The Neuralacoustics Project: Exploring Deep-Learning for Lightweight Numerical Modeling Synthesis

Victor Zappi Kıvanç Tatar

Numerical modeling synthesis is one of the most transformative technologies in the context of embedded digital musical instrument (DMI) design. By means of simulating acoustic phenomena like excitation, wave propagation and resonance, numerical models are capable of synthesizing the sounds radiated from a wide variety of objects, whose shapes, sizes and materials are fully parametrized. The use of embedded technologies permits to run such simulations in low-latency environments, where multiple sensors and actuators may be coupled with acoustic parameters and different oscillatory outputs. This allows for the design of sophisticated control paradigms that leverage feedback loops between synthesis and control, and closely resemble those found in acoustic instruments.

Unfortunately, numerical modeling is characterized by a major drawback, consisting of the high computational requirements of the underlying simulations [1]. The implemented algorithms have to solve partial differential equations across the full simulation domain at audio rate, leading to performance peaks in the order of Giga FLOPS [2]. This is critical in the context of embedded systems, whose computational power is way more limited compared to both laptops' CPUs and graphics cards—the latter being very often used to accelerate the solution of numerical models [1, 2]. As a result, today's embedded DMIs can include only relatively simple linear models, mostly one-dimensional and constrained in terms of either the size or the resolution of the simulated domain. Some examples are the Daïs [3] and the Half-Physler [4], as well as the Pink Trombone project that can be found in Bela's example library [5]. Although innovative and characterized by neat designs, these instruments are only scratching the surface of what is possible to achieve with embedded numerical modeling synthesis.

In this talk, we present a newly started project called *Neuralacoustics*, aimed at using embedded AI to reduce the computational requirements of numerical modeling synthesis. To achieve this goal, we propose a novel design pipeline, in which neural networks are first trained off-line to approximate the behavior of target interactive acoustic systems, and then deployed on embedded and mobile platforms as real-time synthesis engines. The talk is articulated around the description of the three main components of the project (all in development). The first one is a collection of numerical models, written in Python and specifically designed for the rapid generation of large datasets, that exemplify the main acoustic behavior of sound emitting objects/instruments. The second one consists of a PyTorch platform for the training and the testing of neural network architectures; the platform is designed to facilitate the exploration of different network structures and hyperparameter configurations, with the goal of identifying the best fit for each acoustic target. The third component consists of a C++ framework for real-time inference on embedded hardware, based on LibTorch. The code is hosted on a GitHub repository that will be made publicly available in the next months.

The talk includes a practical demonstration of the first steps of the proposed pipeline. The first numerical models included in the repository will be analyzed and used for the generation of sample datasets. Then, the training procedure will be showcased and preliminary results obtained with a state-of-the-art deep neural network for the solution of partial differential equations [6] will be discussed, covering both accuracy and computational load.

References

[1] Zappi, Victor, Andrew Allen, and Sidney S. Fels. "Shader-based physical modelling for the design of massive digital musical instruments." *Proceedings of NIME*. 2017.

[2] Renney, Harri, Benedict R. Gaster, and Tom Mitchell. "OpenCL vs: Accelerated finite-difference digital synthesis." *Proceedings of the International Workshop on OpenCL*. 2019.

[3] Christensen, Pelle J., Dan Overholt, and Stefania Serafin. "The Daïs: A Haptically Enabled NIME for Controlling Physical Modeling Sound Synthesis Algorithms." *Proceedings of NIME. 2020.*

[4] Hofmann, Alex, et al. "The Half-Physler: An oscillating real-time interface to a tube resonator model." *Proceedings of NIME. 2019.*

[5] Bela Pink Trombone example project - https://github.com/BelaPlatform/Bela/tree/master/examples/Salt/pink-trombone

[6] Li, Zongyi, et al. "Fourier neural operator for parametric partial differential equations." *Proceedings of the Ninth International Conference on Learning Representations. 2021.*