BAYESMEM-CD: LEARNING STABLE DYNAMIC GRAPH REPRESENTATIONS WITH BAYESIAN TEMPORAL MEMORY

Anonymous authors
Paper under double-blind review

ABSTRACT

Dynamic community detection is confronted with a fundamental ambiguity: an observed structural change may be indicative of genuine community evolution or merely transient graph noise. This issue is typically resolved by existing methods through deterministic temporal smoothing or snapshot-wise inference, which can either suppress meaningful structural transitions or generate unstable assignments. Additionally, deterministic temporal representations offer restricted insight into the periods during which historical evidence is uncertain. We introduce BayesMem-CD, an uncertainty-aware framework that is based on Bayesian temporal memory. This framework conceptualizes community evolution as a probabilistic temporal inference, rather than a series of deterministic clustering decisions. BayesMem-CD quantifies epistemic uncertainty and regulates temporal adaptation based on the reliability of observed structure by maintaining long-range structural evidence through graph memory and representing community assignments probabilistically. Spurious assignment changes are further discouraged by short-term temporal regularization, which does not preclude persistent evidence from driving genuine community transitions. Experiments conducted on synthetic and real-world dynamic networks have shown that BayesMem-CD is capable of achieving robust community recovery and temporal stability, as well as better-calibrated uncertainty under structural perturbations. The complementary benefits of probabilistic inference, temporal memory, and short-term consistency are further demonstrated by ablation and sensitivity analyzes. These findings indicate that a principled approach to distinguishing structural evolution from transitory variation in dynamic graphs is provided by explicitly modeling uncertainty in temporal evidence.

1 INTRODUCTION

Community detection in dynamic graphs aims to identify cohesive groups of nodes while tracking how these groups evolve over time. It is relevant to applications such as social networks, communication infrastructures, and biological interactions. Unlike static community detection, the dynamic setting requires models to preserve temporal consistency across related graph snapshots while remaining responsive to genuine structural changes Kazemi et al. (2020); Pareja et al. (2020). This is challenging because real-world temporal graphs are often noisy and non-stationary, so small changes in observed connectivity may produce unstable community assignments even when the underlying structure changes only marginally Rossi et al. (2020). At the same time, excessive temporal smoothing can suppress meaningful community evolution Sankar et al. (2020).

Recent temporal graph learning methods have substantially improved the modeling of evolving graph structure. EvolveGCN Pareja et al. (2020), DySAT Sankar et al. (2020), TGAT Xu et al. (2020), and Temporal Graph Networks Rossi et al. (2020) capture temporal dependencies through evolving parameters, temporal attention, continuous-time representations, or memory-based mechanisms. However, these methods are primarily designed for representation learning and downstream prediction tasks rather than community-level inference. Dynamic community detection methods address the task more directly, but temporal consistency is commonly imposed through deterministic clustering, embedding smoothing, or regularization across consecutive snapshots.

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.

A limitation of these approaches is that they provide little information about whether a community assignment is reliable. This becomes important near community boundaries or during periods of rapid structural evolution, where transient perturbations and genuine transitions may induce similar changes in the observed graph. Uncertainty-aware learning provides a natural mechanism for characterizing this ambiguity. Bayesian neural networks represent model parameters probabilistically and enable estimation of epistemic uncertainty, while recent calibration studies show that neural models may remain poorly calibrated despite strong predictive performance Minderer et al. (2021). Yet uncertainty modeling has received comparatively limited attention in dynamic community detection.

We introduce **BayesMem-CD**, a Bayesian temporal memory framework for uncertainty-aware dynamic community detection. Rather than producing only deterministic community labels, BayesMem-CD represents community assignments probabilistically and estimates epistemic uncertainty through Bayesian inference. A graph memory mechanism retains structural information over longer temporal horizons, while Markov regularization constrains excessive changes between consecutive community distributions. Snapshot-level graph representations are therefore combined with historical information before probabilistic community inference, allowing the model to reduce fluctuations caused by transient structural perturbations while remaining responsive to persistent changes in community structure.

We evaluate BayesMem-CD on synthetic and real-world temporal networks against representative static, dynamic, and temporal graph-learning baselines. The experiments examine community recovery, structural quality, temporal stability, uncertainty calibration, robustness to perturbations, and computational efficiency. The results show that combining Bayesian inference with temporal memory provides stable community predictions while preserving sensitivity to meaningful structural evolution, and ablation analyses further examine the roles of uncertainty modeling, memory retention, and temporal regularization.

