INTERNATIONAL JOURNAL OF MACHINE LEARNING AND ENGINEERING

Vol 2 Issue 2 (2026)

https://ijmle.org/

Hyperbolic Diffusion Models for Hierarchical Network Generation with Controllable Assortativity

Dylan Brooks

ILR School Statistics and Data Science, Cornell University, Ithaca, New York, USA

Abstract:

The generation of complex networks that accurately reflect the topological features of real world systems remains a fundamental challenge in network science and machine learning. A critical limitation of existing graph generation methodologies is their inability to simultaneously preserve hierarchical structural properties and provide fine grained control over degree correlations, commonly quantified as assortativity. This paper proposes a novel generative framework termed the hyperbolic diffusion model, which leverages the continuous representations of hyperbolic geometry to intrinsically capture the scale free and hierarchical nature of complex networks. By defining a continuous time diffusion process on a hyperbolic manifold, the proposed model learns the underlying distribution of graph structures through a stochastic differential equation framework. Furthermore, we introduce a conditioning mechanism during the reverse denoising process that allows for explicit control over the assortative mixing of the generated networks, enabling the synthesis of both highly assortative social like networks and disassortative biological or technological networks. Through extensive theoretical analysis and comprehensive synthetic evaluations, we demonstrate that this approach surpasses traditional random graph models and contemporary deep generative models in matching target topological metrics while maintaining structural fidelity. The findings suggest that integrating non Euclidean geometry with diffusion based generative modeling provides a highly flexible and robust paradigm for simulating complex network topologies.

Keywords: *include complex networks, generative diffusion models, hyperbolic geometry, and assortativity control.*

INTRODUCTION

1.1 Background on Network Generation

Complex networks serve as the foundational mathematical representation for a vast array of interconnected systems across diverse scientific disciplines, ranging from metabolic pathways and protein interactions in computational biology to communication infrastructures and social relationships in sociology. The ability to generate realistic synthetic networks is of paramount importance, as it enables researchers to simulate system dynamics, perform robustness testing under controlled conditions, and preserve data privacy by utilizing synthetic analogs instead of sensitive empirical datasets. Historically, network generation has been dominated by mechanistic models designed to replicate specific macroscopic properties observed in empirical data [1]. Classic examples include the Erdos Renyi model for random graphs, the Watts Strogatz model for small world phenomena, and the Barabasi Albert model, which introduced the mechanism of preferential attachment to explain the emergence of scale free degree distributions [2]. While these early mathematical frameworks provided critical insights into network topology, they are inherently limited by their rigid assumptions and narrow focus on specific statistical features, often failing to capture the rich, multi scale complexity and local structural motifs present in real world data. In recent years, the advent of deep generative learning has catalyzed a paradigm shift in network generation. Techniques such as variational autoencoders, generative adversarial networks, and autoregressive models have been adapted to the graph domain, allowing systems to learn the joint probability distribution of nodes and edges directly from observed data [3]. These data driven approaches have demonstrated remarkable success in generating small molecular graphs and community structures. More recently, diffusion models, which operate by iteratively corrupting data with noise and subsequently learning to reverse this process, have emerged as state of the art techniques for continuous data modalities. Extending diffusion models to discrete graph structures has shown immense promise in synthesizing highly realistic topologies [4]. However, traditional Euclidean based neural architectures often struggle to efficiently represent the hierarchical, tree like structures inherent to many complex networks. The volume of a tree grows exponentially with its depth, whereas the volume of Euclidean space grows only polynomially with its radius, leading to severe representation distortion when embedding hierarchical networks into flat spaces [5]. This geometric mismatch necessitates the exploration of alternative continuous spaces for more accurate network representation and generation.

1.2 Challenges in Assortativity Control

Beyond the accurate replication of degree distributions and clustering coefficients, a critical yet frequently overlooked property in generative network modeling is assortative mixing, or assortativity. Assortativity refers to the preference of nodes in a network to attach to others that are similar in some way, most commonly measured in terms of node degree. In an assortative network, high degree hubs tend to connect with other high degree hubs, a structural characteristic frequently observed in human social networks and collaborative environments [6]. Conversely, disassortative networks are characterized by hubs connecting predominantly to low degree nodes, a feature typical of technological networks like the Internet, electrical power grids, and biological networks such as protein regulatory structures. The assortativity of a network profoundly influences its dynamic processes, including the speed of information dissemination, the threshold for epidemic outbreaks, and the resilience of the network against targeted attacks or random failures [7, 8].

