Journal of Electronic Research and Application, 2026, Volume 10, Issue 1
http://ojs.bbwpublisher.com/index.php/JERA
ISSN Online: 2208-3510
ISSN Print: 2208-3502

Overview of the Integration of Large Language Models, Knowledge Graphs, and GraphRAG, along with Research on Their Industrial Applications

Xian Ye*

School of Mine Safety, North China Institute of Science and Technology, Yanjiao 101601, Hebei, China

**Author to whom correspondence should be addressed.*

Abstract: In recent years, Large Language Models (LLMs) have rapidly advanced in language understanding, reasoning, and generation, and are increasingly adopted as the "brain" of industrial intelligent systems. Nevertheless, in high-risk and strongly regulated domains they still exhibit hallucination, weak domain grounding, limited interpretability, and privacy as well as security constraints. Knowledge graphs (KGs) encode domain entities, relations, rules, and events explicitly, providing controllable semantics and an explainable reasoning substrate. Retrieval-augmented generation (RAG) injects external evidence into LLM prompting, while GraphRAG further introduces graph indexing and community-level retrieval to preserve global structure and support multi-hop reasoning. This review summarizes the evolution of LLMs, KG modeling and extraction, GraphRAG mechanisms, and a general fusion framework. Typical industrial applications are surveyed, and a coal mine flood emergency plan generation and evaluation approach is discussed to illustrate the practical value of graph-grounded large models. KG-enhanced retrieval also supports provenance tracking, allowing industrial users to audit the evidence behind model outputs.

Keywords: Large language model; Knowledge graph; GraphRAG; Retrieval-augmented generation; Coal mine flood; Emergency management

Online publication: February 13, 2026

1. Introduction

Large model-driven general AI has shifted industrial systems from pipeline-style NLP modules to unified model-centric architectures. LLMs offer strong zero-/few-shot transfer, natural human-computer interaction, and the ability to synthesize multi-source information, making them suitable as universal interfaces for production, governance, and safety services. However, industrial tasks emphasize factual correctness, timeliness, traceability, and compliance. Pure LLMs cannot reliably store long-tail or rapidly changing professional knowledge, and their black-box generation may lead to unsafe suggestions in critical decision processes. Existing work therefore combines LLMs with external knowledge. RAG improves factuality by retrieving evidence from document stores, but vector-only retrieval may fragment long documents and ignore structural relations, limiting global understanding and multi-hop reasoning. KGs provide structured semantics and explicit causal/procedural connections, yet traditional KG-based QA often depends on brittle symbolic pipelines and heavy manual schema design. GraphRAG integrates graph structure into retrieval and reasoning, enabling global-aware evidence selection and explainable paths, and thus becomes a promising foundation for trustworthy industrial LLM systems. This paper reviews key techniques and proposes an LLM + KG + GraphRAG fusion route for industry scenarios, with a focus on coal mine flood emergency management.

2. Development and industrialization trends of large language model technology

2.1. Evolution of general large models

Since the transformer architecture, pretraining on massive corpora plus instruction tuning and alignment (e.g., RLHF or preference optimization) has produced foundation models with strong reasoning and generation capabilities. Research trajectories include scaling laws, long-context modeling (positional encoding variants, sparse attention, retrieval-memory hybrids), mixture-of-experts for efficiency, distillation/quantization and speculative decoding for deployment, and tool/function calling for structured actions. Inference optimization and hardware acceleration have reduced serving cost, while multimodal models extend LLMs to images, audio, and sensor-like signals, supporting richer industrial inputs such as inspection photos, charts, or monitoring reports. Despite this progress, general models remain weak on domain terminology, localized standards, and task-specific workflows. Their performance degrades under distribution shift (new devices, new policies, unseen hazards), and safety alignment for professional scenarios is still insufficient. Therefore, domain customization and reliable knowledge grounding are indispensable for industrial use.

