information

MDPI

Article

EcoRestore-KG: A Multi-Agent Framework for Knowledge Graph Construction in Territorial Ecological Restoration

Shibin Zhong [1,2], Xiaoji Lan [1,2,*], Wenhao Yi [1] and Shengdong Nie [1]

[1] School of AirSpace Technology, Jiangxi University of Science and Technology, Ganzhou 341000, China; shibinzhong@mail.jxust.edu.cn (S.Z.); 6120240408@mail.jxust.edu.cn (W.Y.); nieshengdong@mail.jxust.edu.cn (S.N.)
[2] Technology Innovation Center for Land Spatial Ecological Protection and Restoration in Great Lakes Basin, Ministry of Natural Resources, Nanchang 330013, China
* Correspondence: lanxiaoji@jxust.edu.cn

Abstract

Territorial ecological restoration involves a continuous chain of tasks, including degradation diagnosis, target identification, zoning-based governance, process monitoring, effectiveness assessment and adaptive management. The relevant knowledge is widely distributed across policy plans, monitoring reports, restoration cases and the academic literature, and is characterized by multi-source heterogeneity, cross-scale associations and dynamic change. Existing knowledge organization approaches mainly rely on textual synthesis, indicator systems or general-purpose knowledge graph construction tools, and therefore struggle to simultaneously handle cross-context implicit relations, domain-rule constraints, inconsistent entity expressions and evidence traceability in ecological restoration knowledge. To address these limitations, this paper proposes EcoRestore-KG, a multi-agent knowledge graph construction framework for territorial ecological restoration. The framework unifies heterogeneous inputs through controlled evidence representation and adaptive context segmentation, and organizes ontology-guided triple mining, cross-context relation inference, graph quality control, entity canonicalization, relation endpoint remapping and evidence binding into a progressive workflow for the automatic extraction, auditing and assembly of ecological restoration knowledge. Experimental results show that EcoRestore-KG outperforms general-purpose large language models and existing knowledge graph construction baselines in relation extraction, entity coverage and semantic-quality evaluation. It achieves relation precision, recall and F1 scores of 71.4% $\pm$ 0.3%, 69.5% $\pm$ 4.4% and 70.3% $\pm$ 2.3%, respectively, improving relation F1 by 9.9 percentage points over the strongest baseline. Its entity F1 reaches 79.8% $\pm$ 0.7%, and its LLM-S score reaches 8.48 $\pm$ 0.11. Single-module and combined ablation experiments further demonstrate that evidence representation, context segmentation, cross-context relation inference, relation quality auditing and entity canonicalization jointly support the performance gains of the framework. This study provides a verifiable methodological pathway for structured organization, quality auditing, evidence tracing and subsequent integration of newly available knowledge.

Keywords: territorial ecological restoration; knowledge graph; large language model; multi-agent; relation extraction; graph evolution

Check for updates

Academic Editor: Charalampos Bratsas

Received: 13 July 2026
Revised: 7 August 2026
Accepted: 11 August 2026
Published: 13 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

In the Chinese land-space governance context, territorial ecological restoration is a foundational task for advancing ecological civilization, supporting the Beautiful China ini-

tiative and aligning environmental governance with green and low-carbon transition [1,2]. At the national level, ecosystem protection and restoration have been incorporated into the optimization of territorial development and protection patterns, the safeguarding of national ecological security and the deep decarbonization of key sectors. Relevant plans emphasize the integrated protection and restoration of mountains, rivers, forests, farmland, lakes, grasslands and deserts through systems thinking, as well as strengthened monitoring, supervision and dynamic implementation [3]. Under this policy orientation, territorial ecological restoration is no longer the local implementation of a single engineering measure. Rather, it is a continuous process spanning degradation identification, risk diagnosis, target setting, zoning-based governance, process monitoring, effectiveness assessment and adaptive management [4,5]. Because these stages involve multiple data sources, governance units, stakeholder groups and objective trade-offs, ecological restoration exhibits pronounced multi-factor, multi-scale, multi-stage and dynamic-feedback characteristics [5,6].

As research on ecological restoration has deepened, substantial progress has been made in theoretical frameworks, indicator systems, restoration zoning, effectiveness assessment and decision support [5,7,8]. However, these achievements are still mainly carried by policy documents, planning reports, review-based knowledge, indicator systems and model outputs. Their forms of knowledge expression remain fragmented, making it difficult to establish stable semantic associations across sources, scales and governance stages. For ecological restoration practice, the key challenge is not only to obtain individual assessment results, but also to identify the continuous relations among restoration objects, ecological problems, driving factors, governance measures, implementation processes and restoration outcomes. Therefore, organizing ecological restoration knowledge scattered across policy texts, planning schemes, survey and monitoring reports, engineering cases and the academic literature into a structured, computable and updatable form has become an important prerequisite for improving intelligent ecological restoration governance.

Knowledge graphs provide a feasible pathway for the structured organization of complex domain knowledge. By using entities, attributes and relations as basic units, a knowledge graph can transform fragmented knowledge into a searchable, inferable and reusable semantic network, and support rule constraints, relation queries and auxiliary analysis [9,10]. With the development of large language models (LLMs), knowledge graph construction has gradually shifted from manual rules, expert annotation and traditional information extraction models to LLM-assisted entity recognition, relation extraction, knowledge completion, graph reasoning and question-answering-enhanced frameworks. From a roadmap perspective, Pan et al. [11] noted the complementary relationship between LLMs and knowledge graphs: knowledge graphs can provide LLMs with external factual knowledge, structural constraints and interpretable reasoning paths, whereas LLMs can be used for knowledge graph construction, completion and natural-language interaction. Zhu et al. [12] further evaluated the capabilities of LLMs in knowledge graph construction and reasoning tasks, including entity-relation extraction, event extraction, link prediction and question answering, and proposed AutoKG based on a multi-agent paradigm, providing a new methodological pathway for LLM-assisted knowledge acquisition.

Recent studies on document knowledge organization and graph-enhanced retrieval have further promoted the automatic conversion of unstructured materials into graph-structured knowledge. GraphRAG constructs an entity knowledge graph from source documents using LLMs and supports global question answering over an entire corpus through graph community summaries, demonstrating that graph structures can compensate for the limitations of conventional vector retrieval in global and synthetic queries [13]. Docs2KG proposes a unified document knowledge graph construction framework for

heterogeneous documents such as emails, webpages, PDFs and tables, enabling key information to be extracted from multiple types of unstructured materials and organized into a queryable unified graph [14]. AutoKG, proposed by Chen and Bertozzi, extracts keywords using LLMs and estimates relation weights among keywords through graph Laplacian learning, thereby constructing lightweight knowledge graphs and enhancing language model retrieval [15]. In addition, iText2KG emphasizes incremental and topic-independent knowledge graph construction, reducing duplicate entities and redundant relations through document distillation, entity extraction, relation extraction and graph integration modules [16]. Meanwhile, LLM application development frameworks such as LangChain provide engineering support for retrieval-augmented generation, agent workflow orchestration, tool invocation and multi-stage task organization, allowing document processing, retrieval, extraction and validation to be organized as reusable workflows [17].

However, these methods do not fully address the requirements of knowledge graph construction for territorial ecological restoration. GraphRAG primarily supports graph-based global retrieval and question answering, rather than fine-grained relation extraction under domain ontology constraints. Docs2KG focuses on constructing unified knowledge graphs from heterogeneous documents, but provides limited support for implicit cross-chunk relation inference and domain-specific relation auditing. AutoKG constructs lightweight knowledge graphs through keyword extraction and relation-weight estimation, with less emphasis on entity types, relation directions, and provenance binding. Although iText2KG supports incremental graph construction and entity integration, it does not explicitly model the process stages and relation rules of ecological restoration. LangChain mainly provides general-purpose capabilities for model invocation, tool integration, and workflow orchestration, but does not itself include domain-specific mechanisms for knowledge extraction, relation auditing, or entity canonicalization. Taken together, the reviewed methods do not provide, within a unified workflow, domain ontology constraints, cross-context relation inference, candidate relation auditing, entity canonicalization, and evidence traceability tailored to ecological restoration.

To address this gap, we propose EcoRestore-KG, a multi-agent framework for constructing territorial ecological restoration knowledge graphs. It integrates multi-stage task orchestration and LLM-assisted extraction using Deep Agents, an open-source multi-agent collaboration framework developed by LangChain AI [18]. The methodological novelty of EcoRestore-KG lies neither in a new foundation language model nor in a standalone information extraction algorithm. Instead, it lies in a domain-ontology-constrained multi-agent coordination strategy that assigns explicit responsibilities to individual agents according to the sources of errors arising during graph construction. Through structured intermediate representations and ontology-based constraints, the strategy coordinates multi-source evidence representation, adaptive context segmentation, in-chunk fact extraction, cross-context relation supplementation, candidate relation auditing, entity canonicalization, and graph assembly. The domain ontology therefore serves not only as a conceptual schema but also as an operational constraint on entity types, relation types, relation directions, and relation endpoints. This design aims to improve relation completeness, structural consistency, and evidence traceability.

The main contributions of this study are threefold. First, we propose a domain-ontology-constrained multi-agent coordination strategy for constructing territorial ecological restoration knowledge graphs. Rather than directly applying a general-purpose LLM or knowledge graph construction tool, the strategy decomposes graph construction into coordinated stages with explicit task boundaries and structured intermediate outputs. Second, we design a progressive knowledge acquisition mechanism that integrates ontology-guided in-chunk fact extraction, cross-context relation supplementation, and candidate relation auditing. In particular, the

cross-context relation inference agent establishes additional relations only within the extracted entity space, whereas the graph quality control agent evaluates relation types, directions, endpoint constraints, and evidence consistency under the domain ontology. Third, we develop an evidence-traceable graph assembly approach that combines entity canonicalization, relation endpoint remapping, and provenance binding. By linking entities and relations to their source evidence, document locations, processing stages, and auditing records, this approach provides a structured basis for knowledge verification and the subsequent integration of newly available materials.

