



KGDRL-HAO: A knowledge reasoning-guided deep reinforcement learning framework for highway alignment optimization

Downloaded from: <https://research.chalmers.se>, 2026-10-15 01:09 UTC

Citation for the original published paper (version of record):

Gao, J., Yu, B., Chen, Y. et al (2026). KGDRL-HAO: A knowledge reasoning-guided deep reinforcement learning framework for highway alignment optimization. *Computer-Aided Civil and Infrastructure Engineering*, 51. <http://dx.doi.org/10.1016/j.cacaie.2026.100231>


N.B. When citing this work, cite the original published paper.



Contents lists available at ScienceDirect

Computer-Aided Civil and Infrastructure Engineering

journal homepage: www.sciencedirect.com/journal/computer-aided-civil-and-infrastructure-engineering



Research Article

KGDL-HAO: A knowledge reasoning-guided deep reinforcement learning framework for highway alignment optimization

Jianqiang Gao^a, Bo Yu^{b,*}, Yuren Chen^b, Kun Gao^c, Shan Bao^{d,e}, Junhua Wang^b

^a School of Highway, Chang'an University, Middle Section of Nan Erhuan Road, Xi'an, Shaanxi, 710064, China

^b Key Laboratory of Road and Traffic Engineering of the Ministry of Education, College of Transportation, Tongji University, 4800 Cao'an Highway, Shanghai, 201804, China

^c Department of Architecture and Civil Engineering, Chalmers University of Technology, Gothenburg, SE-412 96, Sweden

^d Industrial and Manufacturing Systems Engineering Department, University of Michigan-Dearborn, 4901 Evergreen Rd, Dearborn, MI, 48128, USA

^e University of Michigan Transportation Research Institute, 2901 Baxter Rd, Ann Arbor, MI, 48109-2150, USA

ARTICLE INFO

Keywords:

Highway alignment optimization
Knowledge reasoning
Graph neural networks
Large language models
Deep reinforcement learning

ABSTRACT

Highway geometric design is a continuous alignment optimization process involving geometric constraints and multi-objective trade-offs. However, current highway alignment optimization methods are often limited by fragmented and unstructured knowledge representations, and the lack of prior knowledge guidance results in poor exploration efficiency and restricted practical applicability. To address these issues, this study proposes a knowledge reasoning-guided deep reinforcement learning framework for highway alignment optimization, named KGDL-HAO. The KGDL-HAO employs graph neural networks based on multi-hop inference and large language models with contextual understanding to capture semantic dependencies among design variables, constraint conditions, and optimization objectives, outperforming baseline knowledge reasoning models. Then, by incorporating structured prior knowledge into the soft actor-critic learning process, the KGDL-HAO enhances exploration efficiency and decision consistency, reducing the weighted total cost by 6.4% while satisfying design constraints. The findings would facilitate intelligent and interpretable highway geometric design through knowledge-augmented optimization and natural language interaction.

1. Introduction

Highways constitute the backbone of national transportation systems, acting as primary arteries for economic productivity and social development (Han et al., 2023). Highway geometric design is a fundamental phase of highway planning and construction, aimed at determining a three-dimensional alignment that ensures safe, efficient, and sustainable mobility between termini (Findley, 2016). Highway geometric design faces multifaceted challenges arising from diverse terrain conditions, interdependent geometric constraints, and competing design objectives (Casal et al., 2017; Yu et al., 2018). However, since the number of possible alignments between termini is virtually infinite, relying solely on engineers' experience and intuition may lead to the omission of many promising alternatives (Hirpa et al., 2016). Therefore, accurate and efficient alignment decision-making is critical for highway geometric design.

Geometric design knowledge is the practical foundation of highway alignment optimization, defining the design boundaries and decision logic. Design boundaries are primarily defined by explicit knowledge such as design standards and technical specifications, which prescribe permissible value ranges, applicable conditions, and interdependencies among geometric elements (AASHTO, 2018). Design decision-making further relies on expert experience and historical cases to support multi-objective trade-offs among candidate design alternatives within defined performance goals and design constraints (Fitzpatrick et al., 2021). This domain knowledge typically exhibits heterogeneous, fragmented, and unstructured characteristics, which severely hinder its extension, modification, and dynamic updating (Pu et al., 2019). Specifically, design specifications define criteria for curve radius, grade, sight distance, and superelevation, which are often converted into geometric constraints in automated alignment design (Sushma and Maji, 2020). Engineering cases and expert experience involve project-specific

* Corresponding author.

E-mail addresses: 260010@chd.edu.cn (J. Gao), boyu@tongji.edu.cn (B. Yu), chenyr@tongji.edu.cn (Y. Chen), gkun@chalmers.se (K. Gao), shanbao@umich.edu (S. Bao), benwjh@163.com (J. Wang).

<https://doi.org/10.1016/j.cacaie.2026.100231>

Received 22 June 2026; Received in revised form 13 September 2026; Accepted 14 September 2026

Available online 16 September 2026

1093-9687/© 2026 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

constraints and heuristic judgments regarding terrain, land use, earthwork, emissions, cost, and three-dimensional alignment coordination (Gao et al., 2022). Existing studies have commonly translated design standards into explicit IF–THEN rules for compliance checking, but rule-based methods show limited adaptability to complex and context-dependent design situations (Makke & Chawla, 2024; Boden et al., 2022). With formal definitions of domain concepts, properties, and semantic relationships, ontology transforms design standards and experiential knowledge into computable constraints and logic (Du et al., 2025). Building on ontology-based semantic modeling, the knowledge graph integrates symbolic representation and vector-based modeling to further enhance reasoning and analytics, and has gained increasing attention in civil engineering domains (Wang et al., 2021). For example, a bridge maintenance knowledge graph built from maintenance reports is used for component category classification and potential maintenance relationship prediction, thereby improving knowledge retrieval and decision-making efficiency (Gao et al., 2024a). Therefore, leveraging the strengths of semantic integration and machine interpretability, the knowledge graph presents a promising solution for structured representation of highway geometric design knowledge.

Knowledge reasoning bridges the gap between problem-driven design contexts and performance-based decision-making in highway geometric design. The introduction of the performance-based approach has shifted the paradigm of highway geometric design from the mechanical application of prescriptive standards to a problem-driven process, requiring designers to prioritize competing goals and make informed decisions and tradeoffs that best meet the needs of all transportation modes (Easa & Yan, 2019). Previous studies have primarily relied on symbolic reasoning methods to support route optimization decisions, whereas they often lack consideration of the interactions between design context and complex constraints (Liu et al., 2022; Mohammed et al., 2019). For example, a study developed a rule-based expert system that formalizes design specifications and expert knowledge into IF–THEN rules to evaluate and rank alternative road alignments based on factors such as traffic conditions, terrain, and sight distance (Syamsunur et al., 2011). In the study of Jung et al. (Jung et al., 2020), a case-based reasoning framework is employed to retrieve and adapt historical cases, reusing multi-objective weightings to balance safety and efficiency in cost estimation and decision-making. Neural approaches are increasingly used in knowledge reasoning to enhance the efficiency of domain knowledge representation and the robustness of decision support (Chen et al., 2020). Graph neural networks (GNN) have been applied to capture high-order semantic dependencies among entities through multi-hop message passing and perform structured reasoning over a knowledge graph (Wu et al., 2020). Another study leverages the contextual semantic understanding and language generation capability of a large language model (LLM) to infer causal reasoning paths, while the model often produces hallucinated content that is inconsistent with the input facts (Guo et al., 2025). However, geometric design must be understood, taught, executed, and communicated as iterative in nature with performance at the center of all iterations (Fitzpatrick et al., 2021). It is imperative to incorporate interpretable and reliable knowledge reasoning into highway geometric design to provide traceable evidence and informed guidance for iterative design decisions.

Different methods have been employed for highway alignment optimization and can be classified into four major categories, namely analytical and enumerative methods, metaheuristic methods, hybrid optimization methods, and deep reinforcement learning methods. Variational calculus, enumeration methods, and dynamic programming are primarily used for highway alignment optimization problems characterized by small scales and simple terrain conditions (Casal et al., 2017; Zhang et al., 2024). As design complexity and constraints increase, metaheuristic methods such as genetic algorithms, particle swarm optimization, and ant colony optimization are widely adopted to address nonlinear and multi-objective alignment optimization problems

(Sushma et al., 2022). Although these methods offer good global search capabilities, they tend to exhibit strong dependence on initial solutions and are prone to local optima (Li et al., 2016; Pu et al., 2019). A distance-transform-based three-dimensional alignment generation method is proposed and further integrated with genetic algorithms to efficiently optimize highway alignments in mountainous terrain (Li et al., 2017). Additionally, the deep deterministic policy gradient algorithm is applied to construct a policy-learning model that enables agents to learn optimization strategies through interactions with simulated environments (Gao et al., 2022). Recent studies have further advanced intelligent alignment optimization by incorporating sequential exploration, multi-objective reinforcement learning, and fine-grained risk-cost optimization. For example, sequential exploration has been used for horizontal road alignment optimization (Zhang et al., 2023), bi-objective deep reinforcement learning has been applied to low-carbon high-speed railway alignment design (He et al., 2023), and fine-grained optimization has been conducted for railway alignment design considering construction cost and risk near existing railways (Gao et al., 2024b). Due to the absence of explicit prior knowledge guidance, these methods exhibit low search efficiency under complex terrain and multidimensional constraints, often generating alignments that fail to satisfy engineering feasibility requirements. Therefore, incorporating human design knowledge is essential for improving exploration efficiency and decision consistency in highway alignment optimization.

Given the above, the fragmented and unstructured application of highway geometric design knowledge hinders the incremental refinement and evolution of domain knowledge in computational reasoning. In addition, existing numerical optimization methods struggle to incorporate human prior knowledge, thereby limiting the practical applicability of the optimal alignment solutions. To fill the above research gaps, this study aims to propose a knowledge reasoning-guided deep reinforcement learning (DRL) framework for highway alignment optimization, named KGDR-L-HAO. A hierarchical highway geometric design knowledge graph (HGD-KG) is constructed to semantically represent design specifications, engineering cases, and expert experience. Then, using the domain-specific knowledge corpus, a knowledge reasoning model for highway geometric design is built by the graph neural retrieval for LLM reasoning algorithm. The structured inferred knowledge is further incorporated as expert policies into a highway alignment optimization model based on the soft actor-critic algorithm. Finally, a real-world case study is conducted to validate the enhanced exploration efficiency and decision consistency of the proposed framework. The contributions of the study are threefold:

- (1) Through semantic networks and graph-based structures, the proposed HGD-KG integrates heterogeneous knowledge from design specifications, engineering cases, and expert experience, providing a more comprehensive and computable knowledge representation for highway geometric design.
- (2) Based on graph neural retrieval for LLM reasoning, this study proposes an interpretable knowledge reasoning model for highway geometric design. The model effectively captures implicit multi-hop semantic dependencies among design variables, constraint conditions, and optimization objectives, providing traceable and reliable reasoning paths to support highway geometric design decisions.
- (3) A knowledge reasoning-guided deep reinforcement learning framework is developed for highway alignment optimization, in which inferred prior knowledge is incorporated into the optimization process to guide feasible exploration and improve decision consistency, thereby enhancing alignment optimization under complex design situations.

2. Problem formulation

Highway geometric design can be formulated as a highly con-

strained, nonlinear, and non-differentiable multi-objective optimization problem (Wan et al., 2025). The main goal of this work is to determine an optimal highway alignment connecting two given terminal points $a = (x_a, y_a, z_a) \in \mathbb{R}^3$ and $b = (x_b, y_b, z_b) \in \mathbb{R}^3$. As shown in Fig. 1, a highway alignment is essentially a three-dimensional curve, jointly defined by the horizontal alignment, the vertical alignment, and the cross-sections. Fundamentally, this optimization problem consists of four key components, including the design variables that determine the layout, a mathematical model that generates the alignment, the constraints that define feasibility conditions, and the objective functions that measure the achievement of the design objectives.

