FLAMO
An Open-Source Library for Frequency-Domain Differentiable Audio Processing

Gloria Dal Santo[1], **Gian Marco De Bortoli**[1], **Karolina Prawda**[2], **Sebastian J. Schlecht**[3], **and Vesa Välimäki**[1]

[1] Acoustics Lab, Aalto University, Espoo, Finland
[2] AudioLab, University of York, York, United Kingdom
[3] Multimedia Communications and Signal Processing, FAU, Erlangen, Germany

50th Int. Conf. on Acoustics, Speech, and Signal Processing Hyderabad, India, 6-11 April 2025

github.com/gdalsanto/flamo

Abstract

We present **FLAMO**, a **F**requency-sampling **L**ibrary for **A**udio-**M**odule **O**ptimization, designed for implementing and optimizing differentiable linear time-invariant audio systems. The library is open-source and developed with the PyTorch framework for automatic differentiation. **FLAMO** includes a variety of customizable differentiable digital signal processors (DDSPs), which can be arranged in parallel, series, or recursive configurations. These systems can be optimized to match a target response or be integrated into the neural network computation graph.

Main Features

- Frequency-sampling optimization [1] with built-in time-aliasing mitigation method
- Customizable modular DDSP classes
  - flexible parameterization via callable mappings
  - system wrappers that ensure safe chaining
- Training utilities with integrated logging
- Loss analysis utilities, including loss surface exploration
- Constantly expanding with new modules and features

Frequency-domain optimization

- FIR approximation using the frequency-sampling method
- DDSPs are sampled at M linearly-spaced frequencies in $[0, \pi]$
- The system transfer function computed by multiplication

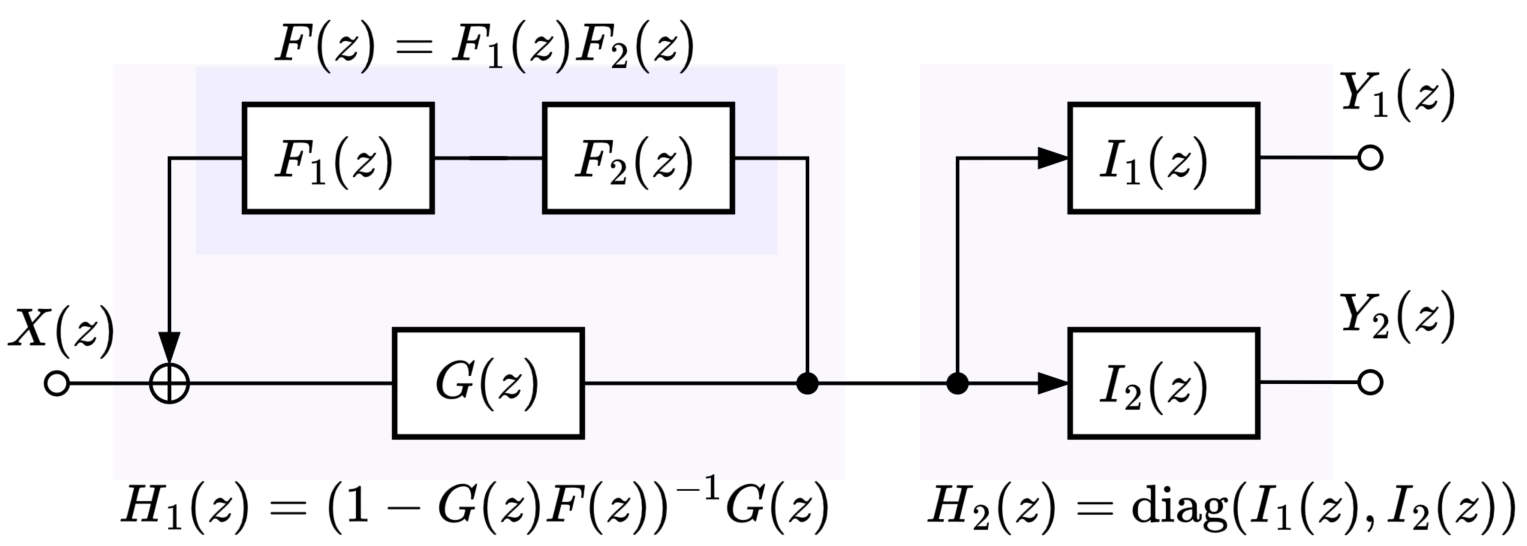

Figure 1. *Examples of module chaining configurations and their corresponding transfer function equations*

