Neural SPH: Improved Neural Modeling of Lagrangian Fluid Dynamics

AP Toshev[1], JA Erbesdobler[1], NA Adams[1,2], J Brandstetter[3,4]

[1] Technical University of Munich, [2] MEP, [3] JKU Linz, [4] NXAI GmbH

TL;DR

Neural SPH is a framework for improved training and inference of graph neural network (GNN)-based simulators for Lagrangian fluid dynamics. Based on the insight that **particle clumping** is one of the main reasons for the failure of learned Lagrangian solvers on long rollouts, we enhance such state-of-the-art GNN-based simulators with ideas from standard smoothed particle hydrodynamics (SPH) solvers. Our contributions are:

- Novel **external force treatment**: excluding external forces from the model target.
- **SPH relaxation** during inference: relaxing particles to a more physical configuration using adopted **pressure** and **viscous** terms from standard SPH.

Strengths

- Simple, interpretable, robust, and efficient approach.
- Does not require a differentiable solver.
- Applicable to all weakly compressible problems, incl. walls and free surfaces.

Smoothed Particle Hydrodynamics (SPH)

- Our approach targets systems governed by the weakly compressible Navier Stokes equations (NSE) with density ρ, velocity vector $\mathbf{u}$, pressure p, external force $\mathbf{g}$, and Reynolds number Re.

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{u}) = \underbrace{-\frac{1}{\rho}\nabla p}_{\text{pressure}} + \underbrace{\frac{1}{Re}\nabla^2 \mathbf{u}}_{\text{viscosity}} + \underbrace{\mathbf{g}}_{\text{ext. force}} \quad \text{(Mom.)}, \qquad p(\rho) = p_{ref}\left(\frac{\rho}{\rho_{ref}} - 1\right) \quad \text{(EoS)}$$

- The term responsible for a homogeneous particle/density distribution in SPH is the pressure gradient term in the momentum equation.

Neural SPH

External Force Treatment ($\square_g$)

- Split the terms on the right-hand side of the momentum equation (Mom.) into [...] + $\mathbf{g}$
- Remove the accumulated external force from the target acceleration, i.e.,

$$\mathbf{a}^{target} = \mathtt{GNN}(\mathbf{X}^{t_{k-H-1}:t_k}, \mathbf{g}) + \mathbf{g}_M^{FD} \,,$$

SPH Relaxation ($\square_p$ and $\square_\nu$)

- Apply SPH relaxation after each learned solver step to improve the particle distribution.
- The SPH relaxation has access only to the particle coordinates and no physical quantities like density and velocity. For the viscous term, we use the effective velocity from the difference in coordinates.
- One update step of the relaxation corresponds to

$$\mathbf{a} = \alpha \frac{-1}{\rho}\nabla p + \alpha\beta\nabla^2 \mathbf{u} \,, \qquad \mathbf{p} = \mathbf{p} + \mathbf{a} \,,$$

  where we hide the time step and the pre-factors in the hyperparameters α and β.
- According to SPH theory, density fluctuations should not exceed $\sim 1\%$. We use density summation and set all $\rho < 0.98\rho_{ref}$ to ρ_{ref}, and all $\rho > 1.02\rho_{ref}$ to $1.02\rho_{ref}$.
- Our relaxation implementation is based on the JAX-SPH code [Toshev et al., 2024b].

SPH Relaxation Parameter Tuning Recipe

We propose a three-step parameter-tuning process while monitoring the **position MSE, Sinkhorn divergence, kinetic energy MSE, MAE of density deviation from the reference ρ_{ref}, Dirichlet energy of the density field, and Chamfer distance**:

1. Tune α while number of relaxation steps $l = 1$ and $\beta = 0$. Typically, $\alpha \in (0.005, 0.05)$.
2. Tune l with optimal α and $\beta = 0$. Typically, $l \in (1, 5)$.
3. Tune β with optimal α and l. Typically, $\beta \in (0.1, 1)$.

