Solving scientific model comparison with Evidence Networks

Niall Jeffrey

ArXiv:2305.11241

(*co-author Ben Wandelt*)

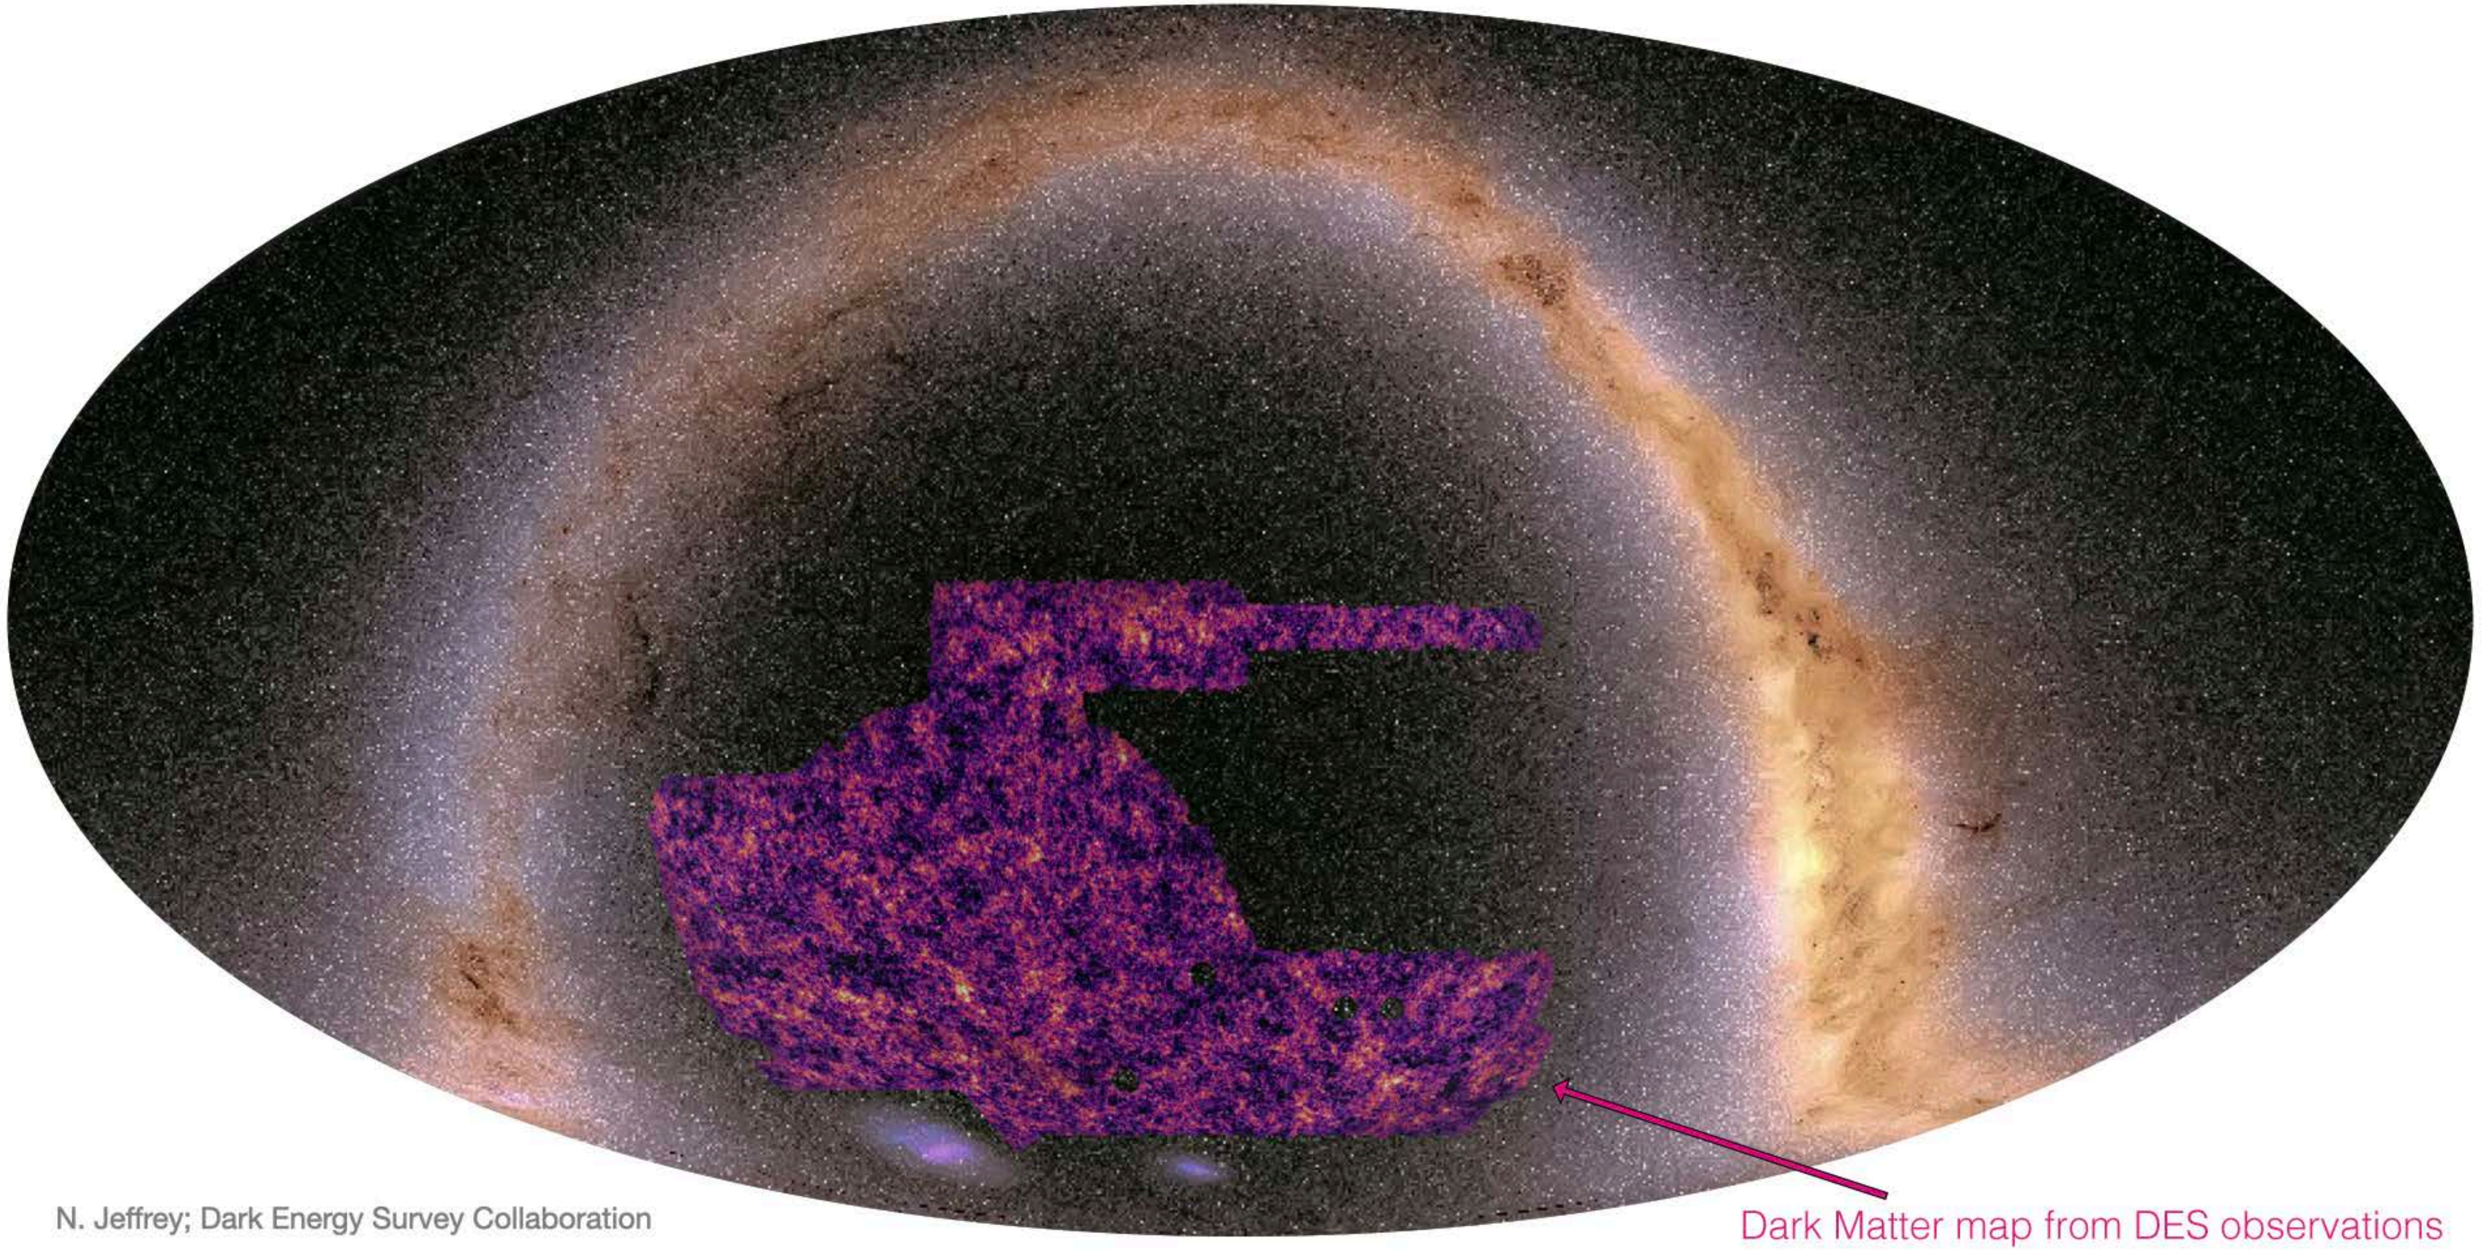
N. Jeffrey; Dark Energy Survey Collaboration
Dark Matter map from DES observations

