# Semantic-Enriched Heterogeneous Embeddings for Enhanced Node Classification

**Silvio Fernando Angonese**[1], **Renata Galante**[1]

[1]Institute of Informatics – Federal University of Rio Grande do Sul (UFRGS)
P.O. Box 15.064 – 91.501-970 – Porto Alegre – RS – Brazil

{sfangonese,galante}@inf.ufrgs.br

***Abstract.*** *Information Systems increasingly rely on heterogeneous and multimodal data, making it challenging to build semantically informative node representations. This paper proposes a process for generating semantic-enriched heterogeneous embeddings by integrating local features, multimodal signals, and ontology-driven metapath design into a unified representation. Experiments on a Person Relationship Graph show that embedding compositions outperform single embeddings, achieving higher F1-macro scores and more stable results. The findings indicate that combining multimodal semantics with ontology-guided metapath reasoning improves representation quality and generalization in heterogeneous environments.*

***Resumo.*** *Sistemas de Informação dependem cada vez mais de dados heterogêneos e multimodais, o que dificulta a construção de representações de nós informativas. Este trabalho propõe um processo para geração de embeddings heterogêneos enriquecidos semanticamente, integrando características locais, representações multimodais e metapaths orientados por ontologia. Experimentos em um grafo de Pessoas demonstram que composições de embeddings superam embeddings individuais, alcançando maiores valores de F1-macro e maior estabilidade. Os resultados indicam que a combinação de semântica multimodal com metapaths orientados por ontologia melhora a qualidade das representações e a generalização em ambientes heterogêneos.*

## 1. Introduction

Information Systems increasingly rely on heterogeneous, interconnected, and multimodal data, where entities such as people, products, and devices interact through text, images, structured attributes, and relational links [Liebenberg and Jarke 2020]. Heterogeneous graphs provide an effective modeling paradigm for representing such environments, enabling the capture of rich semantic context and complex relationships. Learning expressive node representations in these settings is essential to support analytical tasks such as classification, recommendation, and decision support [Wang et al. 2023].

A key challenge in Information Systems Engineering is designing node representations that remain robust and semantically meaningful as systems evolve. In real-world Information System, node types, relationships, and multimodal data vary in availability and quality [Jr. et al. 2013]. Under these conditions, models based solely on structural or feature-based embeddings often struggle. Incorporating multimodal semantic features and higher-order structural patterns improves representation quality and generalization.

The objective of this paper is to propose a unified process for generating semantic-enriched heterogeneous node embeddings for Information Systems. Building on the AGHE approach, it integrates multimodal node content, metapath-based structural patterns, and neighborhood information into representations compatible with standard supervised classifiers. The approach is evaluated under full-data and controlled sparsity conditions to assess robustness and generalization. Results show that combining multimodal semantics with metapath-aware structures produces more stable and discriminative representations. The contributions of this paper can be summarized as follows:

- A unified process for generating semantic-enriched heterogeneous embeddings for Information Systems;
- A multimodal graph construction procedure integrating text-, image-, and attribute-based features;
- The definition of embedding compositions combining local features, neighbor aggregation, and metapath semantics;
- An evaluation protocol assessing robustness and generalization under full-data and feature-sparsity settings;
- An ablation study comparing single and composed embeddings using a GAT-based classifier.

This paper builds upon previous AGHE-based works on heterogeneous embedding composition and multimodal semantic enrichment [Angonese and Galante 2024b, Angonese and Galante 2024a, Angonese and Galante 2025], extending them with ontology-guided metapath design and robustness evaluation under feature sparsity conditions.

The remainder of this paper is organized as follows. Section 2 introduces the theoretical background. Section 3 depicts the AGHE approach that forms the basis of this work. Section 4 presents the process for semantic composition and robustness. Section 5 reports the experimental evaluation. Finally, Section 6 summarizes the conclusions and outlines future research directions.

## 2. Related Work

Research on heterogeneous graph representation learning has advanced significantly, aiming to encode multi-typed structures where node and edge categories play distinct semantic roles. Early inductive approaches such as GraphSAGE [Hamilton et al. 2017] introduced neighborhood aggregation mechanisms that scale to dynamic environments but treat neighbors uniformly, limiting the capture of type-specific semantics. Metapath-based techniques, including MetaPath2Vec [Dong et al. 2017], highlighted how type-guided traversal patterns encode meaningful connectivity, while heterogeneous GNNs such as MAGNN [Fu et al. 2020] and HAGNN [Zhu et al. 2025] extended this idea through intra-metapath aggregation and multi-level attention. These works demonstrate the value of structured semantic signals but remain predominantly focused on graph structure. Another related line of work applies graph neural networks to ontology-driven conceptual models, learning embeddings over OntoUML structures to support ontological category prediction, which is closely aligned with our goal of deriving semantically rich node representations for Information Systems [Ali et al. 2023].

