



DISSERTATION | DOCTORAL THESIS

Titel | Title

Geographic Concept Drift in Research Data Management

verfasst von | submitted by

Meilin Shi

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Doktorin der Naturwissenschaften (Dr.rer.nat.)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree
programme code as it appears on the
student record sheet:

UA 796 605 452

Dissertationsgebiet lt. Studienblatt | Field of
study as it appears on the student record
sheet:

Geographie

Betreut von | Supervisor:

Univ.-Prof. Dr. Krzysztof Janowicz

Acknowledgements

I would like to trace the beginnings of my PhD journey back to my early days at UC Santa Barbara. I took the GIS classes during my second year of college, where I first met my advisor, Jano. It is hard to believe that almost ten years have passed since then. I took every class that Jano offered, and they were absolutely among the most interesting classes I had during my college years, also the reason I later decided to pursue a PhD with him. Knowing how competitive the geography PhD program at UCSB was, I am deeply grateful for the support I received from Prof. Kostas Goulias, Vena Chu, Jennifer King, and, above all, from Jano, when I applied.

After spending eight years in Santa Barbara, starting with my bachelor's, Jano offered me the chance to continue my PhD with him when he took a new position at the University of Vienna. It was the kind of once-in-a-lifetime opportunity, and since I had studied German for a couple of years during my undergrad and was longing for Europe, I accepted his offer right away. From the bottom of my heart, I am grateful to Jano for giving me the chance to live and work on my PhD in Vienna. It was an incredible experience. I cannot imagine any other city could offer the chance for us to have our weekly group meetings in different cafés, where we had endless brainstorming and discussions on research, teaching, and life, accompanied by Wiener Melange. Thanks to Jano for suggesting the coffee meeting idea and making it all so memorable. I realize it may sound cliché to thank Jano as my advisor for his guidance and mentorship during my PhD, but I have learned far more from him over the ten years I have known and worked with him than just research.

Acknowledgements

I am especially grateful to him for giving me the freedom to explore whatever topics that I'm interested in and always providing the most constructive feedback. Most importantly, I am deeply thankful for the school of thinking he instilled in me over the years, which has shaped the way I value ideas and research to this day.

I could not have come this far in my PhD without the help of the STKO members. I would like to thank Song for giving me my very first research experience during the summer of my sophomore year, and Ling, who was both a cohort and mentor in the early stages of my research career. I am also grateful to Bo, Rui, Gengchen, and Blake, who helped build my knowledge of geosemantics and knowledge graphs when I had zero experience upon joining the group. Special thanks to Kitty and Zilong for sharing the same transition from Santa Barbara to Vienna, for being fellow expats, and for all the explorations we enjoyed in Vienna. For the Vienna group, Zilong, Ivan, Mina, and Alex, who were involved in all three papers of my PhD dissertation, thank you for your ideas, discussions, feedback, and polishing. A shout-out to Zilong again for all the co-teaching, research discussions, and for being my neighbor during my time in Vienna. Finally, without Jano and his way of leading the group, I would not have this STKO experience. I would like to express my gratitude to all STKO members, who have become friends and family over the years.

My PhD journey, though a bit longer and more winding than others, has been overall quite enjoyable, just missing some Vitamin D during the winters after moving from California to Vienna. Many thanks to my friends and cohort back in the States who came to visit and supported me along the way. Lastly, to the two most important people in my life, my parents, for the way they raised me, for their vision, for believing in me, for being more like friends and life mentors, and for supporting every decision I made. Having moved from China to the US, then to Austria, and now settled in the UK, I have always felt close to my parents despite the physical distance. I would not be where I am today without them.

Abstract

“One cannot step into the same river twice,” stated Heraclitus (c. 535 - c. 475 BC), emphasizing the perpetual flux of all things. This statement then triggers questions on what makes a river remain “the same” and the nature of change itself. It draws attention to a fundamental duality: change unfolds both in physical space, as with the flowing water, and in semantic space, through evolving conceptual boundaries that define the “river”. This dissertation examines drifting concepts within geographic domains, ranging from abstract keywords in geographic information science (GIScience) research (e.g., spatial reasoning, urban planning) to physical geographic features (e.g., forests, mountains, islands), investigating how their representations evolve across two main approaches: natural language and knowledge graphs. Throughout the dissertation, I focus on three aspects of conceptual changes via three chapters: (1) identifying different types of changes, including concept birth, drift, and split; (2) examining how these changes occur across both space and time; and (3) evaluating the consequences of conceptual changes in geo-knowledge graphs. Using scientometric analysis of GIScience literature and ontology change impact analysis, I demonstrate that geographic concepts exhibit varying drift patterns, that drifts happen across both space and time, and that even minor definitional changes can cascade through spatio-temporal reasoning. In examining concept drift, my dissertation contributes to research data management, FAIR data principles, and the broader domain of knowledge representation and reasoning by providing frameworks that document different types of conceptual changes, evaluate their consequences, and support the management of evolving geographic knowledge systems.

Zusammenfassung

“Man kann nicht zweimal in denselben Fluss steigen”, erklärte Heraklit (ca. 535 – ca. 475 v. Chr.) und betonte damit den ständigen Wandel aller Dinge. Diese Aussage wirft Fragen darüber auf, was einen Fluss “derselben” bleiben lässt und was das Wesen der Veränderung an sich ist. Sie lenkt die Aufmerksamkeit auf eine grundlegende Dualität: Veränderung vollzieht sich sowohl im physischen Raum, wie beim fließenden Wasser, als auch im semantischen Raum, durch sich entwickelnde konzeptuelle Grenzen, die den “Fluss” definieren. Diese Dissertation untersucht driftende Konzepte im geografischen Bereich, von abstrakten Schlüsselbegriffen in der Geoinformationswissenschaft (GIScience) (z. B. räumliches Denken, Stadtplanung) bis hin zu physischen geografischen Merkmalen (z. B. Wälder, Berge, Inseln), und untersucht, wie sich ihre Repräsentationen anhand von zwei Hauptansätzen—natürlicher Sprache und Wissensgraphen—entwickeln. In dieser Dissertation konzentriere ich mich in drei Kapiteln auf drei Aspekte konzeptioneller Veränderungen: (1) Identifizierung verschiedener Arten von Veränderungen, darunter die Entstehung, Drift und Aufspaltung von Konzepten; (2) Untersuchung, wie diese Veränderungen räumlich und zeitlich auftreten; und (3) Bewertung der Folgen konzeptioneller Veränderungen in Geo-Wissensgraphen. Anhand einer scientometrischen Analyse der GIScience-Literatur und einer Analyse der Auswirkungen von Ontologieveränderungen zeige ich, dass geografische Konzepte unterschiedliche Driftmuster aufweisen, dass Drifts sowohl räumlich als auch zeitlich auftreten und dass selbst geringfügige Definitionssänderungen eine Kaskade von räumlich-zeitlichen Schlussfolgerungen auslösen können. Durch die Untersuchung von Concept Drift leistet meine Dissertation einen Beitrag zum

Zusammenfassung

Forschungsdatenmanagement, zu den FAIR-Datenprinzipien und zum breiteren Bereich der Wissensrepräsentation und des Schlussfolgerns, indem sie Frameworks bereitstellt, die Arten von konzeptionellen Veränderungen dokumentieren, deren Folgen bewerten und die Verwaltung sich entwickelnder geografischer Wissenssysteme unterstützen.

Contents

Acknowledgements	i
Abstract	iii
Zusammenfassung	v
1 Introduction	1
1.1 Motivation	2
1.2 Space and Time in Concept Representation	4
1.3 Research Questions	7
1.4 Structure of the Dissertation	11
2 Dissertation Synopsis	15
2.1 Research Design	15
2.2 Key Findings	18
3 Defining Concept Drift and Its Variants in Research Data Management	21
4 Detecting Spatio-Temporal Concept Drift in Scientific Texts	49
5 Quantifying Spatial, Temporal, and Semantic Consequences of Ontology	
Changes	77

6 Conclusions and Future Work	111
6.1 Summary	111
6.2 Research Contributions	114
6.2.1 Theoretical Contributions	114
6.2.2 Practical Contributions	116
6.3 Limitations and Future Work	117
Bibliography	121

1 Introduction

This chapter starts with the motivation for studying concept drift, that is, changes in a concept’s meaning over time [48]. For example, research on changing climate had originally been linked to keywords such as “global warming”, while nowadays “climate change” is preferred [37]. This is not just indicative of changing keywords, but of the conceptual understanding of how the change plays out geographically. However, this kind of transition is not naturally captured in traditional keyword-based search, which means that historical data labeled with *global warming* might be missed in modern searches. Next, I will present background readings on concept representation, more specifically on how space and time are incorporated into it. I will outline the challenges posed by concept drift for research data management (RDM) and identify key research gaps. These lead to the **overarching goal** of this dissertation: to develop methods for handling conceptual changes in knowledge representation and to support the management of evolving knowledge in RDM systems. This goal is addressed through three research questions, forming a progressive inquiry: **(1)** how to define different types of conceptual changes, **(2)** how to detect them across both space and time, and **(3)** how to quantify their impacts in structured knowledge systems. Finally, this chapter outlines the structure of the dissertation.

1.1 Motivation

With advancing technology, information networks continue to expand, not only by the influx of new data but also through active digitization and preservation efforts of historical records [53]. However, in both circumstances, without proper documentation, such data can often not be meaningfully interpreted or made FAIR (Findable, Accessible, Interoperable, and Reusable) [49]. For example, the concept *artificial intelligence* in the 1960s mostly referred to symbolic reasoning systems and rule-based approaches, whereas later its meaning shifted more toward machine learning and nowadays deep learning specifically. When using keyword-based search, cases may arise in which the same keyword has different interpretations over time, or historical records use entirely different sets of keywords describing the same concept. This phenomenon is also referred to as *semantic aging* for long-term digital records preservation [36]. It has received growing attention in the fields of linguistics [20, 27], machine learning [28, 29], the Semantic Web [48, 43, 7], and so on. Note that even state-of-the-art large language models (LLMs) cannot entirely solve these issues, as they are themselves affected by concept drift and contextual dependency, even though they may be able to detect shifts in meaning via attention mechanisms.

Beyond technological progress, political and socio-cultural shifts have also reshaped the meanings of concepts such as *gender*, *marriage*, and *privacy*, as compared to a generation ago. The dynamic nature of language makes RDM no exception: the keywords that are used to describe data evolve over time, e.g., the aforementioned transition from *global warming* to *climate change*, or the coining of new terms like *spatial data science* and *GeoAI* [25]. These evolving concepts pose significant challenges for maintaining data usability and interoperability, especially given the ever-increasing volume of data in diverse formats and modalities. This has motivated the study of concept drift in RDM, which aims to systematically track and handle conceptual changes in research domains and general knowledge systems, thereby supporting the FAIR principles.

While concept drift is relevant across all domains from an RDM perspective, geographic concepts demand special attention because of their spatio-temporal nature, which makes them uniquely dynamic [5, 35]. For instance, there are hundreds of *forest* definitions with substantial differences across countries, organizations, and over time [3, 8, 6]. In other words, what qualifies as a forest in India (e.g., based on minimum area, tree height, and canopy cover) may not meet the definition of a forest in Brazil. Even within the same country, the definition can change because of policies, environmental factors such as deforestation, or simply due to climate change. Such variability makes geographic concepts a compelling case for studying concept drift in RDM, because it would directly affect the dataset discovery, integration, and interpretation over time *and* across space.

Building on these unique spatio-temporal properties, this dissertation also draws motivation from early work in geographic information science (GIScience) to propose formalized geographic knowledge through ontologies [32, 12]. Geographic concepts and categories have also been a topic of interest at the intersection of cognitive science and GIScience, e.g., regarding how they vary across cultural and regional contexts [46]. For example, the differences between *forests* and *woods*, *mountains* and *hills*, *lakes* and *ponds* are not universally agreed on and depend mostly on local conditions [42, 4]. These examples show that to explore geographic concept drift, space and time are two equally important components, whereas prior work on concept drift [48, 43, 7] has focused *almost exclusively on the temporal dimension*.

In this dissertation, the geographic concepts of interest include both physical features such as *mountains*, *forests*, *lakes*, and research fields such as *geographic information science*, *volunteered geographic information*, and *urban planning*. These concepts, *signified* by keywords, are used in RDM either to label datasets or to describe research publications. When these keywords (and their underlying concepts) undergo change, over time and possibly also across space, they may no longer accurately reflect the associated data or publications, thereby hindering FAIR-based retrieval. This dissertation, therefore, aims


1 Introduction

to (1) refine the definition of concept drift in the context of RDM, (2) showcase drifting concepts in the field of GIScience, and (3) systematically detect and quantify geographic concept drift over space and time to support the long-term usability of research data in the rapidly growing and evolving information landscape.


1.2 Space and Time in Concept Representation

Before examining the challenges that concept drift poses for RDM in detail, it is necessary to first understand how concepts are represented and function in knowledge systems. Concepts are mental representations of categories [2]. They are the basis of vocabularies and taxonomies in knowledge systems and the building blocks of a series of processes [39], including information retrieval, reasoning, and integration, to name a few. It is worth noting that there is a rich body of prior work [12, 14, 32, 41, 42] on developing ontological frameworks to systematically organize the conceptualization of geographic phenomena, e.g., through hierarchical structures and semantic relations. However, in such ontologies, the representations are (mostly) static. The modeling of evolving (geographic) knowledge and the representation of space and time have long been a challenge, not only within the geographic domain but also more broadly in KRR [23, 11].

In this dissertation, I include two views of concept representation, as depicted in Figure 1.1. In information science, concepts are usually represented as sets of instances that share common properties [30, 44]. The top view (Figure 1.1a) is mostly used in structured knowledge systems, where a concept is represented by its label, intension (i.e., class-defining properties), and extension (i.e., class instances) [44, 48]. Note that this view is also known as the *classical* theory of concepts, and numerous alternatives, such as prototype theory, exist. Figure 1.1b shows a natural language view, where a concept is represented by its label and the context in which it appears (e.g., surrounding text) [47, 38].



(a) Label-intension-extension structured system view.



(b) Label-context natural language view.

Figure 1.1: Structured versus unstructured views of concept representation.

To give an intuitive example, the concept of *rainforest* in the structured view can be represented as a class, identified by the label “rainforest”, defined by properties of annual rainfall, vegetation coverage, biodiversity, etc., and with instances such as the Amazon Rainforest and the Daintree Rainforest. In formal ontologies, knowledge is structured into two components, namely the TBox (terminology) and the ABox (assertions). The class-defining properties constitute the TBox, which represents schema-level knowledge that defines what a rainforest is. Specific instances like the Amazon Rainforest belong to the ABox, which contains instance-level knowledge about the list of rainforests in

1 Introduction

the world. In reality, things are often more complicated, and instance-level statements can be part of terminological definitions, such as for enumerations (e.g., seasons). In contrast, in the natural language view, the concept of *rainforest* is represented through its context, which can draw on flexible sources such as its Wikipedia entries, related scientific publications, news articles, and so forth.

Building on these two views, the inclusion of space and time in concept representation is also achieved in two different approaches. In natural language, spatial and temporal information can usually be derived directly from the context or associated metadata. For example, to study the concept of *climate change* in Asia during the 2000s, the context can be focused on news or scholarly articles that discuss the topic concerning Asian countries during that time frame. The flexibility of the natural language view provides the opportunity to explore emerging and retiring concepts, as concepts are extracted from text, instead of being predefined in an ontology. It also allows for capturing more fine-grained variations across space.

At the same time, the structured view of concepts (e.g., as formalized in ontologies) offers additional opportunities, as it requires explicit representation of space and time and is well-suited for modeling and analyzing existing knowledge. It is worth noting that, although both ontologies and knowledge graphs (KGs) represent the structured view, they serve different purposes in this dissertation. Ontologies define how concepts are formally represented and related, using standardized vocabularies, hierarchical structures (e.g., `subClassOf`, `partOf`), properties (e.g., `hasLocation`), and other logic constraints. KGs, on the other hand, act as the operational environment, where these definitions are instantiated with real-world data and the actual impacts of concept changes can be measured over space and time. For example, an ontology might define the temporal validity constraints of historical states, whereas a KG populated with real-world data represents concrete instances, such as the *Austro-Hungarian Empire*, with explicit start and end times (`AustroHungarianEmpire hasBeginning 1867`;

`AustroHungarianEmpire hasEnd 1918.`). Spatial information can similarly be represented by attaching geometries and computed topological relations (e.g., `AustroHungarianEmpire Touches Germany`) using GeoSPARQL¹.

Both structured and unstructured views of concepts are adopted throughout this dissertation. The unstructured, natural language view is used in Chapters 3 and 4 to define and detect concept drift across space and time in RDM. The structured, ontology and KG view is used in Chapter 5 to assess the consequences of TBox (schema-level) changes and examine how they propagate through a knowledge system, affecting downstream spatio-temporal reasoning.

1.3 Research Questions

Concept drift has been a topic of interest in the field of knowledge representation and reasoning (KRR) for years [48, 43, 7, 47], yet several important gaps remain. Despite extensive prior work, there is still no widely accepted definition of concept drift or its variants (e.g., split, merge) [15], which makes it difficult for the findings to be adopted or implemented across domains. Furthermore, existing work has mainly focused on the change of a concept’s meaning over *time*, leaving the role of *space* in this process under-explored. These gaps pose challenges for studying concept drift in RDM, which depends on clear conceptual categorization to ensure consistent metadata, effective data discovery, etc. To address these gaps, this dissertation takes a progressive inquiry into geographic concept drift in RDM systems: starting from the definitions, then to the detection of concept drift across both time and space, and finally to the assessment of the impact of these conceptual changes on knowledge systems.

Research Question 1: *What are the various forms of change in concepts, e.g., drift, shift, split, and merge? How can they be defined and measured in research data*

¹<https://www.ogc.org/standards/geosparql/>

1 Introduction

management (RDM) systems?

This research question focuses on establishing definitions and categorizing concept drift and its variants within RDM contexts. The aim here is to foster cross-domain consensus on this topic. Currently, there are different definitions and inconsistent terminologies (e.g., semantic drift, concept shift, and semantic shift) when describing the same phenomena across domains, including history, linguistics, cognitive science, and machine learning communities [15, 43]. Even within the same community, terminological inconsistencies remain. For example, in the Semantic Web community, Wang et al. [48] used “concept drift”, while Verkijk et al. [47] used the term “concept shift” to describe the same underlying idea. Such naming inconsistencies can hinder FAIR usage of research findings and cross-disciplinary collaboration on the topic of concept drift. Addressing this issue is therefore a key priority in this dissertation.

Wang et al., [48], among the earliest to work on this topic in the Semantic Web community, identified a couple of variants of concept drift, including shift (i.e., the meaning of one concept shifts to another) and split (i.e., the meaning of one concept splits into two or more concepts). However, other types of conceptual changes, such as merge, birth, and retirement, were not included, which are equally important for understanding the full spectrum of concept evolution. Gupta et al. [19] offer inspiration for this research question by identifying these additional forms of change; however, their work focuses more on the application and does not provide explicit definitions for these variants.

Besides inconsistent terminology and incomplete drift variants, prior work has largely overlooked concept drift in the context of RDM. When scientific concepts evolve without systematic documentation, it becomes challenging to retrieve datasets, integrate datasets from different time periods, and maintain interoperability when systems cannot recognize semantic relations between different concepts over time (and space). If such changes are not properly documented, both precision and recall in data retrieval will decline, directly lowering the reusability of data in the evolving research landscape. Given these challenges,

in this research question, I will propose definitions for different types of conceptual changes and develop methods to detect and measure them in RDM systems.

Research Question 2: *Can concept drift be observed across both space and time? Namely, does the meaning of a concept change at different rates in different regions? If so, how can spatio-temporal concept drift be detected in scientific texts?*

Having established definitions for concept drift and its variants, this research question extends the theoretical foundation of concept drift from a purely temporal aspect to include spatial and place-based dimensions. In the context of RDM, this addresses an important gap in the current understanding of how scientific concepts evolve—that they evolve not only over time but also across space. To give an intuitive example, *natural disaster* and *natural hazard* were previously used interchangeably [10]. However, when these concepts got refined meanings over the years, i.e., *natural disaster* has a focus on societal disruption causing (economic) losses, while *natural hazard* refers to the actual event [45]. The underlying hypothesis of this research question is that these changes do not happen across all research communities at the same time. Namely, the adoption rate of the refined meanings should vary across locations. Such spatio-temporal differences are overlooked when only temporal variations are considered. The lack of spatial dimensions in concept drift research becomes more problematic given the global nature of contemporary scientific research. Concepts such as *sustainable development*, *smart city*, and *climate resilience* may carry different meanings as well as operational definitions across geographic locations, based on local priorities, policies, and cultural values [16]. These meanings also shift over time. This research question, therefore, aims to capture how concept meanings evolve differently (i.e., at different rates) across geographic regions.

In addition to the theoretical focus described above, this research question also takes a technical focus. Specifically, I aim to develop methods for detecting semantic similarity across both space and time. Previous work has employed word embeddings to compute semantic similarity over time or across space [53, 54, 17], but has not considered these two

1 Introduction

dimensions jointly. Scientific texts provide an ideal data source for examining both spatial and temporal dimensions of concept drift, because they contain rich information about study areas, temporal context, and research topics, especially in the field of GIScience. With the availability of scientific publications that have both temporal and geographic coverage, I aim to systematically analyze how concepts evolve differently across regions. To sum up, in this research question, I will extend the study of concept drift to include the role of space and develop novel methods to detect spatio-temporal concept drift in scientific texts.

Research Question 3: *How can the impact of conceptual changes in structured knowledge systems, e.g., geospatial ontologies, be quantified across space, time, and semantics, and what are the implications for spatio-temporal reasoning?*

In the previous two research questions, I focused on the unstructured view of concepts. I defined different types of conceptual changes and detected concept drift over time and space, taking advantage of the flexibility and rich contextual information of natural language. The findings from these studies have shown that geographic concepts can drift across both space and time. I then hypothesize that once a concept changes its meaning, the consequences will also span both spatial and temporal dimensions, particularly on downstream spatio-temporal reasoning tasks. For instance, when the definition of *forest* changes, say, the minimum canopy coverage requirement decreases from 30% to 25%, this will have a spatial impact on the overall forest coverage, affect time series analyses of areas previously classified as forest, and essentially also complicate the calculation of deforestation rates and related policy. Therefore, in this research question, I shift the focus to the structured view of concepts. I am interested in exploring when conceptual changes occur, how they affect established knowledge systems, and how their impacts can be systematically evaluated over space, time, and semantics.

Although substantial work has focused on ontology evolution [52, 33, 34] and change impact analysis [21, 26, 13] in the broader KRR community, effectively managing and

reasoning over evolving knowledge remains an open challenge [34]. One key issue is that the complexity of real-world changes, which happen across both space and time, is often overlooked. For example, if the Food and Agriculture Organization (FAO) updates its definition of *forest*, how many countries would be affected, and how many historical datasets could no longer be used to infer new knowledge? Another issue to note is that for different types of change, previous work has treated them as the same, without differentiating between TBox (e.g., redefining a concept) and ABox (e.g., updating a specific instance) changes [9]. To address these gaps, this research question aims to develop a framework that incorporates the often missing spatial and temporal perspectives. Such a framework will provide interpretations on how changes affect knowledge systems not only in abstract semantic space but also in the physical world. Furthermore, I am interested in designing a weighting scheme to rank different types of change, helping to identify which changes carry more significance.

1.4 Structure of the Dissertation

So far, this chapter has presented the motivation for studying concept drift in RDM, with a particular focus on geographic concepts. I have then provided the background for this dissertation on the role of space and time in concept representation. In Section 1.3 I first identified the broad, overarching research gaps and then outlined more specific gaps, which shaped the formulation of three research questions to be addressed in a progressive manner.

Chapter 2 presents a synopsis of the dissertation, which serves as a cohesive overview of the included publications. It summarizes the research objectives, the methods applied, the main findings, and the overall contributions. The following three chapters (Chapter 3, 4, and 5) contain the individual publications that address the research questions and form the core of this dissertation.

1 Introduction

- Chapter [3](#) **Shi, M.**, Janowicz, K., Liu, Z., Karimi, M., Majic, I. and Fortacz, A. (2025), Defining Concept Drift and Its Variants in Research Data Management: A Scientometric Case Study on Geographic Information Science. *Transactions in GIS*, 29: e70058. <https://doi.org/10.1111/tgis.70058>
- Chapter [4](#) **Shi, M.**, Janowicz, K., Liu, Z., Karimi, M., Majic, I., and Fortacz, A. What, When, and Where Do You Mean? Detecting Spatio-Temporal Concept Drift in Scientific Texts. In 13th International Conference on Geographic Information Science (GIScience 2025). *Leibniz International Proceedings in Informatics (LIPIcs)*, Volume 346, pp. 16:1-16:18, Schloss Dagstuhl – Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/LIPIcs.GIScience.2025.16>
- Chapter [5](#) **Shi, M.**, Janowicz, K., Liu, Z., Karimi, M., Majic, I., and Fortacz, A. (In revision). The Geography of Evolving Concepts: Quantifying Spatial, Temporal, and Semantic Consequences of Ontology Changes. Submitted to the *International Journal of Digital Earth* on 13 August 2025. Revision submitted on 10 November 2025.

The rest of this dissertation is structured as follows.

Chapter [3](#) starts with an overview of concept drift research across different domains, which outlines how different fields define and interpret its key aspects, and how these relate to RDM. This is followed by a review of structured and unstructured approaches to studying concept drift, more specifically, on KG-based and natural language-based methods, with a focus on their differences in concept extraction and semantic similarity measurement. Definitions of concept drift and its related variants, including birth, shift, split, merge, and retirement, are then proposed. Using a scientometric dataset of GIScience publications as a case study, this chapter presents representative examples of each type of conceptual change, using paper keywords as signifiers of the concepts. Overall, this chapter aims to build cross-domain consensus on the terminology and definitions related

to concept drift. It then emphasizes the need to monitor conceptual changes in RDM systems to support FAIR data principles (e.g., improving precision and recall in data and literature retrieval).

Chapter 4 introduces the notion of spatio-temporal concept drift, which extends previous definitions to account for both temporal and spatial dimensions. This chapter, therefore, has a specific focus on the role of space and time, and how they are represented in concepts through word embeddings from a technical perspective. A novel approach for detecting spatio-temporal concept drift in scientific texts is proposed, which jointly considers both dimensions. Using the same scientometric dataset as Chapter 3, this chapter focuses on the subset from which spatial and temporal information can be extracted for the case study. The results show varying semantic trajectories and drift patterns of representative GIScience concepts (e.g., *urban planning*, *spatial analysis*, and *cellular automaton*). This chapter can be seen as an extension of Chapter 3. It aims to highlight the importance of jointly considering space and time in concept drift detection. This is particularly important for geographic concepts in FAIR-based RDM systems to support the development of more spatio-temporally aware knowledge systems.

Chapter 5 starts with background on the dynamics of geographic concepts and their representation in geospatial ontologies. This is followed by a review of existing methods on change impact analysis in the broader domain of KRR. Based on the gaps identified, a multi-dimensional framework is presented to examine the consequences of conceptual changes in structured knowledge systems across spatial, temporal, and semantic dimensions. A weighting scheme is also proposed to rank relation types by their structural significance. The application of this framework is then demonstrated using a geospatial ontology, which includes details on the construction of a maritime legal ontology, from its design to the population with real-world data in the exemplary knowledge graph. The results show that by introducing a single additional requirement to a concept, the impacts can propagate across spatial, temporal, and semantic dimensions. This chapter aims to support better

1 Introduction

decision-making when handling conceptual changes in structured knowledge systems. More specifically, it helps RDM practitioners identify the changes that require the most attention and determine the most appropriate strategies to manage knowledge systems.

Finally, Chapter 6 concludes this dissertation with a summary of the previous chapters and their contributions to the overarching research goal. I present these contributions from both theoretical and practical perspectives. Following this, I discuss the limitations of this work and suggest potential directions for future research.

2 Dissertation Synopsis

This dissertation focuses on geographic concept drift in research data management (RDM). More concretely, I examine the role of space and time in concept representation, how they shape concept meanings, and how they affect knowledge organization in RDM systems. Concept drift refers to the phenomenon that the meaning of a concept changes over time [48]. It poses challenges for RDM in data integration, reusability, and broader knowledge management of scientific research, because changes in meaning underlying the usage of keywords (e.g., that are used to describe datasets) can lead to inconsistencies in RDM systems. The overarching research question of this dissertation can be summarized as:

How to develop methods for handling conceptual changes in knowledge representation and how to support the management of evolving knowledge over time and space in RDM systems?

This dissertation is motivated by the need to uphold FAIR data principles [49], which ensure that scientific concepts are consistently interpreted and properly linked across datasets and literature when their meanings evolve over both time and space.

2.1 Research Design

The core idea underlying this dissertation is that concepts are dynamic [2]; they evolve with cultural, social, political, or scientific processes rather than being fixed representations [39, 24]. In other words, this dissertation treats the meaning of a concept as context-dependent.

These contexts keep evolving, while the keywords are frozen in time—so to speak—once an article or dataset has been published.

The research design of this dissertation follows a progressive structure, moving from definition to detection and finally to evaluation. It starts by proposing and establishing definitions for concept drift and its related variants, as summarized in Table 2.1. These definitions (concept *birth*, *drift*, *shift*, *split*, *merge*, and *retirement*) show the different forms of conceptual changes and are grounded in a case study using a scientometric dataset covering 30 years of GIScience publications [51]. Building on these definitions, I apply embedding-based methods to capture changes in the semantic similarity of concepts in the field of GIScience. Next, I map them to the different types of conceptual changes, demonstrating their real-world applications for RDM in Chapter 3. This approach is further extended in Chapter 4, where I incorporate both temporal and spatial dimensions (as unique combinations) into concept representation using embeddings. I apply a similar semantic similarity measurement of GIScience concepts as in Chapter 3, but this time examine changes both over time and across geographic space.

After characterizing and identifying conceptual changes, I shift the focus to their implications for structured knowledge systems. I propose a framework to evaluate the consequences of conceptual changes in Chapter 5 i.e., how changes propagate through spatial, temporal, and semantic dimensions within knowledge graphs, as summarized in Table 2.2. This evaluation emphasizes the practical implications of concept drift for RDM, that even minor changes in concept meaning can affect data quality, long-term inference, and even interoperability. By combining the definition, detection, and evaluation of conceptual changes, the research design offers a coherent path to address the challenges such changes pose for RDM, from theory and methods to the practical impact of concept drift on structured knowledge systems, e.g., knowledge graphs.

Table 2.1: Summary of proposed conceptual changes definitions, adapted from Chapter 3 [39] for presentation in this dissertation.