Likelihood-free result…

NJ, Alsing, Lanusse 2009.08459

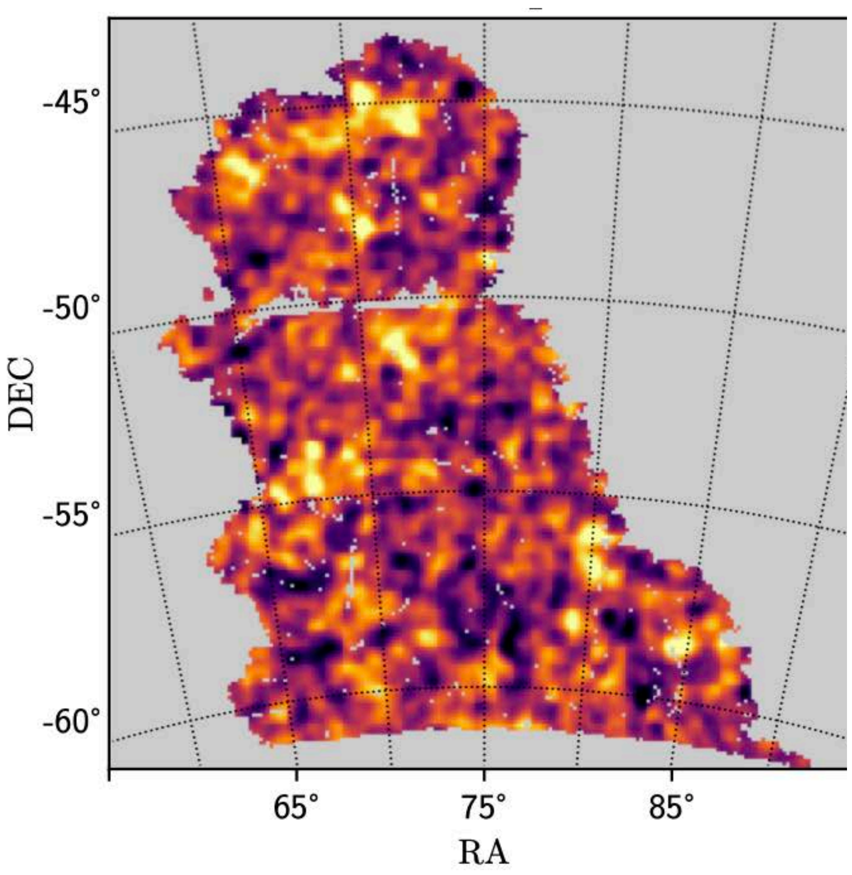

Dark Energy Survey SV data

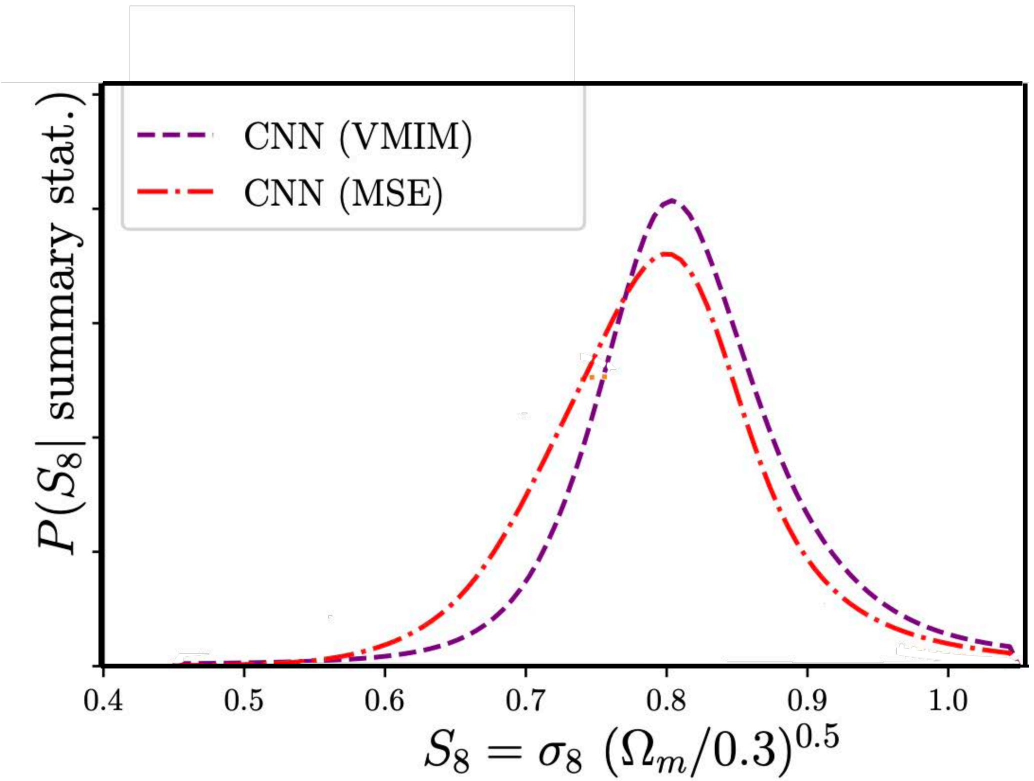

