

LAGRANGIAN FLUID SIMULATION WITH CONTINUOUS CONVOLUTIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

We present an approach to Lagrangian fluid simulation with a new type of convolutional network. Our networks process sets of moving particles, which describe fluids in space and time. Unlike previous approaches, we do not build an explicit graph structure to connect the particles but use spatial convolutions as the main differentiable operation that relates particles to their neighbors. To this end we present a simple, novel, and effective extension of N-D convolutions to the continuous domain. We show that our network architecture can simulate different materials, generalizes to arbitrary collision geometries, and can be used for inverse problems. In addition, we demonstrate that our continuous convolutions outperform prior formulations in terms of accuracy and speed.

1 INTRODUCTION

Understanding physics is a fundamental ability required to reason about our environment and to interact with it. Neural networks have emerged as a particularly promising approach to capture the complexity of natural phenomena from data (Ling et al., 2016; Tompson et al., 2017; Morton et al., 2018). An important aspect of learning physics with neural networks is the choice of representation. Lagrangian representations based on particles are particularly popular and have supported a range of impressive results with rigid bodies, deformable solids, and fluids (Battaglia et al., 2016; Mrowca et al., 2018; Li et al., 2019). Many of these approaches use graph structures to define interactions; the existence of an edge determines in a binary fashion whether two particles interact.

However, a wide range of physical processes such as fluid mechanics are described by continuous partial differential equations rather than discrete graph structures. The continuous, volumetric, and tightly coupled nature of these processes causes inherent difficulties for graph-based approaches, such as a large number of edge connections that must be established, tracked, and disengaged as the particles move. In this work, instead of using graphs as the underlying representation, we adopt a continuous viewpoint. We propose to use convolutional networks (ConvNets) with *continuous convolutions* on particles for learning fluid mechanics. We treat fluids as spatially continuous functions sampled at a finite number of (continuously evolving) positions and process them with a novel continuous convolution layer. This matches the continuous nature of the problem more closely and simplifies the definition of neural networks by abstracting the underlying particle representation.

We naturally extend the grid-based filter representation commonly used for discrete convolutions to the continuous domain by simple linear interpolation. Linear interpolation of the filters allows efficient lookup of spatially varying filter values at arbitrary positions while retaining the compactness and efficiency of the grid representation. We show that our convolutions, despite their simplicity, perform better than more sophisticated representations (Wang et al., 2018; Schenck & Fox, 2018).

With the presented continuous convolution layer, we develop an efficient ConvNet architecture for learning fluid mechanics. The network processes sets of particles. We use dynamic particles to represent the fluid and static particles to describe the boundary of the scene. Modelling the scene boundary with particles makes it easy to apply our network to new scenes and allows the network to learn collision handling in a unified framework. Our network generalizes to arbitrary obstacle configurations and can simulate a range of material behavior. To demonstrate the usefulness of a learned – and hence differentiable – fluid simulator, we show that material properties can be estimated from observed simulation data. Experimental results indicate that the presented approach outperforms a state-of-the-art graph-based framework (Li et al., 2019).

2 RELATED WORK

Fluids encompass a range of materials that are important in everyday life and throughout science and engineering (Wilcox, 2006). A prominent set of methods, known as Smoothed Particle Hydrodynamics (SPH), employs a Lagrangian viewpoint to simulate these phenomena. SPH originated in particle-based models for astrophysics (Gingold & Monaghan, 1977) and has become extremely popular for simulating complex flows in many fields of science (Monaghan, 1988). Among others, it has been highly successful for modeling interface flows (Colagrossi & Landrini, 2003), complex multi-phase phenomena (Hu & Adams, 2006), and even magnetohydrodynamics (Price, 2012).

SPH and its variants have been widely used to model complex real-world phenomena for visual effects. Following the earlier development of physics-based fluid simulation for special effects by Foster & Metaxas (1996), the introduction of SPH (Müller et al., 2003) has led to a large class of powerful algorithms (Solenthaler & Pajarola, 2009; Bender & Koschier, 2015). These applications often involve complex geometries at large scales, and extensions such as FleX (Macklin et al., 2014) and pressure-aware rigid body coupling (Gissler et al., 2018) broaden the framework to encompass many physical phenomena.

