A glowing blue neutron star is the central focus, surrounded by complex, swirling magnetic field lines that create a butterfly-like pattern. The background is a dark space filled with stars.

# Observing and simulating the inaccessible: Neutron stars, from NICER to SKA

**Joanna Berteaud**


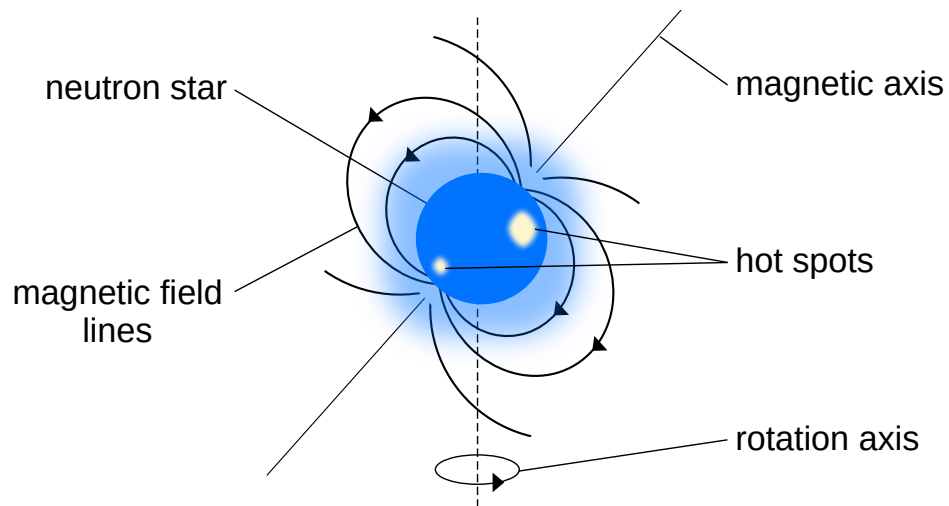
Cole Miller, Ananya Kulkarni (UMD, NASA)

Vanessa Graber, Gary Liu, Rebecca McShane (RHUL)



# Neutron stars...

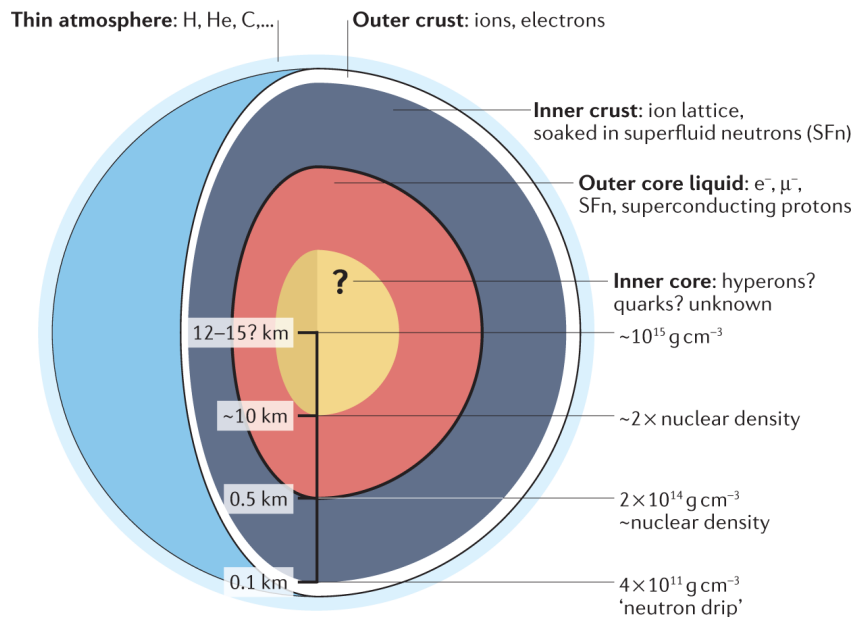
- End state of stellar evolution
- $M \sim 1-2 M_{\odot}$ ,  $R \sim 8-15$  km
- Multi-wavelength emission
- Steady spin down