2.1. Design variables and constraints

2.1.1. Horizontal alignment

The horizontal alignment defines the geometric layout of a highway in the horizontal plane, consisting of a sequence of tangents, circular curves, and transition curves (i.e., spiral curves). As illustrated in Fig. 1 (b), the horizontal alignment is represented by a series of horizontal points of intersection (HPIs) together with the associated lengths of spiral and circular curves. Let $a_{2D} = (x_a, y_a) \in \mathbb{R}^2$ and $b_{2D} = (x_b, y_b) \in \mathbb{R}^2$ denote the two terminal points of the alignment projected on the horizontal plane. When the horizontal alignment consists of N HPIs, indexed by $n = 1, 2, \dots, N$, the horizontal design-variable vector x^N and the corresponding admissible design space X_{ad}^N are defined as:

$$x^N = (x_1, y_1, L_{sf,1}, R_{h,1}, L_{sb,1}, \dots, x_N, y_N, L_{sf,N}, R_{h,N}, L_{sb,N}) \in \mathbb{R}^{5N}, \quad (1)$$

$$X_{ad}^N = \left\{ x^N \in \mathbb{R}^{5N} \left| \begin{array}{l} \theta_n \geq \theta^{\min}, R_{h,n} \geq R_h^{\min}, \\ L_{sf,n} \geq L_s^{\min}, L_{sb,n} \geq L_s^{\min}, \\ x_a < x_n < x_b, y_a < y_n < y_b, \\ \forall n \in \{1, \dots, N\} \end{array} \right. \right\}, \quad (2)$$

where (x_n, y_n) denote the horizontal coordinates of the n -th intersection point, $L_{sf,n}$ and $L_{sb,n}$ denote the lengths of the forward and backward spiral curves, $R_{h,n}$ denotes the radius of circular arc associated with the n -th intersection point, θ_n is the deflection angle between the back and

forward tangents at the n -th intersection point, $L_s^{\min}$ denotes the minimum allowable spiral curve length, and $\theta^{\min}$ and $R_h^{\min}$ denote the minimum deflection angle and circular arc radius, respectively.

Given a set of design variables x^N , the corresponding horizontal alignment curve $C_{2D}^{x^N} \subset \mathbb{R}^2$ can be parameterized as:

$$\sigma_{2D}^{x^N} : s \in [0, L(x^N)] \mapsto \sigma_{2D}^{x^N}(s) = (\sigma_x(s), \sigma_y(s)) \quad (3)$$

where $L(x^N)$ denotes the total length of the horizontal alignment, and σ_x, σ_y denotes the coordinate functions computing the horizontal coordinates of the horizontal alignment with respect to the arc-length parameter s , as detailed in Appendix A1.

2.1.2. Vertical alignment

Once the horizontal alignment is determined, the vertical alignment is used to describe the elevation profile along the highway centerline. A typical vertical alignment is composed of straight segments known as vertical tangents and parabolic curves, commonly referred to as crest and sag curves. As illustrated in Fig. 1(c), the design of the vertical profile must satisfy two design criteria: (1) the vertical alignment should follow a standard geometric sequence of tangent-parabola-tangent; and (2) both the start and end sections of the vertical alignment must be tangent grades. Parabolic curves are inserted between adjacent tangents to ensure smooth elevation transitions and to mitigate abrupt slope changes. Let the vertical alignment consist of M vertical points of intersection (VPIs), indexed by $m = 1, 2, \dots, M$. The vertical design-variable vector y^M and the corresponding admissible design space Y_{ad}^M are defined by:

$$y^M = (s_1, e_1, R_{v,1}, \dots, s_M, e_M, R_{v,M}) \in \mathbb{R}^{3M}, \quad (4)$$

$$Y_{ad}^M = \left\{ y^M \in \mathbb{R}^{3M} \left| \begin{array}{l} 0 < s_1 < \dots < s_M < 1, \\ G_v^{\min} \leq |G_{v,m}| \leq G_v^{\max}, \\ R_{v,m} \geq R_v^{\min}, \\ \forall m \in \{1, \dots, M\} \end{array} \right. \right\}, \quad (5)$$

where $s_m \in (0, 1)$ denotes the normalized arc length along the horizontal alignment, e_m denotes the elevation at the m -th VPI, $G_{v,m}$ denotes the




Fig. 1. Schematic diagram of three-dimensional highway alignment.

vertical grade of the m -th tangent section, and $R_{v,m}$ is the design parameter of the m -th vertical curve. The parameters $G_v^{\min}$ and $G_v^{\max}$ define the allowable range of grades, while $R_v^{\min}$ denotes the minimum permissible vertical-curve radius according to design specifications.

Given a set of vertical alignment design variables y^M , the corresponding elevation profile $C_{2D}^M \subset \mathbb{R}^2$ can be parameterized as:

$$\sigma_{2D}^M : s \in [0, L(x^N)] \mapsto \sigma_{2D}^M(s) = (s, \sigma_z(s)) \quad (6)$$

where $\sigma_z(s)$ denotes the elevation function that computes the vertical elevation along the vertical alignment with respect to the arc-length parameter s , as provided in Appendix A2.

2.1.3. Cross-section

A three-dimensional highway alignment is constructed by superimposing the two-dimensional horizontal and vertical alignments. The cross-section refers to the normal plane of the horizontal alignment at any given station along the highway centerline. As illustrated in Fig. 1 (d), the cross-section is classified into three typical types, including subgrade, bridge, and tunnel, according to the structural form of the highway at the corresponding location. In this study, the geometry of each cross-section is characterized by the pavement width, superelevation rate, and section type. The superelevation is defined as the transverse slope of the roadway surface, while the section type is represented by an integer code to distinguish different structural forms. Given K cross-section control points along the horizontal alignment, indexed by $k = 1, 2, \dots, K$, the cross-sectional design-variable vector z^K and the corresponding admissible design space Z_{ad}^K are formulated as:

$$z^K = (s_1, W_{c,1}, S_{c,1}, T_{c,1}, \dots, s_K, W_{c,K}, S_{c,K}, T_{c,K}) \in \mathbb{R}^{4K}, \quad (7)$$

$$Z_{ad}^K = \left\{ z^K \in \mathbb{R}^{4K} \left| \begin{array}{l} 0 < s_1 < \dots < s_K < 1, \\ h_{f,k} \leq h_f^{\max}, \\ h_{c,k} \leq h_c^{\max}, \\ \forall k \in \{1, \dots, K\} \end{array} \right. \right\}, \quad (8)$$

where $s_k \in (0, 1)$ denotes the normalized arc-length location along the horizontal alignment, $W_{c,k}$ and $S_{c,k}$ represent the pavement width and superelevation rate at the k -th cross-section, $T_{c,k} \in \{0, 1, 2\}$ denotes the cross-section type (subgrade-0, bridge-1, and tunnel-2), $h_{f,k}$ and $h_{c,k}$ denote the fill height and cut depth at the k -th cross-section, and $h_f^{\max}$ and $h_c^{\max}$ denote the maximum allowable fill height and cut depth, respectively.

Given the set of horizontal and vertical alignment design variables $u^{N,M} = (x^N, y^M) \in \mathbb{R}^{5N+3M}$, the three-dimensional highway alignment can be parameterized as:

$$\sigma_{3D}^{u^{N,M}} : s \in [0, L(x^N)] \mapsto \sigma_{3D}^{u^{N,M}}(s) = (\sigma_x(s), \sigma_y(s), \sigma_z(s)) \quad (9)$$

Then, the three-dimensional highway alignment surface $C_{3D}^{u^{N,M,K}}$ is defined as the set of all cross-sectional contours generated along the spatial centerline $\sigma_{3D}^{u^{N,M}}$ (s):

$$C_{3D}^{(u^{N,M,K})} = \left\{ \sigma_{3D}^{u^{N,M}}(s) + \tau^{(z^K)}(s) \mid s \in [0, L(x^N)] \right\} \quad (10)$$

where $\tau^{(z^K)}(s)$ is the function that generates the contour points of the cross-section in the normal plane at station s , as detailed in Appendix A3.

The highway alignment design variables and constraints can be aggregated into a high-dimensional decision vector $u^{N,M,K} = (x^N, y^M, z^K)$, which jointly incorporates the geometric parameters of horizontal alignment, vertical alignment, and cross-section. The set of all feasible highway alignment candidates $U_{ad}^{N,M,K}$ is then expressed in Eq. (11). In this study, mandatory geometric specifications are enforced as hard

constraints through the feasible design space $U_{ad}^{N,M,K}$ to restrict candidate alignments.

$$U_{ad}^{N,M,K} = \{ u^{N,M,K} = (x^N, y^M, z^K) \mid x^N \in X_{ad}^N, y^M \in Y_{ad}^M, z^K \in Z_{ad}^K \} \quad (11)$$

2.2. Optimization objective

Highway geometric design aims to determine an optimal alignment that satisfies all applicable design constraints while minimizing total life-cycle costs. These cost components typically include construction costs (e.g., earthwork, pavement, and structural elements), user costs (e.g., fuel consumption, travel time, and crash risk), social costs (e.g., benefits to local development and accessibility), and environmental costs (e.g., noise, air pollution, and ecological disruption). Social cost is expressed as a benefit-related term, where negative values indicate positive social benefits such as improved accessibility and regional development. Although construction, user, social, and environmental costs are integrated within a unified cost framework, the relative importance of each cost component may vary across project contexts. Therefore, separate weights are assigned to construction, user, social, and environmental costs to represent project-specific decision preferences and trade-offs among construction investment, user benefits, social development, and environmental impacts (Yang et al., 2014). To establish a generalized mathematical formulation, the weighted total cost J is computed by integrating the composite cost $p(x, y)$ at each point (x, y) along the highway centerline:

$$J(u^{N,M,K}) = \int_{C_{3D}^{u^{N,M,K}}} p(x, y) d\sigma_{3D}^{u^{N,M}} = \int_0^{L(x^N)} p(\sigma_{3D}^{u^{N,M}}(s)) ds, \quad (12)$$

$$p(x, y) = w_c \cdot p_c(x, y) + w_u \cdot p_u(x, y)$$

$$-w_s \cdot p_s(x, y) + w_e \cdot p_e(x, y) + 0.5 \cdot \chi_R(x, y), \quad (13)$$

where $C_{3D}^{u^{N,M,K}}$ and $\sigma_{3D}^{u^{N,M}}$ denote the three-dimensional highway alignment surface and the arc-length-parameterized representation determined by the design variables x^N , $L(x^N)$ is the total length of horizontal alignment, p_c , p_u , p_s , p_e denote the construction, user, social, and environmental cost components, w_u , w_e , w_s , w_c are the respective weight coefficients satisfying $w_u + w_e + w_s + w_c = 1$, and $\chi_R(x, y)$ represents a penalty term introduced to penalize alignments that pass through restricted areas. The weight coefficients are determined by combining entropy-based feature weighting with similarity analysis to historical design cases, thereby reflecting project-specific design preferences in multi-objective alignment optimization. The supplementary formulations of user costs, social costs, and environmental costs are provided in Appendix B. The penalty function $\chi_R(x, y)$ is defined as:

$$\chi_R(x, y) = \begin{cases} 0 & \text{if } (x, y) \in \Omega \setminus \varphi \\ p_R & \text{if } (x, y) \in \varphi \end{cases} \quad \text{with } p_R \gg 0 \quad (14)$$

where Ω denotes the design domain, $\varphi \subset \Omega$ denotes the set of restricted areas (e.g., urban zones, protected farmland, environmentally sensitive areas), and p_R is a large penalty value introduced to discourage horizontal alignments from entering infeasible areas.

Therefore, the task of determining the optimal highway alignment between two given terminal points $a = (x_a, y_a, z_a) \in \mathbb{R}^3$ and $b = (x_b, y_b, z_b) \in \mathbb{R}^3$ can be reformulated as the cost-minimization problem given in Equation (15).

$$\begin{aligned} \min_{u^{N,M,K}} J(u^{N,M,K}) &= \int_0^{L(x^N)} p(\sigma_{3D}^{u^{N,M}}(s)) ds \\ \text{s.t. } u^{N,M,K} &= (x^N, y^M, z^K), x^N \in X_{ad}^N, y^M \in Y_{ad}^M, z^K \in Z_{ad}^K \end{aligned} \quad (15)$$

3. Methodology

3.1. Overall framework

As illustrated in Fig. 2, the whole methodology framework consists of three main components, including highway geometric design knowledge graph (HGD-KG) construction, a highway geometric knowledge reasoning model, and a highway alignment optimization model. The purpose of HGD-KG construction is to structurally represent domain knowledge extracted from design specifications, engineering cases, and expert experience. Then, the HGD-KG is used to establish a highway geometric knowledge reasoning model (GNNRAG-H) based on the graph neural retrieval for efficient LLM reasoning algorithm. This kind of interpretable retrieval-augmented reasoning model is designed to capture implicit reasoning paths embedded in highway geometric design. Finally, the prior knowledge inferred by GNNRAG-H, including feasible variable ranges, design constraints, and optimization objective weights, is integrated into the SAC-based highway alignment optimization model (KGDRL-HAO) through feasible action restriction, reward shaping, and knowledge-guided policy regularization. The knowledge-guided learning process constrains alignment design variables, guides feasible exploration, and directs the search toward feasible and high-quality alignment solutions.