Lagrangian flows have also been considered in machine learning. The seminal work of Ladický et al. (2015) demonstrated that flow representations can be learned with regression forests. CNNs were used by Tompson et al. (2017) to accelerate the expensive pressure correction step of grid-based solvers, while other works have focused on super-resolution (Xie et al., 2018), learning time evolution via the Koopman operator (Morton et al., 2018), and learning reduced representations (Kim et al., 2019). Differentiable SPH solvers were proposed to solve control tasks for robotic applications (Schenck & Fox, 2018). More generic physics simulations for Lagrangian rigid and deformable bodies were the target of a series of works that developed graph-based representations (Battaglia et al., 2016; Sanchez-Gonzalez et al., 2018). Such graph-based representations were recently applied directly to fluid simulation (Mrowca et al., 2018; Li et al., 2019). We share with these recent works the goal of modelling Lagrangian fluids with differentiable neural networks, but take a different tack: rather than using graph-based representations, we work with point clouds and continuous convolutions over the spatial domain.

From a technical perspective, our approach is also related to existing works that apply convolutions on point clouds. A number of methods transferred the convolution concept to point clouds in the context of semantic classification and segmentation of 3D objects (Hua et al., 2018; Atzmon et al., 2018; Hermosilla et al., 2018; Li et al., 2018; Su et al., 2018; Wu et al., 2018; Xu et al., 2018; Lei et al., 2019). Particularly notable in our context are the works of Wang et al. (2018), who used continuous convolutions to compute the scene flow between two point clouds, and the aforementioned work of Schenck & Fox (2018), who used convolutions to implement a differentiable version of position-based fluids (Macklin & Müller, 2013). Both works define continuous convolution operators that can be used for regression tasks and we compare with them directly in Section 6.

3 BASICS

Fluids have been studied for centuries, and the Navier-Stokes equations for incompressible fluids are well established (Batchelor, 1967):

$$\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{v} + \mathbf{g}, \quad \text{s.t. } \nabla \cdot \mathbf{v} = 0. \quad (1)$$

A common approach to solve these partial differential equations is to approximate the fluid with a set of smooth particles (Monaghan, 1988). Each particle corresponds to a continuous blob of matter and carries the local properties of the fluid, such as velocity and density, which move with the flow. This is motivated by the fact that a function $A(x)$ can be represented by an integral interpolation

$$A(\mathbf{x}) = \int A(\mathbf{x}') \delta(\|\mathbf{x} - \mathbf{x}'\|_2) dV(\mathbf{x}'), \quad (2)$$

where $\delta(x)$ denotes the Dirac delta function. This equation can be discretized as

$$A(\mathbf{x}) \approx \sum_i V_i A_i W(\|\mathbf{x} - \mathbf{x}'\|_2, h), \quad (3)$$

where V_i is the volume at the given point in space and

$$\lim_{h \rightarrow 0} W(|\mathbf{x} - \mathbf{x}'|, h) = \delta(|\mathbf{x} - \mathbf{x}'|). \quad (4)$$

Here $W(\mathbf{x})$ is a smooth kernel or convolution with radius h , usually in the form of a Gaussian distribution, but more complex functions can also be used. In practice, the kernel is finite and yields localized neighborhoods of particles that interact via interactions weighted by the kernel or its derivatives. In this way, the continuous description for fluids from Equation 1 can be discretized in a Lagrangian fashion and solved numerically. Typically, internal forces are calculated based on the local pressure, viscosity, and surface tension, which give an update for the position of each particle. Below, we adopt the position-based fluids (PBF) method (Macklin & Müller, 2013; Macklin et al., 2014), which likewise is based on SPH, but reformulates the updates as constraints on the positions.

4 CONTINUOUS CONVOLUTIONS

The discrete convolution operator as commonly used in ConvNets is defined as

$$(f * g)(\mathbf{x}) = \sum_{\boldsymbol{\tau} \in \Omega} f(\mathbf{x} + \boldsymbol{\tau})g(\boldsymbol{\tau}), \quad (5)$$

where f and g are the input and the pre-mirrored filter function, \mathbf{x} is the position, $\boldsymbol{\tau}$ is the shift vector, and Ω is the set of shift vectors that defines the support of the filter function. On regular data such as images, the positions \mathbf{x} range over a regular grid and the shift vectors $\boldsymbol{\tau}$ are integer-valued, i.e. $\mathbf{x}, \boldsymbol{\tau} \in \mathbb{Z}^d$ for some dimensionality d . In the continuous domain, the convolution is defined as

$$(f * g)(\mathbf{x}) = \int_{\mathbb{R}^d} f(\mathbf{x} + \boldsymbol{\tau})g(\boldsymbol{\tau})d\boldsymbol{\tau}, \quad (6)$$

where \mathbf{x} and $\boldsymbol{\tau}$ are real-valued vectors, i.e. $\mathbf{x}, \boldsymbol{\tau} \in \mathbb{R}^d$.

