MTAM: Multi-Level Trustworthiness Assessment via Neurosymbolic Synergy for Knowledge Graph Error Detection

Jinzhu Wang[12], Yongchao Gao[12†], Heng Qian[12], Yan Wang[3]
[1] *Key Laboratory of Computing Power Network and Information Security, Ministry of Education, Shandong Computer Science Center (National Supercomputer Center in Jinan), Qilu University of Technology (Shandong Academy of Sciences), Jinan, China*
[2] *Shandong Provincial Key Laboratory of Computing Power Internet and Service Computing, Shandong Fundamental Research Center for Computer Science, Jinan, China*
[3] *Shandong Standard Institute of Emerging Technologies and Innovations Co., Ltd., Jinan, China*
10431230170@stu.qlu.edu.cn, gaoych@sdas.org, qianh@sdas.org, wangyan@sieti.cn

***Abstract*—Automated knowledge graph (KG) construction inevitably introduces noise, resulting in inaccurate or logically inconsistent triples. Existing error detection methods typically rely on either structural embeddings or static constraints, often failing to capture the complex interplay between graph topology and ontological logic. To address this limitation, this paper proposes a Multi-level Trustworthiness Assessment Model (MTAM) via neurosymbolic synergy. Unlike static fusion approaches, our framework dynamically integrates three complementary dimensions: (1) Entity Association Strength, which filters low-confidence links based on local connectivity patterns; (2) Hierarchical Path Reasoning, which validates logical plausibility over multi-hop networks using an adaptive attention mechanism; and (3) Type-Structure Consistency, a novel metric that employs dual hypergraphs to align neural structural embeddings with symbolic type constraints. This dual-hypergraph design effectively detects sophisticated semantic anomalies where structural existence contradicts entity types. Extensive experiments on FB15k and YAGO43k demonstrate that MTAM significantly outperforms state-of-the-art baselines in both error detection precision and ranking correlation, validating its effectiveness in identifying diverse noise patterns.**

***Index Terms*—Knowledge Graph, Knowledge Assessment, Knowledge Trustworthiness, Error Detection.**

I. Introduction

Knowledge Graphs (KGs) serve as a critical component in many knowledge-driven AI systems but are often built via automated pipelines that extract information from unstructured text, inevitably introducing noise [1], such as incorrect entity links or spurious relations. These errors can propagate to downstream tasks, leading to unreliable outcomes in critical domains like healthcare and finance [2]. Therefore, identifying untrustworthy facts is essential for maintaining data quality [3].

Existing assessment approaches generally fall into structure-based or semantic-augmented categories, yet each has key limitations. Embedding methods tend to prioritize geometric proximity while ignoring ontological constraints; although recent works attempt to address this via contrastive learning [4], they still fail to align topology with entity types, often accepting structurally plausible but semantically invalid triples. In contrast, multi-modal methods typically rely on rigid, static fusion weights, lacking the flexibility to handle diverse error patterns—especially in sparse subgraphs where structural signals are weak but semantic constraints remain valid. More critically, most black-box scoring models lack interpretability, offering little insight into whether a rejection stems from weak topological association, logical conflicts, or type violations.

To address these gaps, we propose MTAM, a multi-level trustworthiness learning framework that integrates entity association, hierarchical paths, and type-structure consistency. Our neurosymbolic approach uses dual hypergraphs to align neural structural embeddings with symbolic entity type constraints. By enforcing ontological consistency within the embedding space, the model combines the generalization power of neural networks with the rigor of symbolic logic, enabling detection of sophisticated errors that escape traditional methods.

The main contributions are summarized as follows:

- We propose a comprehensive framework that synergistically fuses signals from entity, relation, and type-structure levels to robustly detect heterogeneous noise patterns.
- We design a hierarchical path modeling module with an adaptive attention mechanism, which dynamically weights multi-hop reasoning evidence to distinguish valid chains from spurious connections.
- We introduce a Type-Structure Consistency metric quantified via dual hypergraphs, effectively identifying semantic anomalies where structural links contradict explicit entity type constraints.

† Corresponding author.

II. Related Work

The field of KG trustworthiness assessment has evolved from statistical outlier detection to sophisticated neural representation learning, uncertainty-aware reasoning [5], and differential testing frameworks [6].

