Toward Scalable and Physically-Grounded Learning Frameworks: Neural Finite Volume Methods for Solving Hyperbolic PDEs

Hyperbolic partial differential equations (PDEs), especially scalar conservation laws, are fundamental tools for modeling wave propagation and transport phenomena in physics and engineering. Unlike parabolic or elliptic PDEs, hyperbolic PDEs often develop discontinuities such as shocks, which makes their numerical solution particularly challenging. While machine learning methods have recently shown success in solving parabolic and elliptic PDEs, their extension to hyperbolic equations remains less developed. This motivates our work: to design learning-based solvers that retain the stability and conservation guarantees of classical numerical methods, while offering the flexibility and efficiency of modern neural networks.

To address this gap, we propose a Neural Finite Volume (NFV) framework that integrates classical numerical methods with trainable neural networks. In NFV, the traditional numerical flux function in finite volume schemes is replaced by a learned neural flux. This design preserves key properties of conservation and stability while introducing data-driven flexibility. We develop both supervised and unsupervised training schemes: supervised learning directly from PDE entropy solution data, and unsupervised learning via weak-form residual loss.

We conduct an extensive numerical study across 7 hyperbolic PDE variants, covering scalar conservation laws with different flux functions. Our approach is benchmarked against 6 classical numerical schemes, including both finite volume and finite element methods. NFV consistently outperforms baselines in terms of accuracy and adaptability, while remaining computationally efficient. It achieves up to 10x lower error than Godunov's method, outperforms ENO/WENO, and rivals discontinuous Galerkin solvers with far less complexity. Beyond synthetic PDE data, we validate the method on real-world traffic datasets, where NFV demonstrates strong predictive performance despite noise and partial conservation. These results highlight the potential of combining physics-preserving architectures with the expressive power of neural networks.

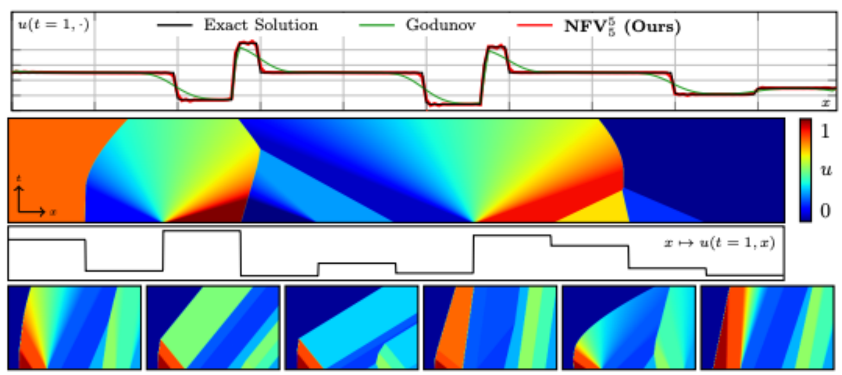

Figure 1: **Prediction of entropy solutions of hyperbolic PDEs. Top:** NFV_5^5 prediction vs. the Godunov scheme for Burgers' equation at a fixed time. **Mid:** Entropic solution $u(t,x)$ for Burgers' equation over domain $(t,x) \in [0,1]^2$, and its corresponding initial condition. **Bottom:** Entropic solution $u(t,x)$ for the LWR equation with different fluxes sharing the same initial condition.

Looking forward, our ongoing work extends this framework along two key directions: 1) Velocity-based PDEs (V-PDEs). In many applied domains such as traffic, velocity is far easier to measure than density or flow. Reformulating conservation laws into V-PDEs makes the learning problem more directly aligned with observable data, increasing real-world applicability. 2) Higher-dimensional PDEs. Realistic systems often evolve over space and time in multiple dimensions. Extending NFV to two- and three-dimensional hyperbolic PDEs requires new architectures and stability guarantees, but offers the opportunity to scale to richer scientific domains.

This research contributes to the growing effort to merge machine learning with scientific computing, emphasizing models that are both physically principled and data-informed. By focusing on the most challenging PDE class, hyperbolic PDEs, we aim to broaden the scope of learning-based solvers, with applications ranging from traffic flow modeling to fluid dynamics and beyond.