


# Semantic Perspectives on the Lake District Writing: Spatial Ontology Modeling and Relation Extraction for Deeper Insights

Erum Haris<sup>1</sup> ✉ 

School of Computing, University of Leeds, UK

Anthony G. Cohn ✉ 

School of Computing, University of Leeds, UK  
The Alan Turing Institute, London, UK

John G. Stell ✉ 

School of Computing, University of Leeds, UK

---

## Abstract

Extracting spatial details from historical texts can be difficult, hindering our understanding of past landscapes. The study addresses this challenge by analyzing the Corpus of the Lake District Writing, focusing on the English Lake District region. We systematically link the theoretical notions from the core concepts of spatial information to provide basis for the problem domain. The conceptual foundation is further complemented with a spatial ontology and a custom gazetteer, allowing a formal and insightful semantic exploration of the massive unstructured corpus. The other contrasting side of the framework is the usage of LLMs for spatial relation extraction. We formulate prompts leveraging understanding of the LLMs of the intended task, curate a list of spatial relations representing the most recurring proximity or vicinity relations terms and extract semantic triples for the top five place names appearing in the corpus. We compare the extraction capabilities of three benchmark LLMs for a scholarly significant historical archive, representing their potential in a challenging and interdisciplinary research problem. Finally, the network comprising the semantic triples is enhanced by incorporating a gazetteer-based classification of the objects involved thus improving their spatial profiling.

**2012 ACM Subject Classification** Computing methodologies → Artificial intelligence

**Keywords and phrases** spatial humanities, spatial narratives, ontology, large language models

**Digital Object Identifier** 10.4230/LIPIcs.COSIT.2024.11

**Supplementary Material** *Dataset:* <https://github.com/UCREL/LakeDistrictCorpus> [24]  
archived at `swh:1:dir:221aa46bd3eff0783f0a471d7b8031a2910b6739`

**Funding** The research is supported by the Economic and Social Research Council (ESRC) under grant ES/W003473/1.

**Acknowledgements** The support of Microsoft under their Accelerating Foundation Models Research initiative in providing Azure credits to AGC is also acknowledged. Finally, we thank the entire team of the Spatial Narratives project for their discussions on the CLDW and this work.

## 1 Introduction

The field of Spatial Humanities thrives on unlocking the intricate connections between literary descriptions and the physical world. In this context, Geographic Information Systems (GIS) offer exciting possibilities [30]. By precisely placing historical information

---

<sup>1</sup> Corresponding author



in space and time, GIS facilitates pattern recognition, comparisons across regions, and enriched perspectives through data visualization [4]. Beyond that, GIS allows exploration of “subjective geographies,” [14] where qualitative or even ambiguous data is spatially represented. However, extracting meaningful spatial information from historical texts presents several challenges. Narratives, with their inherent complexity, can be difficult to translate directly into a GIS framework. Traditional methods often rely on identifying named locations, which can be limited by inconsistent naming conventions, missing information, and the dynamic nature of geographic references over time. Additionally, these approaches struggle to capture the rich nuances of spatial language, such as implicit references and figurative descriptions that paint a vivid picture of a place without explicitly mentioning its name. Despite these challenges, employing GIS with non-modern literary geographies remains a valuable approach. In fact, it serves as a “provocation” [14] to delve deeper into the complexities, pluralities, and ambiguities inherent in how texts represent space and place.

In the realm of analysing spatial narratives, the Corpus of Lake District Writing (CLDW) [24] has received significant attention as a valuable research data source for the exploration of the English Lake District, a region of lakes and mountains. It has long been represented in different genres of travel writing and topographical literature, compiled into the CLDW. The region’s geography embodies a harmonious blend of natural beauty, ecological richness, and cultural heritage, making it a compelling and multifaceted study area. While researchers have applied GIS methods to study the corpus from different perspectives, the first objective of this study is to specifically address the spatial aspects of the CLDW to uncover geographical features of the region. This requires a clear description on the theoretical and conceptual foundations of the problem. Hence, we present conceptual modeling supported with domain-specific ontological and gazetteer-based definition of the narrative taxonomy.

Going beyond conventional methods to extract spatial information that are difficult to be scaled and trained on a large un-annotated corpus, our second contrasting objective is to present a distinctive approach for spatial relations extraction based on pre-trained Large Language Models (LLMs) [6]. The remarkable advancements in Natural Language Processing (NLP) achieved by LLMs have opened doors for their potential application in GIS and spatial analysis [27]. These models, trained on massive amounts of text data, have the ability to adapt to new information even without specific annotations or needing a large number of training examples. Their versatility in handling various NLP tasks has earned them the designation of “foundation models” [5]. This potential for LLMs in the realm of spatial analysis is a promising avenue for further exploration.

This research bridges multiple disciplines. It sits at the intersection of GIS and NLP within the field of Spatial Humanities. Our work is particularly relevant to studies on the CLDW, spatial relation extraction, and the application of deep learning models, especially LLMs. By combining these approaches, we aim to offer a unique perspective on the spatial tapestry of the English Lake District. This multi-faceted and interdisciplinary approach is likely to resonate with a broad audience and contribute valuable insights to the fields involved.

## **2 Related Work**

### **2.1 Analysis of Lake District Narratives in Spatial Humanities**

GIS technology, with its ability to gather, manage, and analyze data systematically, has significantly fostered the development of the Spatial Humanities [12]. This emerging field utilizes GIS tools and methods to examine cultural data, including textual corpora, allowing

digital humanists to answer new and intriguing research questions from an interdisciplinary perspective. In the area of digital literary studies, “Mapping the Lakes: A Literary GIS,” a pioneering project at the University of Lancaster [17] investigated the literary geography of the Lake District. Along this line, Donaldson et al. [11] examine the spatial dimensions linked to significant terms historically used to describe the beauty of the English Lake District. They propose a methodology for analyzing this massive historical archive, combining corpus analysis, automated location identification through geoparsing, and GIS. This approach, named Geographical Text Analysis (GTA), aims to uncover connections between literary descriptions, aesthetic elements, and the physical geography of the English Lake District.

In another initiative, “Geospatial Innovation In the Digital Humanities: A Deep Map of the Lake District,” Rayson et al. [40] highlighted the need for more open-access toponym corpora to support advancements in Named Entity Recognition (NER), Named Entity Disambiguation (NED), and geoparsing. They identified additional concerns relevant to humanities research, such as the limitations of relying solely on proper nouns and the need to account for historical variations in spelling and place names. As a result, they developed the annotated CLDW, a deeply-labeled resource offers a significant leap forward in the available testbeds for researchers. The corpus is annotated more deeply than is currently possible with vanilla NER, disambiguation and geoparsing tools. The annotation scheme links different historical and spelling variants of place-names along with identifying and differentiating geographical features.

Chesnokova et al.[33] investigated how the perception of landscapes in England’s Lake District National Park has changed over time. They analyzed mentions of silence and tranquility in historical texts i.e. the CLDW and contemporary data from the Geograph Project. Using sentiment analysis, the researchers found that mentions of silence and tranquil sounds were more likely to have positive connotations than random text from the same sources. Additionally, mapping the places mentioned in the CLDW texts affirms the lasting influence of famous authors such as Wordsworth on the way silence is described.

Smail et al. [45] argued that relying solely on GTA has limitations and emphasized the need to go beyond it. While valuable, GTA only reveals geographies linked to place names with known coordinates. This essentially provides a quantitative representation of named locations. They proposed complementing GTA with methods that capture the qualitative aspects of geography embedded within texts. Drawing on the CLDW, they present initial ideas along this line, primarily based on techniques from corpus linguistics.

