

Graph of Thrones : Adversarial Perturbations dismantle Aristocracy in Graphs

Adarsh Jamadandi ¹ and Uma Mudenagudi

Computer Vision and Graphics Lab,
KLE Technological University, India.


November 13, 2020




¹adarsh.jamadandi@kletech.ac.in

Graph Data


- 1 Graph Structured Data provide powerful representation.
- 2 Graphs are ubiquitous and can model heterogeneous data from varied fields -



Chemistry:
Molecules



Biology: Protein
Interaction Network



Social Network
(Facebook-Ego)


Protein interaction : <http://jeswcollins.github.io/PPI/>

Fb-Ego : <https://github.com/AnilOsmanTur/ComplexNetworksProjects>




How to learn Graph Structure?


- ① Real-world Graphs often exhibit properties such as -



Scale-Free
(Barabási-Albert
Graph)



Hierarchical
Anatomy




Causality (Citation
Network)

- ② The functional/semantic similarity between entities/nodes should be preserved by the embedding space [Nickel and Kiela(2017)].




Node Embeddings in Euclidean Space

- 1 Map graph nodes $v \in \mathcal{V}$ to low dimensional vector $z_v \in \mathbb{R}^n$.




- 2 Trying to embed graphs in Euclidean space - We quickly run out of space!




Note that, for increasing level of the binary tree, unrelated nodes are forced together, distorting the original tree structure.



Hyperbolic Space (\mathbb{H}^n)

- 1 Hyperbolic space offers exciting alternative to Euclidean space [Nickel and Kiela(2017)].
- 2 Hyperbolic space  - non-Euclidean, constant negative curvature.



Datasets with hierarchical structure can be embedded in low-dimensional hyperbolic space without distortions.

- 3 The space grows **exponentially**, the hierarchical nature of the data is preserved.



When should you use Hyperbolic Space?

- 1 [Gromov(1987)] introduced the notion of **hyperbolicity** (δ), to measure the **Tree-likeness** of a metric space.
- 2 Mathematically,

Let $\{a, b, c, d\}$ be the vertices of the graph $G(\mathcal{V}, \mathcal{E})$. and let $(\mathbb{S} = \{S_1 = d\langle a, b \rangle + d\langle d, c \rangle\}, \{S_2 = d\langle a, c \rangle, d\langle b, d \rangle\}, \text{ and } \{S_3 = d\langle a, d \rangle + d\langle b, c \rangle\})$. The $\delta(a, b, c, d)$ is given by


$$\delta(a, b, c, d) = \frac{1}{2} \max_{\{a,b,c,d\} \in \mathcal{V}(G)} \text{hyp}(a, b, c, d) \quad (1)$$

where, $\text{hyp}(a, b, c, d) =$ Difference of two largest values in \mathbb{S} .



δ – Hyperbolicity and Graph Aristocracy

- 1 A graph $G = (\mathcal{V}, \mathcal{E})$ can be viewed as a metric space with $d\langle \cdot, \cdot \rangle$ measuring the (geodesic) distance between vertices.
- 2 For $\delta \leq 0 \implies$ Hyperbolic also \implies Aristocratic.
- 3 Graph Aristocracy - Small set of vertices controlling the overall aspects of the network
[Borassi et al.(2015)Borassi, Chessa, and Caldarelli].




A subset of WordNet-Noun hierarchy [Miller and Fellbaum(1998)] with $\delta = 0.183$ (Left). A random DAG with $\delta = 6.5$ (Right.)

How Robust are Node Embeddings?

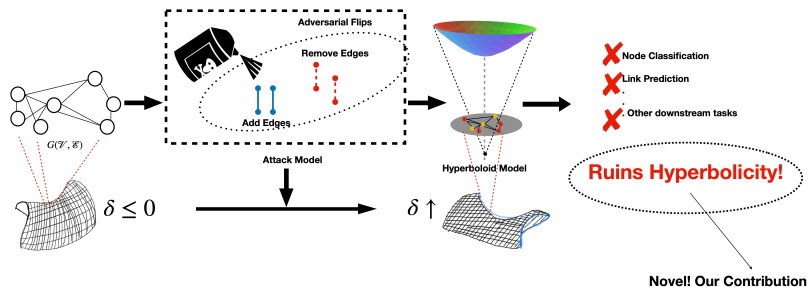
- 1 Adversarial attacks - Deliberate and Random Perturbations injected into the data, that affects the model's performance.

Attack Model [Bojchevski and Günnemann(2019)]



Adversarial Perturbations ruin δ -hyperbolicity


Our Contribution



Adversarial Perturbations ruin δ -hyperbolicity

We compute δ -hyperbolicity of standard graph data sets before and after introducing adversarial perturbations. Its evident from the Table below that hyperbolicity increases for random adversarial attacks rendering embedding in hyperbolic space less effective.

Dataset	# of Edge Flips	δ_{before}	δ_{after}
Cora [McCallum et al.(2000)McCallum, Nigam, Rennie, and Seymore]	1000	2.0	2.5
Citeseer [Giles et al.(1998)Giles, Bollacker, and Lawrence]	1000	2.5	3.0
Polblogs [Adamic and Glance(2005)]	1000	1.0	1.5



Visualization of standard datasets before (first row) and after introducing adversarial edge flips(second row).



How do we navigate adversarial vulnerabilities?

- 1 Lorentzian manifolds are **natural** embedding spaces for data exhibiting properties such as hierarchy and Causality.
- 2 **Lorentzian Manifold** : A Lorentzian manifold (\mathcal{L}, η) is a pseudo-Riemannian manifold equipped with metric signature $\{-, +, +, +, \dots\}$.

