

Graph Structure Learning **with** Interpretable Bayesian Neural Networks

Max Wasserman
Dept. of Computer Science
University of Rochester
mwasser6@ur.rochester.edu
github.com/maxwass

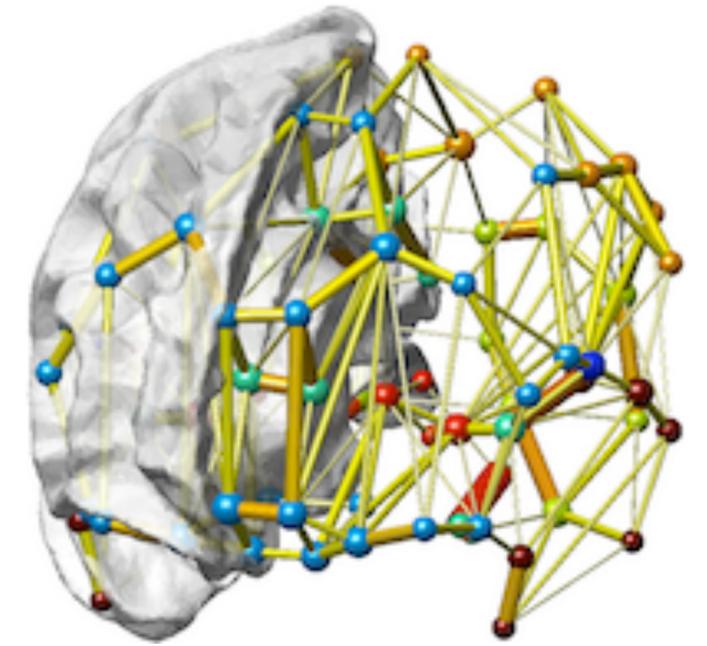
Collaborator: Gonzalo Mateos

Delft, Netherlands, June 25, 2024



What is this talk about?

Edge Predictions with Uncertainty Quantification



- Traditional GSP: study signals/filters with known **graph**
 - But often **graph not observed**
- Graph Structure Learning (GSL): **Learning graphs** from nodal observations. 2 Main Approaches.
 1. Model Based
 - Solve an **optimization problem** [Friedman'08], [Kalofolias'16], [Saboksayr'21]
 2. Unrolling Based
 - Constructs **deep network** using Model Based solution iterations [Pu'21], [Wasserman'22]
- Both only provide **point** estimates of graph structure
- **Goal:** **point** & **uncertainty** estimates of graph structure from nodal observation

Talk Outline

Edge Predictions with Uncertainty Quantification

- Develop a **point estimate function**
 - Pose and solve **inverse** optimization problem to estimate **graph structure** from **nodal observations**
 - ‘Unroll’ solution iterations to form a deep network
- Make it **Bayesian**
 - Parameter **priors**, **inference**, derive predictive **point** & **uncertainty** estimates over unobserved edges

Graph Signal Processing

Notation & Background

- Given graph \mathcal{G} with adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$
- Collect node signals $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{N \times p}$, where $\bar{\mathbf{x}}_i^\top$ denotes its i -th row.
- Form Euclidean Distance Matrix $\mathbf{E} \in \mathbb{R}^{N \times N}$, where $E_{ij} := \|\bar{\mathbf{x}}_i^\top - \bar{\mathbf{x}}_j^\top\|^2$
- Work with undirected graphs without self-loops. Reduce dimensionality.
 - $\mathbf{a} = \text{vec}[\text{triu}[\mathbf{A}]] \in \mathbb{R}^{N(N-1)/2}$, $\mathbf{e} = \text{vec}[\text{triu}[\mathbf{E}]] \in \mathbb{R}^{N(N-1)/2}$
- Total Variation of \mathbf{X} w.r.t. $\mathcal{G} := \text{Trace}(\mathbf{X}^\top \mathbf{L} \mathbf{X}) = \|\mathbf{A} \circ \mathbf{E}\|_1 = 2\mathbf{a}^\top \mathbf{e}$ [Kalofolias'16]

Smoothness \longleftrightarrow Sparsity!

Point Estimate Approaches via Optimization

SOTA Convex Formulation

Nodal Degrees
Vector: $\mathbf{A}\mathbf{1} = \mathbf{S}\mathbf{a}$

$$\mathbf{a}^*(\mathbf{e}, \alpha, \beta) = \arg \min_{\mathbf{a} \in \mathbb{R}^{N(N-1)/2}} \left\{ \underbrace{2\mathbf{a}^\top \mathbf{e}}_{\text{Data Fidelity}} - \underbrace{\alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{a})}_{\text{Regularizer}} + \underbrace{\frac{\beta}{2} \|\mathbf{a}\|_2^2}_{\text{Regularizer}} + \underbrace{\mathbb{I}\{\mathbf{a} \geq 0\}}_{\text{Constraint}} \right\},$$

Data Fidelity
Smoothness

Regularizer
Isolated Nodes

Regularizer
Small Edge

Constraint
Edge Non-negativity

We can **reparameterize** from (α, β) to (θ, δ)

[Kalofolias'16]

$$\begin{aligned} \mathbf{a}^*(\mathbf{e}, \alpha, \beta) &= \sqrt{\frac{\alpha}{\beta}} \mathbf{a}^*\left(\frac{1}{\sqrt{\alpha\beta}} \mathbf{e}, 1, 1\right) = \delta \mathbf{a}^*(\theta \mathbf{e}, 1, 1) \\ &= \delta \arg \min_{\mathbf{a} \in \mathbb{R}^{N(N-1)/2}} \left\{ 2\theta \mathbf{a}^\top \mathbf{e} - \mathbf{1}^\top \log(\mathbf{S}\mathbf{a}) + \frac{1}{2} \|\mathbf{a}\|_2^2 + \mathbb{I}\{\mathbf{a} \geq 0\} \right\} \end{aligned}$$

Point Estimate Approaches via Optimization

DPG's Independent Interpretability

Algorithm 1 Dual Proximal Gradient Descent

Inputs: Fixed parameters $\theta, \delta \in \mathbb{R}$ and data e

Initialize: a_0 and λ_0 at random.

for $k = 1, 2, \dots$ **do**

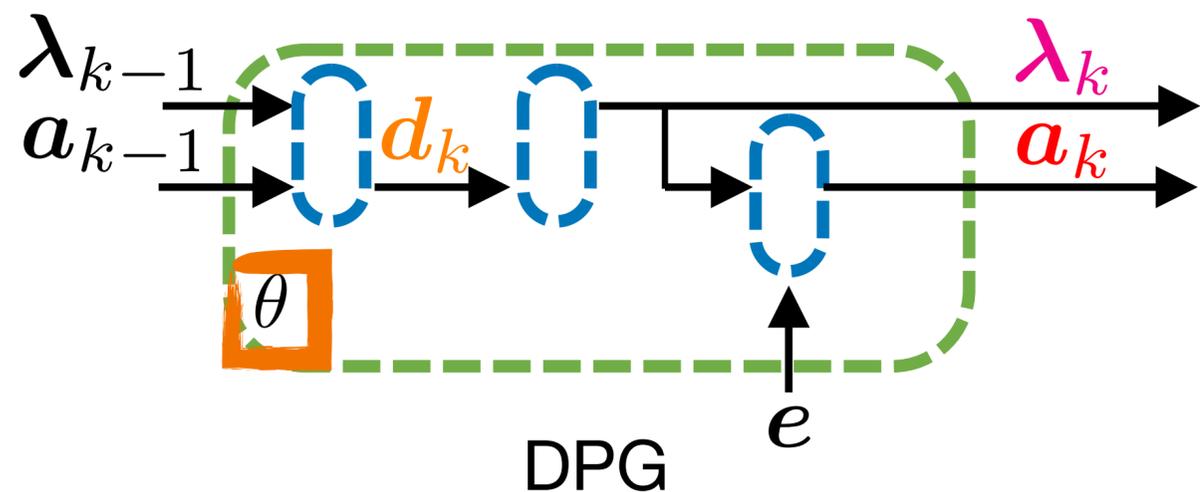
$$d_k = S a_{k-1} - (N - 1) \lambda_{k-1}$$

$$\lambda_k = \frac{-1}{2(N-1)} \left(d_k - \sqrt{d_k^2 + 4(N-1)\mathbf{1}} \right)$$

$$a_k = \max \left(\mathbf{0}, \frac{1}{2} S^\top \lambda_k - \theta e \right)$$

end for

Return: δa_k



Iterations ONLY contain θ

Sparsity pattern of solution determined by θ independently of all other parameters

