

Data-Driven Viscosity Solver for Fluid Simulation

Wonjung Park^a, Hyunsoo Kim^b and Jinah Park^c
Korea Advanced Institute of Science and Technology, Daejeon, Korea

Keywords: Physical Simulation, Machine Learning, Viscous Fluid, Convolutional Neural Network.

Abstract: We propose a data-driven viscosity solver based on U-shaped convolutional neural network to predict velocity changes due to viscosity. Our solver takes velocity derivatives, fluid volume, and solid indicator quantities as input. The traditional marker-and-cell (MAC) grid stores velocities at the edges of the grid, causing the dimensions of the velocity field vary from axis to axis. In our work, we suggest a symmetric MAC grid that maintains consistent dimensions across axes without interpolation or symmetry breaking. The proposed grid effectively transfers spatial fluid quantities such as partial derivatives of velocity, enabling networks to generate accurate predictions. Additionally, we introduce a physics-based loss inspired by the variational formulation of viscosity to enhance the network's generalization for a wide range of viscosity coefficients. We demonstrate various fluid simulation results, including 2D and 3D fluid-rigid body scenes and a scene exhibiting the buckling effect.

1 INTRODUCTION

Numerical methods for fluid simulation have been widely used in various fields such as industrial design, and computer graphics. These methods solve incompressible Navier-Stokes (N-S) equations, the governing equation of fluid dynamics in the form of partial differential equations (PDEs). The process of solving N-S equations is well-known for being computationally expensive and the cost is mostly from solving a series of linear systems by iterative methods.

In recent years, various studies (Kim et al., 2019; Richter-Powell et al., 2022; Tompson et al., 2017) have adopted machine learning to accelerate the projection step of fluid simulation. By replacing iterative methods with highly parallel GPU-optimized neural networks, they were able to greatly accelerate the overall fluid simulation. In addition to the parallelism, these methods exploit the fact that pressure makes the fluid incompressible. Unlike typical machine learning regression, which solely relies on supervised learning from the dataset, they are designed to enforce the incompressibility of the velocity field. Therefore, they can be trained with a small dataset and generalize well to scenes outside of the training data.

Although the data-driven projection solver accel-

erates the simulation of inviscid fluid to the real-time level, we found that it is insignificant when used for a viscous fluid. Traditional viscosity solvers have to solve linear systems that are 4 times larger than the projection stage in 2D (x8 for 3D), since it deals with twice more fluid quantities per each dimension. However, studies on the data-driven viscosity solver have not yet been conducted in depth, since generalization of the viscosity solver is a challenging problem.

The challenge comes from two factors. First, it is difficult for networks to generalize over viscosity coefficients because it varies in a wide range from zero (inviscid) to infinity (solid). Second, the pressure is stored at the cell center whereas the fluid velocity is stored at the cell boundary of the marker-and-cell (MAC) grid for Eulerian simulation. This structure yields an asymmetry of the velocity field; for example, the velocity of the x-direction has the dimension $(n + 1, n)$ but the y-direction has $(n, n + 1)$ for the Eulerian MAC grid $n \times n$. Therefore, unlike general image data for CNNs, fluid velocity requires special care due to their distinct and asymmetric structure.

We design the viscosity solver architecture with these problems in mind. To generalize the solver for viscosity coefficients, we inherit the variational interpretation of the viscosity from the conventional solver (Batty and Bridson, 2008), using its minimization problem as the training loss. In addition, we introduce the *symmetric MAC grid* to enable the sym-

^a <https://orcid.org/0000-0003-1075-612X>

^b <https://orcid.org/0000-0002-0404-1892>

^c <https://orcid.org/0000-0003-4676-9862>

metric representation of fluids regardless of the axis.

To our knowledge, this is the first data-driven viscosity solver, so we focus on elaborating an architecture that generalizes to various scenarios. In this paper, we validate our architecture through ablation studies of neural network configurations. Finally, we demonstrate the universality of our solver in various scenarios, including unseen rigid-fluid scenes, mixed dynamic viscosity scenes, and 3D scene which shows buckling effect of viscous fluid.

2 RELATED WORK

The main subject of this study is to simulate viscous fluids governed by the incompressible Navier-Stokes (N-S) equations with our data-driven viscosity solver. The Navier-Stokes (N-S) equations are:

$$\frac{D\mathbf{v}}{Dt} = -\frac{1}{\rho}\nabla p + \frac{\mu}{\rho}\nabla^2\mathbf{v} + \mathbf{g}, \quad (1)$$

$$\nabla \cdot \mathbf{v} = 0. \quad (2)$$

There are three approaches to numerically simulate the N-S equations: grid-based (Eulerian), particle-based (Lagrangian), and Lagrangian/Eulerian hybrid approaches. Hybrid simulation uses particles to transport fluid and grids to update velocity by external forces, pressure, and viscosity. For spatial discretization of fluid quantities, marker-and-cell method (Harlow and Welch, 1965), which stores velocities on edges of cells, is often used in grid-based simulation.

