I.J. Modern Education and Computer Science, **2025, 6, 48-64**
Published Online on December 8, 2025 by MECS Press (http://www.mecs-press.org/)
DOI: 10.5815/ijmecs.2025.06.04

NeSy-Guidance: A Neuro-Symbolic Knowledge Graph for Academic Recommendations Combining Rule-Based Reasoning and Neural Inference

Zineb Elkaimbillah*
IMS Team, ADMIR Laboratory Rabat IT Center, ENSIAS, Mohammed V University in Rabat; Rabat, Morocco
Email: zineb_elkaimbillah@um5.ac.ma
ORCID iD: https://orcid.org/0009-0008-4061-7714
*Corresponding Author

Zineb Mcharfi
IMS Team, ADMIR Laboratory Rabat IT Center, ENSIAS, Mohammed V University in Rabat; Rabat, Morocco
Email: zineb_mcharfi@um5.ac.ma
ORCID iD: https://orcid.org/0009-0007-6242-6006

Mohamed Khoual
IMS Team, ADMIR Laboratory Rabat IT Center, ENSIAS, Mohammed V University in Rabat; Rabat, Morocco
Email: mohamed_khoual@um5.ac.ma
ORCID iD: https://orcid.org/0009-0005-2499-573X

Bouchra El Asri
IMS Team, ADMIR Laboratory Rabat IT Center, ENSIAS, Mohammed V University in Rabat; Rabat, Morocco
Email: b.elasri@um5s.net.ma
ORCID iD: https://orcid.org/0000-0001-8502-4623

Received: 12 February, 2025; Revised: 17 March, 2025; Accepted: 26 July, 2025; Published: 08 December, 2025

Abstract: The growing complexity of academic and career decision-making requires intelligent systems that can deliver personalized recommendations while ensuring strict compliance with institutional policies and incorporating evolving contextual factors. This paper introduces NeSy-Guidance, a neuro-symbolic recommendation approach that combines symbolic rule reasoning with graph-based neural inference over a dynamic and time-aware academic Knowledge Graph (KG). The model encodes students, programs, and contextual entities including regions and emerging trends while integrating regulatory constraints derived from Moroccan admission policies. It applies a two-stage reasoning pipeline: a symbolic layer enforcing hard eligibility rules and soft preference-based adjustments mined automatically from the knowledge graph, and a Graph Convolutional Network (GCN) layer trained with a weighted loss to address class imbalance and capture latent student–program compatibility. A weighted score fusion mechanism integrates both inference outputs, achieving a balance between interpretability, adaptability, and predictive performance. Evaluated on a real-world dataset of 800 students and 325 academic programs, NeSy-Guidance outperforms three state-of-the-art baselines in both accuracy and policy compliance. It achieves 83.8% accuracy, 74.5% precision@5, 75.4% F1-score, and ensures 100% compliance with institutional eligibility rules. Furthermore, a qualitative survey confirmed positive student satisfaction regarding the clarity and relevance of recommendations. These results demonstrate the effectiveness of hybrid reasoning and validate NeSy-Guidance as a reliable, explainable, and regulation-aware academic guidance system capable of adapting to regional disparities, emerging academic trends, and evolving student preferences.

Index Terms: Neuro-Symbolic Approach, Knowledge Graph, Graph Convolutional Network (GCN), Symbolic Rule Reasoning, Policy-alignment, Academic Recommendation

1. Introduction

The rapid evolution of educational systems [1], combined with the increasing diversity of postsecondary academic pathways, has significantly complicated academic and career decision-making [2]. Today's students must navigate a growing number of programs while balancing institutional admission constraints, personal interests, regional disparities, and emerging academic trends. In this context, intelligent recommendation systems have emerged as essential tools to support learners in making informed, relevant, and policy-aligned choices [3].

However, conventional approaches such as collaborative filtering [4] or content-based recommendation [5] often fall short. They struggle to handle heterogeneous and dynamic data, offer limited interpretability, and fail to comply with formal admission policies. While deep learning models especially Graph Neural Networks (GNNs) like GCN [6] excel at capturing latent patterns in graph-structured data, their black-box nature and lack of symbolic reasoning capabilities make them difficult to apply in domains that require transparency and strict compliance, such as academic guidance [7].

To address these limitations, we introduce NeSy-Guidance, a neuro-symbolic hybrid recommendation framework designed for educational orientation. Our approach leverages a dynamic and time-aware knowledge graph [8], which encodes students, academic programs, regulatory constraints, contextual entities (e.g., regions, skills, trends), and evolving learner preferences. The model integrates two complementary reasoning layers: (1) a symbolic inference module that enforces hard eligibility rules and soft, automatically mined preference-based rules weighted according to statistical reliability; and (2) a graph-based neural inference layer using a GCN trained with a weighted loss function to address the imbalance between positive and negative samples. Final recommendation scores are computed through a weighted fusion of symbolic and neural outputs, striking a balance between predictive performance, transparency, and policy alignment.

The main contributions of this work are as follows:

- We propose a hybrid neuro-symbolic recommendation model that integrates symbolic reasoning via domain-specific rules (hard and soft) with GCN-based inference over a dynamic knowledge graph.
- We construct a dynamic and time-aware knowledge graph simulating the Moroccan academic landscape, encoding admission constraints, student preferences, regional disparities, and emerging academic trends.
- We introduce an automatic soft rule mining and weighting strategy based on support, confidence, and coverage metrics.
- We propose a score fusion mechanism balancing symbolic interpretability and neural predictive power.
- We conduct an extensive experimental evaluation, including a comparative analysis against three state-of-the-art models (KGAT, RuleRec, NER-SemSim), and a qualitative assessment of student satisfaction with the recommendations.

The remainder of this paper is structured as follows: Section 2 reviews related work, focusing on knowledge graphs and neuro-symbolic recommender systems. Section 3 elaborates on the proposed methodology, covering knowledge graph modeling, rule extraction, GCN architecture, and the fusion strategy. Section 4 presents the experimental results, comparative analysis, and qualitative evaluation. Finally, Section 5 concludes the study and discusses potential future research directions.

2. Related Works

2.1. Knowledge Graph-Based Recommender Systems

Recent advances in knowledge graph-based recommender systems (KGRS) have highlighted the effectiveness of structured semantic representations in improving both the accuracy and explainability of recommendations. By organizing users, items, and attributes as interconnected entities, KGRS models can support multi-hop reasoning and capture rich contextual information capabilities often lacking in traditional collaborative filtering methods.

A seminal contribution in this field is the Knowledge Graph Attention Network (KGAT) [9], which introduced an attention mechanism to enhance relation-aware information propagation across the graph. While KGAT significantly improves recommendation accuracy through embedding-based interaction modeling, it does not incorporate rule-based reasoning or support compliance with domain-specific constraints.

In response to the demand for interpretability,[10] proposed a joint learning framework that combines neural inference with explainable symbolic rules extracted from the graph. This approach demonstrates that symbolic reasoning can complement neural models, enhancing transparency and generalization. However, it lacks context-aware adaptability and fails to account for evolving domain knowledge both critical requirements in educational environments.

In the academic domain, [11] proposed a GNN-based recommendation system for heterogeneous scholarly networks, capturing interactions among researchers, institutions, and publications to generate career path suggestions.

While this approach effectively models research collaborations, it omits two key aspects of formal academic guidance: curriculum-aware reasoning and policy compliance both essential for structured educational recommendations.

Recent work by [12] leverages large language models (LLMs) for knowledge graph completion in curriculum and domain modeling. Their approach automates the extraction of latent semantic relationships from unstructured educational materials, thereby enhancing the graph's structural richness. Nevertheless, this methodology remains limited by its disregard for symbolic constraints and its absence of integrated, inference-driven recommendation mechanisms.

In summary, while current knowledge graph-based recommendation systems (KGRS) effectively model relationships and improve personalization, most fail to explicitly reconcile the dual demands of regulatory compliance and dynamic adaptability required in educational settings. This critical gap motivates our proposed NeSy-Guidance framework, which uniquely combines symbolic rule enforcement with graph-based learning to deliver transparent, personalized, and policy-compliant academic recommendations.