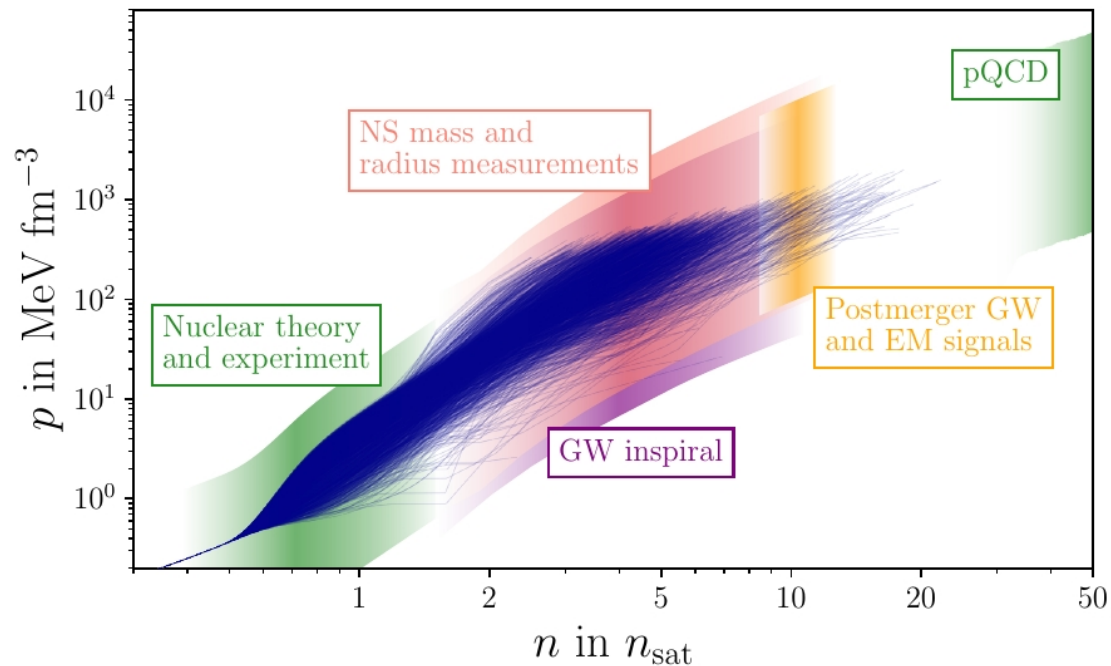
# ...and their equation of state

Koehn et al. (2024)

Unique laboratories to study  
cold, dense matter

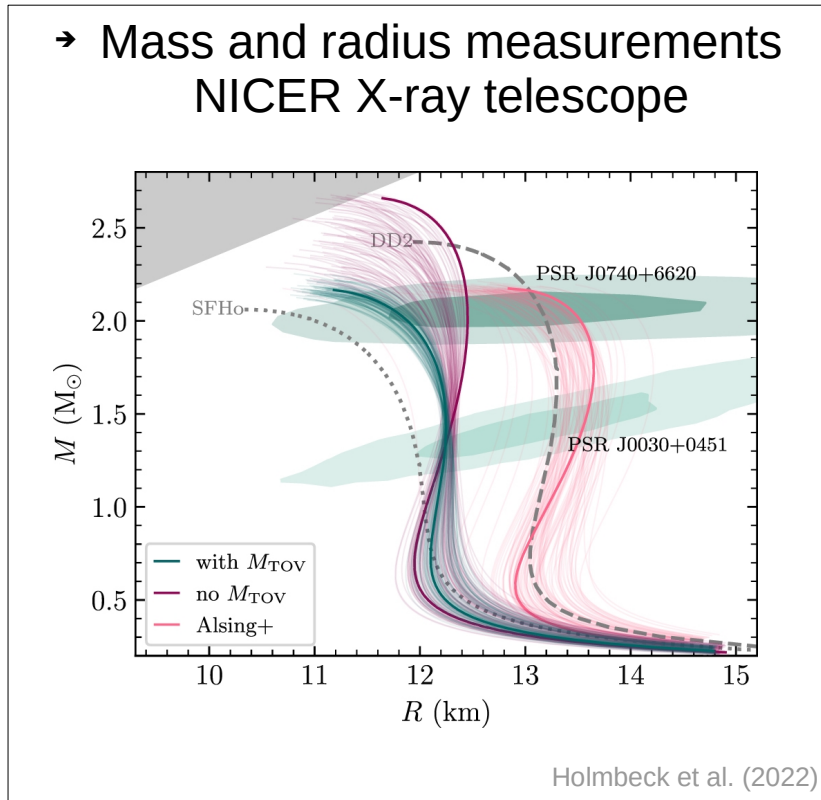


Yunes et al. (2022)




Equation of state (EOS):  
 $P(\rho, T, \dots)$

# Probing the EOS of neutron stars




# Probing the EOS of neutron stars

→ Mass and radius measurements  
NICER X-ray telescope



Holmbeck et al. (2022)


→ Spin-up glitches  
SKA and other radio observatories



Adapted from M. Antonelli


# I. Mass and radius: X-ray emission

Pulsed X-ray emission from neutron stars comes from hot spots on their surface. We can simulate their spectro-temporal emission.



# I. Mass and radius: X-ray emission

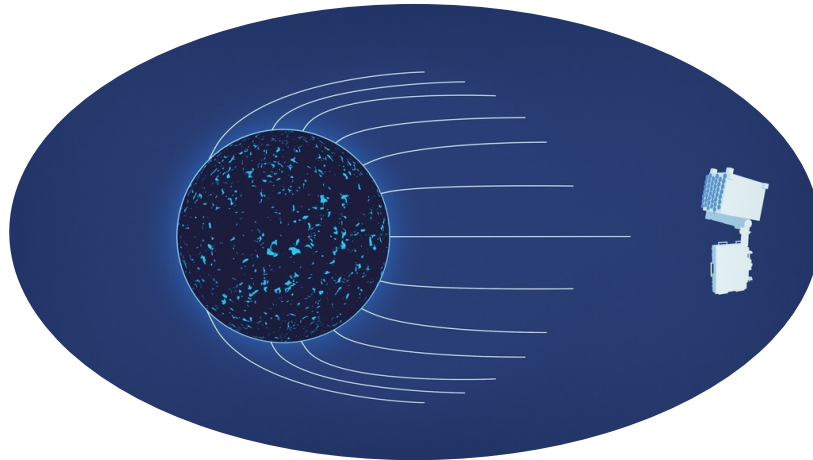
Pulsed X-ray emission from neutron stars comes from hot spots on their surface. We can simulate their spectro-temporal emission.