The projection term in the N-S equations (the first term in the RHS of Equ. 1) enforces the incompressibility as formulated in Equ. 2. Combining the two equations gives Poisson's equation ($\nabla^2 p = f$), whose discrete form can be solved using linear system solvers.

The viscosity term in the N-S equations (second term in RHS of Equ. 1) gives the viscous behavior, accounting for internal friction. With explicit solvers, this term can become numerically unstable especially for high dynamic viscosity μ . Thus, a more stable implicit formulation of the viscosity fluid for hybrid methods is proposed in (Batty and Bridson, 2008). They interpreted the viscosity term as a minimization problem that seeks to minimize the rate of viscous dissipation with the smallest velocity change possible. In 2D, the minimization problem is defined as:

$$\min_{\vec{u}} \iint \rho \left\| \vec{u} - \vec{u}^{old} \right\|^2 + 2\Delta t \iint \mu \left\| \frac{\nabla \vec{u} + (\nabla \vec{u})^T}{2} \right\|_F^2 \quad (3)$$

To reduce the computational cost of fluid simulations, various dimensionality reduction methods have

been proposed. Model reduction for fluid-rigid simulation is discussed in (Treuille et al., 2006). In addition, multi-resolution Eulerian grids are introduced for efficient Poisson equation solving in both inviscid (Chentanez and Müller, 2011; McAdams et al., 2010) and viscous fluids (Aanjaneya et al., 2019; Shao et al., 2022). Furthermore, recent studies applied machine learning techniques including principal component analysis (Xiao et al., 2018b), regression forests (Ladický et al., 2015), Long Short-Term Memory (Wiewel et al., 2020), generative adversarial networks (Cheng et al., 2020), and graph networks (Sanchez-Gonzalez et al., 2020; Li and Farimani, 2022).

Among these machine learning techniques, convolutional neural networks (CNNs) have shown their effectiveness in Eulerian and hybrid fluid simulation. Since the Eulerian grid has a similar structure to the image, many studies applied CNN to fluid simulation (Guo et al., 2016; Obiols-Sales et al., 2020; Thurey et al., 2020; Xiao et al., 2019). Some studies (Kim et al., 2019) proposed generative networks that output the next velocity fields based on reduced parameters. Alternatively, other studies (Tompson et al., 2017; Xiao et al., 2018b; Xiao et al., 2018a; Yang et al., 2016; Gao et al., 2020) maintain the traditional fluid simulation process but replace the most time-consuming projection step with a data-driven solver. The traditional projection step with discretized PDE is computationally expensive and data dependent, since it is solved by iterative methods until it meets the conditions. Thus, the replacement of Poisson's equation with deep learning techniques (Richter-Powell et al., 2022; Tang et al., 2017) have been actively researched to efficiently derive the solution. Despite the efficiency and generalization of these data-driven projection solvers, their impact on viscous fluids is limited as the viscosity step remains a significant bottleneck.

In this study, we propose a CNN-based viscosity solver that replaces the PDE-based solver (Batty and Bridson, 2008) to effectively reduce the computational costs of the viscous fluid.

3 METHOD

3.1 Data-Driven Viscosity Solver

We suggest a data-driven viscosity solver that predicts velocity change by viscosity. As Fig. 1 shows, the overall process is based on the affine particle-in-cell (APIC) (Jiang et al., 2015) simulator, which transports fluid with particles and updates the velocity by external force, viscosity, and pressure on the grid. The

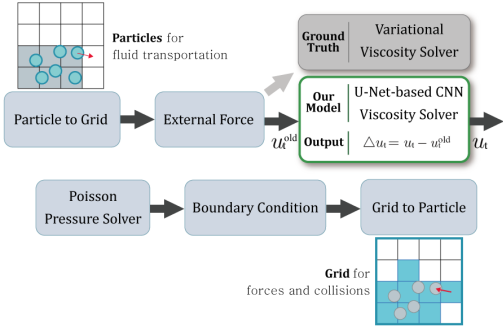


Figure 1: **Fluid simulation with data-driven viscosity solver.** Using the traditional affine particle-in-cell (Jiang et al., 2015) simulator as a baseline, our data-driven solver replaces the viscosity solver to output velocity changes due to viscosity. Training ground truth is generated with the traditional viscosity solver (Batty and Bridson, 2008).

viscosity solver updates the velocity u_t^{old} from the external force step (gravity) to the velocity u_t . We use the velocity difference $\Delta u_t = u_t - u_t^{old}$ solved by the traditional solver (Batty and Bridson, 2008) as the ground truth for training.