2.2. Neuro-symbolic AI for Recommendation

Recent advances in neuro-symbolic artificial intelligence have led to the emergence of hybrid recommender systems that aim to combine the generalization capabilities of neural models with the interpretability and rule compliance of symbolic reasoning. This integration addresses key limitations of traditional deep learning models, particularly in domains requiring transparency and policy alignment, such as academic guidance.

A recent systematic review by [13] highlights the rapid development of neuro-symbolic approaches, with most efforts focusing on learning and inference mechanisms, while challenges related to explainability and trustworthiness remain largely unresolved. These limitations are especially critical in educational settings, where recommendations must be not only effective but also justifiable and policy-compliant.

[14] proposed a knowledge graph-based recommendation system that incorporates first-order logic rules into neural embeddings. Their results confirm that logic-guided representations improve both precision and novelty. However, their work is primarily focused on domains such as e-commerce, and lacks mechanisms to enforce institutional constraints specific to education.

Other efforts, such as those by [15], explore cross-domain symbolic reasoning using Logic Tensor Networks (LTNs), a neuro-symbolic framework that integrates deep learning with logical reasoning from structured knowledge. While powerful and theoretically elegant, these architectures often require significant computational resources and lack explicit mechanisms for handling domain-specific policy constraints, making them challenging to apply effectively in educational recommendation scenarios where strict compliance and interpretability are essential.

In a similar vein, [16] introduced a hybrid model that integrates graph neural networks and propositional logic for recommendation. Although the system achieves strong predictive performance, it does not incorporate domain-specific constraints, such as admission rules or eligibility thresholds, which are crucial in academic orientation.

[17] explored multi-stage symbolic reasoning over knowledge graphs to provide explainable recommendations. Their sequential model enhances transparency but remains rigid and challenging to adapt in dynamic educational environments where rules and student data evolve rapidly.

Finally, [18] proposed the PSYCHIC framework, which bridges symbolic entities and neural embeddings for question-answering tasks over knowledge graphs. Although insightful, this work does not focus on structured recommendation or score fusion strategies, which are central to the decision-making process in academic systems.

2.3. Positioning of Our Work

Compared to these approaches, NeSy-Guidance introduces the following key innovations:

- A dual-layer symbolic reasoning module that combines hard rules validated by academic experts to enforce institutional policies and soft rules, automatically mined from a filtered knowledge graph, capturing student preferences and behavioral patterns.
- A GCN trained exclusively on the policy-compliant graph, ensuring strict adherence to eligibility criteria prior to neural inference.
- A weighted score fusion mechanism that balances symbolic reasoning with neural generalization, offering tunable control over the trade-off between accuracy and rule compliance.

To the best of our knowledge, this work is the first to integrate policy-aligned hard constraints, automatically weighted soft rules, regional and trend-aware contextual modeling, and GCN inference within a unified, explainable framework tailored specifically for academic recommendation. This contribution fills a critical gap in the literature by delivering a neuro-symbolic recommendation architecture that is both adaptable to evolving institutional policies and capable of providing personalized, trustworthy guidance in real-world educational settings.

3. Methodology

3.1. Neuro-Symbolic Architecture

NeSy-Guidance is a hybrid neuro-symbolic approach designed to provide personalized and policy-compliant academic recommendations. The system is structured as a modular, multi-step pipeline that combines symbolic rule reasoning with graph-based neural inference over a dynamic and time-aware academic knowledge graph.

The process begins with data collection and preprocessing, followed by the construction of the initial knowledge graph encoding students, academic programs, and contextual entities such as regions, career goals, and emerging trends. A dedicated update and evolution layer then monitors structured data sources to incrementally enrich the graph with new observations and time-stamped relations.

A hard rule filtering module is subsequently applied to enforce institutional eligibility criteria, resulting in a filtered graph containing only admissible student–program pairs. From this dynamically maintained graph, two parallel inference pathways are activated:

- A symbolic rule layer evaluates soft rules that reflect student interests, aspirations, and contextual coherence. These soft rules are automatically mined and weighted based on statistical measures of confidence and coverage. The resulting symbolic scores are normalized to ensure comparability across candidates.
- In parallel, a GCN is trained on the same filtered graph using encoded node features and a weighted binary cross-entropy loss to handle class imbalance and capture latent compatibility patterns.

The outputs from both reasoning components symbolic and neural are subsequently combined through a score fusion mechanism governed by a tunable parameter α. This integration balances predictive performance, interpretability, and compliance, ultimately producing a Top-K ranked list of academic programs that remain aligned with evolving learner profiles and institutional constraints. The complete architecture and workflow of the NeSy-Guidance system are illustrated in Figure 1, which highlights the modular components, dual-layer reasoning, and update capabilities.

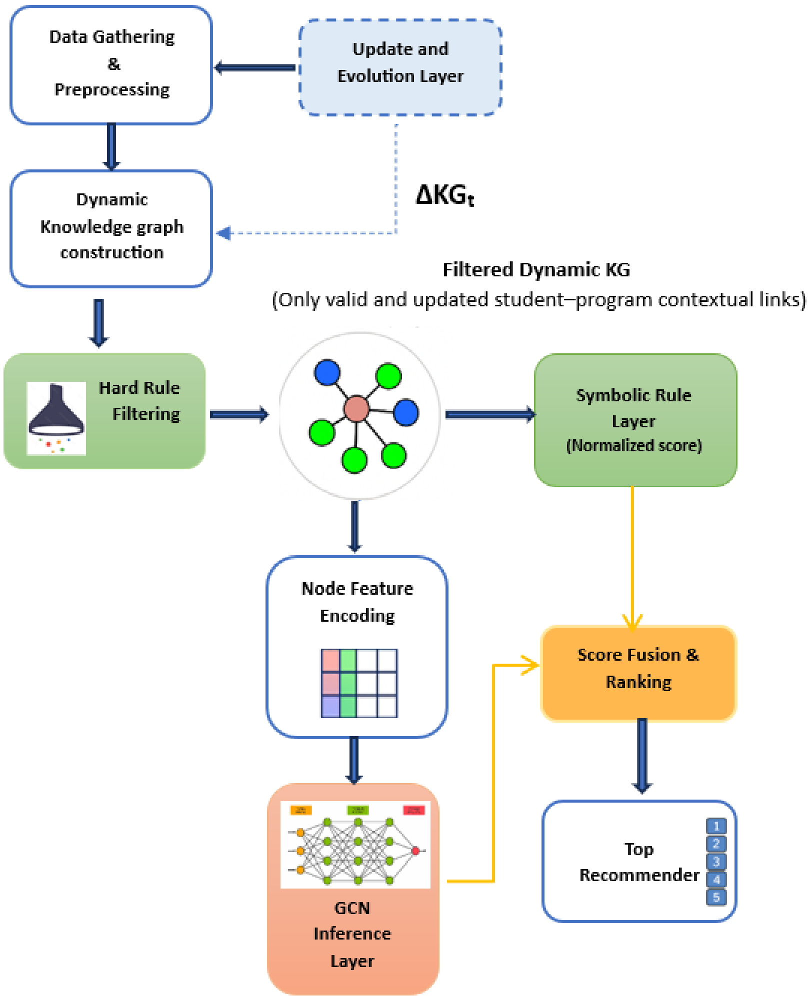

Fig. 1. Modular Workflow of the NeSy-Guidance System

3.2. Knowledge Graph Modeling for Academic Recommendation

We model the academic recommendation problem through a dynamic and time-aware knowledge graph (KG) that encodes students, academic programs, and relevant contextual entities. This flexible representation enables both symbolic reasoning over eligibility constraints and neural inference over evolving relational structures. In contrast to traditional static graphs, our KG supports real-time updates, preference evolution, and regional or emergent academic dynamics.

