

Kernel Latent Space Models for understanding neural connectomes

Eric Jonas and Sрни Turaga

jonas@eecs.berkeley.edu

turagas@janelia.hhmi.org



Abstract

It has now become possible to map the synaptic connectivity of neural circuitry at the cellular resolution using electron microscopy [1]. In this work, we present a new class of models for the analysis of connectomic data. Many theories of neural computation propose specific patterns of neural connectivity tied to the tuning properties of neurons. We propose an extension to traditional latent space models [2] to uncover continuous hidden structure in these connectomes, such as the neural tuning property of a neuron and the function that determines neural connectivity. Our scalable model provides the flexibility to recover structure in both directed and undirected graphs. We demonstrate our model on synthetic connectomes and on the recently published mouse retinal connectome.

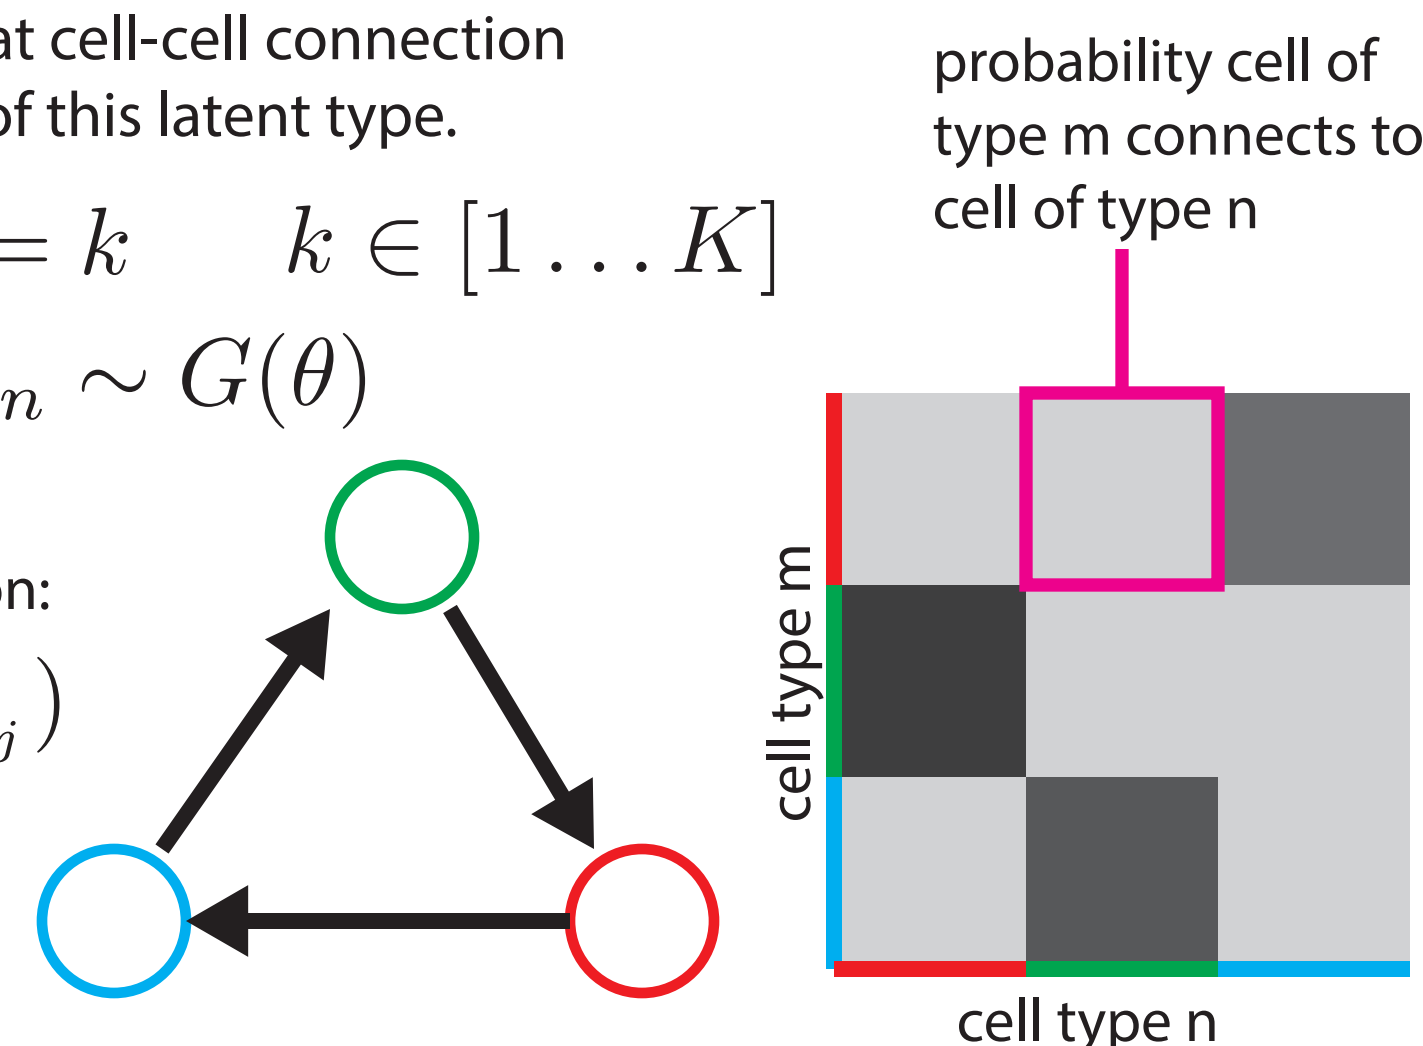
Discrete Cell Types Via Block Models

Classic stochastic block [3] models assume each cell is a discrete $[1..K]$ type and that cell-cell connection probabilities are a function of this latent type.

cell-type assignment: $c_i = k \quad k \in [1..K]$

Type-type connection probability: $\eta_{mn} \sim G(\theta)$

Total probability of connection: $p(R_{ij}) = f(\eta_{c_i, c_j})$



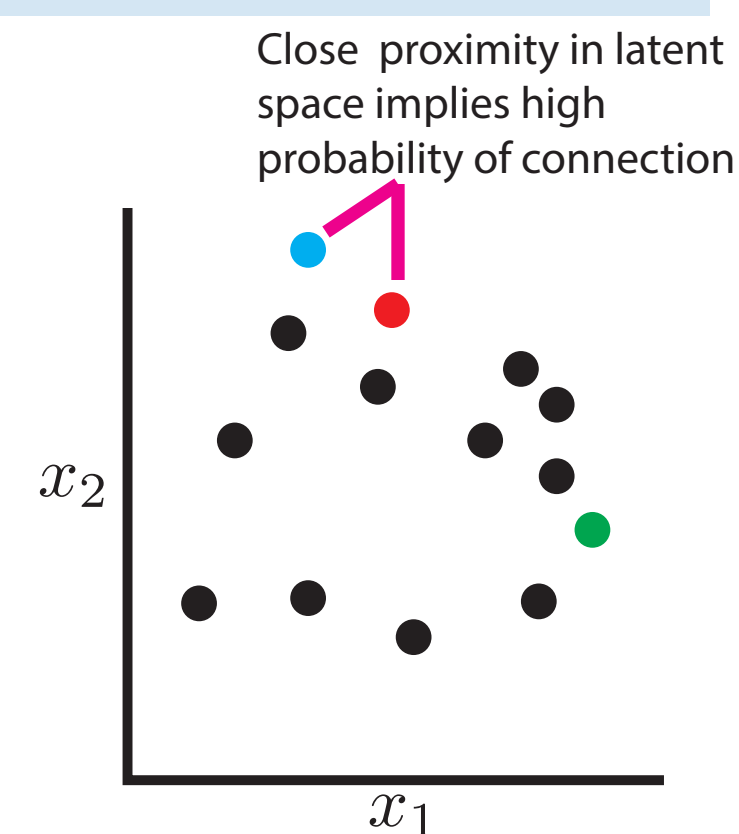
Nonparametric stochastic blockmodels have previously been successfully applied to the mouse retina connectome [4].

Latent Space and Continuous Type

Classic latent space models [2] assume that the probability of a connection between two cells as a function of their position in a continuous latent space.

$p(R_{ij}) = f(\|\mathbf{x}_i - \mathbf{x}_j\|^2) \quad \mathbf{x}_i \in \mathbb{R}^d$

Great for modeling "communities" where you have groups with a large number of intra-group connections but fewer inter-group connections.