3.1.1 Loss Function

We train our model by supervised learning with L2 loss which is the mean squared error between ground truth and model output. In addition, we suggest the *variational loss* L_v which is the variational formulation of the conventional viscosity solver. When training the model only for fixed dynamic viscosity, L2 loss is sufficient to converge the model. However, for a dataset that contains various dynamic viscosity at the same time, we use both losses together to regularize the network for dynamic viscosity μ . Our loss function is composed of the sum of L2 loss and variational loss L_v , with equal weights. We trained the network with Adam optimizer (Kingma and Ba, 2014). The variational loss L_v on the 2D Eulerian grid is derived as follows, where velocity $\mathbf{u} = (u, v)$:

$$\begin{aligned}
 L_v(\mathbf{u}, \Delta t, \rho, \mu) &= \frac{1}{n(n+1)} \sum_{i=1}^{n+1} \sum_{j=1}^n \rho (u_{i-\frac{1}{2},j} - u_{i-\frac{1}{2},j}^{old})^2 + \frac{1}{n(n+1)} \sum_{i=1}^n \sum_{j=1}^{n+1} \rho (v_{i,j-\frac{1}{2}} - v_{i,j-\frac{1}{2}}^{old})^2 \\
 &+ \frac{1}{n^2} 2\Delta t \sum_{i=1}^n \sum_{j=1}^n \mu \left\| \frac{(\nabla u_{i,j} + \nabla u_{i,j}^T)}{2} \right\|_F^2 \\
 \text{where } \nabla u_{i,j} &= \begin{bmatrix} \frac{1}{\Delta x} (u_{i+\frac{1}{2},j} - u_{i-\frac{1}{2},j}) & \frac{1}{\Delta y} (u_{i,j+\frac{1}{2}} - u_{i,j-\frac{1}{2}}) \\ \frac{1}{\Delta x} (v_{i+\frac{1}{2},j} - v_{i-\frac{1}{2},j}) & \frac{1}{\Delta y} (v_{i,j+\frac{1}{2}} - v_{i,j-\frac{1}{2}}) \end{bmatrix}
 \end{aligned} \tag{4}$$

The sum of the first and second term of L_v can be derived by the average of the element-wise squared output because the output is velocity change due to viscosity. The third term denotes the fluid dissipation over the timestep. Through this term, L_v imposes a greater loss for higher viscosity, helping the

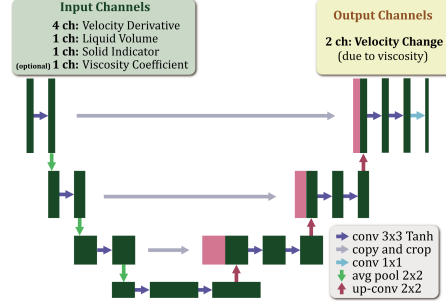


Figure 2: U-Shaped convolutional neural network for our data-driven viscosity solver.

network to create less velocity difference between adjacent fluid elements.

3.2 Architecture

3.2.1 u-Shaped CNN Model

We design the neural network architecture with the following two conditions in mind. First, the network allows input and output to have the same dimensions in width and height since they are defined on the same Eulerian grid. Second, the architecture should grasp important features of viscosity.

Viscosity is the resistance force of a fluid that reduces the velocity difference with neighboring fluids. In addition, fluids in the boundary layer, which is a thin layer near the boundary between the fluid and the solid, are more affected by viscosity. In other words, both the local feature representing adjacent fluid elements and the global feature revealing the location of the fluid element with respect to the fluid-rigid body scene are important for predicting the velocity. Therefore, we set the U-Net (Ronneberger et al., 2015) as a backbone because its multi-resolution feature maps help the network learn both local and global features of fluid-rigid body scene, and generate output of the same size as the input.

Unlike the image segmentation task of the original U-Net, our viscosity solver aims to output the velocity change for each dimension. The velocity change has the distinct property that it has the range $(-\infty, \infty)$. Constructing a network that produces exact velocity changes that could be positive or negative is a challenging task. To find the best architecture, we have experimented with various conditions to explore which model predicts accurate and unbiased output. As a result, our network consists of tangent hyperbolic activation function and average pooling layers as depicted in Fig. 2. We will elaborate the comparison of network configuration in section 4.

