A Spatiotemporal Knowledge Graph Construction and Management System

Xiuli Zhu [1, 2], Wei Zhang [1], Jie Yang [3], Ye Zhang [1, 2], Ping Lu [1, 2], Xiaoyu Liu [4, *], Yunlu Peng [1, 2], Ran Li [1, 2], Tingting Zhao [1, 2], Xinpeng Wang [1, 2], Xinli Di [1, 2]

[1] National Geomatics Center of China, 28 Lianhuachi West Road, Haidian District, Beijing, 100830, China – (zhuxiuli, zhangw, zhangye, luping, pengyunlu, liran, zhaotingting, wangxinpeng, dixinli)@ngcc.cn
[2] Key Laboratory of Spatio-temporal Information and Intelligent Services (LSIIS), MNR, 28 Lianhuachi West Road, Haidian District, Beijing, 100830, China
[3] State Key Laboratory of Resources and Environmental Information System, Institute of Geographic Sciences and Natural Resources Research, Chinese Academy of Sciences, Beijing, 100101, China – yangjie@igsnrr.ac.cn
[4] Chongqing Institute of Surveying and Monitoring for Planning and Natural Resources, Chongqing, 401120, China – 13983754721@139.com

Keywords: Spatiotemporal Knowledge Graph, Spatiotemporal Knowledge Construction, Spatiotemporal Knowledge Extraction and Alignment, Knowledge Visualization and Retrieval, Spatiotemporal Knowledge Review and Quality assessment, Construction and Management System.

Abstract

To fulfill the current lack of scalable and efficient software systems for constructing and managing spatiotemporal knowledge graphs, this research develops an integrated system that supports full lifecycle knowledge management, multimodal visualization, and the construction of both generic and domain-specific knowledge graphs. By collaboratively building, storing, and updating large-scale spatiotemporal knowledge graphs with hundreds of millions of triples, the system systematically resolves four key difficulties: First, it improves construction efficiency through human-machine collaboration and streamlined workflows. Second, with the help of an integrated knowledge update and completion engine, it enables dynamic updates of instances, entities, relationships, and other knowledge graph elements. Third, the system incorporates review and quality assessment mechanisms to ensure the reliability and consistency. Lastly, it offers multi-layered, multi-view visualization for intuitively displaying spatial distribution and temporal evolution. The system not only enables the engineering implementation of large-scale spatiotemporal knowledge graphs but also provides a closed-loop management framework, delivering a comprehensive solution for constructing, updating, quality-controlling, and providing complex spatiotemporal knowledge applications.

1. Introduction

1.1 Background

Recent advances in Earth observation—from high-resolution satellite imagery to widespread sensor networks and geotagged social media—are generating massive multimodal, dynamic spatiotemporal data. While these data sources provide rich detail for analyzing geographic processes analysis, integrating, interpreting, and extracting meaningful knowledge from them remains a significant challenge. Conventional geographic information systems are good at managing and visualizing spatial data, but they cannot fully modelling semantic relationships to support computing reasoning across domains or simulating change trend. This semantic disconnect limits our ability to transform raw geodata into knowledge service, contributing to urban development, environmental assessment, or emergency planning.

Spatiotemporal knowledge applications need frameworks that can systematically organize, link, and reason with diverse spatiotemporal information. By formally representing semantics and relations, knowledge graphs bridge the gap between data services and knowledge services, thereby shifting from merely processing spatiotemporal data toward building intelligent, knowledge-driven systems.

1.2 Related Work

With the deep integration of big data and artificial intelligence technologies, the knowledge graph plays an important role in handling and understanding complex spatiotemporal information (Lu et al., 2023; Yang et al., 2024). When dealing with spatiotemporal information in domains such as natural resource (Chen et al., 2022; Zhu et al., 2024; Deng and Wang, 2025), urban studies (Wu et al., 2024; Zhang et al., 2025), and emergency management (He et al., 2023), traditional knowledge graph management systems often face four significant challenges: 1) The inherent complexity of fine-grained spatiotemporal knowledge modelling needs substantial expert intervention throughout the construction workflow, resulting in access barriers and operational inefficiency. 2) The fundamental challenge in maintaining dynamic spatiotemporal knowledge repositories is the lack of sophisticated algorithms for systematic knowledge updating and renewal. 3) The absence of established frameworks for spatiotemporal knowledge quality assurance directly compromises reliability, weakening the effectiveness of diagnostic, predictive, and decision-support applications. 4) Current visualization methodologies are unable to effectively present the complex interrelationships between spatial distribution and temporal dynamics. To address these issues systematically, this paper designs and implements a comprehensive spatiotemporal knowledge graph construction and management system that integrates full lifecycle knowledge management, multimodal visualization methods, general and thematic knowledge graph construction.

[*] Corresponding author

To overcome the existing gap in scalable and efficient software solutions for spatiotemporal knowledge graph construction and management, this study proposes an integrated system that enables comprehensive lifecycle management, multimodal visualization, and the construction of both general and domain-specific knowledge graphs.