2. EcoRestore-KG Framework

2.1. Overall Framework

Territorial ecological restoration knowledge is widely distributed across texts, figures, tables and remote-sensing imagery. Because ecological restoration objects are multi-scale, multi-stage and highly process-oriented, the relevant knowledge often involves dispersed evidence sources, strong context dependence and complex relation expressions. Meanwhile, general-purpose LLMs still have limitations in long-context understanding, adherence to domain constraints and consistency control over entities and relations, making it difficult for them to directly satisfy the requirements of knowledge graph construction for accuracy, traceability and continuous updating. To this end, this paper proposes EcoRestore-KG based on the Deep Agents multi-agent collaboration framework. EcoRestore-KG is a multi-agent knowledge graph construction framework for territorial ecological restoration that extracts restoration-related entities, relations and evidence sources from heterogeneous materials and forms a traceable, verifiable and updatable domain knowledge graph. The overall workflow of EcoRestore-KG is shown in Figure 1.

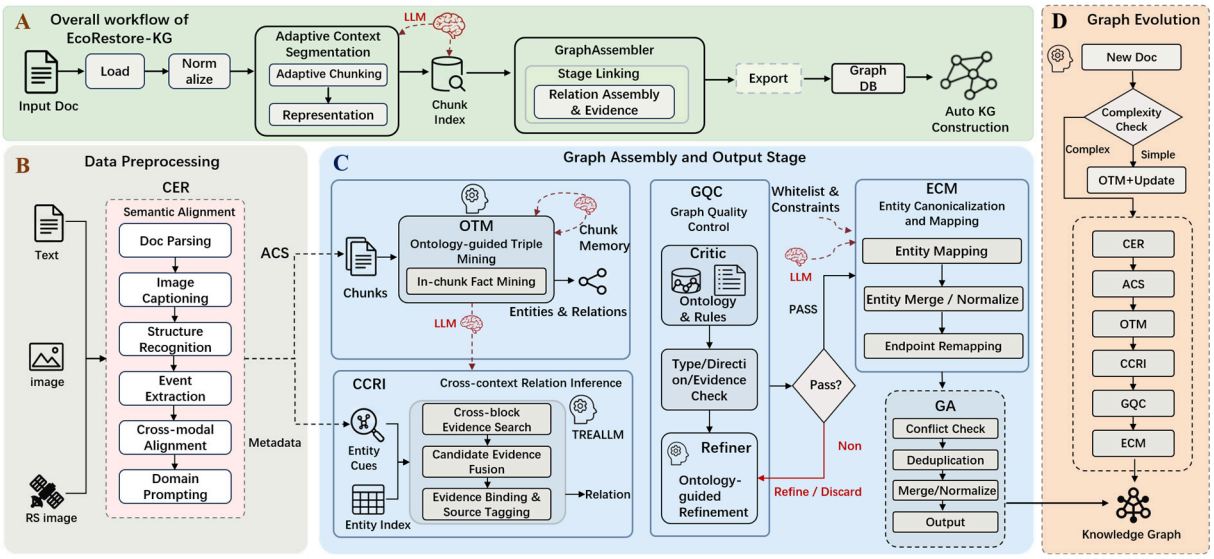


Figure 1. Overall framework of EcoRestore-KG.

The framework consists of four main parts: data preprocessing, knowledge extraction and quality control, graph assembly and output, and graph evolution. EcoRestore-KG first uses the multimodal preprocessing agent CER (Controlled Evidence Representation) to parse and align text, image descriptions, remote-sensing imagery and related materials, transforming multi-source inputs into a unified textual semantic stream. The adaptive context segmentation agent ACS (Adaptive Context Segmentation) then divides the textual semantic stream into semantically coherent evidence chunks and constructs a chunk index, providing length-controlled input units for subsequent extraction. In the knowl-

edge extraction and quality control stage, the ontology-guided triple mining agent OTM (Ontology-guided Triple Mining) identifies entities and explicit in-sentence relations from individual evidence chunks. For relations that depend on cross-sentence or cross-chunk evidence, the cross-context relation inference agent CCRI (Cross-context Relation Inference) further uses entity cues, chunk indexes and contextual evidence to supplement implicit candidate relations. Candidate relations are then passed to the graph quality control agent GQC (Graph Quality Control), where the Critic and Refiner review relation types, directions, entity constraints and evidence consistency, and revise or remove unreliable relations under ontology constraints. The quality-controlled results are sent to the entity canonicalization and mapping agent ECM (Entity Canonicalization and Mapping), which merges and standardizes synonymous entities, aliases and entities expressed at different granularities, and remaps relation endpoints to canonical entities. In the graph assembly and output stage, the graph assembly agent GA (Graph Assembly) organizes entities, relations, processing stages and evidence chains into a structured graph and writes the results to a graph database, forming an automatically constructed ecological restoration knowledge graph. For new materials, EcoRestore-KG first performs a complexity check: simple documents can enter a fast extraction and update workflow, whereas complex documents re-execute preprocessing, segmentation, extraction, inference, quality control and entity canonicalization before integrating new knowledge with the existing graph. Through this evolution mechanism, EcoRestore-KG supports the continuous expansion and dynamic updating of territorial ecological restoration knowledge while maintaining the structural consistency of the existing graph. The detailed execution procedure of the framework is presented in Algorithm 1.

Algorithm 1. EcoRestore-KG Construction Workflow

Input: Multi-source ecological restoration documents D; domain ontology O
Output: Set of document-level knowledge graphs G_set

```
G_set <- empty set
for each document d in D do
    U <- CER(d)
    (C, I) <- ACS(U)
    E_raw <- empty set
    R_local <- empty set
    for each text chunk c in C do
        (E_c, R_c) <- OTM(c, O)
        E_raw <- E_raw union E_c
        R_local <- R_local union R_c
    end for
    R_cross <- CCRI(C, I, E_raw, R_local, O)
    R_reviewed <- GQC(R_local union R_cross, C, O)
    (E_canonical, M) <- ECM(E_raw, R_reviewed, O)
    R_mapped <- RemapEndpoints(R_reviewed, M)
    G_d <- GA(E_canonical, R_mapped, EvidenceMetadata)
    Export(G_d)
    G_set <- G_set union {G_d}
end for
return G_set
```

In the implementation, the agents are orchestrated by a centralized workflow in a predefined sequence and exchange structured data, including text chunks, entities, relations, evidence locations, and auditing results. ACS preferentially segments the input at changes in headings, paragraphs, or ecological restoration stages. The target, maximum, and minimum chunk lengths are set to 900, 1200, and 300 characters, respectively, with an overlap of 120 characters between adjacent chunks. These parameters were determined based on the distribution of paragraph lengths in the non-test corpus and preliminary experiments, with the aim of preserving local semantic integrity while controlling the context length for each knowledge extraction operation. The 120-character overlap, corresponding to approximately 10% of the maximum chunk length, is used to retain entity and relation information near chunk boundaries.

OTM uses structured prompts to restrict its outputs to the entity types, relation types, and relation directions predefined in the domain ontology, with a maximum of five entities and three relations extracted from each sentence. CCRI establishes new relations only among entities already identified by OTM and does not generate new entities. Each cross-chunk candidate relation must cite at least two distinct evidence chunks and satisfy the endpoint-type and relation-type constraints specified in the domain ontology. GQC performs at most one Critic–Refiner review round and terminates when the Critic determines that no revision is required or when the Refiner produces no change to the current result. The complete prompt templates, decision rules, and output schemas for all agents are provided in Appendix A.

2.2. Multimodal Preprocessing and Adaptive Context Segmentation

Territorial ecological restoration materials are typically composed of text, figures, remote-sensing imagery and tables. Information in different modalities is expressed in different ways: texts usually describe restoration objects, governance measures and ecological effects; figures and tables record monitoring indicators and experimental results; and remote-sensing imagery reflects spatial patterns, disturbance ranges and recovery states. Directly feeding these heterogeneous materials into a knowledge extraction agent can easily lead to inconsistent evidence granularity, long-context overload and missing cross-modal information [19]. Therefore, EcoRestore-KG first transforms raw materials into a unified textual semantic stream through the multimodal preprocessing agent, and then generates text chunks suitable for subsequent extraction through the adaptive context segmentation agent.

Given the i-th input material d_i, its content can be represented as a combination of information from different modalities and the associated metadata:

$$d_i = \{T_i, I_i, F_i, R_i, M_i\}, \tag{1}$$

where T_i, I_i, F_i, R_i and M_i denote the body text, image or caption information, tables and structured records, remote-sensing-related information and document metadata, respectively. The multimodal preprocessing agent parses, transcribes and aligns information from different sources to form a unified textual semantic stream:

$$S_i = P(d_i) = \{s_{i1}, s_{i2}, \ldots, s_{in}\}, \tag{2}$$

where $P(\cdot)$ denotes the multimodal preprocessing process, and S_i denotes the textual semantic stream corresponding to the i-th material. Each semantic unit s_{ik} retains its text content, source modality, original location and metadata tags:

$$s_{ik} = (u_{ik}, m_{ik}, p_{ik}, q_{ik}), \tag{3}$$

where u_{ik} is the textual content of the semantic unit, m_{ik} is the source modality, p_{ik} is its position in the original material, and q_{ik} is the source or processing tag. This representation enables subsequent extraction results to be traced back to the original material and provides a basis for the consistent expression of cross-modal information.

After multimodal preprocessing, EcoRestore-KG uses the adaptive context segmentation agent to divide the textual semantic stream into chunks suitable for knowledge extraction:

$$C_i = \mathrm{ACS}(S_i) = \{c_{i1}, c_{i2}, \ldots, c_{im}\}, \tag{4}$$

where C_i denotes the set of text chunks corresponding to the i-th material, and c_{ij} denotes the j-th text chunk. Each text chunk consists of several consecutive semantic units and retains the chunk identifier, source range and adjacency relations. Unlike fixed-length segmentation, ACS preferentially creates chunks at semantic boundaries, such as paragraph endings, title transitions, shifts in experimental procedures or changes in restoration objects. When a text exceeds the preset length, the agent performs secondary splitting based on sentence boundaries; when a chunk is too short and semantically mergeable, it is merged with adjacent semantic units. For relations that may be expressed across sentences or paragraphs, ACS preserves the chunk index and adjacency relations so that the subsequent cross-context relation inference agent can retrieve relevant context.