Change Type	Definition & Description
Concept Birth	$concept_birth_t(C_i) \iff C_i \in \mathbb{C}_t \wedge C_i \notin \mathbb{C}_{t-n} \wedge \max_{C_j \in \mathbb{C}_{t-n}} sim(C_i^t, C_j^{t-n}) < \theta_1$ where concept exists at time t, but does not exist at prior time steps $t - n$, and no sufficiently similar concept exists previously (similarity $< \theta_1$).
Concept Drift	$concept_drift_{t,t+n}(C_i) \iff \underset{m}{\operatorname{argmax}} sim(C_i^t, C_m^{t+n}) \neq i$ where concept exists at both time t and $t + n$, but the most similar concept at $t + n$ is not the concept itself. Another concept C_m^{t+n} is more similar to the original C_i^t.
Concept Shift	$concept_shift_{t,t+1}(C_i) \iff C_i \notin \mathbb{C}_{t+1} \wedge \underset{m}{\operatorname{argmax}} sim(C_i^t, C_m^{t+1}) = k$ where original concept C_i does not exist at $t + 1$, and the most similar concept at $t + 1$ is C_k.
Concept Split	$concept_split_{t,t+1}(C_i) \iff \begin{cases} sim(C_i^t, C_{i,j,k,\dots}^{t+1}) > \theta_2 \\ C_{j,k,\dots}^t \notin \mathbb{C}_t \\ C_i^{t+1} \in \{\emptyset, C_i^{t+1}\} \end{cases}$ where original concept splits into multiple concepts, and new concepts $\{C_j, C_k, \dots\}$ did not exist at time t; original concept may or may not persist, and similarity exceeds threshold θ_2.
Concept Merge	$concept_merge_{t-1,t}(\{C_i, C_j, \dots\}) \iff sim(C_{i,j,\dots}^{t-1}, C_i^t) > \theta_3$ where multiple concepts at $t - 1$ merge into single concept at t, and similarity between original concepts and merged concept exceeds θ_3.
Concept Retirement	$concept_retirement_t(C_i) \iff C_i \in \mathbb{C}_t \wedge C_i \notin \mathbb{C}_{t+n} \wedge \max_{C_j \in \mathbb{C}_{t+n}} sim(C_i^t, C_j^{t+n}) < \theta_4$ where concept exists at time t, but does not exist at later time steps $t + n$, and no sufficiently similar concept exists later (similarity $< \theta_4$).

Notes: $\mathbb{C}_t$ represents the set of all concepts at time t ; $sim(\cdot)$ represents the similarity function between concepts; $\theta_1, \theta_2, \theta_3, \theta_4$ are empirically defined similarity thresholds; $n \in \{1, 2, \dots\}$ represents the time step intervals.

Table 2.2: A multi-dimensional framework for ontology change impact, adapted from Chapter 5 for presentation in this dissertation.

Dimension	Metric & Description
Spatial Impact	Coverage Change: $\Delta S_{\text{norm}} = \frac{S_{\text{after}} - S_{\text{before}}}{A_{\text{before}}}$ measures relative spatial coverage (e.g., length, area, or volume) change.
Spatio-temporal Impact	Temporal Reasoning Validity: $\delta TRV(s) = \frac{N_{\text{affected}}(s)}{N_{\text{total}}}$ measures proportion of historical records that lose inference validity for location s .
Integrated Semantic Impact	Entailment Change: $\Delta E = E_{\text{after}} \oplus E_{\text{before}} $ counts total changes in derivable logical statements across all spatial, temporal, and semantic dimensions.
Structural Weight	Relation Weight: $w_r = \frac{ E_{\text{in}}(r) \times E_{\text{out}}(r) }{ E_{\text{total}} }$ assigns significance weights to relations based on upstream dependencies and downstream entailments.

Notes: $\oplus$ denotes the symmetric difference, capturing both statements lost and gained. More details are provided in Section 3 of Chapter 5.

2.2 Key Findings

The main findings of this dissertation can be summarized into three aspects, corresponding to Chapters 3, 4, and 5.

First, Chapter 3 establishes a shared terminology and mathematical definitions of concept drift and its related variants. Using a scientometric dataset of GIScience publications [51], I identify and exemplify different types of conceptual changes (birth, shift, split, merge, and retirement). The findings reveal that keywords used in GIScience publications have undergone these proposed changes; for example, *big data* has potentially split into *big data* and *urban big data*; *intelligent GIS*, *distributed GIS*, and *web GIS* have potentially merged into *web GIS*. This chapter lays the groundwork for a more consistent way to

document conceptual changes in both research and RDM, calling for more FAIR-aware RDM practices to account for potential concept drift over time.

Chapter 4 extends the work from Chapter 3 by introducing the notion of spatio-temporal concept drift. In this chapter, time and space are incorporated directly into concept representations using word embeddings. This approach allows us to study how concepts evolve differently along both dimensions, as illustrated with two example cases in Figure 2.1. Using a spatio-temporally scoped subset of the same scientometric dataset as in the previous chapter, the findings reveal diverse semantic trajectories of GIScience keywords, such as *urban planning* and *spatial analysis*, across different countries. These findings demonstrate that temporal trends alone do not fully capture the evolution of geographic concepts, which calls for the importance of jointly considering spatial and temporal dimensions in concept drift detection in RDM systems.




Figure 2.1: Representation of concepts drift and split over geographic space and time. The intensity of color indicates concept similarity. Figure originally published in Chapter 4 [38], included here for presentation in this dissertation.

Finally, Chapter 5 develops a framework for evaluating the consequences of conceptual changes in structured knowledge systems. This framework is applied to a simplified geospatial ontology (i.e., a maritime legal ontology), which is then populated with real-world data to form a knowledge graph (KG), including historical claims and maritime disputes. When a prominent concept changes, for example, by adding an additional restriction to the concept of *island*, the results show that approximately 6% of the existing statements in the KG need to be reassessed. Though this percentage is specific to this case study, the underlying mechanism demonstrates that conceptual modifications can propagate across spatial, temporal, and semantic dimensions, especially in interdependent geospatial knowledge. The results also show that, although spatial relations account for the most statement changes in the example KG, their impact does not carry structural significance. Because these relations depend solely on the geometry of the entities, they influence the KG quantitatively rather than qualitatively—they do not alter the ontology’s organizational structure in the way that taxonomic or mereological relations do. These findings help identify which changes are more critical and support better decision-making in the management of structured knowledge systems, such as ontologies and RDM.

The findings of this dissertation provide a progressive view, from proposing definitions and detecting conceptual changes over time, to an extension of detection across both space and time, and finally evaluating their potential impact on structured knowledge systems. Overall, the results contribute to RDM by improving the understanding of how concepts (especially geographic concepts) evolve and by calling for more FAIR-aware RDM practices.

3 Defining Concept Drift and Its Variants in Research Data Management

In this chapter, I propose definitions of concept drift and its related variants, including birth, shift, split, merge, and retirement. To demonstrate the practical application of these definitions, I use a scientometrics dataset covering 30 years of GIScience publications. For each type of conceptual change, I provide representative GIScience concepts that are indicative of it. Furthermore, I incorporate spatial and place-based perspectives as exploratory analysis to examine whether concept drift exhibits geographic patterns in addition to temporal ones alone.

Shi, M., Janowicz, K., Liu, Z., Karimi, M., Majic, I. and Fortacz, A. (2025), Defining Concept Drift and Its Variants in Research Data Management: A Scientometric Case Study on Geographic Information Science. *Transactions in GIS*, 29: e70058. <https://doi.org/10.1111/tgis.70058>

The publication is preceded by a co-author declaration specifying individual contributions, with the PhD candidate's research contribution confirmed by all co-authors' signatures.



Co-author declaration for the following joint paper

This declaration states the research contribution (e.g. research idea and – question; data compilation, manipulation and modelling; design and preparation of graphics, maps and tables; writing of text) of the candidate, the main supervisor (where he/she is an associate author) and the other authors.

If applicable, the contributions from other PhD and / or habilitation candidates who have or intend to include the paper in a thesis should be described. Contributions from master students should also be mentioned.

Research contribution details

Name of candidate:

Meilin Shi

All other authors (Please type the corresponding author in caps):

Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, Alexandra Fortacz

Title of contribution:

Defining Concept Drift and Its Variants in Research Data Management: A Scientometric Case Study on Geographic Information Science

Journal details:

Transactions in GIS

Author's contribution details

Independent contribution of candidate:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

First author's contribution:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

Contributions of other authors:

- Discussion of the research design
- Revision of the manuscript



Contributions from other colleagues and master students

Has this paper been, or will this paper be part of another doctoral degree thesis?

Yes: ☐

No: ☒

If yes, please elaborate:

Contributions from master students:

Yes: ☐

No: ☒

If yes, please elaborate details:

First author's declaration of contribution for candidate

First author's name:

Meilin Shi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable first author's signature:

21.07.2025





First **co-author's declaration of contribution for candidate**

Co-author's name:

Krzysztof Janowicz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

Second **co-author's declaration of contribution for candidate**

Co-author's name:

Zilong Liu

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:



23.07.2025



Third **co-author's declaration of contribution for candidate**

Co-author's name:

Mina Karimi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

Fourth **co-author's declaration of contribution for candidate**

Co-author's name:

Ivan Majic

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025



3 Defining Concept Drift and Its Variants in Research Data Management



Fifth **co-author's declaration of contribution for candidate**

Co-author's name:

Alexandra Fortacz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☐

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

RESEARCH ARTICLE

OPEN ACCESS

Defining Concept Drift and Its Variants in Research Data Management: A Scientometric Case Study on Geographic Information Science

Meilin Shi  | Krzysztof Janowicz | Zilong Liu  | Mina Karimi | Ivan Majic | Alexandra Fortacz

Department of Geography and Regional Research, University of Vienna, Vienna, Austria

Correspondence: Meilin Shi (meilin.shi@univie.ac.at)

Received: 3 April 2025 | Revised: 3 April 2025 | Accepted: 21 April 2025

Keywords: concept drift | knowledge graph | natural language processing | research data management | scientometrics | semantic similarity

ABSTRACT

Just as the language used to express them, concepts constantly evolve, due to changing cultural norms, scientific progress, or major events that (re)shape our understanding of the world. For example, as humans increasingly contribute to disasters, e.g., by altering the landscape, the line between *natural* and *anthropogenic* disasters blurs. This has far-reaching consequences for disaster response, policymaking, and also research. In research data management (RDM), evolving terms like *natural disaster* or *global warming* bring challenges for data retrieval, integration, interpretation, and reuse. This highlights the importance of studying concept drift, defined as the changes in a concept's meaning over time. While concept drift has been explored in numerous domains, including cognitive science, linguistics, semiotics, and ontology evolution, a clear definition within RDM is still lacking. This is significant as it may reduce recall during RDM-based data retrieval. To bridge this gap, this work proposes definitions of concept drift and related phenomena, such as shifts and merges, to foster cross-domain consensus. Using a scientometric dataset in GIScience, we apply these definitions to identify representative cases of concept drift. Our findings reveal that concepts in the GIScience domain may undergo various changes, including birth, drift, and shift, among others. These findings emphasize the need to monitor concept drift in RDM to uphold FAIR data principles, thereby improving precision and recall when searching for relevant literature or data. Lastly, we highlight the potential to measure concept drift over space (place) in addition to time alone.

1 | Introduction

Before embarking on his journey to find a shortcut to Asia, Christopher Columbus is said to have relied on calculations from the medieval Persian astronomer and geographer, Alfraganus, who defined one degree at the equator as 56.67 miles (McCormick 2012). Today, it is believed that Columbus mistook Alfraganus's 7091-ft *Arabic mile* for the shorter *Roman mile* of 4856-ft (Kelley Jr 1987). This misconception (or misdocumentation) of units may have led the expedition to substantially underestimate the Earth's circumference, thereby potentially contributing to the unexpected discovery of the Americas.

The *mile*, a unit of measurement, was defined differently across periods, regions, and cultures. This serves as an intuitive example of how differences in the mapping of a concept label to its intent can have far-reaching consequences. When Columbus applied Alfraganus's calculations made hundreds of years before, he was unaware that the definition of a mile had changed over time and across regions. Beyond units of measurement, evolving concepts—such as the redefinition of *planet* following Pluto's reclassification as *dwarf planet* (IAU 2006) or the rapid change of *gender* to include a broader spectrum of identities (Lindqvist et al. 2021)—illustrate how changes in meaning can hinder interpretation. For example, datasets annotated with outdated keywords may be difficult to find and reuse if the changes are not carefully documented.

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Transactions in GIS* published by John Wiley & Sons Ltd.

Transactions in GIS, 2025; 29:e70058
<https://doi.org/10.1111/tgis.70058>

1 of 18

3 Defining Concept Drift and Its Variants in Research Data Management

Although such dramatic consequences as Columbus's expedition or the 1999 loss of the Mars Climate Orbiter due to a unit mismatch (NASA 1999) are less likely to occur frequently, the underlying issue of varying concept definitions across space and time remains a key challenge. In research data management (RDM), this is also referred to as *semantic aging* (Schlieder 2010), where the meaning and relevance of concepts change over time. As scientific knowledge progresses, new methodologies emerge, and societal goals change, keywords used to describe research data can become outdated or misaligned, thereby impacting data interpretation. In this work, we treat keywords as representations of latent concepts. For example, keywords like *natural disaster* and *natural hazard* were once used interchangeably (Chmutina and Von Meding 2019), but today *natural disaster* more specifically refers to the societal disruption caused by a *natural hazard* (UNISDR 2018). Such changes in the meaning of keywords can make it difficult to retrieve and interpret archival data if users are not aware of these changes.

The challenges posed by semantic aging have brought the study of concept drift into the spotlight. Concept drift in RDM studies how concepts evolve within specific domains (Shi et al. 2024). These changes directly influence what we retrieve via keyword search, how we engage with others' work, and how we reuse their outputs (Pasquetto et al. 2017). The significance of the FAIR principles (Findable, Accessible, Interoperable, and Reusable) becomes evident (Wilkinson et al. 2016) in this context. For example, earlier work on human impact on Earth's climate often used the keyword *global warming*, which is now considered a more specific case of *climate change* (Schuldt et al. 2011). As a result, searching for datasets and articles on *climate change* from today's perspective may overlook earlier literature that used a different set of keywords.

Hence, monitoring and identifying concept drift in RDM is important for preserving the long-term value and impact of research data. However, despite its importance and extensive studies across various fields, there is no widely accepted definition or categorization of concept drift and its variants (Fokkens et al. 2016) that would allow for their implementation in RDM systems. This lack of consensus complicates the development of effective detection methods and limits interoperability among research domains. To address this gap, we propose definitions of concept drift and related changes, such as splits and merges, with a focus on natural language, particularly in applications related to RDM. In doing so, we highlight the difference between approaches to concept drift from a natural language and cognitive science perspective and those from a knowledge graph perspective.

The contributions of this work are as follows:

- We contrast distinct approaches to concept drift that use natural language processing (NLP) and knowledge graph (KG) representations, enhancing cross-domain interoperability and detection methods.
- We propose definitions of concept drift-related variants, including concept *birth*, *drift*, *shift*, *split*, *merge*, and *retirement*, aiming to establish a widely accepted framework across domains.
- We apply these definitions to a scientometric dataset in geographic information science (GIScience), uncovering evolving concepts within this domain. In this context, we also

discuss problems arising from selecting temporal intervals, i.e., the Modifiable Temporal Unit Problem (Cheng and Adepeju 2014).

- Finally, we introduce the notion of concept drifting over both time and space (place), an issue that has received little recognition to date.

The remainder of this article is structured as follows. Section 2 provides the background on the various notions of concept drift across different domains. It highlights how each domain interprets and addresses this phenomenon. In Section 3, we discuss in detail how our work differs from existing approaches that use KGs to study concept drift, including aspects from concept extraction to similarity measurement. Section 4 presents our proposed definitions, which are then applied to a scientometrics dataset in Section 5 to demonstrate their practical applications in the field of GIScience. Section 6 shows the results, followed by a discussion of the challenges and potential directions for future work in Section 7. Finally, we conclude our work in Section 8.

2 | Background on Concept Drift

Concept drift is an active research topic across various domains, including linguistics, cognitive science, machine learning, digital humanities, and the Semantic Web (Li et al. 2021; Meroño-Peñuela et al. 2013; Schlimmer and Granger Jr 1986; Stavropoulos et al. 2019; Wang et al. 2011) or in the broader context of studying the diversification of schemata in knowledge representation (Janowicz et al. 2022). Each field approaches concept drift from a different perspective. In this section, we focus on how these fields define and interpret various aspects of concept drift.

In linguistics, for instance, concept drift, often called *semantic shift* or *semantic change*, studies the evolution of word meanings driven by cultural and societal changes (Hamilton et al. 2016; Li et al. 2021). To give a concrete example, the term *gay*, which initially described joy or happiness, has evolved to now mostly refer to sexual orientation (Wijaya and Yeniterzi 2011). Scholars use historical corpora to track and analyze such linguistic changes, identifying underlying patterns and drivers of language evolution. Such research not only documents the ever-changing nature of languages *in the wild* but also uncovers how linguistic changes reflect broader shifts in society (e.g., cultural norms), thereby offering insights into the connection between language, culture, and society at scale.

Within the machine learning (ML) community, the term concept drift often refers to changes in the statistical properties of a target variable (e.g., a data feature) over time (Schlimmer and Granger Jr 1986; Tsymbal 2004), also known as *concept shift* (Widmer and Kubat 1996) or *dataset shift* (Moreno-Torres et al. 2012). Here, the measured changes typically arise from shifts in the underlying statistical data distributions, influenced by latent variables that cannot be measured directly (Liu et al. 2017). This presents a core challenge to the accuracy of predictive models. Unlike in linguistics, where changes are typically limited to language, concept drift in ML can result from events such as changes in the behavior of users or contributors of collection

such as social media or from technological breakthroughs (Lu et al. 2018). To address these changes, various techniques such as online learning algorithms (You and Lin 2016) and adaptive data stream mining techniques (Gama et al. 2014) can be employed to improve a model's robustness.

Within the Semantic Web and the broader knowledge representation and reasoning communities, concept drift refers to the changes in the semantics (i.e., the formal definition) of a concept over time, possibly also across locations or cultures (Wang et al. 2011). These changes often result from ontological shifts as knowledge domains evolve, changes in cultural and societal norms, technological advancements, or interdisciplinary integration (Capobianco et al. 2020). The main focus is on maintaining schema relevance and ensuring interoperability over time. Hence, predicting ontology evolution is a key application area for studying concept drift within the Semantic web community (Stavropoulos et al. 2019). This contrasts with linguistics, which focuses on the evolution of word meanings, and ML, where the emphasis is on improving the performance and robustness of predictive models over time. Approaches from the ontology modeling and linked data community (Recchia et al. 2017) are key to these efforts as they allow for updating (e.g., versioning) of evolving concepts.

Despite the substantial work across these domains, defining concept drift and related phenomena remains challenging due to multiple overlapping terms such as *semantic drift*, *concept shift*, and *concept change*. Within the Semantic Web field, various frameworks based on different definitions (Verkijk et al. 2023; Wang et al. 2011) further complicate efforts to achieve interoperability and conduct subsequent research. These terminological variations cause risks for researchers attempting to compare findings across studies or apply techniques across fields. To address these issues, this article proposes definitions of concept drift and related variants, with a focus on their role in RDM. By clarifying these definitions, we aim to improve our understanding and facilitate more effective detection and management of concept drift in scientific metadata used to organize and search within RDM systems.

3 | Approaches to Concept Drift: Knowledge Graphs Versus Natural Language Processing

In the Semantic Web literature, concept drift is addressed via two main avenues: one employing a structured, top-down methodology based on knowledge graphs (KGs), and the other utilizing bottom-up techniques from NLP that excel at unstructured data (Verkijk et al. 2023). In this work, we focus on the practical applications of KGs in the context of ontology engineering, where entities (nodes) and relations are defined and structured according to a formal ontology.

KGs enable effective querying, reasoning, and knowledge discovery, but they usually capture knowledge as static snapshots (Tasnim et al. 2019). Given the fast-evolving nature of scientific data and literature, KGs also need to incorporate updates, such as the addition or removal of instances or changes to axioms, to stay relevant (Paulheim 2017). This applies to both individual assertions (e.g., those regarding the population of a country) and

concepts. For example, the evolving concept of the *European Union* is measured by changes in its extension (the number of the member states) and its intension (our perception of it as an economic and/or political entity) (Wang et al. 2011). NLP, in contrast, deals with unstructured data, deriving meaning from context and usage patterns (Camacho-Collados and Pilehvar 2018). The differences between KGs and NLP-based representations shape their approaches to concept drift and (semantic) similarity measurements (Janowicz et al. 2011; Schwering and Raubal 2005). This section examines how these differences impact interoperability.

3.1 | Concept Extraction

Concepts are often defined as mental representations of categories (i.e., sets of entities categorized under a concept) and play a fundamental role in how humans classify objects in their surroundings (Barsalou 1993). However, fields such as the humanities and social sciences have their own requirements, and there is no universally agreed-upon definition of what precisely constitutes a *concept* (Kuukkanen 2008). This lack of consensus complicates the study of concept drift, especially as KGs and NLP employ distinct methods for concept extraction, representation, and alignment, influenced by the nature of the data they handle.

3.1.1 | Concepts in Knowledge Graphs

In KGs, concepts serve as the building blocks for vocabularies and taxonomies that facilitate information retrieval, reasoning, integration, and interoperability. The distinction between a concept and an object in KGs is nicely illustrated by Frege's paradox in *On Concept and Object*: "The concept horse is not a concept" (Frege et al. 1951). Frege's provocative example shows that while "horse" as a concept represents the general notion (e.g., a blueprint) of horses, the phrase "the concept horse" is treated as a specific instance (of the domain of concepts). This paradox highlights the complex nature of concepts in ontologies, separating conceptual frameworks (TBox) from individual instances (ABox). In practice, concepts are often represented as classes containing instances that share common properties (Lv et al. 2018; Stock 2010). For example, as illustrated in Figure 1, a concept would be *hurricane* and an instance would be *Hurricane Katrina*. Note, however, that this classical distinction into the TBox as a set of axioms forming the concepts and their relationships, and the ABox as a set of assertions formed using these axioms, is rather a theoretical guiding principle than actual knowledge engineering practice. For instance, *hurricane* may have also been modeled as an instance of a *disaster* class. Even more, classes can be singleton classes, purely defined by enumeration, or be regarded as both a class or instance (e.g., in punning) (Hitzler et al. 2009).

Concepts are usually understood through their *intension* (class-forming properties) and *extension* (class instances) (Wang et al. 2011). Even though concepts are often regarded as static, they are inherently dynamic. They evolve with scientific and technological advancements, as well as through communication,




FIGURE 1 | An example of concepts and instances in a KG.

where two parties (implicitly) agree on the interpretations of terms (Barsalou 1993; Raubal 2008; Scheider and Kuhn 2015; Stock 2010). That is to say, changes can occur at both the TBox and ABox levels. This includes, for example, the redefinition of *planet* to include specific orbital characteristics, leading to the exclusion of Pluto, and the evolving understanding of COVID-19, where new variants result in the continuous addition of instances to the ABox.

Building on this understanding of concepts in KGs (Hogan et al. 2021; Ji et al. 2021; Kuhn et al. 2014), concept extraction is the process of identifying and delineating entities and their relations in a dataset. When constructing graphs from sources such as Wikipedia or WordNet (Miller 1995), relations like *instanceOf*, *subClassOf*, or *superClassOf* are used to categorize entities (Fabian et al. 2007; Lehmann et al. 2015). In order to identify clusters (or patterns) indicative of conceptual groupings, this construction process also involves methods that analyze a graph's structure and the underlying semantics of its classes and relations (Navigli and Velardi 2004). Conversely, extracting concepts from established KGs typically follows the reverse process, leveraging these hierarchical relations based on predefined ontologies to identify the underlying conceptual structure (Lv et al. 2018; Mai et al. 2024).

3.1.2 | Concepts in Natural Language

In natural language, there is less rigorous structure in determining what constitutes a concept from unstructured text, unlike the more explicitly defined intension and extension of formal concepts in KGs. Therefore, NLP allows for a wider interpretation of what qualifies as a concept, as seen in prior research on concept drift. Examples range from the concept of apple pie derived from recipes in newspaper articles (Tietz et al. 2020) to categories presented in remote sensing classification results (Gupta and Gahegan 2020). All these varying interpretations have made concept extraction and drift measurement in natural language more challenging, as they are more dynamic and context dependent. However, this diverse interpretation also provides the opportunity to include a broader and more flexible range of concepts, rather than being limited to the predefined, often static, ontologies in KGs.

Recognizing the challenges and opportunities of working with natural language, we use existing concepts in this work, instead of extracting them from unstructured text, following a similar approach as proposed by Gupta and Gahegan (2020). In our study, these concepts are signified and represented by keywords associated with academic papers published in conference proceedings or journals. In theirs, the concepts correspond to categories used in remote sensing classifications. For instance, land cover classifications have evolved to accommodate emerging topics of interest, such as forest regeneration, invasive species, fire risk, soil erosion, and water quality. These evolving categories reflect changes in research priorities and potentially also in the environment. In both cases, these concepts, if drifted, indicate a measurable change in our understanding within a particular domain.

3.2 | Similarity Measurement

After identifying and/or extracting concepts, the next step in assessing concept drift is to measure how they evolve over time. Drift measurement can be modeled as a special case of semantic similarity measurement (Lu et al. 2018). This section reviews different methods for computing similarity in the contexts of KGs and natural language. We also clarify the distinction between similarity and relatedness.

3.2.1 | Similarity Measurement in Knowledge Graphs

In KGs, similarity is often determined by comparing entities (nodes) based on their structure, proximity, and shared relations in the graph. This process usually makes use of ontological hierarchies and graph embeddings (Bordes et al. 2013; Jiang and Conrath 1997). Most KG embedding methods, however, tend to overlook the distinctions between concepts and their instances by treating both as equivalent (Hao et al. 2019; Verkijk et al. 2023). Maintaining these hierarchical relations and differentiating between concepts and instances during the embedding process is crucial when measuring concept drift. Several embedding models have been adapted to explicitly address this issue, as summarized in Table 1. Once concepts are converted into vectors via embeddings, we can measure their similarity

TABLE 1 | Summary of KG embedding models for concept-instance differentiation.

Model	Representation	Differentiation approach
TransC (Lv et al. 2018)	Concepts: spheres; Instances: vectors	Hierarchical modeling using spheres and vectors in the same space
JOIE (Hao et al. 2019)	Concepts and instances as vectors in separate embedding spaces	Cross-view association model and hierarchy-aware encoding for ontologies
JECI (Zhou et al. 2020); JECI++ (Wang and Zhou 2021)	Concepts: hierarchical tree; Instances: context vector	Joint embedding in hierarchical tree and context space
CIST (Zhang et al. 2022)	Concepts and instances in different embedding spaces	Spatial transformation with learnable hierarchical range adjustment
Concept2Box (Huang et al. 2023)	Concepts: boxes; Instances: vectors	Dual geometric embeddings with boxes and point vectors

with cosine similarity by calculating the angle between them (Turney and Pantel 2010). For models that represent concepts using box or sphere embeddings, similarity measurement often includes calculating their overlap between shapes, the distance between their centers, or other geometric relations (Cai et al. 2021; Chheda et al. 2021).

3.2.2 | Similarity Measurement in Natural Language

In natural language, where structured information is often missing, similarity measurement relies on statistical and contextual analysis of text, i.e., by situating a word in its co-occurring neighbors (Bordag 2008). Over the past decade, emerging techniques such as word embeddings and large language models have been used to capture semantic similarity effectively (Devlin et al. 2019; Mikolov et al. 2013). Although the methods discussed below serve broader semantic evaluation purposes, here we focus on applying them specifically to our extracted concepts.

Key developments in semantic similarity measurement in NLP began with the early adoption of latent Dirichlet allocation (LDA), which established the foundation for modern statistical text analysis. LDA (Blei et al. 2003) measures semantic similarity by examining word co-occurrence patterns. It groups texts into topics, which enables the identification of documents with similar themes and related content. Word2Vec (Mikolov et al. 2013) later shifted the paradigm by converting words into dense vectors. Semantically similar words (e.g., “king” and “queen”) are positioned near one another within the vector space, providing a more intuitive understanding of word relations based on their contextual usage patterns. Doc2Vec (Le and Mikolov 2014) expanded on this idea by representing entire documents as vectors, which allows for semantic similarity comparisons at the level of documents. More recent advancements in NLP embraced so-called transformer models such as BERT (Devlin et al. 2019) and its many specialized variants (Beltagy et al. 2019; Reimers and Gurevych 2019; Yang et al. 2019), which have significantly improved the ability to dynamically capture contextual differences in text. Once vector representations or embeddings of concepts are generated using these methods, subsequent analysis is similar to that used with KG embeddings. It often involves cosine

similarity measures to evaluate the proximity of concepts in their vector space, thus resulting in similarity scores.

3.2.3 | Similarity Versus Relatedness

It is important to distinguish between computing (or detecting) similarity and relatedness when assessing concept drift via similarity measures, because they capture different aspects of concept interactions (Budanitsky and Hirst 2006). Similarity is often said to focus on shared attributes, whereas relatedness captures broader, contextual connections including meronymy (e.g., car and wheel) or frequent associations (e.g., dog and leash) (Pirró and Euzenat 2010).¹

When comparing KGs and NLP-based approaches, methods from computational linguistics are often more likely to capture relatedness than similarity. KG-specific methods, on the other hand, tend to stress similarity over relatedness (Zhu and Iglesias 2018). As structured representations of knowledge, KGs excel at identifying similarities within hierarchical contexts. For example, as shown in Figure 1, *hurricane* and *tropical cyclone* are subclasses of *storm*. This makes them similar in the broader taxonomy of weather events. KGs are designed to capture these and other explicit relations (e.g., partonomies, topological relations, and causal relations). Therefore, they are ideal for tasks that require a clear understanding of ontological structures and well-defined categories (Ji et al. 2021). NLP models, especially those built on large corpora, inherently capture relatedness by analyzing the co-occurrence of words or phrases across contexts. These include word embedding methods (e.g., Word2Vec (Mikolov et al. 2013), GloVe (Pennington et al. 2014)) and deep learning-based contextual language models (e.g., BERT (Devlin et al. 2019), GPT (Brown et al. 2020)). Unlike KGs, which rely on *predefined* structures, NLP models *learn* relationships from vast amounts of unstructured text. They can, as a result, identify latent connections that traditional similarity metrics, such as path-based methods, might overlook. For example, NLP models naturally capture the relatedness between *sports bar* and *drink* based on how often they appear together in different contexts, even though they do not share common attributes in the traditional sense of similarity (Majic et al. 2017).

3 Defining Concept Drift and Its Variants in Research Data Management

Given these differences, our case study below adopts a natural language approach to address concept drift. We use *existing concepts* represented by keywords from a corpus of academic papers. This approach allows us to focus on established concepts rather than extracting new concepts from scratch, thus avoiding the potential challenge of measuring relatedness rather than similarity.

4 | Definitions of Concept Drift and Its Variants

Building on earlier works (Verkijk et al. 2023; Wang et al. 2011) that defined concept drift and concept shift, our previous research (Shi et al. 2024) introduced the initial definitions of concept drift and its variants. This work refines these definitions and introduces two additional variants—concept birth and retirement—to account for the broader variations observed in real-world datasets. While concept birth and retirement have been previously introduced in the context of conceptual-level changes (Gupta and Gahegan 2020), they were not explicitly defined.

Intuitively, a concept has drifted at some point in the future if its self-similarity is less than 1 when compared to its initial interpretation, i.e., $\text{sim}(C_i, C_{i+n}) < 1$. However, measuring similarity between concepts is notoriously difficult without a reference point (Goodman 1972; Janowicz et al. 2011). Hence, we model phenomena such as drift, split, and merge by comparing the focal concept to the set of all other concepts existing at some point in time, denoted as C_t . In this section, we define six variants—birth, drift, shift, split, merge, and retirement—which collectively describe the life cycle of concepts in dynamic contexts. For each definition, we include an intuitive example to aid understanding. Note that these definitions will rely heavily on the aforementioned notion of similarity and thus application-specific thresholds will be introduced for the use cases later on.