Time-aliasing mitigation

Frequency sampling is affected by time-aliasing

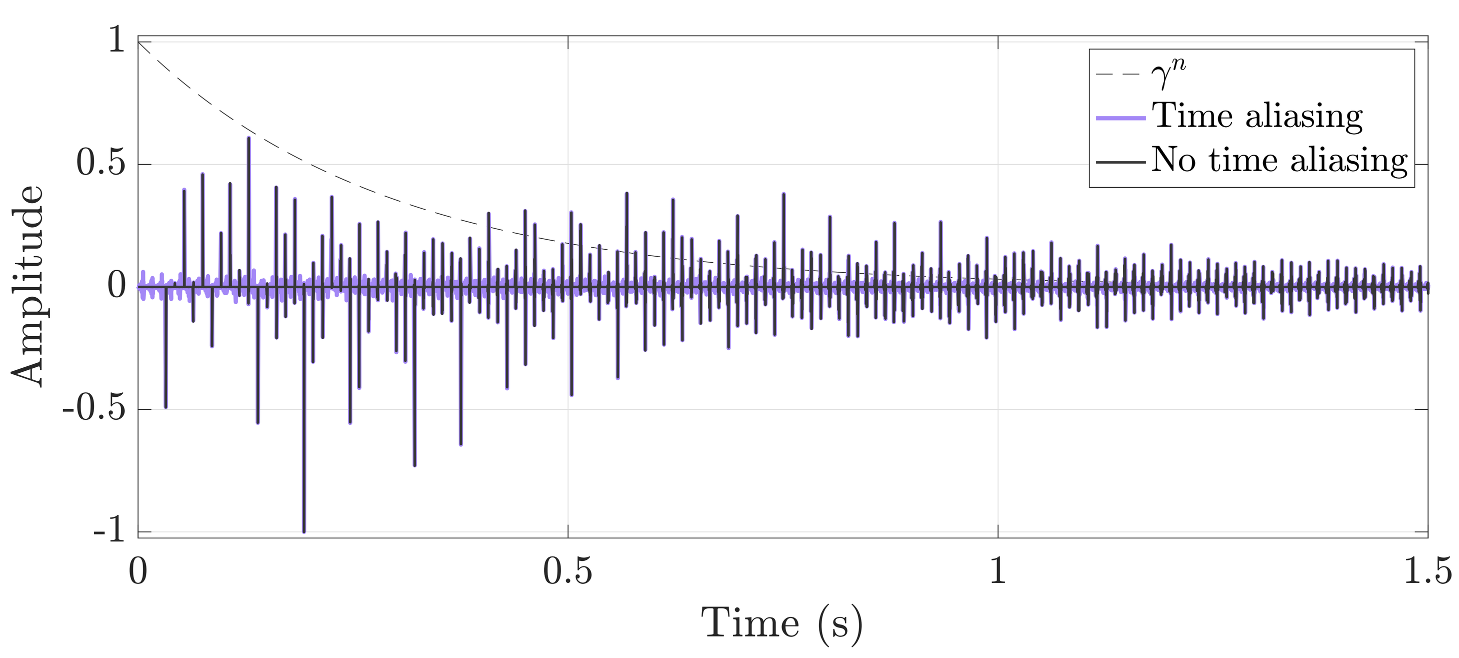

Figure 2. *IR of an artificial reverb decaying at a reverberation time of* $T_{60} = 9$ *s, with (purple) and without (black) time aliasing*

In FLAMO, aliasing is mitigated by applying an exponentially decaying envelope γ^n, for $0 < \gamma \leq 1$, or by sampling the frequency response outside the unit circle, at a radius of $1/\gamma$:

$$\hat{H}(e^{j\omega}) = H(e^{j\omega}/\gamma) \rightarrow \hat{h}[n] = h[n]\gamma^n$$