Through multimodal preprocessing and adaptive context segmentation, EcoRestore-KG converts raw heterogeneous materials into a set of text chunks that are structurally clear, source-traceable and length-controlled. This process reduces multimodal noise and long-text complexity, enabling subsequent agents to perform entity recognition, relation extraction and cross-context inference on more stable input units.

2.3. Knowledge Extraction, Cross-Context Inference and Quality Control

After multimodal preprocessing and context segmentation, EcoRestore-KG feeds evidence chunks into the knowledge extraction stage. This stage is jointly completed by the ontology-guided triple mining agent OTM, the cross-context relation inference agent CCRI and the graph quality control agent GQC. OTM extracts entities and explicit in-sentence relations from individual text chunks, CCRI supplements cross-sentence or cross-chunk relations, and GQC reviews, revises or removes candidate relations. Through this staged design, the framework reduces relation omissions, direction errors and ontology-type conflicts caused by direct long-text extraction.

For text chunk c_{ij}, OTM generates a raw entity set and an in-chunk relation set under the constraints of the domain ontology Ω:

$$\left(E_{ij}^{raw}, R_{ij}^{OTM}\right) = \mathrm{OTM}(c_{ij}, \Omega), \tag{5}$$

where E_{ij}^{raw} denotes the raw entity set identified from text chunk c_{ij}, and R_{ij}^{OTM} denotes the explicit in-chunk relation set. The domain ontology Ω is used to constrain entity types, relation types and relation directions, ensuring domain semantic consistency in the extraction results. Each in-chunk relation can be represented as a triple $r = (h, \rho, t)$, where h and t are relation endpoints and ρ is the relation type.

Some ecological restoration relations depend on cross-sentence or cross-chunk evidence, and relying only on OTM can therefore miss contextual associations. To address this, CCRI generates cross-context relations based on the text chunk set, the raw entity set and the relations already extracted by OTM:

$$R_i^{CCRI} = \mathrm{CCRI}\left(C_i, E_i^{raw}, R_i^{OTM}, \Omega\right), \tag{6}$$

where C_i denotes the text chunk set corresponding to the i-th material, E_i^{raw} denotes all raw entities in that material, R_i^{OTM} denotes the in-chunk relations extracted by OTM, and R_i^{CCRI} denotes the cross-context relation set inferred by CCRI. CCRI establishes new relations only within the existing entity space, thereby avoiding the introduction of new entities that are not supported by the text during cross-chunk inference.

GQC then performs unified quality control on the candidate relations generated by OTM and CCRI:

$$\hat{R} = \mathrm{GQC}\left(R^{OTM} \cup R^{CCRI}, \Omega, \varepsilon\right), \tag{7}$$

where R^{OTM} and R^{CCRI} denote explicit in-chunk relations and cross-context relations, respectively; ε denotes the evidence set bound to candidate relations; and $\hat{R}$ denotes the reviewed and revised relation set. In GQC, the Critic checks relation types, relation directions, entity constraints and evidence consistency, whereas the Refiner revises recoverable erroneous relations under ontology constraints. Candidate relations that lack evidence support or violate ontology constraints are removed.

Accordingly, OTM, CCRI and GQC are responsible for local fact extraction, cross-context relation supplementation and relation quality auditing, respectively. This collaborative mechanism preserves the reliability of explicit in-chunk relations, supplements cross-context relations and reduces ontology-inconsistent relations, direction errors and relation redundancy through graph quality control.

2.4. Entity Canonicalization and Graph Construction

After knowledge extraction, cross-context inference and quality control, EcoRestore-KG obtains a set of candidate entities and reviewed relations. However, because different materials may contain synonymous entities, aliases, abbreviations and descriptions at different granularities, directly writing these results into the graph can easily produce duplicate nodes and inconsistent relation endpoints [20]. Therefore, the framework uses the entity canonicalization and mapping agent ECM to merge and standardize raw entities and remap relation endpoints.

Given the raw entity set E^{raw} and the reviewed relation set $\hat{R}$, ECM maps semantically identical or highly similar entities to a unified canonical entity set:

$$\hat{E} = \mathrm{ECM}(E^{raw}, \Omega), \tag{8}$$

where E^{raw} denotes the raw entity set obtained in the extraction stage, $\hat{E}$ denotes the canonicalized entity set, and Ω denotes the domain ontology. ECM determines whether entities should be merged based on entity names, entity types, contextual semantics and ontology constraints. For example, the same restoration object may appear in different documents as a full name, an abbreviation or a local description. ECM maps these expressions to a unified node, thereby reducing duplicate entities in the graph.

After entity canonicalization, all relation endpoints must be mapped from raw entities to canonical entities. This process can be represented as:

$$\hat{r} = (map(h), \rho, map(t)), \quad r = (h, \rho, t) \in \hat{R}, \tag{9}$$

where r denotes a reviewed raw relation, h and t are the raw head and tail entities, respectively, ρ is the relation type, $map(\cdot)$ denotes the entity mapping function, and $\hat{r}$ denotes the relation after endpoint remapping. If duplicated relations, direction conflicts or ontology violations occur after mapping, deduplication and consistency checks are further performed.

Finally, the graph assembly agent organizes canonicalized entities, remapped relations and their evidence sources into a knowledge graph:

$$G = (V, A, P), \tag{10}$$

where V denotes the entity node set, A denotes the relation edge set, and P denotes the evidence sources, text locations, processing stages and auditing information bound to nodes and relations. Unlike extraction results that store only triples, EcoRestore-KG preserves evidence chains and processing provenance during graph assembly, allowing each relation to be traced back to the corresponding text chunk or multimodal source.

Through the collaboration between ECM and the graph assembly agent, EcoRestore-KG transforms local extraction results into a globally consistent knowledge graph. Entity canonicalization reduces node redundancy caused by synonymous entities and aliases, relation endpoint remapping ensures that edges connect to standardized entities, and graph assembly organizes entities, relations and evidence into a queryable, verifiable and updatable structured graph.

3. Data, Domain Ontology and Evaluation Metrics

3.1. Data Sources and Corpus Construction

Territorial ecological restoration knowledge is characterized by pronounced multi-source heterogeneity. In this study, the corpus was organized along two dimensions: data source and data modality. The academic literature and ecological restoration cases served as the core materials for knowledge extraction and result validation. Policies, regulations and technical documents were used to support domain ontology construction, rule constraints and terminology standardization, whereas remote-sensing imagery provided auxiliary evidence for identifying spatial change and restoration effects. The corpus composition is summarized in Table 1.

Table 1. Composition of the multi-source corpus.

Data Type	Number of Files	Characters [1]	
Academic literature	197	4.5 m+	Used for knowledge extraction
Ecological restoration cases	7	79 k+	Used for case knowledge extraction and result validation
Remote-sensing imagery	16	/	Auxiliary corpus; used to support the identification of spatial change and restoration effects
Policies, regulations and technical documents	236	2.1 m+	Auxiliary corpus; used for ontology construction and rule constraints

[1] The number of extractable characters refers to non-whitespace characters that can be parsed from PDF and DOCX files. Remote-sensing imagery is image-type data and was not included in the textual character statistics. Here, k denotes thousand characters and m denotes million characters.

Table 1 summarizes the data types, corpus scale and functional roles of the materials used in this study. Overall, the corpus covers text, figures, image descriptions and remote-sensing imagery, providing a multi-source data basis for ontology construction, knowledge extraction, evidence tracing and experimental evaluation in EcoRestore-KG.

During corpus construction, textual materials were first parsed and cleaned in a unified manner. Body text, titles, table text, figure captions and paragraph structures were transformed into processable textual semantic streams, while metadata such as material type, document source, text location and paragraph hierarchy were retained. For documents containing images or before-and-after restoration comparisons, figure titles, captions and

related contexts were preserved so that they could be used as textualized evidence in subsequent knowledge extraction. Remote-sensing imagery was linked to textual evidence through descriptive information on spatial location, land-cover type, vegetation change and pre- and post-restoration differences.

After this processing, raw multi-source materials were converted into structured corpora with source identifiers and location information. These corpora provided unified inputs for subsequent adaptive context segmentation, entity recognition, relation extraction, cross-context inference and evidence tracing.

3.2. Domain Ontology Construction

To constrain knowledge extraction and graph assembly in EcoRestore-KG, we constructed a domain ontology for territorial ecological restoration. The ontology specifies entity classes, relation types, relation directions, stage attributes and evidence records. These constraints are intended to reduce entity-boundary inconsistency, relation-direction errors and weak evidence support that may occur when general-purpose LLMs are used for open extraction. The ontology schema is shown in Figure 2.

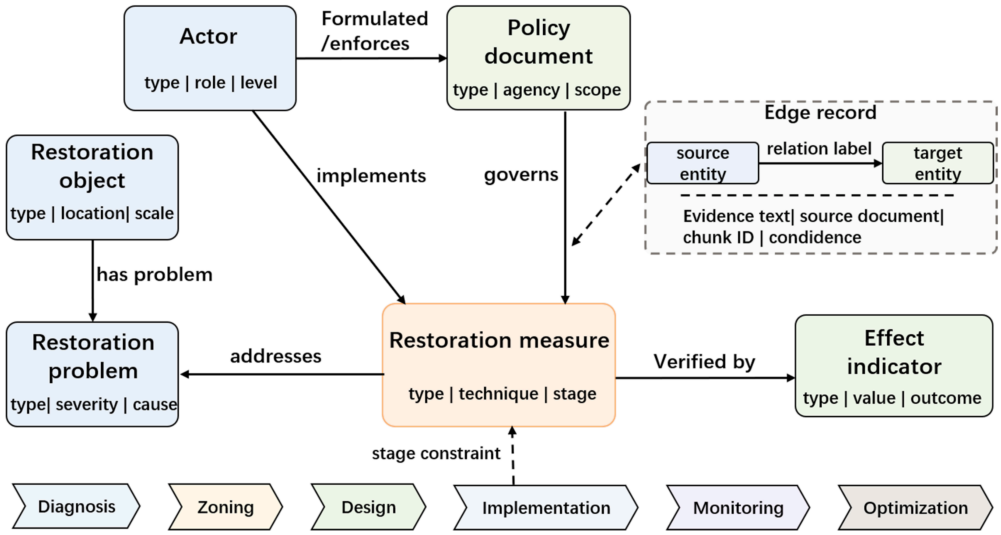


Figure 2. Domain ontology schema of EcoRestore-KG.

