BIOMETAEVO-GNN: BAYESIAN FOURIER MESSAGE PASSING WITH IMMUNE MEMORY FOR ROBUST COMMUNITY DETECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Community detection in real-world graphs is challenged by noisy edges, ambiguous node features, perturbation-induced assignment drift, and unstable community boundaries. Existing GNN-based methods typically optimize clustering accuracy or local robustness, but rarely unify uncertainty modeling, memory preservation, and adaptive defense. We propose BioMetaEvo-GNN, which integrates Bayesian Fourier learning, immune-memory defense, and UCB-based meta-governance. Bayesian Fourier learning identifies spectral instability and drift-prone boundaries, immune memory preserves reliable structural evidence and suppresses unstable transitions, and meta-governance adaptively selects strategies using uncertainty, stability, and historical rewards. Experiments on real-world and synthetic graphs show consistent gains in clustering accuracy, perturbation robustness, uncertainty calibration, assignment stability, and cross-dataset generalization over representative GNN and robust graph learning baselines. These results establish community detection as a joint problem of representation learning, uncertainty estimation, and adaptive robustness.

1 INTRODUCTION

Community detection identifies densely connected node groups and supports applications in social networks, citation analysis, biological systems, recommendation, and fraud detection. With graph neural networks (GNNs), the field has shifted from topology-based partitioning toward representation learning that jointly exploits graph structure and node attributes. Recent methods can recover communities in sparse and attributed graphs without requiring a predefined number of clusters Kojaku et al. (2024); Liu et al. (2024a), while community-aware representations also benefit downstream tasks such as link prediction Liu et al. (2024b). However, real-world community detection remains vulnerable to graph noise, uncertainty, and perturbation-induced assignment drift.

Practical graphs often contain missing or spurious edges, corrupted attributes, low homophily, distribution shifts, and evolving interactions Wu et al. (2025b); Zheng et al. (2025). Message passing can further propagate such noise, and GNN robustness depends strongly on graph topology, node uncertainty, and perturbation type Jin & Zhu (2025); Wang et al. (2024c); Abbahaddou et al. (2024); Gosch et al. (2023); Hou et al. (2024). In community detection, even minor changes to bridge edges or boundary nodes may cause substantial assignment drift under weak community structure Ma et al. (2026).

Uncertainty quantification provides a way to identify these failures and increasingly serves as an active component of graph learning rather than a post-hoc diagnostic Wang et al. (2024a); Kong et al. (2024). Conformal prediction, uncertainty-aware learning, structure refinement, and confidence calibration have improved graph prediction reliability Huang et al. (2023a); Chen et al. (2025); Han et al. (2025); Yang et al. (2024). Yet most approaches focus on node-level prediction and do not jointly model spectral instability, community-boundary drift, historical partition evidence, and adaptive community-level intervention. Existing neural community detection, robust GNN, and uncertainty-aware methods therefore remain insufficient for preserving stable communities under perturbation.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To address this gap, we propose **BioMetaEvo-GNN**, a robust and uncertainty-aware framework for perturbation-stable community detection. Bayesian Fourier learning models spectral uncertainty and identifies unstable components and drift-prone boundaries, while energy-time filtering suppresses unreliable responses. An immune-inspired mechanism combines T-cell defense dynamics with T-SILE memory to reject unstable transitions and retain reliable historical evidence. A lightweight Upper Confidence Bound controller adaptively selects community-level strategies according to uncertainty, stability, and historical rewards. Experiments on real-world and synthetic graphs demonstrate improved clustering accuracy, perturbation robustness, uncertainty calibration, assignment stability, and cross-dataset generalization over representative neural, robust, and stochastic graph learning baselines.

2 Related Work

Neural community detection increasingly combines graph topology and node attributes through learned representations. Neural embeddings can recover spectral community structures, while adaptive generative methods support attributed clustering without a predefined number of communities Kojaku et al. (2024); Liu et al. (2024a). Community-aware GNNs also improve downstream tasks such as link prediction Liu et al. (2024b). However, most methods emphasize clean-graph partition quality, despite evidence that structural and attribute perturbations can substantially alter community assignments Goel et al. (2026).

Robust graph learning addresses feature attacks, structural disturbances, label noise, security threats, and distribution shifts Gosch et al. (2023); Abbahaddou et al. (2024); Wang et al. (2024c); Guan et al. (2024); Jin & Zhu (2025). These methods mainly target node classification or general graph prediction, whereas community detection requires group-level stability because local perturbations near bridge nodes can induce global assignment drift. Uncertainty-aware GNNs further provide conformal prediction, epistemic uncertainty estimation, calibration, and uncertainty-guided structure refinement Huang et al. (2023a); Trivedi et al. (2024); Wang et al. (2024a); Yang et al. (2024); Wang et al. (2024b); Han et al. (2025), but rarely connect uncertainty with spectral boundary instability or community-transition control.

Dynamic GNNs model evolving graph structures and interactions Zheng et al. (2025); Huang et al. (2023b), yet mainly focus on temporal prediction rather than distinguishing genuine community evolution from perturbation-induced drift. Verification and explanation methods provide additional diagnostics for unreliable graph predictions Deng (2025); Jiang et al. (2024), but do not perform adaptive community-level intervention. BioMetaEvo-GNN instead unifies Bayesian Fourier uncertainty modeling, immune-memory transition control, and UCB-based meta-governance to preserve reliable community evidence and suppress unstable assignments under noisy and evolving conditions.