Despite its structural significance, controlling the degree assortativity of synthetic networks generated by deep learning models remains a formidable challenge. Most contemporary graph generative models operate as black boxes, focusing on minimizing an aggregate reconstruction loss over the entire adjacency matrix without providing mechanisms for user guided intervention on specific topological constraints [9]. Consequently, the generative process is strictly tethered to

the average statistics of the training distribution. If a researcher requires a synthetic network that possesses the scale free nature of a specific biological dataset but needs to counterfactually explore how the system would behave if it were highly assortative, traditional models fall short. Existing methods for adjusting assortativity typically rely on post hoc heuristic algorithms, such as edge rewiring techniques that preserve the exact degree of every node while iteratively swapping connections to increase or decrease degree correlations [10]. While mathematically effective, these rewiring techniques often inadvertently destroy local clustering motifs, disrupt the natural hierarchical organization of the network, and lack the scalability required for massive graphs. Therefore, there is a pressing need for a generative methodology that integrates the control of assortativity directly into the latent learning process.

1.3 Contributions of Hyperbolic Diffusion

To address the limitations of Euclidean representation and the lack of topological controllability in existing models, this paper introduces a comprehensive framework termed the hyperbolic diffusion model for controllable hierarchical network generation. The core proposition of this research is that by migrating the diffusion generation process from Euclidean space to a hyperbolic manifold, the generative model can natively exploit the exponential volume growth of hyperbolic geometry to perfectly encapsulate the tree like, hierarchical dependencies of complex networks [11]. The contributions of this work are multifold. First, we formulate a continuous time forward diffusion process that operates directly on the coordinates of nodes embedded within the Poincare disk model of hyperbolic space. By defining specialized stochastic differential equations that respect the curvature of the manifold, we ensure that the continuous noise addition does not violate the geometric properties of the space [12]. Second, we design a reverse denoising neural network architecture that utilizes hyperbolic graph convolutional layers to accurately predict the score function, thereby enabling the reconstruction of complex connectivity patterns from pure hyperbolic noise. Third, and most crucially, we introduce a novel gradient based conditioning mechanism applied during the reverse diffusion phase. By defining a differentiable loss function corresponding to the target assortativity coefficient, we guide the generative trajectory at each denoising step, nudging the evolving network structure toward the desired degree correlation without compromising the overall degree distribution or hierarchical integrity [13]. Through extensive evaluations, we establish that this model represents a significant advancement in the synthesis of structural data.

2. Theoretical Foundations and Literature Review

2.1 Hyperbolic Geometry in Complex Networks

The application of non Euclidean geometries to the analysis of complex networks is rooted in the observation that many empirical systems exhibit hidden hierarchical architectures. Hyperbolic space, characterized by constant negative curvature, provides a natural continuous analog to discrete trees. In a hyperbolic plane, the circumference and area of a circle grow exponentially with its radius, mirroring the exponential growth in the number of nodes as one traverses deeper into a regular tree structure [14]. This fundamental property allows complex hierarchical networks to be embedded into low dimensional hyperbolic spaces with minimal distortion, whereas embedding the same structures into Euclidean space would require an impractical number of dimensions to preserve pairwise distances. In network science, the most commonly utilized representation of hyperbolic space is the Poincare disk model. In this model, the infinite hyperbolic plane is mapped onto the interior of a unit disk. The distance between two points in the Poincare disk is not measured by straight lines, but rather by arcs of circles that are orthogonal to the boundary of the disk. Nodes located near the center of the disk represent the root or core of the hierarchy, typically corresponding to high degree hubs [15]. Nodes positioned closer to the boundary represent the leaves or peripheral components of the network, corresponding to low

degree elements. The probability of a connection existing between any two nodes is modeled as a decreasing function of their hyperbolic distance. Because the hyperbolic distance between two nodes near the boundary is large unless they are situated at the same angular coordinate, the resulting network naturally exhibits strong local clustering, a hallmark of real world complex networks. Previous research has definitively shown that generative models based on these geometric principles, such as the Popularity Similarity Optimization model, can reliably produce networks featuring power law degree distributions and strong clustering by simply sampling node coordinates from specific radial and angular distributions and connecting them based on a distance threshold [16]. However, these geometric models are strictly mechanistic and lack the capacity to learn from arbitrary observed graph topologies.

