Symmetries as a guide for network hyperbolic embedding

S. Lizotte[1], J.-G. Young[1,2], A. Allard[1,2]

1. Université Laval, Québec, Canada
2. University of Vermont, Burlington, VT, USA

Network geometry is a versatile yet simple framework that captures several observed properties of empirical networks, such as non-vanishing clustering, sparsity, and power-law degree distribution [1]. This accurate description is achieved by placing vertices in a metric space (usually hyperbolic) and connecting them according to their proximity. To use this framework inferentially, one must find vertex coordinates that best reproduce the observed topology. Unfortunately, the posterior landscape of this model turns out to be non-convex, which makes the inference of the coordinates of the vertices a challenging task. Common embedding techniques circumvent this difficulty by using a combination of simplifying heuristics, greedy algorithms, and machine learning techniques [2–5].

By approaching the embedding task from a Bayesian perspective, we quantify the uncertainty of the inferred positions and parameters [6], thereby going beyond the pointwise estimates obtained with common embedding techniques. We find that naive Hamiltonian Monte Carlo (HMC)—the state-of-the-art method for sampling continuous random variables—fails to adequately explore the parameter space of this model due to a truly multimodal posterior distribution [7]. As depicted in the Figure, we overcome this technical challenge by improving the naive HMC algorithm with a Metropolis-Hastings (MH) kernel that leverages the natural symmetries of the hyperbolic space to perform adequate global transformations.

The resulting algorithm yields embeddings in proportion to their plausibility, which can be translated into margins of errors and thus enables the identification of multiple viable embeddings for the same graph. As opposed to pointwise estimates, uncertain embeddings prove valuable for deriving error bars on graph and geometric properties that ensue from an embedding. Furthermore, preliminary results suggest that our Markov transition kernel holds great potential for robust optimization when employed alongside a stochastic gradient descent algorithm, effectively trading error margins for pointwise estimation's computational efficiency.

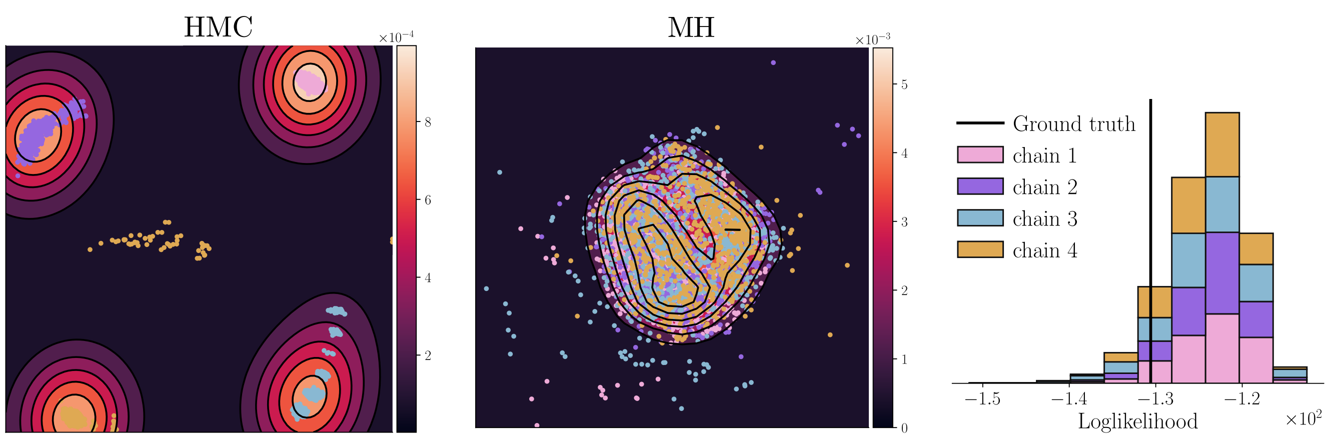

Figure: Multidimensional scaling (MDS) of the angular coordinates sampled from the posterior distribution (left) using HMC and (middle) using our MH algorithm. The right panel shows the loglikelihood values of the sample obtained with our algorithm. The MDS distances are computed as the sum of angular distances between each vertex. Each point on the plot represents an element of the posterior sample, with different colors indicating Markov chains initialized differently. The heatmap overlays the probability density of a Gaussian kernel density estimation. The posterior distribution visualized here encodes the plausible embeddings of a graph of 30 vertices generated from the $\mathbb{S}^1$ model, which is nearly identical to the two-dimensional hyperbolic model $\mathbb{H}^2$. Comparing the left and middle panels, it is evident that the HMC algorithm struggles to reliably move away from its initial position, while our algorithm achieves excellent mixing. The right panel further confirms the effectiveness of the MH algorithm by demonstrating that embeddings sampled using this approach generate the original graph with high probability.

[1] Nat. Rev. Phys. **3**, 114–135 (2021).
[2] New J. Phys. **21**, 123033 (2019).
[3] Phys. Rev. E **104**, 044315 (2021).
[4] Nat. Commun. **8**, 1615 (2017).
[5] IEEE/ACM Trans. Netw. **26**, 920–933 (2018).
[6] Nat. Commun. **13**, 6794 (2022).
[7] J. Stat. Softw. **76**, 1–32 (2017).