Posterior probability for unknown parameters

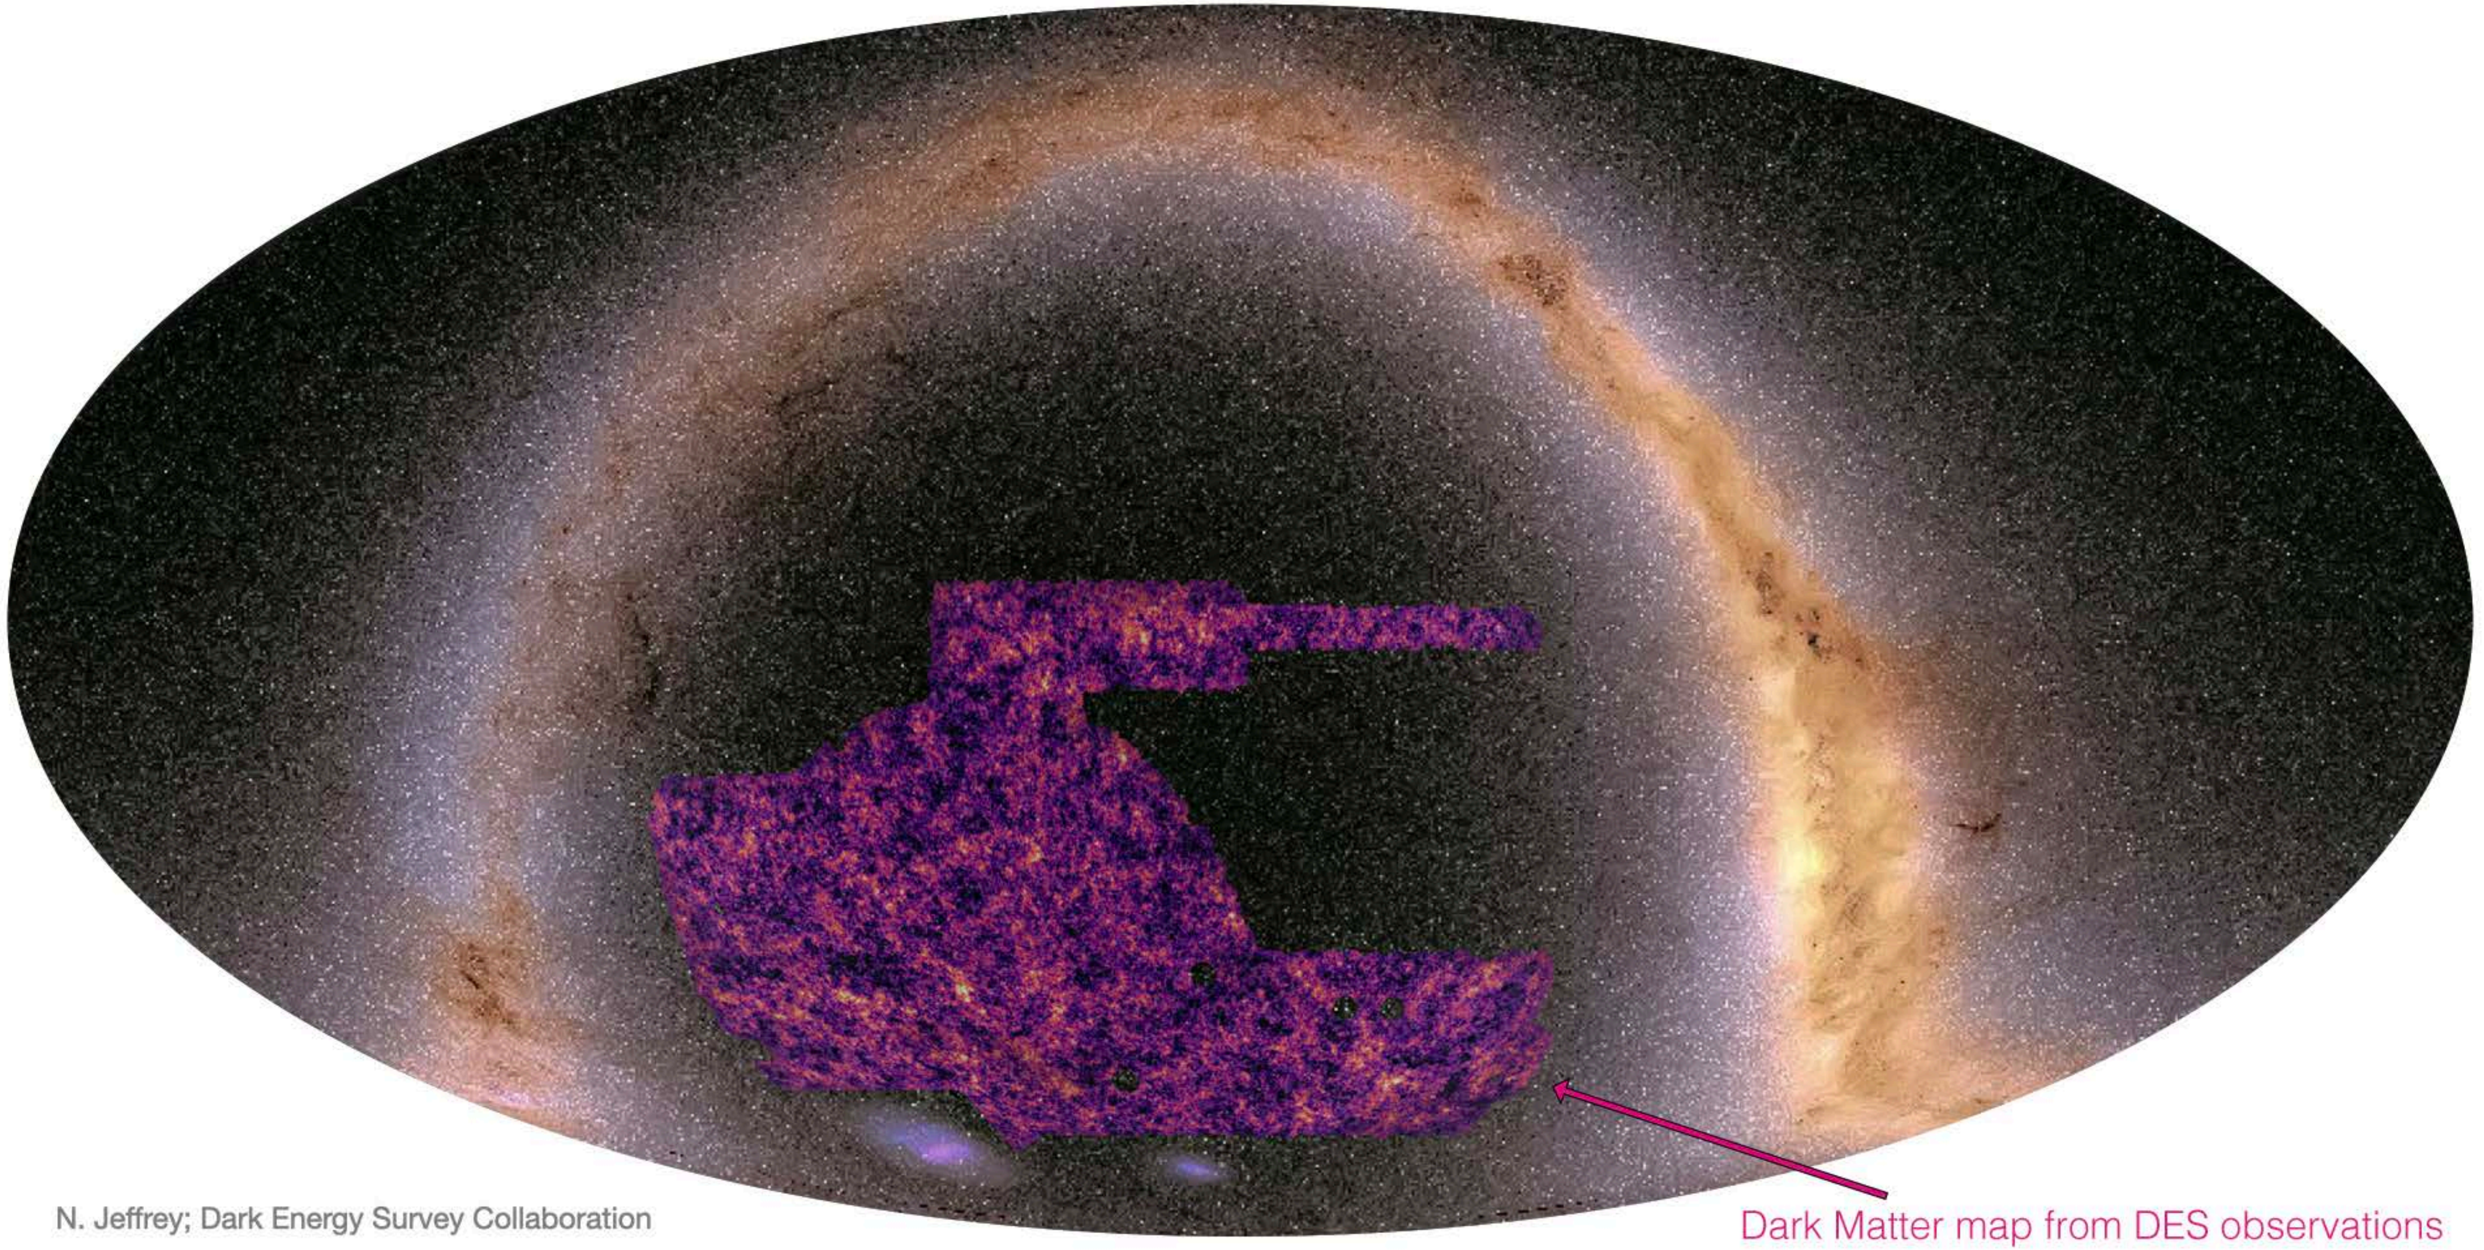
N. Jeffrey; Dark Energy Survey Collaboration
Dark Matter map from DES observations