Multimodal extensions have investigated incorporating textual or visual evidence into graph models, for example via superpixel-based image representations [Avelar et al. 2020] or node-level visual encodings. These methods complement structural signals with external modalities but do not fully integrate multimodal evidence with type-aware structural reasoning.

A complementary line of research has explored embedding composition strategies. Prior work introduced compositions that combine local features, neighborhood information, and metapath-based semantics for node classification [Angonese and Galante 2024b, Angonese and Galante 2024a]. Additional work addressed heterogeneous data quality by deriving single and composed embeddings from diverse node attributes [Angonese and Galante 2025]. Together, these contributions establish the technical foundation for the present paper, which advances this line of research by unifying multimodal semantics and metapath-aware structural cues into a single process and examining their impact on robustness and reliability in Information Systems relevant settings.

## 3. AGHE Approach for Embedding Generation

The AGHE - Generating Enhanced Node Embeddings in Heterogeneous Graphs [Angonese and Galante 2024a] is an approach for constructing enriched node representations in heterogeneous graphs. It defines how multimodal content such as text, images, and structured subgraphs is incorporated into graph construction, and how individual semantic encodings are produced from this information. AGHE establishes a complete pipeline for creating heterogeneous representations that combine local features, neighborhood information, and type-guided patterns. Figure 1 summarizes this pipeline, which unfolds in four stages. The steps in white correspond to components already established in previous publications, while the shaded region highlights the extension introduced in this paper.

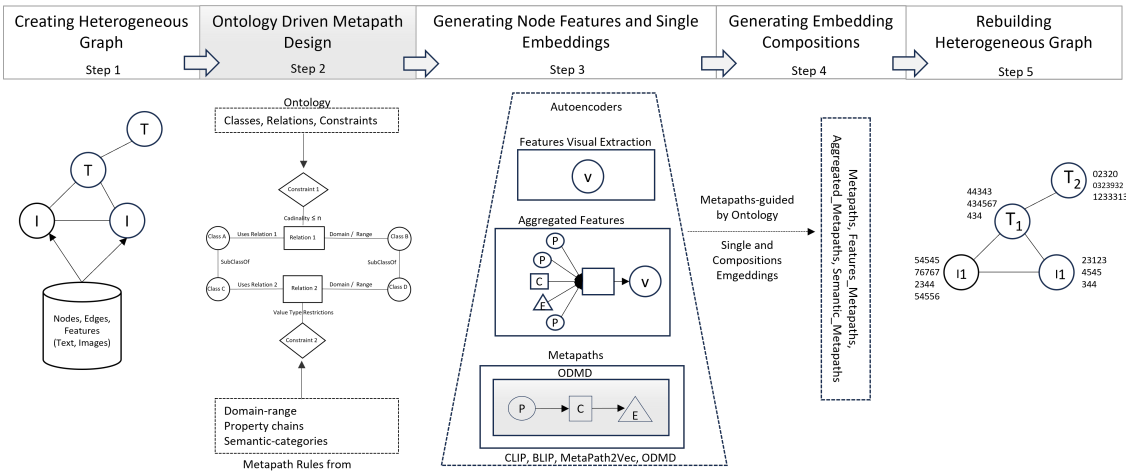


**Figura 1. Steps for graph creation, metapaths definition by ontology, embedding generation, and graph delivery.**

At a high level, the pipeline operates as follows. *Step 1* constructs the heterogeneous graph by defining node types, relations, and multimodal attributes derived from text, images, and structured metadata. *Step 2* introduces ontology-driven metapath design, where traversal patterns are derived from ontology-defined relations to ensure semantic validity and reduce arbitrary graph exploration. In the original AGHE pipeline, metapaths were manually defined based on domain knowledge. *Step 3* generates single embeddings by transforming raw node attributes into semantic features, including textual encodings, visual information, aggregated neighborhood signals, and guided traversal patterns that reflect type-aware connectivity. *Step 4* produces embedding compositions such as *Features_Metapaths*, *Aggregated_Metapaths*, and *Semantic_Metapaths*, combining complementary signals into unified representations. *Step 5* rebuilds the heterogeneous graph by assigning single or composed embeddings to nodes, yielding a transformed representation that can be directly consumed by graph neural network models and other supervised classifiers.