In some recent approaches beyond analysing toponymic geography in the CLDW, Haris et al. [19] propose using qualitative spatial representation (QSR) for extracting and interpreting spatial relationships and perform spatial reasoning for deducing new knowledge. Steiner et al. [46] defined “Spatio-Textual Regions” as “a clustered set of toponyms and a contiguous section of text describing those toponyms.” They applied spatial clustering on CLDW and identified 8 major clusters and an outside region that indicates places outside the Lake District. The authors in [13] propose an extensible framework based on NLP, QSR, and visual analytics to study the CLDW. Their initial work focused on extracting, analyzing, and visualizing elements related to location, locale, and sense of place described in the texts. The presented workflow highlighted two key components – Dataset Creation and Model Training and reported experimental results on spatial entity extraction process that identifies and extracts relevant spatial elements in text using the an “Extractor” based on standard NLP entity extraction libraries.

## 2.2 LLMs for Spatial Information Extraction

The intersection of NLP and spatial information science has opened doors to novel research possibilities [27]. Studies like Hu et al. [22] have delved into understanding relationships between cities by analyzing news articles. Signorini et al. [43] have leveraged the power of social media, specifically tweets, to predict disease outbreaks in specific locations. Miranker and Giordano [28] utilized corpus linguistics and NLP to analyze social and media releases from official sources related to the death of migrants at the Texas-Mexico border. Others, like Sheehan et al. [42], have utilized Wikipedia data to predict poverty levels in different regions. These examples showcase the diverse and impactful applications that arise when NLP and spatial information science work together.

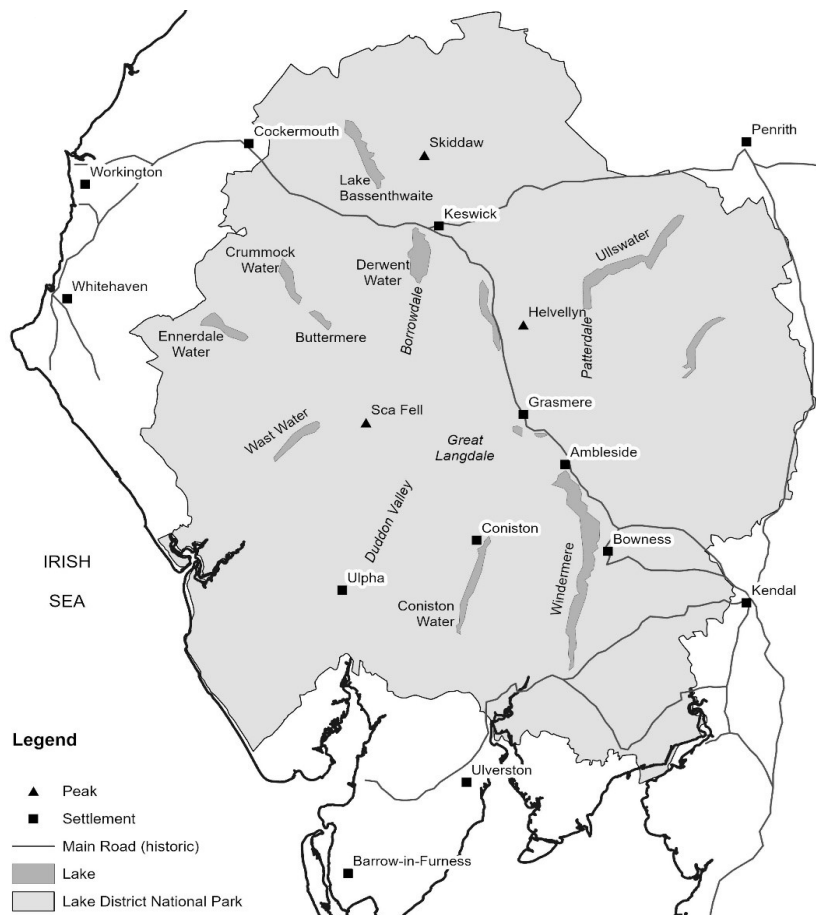
The field of geospatial analysis is seeing a surge of LLMs, particularly the Generative Pre-trained Transformer (GPT) models [6]. Recent research has explored their potential in various tasks. Mai et al. [26] demonstrated their use for tasks like pinpointing addresses, predicting medical trends, and even forecasting urban development. Similarly, Hu et al. [21] integrate geo-knowledge from location texts with a GPT model resulting in a geo-knowledge-guided GPT for extracting precise location information from social media messages about disasters. GeoGPT [54] takes this a step further by automating data collection, processing, and analysis based on natural language instructions. In the domain of geoscience, Deng et al. [10] developed K2, an LLM specifically trained on relevant text corpora, achieving impressive results in tasks like answering questions and summarizing information. Mooney et al. [29] highlighted the gap between human spatial understanding and current LLM training methods. They demonstrate this by testing ChatGPT on a GIS exam, revealing limitations in spatial literacy. Similarly, Ramrakhiyani et al. [38] focus on a specific aspect of spatial relations – extracting border orientations from Wikipedia text. Their research employs LLMs combined with natural language inference techniques to tackle this challenge. The presented work is also an extension of our recent study [20] on harnessing GPT model to extract spatial information from the CLDW and understand the spatial dimensions inherent in this historical archive comprehensively. The proposed framework extracted spatial relations from the CLDW and presented results for relation “near”.

These studies showcase the exciting potential of LLMs in geospatial analysis, while also acknowledging areas for further development. As research continues, LLMs have the ability to become powerful tools for understanding and manipulating spatial information.

## 3 Data and Methods

### 3.1 Study Area: The English Lake District

Located in the northwest of England, the Lake District (Figure 1) encompasses an area of approximately 2,362 square kilometers, making it the largest National Park in the country. Its topography is dominated by glaciated valleys, known locally as “dales,” which were sculpted in the Ice Age, giving rise to the distinctive U-shaped valleys and craggy peaks that define the area. The region has been subsequently shaped by the agro-pastoral land-use system [49]. The blend of natural processes and human activity has produced a harmonious landscape characterized by rugged mountains, ribbon lakes and wild meadows. The Lake District is home to sixteen major lakes, including Windermere, Derwentwater, and Ullswater, which contribute to its breathtaking scenery and recreational opportunities. The highest peak in England, Scafell Pike, stands within the Lake District, reaching an elevation of 978 meters. The region’s diverse habitats support a rich variety of flora and fauna, with ancient woodlands, heather-clad moors and meandering rivers, providing habitats for various species.



■ **Figure 1** Map of the Lake District [18].

### 3.2 Corpus of Lake District Writing

The Lake District region boasts a profound literary legacy, enhancing its cultural significance and has been designated as a UNESCO World Heritage “cultural site.” [32]. The sublime beauty of the landscape has inspired generations of writers and poets, such as William Wordsworth, Beatrix Potter, John Ruskin as well as lesser-known figures [8]. The CLDW is a compilation of 80 texts spanning nearly three centuries, from 1622 to 1900 and can be subdivided into three periods: the Age of Sensibility (1740-1788), the Romantic period (1789-1836) and the Victorian era (1837-1901), each period constituting a pivotal stage in shaping contemporary perspectives of the Lakes’ landscape and culture [11]. These texts, manually transcribed and georeferenced, encompass over 1.5 million words and represent various genres such as guidebooks, travelogues, and topographical literature [11]. The corpus texts, assembled for the Lancaster University Spatial Humanities project, were processed using NER via the Edinburgh Geoparser, linking locations mentioned in the texts to the Unlock gazetteer [40]. The CLDW provides insights into the cultural, historical, and geographical aspects of the Lake District, offering a diverse array of perspectives from various literary periods and genres. A statistical summary of the CLDW can be found in [13, 20].

■ **Table 1** Selected concepts from Kuhn’s foundational ontology.