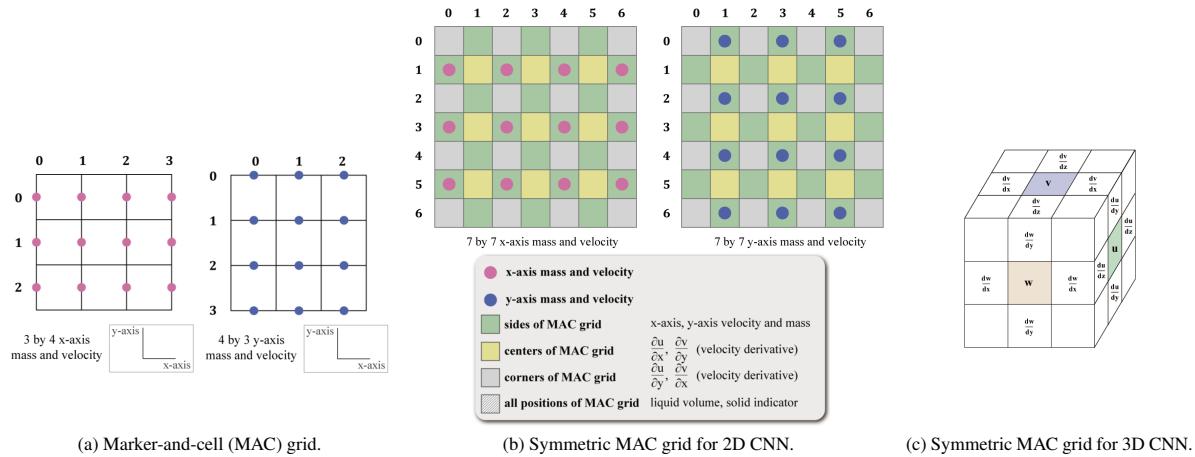


Figure 3: **Extension of the MAC grid for CNN-based viscosity solver** To adjust input channel dimensions without losing information or breaking symmetry, we propose the symmetric MAC grid. This grid ensures symmetry in all input channels.

3.2.2 Configuration of Input Channels

Our data-driven solver takes all necessary fluid quantities as input. This includes velocity derivatives to capture velocity differences with neighboring grid positions. In addition, the conventional solver requires fluid volume for each cell and a phase indicator (solid or not) on the grid positions. In summary, for the 2D grid, the input channels consist of four velocity derivative channels $\frac{\partial u}{\partial x}, \frac{\partial v}{\partial y}, \frac{\partial u}{\partial y}, \frac{\partial v}{\partial x}$ for the velocity field $\mathbf{u} = (u, v)$, one channel for fluid volume, and one channel for the solid indicator.

Velocity Derivative Channels. We use velocity derivatives instead of velocity itself in the construction of input channels because viscosity acts to reduce the velocity derivatives of the neighboring fluid. Moreover, using the velocity itself as input, the output would be varied by the velocity scale. For example, during free fall, there should be no velocity change caused by viscosity, but the network prones to output varying values due to the accelerated velocity. Hence, velocity derivative channels are employed for generalization across different velocity scales.

Fluid Volume Channel. To convey the volume quantity to the CNN, we suggest using a fluid volume channel. First, the signed distance from the fluid surface, constructed from a spherical surface of each particle, is computed for each corner (4 corners for 2D cell, and 8 for 3D cell) of the cell. From signed distances of two adjacent corner d^+ and d^- , edge occupancy of fluid is derived as $V_e = -(d^- / (d^+ - d^-))$. Then, the face (or volume) occupancy is approximated as the average of the occupancy of its 4 edges (or 6 faces in 3D). For the generalization of the fluid volume, we normalized the volume by the factor of maximum volume of one grid cell, so that the fluid volume is in $[0, 1]$.

Solid Indicator Channel. In addition, to deliver the status of the grid position, we suggest solid indicator channel, constructed by the signed distance $D_{i,j}$ from solid on the position $\mathbf{x} = (i, j)$. We set the solid indicator channel as follows:

$$\text{solid indicator}_{i,j} = \begin{cases} 0 & \text{if } D_{i,j} > 0 \\ 1 & \text{if } D_{i,j} \leq 0 \end{cases} \quad (5)$$

Viscosity Coefficient Channel. Finally, we suggest an optional viscosity coefficient channel. While the network converges well without this channel when trained for a fixed viscosity coefficient, it becomes essential for generalization across a wide range of viscosity coefficients. The viscosity coefficient channel, denoted by $\text{coeff}_{i,j}$, is assigned the dynamic viscosity μ in Pa·s. When implementing the channel, only cells with positive fluid volume are set as μ . The viscosity coefficient channel is set as follows:

$$\text{coeff}_{i,j} = \begin{cases} 0 & \text{if cell}_{i,j} \text{ is air or solid} \\ \mu & \text{if cell}_{i,j} \text{ is liquid} \end{cases} \quad (6)$$

The U-shaped architecture requires pooling, thus we expand the dimension by padding to be divisible by 2^n where n is the number of pooling layers. Zero padding is applied for velocity derivative, fluid volume, and viscosity coefficient channels, while one is used for a solid indicator channel. Random padding is applied to both the height and width dimensions to accommodate the various positions of the rigid body.