3 Method

3.1 Framework Overview and Problem Formulation

BioMetaEvo-GNN is designed around three reliability requirements for community detection under noisy and perturbed graphs: detecting unstable spectral and boundary patterns, controlling unreliable community transitions, and adaptively selecting corrective actions. Accordingly, the framework integrates Bayesian Fourier uncertainty modeling, immune-memory transition control, and UCB-based meta-governance into a closed-loop graph learning process.

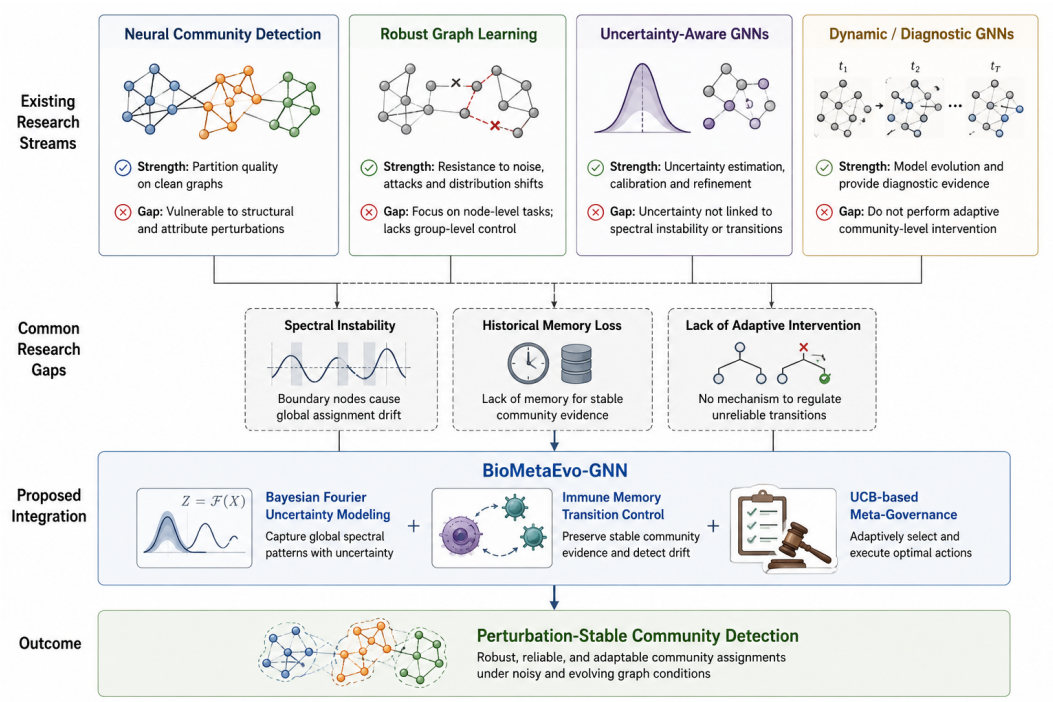

Figure 1: Related-work landscape and research gaps addressed by BioMetaEvo-GNN.

As summarized in Figure 1, prior work treats community detection, robustness, uncertainty estimation, and dynamic graph learning largely as separate problems. BioMetaEvo-GNN bridges these directions by jointly modeling spectral instability, historical community memory, and adaptive community-level intervention.

Let the input attributed graph be

$$G = (V, E, X), \tag{1}$$

where V and E denote the node and edge sets, respectively, and $X \in \mathbb{R}^{N \times d}$ is the node-feature matrix. The objective is to learn a partition $\mathcal{C}^* = \{C_1, \ldots, C_K\}$ through

$$\mathcal{C}^* = \Pi_{\text{meta}}\left(G; \Theta_{\text{GNN}}, \Theta_{\text{BFL}}, \Theta_{\text{immune}}, \Theta_{\text{gov}}\right), \tag{2}$$

where Π_{meta} coordinates representation learning, spectral uncertainty estimation, immune-memory control, and adaptive governance.

3.2 Graph Representation and Bayesian Fourier Learning

The initial node representation is $H^{(0)} = X$. At layer l, the GNN backbone performs normalized message passing:

$$H^{(l+1)} = \sigma\left(\widetilde{D}^{-\frac{1}{2}}\widetilde{A}\widetilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}\right), \tag{3}$$

where $\widetilde{A} = A + I$, $\widetilde{D}$ is its degree matrix, and $W^{(l)}$ contains trainable parameters. After L layers, the node embeddings and preliminary soft assignments are

$$Z = H^{(L)}, \qquad P = \text{softmax}(ZW_c). \tag{4}$$

Spatial message passing alone may amplify noisy edges and fail to expose globally unstable community boundaries. We therefore transform the embeddings into the Fourier domain:

$$\widehat{Z} = \mathcal{F}(Z). \tag{5}$$

Let θ denote the Bayesian Fourier parameters. The predictive distribution is

$$p(Y \mid \widehat{Z}) = \int p(Y \mid \widehat{Z}, \theta) p(\theta \mid \widehat{Z}) \, d\theta. \tag{6}$$

The resulting node-level uncertainty is measured by predictive entropy:

$$U_i = -\sum_{k=1}^{K} P_{ik} \log P_{ik}, \qquad \mathcal{U}(C_k) = \frac{1}{|C_k|} \sum_{v_i \in C_k} U_i. \tag{7}$$

High $\mathcal{U}(C_k)$ indicates ambiguous memberships or a drift-prone community boundary.