Core Concepts of Spatial Information	
Base Concept	Location (Neighbourhood)
Content Concepts	Object, Network
Quality Concepts	Granularity, Accuracy

### 3.3 Proposed Approach

#### 3.3.1 Conceptual Modeling for Problem Domain

In spatial humanities and geographic information science, a conceptual definition of the notion “place” holds significant importance, especially as natural language systems gain prominence [37]. GIS typically represent environments as object-based (using geometric objects with regular boundaries and related attributes) or field-based (continuous representations of properties of the geographic environment, using sampling of attributes) [37, 9, 16]. In general, place information can vary in precision and detail, leveraging existing concepts and theories. Hence, a comprehensive model called the set of core concepts of spatial information [23, 1] is used to formalize the research problem. This foundational ontology bridges the gap between abstract notions of place and concrete spatial representations. Moreover, the associated concepts facilitate transdisciplinary research and offer an alternative to traditional GIS models. Table 1 lists the set of core concepts chosen to model the problem domain.

To begin with, “location” is the fundamental concept in understanding spatial information. It answers the question “where” and allows us to define places using various spatial relations, both geometric and linguistic [37]. Through location, anything can become a “place” we can experience and refer to, such as the western side of the lake, middle of the valley, or even the halfway point of a trip. Assigning identity to these locations transforms them into places, real or imagined, as long as they hold shared meaning. The second concept, an “object” is something distinct and identifiable with specific spatio-temporal features, associated thematic sense and can be related to other objects in different ways, including location, time, and purpose [37]. Both natural and man-made entities as well as people can be classified as objects. In the context of geography, objects are always associated with a specific location. The concept of a “geographic object” combines two different ideas: “Clearly defined objects” applies to things with clear boundaries like buildings; “Features of a surface” applies to things such as a landscape, which may not have a clearly defined boundary [37]. Locations are transformed into objects through naming or describing them, which can refer to different places over time, or even to locations that no longer exist. Hence, the same physical location can hold multiple places, potentially coinciding spatio-temporally [37]. By treating places as objects, we can capture their unique characteristics beyond their physical location. We can use properties and relations to represent all the aspects that make a place more than merely a point on a map. These properties and relations can include things like historical significance, cultural meaning, or personal experiences associated with the place.

The next concept, a “network”, is a collection of objects (nodes) existing in space and time, connected by relationships (the edges) [41]. Places, regarded as objects, can also be parts of networks, acting as either connection points (nodes) or the connections themselves



(edges) [50]. For instance, a city with road links to neighbouring cities becomes a node in the network. Similarly, a bridge connecting two towns across a river can also be considered a place within the network. These places can be positioned based on their connections within the network, without needing precise geographic coordinates. This way of understanding places through connections, rather than relying on geometry and exact locations, is a well-established concept in AI [37].

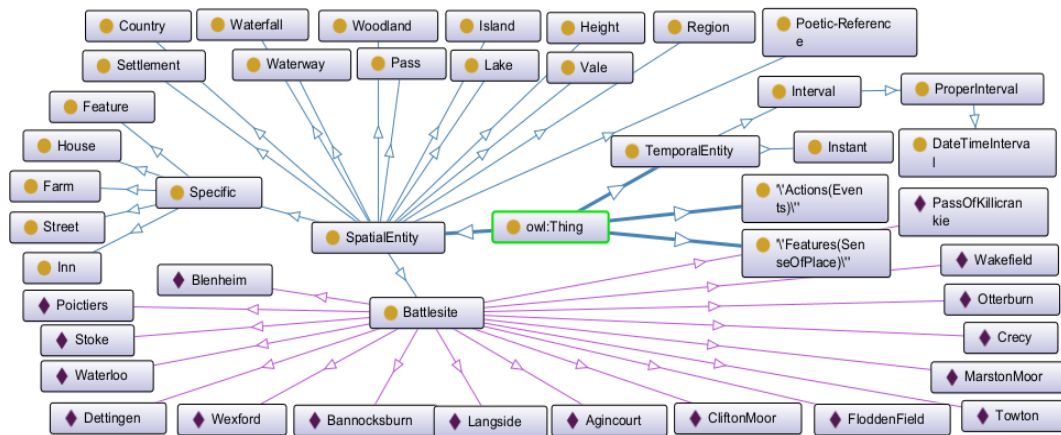
Among quality concepts, “granularity” defines the level of detail of spatial information as the location data can be extracted and analysed at different scales [23]. For example, we can focus on individual buildings or entire cities. The semantic classification of objects in the network can determine the level of granularity. Being a quality concept, the information granularity effects the decisions and their evaluations. Finally, “accuracy” is another key quality concept of spatial information [23]. It reflects the correctness and validity of information in representing the real world and needs relevant measures to assess the reliability and trustworthiness of the spatial data in hand.

### 3.3.2 The Spatial Ontology and the Custom CLDW Gazetteer

Having laid out the foundations of our work, we now elaborate on the ontology design for the domain knowledge representation. An ontology formally defines concepts and relationships within a domain, facilitating knowledge sharing, interoperability, and reasoning. The basic components of an ontology include class, subclass, object property (relates one class to another class), data property (is defined for something having a specific value) and individuals (instances). Standard geospatial ontologies [32, 7, 39, 51] delineating the geographical aspects of place fundamentally define qualitative spatial relations as object properties, for example,  $\langle \text{Entity: Place1, Object Property: left, Entity: Place2} \rangle$ . We have developed an ontology (see Figure 2) specifically tailored for the narrative corpus, encompassing spatial and temporal elements as well as their associated attributes. These attributes can also pertain to actions or events occurring within space and time. By detailing the relationships between feature classes and their sub-classes, we can effectively link the sense of place with specific entities. Our ontology design is rooted in a fundamental taxonomy of the narrative domain [13], integrating standard vocabularies and adhering to conventions found in geo-ontologies. To define spatial relationships as object properties, we have utilized GeoSPARQL [7] and Ordinance Survey ontologies [25], while temporal entity classes and sub-classes are defined using the OWL-Time ontology [35]. For spatial entity classes, we utilized a rich and notable source which is a corpus-specific classification named as the CLDW custom gazetteer classification scheme proposed by James Butler [40], described below. Finally, though the ontology is designed to incorporate spatial, temporal and thematic information; we focus on the spatial aspects of the ontology in this work.

We have employed the Protégé platform [31] to construct the ontology, enabling the recording of facts such as the classification of spatial entity “Pooley Bridge” as a settlement and its relationship recorded as an object property with the spatial entity “River Eamont” categorized as a waterway [13]. In addition to capturing specific entity facts, Protégé’s description logic facilitates the storage of general statements about classes and properties as axioms, allowing reasoning tasks to be performed using associated reasoners and querying the ontology for certain facts, such as to find all villages “near” a specific lake, or to identify the “highest mountain” “visible from” a particular settlement. Hence, these capabilities facilitate the further exploration of travel routes, landscape descriptions, and the relationships between different locations depicted in the writings. Beyond representing static entities and relationships, this structured knowledge base allows the tracing of trajectories or journeys

which are complex spatial entities. Each journey described in a narrative consists of a sequence of spatial entities interacting through spatial relations and actions [13]. For example, a journey might involve crossing the river Eamont, passing through Pooley Bridge, and returning to Penrith. Throughout the journey, experiences like seeing a lake through trees or encountering a rainstorm can be identified as event classes with instances arranged in sequence within the ontology [13]. By representing journeys and paths through landscapes in narratives using ontology, we aim to explore similarities and differences between journeys as experienced rather than simply plotted on a map.