Outline

- Challenge of model comparison
- Evidence Networks
- Demonstration highlights:
    - Time-series data
    - DES data
- Extensions and applications

The challenge of model comparison

The challenge of model comparison

$$p(M_1 | x) \text{ vs } p(M_0 | x)$$

‘Bayes factor’: $$K = \frac{p(x_O|M_1)}{p(x_O|M_0)}$$

What do people usually do?

Marginal likelihood:

$$p(x_O|M_1) = \int p(x_O|\theta, M_1)\ p(\theta|M_1)\ \mathrm{d}\theta$$

Evidence Networks

Evidence Networks

1. Generate/collect data for each model: $x_i \sim p(x \,|\, M_1)$

2. Bespoke loss function: $\mathscr{V}(f(x), m)$

3. Simple network to estimate Bayes factor: $f^*(x_O) = \log K$

How does this work?

$$\mathcal{V}(f(x), m) = e^{\left(\frac{1}{2} - m\right) f(x)}$$

$m \in \{0,1\}$

model label

How does this work?

$$\mathcal{V}(f(x), m) = e^{\left(\frac{1}{2} - m\right) f(x)}$$

$$\Downarrow$$

$$f^*(x_O) = \log K$$

How does this work?

$$\mathcal{V}(f(x), m) = e^{\left(\frac{1}{2} - m\right)\mathcal{J}_\alpha(f(x))}$$