Latent Space Models with Kernels

We can replace the Euclidean distance in a latent space model with an arbitrary kernel function.

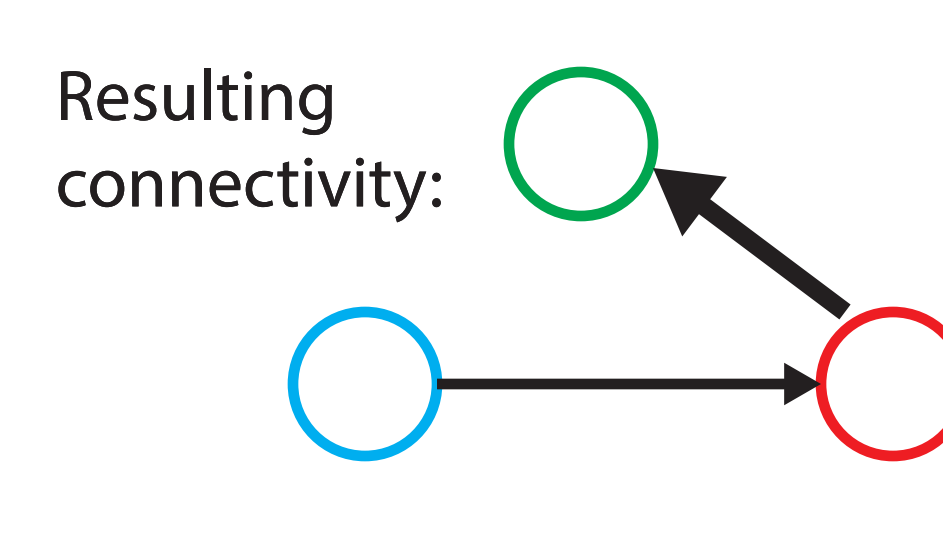
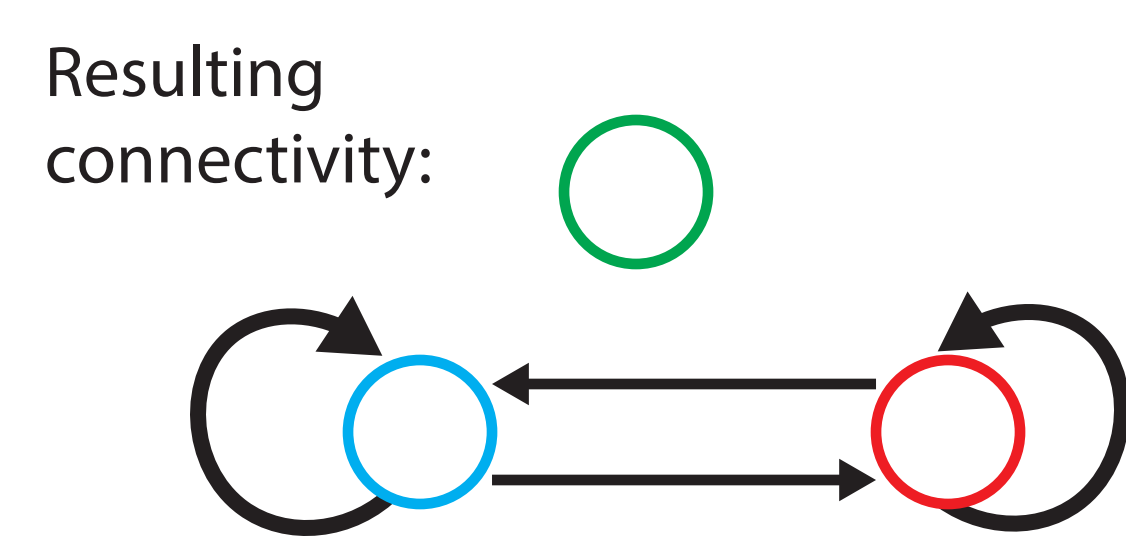
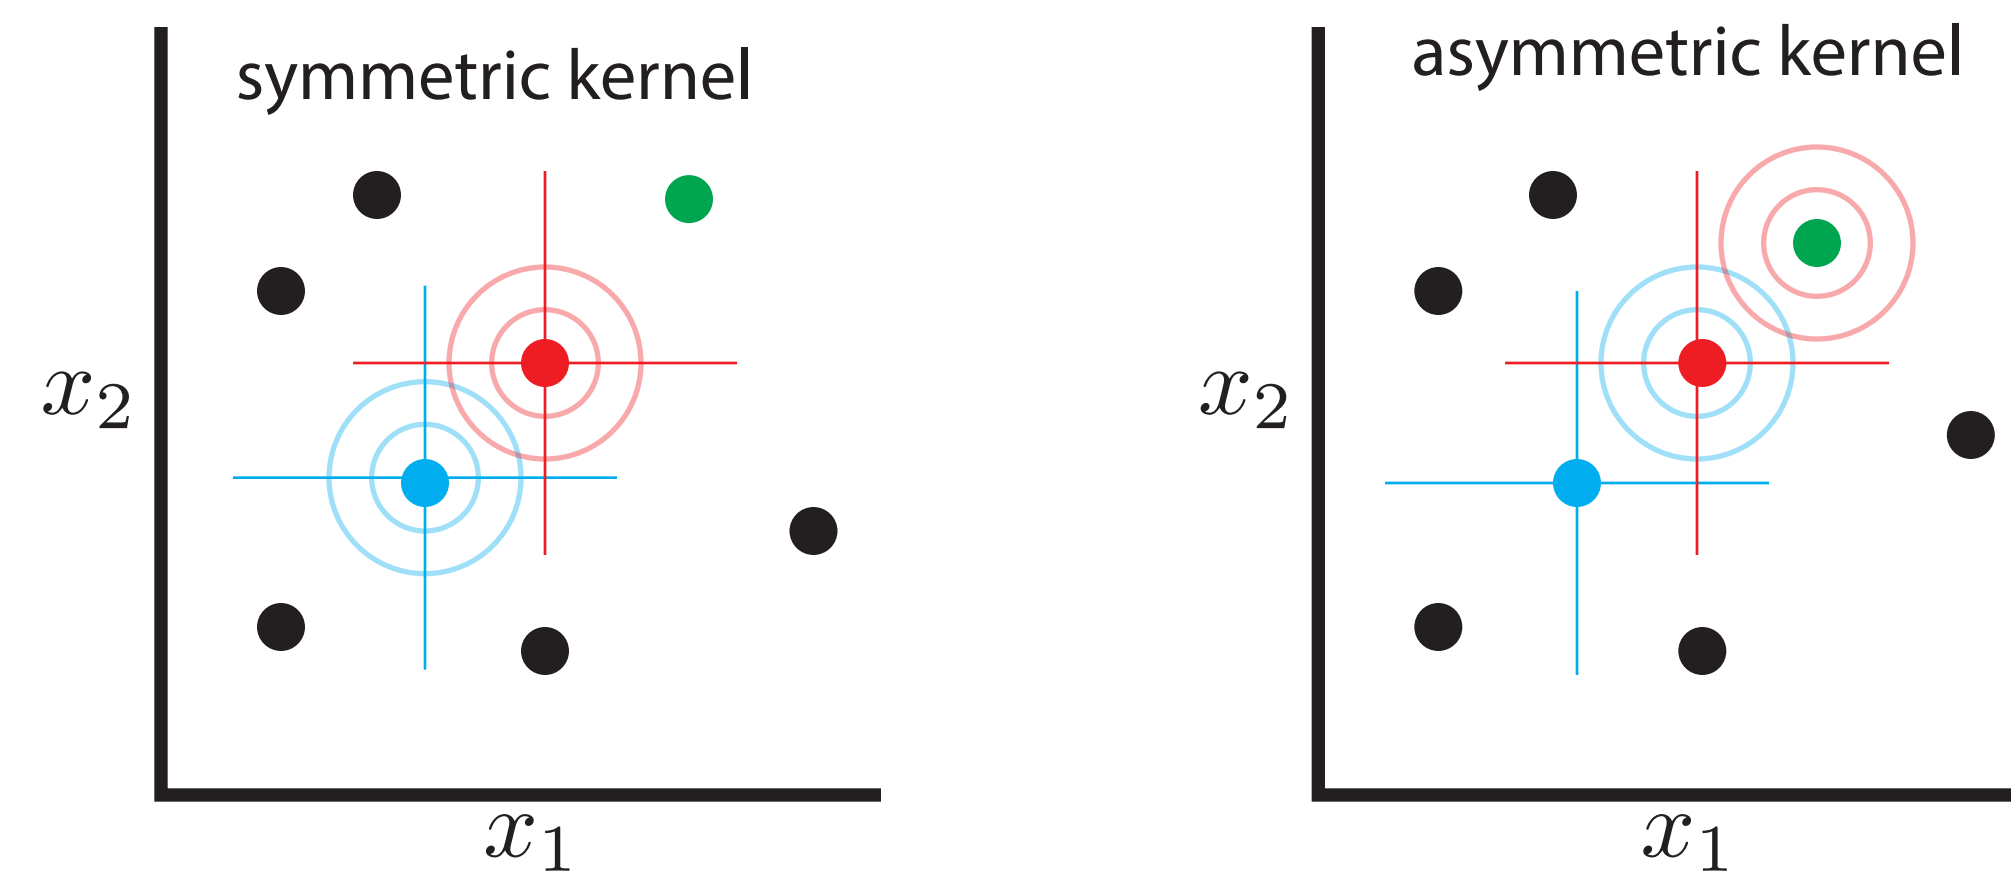
$R_{ij} \propto \kappa(\mathbf{x}_i, \mathbf{x}_j) \quad \mathbf{x}_i \in \mathbb{R}^d$