A. Structure-Based Assessment

Early research primarily utilized statistical distributions, matrix decomposition techniques, and graph topology to identify anomalies. Methods like CKRL [7] advanced this by integrating confidence scores into translation-based embedding models, treating trustworthiness as a learnable weight during training. However, these structure-centric methods rely solely on geometric latent spaces and often fail to capture explicit semantic violations, such as illogical entity type pairings.

B. Semantic and Type-Aware Approaches

To address the lack of semantics, recent works have incorporated entity types and attributes. Zhao et al. [8] and Moon et al. [9] proposed embedding entity hierarchies and type instances to enforce semantic validity. Similarly, Yu et al. [10] utilized entity descriptions to enhance representation richness. While these methods improve performance on specific tasks, they often treat entity types as static features, failing to model the complex, high-order correlations between relations and type constraints.

C. Multi-Signal Fusion and Error Detection

Current state-of-the-art methods aim to fuse multiple signals. Jia et al. [11] introduced a framework integrating entity, relation, and global features. More recently, Graph Neural Networks (GNNs) have been applied to this task; for instance, KGTrust [12] and other GNN-based models [13] aggregate neighborhood information to assess reliability. Zhang et al. [14] further proposed contrastive learning strategies for error detection.

Current fusion methods typically rely on rigid concatenation or fixed weights [15], [16], failing to adaptively prioritize diverse signals.

III. Methodology

A. Overview

In this work, we propose a Multi-level Trustworthiness Assessment Model (MTAM) to rigorously evaluate the validity of knowledge graph triples. As illustrated in Fig. 1, the framework integrates three synergistic modules to capture complementary evidence of reliability:

1) Entity Association (EA): Quantifies the prior probability of a relationship based on global connectivity patterns, positing that topologically proximal entities are more likely to form valid links.
2) Hierarchical Path Reasoning (HP): Evaluates the reasoning chain between head and tail entities by scrutinizing Semantic Consistency, Translational Plausibility, and Transition Smoothness.
3) Type-Structure Consistency (TC): Assesses the alignment between the observed structural context and the latent semantic constraints imposed by entity types via a dual-view attention mechanism.

The outputs of these modules are fused via a non-linear projection to predict the final trustworthiness probability.

B. Problem Formulation and Notation

Let $\mathcal{G} = (\mathcal{E}, \mathcal{R})$ denote a Knowledge Graph (KG), where $\mathcal{E}$ and $\mathcal{R}$ represent the sets of entities and relations, respectively. The task is to compute a trustworthiness score $P(y|x)$ for a given triple $x = (h, r, t) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}$.

We denote the d-dimensional embedding vectors corresponding to entities h, t and relation r as $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{R}^d$. Each entity e is associated with a set of ontological types $\mathcal{C}_e \subseteq \mathcal{C}$. The embedding matrix for the global type set is denoted as $\mathbf{M}^{\mathcal{C}} \in \mathbb{R}^{|\mathcal{C}| \times d}$. Key notations are summarized in Table I.

TABLE I
Summary of Key Notations

Symbol	Description
$\mathcal{G} = (\mathcal{E}, \mathcal{R})$	Knowledge Graph with entity/relation sets
(h, r, t)	Target triple (head, relation, tail)
$\mathbf{h}, \mathbf{r}, \mathbf{t}$	Embeddings for h, r, t
$\mathcal{C}_e$	Set of types associated with entity e
$\mathbf{M}^{\mathcal{C}}$	Global type embedding matrix
$\mathbf{m}_k$	Embedding vector for type $k \in \mathcal{C}$
$\mathcal{P}_{h,t}$	Set of paths connecting h and t
$\eta_{h,t}$	Total number of paths between h and t
s_{ea}, s_{hp}, s_{tc}	Scores from EA, HP, and TC modules

C. Entity Association Strength (EA)

The EA module assesses coarse-grained relational intensity. We extend the resource allocation principle [11] to construct a connectivity feature vector $\mathbf{v}_{ea} \in \mathbb{R}^6$. This vector encompasses:

- *Resource Flow*: Bi-directional resource allocation scores $RA(h \to t)$ and $RA(t \to h)$.
- *Degree Centrality*: In-degree and out-degree for both h and t, reflecting node importance.
- *Path Reachability*: The count of distinct paths $\eta_{h,t}$ connecting h and t.