3.2. HGD-KG construction

Highway geometric design has progressively accumulated domain knowledge from design specifications, engineering cases, and expert experience. However, much of this knowledge remains unstructured or semi-structured, which limits its effectiveness in supporting computational optimization and intelligent decision-making. To address this issue, this study constructs the highway geometric design knowledge graph (HGD-KG) to represent and organize domain knowledge in a structured manner for graph-based reasoning tasks. A knowledge graph is a semantic network that organizes real-world knowledge in the form

of a data graph, where nodes denote entities or semantic units and edges denote the relationships between entities (Li et al., 2020). The proposed HGD-KG can be represented as a triple set $G = \{E, R, F\}$, where E is the set of entities, with each $e \in E$ corresponding to a concept, object, design criterion, design condition, value node, rule instance, or semantic unit in the domain of highway geometric design. R is the set of relationships, with each $r \in R$ describing a semantic dependency between entities, and $F \subseteq E \times R \times E$ is the set of factual triples, representing individual semantic units. In the instance layer, rule instances are further introduced as special semantic nodes to represent condition-dependent specification knowledge, rather than being treated as primary domain classes in the conceptual layer. Each fact is expressed as a triple $(h, r, t) \in F$, where the head entity h is connected to the tail entity t through the relation r , forming a basic semantic unit:

$$(h : \text{Entity}) \xrightarrow{r : \text{Relation}} (t : \text{Entity}) \quad (16)$$

In addition, the HGD-KG supports meta-paths, which are composed of multiple interlinked triples and represent multi-hop semantic chains between entities. A meta-path from h to t is defined as:

$$\mathcal{P}(h, t) = (h, r_1, e_1), (e_1, r_2, e_2), \dots, (e_{n-1}, r_n, t) \quad (17)$$

where $n \in \mathbb{N}$ is the length of the meta-path (i.e., the number of relations or hops, $n \geq 1$), $e_n \in E$ denotes intermediate entities, $r_n \in R$ denotes the relations connecting consecutive entities, and $\mathcal{P}$ represents a multi-hop semantic path from the head entity h to the tail entity t within the decision logic of highway geometric design, typically occurring among semantic units such as a design principle, a design constraint, or a geometric element. Equivalently, the meta-path can be written as an alternating sequence of entities and relations:

$$h \xrightarrow{r_1} e_1 \xrightarrow{r_2} e_2 \xrightarrow{r_3} \dots \xrightarrow{r_n} t \quad (18)$$

Considering the hierarchical structure and semantic granularity inherent in highway domain knowledge, the proposed HGD-KG adopts a hybrid modeling approach that integrates expert-defined ontologies with data-driven knowledge acquisition. Specifically, the top-down




Fig. 2. The methodology framework of this study.

approach establishes a class hierarchy ontology to ensure logical consistency and domain adaptability, while the bottom-up approach enriches the knowledge graph through the extraction of entities, attributes, relations, and rule knowledge from domain data (Shen et al., 2023). As illustrated in Fig. 3, the proposed HGD-KG employs a two-layer semantic architecture based on RDF Schema (RDFS). RDFS is a lightweight ontology language that provides fundamental constructs for defining classes, properties (object/data), and hierarchical relations (e.g., `rdfs:subClassOf`, `rdfs:domain`, `rdfs:range`) (Colucci et al., 2024). This architecture consists of two layers: (1) a conceptual layer (i.e., terminological box, TBox) that defines the ontology's class hierarchy, data properties, and object properties to establish a structured semantic schema, and (2) an instance layer (i.e., assertional box, ABox) that populates the knowledge graph with concrete entity instances, rule instances, and their associated properties.

The conceptual layer constructs the core semantic structure by organizing class hierarchies and defining data and object properties, whereas the instance layer instantiates the conceptual schema with concrete entity instances and rule instances, thereby encoding relational semantics through explicit instantiation. Besides ordinary domain entity instances, rule instances are modeled as special semantic nodes in the ABox to represent condition-dependent specification-based knowledge. Each rule instance serves as an intermediate node that explicitly links a design criterion with its applicable design conditions, prescribed design value, and specification source through factual triples. Through the integration of conceptual and instance layers, the HGD-KG achieves structured knowledge representation and facilitates knowledge

reasoning in highway geometric design.

3.2.1. Class hierarchy modeling

To establish a semantic backbone for the HGD-KG, this study defines a class hierarchy ontology to organize key concept categories and their semantic inheritance relationships within the domain of highway geometric design. Based on domain-specific abstractions in highway geometric design (AASHTO, 2018), the proposed HGD-KG ontology consists of six primary domain classes, including overall controls, design principles, design constraints, design criteria, design treatments, and design elements. Rule instances are ABox-level special semantic nodes used to encode condition-dependent specification rules, rather than primary domain classes in the TBox. Table 1 presents detailed descriptions of each class. These classes and their subclasses capture distinct functional aspects of the geometric design process, ranging from global control parameters and general design principles to specific geometric elements, design criteria, contextual constraints, and design treatments. In addition, rule instances are introduced in the instance layer to represent condition-dependent specification-based knowledge. For example, the geometric elements class contains subclasses such as horizontal alignment, vertical alignment, and cross-section, with horizontal alignment further divided into tangent, spiral curve, and circular curve. These hierarchical relations are encoded using the “`rdfs:subClassOf`” property to represent “is-a” semantics, thereby ensuring the logical consistency, semantic inheritance, and scalability of the ontology.




Fig. 3. The schema of HGD-KG ontology.

Table 1
The class hierarchy of HGD-KG.

Category	Description	Subclasses
Overall controls	Global parameters that determine the design basis	Design phase, functional classification, design speed, design vehicle, traffic volume, level of service, terrain type
Design principles	Abstract principles reflecting design philosophies and objectives	Safety, continuity, consistency, sustainability, economy
Geometric elements	Fundamental components of geometric design	Horizontal alignment (tangent, spiral curve, circular curve, reverse curve), vertical alignment (grade, crest curve, sag curve), cross-section (lane width, shoulder width, cross slope), sight distance (stopping, decision, passing)
Design criteria	Design requirements or control indicators used to determine, check, or evaluate geometric elements under given design conditions	Minimum radius, stopping sight distance, decision sight distance, passing sight distance, maximum grade, vertical curve length, lane width, shoulder width, cross slope, superelevation rate
Design constraints	Specification-based or contextual constraints in highway geometric design	Terrain (level, rolling, mountainous), geology, environmental considerations (protected areas, water sources, noise), climate zones (normal, snow/frost), regulatory limits, vehicle dynamics, driver behavior
Design treatments	Strategies adopted to mitigate or respond to design constraints	Alignment alternatives (avoidance, crossing, approach), structural solutions (tunnel, bridge, embankment), intersection design, superelevation transition, sight distance improvement

3.2.2. Data properties definition

To further enhance the semantic expressiveness of the HGD-KG, data properties are used to annotate both domain entities and rule instances with literal values. These properties provide structured support for validating design constraints, checking parameter ranges, and enforcing specification rules during knowledge reasoning. Based on different semantic roles within the proposed HGD-KG, data properties are categorized into three types: basic, numerical, and semantic properties. Table 2 summarizes the definitions and representative examples of each category. All property values are normalized to a standardized format to ensure semantic consistency and computational comparability across

Table 2
Data property types and representative examples.

Category	Description	Example
Basic property	Fundamental attributes describing an entity or rule instance, such as name, identifier, unit, data type, clause number, and source reference.	A rule instance may have basic properties such as “Name: Rule_Rmin_120”, “Unit: m”, and “Clause number: 7.3.2”.
Numerical property	Quantitative attributes specifying design values, thresholds, ranges, limits, or comparison operators for design criteria.	A rule instance for the general minimum radius of a circular curve may have numerical properties such as “Design value: 1000”, “Unit: m”, and “Constraint operator: $\geq$ ”.
Semantic property	Semantic annotations indicating applicable design conditions, contextual constraints, rule modality, and rule status.	A rule instance may include semantic properties such as “Rule modality: recommended”, “Rule status: valid”, and “Applicability type: general value”.

heterogeneous data sources. For example, length measurements are converted to meters (m), velocities to kilometers per hour (km/h), and slopes are converted to percentages (%). For rule instances, prescribed design values, units, comparison operators, specification sources, clause numbers, and table numbers are represented as data properties, such as *hasValue*, *hasUnit*, *hasOperator*, *hasSource*, and *hasClauseNumber*. In accordance with the RDFS schema, each data property is defined with an associated “*rdfs: domain*” and “*rdfs: range*”, formally specifying the applicable class or rule instance type and the expected data type, respectively.

3.2.3. Object properties linking

Highway geometric design adheres to a well-defined hierarchy of knowledge and decision logic that can be abstracted into a hierarchical decision-making chain extending from overall controls to design principles, geometric elements, design criteria, design constraints, and design treatments. This chain describes the knowledge evolution from macro-level design goals to specific design requirements and actions, while establishing semantic dependencies among different types of entities. By discretizing this decision-making chain, this study defines a comprehensive set of object properties to characterize multiple semantic relation types between different entities. These relations are categorized into eight distinct types, represented by the verbs *include*, *connect*, *influence*, *constrain*, *mitigate*, *guide*, *comply with*, and *compare to* (see Table 3). In addition, two rule-instance object properties, namely *refer to* and *apply to*, are introduced in the instance layer to link rule instances with their corresponding design criteria and applicable design conditions, thereby supporting the representation of condition-dependent specification-based knowledge. These object properties provide essential semantic links that support both structural integration and logical reasoning across heterogeneous entities and rule instances in the

Table 3
Object property types and examples.

Property	Description	Example
Include	Compositional or enumerative relation among design elements or categories.	Horizontal alignment includes tangents, spiral curves, and circular curves.
Connect	Topological or structural relation between geometric elements.	A spiral curve connects a tangent and a circular curve.
Influence	Non-mandatory effect of design constraints or design conditions on geometric elements or design criteria.	Mountainous terrain influences the minimum radius of circular curves.
Constrain	Mandatory restriction imposed by design controls, design constraints, or design conditions on geometric elements or design criteria.	Design speed constrains the required stopping sight distance.
Mitigate	Relation in which design treatments mitigate or respond to specific design constraints.	A bypass mitigates the impact of ecological protection zones.
Guide	Relation in which design principles guide the selection or application of geometric elements, design criteria, or design treatments.	Safety guides the use of transition curves, while economy guides bypass strategies.
Comply with	Relation indicating that a geometric element or design scheme satisfies a specified design criterion, rule instance, or requirement.	An exit segment should comply with decision sight distance requirements.
Compare to	Quantitative comparison relation between design values, design criteria, or design alternatives.	The minimum radius of one design alternative may be compared to that of another alternative.
Refer to	Relation linking a rule instance to its corresponding design criterion.	A rule instance refers to the general minimum radius of a circular curve.
Apply to	Relation linking a rule instance to its applicable design condition.	A rule instance applies to a design speed of 120 km/h.

proposed HGD-KG. By explicitly modeling these relationships, the knowledge graph can capture cross-category and cross-layer dependencies throughout the design process, thereby supporting structured knowledge integration and reasoning tasks.

3.3. Highway geometric design knowledge reasoning model

To better realize the computational embedding and automated reasoning of highway geometric design knowledge, this study develops a domain-specific reasoning model, GNNRAG-H, which integrates query-aware subgraph construction, GNN-based multi-hop retrieval, retrieval augmentation, and LLM-based reasoning over the constructed HGD-KG (Mavromatis and Karypis, 2025). Given a natural language query and the HGD-KG as inputs, the model outputs a structured semantic answer together with its associated reasoning paths. By leveraging query-aware graph retrieval, structured multi-hop reasoning, LLM-assisted semantic path expansion verified against the HGD-KG, and context-aware language generation, GNNRAG-H can improve the accuracy, reliability, and interpretability of knowledge graph question answering in highway geometric design tasks.