This emission depends on their mass, their radius, and other parameters.

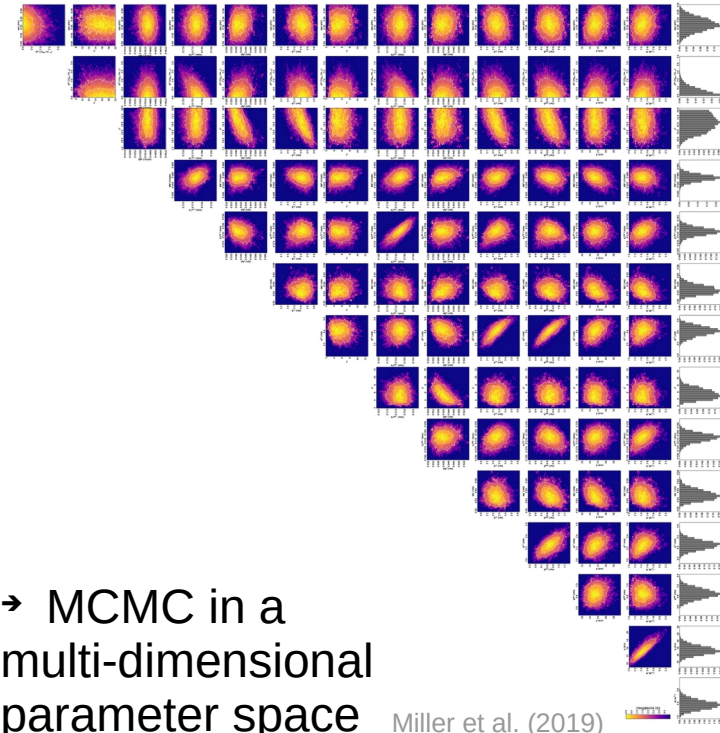
# I. Mass and radius: standard analysis

→ Relativistic ray-tracing



NASA


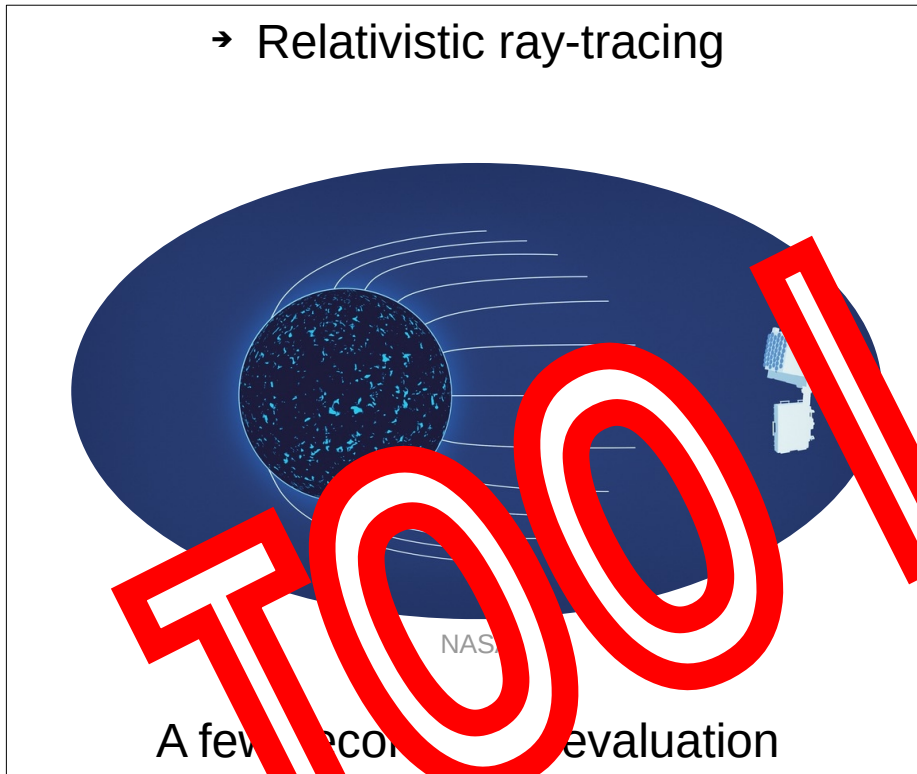
A few seconds per evaluation



→ MCMC in a multi-dimensional parameter space

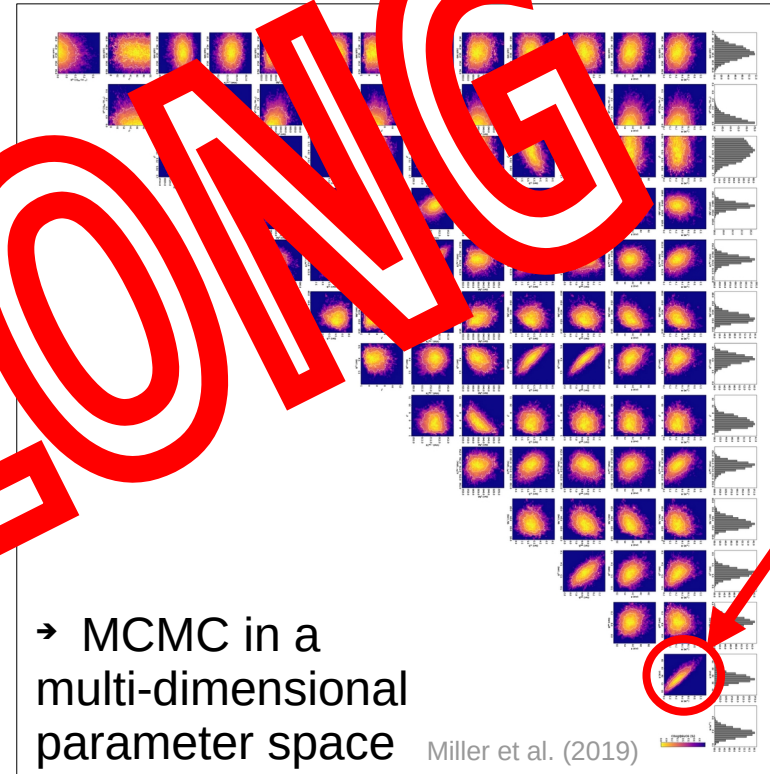

Miller et al. (2019)

# I. Mass and radius: standard analysis




**TOO LONG**

# I. Mass and radius: standard analysis



**TOO LONG**

# I. Mass and radius: WaFLe



# I. Mass and radius: WaFLe

## 3 input parameters

Training: a few  $10^4$  waveforms, days

Model evaluations:

- ✓ ~100× faster
- ✓ Accurate

Inference:

- ✓ ~20× faster
- ✓ Statistically correct
- × Toy model, not usable in real analyzes

# I. Mass and radius: WaFLe

## 3 input parameters

Training: a few  $10^4$  waveforms, days

Model evaluations:

- ✓ ~100× faster
- ✓ Accurate

Inference:

- ✓ ~20× faster
- ✓ Statistically correct
- × Toy model, not usable in real analyzes

## 6 inputs parameters

- ✓ Simplest model usable in analyzes

Model evaluations:

- ✓ ~100× faster
- × Accuracy not sufficient for MCMC


Training:  $\sim 10^8$  waveforms, weeks

# I. Mass and radius:




Simulation-based Inference  
of RPP Useful Parameters

## A. Kulkarni's summer 2025 internship


- ✓ Model evaluations with `waffle`
- ✓ Neural Ratio Estimation with `swyft`
- ✓ (M,R) inference done in seconds
- ✓ Well-behaved Bayesian posteriors
- ✗ Perfect instrument response
- ✗ No background signal