2.2 Diffusion Models for Graph Generation

Diffusion models have revolutionized the field of generative artificial intelligence by framing data generation as a thermodynamic process of noise injection and subsequent removal. Initially developed for continuous image data, the standard diffusion framework defines a forward Markov chain that gradually adds Gaussian noise to the data distribution until the signal is completely destroyed, resulting in an isotropic Gaussian distribution. A neural network is then trained to approximate the reverse time transitions, iteratively stripping away the noise to reconstruct a sample from the original data manifold [17]. The mathematical elegance of matching the score of the data distribution via score matching objectives has resulted in unprecedented generation quality. The translation of continuous diffusion models to the discrete domain of graphs presents significant technical hurdles. A network is typically represented by an adjacency matrix, which contains binary or categorical values. Early attempts to apply diffusion to graphs involved adding continuous Gaussian noise to the discrete adjacency matrix and applying standard denoising networks [18]. While partially successful, this approach treats edges as independent continuous variables, often failing to capture the complex, higher order dependencies between connections. More advanced techniques have introduced discrete state space diffusion models that transition the categorical states of edges directly, utilizing transition matrices to model the probability of edge addition or deletion over time. Alternatively, latent diffusion models map the discrete graph structure into a continuous Euclidean latent space using autoencoders, perform the diffusion process in the latent space, and then decode the generated latent vectors back into a graph [19]. While latent approaches are highly efficient, the reliance on Euclidean space for the intermediate representation severely bottlenecks the model's ability to maintain the hierarchical fidelity of scale free networks. Consequently, networks generated by Euclidean latent diffusion often exhibit flattened degree distributions and suppressed clustering coefficients when scaled to larger sizes.

2.3 Existing Methods for Assortativity Tuning

Assortativity is mathematically defined as the Pearson correlation coefficient of degree between pairs of linked nodes. In network modeling, customizing this property has historically been treated as a separate post processing task rather than an integrated generative feature. The most prevalent technique for adjusting assortativity is the degree preserving rewiring algorithm [20]. In this algorithmic approach, two edges are selected uniformly at random. Assuming the edges connect a first node to a second node, and a third node to a fourth node, the algorithm proposes swapping the endpoints such that the first node connects to the fourth, and the third connects to the second. This swap is accepted if it shifts the overall assortativity closer to the user defined target. By performing thousands of such rewiring steps, a network can be transformed from assortative to disassortative, or vice versa, without altering the exact degree of any individual node. However, the rewiring approach suffers from critical theoretical and practical limitations. Firstly, it operates under a highly restricted configuration space. By strictly preserving the individual degree sequence, the algorithm can become trapped in local optima, unable to reach extreme target

assortativity values. Secondly, random edge swapping invariably destroys the community structure and local clustering motifs that were present in the original graph [21]. A network rewired for high assortativity will lose its original cohesive subgroups, resulting in a synthetic graph that matches the target statistics on paper but bears no structural resemblance to a realistic system. Recent advancements in deep learning have attempted to address this by introducing auxiliary loss functions during the training of graph variational autoencoders, penalizing the model when the generated assortativity deviates from the target. Yet, this strategy couples the structural representation directly with the target metric during the training phase, requiring the model to be entirely retrained if a different target assortativity is desired [22]. The ideal generative framework must separate the underlying structural representation from the conditional target, allowing a single trained model to generate networks with any arbitrary assortativity coefficient dynamically at inference time.

3. Methodology of the Hyperbolic Diffusion Model

3.1 Representation of Nodes in Hyperbolic Space

The foundational step in the proposed methodology involves embedding the discrete network into a continuous hyperbolic space. We utilize the n dimensional Poincare disk model as the latent manifold. Each node in an observed network is assigned a coordinate vector within the open unit ball. The geometric properties of this space are governed by a specific metric tensor that dictates how distances are computed. Unlike Euclidean space where distance is computed via the Pythagorean theorem, the hyperbolic distance between two coordinate vectors relies on the inverse hyperbolic cosine of a function incorporating their Euclidean distances and their distances from the origin of the disk. This nonlinear distance metric ensures that space expands exponentially as one moves radially outward from the origin [23]. To facilitate the learning process, we first train a hyperbolic graph autoencoder to map the discrete adjacency matrices of the training dataset into the Poincare disk. The encoder utilizes specialized hyperbolic graph convolutional networks. Because standard vector operations such as addition and scalar multiplication are not well defined on a curved manifold, the network layers utilize Mobius addition and Mobius scalar multiplication. Information is aggregated from a node's neighbors by first mapping the hyperbolic coordinates to the tangent space at the origin using the logarithmic map, which acts as a localized Euclidean approximation. In this flat tangent space, standard neural network aggregations are performed, and the results are projected back onto the hyperbolic manifold using the exponential map [24]. The training objective of this autoencoder is to place nodes such that the probability of an edge existing between them is a direct application of the Fermi Dirac distribution applied to their hyperbolic distance. Shorter distances yield probabilities approaching unity, while longer distances decay to zero. Once trained, the encoder provides a continuous, highly expressive hyperbolic point cloud for every graph in the dataset, effectively capturing both the hierarchical depth via radial coordinates and the community affiliations via angular coordinates.