Let $S = \{s_1, s_2, ..., s_n\}$ denote the set of students, $D = \{d_1, d_2, ..., d_m\}$ the set of academic programs, and $C = \{c_1, ..., c_k\}$ the set of contextual entities such as skills, careers, modules, universities, or regions. The complete set of entities is defined as:

$$E = S \cup D \cup C \tag{1}$$

Each relation in the graph is defined as a time-aware quadruple, capturing the semantic and temporal dimensions of the interaction:

$$KG_t = \{ (e_h, r, e_t) \mid e_h, e_t \in E, r \in R, t \in T \} \tag{2}$$

Where, e_h denotes the head entity (source node of the relation), e_t denotes the tail entity (target node of the relation). r denotes the relation type, drawn from the set R, and t denotes the timestamp, drawn from the set of temporal indices T.

We distinguish between two main categories of relations:

- R_e: Eligibility relations, encoding hard institutional constraints such as admissibility rules (e.g., (s, eligible_for, d), (d, requires_bac_type, 'SM',t), (d, has_cutoff, 14.5, t)).
- R_p: Preference relations, capturing soft, temporal, and contextual indicators such as declared interests, career goals, or regional preferences (e.g., (s, interested_in, 'Engineering', t), (s, wants_career, 'AI', t)). Hence, the complete relation set is:

$$R = R_e \cup R_p \tag{3}$$

This modeling approach allows the knowledge graph to be continuously enriched over time with new facts and interactions. As new student information, program updates, or contextual data become available, additional timestamped relations are integrated incrementally. This ensures that the system remains responsive to changes in both academic structures and learner behavior, providing a robust foundation for hybrid neuro-symbolic inference. Figure 2 illustrates a simplified view of the constructed dynamic academic knowledge graph. Nodes represent student profiles (red), academic programs (yellow), and contextual entities such as modules, regions, and trends (blue). Edges encode eligibility constraints as well as time-aware preference relations.

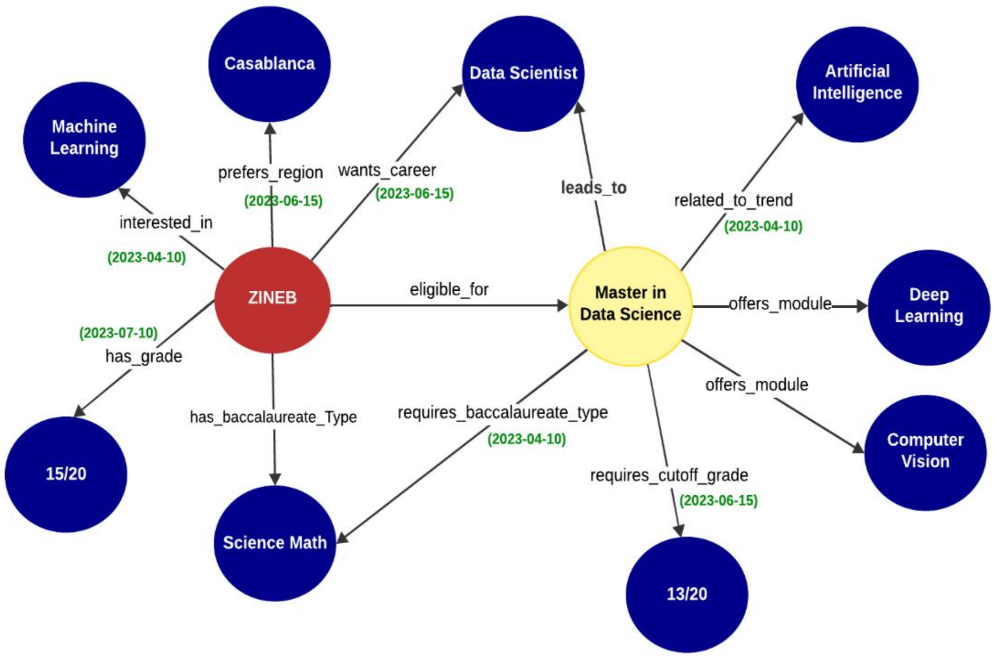

Fig. 2. Visual Representation of the Dynamic Time-Aware Knowledge Graph

3.3. KG Update and evolution layer

To ensure that the knowledge graph remains both temporally consistent and responsive to the evolving academic landscape, we incorporate a dedicated KG Update & Evolution Layer into the model. This layer provides a principled mechanism to integrate new observations and changes without requiring a complete reconstruction of the graph. Formally, the update process is defined by the incremental enrichment rule:

$$KG_{\{t+\Delta t\}} = KG_t \cup \Delta KG_t \tag{4}$$

where ΔKG_t denotes the set of newly introduced or modified nodes, edges, or attributes at time $t + \Delta t$. The update mechanism is implemented through three main components:

1) Change Detection: The system continuously monitors structured data sources including batch CSV files and user interface logs to identify any new or modified records. For each entity, the update layer performs an efficient diff against previous snapshots to extract only the delta (ΔKG_t). This reduces processing overhead and avoids redundant insertions.
2) Incremental Graph Enrichment: For each detected update, the system applies incremental modifications to the graph. These updates are expressed in Cypher MERGE operations, ensuring that entities are either created or updated in place:
   - Student Profile Updates: if a student's aspiration, grade, or region changes, the corresponding node properties or relations are updated with a new timestamp.
   - Program Attribute Updates: changes in admission cutoffs or eligibility constraints are reflected as updated properties or new relations.
   - New Entities: when new programs, skills, or trends are introduced, the system creates the relevant nodes and semantic links.

Each update operation adds a timestamped version of the fact, preserving temporal consistency and enabling reasoning over historical data.

3) Temporal Alignment and Rule Activation: Once the graph is enriched with new elements, the update layer triggers temporal alignment processes. In particular, any updates to preference-related relations such as interests or aspirations activate soft rule recalculation for the affected students. This dynamic weighting ensures that recent behaviors are prioritized over outdated patterns in the recommendation process.

3.4. Symbolic reasoning layer

To ensure that recommendations are both policy-compliant and semantically coherent, NeSy-Guidance integrates a symbolic reasoning layer that operates on domain-specific logical rules. These rules encode institutional constraints as well as educational preferences, and are applied at different stages of the inference pipeline to filter and adjust candidate program recommendations. We distinguish between two complementary categories of rules: hard constraints and soft preference-based rules. This section focuses on the former.

3.4.1. Hard Constraints (Eligibility Filtering)

This module applies strict eligibility criteria derived from national academic regulations and expert recommendations in Morocco. These hard rules (HR) ensure that only students who meet fundamental admission requirements are considered for recommendation.

The rules are encoded as symbolic logical expressions and executed as pre-filtering constraints on the knowledge graph. Each rule checks a specific eligibility aspect (grade threshold, bac compatibility, mention, etc.), thereby guaranteeing semantic validity before any learning or prediction is performed. Table 1 summarizes example of the hard rules implemented in our system.

Table 1. Examples of Hard Eligibility Rules

ID	**Description**	**Logical form**
HR1	Student’s grade must meet or exceed the diploma cutoff	Note(s) ≥ cutoff(d)
HR2	Student’s bac type must be accepted by the diploma	bac_type(s) ∈ accepted_bac_types(d)
HR3	Student’s mention must meet the required level	mention(s) ≥ required_mention(d)
HR4	Desired level must not exceed diploma level	desired_level(s) ≤ level(d)
HR5	Certain bac types are explicitly excluded	NOT (bac_type(s) ∈ excluded_bac_types(d))
HR6	Student must fulfill language prerequisites if applicable	lang_required(d) ⊆ student_languages(s)

3.4.2. Soft Rule Mining

Unlike hard rules that enforce strict eligibility based on institutional policies, soft rules introduce a flexible layer of symbolic reasoning to accommodate student preferences, contextual factors, and observed behavioral tendencies. These rules do not act as binary filters; rather, they adjust the final recommendation scores by reflecting latent patterns discovered in the knowledge graph.

Soft rules are automatically mined from the filtered knowledge graph [19], which has already undergone pre-processing through hard constraints. Inspired by the AMIE+ framework, our system extracts relational patterns that connect student attributes to academic programs through 2-hop semantic paths. The algorithm starts by exploring all outgoing relations from each student node and matches them to programs that share a common entity or attribute [20]. A rule candidate is generated whenever such a link is found, forming an implication of the form: If a student has attribute(s) *A*, then they are likely suited for diploma *D*.