4.1 | Concept Birth

Definition 1. A concept C_i is said to be born at time t if and only if it exists at time t , does not exist in any prior time steps $t - n$ ($n \geq 1$), and no concept C_j from the set of all concepts C_{t-n} is sufficiently similar to C_i at time t , according to a predefined threshold θ_1 . Expressed mathematically:

$$\text{concept_birth}_t(C_i) \Leftrightarrow C_i \in C_t \wedge C_i \notin C_{t-n} \wedge \max_{C_j \in C_{t-n}} \text{sim}(C_i, C_j^{t-n}) < \theta_1, n \in \{1, 2, \dots\} \quad (1)$$

Example. The concept of *anthropocene*, introduced in scientific discourse in the early 2000s (Crutzen 2006), can be seen as a newly born concept at that time. It describes the period of time when human activity becomes the dominant force driving geophysical and biochemical changes on Earth. There was no prior counterpart that specifically captured this perspective before its introduction, apart from broader discussions on human impact on the environment.

4.2 | Concept Drift

Definition 2. A concept C_i present at both time t and $t + n$ ($n \geq 1$), is said to undergo concept drift from time t to $t + n$ if and only if the most similar concept to C_i^t among all concepts in C_{t+n} is not C_i^{t+n} itself. In other words, there exists a concept C_m^{t+n} (where m is an index for the concepts in C_{t+n}) such that the similarity between C_i^t and C_m^{t+n} is higher than that between C_i^t and C_i^{t+n} . Expressed mathematically:

$$\text{concept_drift}_{t+n}(C_i) \Leftrightarrow \arg\max_m \text{sim}(C_i^t, C_m^{t+n}) \neq i, n \in \{1, 2, \dots\} \quad (2)$$

Example. The concept of *disasters* has drifted from being defined purely as natural events that affect humans, such as earthquakes or hurricanes, to a more inclusive view that includes human-induced causes (e.g., climate change driven by human activities) and socio-environmental factors like poverty, inadequate infrastructure, and governance failures (Perry 2018).

4.3 | Concept Shift

Definition 3. A concept C_i is said to undergo concept shift to C_k from time t to $t + 1$ if and only if C_i does not exist at time $t + 1$, and the most similar concept to C_i at time t among all concepts in C_{t+1} is C_k . Namely, k is the index of the concept in C_{t+1} with the maximum similarity value. Expressed mathematically:

$$\text{concept_shift}_{t+1}(C_i) \Leftrightarrow C_i \notin C_{t+1} \wedge \arg\max_m \text{sim}(C_i^t, C_m^{t+1}) = k \quad (3)$$

Example. The concept of *Third World countries* has shifted since the end of the Cold War and is now rarely used by the United Nations or international media. While some aspects of its classification persist, it has shifted to *developing countries* or the *Global South*, thereby highlighting different aspects, e.g., economic conditions (Solarz 2012).

4.4 | Concept Split

Definition 4. A concept C_i is said to undergo concept split into two or more concepts $\{C_j, C_k, \dots\}$ from time t to $t + 1$ if and only if the similarity between C_i at time t and each of the concepts $\{C_j^{t+1}, C_k^{t+1}, C_l^{t+1}, \dots\}$ at time $t + 1$ exceeds a predefined threshold θ_2 . Additionally, there is no prior instance of $\{C_j, C_k, \dots\}$ at time t and C_i may or may not exist at time $t + 1$ (see Figure 2 for illustration). Expressed mathematically:

$$\text{concept_split}_{t+1}(C_i) \Leftrightarrow \begin{cases} \text{sim}(C_i^t, C_{j,k,\dots}^{t+1}) > \theta_2, \\ C_{j,k,\dots}^t \notin C_t, \\ C_i^{t+1} \in \{\emptyset, C_i^{t+1}\} \end{cases} \quad (4)$$




FIGURE 2 | Illustration of different concept split patterns: (a) branch split, where C_i continues to exist at $t+1$ after splitting and (b) fork split, where C_i does not exist at time $t+1$, and splits into new concepts C_j and C_k .

Example. The concept of *artificial intelligence (AI)* has split into sub-concepts such as *machine learning*, and, more recently, *deep learning*. This represents a branch split from the broader concept of *AI* into more specialized directions, while the use of *AI* remains prevalent across domains.

4.5 | Concept Merge

Definition 5. Two or more concepts $\{C_i, C_j, \dots\}$ are said to undergo concept merge into concept C_t from time $t-1$ to t if and only if the similarity between each of the concepts $\{C_i^{t-1}, C_j^{t-1}, \dots\}$ at time $t-1$ and C_t at time t exceeds a predefined threshold θ_3 . Expressed mathematically:

$$\text{concept_merge}_{t-1,t}(\{C_i, C_j, \dots\}) \Leftrightarrow \text{sim}(C_{i,j,\dots}^{t-1}, C_t^t) > \theta_3 \quad (5)$$

Example. The concepts of *spatial analysis*, *spatial modeling*, *spatial statistics*, and others have gradually become collectively referred to as *spatial data science*. Even if *spatial data science* existed at some time before, it could still absorb or merge with the other related concepts over time.

4.6 | Concept Retirement

Definition 6. A concept C_i is said to be retired at time t if and only if it exists at time t , does not exist at any later time steps $t+n$ ($n \geq 1$), and no concept C_j from the set of all concepts C_{t+n} is sufficiently similar to C_i at time t , according to a predefined threshold θ_4 . Expressed mathematically:

$$\text{concept_retirement}_t(C_i) \Leftrightarrow C_i \in C_t \wedge C_i \notin C_{t+n} \wedge \max_{C_j \in C_{t+n}} \text{sim}(C_i^t, C_j^{t+n}) < \theta_4, n \in \{1, 2, \dots\} \quad (6)$$

Example. The concept of *flat Earth* has been retired from scientific discourse since empirical evidence—such as observations by ancient Greek philosophers like Aristotle, later circumnavigation voyages, and advancements in astronomy and mathematics—conclusively proved Earth's spherical shape.

Note on Definitions

1. The identity of a concept depends on how it is operationalized. In top-down knowledge engineering approaches,

concepts (classes) are declared and formally specified via necessary and sufficient conditions and, in theory, identity criteria. They may also be specified through axioms in knowledge representation languages, such as `owl:equivalentClass` or `owl:sameAs`. In bottom-up approaches, concepts (or categories) emerge from data patterns without predefined identity and often lack formal definitions. To some degree, this gives rise to the importance of considering both intension and extension when discussing concepts in the wild.

2. The similarity function $\text{sim}(\cdot)$ in the definitions is not limited to any specific measurement method.
3. The thresholds $\theta_{1,2,3,4}$ determine what is considered as sufficiently similar across different types of concept evolution. Setting them ensures that each concept evolution is based on a *meaningful* degree of similarity. Their values are determined based on empirical evidence, domain-specific requirements, parameter tuning, and so forth.

5 | Case Study in Geographical Information Science

In this section, we describe a case study to show how the definitions we proposed above can be applied in practice. We adopted a scientometric dataset from Wu et al. (2023), which covers research themes in geographical information science (GIScience) from 1991 to 2020 (Wu et al. 2022). Scientometrics studies the quantitative aspects of scholarly literature, such as research trends (Hood and Wilson 2001). This dataset serves as a great example of how documenting concept drift can improve keyword-based RDM retrieval.

5.1 | Keywords as Indicators of Concepts

In this case study, we approach concept drift from a bottom-up NLP perspective, aligning with the NLP view on concepts that meaning is reflected through the use in context. Since conceptual changes are abstract and difficult to measure directly, we use keywords as observable indicators that signal concept drift. Each keyword is identified by its label and serves as the symbolic representation of an underlying concept. The triangle of reference (Figure 3) can be used to illustrate the connection between a concept, its symbol (keyword), and its referent in reality. By monitoring keyword use over time in RDM, we infer how conceptual understandings evolve within GIScience.

3 Defining Concept Drift and Its Variants in Research Data Management

From this dataset, we excluded the journal *Geographical Analysis* due to its limited data volume (only five papers with keywords). Table 2 provides the summary statistics. For readability, we list the full conference and journal names in Appendix A (see Table A1).

All used keywords underwent a series of preprocessing steps: (1) conversion to lowercase, (2) lemmatization, (3) standardization to American spelling, and (4) expansion of abbreviations (e.g., *VGI* to *volunteered geographic information*). For the abbreviation of *GIS*, which can refer to *geographic information system* or *geographic information science*, we retained it as a separate keyword to study its drift, which will be discussed in more detail in Section 6. Additionally, synonymous keywords, e.g., *spatio-temporal* and *spatial temporal*, were merged to ensure consistency and improve sample size. Figure 4 shows




FIGURE 3 | The triangle of reference (Ogden and Richards 1923), illustrating the difference between concepts, (key)words, and referents.

TABLE 2 | Summary statistics of selected conferences and journals.

Type	Name	Time range	Number of papers	Number of keywords
Conference	COSIT	1993–2019	187	900
Conference	GIScience	2006–2020	105	454
Journal	CEUS	1999–2020	1131	5641
Journal	CaGIS	1991–2020	606	2966
Journal	EPB	1998–2020	559	2734
Journal	Geol	1997–2020	516	2410
Journal	IJGIS	1998–2020	1465	6887
Journal	JGS	1996–2020	462	2107
Journal	JOSIS	2010–2020	99	638
Journal	SCC	2003–2020	230	1053
Journal	TGIS	2007–2020	117	539
Total		1991–2020	5477	26,329

the distribution of the top 10 keywords across all journals and conferences.

5.2 | Temporal Interval for Drift Measurement

To detect concept drift over a reasonable temporal range, we selected six time points from the dataset, spaced 5 years apart. Table 3 provides detailed statistics for this subset. The rationale behind this choice is that evaluating concept evolution on a yearly basis would rather yield insignificant results, while using a 10-year interval would be too sparse given the size of our dataset. It is worth noting that the choice of the interval (e.g., 5 years vs. 10 years) can significantly affect the results. This will be discussed in detail in Section 7.1.

By selecting the six specific years, we focused only on the papers from these years for similarity measurement, leaving out those from the years in between. However, when counting keyword occurrences (e.g., for concept birth and retirement), we considered the entire dataset from 1991 to 2020.

5.3 | Keyword Similarity via Embeddings

Following the overview in Section 3.2.2, we decided to use SciBERT (Beltagy et al. 2019) for keyword similarity measurement. SciBERT is a pre-trained language model specifically designed to compute embeddings for scientific texts. While this work uses NLP techniques to approach concept drift, the definitions in Section 4 are intended to be applicable to various similarity measurement techniques. Here, we computed two types of keyword similarity: (1) *textual similarity*, calculated as the cosine similarity between the embeddings of two keywords, and (2) *contextual similarity*, measured by the cosine similarity between embeddings generated from the paper titles and abstracts associated with the keywords. Textual similarity captures direct




FIGURE 4 | Distribution and frequency of top 10 keywords by year.

TABLE 3 | Number of papers and keywords by selected year.

Year	Number of papers	Number of keywords
1995	15	66
2000	91	430
2005	144	682
2010	232	1097
2015	370	1823
2020	419	2058
Total	1271	6156

keyword relationships via label variations. Contextual similarity takes into account how keywords are used in different contexts. Identical keywords within the same year were averaged into a single embedding. We assigned weights of $\alpha = 0.2$ and $\beta = 0.8$ to balance textual and contextual similarities. These weights indicate that while both components are important, contextual clues provide more information. The overall similarity between keyword k_i and k_j at time t can be expressed as a convex combination:

$$\text{sim}(k_i^t, k_j^t) = \alpha \cdot \text{text_sim}(k_i^t, k_j^t) + \beta \cdot \text{context_sim}(k_i^t, k_j^t) \quad (7)$$

To investigate how the results vary with different weightings of textual and contextual components, we performed a sensitivity

analysis using different choices of α and β . Appendix B provides detailed comparisons.

6 | Representative Examples in Geographic Information Science

Building on this case study, we present representative examples to show how concepts in GIScience have evolved according to our definitions. It is worth noting that the thresholds applied in this case study are specific to the analysis within GIScience. However, our proposed definitions can be adapted to study concept evolution in other domains by adjusting these thresholds.

To establish context, we begin by examining the keyword *GIS*. During data pre-processing, we retained the abbreviation to observe how its meaning evolved over time, specifically whether it referred to *geographic information system* or *geographic information science*. Figure 5 presents the three most similar neighbors of *GIS* among these two target keywords over the six selected years.

Initially, in 1995, *GIS* shows similarities to *geographic information system* from the same year and to *geographic information science* in 2000 and 2010. From 2000 onward, it becomes increasingly similar to *geographic information system*, while still maintaining some similarity to *geographic information science* in 2010. This is probably influenced by the rise of relevant software such as ArcGIS and QGIS. Starting from 2005, which marks a transition in its meaning, *GIS* becomes very stable, being most similar only to *geographic information system*. These trends reflect broader themes in the evolution of

3 Defining Concept Drift and Its Variants in Research Data Management




FIGURE 5 | Evolution of GIS meaning: geographic information system vs. science. The network shows the three most similar neighbors of GIS among the keywords of interest (geographic information system and geographic information science) over the six selected years. The edge labels indicate similarity scores derived from Equation (7).

GIS over the past two decades, as also discussed by Egenhofer et al. (2016).

Following this, we present examples and group the results by concept birth and retirement, concept drift and shift, and concept split and merge for better comparison. The similarity scores for these examples are calculated using Equation (7) unless otherwise specified.

6.1 | Concept Birth and Retirement

To identify concept birth, we selected keywords from 2020 as the focal point. Following our definition, we first checked the occurrences of these keywords. Since we used intervals of 5 years, we filtered for keywords that did not appear in or before the previous time step of 2015 within the entire dataset. For concept retirement, we combined keywords from 1995 and 2000 because the number of keywords in 1995 alone was limited. We then identified keywords that did not appear in or after the subsequent time step of 2005. After applying these filters, we calculated the maximum similarity between the selected keywords and those from other years.

Figure 6 reports candidate keywords that may indicate concept birth or retirement, sorted by ascending maximum similarity to keywords from other time steps. To reduce noise, we only included keywords that appear at least three times in the entire dataset. From the figure, we can notice that these candidates all exhibit relatively high similarity to other years, which shows the complexity of determining concept birth or

retirement. Therefore, we only present these candidates as potential cases and leave the final determination to domain experts.

6.2 | Concept Drift and Shift

Figure 7 depicts the evolution of the keyword *agent based model* in 2005 and 2010. For simplicity in illustration, the figure only includes the three most similar keywords and/or the keyword itself. In the left trajectory, the keyword from 2005 indicates a potential concept drift in 2010 and 2015, drifting toward related keywords such as *spatial temporal data model* and *multi agent system*. During this period, its similarity to the original keyword in 2005 decreases, implying that the meaning of this keyword is changing. In contrast, the right trajectory shows that the keyword in 2010 remains most similar to itself in 2015 and 2020. This indicates that the concept represented by the keyword has stabilized during this period, showing a more consistent understanding of agent-based model research within the GIScience field.

Figure 8 presents a potential example of concept shift, as indicated by the keyword *activity based model* in 2015. Unlike the previous drift case, this keyword no longer appears after 2015 (with temporal coverage from 2001 to 2015). According to our definition, we say it shifts to the most similar keyword in the following time step. Note that shift is different from retirement, where we examine multiple time steps rather than just two consecutive ones. In 2020, the keyword shifts to *agent based model*, which may be seen as a replacement for the original keyword, reflecting a potential broader transformation of the concept over time.




FIGURE 6 | Keywords indicating potential concept birth and retirement. The first value in parentheses is the maximum similarity of the keyword compared to all other time steps; the second value is the total occurrences across the entire dataset.




FIGURE 7 | Evolution of *agent based model* in 2005 and 2010. Similarity values in parentheses are computed between the target keyword from 2005 and 2010, and those from other years.




FIGURE 8 | Evolution of *activity based model* in 2015. Similarity values in parentheses are computed between the target keyword from 2015 and those from 2020.

3 Defining Concept Drift and Its Variants in Research Data Management

6.3 | Concept Split and Merge

To measure concept split and merge, we chose a threshold value of 90% of the *maximum* similarity score. We chose a relative threshold, rather than using absolute values, to create an adaptive framework that adjusts to the characteristics of our dataset. Building upon this, our implementation includes an additional assumption that splits and merges would occur between keywords with closely related labels. Specifically, we used an additional 90% of the *maximum* textual similarity between the keywords. This refinement we added during the implementation aims to address limitations observed in previous work (Wang et al. 2011), where the authors identified examples of concept splits, e.g., “free childcare” splitting into “free childcare” and “companies, corporations, business.” While such connections might be detected through certain similarity measurements, they lack semantic coherence from a domain knowledge perspective.

Figure 9a illustrates an example that indicates a potential concept split, following the branch split pattern (see Figure 2). The keyword *big data* splits into *big data* and *urban big data* between 2015 and 2020, with *urban big data* first appearing in 2020. Figure 9b provides an example that indicates a potential concept merge, where multiple keywords, including *intelligent gis*, *distributed gis*, and *web gis*, merge into the unified keyword *web gis* by 2010. These examples showcase the possible restructuring of concepts by demonstrating how related keywords can either diverge into more specialized directions or converge into broader categories over time.

In summary, the examples presented in this section demonstrate the applicability of our definitions and uncover potential concept evolution patterns in the field of GIScience over recent decades. In line with FAIR, researchers need to stay informed and continuously monitor these changes to ensure that archival knowledge remains accessible via keyword search. For example, restricting a search to *urban big data* in 2020 or later would miss earlier work in which the urban component was not separated from general *big data*. Similarly, if searching for *web gis* prior to 2010 literature, one should include the other two keywords shown in Figure 9b. To bridge the gap between past and present research, retrospective works such as those by Egenhofer et al. (2016) and Wu et al. (2023) can help document concept drift. Future retrospective analyses that incorporate the

variations of concept evolution based on our definitions could further strengthen interoperability within the field.

7 | Challenges and Future Directions

This section discusses the sampling challenges posed by the Modifiable Temporal Unit Problem (MTUP) in our case study and points out potential directions for future research on concept drift occurring at varying rates across space.

7.1 | Modifiable Temporal Unit Problem

The MTUP (Cheng and Adepeju 2014) refers to the sensitivity of data analysis to the temporal intervals at which the data were sampled, similar to the Modifiable Areal Unit Problem (Openshaw 1984) in spatial data analysis. In this work, MTUP arises when detecting and measuring concept drift over different time intervals. We sampled keywords over 5-year periods, as discussed in Section 5.2. However, if we use an alternative sampling interval, e.g., yearly, biennial, or over longer periods like 10 years, the results would be different. Shorter intervals capture more subtle and transient changes, whereas longer intervals would feature more prominent changes.

For instance, Figure 10a shows that examining the keyword *spatial reasoning* (with temporal coverage from 1999 to 2020) over a 10-year period (2005–2015) indicates a potential concept drift. However, Figure 10b, which examines the same keyword at a 5-year interval, reveals that a drift may occur between 2005 and 2010 but not from 2010 to 2015. We may observe similar cases for concept splits or merges, where detection is heavily influenced by the sampling interval. This reminds us of the necessity of choosing an appropriate interval that aligns with the specific goals of the analysis and the characteristics of the dataset.

7.2 | Concept Drift Over Space

Previous work (Wang et al. 2011) defined concept drift as changes in the meaning of a concept over time, possibly also across locations or cultures. However, most studies have focused on the temporal aspect, leaving its spatial aspect largely unexplored. To address this gap, we first differentiate it from




FIGURE 9 | Examples of (a) concept split: evolution of *big data* and (b) concept merge: evolution of *web gis*. The first value in parentheses represents overall similarity, and the second represents textual similarity, both computed between the target keyword and other keywords from adjacent time steps.




FIGURE 10 | Evolution of *spatial reasoning*: (a) from 2005 to 2015. Similarity values in parentheses are computed between the target keyword from 2005 and those from other years and (b) from 2005 to 2010 and 2010 to 2015. Similarity values are computed between the target keyword from 2005 and 2010, and those from other years.

spatial variability. *Spatial variability* refers to differences in how a concept is understood across geographic or cultural contexts (Levinson 1996), such as how the concepts “place” and “space” are perceived by scholars from English-, Spanish-, and German-speaking backgrounds (Blaschke et al. 2018), without incorporating a temporal dimension. However, concept drift over space involves changes in the meaning of a concept *both* over time *and* across locations, potentially occurring at different rates in different regions.

To explore potential concept drift across geographic locations, we collected additional author and affiliation data for a subset of 36 papers with the keyword *qualitative spatial reasoning* for exploratory analysis. Among these, 31 papers were affiliated with institutions in four countries: Australia, Germany, the United Kingdom, and the United States. Each paper was assigned to a single country, even if multiple authors from the same country were involved. Due to collaborative authorship, 7 of these 31 papers were counted multiple times. Table 4 summarizes the number of papers associated with each country and their publication time ranges.

Due to the limited size of this subset, we categorized the data into three phases—early, mid, and late—using quantiles based on each country’s time range. Following the method in Section 5, we aggregated paper titles and abstracts, and computed a single contextual embedding for each phase per country. Since this analysis focuses on a single keyword, textual similarity was not applied. The pairwise cosine similarities, visualized in Figure 11, show varying similarities within the same country across phases and notable differences across countries. For example, the mid-phase work on qualitative spatial reasoning from Germany stands out. While investigating what exactly causes the similarities in the keyword to change at different rates and geographic locations is beyond the scope of this article, we would like to provide some

TABLE 4 | Summary statistics of papers with the keyword *Qualitative Spatial Reasoning* by country of affiliations.

Country	Time range	Number of papers
Australia	2004–2016	8
Germany	1999–2011	8
United Kingdom	1997–2018	9
United States	2002–2017	14

background to make the results less abstract. Just as the concept of *democracy* evolves differently over time *and* space due to cultural and political contexts (Fokkens et al. 2016), local scientific communities also shape their terminology. For instance, European literature often associates spatial reasoning with geo-ontologies and geo-semantics, as part of the broader (predominantly) logic-based knowledge representation literature. In contrast, while such research has also been prominent in the United States, much of the earlier work there was motivated by geo-databases and constraints as part of geographic information systems (Bhatt et al. 2011; Cohn and Renz 2008).

This preliminary analysis suggests the need for further research to comprehensively measure concept drift across *both* time *and* space. While the results do not lead to concrete conclusions about drift or shift in certain geographical locations, they point to a potential direction for future work in this area.

8 | Conclusions

As research topics evolve, social norms change, and technology advances, the keywords used to describe records in RDM

3 Defining Concept Drift and Its Variants in Research Data Management




FIGURE 11 | Pairwise cosine similarity heatmap of *Qualitative Spatial Reasoning* across different time periods and countries.

systems experience semantic aging, which poses challenges across the entire FAIR principles chain. For instance, over time, these changes can hinder the efficient search and retrieval of longitudinal archival documents. Identifying and addressing concept drift offers the potential to mitigate some of these issues and improve the accessibility of scientific literature and datasets. In this work, we explored the dynamic nature of concepts, specifically, focusing on how keywords signify the evolution of concepts in the GIScience literature. We compared two main approaches to detect concept drift—knowledge graph (KG) and natural language processing (NLP)—from the perspectives of concept extraction and similarity measurements. Additionally, we introduced definitions for various forms of conceptual change, including concept *birth*, *drift*, *shift*, *split*, *merge*, and *retirement*. We conducted a scientometric case study analyzing 30 years of GIScience literature, based on the dataset by Wu et al. (2023), to illustrate these definitions in practice.

Our proposed definitions are intended to foster greater interoperability among different domains studying concept drift. This

work mostly focuses on the temporal aspect of concept drift. However, our exploratory analysis of concept drift over space demonstrates the potential to uncover how concepts evolve differently across geographical locations. In future work, we will extend the approach to incorporate a spatial (and place-based) dimension. Our current definitions lay the groundwork for identifying these changes, with the broader goal of establishing standardized terminology and improving interoperability in concept drift research. However, further efforts are needed to develop quantitative methods for measuring and modeling them within specific domains and to extend their broader applications in RDM and beyond.

Acknowledgments

Open access funding provided by Universität Wien/KEMÖ.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study were reused from Wu et al. (2022), which are openly available on Figshare at <https://doi.org/10.6084/m9.figshare.19242654.v1>.

Endnotes

¹Another key notion used in this context is analogy.

²Here, k_i and k_j act as proxies for the underlying concepts.

References

- Barsalou, L. 1993. "Concepts and Meaning." In *Chicago Linguistic Society 29: Papers From the Parasession on Conceptual Representations*, edited by L. Barsalou, W. Yeh, B. Luka, K. Olseth, K. Mix, and L. Wu, 23–61. University of Chicago.
- Beltagy, I., K. Lo, and A. Cohan. 2019. "SciBERT: A Pretrained Language Model for Scientific Text." In *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Bhatt, M., H. Guesgen, S. Wölfl, and S. Hazartka. 2011. "Qualitative Spatial and Temporal Reasoning: Emerging Applications, Trends, and Directions." *Spatial Cognition & Computation* 11, no. 1: 1–14.
- Blaschke, T., H. Merschdorf, P. Cabrera-Barona, S. Gao, E. Papadakis, and A. Kovacs-Györfi. 2018. "Place Versus Space: From Points, Lines and Polygons in GIS to Place-Based Representations Reflecting Language and Culture." *ISPRS International Journal of Geo-Information* 7, no. 11: 452.
- Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3, no. Jan: 993–1022.
- Bordag, S. 2008. "A Comparison of Co-Occurrence and Similarity Measures as Stimulations of Context." In *International Conference on Intelligent Text Processing and Computational Linguistics*, 52–63. Springer.
- Bordes, A., N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. 2013. "Translating Embeddings for Modeling Multi-Relational Data." *Advances in Neural Information Processing Systems* 26: 2787–2795.
- Brown, T. B., B. Mann, N. Ryder, et al. 2020. "Language Models Are Few-Shot Learners." In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20, Red Hook, NY, USA*. Curran Associates Inc.
- Budanitsky, A., and G. Hirst. 2006. "Evaluating Wordnet-Based Measures of Lexical Semantic Relatedness." *Computational Linguistics* 32, no. 1: 13–47.
- Cal, L., K. Janowicz, B. Yan, R. Zhu, and G. Ma. 2021. "Time In a Box: Advancing Knowledge Graph Completion With Temporal Scopes." In *Proceedings of the 11th Knowledge Capture Conference*, 121–128. Association for Computing Machinery.
- Camacho-Collados, J., and M. T. Pilehvar. 2018. "From Word to Sense Embeddings: A Survey on Vector Representations of Meaning." *Journal of Artificial Intelligence Research* 63: 743–788.
- Capobianco, G., D. Cavaliere, S. Senatore, et al. 2020. "OntoDrift: A Semantic Drift Gauge for Ontology Evolution Monitoring." In *CEUR Workshop Proceedings*, vol. 2821, 1–10. CEUR-WS.
- Cheng, T., and M. Adepeju. 2014. "Modifiable Temporal Unit Problem (MTUP) and Its Effect on Space-Time Cluster Detection." *PloS One* 9, no. 6: e100465.
- Chheda, T., P. Goyal, T. Tran, et al. 2021. "Box Embeddings: An Open-Source Library for Representation Learning Using Geometric Structures." arXiv preprint arXiv:2109.04997.
- Chmutina, K., and J. Von Meding. 2019. "A Dilemma of Language: 'Natural Disasters' in Academic Literature." *International Journal of Disaster Risk Science* 10: 283–292.
- Cohn, A. G., and J. Renz. 2008. "Qualitative Spatial Representation and Reasoning." *Foundations of Artificial Intelligence* 3: 551–596.
- Crutzen, P. J. 2006. "The 'Anthropocene'." In *Earth System Science in the Anthropocene*, 13–18. Springer.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. 2019. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." In *North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Egenhofer, M. J., K. C. Clarke, S. Gao, et al. 2016. "Contributions of GIScience Over the Past Twenty Years." In *Advancing Geographic Information Science: The Past and Next Twenty Years, Chapter 1*, edited by H. Onsrud and W. Kuhn, 9–34. GSDI Association Press.
- Fabian, M., K. Gjergji, W. Gerhard, et al. 2007. "Yago: A Core of Semantic Knowledge Unifying Wordnet and Wikipedia." In *16th International World Wide Web Conference*, 697–706. WWW.
- Fokkens, A., S. Ter Braake, I. Maks, and D. Ceolin. 2016. *On the Semantics of Concept Drift: Towards Formal Definitions of Semantic Change*. Drift-a-LOD@ EKAW.
- Frege, G., P. T. Geach, and M. Black. 1951. "On Concept and Object." *Mind* 60, no. 238: 168–180.
- Gama, J., I. Žilobaitė, A. Bifet, M. Pechenizky, and A. Bouchachia. 2014. "A Survey on Concept Drift Adaptation." *ACM Computing Surveys (CSUR)* 46, no. 4: 1–37.
- Goodman, N. 1972. *Problems and Projects*. Bobbs-Merrill Company Inc.
- Gupta, P., and M. Gahegan. 2020. "Categories Are In Flux, but Their Computational Representations Are Fixed: That's a Problem." *Transactions in GIS* 24: 291–314.
- Hamilton, W. L., J. Leskovec, and D. Jurafsky. 2016. "Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change." In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, vol. 2016, 2116. NIH Public Access.
- Hao, J., M. Chen, W. Yu, Y. Sun, and W. Wang. 2019. "Universal Representation Learning of Knowledge Bases by Jointly Embedding Instances and Ontological Concepts." In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1709–1719. Association for Computing Machinery.
- Hitzler, P., M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph. 2009. "OWL 2 Web Ontology Language Primer." *W3C Recommendation* 27, no. 1: 123.
- Hogan, A., E. Blomqvist, M. Cochez, et al. 2021. "Knowledge Graphs." *ACM Computing Surveys (CSUR)* 54, no. 4: 1–37.
- Hood, W. W., and C. S. Wilson. 2001. "The Literature of Bibliometrics, Scientometrics, and Informetrics." *Scientometrics* 52, no. 2: 291–314.
- Huang, Z., D. Wang, B. Huang, et al. 2023. "Concept2Box: Joint Geometric Embeddings for Learning Two-View Knowledge Graphs." arXiv preprint arXiv:2307.01933.
- IAU. 2006. "International Astronomical Union 2006 General Assembly: Result of the IAU Resolution Votes." Accessed 5 November, 2024. <https://www.iau.org/news/pressreleases/detail/iau0603/>.
- Janowicz, K., M. Raubal, and W. Kuhn. 2011. "The Semantics of Similarity in Geographic Information Retrieval." *Journal of Spatial Information Science* 2: 29–57.
- Janowicz, K., C. Shmiz, P. Hitzler, et al. 2022. "Diverse Data! Diverse Schemata?" *Semantic Web* 13, no. 1: 1–3.
- Ji, S., S. Pan, E. Cambria, P. Marttinen, and S. Y. Philip. 2021. "A Survey on Knowledge Graphs: Representation, Acquisition, and Applications." *IEEE Transactions on Neural Networks and Learning Systems* 33, no. 2: 494–514.

