

Neural Implicit Reduced Fluid Simulation

Yuanyuan Tao¹, Ivan Puhachov², Derek Nowrouzezahrai¹, and Paul Kry¹

¹ McGill University

² Université de Montréal

Abstract

High-fidelity simulation of fluid dynamics is challenging because of the high dimensional state data needed to capture fine details and the large computational cost associated with advancing the system in time. We present neural implicit reduced fluid simulation (NIRFS), a reduced fluid simulation technique that combines a neural-implicit representation of fluid shapes and a neural ordinary differential equation to model the dynamics of fluid in the reduced latent space. Trajectories for NIRFS can be computed at very little cost in comparison to simulations for generating training data, while preserving many of the fine details. We show that this approach can work well, capturing the shapes and dynamics involved in a variety of scenarios with constrained initial conditions, e.g., droplet-droplet collisions, crown splashes, and fluid slosh in a container. In each scenario, we learn the latent implicit representation of fluid shapes with a deep-network signed distance function, as well as the energy function and parameters of a damped Hamiltonian system, which helps guarantee desirable properties of the latent dynamics. To ensure that latent shape representations form smooth and physically meaningful trajectories, we simultaneously learn the latent representation and dynamics. We evaluate novel simulations for conservation of volume and momentum conservation, discuss design decisions, and demonstrate an application of our method to fluid control.

CCS Concepts

• **Computing methodologies** → Modeling and simulation; Computer graphics; Artificial intelligence;




Figure 1: Neural implicit reduced fluid simulation (NIRFS) performs high-fidelity fluid simulation orders of magnitudes faster than conventional full-order methods. An initial condition encoder (ICE) maps a novel physical initial condition, e.g., droplet position x and velocity v , to a latent initial condition $(\mathbf{q}_0, \mathbf{p}_0)$ (a). In a physically structured latent space, a damped Hamiltonian neural ODE (DHNODE) is integrated (b), generating a latent fluid trajectory (c). From the latent geometry vectors \mathbf{q} in the trajectory, an implicit neural representation (INR) decoder constructs fluid geometries with fine details (a and d) and high-frequency structures (e top) through. The correct steady state is stably reached in the end (e bottom).

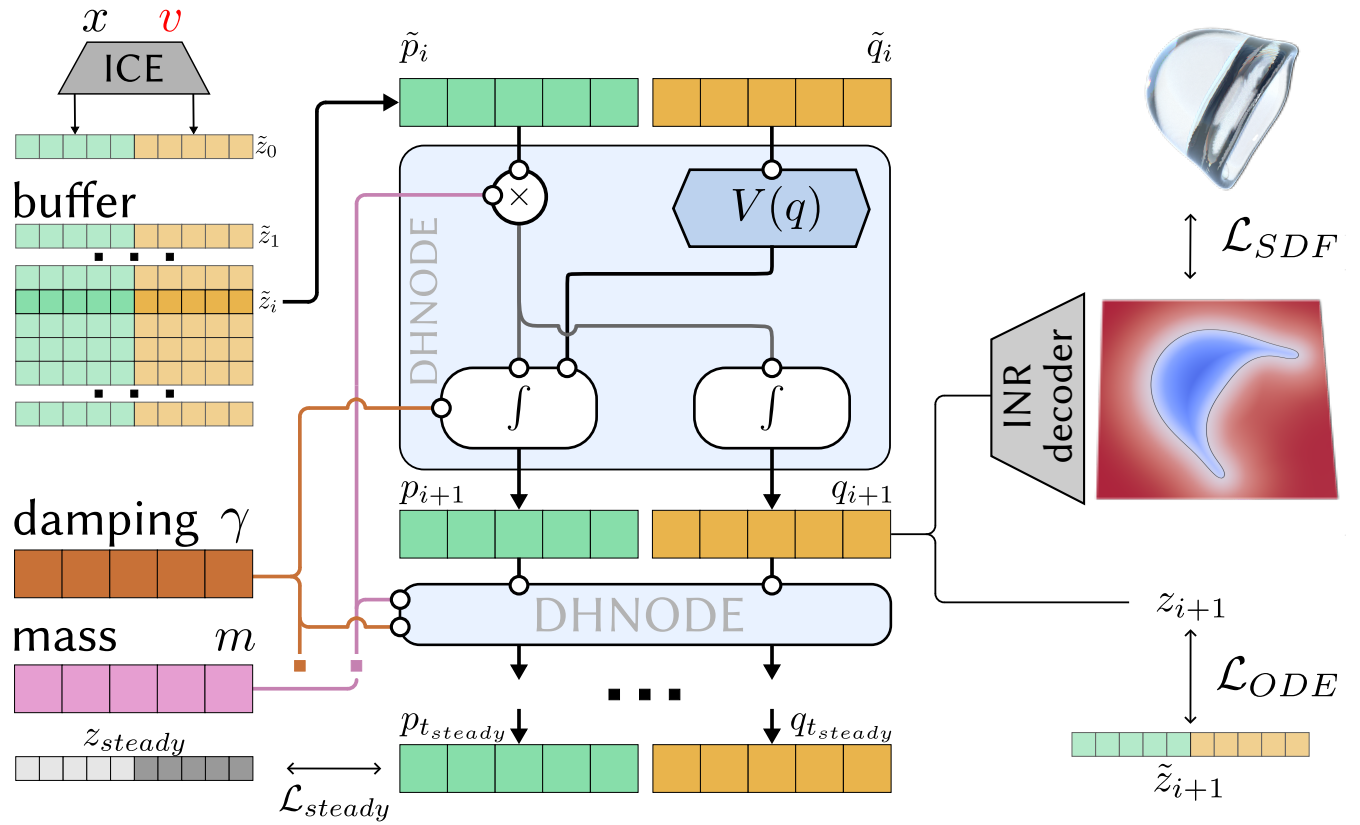


Figure 2: An ICE, a DHNODE, and an INR decoder are jointly trained. To achieve training stability and convergence, we keep latent states for each frame in training simulations in a buffer \tilde{z} . The ICE is trained to map physical initial conditions of each full length simulation to latent initial conditions, which replace the first states of each simulation in the buffer. Each DHNODE learns an energy function, a mass matrix, and a damping matrix for a fluid model. The DHNODE is integrated from uniformly sampled latent states until (after) t_{steady} . For the short initial segment of each integrated trajectory, \mathcal{L}_{ODE} measures the difference between the rollouts and states in the buffer, and \mathcal{L}_{SDF} measures the difference between the decoded and ground truth SDF geometries. For the states rolled out at or beyond t_{steady} , \mathcal{L}_{steady} measures the difference between the rolled out states and the determined latent steady states.