leaky parity odd power
(l-POP) transform

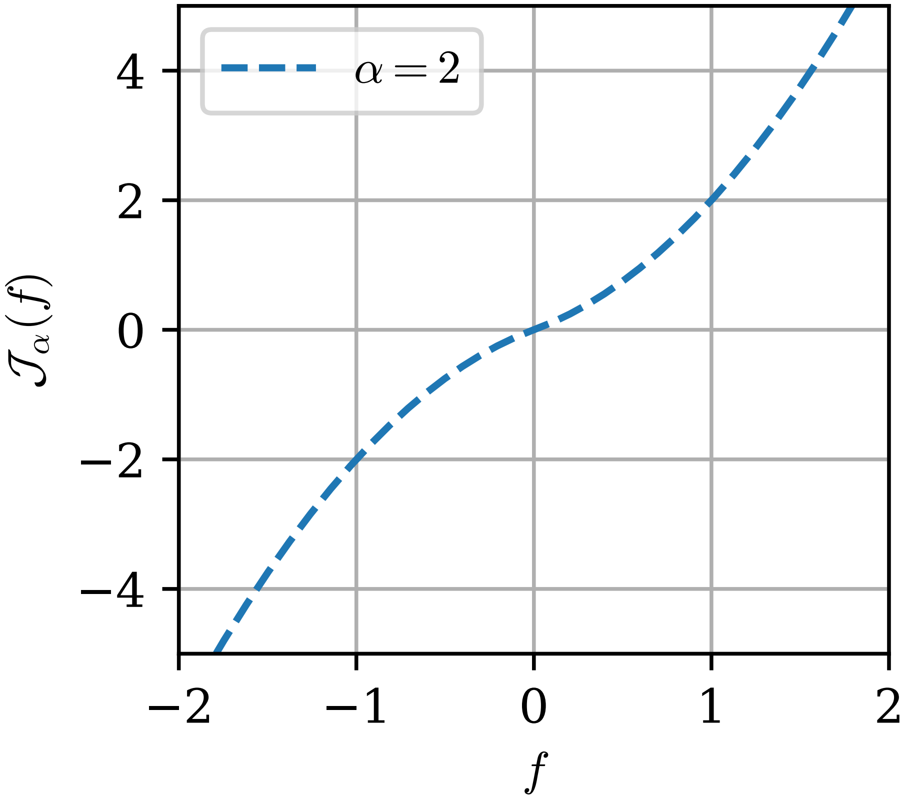

Evidence Network demonstration highlights

Time-series results

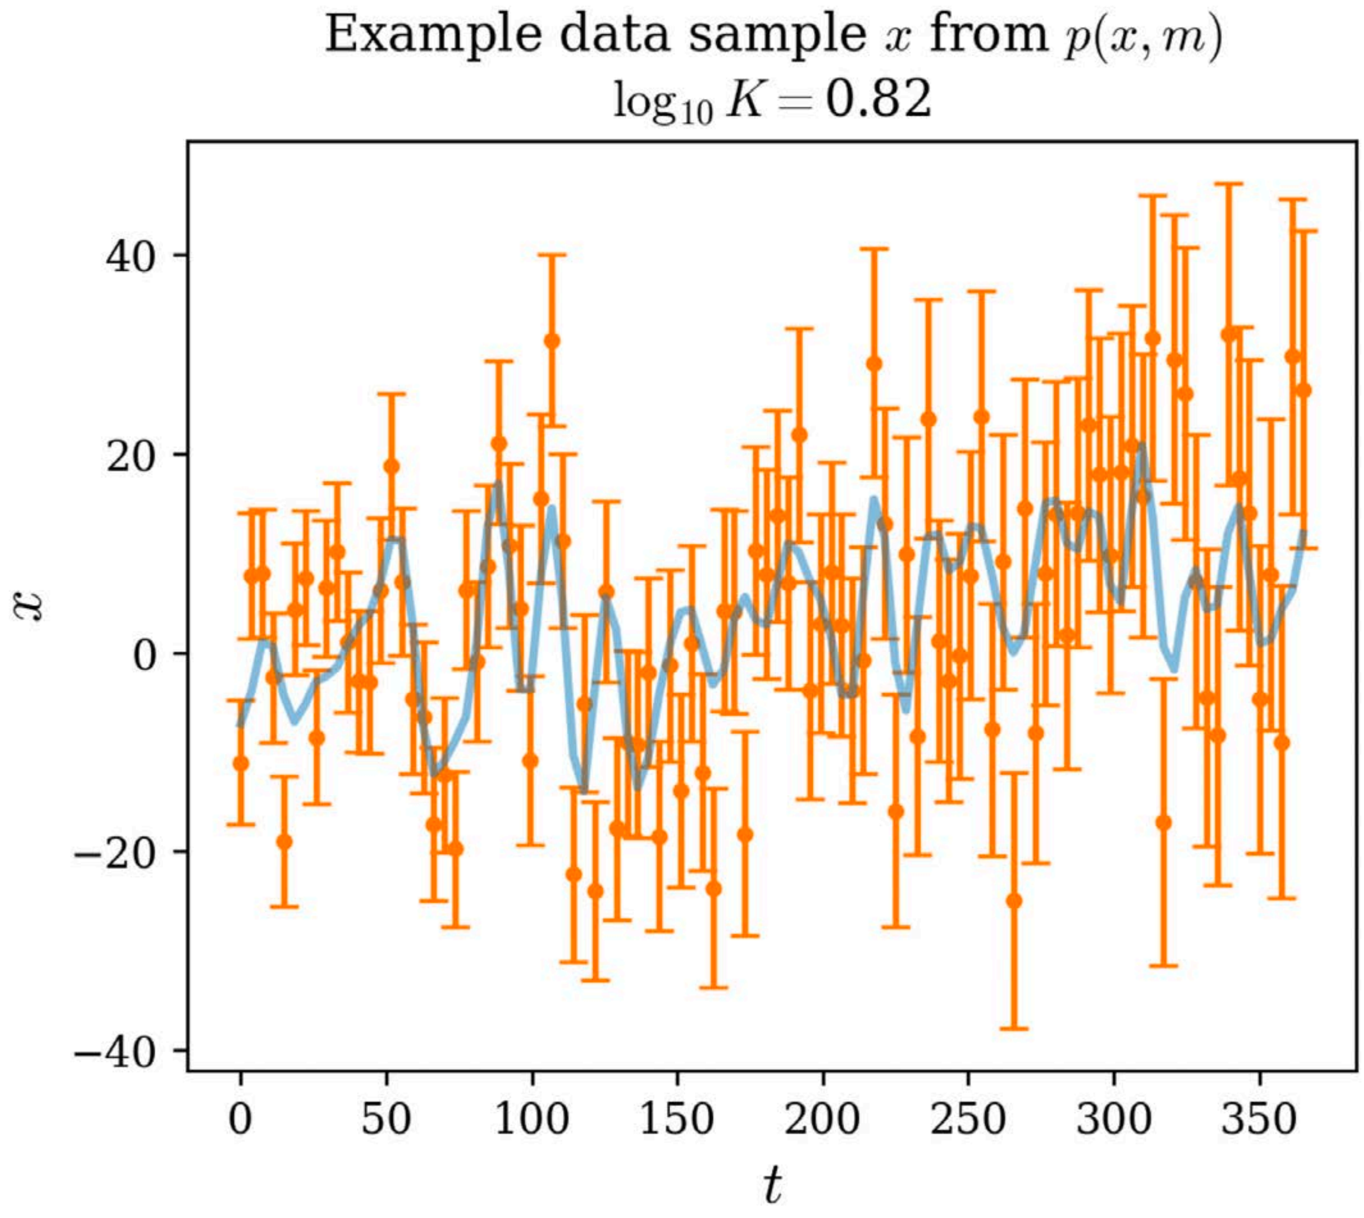

Model 1: Linear growth term

Model 2: No growth term

Time-series results

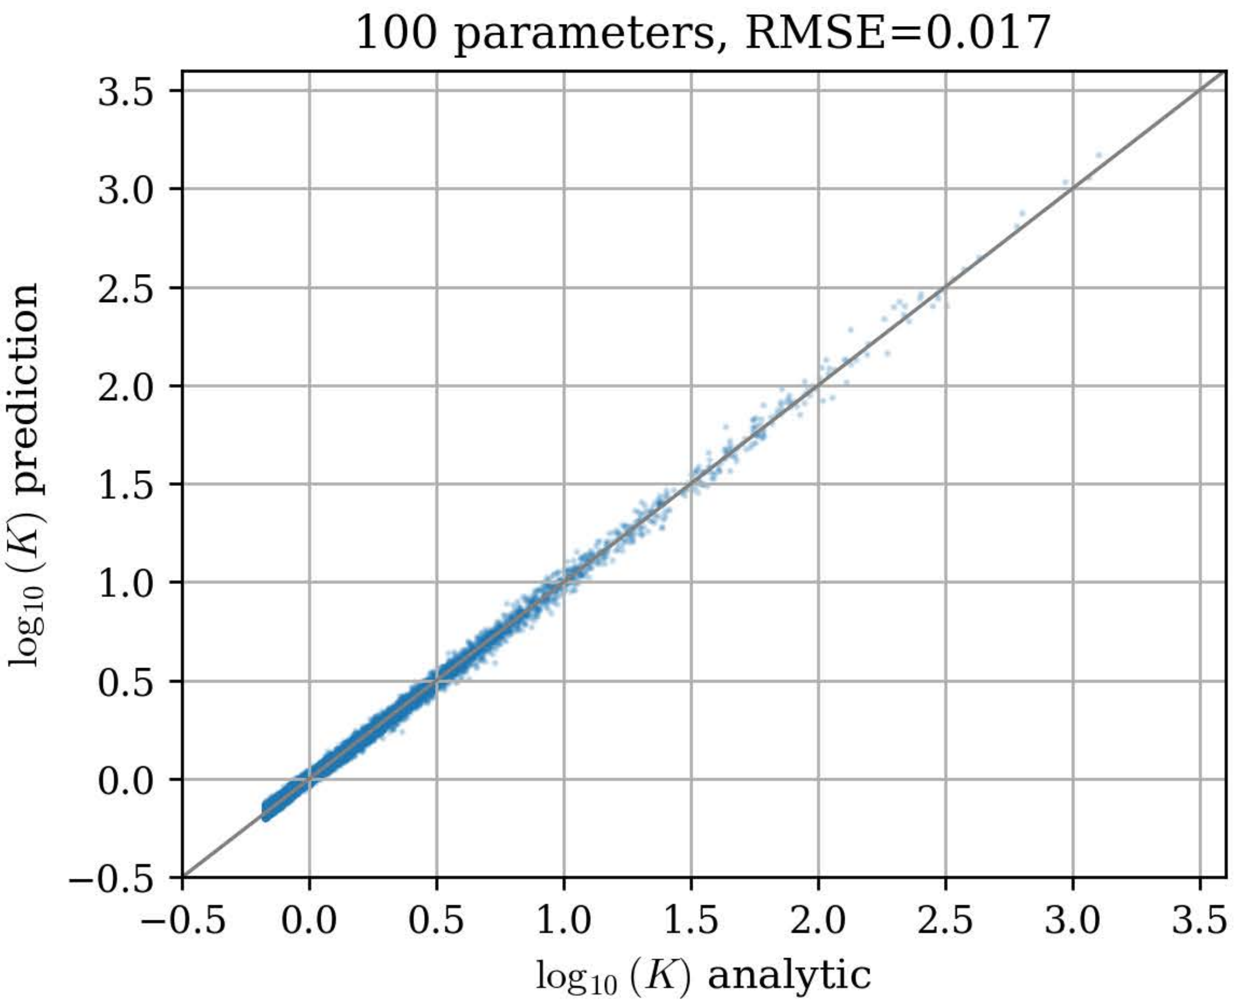