As demonstrated in Fig. 4, GNNRAG-H consists of four components: query-aware subgraph construction, GNN-based retrieval, retrieval augmentation, and LLM-based reasoning. The query-aware subgraph construction module dynamically identifies and extracts query-relevant entities, rule instances, and their associated relations from the HGD-KG. The GNN-based retrieval module performs multi-hop reasoning over the constructed subgraph through message-passing mechanisms to generate candidate answers and graph-grounded reasoning paths. The retrieval augmentation module uses the LLM to generate semantically enriched candidate paths and filters them through HGD-KG verification, retaining only paths supported by existing entities, triples, or rule instances. The LLM-based reasoning module then conducts semantic inference based on the verified and augmented structured reasoning paths and produces natural language answers grounded in knowledge-graph facts. The detailed process is described as follows:

Step 1 (Query-aware subgraph construction): To improve the relevance and efficiency of knowledge graph reasoning, this study adopts a query-aware subgraph construction strategy to dynamically extract a compact reasoning subgraph from the HGD-KG. Specifically, the natural language query is analyzed to identify key design entities, design criteria, design conditions, constraints, and specification-related requirements. Based on these cues, query-relevant entities, rule instances, and associated relations are retrieved from the HGD-KG. Applicable specification subgraphs are then selected by matching the identified design conditions with the *apply to* relations of rule instances, while related design criteria, design values or value nodes, and specification sources are retrieved through the *refer to*, *specify*, and *derive from* relations. Query-aware relevance weights are further computed to capture variable dependencies and reduce redundant reasoning paths. When available, semantically similar historical cases and expert knowledge are incorporated to improve engineering applicability. The resulting query-aware reasoning subgraph G_q^* is defined as:

$$G_q^* = G_S \cup \{(v, r, v') | v, v' \in V_q, r \in R_q\} \cup \{(v, r, v') | (v, r, v') \in G_{q_H}, q_H \in Q_H, \text{sim}(q, q_H) > \delta\} \quad (19)$$

where v and v' denote the head and tail entities in the HGD-KG, which may correspond to domain entities or rule instances, r denotes the semantic relation between two nodes, G_S denotes the specification subgraph matched to the query context, V_q denotes the set of query-relevant nodes retrieved from the HGD-KG, R_q denotes the set of relations associated with these query-relevant nodes, G_{q_H} denotes the subgraph associated with a historical case q_H , Q_H denotes the set of historical cases, $\text{sim}(q, q_H)$ denotes the semantic similarity between the current query q and the historical case q_H , and δ is the similarity threshold for historical case matching.

Step 2 (GNN-based retrieval): Once the query-aware reasoning subgraph G_q^* is obtained, this study introduces a multi-layer GNN to perform structured multi-hop retrieval, aiming to infer candidate answer nodes




Fig. 4. The architecture of GNNRAG-H.

and their associated reasoning evidence. Through message passing across GNN layers, the model iteratively updates node representations and captures multi-hop semantic dependencies among domain entities, rule instances, and value nodes (Gao et al., 2025). At the l -th layer, the embedding $h_v^{(l)}$ of node v is updated by aggregating messages $m_{vv'}^{(l)}$ from its neighbors $v' \in N_v$:

$$h_v^{(l)} = \psi \left(h_v^{(l-1)}, \sum_{v' \in N_v} \alpha_q(v, r, v') \cdot m_{vv'}^{(l)} \right), \quad (20)$$

$$\alpha_q(v, r, v') = \frac{\exp(\phi(q, v, r, v'))}{\sum_{(v, r, u) \in G_q^*} \exp(\phi(q, v, r, u))}, \quad (21)$$

where N_v denotes the neighbor set of nodes v , $\alpha_q(v, r, v')$ is a query-aware relevance indicator for the triple (v, r, v') , $m_{vv'}^{(l)}$ denotes the message passed from neighbor v' to v node at layer l , and $\phi(\cdot)$ is a scoring function that measures the relevance between the query q and the triple (q, v, r, v') , and $\psi(\cdot)$ denotes the aggregation function, such as mean pooling or GRU.

After L layers of message passing, the final node embedding $h_v^{(L)}$ is fed into a multi-layer perceptron (MLP) to produce an answer confidence score p_v . The node with the highest score is selected as the candidate answer a_G :

$$a_G = \arg \max_{v \in V(G_q^*)} p_v, p_v = \frac{\exp(\text{MLP}(h_v^{(L)}))}{\sum_{u \in V(G_q^*)} \exp(\text{MLP}(h_u^{(L)}))} \quad (22)$$

where $V(G_q^*)$ denotes the node set of the query-aware reasoning subgraph, and L denotes the total number of GNN layers.

Then, the top- k reasoning paths from each question entity $e \in E_q$ to the predicted answer node a_G are extracted as the GNN-based retrieval paths $\mathcal{P}_{\text{GNN}}$:

$$\mathcal{P}_{\text{GNN}} = \left\{ \text{TopKPath}(e, a_G, G_q^*) \mid e \in a_G \right\} \quad (23)$$

where E_q denotes the set of entities mentioned in the query, and $\text{TopKPath}(\cdot)$ returns the top- k paths ranked by query relevance and graph connectivity.

Step 3 (Retrieval augmentation): Although $\mathcal{P}_{\text{GNN}}$ provides structurally grounded reasoning signals, its coverage may be limited when the query involves implicit semantic associations or reasoning paths that are not fully captured by GNN-based retrieval. To improve interpretability, each GNN-derived path is first verbalized into natural language, transforming symbolic entity-relation chains into readable reasoning statements:

$$\text{Verbalized path} = \{e_1 \rightarrow r_1 \rightarrow e_2 \rightarrow \dots \rightarrow e_i \rightarrow r_i \rightarrow a_G\} \quad (24)$$

where e_i and r_i denote an entity and its corresponding relation. For example, the path “highway classification” $\rightarrow$ “influences” $\rightarrow$ “design speed” $\rightarrow$ “constrains” $\rightarrow$ “minimum curve radius” can be verbalized as: “In highway geometry design, the highway classification influences the design speed, which in turn constrains the minimum radius of circular curves.”

To enhance semantic diversity and reasoning robustness, an LLM (e.g., Qwen, DeepSeek, or ChatGPT) is then used to generate enriched reasoning paths based on the query, the verbalized GNN paths, and a task-specific prompt. The resulting set of LLM-augmented reasoning paths is denoted as $\mathcal{P}_{\text{LLM}}$:

$$\mathcal{P}_{\text{LLM}} = \text{LLM}(\text{prompt}, \mathcal{P}_{\text{GNN}}, q) \quad (25)$$

where $\mathcal{P}_{\text{GNN}}$ denotes the set of verbalized reasoning paths derived from the GNN, and q denotes the input query, and prompt denotes the instruction template guiding the LLM to generate semantically enriched

reasoning paths.

Only LLM-generated paths grounded in the HGD-KG are retained and integrated with the GNN-derived paths to form the final reasoning-path set $\mathcal{P}_{\text{final}}$.

$$\mathcal{P}_{\text{final}} = \mathcal{P}_{\text{GNN}} \cup \left\{ p \in \mathcal{P}_{\text{LLM}} \mid \text{Ground}(p, G_q^*) = 1 \right\} \quad (26)$$

where $\text{Ground}(\cdot) = 1$ indicates that path p is supported by existing entities, relations, triples, or rule instances in G_q^* . Specifically, each LLM-generated reasoning path is verified against the HGD-KG before final answer generation (Zheng et al., 2025). Paths containing unsupported entities, relations, triples, or rule instances, or inconsistent with applicable rule conditions, are filtered out. When multiple verified paths are available, they are ranked according to query relevance, graph connectivity, and consistency with the identified design conditions. Only verified reasoning evidence is provided to the LLM for final answer generation, thereby improving reasoning reliability and consistency while ensuring that the resulting design suggestions remain grounded in HGD-KG facts.

Step 4 (LLM-based reasoning): The final reasoning paths $\mathcal{P}_{\text{final}}$ are then verbalized and fed into the LLM, together with the original question q and a task-specific prompt. Using task-specific prompt engineering, the LLM produces structured outputs that specify feasible design variables, corresponding design constraints, and normalized optimization objective weights. The generation process can be formulated as:

$$a_{\text{final}} = \text{LLM}(\text{prompt}, \mathcal{P}_{\text{final}}, q) \quad (27)$$

To reflect the structured nature of the output, the final answer a_{final} is represented as:

$$a_{\text{final}} = \left(\left\{ (u_i, [u_i^{\min}, u_i^{\max}]) \mid u_i \in U \right\}, \left\{ (o_j, w_j) \mid o_j \in O \right\}, \left\{ (c_k, c_k) \mid c_k \in C \right\} \right)$$

where U denotes the set of design variables, u_i is the i -th design variable with its recommended lower and upper bounds $[u_i^{\min}, u_i^{\max}]$, O denotes the set of optimization objectives, o_j is the j -th objective associated with a normalized importance weight $w_j \in [0, 1]$, C denotes the set of design constraints, and c_k is the k -th constraint in design constraints.

3.4. Highway alignment optimization model

Highway geometric design is inherently an iterative optimization process that requires the continuous adjustment of design variables to satisfy and balance design specifications, geometric constraints, and multi-objective optimization. To address these challenges, this study integrates knowledge reasoning into the soft actor-critic (SAC) algorithm to develop a knowledge-augmented deep reinforcement learning model for highway alignment optimization, named KGDL-HAO, aiming to transform the design knowledge inferred from reasoning into actionable design decisions and to improve policy learning under engineering constraints. SAC is a deep reinforcement learning algorithm that balances exploration and exploitation via entropy maximization, making it particularly suitable for high-dimensional and continuous optimization tasks (Sun et al., 2022). The model treats horizontal alignment, vertical alignment, and cross-sectional parameters as controllable design variables, while terrain data and design constraints are treated as environmental inputs. The model output is the optimal alignment solution satisfying multiple optimization objectives. By incorporating a maximum entropy-based reinforcement learning framework with a knowledge-guided policy regularization, the proposed model improves the agent's exploration efficiency and decision consistency in complex terrain environments.

Specifically, the prior knowledge inferred by GNNRAG-H is incorporated into the SAC learning process through different mechanisms. First, the inferred design variables and recommended ranges are used to

guide state initialization, while mandatory specification limits and design constraints are used to restrict the feasible action space. Second, the inferred objective weights are incorporated into reward evaluation to guide multi-objective optimization. Third, expert design preferences derived from the verified reasoning results are encoded into an expert-guided action preference distribution, which is introduced into the SAC actor loss through policy regularization.

As shown in Fig. 5, the proposed KGDRL-HAO establishes a closed loop among knowledge reasoning, alignment optimization, and environment feedback. The Markov decision process formalizes the highway geometric design problem as an interaction process defined by state, action, and reward, thereby constructing a computational decision space. The GNNRAG-H model performs structured reasoning over the highway geometric-design knowledge graph to infer structured prior knowledge that guides the agent's exploration toward satisfying design specifications and knowledge-guided design preferences. The SAC algorithm leverages a maximum-entropy mechanism to conduct policy learning and maintain an exploration-exploitation balance, guiding the knowledge-informed agent to iteratively refine the highway alignment and achieve a computable transformation from knowledge reasoning to design decision-making.

3.4.1. Markov decision process modeling

The highway geometric design problem can be formulated as a Markov decision process (MDP). The MDP is described by a five-tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$, where $s \in \mathcal{S}$ is the state space, $a \in \mathcal{A}$ is the action space, $\mathcal{P}(s_{t+1}|s_t, a_t)$ denotes the transition probability from state s_t to s_{t+1} under action a_t , $r \in \mathcal{R}$ is the reward, and $\gamma \in [0, 1]$ is the discount factor that balances the relative importance of immediate and future

rewards. In the MDP, the agent continuously interacts with the environment by selecting an action a_t based on policy $\pi(a|s)$, which transitions the environment from the state s_t to s_{t+1} , while the agent receives the corresponding reward r_t . The goal of the agent is to maximize the expected cumulative discounted reward over all time steps. In this study, the environment is constructed as a virtual three-dimensional geometry design platform, which incorporates realistic terrain, land-use constraints, and design specifications. The optimization task is thus modeled as a single-agent deep reinforcement learning problem in a constrained continuous design space. The state, action, and reward functions are defined as follows:

State: The state space $\mathcal{S}$ captures observations of current highway alignment and terrain features. The state s_t is defined as:

$$s_t = (\mathbf{x}_t^N, \mathbf{y}_t^M, \mathbf{z}_t^K, \mathbb{T}), \quad s_t \in \mathcal{S} \quad (29)$$

where $\mathbf{x}_t^N$, $\mathbf{y}_t^M$, and $\mathbf{z}_t^K$ denote the horizontal alignment, vertical alignment, and cross-section design variables at time t , respectively; $\mathbb{T}$ denotes a terrain matrix that encodes ground elevation and land-use attributes.