The ontology contains six entity classes: restoration object, restoration problem, restoration measure, actor, policy document and effect indicator. Restoration objects include managed objects such as mines, watersheds, land, soil and vegetation. Restoration problems describe pollution, degradation, erosion, subsidence and ecological function decline. Restoration measures include engineering treatment, biological restoration, soil reconstruction, vegetation restoration and monitoring management. Actors include government departments, enterprises, research institutions and implementation units. Policy documents represent plans, specifications and technical guidelines, whereas effect indicators describe evaluation outcomes such as vegetation coverage, soil quality, soil and water conservation effects and ecological function change.

At the relation level, the ontology follows the ecological restoration logic of object identification, problem description, measure response and effect evaluation. Restoration objects, restoration problems, restoration measures and effect indicators form the main semantic chain from problem discovery to treatment verification. Actors, policy documents and restoration measures describe implementation organization and institutional constraints. These relations provide explicit constraints for OTM, CCRI and GQC, allow-

ing candidate triples to maintain type consistency and clear relation directions during generation, supplementation and auditing.

The ontology also represents the restoration process as six stages: diagnosis, zoning, design, implementation, monitoring and optimization. These stages are encoded as key attributes of restoration measures. Each relation edge is further linked to evidence text, source document, text-chunk identifier and confidence score, supporting relation tracing, quality control and graph updating. By jointly defining entity classes, relation constraints, stage attributes and evidence records, the domain ontology provides the structured semantic basis for the automatic construction of EcoRestore-KG.

3.3. Reference Set Construction and Experimental Configuration

The corpus was divided at the document level to prevent text chunks originating from the same document from appearing in both the development and test sets. The non-test corpus was used for ontology construction, prompt design, and parameter determination. The held-out test set was frozen before formal evaluation and was not used to adjust the prompts, ontology, or operational parameters. The domain ontology was developed from ecological restoration policies, technical specifications, the academic literature, and representative cases in the non-test corpus. Its entity classes, relation types, relation directions, and permitted endpoint types were reviewed and finalized before testing.

The reference set was constructed from 197 held-out test documents selected from the complete corpus reported in Table 1 through a combination of manual annotation and dual-model-assisted annotation. First, a subset of the test documents was manually annotated according to unified guidelines covering entity boundaries, entity types, relation types, relation directions, and evidence requirements. For the remaining documents, GPT-5.5 and DeepSeek-V4-Pro independently generated candidate entities and relations under the same ontology constraints. The outputs of the two models were compared after entity-name canonicalization, type alignment, and relation-endpoint mapping. Results agreed upon by both models were retained as candidate annotations, whereas disagreements were manually reviewed against the source evidence. The manual review involved removing unsupported results, correcting entity and relation types and directions, merging synonymous entities, and adding omitted entities and relations. After all identified disagreements had been resolved, the final reference set comprised 197 document-level reference graphs, containing a total of 3185 entity annotations and 1686 relation annotations when aggregated across documents. The same reference set was used for the comparative and the ablation experiments.

All large language models were accessed through cloud-based APIs. GPT-5.5 was accessed through an OpenAI-compatible API, DeepSeek-V4-Pro through the DeepSeek API, and Qwen 3.6-plus through the DashScope OpenAI-compatible API. EcoRestore-KG used temperatures ranging from 0.1 to 0.2 for different agent roles, with a maximum output length of 8192 tokens. Unless otherwise required by their original implementations, the comparison frameworks used a temperature of 0.0 and a maximum output length of 4096 tokens. The LLM-S evaluation model used a temperature of 0.0 and a maximum output length of 1024 tokens.

3.4. Evaluation Metrics

To evaluate the performance of EcoRestore-KG in knowledge graph construction, this paper assesses relation extraction accuracy, entity coverage and triple semantic quality [21]. Relation metrics measure the consistency between predicted relations and the reference graph, entity coverage metrics evaluate the system's ability to identify and cover reference entities, and the LLM-S metric complements these metrics by evaluating the semantic plausibility of extracted triples in the original textual context [22].

For relation extraction, this paper adopts precision, recall and F1. Let N_m^r be the number of matched relations between the predicted and reference graphs, N_p^r be the number of predicted relations, and N_g^r be the number of reference relations. Relation precision, recall and F1 are defined as:

$$Pre_r = \frac{N_m^r}{N_p^r}, \tag{11}$$

$$Recall_r = \frac{N_m^r}{N_g^r}, \tag{12}$$

$$F1_r = \frac{2 \times Pre_r \times Recall_r}{Pre_r + Recall_r}, \tag{13}$$

where Pre_r denotes the proportion of correct relations among predicted relations, $Recall_r$ denotes the proportion of reference relations successfully covered, and $F1_r$ jointly reflects precision and recall in relation extraction. Relation matching requires the head entity, tail entity and relation type to satisfy the matching conditions simultaneously, and semantic similarity is used to alleviate literal inconsistency caused by synonymous expressions [23].

For entity recognition, this paper adopts EACov, ECov and entity F1. EACov denotes entity extraction coverage precision, namely the proportion of predicted entities that can be matched with reference entities. ECov denotes entity coverage recall, namely the proportion of reference entities covered by predicted entities. Let N_m^e be the number of matched entities between predicted and reference entities, N_p^e be the number of predicted entities, and N_g^e be the number of reference entities. Then:

$$\mathrm{EACov} = \frac{N_m^e}{N_p^e}, \tag{14}$$

$$\mathrm{ECov} = \frac{N_m^e}{N_g^e}, \tag{15}$$

$$F1_e = \frac{2 \times \mathrm{EACov} \times \mathrm{ECov}}{\mathrm{EACov} + \mathrm{ECov}}, \tag{16}$$

where EACov reflects the accurate coverage of the predicted entity set, ECov reflects the system's ability to cover the reference entity set, and $F1_e$ jointly evaluates the precision and completeness of entity coverage. Entity matching considers both semantic similarity of entity names and consistency of entity types [24].

In addition, this paper uses the LLM-S metric to evaluate the semantic quality of extracted triples. GLM-5.1, which was not involved in the triple-generation process, was used as an independent judge. For each method, a fixed random seed of 1 was used to sample 30 non-empty predicted graphs, with up to five triples selected from each graph, yielding K = 150 triples for evaluation. The judge model assigns each sampled triple an integer score from 0 to 10 based on consistency with original evidence, semantic plausibility of the relation and relevance to the ecological restoration domain. The evaluation prompt contains the triple to be evaluated, its corresponding source evidence, ontology constraints and scoring criteria, but excludes the method name and reference graph. The temperature of the judge model was set to 0.0, with a maximum output length of 1024 tokens. For K sampled triples, the Q score of LLM-S is defined as:

$$\mathrm{Q} = \frac{1}{K} \sum_{k=1}^{K} score(r_k), \tag{17}$$

where r_k denotes the k-th sampled triple, and $score(r_k)$ denotes the integer score given by the large language model. A higher Q score indicates that extracted triples are more consistent with original evidence, that relation expressions are more reasonable and that

they better conform to the semantics of territorial ecological restoration. LLM-S complements reference-graph-based matching evaluation and is particularly suitable for assessing predicted relations that may be semantically reasonable but are not fully covered by the reference graph. To reduce evaluation bias, all methods used the same sample size, random seed, evaluation prompt and scoring parameters, and the identity of each method was hidden from the judge model. The complete evaluation prompt is provided in Appendix A.

4. Experimental Results and Analysis

4.1. Comparative Experiments

To verify the effectiveness of EcoRestore-KG in territorial ecological restoration knowledge graph construction, this paper compares it with eight categories of baseline methods, including general-purpose LLM methods Qwen 3.6-plus, DeepSeek-V4-Pro [25] and GPT-5.5, as well as general knowledge extraction frameworks LangChain, KG_Builder, GraphRAG, Docs2KG and AutoKG. All methods are evaluated on the same test set and the same reference graph. The evaluation metrics include relation precision, recall and F1; entity EACov, ECov and F1; and the LLM-S Q score. The comparative results are as follows:

Table 2 reports the overall performance of different methods. EcoRestore-KG achieves the best results on all seven metrics. For relation extraction, its precision, recall and F1 reach 71.4%, 69.5% and 70.3%, respectively. Compared with the strongest baseline, Docs2KG, whose relation F1 is 60.4%, EcoRestore-KG improves relation F1 by 9.9 percentage points. This indicates that EcoRestore-KG can extract more reference relations while effectively controlling the introduction of erroneous relations.

Table 2. Performance comparison and statistical significance analysis of different knowledge graph construction methods.

