct machine errors (Snyder et al. 2020; Andrienko et al. 2022). An extension of GeoMovement can make it an intermediary between computational predictions and humans. As more humans use GeoMovement, humans can iteratively correct any errors they encounter, thereby improving the predictions. As an example, in Sect. 3.3, we discuss how an ML model would likely be superior to a rules-based approach detecting impaired movement but requires a large amount of training data. GeoMovement can show initial predictions of normal movement and impaired movement from an ML model that used a small amount of training data. A simple tool can allow users to correct errors in the predictions. Once many users are using GeoMovement, the training data set can proliferate. This technique has also shown success in commercial production mapping systems like Google Maps, where users reach their destination, and Maps asks about driving directions' accuracy.

Two overall ways that GeoMovement can improve the computational predictions are, first, explicitly asking for feedback on any incorrect predictions and second, through implicit user actions. For example, in GeoMovement's statements view, a button can be added to explicitly mark any statements that do not describe movement. Once a sufficient number of users mark a statement to conclude that it is incorrect, the corrected statement can be added to the corpus of statements used to train the model. As an example of implicit actions, the statements' list is currently returned in order by the search engine software that roughly corresponds to how closely they match the search terms entered and how frequently those search terms appear in the statement. Skipped statements can be recorded when users page through results to find what they are looking for. If many users skip certain statements, they can be deemed less valuable and given a weighting that lets them be listed lower in the order they are returned, thereby promoting more critical statements. There are many other possibilities to obtain either explicit feedback (users correct place mentions that

were assigned to the incorrect location using the GIR techniques Karimzadeh and MacEachren 2019) or implicit feedback from users (identifying important locations to highlight based on where previous users navigated to on the map) to improve GeoMovement's sensemaking ability.

Additionally, none of the existing GIR systems have ideal performance metrics. There is often a trade-off between different geoparsers' false positives and false negatives and other ways to rank geocoding results on large datasets. Future work can allow user controls in GeoMovement to choose different geoparsing back-ends. Also, the open-source GIR GeoTxt can be modified to allow GeoMovement users to set a tolerance to allow more false-positives in situations where missing relevant statements are most critical or allow fewer false-positives in situations where the key is a quick overview and having every relevant statement does not matter. Our use of geovisual analytics for GeoMovement adds substantial potential for future improved analysis of movement statements.

Finally, it is important to revisit that GeoMovement only supports English. Therefore, a valuable extension to our research would be the addition of other languages. The higher-level software, programming languages, and APIs (discussed in Sect. B) we used in our research to harvest, process, store, retrieve, and visualize data are all capable of handling many languages and character sets. However, incorporating documents in other languages into GeoMovement would require significant work on the three main computational predictions we used. First, identifying statements that describe movement would require a large amount of labeled training data in the added language. Also, for higher accuracy in predictions, language models like ELMo would need to be re-trained for that language, like the research of Che et al. (2018) and Fares et al. (2017). Second, our geoparsing methods to extract place mentions and geocode them to their correct location on Earth would need to be adapted (Mandl et al. 2009). Third, our negation detection techniques would also require adaptation to support other languages (Morante and Blanco 2021). These changes are possible but need significant work as existing research is less robust than the equivalent in English.

The World is a dynamic place, and understanding geographic-scale movements of things is essential in many domains like business, public health, and environmental science. Leveraging information about movement found in text can complement precise sensor-based movement data. Also, movement described in text is valuable since precise movement data is often unavailable because of high costs or impracticality to deploy sensors. GeoMovement and the combination of computational and visual methods it integrates are steps toward that objective.

Appendix A Terms Related to Movement

Table 3 shows the list of verbs related to movement used in our negation detection techniques as retrieved from <https://archiewahwah.wordpress.com/2019/04/16/movement-verbs-list/> on January 21, 2022.

Table 4 shows the list of adjectives related to movement used in our negation detection techniques as retrieved from https://dSPACE.ut.ee/bitstream/handle/10062/18053/adjectives_and_adverbs_of_movement.html on January 21, 2022.

Appendix B Detailed Description of GeoMovement

B.1 Text Data Sources

First, we purchased news articles from Webhose.io (<https://webhose.io/>). Webhose.io collects text documents from various Web sources and distributes them in a clean semi-structured XML format. To save costs and reduce the number of irrelevant statements, we used their web archive tool to obtain text documents that match the key term *travel*, that is in the *English* language, and that was published on a website categorized as *news*. We downloaded all articles in their archive between August 1, 2019, and May 10, 2020, that fit these filters (397,919 articles).

Second, we acquired and incorporated tweets using the GeoCov19 dataset spanning February 1, 2020, to May 2, 2020 (Qazi et al. 2020). This dataset contains hundreds of millions of tweets related to the Covid-19 pandemic. We harvested only the tweets in English, which resulted in 328 million tweets.