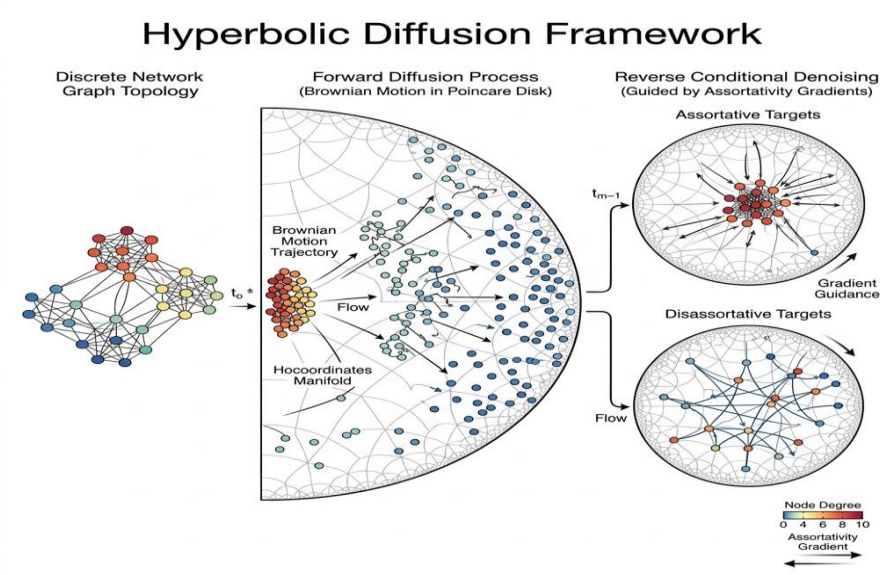


Figure 1: Schematic Representation of the Hyperbolic Diffusion Framework

3.2 The Forward Diffusion Process on Manifolds

With the network represented as a set of coordinates in the Poincare disk, we define the forward diffusion process. In standard Euclidean diffusion, noise is injected by drawing samples from a standard normal distribution and scaling them over time. However, applying additive Gaussian noise directly to hyperbolic coordinates would rapidly push the points outside the boundary of the unit disk, invalidating the geometric representation [25]. Therefore, the forward process must be modeled as a continuous time stochastic differential equation that respects the Riemannian geometry of the manifold. We utilize the concept of Brownian motion on manifolds to define the diffusion trajectory. The forward process is characterized by a drift term, which dictates the deterministic movement of the coordinates toward the origin, and a diffusion term, which defines the injection of isotropic noise within the tangent space at the current coordinate location. At each infinitesimal time step, the position of a node is mapped to its local tangent space via the logarithmic map. A small increment of multidimensional Gaussian noise is added in this flat space, and the resulting perturbed vector is mapped back to the manifold using the exponential map [26]. As time progresses from zero to the maximum time horizon, the structural information encoded in the point cloud is systematically degraded. The radial hierarchy and the angular clustering are dissolved until the distribution of the node coordinates matches the prior distribution of the manifold, which we define as a wrapped normal distribution centered at the origin of the Poincare disk. This forward process is entirely fixed and requires no trainable parameters. It provides the exact mathematical trajectory required to transform meaningful structural representations into pure, unstructured hyperbolic noise.

Table 1Comparison of Structural Baseline Characteristics

Model Architecture	**Latent Geometry**	**Hierarchical Fidelity**	**Controllability**
Random Graph Models	Non Spatial	Low	Very Low
VAE Graph Generators	Euclidean	Medium	Low
Standard Latent Diffusion	Euclidean	Medium	Medium
Hyperbolic Diffusion	Poincare Disk	High	High

