Paying Attention to curvature on surfaces with Graph Convolutions

Baskaran Sripathmanathan, Department of Engineering, University of Oxford

EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines & Systems

Introduction & Applications

Curved surfaces are important in climate science and predicting protein-protein interactions [1]. Sharp curvature sometimes denotes a 'break', like two sides of a folded piece of paper, and sometimes doesn't, like on a key or on certain proteins. As this depends on the task, we ask: how can we learn to use curvature in a task-dependent manner?

Convolutions on Surfaces

Flat paper maps of the world cannot represent the sizes of countries without distortion because curved surfaces cannot be projected onto flat surfaces without distortion. This is the same reason we don't want to flatten surfaces and use standard convolutional neural networks (CNNs).

Imagine a wind blowing around the planet at a constant speed. In 3D, these vectors all look different. If we understand these vectors as living on planes tangent to the surface, and make our convolution translates to its current plane via *parallel transport* (in this case, projection onto the current plane), we get much simpler weights.

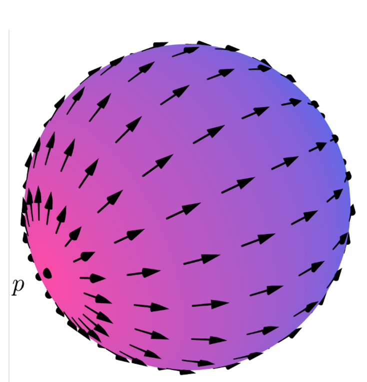

Why pay Attention to curvature?

We discretise all our surfaces to mesh graphs and talk about the graphs.

Attention is a mechanism allows us to learn to reweight the importance of certain edges on a graph. Incorporating curvature into Attention allows for learned curvature-based segmentation.

We try to explicitly precompute construct bases for the tangent space at each point on the surface and use these bases to calculate curvature. This should give significantly better results when the graph isn't segmented evenly with regards to curvature.

On earth, it may be easier to get better measurements near population centers, or data near population centers may be noisier, depending on the task. Regardless, collection of data is irregular and noisy over the surface of the planet. Curvature can act both as a proxy for distance and capture geographical features such as mountain ranges.

Sverrisson et al. [1] takes a point cloud of proteins, infers normals to the surface and uses the curvature to avoid connecting two sides of a bend when using k-nearest neighbours to make the mesh. This more faithfully recovers the shape of complex, bendy parts of the protein surface.

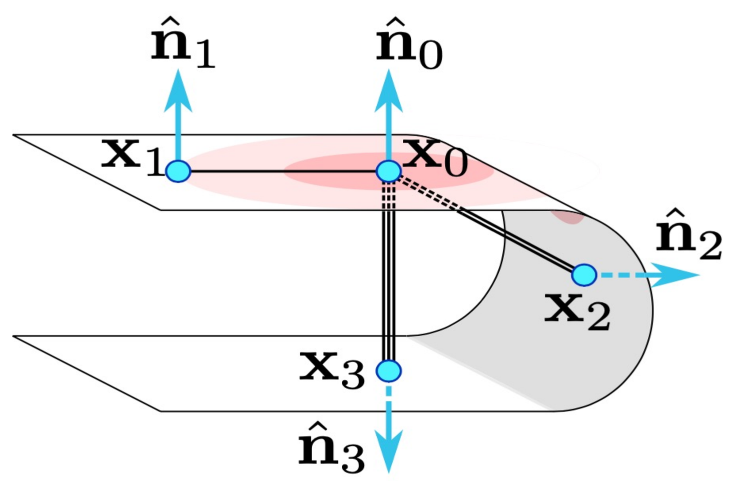

Discussion

Geometric deep learning already has exciting results in protein-protein interactions and drug discovery. While it's obvious how to apply this work to physical surfaces, this would also be interesting for social networks, which can be embedded into hyperbolic space, and for cases where the data has separate symmetries to the base surface – for example colour data on a surface.

Literature cited

1. Sverrisson, Freyr, et al. "Fast end-to-end learning on protein surfaces." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021.
2. 2. Cohen, Taco, et al. "Gauge equivariant convolutional networks and the icosahedral CNN." International Conference on Machine Learning. PMLR, 2019.
3. Singer, Amit, and H-T. Wu. "Vector diffusion maps and the connection Laplacian." Communications on pure and applied mathematics 65.8 (2012)

Acknowledgments

I thank Xiaowen Dong and Michael Bronstein for supervising me, and Francesco Di Giovanni and Ben Chamberlain for their input. This work was supported by the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines & Systems (EP/S024050/1).

Engineering and Physical Sciences Research Council