Building on this pipeline, the present paper introduces a dedicated process for strengthening semantic compositions and assessing their robustness under information sparsity. The shaded region in Figure 1 represents this extension, which systematizes how individual embeddings are combined and evaluated to enhance expressiveness and support robustness analyses in Information Systems Engineering.

## 4. Proposed Process for Semantic Composition and Robustness

This section presents a process for strengthening semantic compositions and assessing their robustness in a heterogeneous Information System. While AGHE defines the overall embedding pipeline, the proposed process focuses on how semantic and structural signals are combined and evaluated under varying information conditions. Figure 2 provides an overview. The main contributions introduced in this paper are the ontology-guided metapath design, the semantic composition strategy integrating multimodal and structural signals, and the robustness evaluation under controlled feature sparsity.

### 4.1. Process Overview and Definitions

The process extends the composition step of AGHE by integrating multimodal, neighborhood-aware, and metapath-guided information into unified representations. It includes the generation of composition embeddings and mechanisms to assess robustness under feature sparsity.

Single embeddings form the basis of composition. Three types are considered: *Feature* (derived from node attributes such as text, images, and structured data), *Aggregated* (incorporating neighborhood information), and *Semantic* (derived from multimodal descriptions using models such as CLIP [Radford et al. 2021] and BLIP [Dai et al. 2023]). The unified representation is defined as:

$$\mathbf{h}_v = f_{\text{single}}(v),$$

where $f_{\text{single}}$ represents the extraction function.

Composition embeddings combine local and structural information. Let $\mathbf{h}_v$ be a single embedding and $\mathbf{h}_v^m$ the metapath-based representation. The composition is defined as:

$$\mathbf{z}_v = \mathbf{h}_v \parallel \mathbf{h}_v^m.$$

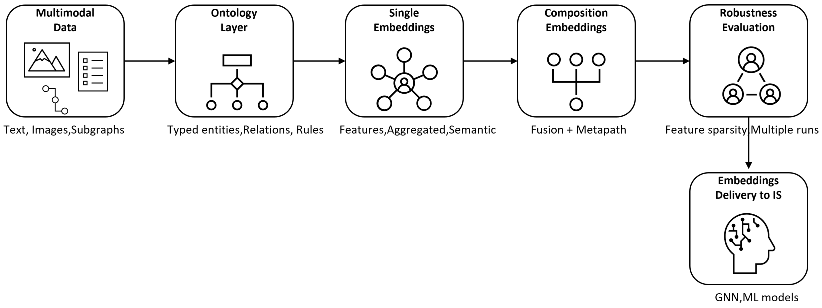


**Figura 2. Conceptual process linking multimodal data with Information Systems applications.**

### 4.2. Proposed Process

The process evaluates how composition embeddings behave under varying data conditions, particularly in the presence of missing or degraded node attributes. Robustness is assessed through controlled degradation of node features, simulating different levels of information loss, combined with repeated $k$-fold cross-validation to ensure statistical reliability. From an engineering perspective, the process is designed to remain compatible with Information Systems requirements. Semantic enrichment is performed upstream, allowing the use of standard classifiers such as GAT without architectural changes. Embeddings are modular, and metapaths are treated as configurable templates, enabling adaptation to evolving systems. Metapath-based representations are derived from ontology-defined relations, ensuring that all structural patterns used in the embedding process remain semantically valid and avoiding arbitrary graph traversals.

## 5. Evaluation from an Information Systems Engineering Perspective

The empirical evaluation validates the proposed process with focus on key Information Systems Engineering qualities: *Reliability*, *Discriminative Capacity*, and *Robustness*.

*Experiment I*, conducted on the *Person Relationships* heterogeneous graph, evaluates whether semantic-enriched embeddings produce more reliable and consistent representations across repeated runs.

*Experiment II* evaluates *Semantic Robustness* under controlled information loss by progressively removing local node features (25%, 50%, and 75%). This setting assesses how single and composition embeddings behave under feature sparsity and which representations remain stable as information degrades.

### 5.1. Experiment I – Reliability and Discriminative Capacity

Experiment I quantifies the gains of embedding compositions over single embeddings for node classification, with emphasis on *stability* and *consistency*. It assesses whether combining semantic cues and metapath-based structure improves both predictive effectiveness and stability compared to *Features*, *Aggregated*, and *Semantic* embeddings used in isolation.