To enhance transparency and reproducibility of the symbolic reasoning layer, we explicitly describe the process of soft rule generation using a graph traversal strategy. For each student node, the system inspects all outgoing links and identifies other nodes (e.g., programs) connected through shared intermediate entities. Figure 3 illustrates the pattern-based algorithm used to extract candidate soft rules from the filtered knowledge graph, systematically exploring two-hop semantic paths and generating conditional rule candidates that link contextual evidence to recommendation targets.

Algorithm 1 Pattern-Based Soft Rule Extraction

Require: Filtered knowledge graph $G = (V, E)$
Ensure: Set of rule candidates $\mathcal{R}$

$\mathcal{R} \leftarrow \emptyset$
for all student $s \in V$ where $type(s) = Student$ **do**
  for all (s, r_1, x) where $(s, r_1, x) \in E$ **do**
    for all (y, r_2, x) where $(y, r_2, x) \in E$ and $type(y) = Program$ **do**
      $Cond \leftarrow \{r_1(s, x), r_2(y, x)\}$
      $Cons \leftarrow suited\_for(s, y)$
      $\mathcal{R} \leftarrow \mathcal{R} \cup \{Cond \Rightarrow Cons\}$
    end for
  end for
end for
return $\mathcal{R}$

Fig. 3. Pattern-Based Algorithm for Soft Rule Extraction from the Knowledge Graph

This mechanism identifies pairs of related facts that can semantically imply a recommendation. For example, if a student is interested in "AI" and a program teaches "AI", the system infers the following candidate rule, IF student hasInterest = AI AND program teaches = AI THEN student is suited_for that program, such rule candidates are then evaluated statistically.

To evaluate the validity and utility of each candidate rule $A \Rightarrow D$, we compute three key statistical metrics:

1. Support: the number of instances where both the antecedent and the consequent are true

$$support(A \Rightarrow D) = |\{s \in G: A(s) \wedge D(s)\}| \tag{5}$$

2. Confidence: the conditional probability of the diploma given the student attribute(s):

$$confidence = P(D \mid A) = \frac{support(A \wedge D)}{support(A)} \tag{6}$$

3. Coverage (head coverage): the proportion of the conclusion instances that are explained by the rule:

$$coverage(D) = \frac{support(A \wedge D)}{support(D)} \tag{7}$$

Only rules that satisfy both of the following thresholds are retained:

$$confidence \geq minConf = 0{,}6 \ and \ coverage \geq minCov = 0{,}3 \tag{8}$$

To enhance consistency and scalability, we compute a symbolic weight automatically for each retained rule, reducing subjectivity and improving adaptability across datasets. The weight is calculated as a linear combination of confidence and coverage metrics, defined as:

$$weight = \alpha \times confidence + \beta \times coverage \tag{9}$$

With empirically set parameters (e.g., α = 0.7, β = 0.3), this scoring reflects both the statistical reliability (confidence) and representativeness (coverage) of the rule. The resulting weight is used as a bonus or penalty within the symbolic reasoning layer:

- Positive weight: encourages the recommendation (bonus)
- Negative weight: penalizes the recommendation (penalty)

This strategy ensures that symbolic reasoning is both data-driven and adaptive, reducing human bias and allowing for automated adjustment as new data enters the knowledge graph. Table 2 provides examples of soft rules (SR) included in the symbolic reasoning layer

Table 2. Representative Sample of Soft Rules

Rule ID	Description
SR1	Bonus if student interest aligns with program modules
SR2	Bonus if professional aspiration matches career outcome
SR3	Bonus if university is in the same region as the student (mobility constraint)
SR4	Penalty if program domain is too distant from student's bac specialty
SR5	Bonus for coherent academic goal (e.g., “Engineering” aspiration + “Maths” background)
SR6	Penalty if selected diploma leads to saturated or low-employability career path
SR7	Bonus if diploma aligns with national or regional job market trends

This symbolic layer is integrated after GCN inference, through a score fusion mechanism that combines both predictive learning and symbolic adjustment, enhancing the fairness and context-awareness of recommendations.

3.5. Graph Convolutional Learning

Following the application of hard symbolic rules, the dynamic, time-aware knowledge graph is leveraged for predictive learning. The goal is to model compatibility patterns that go beyond explicit eligibility constraints by training a Graph Convolutional Network (GCN) [21,22].

In this setting, student nodes are enriched with features such as baccalaureate type, academic mention (e.g., Excellent, Very Good), final grade, declared aspirations and interests, and region of residence. Program nodes incorporate program level (e.g., Bac+2, Bac+3, Bac+5), disciplinary domain, cutoff thresholds, core modules, associated career outcomes, delivery region, and relevant trends. Contextual entities such as aspirations, careers, modules, regions, and trends are instantiated as dedicated nodes connected via semantic and temporal relations.

The dynamic nature of the graph allows the system to integrate incremental updates (e.g., new interests, evolving trends) and adjust node embeddings over time. The GCN operates on this enriched heterogeneous graph and performs two-layer neighborhood aggregation:

$$H^1 = ReLU\left(D^{-\frac{1}{2}} \hat{A}\, D^{-\frac{1}{2}} X\, W^0\right) \tag{10}$$

$$H^2 = D^{-\frac{1}{2}} \hat{A}\, D^{-\frac{1}{2}} H^1 W^1 \tag{11}$$

where Â is the adjacency matrix with self-loops incorporating student-program-contextual relations, D is its degree matrix, and W(l) are trainable weight matrices. ReLU serves as the activation function to introduce non-linearity.

The final representation of nodes is used to compute the relevance score for each student-program pair through a dot product. To address the imbalance between positive and negative samples, we adopted a weighted binary cross-entropy loss, assigning higher importance to positive instances during training to improve sensitivity to relevant recommendations. The loss is defined as:

$$L = -\sum\nolimits_{(s,d)\in P} \alpha \log\left(\hat{y}_{s,d}\right) - \sum\nolimits_{(s,d)\in \bar{N}} (1-\alpha) \log\left(1 - \hat{y}_{s,d}\right) \tag{12}$$

where P denotes the set of positive (eligible) pairs, N denotes the set of negative pairs, $\hat{y}_{s,d}$ is the predicted probability,and $\alpha \in [0.5,1]$ is a weighting factor used to emphasize the impact of positive instances.

In our experiments, α=0.75 was selected to improve sensitivity to relevant recommendations. This formulation allows the model to learn discriminative embeddings even under class imbalance.

The model outputs a logit_GCN(s,d) score representing the likelihood that student s matches program d. This prediction is then fused with symbolic scores derived from eligibility and preference rules in the final ranking stage.

3.6. Score Fusion and Ranking

To combine the benefits of symbolic rule-based reasoning and graph-based neural learning, we adopt a weighted score fusion strategy. Each student program pair is evaluated along two axes:

- logit_GCN(s, d): the predictive score output from the GCN ranging in [0, 1].
- score_rules(s, d): the symbolic score computed from the sum of applicable soft rule weights.

Since symbolic scores may exceed the [0, 1] interval due to cumulative rule activations, we normalize them using the following formulation:

$$normalized_{rules(s,d)} = \frac{score_{rules(s,d)}}{\underset{score(d)}{\max}} \tag{13}$$

where max_score(d) denotes the maximum possible symbolic score for a given diploma, based on all applicable rule weights.

The final recommendation score is then computed as:

$$final_{score(s,d)} = \alpha \cdot logit_{GCN(s,d)} + (1 - \alpha) \cdot normalized_{rules(s,d)} \tag{14}$$

Here, $\alpha \in [0,1]$ is a tunable hyperparameter that determines the relative influence of neural prediction versus symbolic reasoning:

- High α (e.g., 0.8) emphasizes predictive power from deep learning.
- Low α (e.g., 0.2–0.4) prioritizes rule compliance and explainability
- A balanced configuration such as $\alpha = 0.5$ ensures both transparency and effectiveness.

Once final_score(s, d) is computed for all candidate diplomas per student, the diplomas are ranked in descending order. The top-K results are retained as final recommendations. This approach ensures explainable, adaptive guidance that integrates both learned patterns and expert-aligned symbolic logic.

