

Neural network simulating the Lomb-Scargle periodogram

Topinskiy Vladislav, Baluev Roman

Saint Petersburg State University & Special Astrophysical Observatory, RAS

August 20, 2026

Problem: detecting periodic signals in astronomical data (exoplanets, variable stars)

Classical tool: Lomb-Scargle periodogram (least-squares spectral analysis)

Why neural networks?

- Potential for huge speed-up
- Ability to learn complex patterns beyond simple sinusoidal models

This work:

- Build a CNN that emulates the Lomb-Scargle periodogram
- Assess its performance, identify pitfalls, propose solutions

Test case: unevenly sampled time series with gaps + Gaussian white noise

Problem Statement

Time series: uneven time series of fixed length with periodic gaps

Signal model: $s(t) = A \sin(2\pi f_t + \varphi)$

Noise: Gaussian white (series are centred and normalised to remove offset/variance)

Observation gaps: periodic pattern with random parameters

- $\lambda \sim \mathcal{B}(1, 2)$ – fraction of observing window
- $\mu \sim \mathcal{B}(1, 1)$ – duty cycle

Classification: two classes

- ω_1 : signal present
- ω_2 : noise only

Data Preprocessing & Signal Input

Input features per frequency f_k :

- Real and imaginary parts of DFT
- Spectral window $\Omega(f_k)$ and $\Omega(2f_k)$ (related to generalised Lomb–Scargle)

Convolutional design:

- Layer 1: kernel $1 \times 6 \times m$ ($m = 16$ or 32)
- Layer 2: kernel $1 \times m \times 1$
- Fully-connected layer (size = number of frequencies)
- Output: single neuron with sigmoid activation

Advantages: fewer parameters, scalable to longer series




Figure: Network Architecture

A Bayesian approach was used to evaluate the results obtained. For each series, the probability of the presence of a sinusoidal component in the signal and the probability of its absence were calculated. The ratio of these probabilities is called the Bayes factor (R), which determines the maximum possible precision of the determination. The output of the artificial neural network was compared with the posterior probability of detecting the signal $p = \frac{R}{1+R}$. The Bayesian model works with probabilities of this type:

$$E_B(\mathbf{x}|\mathcal{M}, \mathcal{W}) = \int L(\mathbf{x}|\boldsymbol{\theta}, \mathcal{M}) \mathcal{W}(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (1)$$

where L is the likelihood function, $\mathbf{x}$ are the samples of the signal under study, $\mathcal{M}$ is the signal model, $\boldsymbol{\theta}$ are its parameters and $\mathcal{W}$ is the weight function.

The probability E_B is also known as the marginal probability. The ratio of two Bayesian probabilities calculated on two opposing models is called the Bayes factor:

$$R_B = \frac{E_B(x|\mathcal{M}_2, \mathcal{W}_2)}{E_B(x|\mathcal{M}_1, \mathcal{W}_1)}. \quad (2)$$

Weight Initialisation

Training from random weights fails – dead neurons, unstable gradients
Empirical initialisation for the first layer:

- For each frequency block define
 $H_i = \sqrt{C_i^2 + S_i^2}$, $R_i = -B_i/H_i$, $\alpha_i = \text{atan2}(S_i, C_i)$
- Distributions: $H_i \sim \mathcal{N}$, $R_i \sim \mathcal{N}$, $\alpha_i \sim \mathcal{U}(-\pi, \pi)$ output weights D_i half-normal

Additionally, weights handling $\Omega(2f)$ linked to DFT weights:

$$\frac{S''}{C''} = \frac{SC}{C''^2 - S''^2}, S''^2 + C''^2 \approx 0.5(S^2 + C^2)$$

This "hint" dramatically improves convergence;
optimal block size = 16 or 32




Figure: The ratio of the model probability of signal existence to the theoretical probability. Color denotes the density of points in the control sample (40,000 points), two-dimensional histogram

- Test sample: 40,000 series
- Posterior Bayesian probability: 0.3851
- Network BCE ($m = 16$ or 32): $\approx 0.4048/0.4047$
- Root-mean-square error vs Bayesian posterior: ≈ 0.07

- Network inference: 10^6 series processed in ~ 1 minute on a GPU
- Bayesian numerical integration: 40,000 series took ~ 1 week on CPU
- **Speed-up factor:** $10^5 - 10^6$
- Implication: the method is practical for large astronomical surveys and catalogues

- Main limitation: training is extremely sensitive to initial weights – requires empirical ”recipe”
- Planned improvements:
 - Increase K and spectral range
 - Adapt to non-sinusoidal signals (e.g., eccentric orbits)
 - Handle correlated noise (non-white spectra)
- Outlook: this work provides a foundation for using neural networks as fast Bayesian-inspired detectors in real-world astronomical pipelines

- Built a convolutional neural network that emulates the Lomb–Scargle periodogram
- Its output approximates the Bayesian posterior probability for signal presence
- Achieved good accuracy ($\text{BCE} \approx 0.405$, $\text{RMSE} \approx 0.07$) with massive speed-up
- Identified and solved the critical initial-weight problem via empirical distributions
- The approach opens the door for scalable, deep-learning-based periodicity detection in astrophysics