Higher $\theta \rightarrow$ Higher sparsity

Defn: θ is independently interpretable w.r.t. sparsity of recovered graph

A bridge between prior information on sparsity and the value of θ

Point Estimate Approaches via Deep Unrollings

Iterative Algorithm

Algorithm: Input \mathbf{z}^0 , Output \mathbf{z}^L

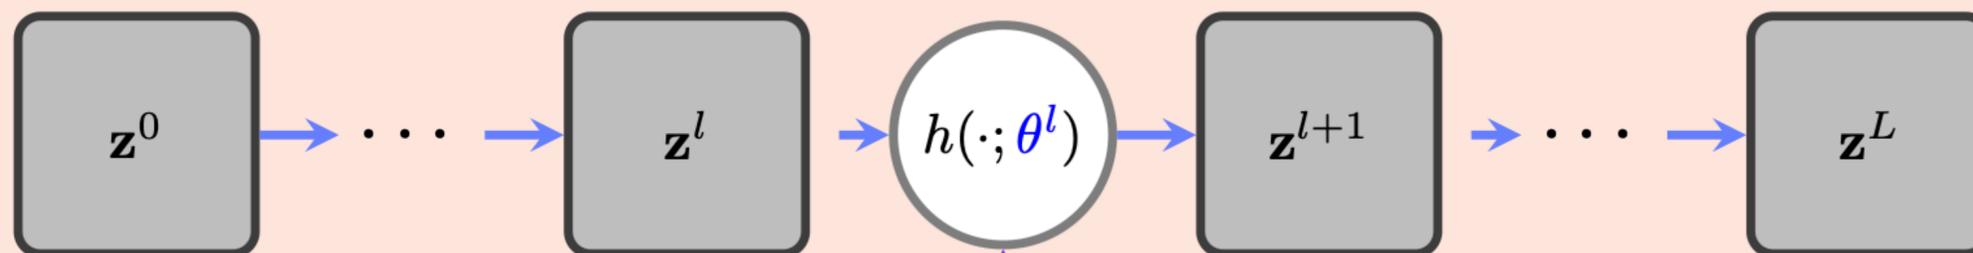
for $l = 0, 1, \dots, L - 1$ do

$\mathbf{z}^{l+1} \leftarrow h(\mathbf{z}^l; \theta^l)$,

end for

Unrolling

Unrolled Deep Network



[Monga'19]

1. Assume **generative** process $\mathbf{X} \sim \mathcal{F}(\mathbf{A})$

- ties data to the graph

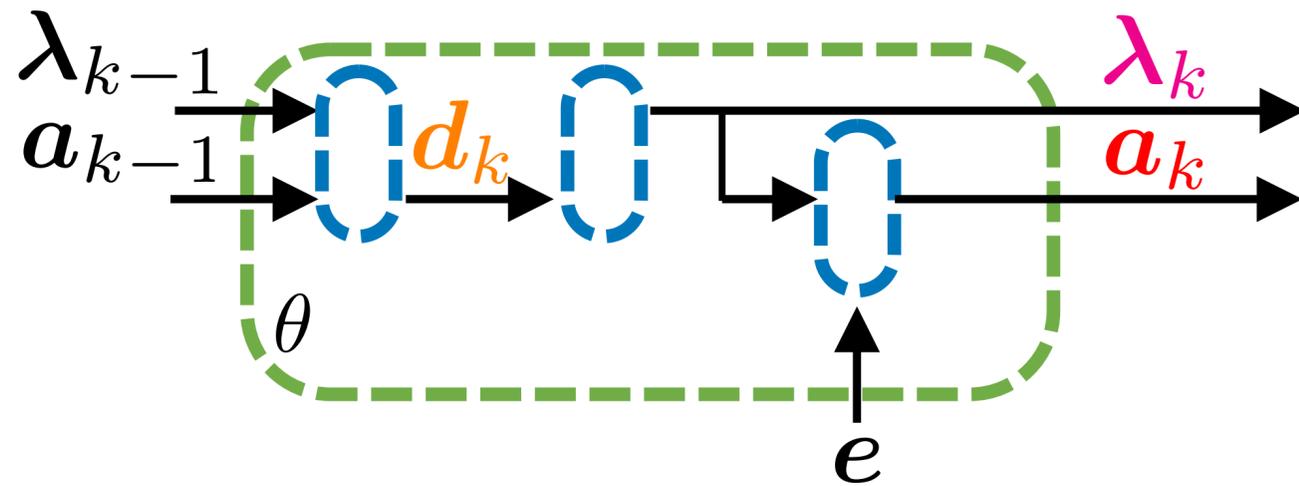
2. Propose optimization framework & **iterative solution procedure** for inverse problem $\mathbf{A} = \mathcal{F}^{-1}(\mathbf{X})$

3. **Unroll** iterative algorithm to **motivate** deep network architecture

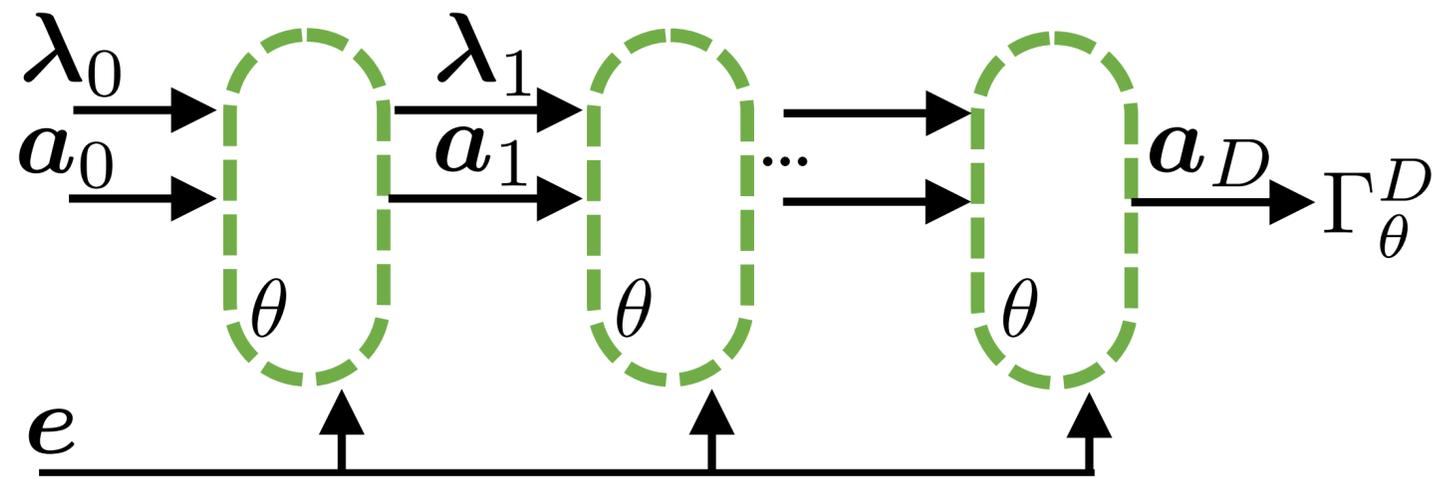
Model
Based
GSL

Point Estimate Approaches via Deep *Neural Network* Unrollings

Single Layer



Unrolled DPG



- Regularization parameters \rightarrow **Learnable** parameters with backprop.
- **Truncate** after D DPG iterations. We now **approximate solutions**.
- The first GSL *Neural Network*!
- Layers of linear transformations and point wise non-linearities.

Unrolled DPG: A Graph Valued NN With Interpretability!

- Unrolled DPG is a neural network **function** $\Gamma_{\theta}^D : e \rightarrow a$
- **Gradient** w.r.t. **parameters** well-defined!
- Dataset $\mathcal{T} = \{\mathcal{T}_e, \mathcal{T}_a\} = \{e^{(t)}, a^{(t)}\}_{t=1}^T$
- We use **unweighted** graphs $a^{(t)} \in \{0, 1\}^{|\mathcal{E}|}$. Subtract mean b . Drive through sigmoid σ .

Final 3 Parameter
GSL Neural Network

$$\sigma(\delta \Gamma_{\theta}^D(e) - b); \Theta = \{\theta, \delta, b\}$$

- Bernoulli **likelihood**: Unrolling encodes the mean.

Bayesian Neural Networks (BNN)

Background

- A Bayesian NN: a NN with **stochastic** weights.
- Posterior Distribution: **Distribution** over weights conditioned on observed data.
- Pushing posterior distribution through the NN produces a **distribution** over **predictions**.
- We can use this distribution to derive a measure of **uncertainty**.
- Key Ingredients

- Weight Prior

$$\underline{p(\Theta)}$$

- Likelihood

$$\underline{p(\mathcal{T}_a | \mathcal{T}_e, \Theta)}$$

- Posterior

$$\underline{p(\Theta | \mathcal{T})}$$

$$\underline{p(\Theta | \mathcal{T})} \propto \underline{p(\mathcal{T}_a | \mathcal{T}_e, \Theta)} \underline{p(\Theta)}$$

$$p(\tilde{\mathbf{a}} | \tilde{\mathbf{e}}, \mathcal{T}) = \int \underline{p(\tilde{\mathbf{a}} | \tilde{\mathbf{e}}, \Theta)} \underline{p(\Theta | \mathcal{T})} d\Theta$$