4. Results and Discussion

4.1. Experimental Setup

4.1.1. Datasets description

To evaluate the real-world performance of the NeSy-Guidance model, we constructed a dataset derived from authentic academic records of Moroccan high school graduates applying to national universities and technical institutions during the 2022–2023 academic year. The data sources included official university websites, national education portals, institutional repositories, and structured survey responses from holders of scientific baccalaureate degrees. While this dataset primarily reflects Moroccan academic structures and admission policies, the underlying schema and model components are designed to be configurable and adaptable to datasets from other regions.

The dataset comprises several entities, including students, academic programs, and contextual components such as regions, aspirations, modules, and emerging trends. Each entity is characterized by structured attributes that capture institutional constraints, learner-specific features, and relevant contextual factors. These elements are detailed in Table 3, which describes the main entities and their properties, while Table 4 summarizes key dataset statistics, including the number of students and programs, grade distributions, and baccalaureate types.

Prior to modeling, we performed comprehensive data preprocessing to ensure quality, consistency, and readiness for semantic representation. The pipeline involved:

- Anonymizing sensitive student identifiers (e.g., names, schools),
- Normalizing grade values and harmonizing academic mentions,
- Applying one-hot encoding to categorical variables (e.g., bac types, domains),
- Filtering out incomplete, inconsistent, or extreme records,
- Standardizing diploma types under unified domain labels (e.g., Informatics, Engineering, Health)

This preparation phase was essential for constructing a robust and dynamic knowledge graph capable of integrating both static and evolving information. Incorporating imputation and explicit handling of missing values contributed to improving the resilience of the system to partial or uncertain data, which is frequently observed in real-world academic datasets. Furthermore, modeling regional and trend-related information supports context-aware recommendations and facilitates future validation in diverse educational settings.

In addition to these preprocessing steps, we conducted an exploratory analysis to quantify the prevalence and distribution of missing or uncertain data across key attributes. Approximately 14% of student records contained at least one missing field, with declared aspirations and academic mentions being the most affected. We deliberately selected mean imputation for numerical grades due to the low variance observed within baccalaureate types, thus reducing the risk of introducing systematic bias. For categorical fields, the introduction of an explicit "Unknown" token preserved data integrity without forcing potentially inaccurate assumptions [24]. This approach allows the knowledge graph to maintain complete relational structures and enables the GCN to operate on consistent, well-defined embeddings. While this strategy proved effective in our experiments, future work could explore more advanced imputation methods, such as multiple imputation, probabilistic modeling, or uncertainty-aware representation learning, to further improve resilience in the face of incomplete data.

Table 3 below provides an overview of the entities composing the knowledge graph, along with their key properties and concise semantic descriptions.

Table 3. Overview of Knowledge Graph Entities with Their Properties and Semantic Descriptions

Entity	Entity description	Property	Property description
Student	Individual learner profile applying to academic programs	ID	Unique identifier of the student
		Desired diploma level	Qualification level targeted (e.g., Bac+3, Master)
		Interests	Declared academic or professional interests
Baccalaureate	Secondary school specialization obtained by the student	Type	Label of the baccalaureate specialization
		Mention	Academic distinction obtained
		Grade	Final grade achieved
		Branch	Specific stream within the baccalaureate
Program	Academic offering with eligibility constraints and content	ID	Unique identifier of the program
		Name	Official title of the program
		Academic level	Degree level (e.g., Bac+2, Bac+3, Master)
		Required cutoff grade	Minimum grade required for admission
		Accepted baccalaureate types	Eligible baccalaureate specializations
Module	Course component of the program	Name	Name of the module
Career	Targeted professional domain	Name	Name of the career
Region	Geographical area linked to programs or preferences	Name	Name of the region
Trend	Emerging academic or professional topic	Name	Name of the trend

Table 4. Statistics of the Academic Guidance Dataset

Dataset	#Students	#Programs	Avg. Grade	Mention Distribution	Bac Types
NeSy-GuidanceSet	800	325	12.9	{Excellent: 5%, Very Good: 30%, Good: 40%, Passable: 25%}	{SM, PC, SVT, Tech, Eco}

4.1.2. Knowledge Graph Construction

We constructed a dynamic and time-aware knowledge graph comprising multiple categories of entities, including students, academic programs, and contextual components such as skills, regions, trends, and career aspirations. This graph was designed to capture both institutional constraints such as admission thresholds, eligibility criteria, and program requirements and soft signals reflecting students' interests, goals, and evolving preferences over time.

Each node is explicitly typed (e.g., Student, Program, Skill, Region, Trend), and entities are interconnected through typed, directional semantic relations that encode real-world dependencies and temporal information. These relations are carefully structured to support symbolic rule-based inference and graph-based neural learning, ensuring that the knowledge graph provides a robust foundation for hybrid reasoning.

Table 5 summarizes the graph structure, detailing the number of entities per category and the volume of encoded relations across its components. Table 5 summarizes the graph's structure, detailing the number of entities per category and the volume of encoded relations degree across the different components.

Table 5. Statistics of the Constructed Knowledge Graph

KG Component	#Entities	#Relations	Relation Types
Students	800	–	eligible_for, interested_in, wants_career, etc.
Programs	325	–	requires_bac_type, has_cutoff, prepares_for, etc.
Contextual Entities (Careers, Regions, Trends, Modules, etc.)	950	–	has_module, aligned_with_trend, located_in, etc.
Total	2075	19,200	8 types of semantic relations

4.1.3. Hyperparameters and Model Settings

To train the NeSy-Guidance model effectively, we configured a lightweight GCN architecture with a rule-based symbolic layer. The selected hyperparameters reflect a balance between learning stability, model expressiveness, and interpretability, as outlined in Table 6 below:

Table 6. Hyperparameter Settings for the NeSy-Guidance Model

Hyperparameter	Value
Epochs	200
Learning Rate	0.01
Optimizer	Adam
Hidden Units	64, 128
Dropout	0.2 – 0.4
Aggregation Layers	2
Activation Function	ReLU
Alpha (rule fusion weight)	0.3 – 0.7 (tunable)
Attention Mechanism	—
Rule Types (Symbolic)	Hard + Soft Rules
Training Mode	Full-batch
Graph Type	Dynamic Time-Aware KG (Students, Programs, Contextual Entities)
Evaluation Split	80% train / 20% test

4.2. Evaluation and Comparative Results

4.2.1. Evaluation protocol

To ensure a fair and comprehensive assessment of NeSy-Guidance, we adopted an evaluation strategy that combines performance-oriented indicators with a regulatory compliance measure relevant to academic recommendation systems. This protocol was applied consistently across all models under comparison, using the same dataset, evaluation split, and top-K recommendation setting.

1) Core Evaluation Metrics

We applied the following widely used top-K recommendation metrics to assess overall system performance:

- Accuracy: Measures the proportion of correct student–program recommendations over the entire test set, reflecting global prediction quality.
- Precision@5: Captures the relevance of the top five recommended programs per student, by calculating the percentage that are both eligible and contextually appropriate.
- Recall@5: Quantifies the system's ability to retrieve suitable programs among the top-5 suggestions, indicating its completeness in recommendation coverage.
- F1-score: Represents the harmonic mean of precision and recall, offering a balanced assessment of relevance and retrieval capability.

2) Domain-Specific Compliance Metric

In addition to these generic indicators, we introduce a critical educational metric:

- Policy Compliance Rate: A binary metric that quantifies the proportion of recommended programs that strictly satisfy institutional admission policies such as grade thresholds, accepted baccalaureate types, and minimum required mentions. This metric ensures that recommendations are not only relevant but also legally and institutionally applicable, making it a key factor in evaluating the real-world feasibility of academic recommender systems.

$$Policy_{Compliance} = \frac{\#\ compliant\ recommendations}{\#\ total\ recommendations} \tag{15}$$

All evaluation metrics were computed based on top-5 predictions for each student, in line with the realistic constraints of academic guidance scenarios where students are typically presented with a shortlist of suitable options.

