Gravitational-wave background detection using machine learning

Hugo Einsle, Marie-Anne Bizouard, Mairi Sakellariadou, Tania Regimbau

October 14th, 2025

Gravitational wave background

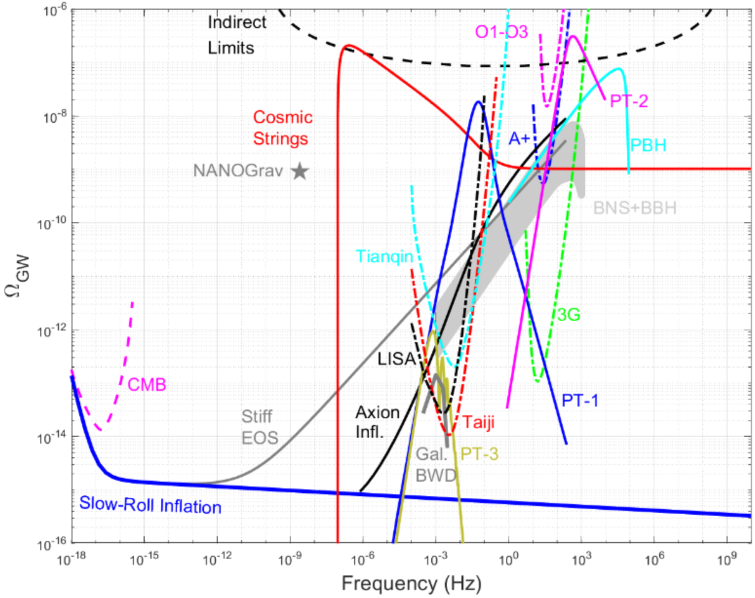

Motivations and proposed solution

Goal: detecting and disentangling GWB components in LVK data.
Target components: astrophysical (CBC) and cosmological.
Proposed solution: hybrid approach combining deep learning and Bayesian inference.

- Custom Multi-Scale Multi-Headed Autoencoder (MSMHAutoencoder) separates noise and GWB.
- Physics-informed training, curriculum learning for low SNR signals.
- MCMC parameter estimation for GWB component inference.

What is an autoencoder?

Autoencoders are neural networks with bottleneck structure trained to compress and reconstruct the input.

- **Structure:** An encoder maps the input data to a compressed latent space; a decoder reconstructs the input from this representation.
- **Objective:** Learn a compact encoding that preserves the essential structure and information, even in the presence of noise.
- **Key idea:** The compressed representation (latent space) acts as a denoising bottleneck, filtering out unstructured noise.

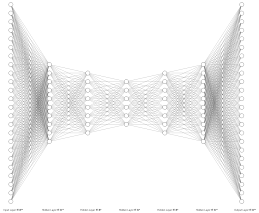

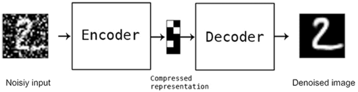

Method - deep learning and MCMC

Inputs: $x \in \mathbb{R}^{M \times N}$, M consecutive H1-L1 cross-corelation spectra (noise + signal), N frequency bins.
Outputs: noise estimation $\hat{n}$, signal estimation $\hat{s}$
Model. $\Omega_{\mathrm{GW}}(f) = \Omega_\alpha (f/f_{\mathrm{ref}})^\alpha$ at $f_{\mathrm{ref}} = 25\,\mathrm{Hz}$. Evidence by $\log_{10}(\mathrm{BF})$.

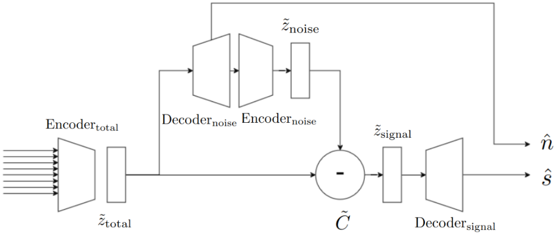

Training: physics-informed loss function

The training objective is a weighted sum of three terms:

$$\mathcal{L}_{\rm total} = \lambda_s \mathcal{L}_{\rm spectral} + \lambda_l \mathcal{L}_{\rm latent} + \lambda_c \mathcal{L}_{\rm consist}$$

- **Spectral loss:** $\mathcal{L}_{\rm spectral}$ combines signal and noise reconstruction accuracy with a spectral smoothness penalty
- **Latent loss:** $\mathcal{L}_{\rm latent}$ ensures internal consistency between predicted and true noise in latent space
- **Consistency loss:** $\mathcal{L}_{\rm consist}$ imposes self-consistency between reconstructed and original input
- **Loss weights:** $\lambda_s, \lambda_l, \lambda_c$