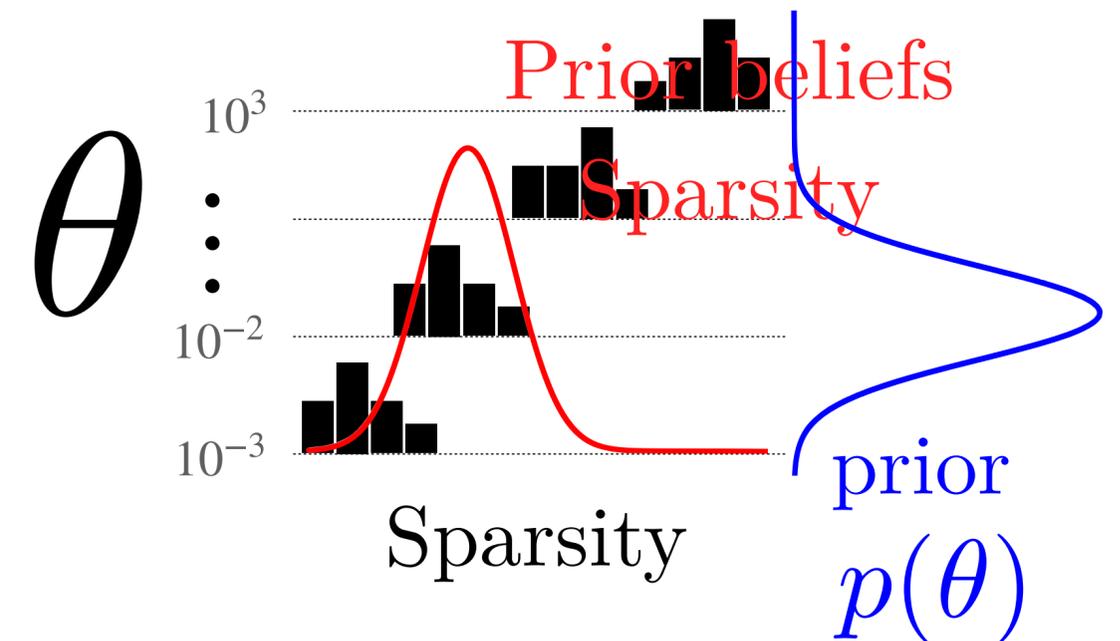
Difficulties

1. How do we set the prior?
2. How do we approximate the posterior?

Producing a Bayesian Neural Network

Informative Prior Over Parameters

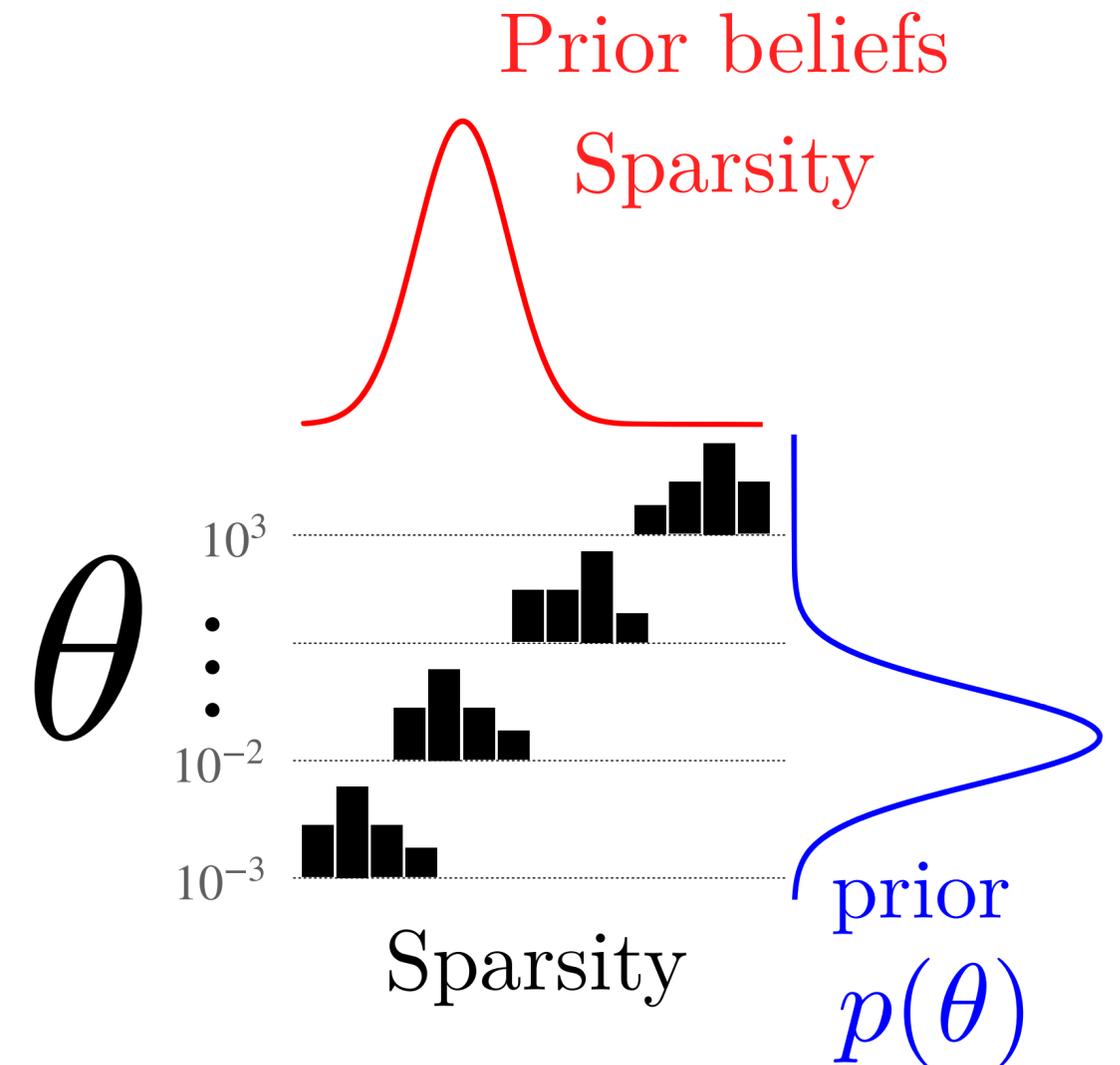
- BNNs require parameter priors $p(\Theta)$.
- To construct $p(\Theta)$:
 - Use **independent interpretability** of θ .
 - Subset of inputs \mathcal{T}_e .
 - **prior beliefs over sparsity** \rightarrow **prior distribution over θ** .
- **Weakly informative prior** for $p(\delta, b)$.
 - Inspect edge weight magnitudes at performant θ 's.



Producing a Bayesian Neural Network

Informative Prior Over Parameters

- BNNs require parameter priors $p(\Theta)$.
- To construct $p(\Theta)$:
 - Use **independent interpretability** of θ .
 - Subset of inputs \mathcal{T}_e .
 - **prior beliefs over sparsity** \rightarrow **prior distribution over θ** .
- **Weakly informative prior** for $p(\delta, b)$.
 - Inspect edge weight magnitudes at performant θ 's.



GSL with Interpretable BNNs

Bayesian Modeling: Posterior Predictive

1. **Inference**: condition on the data
Via Hamiltonian Monte Carlo

$$\underline{\Theta}^{(m)} \sim p(\Theta \mid \mathcal{T})$$

“Posterior Samples”

2. **Marginalize out** parameters $p(\tilde{\mathbf{a}} \mid \tilde{\mathbf{e}}, \mathcal{T}) \approx \frac{1}{M} \sum_{m=1}^M p(\tilde{\mathbf{a}} \mid \tilde{\mathbf{e}}, \underline{\Theta}^{(m)})$

“Posterior Predictive”
over test sample
($\tilde{\mathbf{e}}, \tilde{\mathbf{a}}$)

3. Which we **sample** from ... $\tilde{\mathbf{a}}^{(m)} \sim p(\tilde{\mathbf{a}} \mid \tilde{\mathbf{e}}, \underline{\Theta}^{(m)})$

4. To produce edge-wise **point** and **uncertainty** estimates!

Point (‘pred. mean’)