Action: The action space $\mathcal{A}$ defines the set of allowable modifications to the design variables that the agent can execute at each time step t . The action a_t is given by:

$$a_t = (\Delta \mathbf{x}_t^N, \Delta \mathbf{y}_t^M, \Delta \mathbf{z}_t^K), \quad a_t \in \mathcal{A} \quad (30)$$

where $\Delta \mathbf{x}_t^N$, $\Delta \mathbf{y}_t^M$, and $\Delta \mathbf{z}_t^K$ represent the incremental changes to the corresponding design variables, respectively. All actions are required to keep the updated alignment within the feasible design space $\mathcal{U}_{ad}^{N,M,K}$. Actions violating mandatory geometric specifications are clipped,




Fig. 5. The KGDRL-HAO architecture based on soft actor-critic.

rejected, or resampled before the environment state is updated.

Reward: The reward function evaluates the quality of the current alignment and guides the agent toward satisfying design constraints while improving optimization objectives. To better account for optimization objectives and design constraints, this study employs a hybrid reward function combining a cost-based reward term and a constraint violation penalty term, formulated as:

$$r_t(s_t, a_t) = \lambda_1(1 - J_N(s_t, a_t)) - \lambda_2 \cdot V_N(s_t, a_t) \quad (31)$$

where $J_N(s_t, a_t)$ denotes the normalized weighted total cost derived from the objective function J , and $V_N(s_t, a_t)$ denotes the normalized degree of constraint violation of the candidate action with respect to mandatory geometric specifications before hard-constraint handling. Candidate actions violating mandatory geometric specifications are subsequently clipped, rejected, or resampled before the environment state is updated. All cost components and design constraint terms are normalized (e.g., scaled to the range [0,1]) to avoid scale imbalance. The coefficients λ_1 and λ_2 are weighting factors that balance the trade-off between cost minimization and constraint satisfaction. In the proposed framework, the structured prior knowledge inferred by GNNRAG-H is incorporated into the SAC learning process through state initialization, feasible action restriction, reward evaluation, and knowledge-guided policy regularization.

3.4.2. Soft actor-critic algorithm

As illustrated in Fig. 5, the SAC algorithm consists of one stochastic policy network and two Q-value networks, together with their corresponding target networks. The policy network outputs a Gaussian distribution over actions to produce stochastic geometric design decisions, while the Q-value networks estimate the expected long-term returns of state-action pairs, thereby guiding the agent to explore high-quality alignment solutions. To ensure differentiability and maintain sufficient exploration capability, actions are sampled from the policy network using the reparameterization trick, as expressed in Eq. (32).

$$a_t = \mu_\phi(s_t) + \epsilon_t \cdot \sigma_\phi(s_t), \quad \epsilon_t \sim \mathcal{N}(0, I) \quad (32)$$

where a_t and s_t are the action and state vector at time step t , μ_ϕ and σ_ϕ are the mean and standard deviation of the Gaussian distribution output by the policy network given state s_t , ϵ_t is noise sampled from a standard normal distribution, and I is the identity covariance matrix.

Based on the current design state s_t , GNNRAG-H provides verified reasoning paths relevant to the current geometric design conditions. For policy regularization, the verified reasoning paths are further processed by the LLM using a fixed task-specific prompt and structured output format to infer expert action preferences. The LLM converts the verified reasoning results into a normalized expert preference vector $\eta_t \in [-1, 1]^D$, where the sign of each element indicates the preferred adjustment direction for the corresponding action dimension and the absolute value represents the degree of preference. The expert-guided action preference distribution is then constructed as:

$$\pi_e(\cdot|s_t) = \mathcal{N}(\Delta a_{\max} \odot \eta_t, \Sigma_e), \eta_t = \text{LLM}(\text{prompt}, \mathcal{P}_{\text{final}}, s_t) \in [-1, 1]^D \quad (33)$$

where $\Delta a_{\max}$ denotes the maximum allowable adjustment magnitude of each action dimension, $\odot$ denotes element-wise multiplication, and Σ_e denotes a fixed diagonal covariance matrix used to retain exploration around the knowledge-guided action tendency. During action execution, the feasible action restriction described in Section 3.4.1 remains unchanged. The resulting $\pi_e(\cdot|s_t)$ serves as an auxiliary expert policy under the current state.

To encourage the learned policy $\pi_\phi(\cdot|s_t)$ to remain broadly consistent with the knowledge-guided prior preference, a Jensen-Shannon divergence term is introduced into the SAC objective as an auxiliary policy regularization term (Yang et al., 2021). The knowledge-augmented

maximum-entropy objective J_π^{KG} is formulated as follows:

$$J_\pi^{\text{KG}} = E_{(s_t, a_t) \sim \pi_\phi} \left[\sum_t \left(r(s_t, a_t) + \alpha H(\pi_\phi(\cdot|s_t)) - \beta D_{\text{JS}}(\pi_\phi(\cdot|s_t) || \pi_e(\cdot|s_t)) \right) \right] \quad (34)$$

where ρ_{π_ϕ} denotes the state-action distribution induced by policy π_ϕ , $r(s_t, a_t)$ denotes the immediate reward received after taking action a_t in state s_t , $H(\pi_\phi(\cdot|s_t))$ denotes the entropy of the policy distribution, α is the temperature coefficient controlling the trade-off between exploration and exploitation, β controls the strength of knowledge-guided policy alignment, and $D_{\text{JS}}(\cdot)$ denotes the Jensen-Shannon divergence.

To reduce value overestimation bias, SAC employs two independent Q-networks to estimate the expected return of a given state-action pair (s_{t+1}, a_{t+1}) . The target Q-value is computed as:

$$Q_{\text{target}}(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{a_{t+1} \sim \pi_\phi} \left[\min_{i=1,2} Q_{\theta_i}(s_{t+1}, a_{t+1}) - \alpha \log \pi_\phi(a_{t+1}|s_{t+1}) \right] \quad (35)$$

where $Q_{\text{target}}(s_t, a_t)$ denotes the target expected return for taking action a_t in state s_t , and $Q_{\theta_i}(s_{t+1}, a_{t+1})$ represents the evaluation of the next state-action pair (s_{t+1}, a_{t+1}) by the i -th target Q-network. The two Q-networks are trained by minimizing the mean squared error between the predicted Q-values and the target Q-values:

$$J_Q(\theta_i) = \mathbb{E}_{(s_t, a_t, r_t, s_{t+1}) \sim D} \left[\frac{1}{2} (Q_{\theta_i}(s_t, a_t) - Q_{\text{target}}(s_t, a_t))^2 \right], i = 1, 2 \quad (36)$$

where D is a replay buffer containing tuples (s_t, a_t, r_t, s_{t+1}) . To stabilize training, the target Q-networks are updated by a soft update mechanism. Accordingly, the policy network is updated by minimizing the following knowledge-regularized actor loss $\mathcal{L}_\pi^{\text{KG}}(\phi)$:

$$\mathcal{L}_\pi^{\text{KG}}(\phi) = \mathbb{E}_{s_t \sim D, a_t \sim \pi_\phi} \left[\alpha \log \pi_\phi(a_t|s_t) - \min_{i=1,2} Q_{\theta_i}(s_t, a_t) + \beta D_{\text{JS}}(\pi_\phi(\cdot|s_t) || \pi_e(\cdot|s_t)) \right] \quad (37)$$

Through iterative interaction with the environment, the policy network, Q-networks, and target networks are continuously updated. In the proposed KGDL-HAO, the structured prior knowledge inferred by GNNRAG-H guides the optimization process mainly by constraining feasible action regions and shaping reward evaluation, while the expert-guided regularization term provides an auxiliary preference alignment between the learned policy and knowledge-guided prior action tendencies. As a result, the agent can progressively improve the feasibility, consistency, and engineering rationality of highway alignment optimization.

4. Experiments

4.1. Data preparation

Domain knowledge in highway geometric design is typically embedded in unstructured or semi-structure formats, such as textual descriptions, images, tables, engineering drawings, and vector graphics. In this study, design standards, engineering cases, and expert experience are integrated to construct the HGD-KG, thereby supporting knowledge reasoning tasks. Specifically, the design specifications contain the technical standard of highway engineering (MTPRC, 2014) and the design specification for highway alignment (MTPRC, 2017), which comprehensively cover all fundamental geometric design elements and associated constraints. Engineering case data are collected from 46 highway projects across various regions in China, including alignment alternatives, digital terrain models, land use characteristics, geological hazard zones, and environmentally sensitive areas. Expert knowledge is extracted from domain-related literature, consisting of alignment

optimization strategies and empirical parameter values.

During the data preprocessing stage, all unstructured and semi-structured data are first normalized into textual form. A Transformer-based information extraction model, REBEL (Relation Extraction By End-to-end Language Generation), is employed to automatically perform named entity recognition and relation extraction (Cabot and Navigli, 2021). Specifically, design specification clauses, engineering case descriptions, and expert-experience texts are converted into candidate triples in the form of head entity-relation-tail entity. These candidate triples are then cleaned, normalized, deduplicated, and manually checked to ensure semantic consistency and validity before being incorporated into the HGD-KG. After extraction, the resulting triples are further cleaned, normalized, and manually checked to ensure semantic consistency and validity. As a result, a total of 3,754 high-quality knowledge triples are constructed, covering 6,654 semantic entities and 488 instantiated object-property relationships. Fig. 6 presents the Neo4j-based visualization of the HGD-KG, illustrating the structural organization of highway geometric design knowledge.

Building upon these knowledge triples, this study further constructs a highway geometric question-answering benchmark dataset, termed HGQA-KG, which is specifically designed to evaluate the reasoning capability of the proposed GNNRAG-H model. The HGQA-KG dataset comprises natural language questions centered on feasible design variables, design constraints, and optimization objective weights. To improve the model's discriminative ability between valid and invalid reasoning paths, each ground-truth triple (i.e., positive instance) is paired with a corresponding negative instance generated by controlled semantic perturbations such as entity replacement, relation substitution, or value corruption (Yang et al., 2020). Consequently, the resulting HGQA-KG dataset contains a total of 7,508 question-answer pairs, comprising 3,754 positives and 3,754 semantically perturbed negative instances. The HGQA-KG dataset is then used for training and evaluating the reasoning performance of the proposed GNNRAG-H, while a subset of the dataset is used for fine-tuning the LLM.

4.2. Implementation details

In this study, the overall experimental framework consists of three main components, namely HGD-KG construction, the knowledge-reasoning model GNNRAG-H, and the highway alignment optimization model KGDL-HAO. For the knowledge reasoning model, the GNNRAG-H is proposed as a hybrid reasoning framework that combines structured graph inference with LLM generation to support knowledge-guided highway alignment optimization. The relation-aware relational evidence architecture is adopted as the backbone of the GNN-based retrieval module to enhance relation awareness and facilitate multi-hop message passing for capturing causal dependencies within the HGD-KG (Mavromatis and Karypis, 2022). To semantically align natural-language queries with graph-based knowledge, the GNNRAG-H employs the BGE-m3 encoder to generate dense semantic representations of queries, entities, and relations (Chen et al., 2024). The BGE-m3 model, trained through contrastive representation learning, is further fine-tuned on the HGQA-KG corpus to capture domain-specific semantics and enhance retrieval accuracy in highway geometric design reasoning tasks. For the LLM module, instruction-formatted training samples are constructed from the HGQA-KG dataset and its corresponding verbalized reasoning paths. Then, the Qwen2.5-7B-Instruct model is fine-tuned through quantized low-rank adaptation with 4-bit NormalFloat quantization and low-rank adapters (rank = 16, $\alpha = 32$, dropout = 0.1), using 10549 instruction-formatted question-answer pairs (Hu et al., 2022). The fine-tuning is conducted on a single NVIDIA L40S GPU (48 GB VRAM) under mixed-precision training with gradient accumulation. Overall, the GNN module is trained in a supervised retrieval-oriented manner, whereas the LLM is fine-tuned on retrieval-augmented reasoning tasks to ensure accurate and context-aware response generation.