2 Related Work

Dynamic graph learning has evolved from snapshot-based methods toward temporal message passing, continuous-time modeling, and memory-based architectures Skarding et al. (2021); Barros et al. (2023); Yang et al. (2024). Representative approaches such as EvolveGCN, DySAT, TGAT, and TGN model temporal dependencies through evolving parameters, temporal attention, or persistent memory Pareja et al. (2020); Sankar et al. (2020); Xu et al. (2020); Rossi et al. (2020), while more recent work has addressed scalability, distribution shifts, and temporally faithful evaluation Zheng et al. (2022; 2023); Yuan et al. (2023); Zhu et al. (2024); Huang et al. (2023; 2024); Su & Wu (2025). However, these methods are primarily developed for representation learning, node prediction, or link prediction rather than community-level inference. In parallel, learning-based community detection has incorporated graph autoencoders, attention mechanisms, attributed representations, and overlapping-community objectives He et al. (2022a;b); Guo et al. (2023); Cheng et al. (2023); Zhang et al. (2024); Xie et al. (2025). Although these approaches improve community representation and recovery, many operate in static or weakly temporal settings, while temporal consistency is commonly introduced through deterministic smoothing or regularization. Uncertainty estimation provides a complementary perspective: Bayesian and uncertainty-aware graph models characterize uncertainty arising from model parameters, graph topology, and message passing Wang et al. (2024), and graph neural networks have been shown to exhibit substantial miscalibration despite strong predictive performance Liu et al. (2022); Minderer et al. (2021). Recent graph-specific uncertainty methods further emphasize the importance of explicitly characterizing prediction reliability Yu et al. (2025). Nevertheless, uncertainty modeling remains comparatively underexplored in dynamic community detection, where transient structural perturbations and genuine community transitions may lead to similar changes in observed assignments. BayesMem-CD connects these directions by combining probabilistic community inference with long-range temporal memory and short-term regularization, providing a unified treatment of community structure, temporal consistency, and uncertainty in evolving graphs.

3 METHOD

3.1 PROBLEM FORMULATION

We consider a dynamic graph represented by a sequence of snapshots

$$\mathcal{G} = \{G_t\}_{t=1}^{T}, \qquad G_t = (V_t, E_t, X_t), \tag{1}$$

where V_t and E_t denote the node and edge sets at time t, and X_t is the corresponding node feature matrix. Using a common node indexing scheme across snapshots, we denote the adjacency matrix by A_t.

For each snapshot, the goal is to infer a probabilistic community assignment matrix

$$P_t \in [0,1]^{N\times K}, \qquad \sum_{k=1}^{K} P_t(i,k) = 1, \tag{2}$$

where K is the number of communities. A hard assignment is obtained as

$$\hat{y}_{t,i} = \arg\max_{k} P_t(i,k). \tag{3}$$

Unlike deterministic community detection, we treat P_t itself as the primary prediction, allowing community uncertainty to be represented explicitly. The overall task is therefore

$$\Phi : \{A_t, X_t\}_{t=1}^{T} \mapsto \{P_t, U_t\}_{t=1}^{T}, \tag{4}$$

where U_t denotes uncertainty associated with community assignments.

3.2 BAYESIAN TEMPORAL MEMORY

The proposed BayesMem-CD framework combines snapshot-level structural encoding, long-horizon graph memory, Bayesian community inference, and short-horizon temporal regularization. Its architecture is illustrated in Fig. 1.

At each time step, a topology-aware graph encoder extracts a structural representation

$$H_t = f_\theta(A_t, X_t) \in \mathbb{R}^{N\times h}. \tag{5}$$

The encoder can be instantiated using standard message-passing or attention-based graph layers. Detailed encoder formulations are provided in the Supplementary Material.

To incorporate historical structure, the current representation is fused with the previous memory state:

$$\widetilde{H}_t = \tanh(W_f[H_t \| M_{t-1}] + b_f), \qquad t \geq 2, \tag{6}$$

with $\widetilde{H}_1 = H_1$ and $M_1 = H_1$. The memory is then updated as

$$M_t = \gamma M_{t-1} + (1-\gamma)\widetilde{H}_t, \qquad \gamma \in [0,1]. \tag{7}$$

This recursive update preserves long-range structural information while allowing persistent changes in the current graph to gradually modify the stored representation.