Our architecture readily extends to 3D fluid simulation using the symmetric MAC grid for 3D fluid quantities as shown in Fig. 3c. All layers of U-shaped architecture should be replaced into 3D layers, for example, 3D convolution, 3D average pooling, and 3D transposed convolution layers for upsampling. The input channels are given with $\frac{\partial u}{\partial x}, \frac{\partial v}{\partial y}, \frac{\partial w}{\partial z}, \frac{\partial u}{\partial y}, \frac{\partial u}{\partial z}, \frac{\partial v}{\partial x}, \frac{\partial v}{\partial z}, \frac{\partial w}{\partial x}, \frac{\partial w}{\partial y}$, fluid volume,

and solid indicator on symmetric MAC grid for fluid velocity $\mathbf{u} = (u, v, w)$.

3.3 Symmetric MAC Grid for CNN

Marker-and-cell (MAC) grid (Harlow and Welch, 1965) stores fluid velocity and mass quantities at the edge of the grid as illustrated in Fig. 3(a) so that the dimensions of the quantities vary from axis to axis. For example, for the grid size (n, n) , the velocity dimension of the x-direction is $(n + 1, n)$ and the y-direction is $(n, n + 1)$. Therefore, to input information from the x and y-direction together, modifying the dimensions is required.



Figure 4: **Comparison of input architecture.** (Left) Ground truth, (Middle) Result from the input with simple padding, and (Right) input with our symmetric MAC grid.

For fluid simulation, fluid quantities must be handled with care since tiny biases accumulate errors over successive frames. For the case of CNN for images, interpolation or padding could be a simple solution that rectifies the dimension. However, interpolation yields a loss of fluid information. Furthermore, simple padding causes an asymmetry between input channels that results in a tilted simulation, as Fig. 4 illustrates.

To address this problem we suggest a symmetric MAC grid for CNN that preserves all information without loss and generates symmetric input channels, regardless of the axis. Our symmetric MAC grid extends the original MAC grid including not only the centers but also the corners and edges of the cells. Therefore, the dimension of the symmetric MAC grid becomes $(2n + 1, 2n + 1)$ for the (n, n) grid. The quantities of mass and velocity of the x-direction are defined in pink circles, and the y-direction in blue circles as shown in Fig. 3(b). The network output channels, which are the x-direction and the y-direction velocity change, are also placed in the same position. The other cells without value are filled with zero not to influence the neural network.

In addition, this architecture facilitates the generation of derivative channels. To be more specific, $\frac{\partial u}{\partial x}, \frac{\partial v}{\partial y}$ are placed in the yellow cell and $\frac{\partial u}{\partial y}, \frac{\partial v}{\partial x}$ are in the gray cell. Liquid volume and solid indicator channels are calculated for all cells. Therefore, our input architecture guarantees that all input channels are symmetric and yields symmetric simulation results as shown in Fig. 4.

4 EXPERIMENTS ON U-SHAPED CNN ARCHITECTURE

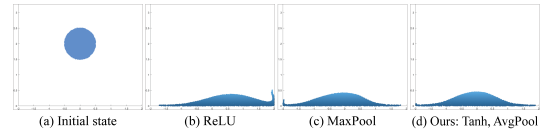


Figure 5: (a) From the initial resting state, the frame with (b) average pooling and ReLU (c) max pooling and tangent hyperbolic (Tanh) (d) average pooling and Tanh (ours).

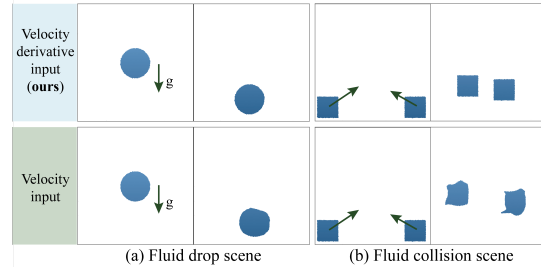


Figure 6: *Top row*: velocity derivative input and *Bottom row*: velocity input for (a) fluid drop and (b) colliding fluids scenes. The scenes are in a training dataset.

Ablation studies on network architecture are demonstrated in this section. We verify our network for various conditions including activation functions, pooling methods, and input channels.

Activation Function and Pooling Layers. The most important factor in network configuration is unbiased learning because bias results in tilted simulation. To explore the best CNN layers, we conducted ablation studies for the activation function between ReLU and tangent hyperbolic, and the pooling method between the max and average pooling. As shown in Fig. 5, ReLU which removes negative features, and max pooling layer which ignores small features produce a distorted result for fluid drop scene. Therefore, we design our network with a tangent hyperbolic activation function and average pooling layers.

Velocity Channel vs. Velocity Derivative Channel. As mentioned in subsection 3.2.2, the network receives velocity derivatives instead of raw velocity inputs. As shown in Fig. 6(a), the velocity derivative maintains the shape of the sphere during free fall, whereas the velocity channel distorts the sphere. This is because the velocity is accelerated, so the CNN solver prones to produce different outputs across the time steps. On the other hand, since the velocity derivative has a constant value of zero during free fall, the consistent output is generated. For the same reason, the velocity channel also produces a strange behavior for a more dynamic scene such as Fig. 6(b).