The association score s_{ea} is computed via a learnable linear transformation followed by a sigmoid activation:

$$s_{ea}(h, t) = \sigma(\mathbf{w}_{ea}^{\top} \mathbf{v}_{ea} + b_{ea}) \tag{1}$$

where $\mathbf{w}_{ea}$ and b_{ea} are trainable parameters.

D. Hierarchical Path Information (HP)

While EA captures global topology, the HP module scrutinizes the *quality* of the top-N paths $\mathcal{P}_{h,t} = \{p_1, \ldots, p_N\}$ connecting h and t. A path p_k of length l consists of a sequence of entities and relations. We evaluate each path using three indicators:

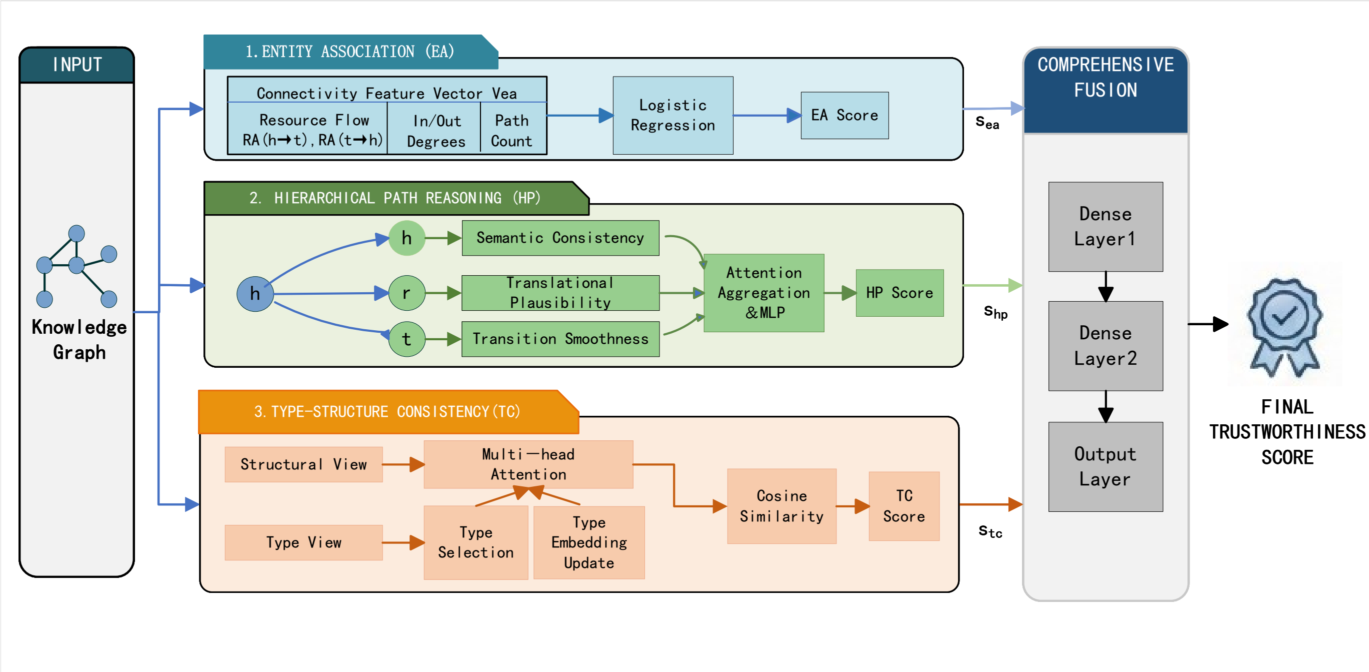

Fig. 1. The proposed framework consists of three modules: Entity Association Strength(EA), Hierarchical Path Reasoning (HP), and Type-Structure Consistency (TC). These modules extract multi-view features which are fused to compute the final trustworthiness score.

1) **Semantic Consistency (ϕ_{sem}):** Measures type compatibility between adjacent entities.

$$\phi_{sem}(p_k) = \frac{1}{l}\sum_{j=0}^{l-1}\cos(\bar{\mathbf{c}}_{e_j}, \bar{\mathbf{c}}_{e_{j+1}}) \tag{2}$$

where $\bar{\mathbf{c}}_e$ is the centroid embedding of entity types.