Results on LagrangeBench Datasets

We demonstrate the efficacy of our approach on all datasets accompanying the LagrangeBench benchmarking suite [Toshev et al., 2024a].

Dam Break

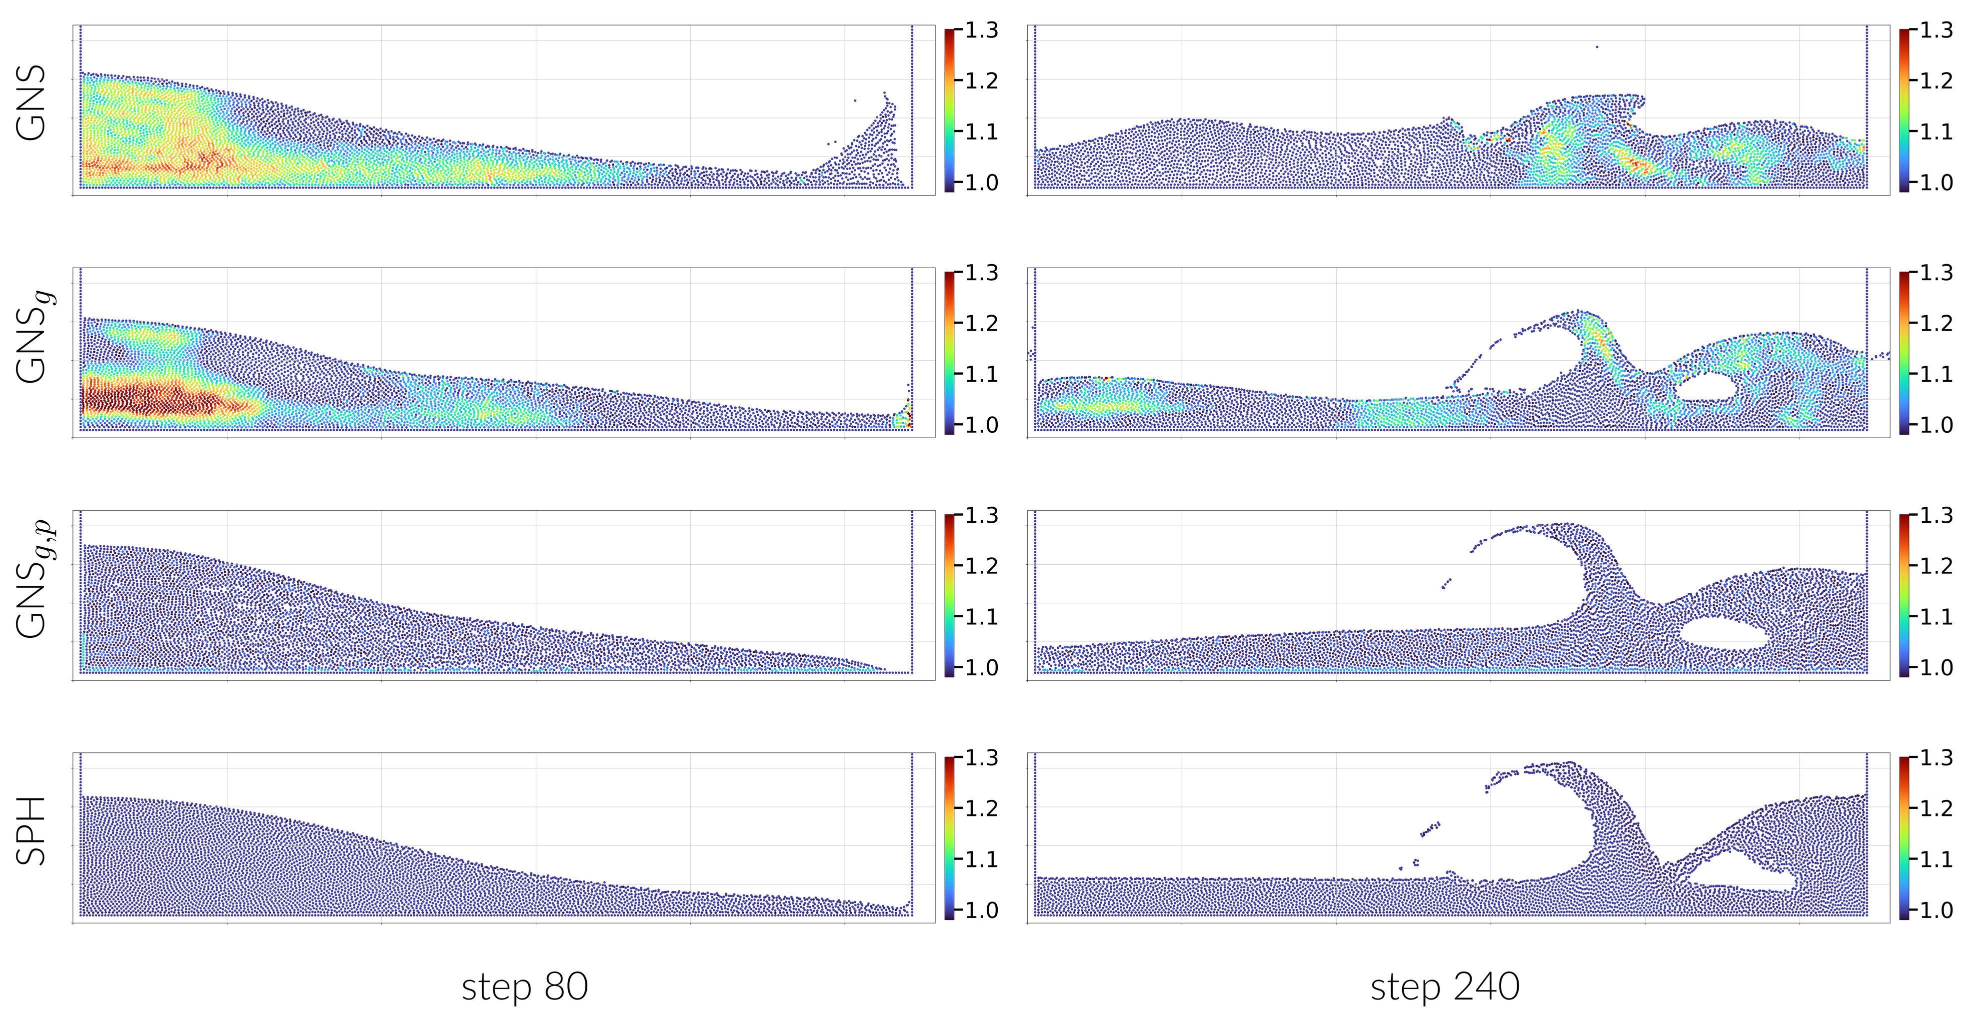

Neural SPH improves Lagrangian fluid dynamics, showcased by physics modeling of the 2D dam break example after 80 (left) and 240 (right) rollout steps. Different models exhibit different physics behaviors. From top to bottom: GNS [Sanchez-Gonzalez et al., 2020], GNS with corrected force only (GNS_g), full SPH enhanced GNS ($GNS_{g,p}$), and the ground truth SPH simulation. The colors correspond to the density deviation from the reference density; the system is considered physical within 0.98-1.02.