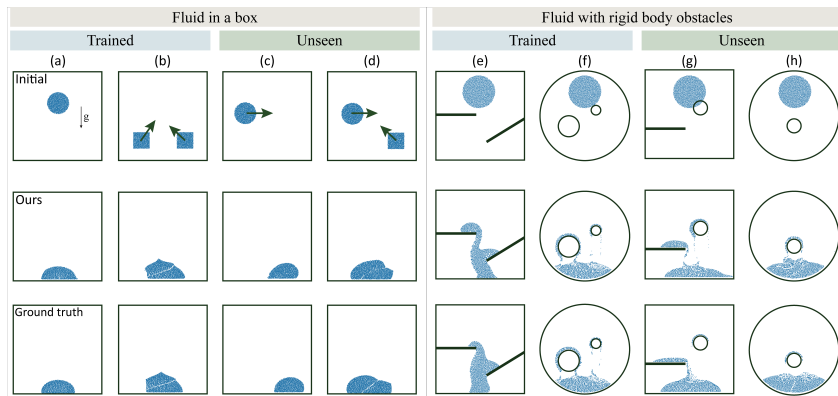


Figure 7: Simulation result comparison between our data-driven solver and ground truth for various scenarios.

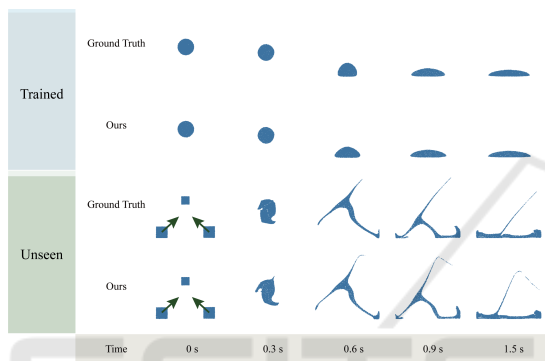


Figure 8: Simulation result comparison with timestamps for both trained and unseen scenes.

5 RESULTS

5.1 Universality of Data-Driven Viscosity Solver

In order to verify the universality of the proposed architecture, we demonstrate fluid simulation with various scenarios: fluids in a box, fluids with rigid body obstacles, and 3D scenes exhibiting buckling effect.

5.1.1 Viscous Fluid with Rigid Bodies

We demonstrate our solver with two scenarios: fluid in a box and fluid with rigid body obstacles in Fig. 7. Both scenes are constructed with 1 Pa·s dynamic viscosity, 300 Hz framerate, $2m \times 2m$ space, and 100×100 Eulerian grid dimension. The training epoch is 2000 and the learning rate is 5×10^{-4} . Our data-driven viscosity solver (with four pooling layers) takes approximately 0.007 seconds per frame.

For the fluid in a box scenario, the neural network is trained using 5-second fluid simulations of Fig. 7(a,b), resulting in a dataset consisting of 1500 frames for each scene. Our solver outputs acceptable

visual fidelity for both trained and unseen(Fig. 7(c,d)) scenarios.

We tested our method in more complex scenes with rigid body obstacles. The network is trained with 10-seconds fluid simulations of Fig. 7(e,f). Our solver also shows similar results to ground truth for unseen scenes of Fig. 7(g,h).

For further validation of our data-driven simulation, we performed a timestamp-based comparison between ours and the ground truth. For both trained and unseen scenarios, our data-driven solver consistently predicts similar velocity changes for every timestamp as visualized in Fig. 8.

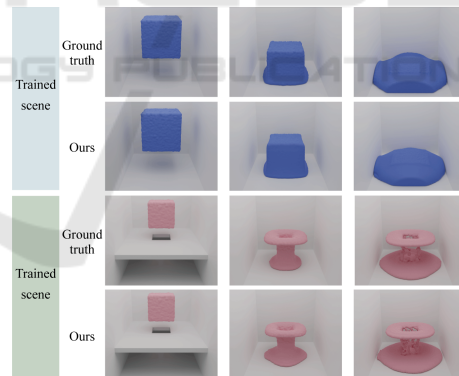


Figure 9: **3D fluid-rigid simulation** Trained solely on the fluid drop scene, our solver similarly simulates the unseen scene with the ground truth. The space is $2m \times 2m \times 2m$ with $80 \times 80 \times 80$ Eulerian grid dimension.

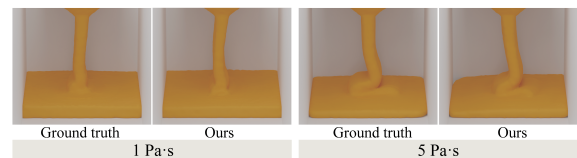


Figure 10: Buckling effect of viscous fluid with 1 and 5 Pa·s dynamic viscosities. The simulated space is $0.6m \times 1m \times 0.6m$ with Eulerian grid dimension $48 \times 80 \times 48$.