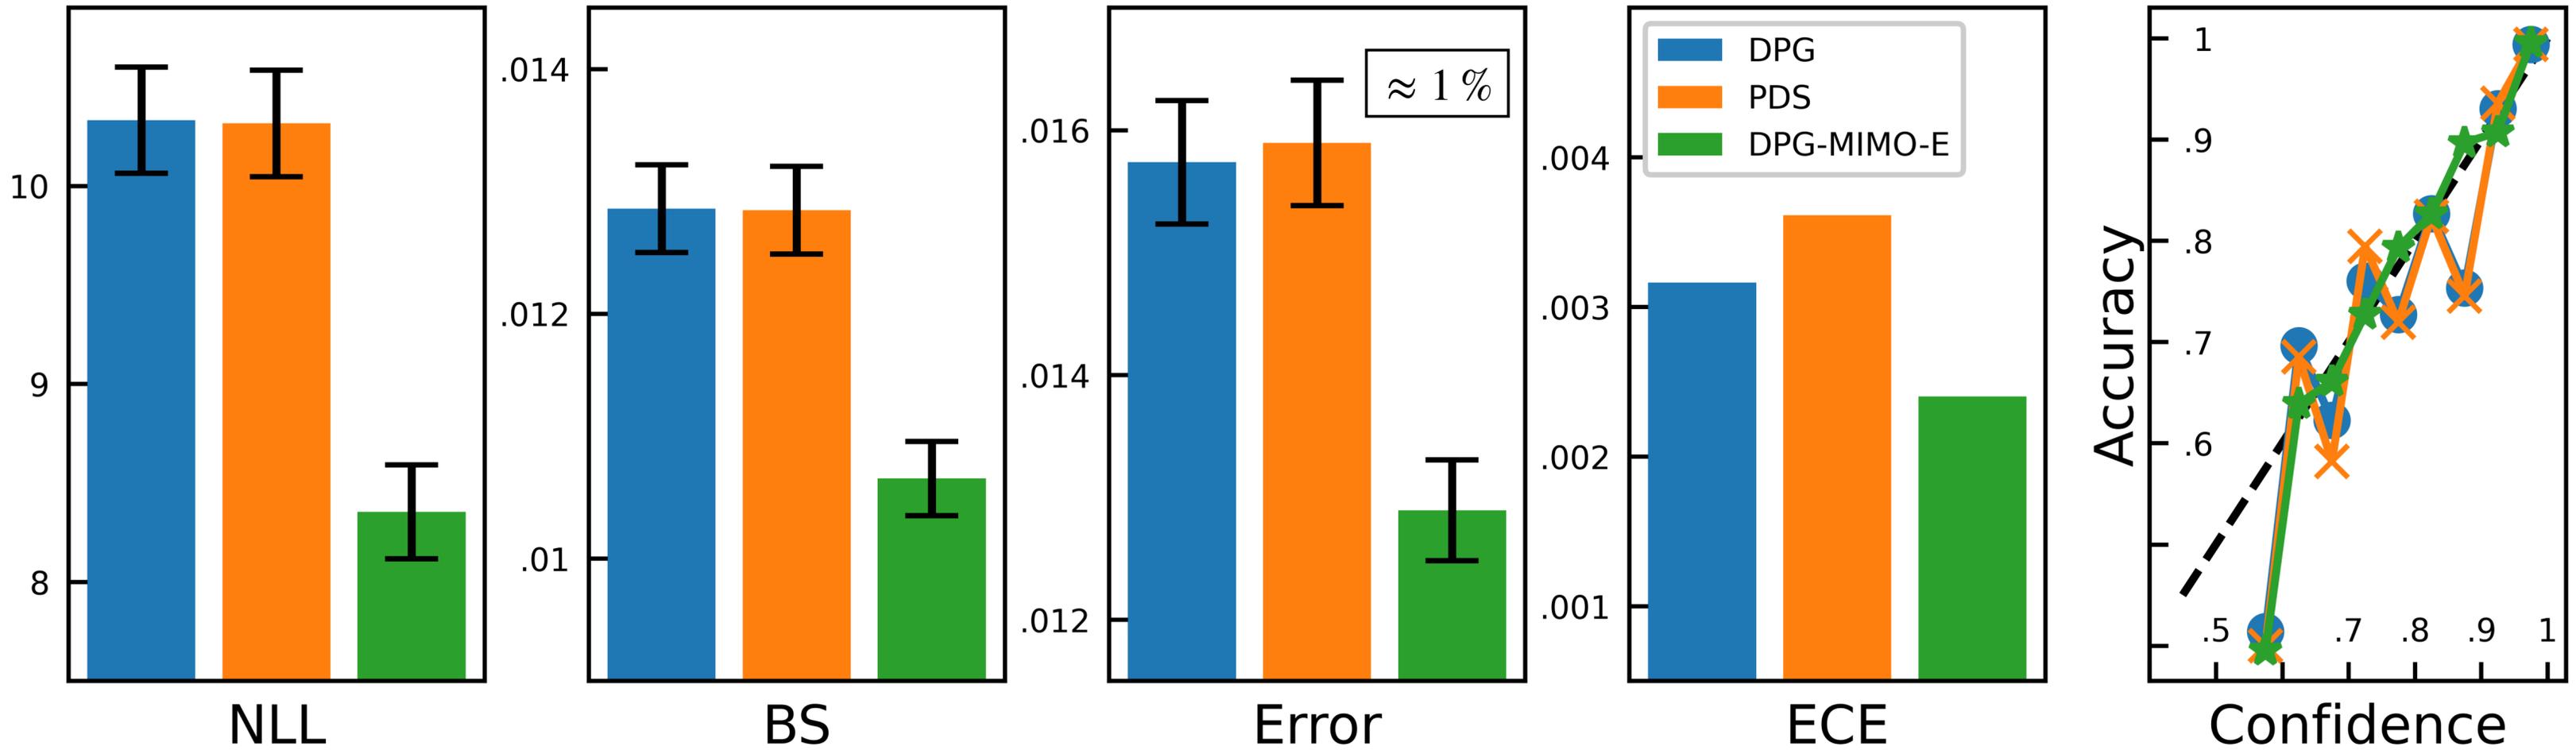
$$\mathbb{E}[\tilde{a}_i \mid \tilde{\mathbf{e}}, \mathcal{T}] \approx \frac{1}{M} \sum_{m=1}^M \tilde{a}_i^{(m)}$$

Uncertainty (‘pred. stdv.’)

$$\text{Var}[\tilde{a}_i \mid \tilde{\mathbf{e}}, \mathcal{T}]^{\frac{1}{2}} \approx \left[\frac{1}{M} \sum_{m=1}^M (\tilde{a}_i^{(m)} - \mathbb{E}[\tilde{a}_i \mid \tilde{\mathbf{e}}, \mathcal{T}])^2 \right]^{\frac{1}{2}}$$