2. Spatiotemporal Knowledge Graph Construction and Management System Design

2.1 System Architecture

As shown in Figure 1, the overall structure of spatiotemporal knowledge graph construction and management system is organized into four hierarchical layers: the infrastructure layer, the data/knowledge repository layer, the software system layer, and the application layer.

The infrastructure layer provides the fundamental operational environment for the entire system. In a cloud-native context, this layer is composed of either public or private cloud platforms. Through efficient resource scheduling, the cloud environment provides elastic scalability for computing, network, and storage resources, ensuring stable and reliable operation of the spatiotemporal knowledge graph construction and management system.

The data/knowledge repository layer encompasses all the involved data and knowledge resources, including spatiotemporal knowledge graphs, crowdsensing data resources, spatiotemporal information corpus, and business data generated during system runtime (e.g., user logs and the metadata of spatiotemporal knowledge graphs). By offering storage, retrieval, and other fundamental operations on these resources, the Data/Knowledge Repository Layer lays the foundation for implementing various functional modules related to spatiotemporal knowledge graphs.

The software system layer constitutes a collection of supported functionalities and workflow. It includes both independently usable engine tools and integrated functional modules. The engine tools comprise the spatiotemporal knowledge graph collaborative construction engine, the update and completion engine, the spatiotemporal knowledge extraction tool, the data retrieval and quality traceability engine, and the knowledge retrieval and computational reasoning engine. Built on the aforementioned engine tools, the system contains some other functional modules, such as knowledge graph repository management, knowledge construction management, knowledge extraction and alignment, knowledge visualization and retrieval, knowledge review and quality assessment, and user permission management.

Supported by the software system layer, the application layer provides customized knowledge services to different user groups. Specifically, typical applications such as spatiotemporal knowledge graph construction for business cases like natural resource survey and monitoring and geological hazard monitoring and early warning knowledge services. Apart from these large-scale spatiotemporal knowledge graph construction and application, the system delivers knowledge base creation and collaborative rule management for administrators; collaborative ontology construction, knowledge extraction and alignment, and knowledge review for knowledge producers; knowledge graph retrieval and visualization for end-users.

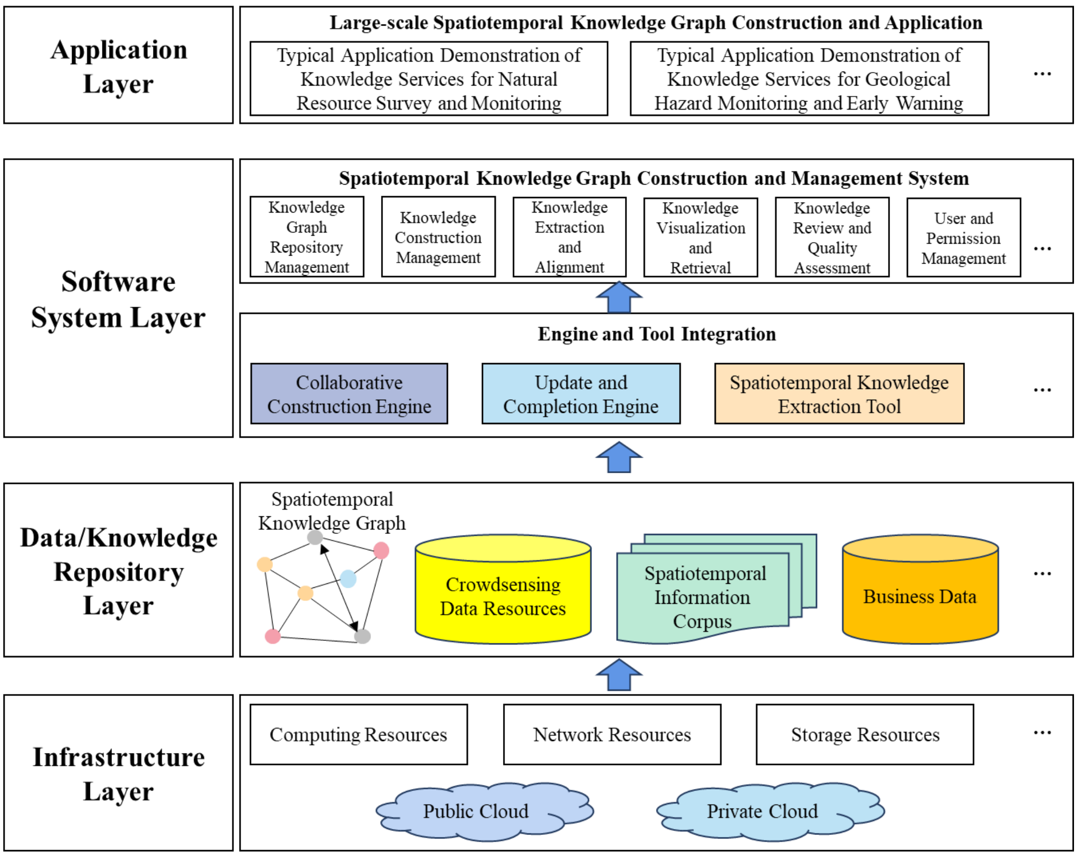

Figure 1. System architecture of the spatiotemporal knowledge graph construction and management system.