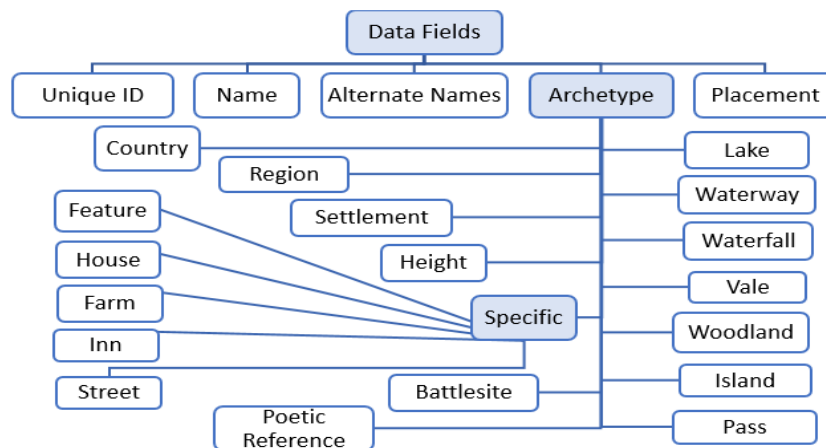


■ **Figure 2** Ontology graph with classes, sub-classes and some instances of the sub-class “BattleSite”.

We mentioned earlier that a classification system [40] has been used to categorize the various place names within the CLDW corpus (see Figure 3). This system was chosen because it offers a comprehensive list of geospatial categories specifically relevant to the corpus content. It is a manually compiled collection of every named place mentioned within the corpus. This manual process allowed the identification of variant forms, spellings, and variations of the same place name. For example, “Goat’s Water = Gates Water = Gait’s Water = Gait’s Tarn”, similarly, “Sparkling Tarn = Low Tarn = Sty Head Tarn.” The actual paper [40] should be studied for detailed understanding of the corpus preparation and annotation process. Nevertheless, we briefly provide important details on the classification system. The gazetteer uses modern spellings as the primary identifier for each location. These identifiers were assigned unique numeric codes, and any alternative spellings were linked back to the main entry. Additionally, each place name recognized as a valid and identifiable location was assigned two data points: its position within (Internal) or outside (External) the Lake District region, and its geographical category. This detailed classification system allowed for the distinction between places sharing similar names but differing types, for example, Ambleside Road [Pass], Ambleside Waterfall [waterfall], Ambleside valley [vale], Ambleside [settlement], or New Ambleside [house], thereby removing ambiguity and enabling a wider range of sophisticated geographical analyses. The system uses nineteen geospatial categories, striking a balance between detailed analysis and clear distinctions for geographical studies. Most categories are self-explanatory. While there is no standardized format for topographic features, these categories capture important and distinct elements within the texts, allowing for potential analysis of how authors describe different location types [40]. The “Specific” category covers named areas of tourist interest and man-made features (castles, stone circles, churches, bridges). This category essentially includes locations that do not fit the typical



environmental classifications. The “Features” category focuses on unique, named elements within larger locations. Examples include Scafell Chimney (on Scafell Pike), the Bowder Stone, and named edges [40].



■ **Figure 3** CLDW custom gazetteer hierarchy.

### 3.3.3 Integrated Framework for Extraction and Visualization

Having laid the foundation with the core concepts, domain ontology and CLDW custom gazetteer, we now present a holistic description of the overall approach. The definition of location, object and network lets us to formalize the research problem and extract places as spatially related to each other and organizing them in the form of a network of semantic triples. The spatial relation extraction process has been outlined with a framework diagram in our recent work [20] with initial results which gives step-wise details. In this paper, we systematically link the theoretical notions to provide a basis for the practical aspects with substantially detailed understanding.

From a technical perspective, formulating effective prompts is challenging and demands substantial effort. The prompt acts as a mechanism to influence the model’s behavior by providing context and direction. When the model receives a prompt, it adjusts its internal weights and parameters during the inference process to generate a response that aligns with the prompt’s objectives [20]. Essentially, the prompt guides the model in focusing its attention on relevant information and adjusting its predictions accordingly to accomplish the intended task more effectively. We formulate zero-shot prompts [53] leveraging the conception of the LLMs of the intended task. We expand the list of target spatial relations and extract LLM-generated list of most recurring proximity or vicinity spatial relations for the top five place names appearing in the CLDW. We also compare the extraction capabilities of three benchmark LLMs for a scholarly significant historical archive, representing their potential in one of the challenging and interdisciplinary research problems. Finally, the network information is enhanced by incorporating gazetteer-based classification of resultant objects improving the spatial profiling of the target place.

## 4 Experiments and Results

In [20], GPT-4 [34] has been used as the primary model for experimental evaluations. In this paper, we still base our extraction results and network generation on GPT-4 and compare the performance with two other LLMs. The selected academic benchmarks have been described below with highlights on their core features and accessibility information.

- GPT-4 [34], an advanced version of the GPT series, with a larger model size and more extensive training data, boasts improved contextual understanding, finer-grained text generation, and better handling of complex linguistic tasks. We evaluated GPT-4 using the Azure OpenAI services [2], which provide access to OpenAI’s powerful language models. The prompts are developed and passed to the model through a Python SDK.
- LLaMA-2 [48] represents a sophisticated language model tailored for continual learning and adaptation. Its autoregressive architecture enables it to generate coherent and contextually relevant text while continually updating its understanding based on incoming data, making it highly versatile for lifelong learning scenarios and adaptive NLP. We utilized the LLaMA-2-70B base model available from Perplexity Labs [36].
- Gemini 1.5 [47] introduces a novel approach to language modeling by leveraging dual-stream architectures, combining the strengths of autoregressive and autoencoding models enabling it to achieve a balance between context-aware generation and feature-based representation learning. We accessed the chat model based on Gemini 1.5 through public interface [15] whose responses are although not a substitute of the full Gemini model, which offers significant improvements in long-context understanding; it still offers acceptable performance in downstream NLP tasks.

### 4.1 Generating Spatial Proximity or Vicinity Prompts

For experimenting with prompts focusing on spatial relation extraction, we adopted an approach to generate customized zero-shot prompts [53]. In general, zero-shot prompting enables the model to generalize its knowledge across a wide range of tasks and domains, even those it has not been explicitly trained on. By utilizing its pre-existing understanding of language and concepts, the model can effectively infer the desired task from the prompt and generate coherent and contextually relevant outputs. Leveraging this ability, we generated input prompts from the three base LLMs. Specifically, for the input instruction “Write a prompt to extract spatial relations between two places or objects in proximity or vicinity. In your prompt, mention the most frequently used proximity relations by name.”, following response prompts (highlighted in grey boxes) have been generated by each of the LLM.

GPT-4: “Describe the spatial relations between two places or objects in terms of their proximity or vicinity. Please include details about their relative positions, distance, direction, and orientation. Use common spatial terms such as \“next to\”, \“near\”, \“far from\”, \“beside\”, \“behind\”, \“in front of\”, \“above\”, \“below\”, \“inside\”, \“outside\”, \“around\”, \“between\”, and \“adjacent to\”.”

**LLaMA2:** Please provide two places or objects that are in close proximity or vicinity to each other. Your response should include the following information:

1. The names or descriptions of the two places or objects.
2. The specific proximity relation that best describes their spatial relationship. Choose from the following options:

Adjacent: Sharing a common boundary or wall. Nearby: Located within a short distance, but not necessarily sharing a common boundary. Close proximity: Located within a close distance, but not necessarily sharing a common boundary or in a direct line of sight. Vicinity: Located in the same general area or neighborhood, but not necessarily in close proximity. Overlapping: Partially or completely sharing the same space.