To prevent transient perturbations from directly changing the partition, the representation is smoothed across iterations:

$$\overline{Z}_t = \alpha \overline{Z}_{t-1} + (1 - \alpha) Z_t, \tag{8}$$

where α controls historical retention. The corresponding community energy reliability is denoted by $\Phi(C_k)$ and is used together with uncertainty and structural quality in the subsequent transition-control stage. Its complete formulation is provided in the Appendix.

3.3 Immune-Memory Transition Control

The immune-inspired component is a computational transition-control mechanism rather than a biological simulation. For each candidate community C_k, its current prototype is

$$q_k^t = \frac{1}{|C_k|} \sum_{v_i \in C_k} \overline{z}_i^t. \tag{9}$$

BioMetaEvo-GNN maintains a memory prototype m_k^{t-1} representing historically reliable community evidence. The deviation from memory is measured as

$$\mathcal{A}(C_k) = 1 - \operatorname{sim}\left(q_k^t, m_k^{t-1}\right). \tag{10}$$

The community defense score jointly considers robustness, structural consistency, energy reliability, uncertainty, and historical deviation:

$$\begin{aligned} D(C_k) = \gamma_1 \mathcal{R}(C_k) + \gamma_2 \mathcal{S}(C_k) + \gamma_3 \Phi(C_k) \\ - \gamma_4 \mathcal{U}(C_k) - \gamma_5 \mathcal{A}(C_k). \end{aligned} \tag{11}$$

The candidate is then activated, monitored, or suppressed:

$$\mathcal{I}(C_k) = \begin{cases} \text{activate}, & D(C_k) \geq \tau_D, \\ \text{monitor}, & \tau_M \leq D(C_k) < \tau_D, \\ \text{suppress}, & D(C_k) < \tau_M. \end{cases} \tag{12}$$

Activated communities participate directly in optimization, monitored communities undergo repair or stabilization, and suppressed communities cannot immediately alter the current partition. Reliable prototypes update the T-SILE memory bank through

$$m_k^t = \beta m_k^{t-1} + (1 - \beta) q_k^t, \tag{13}$$

where β balances historical preservation and adaptation. This mechanism limits perturbation-induced assignment drift without forcing communities to remain static. Detailed Bio-Phase screening, memory loss, and defense–enhancement–transfer operations are provided in the Appendix.

3.4 UCB-Based Adaptive Meta-Governance

A fixed defense strategy may be unsuitable when partition quality, uncertainty, and stability provide conflicting signals. BioMetaEvo-GNN therefore defines the community-level action space

$$\mathcal{A} = \{a_{\text{def}}, a_{\text{enh}}, a_{\text{trans}}, a_{\text{replace}}, a_{\text{grow}}, a_{\text{stabilize}}\}. \tag{14}$$

At iteration t, the controller assigns each action the UCB score

$$\mathrm{UCB}_a(t) = \overline{r}_a(t) + c\sqrt{\frac{\ln t}{n_a(t) + \epsilon}}, \tag{15}$$

where $\overline{r}_a(t)$ is the historical mean reward, $n_a(t)$ is the action-selection count, and c controls exploration.

The reward measures the improvement obtained after applying an action:

$$r_t = w_Q \Delta\mathcal{Q} + w_S \Delta\mathcal{S} + w_R \Delta\mathcal{R} + w_T \Delta\mathcal{T} - w_U \Delta\mathcal{U}. \tag{16}$$

The selected action is

$$a_t = \arg\max_{a \in \mathcal{A}} \mathrm{UCB}_a(t). \tag{17}$$

Thus, actions that consistently improve community quality, robustness, stability, and uncertainty receive greater preference, while exploration remains possible. The detailed stage-transition, incremental-growth, replacement, and meta-policy formulations are reported in the Appendix.

3.5 Joint Optimization

BioMetaEvo-GNN jointly optimizes community quality, spectral consistency, predictive uncertainty, immune-memory stability, temporal smoothness, and governance utility. The overall training objective is

$$\begin{aligned}\mathcal{L} = \mathcal{L}_{\mathrm{comm}} + \lambda_F \mathcal{L}_{\mathrm{Fourier}} + \lambda_U \mathcal{L}_{\mathrm{uncertainty}} \\ + \lambda_I \mathcal{L}_{\mathrm{immune}} + \lambda_T \mathcal{L}_{\mathrm{time}} + \lambda_G \mathcal{L}_{\mathrm{govern}},\end{aligned} \tag{18}$$

where $\mathcal{L}_{\mathrm{comm}}$ optimizes community assignments, $\mathcal{L}_{\mathrm{Fourier}}$ preserves consistency between spatial and spectral representations, and $\mathcal{L}_{\mathrm{uncertainty}}$ penalizes ambiguous predictions. The immune-memory term $\mathcal{L}_{\mathrm{immune}}$ constrains abrupt prototype deviations from historically reliable community states, while $\mathcal{L}_{\mathrm{time}}$ suppresses transient representation drift across iterations. The governance term is defined as $\mathcal{L}_{\mathrm{govern}} = -\mathbb{E}[r_t]$, encouraging actions that improve partition quality, robustness, stability, and uncertainty calibration. The coefficients λ_F, λ_U, λ_I, λ_T, and λ_G control the relative contribution of these objectives.

Training follows a closed-loop procedure. At iteration t, the GNN backbone produces node embeddings and preliminary community assignments. Bayesian Fourier learning estimates spectral uncertainty, and energy-time filtering attenuates transient representation fluctuations. Candidate community transitions are then evaluated using their uncertainty, structural consistency, energy reliability, and deviation from T-SILE memory. Reliable memory prototypes are updated, after which the UCB controller selects and applies an adaptive community-level action. The model parameters are jointly updated by minimizing Eq. equation 18.