The third source for statements is the COVID-19 Open Research Dataset (CORD-19) (Wang et al. 2020b). This dataset contains scientific articles about the COVID-19 pandemic that were harvested and placed in a clean semi-structured XML format. The collection has over 50 thousand articles retrieved from PubMed Central (<https://www.ncbi.nlm.nih.gov/pmc/>), bioRxiv (<https://www.biorxiv.org/>), and medRxiv (<https://www.medrxiv.org/>) that match search terms related to all Coronaviruses and is updated daily. We downloaded the dataset on November 14, 2020, and selected documents from that date back until August 1, 2019, which totaled about 15,600 documents.

B.2 Application Development for Efficient Sensemaking

GeoMovement is an entirely web-based application. We built the web client using the React JavaScript Library (<https://reactjs.org/>) as its framework. We also used Material

UI (<https://material-ui.com/>) React Components to make development efficient and provide a friendly and familiar user experience. We used Mapbox GL JS (<https://docs.mapbox.com/mapbox-gl-js/api/>) in the map portion of GeoMovement and Deck GL Libraries (<https://deck.gl/>) (Wang 2017) to provide advanced mapping and geovisualization capabilities along with state-of-the-art rendering speeds for large amounts of data. The web client application is a visual interface for humans to search the server's data in Elasticsearch. Other software used to pre-process the data includes Postgresql (<https://www.postgresql.org/>) for data storage and PostGIS (<http://postgis.net/>) for geospatial computations. The result of this application design is a web application accessible in any major modern web browser, both desktop, and mobile, that produces rapid results to user searches of hundreds of millions of descriptions of geographic movement. Almost instantaneous responses to searches enable users to more easily forage through data in the sensemaking process and identify important information. Figure 1 shows the user interface of GeoMovement.

Appendix C GeoMovement User Requirements

In Section 5.2, we address the utility of GeoMovement through two case study scenarios that demonstrate the ability of GeoMovement to support information foraging through the large volumes of text coming from multiple kinds of sources. The scenarios make specific assumptions about the required capabilities, skills, and prior knowledge of users.

User Capabilities

Currently, GeoMovement works only with text in English; thus, users are expected to be English speakers. Given the focus on interactive maps and graphical displays, users need manual dexterity sufficient to handle interaction with the map and various controls and adequate visual acuity to read the maps and graphs. Although we did not test GeoMovement on many computer configurations, based on our internal testing, web development expertise, and established guidelines and recommendations for the development technologies, we can estimate that the user needs modest computing capabilities. The user should have a computer with a minimum monitor screen resolution of 1024 × 768 pixels. A modern JavaScript-enabled browser (Chrome, Firefox, Safari, and Microsoft Edge) is required, and minimum computer hardware specifications to support these browsers. GeoMovement was tested and worked on a modern smartphone with a mobile Chrome browser; however, GeoMovement's sensemaking capabilities are restricted because of the small touch screen. Although GeoMovement's design and architecture provide swift search responses from the server software, performance depends on the users' Internet

Table 3 List of verbs related to movement used in our detection of impaired movement

Advance—move forward	Pootle [informal]—proceed casually
Aim (for)—go in the direction of	Pop (in, to, over etc)—quickly visit/pass etc
Amble—walk casually	Potter—move in an unhurried way
Angle (towards etc)—turn one—s steps towards	Pound—proceed with fast, heavy steps
Back (out, away etc)—move in reverse	Prance—move flamboyantly, with effected grace
Barrel—move in a forceful, uncontrolled way	Progress—advance
Beetle—hurry like an insect	Proceed—go forward
Belt—move swiftly	Promenade—take a leisurely walk
Bez [informal, dialect]—zip around	Prowl—move in a shifty or predatory manner.
Bluster—move forcefully yet ungraciously	Race—move quickly, in competition.
Bolt—move swiftly	Ramble—walk far and wide
Bounce—move with elastic motions	Regress—go back
Bound—move quickly with large steps	Return—go back
Bumble—proceed in a clumsy fashion	Roam—proceed with no direction in mind
Canter—move fairly, like a horse	Roll—[literal] proceed in turning motions like a wheel/[figura- tive] move steadily
Careen—speed forward uncontrollably	Rove—wander far and wide
Career—speed forward with little control	Run—proceed quickly, both feet leaving the floor
Charge—move aggressively towards something	Rush—move with haste
Crawl—[literal] go on all fours / [figurative] proceed slowly	Sashay—move in a confident and flamboyant way
Creep—move sneakily or slowly	Saunter—walk arrogantly, confidently
Dance—move rhythmically	Scamper—run like an agitated animal
Dart—go swiftly	Scarper—run away
Dash—run quickly	Scoot—proceed at a fair pace / shuffle to one side
Dawdle—proceed slowly and reluctantly	Scud—move quickly as if blown by the wind
Dive—descend quickly	Scuff—walk in a careless, friction-producing way
Dodder—move unsteadily, as if elderly	Scurry—hurry like a small animal
Dogtrot—move at a brisk, comfortable pace, like a dog	Scuttle—hurry like an insect
Emerge (from)—come out of	Seethe—proceed like oozing liquid
Escape—move out of danger/confinement	Shuffle—walk slowly, without lifting one—s feet
File (in)—[of multiple people] go one-by-one	Skedaddle—depart in haste
Flee—run away from	Skip—proceed bouncing from one foot to the other
Flounce—move in a flamboyant way	Skitter—move hurriedly
Flop—move loosely	Slide—move frictionlessly
Fly—[literal] move through the air / [figurative] proceed swiftly	Slink—go smoothly/sensuously
Footslog—march a long distance	Slip—move frictionlessly / make an accidental movement
Forge (on, ahead)—proceed strongly and steadily	Slither—slide forward like a snake
Gallop—move quickly, like a horse	Slope—as sneak.
Gambol—proceed in a playful, energetic manner	Sneak—Proceed surreptitiously
Glide—move frictionlessly	Speed—move very fast
Go—basic movement verb / depart	Split—depart
Hare—proceed extremely quickly, like the animal	Sprint—run at top speed
Hasten—move with haste	Stagger—move unbalanced, unsteadily
Head (towards, for, to etc)—proceed in the direction of	Stalk—move as though hunting
Hie [archaic]—go quickly	Stampede [multiple people]—progress chaotically /in agitation
Hightail [informal]—move quickly	Steam—power forward
Hike—go a long distance	Step—move with the feet
Hop—[literal] proceed on one foot / [figurative] make a short journey	Streak—move quickly, as if leaving a line of light behind you
Hurtle—move quickly, violently and recklessly	Stride—walk purposefully
Issue (from)—come out of	Stroll—walk in a brisk, leisurely manner
Jog—move at a medium pace/half-run	Strut—walk stiffly / arrogantly