2) **Translational Plausibility (ϕ_{trans}):** Penalizes structural violations based on the translation assumption ($\mathbf{h}+\mathbf{r} \approx \mathbf{t}$):

$$\phi_{trans}(p_k) = -\frac{1}{l}\sum_{j=1}^{l}\|\mathbf{e}_{j-1} + \mathbf{r}_j - \mathbf{e}_j\|_2^2 \tag{3}$$

3) **Transition Smoothness (ϕ_{smt}):** Quantifies entity-relation compatibility via projection:

$$\phi_{smt}(p_k) = \frac{1}{l}\sum_{j=1}^{l}\sigma(\mathbf{e}_{j-1}^\top \mathbf{r}_j)\cdot\sigma(\mathbf{r}_j^\top \mathbf{e}_j) \tag{4}$$

We concatenate the indicators into a feature vector $\mathbf{f}_k$. An attention mechanism aggregates these path features, weighted by their relevance to the query relation. The final HP score incorporates a regularization term for path scarcity:

$$s_{hp}(h,r,t) = \sigma\left(\mathrm{MLP}(\mathbf{z}_{path} \oplus \mathbf{r}) + \beta\log(\eta_{h,t}+1)\right) \tag{5}$$

where $\mathbf{z}_{path}$ is the weighted sum of path representations and $\oplus$ denotes concatenation.

E. Type-Structure Consistency (TC)

This module verifies if the triple's structural context aligns with the semantic constraints of its types. Ideally, valid triples should exhibit consistency across both structural topology and semantic type subspaces. To achieve this, we employ a symmetric **Multi-Head Attention (MHA)** mechanism for both views, where the number of attention heads is set to $H = 8$.

1) Structural View: To capture complex topological dependencies among triples (e.g., distinct structural patterns defined by shared heads vs. shared tails), we employ a *Structure-Aware MHA*.

We construct a dual relation graph where nodes represent triples. Let $\mathbf{u}_i^{(s)} = [\mathbf{h};\mathbf{r};\mathbf{t}]$ be the initial embedding of triple x_i, and $\mathcal{N}_i$ be its set of neighbor triples. We treat $\mathbf{u}_i^{(s)}$ as the *query*, and its neighbors $\{\mathbf{u}_j^{(s)}\}_{j\in\mathcal{N}_i}$ as *keys* and *values*. For the m-th attention head ($m = 1,\dots,H$), the structural attention coefficient is:

$$\beta_{i,j}^{(m)} = \mathrm{Softmax}\left(\frac{(\mathbf{W}_Q^{(s,m)}\mathbf{u}_i^{(s)})^\top(\mathbf{W}_K^{(s,m)}\mathbf{u}_j^{(s)})}{\sqrt{d/H}}\right) \tag{6}$$

The head-specific structural context is aggregated as $\mathbf{z}_m^{(s)} = \sum_{j\in\mathcal{N}_i}\beta_{i,j}^{(m)}(\mathbf{W}_V^{(s,m)}\mathbf{u}_j^{(s)})$. The final structure-aware embedding $\tilde{\mathbf{u}}_i^{(s)}$ is obtained by projecting the concatenated heads:

$$\tilde{\mathbf{u}}_i^{(s)} = \mathrm{LayerNorm}\left(\mathbf{u}_i^{(s)} + \mathbf{W}_O^{(s)}[\mathbf{z}_1^{(s)} \oplus \cdots \oplus \mathbf{z}_H^{(s)}]\right) \tag{7}$$

where residual connection and layer normalization are applied to stabilize training.

2) Type View: To capture diverse semantic dependencies (e.g., a relation may attend to specific type subspaces), we employ a *Type-Aware MHA*.