Library structure

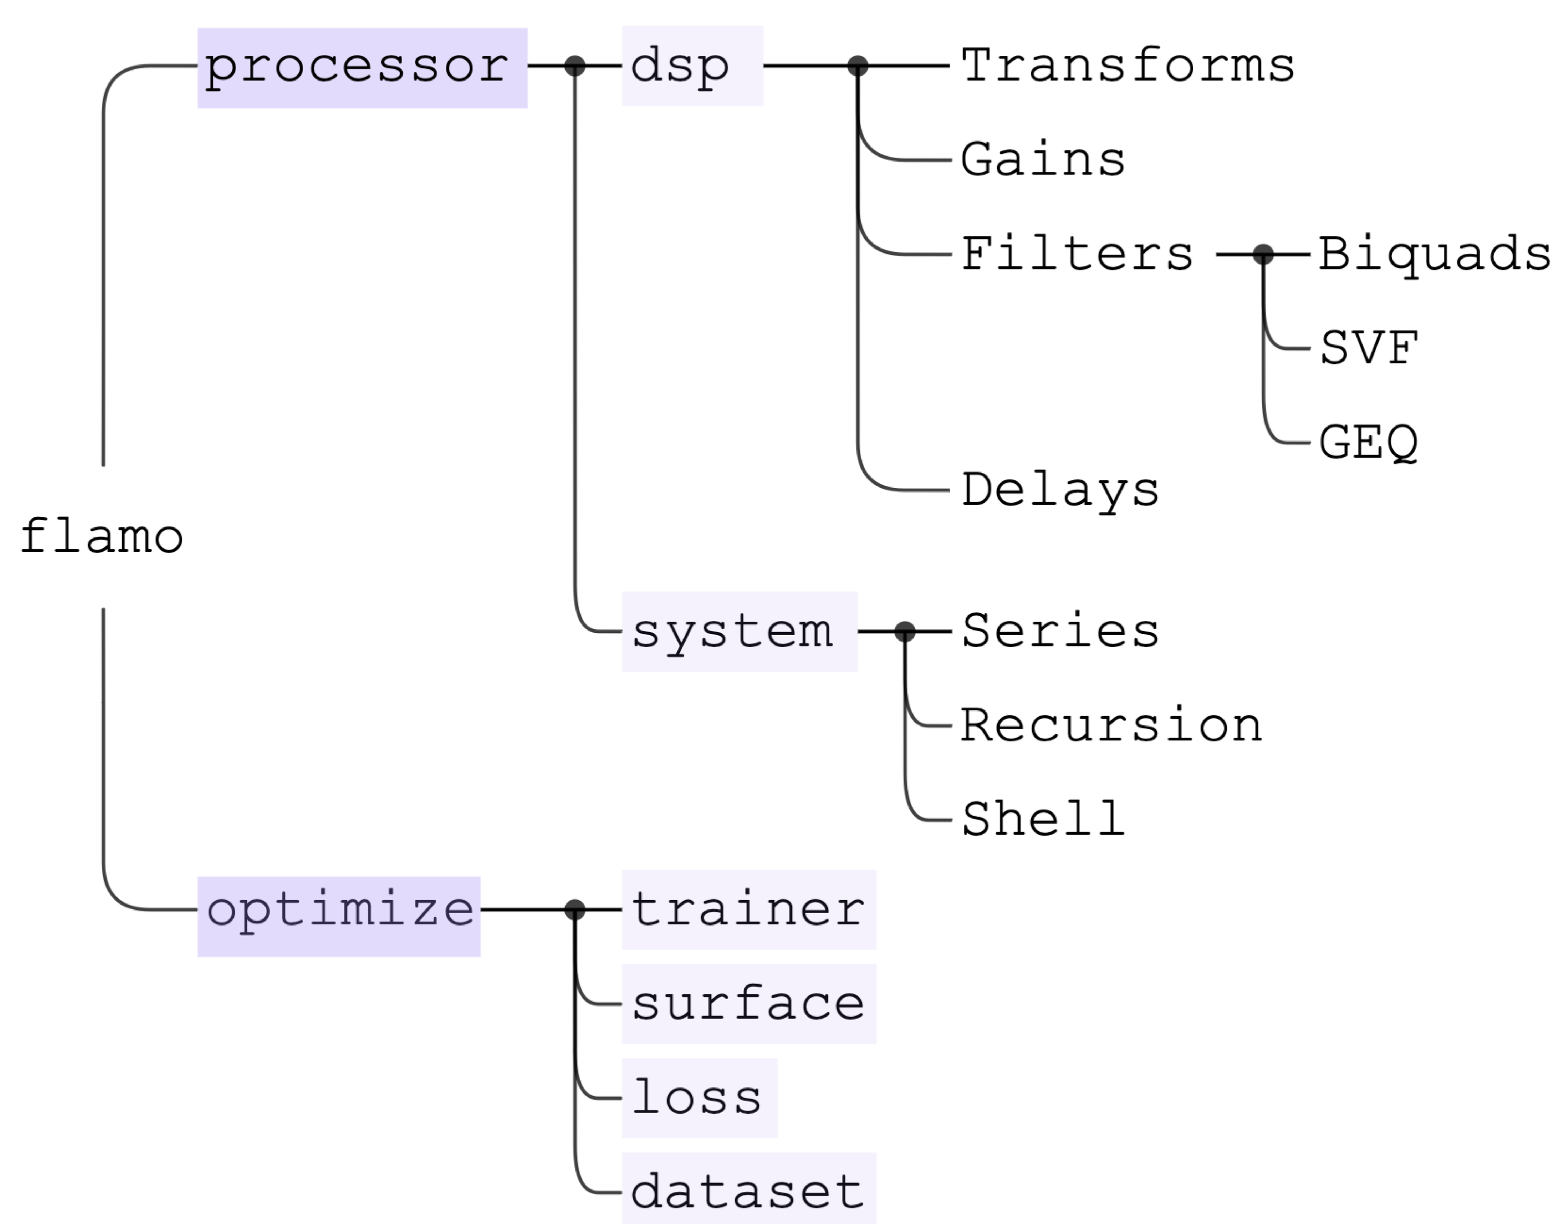

Figure 3. *Overview of the library structure and its modules*

Parameter hyperconditioning

Integration of DDSP modules with neural networks for parameter estimation [2]