Lorentzian Inner Product

$$\langle x, y \rangle_{\mathcal{L}} = -x_0y_0 + x_1y_1 + x_2y_2 + \dots$$

Tangent Space and Vectors


Tangent vectors $v \in$ tangent space $T_x\mathcal{L}$ can be classified as -

- Time-like, if $\eta\langle x, x \rangle_{\mathcal{L}} < 0$
- Light-like, if $\eta\langle x, x \rangle_{\mathcal{L}} = 0$
- Space-like, if $\eta\langle x, x \rangle_{\mathcal{L}} > 0$



Lorentzian Manifolds and Posets

- 1 Let $\mathcal{C} = \{c_i\}_{i=1}^m$, be the set of concepts. Inferring concept hierarchies involves defining a partial order set (\mathcal{C}, \preceq) over the elements of \mathcal{C} - [Nickel and Kiela(2018)].
- 2 Surprisingly, a Lorentz manifold (\mathcal{L}, η) equipped with a Causal structure also forms a partial ordered set (\mathcal{L}, \prec) - [Zeeman(1964)].



Node Embeddings in Lorentzian Manifolds

- 1 We choose the simplest Lorentzian manifold - Minkowski Space \mathbb{M} equipped with metric -

$$d_{\mathbb{M},i,j} = -c^2(x_i^0 - x_j^0)^2 + \sum_{k=1}^d (x_i^k - x_j^k)^2 \quad (2)$$

- 2 The Minkowski spacetime (\mathbb{M}) consists of d spatial dimensions and 1 time dimension. $\{x_i^0, x_j^0\}$ represent the time co-ordinates and $\{x_i^k, x_j^k\}$ represent spatial co-ordinates and c is the speed of light, which indicates the speed of flow of information in this case.
- 3 We embed graphs using Lorentzian-MDS proposed by authors in [Clough and Evans(2017)].



Lorentzian-Multidimensional Scaling (MDS)

- 1 Given points $X = \{x_1, x_2, x_3, \dots, x_n\} \in \mathbb{M}^d$, expressed as a matrix $X \in \mathbb{R}^{n \times r}$.
- 2 We have access to the pair-wise distances $d_{i,j} = d_{\mathbb{M}_{i,j}}$ and not X .


Lorentzian-MDS

Recover $X = \{x_1, x_2, x_3, \dots, x_n\}$, by observing $d_{i,j}$.




Results

- 1 We compare by embedding popular graph datasets - Citeseer, PolBlogs and Cora datasets in both Poincaré disk model and Minkowski space.




Before attack. AUC = 0.63.




After attack. AUC = 0.62.



Results




Before attack. AUC = 0.81.




After attack. AUC = 0.69.

Results




Before attack. AUC = 0.74.




After attack. AUC = 0.64.

Results




Before attack. AUC = 0.74.



After attack. AUC = 0.64.

Results



Lorentzian-MDS algorithm breaks down when the embedding space is chosen to be the Poincaré ball model.



Summary

- 1 We study adversarial perturbations in the context of geometric graph learning for the first time.
- 2 Empirically, we show that unsupervised embeddings in hyperbolic space are susceptible to adversarial attacks, making the embedding space less effective.
- 3 We quantify the inefficacy of hyperbolic spaces by measuring the Gromov hyperbolicity.
- 4 We advocate for the utility of Lorentzian manifolds for learning hierarchical data, as they are more robust to adversarial perturbations.





L. A. Adamic and N. Glance.

The political blogosphere and the 2004 u.s. election: Divided they blog.

In Proceedings of the 3rd International Workshop on Link Discovery, LinkKDD '05, page 36–43, New York, NY, USA, 2005. Association for Computing Machinery.

ISBN 1595932151.

doi: 10.1145/1134271.1134277.

URL <https://doi.org/10.1145/1134271.1134277>.



A. Bojchevski and S. Günnemann.

Adversarial attacks on node embeddings via graph poisoning.

In Proceedings of the 36th International Conference on Machine Learning, ICML, Proceedings of Machine Learning Research. PMLR, 2019.



M. Borassi, A. Chessa, and G. Caldarelli.

Hyperbolicity measures democracy in real-world networks.

Phys. Rev. E, 92:032812, Sep 2015.





J. R. Clough and T. S. Evans.

Embedding graphs in lorentzian spacetime.

PLOS ONE, 12(11):1–14, 11 2017.

doi: [10.1371/journal.pone.0187301](https://doi.org/10.1371/journal.pone.0187301).

URL <https://doi.org/10.1371/journal.pone.0187301>.



C. L. Giles, K. D. Bollacker, and S. Lawrence.

Citeseer: an automatic citation indexing system.

In INTERNATIONAL CONFERENCE ON DIGITAL LIBRARIES,
pages 89–98. ACM Press, 1998.



M. Gromov.

Hyperbolic Groups, pages 75–263.

Springer New York, 1987.



A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore.

Automating the construction of internet portals with machine learning.

Information Retrieval, 3(2):127–163, 2000.

doi: [10.1023/A:1009953814988](https://doi.org/10.1023/A:1009953814988).





G. Miller and C. Fellbaum.

Wordnet: An electronic lexical database.
1998.



M. Nickel and D. Kiela.

Poincaré embeddings for learning hierarchical representations.
In Advances in Neural Information Processing Systems 30. Curran Associates, 2017.



M. Nickel and D. Kiela.

Learning continuous hierarchies in the lorentz model of hyperbolic geometry.
Proceedings of the 35th International Conference on Machine Learning, 2018.



E. C. Zeeman.

Causality implies the lorentz group.
Journal of Mathematical Physics, 5(4):490–493, 1964.
doi: 10.1063/1.1704140.
URL <https://doi.org/10.1063/1.1704140>.