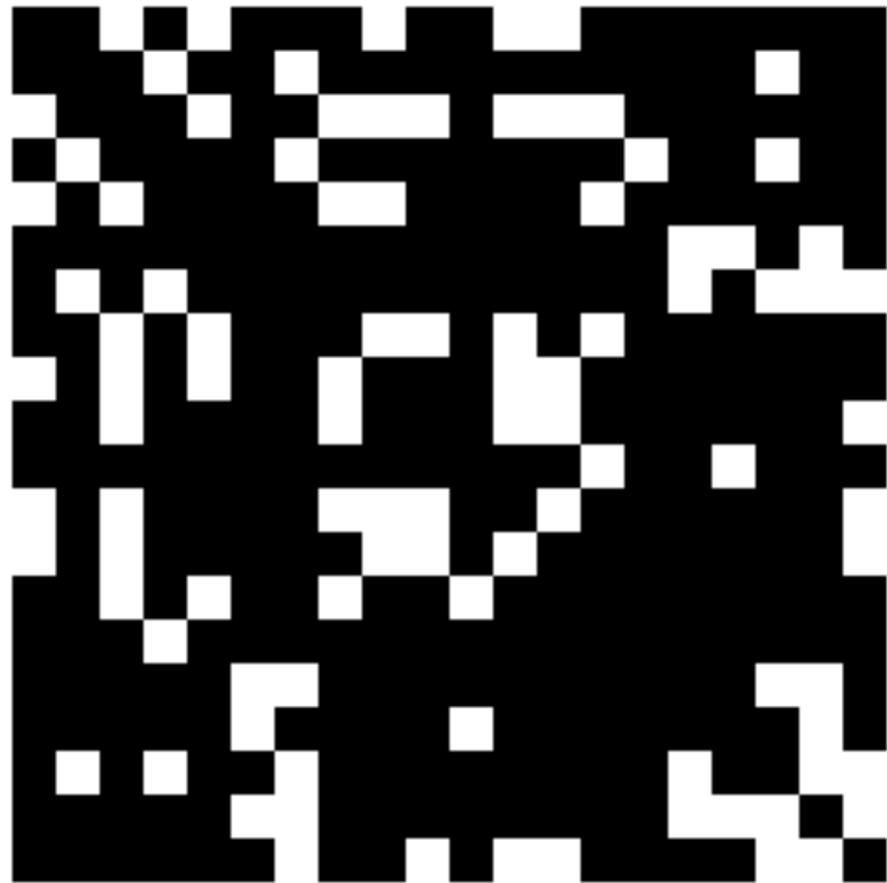
GSL with Interpretable BNNs

Synthetic Evaluation



GSL with Interpretable BNNs

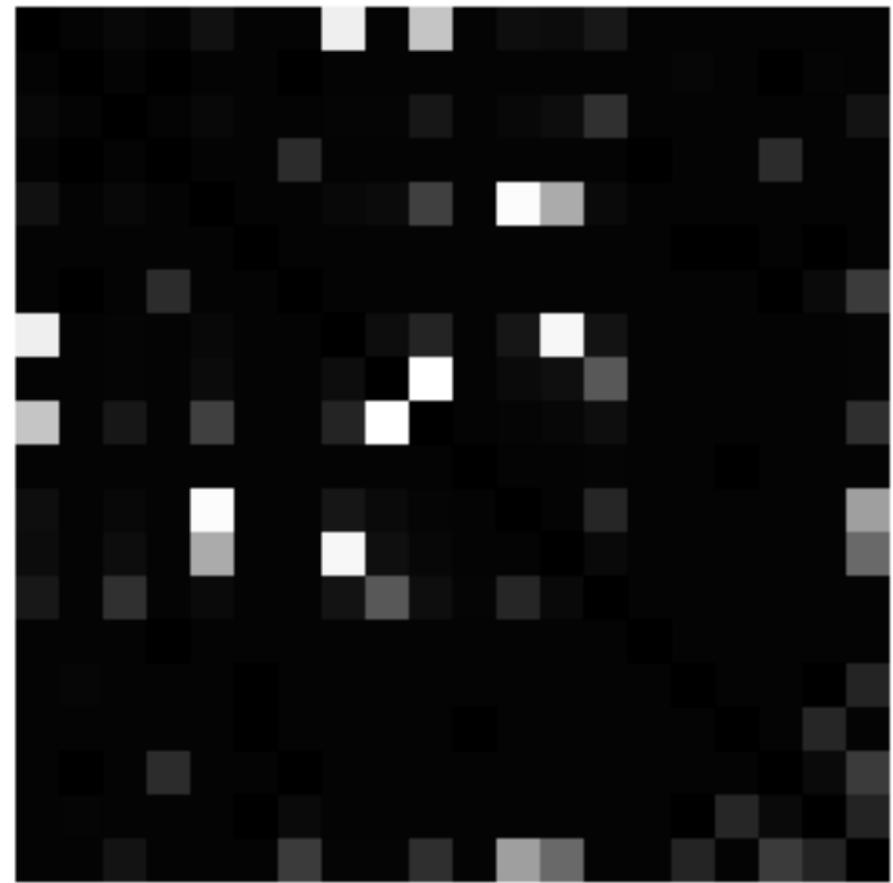
Synthetic Evaluation



label



pred. mean



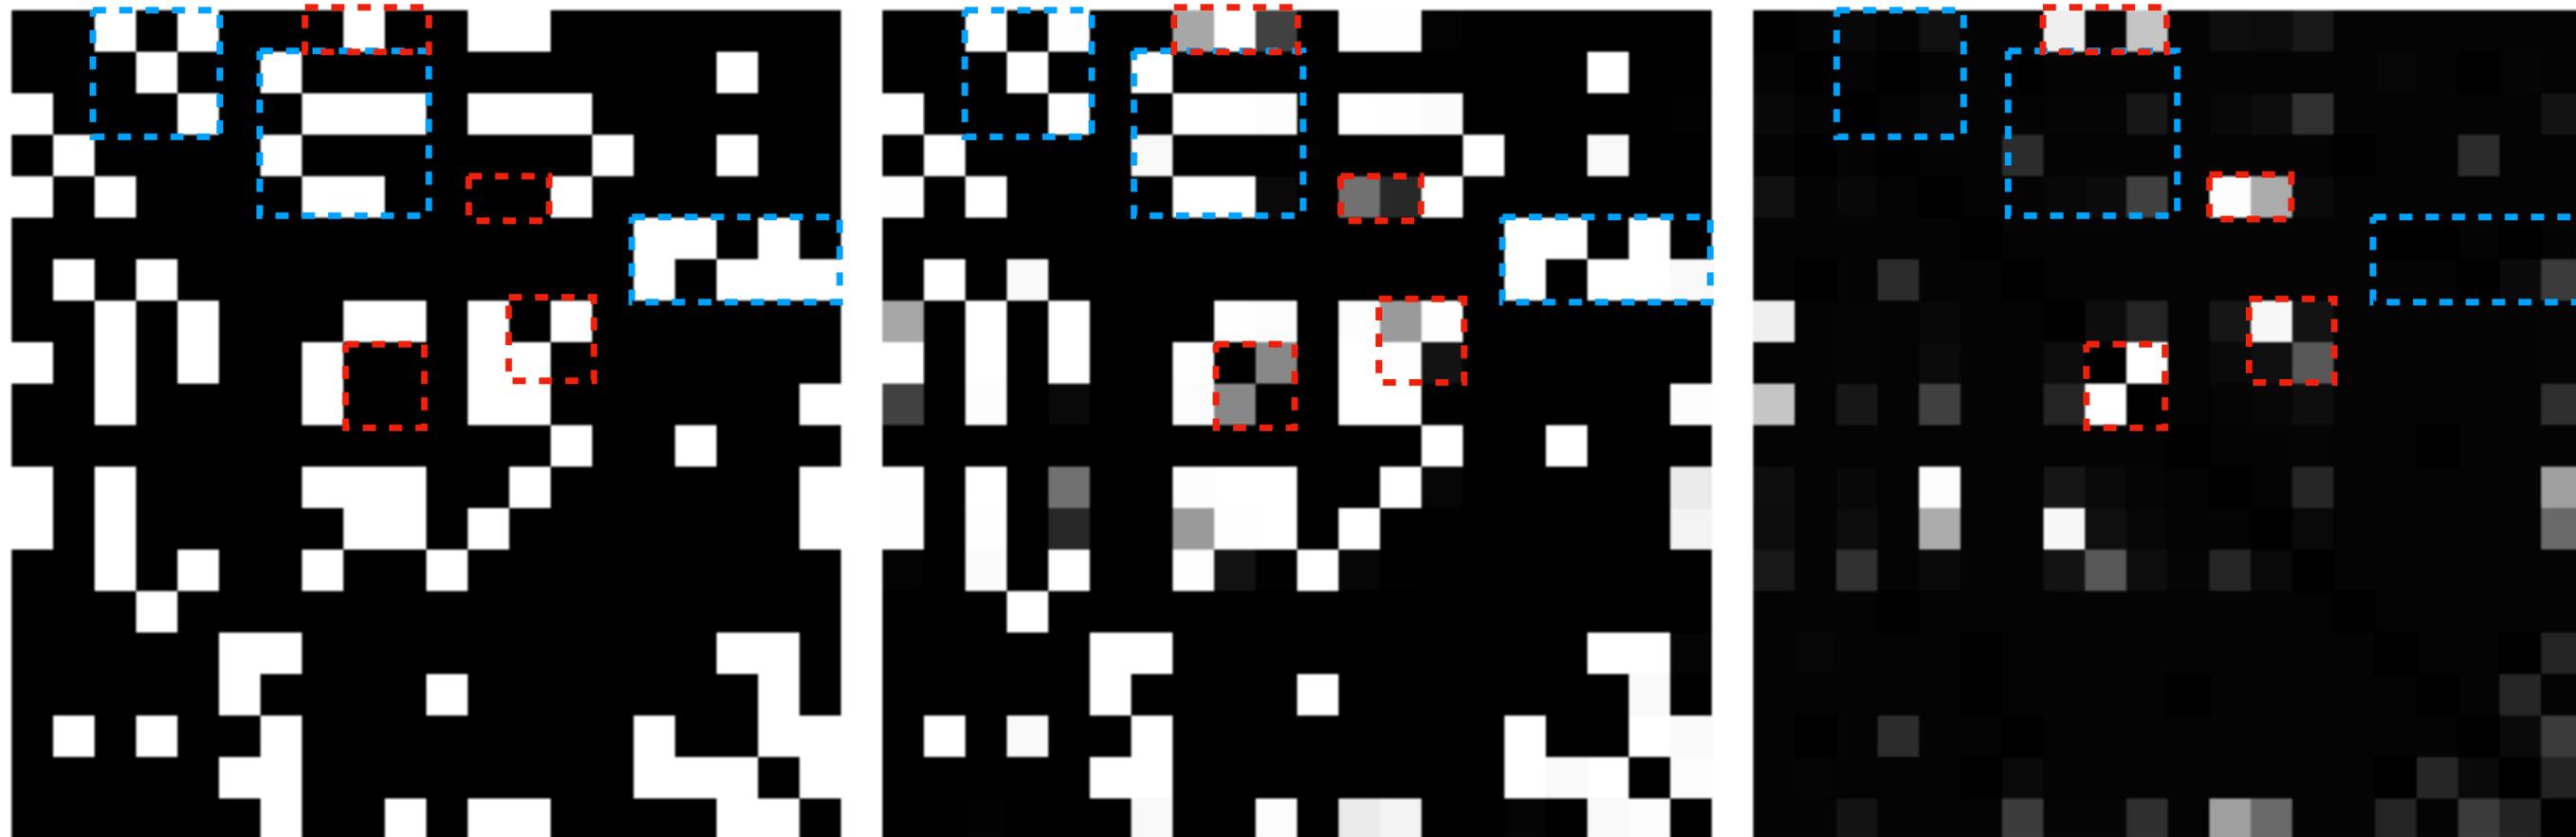