2.2. Industry large models and task customization practice

Industrial adaptation usually follows two intersecting paths: domain fine-tuning and knowledge-grounded augmentation. Fine-tuning uses curated professional corpora and parameter-efficient methods (LoRA, adapters) to strengthen domain vocabulary, response style, and constraint compliance; continual learning is used to follow evolving regulations and incident patterns. Knowledge augmentation connects LLMs to enterprise databases and document stores through RAG, KGs, and agent workflows so that answers are evidence-based and controllable. In industry and smart-city contexts, private data are heterogeneous and sensitive; deployment thus stresses on-premise or secure-cloud serving, access control, and data-quality governance. For mine safety, models are customized to understand emergency plans, accident reports, geological texts, and monitoring indicators, and to generate disposal suggestions aligned with standards. LLMs also assist KG construction by extracting entities/relations from long technical documents, reducing manual labeling cost and enabling faster knowledge updates. For library, information, and scientific services, LLMs act as semantic organizers and report writers on top of curated KGs, improving literature retrieval and thematic analysis.

3. Knowledge graph and structured knowledge extraction technology

3.1. Modeling elements of industry knowledge graph

An industry KG encodes domain entities, attributes, relations, and events under a shared ontology. Entities may represent equipment, hazards, resources, organizations, policies, or cases. Relations include hierarchical, causal, temporal, spatial, procedural, and constraint links. Industrial KGs emphasize multi-source fusion (plans, logs, sensors), spatiotemporal evolution, uncertainty/confidence, and versioned updates, making them suitable for decision support, risk assessment, and post-event auditing. Schema design usually aligns with standards or regulations, and quality assurance relies on constraint rules, consistency checking, and iterative expert validation.

3.2. Knowledge extraction based on pre-trained models and LLMs

KG construction requires NER, relation/event extraction, entity linking, and coreference resolution. Traditional supervised IE provides stable baselines but is costly to label. Pretrained encoders and seq2seq extractors reduce data needs, while LLM prompting enables few-shot or open IE, structured triple/JSON outputs, and automatic schema induction. In practice, hybrid pipelines are common, where statistical or rule extractors provide high-precision seeds; LLMs expand coverage by self-asking and self-correcting; normalization and deduplication fuse results; and human auditing remains the final safety gate. LLMs are also used for knowledge completion and conflict detection, supporting continuous KG evolution.

4. Technical route of 4RAG and GraphRAG

4.1. Vector RAG framework and its limitations

Vector RAG embeds text chunks, retrieves top-k similar units, and feeds them into LLM prompts. It enhances factuality for single-hop questions but struggles with complex industrial corpora as follows:

(1) Fixed chunking breaks long-range coherence, and top-k evidence may miss the global narrative;
(2) Embeddings emphasize topical similarity but may overlook explicit causal or procedural relations;
(3) Retrieved snippets can be redundant or inconsistent, so the model may still hallucinate when integrating evidence.

These limits are evident in tasks requiring cross-document causal tracing, such as policy coordination, equipment fault localization, or emergency decision making.

4.2. Core idea and implementation process of GraphRAG

GraphRAG constructs a graph over text units and/or KG nodes to preserve structure during retrieval. Offline, it does as follows:

(1) Segments documents into base units;
(2) Extracts entities and relations for linking;
(3) Builds graph edges using KG relations and/or semantic similarity;
(4) Applies community detection to obtain topic-level subgraphs;
(5) Summarizes each community to form a compact global representation.

Online, a query retrieves relevant community summaries to locate the correct knowledge region, and then expands into local nodes and paths for fine-grained evidence. By combining global summaries with local facts, GraphRAG improves multi-hop retrieval stability, reduces redundancy, and provides interpretable evidence chains.

4.3. Integration mode of GraphRAG and knowledge graph

Three integration patterns are observed as listed:

(1) Document-graph GraphRAG builds graphs mainly from unstructured corpora for domains without mature

KGs; it is easy to deploy but less controllable;

(2) KG-enhanced GraphRAG uses a domain KG as the core graph and links documents to KG nodes, enabling rule-consistent retrieval and reasoning;

(3) Hybrid GraphRAG merges document graphs and KGs into a heterogeneous graph, combining semantic proximity search with symbolic multi-hop traversal.

The choice depends on KG availability, update frequency, interpretability requirements, and engineering cost.

5. General fusion framework of LLM + KG + GraphRAG

5.1. Offline stage: Multi-source knowledge modeling and graph index construction

The offline stage ingests domain plans, regulations, cases, technical manuals, databases, and logs. Documents are cleaned and segmented into base text units; entities, relations, and events are extracted to build or update the KG. Text units are embedded and linked to KG nodes, producing a heterogeneous retrieval graph. Community detection and summarization create a multi-level index supporting global overview plus local evidence access. Graph and vector indexes are stored for fast online retrieval and are periodically refreshed as new data arrives.