2.2 Technical Architecture

Figure 2 presents that the system adopts a containerized deployment strategy using Docker and Kubernetes. Docker packages the application and its dependencies into a portable image, enabling cross-platform deployment on both Linux and Windows systems. It has many advantages such as ease of installation, robust functionality, and a mature ecosystem. As an open-source container orchestration engine, Kubernetes is able to support automated deployment, elastic scaling, and container management capabilities, thereby ensuring high portability, scalability, and automatic operation.

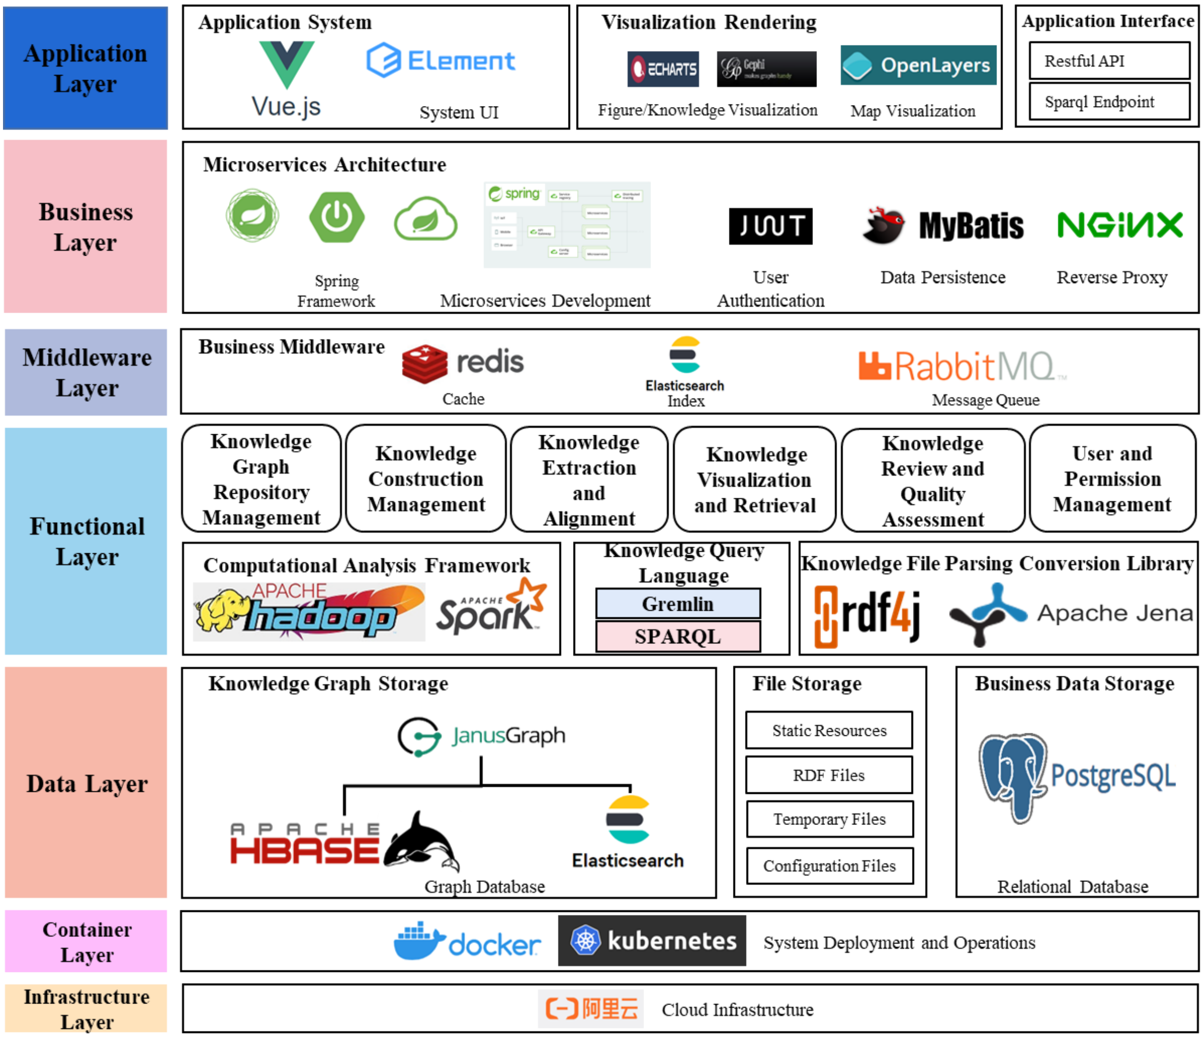

Figure 2. Technical architecture of the spatiotemporal knowledge graph construction and management system.