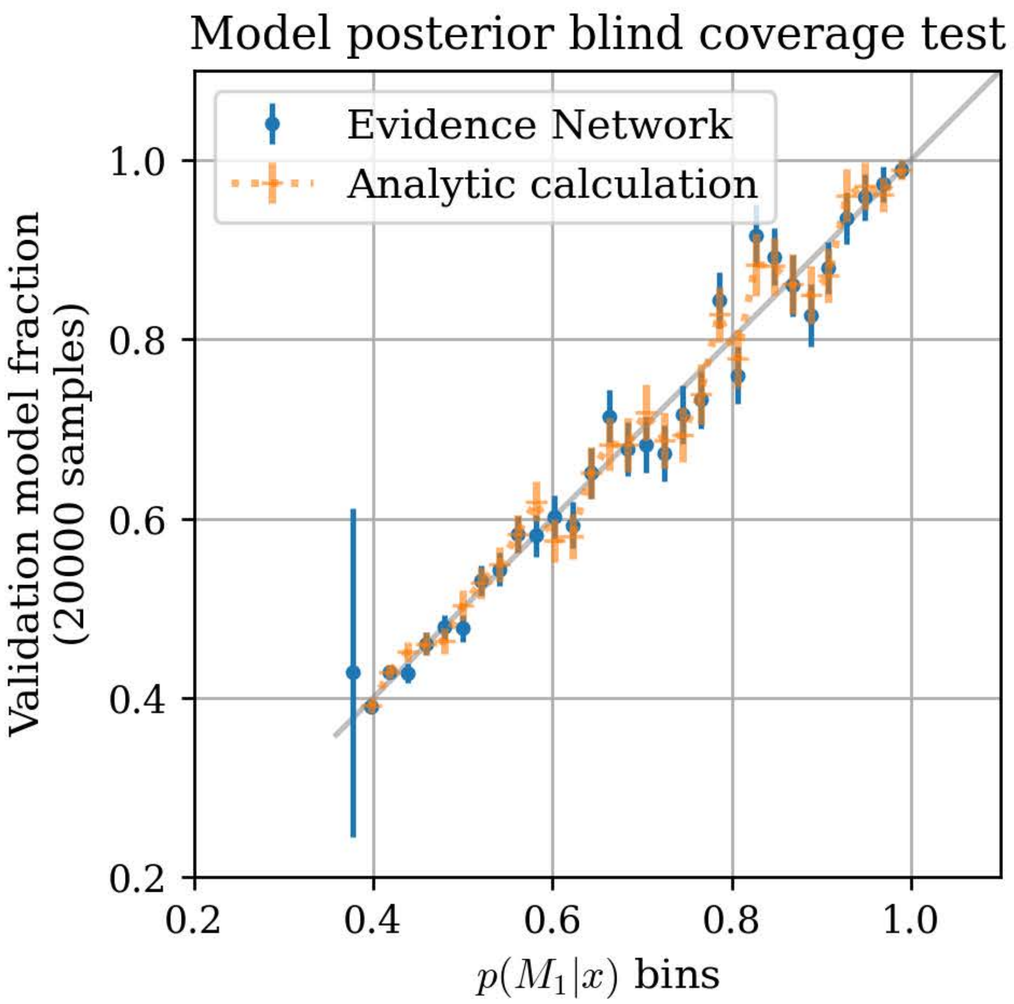

Versus alternative methods

1. Evidence Nets can work if alternatives intractable
2. More accurate than SOTA Nested Sampling with only 1% of the likelihood evaluations/samples
3. Accuracy 10x than neural density $p(x \mid M_1)$ ratios