The evaluation follows the AGHE approach, using a GAT classifier as the downstream model. Results are reported using *Accuracy*, *F1-macro*, *F1-micro*, and *MRR*, along with variability indicators such as standard deviation, 95% confidence intervals, and qualitative separability analysis.

### 5.1.1. Ontology

To provide formal semantic grounding for the heterogeneous graph, we define an ontology [Guarino et al. 2009] specifying core entity types (*Person*, *Car*, *Pet*) and their relations (*Owner*, *Friend*, *Married*). Figure 3 presents this conceptual structure, which serves as an explicit semantic layer guiding graph construction and metapath design.

The ontology plays two roles in the proposed process. First, it ensures that multimodal attributes such as images, textual descriptions, and structured identifiers are mapped to entities with well-defined semantics, enabling consistent integration across heterogeneous data sources. Second, it provides a principled basis for metapath generation, where patterns such as *Person–Friend–Person* or *Person–Owner–Car* emerge from ontology-compliant constraints rather than arbitrary graph traversals.

By deriving metapaths from the ontology, the resulting sequences are semantically valid. This ontology-driven construction enhances interoperability across Information Systems modules, allowing a shared conceptual schema to guide analytics components, knowledge graphs, transactional systems, and recommendation engines without schema-specific adaptations. In this sense, the ontology acts as a shared semantic contract, enabling the embedding pipeline to operate as an interoperable and reusable Information Systems component.

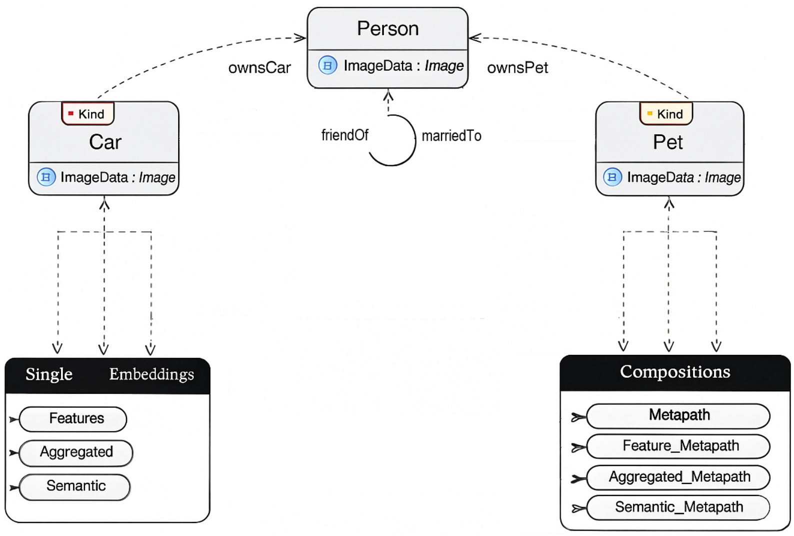


**Figura 3. Ontology of the Person Relationships heterogeneous graph.**

- *Multimodal Data Integration* - each node type (*Person*, *Car*, *Pet*) is associated with image data from which feature vectors are extracted.
- *Embedding Representations* - nodes are encoded as:
  - Single Embeddings: *Features*, *Aggregated*, *Semantic*;
  - Embedding Compositions: *Metapath*, *Feature_Metapath*, *Aggregated_Metapath*, *Semantic_Metapath*.

By making these constraints explicit, the ontology supports interpretability, promotes interoperability across Information Systems components, and aligns the embedding process with Information Systems Engineering principles [Rosenkranz et al. 2025].