5.2. Online stage: Query comprehension, graph retrieval, and generative reasoning

Given a user query, the system performs intent and entity parsing, retrieves relevant community summaries or KG subgraphs, and expands along graph paths to gather supporting text. Evidence is re-ranked and organized into structured context (facts, relations, timelines, procedures). The LLM then generates answers under evidence constraints, optionally using self-verification, rule checks on KG paths, and tool calls (database lookup, calculation, simulation). Outputs can include cited evidence and reasoning paths for expert auditing and downstream execution.

5.3. Coal mine flood emergency plan generation and evaluation

For coal mine flood disasters, knowledge sources include emergency plans, laws, historical accidents, hydrology/geology reports, and monitoring indicators. A safety schema models water-inrush sources, precursor signals, affected equipment, response teams, and disposal actions with causal and temporal links. GraphRAG retrieves scenario-matched subgraphs and similar accident communities, providing grounded context on hazard evolution and proven measures. The LLM produces structured, stepwise emergency plans (monitoring–alarm–confirmation–control–evacuation–rescue–recovery). Plan evaluation can leverage KG rule compliance, evidence coverage, and similarity to validated historical cases.

6. Overview of industry applications

6.1. Analysis of government data governance and policy coordination

Zhu and Qin built a policy knowledge graph and adopted GraphRAG for government data governance and policy coordination [1]. Their method converts policy texts into a unified graph space, uses community division to cluster policy topics, and performs path-based retrieval to trace cross-policy dependencies. GraphRAG locates relevant clauses quickly and provides multi-document evidence chains, while the LLM assists semantic interpretation and policy-effect analysis, improving policy synergy and execution efficiency.

6.2. Industrial private data and smart city knowledge services

Hu *et al.* proposed an ALBERT-XL pipeline for constructing KGs from industrial private data, enabling knowledge extraction without exposing raw sensitive logs [2]. Based on industrial KGs, Li *et al.* implemented a GraphRAG-based enterprise private knowledge base for construction projects, enhancing cross-document retrieval, fault troubleshooting, and experience reuse [3]. In smart-city governance, similar KG + GraphRAG services unify heterogeneous data for infrastructure management, incident response, and citizen Q&A, and LLMs translate retrieved evidence into actionable and understandable recommendations.

6.3. Library and information services and domain knowledge services

Xie *et al.* introduced TTKE-LLM, using LLM prompting plus engineering constraints to extract tourism entities and relations and to accelerate KG construction [4]. Ma *et al.* combined LLMs with GraphRAG to generate graph-guided abstracts for scientific literature, turning scattered papers into structured knowledge and improving literature navigation [5]. Compared with keyword-based intelligence analysis, KG + GraphRAG enables semantic clustering, relationship tracing, and explainable answers to complex queries, while LLMs synthesize readable intelligence reports.

6.4. Agricultural pest and disease control and expert system

Wu *et al.* developed a rice pest-and-disease expert system by coupling agricultural KGs with LLM reasoning [4]. Symptoms, varieties, pesticides, and environmental factors are encoded as KG nodes/relations, and GraphRAG retrieves multi-hop evidence for diagnosis. The LLM serves as a hypothesis generator and interactive expert, revising conclusions with retrieved evidence to improve interpretability and practical usability.

6.5. Power system, autonomous driving and other scenarios

In power systems, Liu *et al.* applied LLMs to electric-vehicle charging and swapping load forecasting [5]. With KG/RAG context on equipment states, user behavior, and grid constraints, their approach yields more robust and interpretable forecasts. In autonomous driving, Song *et al.* summarized large-model decision/planning progress, and Wu *et al.* further integrated vehicle KGs with LLMs to support multi-scenario decision making and safer reasoning [6,7]. Ai *et al.* introduced a GraphRAG-based assistant for spacecraft fault localization, showing that graph-guided retrieval can improve cross-document fault reasoning [8].