Kernel functions can be asymmetric, non-monotonic, non-isotropic



Asymmetric Kernels

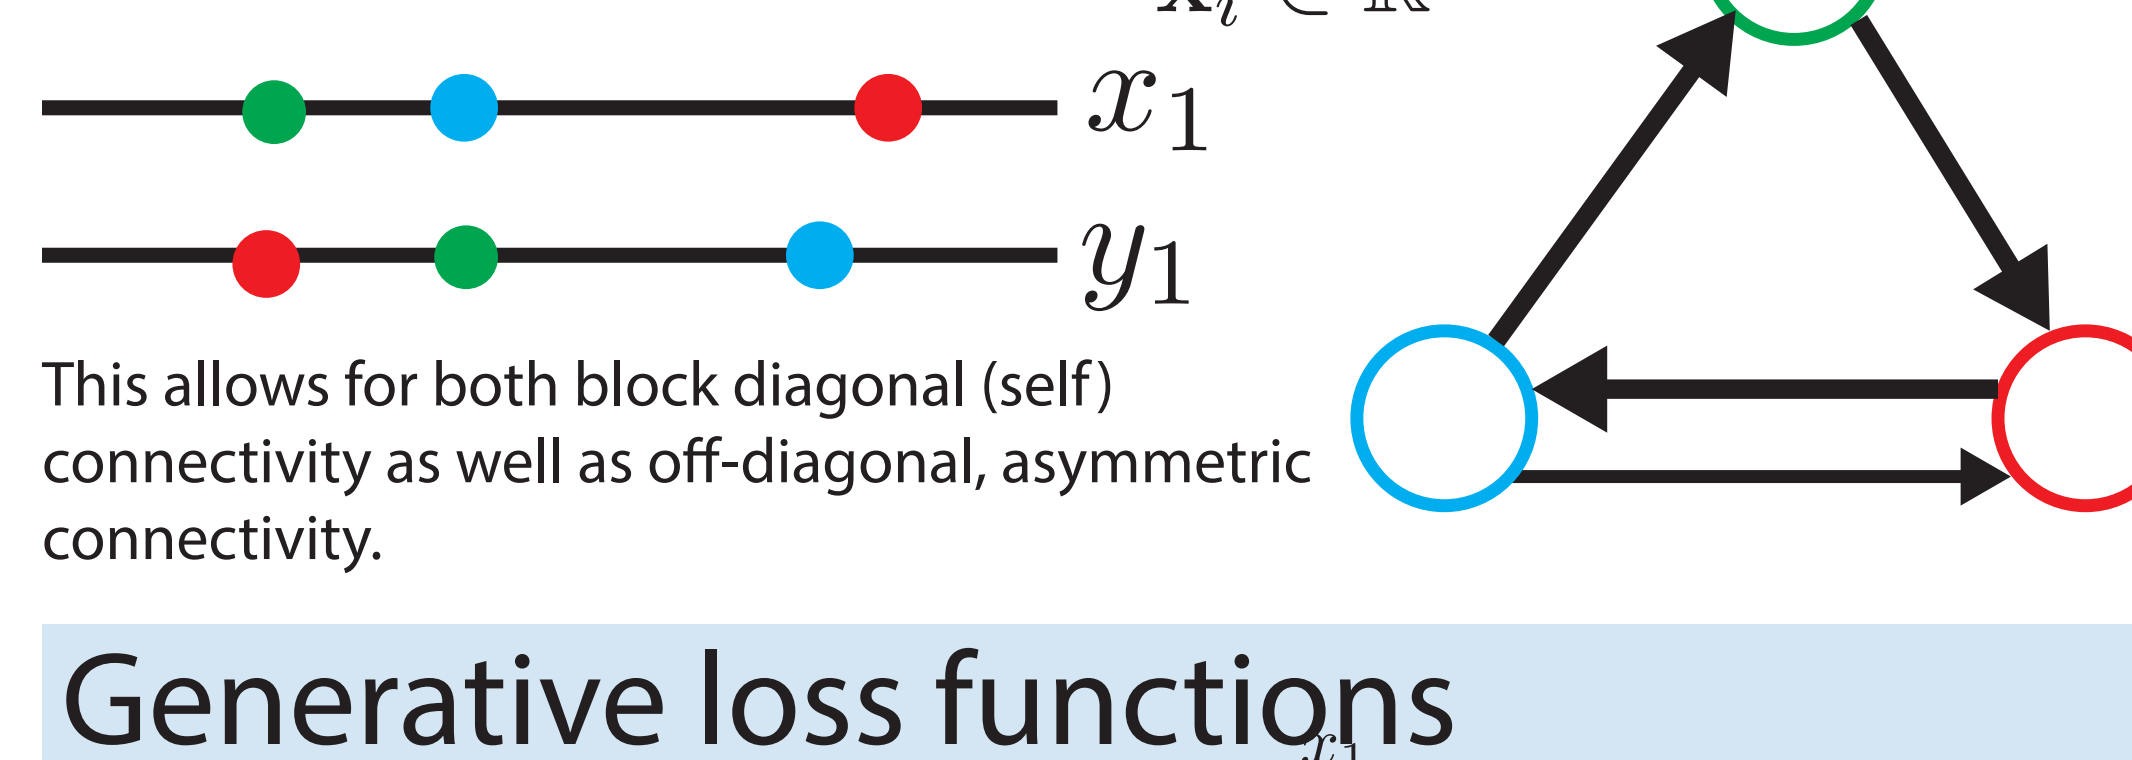
Symmetric kernels privilege self-self connections over all others. This does not accurately model most neural circuits, so we turn to asymmetric kernels whose peak is offset from zero. By breaking the isotropy of kernel space, we allow for directionality and non-block-diagonal structure.



Two-Space Kernels

We can split the latent space into two different spaces where the probability of connection is related to the distance between cell i in the first space and cell j in the second space.

$p(R_{ij}) = k(\mathbf{x}_i, \mathbf{y}_j) \quad \begin{matrix} \mathbf{y}_j \in \mathbb{R}^d \\ \mathbf{x}_i \in \mathbb{R}^d \end{matrix}$



This allows for both block diagonal (self) connectivity as well as off-diagonal, asymmetric connectivity.