For the alignment optimization model, the proposed KGDL-HAO is developed based on the SAC framework. The actor network and critic networks adopt fully connected architectures with hidden layers of 256,




Fig. 6. Schematic illustration of the HGD-KG and the HGQA-KG dataset.

128, 128, and 64 units, using ReLU as the activation function. The primary hyperparameters of the SAC algorithm are summarized in Table 4. The SAC algorithm is implemented in PyTorch and trained using the Adam optimizer. All experiments are conducted on a workstation equipped with an Intel Core Ultra 7 265KF CPU and an NVIDIA GeForce RTX 5070TI GPU.

4.3. Model validation

Three commonly used metrics are adopted to evaluate the accuracy and completeness of the proposed GNNRAG-H in highway geometric design knowledge reasoning, including Hits, Hits@1, and F1 score. Hits measures whether the correct answer appears in the set of predicted candidates. Hits@1 evaluates the accuracy of the top-ranked prediction. F1 reflects the harmonic mean of precision and recall and evaluates the overlap between the predicted answer set and the ground-truth answer set. Higher values of these metrics indicate better reasoning performance. These metrics are calculated as below:

$$\text{Hits} = \frac{1}{Z} \sum_{z=1}^Z 1(y_z \in \hat{Y}_z), \quad (38)$$

$$\text{Hits@1} = \frac{1}{Z} \sum_{z=1}^Z 1(y_z = \hat{Y}_z^{(1)}), \quad (39)$$

$$\text{F1 score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \quad \text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (40)$$

where Z denotes the total number of test samples, y_z is the ground-truth answer for the z -th sample, $\hat{Y}_z$ is the set of predicted answers, $\hat{Y}_z^{(1)}$ is the top-1 prediction, and $1(\cdot)$ is the indicator function that returns 1 if the condition is true and 0 otherwise; TP, FP, and FN denote the counts of true positives, false positives, and false negatives, respectively.

In addition, the optimization performance of the proposed KGDL-HAO in multi-objective highway alignment design is evaluated using the Pareto front and the spacing metric. The Pareto front represents the set of non-dominated solutions that achieve the best possible trade-offs among multiple conflicting objectives (e.g., construction, user, social, and environmental costs). A Pareto-optimal solution is one for which no objective can be improved without deteriorating at least one other objective. The spacing metric measures the uniformity of Pareto-optimal solutions along the Pareto front. A smaller spacing metric indicates a more uniform and diverse distribution of Pareto-optimal solutions along the Pareto front, reflecting stronger global exploration capability and better optimization stability. The spacing metric D_s is calculated as follows:

$$d_i = \min_{j \neq i} \sum_{\lambda=1}^{\Lambda} |f_{\lambda}(x_i) - f_{\lambda}(x_j)|, \quad \bar{d} = \frac{1}{N_p} \sum_{i=1}^{N_p} d_i, \quad (41)$$

Table 4
Hyperparameters of the SAC algorithm.

Parameter	Description	Value
Maximum episodes	Total number of training episodes	1000
Exploration step	Steps for random exploration before learning	$1e^3$
Replay buffer size	Size of experience replay buffer	$1e^5$
Batch size	Number of samples per training step	512
Learning rate	Learning rate for policy and value networks	$3e^{-4}$
Discount factor	Discount factor in the Bellman equation	0.99
Policy deviation tolerance	Policy noise standard deviation	0.2
Target update rate	Soft update coefficient for target networks	0.005

$$D_s = \sqrt{\frac{1}{N_p - 1} \sum_{i=1}^{N_p} (d_i - \bar{d})^2}, \quad (42)$$

where $f_{\lambda}(x_i)$ denotes the value of the λ -th objective for the i -th Pareto-optimal solution, Λ is the total number of optimization objectives, and N_p is the total number of Pareto-optimal solutions.

5. Case study

In this study, a real-world case is employed to validate the proposed framework. The case study involves a section of the Xuancheng-Dongzhi highway in Anhui Province, China. The highway is designed for a speed of 100 km/h, with a four-lane (two lanes in each direction), and spans approximately 24 km. The highway corridor selected for the case study covers approximately 910 km² and is characterized by mountainous terrain. The start and end points are already fixed, with local coordinates specified $a = (17661.588, 12331.379)$ and $b = (54.793, 555.406)$. As shown in Fig. 7, the highway alignment is required to avoid agricultural land, environmentally sensitive areas, and urban development. In addition, the terrain elevation is obtained from contour grid-line data at 2 m intervals using a numerical interpolation approach.

5.1. Analyses of knowledge reasoning performance

To evaluate the reasoning capability and generalization performance of the proposed GNNRAG-H in highway geometric design tasks, this study conducts comparative experiments by comparing four representative models on different types of reasoning problems. The benchmark models include a basic GNN model, a fine-tuned 7-billion-parameter LLM (Qwen2.5-7B-Instruct), a MoE-based large-scale LLM (DeepSeek-V4-Flash, 13B activated parameters), and the proposed GNNRAG-H. The complete HGQA-KG dataset containing 7,508 question-answer pairs is divided into training, validation, and testing sets at a ratio of 80%, 10%, and 10%, respectively, resulting in 6,006 training samples, 751 validation samples, and 751 testing samples. All comparison models are evaluated on the same testing set with identical input questions. All methods are evaluated on the same testing set with identical input questions. For graph-based methods, the same HGD-KG is used as the knowledge source to ensure a fair and controlled comparison.

Table 5 summarizes the knowledge reasoning performance of the four models on the HGQA-KG testing dataset. The Hits, Hits@1, and F1 scores of GNNRAG-H reach 89.1%, 83.7%, and 87.8%, respectively, significantly outperforming both the basic GNN model and the fine-tuned 7B-parameter LLM. Notably, although the language-generation backbone of GNNRAG-H is a fine-tuned 7B-parameter LLM, GNNRAG-H achieves better overall performance than DeepSeek-V4-Flash, with relative gains of 16.0%, 18.9%, and 6.9% in Hits, Hits@1, and F1 score, respectively. In addition, Fig. 8 presents the confusion matrices of different models. The confusion-matrix results further demonstrate the superior reasoning consistency and structural robustness of the proposed GNNRAG-H. The performance advantage of GNNRAG-H can be attributed to the integration of structural graph reasoning and retrieval-augmented language generation, and such integration supports multi-hop inference and contextual semantic understanding, thereby reducing hallucinations and logical discontinuities commonly observed in standalone LLMs.

To evaluate the interpretability of the GNNRAG-H model's knowledge reasoning in highway geometric design tasks, this study conducts a visualization analysis from three perspectives, including causal reasoning path, reasoning path depth, and embedding-based semantic clustering. As shown in Fig. 9, the representative reasoning paths generated by GNNRAG-H demonstrate different reasoning processes, including inference of design criterion values, coordination among geometric elements, selection of design treatments under constraints,




Fig. 7. Terrain elevation and constraint areas along the highway corridor.

Table 5

Performance analysis on multi-hop (hops ≥ 1) and multi-entity (entities ≥ 2) questions on the HGQA testing dataset.

Model	Hits	Hits @1	F1
GNN	61.4%	54.6%	67.2%
Fine-tuned 7B LLM (Qwen2.5-7B-Instruct)	66.3%	60.8%	74.6%
MoE-based large-scale LLM (DeepSeek-V4-Flash, 13B activated)	76.8%	70.4%	82.1%
GNNRAG-H	89.1%	83.7%	87.8%

and preference reasoning for objective weights. Specifically, paths (a) and (b) infer prescribed geometric design values under a given design speed, path (c) reflects the coordination among terrain conditions, horizontal alignment, vertical alignment, and cross-section design, paths (d) and (e) describe the selection of design treatments under terrain or obstacle constraints, and path (f) illustrates how safety, economic efficiency, environmental impact, and farmland protection jointly guide the determination of objective weights. These causal chains typically involve overall controls, design principles, geometric elements, design criteria, design constraints, and design treatments, thereby reflecting the structured semantic dependencies among entities in the knowledge graph.

Additionally, 93.3 % of the valid reasoning paths fall within a depth of four hops or fewer (Fig. 10(a)), indicating that the model achieves a desirable balance between semantic depth and path stability. Fig. 10(b) presents the t-SNE visualization of node embeddings learned by the GNNRAG-H model. In the low-dimensional embedding space, entities belonging to the same semantic category form compact clusters with clear inter-cluster boundaries, with an average intra-cluster Euclidean distance of 0.91, which is substantially lower than the global average distance of 1.80. These findings demonstrate that the proposed model can effectively distinguish semantic categories and capture structured causal dependencies during multi-hop graph inference. A standalone

LLM often produces fragmented or hallucinatory reasoning chains due to the absence of explicit factual grounding and structural constraints. Compared with a standalone LLM, the GNNRAG-H grounds the reasoning process in explicit knowledge-graph facts and integrates symbolic reasoning with contextual generation, thereby generating transparent and traceable reasoning paths that support engineer-interpretable decision logic and enhance overall interpretability and credibility.

5.2. A comparison of DRL algorithms in highway alignment optimization

Using the SAC as the foundation algorithm, the proposed KGDRL-HAO is established for knowledge reasoning-guided highway alignment optimization, which realizes knowledge-driven exploration and decision-making. For comparison, this study also implements several classical DRL algorithms under the actor-critic architecture, including twin delayed deep deterministic policy gradient (TD3), deep deterministic policy gradient (DDPG), and proximal policy optimization (PPO). Different from the KGDRL-HAO that incorporates prior knowledge inferred by the GNNRAG-H model as an expert policy to guide exploration and decision-making, these baseline algorithms rely solely on environment interactions without any external domain knowledge. To ensure fairness in comparison, all algorithms are implemented under the same terrain conditions, design constraints, and optimization objective weights.

Fig. 11 illustrates the weighted total cost normalized to [0,1] and the constraint violation curves of different DRL algorithms during training. The constraint-violation curve represents the normalized violation degree of candidate actions before hard-constraint handling, rather than violations in the alignment after hard-constraint handling. The normalized values are used to evaluate the convergence efficiency and stability of the algorithms rather than represent actual engineering costs. As shown in Fig. 11(a), the KGDRL-HAO's weighted total cost decreases sharply after 100 episodes and stabilizes around 0.48. In




Fig. 8. Confusion metrics of different models on structured triple reasoning (Testing dataset).

contrast, SAC converges more slowly to approximately 0.50, while TD3, DDPG, and PPO exhibit limited learning and tend to converge early to local optima. As illustrated in Fig. 11(b), the KGDL-HAO achieves rapid convergence within approximately 150 episodes and gradually reduces constraint violations to near zero. The sharp reduction in constraint violation after approximately 400 episodes is attributed to the reasoning-guided correction provided by GNNRAG-H. However, SAC, TD3, DDPG, and PPO stabilize after 300 episodes while still maintaining constraint violation between 0.15 and 0.20. These results suggest that KGDL-HAO leverages expert policy to accelerate convergence and ensure strict constraint satisfaction by prioritizing low-cost and feasible design regions. Overall, the proposed KGDL-HAO outperforms other DRL algorithms in terms of exploration efficiency, decision robustness, and solution quality in highway alignment optimization.

Fig. 12 illustrates the Pareto front distributions and the corresponding spacing metrics obtained by different DRL algorithms. As shown in Fig. 12, the Pareto front generated by KGDL-HAO exhibits superior convergence and broader coverage, achieving the lowest spacing metric (0.022), which significantly outperforms SAC (0.061), TD3 (0.073), DDPG (0.098), and PPO (0.169). Table 6 lists the construction, user, social, and environmental costs associated with representative Pareto-optimal solutions obtained by each algorithm. Different from the normalized training cost shown in Fig. 11, Table 6 reports the cost components and weighted total cost of representative Pareto-optimal solutions in engineering monetary units ($\times 10^8$ ¥). Compared

with SAC, KGDL-HAO reduces the weighted total cost by approximately 6.4%, with more pronounced advantages over TD3, DDPG, and PPO. KGDL-HAO also achieves lower construction, user, and environmental costs and a higher equivalent social benefit, demonstrating superior trade-off performance among multiple cost objectives. These results indicate that KGDL-HAO achieves superior trade-off performance in multi-objective highway alignment optimization. This advantage may be attributed to the integration of expert policy guidance and entropy maximization, which helps maintain solution diversity, avoid premature convergence to suboptimal regions, and improve global exploration capability and decision robustness.