We reconstruct the triple representation using type information. For an entity h, we treat the relation embedding $\mathbf{r}$ as the *query*, and its type embeddings $\{\mathbf{m}_k\}_{k\in\mathcal{C}_h}$ as *keys* and *values*. For the m-th attention head, the semantic relevance is measured as:

$$\alpha_{h,k}^{(m)} = \text{Softmax}\left(\frac{(\mathbf{W}_Q^{(t,m)}\mathbf{r})^\top(\mathbf{W}_K^{(t,m)}\mathbf{m}_k)}{\sqrt{d/H}}\right) \tag{8}$$

The type-refined entity representation for the m-th head is $\mathbf{z}_m^{(t)} = \sum_k \alpha_{h,k}^{(m)}(\mathbf{W}_V^{(t,m)}\mathbf{m}_k)$. The comprehensive type-based entity embedding $\hat{\mathbf{h}}$ is:

$$\hat{\mathbf{h}} = \mathbf{W}_O^{(t)}[\mathbf{z}_1^{(t)} \oplus \cdots \oplus \mathbf{z}_H^{(t)}] \tag{9}$$

The tail embedding $\hat{\mathbf{t}}$ is computed analogously. Finally, the type-view triple embedding $\tilde{\mathbf{u}}_i^{(t)} = [\hat{\mathbf{h}};\mathbf{r};\hat{\mathbf{t}}]$ serves as the semantic counterpart to the structural embedding.

3) Consistency Calculation: The TC score is defined as the cosine similarity between the two views, quantifying the semantic alignment:

$$s_{tc}(h,r,t) = \frac{(\tilde{\mathbf{u}}_i^{(s)})^\top\tilde{\mathbf{u}}_i^{(t)}}{\|\tilde{\mathbf{u}}_i^{(s)}\|_2\|\tilde{\mathbf{u}}_i^{(t)}\|_2} \tag{10}$$

F. Model Fusion and Training

The final trustworthiness probability is computed by fusing the module scores $\mathbf{s} = [s_{ea}, s_{hp}, s_{tc}]$ via a multi-layer perceptron:

$$P(y=1|x) = \sigma(\mathbf{w}_{out}^\top\text{ReLU}(\mathbf{W}_f\mathbf{s} + \mathbf{b}_f)) \tag{11}$$

The model is trained using binary cross-entropy loss over positive observed triples $\mathcal{D}^+$ and negative sampled triples $\mathcal{D}^-$:

$$\mathcal{L} = -\sum_{x\in\mathcal{D}^+}\log P(1|x) - \sum_{x'\in\mathcal{D}^-}\log(1-P(1|x')) \tag{12}$$

The optimization is performed using the Adam algorithm [17] to adaptively adjust the learning rates.

IV. Experiments

In this section, we conduct comprehensive experiments on multiple real-world knowledge graphs to verify the effectiveness of the proposed method. We specifically focus on the following core questions: **Q1:** Can our method effectively distinguish between noisy triples and valid triples? **Q2:** How does the learned KG embedding perform in error detection, and does it demonstrate a significant advantage over state-of-the-art baselines? **Q3:** What are the independent contributions of the individual modules? **Q4:** For the three path factors in the hierarchical path module, what are their specific contributions to the trustworthiness assessment?

A. Datasets

Experiments were conducted on the benchmark datasets FB15k [18] and YAGO43k [9]. We also utilized their corresponding entity type datasets constructed in the same study [9]. Following the procedure described in [7], we generated synthetic noisy versions by randomly replacing the head entity, tail entity, or relation in triples under varying noise ratios.

B. Implementation Details

Experiments were conducted on NVIDIA A100 GPUs, with entity, relation, and type embeddings uniformly set to dimension $d = 100$. A batch size of 50 was adopted during training, and parameters were optimized using Adam [17] with an initial learning rate of 10^{-3}. To mitigate overfitting, we applied Dropout regularization [19] and employed early stopping based on validation performance [20].

C. Baselines and Evaluation Metrics

The proposed detection method is compared against several state-of-the-art KG embedding models, including TransE [18], DistMult [21], ComplEx [22], Bilinear [23], MLP [24], TransH [25], TransR [26], TransD [27], and PTransE [28]. Additionally, we compare with specialized error detection approaches: CKRL [7], TTMF [11], KGIst [29], and CAGED [14].

We employ Accuracy and F1 Score for classification tasks on balanced datasets (FB15k-N5), and Precision@K and Recall@K for ranking tasks on datasets with a 5% error rate.