Generative loss functions

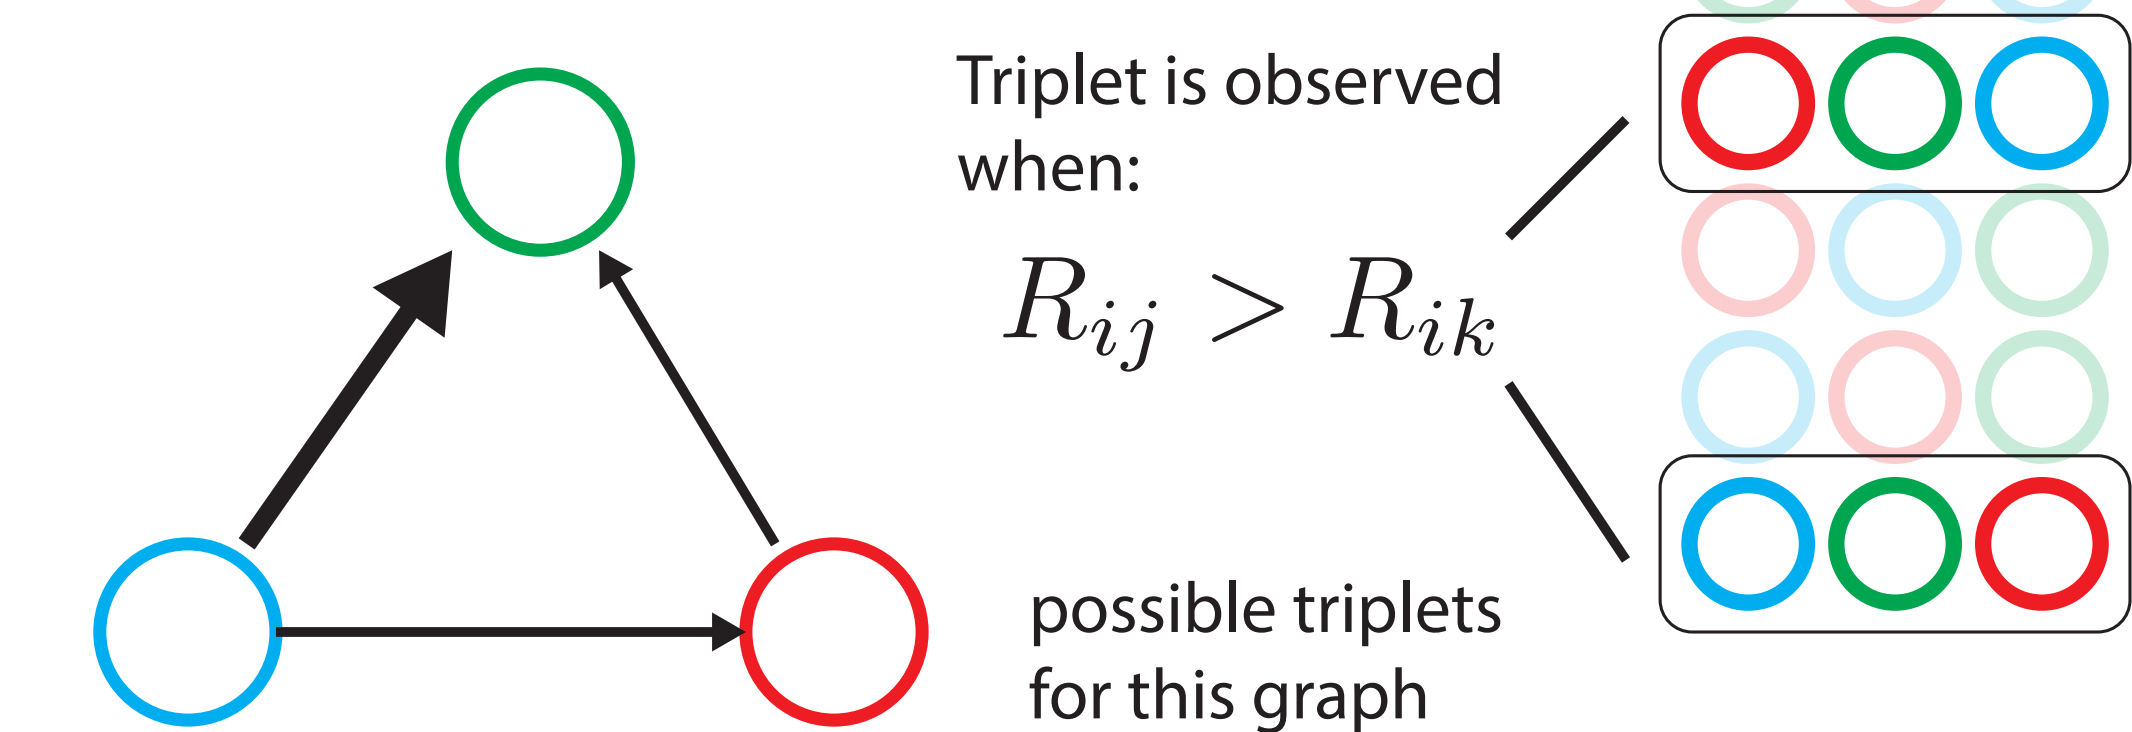
We can explicitly model the probability of a connection via a parametric, generative function such as a Poisson distribution or via a Bernoulli-Logistic function. Finding the right parametric form and hyperparameter sensitivity continue to be issues.

$P(R_{i,j}) \propto \text{Logistic}(\kappa(\mathbf{x}_i, \mathbf{x}_j)) \quad P(R_{i,j}) \propto \text{Poisson}(\kappa(\mathbf{x}_i, \mathbf{x}_j))$

Ranking loss

Modern connectomics methods can provide estimates of synaptic strength for both directed and undirected models, but the underlying (parametric) distribution of synapses is unknown. We adopt a ranking loss function [6] to fit our connectivity

$p_{ijk} \propto \frac{\kappa(\mathbf{x}_i, \mathbf{x}_j)}{\kappa(\mathbf{x}_i, \mathbf{x}_j) + \kappa(\mathbf{x}_i, \mathbf{x}_k)}$




Fitting models

We fit our non-convex models using stochastic gradient descent, using Theano for symbolic (exact) evaluation of the gradient. Adagrad is used when the stepsize could not otherwise be determined.