Table 3 (continued)

Advance—move forward	Pootle [informal]—proceed casually
Jump—propel oneself through the air	Swagger—move arrogantly
Jaunt—go on a short trip	Sweep—proceed swiftly
Journey—travel a distance	Tank—progress swiftly and forcefully
Labour—move with difficulty, requiring force	Tiptoe—proceed lightly, silently on the toes
Leap—jump far	Traipse—walk a distance
Leg (it)—run (away)	Tramp—walk a distance
Limp—proceed unevenly / with an injured leg	Trample—walk without precision or care
Lollop—proceed in ungainly bounds	Travel—move a distance
Lope—move in large strides	Tread—move using the feet
Lunge—jump forward to attack	Trek—travel a long time / distance
March—move steadily/forcefully/with purpose	Trip—proceed lightly, gaily
Meander—proceed in an indirect way	Tromp—walk heavy-footed
Mooch [informal]—go around in a skulking manner	Troop—march with effort
Mosey—walk in a leisurely manner	Trot—move briskly like a horse
Move—basic verb of movement	Trundle—move arduously like a cart
Nip (into, across, over etc)—quickly go	Tumble—fall / spiral forward
Pace—walk steadily	Undulate—proceed in wavy motions
Pad—walk casually/softly/steadily like an animal	Waddle—Walk in ungainly fashion, from side to side
Parade—proceed in an extrovert manner	Walk—go by foot
Patrol—walk around to guard	Wander—travel without a direction in mind
Patter—go with a light tripping sound	Wend—travel by a circuitous route
Pass—move beyond	Whizz—go speedily
Pelt—move quickly, like a hurled stone	Wobble—move unsteadily
Perambulate [formal, rare]—walk	Zip—move swiftly
Plod—move with heavy, laborious motions	

Table 4 List of adjectives related to movement used in our detection of impaired movement

In the text, some nouns are qualified with adjectives (sharp, sudden) and some verbs with adverbs (dramatically, significantly)
Adjectives used to describe movement include: slow, slight, moderate, gradual, steady, quick, rapid, significant, sharp, substantial, dramatic
Used to show a small change: slight
Used to show a regular movement: gradual, steady
Used to show considerable, striking or unexpected change: significant, substantial, dramatic (both positive and negative change), sharp, sudden
Adverbs are formed by adding -ly to the adjective, and sometimes one or two other letters change as well
Degree of change:
dramatically, considerably, significantly, substantially, sharply, moderately, slightly
Note that “dramatically” can refer to both good and bad changes
Speed of change:
rapidly, quickly, suddenly, gradually, steadily, slowly

connection speed because it is web-based. Therefore, GeoMovement will function, but user Internet connection speeds under the typical broadband minimum of 25 Mbps will likely produce a lag in user search results that inhibits the analysis.

User Skills

GeoMovement expects familiarity with web browsers, graphs, and maps. We assume some experience with thematic maps (maps that depict data geographically), but no

particular expertise is required. Thus, users familiar with the kinds of maps routinely found on news sites like those of the New York Times or the Guardian are sufficient. An ability to understand simple statistical graphs (bar charts) is assumed.