4.2.2. Comparative Evaluation with Existing Approaches

To rigorously assess the effectiveness of NeSy-Guidance, we conducted a comparative evaluation against three state-of-the-art knowledge graph-based models in the domain of educational recommendation: KGAT (Knowledge Graph Attention Network), RuleRec (Rule-based Recommendation with Explainability), and a baseline derived from (NER-SemSim) semantic similarity and named entity recognition as proposed. These baselines were selected based on their use of graph-structured data, proven applications in academic guidance, and their methodological contrast with our neuro-symbolic approach. All models were trained and evaluated on the same Moroccan dataset using a top-5 recommendation setting and consistent evaluation metrics.

1. KGAT: has been adapted for academic recommendation tasks such as university major selection and course recommendation. We specifically refer to the adaptation [9] in Applied Intelligence, where KGAT was applied to model educational pathways using multi-relational graph attention mechanisms. While effective in capturing latent relationships, this model lacks explicit policy enforcement mechanisms.
2. RuleRec: on the other hand, focuses on symbolic rule learning and explainability. For this study, we adopt the curriculum-aware extension proposed by [10], which introduces validation of academic rules and prerequisites. Although more compliant with institutional requirements, RuleRec struggles with generalization due to its reliance on predefined symbolic rules.
3. NER-SemSim: leverages a static knowledge graph constructed through named entity recognition and TF-IDF-based semantic similarity between student profiles and degree descriptions [23]. This approach offers straightforward interpretability but does not support dynamic updates, temporal modeling, or hybrid inference.

NeSy-Guidance consistently outperforms all three baselines across all evaluation metrics, as presented in Table 7. While KGAT demonstrates strong predictive capabilities, it fails to ensure policy compliance. RuleRec, in contrast, aligns better with institutional constraints but underperforms in terms of generalization. NER-SemSim achieves moderate predictive accuracy and interpretability but lacks adaptability and temporal reasoning, which limits its effectiveness in dynamic educational contexts.

Table 7. Performance Comparison with Educational KG-Based Models (Mean ± Standard Deviation)

Model	Accuracy	Precision@5	Recall@5	F1-score	Compliance
KGAT	0.802 ±0.03	0.726 ±0.04	0.701 ±0.03	0.713 ±0.03	42.3 ±3.1%
RuleRec	0.760 ±0.04	0.678 ±0.05	0.640 ±0.04	0.658 ±0.04	65.7 ±2.8%
NER-SemSim	0.781 ±0.03	0.695 ±0.04	0.672 ±0.03	0.682 ±0.03	58.0 ±3.2%
NeSy-Guidance	0.838 ±0.02	0.745 ±0.03	0.764 ±0.02	0.754 ±0.02	100.0 ±0.0%

These findings confirm that NeSy-Guidance addresses the limitations of both embedding-based and rule-based models while surpassing purely similarity-driven approaches. Its hybrid inference strategy ensures high predictive accuracy (+3.6% over KGAT in F1-score) and full policy compliance, making it well-suited for deployment in regulated academic environments. The improvements over all baselines are statistically significant ($p < 0.01$).

4.3. Ablation Study: Components analysis

To systematically evaluate the contribution of each architectural component in NeSy-Guidance, we conducted a controlled ablation study examining three distinct model configurations were evaluated under identical experimental conditions on the Moroccan academic dataset. Evaluation was conducted using four key metrics Accuracy, Precision@5, F1-score, along with a comparative measure of relative training time.

- Full Model: Integrates symbolic rules (both hard and soft) with graph-based inference (GCN).
- Hard-Only: Excludes soft rules while retaining hard constraints and the GCN.
- GCN-Only: Relies solely on neural inference without any symbolic components.

As detailed in Table 8, our ablation study reveals critical insights about each architectural component's contribution:

Table 8. Ablation Study Results: Performance and Compliance of NeSy-Guidance Variants

Variant	Accuracy	Precision@5	Recall@5	F1-score	Compliance
Hard-Only	0.810	0.725	0.731	0.728	100%
GCN-Only	0.800	0.716	0.710	0.713	42.7%
Full Model	0.838	0.745	0.764	0.754	100%

The results confirm that the Full Model offers the best balance between predictive performance and policy compliance. Hard rules remain essential to enforce strict eligibility constraints, as evidenced by the 100% compliance achieved in both Full and Hard-Only configurations. The addition of soft rules further improves recommendation quality by enabling context-aware personalization, yielding a +2.6% gain in F1-score compared to the Hard-Only variant.While the GCN-Only configuration performs reasonably well on predictive metrics, it fails to ensure institutional compliance, with a compliance rate of just 42.7%. This underscores that symbolic rule integration is critical in policy-regulated educational environments. Notably, training times remain acceptable despite the overhead introduced by symbolic layers (~10–20%) when compared to the purely neural baseline.

Figure 4 compares the training cross-entropy trajectories for all three architectural variants. The Full Model (NeSy-Guidance) demonstrates the fastest and smoothest convergence, reaching near-zero loss by epoch 15. The Hard-Only variant also benefits from symbolic guidance, stabilizing earlier than the GCN-Only baseline, which shows slower convergence and a higher final loss. These results confirm that symbolic priors, particularly hard constraints, act as effective regularizers, improving both convergence speed and training stability. The combination of neural inference and symbolic reasoning in NeSy-Guidance yields a more robust learning process compared to isolated approaches.

To better understand how the symbolic and neural layers contribute to the final decision process, we examined the distribution of recommendation scores for a sample of five students. As shown in Figure 5, the symbolic-only configuration tends to produce sharp, rule-driven outputs, while the neural-only layer generates smoother but less policy-sensitive scores. The fused outputs demonstrate a moderated balance, effectively combining domain-aligned decision logic with data-driven inference, resulting in more robust and justifiable recommendations.