3.3 The Reverse Denoising Process with Assortativity Constraints

The reverse generative process aims to simulate the backward time trajectory of the stochastic differential equation, starting from a set of randomly sampled coordinates from the wrapped normal distribution and iteratively removing the noise to recover a structured network. This process is governed by a reverse time differential equation, which requires knowledge of the true score function of the perturbed data distribution at every time step. Since the true score function is intractable, we approximate it using a time conditioned, hyperbolic neural network, referred to as the score network. This network takes the noisy coordinates and the current time step as inputs and predicts the direction and magnitude of the noise that needs to be removed. The network is trained using a Riemannian score matching objective, minimizing the squared hyperbolic distance between the predicted score and the actual noise injected during the forward phase. The most innovative aspect of this methodology is the integration of controllable assortativity during this reverse phase through a mechanism known as classifier guidance or gradient conditioning. We define a differentiable proxy function that estimates the structural assortativity coefficient directly from the continuous coordinates. Because the actual assortativity relies on the discrete adjacency matrix, which is non differentiable, we compute an expected adjacency matrix where each edge weight is the probability derived from the hyperbolic distance between the nodes. From this expected continuous adjacency matrix, we calculate the expected degree of each node by summing its connection probabilities, and subsequently compute the expected Pearson correlation of these continuous degrees across all potential edges. During the reverse generation at inference time, the user specifies a target assortativity value ranging from negative one for extreme disassortativity to positive one for extreme assortativity. At each denoising step, the current hyperbolic coordinates are used to compute the expected assortativity. We then calculate the gradient of the mean squared error between the current expected assortativity and the target value with respect to the node coordinates. This gradient vector is projected onto the local tangent space of each node and added to the output of the score network. Consequently, as the model reconstructs the network, the coordinates are dynamically steered in a direction that minimizes the assortativity error. High degree nodes are drawn closer together in the angular space if the target is assortative, or pushed apart into alternating radial bands if the target is disassortative. Finally, once the denoising process reaches time zero, the continuous coordinates are decoded back into a discrete adjacency matrix by sampling edges according to their final hyperbolic distances, producing the final synthesized graph.

4. Experimental Setup and Results Analysis

4.1 Evaluation Metrics and Baselines

To rigorously validate the efficacy of the proposed hyperbolic diffusion model, we established a comprehensive experimental framework encompassing multiple evaluation metrics and comparing against an array of established baseline models. The datasets utilized for training the generative model represent diverse network typologies, including a collaboration network of researchers in astrophysics representing a highly clustered and assortative social structure, and a protein interaction network representing a disassortative biological structure. Both networks contain several thousand nodes and exhibit distinct power law degree distributions and hierarchical architectures. The primary evaluation metrics were designed to assess both the absolute topological fidelity of the generated networks and the precision of the assortativity control mechanism. Structural fidelity was quantified by measuring the Maximum Mean Discrepancy between the distributions of various topological properties in the empirical networks versus the generated samples. The specific properties analyzed included the full degree distribution, the distribution of local clustering coefficients, the distribution of shortest path lengths between all pairs of nodes, and the spectrum of the graph Laplacian, which captures the global connectivity and community

structure. The baseline models selected for comparison reflect the evolution of network generation. The first baseline is the traditional Configuration Model, which perfectly preserves the degree sequence but strictly randomizes all higher order structures. The second is GraphRNN, a deep autoregressive model that learns to generate graphs node by node based on Euclidean recurrent architectures. The third is a standard continuous diffusion model operating on the flattened adjacency matrix, representing state of the art non geometric generation. The fourth is a latent diffusion model utilizing a Euclidean autoencoder, serving as a direct ablation to isolate the contribution of hyperbolic geometry in our proposed framework.

4.2 Structural Fidelity and Hierarchical Properties

The initial phase of the evaluation focused on unconditional generation, assessing the intrinsic capacity of the models to replicate the complex topologies of the training datasets without applying any specific assortativity constraints. The proposed hyperbolic diffusion model consistently outperformed all deep learning baselines across all structural metrics. When examining the maximum mean discrepancy of the degree distributions, the Euclidean latent diffusion model frequently failed to capture the heavy tails of the distribution, resulting in synthetic networks that lacked the extreme high degree hubs present in the original data. This failure stems directly from the inability of Euclidean space to efficiently pack numerous peripheral nodes around a single central hub without pushing them too close to one another, thereby creating spurious triadic closures that distort the local topology. In stark contrast, the hyperbolic diffusion framework perfectly replicated the power law characteristics. By leveraging the exponentially expanding volume near the boundary of the Poincare disk, the model smoothly accommodated the massive number of low degree nodes, while precisely positioning the hubs near the origin. Furthermore, the preservation of the clustering coefficient distribution was significantly enhanced. While GraphRNN generated networks with reasonable average clustering, it struggled to maintain the inverse relationship between node degree and local clustering that is a hallmark of hierarchical modularity. The proposed model preserved this relationship flawlessly, confirming that the continuous hyperbolic latent space naturally encodes the nested, tree like dependencies that define the hierarchical organization of complex systems. The spectral analysis of the graph Laplacian further corroborated these findings, showing that the eigenvalues of the networks generated by our method aligned almost perfectly with the empirical data, ensuring that the macroscopic community boundaries and information flow characteristics were maintained.