User prior knowledge

Since GeoMovement is designed for user-led information foraging, it does not provide suggestions to users about what to explore or specifically what to look for. The assumption

made is that each user has some knowledge of the topic they are interested in that they can apply to leverage the visual interface and computational methods included in GeoMovement. That is, in fact, the point. This is a human-in-the-loop system intended for users who have domain expertise that is complemented by the ability of GeoMovement to process very large volumes of text and the potential for interactive filtering tools to support quick narrowing in on items of interest through the use of that domain knowledge to decide what is interesting and what is not interesting. Beyond domain knowledge, however, users also need some understanding of the limitations of the computational methods that underlie GeoMovement. Specifically, they should understand that any system like this will generate some false positives and some false negatives; thus, it will show them some irrelevant information (which their domain expertise is likely to allow them to ignore) and miss some relevant information. Therefore, they need to understand that absence of evidence from GeoMovement does not necessarily mean the absence of the phenomena they are interested in.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s42489-022-00098-3>.

Acknowledgements We would like to thank the Information Technology group in the College of Information Sciences and Technology at The Pennsylvania State University for providing computing resources to host GeoMovement. Specifically, Adam McMillen, a Systems Administrator in the group, provided his expertise in establishing a virtual server and deploying GeoMovement to it.

Author Contributions Conceptualization: SP, PM; Methodology: SP, PM, AMM; Software: SP; Data curation: SP; Formal analysis: SP; Validation: SP, PM; Resources: PM; Visualization Design: SP, AMM; Visualization implementation: SP; Project administration: SP; Writing—original draft: SP; Writing—review and editing: SP, PM, AMM; Supervision: PM, AMM.

Funding No funds, grants, or other support was received.

Declarations

Competing interests The authors have no competing interests to declare that are relevant to the content of this article.

References

- Al-Dohuki S, Wu Y, Kamw F et al (2017) SemanticTraj: a new approach to interacting with massive taxi trajectories. *IEEE Trans Visual Comput Graph* 23(1):11–20. <https://doi.org/10.1109/TVCG.2016.2598416>
- Andrienko N, Andrienko G, Fuchs G et al (2020) *Visual analytics for data scientists*. Springer, Cham. <https://doi.org/10.1007/978-3-030-56146-8>
- Andrienko N, Andrienko G, Adilova L et al (2022) Visual analytics for human-centered machine learning. *IEEE Comput Graph Appl* 42(1):123–133. <https://doi.org/10.1109/MCG.2021.3130314>
- Angelini M, Fazzini V, Ferro N et al (2018) CLAIRe: a combinatorial visual analytics system for information retrieval evaluation. *Inform Process Manag* 54(6):1077–1100. <https://doi.org/10.1016/j.ipm.2018.04.006><https://www.sciencedirect.com/science/article/pii/S0306457317308221>
- Bekele M, de By R, Singh G (2016) Spatiotemporal information extraction from a historic expedition gazetteer. *ISPRS Int J Geo-Inform* 5(12):221. <https://doi.org/10.3390/ijgi5120221>
- Berners-Lee T, Hendler J, Lassila O (2001) The semantic web. *Sci Am* 284(5):34–43. <http://www.jstor.org/stable/26059207>
- Birch CP, Oom SP, Beecham JA (2007) Rectangular and hexagonal grids used for observation, experiment and simulation in ecology. *Ecol Model* 206(3):347–359. <https://doi.org/10.1016/j.ecolmodel.2007.03.041><https://www.sciencedirect.com/science/article/pii/S0304380007001949>
- Blank D, Henrich A (2015) Geocoding place names from historic route descriptions. In: *Proceedings of the 9th Workshop on Geographic Information Retrieval*. ACM, New York, NY, USA, GIR '15, pp 9:1–9:2. <https://doi.org/10.1145/2837689.2837698>
- Bosch H, Thom D, Wörner M, et al (2011) ScatterBlogs: geo-spatial document analysis. In: *VAST 2011—IEEE Conference on Visual Analytics Science and Technology 2011*, Proceedings pp 309–310. <https://doi.org/10.1109/VAST.2011.6102488>
- Bosch H, Thom D, Heimerl F et al (2013) ScatterBlogs2: real-time monitoring of microblog messages through user-guided filtering. *IEEE Trans Visual Comput Graph* 19(12):2022–2031. <https://doi.org/10.1109/TVCG.2013.186>
- Brewer CA, Hatchard GW, Harrower MA (2003) Colorbrewer in print: a catalog of color schemes for maps. *Cartogr Geograph Inform Sci* 30(1):5–32. <https://doi.org/10.1559/152304003100010929>
- Che W, Liu Y, Wang Y, et al (2018) Towards better UD parsing: deep contextualized word embeddings, ensemble, and treebank concatenation. In: *Proceedings of the CoNLL 2018 shared task: multilingual parsing from raw text to universal dependencies*. Association for Computational Linguistics, Brussels, Belgium, pp 55–64. <http://www.aclweb.org/anthology/K18-2005>
- Chu D, Sheets DA, Zhao Y, et al (2014) Visualizing hidden themes of taxi movement with semantic transformation. In: *2014 IEEE Pacific Visualization Symposium*, pp 137–144. <https://doi.org/10.1109/PacificVis.2014.50>
- Dodge S (2016) From observation to prediction: the trajectory of movement research in GIScience. In: Kuhn W, Onsrud H (eds) *The past and next twenty years. Advancing geographic information science*. GSDI Association Press, Taipei, pp 123–136
- Dodge S (2016) From observation to prediction: the trajectory of movement research in GIScience. In: *The Past and Next Twenty Years, Advancing Geographic Information Science*, pp 123–136
- Dodge S, Laube P, Weibel R (2012) Movement similarity assessment using symbolic representation of trajectories. *Int J Geograph Inform Sci* 26(9):1563–1588. <https://doi.org/10.1080/13658816.2011.630003>
- Dodge S, Bohrer G, Bildstein K, et al (2014) Environmental drivers of variability in the movement ecology of turkey vultures (*Cathartes aura*) in North and South America. *Philos Trans R Soc B Biol Sci* 369(1643):20130195. <https://doi.org/10.1098/rstb.2013.0195>
- Dodge S, Weibel R, Ahearn SC et al (2016) Analysis of movement data. *Int J Geograph Inform Sci* 30(5):825–834. <https://doi.org/10.1080/13658816.2015.1132424>
- Drymonas E, Pfoser D (2010) Geospatial route extraction from texts. In: *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Data Mining for Geoinformatics*. ACM, New York, NY, USA, DMG '10, pp 29–37. <https://doi.org/10.1145/1869899.90.1869894>