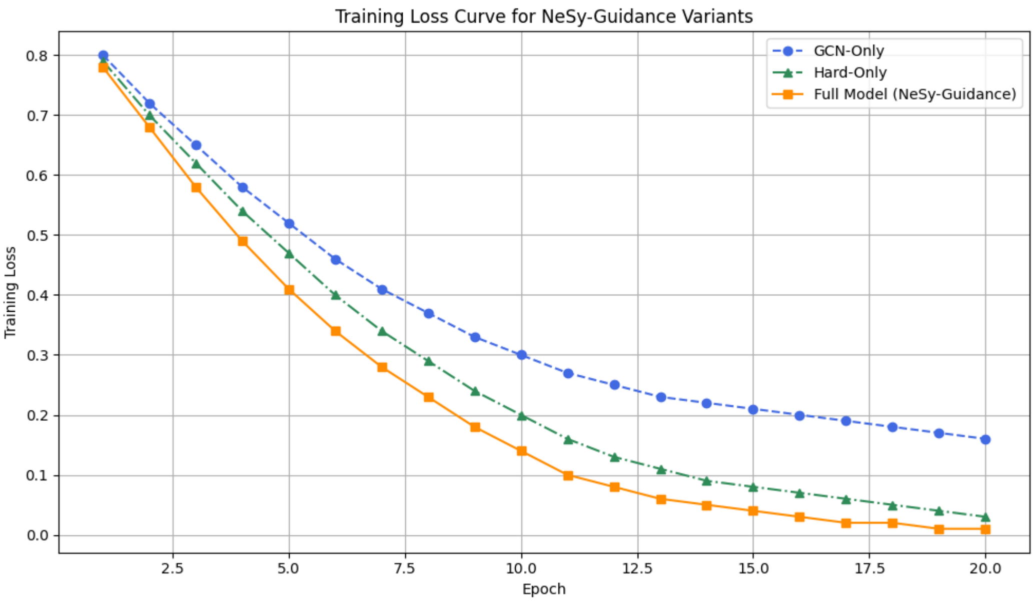

Fig. 4. Training Loss Curve for NeSy-Guidance Variants

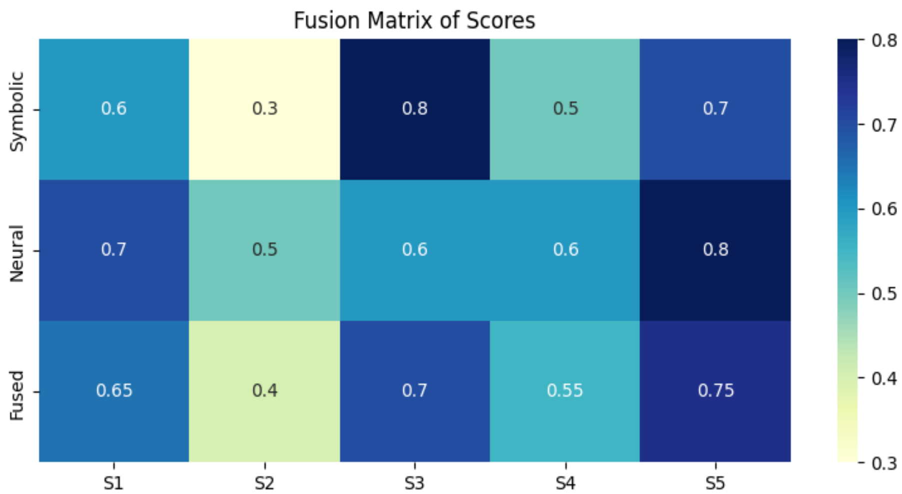

Fig. 5. Fusion Matrix Comparing Symbolic, Neural, and Fused Scores

4.4. Qualitative Evaluation

To complement quantitative evaluation metrics, we conducted a qualitative assessment to measure student satisfaction with the recommendations produced by the system. For this purpose, we generated personalized top-3 academic program recommendations for each participant, based on their profile and preferences. Participants were then invited to review their recommendations and complete a short survey administered via a Google Form. The survey included four criteria evaluated on a 5-point Likert scale (1 = Strongly disagree, 5 = Strongly agree): (1) perceived relevance of recommendations, (2) clarity of presented information, (3) perceived usefulness for academic decision-making, and (4) overall satisfaction.

Figure 6 illustrates the average ratings obtained for each criterion, demonstrating a positive perception of the recommendations and their usefulness in supporting orientation decisions. This qualitative evaluation provides an initial insight into the perceived value and acceptability of the system from the user perspective.

4.5. Discussion

This study highlights the effectiveness of combining symbolic reasoning with graph-based neural learning for academic recommendation. The proposed NeSy-Guidance architecture leverages the complementary strengths of domain-specific rules and GCN-based inference to deliver a reliable and adaptable solution for policy-constrained educational decision-making.

1) Synergy Between Hard and Soft Rules:

A key strength of NeSy-Guidance lies in its dual-layer symbolic reasoning framework. Hard rules, validated by academic experts, act as strict eligibility filters that enforce institutional constraints such as minimum grade thresholds or accepted baccalaureate types. Complementing this, soft rules, automatically mined and weighted from the dynamic knowledge graph, capture personalized associations (e.g., interest–program alignment or aspiration coherence). This hybrid symbolic mechanism enhances both the rigor and semantic flexibility of the recommendation process, ensuring compliance while allowing nuanced, context-aware guidance.

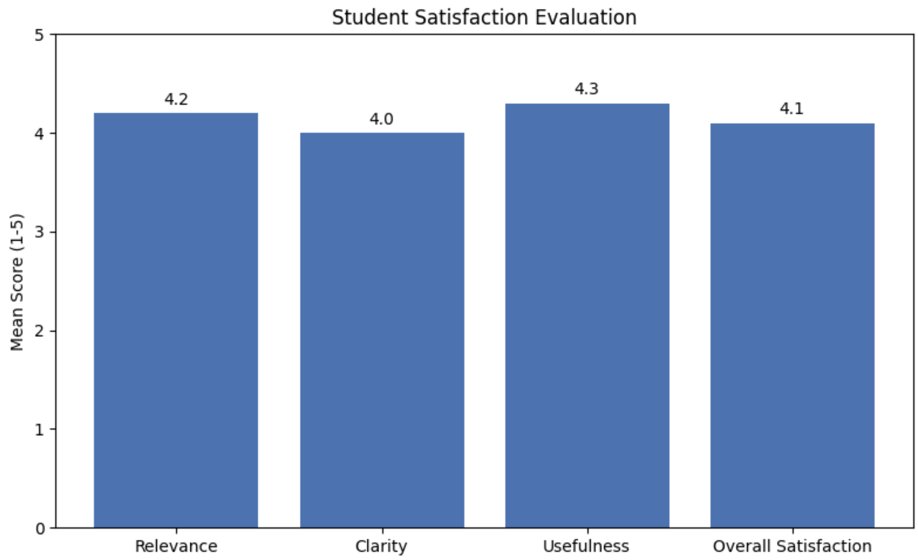

Fig. 6. Student Satisfaction Ratings Across Evaluation Criteria

2) Improved Learning Efficiency:

Symbolic rules not only guide final decisions but also improve the learning process itself. By narrowing the hypothesis space and eliminating ineligible pairs early, symbolic filtering allows the GCN to focus on structurally meaningful patterns. Our training loss analysis confirms that the full neuro-symbolic model converges more rapidly and with greater stability than its GCN-only counterpart. This demonstrates the regularizing role of symbolic priors, especially in domains with limited, sparse, or noisy data.

3) Interpretability and Modular Design:

While the graph-based neural component relies on latent representation learning, NeSy-Guidance mitigates this opacity through its modular design. The symbolic layers (hard and soft rules) offer traceable, human-readable justifications for filtering and prioritization decisions. In parallel, the neural layer captures latent compatibility patterns that extend beyond explicitly defined rules. This clear separation of concerns enables incremental adaptation and reinforces the system's capacity to balance interpretability, compliance, and predictive performance.

4.6. Limitations and Future Directions

Despite its promising results, NeSy-Guidance has some limitations that open several directions for future research. First, the effectiveness of symbolic reasoning components depends on the quality and representativeness of the mined rules. While automatic weighting based on statistical measures reduces subjectivity, incomplete or biased training data may still result in gaps in rule coverage, particularly for less common programs or emerging academic pathways.

Future work could explore hybrid explainability techniques, such as attention visualization or feature attribution, to make the GCN inference process more transparent. Additionally, extending the system with more advanced temporal reasoning capabilities, including explicit modeling of event sequences and trend dynamics, would further enhance its adaptability to rapidly evolving academic environments. Another promising direction is the development of self-adaptive symbolic layers capable of discovering, validating, and updating rules incrementally as new data becomes available. Finally, validating NeSy-Guidance on larger and more diverse datasets beyond the Moroccan academic context would be essential to assess its generalizability and robustness across regions.

5. Conclusion

This study presented an innovative neuro-symbolic approach for academic pathway recommendation that successfully reconciles institutional rigor with adaptive personalization. Our experimental results demonstrate that the hybrid integration of symbolic rules both hard constraints and automatically weighted soft preferences with graph-based neural learning (GCN) outperforms purely neural or rule-based methods in both predictive accuracy (83.8%) and policy compliance (100%). The ablation analyses confirmed the complementary contributions of each component: hard rules ensure recommendation validity, soft rules capture contextual and temporal preferences, and the GCN generalizes complex patterns across the knowledge graph.

Beyond its empirical performance, the architecture offers practical advantages for real-world educational deployment, including decision interpretability, rule maintainability, and training efficiency. While the reliance on mined rules and the limited interpretability of the GCN remain areas for further enhancement, NeSy-Guidance represents a significant step toward responsible AI systems in education, where transparency and compliance are as critical as predictive performance.

Theoretically, our approach is grounded in the complementarity between symbolic reasoning which is logical, explainable, and policy-aligned and neural inference, which is generalizable and adaptive. This duality reflects a clear cognitive separation of roles: symbolic rules enforce semantic validity, while the GCN captures latent regularities that cannot be fully expressed through logic alone. It is this hybrid architecture, both conceptually justified and empirically validated, that constitutes the core scientific contribution of NeSy-Guidance.