3 Defining Concept Drift and Its Variants in Research Data Management

- Jiang, J. J., and D. W. Conrath. 1997. "Semantic Similarity Based on Corpus Statistics and Lexical Taxonomy." arXiv preprint [cmp-lg/9709008](https://arxiv.org/abs/1909.0008).
- Kelley Jr, J. E. 1987. "The Navigation of Columbus on His First Voyage to America." In *Columbus and His World: Proceedings of the First San Salvador Conference*, College Center of the Finger Lakes.
- Kuhn, W., T. Kauptnen, and K. Janowicz. 2014. "Linked Data-A Paradigm Shift for Geographic Information Science." In *Geographic Information Science: 8th International Conference, GIScience 2014, Vienna, Austria, September 24–26, 2014. Proceedings 8*, 173–186. Springer.
- Kuukkanen, J.-M. 2008. "Making Sense of Conceptual Change." *History and Theory* 47, no. 3: 351–372.
- Le, Q., and T. Mikolov. 2014. "Distributed Representations of Sentences and Documents." In *International Conference on Machine Learning*, 1188–1196. PMLR.
- Lehmann, J., R. Isele, M. Jakob, et al. 2015. "Dbpedia—A Large-Scale, Multilingual Knowledge Base Extracted From Wikipedia." *Semantic Web* 6, no. 2: 167–195.
- Levinson, S. C. 1996. "Language and Space." *Annual Review of Anthropology* 25, no. 1: 353–382.
- Li, R., P. Tian, and S. Wang. 2021. "Study Concept Drift in 150-Year English Literature." In *AII@ IConference*, 153–163. CEUR-WS.org.
- Lindqvist, A., M. G. Sendén, and E. A. Renström. 2021. "What Is Gender, Anyway: A Review of the Options for Operationalising Gender." *Psychology & Sexuality* 12, no. 4: 332–344.
- Liu, A., Y. Song, G. Zhang, and J. Lu. 2017. "In IJCAI International Joint Conference on Artificial Intelligence. AAAI Press.
- Lu, J., A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang. 2018. "Learning Under Concept Drift: A Review." *IEEE Transactions on Knowledge and Data Engineering* 31, no. 12: 2346–2363.
- Ly, X., L. Hou, J. Li, and Z. Liu. 2018. "Differentiating Concepts and Instances for Knowledge Graph Embedding." arXiv preprint [arXiv:1811.04588](https://arxiv.org/abs/1811.04588).
- Ma, H. T., C. X. Chu, and H. Paulheim. 2024. "Do LLMs Really Adapt to Domains? An Ontology Learning Perspective." arXiv preprint [arXiv:2407.19998](https://arxiv.org/abs/2407.19998).
- Majic, I., S. Winter, and M. Tomko. 2017. "Finding Equivalent Keys in Openstreetmap: Semantic Similarity Computation Based on Extensional Definitions." In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, 24–32. Association for Computing Machinery.
- McCormick, D. 2012. "Columbus's Geographical Miscalculations." *IEEE Spectrum*. Accessed 31 July, 2024.
- McInnes, L., J. Healy, S. Astels, et al. 2017. "hdbscan: Hierarchical Density Based Clustering." *Journal of Open Source Software* 2, no. 11: 205. <https://doi.org/10.21105/joss.00205>.
- Meroño-Peñuela, A., C. Guéret, R. Hoekstra, and S. Schlobach. 2013. "Detecting and Reporting Extensional Concept Drift in Statistical Linked Data." In *Proceedings of the 1st International Workshop on Semantic Statistics (SemStats 2013), ISWC2013*. CEUR-WS.org.
- Mikolov, T., K. Chen, G. S. Corrado, and J. Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." In *International Conference on Learning Representations*.
- Miller, G. A. 1995. "Wordnet: A Lexical Database for English." *Communications of the ACM* 38, no. 11: 39–41.
- Moreno-Torres, J. G., T. Raeder, R. Alatz-Rodríguez, N. V. Chawla, and F. Herrera. 2012. "A Unifying View on Dataset Shift in Classification." *Pattern Recognition* 45, no. 1: 521–530.
- NASA. 1999. "Mars Climate Orbiter Mishap Investigation Board Phase I Report." https://llis.nasa.gov/llis_lib/pdf/1009464main1_0641-mr.pdf.
- Navigli, R., and P. Velardi. 2004. "Learning Domain Ontologies From Document Warehouses and Dedicated Web Sites." *Computational Linguistics* 30, no. 2: 151–179.
- Ogden, C. K., and I. A. Richards. 1923. *The Meaning of Meaning: A Study of the Influence of Language Upon Thought and of the Science of Symbolism*. Routledge & Kegan Paul.
- Openshaw, S. 1984. "The Modifiable Areal Unit Problem." In *Concepts and Techniques in Modern Geography*. GeoBooks.
- Pasquetto, I. V., B. M. Randles, and C. L. Borgman. 2017. "On the Reuse of Scientific Data." *Data Science Journal* 16, no. 8: 1–9.
- Paulheim, H. 2017. "Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods." *Semantic Web* 8, no. 3: 489–508.
- Pennington, J., R. Socher, and C. D. Manning. 2014. "Glove: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Association for Computational Linguistics.
- Perry, R. W. 2018. "Defining Disaster: An Evolving Concept." In *Handbook of Disaster Research*, 3–22. Springer International Publishing.
- Pirró, G., and J. Euzenat. 2010. "A Feature and Information Theoretic Framework for Semantic Similarity and Relatedness." In *The Semantic Web-ISWC 2010: 9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7–11, 2010, Revised Selected Papers, Part I* 9, 615–630. Springer.
- Raubal, M. 2008. "Representing Concepts in Time." In *Spatial Cognition VI. Learning, Reasoning, and Talking About Space: International Conference Spatial Cognition 2008, Freiburg, Germany, September 15–19, 2008. Proceedings 6*, 328–343. Springer.
- Recchia, G., E. Jones, P. Nulty, J. Regan, and P. deBolla. 2017. "Tracing Shifting Conceptual Vocabularies Through Time." In *Knowledge Engineering and Knowledge Management: EKAW 2016 Satellite Events, EKM and Drift-an-LOD, Bologna, Italy, November 19–23, 2016, Revised Selected Papers*, 19–28. Springer.
- Reimers, N., and I. Gurevych. 2019. "Sentence-BERT: Sentence Embeddings Using Siamese Bert-Networks." In *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Scheider, S., and W. Kuhn. 2015. "How to Talk to Each Other via Computers: Semantic Interoperability as Conceptual Imitation." In *Applications of Conceptual Spaces: The Case for Geometric Knowledge Representation*, 97–122. Springer.
- Schlieder, C. 2010. "Digital Heritage: Semantic Challenges of Long-Term Preservation." *Semantic Web* 1, no. 1–2: 143–147.
- Schlimmer, J. C., and R. H. Granger Jr. 1986. In *Proceedings of the Fifth AAAI National Conference on Artificial Intelligence*, 502–507. AAAI Press.
- Schuldt, J. P., S. H. Konrath, and N. Schwarz. 2011. "Global Warming or 'Climate Change'? Whether the Planet is Warming Depends on Question Wording." *Public Opinion Quarterly* 75, no. 1: 115–124.
- Schwering, A., and M. Raubal. 2005. "Measuring Semantic Similarity Between Geospatial Conceptual Regions." In *International Conference on GeoSpatial Semantics*, 90–106. Springer.
- Shi, M., K. Janowicz, Z. Liu, and K. Currier. 2024. "Towards Formalizing Concept Drift and Its Variants: A Case Study Using Past COSIT Proceedings." In *16th International Conference on Spatial Information Theory (COSIT 2024)*, vol. 315, 23:1–23:8. Schloss Dagstuhl – Leibniz-Zentrum.
- Solarz, M. W. 2012. "Third World: The 60th Anniversary of a Concept That Changed History." *Third World Quarterly* 33, no. 9: 1561–1573.
- Stavropoulos, T. G., S. Andreadis, E. Kontopoulos, and I. Kompatsaris. 2019. "SemaDrift: A Hybrid Method and Visual Tools to Measure Semantic Drift in Ontologies." *Journal of Web Semantics* 54: 87–106.

- Stock, W. G. 2010. "Concepts and Semantic Relations in Information Science." *Journal of the American Society for Information Science and Technology* 61, no. 10: 1951–1969.
- Tasnim, M., D. Collarana, D. Graux, F. Orlandi, and M.-E. Vidal. 2019. "Summarizing Entity Temporal Evolution in Knowledge Graphs." In *Companion Proceedings of the 2019 World Wide Web Conference*, 961–965. Association for Computing Machinery.
- Tietz, T., M. Alam, H. Sack, and M. van Erp. 2020. "Challenges of Knowledge Graph Evolution From an NLP Perspective." In *WHLSe@ESWC*, 71–76. CEUR-WS.org.
- Tsymbal, A. 2004. *The Problem of Concept Drift: Definitions and Related Work*. Computer Science Department, Trinity College Dublin.
- Turney, P. D., and P. Pantel. 2010. "From Frequency to Meaning: Vector Space Models of Semantics." *Journal of Artificial Intelligence Research* 37: 141–188.
- UNISDR. 2018. "United Nations International Strategy for Disaster Reduction. Disaster Terminology." Accessed 7 November, 2024. <https://www.undrr.org/terminology/disaster>.
- Verkijk, S., R. Roothaert, R. Pernisch, and S. Schlobach. 2023. "Do You Catch My Drift? On the Usage of Embedding Methods to Measure Concept Shift in Knowledge Graphs." In *Proceedings of the 12th Knowledge Capture Conference 2023*, 70–74. Association for Computing Machinery.
- Wang, P., and J. Zhou. 2021. "JECI++: A Modified Joint Knowledge Graph Embedding Model for Concepts and Instances." *Big Data Research* 24: 100160.
- Wang, S., S. Schlobach, and M. Klein. 2011. "Concept Drift and How to Identify It." *Journal of Web Semantics* 9, no. 3: 247–265.
- Widmer, G., and M. Kubat. 1996. "Learning in the Presence of Concept Drift and Hidden Contexts." *Machine Learning* 23: 69–101.
- Wijaya, D. T., and R. Yenlier. 2011. "Understanding Semantic Change of Words Over Centuries." In *Proceedings of the 2011 International Workshop on DETecting and Exploiting Cultural Diversity on the Social Web*, 35–40. Association for Computing Machinery.
- Wilkinson, M. D., M. Dumontier, I. J. Aalbersberg, et al. 2016. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data* 3, no. 1: 1–9. <https://doi.org/10.1038/sdata.2016.18>.
- Wu, X., W. Dong, L. Wu, and Y. Liu. 2022. "Data and Code for "Research Themes of Geographical Information Science During 1991–2020: A Retrospective Bibliometric Analysis"." Dataset. <https://doi.org/10.6084/m9.figshare.19242654.v1>.
- Wu, X., W. Dong, L. Wu, and Y. Liu. 2023. "Research Themes of Geographical Information Science During 1991–2020: A Retrospective Bibliometric Analysis." *International Journal of Geographical Information Science* 37, no. 2: 243–275.
- Yang, Z., Z. Dai, Y. Yang, J. G. Carbonell, R. Salakhutdinov, and Q. V. Le. 2019. "XLNet: Generalized Autoregressive Pretraining for Language Understanding." In *Neural Information Processing Systems*. Curran Associates Inc.
- You, S.-C., and H.-T. Lin. 2016. "A Simple Unlearning Framework for Online Learning Under Concept Drifts." In *Advances in Knowledge Discovery and Data Mining: 20th Pacific-Asia Conference, PAKDD 2016, Auckland, New Zealand, April 19–22, 2016, Proceedings, Part I* 20, 115–126. Springer.
- Zhang, P., D. Chen, Y. Fang, X. Zhao, and W. Xiao. 2022. "CIST: Differentiating Concepts and Instances Based on Spatial Transformation for Knowledge Graph Embedding." *Mathematics* 10, no. 17: 3161.
- Zhou, J., P. Wang, Z. Pan, and Z. Xu. 2020. "JECI: A Joint Knowledge Graph Embedding Model for Concepts and Instances." In *Semantic Technology: 9th Joint International Conference, JIST 2019, Hangzhou, China, November 25–27, 2019, Proceedings* 9, 82–98. Springer.
- Zhu, G., and C. A. Iglesias. 2018. "Exploiting Semantic Similarity for Named Entity Disambiguation in Knowledge Graphs." *Expert Systems With Applications* 101: 8–24.

3 Defining Concept Drift and Its Variants in Research Data Management

Appendix A

Conference and Journal Abbreviation Reference

See Table A1.

TABLE A1 | Full names and abbreviations of selected conferences and journals.

Conference/journal name	Abbreviation
International Conference on Spatial Information Theory	COSIT
International Conference on Geographic Information Science	GIScience
Computers, Environment and Urban Systems	CEUS
Cartography and Geographic Information Science	CaGIS
Environment and Planning B: Urban Analytics and City Science	EPB
Geoinformatica	GeoI
International Journal of Geographical Information Science	IJGIS
Journal of Geographical Systems	JGS
Journal of Spatial Information Science	JOSIS
Spatial Cognition & Computation	SCC
Transactions in GIS	TGIS

Appendix B

Sensitivity Analysis of Weight Combinations

To better understand the relation between keywords in their embedding space, we applied HDBSCAN clustering (McInnes et al. 2017) and evaluated the sensitivity to different weightings. We selected HDBSCAN because of its ability to identify clusters of varying densities, making it well-suited for our scientometrics dataset, where the structure of keyword clusters may vary substantially. The resulting clusters using our preferred weighting of $\alpha = 0.2$ and $\beta = 0.8$, which places more emphasis on contextual information, are compared with the swapped weighting scenario in Table B1. We kept the parameters consistent for both cases, with the minimum cluster size of 5 and the minimum samples of 5.

TABLE B1 | Evaluation metrics for keyword clustering using different weight combinations with HDBSCAN.

Weights	Silhouette score	Davies-Bouldin score	Number of clusters
$\alpha = 0.2, \beta = 0.8$	0.5487	0.9003	123
$\alpha = 0.8, \beta = 0.2$	0.2774	1.3198	62

Note: α is the weight assigned to textual information, and β is the weight for contextual information. The Silhouette score ranges from -1 to 1, with higher values indicating well-separated clusters and lower values indicating poorer clustering. The Davies-Bouldin score ranges from 0 to positive values, where lower scores mean better clustering. The number of clusters reported here excludes the noise cluster.

TABLE B2 | Selected cluster results using different weight combinations.

Cluster ($\alpha = 0.2, \beta = 0.8$)	Cluster ($\alpha = 0.8, \beta = 0.2$)
<i>One large cluster:</i> - gis (2005) - visualization (2005) - agent based model (2010) - ldar (2010) - uncertainty (2010) - geographic information system (2010) - cellular automaton (2010) - gis (2010) - spatial analysts (2010) - web gis (2010) - gis (2015) - accessibility (2015) - urban planning (2015) - volunteered geographic information (2015) - geographic information system (2015) ...	<i>Cluster 1—GIS:</i> - gis (1995) - gis (2000) - gis (2005) - modis (2010) - raster (2010) - gis (2010) - gis (2015) - prism (2015) - gis (2020) <i>Cluster 2—Agent-based modeling:</i> - agent based model (2000) - agent based model (2005) - agent based model (2010) - agent based simulation (2015) - agent based model (2015) - agent based model (2020) <i>Cluster 3—Land use/cover:</i> - land use (1995) - land use (2000) - land use (2005) - land cover (2005) - land use (2010) - land cover (2010) - land use (2015) - land cover (2015) - land cover (2020) - land use (2020) - urban land use (2020)

Table B2 provides example keywords from clusters generated using the two weight combinations. We can observe that swapping the weights (i.e., having more emphasis on textual information) leads to keywords with identical labels clustering together, which is not ideal for studying concept drift because it fails to capture contextual variations in meaning.

4 Detecting Spatio-Temporal Concept Drift in Scientific Texts

In this chapter, I extend the definition of concept drift to incorporate a spatial dimension, emphasizing that concepts evolve not only over time but also across space, i.e., at varying rates across regions. I therefore introduce the notion of spatio-temporal concept drift. Leveraging the rich spatial and temporal contextual information in scientific texts, this chapter continues the analysis of the GIScience scientometric dataset from the previous chapter to detect spatio-temporal concept drift. The results reveal different drifting patterns of GIScience concepts, which demonstrate that incorporating spatial perspectives provides a more comprehensive understanding of conceptual changes.

Shi, M., Janowicz, K., Liu, Z., Karimi, M., Majic, I., and Fortacz, A. What, When, and Where Do You Mean? Detecting Spatio-Temporal Concept Drift in Scientific Texts. In 13th International Conference on Geographic Information Science (GIScience 2025). Leibniz International Proceedings in Informatics, Volume 346, pp. 16:1-16:18. <https://doi.org/10.4230/LIPICs.GIScience.2025.16>

The publication is preceded by a co-author declaration specifying individual contributions, with the PhD candidate's research contribution confirmed by all co-authors' signatures.



Department of Geography and Regional Research



universität
wien

Co-author declaration for the following joint paper

This declaration states the research contribution (e.g. research idea and – question; data compilation, manipulation and modelling; design and preparation of graphics, maps and tables; writing of text) of the candidate, the main supervisor (where he/she is an associate author) and the other authors.

If applicable, the contributions from other PhD and / or habilitation candidates who have or intend to include the paper in a thesis should be described. Contributions from master students should also be mentioned.

Research contribution details

Name of candidate:

Meilin Shi

All other authors (Please type the corresponding author in caps):

Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, Alexandra Fortacz

Title of contribution:

What, When, and Where Do You Mean? Detecting Spatio-Temporal Concept Drift in Scientific Texts

Journal details:

Thirteenth International Conference on Geographic Information Science

Author's contribution details

Independent contribution of candidate:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

First author's contribution:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

Contributions of other authors:

- Discussion of the research design
- Revision of the manuscript



Contributions from other colleagues and master students

Has this paper been, or will this paper be part of another doctoral degree thesis?

Yes: ☐

No: ☒

If yes, please elaborate:

Contributions from master students:

Yes: ☐

No: ☒

If yes, please elaborate details:

First author's declaration of contribution for candidate

First author's name:

Meilin Shi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable first author's signature:

21.07.2025





Department of Geography and Regional Research



universität
wien

First **co-author's declaration of contribution for candidate**

Co-author's name:

Krzysztof Janowicz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

Second **co-author's declaration of contribution for candidate**

Co-author's name:

Zilong Liu

Do you verify that the candidate has contributed to this joint paper as described above?


Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025





Department of Geography and Regional Research



universität
wien

Third **co-author's declaration of contribution for candidate**

Co-author's name:

Mina Karimi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

Fourth **co-author's declaration of contribution for candidate**

Co-author's name:

Ivan Majic

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025





Fifth co-author's declaration of contribution for candidate

Co-author's name:

Alexandra Fortacz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☐

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

What, When, and Where Do You Mean? Detecting Spatio-Temporal Concept Drift in Scientific Texts

Meilin Shi¹ ✉ 

Department of Geography and Regional Research, University of Vienna, Austria

Krzysztof Janowicz ✉

Department of Geography and Regional Research, University of Vienna, Austria

Zilong Liu ✉ 

Department of Geography and Regional Research, University of Vienna, Austria

Mina Karimi ✉ 

Department of Geography and Regional Research, University of Vienna, Austria

Ivan Majic ✉ 

Department of Geography and Regional Research, University of Vienna, Austria

Alexandra Fortacz ✉ 

Department of Geography and Regional Research, University of Vienna, Austria

Abstract

Inundated by the rapidly expanding AI research nowadays, the research community requires more effective research data management than ever. A key challenge lies in the evolving nature of concepts embedded in the growing body of research publications. As concepts evolve over time (e.g., keywords like *global warming* become more commonly referred to as *climate change*), past research may become harder to find and interpret in a modern context. This phenomenon, known as *concept drift*, affects how research topics and keywords are understood, categorized, and retrieved. Beyond temporal drift, such variations also occur across geographic space, reflecting differences in local policies, research priorities, and so forth. In this work, we introduce the notion of *spatio-temporal concept drift* to capture how concepts in scientific texts evolve across both space and time. Using a scientometric dataset in geographic information science, we detect how research keywords drifted across countries and years using word embeddings. By detecting spatio-temporal concept drift, we can better align archival research and bridge regional differences, ensuring scientific knowledge remains findable and interoperable within evolving research landscapes.

2012 ACM Subject Classification Information systems → Digital libraries and archives; Computing methodologies → Information extraction; Information systems → Ontologies

Keywords and phrases Concept Drift, Ontology, Large Language Models, Research Data Management

Digital Object Identifier 10.4230/LIPICs.GIScience.2025.16

Supplementary Material *Software (Source Code and Data)*:

<https://github.com/meilinshi/Spatio-temporal-Concept-Drift-in-Scientific-Texts>

archived at `swb:1:dir:928946b53e1ebade1bc2cd2ddfe8a0b23a409111`

1 Introduction

The questions of *what*, *when*, and *where* permeate our daily conversations. When scheduling a group meeting, for instance, we agree on the topic of discussion (e.g., a proposal, *what*), the time (e.g., 10 a.m., *when*), and the location (e.g., a café, *where*). In communicating such

¹ Corresponding author



© Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz;
licensed under Creative Commons License CC-BY 4.0

13th International Conference on Geographic Information Science (GIScience 2025).

Editors: Katarzyna Sila-Nowicka, Antoni Moore, David O'Sullivan, Benjamin Adams, and Mark Gahegan;

Article No. 16; pp. 16:1–16:18



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

16:2 Spatio-Temporal Concept Drift in Scientific Texts

information, we implicitly agree on a particular reference system. For time, we have temporal reference systems such as the Gregorian calendar, the yyyy-mm-dd date format, the 24-hour clock, and so on. For space, we have various geodetic datums, such as WGS 84 and NAD 83, as well as known place names we can refer to. For the *what* question, namely the thematic information, we also need a semantic reference system [23]. In this reference system, an ontology can help ensure that, by “proposal”, we are referring to a research proposal rather than a marriage proposal.

When it comes to ontology modeling and engineering, concepts are often represented as static entities [17]. For example, this is common in a foundational ontology (e.g., DOLCE [15]) to ensure a coherent view and interoperability across domains. In the real world, however, concepts are constantly evolving [19], and their meanings can vary across different social contexts and locations, as seen in the evolving sociocultural definitions of *gender* nowadays. Research in the Semantic Web and the broader knowledge representation and reasoning (KRR) communities has focused on *concept drift* to capture the dynamics of evolving concepts. In this respect, previous work in KRR [49, 14, 44, 8] focused mainly on the temporal aspect of a concept, i.e., the changing meaning of a concept over *time*, and overlooked the spatial perspective that often accompanies it.

In geographic information science (GIScience), constructing an ontology that maps geospatial concepts has always been challenging because of their unique spatio-temporal properties [12, 9]. Geospatial concepts, such as *Mountain* and *Forest*, are different from other general concepts because they do not have clearly defined boundaries nor can they be distinguished in bona fide fashion from neighboring concepts, e.g., *Hill* and *Woods* [42, 43]. For instance, the difference between *Lake* and *Pond* can be affected by seasonal water level variations [28]. This would make downstream tasks, e.g., question answering [33], more challenging. Furthermore, the conceptualization of such landscape concepts may also vary and evolve across languages, cultures, and regions [47, 11].

These challenges are not limited to modeling the aforementioned concepts that are vague geographic features. They also extend to research topics and keywords (e.g., *urban planning*, *climate change*), which we see as signifiers of concepts (i.e., mental representations that categorize areas of research [4]). Although many concepts in this regard exhibit concept drift, geospatial concepts are particularly susceptible due to their inherent dependence on physical, environmental, and sociopolitical contexts. To give a concrete example, the definition of *disasters* could vary significantly depending on local environmental conditions, infrastructure, and response systems. What qualifies as a natural *disaster* in one region (e.g., an earthquake with a magnitude of 6.0 in Haiti) may be labeled differently in another region (e.g., the United States) because of differences in local resilience. In comparison, concepts like *the speed of light* exhibit less spatial variability because of their underlying physical principles.

As Kuhn et al. [24] suggested, we should move space and time from merely being in application domains to becoming foundational aspects of ontologies. While the inherent vagueness in geographic features is not fully resolved, ontologies such as the GeoNames ontology² offer structured representations for these features. However, these efforts have yet to cover more abstract geospatial concepts embedded in research, such as those represented by keywords. Same as geographic features, these concepts are dynamic and spatially grounded, yet they are even more susceptible to societal changes and technological advancements (e.g., the coining of the concept *GeoAI* [21]). In this work, we look into these concepts in scientific research with explicit study areas. We propose an approach to quantify their fluidity and context-dependence across space *and* time via word embeddings. Our long-term goal is to

² <https://www.geonames.org/ontology>

develop an ontology that can address the dynamic nature of these concepts in scientific texts. This contributes to the FAIR principles (Findability, Accessibility, Interoperability, and Reusability) [50] by improving retrieval, reuse, and ensuring the long-term relevance of research [38].

The contributions of this work are as follows:

- We introduce spatio-temporal concept drift, which expands the previous notions of concept drift that focused mainly on temporal changes by incorporating both space *and* time.
- We propose a novel approach to detecting spatio-temporal concept drift in scientific texts via word embeddings.
- We demonstrate how accounting for spatio-temporal concept drift enhances the understanding of concepts in scientific texts, improves recall in FAIR-based research data management systems, and lays the groundwork for ontology learning with large language models (LLMs) in dynamic contexts.

The remainder of this paper is structured as follows. Section 2 introduces the theoretical background of our work. Section 3 reviews related work regarding word embeddings, with a focus on their ability to capture and quantify spatio-temporal variations in semantics. We describe our case study in Section 4, where we use a scientometric dataset in the field of GIScience to assess spatio-temporal concept drift. Section 5 presents the results. Section 6 discusses geographic bias within embeddings and future directions in ontology learning with LLMs. Finally, we conclude our work in Section 7.

2 Background

This section provides the theoretical background for the study on concept drift in knowledge representation. Here, we introduce the notion of spatio-temporal concept drift. This notion adds a spatial dimension to existing definitions of concept drift in the literature. In addition, we discuss the broader implications of spatio-temporal concept drift for research data management (RDM).

2.1 Concepts and Their Representation

The notion of *concept* has many different definitions across or even within domains, such as in linguistics, psychology, computer science, and cognitive science [35]. In this work, we adopt the definition by Stock [45] in information science, which defines a concept as a class containing objects that share certain properties.³

Concepts are fundamental units of meaning and serve as the building blocks of ontologies that help structure knowledge, enable reasoning, and facilitate interoperability. In terms of representation, previous work [49, 44] typically characterized a concept by its label (i.e., name), intension (i.e., defining properties), and extension (i.e., instances that fall under it), in the form $(label(C), int(C), ext(C))$ for a concept C . However, this would only apply to concepts that already existed in predefined ontologies. Verkijk et al. [48] proposed to use embedding techniques to derive vector representations of concepts. While their work focuses on knowledge-graph data, they showcased the ability of embedding techniques to capture flexible, context-aware representations of concepts for natural language data as well, in the form $(label(C), context(C))$.

³ It is worth noting that this definition of *concept* is closer to what cognitive scientists would call a *category*.

16:4 Spatio-Temporal Concept Drift in Scientific Texts

In this work, we treat keywords in research publications as representations of underlying concepts. Unlike established ontological categories, research keywords are rapidly evolving as science and society change. This makes them particularly relevant for studying spatio-temporal concept drift. Here, we focus on research keywords also with the aim of developing a structured ontology that can capture their changing meanings over time and space. Such an ontology could contribute to RDM by improving metadata organization, literature retrieval, question answering, and knowledge graph construction in scientific databases.

2.2 Spatio-Temporal Concept Drift

Adding a temporal dimension to concept representation accounts for changes in their meaning over time. The study of *concept drift*, as defined by Wang et al. [49], aims to capture these changes in concept meaning over time. For example, the keyword *global warming* was once the dominant term in research publications, referring to the rise in Earth's temperature. Over time, *climate change* became more widely used to capture broader climate-induced impacts [26] and account for the fact that *warming* is not uniform. To model a concept C with a temporal component, it can be represented in the form $(label_t(C), int_t(C), ext_t(C))$ or $(label_t(C), context_t(C))$ at time t . Extending this notion to a spatio-temporal dimension means that a concept's meaning may change both over time and space, e.g., at different rates. Here, we define this phenomenon as *spatio-temporal concept drift*. In this case, a concept can be represented as $(label_{t,s}(C), int_{t,s}(C), ext_{t,s}(C))$ or $(label_{t,s}(C), context_{t,s}(C))$ for a concept C at time t and region s .

Figure 1 illustrates how a concept moves in both time and geographic space. More abstractly, this can be thought of as the trajectory of a concept in a space-time prism [35]. The color gradient indicates (*semantic*) *concept similarity* [36, 39, 34, 22], thereby reflecting changes in its thematic dimension over time and space. Take the keywords *global warming* and *climate change* as an illustrative example. Region s_1 may have already adopted the use of *climate change* since time t_1 , while s_2 still has a mixed use of both terms. Note that, here at t_1 , the variation in the thematic dimension between s_1 and s_2 represents spatial variability, which differs from spatio-temporal concept drift, as it captures regional differences without a temporal dimension [25]. Later by t_3 , s_2 adopts the distinct use of *climate change* and *global warming*, aligning its concept representation with s_1 . Over time, as concepts evolve, their meanings may change gradually, showing a concept drift from t_2 to t_3 in s_1 , or change so much that they diverge into two, showing a concept split⁴ in region s_2 . Ultimately, the two concepts may converge into a shared understanding for both regions (indicated by the semantic similarity between C_1 and C'_1 at t_3).

Even with the advent of semantic search [18], which allows for more flexible query interpretation, concept drift remains a challenge, particularly in RDM and other archival systems. If the past and present keywords are not properly linked, search results may still be skewed toward more recent ones, simply because of their pertinence. This would lead to either incomplete retrieval results or misinterpretation of archival documents. Addressing spatio-temporal concept drift helps ensure that evolving knowledge remains accessible and meaningful across different time periods and regions. It could therefore enhance semantic interoperability overall and support the FAIR principles.

⁴ Definitions of *concept drift* and *concept split* are provided in our earlier work [40].




Figure 1 Representation of concepts drift and split over geographic space and time. The intensity of color indicates concept similarity.

3 Related Work

This section reviews existing work that provides means for measuring the latent semantics underlying words in their embeddings, with a spatio-temporal focus.

3.1 Spatial and Temporal Information in Word Embeddings

With the introduction of Word2Vec [29], word embeddings have revolutionized representation by converting words into dense vectors in a high-dimensional space,⁵ where semantically similar words are close to each other. Such representation enables a more flexible study of temporal and spatial variations in lexical semantics, offering an advantage over directly comparing different ontology versions. Early pre-trained word embeddings, such as GloVe [32], provide static representations, where each word is assigned a single vector, independent of context. Later, more advanced models like BERT [10] provide context-aware embeddings that capture more variations in meaning based on surrounding text using attention mechanisms.