Training : Curriculum learning for low SNR

- **Motivation:** realistic GWBs have amplitudes $\ll$ detector noise $\rightarrow$ challenging separation for low SNR GWB.
- **Solution:** curriculum learning : start with easily separable, high-amplitude signals and progressively reduce amplitude to reach realistic low SNR.
- **Impact:** smooth transition from easy (high SNR) to hard (realistic SNR), avoiding model collapse at low amplitudes and boosting final sensitivity.

Astrophysical and cosmological backgrounds

Models :

- Astrophysical GWB: compact binary coalescences, modeled as power law:

$$\Omega_{\mathrm{GW}}(f) = \Omega_\alpha \left(\frac{f}{f_{\mathrm{ref}}}\right)^\alpha$$

  Typical $\alpha = 2/3$ for binary sources.
- Cosmological GWB: early universe processes $\Omega_{GW}(f) = \Omega_0$

Simulations :

- Noise: Gaussian & uncorrelated
- Astrophysical GWB: CBC population-synthesis catalogs IMRPhenomXPHM
- Cosmological GWB: correlated between detectors, generated from flat spectrum

LIGO A+ design sensitivity : Training

Datasets:

- Training and validation sets: 47.4 days of data, mixture of pure noise, astrophysical and cosmological GWB injections.
- Testing dataset: 23.7 days of data held-out segments with known injections, astrophysical and cosmological GWB injections.
- Segments of 2048 s sampled at 512 Hz.
- Inputs stack $M = 12$ spectra.

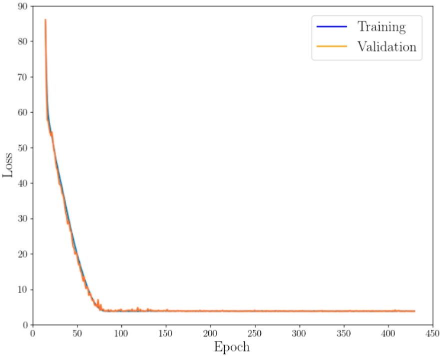

LIGO A+ design sensitivity : Training - Sensitivity

Sensitivity: the lowest GWB amplitude that the model can reliably distinguish from noise.
At each training epoch, varying amplitude Ω_α are injected and the mean squared error (MSE) between the reconstructed and true signal is computed:

$$\mathrm{MSE}_{\hat{s}} = \frac{1}{N}\sum_{j=1}^{N}\left[\hat{s}(f_j|\Omega_\alpha) - s(f_j|\Omega_\alpha)\right]^2 < 0.01 \ .$$

where s is the true spectrum and $\hat{s}$ is the predicted spectrum.

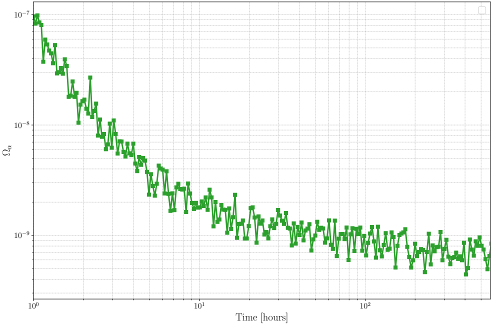

LIGO A+ design sensitivity : results - bayesian inference

- **Objective:** Quantify evidence for a GWB in the reconstructed spectrum.
- **Models:**
  - H_0 (null): Data consists of noise only.
  - H_1 (signal): astrophysical (CBC) signal.
  - H_2 (signals): astrophysical and cosmological signals.
- **Bayes factors:**

$$\mathrm{BF} = \frac{Z(H_1)}{Z(H_0)} \qquad \log_{10}(\mathrm{BF}) \text{ interpreted as detection significance}$$

$$\mathrm{BF}_{\mathrm{Cosmo}} = \frac{Z(H_2)}{Z(H_1)} \qquad \log_{10}(\mathrm{BF}) \text{ interpreted as detection significance}$$

- **Data** : LLO and LHO

LIGO A+ design sensitivity : results - bayesian inference

The Bayes factor $\log_{10}(\mathrm{BF})$ quantifies support for the presence of a GWB. We set $\log_{10}(\mathrm{BF}) > 3$ as **decisive evidence** for GWB detection. We inject and retrieve a **BBH-only** GWB with varying amplitudes

$$\Omega_{\mathrm{GWB}}(f|\theta) = \Omega_{\mathrm{BBH}}(f|\Omega_\alpha, \alpha = 2/3)$$