Lid-Driven Cavity

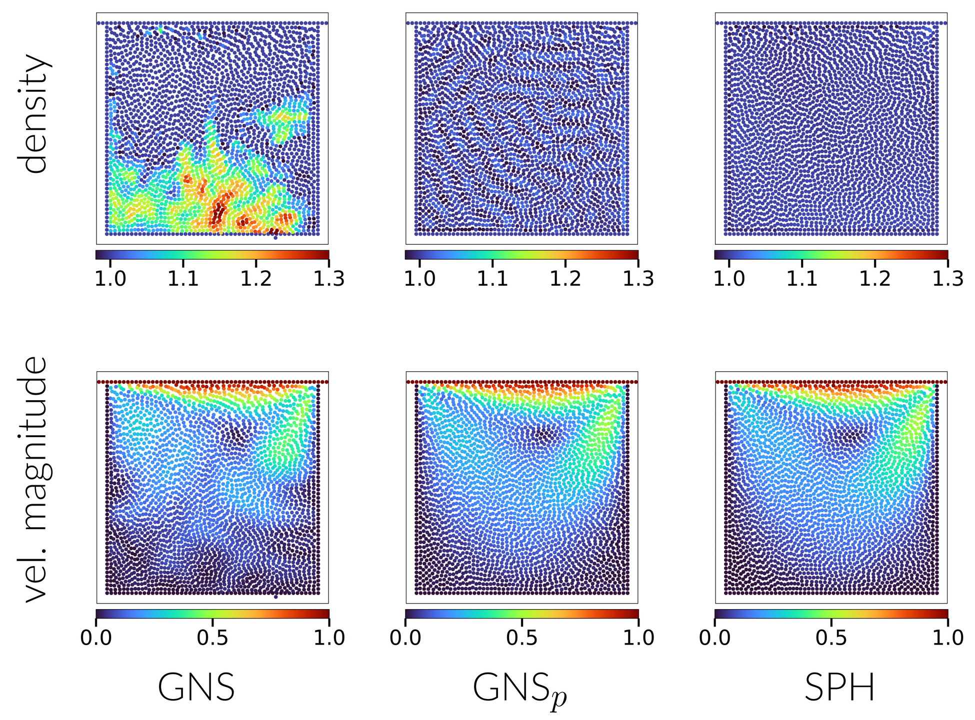

Density and velocity magnitude of 2D lid-driven cavity after 400 rollout steps (left to right): GNS, GNS_p, SPH.

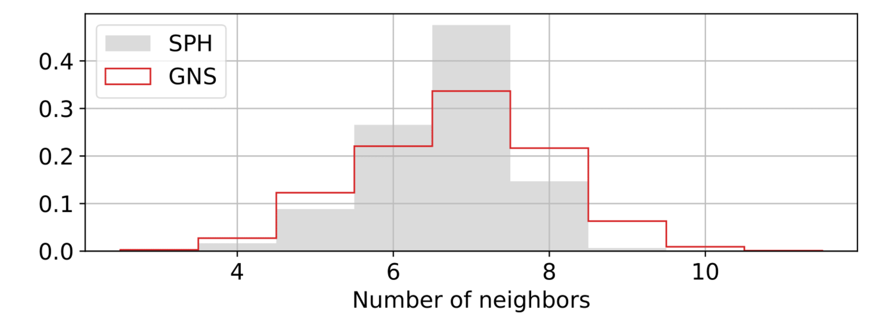

Histogram of the number of neighbors of the 2D lid-driven cavity experiment after 400 rollout steps.