pred. stdv

GSL with Interpretable BNNs

Synthetic Evaluation

Strong Recovery

Some Mistakes



label

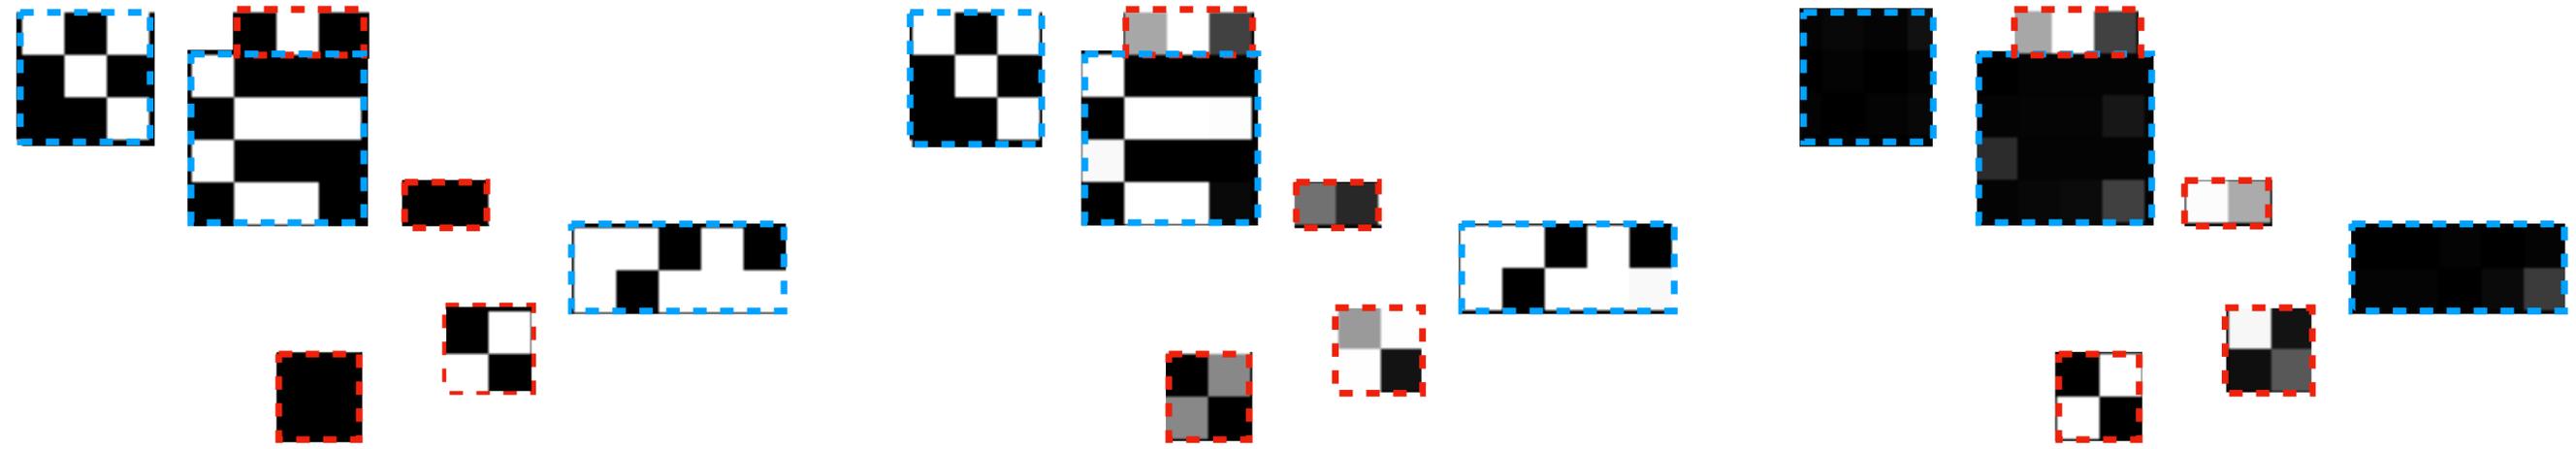
pred. mean

pred. stdv

GSL with Interpretable BNNs

Synthetic Evaluation

Strong Recovery
Some Mistakes