The complete definitions of all loss terms are provided in the Appendix. The full training and inference workflow, output definitions, and computational complexity are reported in the Appendix.

4 Experiments

4.1 Experimental Setup

We evaluate BioMetaEvo-GNN on 14 real-world attributed graphs spanning citation, co-authorship, social, and co-purchase networks, together with 12 synthetic stochastic-block-model variants that control spectral gap, degree heterogeneity, overlapping membership, feature corruption, and heterophily. The comparison includes classical community detection, spectral and random-walk methods, deterministic and probabilistic GNNs, self-supervised graph representation models, deep graph clustering methods, differentiable pooling, and robust GNN baselines. Small and medium graphs use repeated 60/20/20 splits, whereas official splits are used for large-scale benchmarks. All neural models are trained with Adam and validation-based early stopping. Each experiment is repeated with five random seeds, and each perturbation setting uses five independently generated realizations.

Table 1: Assignment drift under general and spectral-risk perturbations. Lower values indicate greater partition stability.

Method	Assign. Drift↓	Spectral Drift↓
GCN	0.241	0.329
GAT	0.224	0.306
VGAE	0.198	0.278
GRACE	0.164	0.232
DMoN	0.158	0.224
Pro-GNN	0.127	0.186
GNNGuard	0.116	0.174
GRAND	0.104	0.159
BioMetaEvo-GNN	**0.067**	**0.101**

Models are selected on clean validation graphs and evaluated without retraining under six perturbation regimes: feature noise, feature masking, edge dropout, edge insertion, edge rewiring, and spectral-risk perturbation. We report NMI and ARI for clean community quality; robustness retention and assignment drift for perturbation stability; and ECE and uncertainty–drift correlation for uncertainty quality. Complete dataset statistics, preprocessing rules, baseline configurations, hyperparameter grids, threshold selection, perturbation definitions, and metric formulations are reported in the Appendix.

4.2 Clean Community Detection

Table 4.2 reports representative family-level results. BioMetaEvo-GNN consistently achieves the strongest NMI and ARI across citation, co-authorship, social, and co-purchase graphs. The gains are larger on social and co-purchase networks, where boundary ambiguity and structural noise are more pronounced.

table Representative family-level clean community detection results. Values are averaged over the datasets in each graph family. The complete comparison with all baselines and standard deviations is reported in the Appendix.

Method	Citation		Co-authorship		Social		Co-purchase	
	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
Leiden	0.386	0.293	0.416	0.316	0.350	0.263	0.405	0.300
Node2Vec	0.446	0.355	0.471	0.370	0.404	0.317	0.460	0.360
GAT	0.497	0.405	0.525	0.431	0.458	0.373	0.515	0.420
VGAE	0.506	0.416	0.534	0.442	0.467	0.385	0.526	0.431
GRACE	0.523	0.432	0.552	0.461	0.486	0.402	0.544	0.451
DMoN	0.531	0.442	0.562	0.471	0.495	0.414	0.554	0.463
GNNGuard	0.525	0.435	0.557	0.467	0.496	0.417	0.556	0.466
GRAND	0.529	0.440	0.561	0.470	0.500	0.421	0.559	0.470
BioMetaEvo-GNN	**0.558**	**0.473**	**0.592**	**0.503**	**0.536**	**0.455**	**0.594**	**0.507**

4.3 Denoising Analysis Under Perturbation Dynamics

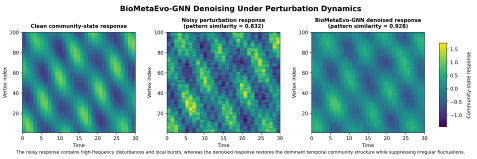

Figure 2: Denoising analysis under perturbation dynamics.

Figure 2 visualizes the denoising behaviour of BioMetaEvo-GNN under perturbation dynamics. The clean panel provides the reference community-state trajectory, whereas the noisy panel contains high-frequency disturbances and localized burst perturbations. After denoising, the recovered response preserves the dominant temporal community pattern while substantially reducing irregular

fluctuations. This result supports the robustness of BioMetaEvo-GNN in suppressing perturbation-induced instability without destroying the main structural evolution of the community signal.

4.4 Perturbation Robustness

Figure 3 and 1summarizes structural, attribute, and spectral-risk robustness. BioMetaEvo-GNN preserves 93.3% of clean NMI under structural perturbations and 92.2% under attribute perturbations, while reducing assignment drift to 0.067. Under spectral-risk perturbation, which preferentially modifies fragile boundary regions, its drift remains 0.101, compared with 0.159 for GRAND and 0.174 for GNNGuard. These results show that the improvement is concentrated on the failure modes targeted by spectral uncertainty estimation and memory-based transition control.

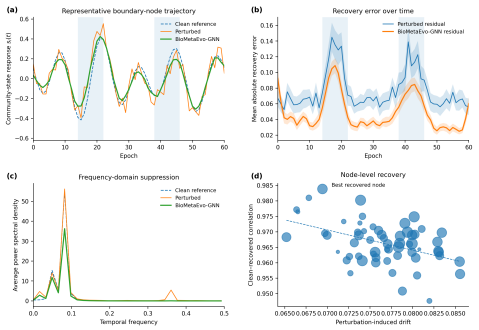

Figure 3: Perturbation-recovery analysis of BioMetaEvo-GNN.