Reverse Poiseuille Flow

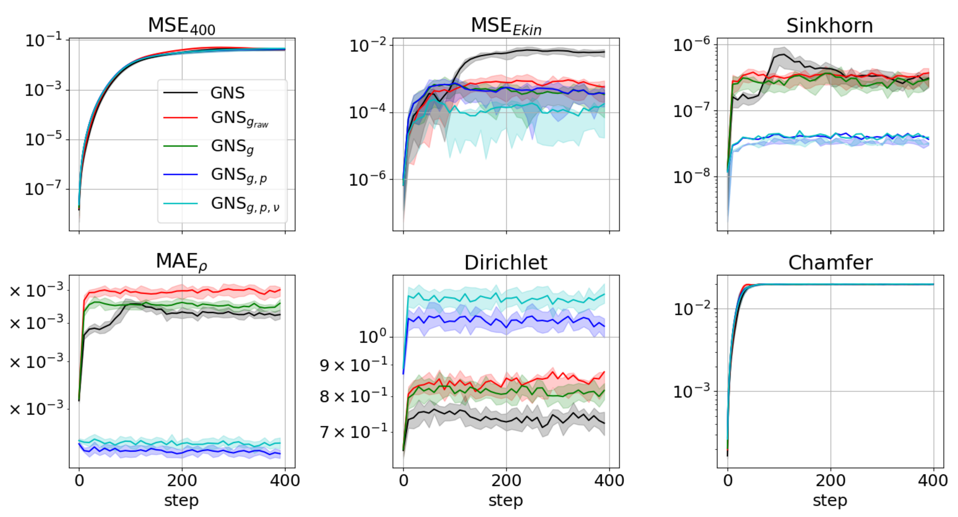

Ablations on RPF 2D with GNS-10-128 over the simulation length.

Forcing step function of the 2D reverse Poiseuille flow before (blue) and after convolution with normal distribution $\mathcal{N}(0, 0.025^2)$ (orange).

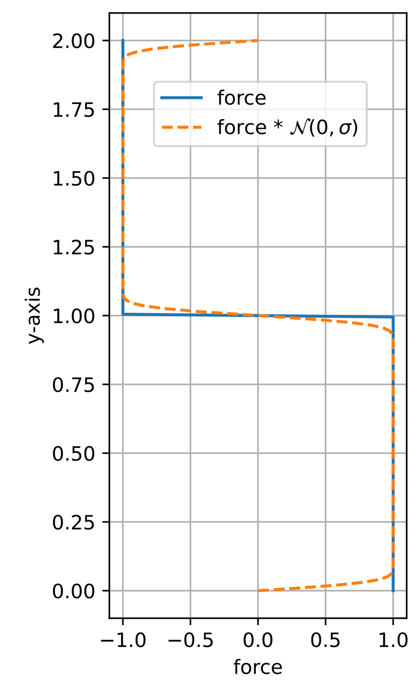

Results in Numbers GNS & SEGNN

	Model	MSE_{400}	Sinkhorn	MSE_{Ekin}
2D TGV	GNS	$5.3e-4$	$5.4e-7$	$5.6e-7$
	GNS_p	$4.8e-4$	$1.7e-8$	$4.8e-7$
2D RPF	GNS	$2.7e-2$	$3.6e-7$	$4.3e-3$
	GNS_g	$2.7e-2$	$2.7e-7$	$3.7e-4$
	$GNS_{g,p}$	$2.7e-2$	$2.9e-8$	$4.1e-4$
	$GNS_{g,p,\nu}$	$2.7e-2$	$3.0e-8$	$1.4e-4$
2D LDC	GNS	$3.3e-2$	$3.1e-4$	$1.1e-4$
	GNS_p	$1.6e-2$	$2.8e-7$	$1.2e-6$
2D DAM	GNS	$1.9e-1$	$3.8e-2$	$4.6e-2$
	GNS_g	$8.0e-2$	$1.3e-2$	$9.4e-3$
	GNS_p	$9.7e-2$	$7.1e-3$	$5.8e-3$
	$GNS_{g,p}$	$8.4e-2$	$7.5e-3$	$2.1e-3$
3D TGV	GNS	$4.8e-2$	$4.1e-6$	$3.6e-2$
	GNS_p	$4.6e-2$	$9.0e-7$	$4.2e-2$
3D RPF	GNS	$2.3e-2$	$4.4e-7$	$1.7e-5$
	GNS_g	$2.3e-2$	$4.4e-7$	$4.1e-5$
	GNS_p	$2.3e-2$	$1.0e-7$	$1.5e-5$
	$GNS_{g,p}$	$2.3e-2$	$1.3e-7$	$4.1e-5$
3D LDC	GNS	$3.2e-2$	$2.0e-5$	$1.3e-7$
	GNS_p	$3.2e-2$	$1.1e-6$	$2.9e-8$