For the data layer, knowledge graph is stored in the JanusGraph graph database—an open-source, distributed graph database adapted to support the physical storage model of the spatiotemporal knowledge graph. Leveraging its distributed cluster architecture, JanusGraph can store and query hundreds of billions of nodes and edges, meeting the storage requirements for spatiotemporal knowledge graphs containing no less than 50 million triples. Combining YARN integration and the Gremlin query language, its bulk import mechanism achieves update rates exceeding 100,000 triples per minute. JanusGraph employs a decoupled design with separate storage and indexing backends. The storage backend utilizes Apache HBase, a distributed, non-relational database built on the Hadoop ecosystem, which provides persistent large-scale data storage and analytical support and serves as the bridge between JanusGraph and big data computation frameworks. Elasticsearch is a distributed search and analytics engine that supports spatiotemporal conditional queries and full-text search with customizable tokenizers. Powered by it, the indexing backend significantly enhancing query efficiency. The stored files can be classified as static resources, knowledge files, temporary files, and configuration files. Knowledge files are stored in standard formats such as RDF and OWL. Business datasuch as geospatial data and user logs, is stored in PostgreSQL, extended with the PostGIS plugin.

All modules in the functional layer are developed in Java. The system leverages RDF4J and Jena libraries for reading and writing knowledge files, as they provide rich operation function for handling RDF and OWL data. SPARQL is the standard query language for general knowledge graphs, with Gremlin used for graph traversal in JanusGraph. Built on Hadoop and Spark, the knowledge computation and analysis framework has been seamlessly integrated with the underlying HBase storage to support large-scale knowledge processing.

The business layer utilizes Redis for caching, Elasticsearch for indexing, and RabbitMQ for message queuing—all supporting distributed clustering and elastic scalability. Built on the Spring-based microservices architecture, functional modules are encapsulated as independent microservices via Spring Boot, while Spring Cloud provides service registry and discovery, load balancing, load fusing, service degradation, authentication, and monitoring capabilities. Additionally, JWT is employed for user authentication, MyBatis for object-relational mapping and persistence, and Nginx as a reverse proxy to ensure secure and efficient system operation.

The application layer employs Vue.js and ElementUI for the front-end interface, Gephi and ECharts for knowledge layout computation and visualization rendering, OpenLayers for geospatial data visualization. All the system interfaces follow RESTful standards, with an interactive SPARQL Endpoint provided for knowledge retrieval services.

2.3 Interface Design

Figure 3 illustrates that the system's interfaces have been designed into four categories: self-descriptive interfaces, knowledge storage/update interfaces, knowledge detection interfaces, and knowledge acquisition interfaces.

Self-descriptive interfaces are used to explain the available interfaces in the spatiotemporal knowledge graph construction and management system and their structures, thereby improving overall system usability. Knowledge storage/update interfaces are mainly responsible for importing and updating of knowledge graph data, ensuring all the incoming knowledge adheres to the unified spatiotemporal framework. Knowledge detection interfaces are designed to determine and retrieve values of specified indicators in the spatiotemporal knowledge graph, supporting spatiotemporal knowledge graph fusion decisions, quality assessment, and further analysis. Knowledge acquisition interfaces are used to query and retrieve knowledge content according to specified conditions and formats.

All system service interfaces are encapsulated following RESTful standards.

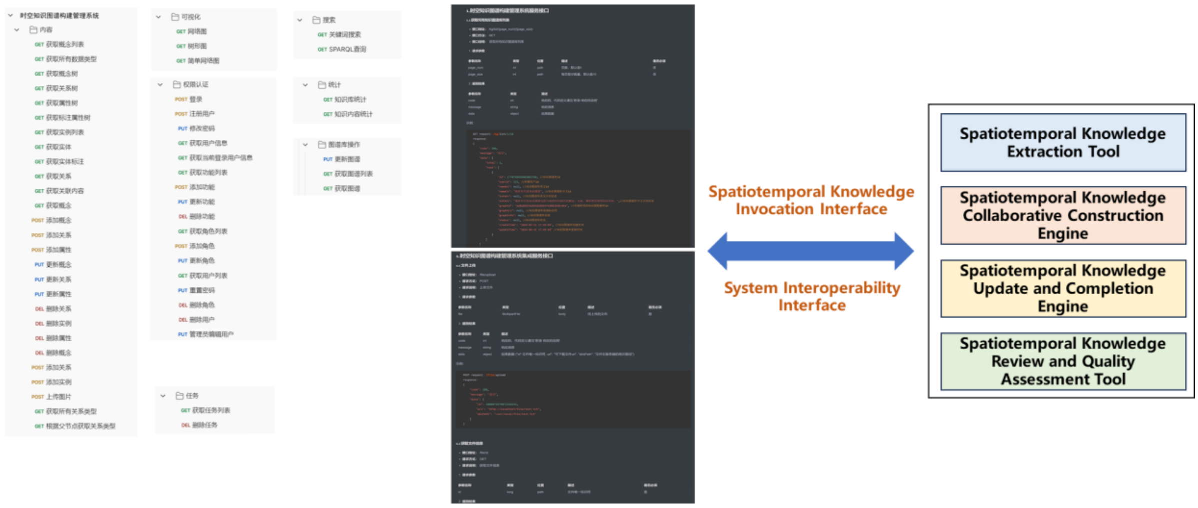

Figure 3. Interface design of the spatiotemporal knowledge graph construction and management system.