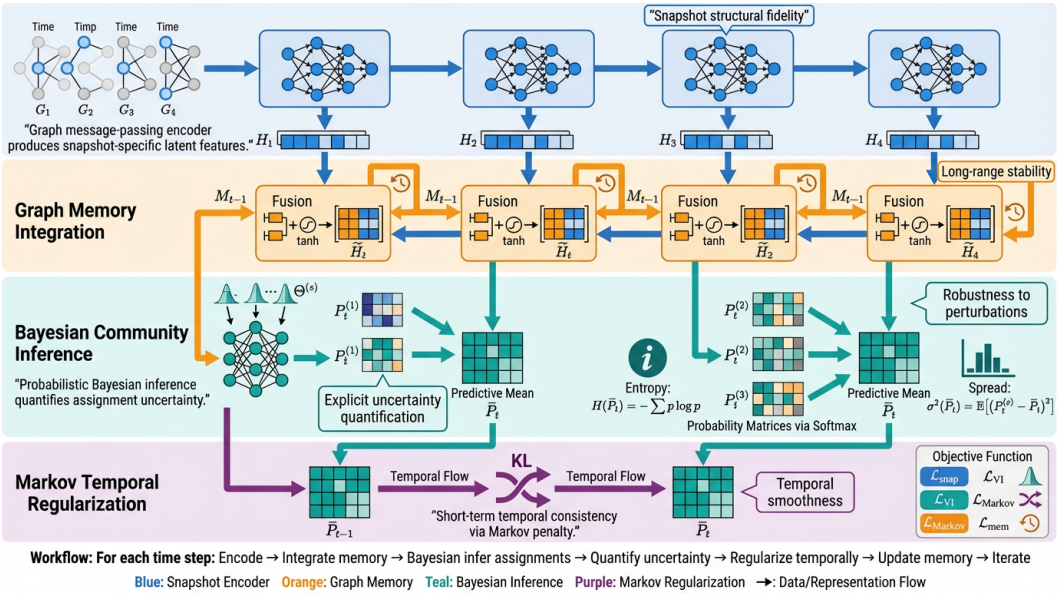


Figure 1: Overview of BayesMem-CD. Each graph snapshot is encoded into a structural representation, integrated with historical graph memory, and passed to Bayesian community inference. Posterior sampling yields probabilistic community assignments and uncertainty estimates, while temporal regularization constrains excessive assignment changes across consecutive snapshots.

3.3 Bayesian Community Inference

Let $\Theta = \{W, b\}$ denote the parameters of the community predictor. Rather than learning a deterministic parameter estimate, we place a Gaussian prior

$$p(\Theta) = \mathcal{N}(0, \sigma_0^2 I) \tag{8}$$

and approximate the intractable posterior using

$$q_\phi(\Theta) = \mathcal{N}\left(\mu_\phi, \mathrm{diag}(\sigma_\phi^2)\right). \tag{9}$$

Posterior samples are obtained through reparameterization:

$$\Theta^{(s)} = \mu_\phi + \sigma_\phi \odot \epsilon^{(s)}, \qquad \epsilon^{(s)} \sim \mathcal{N}(0, I), \tag{10}$$

for $s = 1, \ldots, S$. Given a sampled predictor, the corresponding community distribution is

$$P_t^{(s)} = \mathrm{softmax}\left(\widetilde{H}_t W^{(s)} + \mathbf{1} b^{(s)\top}\right). \tag{11}$$

The predictive community distribution is obtained by Monte Carlo averaging:

$$\bar{P}_t = \frac{1}{S} \sum_{s=1}^{S} P_t^{(s)}. \tag{12}$$

We quantify assignment uncertainty using predictive entropy

$$U_t(i) = -\sum_{k=1}^{K} \bar{P}_t(i,k) \log \bar{P}_t(i,k), \tag{13}$$

and posterior predictive variance

$$V_t(i) = \sum_{k=1}^{K} \frac{1}{S} \sum_{s=1}^{S} \left(P_t^{(s)}(i,k) - \bar{P}_t(i,k)\right)^2. \tag{14}$$

Entropy measures ambiguity in the averaged community distribution, whereas predictive variance captures disagreement across posterior samples.

The posterior parameters are learned using the variational objective

$$\mathcal{L}_{\mathrm{VI}}^{(t)} = \mathrm{KL}\left(q_\phi(\Theta)\|p(\Theta)\right) - \mathbb{E}_{q_\phi(\Theta)}\left[\log p(Y_t \mid \widetilde{H}_t, \Theta)\right]. \tag{15}$$

The complete ELBO derivation and the closed-form Gaussian KL term are provided in the Supplementary Material.

3.4 Structural and Temporal Objectives

To ensure that probabilistic assignments remain consistent with the observed graph structure, we construct the soft co-membership matrix

$$C_t = \bar{P}_t \bar{P}_t^\top. \tag{16}$$

The snapshot-level objective is denoted by $\mathcal{L}_{\mathrm{snap}}^{(t)}$. A reconstruction-based formulation is

$$\mathcal{L}_{\mathrm{snap}}^{(t)} = \|A_t - C_t\|_F^2. \tag{17}$$

Additional derivations and the corresponding modularity-based formulation are provided in the Supplementary Material.