Several studies [3, 20, 37] have explored the enrichment of word embeddings with temporal and spatial information. For example, Zhang et al. [54] focused on temporal counterpart search that detects semantically similar terms over time. The authors later also investigated the geographic variations in lexical semantics [55], e.g., showing that *typhoon* in Japan would be the most similar term to *hurricane* in the United States. Gong et al. [16] further extended this idea and proposed a model that conditions word embeddings on time *or* location (i.e., generating time- and location-specific embeddings). Their findings included word similarities over time (e.g., *bitcoin* in 2015 and *stocks* in 1992) and locations (e.g., *president* in the United States and *prime minister* in Canada). A few other studies in GIScience explored

⁵ Note that some literature uses the term “low-dimensional space” here when comparing the dimensionality to a one-hot encoding. We use the term here to signify that the resulting embeddings are in a, say, 300-dimensional vector space.

16:6 Spatio-Temporal Concept Drift in Scientific Texts

the representation learning of *places* via word embeddings, thereby using spatial information alone. Yan et al. [52] applied word-embedding techniques to learn embeddings of places based on their types and distances. Later, Zhai et al. [53] extended this approach to the representation learning of functional regions.

While these approaches captured variations in lexical semantics along one dimension effectively, they did not *jointly* consider spatial and temporal dimensions. As a result, they would fail to capture spatial-temporal lexical similarity, such as **chancellor** in Germany in 2010 and **prime minister** in the United Kingdom in 1980. Such similarity is centered in our study on spatio-temporal concept drift, which accounts for both dimensions at the same time.

3.2 Word Embedding Association Test

The Word Embedding Association Test (WEAT) [7], inspired by the Implicit Association Test in psychology, is a widely used method for quantifying semantic associations in word embeddings. It calculates association scores by comparing cosine similarities between two sets of target words (e.g., man and woman) and two sets of attribute words (e.g., doctor and gynecologist). For example, associations between $\vec{man}$ and $\vec{doctor}$ versus $\vec{woman}$ and $\vec{gynecologist}$ can be assessed using vector arithmetic [6, 30], expressed as $\vec{man} - \vec{woman} \approx \vec{doctor} - \vec{gynecologist}$. The resulting WEAT score indicates the degree of association between the two groups in the embedding space. WEAT provides a standardized measure and allows for statistical significance testing of observed changes. By adapting this method, we can compute the cosine similarities between different temporal snapshots and geographical regions, e.g., (*hurricane* to *Mexico*, 2005) and (*typhoon* to *China*, 2015). Note that 2005 and 2015 are not treated as vectors themselves but rather indicate the time periods associated with these concept-region pairs. If in a vector space, $\vec{hurricane} - \vec{typhoon} \approx \vec{Mexico} - \vec{China}$, this would allow us to quantify concept changes across space and time and reveal geographic prototypes underlying word embeddings. However, applying WEAT to spatio-temporal analysis also presents challenges, particularly in maintaining statistical power when data is sparse across certain regions or time periods.

4 Case Study: Geographical Information Science Publications

We employ a scientometric dataset from Wu et al. [51] to detect the spatio-temporal concept drift in a real-world dataset. This dataset includes research publications in the field of GIScience from 1991 to 2020, sourced from Scopus⁶. As the dataset focuses on international journals and conferences that publish exclusively in English, all included publications are in English. Here, we explicitly focus on papers that mention locations in their abstracts, including geopolitical entities (GPEs) – such as countries, states, and cities – as well as nationalities (NORP), using the spaCy transformer-based named entity recognition pipeline⁷. Table 1 presents the summary statistics of the dataset after filtering for these papers. For reference, we provide the full names of conference and journal abbreviations in Appendix A.

⁶ <https://www.scopus.com/>

⁷ <https://spacy.io/models>

■ **Table 1** Summary statistics of research publications with location mentions in abstracts.

Type	Name	Time Range	Number of Papers	Number of Keywords
Conference	COSIT	1993-2019	22	100
Conference	GIScience	2006-2020	18	78
Journal	CEUS	1999-2020	622	3177
Journal	CaGIS	1991-2020	196	1002
Journal	EPB	1998-2020	352	1744
Journal	GeoI	1997-2020	61	305
Journal	IJGIS	2005-2020	550	2640
Journal	JGS	1996-2020	196	922
Journal	JOSIS	2010-2020	21	133
Journal	SCC	2003-2020	19	92
Journal	TGIS	2007-2020	50	237
Total		1991-2020	2107	10430

4.1 Spatio-Temporal Dimensions of Concepts

As with the *what*, *when*, and *where* questions, we argue that each concept has thematic, temporal, and spatial dimensions. In this dataset, we treat each keyword as signifying an individual concept⁸ and represent these three dimensions accordingly.

To represent the (1) **thematic dimension**, we use the associated abstract, which provides contextual information of each concept. Since Scopus is an abstract and citation database without guaranteed full-text access, abstracts – being more consistently available across publications – are a practical choice for large-scale analysis. For the (2) **temporal dimension**, we use the publication year of the paper associated with each concept. Lastly, for the (3) **spatial dimension**, we use OpenStreetMap Nominatim⁹ to geocode the identified locations and extract the corresponding country for sub-national locations (e.g., cities). For location mentions like “East African”, we retain them at the continent level. If multiple countries are mentioned in an abstract, we document all of them. After retrieving 2,112 publications with location mentions, we manually reviewed 16 unidentified cases, assigning them to the country level or removing them where necessary. This resulted in a final dataset of 2,107 publications.

Table 2 includes examples of abstracts with location mentions and the extracted countries (or regions). The distribution of the 10 most mentioned countries in publications within our dataset is visualized in the heatmap in Figure 2. From this heatmap, we can observe that the United States and China lead in the number of publications, followed by other English-speaking countries (e.g., the United Kingdom, Canada, and Australia) and several European countries. Along the temporal axis, we also observe a notable increase in publications since 2005 in this scientometric dataset.

4.2 Spatio-Temporal Concept Drift in Embedding Space

With the defined spatial and temporal dimensions of each concept, we leverage word embeddings to capture their variations in the thematic dimension. We employ the pre-trained SciBERT model [5], which is designed for scientific texts, to compute embeddings. We use

⁸ The distinction between a symbol and a concept can be explained using the triangle of reference [31].


⁹ <https://nominatim.org>

4 Detecting Spatio-Temporal Concept Drift in Scientific Texts

16:8 Spatio-Temporal Concept Drift in Scientific Texts

■ **Table 2** Exemplar location mentions and extracted countries in abstracts.

Year	Abstract Excerpt	Location Mentions	Country
1999	“We illustrate...based on the street pattern of a small French town.”	[French]	[France]
2008	“A dataset describing...in New York City is analyzed to...the technique.”	[New York City]	[United States]
2017	“Using Austria and Slovenia as a study area,...modified IL.”	[Austria, Slovenia]	[Austria, Slovenia]



■ **Figure 2** Distribution of publications for the 10 most mentioned countries over time.

context-aware representations of concepts in natural language, i.e., $(label(C), context(C))$ for a concept C , as discussed in Section 2.1. Here, we compute these two types of embeddings for concepts (in this case, keywords): (1) **label embedding**, which is static and derived from the keywords themselves, and (2) **context embedding**, which is context-aware and based on their associated abstracts with location mentions.

Additionally, we investigate the sensitivity of context embeddings to location mentions by computing (3) **context embedding without locations** as well. This embedding is derived from associated abstracts of a concept, where each identified location mention is replaced with the placeholder “[Location]”. For instance, in the first example in Table 2, the sentence would become “We illustrate...based on the street pattern of a small [Location] town.” This helps reveal whether explicit geospatial references influence the context-aware representation of concepts.

For each keyword/concept C in a given year t and country (or region) s , we first average its context embeddings across all relevant abstracts, to ensure a single embedding for each unique keyword-year-country combination. We then integrate this with the label embedding through a convex combination to obtain the composite embedding $C_{t,s}$, formulated as:

$$C_{t,s} = \alpha \cdot label(C) + (1 - \alpha) \cdot \frac{1}{|D_{t,s}|} \sum_{d \in D_{t,s}} context(C, d) \quad (1)$$

where $label(C)$ is the embedding of the keyword itself through its label; $context(C, d)$ is the embedding of the abstract in document d containing the keyword; $D_{t,s}$ is the set of documents that contain the keyword from year t and country s ; and $|D_{t,s}|$ is the number of such documents. The parameter α determines the weight assigned to the label versus context embeddings.

To quantify how a concept C drifts across different space-time combinations, we use cosine similarity between their respective composite embeddings. Given two concept representations, e.g., C_{t_1,s_1} and C_{t_2,s_2} , their similarity is computed as:

$$sim(C_{t_1,s_1}, C_{t_2,s_2}) = \frac{C_{t_1,s_1} \cdot C_{t_2,s_2}}{\|C_{t_1,s_1}\| \cdot \|C_{t_2,s_2}\|} \quad (2)$$

5 Results


To start, we visualize these keywords in the embedding space. Figure 3 shows the distribution of keyword embeddings across countries and selected years (2000, 2005, 2010, 2015, and 2020), generated using t-Distributed Stochastic Neighbor Embedding (t-SNE) [46]. The label embedding weight α is empirically set to 0.3 to place greater emphasis on the context embedding while retaining sufficient label information. Higher values of α tend to produce overly label-driven clusters, whereas lower values may cause semantically related keywords to diverge (see Appendix B for examples).

From the figure, we can observe a cluster of keywords with location mentions of China over the years (represented by triangles of different colors in the middle left of the figure), and those with location mentions of European countries like Germany and the Netherlands appear closer to each other.

The t-SNE visualization provides an overview of keyword distributions; we then look into how individual keywords move along their semantic trajectories across different countries over the years. Note that all keywords are standardized to lowercase and American spelling. They are also lemmatized and expanded to their full forms (e.g., *DEM* to *digital elevation model*), with the exception of *GIS*, which we retain as an abbreviation due to its ambiguous reference to *GI Science* or *GI System*. For each keyword, we quantify its spatio-temporal coverage by multiplying its time span (in years) by the number of unique countries it is associated with, yielding a coverage score to reflect both its temporal persistence and geographic distribution. Table 3 presents the top 10 keywords ranked by their spatio-temporal coverage. From these, we plot the semantic trajectories for selected keywords – *GIS*, *urban planning*, *spatial analysis*, and *cellular automaton* – in Figure 4, using principal component analysis (PCA) [1] to reduce the dimensionality of their embeddings.

From these trajectories, we can notice that the extracted embeddings of *GIS* (Figure 4a) are quite consistent across Italy, Germany, Australia, and the UK in the early years of 1999 and 2000. Afterward, English-speaking countries, including the UK, the US, Canada, and Australia, along with China, have their embeddings clustered together. In contrast, European countries, e.g., Spain, France, and Italy, form a separate cluster between 1999 and 2015. This reflects that these countries might take different approaches to GIS theories and applications. Contrary to *GIS*, the semantic trajectories for *urban planning* (Figure 4b) and *spatial analysis* (Figure 4c) vary significantly across different countries. This indicates that these two keywords exhibit strong region-specific embeddings, reflecting that the research under these two keywords in our scientometric dataset is potentially more influenced by local policies, socioeconomic conditions, and so on. Their variations across countries also suggest that, even based on the same theoretical foundation, the practical applications of

16:10 Spatio-Temporal Concept Drift in Scientific Texts



■ **Figure 3** A t-SNE visualization of keyword embeddings with selected years and countries.

these concepts can vary and lead to country-dependent interpretations. Compared with *urban planning* and *spatial analysis*, *cellular automaton* (Figure 4d) shows similar semantic trajectories and clustered embeddings across countries. This country-wise consistency is likely attributed to the stronger mathematical and computational foundations of *cellular automaton*, which makes it potentially less influenced by local policies or conditions. This observation also indicates a more widely shared understanding and development of theories and applications in cellular automaton.

■ **Table 3** The top 10 keywords by spatio-temporal coverage.

Keyword	Time Span	Unique Years	Unique Countries	Total Count	Coverage Score
GIS	1993-2020 (27)	26	45	138	1215
Geographic Information System	1994-2020 (26)	20	36	72	936
Remote Sensing	1995-2020 (25)	17	29	48	725
Land Use	1995-2020 (25)	20	26	55	650
Model	1999-2020 (21)	12	28	41	588
Urban Planning	1998-2020 (22)	15	25	47	550
Spatial Analysis	1998-2020 (22)	18	25	54	550
Cellular Automaton	2000-2020 (20)	19	25	63	500
Visualization	1997-2019 (22)	14	20	39	440
Cadastre	2001-2020 (19)	8	20	24	380




Figure 4 PCA visualization of the semantic trajectory of keywords by country over time. Only the top 10 countries are included for visual clarity. We use the first and the second principle components for PCA visualization.

While semantic trajectories trace how a single keyword (used as a proxy for an underlying concept) evolves over time and across countries, they do not capture how it relates to other keywords in semantic space. Table 4 presents selected examples of keywords in different country-year combinations and how their meanings evolve. For each unique keyword-country-year combination, we identify its three most similar keywords from the other years, calculated based on Equation 2. For *GIS* in Austria (1999), its closest semantic matches are found in the UK and Australia to itself and *integrated model* in the Netherlands. The temporal distance between these matches is small, suggesting that the conceptualization of GIS remained relatively stable across these countries during this period. In contrast, *machine learning* in France (2003) follows a different pattern. It shares the strongest similarities with *ontology* and *temporal management* in Spain (2012) but also aligns with *machine learning* in Czechia (2020). This indicates that the early machine learning concept was probably integrated into various domains of GIScience over time. Lastly, *urban planning* in China (2011 and 2019) show a rather location-stable pattern that the strongest similarities are all found within China across different years and with related planning keywords. This suggests a more internally consistent evolution of urban planning concepts within the Chinese research community. In contrast, *urban planning* in Australia (2004 and 2018) shows similarities across many

4 Detecting Spatio-Temporal Concept Drift in Scientific Texts

16:12 Spatio-Temporal Concept Drift in Scientific Texts

countries (the US, the UK, Netherlands, and Finland). This indicates a more dynamic and globally connected evolution of the underlying concept. We can infer from this comparison that the concept of *urban planning* may drift slower and more localized in China, while at a faster rate and more international in Australia.

■ **Table 4** Selected cases of the top three similar keywords and their similarity scores across countries and years.

Query Keyword	Top 3 Similar Keywords	Sim.
GIS (Austria, 1999)	GIS (United Kingdom, 2000)	0.931
	Integrated Model (Netherlands, 2000)	0.894
	GIS (Australia, 2000)	0.892
Machine Learning (France, 2003)	Ontology (Spain, 2012)	0.916
	Temporal Management (Spain, 2012)	0.916
	Machine Learning (Czechia, 2020)	0.909
Urban Planning (China, 2011)	Urban Planning (China, 2019)	0.942
	Geographic Information System (China, 2016)	0.939
	Planning Support System (China, 2020)	0.938
Urban Planning (China, 2019)	Urban Spatial Dynamic (China, 2020)	0.953
	Scenario Planning (China, 2020)	0.948
	Urban Land Use (China, 2020)	0.948
Urban Planning (Australia, 2004)	Urban Planning (Finland, 2020)	0.947
	Urban Planning (United States, 1998)	0.937
	Urban Planning (Netherlands, 2016)	0.936
Urban Planning (Australia, 2018)	Urban Data (United States, 2019)	0.928
	Urban Land Use Change (United Kingdom, 2014)	0.927
	Urban Scaling Law (Europe, 2020)	0.924

Finally, we perform a sensitivity analysis on countries with at least 100 associated keywords to evaluate the impact of explicit location mentions on the extracted context embeddings. Using a two-tailed permutation test (with 1,000 permutations), we find that the average cosine distance ($1 - \text{cosine similarity}$) between embeddings with and without location mentions (0.010) is significantly smaller than what would be expected by chance (permutation mean = 0.242, std = 0.0005, $p < 0.001$). This indicates that explicit location mentions have minimal impact on the semantic representation of concepts in our case study. This is likely because geographical context is implicitly encoded in the text. We discuss this in more detail in the discussion section.

6 Discussion

This section discusses the challenges of defining spatio-temporal dimensions of concepts and the potential biases introduced in our case study. We also discuss findings from the sensitivity analysis, outline future research directions and the implications of spatio-temporal concept drift for ontology learning with large language models (LLMs).

6.1 Challenges in Defining Spatio-Temporal Dimensions of Concepts

Understanding spatio-temporal concept drift in scientific texts requires linking keywords (and their underlying concepts) to geographic locations and time, but this process inherently introduces biases.

In our case study, we use publication year as the temporal dimension of a concept. However, the publication year could be different from the actual study period (e.g., a paper published in 2010 on East Africa in the 1970s). We attribute spatial dimensions of concepts (keywords) based on location mentions in their associated abstracts. However, not all location mentions correspond to the actual study area; some may appear as examples, comparisons, or even counterexamples. We then aggregate these locations to use the country as the spatial unit of concept drift, overlooking regional variations within the country level. Take the concept of *urban planning* as an example. Its interpretation could differ significantly for New York City (e.g., a walkable, transit-oriented city) versus Los Angeles (e.g., a car-centric city) over the years. These regional disparities would become particularly pronounced in larger countries (e.g., the United States) with diverse geographic and socioeconomic conditions. Our spatial aggregation approach implicitly assumes concept homogeneity at the country level, which introduces biases into the learned embeddings.

Future work could explore improved spatio-temporal scoping techniques to capture study periods and areas more accurately; it should also include different spatial levels – cities, countries, and continents – to measure variations.

6.2 Sensitivity to Location Removal in Context Embeddings

Our sensitivity analysis reveals varying effects across countries when removing location mentions. For example, keywords associated with Japan show slightly larger differences in their embeddings, though the overall differences remained small. Here, several factors may complicate the interpretation of these results. First, the dataset has a substantial imbalance, with publications mentioning the United States far outnumbering those that mention other countries. When extracting unique country-year combinations, this imbalance leads to sparse samples for less represented countries, thus masking meaningful patterns. Second, we did not account for the proportion of location mentions within abstracts, e.g., some abstracts contain a list of study areas, while others mention one location briefly. The observed differences in embeddings with and without locations may be due to the removal of more contextual information rather than an inherent sensitivity to geographic reference. These factors need a more fine-grained analysis to quantify the impact of explicit location mentions on embedding representations in future work.

6.3 Implication for Ontology Learning with Large Language Models

Since large language models (LLMs) become more commonly used for ontology learning tasks [27, 2, 41], we need to ensure geographic and temporal variations in knowledge are accounted for to have more context-aware representations. Current LLMs are trained on vast corpora of text that usually lack explicit spatial and temporal structuring [13], which would likely overlook the spatio-temporal variations in concept representations unless specified in the prompt. Our findings show that concepts in scientific texts evolve differently across geographic space and over time. This suggests that ontologies derived from LLMs may inherit hidden geographic biases. For instance, when an LLM processes the concept of *smart city*, its interpretation might be overly influenced (and represented) by temporally and regionally dominant implementations (e.g., the recent decade in Singapore) if the training data is dominated by publications from this region and time period.

Our observations show that concepts in scientific texts can vary across geographic space and time, and suggest the need for a more context-aware mechanism when using LLMs for ontology learning. This could be achieved with region-specific knowledge validation

16:14 Spatio-Temporal Concept Drift in Scientific Texts

and/or the development of geographically aware prompting strategies. To capture the spatio-temporal dynamics in scientific concepts, future work could include the design of more few-shot learning approaches, where examples are carefully selected to represent diverse temporal and geographic interpretations of concepts.

7 Conclusions

Space and time are central to the study of geography and GIScience. These dimensions not only shape our daily physical interactions of *when* and *where*, but also influence the abstract representation of concepts in scientific knowledge. With the ever-increasing volume of research publications, we need methods to better structure concepts embedded in scientific research for organization and retrieval purposes. Encoded in an ontology, we could also account for the spatio-temporal dynamics of concepts, which are constantly evolving – often at varying rates across regions – due to technology and societal changes, for more effective research data management (RDM).

In this work, we introduce the notion of spatio-temporal concept drift. We complement previous work on concept drift by including the spatial dimension, and propose a novel approach using word embedding techniques to capture this drift over space and time. Using a scientometric dataset in the field of GIScience, we demonstrate that keywords (used as proxies for underlying concepts) show varying drift patterns over time and across countries. Spatially grounded concepts, such as *urban planning* (as compared to *cellular automaton*), can have substantial differences in meanings for different countries and over time.

The implications of this work extend beyond improving the understanding of concepts in scientific texts to enhancing FAIR-based RDM systems. This reminds us, for example, that concepts like *cellular automaton* may require less user intervention, while concepts like *urban planning* may need query enrichment to account for local and temporal variation to better match a user’s keyword. Given the observed spatio-temporal concept drift and the increasing use of ontology learning with large language models (LLMs), we also suggest that LLM-based ontology learning mechanisms should explicitly account for the spatial and temporal dimensions of concept representation. Making RDM systems and ontology learning approaches more sensitive to these variations will help improve retrieval and maintain the relevance of knowledge in scientific texts.

References

- 1 Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010. doi:10.1002/wics.101.
- 2 Hamed Babaei Giglou, Jennifer D’Souza, and Sören Auer. LLMs4OL: Large language models for ontology learning. In *International Semantic Web Conference*, pages 408–427. Springer, 2023. doi:10.1007/978-3-031-47240-4_22.
- 3 David Bamman, Chris Dyer, and Noah A Smith. Distributed representations of geographically situated language. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 828–834, 2014. doi:10.3115/v1/p14-2134.
- 4 Lawrence Barsalou. Concepts and meaning. In L. Barsalou, W. Yeh, B. Luka, K. Olseth, K. Mix, and L. Wu, editors, *Chicago Linguistic Society 29: Papers From the Parasession on Conceptual Representations*, pages 23–61. University of Chicago, 1993.
- 5 Iz Beltagy, Kyle Lo, and Arman Cohan. SciBERT: A pretrained language model for scientific text. In *Conference on Empirical Methods in Natural Language Processing*, 2019. doi:10.18653/v1/D19-1371.

- 6 Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29, 2016. doi:10.48550/arXiv.1607.06520.
- 7 Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 2017. doi:10.1126/science.aal4230.
- 8 Giuseppe Capobianco, Danilo Cavaliere, Sabrina Senatore, et al. Ontodrift: a semantic drift gauge for ontology evolution monitoring. In *CEUR Workshop Proceedings*, volume 2821, pages 1–10. CEUR-WS, 2020. URL: <https://ceur-ws.org/Vol-2821/paper1.pdf>.
- 9 Christophe Claramunt. Ontologies for geospatial information: Progress and challenges ahead. *Journal of Spatial Information Science*, 20:35–41, 2020. doi:10.5311/JOSIS.2020.20.666.
- 10 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. doi:10.18653/v1/N19-1423.
- 11 Stephanie Duce and Krzysztof Janowicz. Microtheories for spatial data infrastructures-accounting for diversity of local conceptualizations at a global level. In *Geographic Information Science: 6th International Conference, GIScience 2010, Zurich, Switzerland, September 14-17, 2010. Proceedings 6*, pages 27–41. Springer, 2010. doi:10.1007/978-3-642-15300-6_3.
- 12 Max J Egenhofer and David M Mark. Naive geography. In *Spatial Information Theory A Theoretical Basis for GIS: International Conference COSIT'95 Semmering, Austria, September 21-23, 1995 Proceedings 2*, pages 1–15. Springer, 1995. doi:10.1007/3-540-60392-1_1.
- 13 Fahim Faisal and Antonios Anastasopoulos. Geographic and geopolitical biases of language models. In Duygu Ataman, editor, *Proceedings of the 3rd Workshop on Multi-lingual Representation Learning (MRL)*, pages 139–163, Singapore, December 2023. Association for Computational Linguistics. doi:10.18653/v1/2023.mrl-1.12.
- 14 Antske Fokkens, Serge Ter Braake, Isa Maks, Davide Ceolin, et al. On the semantics of concept drift: Towards formal definitions of semantic change. *Drift-a-LOD@ EKAW*, 2016. URL: https://ceur-ws.org/Vol-1799/Drift-a-LOD2016_paper_2.pdf.
- 15 Aldo Gangemi, Nicola Guarino, Claudio Masolo, Alessandro Oltramari, and Luc Schneider. Sweetening ontologies with DOLCE. In *International conference on knowledge engineering and knowledge management*, pages 166–181. Springer, 2002. doi:10.1007/3-540-45810-7_18.
- 16 Hongyu Gong, S. Bhat, and Pramod Viswanath. Enriching word embeddings with temporal and spatial information. In *Conference on Computational Natural Language Learning*, 2020. doi:10.18653/v1/2020.conll-1.1.
- 17 Nicola Guarino. Formal ontology, conceptual analysis and knowledge representation. *International journal of human-computer studies*, 43(5-6):625–640, 1995. doi:10.1006/ijhc.1995.1066.
- 18 Ramanathan Guha, Rob McCool, and Eric Miller. Semantic search. In *Proceedings of the 12th international conference on World Wide Web*, pages 700–709, 2003. doi:10.1145/775152.775250.
- 19 Prashant Gupta and Mark Gahegan. Categories are in flux, but their computational representations are fixed: That's a problem. *Transactions in GIS*, 24(2):291–314, 2020. doi:10.1111/tgis.12602.
- 20 William L. Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, 2016. doi:10.18653/v1/P16-1141.
- 21 Krzysztof Janowicz, Song Gao, Grant McKenzie, Yingjie Hu, and Budhendra Bhaduri. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond, 2020. doi:10.1080/13658816.2019.1684500.

4 Detecting Spatio-Temporal Concept Drift in Scientific Texts

16:16 Spatio-Temporal Concept Drift in Scientific Texts

- 22 Krzysztof Janowicz, Martin Raubal, and Werner Kuhn. The semantics of similarity in geographic information retrieval. *Journal of Spatial Information Science*, 2:29–57, 2011. doi:10.5311/JOSIS.2011.2.3.
- 23 Werner Kuhn. Semantic reference systems. *International Journal of Geographical Information Science*, 17(5):405–409, 2003. doi:10.1080/1365881031000114116.
- 24 Werner Kuhn, Martin Raubal, and Peter Gärdenfors. Cognitive semantics and spatio-temporal ontologies. *Spatial Cognition & Computation*, 7(1):3–12, 2007. doi:10.1080/13875860701337835.
- 25 Stephen C Levinson. Language and space. *Annual review of Anthropology*, 25(1):353–382, 1996. doi:10.1146/annurev.anthro.25.1.353.
- 26 Maurice Lineman, Yuno Do, Ji Yoon Kim, and Gea-Jae Joo. Talking about climate change and global warming. *PloS one*, 10(9):e0138996, 2015. doi:10.1371/journal.pone.0138996.
- 27 Huu Tan Mai, Cuong Xuan Chu, and Heiko Paulheim. Do LLMs really adapt to domains? an ontology learning perspective. In *International Semantic Web Conference*, pages 126–143. Springer, 2024. doi:10.1007/978-3-031-77844-5_7.
- 28 David M. Mark, Barry Smith, and Barbara Tversky. Ontology and geographic objects: An empirical study of cognitive categorization. In *Conference On Spatial Information Theory*, pages 283–298, 1999. doi:10.1007/3-540-48384-5_19.
- 29 Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*, 2013. doi:10.48550/arXiv.1301.3781.
- 30 Malvina Nissim, Rik van Noord, and Rob van der Goot. Fair is better than sensational: Man is to doctor as woman is to doctor. *Computational Linguistics*, 46(2):487–497, 2020. doi:10.1162/coli_a_00379.
- 31 Charles Kay Ogden and Ivor Armstrong Richards. *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*. Harcourt, Brace & World, Inc., 1923. doi:10.2307/2015195.
- 32 Jeffrey Pennington, Richard Socher, and Christopher D Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. doi:10.3115/v1/d14-1162.
- 33 Dharmen Punjani, Kuldeep Singh, Andreas Both, Manolis Koubarakis, Iosif Angelidis, Konstantina Bereta, Themis Beris, Dimitris Bilidas, Theofilos Ioannidis, Nikolaos Karalis, et al. Template-based question answering over linked geospatial data. In *Proceedings of the 12th workshop on geographic information retrieval*, pages 1–10, 2018. doi:10.1145/3281354.3281362.
- 34 Martin Raubal. Formalizing conceptual spaces. In *Formal ontology in information systems, proceedings of the third international conference (FOIS 2004)*, volume 114, pages 153–164. Citeseer, 2004.
- 35 Martin Raubal. Representing concepts in time. In *Spatial Cognition VI. Learning, Reasoning, and Talking about Space: International Conference Spatial Cognition 2008, Freiburg, Germany, September 15-19, 2008. Proceedings 6*, pages 328–343. Springer, 2008. doi:10.1007/978-3-540-87601-4_24.
- 36 M Andrea Rodriguez and Max J. Egenhofer. Determining semantic similarity among entity classes from different ontologies. *IEEE transactions on knowledge and data engineering*, 15(2):442–456, 2003. doi:10.1109/TKDE.2003.1185844.
- 37 Maja Rudolph and David Blei. Dynamic embeddings for language evolution. In *Proceedings of the 2018 world wide web conference*, pages 1003–1011, 2018. doi:10.1145/3178876.3185999.
- 38 Christoph Schlieder. Digital heritage: Semantic challenges of long-term preservation. *Semantic Web*, 1(1-2):143–147, 2010. doi:10.3233/SW-2010-0013.
- 39 Angela Schwering. Approaches to semantic similarity measurement for geo-spatial data: a survey. *Transactions in GIS*, 12(1):5–29, 2008. doi:10.1111/j.1467-9671.2008.01084.x.

- 40 Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz. Defining concept drift and its variants in research data management: A scientometric case study on geographic information science. *Transactions in GIS*, 29(3):e70058, 2025. doi:10.1111/tgis.70058.
- 41 Cogan Shimizu and Pascal Hitzler. Accelerating knowledge graph and ontology engineering with large language models. *Journal of Web Semantics*, page 100862, 2025. doi:10.1016/j.websem.2025.100862.
- 42 Barry Smith and David M Mark. Geographical categories: an ontological investigation. *International journal of geographical information science*, 15(7):591–612, 2001. doi:10.1080/13658810110061199.
- 43 Barry Smith and David M Mark. Do mountains exist? towards an ontology of landforms. *Environment and Planning B: Planning and Design*, 30(3):411–427, 2003. doi:10.1068/b12821.
- 44 Thanos G Stavropoulos, Stelios Andreadis, Efstratios Kontopoulos, and Ioannis Kompatsiaris. SemaDrift: A hybrid method and visual tools to measure semantic drift in ontologies. *Journal of Web Semantics*, 54:87–106, 2019. doi:10.1016/j.websem.2018.05.001.
- 45 Wolfgang G Stock. Concepts and semantic relations in information science. *Journal of the American Society for Information Science and Technology*, 61(10):1951–1969, 2010. doi:10.1002/asi.21382.
- 46 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- 47 Saskia Van Putten, Carolyn O'Meara, Flurina Wartmann, Joanne Yager, Julia Vilette, Claudia Mazzuca, Claudia Bieling, Niclas Burenhult, Ross Purves, and Asifa Majid. Conceptualisations of landscape differ across european languages. *Plos one*, 15(10):e0239858, 2020. doi:10.1371/journal.pone.0239858.
- 48 Stella Verkijk, Ritten Roothaert, Romana Pernisch, and Stefan Schlobach. Do you catch my drift? on the usage of embedding methods to measure concept shift in knowledge graphs. In *Proceedings of the 12th Knowledge Capture Conference 2023*, pages 70–74, 2023. doi:10.1145/3587259.3627555.
- 49 Shenghui Wang, Stefan Schlobach, and Michel Klein. Concept drift and how to identify it. *Journal of Web Semantics*, 9(3):247–265, 2011. doi:10.1016/j.websem.2011.05.003.
- 50 Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016. doi:10.1038/sdata.2016.18.
- 51 Xiaohuan Wu, Weihua Dong, Lun Wu, and Yu Liu. Data and Code for "Research Themes of Geographical Information Science during 1991–2020: A Retrospective Bibliometric Analysis", 2022. doi:10.6084/m9.figshare.19242654.v1.
- 52 Bo Yan, Krzysztof Janowicz, Gengchen Mai, and Song Gao. From ITDL to Place2Vec: Reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In *Proceedings of the 25th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 1–10, 2017. doi:10.1145/3139958.3140054.
- 53 Wei Zhai, Xueyin Bai, Yu Shi, Yu Han, Zhong-Ren Peng, and Chaolin Gu. Beyond Word2vec: An approach for urban functional region extraction and identification by combining Place2vec and POIs. *Computers, Environment and Urban Systems*, 74:1–12, 2019. doi:10.1016/j.compenvurbsys.2018.11.008.
- 54 Yating Zhang, Adam Jatowt, Sourav S Bhowmick, and Katsumi Tanaka. The past is not a foreign country: Detecting semantically similar terms across time. *IEEE Transactions on Knowledge and Data Engineering*, 28(10):2793–2807, 2016. doi:10.1109/TKDE.2016.2591008.
- 55 Yating Zhang, Adam Jatowt, and Katsumi Tanaka. Is tofu the cheese of asia?: Searching for corresponding objects across geographical areas. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1033–1042, 2017. doi:10.1145/3041021.3055132.