- Ender A, Hossain MS, Ramakrishnan N et al (2014) The human is the loop: new directions for visual analytics. *J Intell Inform Syst* 43(3):411–435. <https://doi.org/10.1007/s10844-014-0304-9>
- Esri (2021) Why hexagons? <https://pro.arcgis.com/en/pro-app/latest/tool-reference/spatial-statistics/h-whyhexagons.htm>
- Fares M, Kutuzov A, Oepen S, et al (2017) Word vectors, reuse, and replicability: towards a community repository of large-text resources. In: Proceedings of the 21st Nordic Conference on Computational Linguistics. Association for Computational Linguistics, Gothenburg, Sweden, pp 271–276 <http://www.aclweb.org/anthology/W17-0237>
- Felmler DH, Blanford JI, Matthews SA et al (2020) The geography of sentiment towards the Women’s March of 2017. *PLOS One* 15(6):1–21. <https://doi.org/10.1371/journal.pone.0233994>
- Fialho P, Coheur L, Quaresma P (2020) To bert or not to bert dealing with possible bert failures in an entailment task. In: Lesot MJ, Vieira S, Reformat MZ et al (eds) Information processing and management of uncertainty in knowledge-based systems. Springer International Publishing, Cham, pp 734–747
- Garber J (2019) Qantas airways testing ‘final frontier in aviation’. <https://www.foxbusiness.com/industrials/qantas-airways-testi-ng-final-frontier-in-aviation>
- Graser A, Dragaschnig M, Widhalm P, et al (2020) Exploratory trajectory analysis for massive historical ais datasets. In: 2020 21st IEEE International Conference on Mobile Data Management (MDM), pp 252–257 <https://doi.org/10.1109/MDM48529.2020.00059>
- Graser A, Widhalm P, Dragaschnig M (2020) The m³ massive movement model: a distributed incrementally updatable solution for big movement data exploration. *Int J Geograph Inform Sci* 34(12):2517–2540. <https://doi.org/10.1080/13658816.2020.1776293>
- Graser A, Dragaschnig M, Koller H (2021) Exploratory analysis of massive movement data. In: Werner M, Chiang YY (eds) Handbook of big geospatial data. Springer, Cham, p 285–319. https://doi.org/10.1007/978-3-030-55462-0_12
- Gritta M, Pilehvar MT, Limsopatham N et al (2018) What’s missing in geographical parsing? *Lang Resour Eval* 52(2):603–623. <https://doi.org/10.1007/s10579-017-9385-8>
- Hiremath BN, Patil MM (2021) Sarcasm detection using cognitive features of visual data by learning model. *Expert Syst Appl* 184(115):476 <https://doi.org/10.1016/j.eswa.2021.115476><https://www.sciencedirect.com/science/article/pii/S0957417421008873>
- Hu Y, Ye X, Shaw SL (2017) Extracting and analyzing semantic relatedness between cities using news articles. *Int J Geograph Inform Sci* 31(12):2427–2451. <https://doi.org/10.1080/13658816.2017.1367797>
- Huang Q (2017) Mining online footprints to predict user’s next location. *Int J Geograph Inform Sci* 31(3):523–541. <https://doi.org/10.1080/13658816.2016.1209506>
- Huang Z, Zhao Y, Chen W et al (2020) A natural-language-based visual query approach of uncertain human trajectories. *IEEE Trans Visual Comput Graph* 26(1):1256–1266. <https://doi.org/10.1109/TVCG.2019.2934671>
- Hultquist C, Cervone G (2019) Comparison of simulated radioactive atmospheric releases to citizen science observations for the Fukushima nuclear accident. *Atmos Environ* 198:478–488. <https://doi.org/10.1016/j.atmosenv.2018.10.018><https://www.sciencedirect.com/science/article/pii/S1352231018307088>
- Hultquist C, Cervone G (2020) Integration of crowdsourced images, USGS networks, remote sensing, and a model to assess flood depth during hurricane florence. *Remote Sens* 12(5). <https://doi.org/10.3390/rs12050834><https://www.mdpi.com/2072-4292/12/5/834>
- Ishikawa T (2016) Spatial thinking in geographic information science: students’ geospatial conceptions, map-based reasoning, and spatial visualization ability. *Ann Am Assoc Geogr* 106(1):76–95. <https://doi.org/10.1080/00045608.2015.1064342>