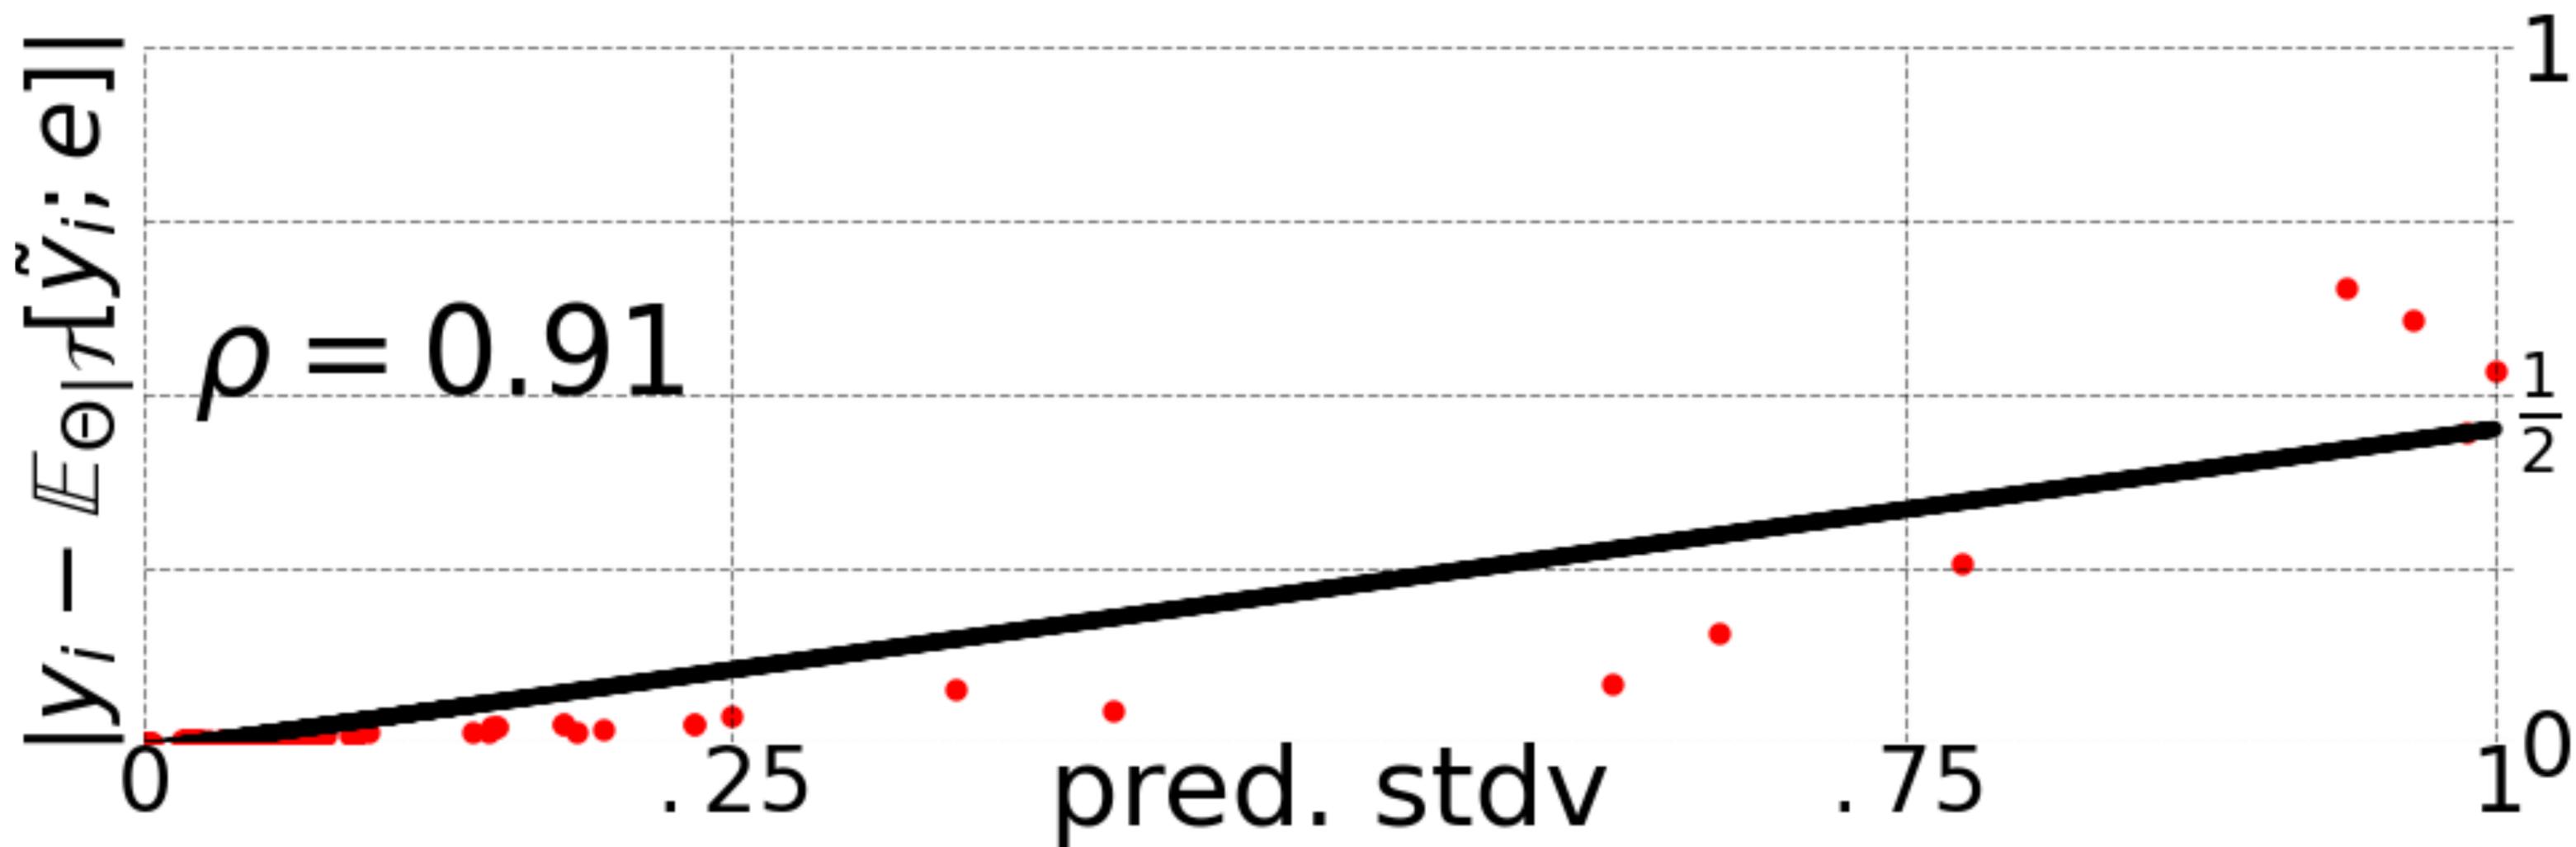
label

pred. mean

pred. stdv

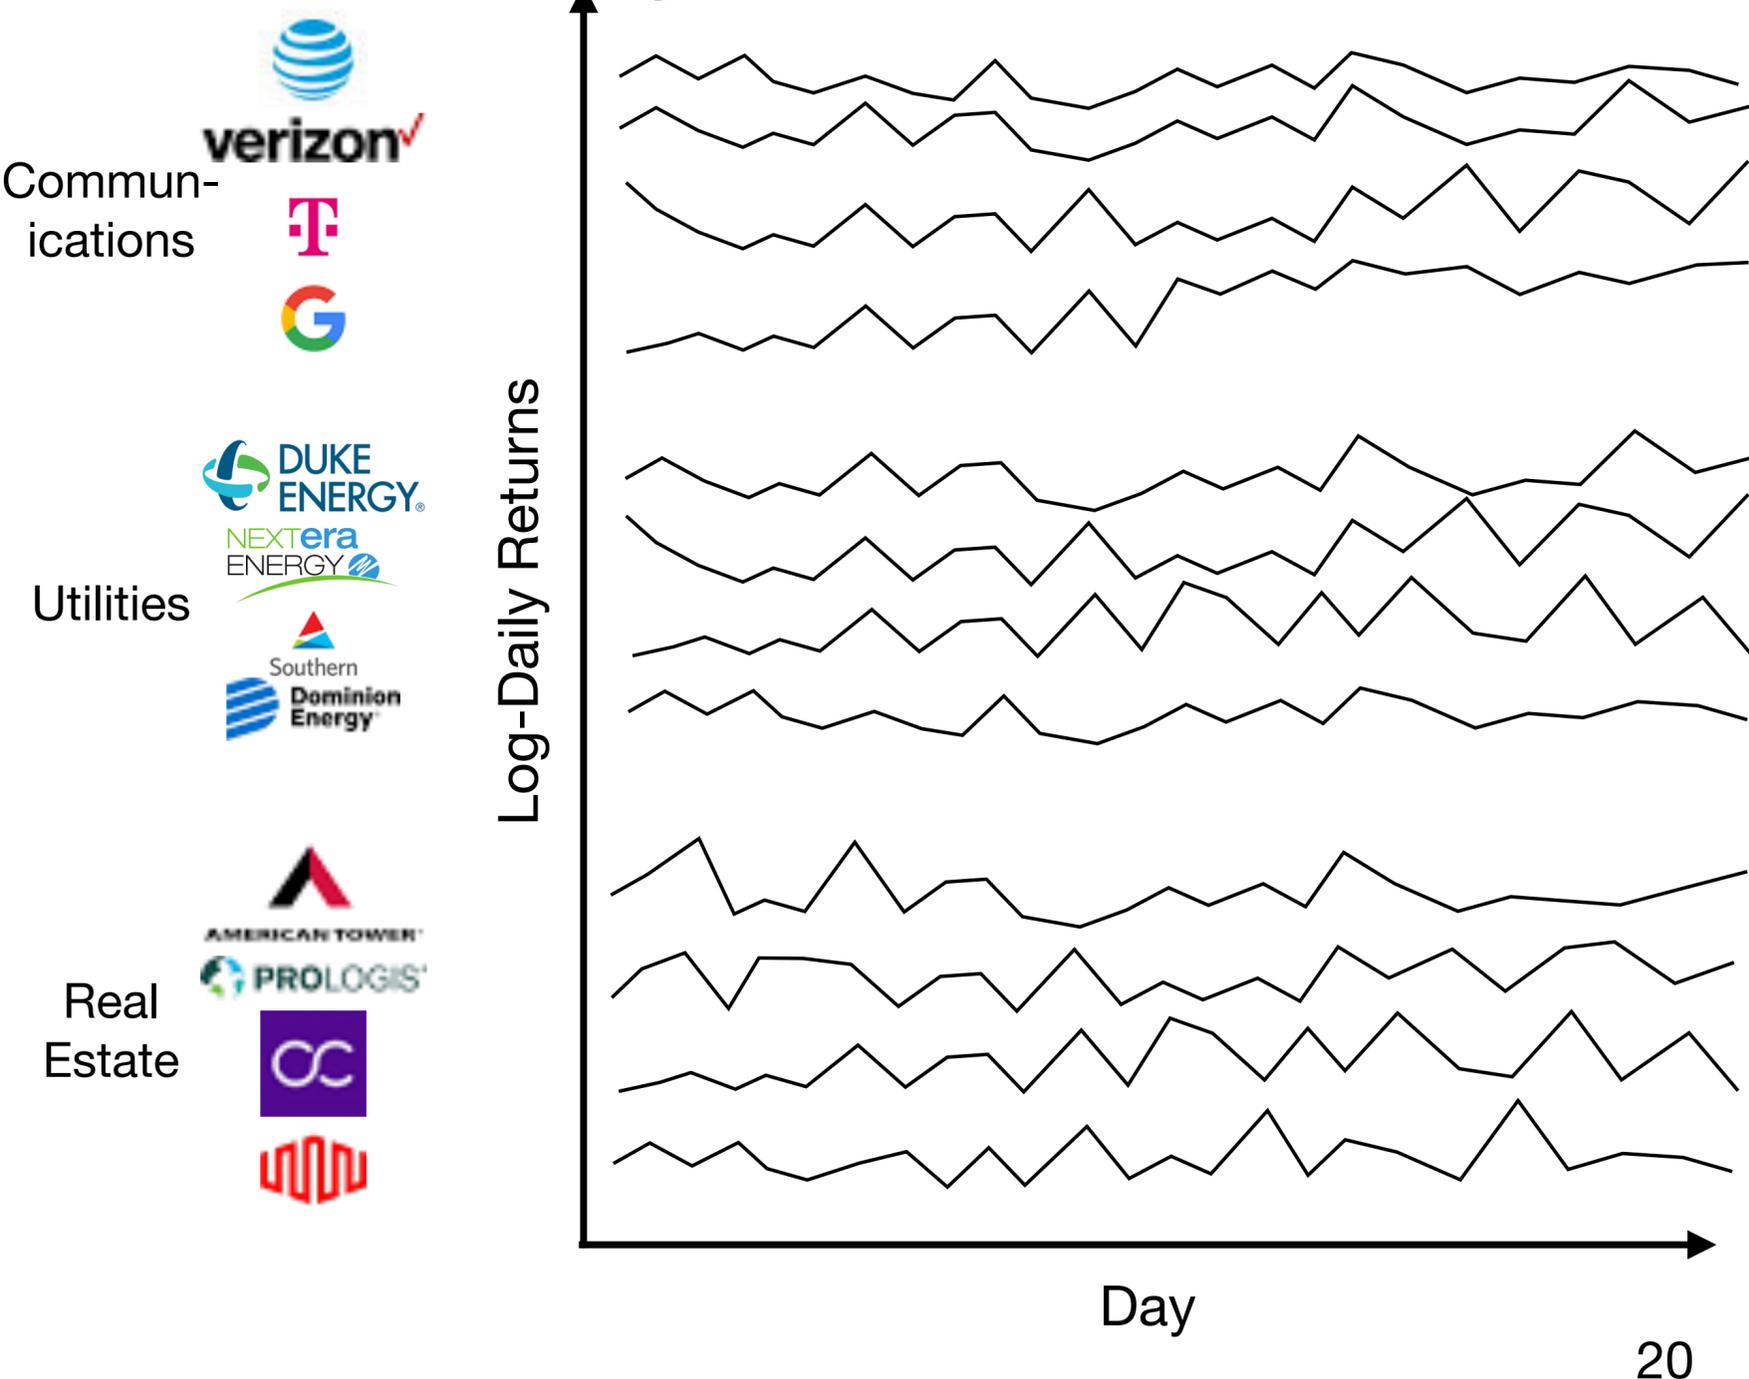
GSL with Interpretable BNNs

Synthetic Evaluation



GSL with Interpretable BNNs

Recovering Sector Graphs from SP500 stock time series [Cardoso'23]