Graph memory captures long-term dependencies, but consecutive community predictions may still vary excessively under local perturbations. We therefore regularize the predictive distributions directly:

$$\mathcal{L}_{\mathrm{M}}^{(t)} = \mathrm{KL}\left(\bar{P}_{t-1}\|\bar{P}_t\right), \qquad t \geq 2. \tag{18}$$

This term acts on community assignments rather than latent embeddings, providing short-horizon consistency at the prediction level.

An auxiliary memory-alignment objective is defined as

$$\mathcal{L}_{\mathrm{mem}}^{(t)} = \|M_t - \widetilde{H}_t\|_F^2. \tag{19}$$

The complete objective is

$$\begin{aligned}\mathcal{L}_{\mathrm{total}} = &\sum_{t=1}^{T} \mathcal{L}_{\mathrm{snap}}^{(t)} + \lambda_{\mathrm{VI}} \sum_{t=1}^{T} \mathcal{L}_{\mathrm{VI}}^{(t)} \\ &+ \lambda_{\mathrm{M}} \sum_{t=2}^{T} \mathcal{L}_{\mathrm{M}}^{(t)} + \lambda_{\mathrm{mem}} \sum_{t=2}^{T} \mathcal{L}_{\mathrm{mem}}^{(t)}.\end{aligned} \tag{20}$$

All model components are optimized jointly through backpropagation using the reparameterized posterior samples. Additional optimization details, gradient expressions, and the complete training and inference procedure are provided in the Supplementary Material.

4 Experiments

4.1 Experimental Setup

We evaluate BayesMem-CD on four synthetic dynamic-network benchmarks (DSBM, Dynamic LFR, Temporal GN, and Event-Driven Synthetic) and four real-world temporal networks (Enron, DBLP, Reddit, and Wikipedia). Synthetic datasets provide reference community labels, whereas the real-world networks are evaluated without assuming ground-truth communities. The evaluation covers community recovery, structural quality, temporal stability, and uncertainty calibration. We compare BayesMem-CD with classical and neural dynamic-graph baselines, including Louvain, Leiden, FacetNet, DynGEM, DySAT, EvolveGCN, TGN, and TGAT. Results are averaged over five independent runs. Complete dataset statistics, preprocessing, hyperparameter configurations, implementation details, and additional metrics are provided in the Supplementary Material.

Table 1: Summary of community recovery, calibration, structural quality, and temporal stability. Synthetic and real-world values are averaged over four datasets. Higher is better except ECE (↓).

	Synthetic			Real World	
Method	**NMI ↑**	**ARI ↑**	**ECE ↓**	**Mod. ↑**	**CSI ↑**
DySAT	0.703	0.666	0.073	**0.538**	0.723
EvolveGCN	0.702	0.689	0.067	0.535	0.729
TGN	–	–	–	0.530	0.765
TGAT	–	–	–	**0.538**	0.768
BayesMem-CD	**0.714**	**0.703**	**0.051**	0.537	**0.769**

Table 2: Component ablation. Higher is better for modularity, NMI, and CSI; lower is better for ECE.

Variant	**Enron Mod.**	**DSBM NMI**	**LFR CSI**	**DSBM ECE**
Full BayesMem-CD	**0.522**	**0.721**	**0.829**	**0.052**
w/o Memory	0.487	0.692	0.781	0.064
w/o Bayesian	0.506	0.704	0.802	0.073
w/o Markov	0.481	0.697	0.758	0.057
w/o Fusion	0.493	0.688	0.774	0.066
w/o Uncertainty Avg.	0.511	0.709	0.811	0.069
Deterministic Predictor	0.498	0.681	0.786	0.078

4.2 Overall Performance

Table 1 summarizes the principal results. For synthetic networks, BayesMem-CD achieves the strongest average community recovery, with an average NMI of 0.714 and ARI of 0.703. The improvement is accompanied by substantially better uncertainty calibration: its average ECE is 0.051, compared with 0.067 for EvolveGCN and 0.073 for DySAT. On real-world networks, BayesMem-CD maintains essentially the same average modularity as the strongest temporal baselines while achieving the strongest average CSI among the compared methods. These results indicate that improved temporal stability and uncertainty estimation do not require a substantial loss in snapshot-level community quality.

4.3 Component Analysis

We next isolate the contributions of the major components of BayesMem-CD. Table 2 shows complementary effects across structural, temporal, and uncertainty-aware criteria. Removing graph memory decreases DSBM NMI from 0.721 to 0.692 and Dynamic LFR CSI from 0.829 to 0.781, demonstrating the importance of long-range structural information. Removing Bayesian inference increases ECE from 0.052 to 0.073, producing the largest calibration degradation. Removing Markov regularization has the strongest effect on temporal stability, reducing CSI from 0.829 to 0.758. Removing representation fusion or posterior averaging also consistently degrades performance, while replacing the Bayesian predictor with a deterministic predictor produces the weakest overall results.