5.3. Comparison with manually designed alternatives

To further evaluate the performance advantage of the proposed KGDL-HAO, this study compares the alignment generated by the model (south alignment) with a manually designed alternative (north alignment). Both alignment schemes are developed under identical terrain conditions, design specifications, and geometric constraints to ensure fairness and comparability. The north alignment is produced by a senior engineering team with more than ten years of professional experience using conventional highway geometric design methods. In contrast, the southern alignment is generated by the proposed KGDL-HAO, a knowledge reasoning-guided deep reinforcement learning framework that iteratively searches for optimal alignment solutions.




Fig. 9. Representative semantic reasoning paths inferred by the GNNRAG-H.




Fig. 10. The reasoning path depth and semantic clustering of GNNRAG-H.

Fig. 13 illustrates the geometric design differences between the two alignments. The north alignment employs a terrain-adaptive design principle, using horizontal detours to minimize environmental disruption and construction cost. In contrast, the south alignment relies on a

higher proportion of bridge-tunnel structures to overcome terrain barriers and ensure coordinated horizontal and vertical alignment. Additionally, the south alignment improves access to nearby towns, reducing the maximum town distance from 6.8 km to 4.9 km, thereby supporting




Fig. 11. The weighted total cost and constraint violations curves.




Fig. 12. Pareto front obtained by different DRL algorithms.

Table 6

Costs of Pareto solutions from DRL algorithms (Unit: $\times 10^8$ ¥).

Algorithm	Construction cost	User cost	Social cost	Environmental cost	Weighted total cost
KGDRL-HAO	24.23	60.84	-91.18	4.22	5.44
SAC	24.57	61.23	-90.48	4.27	5.82
TD3	26.28	63.07	-88.93	4.53	7.16
DDPG	26.87	63.82	-88.16	4.74	7.75
PPO	26.96	64.03	-87.84	4.73	7.88

Note: Positive cost values denote expenditures, whereas negative cost values represent benefits.

local communities and regional development. Table 7 summarizes the main design variables and cost components between the north and south alignments. Compared to the north alignment, the south alignment reduces alignment length, earthwork volume, and agricultural land

occupation by 8.9%, 9.6%, and 58.5%, respectively, while maintaining a maximum vertical grade of 3.5% and a minimum horizontal curve radius of 2200 m. Although the construction cost of the south alignment increases by 7.7%, the user and environmental costs decrease by 8.4% and 12.5%, respectively, while the social benefit increases by 1.8%. Overall, the direct-sum total cost of the south alignment is approximately 6.06×10^8 ¥ lower than that of the north alignment, indicating that the south alignment offers superior cost-effectiveness, social benefits, and environmental sustainability.

To better analyze the preference differences between the two alignments in multi-objective optimization, this study compares the weight distribution across five objectives, including construction cost (CC), user cost (UC), social cost (SC), environmental cost (EC), and restricted-area cost (RC). The objective weights for both alignments are computed by integrating entropy-based feature weighting with similarity analysis based on Euclidean distance to historical design cases (Jin et al., 2021). The weight of restricted-area cost remains fixed at 0.5 in both alignments. As shown in Fig. 14, the north alignment assigns the highest weight to construction cost at 0.36, while user, environmental, and




Fig. 13. Geometric design comparison between the north and south alignments.

Table 7
Design variables and cost components for the two alignments.

Items	Unit	North alignment	South alignment	Improvement (%)
Alignment length	m	24158.525	22001.243	8.9
Agricultural land occupation	$\times 10^4 \text{ m}^2$	2.46	10.2	58.5
Environmental encroachment length	m	0	0	-
Bridge length	km	2.620	1.840	29.8
Tunnel length	km	4.825	6.593	-36.5
Minimum radius of horizontal curve	m	1350	2200	62.9
Maximum vertical grade	%	3.8	3.5	7.9
Embankment volume	m^3	1109266	1413636	-27.4
Cut volume	m^3	1414050	868306	38.6
Construction cost	$\times 10^8 \text{ ¥}$	2.251	24.25	-7.7
User cost	$\times 10^8 \text{ ¥}$	6.640	60.80	8.4
Social cost	$\times 10^8 \text{ ¥}$	-89.60	-91.20	1.8
Environmental cost	$\times 10^8 \text{ ¥}$	4.80	4.20	12.5
Direct-sum total cost	$\times 10^8 \text{ ¥}$	4.11	-1.95	-

Note: Positive cost values denote expenditures, whereas negative cost values represent benefits. User, social, and environmental costs are calculated over a 20-year analysis period. Improvement (%) denotes the relative change of the south alignment with respect to the north alignment.

social costs are weighted at 0.22, 0.24, and 0.18, respectively. In contrast, the south alignment exhibits a more balanced distribution, with weights of 0.27, 0.28, 0.21, and 0.24, respectively. These results

suggest that the south alignment achieves better-coordinated trade-offs among competing objectives, whereas the north alignment is dominated by construction cost considerations. This may be because the




Fig. 14. Multi-objective cost weight distributions for the north and south alignments.

KGDR-LHAO's ability to learn and reason over prior knowledge from historical cases and expert experience, thereby supporting context-aware and knowledge-guided decision-making.

6. Conclusion and discussions

This study fills the research gap by proposing a knowledge reasoning-guided deep reinforcement learning framework for three-dimensional highway alignment optimization from the following aspects. (1) A hierarchical HGD-KG is constructed to achieve the integration and semantic representation of heterogeneous knowledge from design specifications, engineering cases, and expert experience. (2) The graph neural retrieval for large language model reasoning is applied to establish a knowledge reasoning model for highway geometric design, which is designed to capture implicit multi-hop and multi-entity semantic dependencies and provides interpretable inferred paths. (3) The prior knowledge is incorporated as an expert policy into a soft actor-critic-based highway alignment optimization model, thereby further improving the exploration efficiency and decision consistency of the proposed KGDR-LHAO.

This study constructs a domain-specific knowledge graph for highway geometric design, providing a structured foundation for semantic representation and reasoning. Compared with existing domain-specific knowledge graphs, the HGD-KG constructed in this study adopts a hybrid modeling strategy that combines top-down ontology design with bottom-up data extraction, thereby improving both logical consistency and coverage breadth. For example, the pavement performance ontology proposed by Yang et al. (2022) primarily focuses on static attribute modeling, whereas the HGD-KG introduces a class hierarchy structure and cross-level relational to support more complex reasoning tasks. The customized knowledge graph for railway alignment also demonstrates the feasibility of ontology-driven approaches (Pu et al., 2024a), while the HGD-KG further extends this paradigm by establishing a more comprehensive semantic hierarchy and relational structure to support causal reasoning from overall control to geometric elements. Unlike medical knowledge graphs that prioritize entity recognition over relational complexity (Li et al., 2020), the HGD-KG improves representational granularity by explicitly modeling semantic links between design constraints and treatments.

The integration of GNN-based structural reasoning and LLM-based semantic generation demonstrates significant advantages in knowledge reasoning for highway geometric design. On the HGQA-KG testing dataset, the proposed GNNRAG-H achieves Hits@1 and F1 scores of 83.7% and 87.8%, respectively, showing relative gains of 53.3% and 30.7% over the baseline GNN, and 18.9% and 6.9% over DeepSeek-V4-Flash. These results confirm the effectiveness of the GNN-LLM hybrid

reasoning paradigm in combining structured retrieval with semantic generation, aligning with previous findings that GNN-RAG models outperform GPT-4 on multi-hop question-answering tasks (Mavromatis & Karypis, 2022). Prior studies have reported that an LLM exhibits hallucination rates of 30-40% in multi-hop question-answering tasks, often resulting in logical inconsistencies and fragmented inference chains (Ji et al., 2023), which is unacceptable in engineering applications that require robust causal reasoning and traceable decision-making (Zhan et al., 2024). To address these limitations, the GNN constructs stable reasoning paths through multi-layer message passing and provides structured contextual guidance for the LLM, thereby mitigating factual hallucination and enhancing interpretability (Yang et al., 2022). This mechanism not only enhances model robustness, but also aligns with practical engineering requirements for design review and road-safety audit procedures, thereby ensuring decision traceability and design transparency (Su et al., 2025).

Prior knowledge guidance is another novel addition in this study compared to prior research on highway alignment optimization. The results of this study suggest that expert-guided policy can significantly improve the exploration efficiency and decision consistency of the agent in multi-objective highway alignment optimization. Similar knowledge-guided or expert-informed policies have also demonstrated remarkable benefits in other engineering optimization domains. It has been reported that guided by interactive learning and heuristic rules, the DDPG-based railway alignment method reduced construction costs by 17.65% compared with alignments manually designed by experienced highway engineers and outperformed the distance transform method by 7.98% in mountainous areas (Gao et al., 2022). Using a customized knowledge graph with condition-relation-treatment triples and karst hazard priors, the NSGA-II-based alignment optimization method reduces the weighted total cost by 8.7% and improves the optimization efficiency by 2.5% (Pu et al., 2024b). Another study demonstrates that employing an LLM to generate high-level planning instructions for guiding downstream decision-making, followed by execution via a DRL agent in an embodied intelligence architecture, can improve the success rate of robotic manipulation tasks by over 30% compared to conventional end-to-end DRL methods (Dalal et al., 2024). Overall, growing evidence underscores the importance of prior knowledge guidance for achieving substantial improvements in both exploration efficiency and solution quality in highway and railway alignment optimization.

The findings of this study have several practical implications. (1) The proposed HGD-KG effectively integrates multi-source heterogeneous knowledge, which provides a novel and universal solution for the semantic modeling and logical representation of highway geometric design knowledge. The KGDR-LHAO proposed in this study can achieve end-to-end generation of optimal alignment solutions and offer

profound insights into improving the interpretability and traceability of the geometric design decision process. (2) By incorporating knowledge reasoning and DRL, the proposed KGDL-HAO improves the agent's exploration efficiency and decision consistency. This method is helpful to eliminate repetitive manual trial-and-error processes while ensuring the practical applicability of the generated alignment alternatives. In practical engineering workflows, KGDL-HAO can serve as a decision-support tool for preliminary highway alignment design by generating and evaluating alternative alignments under specified terrain, design constraints, and optimization objectives. The case study demonstrates the applicability of KGDL-HAO under mountainous expressway conditions. Application to other terrain conditions and road classes can be further investigated by updating the design specifications, terrain data, engineering cases, and expert knowledge represented in HGD-KG. Additionally, by leveraging the powerful semantic understanding and natural language generation capabilities of LLM, the proposed KGDL-HAO paves the way for the development of more intelligent human-in-the-loop highway geometric design technologies that support natural language interaction (i.e., design by chat).

One limitation of this study is that the performance of the proposed KGDL-HAO framework has not been extensively validated across different geographical regions and highway configurations. Terrain characteristics and highway design principles vary significantly across regions, so in future studies, more engineering cases and expert knowledge from diverse geographical areas and highway configurations should be collected to further expand the domain completeness of the HGD-KG and to validate the robustness and generalizability of the proposed framework. In addition, the inferred knowledge is introduced as deterministic prior input into the highway alignment optimization model, without effectively identifying potential conflicts that may exist among different knowledge sources (Zhai et al., 2024). In future work, this study will explore uncertainty-aware modeling methods to enhance the robustness of agent-based policy learning under large-scale domain knowledge guidance.

CRediT authorship contribution statement

Jianqiang Gao: Writing – original draft, Methodology, Data curation, Conceptualization. **Bo Yu:** Writing – original draft, Methodology, Data curation, Conceptualization. **Yuren Chen:** Writing – original draft, Funding acquisition, Conceptualization. **Kun Gao:** Writing – original draft, Methodology, Conceptualization. **Shan Bao:** Data curation, Conceptualization. **Junhua Wang:** Methodology, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Bo Yu reports financial support was provided by National Natural Science Foundation of China. Jianqiang Gao reports financial support was provided by China Postdoctoral Science Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This study is supported by the National Natural Science Foundation of China (No. 52572381) and the China Postdoctoral Science Foundation (No. 2026M791775).

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.cacaie.2026.100231](https://doi.org/10.1016/j.cacaie.2026.100231).