Figure 3 provides a compact multi-perspective analysis of the perturbation-recovery behaviour of BioMetaEvo-GNN. BioMetaEvo-GNN achieves high clean–recovered correlation even under noticeable perturbation-induced drift.

4.5 Mechanism and Minimal-Design Validation

We compare the complete framework with progressively enhanced alternatives to determine whether robust propagation, calibration, or Bayesian uncertainty alone is sufficient. As shown in Table 3 and 2, robust propagation substantially improves retention and drift, whereas calibration mainly lowers ECE. Bayesian uncertainty provides further improvements but still lacks explicit historical memory and community-level transition control. BioMetaEvo-GNN achieves the strongest joint performance across accuracy, robustness, stability, calibration, and drift.

Table 2: Minimal-design comparison in community quality and robustness.

Model	**NMI↑**	**ARI↑**	**Retention↑**
GNN backbone only	0.485	0.391	0.750
Robust GNN	0.529	0.440	0.878
Robust GNN + calibration	0.531	0.444	0.881
Robust GNN + Bayesian uncertainty	0.536	0.454	0.887
BioMetaEvo-GNN	**0.558**	**0.473**	**0.922**

Table 3: Minimal-design comparison in stability, calibration, and assignment drift. Lower ECE and drift values are better.

Model	**Stability↑**	**ECE↓**	**Drift↓**
GNN backbone only	0.731	0.124	0.241
Robust GNN	0.835	0.096	0.104
Robust GNN + calibration	0.842	0.076	0.101
Robust GNN + Bayesian uncertainty	0.852	0.081	0.091
BioMetaEvo-GNN	**0.891**	**0.052**	**0.067**

Parameter sensitivity. We further examine whether BioMetaEvo-GNN depends on narrowly tuned memory parameters. Figure 4 visualizes the sensitivity of the Energy-Time coefficient α and the T-SILE memory-retention coefficient β. The best observed configuration is obtained at $\alpha = \beta = 0.7$, yielding a stability score of 0.891. Minimal-design comparison with progressively enhanced robust GNN alternatives.

The Bayesian Fourier analysis provides direct evidence for the uncertainty mechanism. Fourier learning increases NMI from 0.485 to 0.512, Bayesian inference further increases it to 0.536, and ECE decreases from 0.124 to 0.081. The complete framework reaches an ECE of 0.052 and an uncertainty–drift correlation of 0.681, indicating that uncertainty is both better calibrated and more strongly aligned with actual assignment instability. The extended immune-response, UCB-governance, stage-transition, ablation, pairwise interaction, transfer, incremental-growth, scalability, failure-case, and significance analyses are reported in the Appendix.

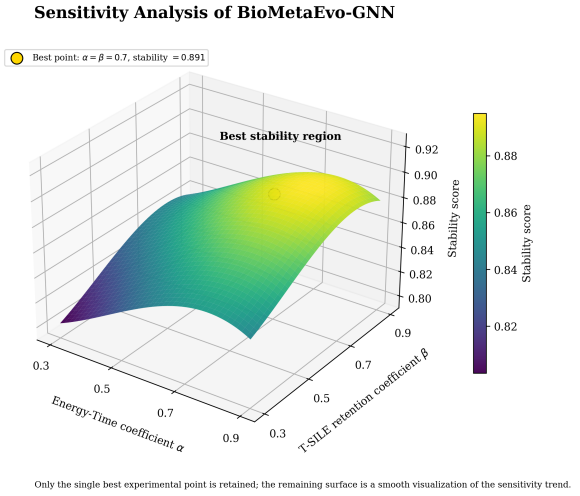

Figure 4: Sensitivity analysis of the Energy-Time coefficient α and T-SILE memory-retention coefficient β. The highlighted point denotes the best observed configuration at $\alpha = \beta = 0.7$, yielding a stability score of 0.891.

As shown in Figure 4, stability improves when both coefficients increase from low values toward the balanced memory configuration. Small values provide insufficient temporal smoothing and historical retention, whereas excessively large values slightly reduce stability by delaying adaptation to genuine community changes. The broad high-stability region around the optimum indicates that BioMetaEvo-GNN does not rely on a narrowly tuned parameter combination.

5 Discussion

From a theoretical perspective, BioMetaEvo-GNN formulates community detection as a multi-objective reliability problem rather than a partition-quality objective alone. It jointly considers modularity, structural consistency, robustness, temporal stability, uncertainty, and governance utility, providing a unified basis for Bayesian Fourier uncertainty estimation, immune-memory transition control, and UCB-based adaptive governance. The central objective is to suppress unstable community transitions while preserving accurate and robust partitions.

This perspective extends recent advances in neural and attributed community detection Kojaku et al. (2024); Liu et al. (2024a;b), robust GNN learning under feature, structural, semantic, and label perturbations Gosch et al. (2023); Abbahaddou et al. (2024); Wang et al. (2024c); Jin & Zhu (2025); Goel et al. (2026), and uncertainty-aware graph learning emphasizing reliable confidence estimation and uncertainty-guided refinement Bergna et al. (2025). Dynamic graph studies further highlight the importance of stability and transferability under evolving structures and domains Huang et al. (2023b); Zheng et al. (2025). BioMetaEvo-GNN connects these directions within a closed-loop

reliability framework: Bayesian Fourier Learning detects spectral instability and drift-prone boundaries, Energy-Time Filtering suppresses transient fluctuations, T-cell Defense and T-SILE memory reject unreliable transitions while preserving stable evidence, and UCB governance adaptively selects community-level actions according to uncertainty, reward, and stability.