4.4 Temporal Signal Recovery and Posterior Smoothing

To examine how BayesMem-CD behaves under temporally corrupted observations, we visualize representative node-level trajectories under clean, noisy, and posterior-recovered signals. This analysis complements the aggregate community metrics by directly showing whether the model suppresses transient noise while preserving meaningful temporal variation.

As shown in Fig. 2, the posterior trajectories remain close to the underlying clean signals despite substantial local perturbations. The recovered signals suppress high-frequency fluctuations while preserving slower structural changes across heterogeneous vertices. This behavior provides direct evidence that the temporal memory and probabilistic inference mechanisms stabilize noisy observations without collapsing genuine temporal variation.

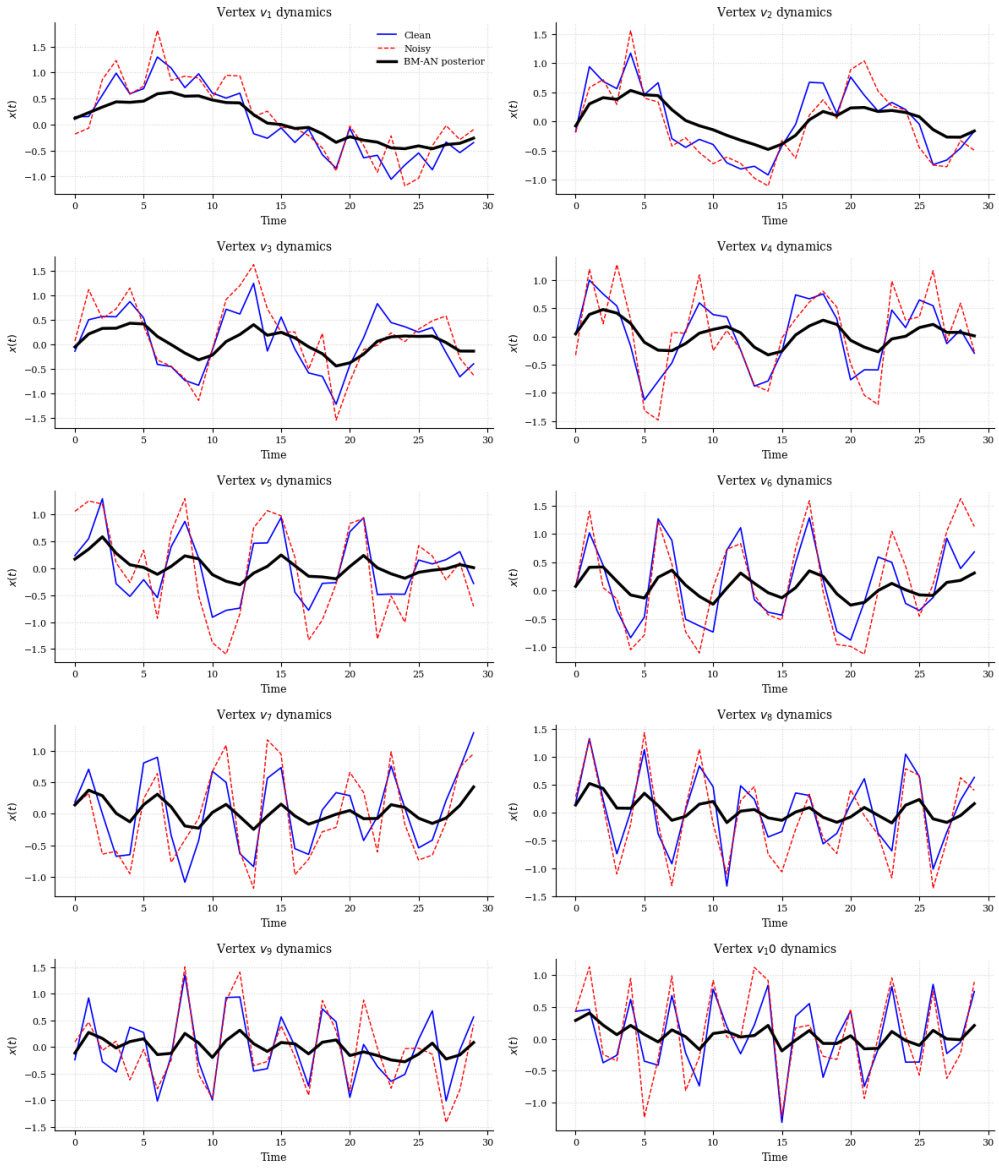


Figure 2: Vertex-level temporal dynamics under clean, noisy, and BayesMem-CD posterior signals. Blue curves denote clean temporal signals, red dashed curves indicate noise-corrupted observations, and black curves represent posterior estimates. BayesMem-CD suppresses short-term perturbations while retaining the underlying temporal structure.