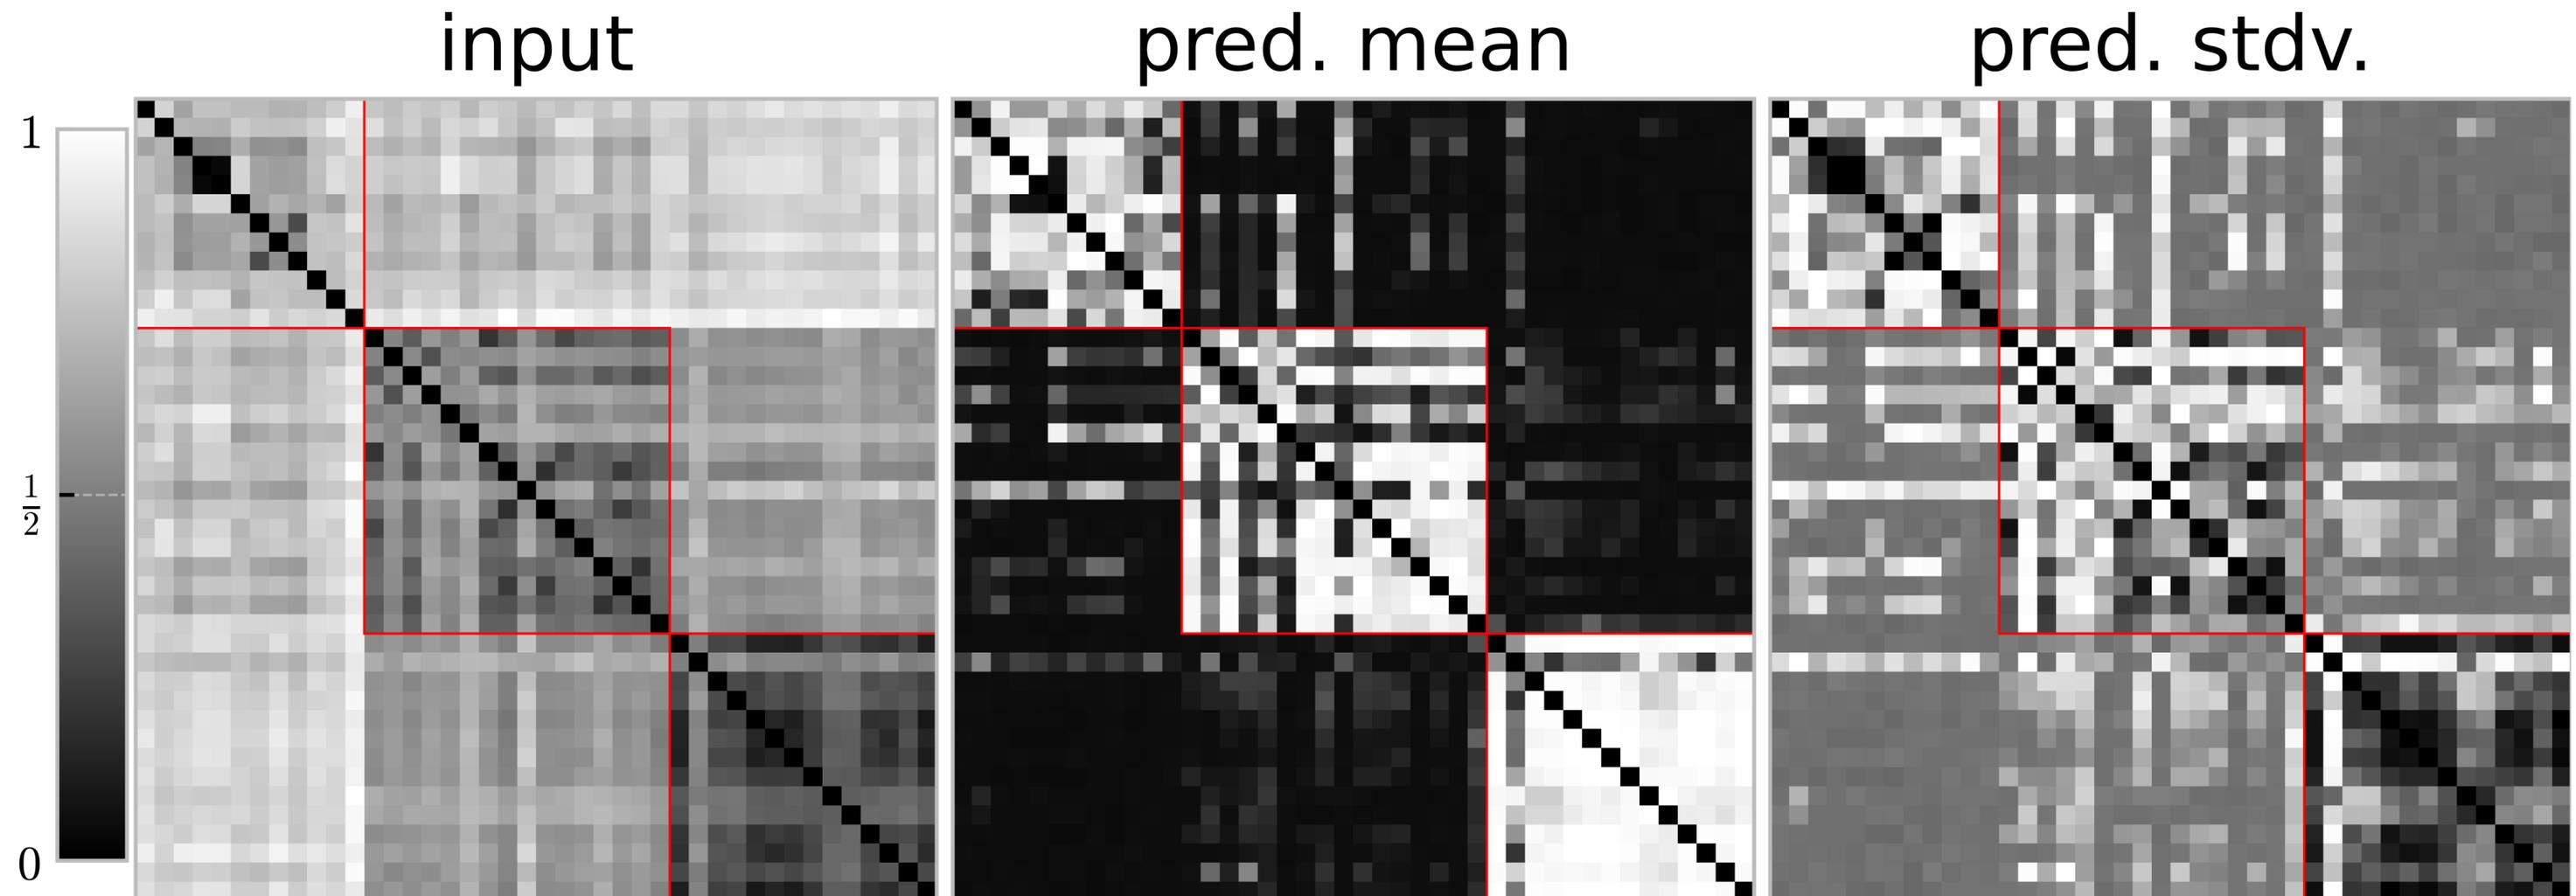


- Split stocks evenly into train/test
- For each split
 - **Input:** $1 - |\text{Pearson correlation matrix}|$
 - **Label:** $A_{ij} = 1$ if stocks i & j in same sector, 0 otherwise.
 - = Block Diagonal with 3 blocks
- Run inference on DPG using this single training sample

GSL with Interpretable BNNs

Evaluation: Recovering Sector Graphs from SP500 stock time series

- Error and Uncertainty Pearson Correlation: **0.70**



Closing Remarks

- We analyze recovering graphs from nodal observations
- DPG: simple iterations which produce a **true neural network** when unrolled
- **Structure** of DPG makes a parameter **independently interpretable**
- Use structure to construct **informative parameter priors** using **priors on sparsity**
- Conditioning on observed data \rightarrow posterior distribution $p(\Theta \mid \mathcal{T})$
- On unseen nodal observations X , **push posterior distribution** through Unrolled DPG to produce a **distribution over edges** $p(\tilde{a} \mid \mathcal{T}; \tilde{e})$
 - From which we recover **point** and **uncertainty** estimates
- Future Work: Scaling & Variational Inference

For more...

- ArXiv
 - <https://arxiv.org/abs/2406.14786>
- Github repo
 - github.com/maxwass/gsl-bnn

Graph Structure Learning with Interpretable Bayesian Neural Networks

Max Wasserman
Department of Computer Science
University of Rochester

mwasser6@cs.rochester.edu

Gonzalo Mateos
Department of Electrical and Computer Engineering
University of Rochester

gmateosb@ece.rochester.edu

Abstract

Graphs serve as generic tools to encode the underlying relational structure of data. Often this graph is not given, and so the task of inferring it from nodal observations becomes