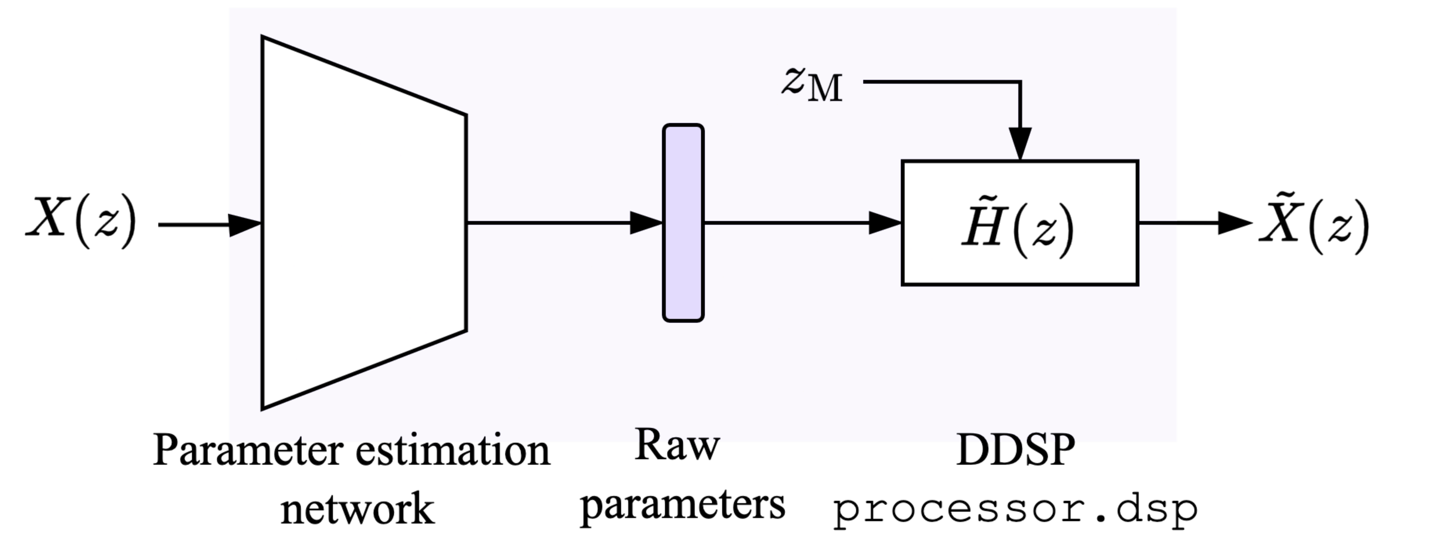

Figure 3. *Example of parameter hyperconditioning via an estimation network. The vector* z_M *indicates the sampling frequencies at which the system* $\tilde{H}(z)$ *is computed*