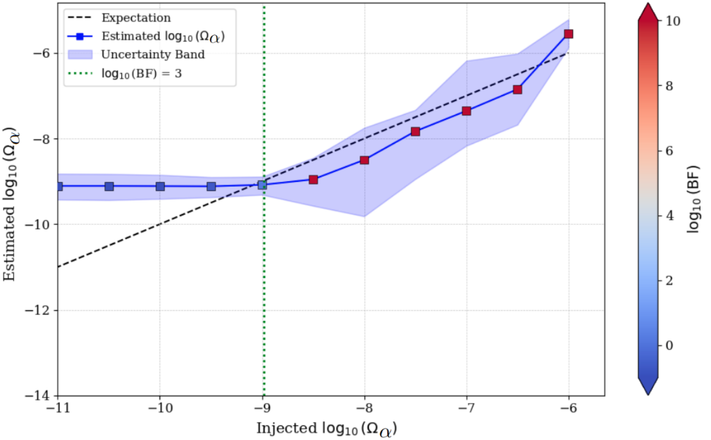

LIGO A+ design sensitivity : Results - Bayesian inference

We inject and retrieve a **Cosmo** GWB with varying amplitudes in the presence of a BBH GWB such that $\Omega_{\mathrm{BBH}} = 10^{-9}$ at 25Hz

$$\mathrm{H}_0 : \quad \Omega_{\mathrm{GWB}}(f \mid \theta) = \Omega_{\mathrm{BBH}}(f \mid \Omega_\alpha, \alpha = 2/3)$$

$$\mathrm{H}_1 : \quad \Omega_{\mathrm{GWB}}(f \mid \theta) = \Omega_{\mathrm{BBH}}(f \mid \Omega_\alpha, \alpha = 2/3) + \Omega_{\mathrm{Cosmo}}(f \mid \Omega_0, 0)$$

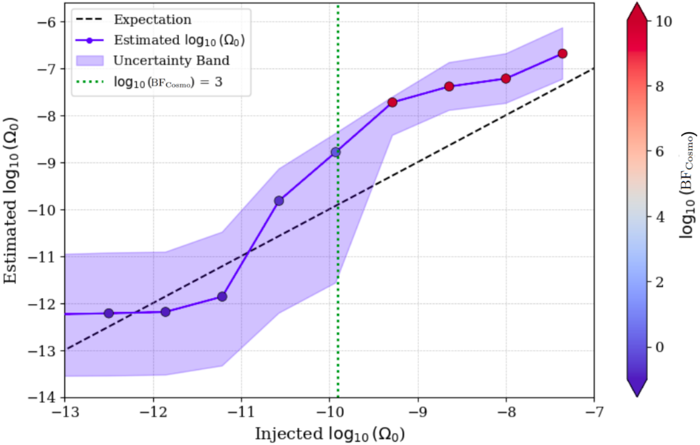

Comparison of posterior distributions for $\log_{10}(\Omega_0)$ and $\log_{10}(\Omega_\alpha)$ at the decisive detection threshold $\log_{10}(\mathrm{BF}) = 3$. **Left:** Fit with α relaxed (free parameter). **Right:** Fit with α fixed to the injected value. The credible intervals for Ω_0 and Ω_α remain consistent; relaxing α shifts the posterior within one σ uncertainty of fixed α hypothesis.