The experimental results show consistent improvements in NMI, ARI, robustness retention, calibration, and assignment stability, supported by perturbation, uncertainty, immune-response, ablation, governance, scalability, and failure analyses. Residual errors remain concentrated in overlapping communities, low-homophily regions, noisy bridge edges, sparse attributes, and ambiguous boundary nodes, consistent with evidence that graph reliability depends strongly on structural, semantic, and explanatory factors Deng (2025); Jiang et al. (2024). BioMetaEvo-GNN should therefore be viewed as a reliability-oriented framework that improves the resilience, interpretability, and controllability of community detection under noisy and dynamic graph conditions rather than eliminating all sources of community-detection error.

6 Limitations and Future Work

BioMetaEvo-GNN still has several limitations. The current framework primarily operates in batch settings and requires further adaptation to fully streaming graphs in which nodes, edges, attributes, and communities evolve continuously Barros et al. (2021); He et al. (2024); Li et al. (2023); Yang et al. (2025). Its Bayesian Fourier module could also be extended with graph transformers, foundation models, wavelets, localized spectral bases, and multi-scale structural encodings to better capture long-range dependencies and fine-grained community boundaries Khemani et al. (2024); Wu et al. (2025a); Franks et al. (2025).

The immune-defense and T-SILE memory mechanisms should further be evaluated on heterogeneous, signed, multiplex, and hypergraph settings involving multiple relation types and higher-order interactions Kim et al. (2024); Yang & Xu (2025); Melton et al. (2023); Jiao et al. (2025). Future work may also incorporate conformal prediction, improved uncertainty calibration, stochastic uncertainty modeling, and uncertainty-aware structure learning to strengthen reliability under noisy and out-of-distribution conditions Cai et al. (2023); Liu et al. (2025); Zhou et al. (2024); Dai et al. (2024).

Finally, broader trustworthiness requires formal verification, privacy preservation, fairness analysis, explainable governance, and standardized robustness evaluation Guo et al. (2025); Zhang et al. (2024); Zheng et al. (2025); Gu et al. (2021); Xiao et al. (2024). Validation across fraud detection, scientific collaboration, biological networks, recommender systems, and social-risk analysis will further determine whether BioMetaEvo-GNN can provide reliable and interpretable community-level decisions in practical dynamic environments.

7 Conclusion

This paper introduced **BioMetaEvo-GNN**, a bio-inspired and meta-evolutionary framework that formulates community detection as a closed-loop reliability process rather than static partition optimization. By integrating Bayesian Fourier uncertainty modeling, Energy-Time Filtering, immune-inspired transition control and memory, and UCB-based adaptive governance, the framework identifies unstable spectral patterns, suppresses unreliable community transitions, preserves trustworthy historical evidence, and adapts community-level decisions under noisy and dynamic graph conditions.

Experiments on real-world and synthetic graphs demonstrated consistent improvements in clustering quality, perturbation robustness, uncertainty calibration, assignment stability, and cross-dataset transfer. Ablation and diagnostic analyses further showed that these gains arise from the coordinated interaction of spectral diagnosis, memory-based control, and adaptive governance. Residual errors remain concentrated around overlapping communities, low-homophily regions, noisy bridge edges, sparse attributes, and ambiguous boundary nodes. Future work will extend BioMetaEvo-GNN to streaming, heterogeneous, signed, multiplex, and hypergraph settings while exploring localized spectral representations and more advanced governance strategies.

AI USE STATEMENT

Generative AI tools were used to assist with language editing, manuscript organization, LaTeX formatting, and improving the clarity and conciseness of the presentation. All technical formulations, methodological descriptions, experimental analyses, tables, figures, claims, and citations were from the authors. Generative AI was not relied upon as an independent source of scientific evidence or experimental validation. The authors take full responsibility for the accuracy, originality, and final content of the manuscript, including any text or artifacts prepared with the assistance of generative AI.

ETHICS STATEMENT

This work focuses on methodological research in dynamic community detection and does not involve human-subject experiments, clinical interventions, or direct interaction with individuals. The experiments use synthetic benchmarks and established real-world network datasets for aggregate graph analysis. We do not attempt to identify individuals, infer sensitive personal attributes, or make decisions about specific users represented in these networks. Nevertheless, community-detection and uncertainty-estimation techniques could potentially be applied to sensitive social or behavioral data. Such deployment should therefore consider privacy, data governance, potential biases in the underlying graph, and the consequences of erroneous or overconfident community assignments. BayesMem-CD is intended as a research framework for studying uncertainty-aware temporal graph structure rather than for automated high-stakes decision making.

REPRODUCIBILITY STATEMENT

We provide the core formulation of BayesMem-CD, including the temporal-memory mechanism, Bayesian community inference, uncertainty estimation, and learning objective, in the main paper. The Supplementary Material contains the complete graph formulation, encoder details, variational derivations, optimization procedure, and training pseudocode. It also provides full dataset statistics and preprocessing procedures, baseline configurations, hyperparameter search ranges, implementation settings, computational analysis, and additional experimental results. All reported results are based on five independent runs with different random seeds, with statistical testing procedures described in the experimental protocol. Hardware and software environments, including the major library versions used in the experiments, are also documented in the Supplementary Material. Together, these materials are intended to provide the information necessary to reproduce the model implementation, training procedure, and empirical evaluation.

REFERENCES

Yassine Abbahaddou, Sofiane Ennadir, Johannes F. Lutzeyer, Michalis Vazirgiannis, and Henrik Boström. Bounding the expected robustness of graph neural networks subject to node feature attacks. In *International Conference on Learning Representations*, 2024.