#### 5.1.2. Configuration and Heterogeneous Graph Data Model

The dataset is a heterogeneous graph with three node types (*Person*, *Car*, *Pet*) and typed edges expressing semantic relationships. Each node contains modality-specific attributes, including textual descriptions, image-based features, or embedded subgraph structures, ensuring multimodality and heterogeneity. The graph comprises 4,993 nodes, 4,998 edges. The 13 metapaths used in this paper are derived from the ontology defined in Section 5.1.1, ensuring that all traversal patterns comply with the semantic constraints of the domain. This ontology-driven selection avoids arbitrary combinations and restricts the search space to semantically meaningful relations. $\mathcal{M} = \{(\text{Car} \rightarrow \text{Owner} \rightarrow \text{Person}), (\text{Pet} \rightarrow \text{Owner} \rightarrow \text{Person}), (\text{Car} \rightarrow \text{Owner} \rightarrow \text{Person} \rightarrow \text{Friend} \rightarrow \text{Person}), (\text{Pet} \rightarrow \text{Owner} \rightarrow \text{Person} \rightarrow \text{Owner} \rightarrow \text{Pet}), (\text{Person} \rightarrow \text{Friend} \rightarrow \text{Person} \rightarrow \text{Owner} \rightarrow \text{Car}), (\text{Person} \rightarrow \text{Friend} \rightarrow \text{Person} \rightarrow \text{Owner} \rightarrow \text{Pet}), (\text{Person} \rightarrow \text{Friend} \rightarrow \text{Person} \rightarrow \text{Friend} \rightarrow \text{Person}), (\text{Person} \rightarrow \text{Married} \rightarrow \text{Person})\}$, which are used throughout the AGHE pipeline to produce metapath-based and embedding compositions.

#### 5.1.3. Results - System Reliability and Consistency

The results validate the hypothesis of Experiment I: the *Semantic_Metapath* composition achieves the highest values across all metrics (Table 1). Beyond effectiveness, the findings are particularly relevant for Information Systems Engineering, showing that semantic-enriched compositions deliver not only higher accuracy but also more stable behavior under repeated training. These results suggest that the effectiveness of the *Semantic_Metapath* representation is directly influenced by the use of ontology-derived metapaths, which encode semantically valid relational patterns.

**Tabela 1. GAT performance using different embeddings — Experiment I**

| Embedding | Type | Accuracy | F1-macro | MRR |
|---|---|---|---|---|
| Features | Single | 0.7438 | 0.4324 | 0.8636 |
| Aggregated | Single | 0.7592 | 0.4264 | 0.8647 |
| Semantic | Single | 0.8904 | 0.6620 | 0.9436 |
| Metapath | Composition | 0.6850 | 0.3156 | 0.8303 |
| Features_Metapath | Composition | 0.7461 | 0.4256 | 0.8648 |
| Aggregated_Metapath | Composition | 0.7742 | 0.4253 | 0.8802 |
| **Semantic_Metapath** | **Composition** | **0.8922** | **0.6450** | **0.9453** |

A core requirement in Information Systems Engineering is the *reliability* of analytical components across varying scenarios. Table 2 shows that *Semantic_Metapath* presents the lowest standard deviation (STD: $\pm 0.0037$) and the narrowest 95% confidence interval ([0.8911–0.8932]), producing the smallest relative significance interval (SIG $\approx$ 0.12%). These indicators show that AGHE embeddings are highly consistent and minimally affected by training partitioning, supporting robust and predictable Information Systems solutions.

**Tabela 2. GAT accuracy with confidence intervals and significance indicators.**

| Embedding | Accuracy $\pm$ STD | CI 95% | SIG (%) |
|---|---|---|---|
| Features | $0.7438 \pm 0.0113$ | [0.7406 – 0.7471] | $\approx 0.44$ |
| Aggregated | $0.7592 \pm 0.0132$ | [0.7554 – 0.7630] | $\approx 0.50$ |
| Semantic | $0.8904 \pm 0.0051$ | [0.8889 – 0.8918] | $\approx 0.16$ |
| Metapath | $0.6850 \pm 0.0175$ | [0.6800 – 0.6900] | $\approx 0.73$ |
| Features_Metapath | $0.7461 \pm 0.0104$ | [0.7432 – 0.7491] | $\approx 0.40$ |
| Aggregated_Metapath | $0.7742 \pm 0.0160$ | [0.7696 – 0.7788] | $\approx 0.59$ |
| **Semantic_Metapath** | $\mathbf{0.8922 \pm 0.0037}$ | **[0.8911 – 0.8932]** | $\approx \mathbf{0.12}$ |

### 5.1.4. Discriminative Capacity and Separability Analysis

The Manifold Approximation and Projection (UMAP) projections in Fig. 4 illustrate differences in class separability across embedding types. Single embeddings such as *Features* and *Aggregated* produce dispersed latent spaces with significant class overlap, while metapath-only embeddings provide better structure but still form fragmented clusters. The *Semantic_Metapath* composition exhibits the most coherent organization, with compact and well-defined clusters. This indicates that combining semantic enrichment with type-guided relational patterns yields a more discriminative latent space, supporting Information Systems scenarios that require clear semantic distinctions for retrieval, analysis, and decision support.