5.1.2 Extension to 3D Fluid Simulation

As depicted in Fig. 9, our solver trained on a fluid drop scene successfully performs 3D fluid simulations for both trained and unseen scenes. The network effectively learns viscosity characteristics using data from a simple scene and extends its capability in more complex scenes with rigid bodies.

Our viscosity solver (with two pooling layers) takes approximately 0.39 seconds per frame with an $80 \times 80 \times 80$ Eulerian grid for scenes in Fig. 9.

Moreover, our solver simulates the buckling effect of viscous fluid, demonstrated in Fig. 10. Similar to the ground truth, the fluid with a lower viscosity coefficient (1 Pa·s) exhibits reduced buckling compared to the 5 Pa·s case.

5.2 Generalization of Dynamic Viscosity

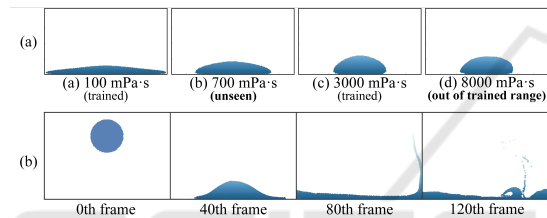


Figure 11: (a) Fluid drop scenes with one neural network model for various viscosity coefficients. (b) Simulation result with viscosity coefficient channel set by left side as 1000 and right side as 1 mPa·s.

Our network can be generalized to various dynamic viscosities μ . To provide dynamic viscosity information, the viscosity coefficient input channel is optionally given to the network as mentioned in subsection 3.2.2. U-shaped viscosity solver with four pooling layers is trained by 12000 frames consisting of dynamic viscosity 1, 100, 500, 1000, 2000, and 3000 mPa·s equally. The trained model is well generalized for unseen coefficients as shown in Fig. 11(a). For example, although 700 mPa·s is not learned, it is spread narrower than 100 and wider than 3000.

However, as shown in Fig. 11, our data-driven solver cannot extrapolate the viscosity simulations outside the trained dynamic viscosity range. The simulation is tilted to the right and upside because a large dynamic viscosity affects the network to output large x-direction (right) and y-direction (upward) velocity changes. Therefore, a corresponding dataset is required to generate the model for a wide range of dynamic viscosity.

Moreover, mixed dynamic viscosity can be simulated without additional training. For example, we assigned dynamic viscosity 1000 to the left and 1 mPa·s to the right side in Fig. 11(b). The left side is less

spread out due to the large viscosity, whereas the right side fluid spreads like water. We anticipate that our model can be applied to complex scenes with multiple viscosity coefficients. In-depth exploration of multiple viscosity coefficients remains a future work.

5.3 Real-Time Application of Rigid-Fluid Interaction

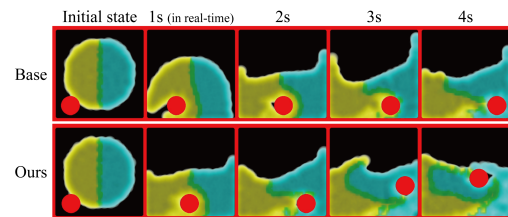


Figure 12: Progression of real-time 2D rigid-fluid interactive application. *Top row*: Iterative viscosity solver (Batty and Bridson, 2008). *Bottom row*: Our U-shaped CNN viscosity solver which is nearly twice faster.

To demonstrate the efficiency of our solver, we applied our solver to a real-time interactive application of paint mixing Fig. 12 where a rigid pointer (red circle) perturbs fluid of two different colors. The scene consists of (25, 25) grid dimension and 848 particles. On average, the baseline (Batty and Bridson, 2008) runs at 14 Hz (71 ms) with 32ms iterative viscosity solver, and ours takes 23Hz (42 ms) with 4ms data-driven viscosity solver.

6 CONCLUSION

In this paper, we propose a data-driven viscosity solver for a hybrid Lagrangian/Eulerian fluid simulator, utilizing a symmetric MAC grid to represent fluid quantities without symmetry breaking or interpolation. Our U-shaped CNN, with input channels of velocity derivatives, fluid volume, and solid indicators, predicts velocity changes with reasonable error in both trained and unseen scenes. To show the effectiveness of our architecture, we demonstrate various scenes, including fluid with rigid bodies, mixed dynamic viscosity, and 3D fluid with a buckling effect.

We believe that the proposed data-driven viscosity solver can be utilized for diverse fluid simulation, as demonstrated in a real-time interactive application. Our solver's generalizability across a wide viscosity range is achieved through a viscosity coefficient channel and a physics-based variational loss. While we primarily analyze our network in 2D space, we plan to further explore and optimize our method for 3D viscous fluid simulation.

ACKNOWLEDGEMENTS

This work was supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (No.00223446, Development of object-oriented synthetic data generation and evaluation methods).