3. Functional Modules

Based on an adaptive expression model for spatiotemporal knowledge, utilizing multi-modal and multi-type crowdsensing data, the primary objective of the spatiotemporal knowledge graph construction and management system is to provide a comprehensive software toolkit for the collaborative construction, updating, completion, unified management, retrieval, browsing, and quality tracing of large-scale, high-quality spatiotemporal knowledge graphs. It can be seen from Figure 4 that the functional modules of the spatiotemporal knowledge graph construction and management system include knowledge graph repository construction and management, spatiotemporal knowledge construction and management, spatiotemporal knowledge extraction and alignment, spatiotemporal knowledge visualization and retrieval, spatiotemporal knowledge review and quality assessment, and user and permission management, as illustrated in the Figure 4.

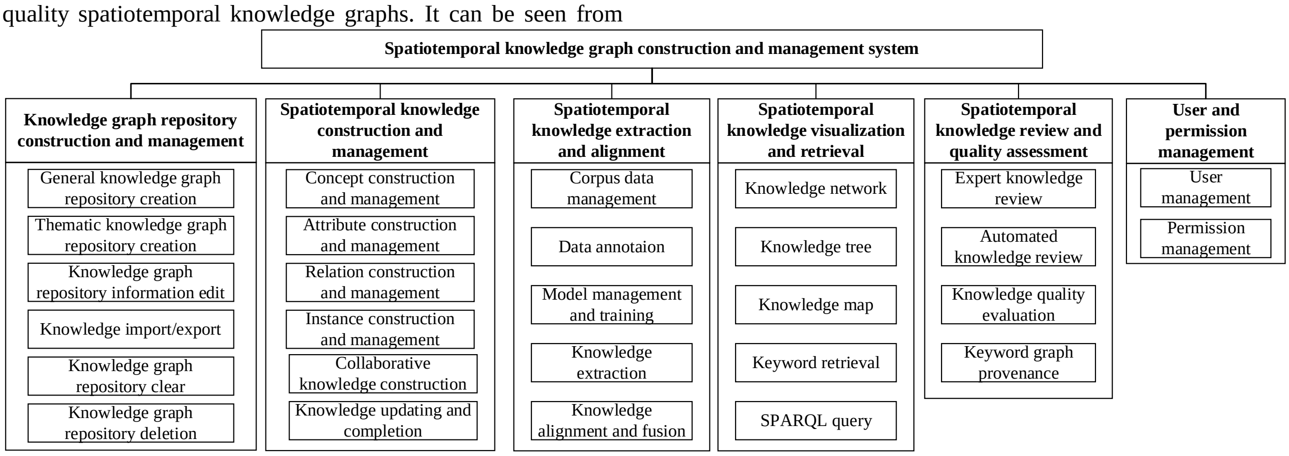

Figure 4. Functional modules of the spatiotemporal knowledge graph construction and management system.

3.1 Knowledge Graph Repository Construction and Management Functional Module

The knowledge graph repository construction and management module enables users to organize spatiotemporal knowledge into a dedicated knowledge graph repository. Figure 5 shows the main interface of it. It provides essential management functionalities at the repository level, facilitating efficient batch management of spatiotemporal knowledge collections for specific domains. This module encompasses operations such as creating new knowledge graph repositories, establishing thematic repositories, editing repository information, importing/exporting knowledge, clearing repository contents, and deleting repositories.

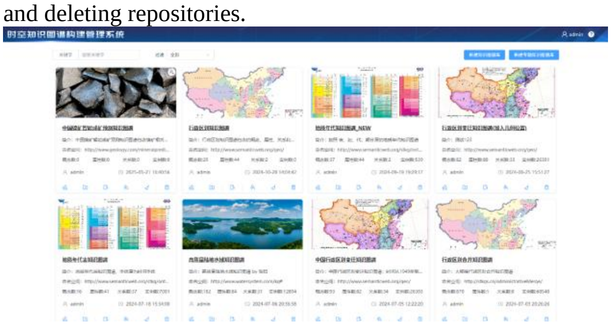

Figure 5. Main interface of spatiotemporal knowledge graph repository management functional module.