Time-series results

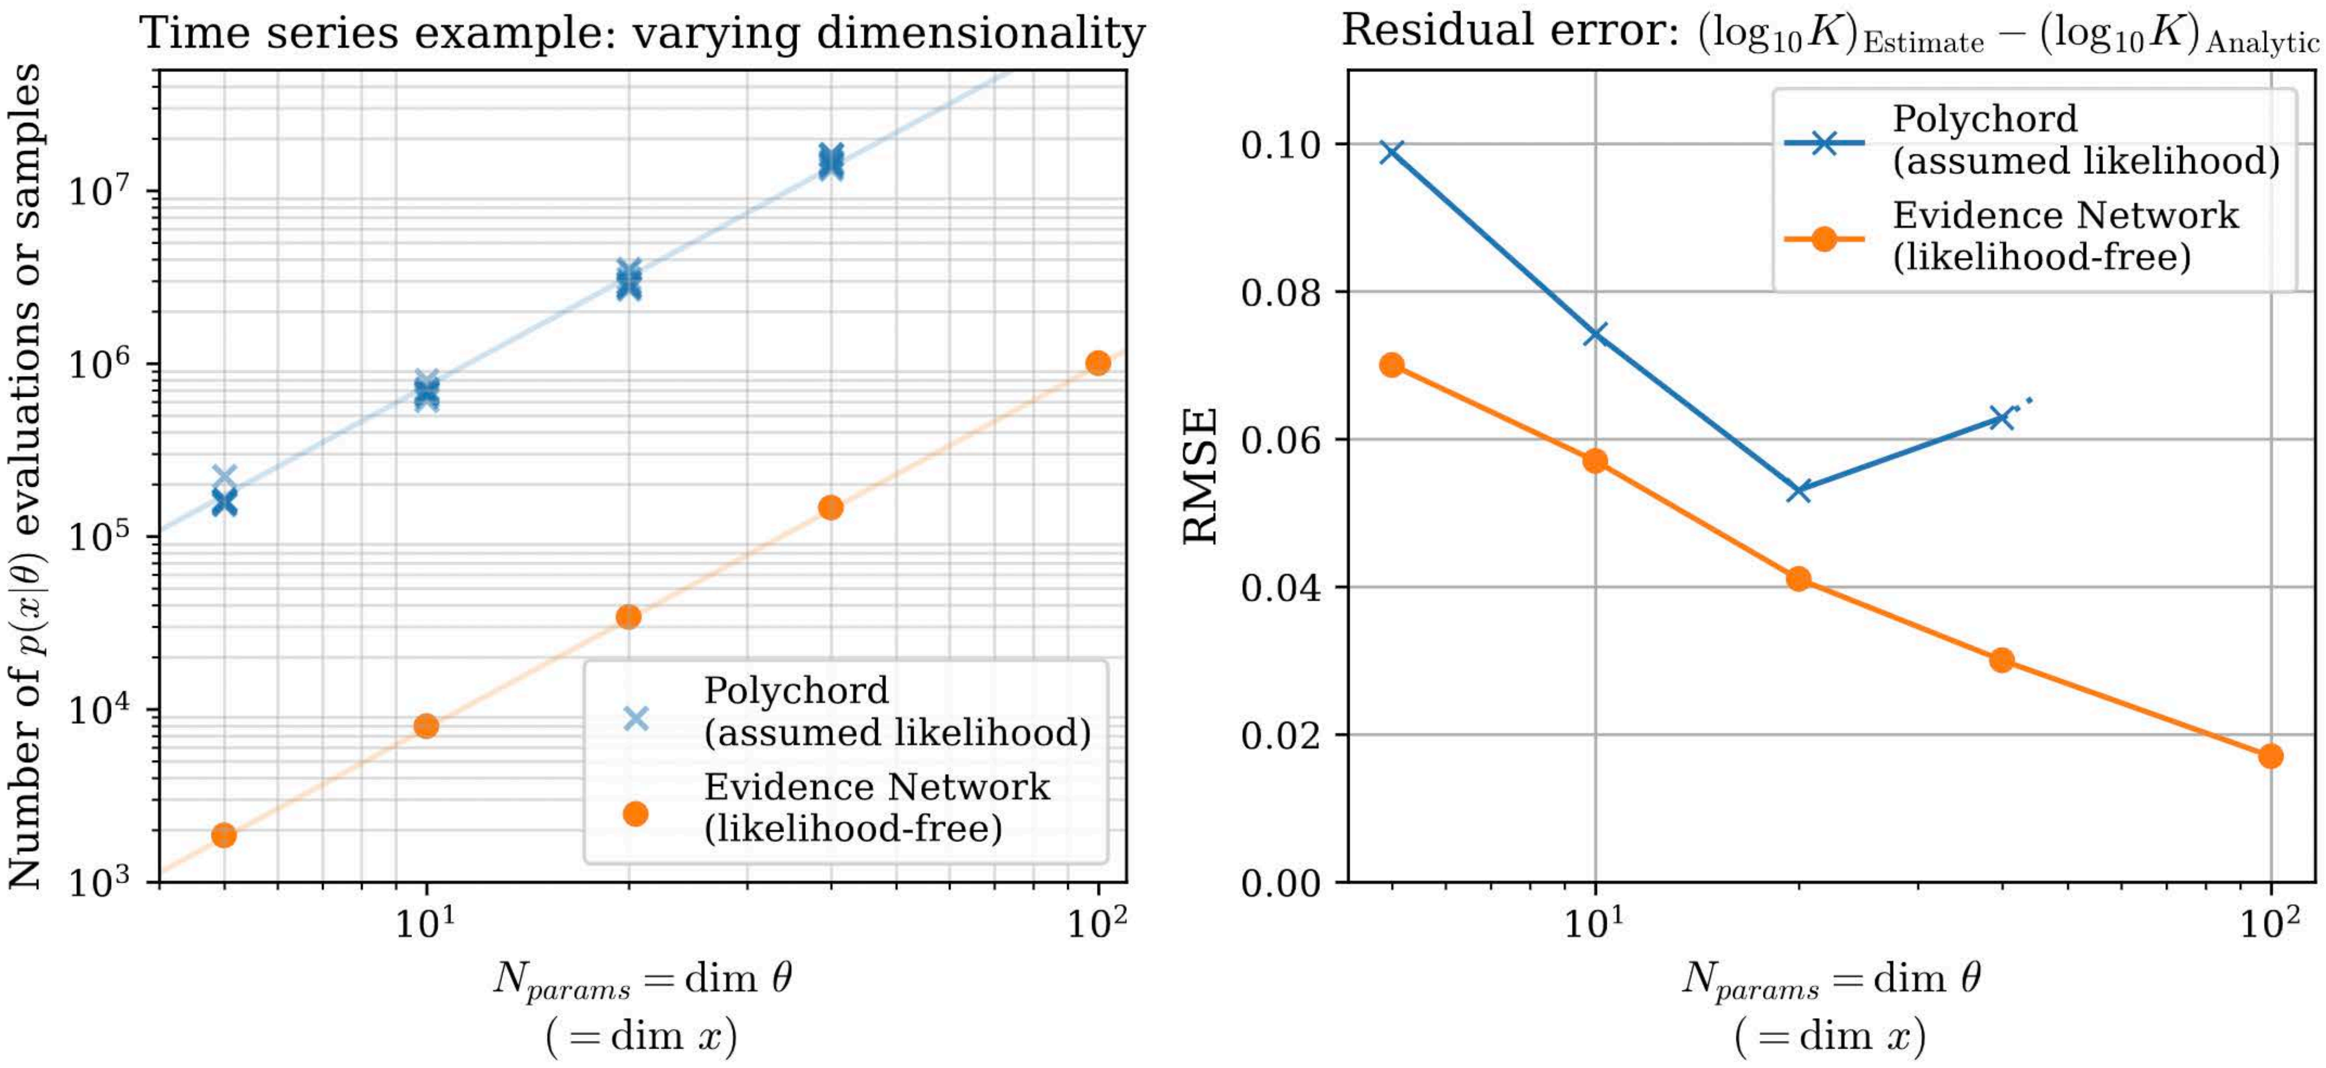

This is brave….

Dark Energy Survey data application

Model 1: galaxies are intrinsically aligned

Simple Evidence Network result:

$$\log_{10} K = -0.8\ (\pm 0.3)$$

Bonus section!