- Jaiswal A, Pezanowski S, Mitra P, et al (2010) GeoCAM: a geovisual analytics workspace to contextualize and interpret statements about movement. *J Spat Inform Sci*: 279–294. <https://doi.org/10.5311/JOSIS.2011.3.55>
- Jamonnak S, Zhao Y, Curtis A et al (2020) GeoVisuals: a visual analytics approach to leverage the potential of spatial videos and associated geonarratives. *Int J Geogr Inform Sci* 34(11):2115–2135. <https://doi.org/10.1080/13658816.2020.1737700>
- Janowicz K, Scheider S, Pehle T et al (2012) Geospatial semantics and linked spatiotemporal data—Past, present, and future. *Semantic Web* 3:321–332. <https://doi.org/10.3233/SW-2012-0077>
- Janowicz K, McKenzie G, Hu Y, et al (2019) Using semantic signatures for social sensing in urban environments. *Mob Patterns Big Data Trans Anal*: 31–54. <https://doi.org/10.1016/b978-0-12-812970-8.00003-8>
- Karimzadeh M, MacEachren A (2019) GeoAnnotator: a collaborative semi-automatic platform for constructing geo-annotated text corpora. *ISPRS Int J GeoInform* 8(4):161. <https://doi.org/10.3390/ijgi8040161><https://www.mdpi.com/2220-9964/8/4/161>
- Karimzadeh M, Huang W, Banerjee S, et al (2013) GeoTxt: a web API to leverage place references in text. In: Proceedings of the 7th Workshop on Geographic Information Retrieval - GIR ’13 pp 72–73. <https://doi.org/10.1145/2533888.2533942>
- Karimzadeh M, Pezanowski S, MacEachren AM, et al (2019a) GeoTxt: a scalable geoparsing system for unstructured text geolocation. *Transactions GIS* 23(1):118–136. <https://doi.org/10.1111/tgis.12510>
- Karimzadeh M, Snyder LS, Ebert DS (2019b) Geovisual analytics and interactive machine learning for situational awareness. *arXiv*: 1910.05441
- Khandelwal A, Britto BK (2020) Multitask learning of negation and speculation using transformers. In: Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis. Association for Computational Linguistics, Online, pp 79–87. <https://doi.org/10.18653/v1/2020.louhi-1.9>. <https://aclanthology.org/2020.louhi-1.9>
- Khandelwal A, Sawant S (2020) NegBERT: a transfer learning approach for negation detection and scope resolution. In: Proceedings of The 12th Language Resources and Evaluation Conference. European Language Resources Association, Marseille, France, pp 5739–5748. <https://www.aclweb.org/anthology/2020.lrec-1.704>
- Klippel A, MacEachren A, Mitra P et al (2008) Geographic analysis of linguistically encoded movement patterns—A contextualized perspective. Extended abstracts for the 5th International Conference GIScience. Park City, Utah, pp 113–117
- Lai K, Porter JR, Amodeo M, et al (2022) A natural language processing approach to understanding context in the extraction and geocoding of historical floods, storms, and adaptation measures. *Inform Process Manag* 59(1):102–735. <https://doi.org/10.1016/j.ipm.2021.102735>. <https://www.sciencedirect.com/science/article/pii/S030645732100217X>
- Li Z, Xiong G, Wei Z, et al (2021) A semi-supervised end-to-end framework for transportation mode detection by using gps-enabled sensing devices. *IEEE Internet Things J*: 1–1. <https://doi.org/10.1109/JIOT.2021.3115239>
- Liben LS, Downs RM (1993) Understanding person-space-map relations: cartographic and developmental perspectives. *Dev Psychol* 29(4):739–752. <https://psycnet.apa.org/record/1993-44885-001>
- Ma C, Zhao Y, Curtis A, et al (2020) CLEVis: A semantic driven visual analytics system for community level events. *IEEE Comput Graph Appl*: 1. <https://doi.org/10.1109/MCG.2020.2973939>
- MacEachren AM (2017) Spatial data handling in big data era: 139–155. <https://doi.org/10.1007/978-981-10-4424-3>