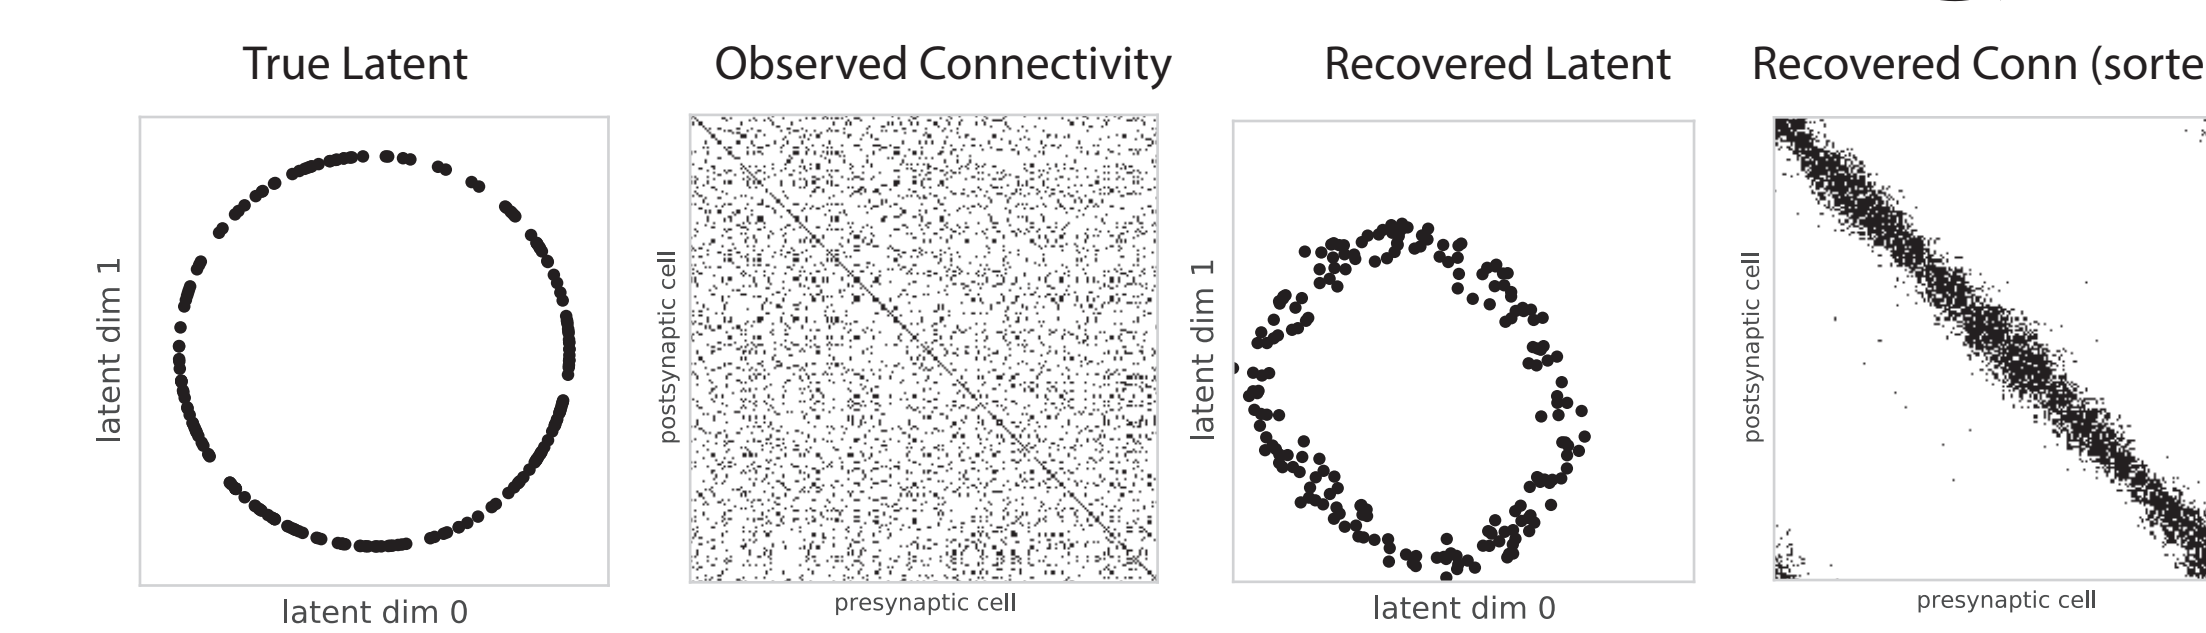
Synfire Chain

Linear chain structured latent space and connection matrix corresponding to a synfire chain. Again, our latent-space embedding discovers the 1-dimensional structure of the latent space, and the asymmetric, upper block-diagonal structure of the synfire chain.



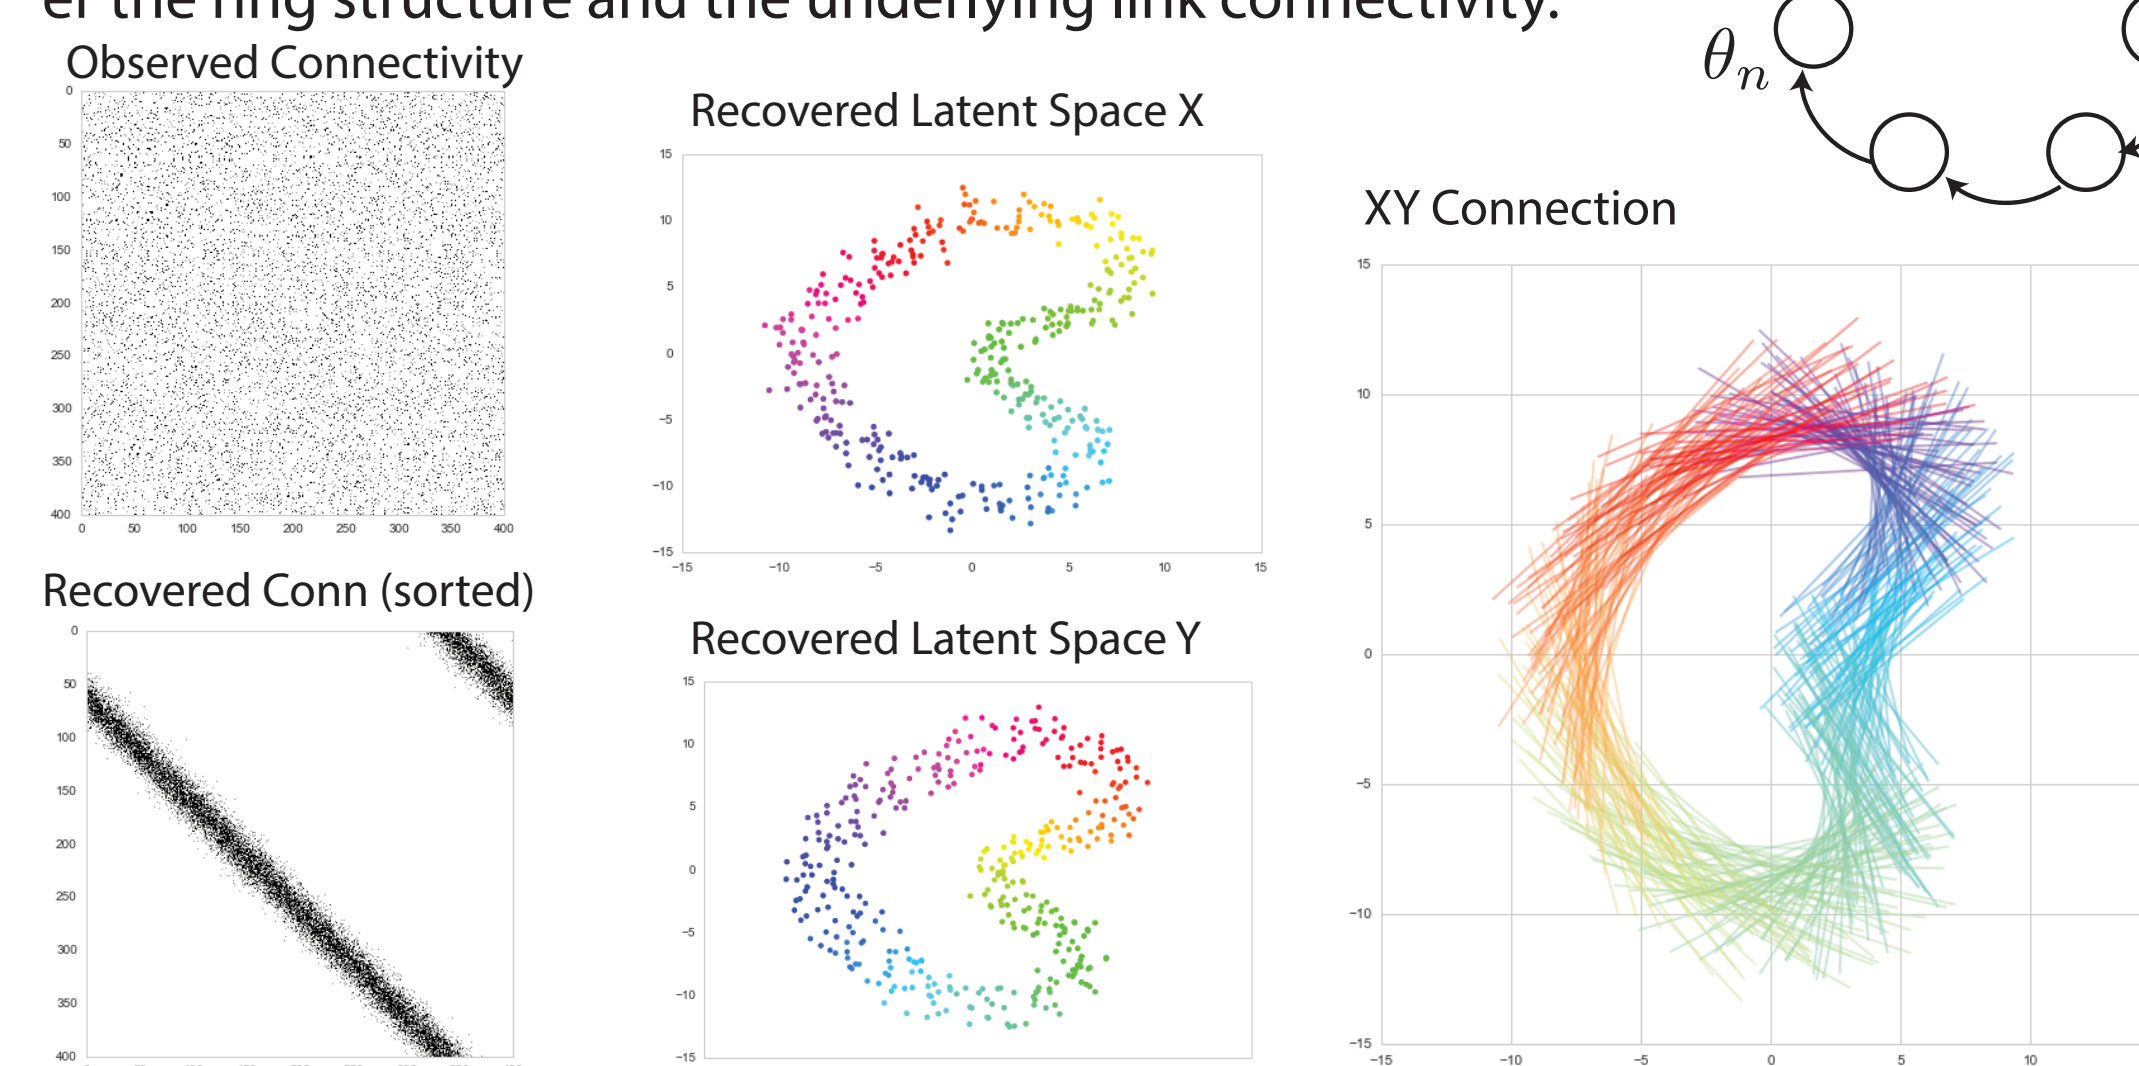
Symmetric Bidirectional Ring

Ring-shaped latent space and connection matrix corresponding to the classic ring attractor model for orientation tuning in visual cortex [5]. Our model accurately recovers the ring structured latent space from the connection matrix. Re-ordering the weight matrix using the latent structure reveals symmetric like-to-like connectivity preference.