D. Validity of the Method (Q1)

To validate the effectiveness of our model in distinguishing noisy triples from correct ones, we constructed datasets FB15k-N3 and YAGO43k-N3 with a 30% noise rate based on the FB15k and YAGO43k datasets, respectively. We plotted the distribution of the triple trustworthiness scores within a coordinate system where the y-axis ranges from 0 to 1, as illustrated in Fig. 2. The results indicate that the trustworthiness scores of positive examples are mainly concentrated above 0.5, whereas those of negative examples are mostly distributed below 0.5. This distribution pattern aligns with the intuitive understanding of triple trustworthiness, confirming the rationality and effectiveness of the scores generated by our model.

E. Model Comparison in Error Detection (Q2)

We conducted a comprehensive comparison with existing KG embedding methods and state-of-the-art error detection baselines to verify the superiority of the proposed approach. These comparisons cover key evaluation metrics including Precision, Recall, and F1-score, ensuring a holistic assessment of model performance across different dimensions. The detailed experimental results are presented in Tables II and III, where the best results are highlighted in bold and the second-best results are underlined. The results confirm that our method outperforms existing approaches across key metrics, validating the effectiveness of our design.

TABLE II
PERFORMANCE COMPARISON OF DIFFERENT MODELS ON FB15K-N5

Model	MLP	Bilinear	TransE	TransH	TransD	TransR	PTransE	Our
Accuracy	0.833	0.861	0.868	0.912	0.913	0.902	0.941	**0.971**
F1-score	0.846	0.869	0.876	0.913	0.913	0.904	0.942	**0.972**

TABLE III
ERROR DETECTION PERFORMANCE (PRECISION@K AND RECALL@K) ON FB15K WITH 5% ANOMALY RATIO

Model	$K = 1\%$		$K = 2\%$		$K = 3\%$		$K = 4\%$		$K = 5\%$	
	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall
TransE	0.756	0.151	0.674	0.270	0.605	0.363	0.552	0.442	0.503	0.503
ComplEx	0.718	0.143	0.651	0.260	0.590	0.354	0.538	0.430	0.491	0.491
DistMult	0.709	0.141	0.646	0.258	0.582	0.349	0.531	0.425	0.484	0.484
CKRL	0.789	0.158	0.736	0.294	0.684	0.410	0.624	0.499	0.569	0.569
TTMF	0.815	0.163	0.767	0.307	0.713	0.428	0.651	0.521	0.593	0.593
KGIst	0.825	0.165	0.754	0.302	0.703	0.422	0.642	0.513	0.585	0.585
CAGED	0.852	0.171	0.796	0.318	0.735	0.441	0.671	0.537	0.612	0.612
Our	**0.883**	**0.177**	**0.818**	**0.327**	**0.753**	**0.451**	**0.692**	**0.553**	**0.628**	**0.628**

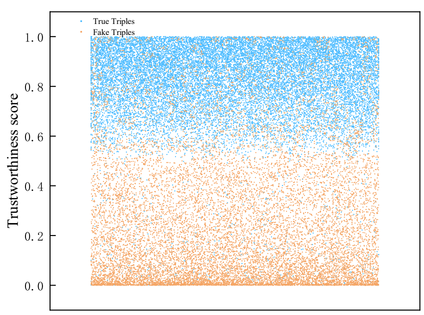

(a) FB15k-N3

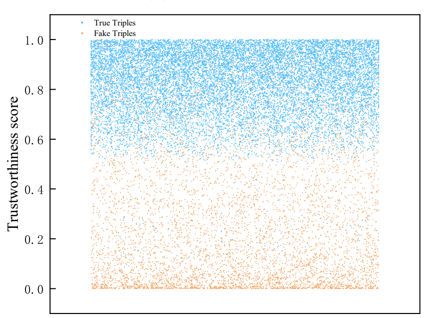

(b) YAGO43k-N3

Fig. 2. Scatter plot of triple trustworthiness scores distribution.

F. Ablation Experiment on Core Module Contributions (Q3)

To assess the contributions of the components in our model, we conducted an ablation study. Specifically, we designed three ablation variants: (1) Our_NEA: No Entity Association (NEA), excluding the entity association strength module; (2) Our_NHP: No Hierarchical Path (NHP), excluding the Hierarchical Path Information module; (3) Our_NTC: No Type Structural Consistency (NTC), excluding the Entity Type and Structural Consistency modules. The results on the FB15k dataset are shown in Table IV. Our full model consistently outperforms all variants in terms of precision and recall. This demonstrates that each component is essential and contributes to the model's superior performance.