Data availability

Data will be made available on request.

References

- American Association of State Highway and Transportation Officials. (2018). *A policy on geometric design of highways and streets (7th edition)*. Washington, D.C: American Association of State Highway and Transportation Officials.
- Boden, B., Cabac, Y., Burschky, T., & Nagel, B. (2022). Rule-based verification of a geometric design using the codex framework. In , 2022. *33rd congress of the international council of the aeronautical sciences*. ICAS.
- Cabot, P. L. H., & Navigli, R. (2021). REBEL: Relation extraction by end-to-end language generation. *Findings of the association for computational linguistics: EMNLP 2021* (pp. 2370–2381).
- Casal, G., Santamarina, D., & Vázquez-Méndez, M. E. (2017). Optimization of horizontal alignment geometry in road design and reconstruction. *Transportation Research Part C: Emerging Technologies*, 74, 261–274.
- Chen, J., Xiao, S., Zhang, P., Luo, K., Lian, D., & Liu, Z. (2024). M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *Findings of the Association for Computational Linguistics ACL 2024*, 2318–2335.
- Chen, X., Jia, S., & Xiang, Y. (2020). A review: Knowledge reasoning over knowledge graph. *Expert Systems with Applications*, 141, Article 112948.
- Colucci, S., Donini, F. M., & Di Sciascio, E. (2024). A review of reasoning characteristics of RDF-based semantic web systems. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 14(4), Article e1537.
- Dalal, M., Chiruvolu, T., Chaplot, D., & Salakhutdinov, R. (2024). Plan-seq-learn: Language model guided rl for solving long horizon robotics tasks. *arXiv preprint. arXiv:2405.01534*.
- Du, J., Yu, A., Sugumaran, V., Hu, Y., Yu, G., & Zhang, J. (2025). Ontological reasoning for automated tunnel defect diagnosis and root cause identification. *Automation in Construction*, 178, Article 106362.
- Easa, S. M., & Yan, W. Y. (2019). Performance-based analysis in civil engineering: Overview of applications. *Infrastructures*, 4(2), 28.
- Findley, D. J. (2016). Highway geometric design. *Highway Engineering: Planning, Design, and Operations*, 167–253.
- Fitzpatrick, K., Das, S., Gates, T., Dixon, K. K., & Park, E. S. (2021). Considering roadway context in setting posted speed limits. *Transportation Research Record*, 2675(8), 590–602.
- Gao, J., Yu, B., Chen, Y., Gao, K., & Bao, S. (2025). A multi-perspective fusion model for operating speed prediction on highways using knowledge-enhanced graph neural networks. *Computer-Aided Civil and Infrastructure Engineering*, 40(8), 1004–1027.
- Gao, T., Li, Z., Gao, Y., Schonfeld, P., Feng, X., Wang, Q., & He, Q. (2022). A deep reinforcement learning approach to mountain railway alignment optimization. *Computer-Aided Civil and Infrastructure Engineering*, 37(1), 73–92.
- Gao, Y., Gao, T., Wu, Y., Wang, P., & He, Q. (2022). Low-construction-emission cross-section optimization for mountainous highway alignment designs. *Transportation Research Part D: Transport and Environment*, 105, Article 103249.
- Gao, Y., Xiong, G., Li, H., & Richards, J. (2024a). Exploring bridge maintenance knowledge graph by leveraging GraphSAGE and text encoding. *Automation in Construction*, 166, Article 105634.
- Gao, Y., Zhang, T., Zhu, C., Yang, S., Schonfeld, P., Zou, K., Zhang, J., Zhu, Y., Wang, P., & He, Q. (2024b). Biobjective optimization for railway alignment fine-grained designs with parallel existing railways. *Computer-Aided Civil and Infrastructure Engineering*, 39(3), 438–457.
- Guo, P., Xue, H., Ma, J., & Cheng, J. C. P. (2025). Advancing BIM information retrieval with an LLM-based query-domain-specific language and library code function alignment system. *Automation in Construction*, 178, Article 106374.
- Han, C., Han, T., Ma, T., Tong, Z., Wang, S., & Hei, T. (2023). End-to-end BIM-based optimization for dual-objective road alignment design with driving safety and construction cost efficiency. *Automation in Construction*, 151, Article 104884.
- Hirpa, D., Hare, W., Lucet, Y., Pushak, Y., & Tesfamariam, S. (2016). A bi-objective optimization framework for three-dimensional road alignment design. *Transportation Research Part C: Emerging Technologies*, 65, 61–78.
- He, Q., Gao, T., Gao, Y., Li, H., Schonfeld, P., Zhu, Y., Li, Q., & Wang, P. (2023). A bi-objective deep reinforcement learning approach for low-carbon-emission high-speed railway alignment design. *Transportation Research Part C: Emerging Technologies*, 147, Article 104006.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., ... Chen, W. (2022). Lora: Low-rank adaptation of large language models. *ICLR*, 1(2), 3.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., ... Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38.
- Jin, D., Yang, M., Qin, Z., Peng, J., & Ying, S. (2021). A weighting method for feature dimension by semisupervised learning with entropy. *IEEE Transactions on Neural Networks and Learning Systems*, 34(3), 1218–1227.
- Jung, S., Pyeon, J. H., Lee, H. S., Park, M., Yoon, I., & Rho, J. (2020). Construction cost estimation using a case-based reasoning hybrid genetic algorithm based on local search method. *Sustainability*, 12(19), 7920.
- Li, L., Wang, P., Yan, J., Wang, Y., Li, S., Jiang, J., ... Liu, Y. (2020). Real-world data medical knowledge graph: construction and applications. *Artificial Intelligence in Medicine*, 103, Article 101817.

- Li, W., Pu, H., Schonfeld, P., Yang, J., Zhang, H., Wang, L., & Xiong, J. (2017). Mountain railway alignment optimization with bidirectional distance transform and genetic algorithm. *Computer-Aided Civil and Infrastructure Engineering*, 32(8), 691–709.
- Li, W., Pu, H., Schonfeld, P., Zhang, H., & Zheng, X. (2016). Methodology for optimizing constrained 3-dimensional railway alignments in mountainous terrain. *Transportation Research Part C: Emerging Technologies*, 68, 549–565.
- Liu, C., Wang, Q., & Cao, Z. (2022). An intelligent optimization method for highway route selection based on comprehensive weight and TOPSIS. *PloS One*, 17(2), Article e0262588.
- Makke, N., & Chawla, S. (2024). Interpretable scientific discovery with symbolic regression: a review. *Artificial Intelligence Review*, 57(1), 2.
- Mavromatis, C., & Karypis, G. (2022). ReaRev: Adaptive reasoning for question answering over knowledge graphs. *arXiv preprint. arXiv:2210.13650*.
- Mavromatis, C., & Karypis, G. (2025). Gnn-rag: Graph neural retrieval for efficient large language model reasoning on knowledge graphs. *Findings of the Association for Computational Linguistics: ACL 2025*, 16682–16699.
- Ministry of Transport of the People's Republic of China. (2014). *Technical standard of highway engineering (JTG B01-2014)*. Beijing: China Communications Press.
- Ministry of Transport of the People's Republic of China. (2017). *Design specification for highway alignment (JTG D20-2017)*. Beijing: China Communications Press.
- Mohammed, A. A., Ambak, K., Mosa, A. M., & Syamsunur, D. (2019). Expert system in engineering transportation: A review. *Journal of Engineering Science and Technology*, 14(1), 229–252.
- Pu, H., Hu, T., Song, T., Schonfeld, P., Li, W., & Peng, L. (2024a). Knowledge graph-driven mountain railway alignment optimization integrating karst hazard assessment. *Applied Soft Computing*, 167, Article 112421.
- Pu, H., Hu, T., Song, T., Schonfeld, P., Wan, X., Li, W., & Peng, L. (2024b). Modeling and application of a customized knowledge graph for railway alignment optimization. *Expert Systems with Applications*, 244, Article 122999.
- Pu, H., Song, T., Schonfeld, P., Li, W., Zhang, H., Wang, J., ... Peng, X. (2019). A three-dimensional distance transform for optimizing constrained mountain railway alignments. *Computer-Aided Civil and Infrastructure Engineering*, 34(11), 972–990.
- Shen, X., Li, X., Zhou, B., Jiang, Y., & Bao, J. (2023). Dynamic knowledge modeling and fusion method for custom apparel production process based on knowledge graph. *Advanced Engineering Informatics*, 55, Article 101880.
- Su, C., Jiang, Q., Han, Y., Wang, T., & He, Q. (2025). Knowledge graph-driven decision support for manufacturing process: A graph neural network-based knowledge reasoning approach. *Advanced Engineering Informatics*, 64, Article 103098.
- Sun, W., Zou, Y., Zhang, X., Guo, N., Zhang, B., & Du, G. (2022). High robustness energy management strategy of hybrid electric vehicle based on improved soft actor-critic deep reinforcement learning. *Energy*, 258, Article 124806.
- Sushma, M. B., & Maji, A. (2020). A modified motion planning algorithm for horizontal highway alignment development. *Computer-Aided Civil and Infrastructure Engineering*, 35(8), 818–831.
- Sushma, M. B., Roy, S., & Maji, A. (2022). Exploring and exploiting ant colony optimization algorithm for vertical highway alignment development. *Computer-Aided Civil and Infrastructure Engineering*, 37(12), 1582–1601.
- Syamsunur, D., Ismail, A., Atiq, R., & Karim, O. A. (2011). Knowledge-based expert system for route selection of road alignment. *Australian Journal of Basic and Applied Sciences*, 5(5), 208–213.
- Wan, X., Pu, H., Schonfeld, P., Ran, Y., Song, T., Li, W., & Hu, J. (2025). Tunnel location optimization for railway alignment design in complex mountainous regions. *Automation in Construction*, 174, Article 106132.
- Wang, X., Chen, L., Ban, T., Usman, M., Guan, Y., Liu, S., ... Chen, H. (2021). Knowledge graph quality control: A survey. *Fundamental Research*, 1(5), 607–626.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.
- Yang, Z., Bi, Y., Wang, L., Cao, D., Li, R., & Li, Q. (2022). Development and application of a field knowledge graph and search engine for pavement engineering. *Scientific Reports*, 12(1), 7796.
- Yang, Z., Ding, M., Zhou, C., Yang, H., Zhou, J., & Tang, J. (2020). Understanding negative sampling in graph representation learning. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 1666–1676).
- Yang, Z., Qu, H., Fu, M., Hu, W., & Zhao, Y. (2021). A maximum divergence approach to optimal policy in deep reinforcement learning. *IEEE Transactions on Cybernetics*, 53(3), 1499–1510.
- Yang, N., Kang, M. W., Schonfeld, P., & Jha, M. K. (2014). Multi-objective highway alignment optimization incorporating preference information. *Transportation Research Part C: Emerging Technologies*, 40, 36–48.
- Yu, B., Chen, Y., Bao, S., & Xu, D. (2018). Quantifying drivers' visual perception to analyze accident-prone locations on two-lane mountain highways. *Accident Analysis & Prevention*, 119, 122–130.
- Zheng, C., Wong, S., Su, X., Tang, Y., Nawaz, A., & Kassem, M. (2025). Automating construction contract review using knowledge graph-enhanced large language models. *Automation in Construction*, 175, Article 106179.
- Zhai, S., Lu, X., Wang, C., Hong, Z., Shan, J., & Ma, Z. (2024). Correcting inconsistencies in knowledge graphs with correlated knowledge. *Big Data Research*, 36, Article 100450.
- Zhan, J., Fang, W., Love, P. E., & Luo, H. (2024). Explainable artificial intelligence: Counterfactual explanations for risk-based decision-making in construction. *IEEE Transactions on Engineering Management*, 71, 10667–10685.
- Zhang, T., Gao, Y., Gao, T., Schonfeld, P., Wu, Y., Zhu, Y., Yang, S., Wang, P., & He, Q. (2023). A sequential exploration algorithm for the design optimization of horizontal road alignment. *Computer-Aided Civil and Infrastructure Engineering*, 38(15), 2049–2071.
- Zhang, T., Xu, S., Gao, Y., Wang, P., Schonfeld, P., Zou, Y., & He, Q. (2024). 3D constrained Hybrid A*: Improved vehicle path planning algorithm for cost-effective road alignment design. *Automation in Construction*, 166, Article 105645.