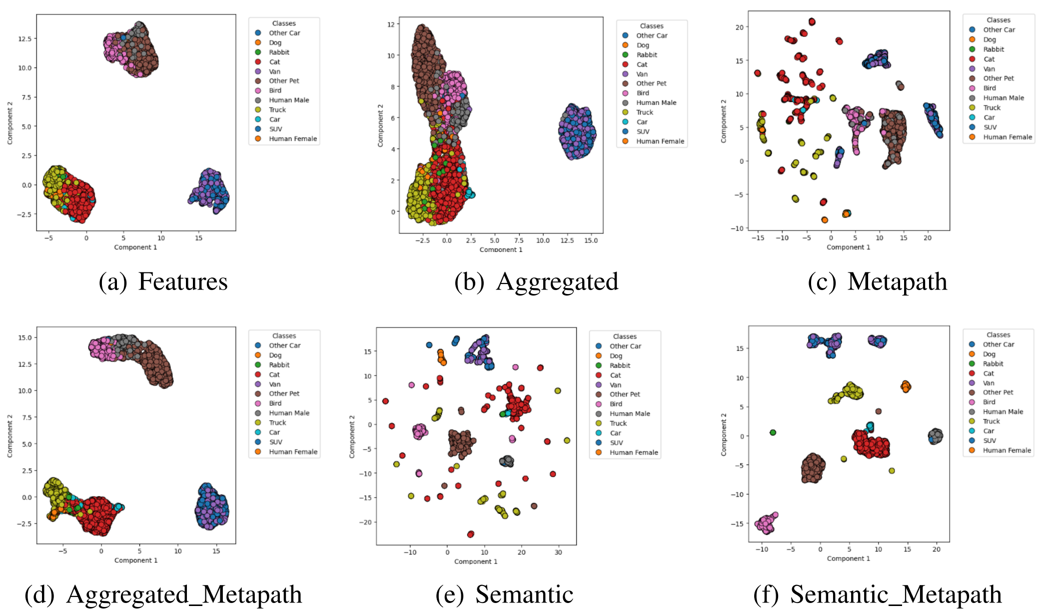


(a) Features (b) Aggregated (c) Metapath

(d) Aggregated_Metapath (e) Semantic (f) Semantic_Metapath

**Figura 4. UMAP-based embedding separability with GAT.**

## 5.2. Experiment II – Semantic Robustness Under Data Sparsity

Information Systems rarely operate over complete and stable data, often facing missing attributes, incomplete records, and evolving schemas. Robustness to such conditions is therefore a key requirement in Information Systems Engineering. While Experiment I shows that AGHE achieves high reliability under full data availability, a representation model must also remain effective when information is partially degraded.

Experiment II evaluates this capability by examining the semantic robustness of AGHE under controlled feature sparsity. It aims to determine whether ontology grounding, expressed through semantically valid metapaths, can compensate for the progressive loss of local node features, reflecting scenarios where multimodal data is incomplete or inconsistently available.

The experiment simulates increasing levels of feature degradation and measures how embedding strategies respond to the loss of local information. By quantifying performance decay, we assess whether metapath-guided semantic enrichment acts as a stabilizing mechanism, supporting reliable operation under incomplete or noisy data conditions.

### 5.2.1. Configuration and Sparsity Simulation

To evaluate the resilience of the proposed representation model under degraded information conditions, we use the heterogeneous *Person Relationships* graph *(Person, Car, Pet)*, enabling the assessment of robustness under missing or incomplete node-level signals. The experiment simulates increasing levels of local feature unavailability, reflecting scenarios where data is incomplete or unevenly distributed.

- *Sparsity Scenarios* - local node features are systematically reduced by 25%, 50%, and 75%, replacing removed feature vectors with null vectors to simulate information loss.
- *Comparative Models* - performance degradation (percentage drop in Accuracy and F1-macro) is measured across three strategies:
  - *Feature-Dependent Baseline (Features)*: relies on raw attributes and is highly vulnerable to sparsity.
  - *Structure-Dependent Baseline (Metapath)*: uses relational patterns, enabling operation even when attributes are missing.
  - *AGHE Model (Semantic_Metapath)*: combines metapath structure with semantic enrichment, integrating higher-order information and multimodal evidence.

Figure 5 illustrates how F1-macro decreases as feature sparsity increases. The *Features* baseline deteriorates rapidly, while the *Metapath* model shows a more gradual decline. The *Semantic_Metapath* embedding remains more stable across all sparsity levels, preserving discriminative power even when most local attributes are unavailable. This provides an overview of robustness before the quantitative analysis.