Model	Relations(%)			Entities			LLM-S (1–10)	p_{Holm}
	Pre	Recall	F1	EACov	Ecov	F1	Q	R/E
Qwen 3.6-plus	39.9 ± 7.9	38.8 ± 4.9	39.3 ± 6.4	76.6 ± 0.3	53.7 ± 4.3	63.1 ± 3.0	7.14 ± 0.30	<0.001
DeepSeek V4 Pro	35.2 ± 7.4	36.7 ± 2.0	35.8 ± 4.7	76.3 ± 1.7	54.6 ± 0.9	63.6 ± 0.3	7.09 ± 0.53	<0.001
GPT-5.5	44.7 ± 5.0	58.0 ± 1.3	50.4 ± 3.6	75.8 ± 2.4	69.6 ± 6.6	72.4 ± 2.7	7.44 ± 0.19	<0.001
LangChain	56.8 ± 0.4	53.5 ± 1.3	55.1 ± 0.6	78.2 ± 0.5	67.8 ± 1.1	72.6 ± 0.6	8.20 ± 0.23	<0.001
KG_Builder	33.4 ± 1.5	48.2 ± 3.1	39.4 ± 2.1	69.0 ± 1.1	70.0 ± 0.9	69.5 ± 0.9	6.51 ± 0.19	<0.001
GraphRAG	30.0 ± 1.2	48.3 ± 0.9	37.0 ± 1.2	61.5 ± 0.5	74.5 ± 0.6	67.4 ± 0.4	6.56 ± 0.18	<0.001
Docs2KG	57.0 ± 0.9	64.2 ± 2.1	60.4 ± 1.4	75.8 ± 0.4	63.9 ± 1.8	69.4 ± 1.1	8.31 ± 0.12	<0.001
AutoKG	47.7 ± 1.3	53.9 ± 0.7	50.6 ± 1.1	75.5 ± 1.2	60.1 ± 0.8	66.9 ± 0.9	7.75 ± 0.07	<0.001
EcoRestore-KG	**71.4 ± 0.3**	**69.5 ± 4.4**	**70.3 ± 2.3**	**81.7 ± 0.7**	**78.1 ± 2.0**	**79.8 ± 0.7**	**8.48 ± 0.11**	–

Note: Values are reported as mean ± standard deviation over five runs. Relation and entity metrics are expressed as percentages, whereas LLM-S is evaluated on a 0–10 scale. Bold values indicate the best performance. R/E denotes the Holm-adjusted *p*-values for relation F1 and entity F1, respectively. *p*-values were obtained using two-sided document-level Wilcoxon signed-rank tests comparing each baseline with EcoRestore-KG. EcoRestore-KG was used as the reference method and is therefore marked with "–".

For entity coverage, the EACov, ECov and entity F1 of EcoRestore-KG are 81.7%, 78.1% and 79.8%, respectively. Entity F1 is 7.2 percentage points higher than that of the strongest baseline, LangChain, which achieves 72.6%. This result suggests that the entity canonicalization and mapping agent in EcoRestore-KG can reduce entity boundary inconsistency and synonymous-entity duplication while maintaining entity coverage.

To further examine how different methods perform across specific relation functions, the six ontology relations were grouped into four functional categories: problem identification, measure response, effect verification, and governance and implementation. The F1 scores of different methods in these categories are shown in Figure 3.

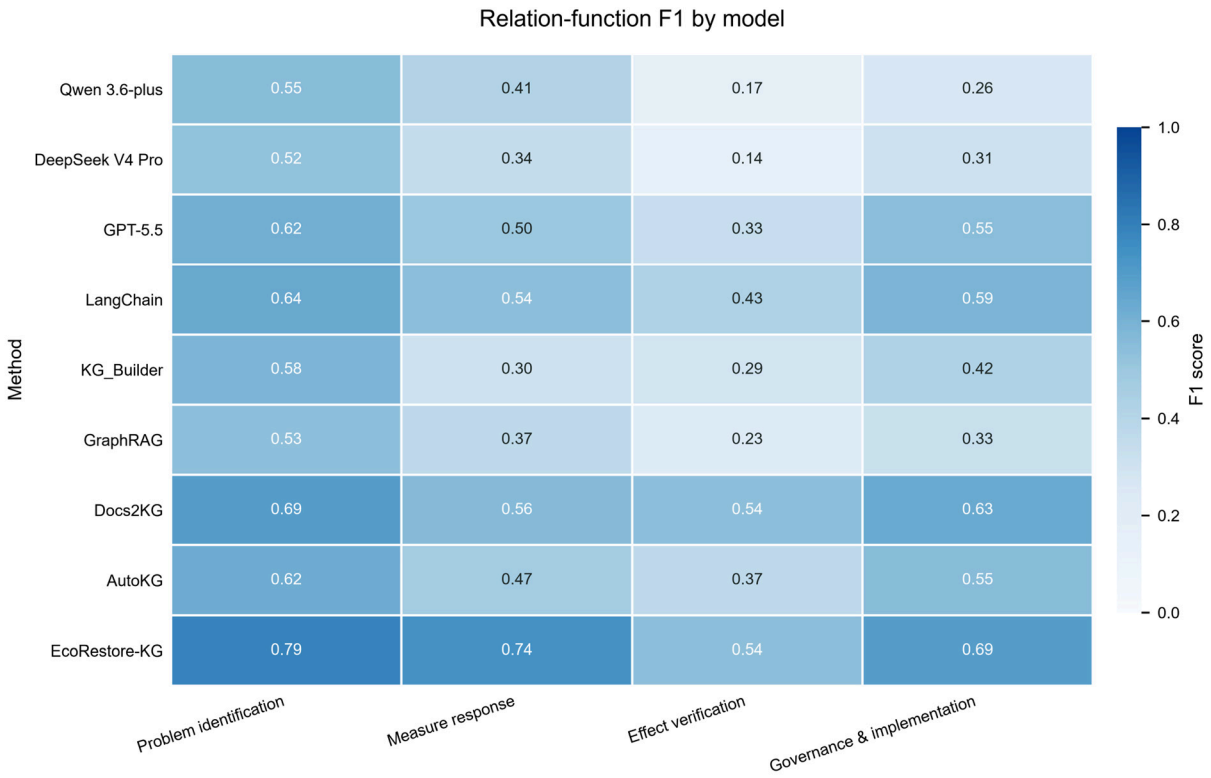


Figure 3. Relation-function F1 scores across different methods.

Consistent with the overall results in Table 2, EcoRestore-KG achieves the highest or near-highest performance across all four relation functions. Its F1 scores reach 0.79, 0.74 and 0.69 for problem identification, measure response, and governance and implementation, respectively, clearly outperforming most baseline methods. This indicates that ontology constraints, cross-context relation inference and graph quality control can effectively improve the extraction quality of both the main ecological restoration relation chain and management-related relations. Although the F1 score for effect verification is relatively low across all methods, EcoRestore-KG still achieves the best result, suggesting that identifying effect relations from monitoring indicators, treatment outcomes and effect descriptions remains a major challenge in this task.

Figure 4 shows the LLM-S scores, relation counts and entity counts of different methods. EcoRestore-KG achieves an LLM-S Q score of 8.48, higher than that of LangChain (8.20) and Docs2KG (8.31), indicating that its generated triples are more stable in terms of consistency with original evidence, relation semantic plausibility and domain relevance. Meanwhile, EcoRestore-KG generates 1676 relations, a number close to those of GPT-5.5 and KG_builder, but with a substantially higher relation F1. This suggests that its performance gains do not come from more aggressive relation expansion, but from more effective candidate relation filtering, quality control and ontology constraints.

Overall, EcoRestore-KG achieves favorable results in relation accuracy, entity coverage and semantic quality. Its performance advantages may be associated with its progressive knowledge-processing mechanism. CER and ACS provide semantically coherent and length-controlled text chunks for subsequent extraction through unified evidence representation and semantic segmentation. OTM and CCRI combine in-chunk explicit fact extraction with cross-context relation completion, thereby improving relation discovery and coverage. GQC examines relation types, directions and endpoints against the domain ontology and source evidence, reducing semantically inconsistent triples. ECM further improves entity coverage and graph consistency through entity canonicalization

and relation-endpoint remapping. Therefore, progressive coordination among multiple agents may be an important reason for the favorable relation F1, entity F1 and LLM-S results achieved by EcoRestore-KG. It also makes the framework more suitable for territorial ecological restoration knowledge characterized by coexisting multi-source evidence, complex cross-paragraph relations and explicit domain constraints.

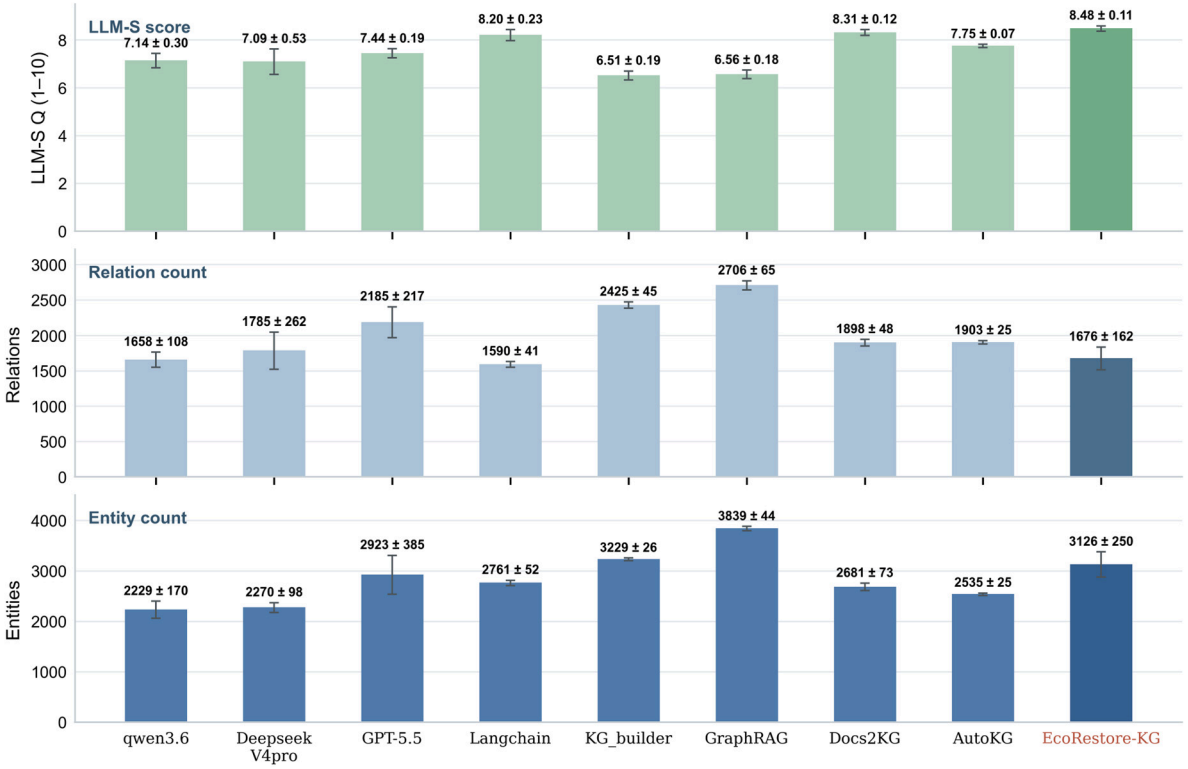


Figure 4. Comparison of LLM-S scores, relation counts and entity counts across different methods.