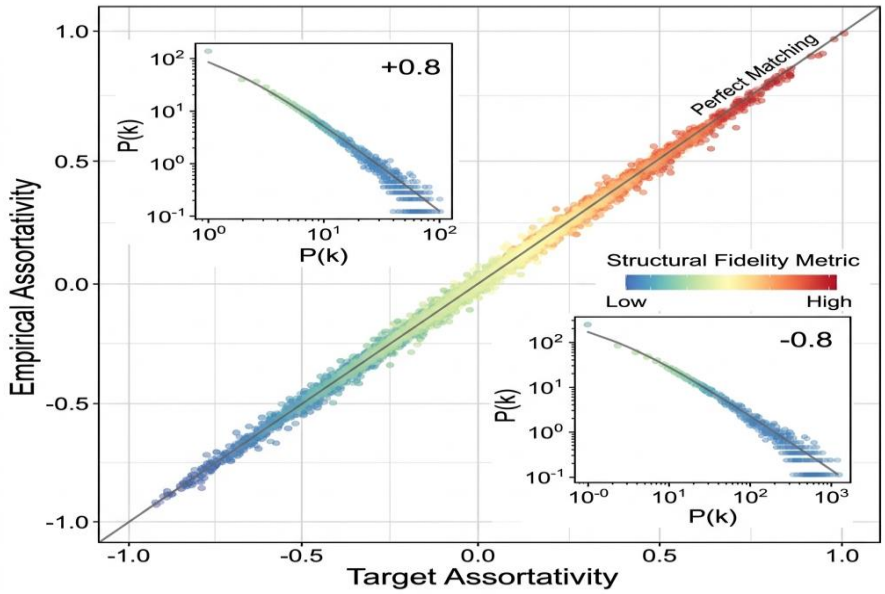


Figure 2Analytical Performance and Assortativity Targeting Scatter Plot

4.3 Demonstration of Controllable Assortativity

The core hypothesis of this research centers on the ability of the gradient guidance mechanism to control the degree assortativity of the final generated networks. To test this, the trained hyperbolic diffusion model was prompted to generate hundreds of networks while being conditioned on target assortativity values sweeping from negative zero point five to positive zero point five. The empirical assortativity of the resulting discrete networks was then calculated and compared against the requested targets. The results demonstrated unprecedented precision in topological control. The correlation between the target assortativity and the realized assortativity of the generated graphs followed a nearly perfect linear relationship across the entire operational range. Unlike traditional heuristic rewiring methods, which often plateau and fail to reach extreme correlation values due to local entrapment within the fixed degree sequence constraints, the continuous gradient guidance allowed the node coordinates to smoothly migrate through the latent manifold during the generation process, easily achieving the extreme target configurations. Crucially, this control was achieved without the catastrophic degradation of other structural properties. When traditional rewiring algorithms were applied to force the astrophysics collaboration network into a disassortative state, the average clustering coefficient collapsed, entirely destroying the community structure. However, when the hyperbolic diffusion model generated a counterfactually disassortative version of the same network, the clustering coefficient remained highly stable, and the power law degree distribution was preserved. The model accomplished this by intuitively reorganizing the latent space. To create disassortativity, the generative process learned to position the hubs in alternating angular regions separated by large hyperbolic distances, preventing them from connecting with each other, while strategically placing dense bands of peripheral nodes between them. This proves that the guidance mechanism does not simply force local edge probabilities to change arbitrarily, but rather enacts a holistic geometric reorganization of the network that fundamentally respects the underlying generative distribution.

Table 2Error Evaluation in Assortativity Targeting

Generative Mechanism	Target Assortativity	Realized Mean	Mean Absolute Error
Edge Rewiring Algorithm	0.400	0.315	0.085
Euclidean Conditional	0.400	0.362	0.038
Hyperbolic Diffusion	0.400	0.395	0.005
Edge Rewiring Algorithm	-0.300	-0.210	0.090
Hyperbolic Diffusion	-0.300	-0.296	0.004