3. A brief description of the context in which the two places or objects are located.

Example:

Place 1: A coffee shop. Place 2: A bookstore. Proximity relation: Adjacent. Context: Both places are located on the same street, with the coffee shop located next door to the bookstore. Please provide your response in the following format:

Place 1: [Name or description] Place 2: [Name or description] Proximity relation: [Adjacent, Nearby, Close proximity, Vicinity, Overlapping] Context: [Brief description of the context]

**Gemini-1.5:** Analyze the text and identify instances where two places or objects are mentioned in close proximity. Look for keywords and phrases that indicate these spatial relations, including:

Near: This is the most common term and indicates general closeness without specifying the distance. (“The school is located near the park.”) Beside/Next to: These terms signify immediate adjacency where the two entities share a common border. (“The cafe is located beside the bookstore.”) Across from: This term indicates that the entities are positioned opposite each other, usually with some separation. (“The library is located across from the street from the museum.”) Behind/In front of: These terms denote the relative position along a linear axis, specifying which entity is further back or further ahead. (“The house is located behind the trees.”) Within/Inside: These terms indicate that one entity is completely contained within the other. (“The office is located within the building.”)

Additionally, consider synonyms and paraphrases for these keywords to capture a wider range of spatial descriptions in the text.

Output: For each identified instance, provide the following information:

Entity 1: The first place or object mentioned. Spatial Relation: The identified spatial relation between the entities (e.g., near, beside, across from). Entity 2: The second place or object mentioned.

Example Output:

Entity 1: Eiffel Tower

Spatial Relation: near

Entity 2: Seine River

The contrasting difference between the prompt generation ability of the three models can be observed. GPT-4 response is succinct yet comprehensive. However, as opposed to the other two models, it did not explicitly mention the target format for the output relation triplet. Besides, the response covers a range of spatial keywords relevant to the query. The other two models attempt to formulate a descriptive prompt with examples included, giving the perception of a few-shot prompt. The response of LLaMA-2 provides a brief yet concise list of relation terms and highlights the extraction of relation context in addition to the required arguments. Both LLaMA-2 and Gemini 1.5 provide definitions for the selected relation terms. Finally, the mention of “synonyms and paraphrases” in Gemini-1.5 response indicates that the model tries to generate an all-inclusive prompt.

## 4.2 Spatial Relation Extraction and Network Visualization

Using the generated prompts, we selected the spatial relation terms as shown in Figure 4 with their frequency and appearance trend in the 80 text files of the CLDW [44]. It can be seen that the term “nearby” does not appear in the corpus at all. The terms with \* represent the root word with other derived terms included in the count. It can also be observed that we have added the term “surrounded” which has not been listed in any of the LLM’s response. We added this relation considering the fact that the word “surrounded” provides important information on objects in vicinity and its frequency value reflects its appearance more than the other terms in the list. Table 2 represents corpus frequency of the top five place names selected for analysis. Another level of filtering has been applied on each place name to count its contextual occurrence with the selected spatial relations of Figure 4.

The generated prompts in Section 4.1 provide information on the cognitive ability of the LLMs on spatial relation extraction task definition. The responses let us to create a generalized list of spatial relation terms. However, there is discrepancy among the prompt texts. GPT-4 does not mention the output format which should be a basic relation triple, though we did not mention it in the input prompt, the other two models explicitly state the output format. The second model Llama-2 generated a prompt where it mentions the extraction of “context” which creates ambiguity whether the context is the actual description or inference. Similarly, the mention of “synonyms and paraphrases” could possibly complicate the extraction task for a model. Keeping these observations in view and avoiding any kind of bias, we created a generalized prompt for spatial relations extraction in the triplet form as follows: “From the passages numbered in brackets (), extract spatial relation triplet only when the entity «place name» is found to be in a relation with other entities. The output should be in the form (subject, spatial relation, object). The list of spatial relations keywords to look for: ‘near, next to, adjacent, besides, within, inside, behind, in front of, vicinity, surrounding’.” «input passages».

	Term	Count
1	near	1611
2	within*	699
3	behind	613
4	surrounded	353
5	beside*	314
6	vicinity*	117
7	adjacent*	79
8	"in front of"	67
9	inside*	54
10	"next to"	40
11	nearby*	0

■ **Figure 4** Frequency of selected relation terms for the CLDW.

We compared the extraction performance of the three models on a subset of labeled input passages for spatial relation triples. Table 3 highlights the results for the selected place names. It can be observed that GPT-4 outperforms both of the models in zero-shot setting

with the same prompt instructions. It is imperative to describe necessary observations on their comparative performance. Though we highlighted the list of relation terms and output format explicitly, the generated output has variations similar to the way models generated proximity prompts for the same instruction. This entails that a less strict evaluation should be made focusing on the correctness and completeness in general. For instance, below is a snippet of the extraction results of the models for the place name “Ambleside”. It is clear that GPT-4 has exactly followed the output format required. However, the model extracts spatial relations other than the list provided, such as “distant from”. Moving to LLaMA-2, the model only extracted relations for the keyword “near” ignoring the other relations in the provided list. The model also is not precise in locating the boundary of the object entity. Another weak point of Llama-2’s result is the high redundancy of extractions, such as the last triple in the displayed output subset has been extracted twenty times. Finally, Gemini-1.5 generated a mixed response. The result marks the number of passages in which a certain relation appears which is acceptable when a relation triple appears in different contextual passages. However, since text mining methods have been used to generate passages, some of them have been generated more than once due to different query keywords appearing in a single passage. GPT-4 handled this redundancy and generated only one triple for any of such input passages as opposed to Gemini-1.5. However, similar to GPT-4, Gemini-1.5 also extracted relations that have not been asked in the instruction, such as the distance in miles, along with some triples with more than one relation which requires further processing. Moreover, as we mentioned the case of LLaMA-2, Gemini-1.5 extractions are also imprecise with respect to the object boundary. It can be concluded that both of the models are not robust and specific enough in their extraction strategy compared to GPT-4. Lastly, the models assumed the entity “Ambleside” as the subject of the triple which in general can be considered at this level since we did not explicitly specify in the instruction. Owing to the vagueness of certain spatial concepts, distinguishing the subject and object in a prompt is necessary based on the formal theories [52]. Such as nearness is not symmetric. Imagine “X” is a big city, such as “London”, and “Y” is a tiny village very close to the boundary of the city “X”. We say “Y is near X” but we do not say “X is near Y”.

GPT-4:

```
<Ambleside, near, banks of the murmuring Rothay>\n
<Ambleside, distant from, Hawkshead and Esthwaite Water>\n
<Ambleside, near, road>\n
<Ambleside, in front of, Rydal Hall>\n
<Ambleside, behind, vale of Lune>\n
<Ambleside, behind, village>\n
```

LLama-2: Near: Ambleside is near the banks of the River Rothay, which of the fells crowd over the scene.

Near: Ambleside is near the head of Windermere.

Near: Ambleside is near the road to Clappersgate, winding on the banks of the Brathay.

Near: Ambleside is near the Roman camp at the mouth of the River Brathay.

Near: Ambleside is near the Roman station at the head of Windermere.

Near: Ambleside is near the Roman fortress at the head of the finest ravine in the country.