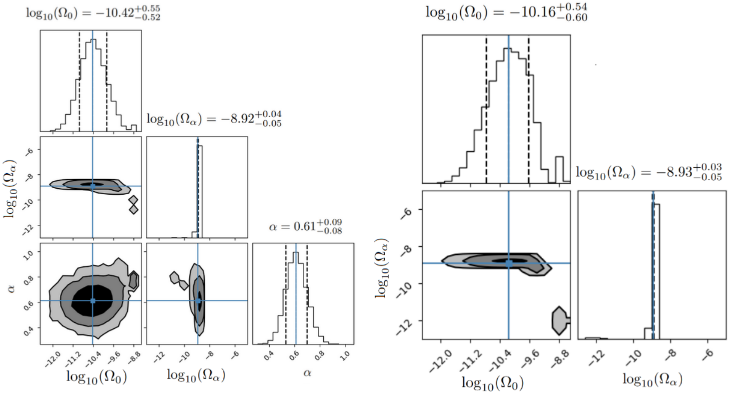

Discussion : A+ design sensitivity

- **MSMHAutoencoder successfully estimates a GWB** in simulated design LVK data and disentangle components.
- **Confident astrophysical GWB detection**.
  - Comparable to expected BNS+BBH GWB upper range.
  - Achieves confident detection at $\Omega_\alpha \sim 10^{-9}$ (25 Hz) ($\log_{10}(\text{BF}) = 3$) with 47.4 days of training data
  - To achieve a 5σ significance sensitivity, with cross-correlation method one would need 1.6 years of coincident data from LLO-LHO network at design sensitivity.
- **GWB components disentanglement**: Sensitive to CBC GWB of 10^{-9} and cosmological GWB of 1.3×10^{-10} with 47.4 days of training data.
- **Advantages**:
  - Relaxed constraint on CBC spectral index does not affect accuracy.
  - Good GWB measurement accuracy obtained significantly quicker than cross-correlation.

O4a sensitivity : datasets - noise

Two mock datasets based on O4a PSDs :

- **LLO and LHO search**
- **All lines**: O4a PSD with spectral lines.
- **No lines**: de-lined PSD for idealized noise.

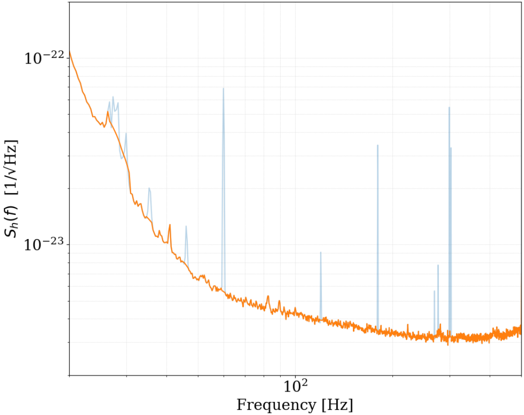

O4a sensitivity : training - datasets and parameters

Datasets.

- Training/validation: ≈ 108 days. Mixtures of pure noise, CBC (BBHs, BNSs, NSBHs) and cosmological ($\alpha = 0$) injections.
- Test (hold-out): ≈ 108 days.
- Segments of $2048\,\mathrm{s}$ sampled at $512\,\mathrm{Hz}$.
- Inputs stack $M = 12$ spectra.
- Two models trained for each noise condition : *with lines* vs *without lines*.

O4a sensitivity : results - CBC sensitivity (with lines)

O4a PSD with lines. $\Omega_\alpha = 4.57 \times 10^{-9}$ at 25 Hz for $\log_{10}(\mathrm{BF}) = 3$.

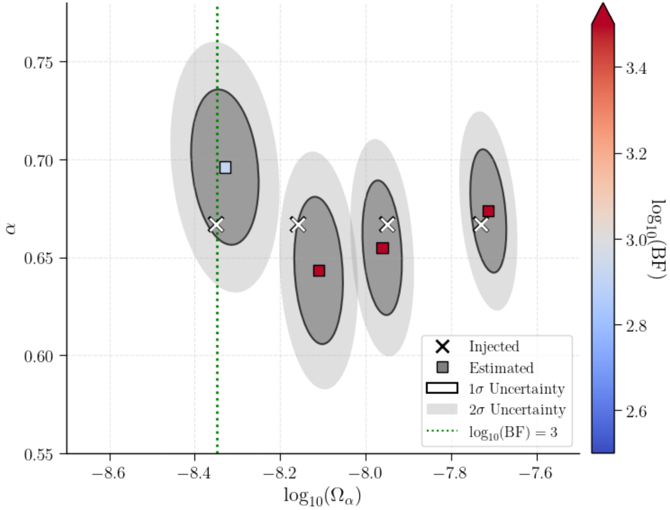

O4a sensitivity : results - CBC sensitivity (no lines)

O4a PSD without lines. Sensitivity improves to $\Omega_\alpha = 1.73 \times 10^{-9}$ at 25 Hz for $\log_{10}(\mathrm{BF}) = 3$.