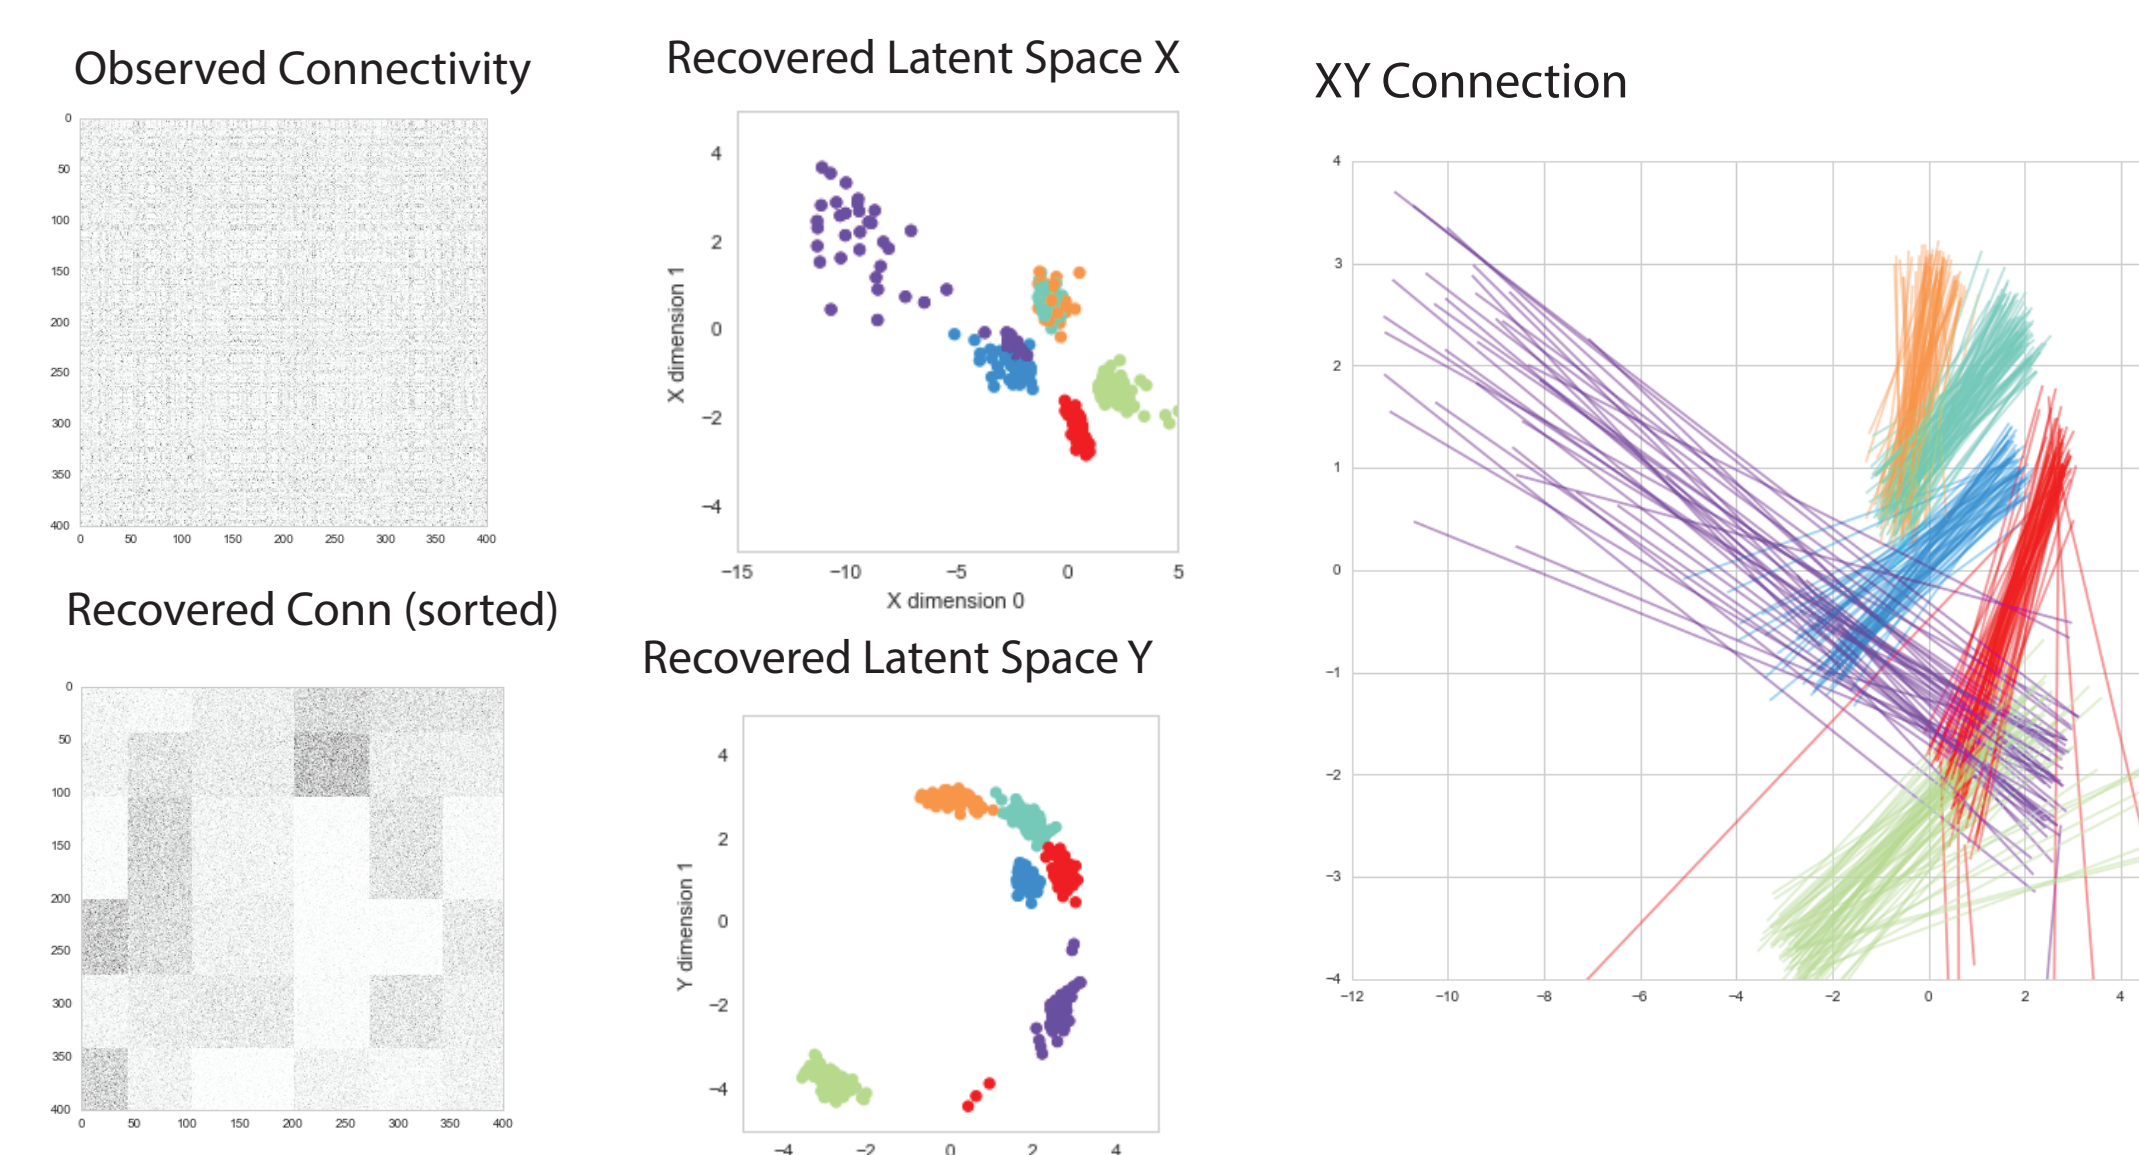
Asymmetric Ring

An asymmetric ring has each cell directionally connected to some of its neighbors in a consistent (here, clockwise) direction. By using a split 2D space with triplet loss, we can recover the ring structure and the underlying link connectivity.