## 11:14 Semantic Perspectives on the Lake District Writing

Gemini-1.5: Here are the spatial relations found in the passages where “Ambleside” is the subject:

near the banks of the murmuring Rothay [1, 3, 4]

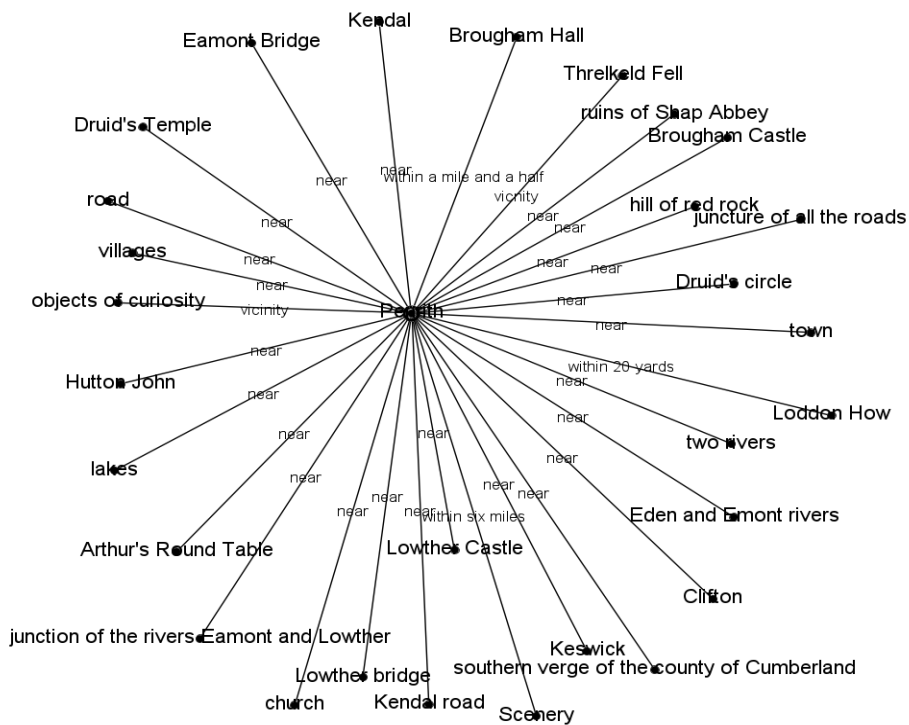
below Ambleside, in a meadow near the River Brathay [5, 6]

near the source of the Brathay [11, 12]

near the Woolen Mill [15, 16]

a few hundred yards above Skelwith Bridge( three miles from Ambleside) [18]

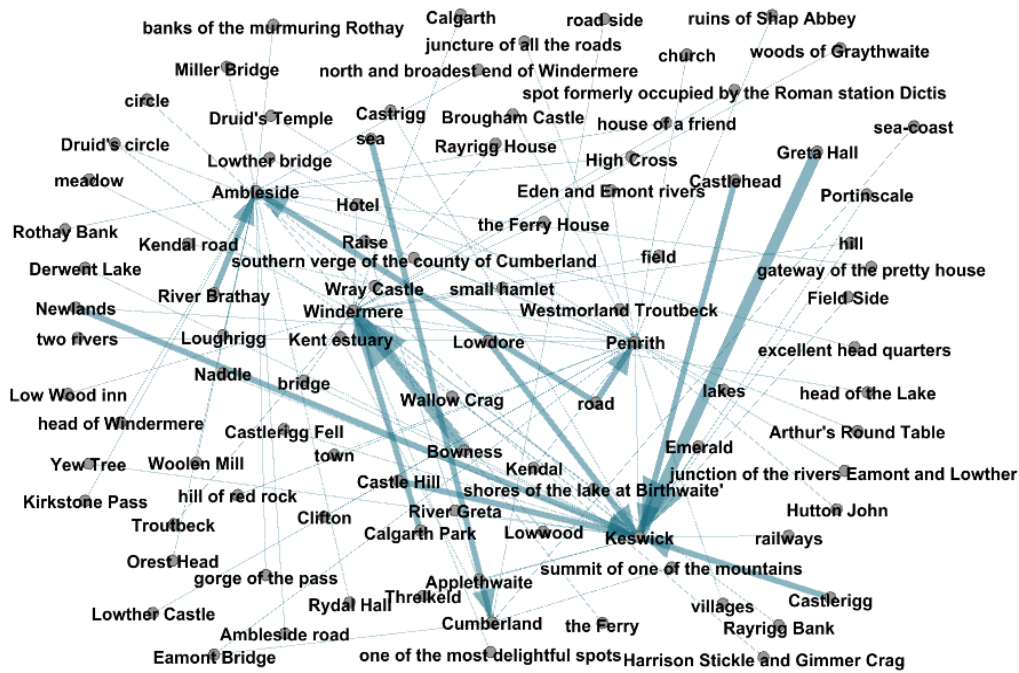
near the head of Windermere [22, 35, 36]



■ **Figure 5** Network visualization for “Penrith” displaying GPT-4 extracted triples before post-processing.

Figure 5 [3] displays semantic triples extracted for the place “Penrith” before applying any post-processing including custom gazetteer-based toponym linking and categorization. Another set of triples in Figure 6 show the network of selected five place names (Table 2) for the relation “near”. The CLDW custom gazetteer classification has been incorporated in the spatial ontology as described in Section 3.3.2 and the entities in the extracted spatial triplets are now assigned their respective geonoun categories. The network of semantic triples in Figure 7 containing the qualitatively located entities is displayed for one of the cases (Subject Entity-Near-Keswick). We distinctively keep “Keswick” as the target entity in the network visualization. The thickness of edges represent the frequency of spatial connection in the semantic triples list. The color coding represents the categories of different places based on





■ **Figure 6** Network visualization of five major place names for the relation “near”.

■ **Table 2** Most frequently occurring place names and the frequency of their involvement in the selected relations.

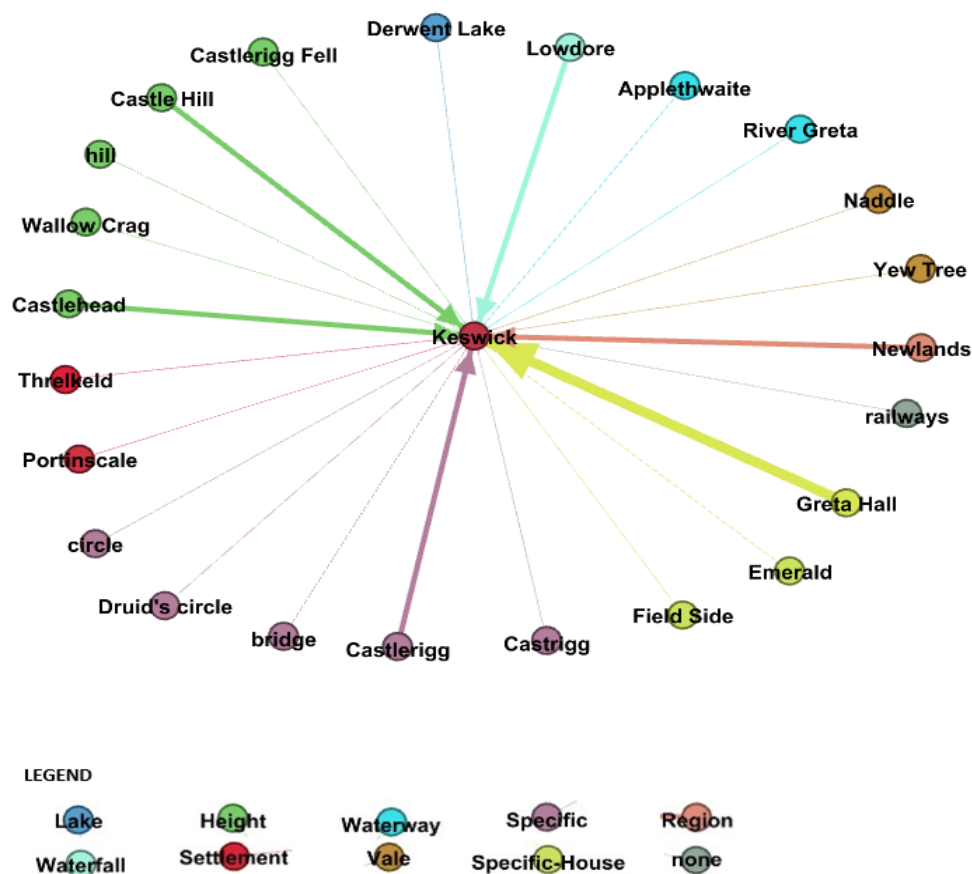
Place	Corpus Frequency	Context Count with selected relations
Keswick	1452	258
Ambleside	900	122
Windermere	877	135
Penrith	715	119
Cumberland	694	131

■ **Table 3** Comparative results of spatial relation extraction.