# II. Glitches: radio timing



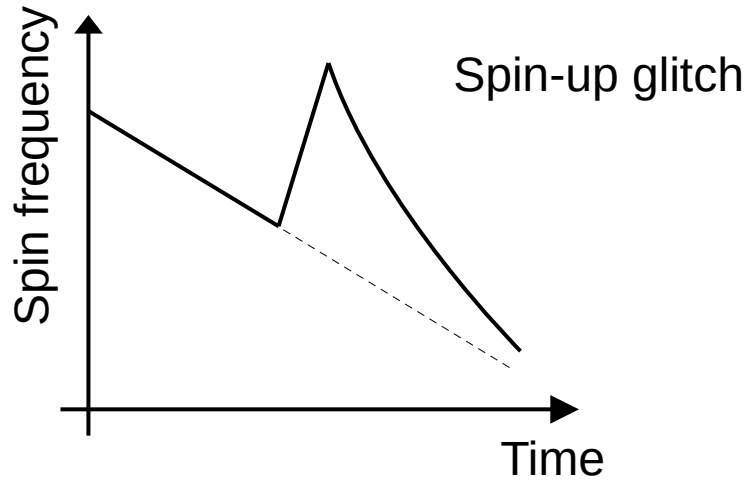
# II. Glitches: radio timing





x Timing noise,  
partly understood




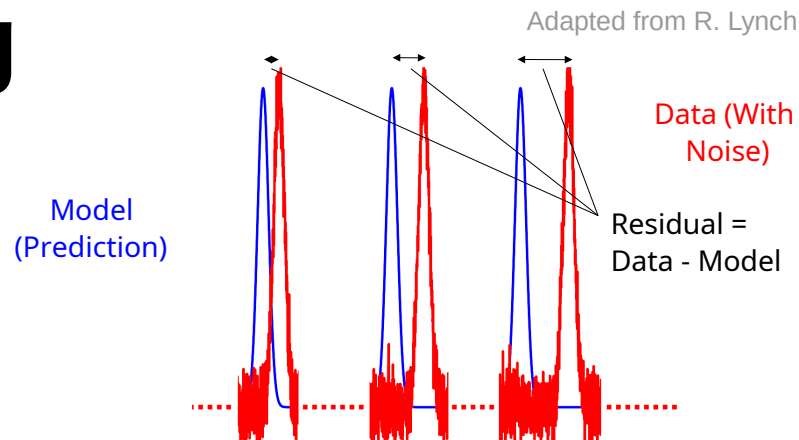
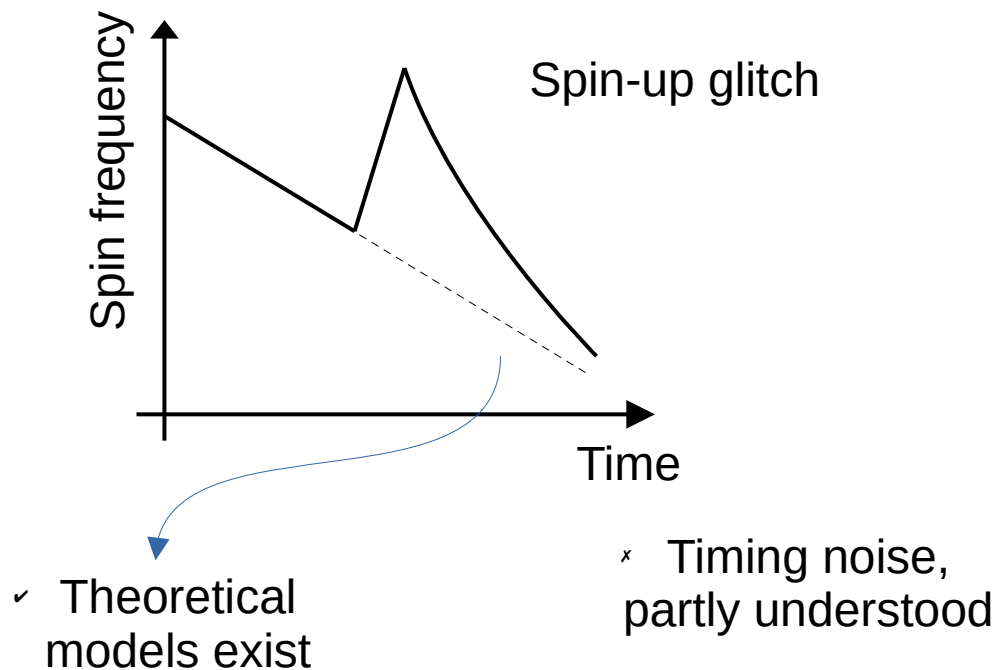
# II. Glitches: radio timing



x Timing noise, partly understood

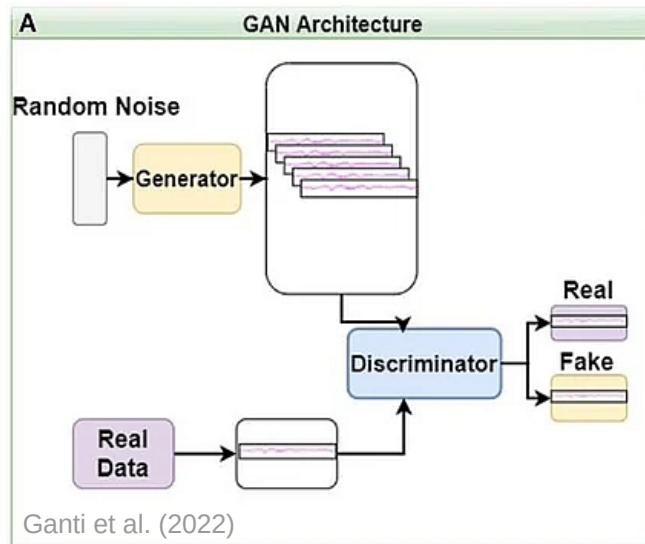


# II. Glitches: radio timing



# II. Glitches: Generative Adversarial Net.

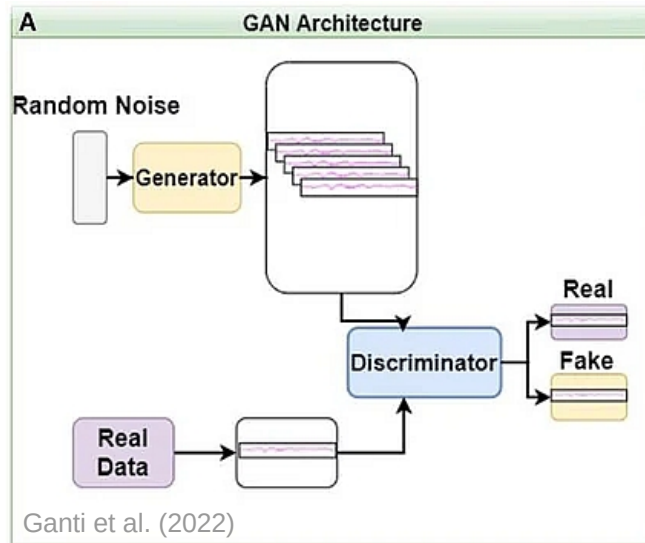
- Generative Adversarial Networks (GAN)



Simulate noise from noise?