4.5 Interpreting Temporal Community Signals

To examine the structure of the learned community signal beyond aggregate metrics, we perform a post-hoc spectral analysis on a representative DSBM snapshot. Figure 3 decomposes the posterior community signal into graph-frequency components. The analysis shows that the learned posterior is dominated by low-frequency components aligned with coherent community structure, whereas much of the injected perturbation is concentrated in higher-frequency components. This analysis is used only for interpretation and is not part of the BayesMem-CD training objective.

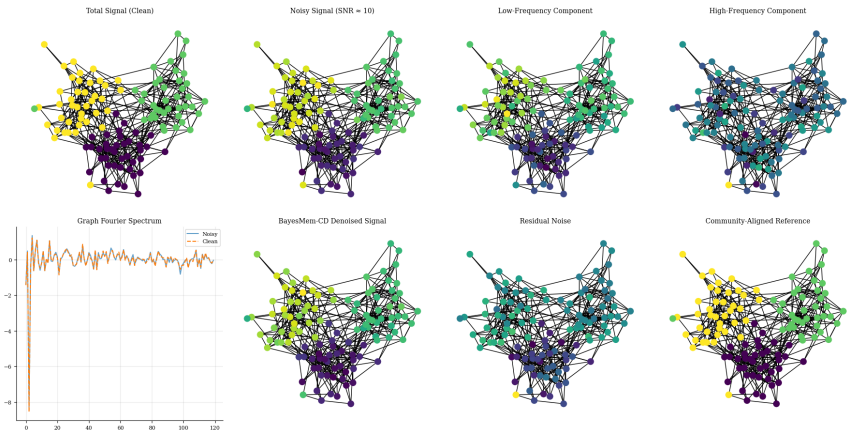


Figure 3: Post-hoc spectral analysis of a BayesMem-CD posterior community signal on DSBM at $t = 25$. The visualization compares the posterior and noisy signals with their graph-frequency decomposition and a Tikhonov-denoised reference. Low-frequency components correspond more closely to coherent community structure, whereas high-frequency components capture local perturbations.

4.6 Sensitivity to Joint Regularization

Finally, we examine whether BayesMem-CD depends on narrowly tuned regularization strengths. We vary the temporal and Bayesian regularization weights jointly and visualize the normalized training objective in Fig. 4. The resulting surface contains a broad low-loss region at moderate parameter values rather than an isolated optimum, indicating that the model remains stable across a non-trivial range of regularization strengths. Additional one-dimensional sensitivity results for the Monte Carlo sample count, memory horizon, and Markov coefficient are provided in the Supplementary Material.

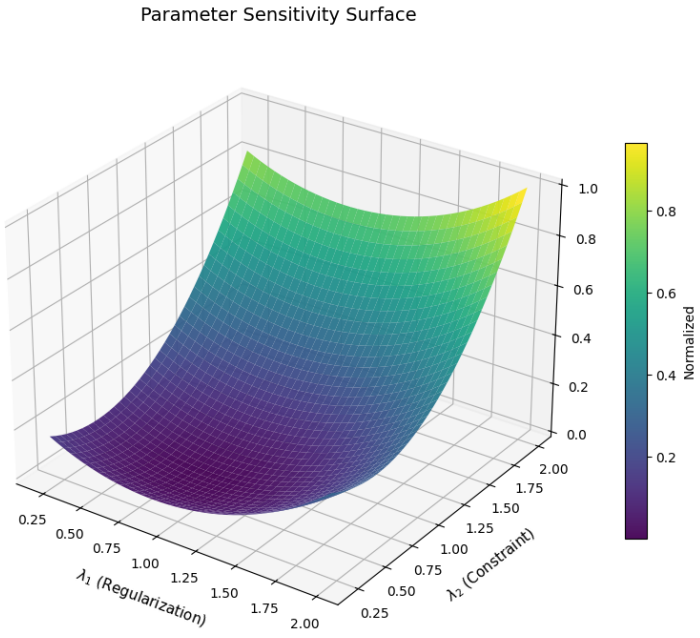


Figure 4: Joint sensitivity of the BayesMem-CD training objective to temporal regularization and Bayesian uncertainty weighting. A broad low-loss basin appears around moderate regularization strengths.

Complete per-dataset results, Macro- and Micro-F1 scores, full modularity and CSI comparisons, evolution consistency, perturbation robustness, computational efficiency, scalability, and additional sensitivity analyses are reported in the Supplementary Material.