TABLE IV
ABLATION EXPERIMENT ON CORE MODULE CONTRIBUTIONS ON FB15K WITH 5% ANOMALY RATIO

		FB15k				
		K=1%	K=2%	K=3%	K=4%	K=5%
Pre@K	Our_NTC	0.869	0.806	0.740	0.679	0.616
	Our_NHP	0.872	0.810	0.742	0.683	0.619
	Our_NEA	0.876	0.809	0.747	0.685	0.620
	Our	**0.883**	**0.818**	**0.753**	**0.692**	**0.628**
Rec@K	Our_NTC	0.174	0.322	0.444	0.543	0.616
	Our_NHP	0.174	0.324	0.445	0.546	0.619
	Our_NEA	0.175	0.323	0.447	0.548	0.620
	Our	**0.177**	**0.327**	**0.451**	**0.553**	**0.628**

G. Ablation Study on Hierarchical Path Sub-Factors (Q4)

To assess the independent contributions of the path validity indicators defined in Section III-D (SC: Semantic Consistency ϕ_{sem}; TP: Translational Plausibility ϕ_{trans}; TS: Transition Smoothness ϕ_{smt}), we conducted an ablation study on the FB15k dataset with a 5% anomaly ratio. Three variants

(Our_NSC, Our_NTP, Our_NTS) were constructed by removing one factor at a time, with the full model (Our) serving as the baseline. Evaluation metrics included Precision@K and Recall@K ($K = 1\%$–5%). Results (see Table V) show that all variants underperform the full model, with consistently lower Precision@K and Recall@K values. This confirms that each indicator provides unique, irreplaceable signals, and all are essential to the module's error-detection performance.

TABLE V
ABLATION STUDY OF HIERARCHICAL PATH SUB-FACTORS ON FB15K WITH 5% ANOMALY RATIO

		FB15k				
		K=1%	K=2%	K=3%	K=4%	K=5%
Pre@K	Our_NSC	0.874	0.811	0.744	0.684	0.621
	Our_NTP	0.881	0.817	0.752	0.690	0.627
	Our_NTS	0.879	0.815	0.748	0.688	0.624
	Our	**0.883**	**0.818**	**0.753**	**0.692**	**0.628**
Rec@K	Our_NSC	0.175	0.324	0.446	0.547	0.621
	Our_NTP	0.176	0.326	0.451	0.553	0.627
	Our_NTS	0.175	0.325	0.449	0.550	0.624
	Our	**0.177**	**0.327**	**0.451**	**0.553**	**0.628**

V. CONCLUSION

In this paper, we have presented MTAM, a novel trustworthiness assessment framework that moves beyond static feature aggregation to achieve neurosymbolic synergy. By integrating Entity Association, Hierarchical Path Reasoning, and Type-Structure Consistency, our model effectively captures the complex interplay between graph topology and semantic constraints. A key innovation lies in the dual-hypergraph architecture equipped with Multi-Head Attention, which aligns neural structural embeddings with symbolic type rules to detect sophisticated errors. Extensive experiments on FB15k and YAGO43k demonstrate that MTAM significantly outperforms state-of-the-art baselines in both error detection precision and ranking correlation. For future work, we plan to extend this framework to generative knowledge graph construction and explore the integration of Large Language Models (LLMs) to further enhance reasoning capabilities in sparse data scenarios.

REFERENCES

[1] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Trans. Knowl. Data Eng.*, vol. 29, no. 12, pp. 2724–2743, 2017.

[2] J. Z. Pan, S. Pavlova, C. Li, N. Li, Y. Li, and J. Liu, "Content based fake news detection using knowledge graphs," in *Proc. ISWC*, 2018, pp. 669–683.

[3] Y. Yang, J. Chen, and Y. Xiang, "A review on the reliability of knowledge graph: from a knowledge representation learning perspective," *World Wide Web*, vol. 28, no. 1, p. 4, 2025.

[4] X. Liu, Y. Liu, and W. Hu, "Knowledge graph error detection with contrastive confidence adaption," in *Proc. AAAI*, vol. 38, no. 8, 2024, pp. 8824–8831.

[5] B. Ni, Y. Wang, L. Cheng, E. Blasch, and T. Derr, "Towards trustworthy knowledge graph reasoning: An uncertainty aware perspective," in *Proc. AAAI*, vol. 39, no. 12, 2025, pp. 12 417–12 425.