3.2 Spatiotemporal Knowledge Construction and Management Functional Module

The spatiotemporal knowledge construction and management module helps users in constructing and managing spatiotemporal knowledge content, particularly at the ontology/schema level. It supports the development of high-quality spatiotemporal knowledge with rich descriptive information and well-defined structures. As we can see from Figure 6, key features of this module include the management of concepts, attributes, relations, and instances, alongside capabilities for collaborative knowledge construction and knowledge updating/completion.

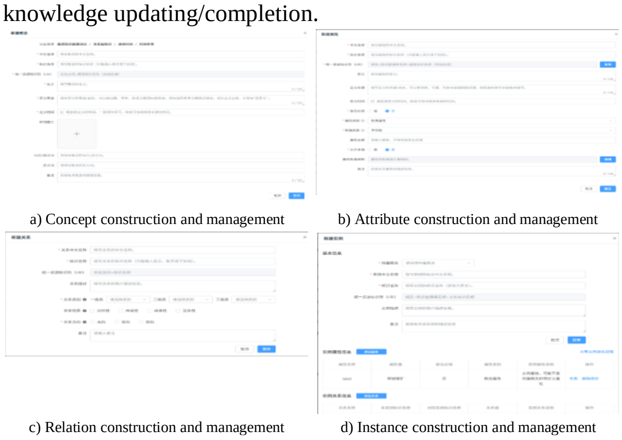

a) Concept construction and management
b) Attribute construction and management
c) Relation construction and management
d) Instance construction and management

Figure 6. Spatiotemporal knowledge construction and management functional module.

3.3 Spatiotemporal Knowledge Extraction and Alignment Functional Module

The spatiotemporal knowledge extraction and alignment module focuses on large-scale extraction of spatiotemporal knowledge instances, thereby enhancing the application value of spatiotemporal knowledge graph repositories. This comprehensive module covers all stages of the knowledge extraction pipeline, including: corpus data management, data annotation, model management and training, knowledge extraction, and knowledge alignment and fusion (Figure 7).

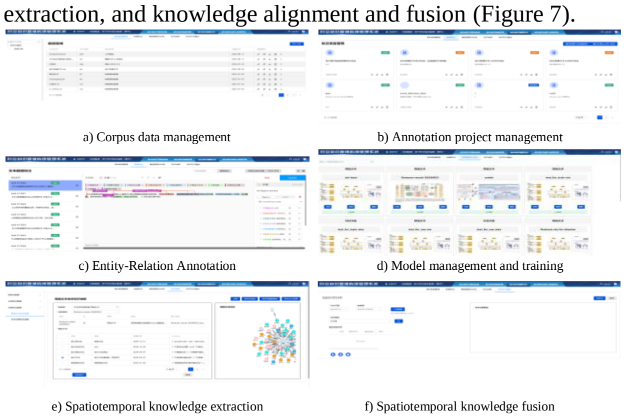

a) Corpus data management
b) Annotation project management
c) Entity-Relation Annotation
d) Model management and training
e) Spatiotemporal knowledge extraction
f) Spatiotemporal knowledge fusion

Figure 7. Spatiotemporal knowledge extraction and alignment functional module.

3.4 Spatiotemporal Knowledge Visualization and Retrieval Functional Module

The spatiotemporal knowledge visualization and retrieval module allows users to retrieve target spatiotemporal knowledge based on customized criteria and intuitively visualize the knowledge and its interrelationships through multiple visualization modes, such as knowledge network, knowledge tree, knowledge map (Figure 8).

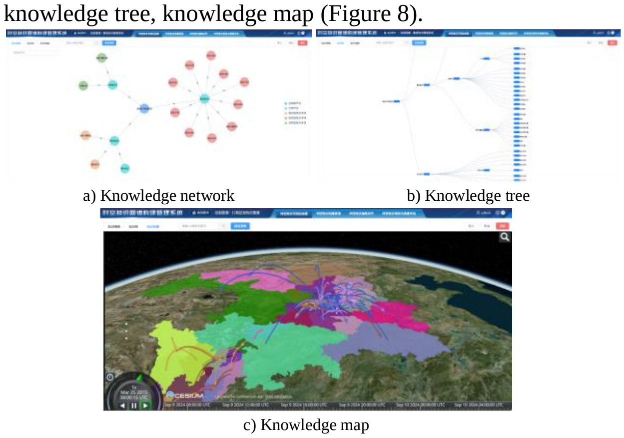

a) Knowledge network
b) Knowledge tree
c) Knowledge map

Figure 8. Spatiotemporal knowledge visualization functional module.

3.5 Spatiotemporal Knowledge Review and Quality Assessment Functional Module

The spatiotemporal knowledge review and quality assessment module delivers systematic workflow support for the evaluation of spatiotemporal knowledge content. It encompasses the following functionalities: expert knowledge review, automated knowledge review, knowledge quality evaluation, and knowledge graph provenance. Two spatiotemporal knowledge retrieval methods is demonstrated in Figure 9.

