Bayesian Recovery of Dependence Structures on Networks

Robbert van der Burg

joint work with Alessandro Zocca and Frank van der Meulen

NetSci Summer Symposium 2026

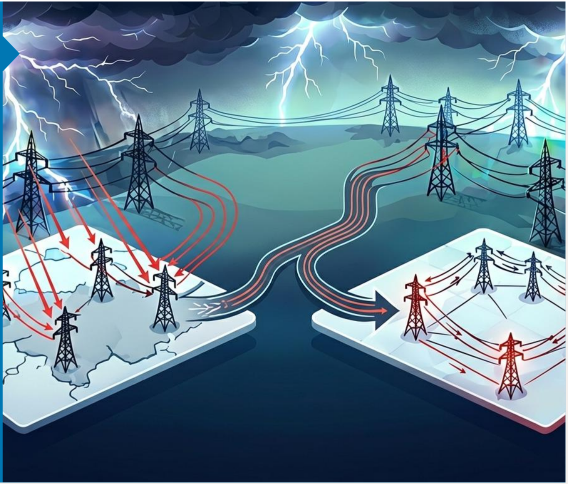

Can we capture complex network failure modes with a parsimonious model?

Data. k binary snapshots of the state of an n-node network:

$$\boldsymbol{X} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(k)}\} \in \{0, 1\}^{k \times n}$$

Physical network $\mathcal{G}$

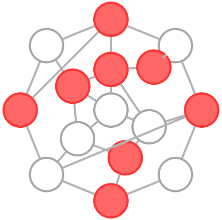

Can we capture complex network failure modes with a parsimonious model?

Data. k binary snapshots of the state of an n-node network:

$$\boldsymbol{X} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(k)}\} \in \{0,1\}^{k \times n}$$

Ising model on $\mathcal{G} = (\mathcal{V}, \mathcal{E})$:

$$p(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) \propto \exp\Big(\sum_i x_i \mu_i + \sum_{i<j} x_i x_j \sigma_{ij}\Big)$$

Physical network $\mathcal{G}$

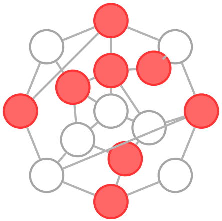

Can we capture complex network failure modes with a parsimonious model?

Data. k binary snapshots of the state of an n-node network:

$$\boldsymbol{X} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(k)}\} \in \{0,1\}^{k \times n}$$

Ising model on $\mathcal{G} = (\mathcal{V}, \mathcal{E})$:

$$p(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) \propto \exp\Big(\sum_i x_i \mu_i + \sum_{i<j} x_i x_j \sigma_{ij}\Big)$$

Goal. Estimate $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$, recover

which edges correctly capture failure correlations?

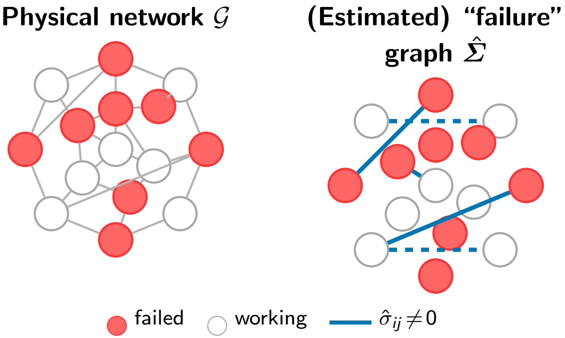

The missing piece: (global) covariates

Introduce covariates. Replace the scalar threshold μ_i by a covariate-informed linear predictor:

$$h_i(\boldsymbol{z}_i) = \alpha_{t(i)} + \boldsymbol{\beta}_{t(i)}^\top \boldsymbol{z}_i$$

The missing piece: (global) covariates

Introduce covariates. Replace the scalar threshold μ_i by a covariate-informed linear predictor:

$$h_i(\boldsymbol{z}_i) = \alpha_{t(i)} + \boldsymbol{\beta}_{t(i)}^\top \boldsymbol{z}_i$$

so the joint distribution becomes

$$p(\boldsymbol{x} \mid \Theta, \boldsymbol{Z}) \propto \exp\Big(\sum_i h_i(\boldsymbol{z}_i)\, x_i + \sum_{i<j} \sigma_{ij}\, x_i x_j\Big),$$

where $\Theta = \{\alpha, \boldsymbol{\beta}, \boldsymbol{\Sigma}\}$.

The missing piece: (global) covariates

Introduce covariates. Replace the scalar threshold μ_i by a covariate-informed linear predictor:

$$h_i(\boldsymbol{z}_i) = \alpha_{t(i)} + \boldsymbol{\beta}_{t(i)}^\top \boldsymbol{z}_i$$

so the joint distribution becomes

$$p(\boldsymbol{x} \mid \Theta, \boldsymbol{Z}) \propto \exp\Big(\sum_i h_i(\boldsymbol{z}_i)\, x_i + \sum_{i<j} \sigma_{ij}\, x_i x_j\Big),$$

where $\Theta = \{\alpha, \boldsymbol{\beta}, \boldsymbol{\Sigma}\}$.

Concrete example: power grid

$x_i \in \{0, 1\}$	component failed?
$\boldsymbol{z}_i \in \mathbb{R}^p$	age, load, temperature
$\alpha_{t(i)}$	base failure rate by type
$\boldsymbol{\beta}_{t(i)}$	covariate effect by type
σ_{ij}	residual coupling

From model to discovery

Approach: covariate-dependent Ising model with spike-and-slab prior $\rightarrow$ EM edge-screening $\rightarrow$ Gibbs sampler for posterior uncertainty.

From model to discovery

Approach: covariate-dependent Ising model with spike-and-slab prior $\rightarrow$ EM edge-screening $\rightarrow$ Gibbs sampler for posterior uncertainty.

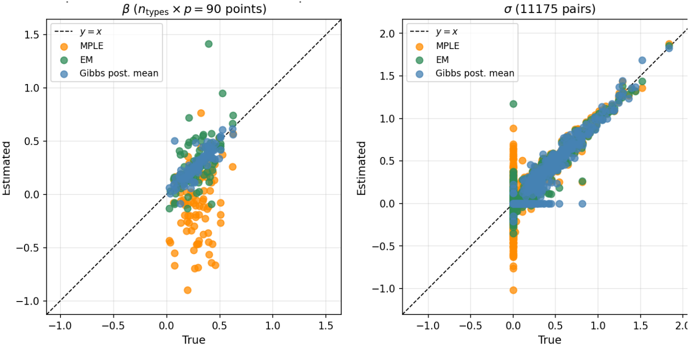

Recovered edge weights $\hat{\sigma}_{ij}$ vs. truth

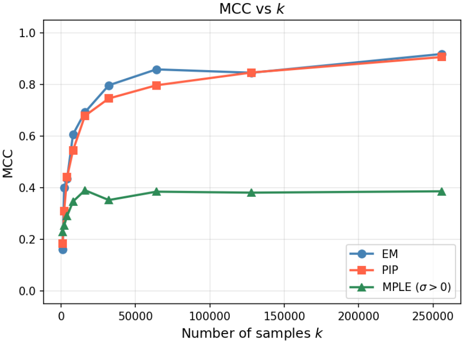

Structure recovery (MCC) vs. sample size