We now adapt this definition to unstructured point clouds. In this setting we have a finite number of points that sample the function f but do not lie on a grid. For a point cloud with $i = 1, \dots, N$ points with values f_i at positions \mathbf{x}_i , we define the convolution at a position \mathbf{x} as

$$(f * g)(\mathbf{x}) = \frac{1}{\psi(\mathbf{x})} \sum_{i \in \mathcal{N}(\mathbf{x}, R)} a(\mathbf{x}_i, \mathbf{x}) f_i g(\Lambda(\mathbf{x}_i - \mathbf{x})). \quad (7)$$

$\mathcal{N}(\mathbf{x}, R)$ is the set of points within a radius R around \mathbf{x} . a is a scalar function that can be used for density normalization specific to the points \mathbf{x}_i and \mathbf{x} as in Hermosilla et al. (2018). In the simplest case, a can be a constant function: $a = 1$. In our case we want to ensure a smooth response of our convolution under varying particle neighborhoods, therefore we define a as a window function:

$$a(\mathbf{x}_i, \mathbf{x}) = \begin{cases} (R^2 - \|\mathbf{x}_i - \mathbf{x}\|_2^2)^3 & \text{for } \|\mathbf{x}_i - \mathbf{x}\|_2 < R \\ 0 & \text{else.} \end{cases} \quad (8)$$

A similar function has been used by Müller et al. (2003) in the SPH framework. ψ is another scalar function for normalization, which can be set in our implementation as either

$$\psi(\mathbf{x}) = 1 \quad \text{or} \quad \psi(\mathbf{x}) = \sum_{i \in \mathcal{N}(\mathbf{x}, R)} a(\mathbf{x}_i, \mathbf{x}). \quad (9)$$

We use $\psi(\mathbf{x}) = 1$ as changes in the density of particles are an important feature for simulating fluids.

For the filter function g we simply use a regular grid to store the filter values but use linear interpolation to make g a continuous function. In addition, we use a mapping $\Lambda(\mathbf{r})$ of a unit ball to a unit cube to implement spherical filters as shown in Figure 1. We use the mapping described by Griepentrog et al. (2008). The intermediate coordinate mapping Λ provides the flexibility to implement different spatial shapes while keeping the advantages of a regular grid for the storage and lookup of filter values.

Note that Equation 7 uses a similar approximation as in the SPH framework (Equation 3). Assuming that each point represents the same volume, V_i is a constant factor in Equation 3, which we drop in our definition.




Figure 1: We use a regular grid to store the filter values of our continuous convolutions. To look up filter values for spherical filter shapes, we map the relative position \mathbf{r} to normalized coordinates, which define the lookup with trilinear interpolation in the regular grid.

5 LEARNING FLUID MECHANICS WITH CONVOLUTIONAL NETWORKS

Our goal is to learn fluid mechanics from observing the motion of particles, thus the input to our ConvNet is a set of particles with corresponding features. Note that we must assign a feature vector to each particle as the position itself is not a feature but defines the particle’s position in space. As a feature vector we use a constant scalar 1 accompanied by the velocity \mathbf{v} and the viscosity ν . A particle p_i with its position and input feature vector at timestep n is thus a tuple $(\mathbf{x}_i^{n*}, [1, \mathbf{v}_i^{n*}, \nu_i])$. Defining the velocity explicitly as an input feature allows us to compute intermediate velocities and positions as in Ladický et al. (2015) and to apply external forces and pass this information to the network. We compute the intermediate positions \mathbf{x}_i^{n*} and velocities \mathbf{v}_i^{n*} beginning with timestep n with the midpoint method as

$$\mathbf{v}_i^{n*} = \mathbf{v}_i^n + \Delta t \mathbf{a}_{\text{ext}} \quad (10)$$

$$\mathbf{x}_i^{n*} = \mathbf{x}_i^n + \Delta t \frac{\mathbf{v}_i^n + \mathbf{v}_i^{n*}}{2}. \quad (11)$$