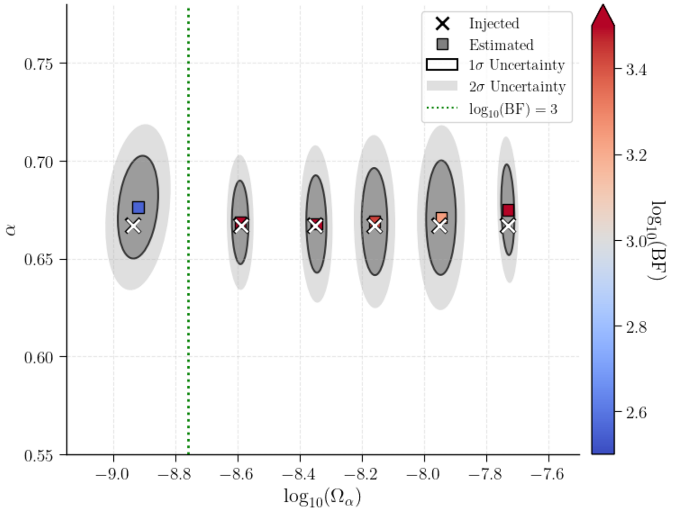

O4a sensitivity : results - cosmological component with CBC present

Fix CBC at $\Omega_{CBC} = 0.9 \times 10^{-9}$ from GTC-4.0 CBC rates (arxiv:2508.20721).
Posterior distribution of Ω_0 with $\log_{10}(BF_{\mathrm{cosmo}}) = 3$:
$\Omega_0 = 3.91 \times 10^{-10}$ (no lines), $\Omega_0 = 5.4 \times 10^{-10}$ (with lines).

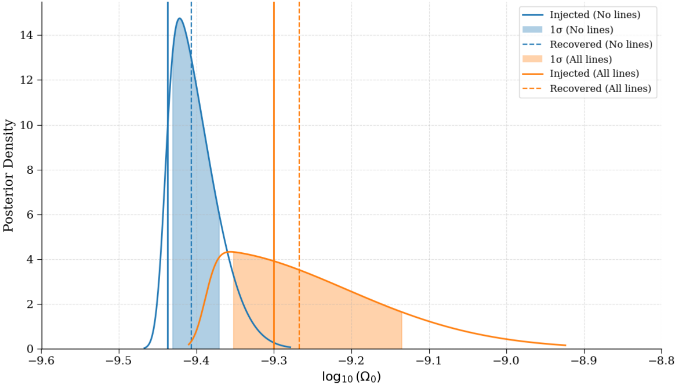

Discussion

Discussion : O4a sensitivity

- MSMH-Autoencoder yields decisive CBC evidence at $\Omega_\alpha = 4.57 \times 10^{-9}$ (with lines) and 1.73×10^{-9} (no lines).
- Given a CBC GWB with amplitude 0.9×10^{-9} we detect a cosmological GWB down to $\Omega_0 \simeq 3.9 \times 10^{-10}$ (no lines) and 5.4×10^{-10} (with lines).
- Sensitivity degradation 30%–60%; 60 Hz mains dominates, harmonics contribute.
- Lines persist across training, so the network treats them as stationary noise; losses weakly penalize them.
- Cross-correlation method can reach a 5σ sensitivity at $25\,\mathrm{Hz}$ for 108 days of $\Omega_{\mathrm{CBC}} \approx 1.8 \times 10^{-8}$ and $\Omega_0 \approx 2.4 \times 10^{-8}$. Assuming "$\log_{10} \mathrm{BF} = 3 \approx \mathrm{SNR} \sim 3$", the MSMHAutoencoder method results are a factor 2 better than cross-correlation. This is equivalent to a reduction by a factor 5 of the observing time required by the cross-correlation method

Discussion - outlook

- proof-of-concept promising results, but further deep learning exploration needed (architecture, loss functions, hyperparameter tuning, more realistic data/noise).
- Pre-notch spectral lines before analysis or add line-aware modules within the MSMH-Autoencoder (narrowband attention, adaptive notch layer).
- Strengthen spectral losses around flagged bins; augment with synthetic line perturbations.
- Next: apply to real O4a segments with non-stationarity and correlated artifacts.

Thank you for listening !
Questions ?