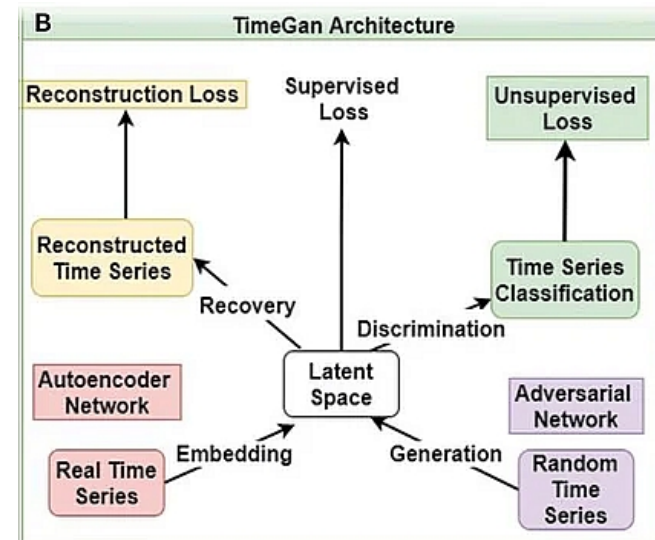
# II. Glitches: Generative Adversarial Net.

- Generative Adversarial Networks (GAN)



Simulate noise from noise?

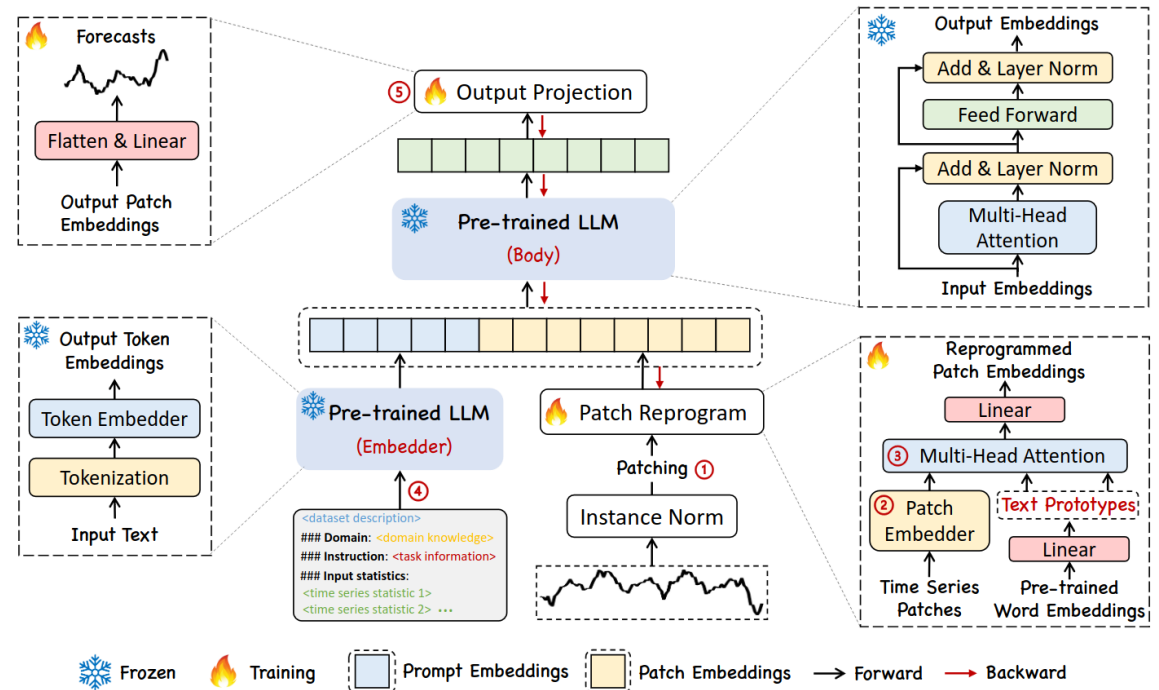
- TimeGAN for time series  
Yoon et al. (2019)



Issues: sparse and small dataset, irregular sampling, smallest time bin?