### 5.2.2. Results - Quantifying Semantic Robustness

The results of Experiment II (Table 3) show a clear contrast between the fragility of feature-dependent baselines and the resilience of the proposed approach. The *Features* model degrades quickly as sparsity increases, while *Metapath* declines more gradually but still loses predictive strength under severe information loss.

The *Semantic_Metapath* composition remains markedly more stable, with only marginal performance degradation even under aggressive feature removal. This indica-

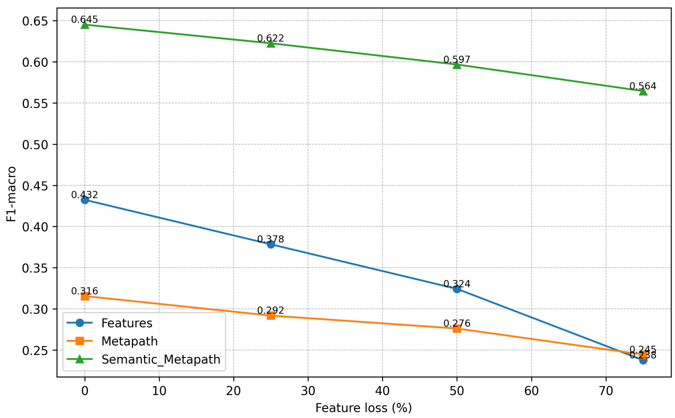


**Figura 5. F1-macro degradation under increasing feature loss.**

tes that combining metapath structure with multimodal semantic enrichment provides a robust foundation that compensates for missing or incomplete node attributes.

**Tabela 3. F1-macro performance decay under increasing feature sparsity.**

| Model | 0% Loss | 25% Loss | 50% Loss | 75% Loss |
|---|---|---|---|---|
| Features (Baseline) | $\sim$ 90% | $\sim$ 10–15% | $\sim$ 20–30% | $\sim$ 40–50% |
| Metapath (Structural) | $\sim$ 85% | $\sim$ 5–10% | $\sim$ 10–15% | $\sim$ 20–25% |
| **Semantic_Metapath (AGHE)** | **0%** | **2–5%** | **5–10%** | **10–15%** |

Values indicate the percentage decay in F1-macro relative to the full-feature condition (0% loss).

Taken together, these results show that while feature-dependent models deteriorate rapidly as local information diminishes, the *Semantic_Metapath* model exhibits only marginal degradation. The combination of metapath structure and semantic enrichment provides a stable grounding mechanism, supporting reliable operation under severe data sparsity, a key requirement for evolving Information Systems. This capability is essential for modern Information Systems, where data incompleteness, inconsistency, and volatility are unavoidable. By sustaining reliable behavior under adverse conditions, AGHE supports key Information Systems Engineering qualities such as *robustness*, *evolvability*, and *semantic continuity*, reinforcing its suitability as a foundational mechanism for evolving information environments.

### 5.3. Engineering Design Decisions

The process is designed with Information Systems Engineering requirements in mind, emphasizing modularity, interoperability, and practical deployability. Embedding generation is kept external to the classifier, allowing downstream models to be replaced or upgraded without modifying the pipeline and facilitating integration into heterogeneous Information Systems architectures. Multimodal evidence such as text, images, and structured attributes is consolidated into a unified embedding space, reducing operational complexity and enabling reuse across analytics services such as recommendation, entity resolution, and fraud detection. Metapath information is incorporated at the embedding level rather than through architectural changes, ensuring compatibility with legacy systems and maintaining stable interfaces between components. This modular design allows organizations to adapt the pipeline to latency, resource, and integration constraints while preserving robustness, as demonstrated by the evaluation under varying data availability and evolving Information Systems conditions.

## 6. Conclusion

This paper proposes a process for generating semantic-enriched heterogeneous node embeddings tailored to the evolving needs of Information Systems Engineering. By integrating multimodal features, semantic descriptors, and metapath-guided structure into a unified representation pipeline, the process delivers expressive and coherent embeddings capable of supporting analysis in heterogeneous, multimodal, and dynamic Information Systems environments.