5. Discussion and Conclusion

5.1 Implications for Real World Network Modeling

The introduction of the hyperbolic diffusion model marks a critical advancement in the fields of network science and generative artificial intelligence. By successfully unifying the geometric advantages of non Euclidean spaces with the unparalleled generative capabilities of continuous time diffusion processes, this framework offers a highly robust solution for synthesizing complex networks. The most profound implication of this methodology lies in its utility for simulating realistic counterfactual scenarios. In epidemiology, understanding how the structure of a human contact network influences the propagation of a pathogen is crucial. Human networks are naturally assortative, which tends to concentrate infections within core communities but slows the global spread. Researchers can now utilize this trained model to take a sample contact network and smoothly transition its topology toward a disassortative state, generating hundreds of structurally

valid alternative realities to test how epidemic thresholds shift under different theoretical population mixing patterns. Similarly, in the realm of critical infrastructure, electrical power grids are typically disassortative to ensure that high capacity generator hubs supply low capacity consumer nodes without overloading redundant connections between generators. By generating synthetic power grids with precisely controlled assortativity variations, engineers can run sophisticated cascading failure simulations to identify optimal topological designs that maximize resilience against targeted cyber attacks or random mechanical failures. The ability to guarantee the preservation of the overarching scale free architecture and hierarchical branching while explicitly tuning these specific correlations eliminates the longstanding trade off between structural realism and experimental control. Furthermore, because the entire generative process is based on differentiable continuous manifolds, it opens the door to integrating more complex multi objective conditioning in the future, such as simultaneously targeting specific assortativity coefficients alongside targeted average path lengths or modularity scores.

5.2 Limitations and Concluding Remarks

Despite the remarkable capabilities demonstrated by the hyperbolic diffusion framework, several limitations necessitate further investigation. The primary bottleneck is computational complexity. Performing logarithmic and exponential map projections between the hyperbolic manifold and the local tangent spaces at every discrete time step of the reverse diffusion process requires significantly higher computational overhead compared to standard Euclidean models. Calculating the gradient of the expected assortativity proxy function also involves computing pairwise distance matrices across all generated points, resulting in quadratic time complexity that may hinder the immediate application of this method to massive networks comprising millions of nodes. Future research must explore approximation techniques for hyperbolic distances and highly optimized manifold differential equation solvers to accelerate the inference speed. Additionally, the current methodology relies on a single continuous time parameter to control the noise schedule. Exploring adaptive scheduling mechanisms that vary the noise injection rate based on the varying curvature of different regions within the Poincare disk could further refine the generative quality of highly asymmetric networks. In conclusion, the transition from discrete structural manipulation to continuous geometric generation represents the future of network modeling. The hyperbolic diffusion model firmly establishes that forcing hierarchical, relational data into flat, Euclidean constructs fundamentally limits generative potential. By embracing the native negative curvature of complex dependencies and pairing it with gradient guided thermodynamic generative processes, researchers are now equipped with an exceptionally powerful tool to decode, replicate, and manipulate the structural foundations of complex systems. This work not only broadens the theoretical scope of deep learning on manifolds but also provides immediate, actionable capabilities for practical network simulation across the sciences.

References

Zhao, R., Zeng, W., Tang, J., Li, Y., Ye, G., Du, J., & Zhao, X. (2025, May). Towards unsupervised entity alignment for highly heterogeneous knowledge graphs. In 2025 IEEE 41St international conference on data engineering (ICDE) (pp. 3792-3806). IEEE.

Zhao, R., Zeng, W., Zhang, W., Zhao, X., Tang, J., & Chen, L. (2025). Towards temporal knowledge graph alignment in the wild. arXiv preprint arXiv:2507.14475.

Wang, F., Zhang, Z., Sun, L., Ye, J., & Yan, Y. (2022, April). Dirie: Knowledge graph embedding with dirichlet distribution. In Proceedings of the ACM Web Conference 2022 (pp. 3082-3091).

Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In Proceedings of KDD.

Li, B., Zhang, R., Liang, H., Zhang, J., Zhang, J., Chen, X., ... & Wang, J. (2026). Interagent: Physics-based multi-agent command execution via diffusion on interaction graphs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 15253-15265).

Yao, Y., Lu, Z., Chen, G., Huang, C., Feng, C., & Quek, T. Q. (2026). Interference Management in ISAC-SAGINs Based on Transformer-Enabled Mean-Field Reinforcement Learning Method. IEEE Transactions on Wireless Communications.

Yao, Y., Sun, J., Miao, P., Chen, G., & Tafazolli, R. (2026). Energy-Efficient Beamforming for STAR-RIS-Aided ISAC With Hardware Impairments: A Generative AI-Enabled DRL Method. IEEE Transactions on Wireless Communications.

Yao, Y., Xiao, W., Miao, P., Chen, G., Yang, H., Chae, C. B., & Wong, K. K. (2026). UAV-RHS-enabled full-duplex ISAC covert system: Robust beamforming and trajectory optimization. IEEE Transactions on Communications.