Asymmetric Block Model

The asymmetric block model assumes the existence of discrete latent types and connection probabilities depending solely on those types. By using a triplet loss and 2d split space embedding, we can recover the clustered cell types.



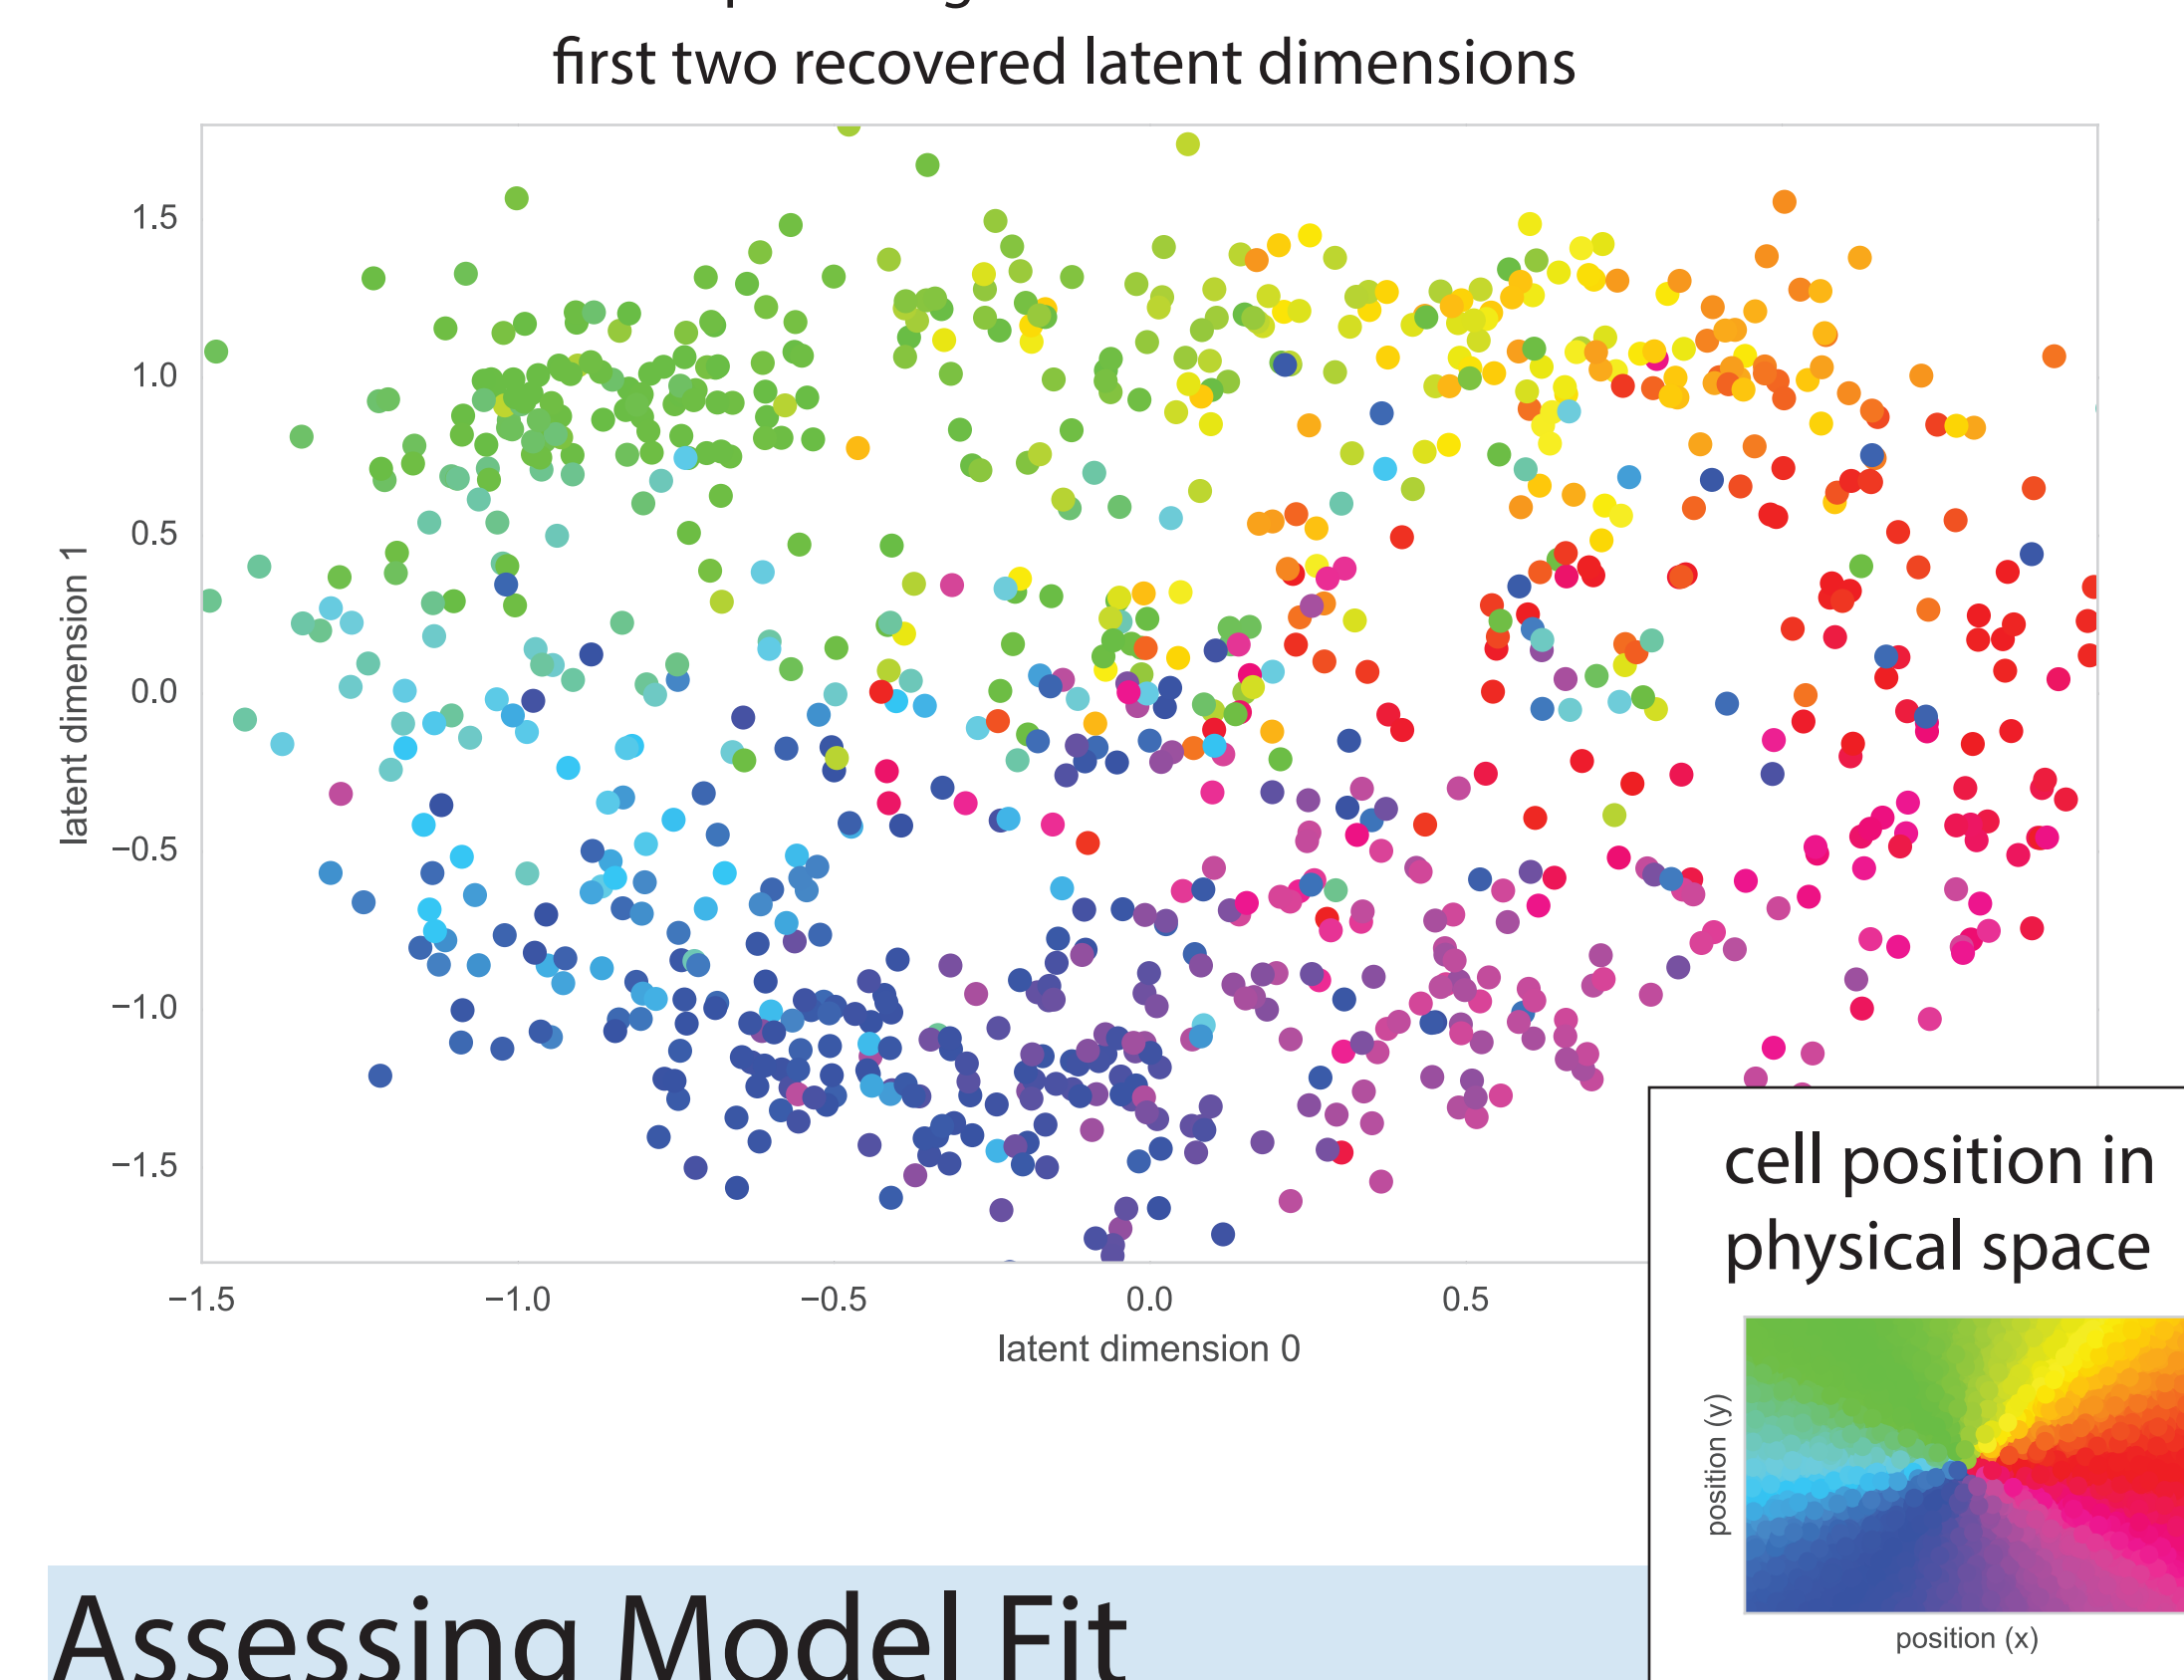
Mouse Retina Connectome

Dense serial electron microscopy of a $5114\mu\text{m} \times 80\mu\text{m}$ area in the mouse retina [1] yielded a listing of places where neurons came into contact. There were over 1000 cells originally, and selected the 950 for which the location of the soma could be reconstructed. Ultimately this left a matrix between the total synapse-like contact area between all pairs of 950 cells. Area was thresholded at $0.1\mu\text{m}$, determined by hand, to yield a 950×950 entry matrix that served as input to our algorithm.



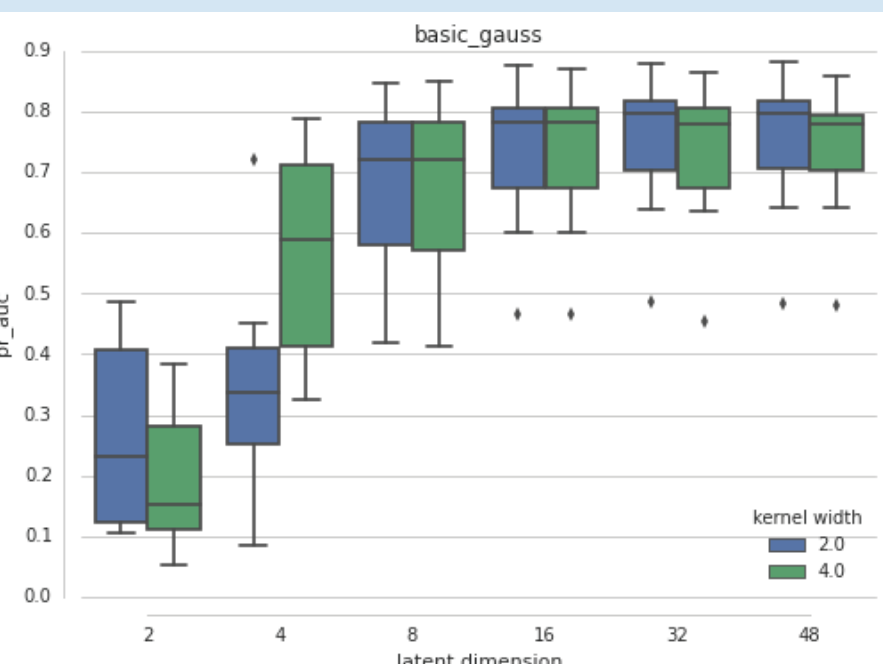
Latent Embedding of Retina

We fit a $D=8$ latent kernel single-space model to the mouse retina connectome with a symmetric rational quadratic kernel. The first two dimensions recover the intralaminar spatial organization of cells in the retina.



Assessing Model Fit

We can hold out connections from the training set and predict the missing connections and compute the area under the resulting precision-recall curve to assess model fit.



Next Steps

Early results are promising. Our next goal is to try a wider variety of kernels and space configurations on real connectomes. While the triplet embedding has shown promise, we need better ways of assessing model fit.

References

- [1] Helmstaedter, M., Briggman, K. L., Turaga, S. C., Jain, V., Seung, H. S., & Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature*, 500(7461), 168–174. doi:10.1038/nature12346
- [2] Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent Space Approaches to Social Network Analysis. *Journal of the American Statistical Association*, 97(460), 1090–1098. doi:10.1198/016214502388618906
- [3] Nowicki, K., & Snijders, T. a. B. (2001). Estimation and Prediction for Stochastic Blockstructures. *Journal of the American Statistical Association*, 96(455), 1077–1087. doi:10.1198/016214501753208735
- [4] Jonas, E., & Kording, K. (2015). Automatic discovery of cell types and microcircuitry from neural connectomics. *eLife*, 4, e04250. doi:10.7554/eLife.04250
- [5] Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 92(9), 3844–3848. doi:10.1073/pnas.92.9.3844
- [6] van der Maaten, L., & Weinberger, K. (2012). Stochastic triplet embedding. 2012 IEEE International Workshop on Machine Learning for Signal Processing, 1–6. doi:10.1109/MLSP.2012.6349720