4.2. Ablation Experiments

To analyze the effects of different agents in EcoRestore-KG on graph construction performance, this paper further conducts ablation experiments. The experiments remove the controlled evidence representation agent CER, the adaptive context segmentation agent ACS, the cross-context relation inference agent CCRI, the graph quality control agent GQC and the entity canonicalization and mapping agent ECM, and compare them with the complete EcoRestore-KG. All ablation settings are evaluated on the same test set, reference graph and metrics.

Table 3 reports the ablation results. The complete EcoRestore-KG reaches 70.3% relation F1, 79.8% entity F1 and an LLM-S Q score of 8.48, outperforming all ablation settings. For both relation F1 and entity F1, all ablated configurations differed significantly from the complete framework (Holm-adjusted $p < 0.001$), demonstrating that each module contributes significantly to the overall performance.

Among the single-module ablations, removing GQC caused the largest decrease in relation F1, from 70.3% to 52.2%, highlighting the importance of relation auditing and evidence-consistency checking for filtering unreliable relations. Removing CCRI reduced relation recall to 53.8%, indicating that cross-context inference supplements relations that cannot be adequately captured within individual text chunks. Removing CER, ACS, and ECM decreased relation F1 to 60.3%, 59.2%, and 59.0%, respectively. Removing ACS and ECM also reduced entity F1 to 74.9% and 74.4%, respectively. These results suggest that evidence representation, context segmentation, and entity canonicalization jointly improve

knowledge extraction by stabilizing input semantics, preserving the semantic integrity of text chunks, and standardizing entities and relation endpoints.

Table 3. Single-module and combined ablation results for EcoRestore-KG.

Model	Relations (%)			Entities (%)			LLM-S (1–10)	p_{Holm}
	Pre	Recall	F1	EACov	ECov	F1	Q	R/E
EcoRestore-KG	**71.4 ± 0.3**	**69.5 ± 4.4**	**70.3 ± 2.3**	**81.7 ± 0.7**	**78.1 ± 2.0**	**79.8 ± 0.7**	**8.48 ± 0.11**	–
w/o CER	61.8 ± 6.0	58.8 ± 7.0	60.3 ± 6.5	79.7 ± 2.2	68.0 ± 0.6	73.3 ± 0.7	8.29 ± 0.18	<0.001
w/o ACS	64.0 ± 0.9	55.2 ± 7.7	59.2 ± 4.8	78.5 ± 0.9	71.5 ± 4.0	74.9 ± 2.5	8.28 ± 0.11	<0.001
w/o CCRI	62.9 ± 0.5	53.8 ± 4.3	57.9 ± 2.3	78.1 ± 1.1	70.2 ± 3.3	73.9 ± 2.3	8.24 ± 0.08	<0.001
w/o GQC	51.1 ± 9.3	53.2 ± 9.4	52.2 ± 9.3	73.7 ± 4.7	70.1 ± 4.1	71.9 ± 4.4	7.77 ± 0.29	<0.001
w/o ECM	64.2 ± 2.9	54.9 ± 8.2	59.0 ± 5.6	78.0 ± 2.2	71.2 ± 6.2	74.4 ± 4.2	8.15 ± 0.14	<0.001
w/o ACS&CCRI	63.3 ± 1.2	54.2 ± 5.4	58.2 ± 2.5	78.2 ± 0.7	70.7 ± 3.3	74.2 ± 2.2	8.23 ± 0.03	<0.001
w/o CCRI&GQC	57.5 ± 5.3	54.4 ± 6.5	55.9 ± 5.9	75.9 ± 3.2	69.9 ± 3.4	72.7 ± 3.3	8.19 ± 0.23	<0.001
w/o GQC&ECM	52.5 ± 6.1	49.5 ± 6.9	50.9 ± 6.5	74.9 ± 3.2	69.3 ± 4.7	72.0 ± 3.9	8.06 ± 0.21	<0.001

Note: Values are reported as mean ± standard deviation over five runs. Relation and entity metrics are expressed as percentages, whereas LLM-S is evaluated on a 1–10 scale. Bold values indicate the best performance. R/E denotes the Holm-adjusted *p*-values for relation F1 and entity F1, respectively. The *p*-values were obtained using two-sided, document-level Wilcoxon signed-rank tests comparing each ablated configuration with the complete EcoRestore-KG framework. EcoRestore-KG served as the reference configuration and is therefore marked with "–".

The combined ablations further reveal the coordination among the modules. Removing ACS and CCRI, CCRI and GQC, and GQC and ECM resulted in relation F1 scores of 58.2%, 55.9%, and 50.9%, and entity F1 scores of 74.2%, 72.7%, and 72.0%, respectively. Simultaneously removing GQC and ECM produced the lowest relation and entity F1 scores, indicating strong functional interdependence between relation-quality auditing and entity-endpoint canonicalization. The combined ablations also differed significantly from the complete framework, further demonstrating that the performance of EcoRestore-KG cannot be attributed to any single component alone.

The LLM-S results further illustrate the influence of different agents on semantic quality. The complete EcoRestore-KG obtains a Q score of 8.48, higher than all ablation settings. Removing GQC reduced the Q score to 7.77, the largest decrease among the single-module ablations, demonstrating the importance of relation auditing and evidence-consistency checking for triple quality. Removing ECM reduced the Q score to 8.15, indicating that inconsistent relation endpoints and duplicate entities also impair the readability and semantic consistency of triples. Removing CER, ACS, and CCRI yielded Q scores of 8.29, 8.28, and 8.24, respectively, although their relation F1 scores decreased more markedly. This suggests that these modules primarily affect evidence organization, relation coverage, and structural completeness, while having a comparatively smaller effect on the local semantic plausibility of the extracted triples.

To further illustrate the specific effects of individual components on knowledge extraction, this study selected the document entitled "*Connotation, Logic, Path and Practice of Mine Ecological Restoration Technology Integration*" as a representative case and compared the outputs of the complete EcoRestore-KG with those of selected ablation settings. This document covers restoration objects, ecological problems, policy provisions, integrated technical measures and restoration effects, making it suitable for examining the framework's ability to preserve the semantics of coordinated measures, refine relation endpoints and standardize entity expressions. Table 4 presents representative outputs generated from the same source evidence under different configurations. The first two contrasts were observed in both repeated runs. Because the entity expressions generated without ECM varied between runs, the third contrast presents a representative result from one run.

Table 4. Representative cross-configuration output contrasts for the case document.

Source Evidence	Ablated Output	Complete EcoRestore-KG Output	Observed Improvement
Physical, chemical and biological remediation were jointly used to address soil and water pollution.	w/o GQC: Physical remediation → address → soil pollution; chemical remediation → address → soil pollution; biological remediation → address → soil pollution.	Integrated physical-chemical-biological remediation → address → soil pollution (and water pollution).	Preserves the coordinated-measure semantics instead of decomposing the joint intervention into unsupported isolated actions.
The Xingdong mining area used high-water backfilling, mine-water treatment and revegetation to address goaf and mine-water problems in an underground mine.	w/o GQC: Integrated measure → address → goaf problem.	Integrated measure → address → goaf problem in underground mine.	Refines a broad endpoint to the mine-specific problem explicitly supported by the source evidence.
Physical, chemical and biological remediation were jointly used for pollution control.	w/o ECM: Physical, chemical and biological remediation jointly → address → soil pollution.	Integrated physical-chemical-biological remediation → address → soil pollution.	Canonicalizes an incomplete surface form to a stable entity name, which supports consistent node reuse and relation-endpoint mapping.

The case results further explain the numerical changes observed in the ablation experiments. Without GQC, the coordinated restoration measure was decomposed into multiple independent entities, resulting in isolated restoration meanings that were not explicitly supported by the source text. In addition, some relation endpoints changed from the specific "goaf problem in underground mines" to the broader "goaf problem". By examining relation types, directions, endpoint types and consistency with source evidence, the complete framework better preserved the original semantics. Without ECM, entity names could be incomplete or inconsistent, affecting node reuse and relation-endpoint mapping for the same entity. Entity canonicalization standardizes these expressions before graph assembly, thereby reducing node redundancy and relation-endpoint inconsistency.

Taken together, the quantitative ablation results and case analysis indicate that the performance improvements of EcoRestore-KG are associated with the coordination of multiple agent components. CER provides a unified representation of multi-source evidence; ACS generates semantically coherent and length-controlled text chunks for subsequent extraction; CCRI supplements candidate relations that depend on cross-sentence or cross-chunk evidence; GQC reduces triples with semantic inconsistencies or inaccurate endpoints through relation auditing; and ECM improves graph consistency through entity canonicalization and relation-endpoint remapping. The case comparison further demonstrates that the complete framework can better preserve the overall semantics of coordinated restoration measures, refine relation endpoints and standardize entity expressions. Together, the quantitative results and case evidence show that these components affect not only the overall evaluation metrics but also the semantic completeness of extracted triples and the quality of final graph assembly.

5. Results and Discussion

The experimental results show that EcoRestore-KG achieves strong overall performance in territorial ecological restoration knowledge graph construction. Compared with general-purpose LLMs and existing knowledge graph construction tools, the main advantage of EcoRestore-KG is not that it generates more entities or relations, but that it maintains high relation accuracy, entity coverage and semantic consistency under relatively high recall. This indicates that multi-agent collaboration and domain ontology constraints can effectively alleviate construction errors caused by long texts, multimodal evidence and cross-context relations in ecological restoration materials.