	Model	MSE_{400}	Sinkhorn	MSE_{Ekin}
2D TGV	SEGNN	$4.0e-4$	$4.4e-7$	$3.9e-7$
	$SEGNN_p$	$3.8e-4$	$1.5e-8$	$2.8e-7$
2D RPF	SEGNN	$2.7e-2$	$3.3e-7$	$4.3e-3$
	$SEGNN_g$	$2.8e-2$	$3.3e-7$	$1.2e-4$
	$SEGNN_{g,p}$	$2.8e-2$	$3.5e-8$	$1.6e-4$
	$SEGNN_{g,p,\nu}$	$2.8e-2$	$3.8e-8$	$7.3e-4$
2D LDC	SEGNN	$7.6e-2$	$2.3e-3$	$9.1e+0$
	$SEGNN_p$	$1.8e-2$	$5.8e-7$	$1.6e-5$
2D DAM	SEGNN	$1.5e-1$	$3.4e-2$	$1.9e-2$
	$SEGNN_g$	$1.6e-1$	$2.1e-2$	$1.9e+1$
	$SEGNN_p$	$1.2e-1$	$9.4e-3$	$1.2e-2$
	$SEGNN_{g,p}$	$8.6e-2$	$4.9e-3$	$2.6e-3$
3D TGV	SEGNN	$4.2e-2$	$6.1e-6$	$2.4e-2$
	$SEGNN_p$	$4.1e-2$	$6.0e-7$	$2.7e-2$
3D RPF	SEGNN	$1.2e-1$	$1.0e-4$	$1.5e+3$
	$SEGNN_p$	$2.6e-2$	$1.3e-5$	$1.8e-2$
	$SEGNN_g$	$2.7e-2$	$2.6e-6$	$9.5e-3$
	$SEGNN_{g,p}$	$2.6e-2$	$7.9e-7$	$5.7e-3$
3D LDC	SEGNN	$3.3e-2$	$2.3e-5$	$1.7e-7$
	$SEGNN_p$	$3.3e-2$	$2.0e-6$	$1.8e-7$

Performance measures averaged over a rollout of 400-steps for GNS [Sanchez-Gonzalez et al., 2020] (left) and SEGNN [Brandstetter et al., 2022] (right). An additional subscript g indicates that external forces are removed from the model outputs, subscript p indicates that the SPH relaxation has a pressure term, and subscript ν that the viscous term is added to the SPH relaxation.

SPH relaxation hyperparameters used in our experiments. These hyperparameters were tuned on the GNS-10-128 model.

Dataset	loops	α	β
2D TGV	5	0.02	–
2D RPF	3	0.02	0.2
2D LDC	5	0.03	–
2D DAM	3	0.03	–
3D TGV	1	0.01	–
3D RPF	1	0.005	–
3D LDC	1	0.02	–

Outlook

Limitations

- Our external force treatment requires information on the time step and the temporal coarsening level.
- As proposed, the SPH relaxation is not directly applicable to compressible fluids.

Future Work

- Simplify parameter tuning recipe.
- Define universal thresholds to determine whether a simulation is physical.
- Explore combinations of learned solvers and other terms from classical numerics.

References

[Brandstetter et al., 2022] Brandstetter, J., Hesselink, R., van der Pol, E., Bekkers, E. J., and Welling, M. (2022). Geometric and physical quantities improve e(3) equivariant message passing. In *ICLR*.

[Sanchez-Gonzalez et al., 2020] Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. (2020). Learning to simulate complex physics with graph networks. In *ICML*, pages 8459–8468. PMLR.

[Toshev et al., 2024a] Toshev, A., Galletti, G., Fritz, F., Adami, S., and Adams, N. (2024a). Lagrangebench: A lagrangian fluid mechanics benchmarking suite. *Advances in Neural Information Processing Systems*, 36.

[Toshev et al., 2024b] Toshev, A. P., Ramachandran, H., Erbesdobler, J. A., Galletti, G., Brandstetter, J., and Adams, N. A. (2024b). Jax-sph: A differentiable smoothed particle hydrodynamics framework. *arXiv preprint arXiv:2403.04750*.

tumaer/neuralsph

2402.06275

@ArturToshev

artur.toshev@tum.de