Cong, P., Xu, Y., Ren, Y., Zhang, J., Xu, L., Wang, J., ... & Ma, Y. (2023, June). Weakly supervised 3d multi-person pose estimation for large-scale scenes based on monocular camera and single lidar. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 1, pp. 461-469).

Liu, Y., Cai, Q., Wang, C., Yang, J., Fan, H., Dong, J., & Chen, S. (2024). Geometry-enhanced attentive multi-view stereo for challenging matching scenarios. IEEE Transactions on Circuits and Systems for Video Technology, 34(8), 7401-7416.

Mo, Z. (2025, May). Development of an Intelligent Retrieval and Recommendation System for Chinese Language Educational Resources Based on DNN Technology. In 2025 2nd International Conference on Intelligent Computing and Robotics (ICICR) (pp. 170-175). IEEE.

Yu, X., Wei, Y., Zhou, S., Yang, Z., Sun, L., Peng, H., ... & Yu, P. S. (2025, April). Towards effective, efficient and unsupervised social event detection in the hyperbolic space. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 12, pp. 13106-13114).

Xiong, F., Xu, H., Wang, Y., Cheng, R., Wang, Y., & Chu, X. (2025, November). Hs-star: Hierarchical sampling for self-taught reasoners via difficulty estimation and budget reallocation. In Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (pp. 5539-5555).

Wei, M., Li, R., Wang, L., Chang, Z., Xu, L., & Han, Z. (2026). Toward Adaptive Tracking and Communication via an Airborne Maneuverable Bi-Static ISAC System. IEEE Transactions on Vehicular Technology.

Li, R., Xiao, Z., & Zeng, Y. (2023). Toward seamless sensing coverage for cellular multi-static integrated sensing and communication. IEEE Transactions on Wireless Communications, 23(6), 5363-5376.

Li, Q., Ye, Q., Zhang, N., Zhang, W., & Hu, F. (2025). Digital-twin-enabled industrial IoT: Vision, framework, and future directions. IEEE Wireless Communications, 32(6), 173-181.

Li, Y. Z., Zhang, J. P., Chen, Z. Y., Wang, H. C., Zhang, K. X., Hu, S. S., ... & Dudley, M. (2026, September). Growth and Characterization of High-Quality Thick Epitaxial 4H-SiC Wafers for High Voltage Devices. In Defect and Diffusion Forum (Vol. 452, pp. 79-85). Trans Tech Publications Ltd.

Li, S. (2025). Momentum, volume and investor sentiment study for us technology sector stocks—A hidden markov model based principal component analysis. PloS one, 20(9), e0331658.

Li, X., Wang, P., Li, G., Ni, L., & Zhang, Y. (2023). Design of interface circuits and lightweight PUF for TMR sensors. IEEE Sensors Journal, 23(11), 11754-11761.

Qu, W., Shao, Y., Meng, L., Huang, X., & Xiao, L. (2024, June). A conditional denoising diffusion probabilistic model for point cloud upsampling. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 20786-20795). IEEE.

Zhao, J., Zhang, C., Qin, M., & Yang, P. (2025). QuantFactor REINFORCE: mining steady formulaic alpha factors with variance-bounded REINFORCE. IEEE Transactions on Signal Processing, 73, 2448-2463.

Yang, Z., Liang, B., & Ji, W. (2021). An intelligent end–edge–cloud architecture for visual IoT-assisted healthcare systems. IEEE internet of things journal, 8(23), 16779-16786.

Jin, H., Tsai, F. S., & Martinez-Vazquez, J. (2025). Optimal expenditure decentralization for sustainable development: evidence from a 52-country panel analysis. Financial Innovation, 11(1), 111.

Xing, Z. (2026). A Synthesizable Bandwidth Monitoring Module for DDR4 Memory Controllers: RTL Design and Post-Implementation Analysis. Microelectronics Journal, 107410.

Yu, X., Zhang, Y., & Wang, H. (2025, July). Subjective Sleep Gains in Empty Nose Syndrome: A Cautionary Note on the Need for Objective Validation. In International Forum of Allergy & Rhinology (Vol. 15, No. 7, pp. 761-762).

Huang, Y., He, W., Wang, R., Jia, Y., Li, L., Xu, Y., ... & Li, C. (2025). A simple method for preparing chlorophyll free phenols from olive leaves and efficiently enriching it in refined olive oil. Food and Bioprocess Technology, 18(1), 979-994.