Extensions and intuition

- Extensions:
  - ✳ Absolute evidence calculation
  - ✳ Frequentist hypothesis testing
  - ✳ Posterior predictive testing

Pedagogical Example: Rastrigin Posterior

Pedagogical Example: Rastrigin Posterior

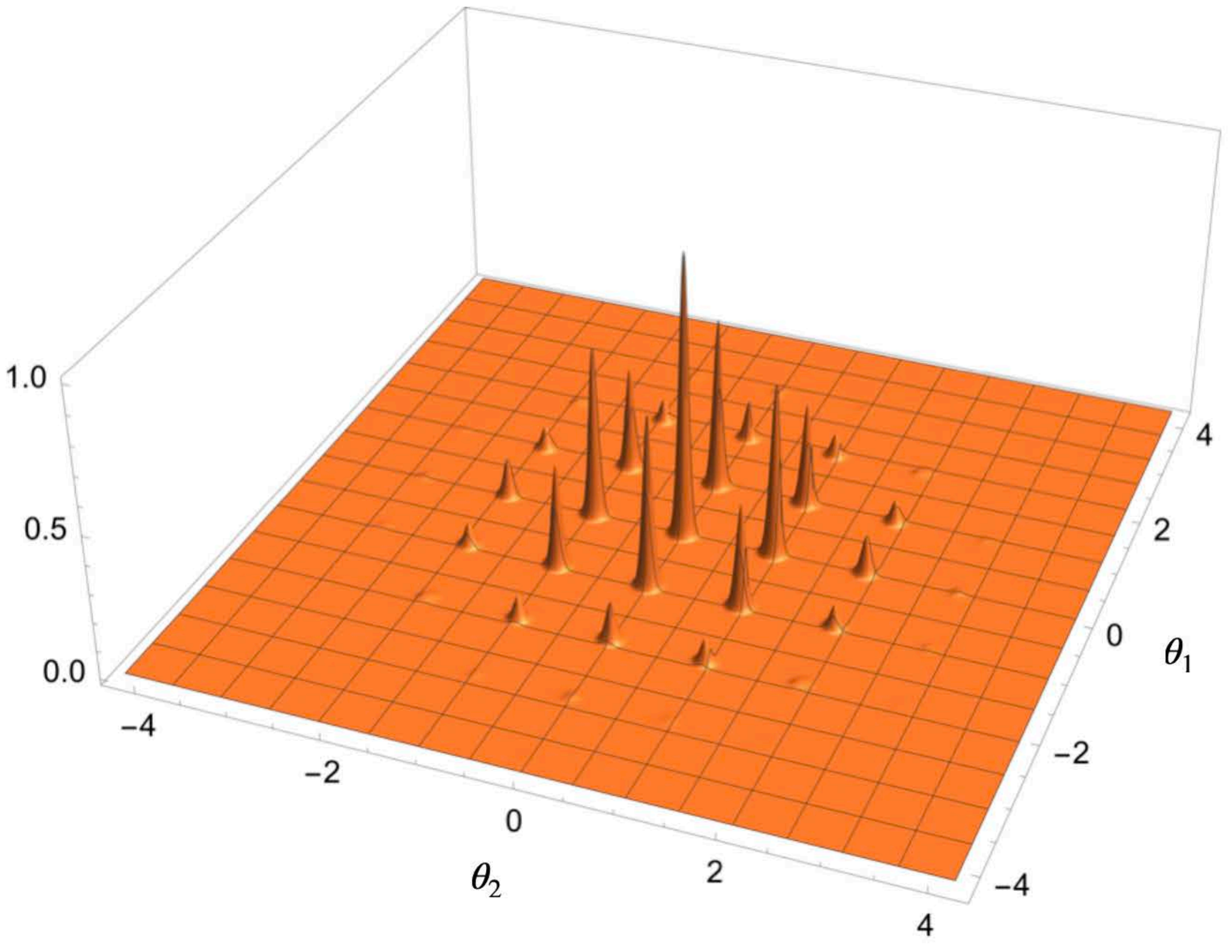

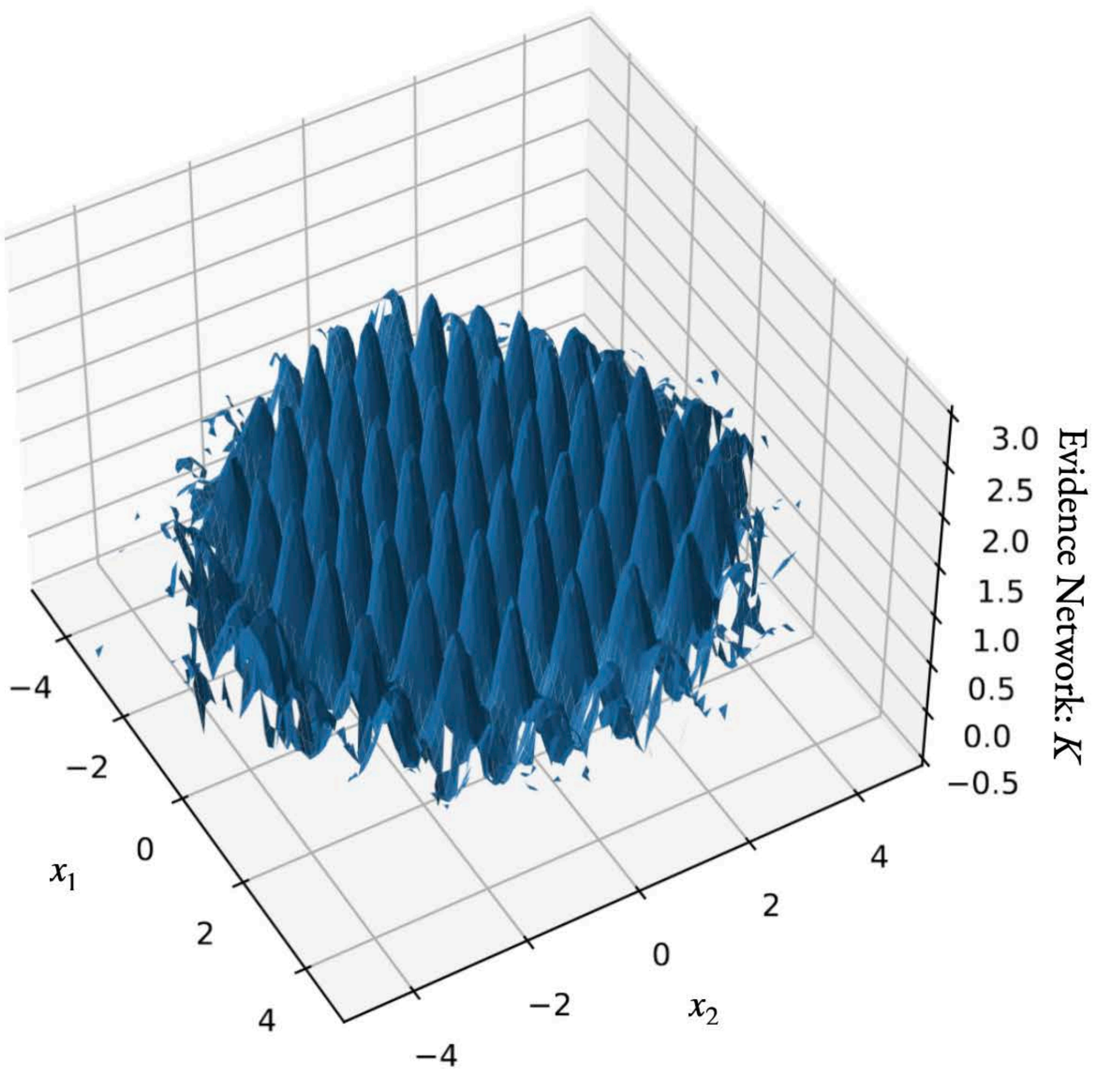

$$p(\theta_1, \theta_2 \,|\, x)$$

Evidence Networks: Extensions and applications

- Extensions:
  - ✳ Absolute evidence calculation
  - ✳ Frequentist hypothesis testing
  - ✳ Posterior predictive testing
- Fast, accurate, simple and work with previously intractable problems
- Which of your model comparison problems do you want to solve?

ArXiv: 2305.11241