# II. Glitches: Large Language Model

Time-LLM (Jin et al. 2023):  
 → Time Series Forecasting  
 → Reprogramming of LLMs





# Summary

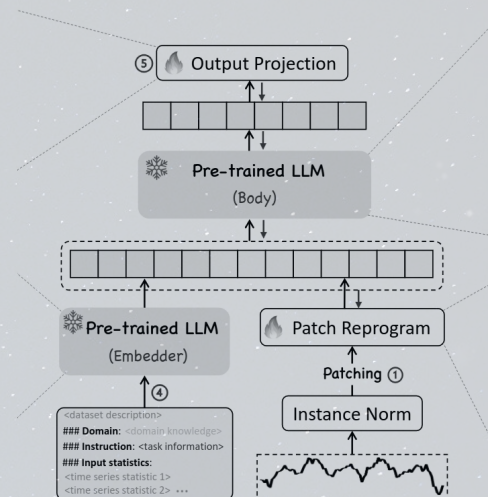
WaFLe, a 2D Res-Net



Simulation-based inference pipeline





GAN for time series

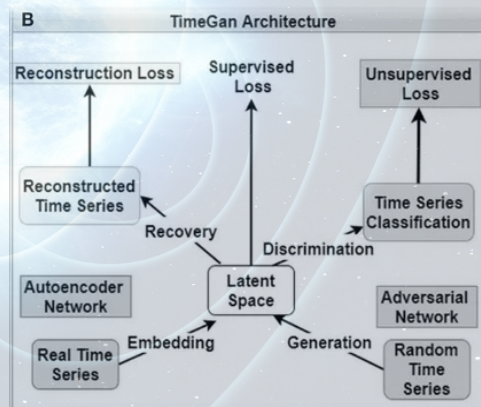



LLM reprogramming

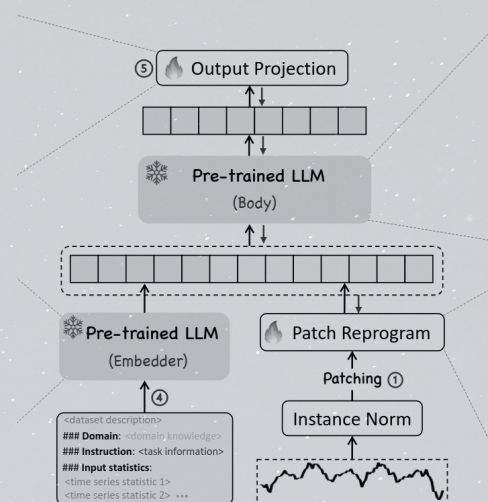
# Summary and perspectives



WaFLe, a 2D Res-Net

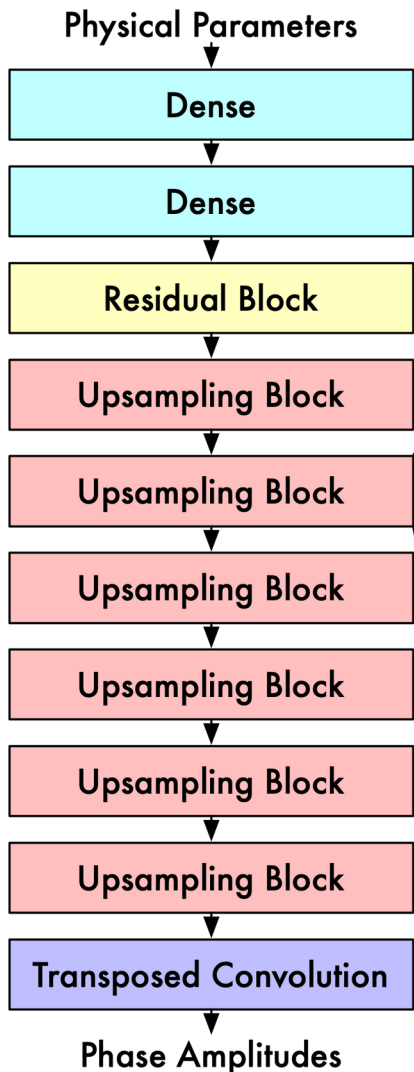


GAN for time series

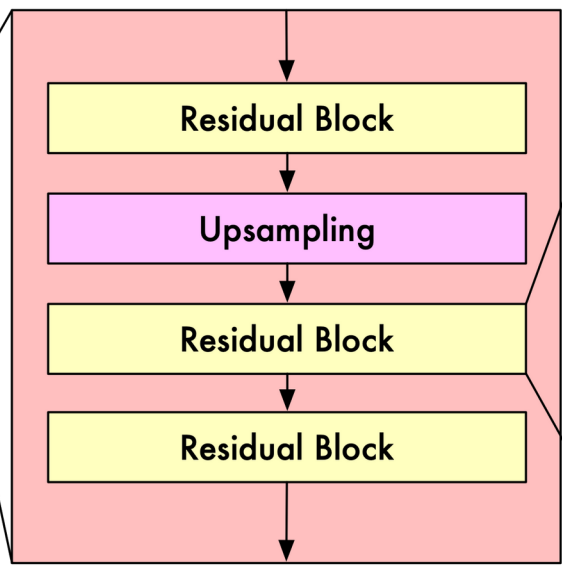


Gaussian process?  
Is AI needed?