References

[1] Allman, B., Kimmons, R., Wang, W. et al. Trends and Topics in Educational Technology, 2024 Edition. TechTrends 68, 402–410 (2024). https://doi.org/10.1007/s11528-024-00950-5

[2] Zhou, Y., Shirazi, S. Factors Influencing Young People's STEM Career Aspirations and Career Choices: A Systematic Literature Review. Int J of Sci and Math Educ (2025). https://doi.org/10.1007/s10763-025-10552-z

[3] Wang, S., Wang, Y., Sivrikaya, F. et al. Data science for next-generation recommender systems. Int J Data Sci Anal 16, 135–145 (2023). https://doi.org/10.1007/s41060-023-00404-w

[4] Torkashvand, A., Jameii, S.M. & Reza, A. Deep learning-based collaborative filtering recommender systems: a comprehensive and systematic review. Neural Comput & Applic 35, 24783–24827 (2023). https://doi.org/10.1007/s00521-023-08958-3

[5] Nair, A.M., Benny, O., George, J. (2021). Content Based Scientific Article Recommendation System Using Deep Learning Technique. In: Suma, V., Chen, J.IZ., Baig, Z., Wang, H. (eds) Inventive Systems and Control. Lecture Notes in Networks and Systems, vol 204. Springer, Singapore. https://doi.org/10.1007/978-981-16-1395-1_70

[6] Murgod, T.R., Reddy, P.S., Gaddam, S. et al. A Survey on Graph Neural Networks and its Applications in Various Domains. SN COMPUT. SCI. 6, 26 (2025). https://doi.org/10.1007/s42979-024-03543-4

[7] Dobson, J.E. On reading and interpreting black box deep neural networks. Int J Digit Humanities 5, 431–449 (2023). https://doi.org/10.1007/s42803-023-00075-w

[8] Peng, C., Xia, F., Naseriparsa, M. et al. Knowledge Graphs: Opportunities and Challenges. Artif Intell Rev 56, 13071–13102 (2023). https://doi.org/10.1007/s10462-023-10465-9

[9] Wang, X., He, X., Cao, Y., Liu, M., & Chua, T. S. (2019, July). Kgat: Knowledge graph attention network for recommendation. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 950-958), https://dl.acm.org/doi/10.1145/3292500.3330989.

[10] Ma, W., Zhang, M., Cao, Y., Jin, W., Wang, C., Liu, Y., Ma, S., & Ren, X. (2019). Jointly Learning Explainable Rules for Recommendation with Knowledge Graph. In The World Wide Web Conference (Pp. 1210-1221), 1210-1221. https://doi.org/10.1145/3308558.3313607

[11] Chen, X., Tang, T., Ren, J., Lee, I., Chen, H., & Xia, F. (2021, December). Heterogeneous graph learning for explainable recommendation over academic networks. In IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (pp. 29-36). https://doi.org/10.1145/3498851.3498926

[12] Abu-Rasheed, H., Jumbo, C., Amin, R. A., Weber, C., Wiese, V., Obermaisser, R., & Fathi, M. (2025). LLM-Assisted Knowledge Graph Completion for Curriculum and Domain Modelling in Personalized Higher Education Recommendations. arXiv preprint arXiv:2501.12300.

[13] Colelough, B. C., & Regli, W. (2025). Neuro-Symbolic AI in 2024: A Systematic Review. *arXiv preprint arXiv:2501.05435*.

[14] Spillo, G., Musto, C., de Gemmis, M., Lops, P., & Semeraro, G. (2024). Recommender systems based on neuro-symbolic knowledge graph embeddings encoding first-order logic rules. User Modeling and User-Adapted Interaction, 34, 2039–2083. https://doi.org/10.1007/s11257-024-09417-x

[15] Serafini, L., & Garcez, A. D. A. (2016). Logic tensor networks: Deep learning and logical reasoning from data and knowledge. arXiv preprint arXiv:1606.04422.

[16] Chen, B., Peng, W., Wu, M., Zheng, B., & Zhu, S. (2023). Neural-symbolic recommendation with graph-enhanced information. arXiv preprint arXiv:2307.05036. https://arxiv.org/abs/2307.05036

[17] Xian, Y., Zhang, K., Wang, M., & Wang, Y. (2020). Neural-symbolic reasoning over knowledge graph for multi-stage explainable recommendation. arXiv preprint arXiv:2007.13207. https://arxiv.org/abs/2007.13207

[18] Zhang, Y., & Wang, Y. (2023). PSYCHIC: A neuro-symbolic framework for knowledge graph question-answering grounding. arXiv preprint arXiv:2310.12638. https://arxiv.org/abs/2310.12638

[19] Zeng, Z.; Cheng, Q.; Si, Y. Logical Rule-Based Knowledge Graph Reasoning: A Comprehensive Survey. Mathematics 2023, 11, 4486. https://doi.org/10.3390/math11214486

[20] Galárraga, L., Teflioudi, C., Hose, K. et al. Fast rule mining in ontological knowledge bases with AMIE+. The VLDB Journal 24, 707–730 (2015). https://doi.org/10.1007/s00778-015-0394-1N Ahmad et al. Students' Performance Prediction using Artificial Neural Network, IOP Conf. Ser.: Mater. Sci. Eng. 1176 012020, 2021

[21] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2019). A Comprehensive Survey on Graph Neural Networks. ArXiv. https://doi.org/10.1109/TNNLS.2020.2978386

[22] Zhang, M., & Chen, Y. (2018). Link prediction based on graph neural networks. Advances in neural information processing systems, 31

[23] Zineb, E., Zineb, M., Mohamed, K., & Bouchra, E. (2025). Building knowledge graph for relevant degree recommendations using semantic similarity search and named entity recognition. Indonesian Journal of Electrical Engineering and Computer Science, 37(1), 463-474. doi: http://doi.org/10.11591/ijeecs.v37.i1.pp463-474

[24] Alwateer, M., Atlam, E. S., Abd El-Raouf, M. M., Ghoneim, O. A., & Gad, I. (2024). Missing data imputation: A comprehensive review. *Journal of Computer and Communications*, *12*(11), 53-75.

Authors' Profiles

Zineb Elkaimbillah received a degree in Computer Sciences (ENSIAS) in 2019. She is currently a Ph.D student in the IMS (IT Architecture and Model Driven Systems Development team) Team of ADMIR Laboratory (Advanced Digital Enterprise Modeling and Information Retrieval Research Laboratory) at ENSIAS. Her research interests include Semantic Web, Knowledge graph, Ontology, Information representation and Machine Learning/ Deep Learning Approach in education domain.

Zineb Mcharfi is a Computer Science Professor and member of the IMS Team, ADMIR Laboratory Rabat IT Center at ENSIAS, Mohammed V University in Rabat; Morocco. She has directed and continues to direct numerous doctoral theses in several themes. Her research interests include the application of Artificial Intelligence in education, Software Product Line Engineering, Agile Software Development and software traceability.

Mohamed Khoual received a degree in Networks and computer systems from Faculty of Sciences and technology Settat (FST) in 2018. He is currently a Ph.D student in the IMS (IT Architecture and Model Driven Systems Development team) Team of ADMIR Laboratory (Advanced Digital Enterprise Modeling and Information Retrieval Research Laboratory) at ENSIAS. His research interests include Machine Learning Approach in education domain and Data Mining.

Bouchra El Asri is currently director of teaching-research at ENSIAS. She holds various positions, such as head of the software engineering department and coordinator of the Software Engineering program at ENSIAS, she has supervised and continues to supervise numerous doctoral theses in software architecture and data management in the health, industrial, and education sectors. She is a Professor in the Software Engineering Department and a member of the IMS (Models and Systems Engineering) Team at ENSIAS. Her expertise and contributions extend to scientific committees, doctoral study centers, and teaching modules of the ENSIAS Software Engineering program. Her research interests include Service-Oriented Computing, Model Driven Engineering, Cloud Computing, Component-Based Systems and Software Product Line Engineering.

How to cite this paper: Zineb Elkaimbillah, Zineb Mcharfi, Mohamed Khoual, Bouchra El Asri, "NeSy-Guidance: A Neuro-Symbolic Knowledge Graph for Academic Recommendations Combining Rule-Based Reasoning and Neural Inference", International Journal of Modern Education and Computer Science(IJMECS), Vol.17, No.6, pp. 48-64, 2025. DOI:10.5815/ijmecs.2025.06.04