4 Detecting Spatio-Temporal Concept Drift in Scientific Texts

16:18 Spatio-Temporal Concept Drift in Scientific Texts

A Conference and Journal Abbreviation Reference

Table 5 Full names and abbreviations of selected conferences and journals.

Conference/Journal Name	Abbreviation
International Conference on Spatial Information Theory	COSIT
International Conference on Geographic Information Science	GIScience
Computers, Environment and Urban Systems	CEUS
Cartography and Geographic Information Science	CaGIS
Environment and Planning B: Urban Analytics and City Science	EPB
GeoInformatica	GeoI
International Journal of Geographical Information Science	IJGIS
Journal of Geographical Systems	JGS
Journal of Spatial Information Science	JOSIS
Spatial Cognition & Computation	SCC
Transactions in GIS	TGIS

B Sensitivity Analysis of the Label Embedding Weight

Table 6 Top similar keywords retrieved under different values of the label embedding weight α , along with their similarity scores. The value of $\alpha = 0.3$ is used in the case study in this paper. Note that the examples are included post hoc to illustrate the qualitative effects of different α values.

Weight	Query Keyword	
	Urban Planning (Australia, 2018)	Climate Change (US, 2020)
$\alpha = 0.1$	urban scaling law (Europe, 2020): 0.925	climate change (US, 2014): 0.913
	zipf's law for city (Europe, 2020): 0.924	sea level rise (US, 2019): 0.910
	land use (Europe, 2020): 0.923	storm surge inundation (US, 2019): 0.909
	population density (Europe, 2020): 0.922	lidar (US, 2008): 0.909
	radial analysis (Europe, 2020): 0.922	greening scenario (US, 2018): 0.909
$\alpha = 0.3$	urban data (US, 2019): 0.928	climate change (US, 2014): 0.938
	urban land use change (UK, 2014): 0.927	climate change (UK, 2018): 0.919
	urban scaling law (Europe, 2020): 0.924	sea level rise (US, 2019): 0.912
	urban planning (US, 2019): 0.924	urban heat island (US, 2018): 0.911
	residential mobility (UK, 2014): 0.921	seasonal impact (US, 2018): 0.911
$\alpha = 0.5$	urban planning (US, 2019): 0.959	climate change (US, 2014): 0.967
	urban planning (Brazil, 2003): 0.948	climate change (UK, 2018): 0.956
	urban planning (Poland, 2017): 0.948	climate change (UK, 2012): 0.947
	urban planning (Spain, 2017): 0.947	climate change (US, 2013): 0.947
	urban planning (Netherlands, 2016): 0.943	climate change (US, 2015): 0.946

5 Quantifying Spatial, Temporal, and Semantic Consequences of Ontology Changes

In this chapter, I focus on quantifying the consequences of conceptual changes in structured knowledge systems, such as geospatial ontologies. I propose a multi-dimensional framework to quantify their impact across space, time, and semantics. I also introduce a weighting scheme, which enables a ranked evaluation of changes based on relation types (e.g., taxonomic, partonomic) and their associated structural significance. To demonstrate the applicability of the framework, I construct a maritime legal ontology and populate it with real-world data to form an exemplary knowledge graph (KG). The results show that a single schema-level change can have cascading effects within the KG and impact downstream spatio-temporal reasoning.

The publication is preceded by a co-author declaration specifying individual contributions, with the PhD candidate's research contribution confirmed by all co-authors' signatures.



Department of Geography and Regional Research



universität
wien

Co-author declaration for the following joint paper

This declaration states the research contribution (e.g. research idea and – question; data compilation, manipulation and modelling; design and preparation of graphics, maps and tables; writing of text) of the candidate, the main supervisor (where he/she is an associate author) and the other authors.

If applicable, the contributions from other PhD and / or habilitation candidates who have or intend to include the paper in a thesis should be described. Contributions from master students should also be mentioned.

Research contribution details

Name of candidate:

Meilin Shi

All other authors (Please type the corresponding author in caps):

Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, Alexandra Fortacz

Title of contribution:

The Geography of Evolving Concepts: Quantifying Spatial, Temporal, and Semantic Consequences of Ontology Changes

Journal details:

International Journal of Digital Earth

Author's contribution details

Independent contribution of candidate:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

First author's contribution:

- Writing the paper
- Research idea and questions
- Data compilation
- Experiment design and implementation

Contributions of other authors:

- Discussion of the research design
- Revision of the manuscript



Contributions from other colleagues and master students

Has this paper been, or will this paper be part of another doctoral degree thesis?

Yes: ☐

No: ☒

If yes, please elaborate:

Contributions from master students:

Yes: ☐

No: ☒

If yes, please elaborate details:

First author's declaration of contribution for candidate

First author's name:

Meilin Shi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable first author's signature:

21.07.2025





Department of Geography and Regional Research



universität
wien

First **co-author's declaration of contribution for candidate**

Co-author's name:

Krzysztof Janowicz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025



Department of Geography and Regional Research



universität
wien

Second **co-author's declaration of contribution for candidate**

Co-author's name:

Zilong Liu

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025



Department of Geography and Regional Research



universität
wien

Third **co-author's declaration of contribution for candidate**

Co-author's name:

Mina Karimi

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

Fourth **co-author's declaration of contribution for candidate**

Co-author's name:

Ivan Majic

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☒

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025





Fifth co-author's declaration of contribution for candidate

Co-author's name:

Alexandra Fortacz

Do you verify that the candidate has contributed to this joint paper as described above?

Yes: ☐

No: ☐

If no, please specify:

If applicable co-author's signature:

23.07.2025

The Geography of Evolving Concepts: Quantifying Spatial, Temporal, and Semantic Consequences of Ontology Changes

Meilin Shi^a, Krzysztof Janowicz^a, Zilong Liu^a, Mina Karimi^a, Ivan Majic^a, and Alexandra Fortacz^a

^aDepartment of Geography and Regional Research, University of Vienna, Austria

ARTICLE HISTORY

Compiled November 11, 2025

ABSTRACT

Concepts, shaped by human perception and cognition of the world, are constantly evolving. Among these, geographic concepts (e.g., *forest*, *mountain*, *island*) are particularly dynamic and context-dependent. These concepts have a key role in inter-linked knowledge systems, providing rich spatio-temporal reasoning in applications such as disaster response, land use classification, and urban planning. However, their evolving nature makes them susceptible to reinterpretation. Even minor changes in their definition, spatial extent, or temporal scope can cascade through ontology-driven reasoning systems, disrupting data reuse, integration, inference, and decision-making processes. In this work, we aim to quantify how the impact of ontology changes propagates through space, time, and semantics using a multi-dimensional framework. We focus on conceptual changes—rather than changes to individual instances—and assess their downstream impacts within a geospatial ontology. To showcase the impacts of such concept drifts, we construct a maritime legal ontology as a real-world application of our framework. We demonstrate that by adding just one additional criterion to the definition of *rocks*, 6.26% of statements in our exemplary knowledge graph would require revalidation. By quantifying the impact, we seek to inform more robust research data management practices and ensure the long-term usability of spatio-temporal knowledge systems in evolving domain semantics.

KEYWORDS

geospatial ontologies, knowledge representation, change impact analysis, research data management

1. Introduction

On January 20, 2025, the U.S. White House issued an executive order to rename the Gulf of Mexico to the Gulf of America in federal documents. One may wonder: what would such a change bring? How many systems might be affected by this sudden naming change? Beyond the map label updates (e.g., renaming or displaying dual names) lie potential cascading consequences—outdated database links, failed historical queries, confused AI systems trained on old toponyms, inconsistent educational materials, and challenges in international communication, to name a few. A slight change in geographic naming can trigger updates across countless interconnected systems, yet

CONTACT Meilin Shi. Email: meilin.shi@univie.ac.at

we lack a systematic way to quantify the cost and complexity of such changes (Polleres et al., 2023).

In a formal ontology, we can count the number of assertions that are affected by the renaming of the Gulf of Mexico. However, it is important to note that this reflects only an instance-level change and, more specifically, a *relabeling* of an individual instance. The consequences can become substantially more problematic when changes occur at the level of concept definitions. For example, when the U.S. government adjusted gender classification from a spectrum to binary biological sex in federal systems through executive order (White House, 2025), some medical databases have required structural adjustments (U.S. Department of Health and Human Services, 2025). These concept-level (or schema-level) modifications, unlike simple value updates, reshape the semantic architecture of our knowledge systems.

While the above examples serve as anecdotal cases, they remind us of the need to systematically evaluate the cost of changes to ontologies—whether for anticipating technical failure, ensuring system robustness, or informing policy decisions. This need applies to all domains of knowledge representation. However, it is especially important for geography-related contexts, such as disaster response, urban planning, land cover, land use policies, and so on (Gupta and Gahegan, 2020). In these domains, outdated or inconsistent information can have tangible real-world consequences. For instance, if what qualifies as a *fire-prone area* is redefined based on new prediction models or regulatory standards, then follow-up insurance coverage and evacuation priorities must be updated accordingly; otherwise, infrastructure and human lives may be put at risk.

Likewise, when terminology or classification schemes change in research data management (RDM), they can hinder data discovery, retrieval, integration, interoperability (Fonseca et al., 2006; Harvey et al., 1999), and ultimately, reuse (Schlieder, 2010). Updates to shared vocabularies, metadata standards, etc., can make it difficult for researchers to retrieve datasets, reproduce experiments, and interpret historical records (Shi et al., 2025). These challenges are complex in geography-related domains, where categories often evolve not only over time but also across space. For instance, what qualifies as a *forest* varies across regions (Bennett, 2001; Janowicz, 2010) and may evolve at different rates because of local environmental conditions, making it more difficult to manage the conceptual changes. Measuring the impact of such changes is necessary for anticipating system disruptions and developing follow-up response strategies (Abgaz et al., 2012). Notably, not all changes carry the same weight (Kong et al., 2025): a simple synonym update (e.g., replacing the label of *wildfire* with *bushfire*) may have minimal impact, whereas the impact of a fundamental redefinition of a core concept (e.g., reclassifying *natural disasters* from *natural hazards*) can change the entire organization of the knowledge base. By assigning different weights to different types of changes—for instance, based on their entailment (i.e., the extent to which one change logically triggers other related changes) (de Salvo Braz et al., 2005)—we can better evaluate their consequences and prioritize which updates require immediate attention. This is conceptually similar to how (online) vandalism in user-generated content such as Wikipedia is managed semi-automatically (Potthast et al., 2008).

To measure the consequences of conceptual changes, our work focuses on schema-level changes (i.e., TBox-level changes), which are modifications to the definitions and relations between concepts, rather than changes to individual data instances (i.e., ABox-level changes). Although there is substantial research on ontology evolution (Pernisch et al., 2021; Zablith et al., 2015), several challenges remain: it is common to (1) overlook the complexity of changes in the real world from a spatio-temporal perspective (Claramunt, 2020) and (2) treat different types of changes identically, without

differentiating between changes to the TBox and ABox (Chen et al., 2021). To address these gaps, we introduce a framework for assessing the impact of conceptual changes across space, time, and semantics. This multi-dimensional view allows us to interpret how conceptual changes affect knowledge systems—not only in the abstract semantic space but also in the tangible physical space. When estimating the impact, we also incorporate a weighting scheme to rank changes by their structural significance, in other words, their potential to trigger cascading effects within the ontology, and, hence, the systems designed on top, such as knowledge graphs, RDM systems, and scientific workflows more broadly (Ludašcher et al., 2006).

The contributions of this work are as follows:

- We propose a multi-dimensional framework to quantify the impact of conceptual changes in ontologies—particularly geospatial ontologies—across space, time, and semantics.
- We develop a weighting scheme for change impact analysis to rank relation types (e.g., taxonomic, partonomic, spatial relations) based on their structural significance within the ontology.
- We construct a maritime legal ontology with real-world data and demonstrate the applicability of our framework to evaluate the cost of changes (e.g., reclassifying maritime zones) in practical scenarios.

The remainder of this paper is structured as follows. Section 2 provides background on the uniqueness of geographic concepts, geospatial ontologies, and their spatio-temporal representations. We then review existing quantification methods for ontology evolution and change impact analysis. In Section 3, we introduce our framework, which progressively examines spatial, spatio-temporal, and integrated semantic impacts of conceptual changes. Section 4 focuses on the application use case of the framework. We present our constructed maritime legal ontology and summarize the results. Section 5 discusses the applicability of our framework to both geographic and non-geographic concepts, and its further implications for RDM and ontology management. Finally, we conclude our work in Section 6.

2. Background and Related Work

In this section, we provide background on the evolving nature of geographic concepts and their spatio-temporal representation in geospatial ontologies (Kuhn, 2005). We then review existing approaches from the broader fields of ontology evolution and change impact analysis, with a focus on gaps in capturing schema-level changes in geographic concepts.

2.1. Dynamics of Geographic Concepts

Geographic concepts (e.g., *forest*, *mountain*, *river*, *lake*), as noted by Goodchild et al. (2000), are special in that they bridge physical space-time and social interactions. They are grounded not just in coordinates, but also in context—locations on a map and meanings assigned to those locations (Mark et al., 1999). Although embedded in diverse systems like disaster response and remote sensing, these concepts remain open to reinterpretation and redefinition over time (Grenon and Smith, 2004). The vagueness and dynamics of geographic concepts make them difficult to categorize. Approaches to address this include calling for semantic flexibility (e.g., adding a threshold

parameter) (Bennett et al., 2008), work on semantic similarity (Adams and Raubal, 2009; Schwering and Raubal, 2005), and event recognition for continuous monitoring to accommodate seasonal variations, agricultural cycles, etc. (Camara, 2020).

The meaning of geographic concepts can evolve under numerous conditions (Pike and Gahegan, 2007), e.g., when administrative boundaries are redrawn, land use classification systems are revised, natural features are reclassified, new environmental phenomena emerge, or when cultural, societal, or scientific (paradigm) shifts change our perspective. A redefinition of *forest*, for instance, can disrupt time series analyses, leading to inaccurate results of deforestation rates and, therefore, potentially misleading restoration policies (Zalles et al., 2024). Emerging concepts like *climate refuge*—referring to locations expected to remain habitable under extreme weather (Amorim-Maia et al., 2023)—illustrate how new environmental realities can prompt the coining of new geographic concepts. These changes are beyond simple adjustments to metadata; they can alter the semantics that underpin spatial analysis and policy decisions. Terminologies need to be revised, previously inferred results revalidated, and affected datasets reannotated. When such changes in geographic concepts are overlooked, the consequences are not only semantic inconsistencies in RDM (e.g., introducing ambiguities in data integration) but also tangible impacts on how resources are allocated and responses are prioritized in the physical world.

2.2. Space and Time in Geospatial Ontologies

In information science, an ontology is a shared and formal specification of concepts and relations abstracted from reality (Gruber, 1993), which is designed to enable interoperability and knowledge sharing across different applications (Guarino, 1998). A geospatial ontology builds on this by including more information on space and time, so as to support structured reasoning about geographic information (Claramunt, 2020; Galton, 2003; Kuhn, 2003).

Unlike general-purpose ontologies used by knowledge bases such as DBpedia (Auer et al., 2007), which typically represent locations with latitude-longitude pairs, geospatial ontologies contain richer spatial semantics. For example, YAGO2geo (Karalis et al., 2019) contains diverse geometries (line, polygon, multipolygon, etc.), which enables more precise spatial representation than point-based approaches only. It supports the computation of topological relations grounded in formal logics such as the Region Connection Calculus (RCC) (Randell et al., 1992) and Dimensionally Extended 9-Intersection Model (DE-9IM) (Clementini et al., 1994), capturing relations like containment, overlap, and touching. With GeoSPARQL (<https://www.ogc.org/standards/geosparql/>), geospatial ontologies enable queries via functions like `geo:sfWithin` and `geo:sfIntersects` to answer spatial questions such as “Which lakes form the borders between countries?” or “How many countries does the Danube River flow through?”. Such spatial reasoning is not supported by general-purpose ontologies due to their limited geometric representations and lack of spatial reasoning capabilities.

In addition to the spatial dimension, geospatial ontologies are also concerned with enhancing the temporal representation of how entities evolve over time (Hoffart et al., 2013; Janowicz, 2010; Keßler and Farmer, 2015). Since many entities (e.g., the borders of administrative areas) undergo changes, temporal reasoning becomes an important challenge in geospatial knowledge representation (Kuhn et al., 2014). For example, a triple like `<Prague, isPartOf, Austro-HungarianEmpire> [validDuring:`

1867–1918] expresses not only a spatial relation but also its temporal validity. This spatio-temporal focus makes geospatial ontologies particularly sensitive to change. Any update—whether a boundary redefinition, renaming of a geographic entity, or the emergence of a new temporal scope—can impact both spatial and temporal reasoning, e.g., when querying overlaps between historical and contemporary regions (Kauppinen and Hyvönen, 2007). Therefore, more effort is required to maintain consistency and ensure correct inference, because even minor changes can cascade through the spatio-temporal reasoning process.

2.3. *Ontology Evolution and Change Impact Analysis*

Ontology evolution has long been a topic of interest in the knowledge representation and reasoning (KRR) community (Noy and Klein, 2004), yet effectively managing and reasoning over evolving ontologies is still an open challenge (Polleres et al., 2023). Research in this area covers many topics, such as change detection algorithms (Hartung et al., 2013), drifts (Wang et al., 2011), and change documentation (Khattak et al., 2013). Here, we focus on two key areas most relevant to this work: (1) approaches for measuring the impact of ontological changes, and (2) methods for assigning differential weights to concept relations in an ontology.

2.3.1. *Impact Measurement*

There are many ways to quantify the impact of ontology changes, each reflecting different aspects toward downstream systems. For instance, Abgaz et al. (2012) propose a bottom-up approach to analyze both structural (e.g., adding new entities) and semantic impacts (e.g., relaxing axioms, which can lead to invalid instances) of change operations. Their approach provides what-if analysis by calculating the number of entity changes needed under different implementation strategies, so as to maintain consistency within local ontologies. Dragoni and Ghidini (2012) evaluate the impact through application-level performance, showing different operations (rename, delete, move) have different impacts on the effectiveness of information retrieval systems. More recently, Pernisch et al. (2021) develop impact measures to analyze how ontology changes affect materialization (i.e., the set of precomputed entailments derived from assertions) through experiments on biomedical ontologies with traceable version histories. Although they do not differentiate between TBox and ABox changes, their findings show that changes to class inheritance have the greatest impact on materialization. Other efforts, such as the OCP2KG framework proposed by Conde-Herreros et al. (2024), represent ontology changes (e.g., AddClass, RemoveClass, AddObject-Property) and analyze how they propagate through mapping rules used for the construction of knowledge graphs (KGs) (Hogan et al., 2021). However, their framework tends to support semi-automated data transformation pipelines after ontology updates, rather than directly measuring the impact of conceptual changes during reasoning processes.

2.3.2. *Semantic Significance*

When measuring the impact of changes in an ontology, we also need to consider their semantic significance. To give an intuitive example, changing the sovereignty relation between a city and a state has a greater impact than modifying a city’s property, such as its population. In a different context of facilitating path computation, Fadnis

et al. (2019) also argue that not all semantic relations should be treated equally, and call for assigning weights to edges in KGs such as ConceptNet. El Ghosh et al. (2022) extend this idea by differentiating the weights for hierarchical (e.g., `subClassOf`) versus non-hierarchical (e.g., `partOf`) relations in topic ontologies, so as to overcome the limitation of treating all edges as equivalent in path-based approaches. In this work, we hold a similar assumption that not all ontological relations carry equal structural significance. For example, taxonomic relations like `subClassOf` carry more weight than data property relations such as `hasPopulation`.

With the weighting scheme, changes can more accurately reflect the significance of conceptual connections when assessing the impact of evolving relations in an ontology. Similarly, Chen et al. (2021) develop weighted KG embeddings that assign different importance scores to different entailments when analyzing ontology stream evolution. However, their approach assumes static TBox structures and considers only ABox-level changes, especially focusing on preserving machine learning model performance rather than assessing the conceptual impact.

2.3.3. Challenges in Geospatial Ontologies

Together, the aforementioned studies provide methods for analyzing the impact of ontology changes within an abstract semantic space. They, however, largely overlook the challenges posed by spatio-temporal dynamics, that concepts may evolve across both space and time (i.e., at different rates) (Shi et al., 2025). This limitation is especially important in geographic domains, where changes in conceptual definitions lead to direct consequences in physical space. For example, by manipulating the definition of *forest*, we can either reduce deforestation rates by narrowing it, or help achieve carbon offsetting goals by relaxing the definition (Gupta and Gahegan, 2020). These changes also make temporal reasoning more difficult, since records based on previous definitions may no longer align with the current interpretations.

In geospatial ontologies, space and time are fundamental components, and both are subject to change. When the conceptual schema evolves, it has cascading effects on downstream reasoning processes, such as the computation of spatial relations. This prompts the need for specialized change impact analysis that can handle the unique characteristics of geospatial knowledge representation.

3. A Multi-dimensional Framework for Ontology Change Impact

In this section, we propose a framework to quantify the impact of ontology changes across three interconnected dimensions, grounded in the fundamental spatial (where), temporal (when), and semantic (what) reference systems for geographic information (Kuhn, 2003; Scheider, 2012), as depicted in Figure 1. Our approach evaluates the impact of changes to geographic concepts in ontologies, particularly how they affect reasoning—progressing from purely spatial to spatio-temporal aspects, and finally to the integration of all three dimensions.

3.1. Spatial Impact

Spatial impact captures how ontological changes affect the spatial extent of a geographic concept and its associated spatial relations. Here, we consider two main components: (1) the physical coverage where the concept applies, e.g., which areas




Figure 1.: Reference systems for geographic information proposed by Kuhn (2003).

are included based on the definitions, and (2) spatial relations with other geographic entities and types. For example, if the criteria for distinguishing *cities* from *towns* change—say, by adjusting the minimum population threshold—sets of places may be reclassified, reshaping their spatial boundaries as well as their relations with surrounding areas. These relations are at the core of spatial reasoning, affecting queries and decision-making processes that rely on well-defined topological relations.

3.1.1. Coverage Change

Let S_{before} and S_{after} represent the spatial measure (e.g., area, length, or volume) of a geographic concept before and after ontology change, respectively. The choice of spatial measure depends on the type of concept, e.g., area for regions, length for linear features like roads, and volume for waterbodies. The resulting change in physical coverage can be measured as:

$$\Delta S = S_{\text{after}} - S_{\text{before}} \quad (1)$$

For instance, due to climate change impacts, an ecology ontology decides to update its definition of *tropical rainforest*. Originally, it was defined as regions with an average annual rainfall $\geq 1,680$ mm and is now revised to $\geq 1,850$ mm. The resulting changes led to decreased area (say, $\Delta S = -2.5$ million km^2) classified as tropical rainforest. A normalized version¹ can be expressed as:

$$\Delta S_{\text{norm}} = \frac{S_{\text{after}} - S_{\text{before}}}{S_{\text{before}}} \quad (2)$$

3.1.2. Spatial Relation Change

For the set $R = \{s_1, s_2, \dots, s_n\}$ of spatial relation statements between concepts, let R_{before} and R_{after} denote the subsets of statements before and after the ontology change, respectively. The degree of spatial relation change can be quantified as:

$$\delta R = \frac{|R_{\text{before}} \oplus R_{\text{after}}|}{|R_{\text{before}} \cup R_{\text{after}}|} \quad (3)$$

where $\oplus$ denotes the symmetric difference, which accounts for both statements lost and gained. δR ranges from 0 (no spatial relation change) to 1 (complete spatial relation change). Consider a transportation ontology that introduces new road classifications to support sustainable mobility. For example, two new concepts—*CarpoolLane* and *TransitOnlyLane*—are added to the existing road types, together with new axioms. This addition introduces new spatial relations such as *sfTouches*(*CarpoolLane*, *RegularLane*), meanwhile some previous inferred relations may no longer hold because of the updated classifications. Note that under the open-world assumption (Baader, 2003), this does not mean previous inferences were incorrect; it simply reflects reasoning over the updated ontology.

3.2. Spatio-temporal Impact

Beyond spatial impact, ontology changes can also disrupt temporal reasoning. Historical records that were previously valid for reasoning about a location s may no longer support inference for current and future analysis. Let $N_{\text{affected}}(s)$ represent the number of affected statements related to location s and N_{total} represent the total number of records relevant to the concept. The temporal reasoning validity delta can be measured as:

$$\delta TRV(s) = \frac{N_{\text{affected}}(s)}{N_{\text{total}}} \quad (4)$$

Consider the UK’s withdrawal from the European Union (EU), i.e., Brexit, which formally took effect in early 2020. Following Brexit, the concept of *freedom of movement within EU member states* became invalid for UK-related trips. Historical records remain true but can no longer support inferences about current or future mobility. Suppose a legal database contains 1,000 historical records about immigration within the EU, of which 200 UK-related records will lose their inference ability after Brexit. ΔTRV about the UK becomes $\frac{200}{1000} = 0.2$, indicating that 20% of historical records become “dead-end” data under the updated ontology.

3.3. Integrated Semantic Impact

The integrated semantic impact captures all cascading effects of ontology changes across spatial, temporal, and semantic dimensions via entailment propagation. When a geographic concept is redefined—such as reclassifying an *island* as a *rock*, where the former has to be capable of sustaining human habitation and the latter is not, according to the United Nations (1982). While both islands and rocks are entitled to 12 nautical miles of territorial sea, islands are additionally entitled to 200 nautical miles of exclusive economic zone. Thus, the semantic impact of such reclassification has historically led to disputed state borders and maritime zone claims. The integrated semantic impact takes into account the direct consequences of changing a concept and how such changes affect related entities and inferences.

3.3.1. Entailment Change

Logical entailment—the set of statements that can be logically inferred from existing knowledge—lies at the core of how ontology changes propagate through an information system. The *island* versus *rock* case serves as an important example to show how a

single conceptual change can lead to downstream impacts across multiple dimensions: (1) *spatial*: updating geographic extent of maritime zones, (2) *temporal*: revising legal validity of existing claims, and (3) *semantic*: modifying inferred statements related to sovereignty, habitation and resource rights, etc. Since entailments are derived under the new definition, each updated statement has the potential to trigger additional changes throughout the system. The entailment change (given a certain entailment regime, e.g., OWL 2 RL) delta can be measured as:

$$\Delta E = |E_{\text{after}} \oplus E_{\text{before}}| \quad (5)$$

where E_{before} and E_{after} represent materialized entailment sets before and after the change, respectively, and $\oplus$ represents the symmetric difference. Both newly derivable entailments and those that become invalid are captured.

3.3.2. Structural Weight of Relations

Although the raw count of entailment changes provides a starting point, the actual impact varies across different relation types. Some relations play a more critical role in the structure of a knowledge system. Relations that serve as the backbone of reasoning and inference—such as taxonomic (e.g., `subClassOf` or `isA`) and compositional (e.g., `partOf`) relations—carry greater weight than associative connections (e.g., `relatedTo`). The structural weight of each relation type can be calculated based on its dependency pattern:

$$w_r = \frac{|E_{\text{in}}(r)| \times |E_{\text{out}}(r)|}{|E_{\text{total}}|} \quad (6)$$

where $E_{\text{in}}(r)$ represents the set of entailments that are prerequisites for relation r (i.e., its upstream dependencies), and $E_{\text{out}}(r)$ represents the set of entailments that are inferred from relation r (i.e., its downstream entailments). E_{total} is the complete set of materialization in the knowledge base, serving as a normalization factor. Consider the relation `subClassOf(AutonomousTerritory, PoliticalEntity)` as an example. An intuitive instance-level example of this would be `partOf(Greenland, KingdomOfDenmark)`. If this relation depends on three prerequisite statements on autonomous status, constitutional arrangements, etc., then $|E_{\text{in}}(r)| = 3$. If this relation leads to five downstream inferences related to territorial waters and fishing rights, then $|E_{\text{out}}(r)| = 5$. Although this example is at the instance-level (ABox) to illustrate the calculation, the structural weight is general and can be applied across the entire ontology (TBox) to capture the global significance of relations.

These weights can then be applied to rank different types of entailment changes. Relations with higher structural significance should receive priority attention during expert review and validation processes, compared to descriptive or associative relations. When considered together with raw entailment change counts, these weights provide a more balanced view of the overall semantic impact of changes across the knowledge base.

4. Application

Having established the framework for measuring ontology change impact, in this section, we apply it to real-world cases on maritime territorial law. Maritime boundaries provide an ideal use case here because they have undergone measurable evolution and integrate spatial, temporal, and semantic dimensions. Although the United Nations Convention on the Law of the Sea (UNCLOS) has standardized many aspects over the years, there are still definitional challenges that continue to create disputes (Østhaugen, 2021), such as the difference between a *rock* and an *island* or the classification of artificial structures. Inconsistencies across jurisdictions and over time have led to cascading effects on territorial claims, fishing and resource rights, etc. (Prescott and Schofield, 2005). These inconsistencies also create barriers from an RDM perspective on data integration, e.g., when researchers are working with multiple datasets spanning different time periods. Our change impact analysis provides systematic methods to identify and track the consequences of these changes. Through the following maritime cases, we demonstrate how the proposed framework reveals the implications of evolving definitions on spatio-temporal reasoning.

4.1. Maritime Legal Ontology Design

We construct a simplified maritime legal ontology for this use case to capture both historical evolution and contemporary definitional challenges. The ontology simulates the progression from pre-UNCLOS territorial breadth claims to the current standardized regime, while incorporating ongoing classification disputes. Figure 2 depicts the schema diagram showing the ontology’s structure. We include core maritime classes (e.g., `TerritorialSea`, `ContiguousZone`), geographic features (e.g., `Island`, `Rock`), and legal and spatial relations (e.g., `claimedBy` and standard GeoSPARQL relations) that capture the definitional dependencies central to maritime law.