REFERENCES

- Aanjaneya, M., Han, C., Goldade, R., and Batty, C. (2019). An efficient geometric multigrid solver for viscous liquids. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 2(2):1–21.
- Batty, C. and Bridson, R. (2008). Accurate viscous free surfaces for buckling, coiling, and rotating liquids. In *In Proc. Symposium on Computer Animation*. ACM/Eurographics, page 219–228.
- Cheng, M., Fang, F., Pain, C. C., and Navon, I. (2020). Data-driven modelling of nonlinear spatio-temporal fluid flows using a deep convolutional generative adversarial network. *Computer Methods in Applied Mechanics and Engineering*, 365:113000.
- Chentanez, N. and Müller, M. (2011). Real-time eulerian water simulation using a restricted tall cell grid. In *ACM Siggraph 2011 Papers*, pages 1–10.
- Gao, Y., Zhang, Q., Li, S., Hao, A., and Qin, H. (2020). Accelerating liquid simulation with an improved data-driven method. In *Computer Graphics Forum*, volume 39, pages 180–191. Wiley Online Library.
- Guo, X., Li, W., and Iorio, F. (2016). Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 481–490.
- Harlow, F. H. and Welch, J. E. (1965). Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *The physics of fluids*, 8(12):2182–2189.
- Jiang, C., Schroeder, C., Selle, A., Teran, J., and Stomakhin, A. (2015). The affine particle-in-cell method. *ACM Transactions on Graphics (TOG)*, 34(4):1–10.
- Kim, B., Azevedo, V. C., Thurey, N., Kim, T., Gross, M., and Solenthaler, B. (2019). Deep fluids: A generative network for parameterized fluid simulations. In *Computer graphics forum*, volume 38, pages 59–70. Wiley Online Library.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ladický, L., Jeong, S., Solenthaler, B., Pollefeys, M., and Gross, M. (2015). Data-driven fluid simulations using regression forests. *ACM Transactions on Graphics (TOG)*, 34(6):1–9.
- Li, Z. and Farimani, A. B. (2022). Graph neural network-accelerated lagrangian fluid simulation. *Computers & Graphics*, 103:201–211.
- McAdams, A., Sifakis, E., and Teran, J. (2010). A parallel multigrid poisson solver for fluids simulation on large grids. In *Symposium on Computer Animation*, pages 65–73.
- Obiols-Sales, O., Vishnu, A., Malaya, N., and Chandramowlishwaran, A. (2020). Cfdnet: A deep learning-based accelerator for fluid simulations. In *Proceedings of the 34th ACM international conference on supercomputing*, pages 1–12.
- Richter-Powell, J., Lipman, Y., and Chen, R. T. (2022). Neural conservation laws: A divergence-free perspective. *arXiv preprint arXiv:2210.01741* (2022).
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.
- Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. (2020). Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR.
- Shao, H., Huang, L., and Michels, D. L. (2022). A fast unsmoothed aggregation algebraic multigrid framework for the large-scale simulation of incompressible flow. *ACM Transactions on Graphics (TOG)*, 41(4):1–18.
- Tang, W., Shan, T., Dang, X., Li, M., Yang, F., Xu, S., and Wu, J. (2017). Study on a poisson’s equation solver based on deep learning technique. In *2017 IEEE Electrical Design of Advanced Packaging and Systems Symposium (EDAPS)*, pages 1–3.
- Thurey, N., Weißenow, K., Prantl, L., and Hu, X. (2020). Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA Journal*, 58(1):25–36.
- Tompson, J., Schlachter, K., Sprechmann, P., and Perlin, K. (2017). Accelerating eulerian fluid simulation with convolutional networks. In *International Conference on Machine Learning*, pages 3424–3433. PMLR.
- Treuille, A., Lewis, A., and Popović, Z. (2006). Model reduction for real-time fluids. *ACM Transactions on Graphics (TOG)*, 25(3):826–834.
- Wiewel, S., Kim, B., Azevedo, V. C., Solenthaler, B., and Thurey, N. (2020). Latent space subdivision: stable and controllable time predictions for fluid flow. In *Computer Graphics Forum*, volume 39, pages 15–25. Wiley Online Library.
- Xiao, X., Wang, H., and Yang, X. (2019). A cnn-based flow correction method for fast preview. In *Computer Graphics Forum*, volume 38, pages 431–440. Wiley Online Library.
- Xiao, X., Yang, C., and Yang, X. (2018a). Adaptive learning-based projection method for smoke simulation. *Computer Animation and Virtual Worlds*, 29(3-4):e1837.
- Xiao, X., Zhou, Y., Wang, H., and Yang, X. (2018b). A novel cnn-based poisson solver for fluid simulation. *IEEE transactions on visualization and computer graphics*, 26(3):1454–1465.
- Yang, C., Yang, X., and Xiao, X. (2016). Data-driven projection method in fluid simulation. *Computer Animation and Virtual Worlds*, 27(3-4):415–424.