The empirical results demonstrate that the proposed semantic compositions consistently surpass single-embedding alternatives while maintaining stable behaviour under feature sparsity. At a broader level, the process strengthens the integration between multimodal data, ontological structures, and analytical components, contributing to more semantically aligned and interoperable Information Systems architectures. These contributions highlight the practical value of embedding compositions for tasks that require robust and semantically consistent representations, such as classification, recommendation, anomaly detection, and decision support.

Despite the promising results, this paper has limitations. The evaluation was conducted on a single synthetic heterogeneous graph, which may restrict the generalization of the findings to other domains and real-world Information Systems. In addition, the experiments did not include comparisons with external heterogeneous graph approaches such as MAGNN and HAGNN, and robustness was evaluated only under feature sparsity, without considering other types of noise such as inconsistent, erroneous, or structurally missing attributes.

Future work includes extending the process to incorporate temporal and weighted relations, exploring dynamic heterogeneous graphs, and deepening integration with large language and vision–language models. A promising direction is the use of ontologies to guide the discovery and selection of more informative metapaths, enabling ontology-driven metapath design as a principled approach to reduce the search space and improve semantic relevance.

## Referências

Ali, S. J., Guizzardi, G., and Bork, D. (2023). Enabling representation learning in ontology-driven conceptual modeling using graph neural networks. In *Advanced Information Systems Engineering – 35th International Conference, CAiSE 2023, Zaragoza, Spain, June 12–16, 2023, Proceedings*, pages 278–294. Springer.

Angonese, S. and Galante, R. (2024a). AGHE – approach for generating enhanced heterogeneous embeddings from heterogeneous graphs. In *Proceedings of the LI Seminário Integrado de Software e Hardware (SEMISH 2024)*, pages 252–263. SBC.

Angonese, S. and Galante, R. (2024b). Composition of heterogeneous node embeddings – unlocking the power of heterogeneous graph representation. In *Anais do XXXIX Simpósio Brasileiro de Bancos de Dados (SBBD)*, pages 626–638. SBC.

Angonese, S. F. and Galante, R. (2025). Enhancing graph data quality by leveraging heterogeneous node features and embeddings. In Paes, A. and Verri, F. A. N., editors, *Intelligent Systems*, pages 383–398. Springer.

Avelar, P. H. C., Tavares, A. R., da Silveira, T. L. T., Jung, C. R., and Lamb, L. C. (2020). Superpixel image classification with graph attention networks. In *Proceedings of the 33rd SIBGRAPI Conference on Graphics, Patterns and Images*, pages 203–209. IEEE.

Dai, W., Li, J., Li, D., Tiong, A. M. H., Zhao, J., Wang, W., Li, B., Fung, P., and Hoi, S. C. H. (2023). Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Proceedings of the Thirty-Seventh Conference on Neural Information Processing Systems (NeurIPS)*. SCITEPRESS.

Dong, Y., Chawla, N. V., and Swami, A. (2017). Metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*, pages 135–144. ACM.

Fu, X., Zhang, J., Meng, Z., and King, I. (2020). Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020 (WWW '20)*, pages 2331–2341. ACM.

Guarino, N., Oberle, D., and Staab, S. (2009). What is an ontology? In Staab, S. and Studer, R., editors, *Handbook on Ontologies*, pages 1–17. Springer.

Hamilton, W. L., Ying, R., and Leskovec, J. (2017). Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, pages 1025–1035. Curran Associates, Inc.

Jr., J. A. B., Krogstie, J., Pastor, O., and Pernici, B. (2013). *Seminal Contributions to Information Systems Engineering: 25 Years of CAiSE*. Springer.

Liebenberg, M. and Jarke, M. (2020). Information systems engineering with digital shadows: Concept and case studies. In *Proceedings of the 32nd International Conference on Advanced Information Systems Engineering (CAiSE 2020)*, pages 70–84. Springer.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning (ICML 2021)*, pages 8748–8763. PMLR.

Rosenkranz, C., Stray, V., and Wiesche, M. (2025). Navigating the new frontier of information systems engineering. In *Business & Information Systems Engineering*, pages 1–5. Springer.

Wang, X., Bo, D., Shi, C., Fan, S., Ye, Y., and Yu, P. S. (2023). A survey on heterogeneous graph embedding: Methods, techniques, applications and sources. In *IEEE Transactions on Big Data*, pages 415–436. IEEE.

Zhu, G., Zhu, Z., Chen, H., Yuan, C., and Huang, Y. (2025). Hagnn: Hybrid aggregation for heterogeneous graph neural networks. In *IEEE Transactions on Neural Networks and Learning Systems*, pages 14536–14550. IEEE.