'Gravitational-wave background detection using machine learning'
https://doi.org/10.1103/hs9b-drwx

Training: Physics-Informed Loss Function

- The training objective is a weighted sum of three terms:

$$\mathcal{L}_{\text{total}} = \lambda_s \mathcal{L}_{\text{spectral}} + \lambda_l \mathcal{L}_{\text{latent}} + \lambda_c \mathcal{L}_{\text{consist}}$$

- **Spectral loss:** Combines signal and noise reconstruction accuracy with a spectral smoothness penalty:

$$\mathcal{L}_{\text{spectral}} = \frac{1}{B} \sum_{i=1}^{B} \left[\|\hat{n}_i - n_i\|_2^2 + w(k_i) \, \|\hat{s}_i - s_i\|_2^2 \right] + \|\nabla^2 \hat{s}_i\|_2^2$$

  Where:
  - $\hat{n}_i, \ n_i$ = predicted / true noise spectrum
  - $\hat{s}_i, \ s_i$ = predicted / true signal spectrum
  - $w(k_i)$ = reconstruction weight
  - $\|\nabla^2 \hat{s}_i\|_2^2$ = second-derivative penalty (enforces spectral smoothness)
- **Latent loss:** Ensures internal consistency between predicted and true noise in latent space:

$$\mathcal{L}_{\text{latent}} = \|\tilde{z}_{\text{noise}}(\hat{n}) - \tilde{z}_{\text{noise}}(n)\|_F^2$$

- **Consistency loss:** Imposes self-consistency between reconstructed and original input:

$$\mathcal{L}_{\text{consist}} = \|\tilde{z}_{\text{total}}(\hat{x}) - \tilde{z}_{\text{total}}(x)\|_F^2$$

- **Loss weights:** $\lambda_s, \lambda_l, \lambda_c$

Training : Curriculum Learning for Low SNR

- **Motivation:** Realistic GWBs have amplitudes $\ll$ detector noise $\rightarrow$ hard to learn separation at low SNR.
- **Solution:** Curriculum learning : start with easily separable, high-amplitude signals and progressively reduce amplitude to reach realistic SNR.
- **Schedule definition:**
  - Epoch progression variable:

$$\tau(\mathcal{E}) = \min\left(1, \frac{\mathcal{E} - 1}{\max(1, E_{\mathrm{curr}} - 1)}\right)$$

    where $\mathcal{E}$ = epoch, E_{curr} = total curriculum epochs.
  - Log-amplitude mean evolves as:

$$\mu_{\ln}(\mathcal{E}) = (1 - \tau(\mathcal{E}))\, \mu_{\ln,\mathrm{start}}(\mathcal{E}) + \tau(\mathcal{E})\, \mu_{\ln,\mathrm{end}}(\mathcal{E})$$

- **At each epoch E:**
  - Draw $\ln A \sim \mathcal{N}(\mu_{\ln}(\mathcal{E}), \sigma_{\ln}(E)^2)$, clip to $[A_{\min}(\mathcal{E}), A_{\max}]$
  - Use A to scale injected signal amplitude in training samples.
- **Impact:** Smooth transition from easy (high SNR) to hard (realistic SNR), avoiding model collapse at low amplitudes and boosting final sensitivity.

O4a PSD de-lining

Inputs. O4a ASD $S_h(f)$; PSD $P(f) = S_h(f)^2$.

Log space. $x(f) = \log_{10} S_h(f)^2$.

Baseline estimate. Median filter, window $K_m = 51$ bins:

$$x_{\mathrm{med}}(f) = \mathrm{median}\{x(f') \mid |f' - f| \leq K_m/2\}$$

then Savitzky–Golay filter of order $p = 5$, length $K_s = 51$ bins:

$$\widehat{x}(f) = \sum_{j=-K_s/2}^{K_s/2} c_j \, x_{\mathrm{med}}(f + j\Delta f), \qquad \widehat{P}(f) = 10^{\,\widehat{x}(f)}$$

Peak selection. Residual $r(f) = x(f) - \widehat{x}(f)$; mark peaks where $r(f) > \tau$ with $\tau = 0.25$ dex ($\sim 80\%$ amplitude excess).

Inpainting around each peak f_k. Define $\Delta N = 2$ bins and interval $[f_{k-\Delta N}, f_{k+\Delta N}]$. Replace by linear interpolation of the baseline:

$$P_{\mathrm{clean}}(f) = \begin{cases} \dfrac{f_{k+\Delta N} - f}{f_{k+\Delta N} - f_{k-\Delta N}} \widehat{P}(f_{k-\Delta N}) + \dfrac{f - f_{k-\Delta N}}{f_{k+\Delta N} - f_{k-\Delta N}} \widehat{P}(f_{k+\Delta N}), & f \in [f_{k-\Delta N}, f_{k+\Delta N}], \\ P(f), & \text{otherwise.} \end{cases}$$

Output. $S_h^{\mathrm{clean}}(f) = \sqrt{P_{\mathrm{clean}}(f)}$. Preserves broadband noise, removes narrow lines (calibration, violin modes, mains harmonics).