The comparative experiments show that differences among methods mainly arise from two aspects. First, although general-purpose LLMs have strong text understanding capabilities, they lack structured constraints tailored to ecological restoration ontologies. As a result, they can generate relations that are semantically plausible but difficult to match precisely with the reference graph. Second, some general-purpose knowledge graph construction frameworks can generate many entities or relations, but lack consistency checks for domain relation types, entity types and relation endpoints. This leads to redundant nodes, weakly relevant relations and direction errors in the graph. EcoRestore-KG constrains candidate knowledge within an interpretable domain semantic space through ontology-guided extraction, cross-context relation inference and graph quality control, thereby improving the structural consistency of the final graph.

The ablation experiments further show that the performance of EcoRestore-KG does not come from a single agent, but from complementary collaboration among multiple agents. Multimodal evidence representation and adaptive context segmentation improve input evidence quality, enabling subsequent extraction agents to process clearer contextual units. Ontology-guided triple mining captures explicit in-chunk facts; cross-context relation inference supplements relations that are difficult to cover within a single text chunk; graph quality control reduces ontology-inconsistent relations and direction errors; and entity canonicalization and mapping reduce the impact of synonymous entities and endpoint inconsistency on graph structure. This staged design allows the framework to address different sources of error, including input noise, relation omissions, candidate relation errors and entity redundancy.

It is worth noting that LLM-S is not fully equivalent to reference-graph-based matching metrics. Some baseline methods have high LLM-S scores but low relation F1, indicating that their generated triples may be locally semantically fluent but may not cover key relations in the reference graph. Conversely, matching metrics can reflect the consistency of entities and relations with the reference graph, but may underestimate reasonable relations that are not covered by the reference graph [26]. Therefore, in ecological restoration knowledge graph construction, jointly using matching metrics and semantic scoring provides a more comprehensive evaluation of model performance.

The ablation results indicate that the performance improvements of EcoRestore-KG do not arise simply from increasing the number of agents, but from their complementary functions in evidence processing, relation completion, quality control and entity canonicalization. Removing CER produces the largest decrease in relation F1, whereas removing GQC and ECM mainly affects triple semantic quality and entity consistency. This finding is consistent with the study by Huang et al., who constructed a domain ontology and knowledge graph from multi-source hazard data, highlighting the importance of high-quality evidence organization and ontology constraints for domain knowledge graph construction [27]. Pu et al. proposed MA-APKG, demonstrating that role specialization among multiple agents can support knowledge graph construction for complex processes [28]. Borrelli et al. further applied multi-agent collaboration to knowledge graph construction and enrichment over the scientific literature in MAKES-QA [29]. Compared with these studies, our results further suggest that role specialization alone is insufficient for territorial ecological restoration. Cross-context relation inference, candidate relation auditing, entity canonicalization and source-evidence binding are also required. Because these studies use different datasets, domain ontologies and evaluation metrics, we do not directly compare their numerical results. Instead, they provide relevant evidence for interpreting the roles of multi-source evidence organization, domain constraints and multi-agent collaboration in the present study.

From an application perspective, EcoRestore-KG is more suitable for ecological restoration knowledge scenarios with clear domain ontologies, complex evidence sources and

continuous updating requirements. The framework preserves information on entities, relations, source texts and processing stages, enabling the graph to support not only knowledge querying but also subsequent evidence tracing and graph evolution. It should be noted that the comparative experiments in this study represent end-to-end system-level comparisons among different methods, and the observed differences cannot be attributed exclusively to the proposed architecture. All methods were evaluated using the same test set, reference graph and evaluation metrics, while major inference parameters, such as temperature and maximum output length, were standardized where compatible. Nevertheless, differences in prompting strategies, model backends, task orchestration and original implementations may also affect the final results. Although the ablation experiments provide additional evidence for the effectiveness of the main EcoRestore-KG components, they cannot completely exclude the influence of prompting and experimental settings. Therefore, the current results should be interpreted as evidence of the performance advantages of the overall EcoRestore-KG workflow under the experimental conditions of this study, rather than as an independent causal effect of the architecture alone. However, the framework still depends on the coverage of the domain ontology and the stability of LLM outputs [30]. When the ontology does not cover some reasonable relations, graph quality control may classify them as unsupported or ontology-inconsistent combinations; when input materials are of low quality or relations are highly implicit, cross-context relation inference may still miss relevant relations [31]. Future research can further expand the ecological restoration ontology, optimize relation auditing strategies and reduce the computational cost introduced by multi-agent invocation [32]. In addition, stricter controlled experiments could be conducted by standardizing the model backend, prompting budget and major inference parameters across framework-based methods, thereby further distinguishing the effects of framework architecture, prompt design and experimental settings on final performance.

Author Contributions: Conceptualization, S.Z. and X.L.; methodology, S.Z. and X.L.; software, S.Z.; validation, S.Z., W.Y. and S.N.; formal analysis, S.Z.; investigation, S.Z.; resources, X.L.; data curation, S.Z.; writing—original draft preparation, S.Z.; writing—review and editing, S.Z., X.L., W.Y. and S.N.; visualization, S.Z.; supervision, X.L.; project administration, X.L.; funding acquisition, X.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Open Fund of the Technology Innovation Center for Territorial Ecological Protection and Restoration in the Great Lake Basin of Jiangxi Province, grant number JXCXZX2025003 and the project Research on a Geographic Entity Data Model for Urban Management, grant number 20244BCE52216.

Data Availability Statement: The data presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

For readability, executable prompts are shown as system_prompt and user_prompt, although the implementation concatenates them into a single request before model invocation. Braced terms denote runtime variables. Repeated ontology labels and few-shot examples are abbreviated without changing the operative decision rules. ACS, ECM, and GA are deterministic modules and therefore do not invoke an LLM prompt.

Controlled Evidence Representation Agent (CER)

system_prompt = """You are an expert in controlled evidence representation for ecological-restoration knowledge graphs. Rewrite the source into concise, coherent, low-ambiguity evidence paragraphs; do not summarize the study or extract triples. Resolve pronouns, recover omitted arguments, preserve names, locations, years, values, and thresholds, and never introduce facts absent from the source. Use natural prose, with no ontology tags, JSON, lists, or explanations. Keep the output within 1,000 characters and express no more than two ontology relations per sentence."""

user_prompt = """Allowed schema: {ontology_schema}
Generation rules: {generation_rules}
Source text: {source_text}
Return only the rewritten evidence text."""

parameters = {"max_input_chars": 80,000,
"max_output_chars": 1000,
"max_relations_per_sentence": 2}

Adaptive Context Segmentation Agent (ACS): Deterministic Rules

llm_prompt = None

implementation_rule = """Parse headings and paragraphs into semantic units. Start a new chunk when the maximum length would be exceeded, or when the title/restoration-stage label changes after the target length is reached. Preserve source spans, section titles, stage candidates, ontology hints, chunk identifiers, and adjacent overlap."""

parameters = {"target_chunk_chars": 900,
"max_chunk_chars": 1200,
"min_chunk_chars": 300,
"overlap_chars": 120}

Ontology-Guided Triple Mining Agent (OTM)

system_prompt = """You are an ecological-restoration knowledge graph extraction expert. Extract only entities and sentence-local relations supported by the source chunk. Use only the supplied ontology schema; do not create entity types, relation types, labels, or output fields. Normalize entity names and split separable coordinated entities. Process the text sentence by sentence, with at most five entities and three relations per sentence. Do not infer cross-sentence relations. Return valid JSON only."""

user_prompt = """Document type: {document_type}
Section title: {section_title}
Stage candidates: {stage_candidates}
Ontology hints: {ontology_hints}
Allowed relations: {relation_schema}
Source chunk: {chunk_text}

Output: {"entities": [{"id": "e1", "entity_type": "...", "name": "...",
"mention": "...", "attributes": {}}],
"relations": [{"type": "relation_type(relation_label)",
"from": "e1", "to": "e2"}]}"""

Cross-Context Relation Inference Agent (CCRI)

system_prompt = """You are a relation-only cross-context inference agent. Infer only relations whose evidence is distributed across two or more chunks. Do not add, rename, or delete entities. Use only the supplied entities and ontology relations. Do not output speculative or single-chunk relations. Every relation must cite at least two distinct valid chunk IDs and set evidence_type to cross_chunk_explicit. Return JSON only."""

user_prompt = """Allowed relations: {relation_schema}
Chunks: {chunks_with_ids_and_text}
Existing entities: {entity_list}
Existing relations: {relation_list}

Output: {"relations": [{"type": "relation_type(relation_label)",
"from": "entity_name_or_id", "to": "entity_name_or_id",
"evidence_chunk_ids": ["chunk#0", "chunk#1"],
"evidence_type": "cross_chunk_explicit", "reason": "..."}]}"""

Graph Quality Control Agent (GQC): Critic

system_prompt = """You are a strict knowledge graph extraction critic. Review only the source text, current extraction, and ontology rules; never use or request a gold/reference graph. Identify evidence-verifiable entity-boundary errors, over-extraction, missing sentence-local relations, reversed directions, wrong relation types, unsupported cross-sentence relations, duplicates, or schema violations. Do not rewrite the graph. If no correction is needed, return should_refine=false."""

user_prompt = """Source context: {document_metadata}
Source text: {chunk_text}
Current extraction: {current_extraction_json}

Output: {"should_refine": true,
"issues": [{"type": "...", "severity": "major", "target": "...",
"evidence": "...", "reason": "...", "suggestion": "..."}]}"""

Graph Quality Control Agent (GQC): Refiner

system_prompt = """You are a strict knowledge graph extraction refiner. Revise the current extraction using only the source text, critic report, and ontology schema. Preserve correct items; add only directly supported items; remove unsupported, duplicate, or over-broad items; and use normalized entity names. Return the complete revised extraction as valid JSON, with no explanations or extra fields."""

user_prompt = """Allowed relations: {relation_schema}
Source context: {document_metadata}
Source text: {chunk_text}
Current extraction: {current_extraction_json}
Critic report: {critique_json}
Output: {"entities": [...], "relations": [...]}"""

review_policy = {"max_rounds": 1, "stop_if_should_refine_is_false": true,
"stop_if_output_is_unchanged": true}

Entity Canonicalization and Graph Assembly (ECM and GA): Deterministic Rules

llm_prompt = None

ECM_rule = """Merge entity records sharing the same normalized entity type and canonical name; combine attributes and evidence/source identifiers while preserving insertion order; and remap relation endpoints to canonical entity identifiers."""