Parameter	GPT-4	LlaMA-2	Gemini-1.5
Precision	0.70	0.61	0.63
Recall	0.71	0.30	0.54
F-measure	0.70	0.40	0.58

## 11:16 Semantic Perspectives on the Lake District Writing

their classification. These categories reflect the granularity of spatial objects considered for relation extraction. Moreover, they also represent the overall spatial profile of the target place which may lead to meaningful inferences. For instance, these categories can be used to analyze linguistic patterns, including how often certain environmental aspects are compared within the text which could be illustrated through a network. This analysis might extend to comparisons with local sites or more distant, well-known places, whether within the same sentence, paragraph, or any designated span [40].



■ **Figure 7** Network visualization for (Subject Entity-Near-Keswick) generated for GPT-4 results with ontology-mapped categorization.

## 5 Conclusion

This research is part of a broader effort to better understand the spatial and temporal semantics of narratives [13]. The overall goal is to develop methods for analyzing extensive historical archives beyond conventional GTA. In this paper, we outlined a theoretical foundation for modelling notions of place and subsequently showed how such notions can be extracted automatically. We propose an extended approach for extracting spatial relations from the CLDW and demonstrate its effectiveness by focusing on a set of relations defining the proximity or vicinity of located objects. The results are visualized as a network that shows the target place and its surrounding spatial entities. Our approach offers a distinctive

computational perspective on the notion of place, which can enhance the ability of social scientists and humanists to interpret narrative descriptions of location. The underlying conceptual and ontological modeling allows a formal and insightful exploration of the massive unstructured corpus. The custom CLDW gazetteer facilitated the analysis in a number of ways including entity classification in ontology, toponym extraction, matching and an enhanced network visualization. The extraction results show the effectiveness of zero-shot prompting, which holds promise for enabling more flexible and adaptive language models that can seamlessly handle new tasks and domains without the need for extensive retraining.

We aim to extend this work in a number of directions. In particular, we intend to extend the ontology schema for the feature/sense of place and actions/events classes, populate existing ontology classes and properties to model journey trajectories and perform spatial reasoning. Additionally, we will work towards improving the extraction performance by experimenting with few-shot and chain-of-thought prompting methods, along with expanding the list of entity and spatial relations, linking geonoun references and disambiguate entities.

---

## References

- 1 Christopher Allen, Thomas Hervey, Sara Lafia, Daniel W. Phillips, Behzad Vahedi, and Werner Kuhn. Exploring the notion of spatial lenses. In Jennifer A. Miller, David O’Sullivan, and Nancy Wiegand, editors, *Geographic Information Science*, pages 259–274, Cham, 2016. Springer International Publishing.
- 2 Azure openai service models (GPT-4), 2023. URL: <https://learn.microsoft.com/en-us/azure/aiservices/openai>.
- 3 Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: an open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, ICWSM, pages 361–362. AAAI, 2009. doi:10.1609/icwsm.v3i1.13937.
- 4 David J Bodenhamer, John Corrigan, Trevor M Harris, and et al. *The Spatial Humanities: GIS and the Future of Humanities Scholarship*. Indiana University Press, 2010. URL: <http://www.jstor.org/stable/j.ctt16gzj77>.
- 5 Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, and et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- 6 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and et al. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS’20, Red Hook, NY, 2020. Curran Associates, Inc.
- 7 Nicholas J. Car and Timo Homburg. Geosparql 1.1: Motivations, details and applications of the decadal update to the most important geospatial lod standard. *ISPRS International Journal of Geo-Information*, 11(2), 2022. doi:10.3390/ijgi11020117.
- 8 David Cooper. The poetics of place and space: Wordsworth, norman nicholson and the lake district. *Literature Compass*, 5(4):807–821, 2008.
- 9 Helen Couclelis. People manipulate objects (but cultivate fields): Beyond the raster-vector debate in gis. In A. U. Frank, I. Campari, and U. Formentini, editors, *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, pages 65–77, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.
- 10 Cheng Deng, Tianhang Zhang, Zhongmou He, Qiyuan Chen, Yuanyuan Shi, Le Zhou, Luoyi Fu, Weinan Zhang, Xibing Wang, Chenghu Zhou, and et al. Learning A foundation language model for geoscience knowledge understanding and utilization. *CoRR*, abs/2306.05064, 2023. doi:10.48550/arXiv.2306.05064.

- 11 Christopher Donaldson, Ian N. Gregory, and Joanna E. Taylor. Locating the beautiful, picturesque, sublime and majestic: spatially analysing the application of aesthetic terminology in descriptions of the english lake district. *Journal of Historical Geography*, 56(1):43–60, 2017. doi:10.1016/j.jhg.2017.01.006.
- 12 Stuart Dunn, Graeme Earl, Anna Foka, and Will Wootton. Spatial narratives in museums and online: The birth of the digital object itinerary. *Museums and Digital Culture: New Perspectives and Research*, pages 253–271, 2019.
- 13 Ignatius Ezeani, Paul Rayson, Ian N. Gregory, Erum Haris, Anthony G. Cohn, John G. Stell, Tim Cole, Joanna E. Taylor, David Bodenhamer, Neil Devadasan, and et al. Towards an extensible framework for understanding spatial narratives. In *Proceedings of the 7th ACM SIGSPATIAL International Workshop on Geospatial Humanities*, GeoHumanities '23, pages 1–10, New York, NY, 2023. ACM Press. doi:10.1145/3615887.3627761.
- 14 Anna Foka, Elton Barker, Kyriaki Konstantinidou, Nasrin Mostofian, O. Cenk Demiroglu, Brady Kiesling, and Linda Talatas. Semantically geo-annotating an ancient greek "travel guide" itineraries, chronotopes, networks, and linked data. In *Proceedings of the 4th ACM SIGSPATIAL Workshop on Geospatial Humanities*, GeoHumanities '20, pages 1–9, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3423337.3429433.
- 15 Gemini, 2024. URL: <https://gemini.google.com/>.
- 16 Michael F. Goodchild, May Yuan, and Thomas J. Cova. Towards a general theory of geographic representation in gis. *International Journal of Geographical Information Science*, 21(3):239–260, 2007. doi:10.1080/13658810600965271.
- 17 Ian Gregory, Christopher Donaldson, Patricia Murrieta-Flores, and Paul Rayson. Geoparsing, gis, and textual analysis: current developments in spatial humanities research. *International Journal of Humanities and Arts Computing*, 9(1):1–14, 2015.
- 18 Ian Gregory, Ian Smail, Joanna Taylor, and James Butler. Exploring qualitative geographies in large volumes of digital text: Placing tourists, travellers and inhabitants in the english lake district. *Annals of the American Association of Geographers*, May 2024.
- 19 Erum Haris, Anthony G. Cohn, and John G. Stell. Understanding the Spatial Complexity in Landscape Narratives Through Qualitative Representation of Space. In Roger Beecham, Jed A. Long, Dianna Smith, Qunshan Zhao, and Sarah Wise, editors, *12th International Conference on Geographic Information Science (GIScience 2023)*, volume 277 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 37:1–37:6, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2023/18932>.
- 20 Erum Haris, Anthony G. Cohn, and John G. Stell. Exploring spatial representations in the historical lake district texts with llm-based relation extraction. In Xuke Hu, Ross Purves, Ludovic Moncla, Jens Kersten, and Kristin Stock, editors, *Proceedings of The GeoExT 2024: Geographic Information Extraction from Texts Workshop co-located with The 46th European Conference on Information Retrieval (ECIR), Glasgow, Scotland, March 24, 2024*, volume 3683 of *CEUR Workshop Proceedings*, pages 63–73. CEUR-WS.org, 2024. URL: <https://ceur-ws.org/Vol-3683/paper9.pdf>.
- 21 Yingjie Hu, Gengchen Mai, Chris Cundy, Kristy Choi, Ni Lao, Wei Liu, Gaurish Lakhnopal, Ryan Zhenqi Zhou, and Kenneth Joseph. Geo-knowledge-guided gpt models improve the extraction of location descriptions from disaster-related social media messages. *International Journal of Geographical Information Science*, 37(11):2289–2318, 2023. doi:10.1080/13658816.2023.2266495.
- 22 Yingjie Hu, Xinyue Ye, and Shih-Lung Shaw. Extracting and analyzing semantic relatedness between cities using news articles. *International Journal of Geographical Information Science*, 31(12):2427–2451, 2017. doi:10.1080/13658816.2017.1367797.
- 23 Werner Kuhn. Core concepts of spatial information for transdisciplinary research. *International Journal of Geographical Information Science*, 26(12):2267–2276, 2012. doi:10.1080/13658816.2012.722637.