4.1.1. Historical Claims Evolution

For historical claims, we populate our maritime ontology by retrieving territorial water claims of 53 coastal states from the 1960 UNCLOS summary records (United Nations, 1960).² To prepare these data for analysis, we manually extracted the relevant statements from the original 1960 document, cross-checked them against another meeting record (United Nations, 1968), and standardized states name to ensure consistent historical mapping using the CShapes (Schvitz et al., 2022) dataset. Figure 3 maps how different legal regimes defined territorial water boundaries before the standardization by United Nations (1982).

To capture temporal evolution in our ontology, we use concept reification, which treats statements as objects so that additional information, such as time or provenance, can be attached to a triple. We model three time periods here: (1) the pre-UNCLOS customary law from 1945 to 1982, (2) the transition phase from 1982 to 1994, and (3) the contemporary standardized regime from 1994 to the present. Each temporal layer has its own claims, allowing us to evaluate the consequences of evolving legal definitions. For example, to identify which states claimed 12 nautical miles of territorial



Figure 2.: Schema diagram of the maritime legal ontology. Yellow rectangles represent proposed classes; blue rectangles represent reused classes from existing ontologies; gray rectangles in dashed lines represent subclasses, and yellow ovals indicate datatype properties. The generic `geo:spatialRelation` is used to denote any of the standard GeoSPARQL spatial relations (e.g., `geo:sfContains`, `geo:sfWithin`) for simplicity.

seas during the pre-UNCLOS period, we can query:

```
PREFIX mo: <http://example.org/maritime-ontology#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX qudt: <http://qudt.org/schema/qudt/>
PREFIX unit: <http://qudt.org/vocab/unit/>

SELECT ?stateName
WHERE {
  ?zone rdf:type mo:TerritorialSea ;
        mo:hasExtentClaim ?claim .

  ?claim rdf:type mo:ExtentClaim ;
        mo:hasExtent ?extentValue ;
        mo:temporalPeriod "Pre-UNCLOS" ;
        mo:claimedBy ?state .

  ?extentValue qudt:numericValue "12"^^xsd:integer ;
              qudt:unit unit:MI_N .

  ?state rdfs:label ?stateName .
}
ORDER BY ?stateName
```




Figure 3.: Territorial waters claims by coastal states based on traceable 1960 UNCLOS records (United Nations, 1960). State boundaries are from CShapes (Schvitz et al., 2022), using the latest configuration available for each state prior to or at 1960.

4.1.2. Contemporary Definitional Challenges

For contemporary challenges, we construct a separate dataset focusing particularly on the classification of *rocks* versus *islands* under UNCLOS Article 121 (United Nations, 1982), where the same geographic entity receives different classifications under different national interpretations over time. This captures how the definitional uncertainty of a single concept can propagate through dependent relations, affecting territorial boundaries, exclusive economic zones, fishing rights, resource extraction rights, and so on. Our dataset for this analysis models five uninhabited geographic features, along with all maritime zones associated with their claiming states, as summarized in Table 1. Most of these features are also involved in maritime or territorial boundary disputes, and their areas were manually collected from Wikipedia and converted to hectares (ha) for clearer visualization.

Table 1.: Disputed geographic features in the maritime legal ontology

Geographic Feature	Area (ha)	Claiming States
Okinotorishima	0.0009	Japan
Rockall	0.078	United Kingdom, Ireland
Machias Seal Island	8	Canada, United States
Hans Island	130	Canada, Denmark
Uotsuri/Diaoyu Island	432	Japan, China

We also compute spatial relations between these geographic features and maritime zones. To do so, we use the latest maritime zone geometries of the claiming states listed in Table 1, sourced from Marine Regions,³ and geographic feature geometries from the Global Islands Dataset (Sayre, 2023). To address minor mismatches between datasets, we apply a 250-meter buffer to geographic feature geometries. Using DE-9IM,

we derive spatial relations between geographic features and maritime zones, including reverse relations to support bidirectional queries and reasoning. Examples of the spatial relation statements include:

```
mo:Okinotorishima geo:sfOverlaps mo:Japan_EEZ_102 .
mo:Rockall geo:sfOverlaps mo:United_Kingdom_EEZ_110 .
mo:Machias_Seal_Island geo:sfWithin mo:United_States_EEZ_115 .
mo:Hans_Island geo:sfOverlaps mo:Canada_EEZ_90 .
mo:Uotsuri_Diaoyu_Island geo:sfOverlaps mo:China_EEZ_96, mo:Japan_EEZ_95 .
```

Together, these datasets demonstrate our ontology’s ability to integrate geometric data (using GeoSPARQL), temporal claims (via reification), and computed spatial and legal relations to represent maritime legal development across space and time.

4.2. Impact Assessment

The impact assessment includes two parts. We first examine the standardization of maritime zone definitions established by the UNCLOS in 1982, focusing specifically on the definitional changes to the territorial sea during treaty negotiations. For the second part, we focus on contemporary maritime disputes, involving the four key maritime zones: territorial sea, contiguous zone, exclusive economic zone (EEZ), and continental shelf (see Figure 4). We evaluate the legal implications of classifying disputed geographic features as either rocks or islands, as defined in UNCLOS Article 121, and demonstrate how the changes would affect the spatial relations within maritime zones. Note that these changes pertain to the TBox from a definitional perspective, instead of single-value updates. Together, these two assessments illustrate how our framework evaluates spatial, spatio-temporal, and integrated semantic impacts.




Figure 4.: Standardized maritime zones as defined by United Nations (1982). The four zones of interest are measured in nautical miles (nmi) from the coastal baseline: territorial sea, contiguous zone, exclusive economic zone, and continental shelf (with possible extension).

4.2.1. Maritime Zone Redefinition

To start with, we compute how spatial coverage changed when UNCLOS standardized customary law of territorial seas across 53 coastal states, with previous claims ranging

from 3 to 200 nmi. These changes involve: (1) a revised definition of the territorial sea, i.e., standardized to 12 nmi, and (2) formal classification of three additional maritime zones, as illustrated in Figure 4. From a knowledge representation perspective, this represents a TBox modification where *territorial sea* becomes a subclass of the broader *maritime zone* concept, with standardized boundaries. Using reified extent claims, we query the change in territorial sea extent as a result of this legal standardization.⁴ The following query returns a 16.7% overall coverage increase across the 53 coastal states.

```
SELECT (SUM(?postSum - ?preSum) / SUM(?preSum) * 100 AS ?overallPercentChange)
WHERE {
  {
    SELECT ?state (SUM(?preValue) AS ?preSum)
    WHERE {
      ?preClaim rdf:type mo:ExtentClaim ;
        mo:temporalPeriod "Pre-UNCLOS" ;
        mo:claimedBy ?state ;
        mo:hasExtent ?extentNode .
      ?extentNode qudt:numericValue ?preValue ;
        qudt:unit unit:MI_N .
    } GROUP BY ?state
  }
  {
    SELECT ?state (SUM(?postValue) AS ?postSum)
    WHERE {
      ?postClaim rdf:type mo:ExtentClaim ;
        mo:temporalPeriod "Contemporary" ;
        mo:claimedBy ?state ;
        mo:hasExtent ?extentNode .
      ?extentNode qudt:numericValue ?postValue ;
        qudt:unit unit:MI_N .
    } GROUP BY ?state
  }
}
```

Beyond spatial coverage, we investigate the temporal reasoning implications of the standardization. Historical extent claims may lose their inference validity when the definitions change. Here, we identify two categories of affected claims: (1) pre-UNCLOS claims that differ from the standardized 12 nmi territorial sea, and (2) all transition period claims, which represent legal uncertainty. Although one may argue that all pre-UNCLOS claims fall under different legal frameworks, those with 12 nmi claims remain spatially compatible over time, given their geometric consistency with contemporary boundaries. The following query calculates the proportion of extent claims that lose their temporal reasoning validity. The result reveals that 58.5% of total territorial sea extent claims cannot be reliably used for contemporary maritime analysis.

```
SELECT (SUM(IF(?period = "TransitionPeriod" ||
  (?period = "Pre-UNCLOS" && ?extentValueNumeric != 12), 1, 0)) /
  COUNT(?totalClaim) * 100 AS ?overallTemporalLoss)
WHERE {
  ?terrSea rdf:type mo:TerritorialSea ;
    mo:hasExtentClaim ?totalClaim .

  ?totalClaim rdf:type mo:ExtentClaim ;
    mo:temporalPeriod ?period ;
    mo:hasExtent ?extentValue .

  ?extentValue qudt:numericValue ?extentValueNumeric ;
    qudt:unit unit:MI_N .
}
```

4.2.2. Island or Rock?

We then apply the framework to evaluate potential future changes arising from contemporary disputes over geographic features, as listed in Table 1. These uninhabited features are claimed by different states, and may also be challenged for being classified as *islands* or *rocks*. According to UNCLOS Article 121, an *island* is defined as follows:

$$\text{Island} \sqsubseteq \text{NaturallyFormed} \sqcap \text{SurroundedByWater} \sqcap \text{AboveHighTide}$$

However, the definition of a *rock*—a disadvantaged sub-category of island with fewer maritime entitlements (Carleton and Schofield, 2002)—has more interpretational ambiguity. Notably, it lacks size requirements and is defined through negative criteria:

$$\text{Rock} \sqsubseteq \text{Island} \sqcap (\neg \text{SupportsHabitation} \sqcup \neg \text{HasEconomicActivity})$$

Under Article 121, islands are entitled to generate all four maritime zones, whereas rocks are limited to only the territorial sea and contiguous zone. Here, we investigate how reclassifying these disputed features would affect existing legal and spatial relation statements. We consider a hypothetical amendment to UNCLOS Article 121 (i.e., to include an additional size requirement): features smaller than 100 ha (1 km²) would be classified as rocks regardless of their habitability or economic life. Under this amended definition, rocks would be classified as:

$$\begin{aligned} \text{Rock} \sqsubseteq & \text{Island} \sqcap (\exists \text{hasAreaHa.xsd:double} [< 100] \\ & \sqcup \neg \text{SupportsHabitation} \sqcup \neg \text{HasEconomicActivity}) \end{aligned}$$

It is important to note several spatial representations in our dataset: (1) EEZ polygons extend from the baseline rather than from the outer edge of the territorial sea, unlike the depiction in Figure 4. Therefore, features within the 12 nautical mile limit intersect with both the territorial sea and EEZ polygons; (2) continental shelf zones are represented by EEZ polygons since both extend to 200 nautical miles, with separate polygons only for extended continental shelf claims beyond this limit. We present the results as recorded in the dataset while acknowledging these spatial overlaps.

By default, the five geographical features in our dataset were classified as islands and associated with their entitled maritime zones. Now, three of them have been downgraded to rocks, therefore potentially losing their EEZs. This taxonomic change cascades through 103 ontological statements (6.26% of the total). However, our measurement of ΔE remains incomplete in this experimental setting because the new spatial relations would depend on adjusted maritime boundaries, which cannot be computed automatically. For islands or rocks near the coast, the EEZ boundary may only require minor adjustments. For those farther offshore, the EEZ could be entirely eliminated. Expert review is therefore needed to ensure that the updated boundaries and the associated legal interpretations by the relevant claiming states are correctly applied in the new context. Table 2 provides more details on the measurable portion of the entailment change.⁵ Our reported impact thus represents a lower bound on the overall entailment change, with the addition operations pending expert determination. We observe that spatial relations constitute the majority (83.5%) of affected statements, which highlight both the interconnectedness of maritime boundaries and the ontology itself. Note that these spatial relations include forward (e.g., `mo:Island geo:sfWithin mo:EEZ`) and reverse (e.g., `mo:EEZ geo:sfContains mo:Island`) as-

sertions to support bidirectional reasoning.

Table 2.: Affected statements after reclassifying three features from Table 1 with area smaller than 100 ha as rocks due to the definitional change in our ontology.

Statement Type	Count	%	% of Total	Impact Description
Direct Reclassifications				
rdf:type	3	2.9	0.18	feature type declarations
Primary Spatial Relations (Feature-Zone)				
geo:sfOverlaps	4	3.9	0.24	spatial relation revalidation
geo:sfWithin	1	1.0	0.06	spatial relation revalidation
geo:sfContains	1	1.0	0.06	spatial relation revalidation
Secondary Spatial Relations (Zone-Zone)				
geo:sfOverlaps	64	62.1	3.89	spatial relation revalidation
geo:sfTouches	16	15.5	0.97	spatial relation revalidation
Maritime Zones				
mo:hasMaritimeZone	3	2.9	0.18	state-zone relations
rdf:type	3	2.9	0.18	zone type declarations
geo:hasGeometry	3	2.9	0.18	zone geometries
Sovereignty Claims				
mo:claimedBy	5	4.9	0.30	sovereignty claim revalidation
Total	103	100.0	6.26	

Total statements (n=1645), with rdfs:label annotation statements excluded.

These affected statements require expert review for modification (e.g., redrawing EEZ geometries) or complete removal.

To illustrate the relative impact of different types of relations, we focus on the specific taxonomic relation `rdf:type` that differentiates rocks from islands as the most representative example for computing the structural weight. Its dependencies (E_{in}) include relevant classification criteria (e.g., `mo:supportsHabitation`), while its dependents (E_{out}) include all statements that rely on type classifications, such as spatial relations and maritime zone assertions. Table 3 lists the detailed relations involved. Using Equation 6, the structural weight of `rdf:type` is calculated as $\frac{15 \times 177}{1645} = 1.61$.⁶ In contrast, spatial relations like `geo:sfOverlaps` typically depend only on geometries and lack structural dependencies. This comparison shows that, while spatial relations make up the majority of the affected statements, they are not structurally significant like declarative or taxonomic relations. With these findings, we aim to call for impact assessments to consider not just the quantity of changes, but also their structural importance, to support more informed decisions in ontology maintenance.

In summary, we demonstrate the practical use of our proposed framework through two TBox changes: (1) historical maritime zone redefinition and standardization, and (2) contemporary island vs. rock classification. In the first case, we show how definitional standardization affects spatial and temporal reasoning, revealing a 16.7% increase in territorial sea coverage and a 58.5% loss in temporal reasoning validity for extent claims. In the second case, we show how adding classification criteria cascades through multiple statement types, affecting 6.26% of the ontology statements. Spatial relations dominate in quantity, while taxonomic relations like `rdf:type` carry higher structural significance and should be prioritized for expert review. Together, these

Table 3.: Semantic dependency breakdown for a representative `rdf:type` relation differentiating `mo:Island` and `mo:Rock` in the ontology.

Component	Relations	Count
Dependencies (E_{in})	<code>mo:supportsHabitation</code>	5
	<code>mo:hasEconomicActivity</code>	5
	<code>mo:hasArea</code>	5
	Subtotal	15
Dependents (E_{out})	<code>geo:sfContains</code>	1
	<code>geo:sfEquals</code>	4
	<code>geo:sfOverlaps</code>	116
	<code>geo:sfTouches</code>	34
	<code>geo:sfWithin</code>	1
	<code>mo:hasMaritimeZone</code>	6
	<code>geo:hasGeometry</code>	6
	<code>mo:claimedBy</code>	9
	Subtotal	177
Weight (w_r)		1.61

$|E_{total}| = 1645$ is the total number of statements used for normalization (excluding annotation statements).

real-world cases remind us that the impact of ontology changes (especially geospatial ontology) is spatial, temporal, and semantic in nature, and requires multidimensional assessment approaches. Although our case study focuses on maritime concepts, the framework and the proposed structural weight are generalizable (discussed in more detail in Section 5.2) to other domains where definitional changes and cascading dependencies exist, supporting similar multi-dimensional change impact assessments.

5. Discussion

In this section, we describe how our proposed framework complements existing schema evolution practices and discuss its implications for research data management (RDM) and the maintenance of evolving ontologies. We then discuss the generalizability of the framework, specifically, how it can be applied beyond geospatial ontologies to include non-geographic concepts, where spatio-temporal reasoning also plays a key role.

5.1. Implications for Ontology and Research Data Management

Under the overarching challenge of evolving ontologies, our proposed framework for assessing change impact represents only a first step contributing to more systematic and informed solutions. Traditional approaches (Hartung et al., 2013; Kondylakis et al., 2017) mostly focus on detecting structural differences (i.e., added or removed classes) between ontology versions, and do not explicitly quantify their impact. However, change mapping often requires complex ontology design and alignment decisions, which may compromise data consistency or interoperability. Notably, there is no single “correct” mapping strategy. Multiple mapping approaches can be equally valid, which sometimes depend on domains or intended use (Yamamoto et al., 2023). Before the actual implementation, it is important to assess the consequences of changes, prioritize updates by structural significance, simulate their downstream-reasoning effects, and

conduct informed validation.

Besides our proposed framework and ontologies, emerging data management technologies offer another promising direction. Apache Iceberg (<https://iceberg.apache.org/>) introduces schema evolution capabilities tailored to tables, that allow changes to be made without rewriting historical data. For example, when a new column is added, Iceberg ensures it will not inherit values from any existing columns in the table. Its non-destructive design and partitioning technique provide a robust change-tracking mechanism, therefore improving the maintainability and reliability of evolving data models. Although Apache Iceberg is designed for tabular datasets rather than ontologies (and, thus, its definition of *schema evolution* differs from *ontology evolution*), applying a similar principle to ontology maintenance may offer a more resilient approach for handling ontology evolution.

Together with these emerging techniques, our multi-dimensional framework could contribute to RDM to ensure FAIR data principles (Wilkinson et al., 2016) as classification systems evolve. For example, if natural disaster classifications are expanded to include human-induced causes (e.g., those related to climate change), such updates could affect how past datasets are categorized and retrieved. By quantifying the impact of conceptual changes across space, time, and semantics, our framework helps RDM practitioners to identify which updates may, at a minimum, break existing annotations, invalidate search results, or disrupt interoperability across datasets and systems. Evaluating this impact also helps prioritize maintenance efforts, improve metadata versioning, and track concept drift over time (Pernisch et al., 2021). When combined with further version-aware tools like Apache Iceberg, we can support traceable, non-destructive evolution of research classifications—ensuring that past data remains interpretable and reproducible as ontologies evolve.

5.2. From Geographic to Non-Geographic Concepts

While we initially highlighted the dynamic nature of geographic concepts in Section 2, i.e., their susceptibility to reinterpretation and redefinition over time and across space, many non-geographic concepts have the same fundamental challenges. Therefore, our framework can readily extend beyond geographic domains to measure change impacts across a broader range of concepts that share similar characteristics of spatio-temporal variability. For example, consider the case of transgender athlete classification in sports. The criteria of an *eligible female athlete* have undergone substantial changes: from surgical requirements in 2004 to hormonal threshold standards in 2016, and more recently to sport-specific determinations in 2021 (Harper, 2022). Even within the hormonal threshold era, the International Olympic Committee (2015) required testosterone levels to be ≤ 10 nmol/L, while World Athletics required a stricter limit of ≤ 5 nmol/L in 2019 (Bowman-Smart et al., 2024). These definitional changes do not propagate uniformly (but rather at different rates) across all sports committees or jurisdictions, creating a fragmented conceptual space.

Such cases create substantial challenges for conducting retrospective research, as well as for maintaining semantic consistency and interoperability. The impacts of each definitional change involve multiple dimensions and call for a systematic evaluation. Spatial impacts, for instance, can be quantified by measuring the spatial coverage where different definitions apply at the same time. Temporal impacts become relevant when researchers trace the eligibility of historical performance data, e.g., determining whether records set by a certain athlete before policy changes remain valid or

comparable under new definitions. Spatio-temporal impacts come into play when the temporal validity of classifications becomes location-dependent (i.e., bounded by both location and time). Beyond the technical perspective, these changes can also give rise to real-world ethical consequences, e.g., fairness concerns for the athletes involved. Our multi-dimensional framework provides a structured way to explicitly map these dynamics. When multiple competing definitions exist, it can help identify who is affected, when the changes take effect, where they apply, and under which institutional authority, thereby supporting more responsible and evidence-based decision-making.

Even without explicit spatial relations, these impacts mirror those we observe in geospatial ontologies. The transgender athlete example showcases that definitional changes—even in non-geographic domains—can produce spatially and temporally bounded effects. Moreover, our framework can be extended to systematically handle conflicting information, where researchers can simulate alternative definitional assumptions and assess the extent to which they propagate through knowledge bases and affect downstream reasoning.

6. Conclusions

Whether embedded in human cognition or encoded in formal data management systems, geographic concepts (such as *forest*, *mountain*, and *lake*) hold a unique position. They are constantly evolving across space and time as our understanding of the world changes, be it due to scientific progress or sociocultural factors. The way we define them has real-world consequences. For example, what we define as a *forest* influences extent estimates of afforested areas, deforestation rates, and ultimately, climate policy decisions. In this increasingly interconnected digital age, even a seemingly minor update to an established concept (e.g., in a geospatial ontology) can cascade through entire knowledge systems (e.g., knowledge graphs), affecting data interpretation, interoperability, and eventually decision-making.

To quantify the impact of such changes, we propose a multi-dimensional framework across space, time, and semantics. In this framework, we take into account spatio-temporal impact (e.g., spatial coverage, temporal reasoning validity), which are often overlooked in existing approaches. We also introduce a weighting scheme to distinguish between changes in spatial relations (e.g., `sfOverlaps`) and taxonomic relations (e.g., `isA`) based on their structural significance within the ontology. To apply this framework in a real-world scenario, we construct a maritime legal ontology that captures both the historical standardization of maritime zones under UNCLOS and contemporary disputes over island classifications. Among our findings, we show that the UNCLOS standardization causes 58.5% of total territorial sea extent claims to lose temporal reasoning validity, making them unreliable for contemporary maritime analysis. Additionally, introducing an additional size requirement for the definition of *rock* leads to 6.26% of total statements pending revalidation.

Using a geospatial ontology of maritime legal concepts, our framework can be adapted and broadly applied to other domains where ontologies evolve and spatio-temporal reasoning plays a critical role. With this framework, we aim to provide a systematic way to identify potentially more disruptive updates, simulate their costs, and prioritize remediation efforts—ultimately contributing to more informed and resilient ontology management practices. The implications also extend naturally to research data management. By anticipating the downstream effects of evolving classification systems, we can help ensure data maintains its semantic clarity, interoperability, and

uphold the FAIR principles within spatio-temporal knowledge systems.

Author Contribution Statement

Meilin Shi: Writing – original draft, conceptualization, methodology, data curation, software. **Krzysztof Janowicz:** Writing – review & editing, conceptualization, supervision. **Zilong Liu:** Writing – review & editing. **Mina Karimi:** Writing – review & editing. **Ivan Majic:** Writing – review & editing. **Alexandra Fortacz:** Writing – review & editing. All authors have read and agreed to the published version of the manuscript.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Open access funding provided by Universitat Wien/KEMÖ.

Data Availability Statement

The data and code that support the findings of this study are available on Figshare (<https://figshare.com/s/2a5bb6db30306e66db52>).

Notes

¹When $A_{\text{before}} = 0$, which indicates a newly emerged area, the normalized change is undefined. In such cases, we report only the raw area increase.

²These claims correspond to the geopolitical context of the time, including states such as the Soviet Union and Yugoslavia that later dissolved, and Ceylon, which was later renamed Sri Lanka.

³We use four geometries from Marine Regions: territorial sea, contiguous zone, and EEZ (all dated 2023-10-25), and extended continental shelf (dated 2024-10-10). All data retrieved from <https://www.marinerregions.org/>.

⁴In this calculation, we consider a simplified case and do not account for post-1960 geopolitical changes, such as the split of Yugoslavia or the dissolution of the Soviet Union, using the same 53 states throughout.

⁵Note that the *actual* results depend not only on the ontology but also on the entailment regime selected for a given knowledge graph and triple store, e.g., the generated inferences depend on whether GeoSPARQL rules are enabled, whether OWL2 RL or RDFS is selected as profile, and so on.

⁶Our structural weight calculation assumes a fully materialized ontology, where all dependencies can be observed. In an incomplete ontology, the weights reflect only the existing entailments, thus providing a partial view of the structural significance.

References

Abgaz, Y., Javed, M., and Pahl, C. (2012). Analyzing impacts of change operations in evolving ontologies.

- Adams, B. and Raubal, M. (2009). A metric conceptual space algebra. In *International Conference on Spatial Information Theory*, pages 51–68. Springer.
- Amorim-Maia, A. T., Anguelovski, I., Connolly, J., and Chu, E. (2023). Seeking refuge? the potential of urban climate shelters to address intersecting vulnerabilities. *Landscape and Urban Planning*, 238:104836.
- Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. (2007). DBpedia: A nucleus for a web of open data. In *The Semantic Web*, pages 722–735. Springer.
- Baader, F. (2003). *The description logic handbook: Theory, implementation and applications*. Cambridge university press.
- Bennett, B. (2001). What is a forest? on the vagueness of certain geographic concepts. *Topoi*, 20(2):189–201.
- Bennett, B., Mallenby, D., and Third, A. (2008). An ontology for grounding vague geographic terms. In *Proceedings of the 2008 Conference on Formal Ontology in Information Systems: Proceedings of the Fifth International Conference (FOIS 2008)*, pages 280–293. IOS Press.
- Bowman-Smart, H., Savulescu, J., O’Connell, M., and Sinclair, A. (2024). World athletics regulations unfairly affect female athletes with differences in sex development. *Journal of the Philosophy of Sport*, 51(1):29–53.
- Camara, G. (2020). On the semantics of big earth observation data for land classification. *Journal of Spatial Information Science*, 2020(20):21–34.
- Carleton, C. and Schofield, C. (2002). Developments in the technical determination of maritime space: Delimitation, dispute resolution, geographical information systems and the role of the technical expert. *Maritime Briefings*, 3(4).
- Chen, J., Lecue, F., Pan, J. Z., Deng, S., and Chen, H. (2021). Knowledge graph embeddings for dealing with concept drift in machine learning. *Journal of Web Semantics*, 67:100625.
- Claramunt, C. (2020). Ontologies for geospatial information: Progress and challenges ahead. *Journal of Spatial Information Science*, 20:35–41.
- Clementini, E., Sharma, J., and Egenhofer, M. J. (1994). Modelling topological spatial relations: Strategies for query processing. *Computers & graphics*, 18(6):815–822.
- Conde-Herreros, D., Stork, L., Pernisch, R., Poveda-Villalón, M., Corcho, O., and Chaves-Fraga, D. (2024). Propagating ontology changes to declarative mappings in construction of knowledge graphs. In *Proceedings of the 5th International Workshop on Knowledge Graph Construction co-located with 21th Extended Semantic Web Conference*, pages 1–16. CEUR.
- de Salvo Braz, R., Girju, R., Punyakanok, V., Roth, D., and Sammons, M. (2005). An inference model for semantic entailment in natural language. In *Machine learning challenges workshop*, pages 261–286. Springer.
- Dragoni, M. and Ghidini, C. (2012). Evaluating the impact of ontology evolution patterns on the effectiveness of resources retrieval. In *EvoDyn@ ISWC*. CEUR-WS.org.
- El Ghosh, M., Delestre, N., Kotowicz, J.-P., Zanni-Merk, C., and Abdulrab, H. (2022). Rel topic: A graph-based semantic relatedness measure in topic ontologies and its applicability for topic labeling of old press articles. *Semantic Web*, 14(2):293–321.
- Fadnis, K., Talamadupula, K., Kapanipathi, P., Ishfaq, H., Roukos, S., and Fokoue, A. (2019). Path-based contextualization of knowledge graphs for textual entailment. *arXiv preprint arXiv:1911.02085*.
- Fonseca, F., Câmara, G., and Miguel Monteiro, A. (2006). A framework for measuring the interoperability of geo-ontologies. *Spatial Cognition and Computation*, 6(4):309–331.
- Galton, A. (2003). Desiderata for a spatio-temporal geo-ontology. In *International Conference on Spatial Information Theory*, pages 1–12. Springer.
- Goodchild, M. F., Anselin, L., Appelbaum, R. P., and Harthorn, B. H. (2000). Toward spatially integrated social science. *International regional science review*, 23(2):139–159.
- Grenon, P. and Smith, B. (2004). SNAP and SPAN: Towards dynamic spatial ontology. *Spatial cognition and computation*, 4(1):69–104.
- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge acquisition*, 5(2):199–220.
- Guarino, N. (1998). *Formal ontology in information systems: Proceedings of the first interna-*

- ational conference (FOIS'98), June 6-8, Trento, Italy, volume 46. IOS press.
- Gupta, P. and Gahegan, M. (2020). Categories are in flux, but their computational representations are fixed: That's a problem. *Transactions in GIS*, 24(2):291-314.
- Harper, J. (2022). Transgender athletes and international sports policy. *Law & Contemp. Probs.*, 85:151.
- Hartung, M., Groß, A., and Rahm, E. (2013). Conto-diff: generation of complex evolution mappings for life science ontologies. *Journal of biomedical informatics*, 46(1):15-32.
- Harvey, F., Kuhn, W., Pundt, H., Bishr, Y., and Riedemann, C. (1999). Semantic interoperability: A central issue for sharing geographic information. *The annals of regional science*, 33(2):213-232.
- Hoffart, J., Suchanek, F. M., Berberich, K., and Weikum, G. (2013). YAGO2: A spatially and temporally enhanced knowledge base from wikipedia. *Artificial intelligence*, 194:28-61.
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G. D., Gutierrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., et al. (2021). Knowledge graphs. *ACM Computing Surveys (CSUR)*, 54(4):1-37.
- International Olympic Committee (2015). Ioc consensus meeting on sex reassignment and hyperandrogenism. Online. https://stillmed.olympic.org/Documents/Commissions_PDFfiles/Medical_commission/2015-11_ioc_consensus_meeting_on_sex_reassignment_and_hyperandrogenism-en.pdf.
- Janowicz, K. (2010). The role of space and time for knowledge organization on the semantic web. *Semantic Web*, 1(1-2):25-32.
- Karalis, N., Mandilaras, G., and Koubarakis, M. (2019). Extending the YAGO2 knowledge graph with precise geospatial knowledge. In *The Semantic Web-ISWC 2019: 18th International Semantic Web Conference*, pages 181-197. Springer.
- Kauppinen, T. and Hyvönen, E. (2007). Modeling and reasoning about changes in ontology time series. In *Ontologies: A handbook of principles, concepts and applications in information systems*, pages 319-338. Springer.
- Kefler, C. and Farmer, C. J. (2015). Querying and integrating spatial-temporal information on the web of data via time geography. *Journal of Web Semantics*, 35:25-34.
- Khattak, A. M., Latif, K., and Lee, S. (2013). Change management in evolving web ontologies. *Knowledge-Based Systems*, 37:1-18.
- Kondylakis, H., Despoina, M., Glykokokalos, G., Kalykakis, E., Karapiperakis, M., Lasithiotakis, M.-A., Makridis, J., Moraitis, P., Panteri, A., Plevraki, M., et al. (2017). EvoRDF: A framework for exploring ontology evolution. In *The Semantic Web: ESWC 2017 Satellite Events*, pages 104-108. Springer.
- Kong, W.-K., Liu, X., Racharak, T., Sun, G., Ma, Q., and Nguyen, L.-M. (2025). Weight-aware tasks for evaluating knowledge graph embeddings. *Knowledge-Based Systems*, 320:113596.
- Kuhn, W. (2003). Semantic reference systems. *International Journal of Geographical Information Science*, 17(5):405-409.
- Kuhn, W. (2005). Geospatial semantics: why, of what, and how? In *Journal on data semantics III*, pages 1-24. Springer.
- Kuhn, W., Kauppinen, T., and Janowicz, K. (2014). Linked data-a paradigm shift for geographic information science. In *Geographic Information Science: 8th International Conference, GIScience 2014, Vienna, Austria, September 24-26, 2014. Proceedings 8*, pages 173-186. Springer.
- Ludašcher, B., Lin, K., Bowers, S., Jaeger-Frank, E., Brodaric, B., and Baru, C. (2006). Managing scientific data: From data integration to scientific workflows. In *Geoinformatics: Data to Knowledge*, pages 1-16. Geological Society of America.
- Mark, D. M., Smith, B., and Tversky, B. (1999). Ontology and geographic objects: An empirical study of cognitive categorization. In *Conference On Spatial Information Theory*, pages 283-298.
- Noy, N. F. and Klein, M. (2004). Ontology evolution: Not the same as schema evolution. *Knowledge and information systems*, 6:428-440.
- Østhagen, A. (2021). Troubled seas? the changing politics of maritime boundary disputes.