5 Discussion

The experiments show that BayesMem-CD provides a balanced improvement in community recovery, temporal stability, and predictive uncertainty rather than optimizing a single metric in isolation. On synthetic networks, the model achieves the strongest average NMI and ARI together with substantially lower calibration error, while on real-world networks it maintains competitive modularity and strong temporal consistency. These results suggest that uncertainty-aware temporal modeling can improve the reliability of dynamic community assignments without sacrificing snapshot-level structural quality.

The ablation results clarify the complementary roles of the main components. Graph memory preserves long-range structural information and improves community recovery under noisy or incomplete snapshots, whereas Markov regularization primarily controls short-term fluctuations between consecutive community distributions. Bayesian inference contributes differently by representing ambiguous assignments probabilistically and improving calibration, which is particularly important for nodes near community boundaries or during birth, split, merge, and drift events. This distinction supports a dual-horizon interpretation of BayesMem-CD: long-range persistence is maintained at the representation level, while short-range consistency is controlled directly at the output-distribution level.

The additional analyses provide further evidence for this behavior. The post-hoc spectral analysis in Fig. 3 shows that the learned posterior community signals are more strongly associated with low-frequency graph components corresponding to coherent community structure, whereas transient perturbations contribute more strongly to high-frequency components. This analysis is diagnostic rather than part of the training procedure. In addition, the joint regularization landscape in Fig. 4 contains a broad low-loss region around moderate Bayesian and temporal regularization strengths, indicating that the model does not depend on a narrowly tuned parameter configuration.

6 Limitations and Future Work

BayesMem-CD has several limitations. First, the current Bayesian formulation primarily captures epistemic uncertainty and does not explicitly separate it from aleatoric uncertainty arising from noisy or incomplete observations. Second, the common node-indexing scheme requires explicit treatment of node arrival, disappearance, and re-entry in highly open dynamic systems. Extending the framework toward multi-source uncertainty modeling and inductive, arrival-aware memory updates would improve its applicability to such settings Trivedi et al. (2024); Damke & Hüllermeier (2024); Cirillo et al. (2023); Zhang et al. (2023).

The current Markov regularizer also models only consecutive community distributions, while posterior sampling increases computational cost. Higher-order or continuous-time temporal models could capture delayed and non-Markovian dependencies, while efficient posterior approximations or adaptive sampling could improve scalability Gravina et al. (2024); Gao et al. (2024); Song et al. (2024). Evaluation under stronger distribution shifts and temporally faithful benchmarks remains an important direction for future work.

7 Conclusion

We introduced BayesMem-CD, a Bayesian temporal-memory framework for uncertainty-aware dynamic community detection that combines long-range graph memory, probabilistic community inference, and short-range temporal regularization. Experiments on four synthetic and four real-world dynamic networks demonstrate strong community recovery, improved uncertainty calibration, and stable temporal behavior. The results show that memory, Bayesian inference, and temporal regularization provide complementary mechanisms for reliable community inference in noisy and evolving graphs.

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

Claudio Daniel Tenorio Barros, Matheus R. F. Mendonça, Alex Borges Vieira, and Artur Ziviani. A survey on embedding dynamic graphs. *ACM Computing Surveys*, pp. 10:1–10:37, 2023.

Junwei Cheng, Chaobo He, Kunlin Han, Wenjie Ma, and Yong Tang. How significant attributes are in the community detection of attributed multiplex networks. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2057–2061, 2023.

Michele Cirillo, Vincenzo Matta, and Ali H. Sayed. Learning dynamic graphs under partial observability. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2023.

Clemens Damke and Eyke Hüllermeier. Linear opinion pooling for uncertainty quantification on graphs. In *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence*, pp. 919–929, 2024.

Shihong Gao, Yiming Li, Yanyan Shen, Yingxia Shao, and Lei Chen. Etc: Efficient training of temporal graph neural networks over large-scale dynamic graphs. *Proceedings of the VLDB Endowment*, 17(5):1060–1072, 2024.

Alessio Gravina, Giulio Lovisotto, Claudio Gallicchio, Davide Bacciu, and Claas Grohnfeldt. Long range propagation on continuous-time dynamic graphs. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 16206–16225, 2024.

Kun Guo, Zizheng Zhao, Zhiyong Yu, Wenzhong Guo, Ronghua Lin, Yong Tang, and Ling Wu. Network embedding based on biased random walk for community detection in attributed networks. *IEEE Transactions on Computational Social Systems*, 10(5):2279–2290, 2023.