- 24 Corpus of lake district writing, 1622-1900, 2017. URL: <https://github.com/UCREL/LakeDistrictCorpus>.
- 25 Spatial relations ontology (osspr), 2013. URL: <https://lov.linkeddata.es/dataset/lov/vocabs/osspr>.
- 26 Gengchen Mai, Weiming Huang, Jin Sun, Suhang Song, Deepak Mishra, Ninghao Liu, Song Gao, Tianming Liu, Gao Cong, Yingjie Hu, and et al. On the opportunities and challenges of foundation models for geoai (vision paper). *ACM Trans. Spatial Algorithms Syst.*, March 2024. Just Accepted. doi:10.1145/3653070.
- 27 Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David Lobell, and Stefano Ermon. Geollm: Extracting geospatial knowledge from large language models. *arXiv preprint arXiv:2310.06213*, 2023.
- 28 Molly Miranker and Alberto Giordano. Text mining and semantic triples: Spatial analyses of text in applied humanitarian forensic research. *Digital Geography and Society*, 1:100005, 2020. doi:10.1016/j.diggeo.2020.100005.
- 29 Peter Mooney, Wencong Cui, Boyuan Guan, and Levente Juhász. Towards understanding the geospatial skills of chatgpt: Taking a geographic information systems (gis) exam. In *Proceedings of the 6th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, GeoAI '23, pages 85–94, New York, NY, 2023. ACM Press. doi:10.1145/3615886.3627745.
- 30 Patricia Murrieta-Flores and Bruno Martins. The geospatial humanities: past, present and future. *International Journal of Geographical Information Science*, 33(12):2424–2429, 2019. doi:10.1080/13658816.2019.1645336.
- 31 Mark A. Musen. The protégé project: A look back and a look forward. *AI Matters*, 1(4):4–12, June 2015. doi:10.1145/2757001.2757003.
- 32 Vatsala Nundloll, Robert Smail, Carly Stevens, and Gordon Blair. Automating the extraction of information from a historical text and building a linked data model for the domain of ecology and conservation science. *Heliyon*, 8(10):e10710, 2022. doi:10.1016/j.heliyon.2022.e10710.
- 33 Ian N. Gregory Olga Chesnokova, Joanna E. Taylor and Ross S. Purves. Hearing the silence: finding the middle ground in the spatial humanities? extracting and comparing perceived silence and tranquillity in the english lake district. *International Journal of Geographical Information Science*, 33(12):2430–2454, 2019. doi:10.1080/13658816.2018.1552789.
- 34 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, and et al. Gpt-4 technical report, 2024. arXiv:2303.08774.
- 35 Feng Pan and Jerry R Hobbs. Time ontology in owl. *W3C working draft, W3C*, 1(1):1, 2006.
- 36 Perplexity labs, 2024. URL: <https://labs.perplexity.ai/>.
- 37 Ross S. Purves, Stephan Winter, and Werner Kuhn. Places in information science. *Journal of the Association for Information Science and Technology*, 70(11):1173–1182, 2019. doi:10.1002/asi.24194.
- 38 Nitin Ramrakhiani, Vasudeva Varma, and Girish Keshav Palshikar. Extracting orientation relations between geo-political entities from their wikipedia text. In *Proceedings of the First International Workshop on Geographic Information Extraction from Texts*, GeoExT '23, pages 44–50. CEUR Workshop Proceedings, 2023.
- 39 Babak Ranjgar, Abolghasem Sadeghi-Niaraki, Maryam Shakeri, and Soo-Mi Choi. An ontological data model for points of interest (poi) in a cultural heritage site. *Heritage Science*, 10(1):13, 2022.
- 40 Paul Rayson, Alex Reinhold, James Butler, Chris Donaldson, Ian N. Gregory, and Joanna Taylor. A deeply annotated testbed for geographical text analysis: The corpus of lake district writing. In *Proceedings of the 1st ACM SIGSPATIAL Workshop on Geospatial Humanities*, GeoHumanities '17, pages 9–15, New York, NY, 2017. ACM Press. doi:10.1145/3149858.3149865.

- 41 Simon Scheider and Tom de Jong. A conceptual model for automating spatial network analysis. *Transactions in GIS*, 26(1):421–458, 2022. doi:10.1111/tgis.12855.
- 42 Evan Sheehan, Chenlin Meng, Matthew Tan, Burak Uzkent, Neal Jean, Marshall Burke, David Lobell, and Stefano Ermon. Predicting economic development using geolocated wikipedia articles. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 2698–2706, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3292500.3330784.
- 43 Alessio Signorini, Alberto Maria Segre, and Philip M Polgreen. The use of twitter to track levels of disease activity and public concern in the us during the influenza a h1n1 pandemic. *PloS one*, 6(5):e19467, 2011.
- 44 Stéfan Sinclair and Geoffrey Rockwell. Voyant tools, 2016. URL: <http://voyant-tools.org/>.
- 45 Robert Smail, Ian Gregory, and Joanna Taylor. Qualitative geographies in digital texts: Representing historical spatial identities in the lake district. *International Journal of Humanities and Arts Computing*, 13(1-2):28–38, 2019.
- 46 Erik Steiner, Zephyr Frank, Ian Gregory, David Bodenhamer, and Ignatius Ezeani. Spatio-textual Regions: Extracting Sense of Place from Spatial Narratives. In R Westerholt and F. B. Mocnik, editors, *Proceedings of the Fourth International Symposium on Platial Information Science (PLATIAL'23)*, pages 1–8, Dortmund, Germany;, 2023. PLATIAL'X.
- 47 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, and et al. Gemini: A family of highly capable multimodal models, 2024. arXiv:2312.11805.
- 48 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and et al. Llama: Open and efficient foundation language models, 2023. arXiv:2302.13971.
- 49 The english lake district, 2017. URL: <https://whc.unesco.org/en/list/422/>.
- 50 Maria Vasardani, Sabine Timpf, Stephan Winter, and Martin Tomko. From descriptions to depictions: A conceptual framework. In Thora Tenbrink, John Stell, Antony Galton, and Zena Wood, editors, *Spatial Information Theory*, pages 299–319, Cham, 2013. Springer International Publishing.
- 51 Barney Warf and Daniel Sui. From gis to neogeography: ontological implications and theories of truth. *Annals of GIS*, 16(4):197–209, 2010.
- 52 Michael F. Worboys. Nearness relations in environmental space. *Intelligent Systems with Applications*, 15(7), 2001. doi:10.1080/13658810110061162.
- 53 Chenhan Yuan, Qianqian Xie, and Sophia Ananiadou. Zero-shot temporal relation extraction with chatgpt. In *The 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, pages 92–102. ACL, 2023. doi:10.18653/v1/2023.bionlp-1.7.
- 54 Yifan Zhang, Cheng Wei, Shangyou Wu, Zhengting He, and Wenhao Yu. Geogpt: Understanding and processing geospatial tasks through an autonomous gpt, 2023. arXiv:2307.07930.