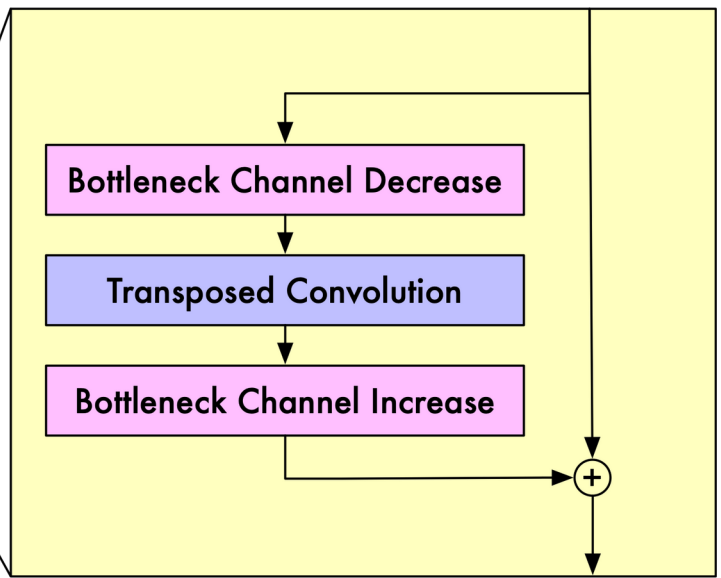
**Back-up slides**



Network




Upsampling Block

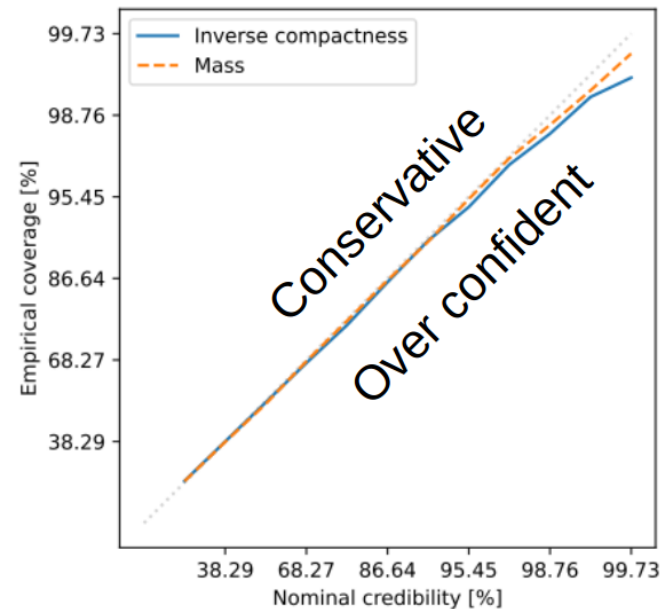


Residual Block

Good accuracy across the entire (M, R/M) parameter space



*“The  $x\%$  credible interval covers the true value in  $x\%$  of the cases”*



# Neural Ratio Estimation

Strategy: We estimate posteriors-to-prior ratio by training a binary classifier to discriminate between matching and scrambled (data, parameter) pairs.

$$r(\mathbf{x}; \mathbf{z}) \equiv \frac{p(\mathbf{x} | \mathbf{z})}{p(\mathbf{x})} = \frac{p(\mathbf{z} | \mathbf{x})}{p(\mathbf{z})} = \frac{p(\mathbf{x}, \mathbf{z})}{p(\mathbf{x})p(\mathbf{z})}$$

$\mathbf{x}, \mathbf{z} \sim p(\mathbf{x}, \mathbf{z})$

**Class 1: Matching (data, parameter) pairs**

(🐱, cat)

(🐶, dog)

(🏠, house)

(🐒, monkey)

(★, star)

$\mathbf{x}, \mathbf{z} \sim p(\mathbf{x})p(\mathbf{z})$

**Class 0: Scrambled (data, parameter) pairs**

(🐱, dog)

(🐶, cat)

(★, star)

(🐒, house)

(🏠, monkey)

Data:  $\mathbf{x}$

Parameter:  $\mathbf{z}$

# The maths

1. Introduce auxiliary classification variable  $Y$ :

$Y = 1$  : Joined draws

$$p(x, z | Y = 0) = p(x, z)$$

$Y = 0$  : Marginal draws

$$p(x, z | Y = 1) = p(x)p(z)$$

2. Assume flat prior for  $Y$ :

$$p(Y = 0) = p(Y = 1) = \frac{1}{2} \quad \Rightarrow \quad p(x, z | Y)p(Y) = \frac{1}{2} (\delta_{Y1}p(x, z) + \delta_{Y0}p(x)p(z))$$

3. Bayes theorem gives the posterior class probability given an  $(x, z)$  pair.

$$p(Y = 1 | x, z) = \frac{p(x, z)}{p(x, z) + p(x)p(z)} = \frac{r(x; z)}{r(x; z) + 1} \quad \text{and} \quad p(Y = 0 | x, z) = \frac{1}{r(x; z) + 1}$$

$\Rightarrow$  Training a binary classifier for  $Y$  is equivalent to learning the ratio  $r(\mathbf{x}, \mathbf{z})$ .