This contribution has been peer-reviewed.
https://doi.org/10.5194/isprs-archives-XLIX-B4-2026-103-2026 | © Author(s) 2026. CC BY 4.0 License.

a) Keyword retrieval b) SPARQL query

Figure 9. Spatiotemporal knowledge retrieval functional module.

3.6 User and Permission Management Functional Module

The user and permission management functional module facilitates the comprehensive administration of users and their associated access rights, encompassing user management and permission management functionalities.

4. Performance Benchmark Testing and Case Study

4.1 Performance Benchmark Testing of the System

The performance benchmark testing of this system was conducted by the China Software Testing Center. The testing environment and results are detailed in the report: ZT102511361.

4.1.1 Testing Environment:

(1) The server is equipped with an Intel Core i9-14900K CPU (up to 6.0 GHz), an NVIDIA RTX 4090 GPU, 128 GB of RAM, and 2 TB of storage. The software environment comprised Ubuntu 20.04 LTS, PostgreSQL 14, and a set of supporting tools including Java JDK 11.0.16.1, Hadoop 3.4.1, HBase 2.6.2, JanusGraph 1.1, Nginx 1.24.0, Redis 6.2.6, Elasticsearch 7.13.4, Docker 28.0.1, and Docker Compose 2.33.1.

(2) The test machine uses the ThinkPad T14, featuring an Intel Core i7-1165G7 CPU @ 2.80 GHz, 32 GB of RAM, and 1 TB of storage. The client ran Windows 11 Home (Chinese Version) with Microsoft Edge 142.0.3595.94 as the browser.

4.1.2 Testing Results

(1) The spatiotemporal knowledge graph construction and management system includes functional modules such as the homepage, spatiotemporal knowledge construction and management, spatiotemporal knowledge visualization and retrieval, spatiotemporal knowledge extraction and alignment, spatiotemporal knowledge review and quality assessment, permission management, and a unified spatiotemporal ontology repository. In addition to creating new general or thematic spatiotemporal knowledge graph repositories, the homepage supports importing, editing, exporting, clearing, and deleting existing repositories. After constructing concepts, attributes, relations, and instances, the spatiotemporal knowledge construction and management module enables collaborative knowledge construction, update, completion and the adaptive expression model for spatiotemporal knowledge graphs. The spatiotemporal knowledge visualization and retrieval module not only supports keyword and SPARQL-based querying, but also provides visual representations of knowledge networks, knowledge trees, and knowledge maps within repositories. The spatiotemporal knowledge extraction and alignment module achieves corpus data management, data annotation, model management and training, knowledge extraction, and knowledge alignment and fusion. The spatiotemporal knowledge review and quality assessment module implements functions for knowledge review, quality evaluation, and graph provenance.

(2) The spatiotemporal knowledge graph construction and management system is capable of supporting the collaborative construction, storage management, and updating of large-scale spatiotemporal knowledge graphs with no fewer than 50 million triples.

During testing, under a data scale of 30 spatiotemporal knowledge graph repositories, 1,902 concepts, 2,598 attributes, 679 relations, 12,927,336 instances, and 77,742,855 triples, the system performed normally in operations related to collaborative construction, storage management, and updating of spatiotemporal knowledge graphs.

4.2 Case Study

The geographical name dataset utilized in this study was primarily sourced from the national database for geographical names of China. After data preprocessing involving data cleaning, deduplication, manual inspection, and correction, a total of 194,000 toponym entries were acquired for the geographical name knowledge graph construction of Chongqing. All spatial data were referenced to the China Geodetic Coordinate System 2000 and the 1985 National Height Datum, in accordance with the specifications outlined in MZ/T 163-2020 for toponym databases (Ministry of Civil Affairs of the People’s Republic of China, 2020).

The classification of geographical entities adheres to the national standard "Rules for classification of geographical names and code representation" (GB/T 18521-2001) (China National Institute of Standardization, 2002). Following this national standard, geographical entities are categorized into two primary classes: natural geographical entities and human geographical entities. Natural geographical entities are further classified into three subcategories, like marine areas, inland water systems, and terrestrial landforms. Human geographical entities encompass a range of subcategories with toponymic significance, including administrative divisions and other regions, residential settlements, transportation facilities, water conservancy, power, telecommunications facilities, monuments, tourist attractions, and places of historic interest, organizations and institutions, as well as buildings and constructions. A part of the geographical name knowledge graph for Chongqing is presented in Figure 10.