[6] J. Tan, D. Wang, J. Sun, Z. Liu, X. Li, and Y. Feng, "Towards assessing the quality of knowledge graphs via differential testing," *Inf. Softw. Technol.*, vol. 174, p. 107521, 2024.

[7] R. Xie, Z. Liu, F. Lin, and L. Lin, "Does william shakespeare really write hamlet? knowledge representation learning with confidence," in *Proc. AAAI*, 2018, pp. 4954–4961.

[8] Y. Zhao, Z. Li, W. Deng, R. Xie, and Q. Li, "Learning entity type structured embeddings with trustworthiness on noisy knowledge graphs," *Knowl.-Based Syst.*, vol. 215, p. 106630, 2021.

[9] C. Moon, P. Jones, and N. F. Samatova, "Learning entity type embeddings for knowledge graph completion," in *Proc. CIKM*, 2017, pp. 2215–2218.

[10] Y. Zhao, H. Feng, and P. Gallinari, "Embedding learning with triple trustiness on noisy knowledge graph," *Entropy*, vol. 21, no. 11, p. 1083, 2019.

[11] S. Jia, Y. Xiang, X. Chen, and K. Wang, "Triple trustworthiness measurement for knowledge graph," in *Proc. WWW*, 2019, pp. 2865–2871.

[12] Z. Yu *et al.*, "Kgtrust: Evaluating trustworthiness of siot via knowledge enhanced graph neural networks," in *Proc. WWW*, 2023, pp. 727–736.

[13] Z. Shi, H. Li, D. Zhao, and C. Pan, "Research on quality assessment methods for cybersecurity knowledge graphs," *Comput. Secur.*, vol. 142, p. 103848, 2024.

[14] Q. Zhang, J. Dong, K. Duan, X. Huang, Y. Liu, and L. Xu, "Contrastive knowledge graph error detection," in *Proc. CIKM*, 2022, pp. 2590–2599.

[15] Q. Zhang, J. Dong, Q. Tan, and X. Huang, "Integrating entity attributes for error-aware knowledge graph embedding," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 4, pp. 1667–1682, 2023.

[16] Z. Zhang, F. Zhang, F. Zhuang, and Y. Xu, "Knowledge graph error detection with hierarchical path structure," in *Proc. CIKM*, 2023, pp. 4430–4434.

[17] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[18] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *Proc. NIPS*, 2013.

[19] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.

[20] L. Prechelt, "Early stopping-but when?" in *Neural Networks: Tricks of the Trade*. Springer, 1998, pp. 55–69.

[21] B. Yang, W. Yih, X. He, J. Gao, and L. Deng, "Embedding entities and relations for learning and inference in knowledge bases," in *Proc. ICLR*, 2015.

[22] T. Trouillon, J. Welbl, S. Riedel, E. Gaussier, and G. Bouchard, "Complex embeddings for simple link prediction," in *Proc. ICML*, 2016, pp. 2071–2080.

[23] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor, "Freebase: a collaboratively created graph database for structuring human knowledge," in *Proc. SIGMOD*, 2008, pp. 1247–1250.

[24] X. Dong *et al.*, "Knowledge vault: A web-scale approach to probabilistic knowledge fusion," in *Proc. SIGKDD*, 2014, pp. 601–610.

[25] Z. Wang, J. Zhang, J. Feng, and Z. Chen, "Knowledge graph embedding by translating on hyperplanes," in *Proc. AAAI*, 2014.

[26] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu, "Learning entity and relation embeddings for knowledge graph completion," in *Proc. AAAI*, vol. 29, no. 1, 2015.

[27] G. Ji, S. He, L. Xu, K. Liu, and J. Zhao, "Knowledge graph embedding via dynamic mapping matrix," in *Proc. ACL*, 2015, pp. 687–696.

[28] Y. Lin, Z. Liu, H. Luan, M. Sun, S. Rao, and S. Liu, "Modeling relation paths for representation learning of knowledge bases," in *Proc. EMNLP*, 2015.

[29] C. Belth, X. Zheng, J. Vreeken, and D. Koutra, "What is normal, what is strange, and what is missing in a knowledge graph: Unified characterization via inductive summarization," in *Proc. WWW*, 2020.