Claudio D. T. Barros, Matheus R. F. Mendonça, Alex B. Vieira, and Artur Ziviani. A survey on embedding dynamic graphs. *ACM Computing Surveys*, 55(1):1–37, 2021. doi: 10.1145/3483595.

Richard Bergna, Sergio Calvo-Ordoñez, Felix L. Opolka, Pietro Liò, and José Miguel Hernández-Lobato. Uncertainty modeling in graph neural networks via stochastic differential equations. In *International Conference on Learning Representations*, 2025.

Xin Cai et al. A graph convolutional fusion model for community detection in multiplex networks. *Data Mining and Knowledge Discovery*, 2023. doi: 10.1007/s10618-023-00932-w.

Shuyi Chen, Kaize Ding, and Shixiang Zhu. Uncertainty-aware robust learning on noisy graphs. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025.

Enyan Dai, Tianxiang Zhao, Huaisheng Zhu, Junjie Xu, Zhimeng Guo, Hui Liu, Jiliang Tang, and Suhang Wang. A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. *Machine Intelligence Research*, 21(6):1011–1061, 2024.

Minghua Deng. Early detection of malicious accounts on social platforms based on temporal graph feature learning. In *Proceedings of the 2025 8th International Conference on Computer Information Science and Artificial Intelligence*, pp. 1320–1328, 2025.

Billy Joe Franks, Moshe Eliasof, Semih Cantürk, Guy Wolf, Carola-Bibiane Schönlieb, Sophie Fellenz, and Marius Kloft. Towards graph foundation models: A study on the generalization of positional and structural encodings. *Transactions on Machine Learning Research*, 2025.

Jaidev Goel, Pablo Moriano, Ramakrishnan Kannan, and Yulia R. Gel. Community detection robustness of graph neural networks. *Physical Review Research*, 8(2):023123, 2026. doi: 10.1103/v6nq-hcm8.

Lukas Gosch, Daniel Sturm, Simon Geisler, and Stephan Günnemann. Revisiting robustness in graph machine learning. In *International Conference on Learning Representations*, 2023.

Xiaotong Gu, Zehong Cao, Alireza Jolfaei, Peng Xu, Dongrui Wu, Tzyy-Ping Jung, and Chin-Teng Lin. Eeg-based brain-computer interfaces (bcis): A survey of recent studies on signal sensing technologies and computational intelligence approaches and their applications. *IEEE/ACM transactions on computational biology and bioinformatics*, 18(5):1645–1666, 2021.

Feng Guan et al. Graph neural networks: A survey on the links between privacy and security. *Artificial Intelligence Review*, 57, 2024. doi: 10.1007/s10462-023-10656-4.

Fei Guo, Renchu Guan, Yaohang Li, Qi Liu, Xiaowo Wang, Can Yang, and Jianxin Wang. Foundation models in bioinformatics. *National science review*, 12(4):nwaf028, 2025.

Shen Han, Zhiyao Zhou, Jiawei Chen, Zhezheng Hao, Sheng Zhou, Gang Wang, Yan Feng, Chun Chen, and Can Wang. Uncertainty-aware graph structure learning. In *Proceedings of the ACM Web Conference*, 2025. doi: 10.1145/3696410.3714927.

Jian He, Liang Qiao, Junlin Li, Zhouping Lu, Jiafang Fu, Xiangchun Chen, Xiulin Zhang, and Xulin Hu. Advanced strategies for 3d-printed neural scaffolds: Materials, structure, and nerve remodeling. *Bio-Design and Manufacturing*, 7(5):747–770, 2024.

Zhichao Hou, Minhua Lin, MohamadAli Torkamani, Suhang Wang, and Xiaorui Liu. Adversarial robustness in graph neural networks: Recent advances and new frontier. In *2024 IEEE 11th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 1–2. IEEE, 2024.

Kexin Huang, Ying Jin, Emmanuel Candès, and Jure Leskovec. Uncertainty quantification over graph with conformalized graph neural networks. In *Advances in Neural Information Processing Systems*, volume 36, pp. 26699–26721, 2023a.

Shenyang Huang, Farimah Poursafaei, Jacob Danovitch, Matthias Fey, Weihua Hu, Emanuele Rossi, Jure Leskovec, Michael Bronstein, Guillaume Rabusseau, and Reihaneh Rabbany. Temporal graph benchmark for machine learning on temporal graphs. In *Advances in Neural Information Processing Systems*, volume 36, 2023b.

Jinhao Jiang et al. Quantifying uncertainty in graph neural network explanations. *Frontiers in Big Data*, 7, 2024. doi: 10.3389/fdata.2024.1392662.

Ziheng Jiao, Hongyuan Zhang, and Xuelong Li. Cnn2gnn: How to bridge cnn with gnn. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Ying Jin and Xiaojin Zhu. A systematic study and analysis of graph neural networks under noise. *ACM Transactions on Knowledge Discovery from Data*, 2025. doi: 10.1145/3733605.

Bharti Khemani, Shruti Patil, Ketan Kotecha, and Sudeep Tanwar. A review of graph neural networks: Concepts, architectures, techniques, challenges, datasets, applications, and future directions. *Journal of Big Data*, 11(18), 2024. doi: 10.1186/s40537-023-00876-4.

Sunwoo Kim, Soo Yong Lee, Yue Gao, Alessia Antelmi, Mirko Polato, and Kijung Shin. A survey on hypergraph neural networks: An in-depth and step-by-step guide. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6534–6544, 2024. doi: 10.1145/3637528.3671457.