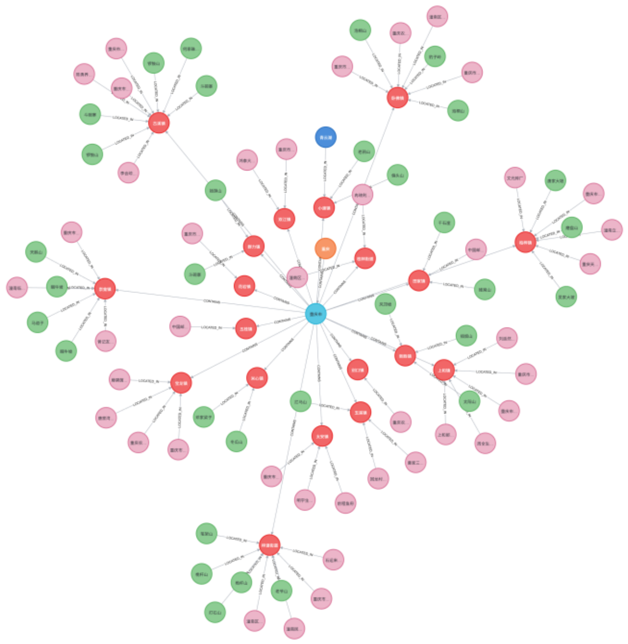

Figure 10. The geographical name knowledge graph for Chongqing.

5. Conclusions

This study presents an integrated system for constructing and managing spatiotemporal knowledge graphs, aiming to address current limitations in transforming geospatial data into actionable knowledge due to the absence of efficient and scalable tools. Our system tackles key challenges in large-scale spatiotemporal knowledge management through a unified approach that combines full lifecycle management, multimodal visualization, and collaborative engines. It enhances construction efficiency via human-machine collaboration and optimized workflows, supports ongoing knowledge refinement through an adaptive update and completion engine, maintains reliability with embedded quality control, and delivers intuitive visual exploration of spatiotemporal patterns. Beyond supporting knowledge graphs at the scale of hundreds of millions of triples, the system establishes a closed-loop "data–knowledge–service" framework, effectively transforming diverse geospatial sources into interpretable and actionable spatial intelligence.

In future, we plan to extend this work in several directions. First, we will incorporate advanced deep learning and pretrained models to improve automated, high-precision extraction from remote sensing imagery, textual content, and other multi-source inputs. Second, we will explore tighter integration with digital twin city platforms, allowing knowledge graphs to interact dynamically with 3D urban models in real time, strengthening decision support in smart city contexts. Finally, to enable cross-institutional collaboration without compromising data security, we will examine federated knowledge graph architectures that support distributed, privacy-aware sharing and joint computation of spatiotemporal knowledge.

Acknowledgements

The comments from two anonymous reviewers are greatly appreciated. The work presented in this paper is supported by the National Key Research and Development Program of China (Grant No. 2022YFB3904201), Research Project of Chongqing Planning and Natural Resources Bureau (No. KJ-2024027).

References

Chen, J., Wu, H., Zhang, Jixian, Wang, D., Liao, A., Liu, W., Zhang, Jun, Miao, Q., Feng, W., Lu, W., 2022. Building natural resources surveying and monitoring technological system: Direction and research agenda. *Acta Geogr. Sin*. 77, 1041–1055. https://doi.org/10.11821/dlxb202205001

China National Institute of Standardization. 2002. Rules for classification and coding of geographical names: GB/T 18521-2001. Standards Press of China, Beijing.

Deng, M., Wang, D., 2025. Research progress in spatio-temporal big data mining and knowledge services for natural resources. *J. Geo-information Sci*. 27. https://doi.org/10.12082/dqxxkx.2025.240625

He, H., Liu, T., Du, P., 2023. Study on the construction of knowledge graph of earthquake disaster emergency management. *J. Lanzhou Jiaotong Univ*. 42, 113–123.

Lu, F., Zhu, Y., Zhang, X., 2023. Spatiotemporal knowledge graph: Advances and perspectives. *J. Geo-Information Sci*. 25, 1091–1105. https://doi.org/10.12082/dqxxkx.2023.230154

Ministry of Civil Affairs of the People's Republic of China. 2020. Specification for database of geographical names: MZ/T 163—2020, Beijing.

Wu, J., Wang, Y., Wang, S., Wang, P., Chen, W., Cui, J., 2024. Knowledge modeling for comprehensive evaluation of urban physical examination: A case study of Guangzhou. *J. Spat. Inf*. 31, 455–468. https://doi.org/10.20117/j.jsti.202402009

Yang, R., Salim, F.D., Xue, H., 2024. SSTKG: Simple Spatio-Temporal Knowledge Graph for Intepretable and Versatile Dynamic Information Embedding. WWW 2024 - Proc. ACM Web Conf. 551–559. https://doi.org/10.1145/3589334.3645441

Zhang, K., Cao, K., Li, W., 2025. A knowledge graph-based approach for recognizing older adults' behavior in urban parks. *Int. J. Geogr. Inf. Sci*. https://doi.org/10.1080/13658816.2025.2499084

Zhu, X.L., Zhu, Xinzhou, Liu, W., Cheng, D., Zhang, X., Zhang, Y., 2024. Knowledge modeling of spatio-temporal changes of land cover in the Amazon Basin. *Bull. Surv. Mapp*. 0, 14–18.