GA_rule = """Create stable entity and restoration-stage node identifiers; construct ontology-compliant edges; attach evidence chunk IDs, source document IDs, stage links, relation labels, and confidence fields; reject unresolved endpoints; and export the document-level graph as structured JSON."""

LLM-S Semantic Quality Evaluation Prompt

system_prompt = """You are an independent evaluator of ecological-restoration knowledge graph triples. Evaluate each triple only against the supplied source evidence and ontology constraints. Do not infer the identity of the extraction method, do not favor any model, and do not use a reference graph or external answer key.

Assign integer scores from 0 to 10 for:
1. evidence_consistency: whether the triple is supported by the source evidence;
2. semantic_reasonableness: whether the head entity, relation, tail entity, and relation direction form a coherent fact;
3. domain_relevance: whether the triple is meaningful and plausible in territorial ecological restoration.

Assign an overall score from 0 to 10 using the following anchors:
- 10: directly supported by the source; entity boundaries and relation direction are accurate;
- 8–9: explicitly supported, with only minor generalization of an entity boundary or wording;
- 6–7: indirectly inferable from the source but not explicitly stated;
- 4–5: topically related, but the supporting evidence is insufficient or ambiguous;
- 1–3: unsupported, semantically mismatched, or contains an evidently incorrect endpoint or direction;
- 0: contradicted by the source evidence.

Evaluate every supplied sample_index exactly once. Return valid JSON only, with integer scores and a concise reason of no more than 30 words. Do not output Markdown or explanatory text outside the JSON object."""

user_prompt = """Ontology constraints:
{relation_schema}

Source evidence:
{source_context}

Triples to evaluate:
{indexed_triples}

Return the following structure:
{"scores": [{"sample_index": 1, "evidence_consistency": 0, "semantic_reasonableness": 0, "domain_relevance": 0, "score": 0, "reason": "..."}]}"""

Runtime settings: judge_model = "GLM-5.1"; graph_sample_size = 30; triples_per_graph <= 5; K = 150; random_seed = 1; temperature = 0.0; max_tokens = 1024; source_context <= 4000 characters; method identity and reference-graph triples are excluded.

References

1. Jin, X.B.; Han, B.; Li, Q.Q.; Zhou, Y.K. Territorial ecological restoration planning towards Land Degradation Neutrality: Logical framework and key issues. *J. Nat. Resour.* **2025**, *40*, 287–301. [CrossRef]
2. Central Committee of the Communist Party of China; State Council. Opinions of the Central Committee of the Communist Party of China and the State Council on Comprehensively Advancing the Construction of a Beautiful China. Available online: https://www.mee.gov.cn/zcwj/zyygwj/202401/t20240111_1063458.shtml (accessed on 13 June 2026).
3. National Development and Reform Commission; Ministry of Natural Resources. Notice on Issuing the Master Plan for Major Projects for the Protection and Restoration of National Important Ecosystems (2021–2035). Available online: https://www.ndrc.gov.cn/xxgk/zcfb/tz/202006/t20200611_1231112.html (accessed on 13 June 2026).
4. Dai, X.; Webb, J.A.; Horne, A.C.; Rumpff, L.; Palmer, M.A. Linking the multiple types of monitoring to the adaptive management cycle to support environmental flows. *J. Environ. Manag.* **2025**, *373*, 123694. [CrossRef] [PubMed]
5. Wang, C.; Liu, Y.; Yu, C.; Liu, X. Research progress on the arrangement of territorial ecological restoration. *Prog. Geogr.* **2021**, *40*, 1925–1941. [CrossRef]
6. Tang, H.; Peng, J.; Xu, D.M.; Wu, J.S. Ecological restoration reference of terrestrial space: Theoretical recognition, identification framework and key issues. *J. Nat. Resour.* **2024**, *39*, 2768–2782. [CrossRef]
7. Gu, T.W.; Peng, J.; Jiang, H.; He, C.F. Watershed-based territorial ecological restoration: Theoretical cognition and key planning issues. *J. Nat. Resour.* **2023**, *38*, 2464–2474. [CrossRef]
8. Peng, J.; Xu, D.M.; Xu, Z.H.; Tang, H.; Jiang, H.; Dong, J.Q.; Liu, Y.X. Ten key issues for ecological restoration of territorial space. *Natl. Sci. Rev.* **2024**, *11*, nwae176. [CrossRef] [PubMed]
9. Ma, X.G. Knowledge graph construction and application in geosciences: A review. *Comput. Geosci.* **2022**, *161*, 105082. [CrossRef]
10. Shimizu, C.; Stephen, S.; Barua, A.; Cai, L.; Christou, A.; Currier, K.; Dalal, A.; Fisher, C.K.; Hitzler, P.; Janowicz, K.; et al. The KnowWhereGraph ontology. *J. Web Semant.* **2025**, *84*, 100842. [CrossRef]
11. Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; Wu, X. Unifying large language models and knowledge graphs: A roadmap. *IEEE Trans. Knowl. Data Eng.* **2024**, *36*, 3580–3599. [CrossRef]
12. Zhu, Y.; Wang, X.; Chen, J.; Qiao, S.; Ou, Y.; Yao, Y.; Deng, S.; Chen, H.; Zhang, N. LLMs for knowledge graph construction and reasoning: Recent capabilities and future opportunities. *World Wide Web* **2024**, *27*, 58. [CrossRef]
13. Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitansky, D.; Ness, R.O.; Larson, J. From local to global: A Graph RAG approach to query-focused summarization. *arXiv* **2024**, arXiv:2404.16130. [CrossRef]
14. Sun, Q.; Luo, Y.; Zhang, W.; Li, S.; Li, J.; Niu, K.; Kong, X.; Liu, W. Docs2KG: Unified knowledge graph construction from heterogeneous documents assisted by large language models. *arXiv* **2024**, arXiv:2406.02962. [CrossRef]
15. Chen, B.; Bertozzi, A.L. AutoKG: Efficient automated knowledge graph generation for language models. In *Proceedings of the 2023 IEEE International Conference on Big Data (BigData), Sorrento, Italy, 15–18 December 2023*; IEEE: Piscataway, NJ, USA, 2023; pp. 3117–3126. [CrossRef]
16. Lairgi, Y.; Moncla, L.; Cazabet, R.; Benabdeslem, K.; Cleau, P. iText2KG: Incremental knowledge graphs construction using large language models. In *Web Information Systems Engineering—WISE 2024*; Lecture Notes in Computer Science; Springer: Singapore, 2025; Volume 15439, pp. 214–229. [CrossRef]
17. Mavroudis, V. LangChain v0.3. *Preprints* **2024**, 202411.0566. [CrossRef]
18. LangChain AI. Deep Agents. Available online: https://github.com/langchain-ai/deepagents (accessed on 5 August 2026).
19. Liu, N.F.; Lin, K.; Hewitt, J.; Paranjape, A.; Bevilacqua, M.; Petroni, F.; Liang, P. Lost in the middle: How language models use long contexts. *Trans. Assoc. Comput. Linguist.* **2024**, *12*, 157–173. [CrossRef]
20. Elmagarmid, A.K.; Ipeirotis, P.G.; Verykios, V.S. Duplicate record detection: A survey. *IEEE Trans. Knowl. Data Eng.* **2007**, *19*, 1–16. [CrossRef]
21. Xue, B.; Zou, L. Knowledge graph quality management: A comprehensive survey. *IEEE Trans. Knowl. Data Eng.* **2023**, *35*, 4969–4988. [CrossRef]
22. Liu, Y.; Iter, D.; Xu, Y.; Wang, S.; Xu, R.; Zhu, C. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Singapore, 6–10 December 2023*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2023; pp. 2511–2522. [CrossRef]
23. Reimers, N.; Gurevych, I. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, 3–7 November 2019*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 3982–3992. [CrossRef]
24. Zhang, R.; Trisedya, B.D.; Li, M.; Jiang, Y.; Qi, J. A benchmark and comprehensive survey on knowledge graph entity alignment via representation learning. *VLDB J.* **2022**, *31*, 1143–1168. [CrossRef]
25. DeepSeek-AI. DeepSeek-V4: Towards highly efficient million-token context intelligence. *arXiv* **2026**, arXiv:2606.19348. [CrossRef]

26. Nickel, M.; Murphy, K.; Tresp, V.; Gabrilovich, E. A review of relational machine learning for knowledge graphs: From multi-relational link prediction to automated knowledge graph construction. *Proc. IEEE* **2016**, *104*, 11–33. [CrossRef]
27. Huang, J.; Hou, W.; Zhang, S.; Du, Z.; Yin, J. Dam defects ontology and knowledge graph construction from multi-source hazard data using large language models. *Adv. Eng. Inform.* **2026**, *76*, 104987. [CrossRef]
28. Pu, J.; Guo, Y.; Huang, S.; Zhou, C.; Tian, J.; Ma, L. MA-APKG: A multi-agent LLM-based method for construction aviation product assembly process knowledge graph. *J. Manuf. Syst.* **2026**, *86*, 1112–1138. [CrossRef]
29. Borrelli, A.; De Angelis, V.; Izzo, S.; Giampaolo, F.; Piccialli, F. MAKES-QA: A multi-agent framework for knowledge graph construction and enrichment over scientific literature for question answering. *Expert Syst. Appl.* **2026**, *320*, 132156. [CrossRef]
30. Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; et al. A survey on evaluation of large language models. *ACM Trans. Intell. Syst. Technol.* **2024**, *15*, 39. [CrossRef]
31. Wilson, R.S.I.; Goonetillake, J.S.; Indika, W.A.; Ginige, A. A conceptual model for ontology quality assessment: A systematic review. *Semant. Web* **2023**, *14*, 1051–1097. [CrossRef]
32. Du, S.; Zhao, J.; Shi, J.; Xie, Z.; Jiang, X.; Bai, Y.; He, L. A survey on the optimization of large language model-based agents. *ACM Comput. Surv.* **2026**, *58*, 223. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.