- Ocean & Coastal Management*, 205:105535.
- Pernisch, R., Dell’Aglío, D., and Bernstein, A. (2021). Beware of the hierarchy—an analysis of ontology evolution and the materialisation impact for biomedical ontologies. *Journal of Web Semantics*, 70:100658.
- Pike, W. and Gahegan, M. (2007). Beyond ontologies: Toward situated representations of scientific knowledge. *International Journal of Human-Computer Studies*, 65(7):674–688.
- Polleres, A., Pernisch, R., Bonifati, A., Dell’Aglío, D., Dobriy, D., Dumbrava, S., Etcheverry, L., Ferranti, N., Hose, K., Jiménez-Ruiz, E., et al. (2023). How does knowledge evolve in open knowledge graphs? *Transactions on Graph Data and Knowledge*, 1(1):11–1.
- Potthast, M., Stein, B., and Gerling, R. (2008). Automatic vandalism detection in wikipedia. In *European conference on information retrieval*, pages 663–668. Springer.
- Prescott, V. and Schofield, C. (2005). *The maritime political boundaries of the world*. Brill.
- Randell, D. A., Cui, Z., and Cohn, A. G. (1992). A spatial logic based on regions and connection. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, KR’92, page 165–176. Morgan Kaufmann Publishers Inc.
- Sayre, R. (2023). Global islands. USGS data release.
- Scheider, S. (2012). *Grounding geographic information in perceptual operations*, volume 244. IOS Press.
- Schlieder, C. (2010). Digital heritage: Semantic challenges of long-term preservation. *Semantic Web*, 1(1-2):143–147.
- Schvitz, G., Rüegger, S., Girardin, L., Cederman, L.-E., Weidmann, N., and Gleditsch, K. S. (2022). Mapping The International System, 1886-2017: The CShapes 2.0 Dataset. *Journal of Conflict Resolution*, 66(1):144–161.
- Schwering, A. and Raubal, M. (2005). Measuring semantic similarity between geospatial conceptual regions. In *International Conference on GeoSpatial Semantics*, pages 90–106. Springer.
- Shi, M., Janowicz, K., Liu, Z., Karimi, M., Majic, I., and Fortacz, A. (2025). Defining concept drift and its variants in research data management: A scientometric case study on geographic information science. *Transactions in GIS*, 29(3):e70058.
- United Nations (1960). Official records: Summary records of plenary meetings and of meetings of the committee of the whole, annexes and final act. https://digitallibrary.un.org/record/857984/files/A_CONF_19_8.pdf.
- United Nations (1968). Survey of national legislation concerning the sea-bed and the ocean floor, and the subsoil thereof, underlying the high seas beyond the limits of present national jurisdiction. <https://digitallibrary.un.org/record/1656535>.
- United Nations (1982). United Nations Convention on the Law of the Sea (UNCLOS), Article 121(3). https://www.un.org/depts/los/convention_agreements/texts/unclos/unclos_e.pdf.
- U.S. Department of Health and Human Services (2025). HHS takes action on President Trump’s executive orders defending women and children. <https://www.hhs.gov/pres-s-room/eo-defending-women-and-children.html>. Accessed: 2025-05-26.
- Wang, S., Schlobach, S., and Klein, M. (2011). Concept drift and how to identify it. *Journal of Web Semantics*, 9(3):247–265.
- White House (2025). Executive order 14173: Defending women from gender ideology extremism and restoring biological truth to the federal government. <https://www.whitehouse.gov/presidential-actions/2025/01/defending-women-from-gender-ideology-extremism-and-restoring-biological-truth-to-the-federal-government/>.
- Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., et al. (2016). The FAIR guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9.
- Yamamoto, V. E., Destro, J. M., and Dos Reis, J. C. (2023). Context-based refinement of mappings in evolving life science ontologies. *Journal of biomedical semantics*, 14(1):16.
- Zablith, F., Antoniou, G., d’Aquin, M., Flouris, G., Kondylakis, H., Motta, E., Plexousakis, D., and Sabou, M. (2015). Ontology evolution: a process-centric survey. *The knowledge*

5 Quantifying Spatial, Temporal, and Semantic Consequences of Ontology Changes

engineering review, 30(1):45–75.

Zalles, V., Harris, N., Stolle, F., and Hansen, M. C. (2024). Forest definitions require a re-think. *Communications Earth & Environment*, 5(1):620.

6 Conclusions and Future Work

This chapter concludes the dissertation and points out directions for future work. It starts with a chapter-by-chapter summary, where I revisit the motivation and research questions, and follows with how each chapter contributes to the overarching research goal. I then present both the theoretical and practical contributions of this dissertation. Finally, I acknowledge the limitations of this work and suggest corresponding research directions that are worth exploring in the future.

6.1 Summary

In 2006, the International Astronomical Union redefined what qualifies as a *planet*, adding the requirement that a planet must have cleared its orbital neighborhood [22]. Pluto, known as the ninth planet for 76 years, was reclassified as a *dwarf planet* under the new definition. Notably, cases like this are far from isolated. Concepts are constantly evolving because of scientific, technological, or social and cultural development, and humans are also constantly updating their mental models [2, 19, 39]. However, these evolving concepts pose challenges for knowledge representation, especially in research data management (RDM), where they directly impact data retrieval, integration, and interoperability, to name a few. Building on this motivation, this dissertation aims to develop methods that can document, detect, and assess conceptual changes, thereby supporting more effective management of evolving knowledge in RDM systems. This goal is examined through three research questions, each addressed in a dedicated chapter of this dissertation.

6 Conclusions and Future Work

Chapter 3 was centered on the definition of concept drift and its related variants. Before proposing the definitions, I conducted a literature review focusing on how structured (i.e., knowledge graph) and unstructured (i.e., natural language processing) approaches have been used in concept drift research, and explained the rationale for choosing the unstructured view for its detection. Within the RDM context, I then proposed definitions for a series of conceptual changes, including concept *birth*, *drift*, *shift*, *split*, *merge*, and *retirement*. Using a scientometric dataset of GIScience publications, each type of conceptual change was detected by measuring semantic similarity between keywords over a chosen time interval. The results showed that keywords can serve as practical signifiers of concepts. Representative examples are provided to illustrate how concepts in GIScience have undergone changes according to the proposed definitions. This chapter lays the groundwork to address the challenge of non-standardized terminology of concept drift as well as inconsistent definitions. The proposed definitions, operationalized in an RDM setting, also highlighted the necessity to monitor and document different types of conceptual changes to effectively manage evolving concepts.

Chapter 4 extended previous definitions of concept drift (i.e., the meaning of a concept changes over time) to include a spatial dimension. More specifically, I proposed the notion of spatio-temporal concept drift, which jointly considers space and time in conceptual changes. This chapter examined concept representation from a technical perspective using word embeddings and discussed how space and time can be incorporated into these representations. The same dataset from the previous chapter was used for the case study, but only those with locations mentioned in their abstracts were selected. To measure semantic similarity across both space and time, context was aggregated for each space-time combination of concepts. The results revealed different semantic trajectories for GIScience concepts and identified varying drifting patterns across countries, particularly for spatially grounded concepts such as *urban planning*. This chapter, as an extension of the definitions proposed in Chapter 3, provides a first exploration of concept

drift across both space and time. The proposed detection method for jointly integrating space and time into embeddings also demonstrated the ability to capture such dynamics. This chapter emphasizes that FAIR-based RDM systems should require spatio-temporal awareness for better knowledge organization.

Compared to the previous two chapters, Chapter 5 took a different direction by adopting a structured view of concept representation. The previous two chapters focused on the definitions and detection of conceptual changes, which required more flexibility in concept representation. This chapter examined what would happen within an existing knowledge system when conceptual changes occur and how to quantify their impact. This chapter proposed a multi-dimensional framework for evaluating the consequences of conceptual changes, more specifically, across spatial, temporal, and semantic dimensions. Here, I argued that for geographic concepts, their inherent spatio-temporal features mean that changes in their meanings can lead to tangible consequences in the real world. To apply this framework in a practical setting, I designed a geospatial ontology that focused on maritime boundaries and their associated legal entitlements. This ontology was populated with both geometries and historical maritime claim statements to enable spatio-temporal reasoning. The results showed that schema-level changes in a concept can affect all three dimensions. A ranking of different types of changes based on their structural significance was also developed to support better decision-making, more specifically, in identifying changes with cascading effects that demand more attention.

Together, the three chapters contribute to the overarching research goal: *supporting more effective management of evolving knowledge organization in RDM systems*. Chapter 3 sets the groundwork by providing definitions of concept drift and its related variants, which are required for systematically monitoring and documenting the different types of conceptual changes in RDM. Chapter 4 extends the notion of concept drift to include both spatial and temporal perspectives. This is particularly important for geographic concepts, because they most likely evolve across both space and time. Finally, Chapter 5 focuses on

6 Conclusions and Future Work

evaluating the consequences of conceptual changes within structured knowledge systems, which allows for the anticipation of the extent of their impact and supports the design of future systems to handle more spatio-temporally aware consequences. The three chapters present a progressive line of work: first defining and detecting conceptual changes, then adding spatio-temporal awareness, and finally assessing their impact to support better RDM practices for managing evolving knowledge.

6.2 Research Contributions

This dissertation contributes to the management of evolving concepts in knowledge representation. In RDM systems, these conceptual changes can potentially hinder FAIR data practices such as retrieval and integration tasks. Given the increasing volume and complexity of data and knowledge nowadays, addressing such challenges becomes an imminent task. This dissertation takes approaches ranging from unstructured text analysis and spatio-temporal modeling to structured ontology reasoning. The core contribution of this dissertation lies in proposing definitions and developing methods to detect and assess conceptual changes in RDM systems. The overall goal is broken down into three research questions, each building progressively toward more robust and spatio-temporally aware knowledge management practices. The following subsections outline the detailed contributions in theory and practice, respectively.

6.2.1 Theoretical Contributions

This dissertation makes several theoretical contributions to the study of concept drift in the broader field of knowledge representation. The main contributions are summarized in the following paragraphs.

Toward more formalized definitions of concept drift and its related variants. Concept drift has been studied across many domains, such as machine learning, linguistics, and

cognitive science. In the Semantic Web context, it can be understood as changes in the meaning of concepts over time [48]. However, its definitions remain inconsistent, and related phenomena lack a clear terminology. This dissertation contributes toward a more formalized understanding of concept drift and proposes definitions for five other related variants of conceptual changes: birth, shift, split, merge, and retirement. These definitions lay a conceptual foundation for systematically documenting and monitoring conceptual changes, with a specific focus on enabling more temporally aware knowledge representation and management in RDM systems.

A definition of spatio-temporal concept drift. Studied only along the dimension of time in previous research, this dissertation extends the definition of concept drift to include both spatial and temporal dimensions. Geographic concepts, among others, evolve not only over time but also across regions, and capturing this dual dynamic is important for semantically richer knowledge representation. This dissertation introduces the definition of spatio-temporal concept drift, defined as changes in the meaning of concepts that vary across both space and time. By considering both dimensions in concept representations, this definition complements the current understanding of how concepts evolve, while also contributing to more spatio-temporally aware approaches to knowledge organization in RDM systems.

A framework for multi-dimensional assessment of conceptual changes. To illustrate the significance of conceptual changes, this dissertation proposes a multi-dimensional framework to evaluate their impact on structured knowledge systems. This framework considers not only the impact on the statements but also the consequences that may extend to the physical spatio-temporal world. Having also introduced a weight that reflects the structural significance of different types of changes (e.g., renaming, adding defining properties, or removing a subclass), it provides a principled basis for ranking the priorities and managing changes in knowledge systems. This framework supports

6 Conclusions and Future Work

the design of knowledge management practices that are sensitive to the cascading effects of conceptual changes and contribute to more informed decisions on modeling evolving concepts.

6.2.2 Practical Contributions

The main practical contributions of this dissertation focus on RDM systems; more detailed aspects are summarized as follows.

Semantic search for evolving concepts. This dissertation contributes to improving semantic search in RDM systems by explicitly accounting for concept drift. Because concepts evolve over time, the keywords used to describe data also need to be updated accordingly [39]. Traditional keyword-based searches are therefore insufficient to retrieve relevant data or literature, highlighting the need for more dynamic cataloging and semantic organization. By incorporating temporal and, where applicable, spatial dynamics into concept representations, the definitions proposed and detection methods developed in this dissertation enable semantic search to capture more underlying evolving meanings. The main contribution is to increase the precision and recall of search results in RDM systems, thereby supporting FAIR data principles.

Spatio-temporal knowledge organization. This dissertation calls for the practical importance of capturing the evolution of concepts across both space and time in RDM systems. By embedding spatial and temporal information directly into concept representations, e.g., linking concepts and related metadata to regions and time periods, research data can be organized to reflect what, where, and when concepts apply. In this way, it enables more accurate querying, retrieval, and integration of datasets, especially for geographic and longitudinal data. This dissertation, more concretely, contributes to RDM by advocating for the explicit consideration of space and time in knowledge organization, enabling systematic monitoring and management of conceptual changes.

Impact-aware ontology management. Besides directly contributing to RDM systems, this dissertation contributes to practical ontology management by proposing a framework to evaluate the consequences of conceptual changes across space, time, and semantics. This framework supports more informed decision-making when updating schema-level changes in structured knowledge systems (e.g., knowledge graphs). In addition to quantifying affected statements, it also evaluates the structural significance of changes, which provides guidance on which modifications have the most critical impact under different ontology modeling choices. These estimates help identify changes with cascading effects and also inform the prioritization of updates. Overall, it contributes to more impact-aware ontology/research data management practices when handling conceptual changes in structured systems.

6.3 Limitations and Future Work

In this section, I discuss the limitations of this dissertation and outline directions for future research.

The choice of concept identity. The question of what constitutes the *identity* of a concept has been discussed extensively in prior work on concept drift [48, 43, 47], and is also brought up in philosophical debate [18]. Among them, Wang et al. [48] distinguish between (1) identity-based drift, where concept identity is determined through explicit mapping by domain experts, and (2) morphing-based drift, which rejects a fixed identity and views concepts at different time points as distinct. In this dissertation, within the RDM context, I treat publication keywords as indicators of underlying concepts and use their labels as operational markers of a concept’s identity. In Chapter 3, I argued that the identity of a concept depends on how it is operationalized, e.g., via top-down or bottom-up approaches. Using the bottom-up (i.e., unstructured) view of concepts allows them to emerge from data patterns without a predefined identity. The choice of

6 Conclusions and Future Work

concept identity in this dissertation remains a pragmatic one, which cannot help resolve the philosophical question of what it means for a concept to remain “the same” over time. Future work could benefit from more principled methods to formalize and implement concept identity, especially for bottom-up, natural language-based approaches.

Spatio-temporal scoping of text. To associate concepts with space and time for concept drift detection, in Chapter 4, I extracted spatial information from publication abstracts through geoparsing and approximated temporal information using publication year. However, these proxies, relying solely on abstracts and limited metadata, may not capture the actual study period or geographic focus of a publication. Additionally, the extracted spatial and temporal references were not systematically evaluated, given that there is no ground truth. Future work could be improved by extracting richer spatio-temporal information directly from full texts, as well as by developing methods to differentiate incidental location mentions from the actual study areas. Establishing evaluation frameworks, such as sampling-based validation, would further improve the accuracy and reliability of spatio-temporal scoping for concept drift detection.

Human-in-the-loop knowledge management. Across all three research questions and experiment designs in this dissertation, I aimed to minimize human intervention in knowledge management processes. However, fully eliminating human input is not yet possible and may also be undesirable in the long term. There are still many tasks that require expert judgement, such as deciding the (dis)similarity threshold of two concepts to be considered *born* or *retired*, modeling conflicting or controversial statements such as maritime boundary disputes, or reviewing pending modeling decisions. Future work could focus on improving human-in-the-loop workflows, for example, by developing semi-automated validation pipelines in RDM systems that allow domain experts to review changes and make design choices. Addressing these challenges in the future would make knowledge management systems more capable of systematically monitoring and handling

conceptual changes.

Ontology learning with large language models. Looking toward a long-term direction of managing conceptual changes in RDM systems, a promising research area is ontology learning (i.e., extracting high-level concepts, relations, and sometimes axioms from unstructured text to form an ontology [50]). Recent work [31, 1, 40] has demonstrated that large language models (LLMs) can be used to support ontology engineering tasks (e.g., taxonomy discovery, relation extraction) across different knowledge domains. I see this as a process to effectively turn the unstructured view of concepts into a structured view, combining the flexibility and rich context of natural language with the reasoning capabilities of structured knowledge. With LLM-powered ontology learning techniques, when new datasets or publications become available, RDM systems could adapt the underlying ontology (e.g., by adding or removing classes) on a regular basis in response to detected conceptual changes. I believe such adaptive ontologies would improve search and retrieval, thereby contributing to the FAIR principles.

Bibliography

- [1] Hamed Babaei Giglou, Jennifer D'Souza, and Sören Auer. LLMs4OL: Large language models for ontology learning. In *International Semantic Web Conference*, pages 408–427. Springer, 2023. doi: 10.1007/978-3-031-47240-4_22.
- [2] Lawrence Barsalou. Concepts and meaning. In L. Barsalou, W. Yeh, B. Luka, K. Olseth, K. Mix, and L. Wu, editors, *Chicago Linguistic Society 29: Papers From the Parasession on Conceptual Representations*, pages 23–61. University of Chicago, 1993.
- [3] Brandon Bennett. What is a forest? on the vagueness of certain geographic concepts. *Topoi*, 20(2):189–201, 2001. doi: 10.1023/A:1017965025666.
- [4] Brandon Bennett, David Mallenby, and Allan Third. An ontology for grounding vague geographic terms. In *Proceedings of the 2008 Conference on Formal Ontology in Information Systems: Proceedings of the Fifth International Conference (FOIS 2008)*, pages 280–293. IOS Press, 2008. doi: 10.3233/978-1-58603-923-3-280.
- [5] Boyan Brodaric and Mark Gahegan. Experiments to examine the situated nature of geoscientific concepts. *Spatial Cognition & Computation*, 7(1):61–95, 2007. doi: 10.1080/13875860701337934.
- [6] Gilberto Camara. On the semantics of big earth observation data for land classification. *Journal of Spatial Information Science*, 2020(20):21–34, 2020. doi: 10.5311/JOSIS.2020.20.645.

Bibliography

- [7] Giuseppe Capobianco, Danilo Cavaliere, Sabrina Senatore, et al. OntoDrift: a semantic drift gauge for ontology evolution monitoring. In *CEUR Workshop Proceedings*, volume 2821, pages 1–10. CEUR-WS, 2020.
- [8] Robin L Chazdon, Pedro HS Brancalion, Lars Laestadius, Aoife Bennett-Curry, Kathleen Buckingham, Chetan Kumar, Julian Moll-Rocek, Ima Célia Guimarães Vieira, and Sarah Jane Wilson. When is a forest a forest? forest concepts and definitions in the era of forest and landscape restoration. *Ambio*, 45(5):538–550, 2016. doi: 10.1007/s13280-016-0772-y.
- [9] Jiaoyan Chen, Freddy Lecue, Jeff Z Pan, Shumin Deng, and Huajun Chen. Knowledge graph embeddings for dealing with concept drift in machine learning. *Journal of Web Semantics*, 67:100625, 2021. doi: 10.1016/j.websem.2020.100625.
- [10] Ksenia Chmutina and Jason Von Meding. A dilemma of language: “natural disasters” in academic literature. *International Journal of Disaster Risk Science*, 10:283–292, 2019. doi: 10.1007/s13753-019-00232-2.
- [11] Christophe Claramunt. Ontologies for geospatial information: Progress and challenges ahead. *Journal of Spatial Information Science*, 20:35–41, 2020. doi: 10.5311/JOSIS.2020.20.666.
- [12] Helen Couclelis. Ontologies of geographic information. *International Journal of Geographical Information Science*, 24(12):1785–1809, 2010. doi: 10.1080/13658816.2010.484392.
- [13] Mauro Dragoni and Chiara Ghidini. Evaluating the impact of ontology evolution patterns on the effectiveness of resources retrieval. In *EvoDyn@ ISWC*. CEUR-WS.org, 2012.
- [14] Max J Egenhofer and David M Mark. Naive geography. In *Spatial Information Theory A Theoretical Basis for GIS: International Conference COSIT’95 Semmering*,

- Austria, September 21–23, 1995 Proceedings 2*, pages 1–15. Springer, 1995. doi: 10.1007/3-540-60392-1_1.
- [15] Antske Fokkens, Serge Ter Braake, Isa Maks, Davide Ceolin, et al. On the semantics of concept drift: Towards formal definitions of semantic change. *Drift-a-LOD@EKAW*, 2016.
- [16] Pantea Foroudi, Reza Marvi, Maria Teresa Cuomo, and Antonio D’Amato. Sustainable development goals in a regional context: conceptualising, measuring and managing residents’ perceptions. *Regional Studies*, 59(1):2373871, 2025. doi: 10.1080/00343404.2024.2373871.
- [17] Hongyu Gong, S. Bhat, and Pramod Viswanath. Enriching word embeddings with temporal and spatial information. In *Conference on Computational Natural Language Learning*, 2020. doi: 10.18653/v1/2020.conll-1.1.
- [18] Nicola Guarino and Christopher Welty. A formal ontology of properties. In *International Conference on Knowledge Engineering and Knowledge Management*, pages 97–112. Springer, 2000. doi: 10.1007/3-540-39967-4_8.
- [19] Prashant Gupta and Mark Gahegan. Categories are in flux, but their computational representations are fixed: That’s a problem. *Transactions in GIS*, 24:291–314, 2020. doi: 10.1111/tgis.12602.
- [20] William L Hamilton, Jure Leskovec, and Dan Jurafsky. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the conference on empirical methods in natural language processing. Conference on empirical methods in natural language processing*, volume 2016, page 2116. NIH Public Access, 2016. doi: 10.18653/v1/D16-1229.
- [21] Michael Hartung, Anika Groß, and Erhard Rahm. Conto–diff: generation of complex

Bibliography

- evolution mappings for life science ontologies. *Journal of biomedical informatics*, 46(1):15–32, 2013. doi: 10.1016/j.jbi.2012.04.009.
- [22] IAU (International Astronomical Union). Pluto and the Developing Landscape of Our Solar System. <https://iauarchive.eso.org/public/themes/pluto/>, 2006. Accessed: 10 September 2025.
- [23] Krzysztof Janowicz. The role of space and time for knowledge organization on the semantic web. *Semantic Web*, 1(1-2):25–32, 2010. doi: 10.3233/SW-2010-0001.
- [24] Krzysztof Janowicz and Pascal Hitzler. The digital earth as knowledge engine. *Semantic Web*, 3(3):213–221, 2012. doi: 10.3233/SW-2012-0070.
- [25] Krzysztof Janowicz, Song Gao, Grant McKenzie, Yingjie Hu, and Budhendra Bhaduri. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, 34(4):625–636, 2020. doi: 10.1080/13658816.2019.1684500.
- [26] Asad Masood Khattak, Khalid Latif, and Sungyoung Lee. Change management in evolving web ontologies. *Knowledge-Based Systems*, 37:1–18, 2013. doi: 10.1016/j.knsys.2012.05.005.
- [27] Ruiyuan Li, Pin Tian, and Shenghui Wang. Study concept drift in 150-year english literature. In *AI@ iConference*, pages 153–163, 2021.
- [28] Anjin Liu, Yiliao Song, Guangquan Zhang, and Jie Lu. Regional concept drift detection and density synchronized drift adaptation. In *IJCAI International Joint Conference on Artificial Intelligence*, 2017. doi: 10.24963/ijcai.2017/317.
- [29] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. Learning under concept drift: A review. *IEEE transactions on knowledge and data engineering*, 31(12):2346–2363, 2018. doi: 10.1109/TKDE.2018.2876857.

- [30] Xin Lv, Lei Hou, Juanzi Li, and Zhiyuan Liu. Differentiating concepts and instances for knowledge graph embedding. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1971–1979. Association for Computational Linguistics, 2018. doi: 10.18653/v1/D18-1222.
- [31] Huu Tan Mai, Cuong Xuan Chu, and Heiko Paulheim. Do LLMs really adapt to domains? an ontology learning perspective. In *International Semantic Web Conference*, pages 126–143. Springer, 2024. doi: 10.1007/978-3-031-77844-5_7.
- [32] David M Mark, Barry Smith, and Barbara Tversky. Ontology and geographic objects: An empirical study of cognitive categorization. In *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science: International Conference COSIT’99 Stade, Germany, August 25–29, 1999 Proceedings 4*, pages 283–298. Springer, 1999. doi: 10.1007/3-540-48384-5_19.
- [33] Romana Pernisch, Daniele Dell’Aglío, and Abraham Bernstein. Beware of the hierarchy—an analysis of ontology evolution and the materialisation impact for biomedical ontologies. *Journal of Web Semantics*, 70:100658, 2021. doi: 10.1016/j.websem.2021.100658.
- [34] Axel Polleres, Romana Pernisch, Angela Bonifati, Daniele Dell’Aglío, Daniil Dobriy, Stefania Dumbrava, Lorena Etcheverry, Nicolas Ferranti, Katja Hose, Ernesto Jiménez-Ruiz, et al. How does knowledge evolve in open knowledge graphs? *Transactions on Graph Data and Knowledge*, 1(1):11–1, 2023. doi: 10.4230/TGDK.1.1.11.
- [35] Martin Raubal. Representing concepts in time. In *Spatial Cognition VI. Learning, Reasoning, and Talking about Space: International Conference Spatial Cognition 2008, Freiburg, Germany, September 15-19, 2008. Proceedings 6*, pages 328–343. Springer, 2008. doi: 10.1007/978-3-540-87601-4_24.

Bibliography

- [36] Christoph Schlieder. Digital heritage: Semantic challenges of long-term preservation. *Semantic Web*, 1(1-2):143–147, 2010. doi: 10.3233/SW-2010-0013.
- [37] Jonathon P Schuldt, Sara H Konrath, and Norbert Schwarz. “Global warming” or “climate change”? whether the planet is warming depends on question wording. *Public opinion quarterly*, 75(1):115–124, 2011. doi: 10.1093/poq/nfq073.
- [38] Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz. What, when, and where do you mean? detecting spatio-temporal concept drift in scientific texts. In *13th International Conference on Geographic Information Science (GIScience 2025)*, pages 16–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025. doi: 10.4230/LIPIcs.GIScience.2025.16.
- [39] Meilin Shi, Krzysztof Janowicz, Zilong Liu, Mina Karimi, Ivan Majic, and Alexandra Fortacz. Defining concept drift and its variants in research data management: A scientometric case study on geographic information science. *Transactions in GIS*, 29(3):e70058, 2025. doi: 10.1111/tgis.70058.
- [40] Cogan Shimizu and Pascal Hitzler. Accelerating knowledge graph and ontology engineering with large language models. *Journal of Web Semantics*, page 100862, 2025. ISSN 1570-8268. doi: 10.1016/j.websem.2025.100862.
- [41] Barry Smith and David M Mark. Geographical categories: an ontological investigation. *International journal of geographical information science*, 15(7):591–612, 2001. doi: 10.1080/13658810110061199.
- [42] Barry Smith and David M Mark. Do mountains exist? towards an ontology of landforms. *Environment and Planning B: Planning and Design*, 30(3):411–427, 2003. doi: 10.1068/b12821.
- [43] Thanos G Stavropoulos, Stelios Andreadis, Efstratios Kontopoulos, and Ioannis Kompatsiaris. SemaDrift: A hybrid method and visual tools to measure semantic

- drift in ontologies. *Journal of Web Semantics*, 54:87–106, 2019. doi: 10.1016/j.webssem.2018.05.001.
- [44] Wolfgang G Stock. Concepts and semantic relations in information science. *Journal of the American Society for Information Science and Technology*, 61(10):1951–1969, 2010. doi: 10.1002/asi.21382.
- [45] UNISDR. United Nations International Strategy for Disaster Reduction. Disaster Terminology. <https://www.undrr.org/terminology/disaster>, 2018. Accessed: 7 November 2024.
- [46] Saskia Van Putten, Carolyn O’Meara, Flurina Wartmann, Joanne Yager, Julia Vilette, Claudia Mazzuca, Claudia Bieling, Niclas Burenhult, Ross Purves, and Asifa Majid. Conceptualisations of landscape differ across european languages. *Plos one*, 15(10):e0239858, 2020. doi: 10.1371/journal.pone.0239858.
- [47] Stella Verkijk, Ritten Roothaert, Romana Pernisch, and Stefan Schlobach. Do you catch my drift? on the usage of embedding methods to measure concept shift in knowledge graphs. In *Proceedings of the 12th Knowledge Capture Conference 2023*, pages 70–74, 2023. doi: 10.1145/3587259.3627555.
- [48] Shenghui Wang, Stefan Schlobach, and Michel Klein. Concept drift and how to identify it. *Journal of Web Semantics*, 9(3):247–265, 2011. doi: 10.1016/j.websem.2011.05.003.
- [49] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016. doi: 10.1038/sdata.2016.18.
- [50] Wilson Wong, Wei Liu, and Mohammed Bennamoun. Ontology learning from text:

Bibliography

- A look back and into the future. *ACM computing surveys (CSUR)*, 44(4):1–36, 2012. doi: 10.1145/2333112.2333115.
- [51] Xiaohuan Wu, Weihua Dong, Lun Wu, and Yu Liu. Data and Code for "Research Themes of Geographical Information Science during 1991–2020: A Retrospective Bibliometric Analysis", 2022.
- [52] Fouad Zablith, Grigoris Antoniou, Mathieu d'Aquin, Giorgos Flouris, Haridimos Kondylakis, Enrico Motta, Dimitris Plexousakis, and Marta Sabou. Ontology evolution: a process-centric survey. *The knowledge engineering review*, 30(1):45–75, 2015. doi: 10.1017/S0269888913000349.
- [53] Yating Zhang, Adam Jatowt, Sourav S Bhowmick, and Katsumi Tanaka. The past is not a foreign country: Detecting semantically similar terms across time. *IEEE Transactions on Knowledge and Data Engineering*, 28(10):2793–2807, 2016. doi: 10.1109/TKDE.2016.2591008.
- [54] Yating Zhang, Adam Jatowt, and Katsumi Tanaka. Is tofu the cheese of asia?: Searching for corresponding objects across geographical areas. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1033–1042, 2017. doi: 10.1145/3041021.3055132.