Application in Artificial Reverb

Optimization of a feedback delay network (FDN) for natural-sounding artificial reverberation [3]

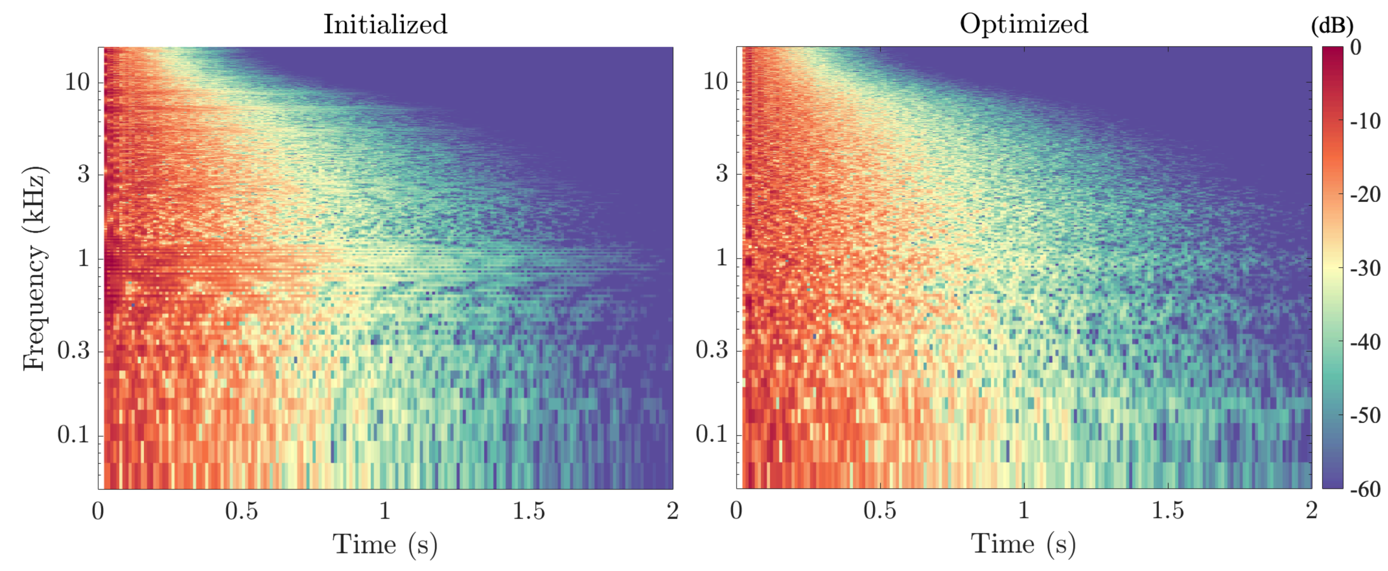

Figure 4. Effect of the optimization of an FDN of size $N = 6$ with gain coefficient optimized to improve smoothness

Resources

The library is open-source and available on GitHub, together with its documentation and a set of introductory examples. The library is also available on the Python package index (PyPI).

References

[1] S. Nercessian, A. Sarroff, and K. J. Werner, "Lightweight and interpretable neural modeling of an audio distortion effect using hyperconditioned differentiable biquads," in *IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, pp. 890–894, IEEE, 2021.

[2] S. Lee, H.-S. Choi, and K. Lee, "Differentiable artificial reverberation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 2541–2556, 2022.

[3] G. Dal Santo, K. Prawda, S. J. Schlecht, and V. Välimäki, "Differentiable feedback delay network for colorless reverberation," in *Proc. DAFx*, (Copenhagen, Denmark), pp. 244–251, Sep. 2023.