- MacEachren AM, Jaiswal A, Robinson AC, et al (2011) SensePlace2: GeoTwitter analytics support for situational awareness. In: VAST 2011 - IEEE Conference on Visual Analytics Science and Technology 2011, Proceedings pp 181–190. <https://doi.org/10.1109/VAST.2011.6102456>
- Mandl T, Carvalho P, Di Nunzio GM et al (2009) Geoclef 2008: the clef 2008 cross-language geographic information retrieval track overview. In: Peters C, Deselaers T, Ferro N et al (eds) Evaluating systems for multilingual and multimodal information access. Springer, Berlin Heidelberg, pp 808–821
- Miller HJ, Dodge S, Miller J et al (2019) Towards an integrated science of movement: converging research on animal movement ecology and human mobility science. *Int J Geogr Inform Sci* 33(5):855–876. <https://doi.org/10.1080/13658816.2018.1564317>
- van Miltenburg E, Morante R, Elliott D (2016) Pragmatic factors in image description: the case of negations. In: Proceedings of the 5th Workshop on Vision and Language, pp 54–59. <https://www.aclweb.org/anthology/W16-3207.pdf>
- Moncla L, Gaio M, Mustière S (2014) Automatic Itinerary Reconstruction from Texts. In: Duckham M, Pebesma E, Stewart K et al (eds) Geographic Information Science. Springer International Publishing, Cham, pp 253–267
- Moncla L, Renteria-Agualimpia W, Noguera-Iso J, et al (2014b) Geocoding for texts with fine-grain toponyms: an experiment on a geoparsed hiking descriptions corpus. In: Proceedings of the 22Nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. ACM, New York, NY, USA, SIGSPATIAL '14, pp 183–192. <https://doi.org/10.1145/2666310.2666386>
- Moncla L, Gaio M, Noguera-Iso J et al (2016) Reconstruction of itineraries from annotated text with an informed spanning tree algorithm. *Int J Geogr Inform Sci* 30(6):1137–1160. <https://doi.org/10.1080/13658816.2015.1108422>
- Morante R, Blanco E (eds) (2012) Resolving the scope and focus of negation—SEM Shared Task, Montreal, Canada, <https://www.clips.uantwerpen.be/sem2012-st-neg/>
- Morante R, Blanco E (2021) Recent advances in processing negation. *Natl Lang Eng* 27(2):121–130. <https://doi.org/10.1017/S1351324920000534>
- Ninkov A, Sedig K (2019) VINCENT: a visual analytics system for investigating the online vaccine debate. *Online J Pub Health Inform* 11(2). <https://doi.org/10.5210/ojphi.v11i2.10114>
- Pezanowski S (2022) GeoMovement source code <https://doi.org/10.5281/zenodo.58888606> <https://github.com/scottpez/geomovement>
- Pezanowski S, Mitra P (2020) Recognition of implicit geographic movement in text. In: Proceedings of The 12th Language Resources and Evaluation Conference. European Language Resources Association, Marseille, France, pp 2047–2056. <https://www.aclweb.org/anthology/2020.lrec-1.251>
- Pezanowski S, MacEachren AM, Savelyev A, et al (2017) SensePlace3: a geovisual framework to analyze place-time-attribute information in social media. *Cartogr Geogr Inform Sci*:1–18. <https://doi.org/10.1080/15230406.2017.1370391>
- Pezanowski S, MacEachren AM, Mitra P (2022) Differentiating geographic movement described in text documents. *Trans GIS*: 1–26. <https://doi.org/10.1111/tgis.12893>
- Piotrowski M, Läubli S, Volk M (2010) Towards mapping of alpine route descriptions. In: Proceedings of the 6th Workshop on Geographic Information Retrieval pp 2–3. <https://doi.org/10.1145/1460007.1460013>. <http://www.zora.uzh.ch>
- Qazi U, Imran M, Ofli F (2020) GeoCoV19: a dataset of hundreds of millions of multilingual COVID-19 Tweets with location information. <https://doi.org/10.21227/et8d-w881>
- Reyes A, Rosso P (2014) On the difficulty of automatically detecting irony: beyond a simple case of negation. *Knowl Inform Syst* 40(3):595–614. <https://doi.org/10.1007/s10115-013-0652-8>
- Robinson AC (2017) Geovisual analytics. In: Wilson JP (ed) The geographic information science and technology body of knowledge, 3rd edn. <https://doi.org/10.22224/gistbok/2017.3.6>
- Robinson AC, Savelyev A, Pezanowski S, et al (2013) Understanding the utility of geospatial information in social media. In: ISCRAM 2013 Conference Proceedings - 10th International Conference on Information Systems for Crisis Response and Management. Karlsruhe Institut für Technologie (KIT), pp 918–922. <http://www.scopus.com/inward/record.url?eid=2-s2.0-84905656005&partnerID=tZOtx3y1>
- Samet H, Han Y, Kastner J, et al (2020) Using animation to visualize spatio-temporal varying COVID-19 data. In: Proceedings of the 1st ACM SIGSPATIAL International Workshop on Modeling and Understanding the Spread of COVID-19. Association for Computing Machinery, New York, NY, USA, COVID-19, pp 53–62. <https://doi.org/10.1145/3423459.3430761>
- Shneiderman B (1996) The eyes have it: a task by data type taxonomy for information visualizations. In: Proceedings 1996 IEEE Symposium on Visual Languages, IEEE, pp 336–343
- Slater LT, Bradlow W, Motti DFA et al (2021) A fast, accurate, and generalisable heuristic-based negation detection algorithm for clinical text. *Comput Biol Med* 130(104):216 <https://doi.org/10.1016/j.combiomed.2021.104216>. <https://www.sciencedirect.com/science/article/pii/S001048252100010X>
- Snyder LS, Lin YS, Karimzadeh M et al (2020) Interactive learning for identifying relevant tweets to support real-time situational awareness. *IEEE Trans Visual Comput Graph* 26(1):558–568. <https://doi.org/10.1109/TVCG.2019.2934614>
- Soares Junior A, Renso C, Matwin S (2017) ANALYTic: an active learning system for trajectory classification. *IEEE Comput Graph Appl* 37(5):28–39. <https://doi.org/10.1109/MCG.2017.3621221>. URL <http://ieeexplore.ieee.org/document/8047427/>, <http://www.ncbi.nlm.nih.gov/pubmed/28945577>
- van Son C, van Miltenburg E, Morante R (2016) Building a dictionary of affixal negations. In: Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics (ExProM), pp 49–56. <https://www.aclweb.org/anthology/W16-5007.pdf>
- Stadler C, Lehmann J, Höffner K et al (2012) LinkedGeoData: a core for a web of spatial open data. *Semantic Web* 3(4):333–354. <https://doi.org/10.3233/SW-2011-0052>
- Tapia-McClung R, Silván-Cárdenas JL (2021) Exploring spatiotemporal urbanization through a hybrid remote sensing-geovisual analytics approach. In: 2021 Mexican International Conference on Computer Science (ENC), pp 1–8. <https://doi.org/10.1109/ENC53357.2021.9534805>
- Teitler BE, Lieberman MD, Panozzo D, et al (2008) NewsStand: a new view on News. In: Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. Association for Computing Machinery, New York, NY, USA, GIS '08, <https://doi.org/10.1145/1463434.1463458>
- Thom D, Bosch H, Koch S, et al (2012) Spatiotemporal anomaly detection through visual analysis of geolocated Twitter messages. In: IEEE Pacific Visualization Symposium 2012, PacificVis 2012 - Proceedings pp 41–48. <https://doi.org/10.1109/PacificVis.2012.6183572>
- Tomaszewski B, Blanford J, Ross K et al (2011) Supporting geographically-aware web document foraging and sensemaking. *Comput Environ Urban Syst* 35(3):192–207 <https://doi.org/10.1016/j.compenurbysys.2011.01.003>. <https://www.sciencedirect.com/science/article/pii/S0198971511000068>
- Vincze V, Szarvas G, Farkas R et al (2008) The BioScope corpus: biomedical texts annotated for uncertainty, negation and their