Sadamori Kojaku, Filippo Radicchi, Yong-Yeol Ahn, and Santo Fortunato. Network community detection via neural embeddings. *Nature Communications*, 15(1):9446, 2024. doi: 10.1038/s41467-024-52355-w.

Lingkai Kong, Haotian Sun, Yuchen Zhuang, Haorui Wang, Wenhao Mu, and Chao Zhang. Two birds with one stone: Enhancing uncertainty quantification and interpretability with graph functional neural process. In *International Conference on Artificial Intelligence and Statistics*, pp. 4582–4590. PMLR, 2024.

Xianghua Li, Xiyuan Zhen, Xin Qi, Huichun Han, Long Zhang, and Zhen Han. Dynamic community detection based on graph convolutional networks and contrastive learning. *Chaos, Solitons & Fractals*, 176:114157, 2023. doi: 10.1016/j.chaos.2023.114157.

Chang Liu, Yuwen Yang, Yue Ding, Hongtao Lu, Wenqing Lin, Ziming Wu, and Wendong Bi. Dag: Deep adaptive and generative k-free community detection on attributed graphs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024a. doi: 10.1145/3637528.3671615.

Chunjiang Liu, Yikun Han, Haiyun Xu, Shihan Yang, Kaidi Wang, and Yongye Su. A community detection and graph-neural-network-based link prediction approach for scientific literature. *Mathematics*, 12(3):369, 2024b. doi: 10.3390/math12030369.

Ling Liu, Fangyan Wang, Xiang Chen, Lanjie Liu, Yibing Wang, Jingyi Bei, Lanjie Lei, Zhangwei Zhao, and Chengxuan Tang. Designing multifunctional microneedles in biomedical engineering: Materials, methods, and applications. *International Journal of Nanomedicine*, pp. 8693–8728, 2025.

Yanfei Ma, Daozheng Qu, and Mykhailo Pyrozhenko. Bio-regnet: A meta-homeostatic bayesian neural network framework integrating treg-inspired immunoregulation and autophagic optimization for adaptive community detection and stable intelligence. *Biomimetics*, 11(1), 2026. ISSN 2313-7673. doi: 10.3390/biomimetics11010048. URL https://www.mdpi.com/2313-7673/11/1/48.

Jeremy Melton et al. Multiplex graph neural network for heterogeneous graphs. *IEEE Access*, 2023. doi: 10.1109/ACCESS.2023.3266823.

Puja Trivedi, Mark Heimann, Rushil Anirudh, Danai Koutra, and Jayaraman J. Thiagarajan. Accurate and scalable estimation of epistemic uncertainty for graph neural networks. In *International Conference on Learning Representations*, 2024.

Fangxin Wang, Yuqing Liu, Kay Liu, Yibo Wang, Sourav Medya, and Philip S. Yu. Uncertainty in graph neural networks: A survey. *Transactions on Machine Learning Research*, 2024a.

Min Wang, Hao Yang, Jincai Huang, and Qing Cheng. Moderate message passing improves calibration: A universal way to mitigate confidence bias in graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024b.

Zhonghao Wang, Danyu Sun, Sheng Zhou, Haobo Wang, Jiapei Fan, Longtao Huang, and Jiajun Bu. Noisygl: A comprehensive benchmark for graph neural networks under label noise. In *Advances in Neural Information Processing Systems*, volume 37, 2024c.

Lingfei Wu et al. Graph foundation models: Challenges, methods, and applications. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2025a. doi: 10.1145/3711896.3736568.

Tao Wu, Canyixing Cui, Xingping Xian, Shaojie Qiao, Chao Wang, Lin Yuan, and Shui Yu. Understanding the robustness of graph neural networks against adversarial attacks. *Knowledge-Based Systems*, 323:113714, 2025b.

Heng Xiao, Xin Wang, Jun Wang, Jing-Ye Cai, Jian-Hua Deng, Jing-Ke Yan, and Yi-Dong Tang. Single image super-resolution with denoising diffusion gans. *Scientific Reports*, 14(1):4272, 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-52370-3. URL https://doi.org/10.1038/s41598-024-52370-3.

Hao Yang, Min Wang, Qi Wang, Mingrui Lao, and Yun Zhou. Balanced confidence calibration for graph neural networks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3747–3757, 2024. doi: 10.1145/3637528.3671741.

Mu-Rong Yang and Xin-Jian Xu. Recent advances in hypergraph neural networks. *Journal of Systems Science and Complexity*, 2025. doi: 10.1007/s40305-025-00630-y.

Shiyu Yang, Guanyu Ding, Zifan Chen, and Jie Si Yang. Gart: Graph neural network-based adaptive and robust task scheduler for heterogeneous distributed computing. *IEEE Access*, 13:200196–200216, 2025.

Yuhang Zhang et al. A survey on privacy in graph neural networks: Attacks, preservation, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 36(12):7497–7515, 2024. doi: 10.1109/TKDE.2024.3454328.

Yanping Zheng, Lu Yi, and Zhewei Wei. A survey of dynamic graph neural networks. *Frontiers of Computer Science*, 19:196323, 2025. doi: 10.1007/s11704-024-3853-2.

Andy Zhou, Xiaojun Xu, Ramesh Raghunathan, Alok Lal, Xinze Guan, Bin Yu, and Bo Li. Knowgraph: Knowledge-enabled anomaly detection via logical reasoning on graph data. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 168–182, 2024.