6.6. Mine accidents and safety management

For mining, Zhang *et al.* constructed a mine-accident KG using LLM-assisted extraction, capturing environment, causal factors, losses, and response measures, which supports case retrieval and risk pattern mining [9]. Xu *et al.* designed an LLM-based coal-mine safety assistant; combined with KG/GraphRAG retrieval, it assists hazard identification, case recall, and on-site decision suggestions with traceable evidence [9]. These results indicate that graph-grounded LLM systems can improve the safety training effectiveness and emergency disposal efficiency.

7. Key challenges and development trends

7.1. Challenges at the data and graph level

Industrial corpora are noisy and heterogeneous, and KG construction faces schema inconsistency, entity ambiguity, sparse labeling, and high expert cost. Continuous updates, temporal reasoning, and uncertainty modeling remain difficult, and extraction errors may propagate into retrieval and generation. Future systems need automatic quality assessment, conflict resolution, and efficient human-in-the-loop editing to maintain graph reliability.

7.2. Challenges at the model and system levels

GraphRAG introduces extra latency, storage, and engineering complexity. Retrieval quality depends on graph construction, community division, and linking accuracy, while LLMs may still hallucinate when evidence is weak or contradictory. Efficient deployment therefore requires lightweight heterogeneous graph indexing, caching, adaptive retrieval depth, and model compression, together with strict privacy/access control for sensitive data.

7.3. Evaluation, compliance and future directions

Benchmarks for graph-grounded generation are scarce. Evaluation should measure factuality, evidence faithfulness, reasoning-path correctness, robustness, and expert trust, ideally via scenario-based tests. Compliance concerns cover data security, copyright, and domain safety regulations. Future research will likely focus on automated schema learning, spatiotemporal and multimodal GraphRAG, causal/functional graphs for decision support, joint graph-LLM training, and agentic multi-step planning-retrieval-generation with human oversight.

8. Conclusion

LLMs provide powerful language interfaces, KGs supply explicit and explainable domain knowledge, and GraphRAG enables global-aware multi-hop retrieval and path-based reasoning. Their fusion improves grounding, interpretability, and robustness for industrial intelligence. In coal mine flood emergency management, a KG-enhanced GraphRAG framework can generate and evaluate structured emergency plans aligned with regulations and verified cases, offering a feasible direction for intelligent emergency disposal. Overall, the LLM + KG + GraphRAG paradigm offers a balanced path between neural flexibility and symbolic controllability for future safety-critical industrial AI.

Disclosure statement

The authors declare no conflict of interest.

References

[1] Zhu X, Qin D, 2025, From Evolution to Coordination: Exploring Policy Effectiveness-Enhancement Paths for Government Data Governance under the GraphRAG Framework. Information and Documentation Services, 46(5): 89–101.

[2] Hu D, Wu X, Liu X, et al., 2025, Knowledge Graph Construction Method for Industrial Private Data Based on ALBERT-XL. Smart Cities, 11(7): 10–14.

[3] Li G, Wu Z, Huang J, 2025, Research and Implementation of an Enterprise-Level Private Knowledge Base in the Construction Domain Based on GraphRAG. Computer Knowledge and Technology, 21(20): 23–25.

[4] Wu H, Chen M, Guo Q, 2025, Rice Pest and Disease Expert System Based on Large Language Models and Knowledge Graphs. Transactions of the Chinese Society of Agricultural Engineering, 41(22): 244-255.

[5] Song Z, Chen J, Jiang L, et al., 2025, Research Progress Review of Large-Model-Based Autonomous Driving

Decision-Making and Planning. Automotive Technology, 2025(10): 21–31.

[6] Wu S, Jiang H, Huang K, et al., 2025, Vehicle Decision-Making System Based on Knowledge Graphs and Large Language Models. Journal of Beijing Institute of Technology, 6–12.

[7] Ai S, He Y, Zhang W, et al., 2025, A GraphRAG-Based Method for Assisted Spacecraft Fault Localization. Spacecraft Engineering, 34(4): 84–90.

[8] Zhang P, Sheng L, Wang W, et al., 2025, Construction of a Mine Accident Knowledge Graph Based on Large Language Models. Industry and Mine Automation, 51(2): 76–83.

[9] Xu J, Gong Q, Xiang Y, et al., 2024, Design and Application of a Coal Mine Safety Assistant based on LLM. Inner Mongolia Coal Economy, 2024(1): 133–135.