Chaobo He, Yulong Zheng, Junwei Cheng, Yong Tang, Guohua Chen, and Hai Liu. Semi-supervised overlapping community detection in attributed graph with graph convolutional autoencoder. *Information Sciences*, 608:1464–1479, 2022a.

Chaobo He, Yulong Zheng, Xiang Fei, Hanchao Li, Zeng Hu, and Yong Tang. Boosting nonnegative matrix factorization based community detection with graph attention auto-encoder. *IEEE Transactions on Big Data*, 8(4):968–981, 2022b.

Qiang Huang, Xin Wang, Susie Xi Rao, Zhichao Han, Zitao Zhang, Yongjun He, Quanqing Xu, Yang Zhao, Zhigao Zheng, and Jiawei Jiang. Benchtemp: A general benchmark for evaluating temporal graph neural networks. In *IEEE International Conference on Data Engineering*, pp. 4044–4057, 2024.

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

Seyed Mehran Kazemi, Rishab Goel, Kshitij Jain, Ivan Kobyzev, Akshay Sethi, Peter Forsyth, and Pascal Poupart. Representation learning for dynamic graphs: A survey. *Journal of Machine Learning Research*, 21(70):1–73, 2020.

Tong Liu, Yushan Liu, Marcel Hildebrandt, Mitchell Joblin, Hang Li, and Volker Tresp. On calibration of graph neural networks for node classification. In *International Joint Conference on Neural Networks*, pp. 1–8, 2022.

Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pp. 15682–15694, 2021.

Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao B. Schardl, and Charles E. Leiserson. Evolvegcn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 5363–5370, 2020.

Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael M. Bronstein. Temporal graph networks for deep learning on dynamic graphs. *CoRR*, abs/2006.10637, 2020.

Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. Dysat: Deep neural representation learning on dynamic graphs via self-attention networks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pp. 519–527, 2020.

Joakim Skarding, Bogdan Gabrys, and Katarzyna Musial. Foundations and modeling of dynamic networks using dynamic graph neural networks: A survey. *IEEE Access*, 9:79143–79168, 2021.

Yifan Song, Xiaolong Chen, Wenqing Lin, Jia Li, Chen Zhang, Yan Zhou, Lei Chen, and Jing Tang. Efficient graph embedding generation and update for large-scale temporal graph. *Proceedings of the VLDB Endowment*, 18(4):929–942, 2024.

Junwei Su and Shan Wu. Temporal-aware evaluation and learning for temporal graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 20628–20636, 2025.

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*, 2024.

Hongkai Xie, Xinyi Ying, Xiaofeng Wang, Yuanyuan Qi, Wei Chen, Xiaofeng Huang, Junzheng Jiang, and Daying Quan. Redundancy-aware masked graph autoencoder for overlapping community detection in attributed networks. *Engineering Applications of Artificial Intelligence*, 162: 112821, 2025.

Da Xu, Chuanwei Ruan, Evren Körpeoğlu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. In *International Conference on Learning Representations*, 2020.

Leshanshui Yang, Clément Chatelain, and Sébastien Adam. Dynamic graph representation learning with neural networks: A survey. *IEEE Access*, 12:43460–43484, 2024.

Linlin Yu, Kangshuo Li, Pritom Kumar Saha, Yifei Lou, and Feng Chen. Evidential uncertainty probes for graph neural networks. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, pp. 2845–2853, 2025.

Haonan Yuan, Qingyun Sun, Xingcheng Fu, Ziwei Zhang, Cheng Ji, Hao Peng, and Jianxin Li. Environment-aware dynamic graph learning for out-of-distribution generalization. In *Advances in Neural Information Processing Systems*, 2023.

Peiyan Zhang, Yuchen Yan, Chaozhuo Li, Senzhang Wang, Xing Xie, Guojie Song, and Sunghun Kim. Continual learning on dynamic graphs via parameter isolation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 601–611, 2023.

Yang Zhang, Ting Yu, Shengqiang Chi, Zhen Wang, Yue Gao, Ji Zhang, and Tianshu Zhou. Attribute diversity aware community detection on attributed graphs using three-view graph attention neural networks. *ACM Transactions on Knowledge Discovery from Data*, 18(8):196:1–196:24, 2024.

Yanping Zheng, Hanzhi Wang, Zhewei Wei, Jiajun Liu, and Sibo Wang. Instant graph neural networks for dynamic graphs. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2605–2615, 2022.

Yanping Zheng, Zhewei Wei, and Jiajun Liu. Decoupled graph neural networks for large dynamic graphs. *Proceedings of the VLDB Endowment*, 16(9):2239–2247, 2023.

Zulun Zhu, Kai Wang, Haoyu Liu, Jintang Li, and Siqiang Luo. Topology-monitorable contrastive learning on dynamic graphs. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4700–4711, 2024.