- scopes. *BMC Bioinform* 9(11):S9. <https://doi.org/10.1186/1471-2105-9-S11-S9>
- Wang F, Lei Y, Liu Z, et al (2020a) Fast and parameter-light rare behavior detection in maritime trajectories. *Inform Process Manag* 57(5):102–268. <https://doi.org/10.1016/j.ipm.2020.102268>. <https://www.sciencedirect.com/science/article/pii/S0306457319315481>
- Wang LL, Lo K, Chandrasekhar Y, et al (2020b) COVID-19: the COVID-19 open research dataset. In: Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020. Association for Computational Linguistics, Online. <https://www.aclweb.org/anthology/2020.nlpCOVID19-acl.1>
- Wang Y (2017) Deck.gl: large-scale web-based visual analytics made easy. In: IEEE 2017 Workshop on Visualization in Practice (VIP), Phoenix, Arizona
- Xu J, Tao Y, Yan Y et al (2018) VAUT: a visual analytics system of spatiotemporal urban topics in reviews. *J Visual* 21(3):471–484. <https://doi.org/10.1007/s12650-017-0464-0>
- Yuan J, Chen C, Yang W et al (2021) A survey of visual analytics techniques for machine learning. *Comput Vis Med* 7(1):3–36. <https://doi.org/10.1007/s41095-020-0191-7>
- Zhao L, Liu Y, Zhang M, et al. (2021) Modeling label-wise syntax for fine-grained sentiment analysis of reviews via memory-based neural model. *Inform Process Manag* 58(5):102–641. <https://doi.org/10.1016/j.ipm.2021.102641>. <https://www.sciencedirect.com/science/article/pii/S030645732100131X>
- Zhu Y, Liu Y, Yu JJQ, et al (2021) Semi-supervised federated learning for travel mode identification from gps trajectories. *IEEE Trans Intell Transp Syst* 1–12. <https://doi.org/10.1109/TITS.2021.3092015>