The vector \mathbf{a}_{ext} is an acceleration through which we can apply external forces to control the fluid or to simply apply gravity. The intermediate positions and velocities lack any interactions between particles or the scene, which we are going to implement with a ConvNet. To enable the network to handle collisions with the environment we define a second set of static particles s_j . We sample particles on the boundaries of the scene with normals \mathbf{n}_j as the feature vectors, i.e. $s_j = (\mathbf{x}_j, [\mathbf{n}_j])$. Our network implements the function

$$[\Delta \mathbf{x}_1, \dots, \Delta \mathbf{x}_N] = \text{ConvNet}(\{p_1, \dots, p_N\}, \{s_1, \dots, s_M\}), \quad (12)$$

which uses convolutions to combine features from both particle sets. $\Delta \mathbf{x}$ is a correction of the position which accounts for all particle interactions including the collision handling with the scene. Finally, we apply the correction to update positions and velocities for $n + 1$ as

$$\mathbf{x}_i^{n+1} = \mathbf{x}_i^{n*} + \Delta \mathbf{x}_i \quad (13)$$

$$\mathbf{v}_i^{n+1} = \frac{\mathbf{x}_i^{n+1} - \mathbf{x}_i^n}{\Delta t}. \quad (14)$$

Note that the updated position \mathbf{x}_i^{n+1} depends on the networks output vector $\Delta \mathbf{x}_i$ and allows us to directly define our learning objective on the particle positions.

5.1 NETWORK ARCHITECTURE

We use a simple convolutional architecture with an effective depth of four. An overview of the network is shown in Figure 2. Since we want to compute the correction for all dynamic particles in our scene, we compute convolutions for the intermediate positions defined in Equation 11. Our network is a sequence of continuous convolutions (CConv), which are defined by an input particle set, the positions at which we want to evaluate the convolution, its filters G and the radius R . For instance, to describe a convolution on the static particles s_i at the intermediate positions \mathbf{x}_i^{n*} as used in the beginning of our network we can write

$$[\mathbf{f}_1, \dots, \mathbf{f}_N] = \text{CConv}(\{s_1, \dots, s_M\}, [\mathbf{x}_1^{n*}, \dots, \mathbf{x}_N^{n*}], G, R), \quad (15)$$

where \mathbf{f}_i are the computed output features for each position \mathbf{x}_i^{n*} . G is a 5D array storing all filters in the layout [width, height, depth, ch_{in} , ch_{out}]. In contrast to discrete convolutions on a regular grid,




Figure 2: Our network has a depth of four. In the first depth level we compute convolutions at each dynamic particle location with the set of static particles that define the environment as well as the dynamic particle set. We also directly process the features of each particle via a fully-connected stream. In the following levels, we compute convolutions only on the dynamic particles. At each level we use addition to aggregate the features computed by convolutions and fully-connected layers. Between the second and third level we also include a residual connection. The final level generates the correction of the position $\Delta \mathbf{x}$.

the spatial filter dimensions here do not define the receptive field but the resolution of the filters. The receptive field depends only on the radius R describing the spatial extent. Throughout our network we use filters with a spatial resolution of $[4, 4, 4]$ and a radius of 4.5 times the particle radius.

For convolutions within the dynamic particles we exclude the particle at which we evaluate the convolution and instead process the particle’s own features in a stream of fully-connected layers. After each depth level we then combine the result from the convolutions and the fully-connected layers by addition. This can be interpreted as a convolution with a spatial resolution of $4 \times 4 \times 4 + 1$.

5.2 TRAINING PROCEDURE

We train our fluid simulation network in a supervised fashion based on the observed particle positions from a ground-truth simulation. Our loss is defined as follows:

$$\mathcal{L}^{n+1} = \sum_{i=1}^N \phi_i \|\mathbf{x}_i^{n+1} - \hat{\mathbf{x}}_i^{n+1}\|_2^\gamma. \quad (16)$$

The ground-truth position at timestep $n + 1$ is denoted by $\hat{\mathbf{x}}_i^{n+1}$ and the predicted position from the network is denoted by $\mathbf{x}_i^{n+1} = \mathbf{x}_i^{n*} + \Delta \mathbf{x}_i$, where $\Delta \mathbf{x}_i$ is provided by the network. ϕ_i is an individual weight for each point to emphasize the loss for particles with a small number of neighbors, which we express as $\phi_i = \exp(-\frac{1}{c}|\mathcal{N}(\mathbf{x}_i^{n*})|)$. We choose $c = 40$, which corresponds to the average number of neighbors across our experiments. Particles with a small number of neighbors are close to the surface or interact with the scene boundary. Both cases are important for fluid simulation because particles near the surface define the liquid-air interface, which is particularly salient, and particles near the scene boundary require collision handling. The parameter $\gamma = 0.5$ makes our loss function more sensitive to small particle motions, which is important for increasing the accuracy and visual fidelity for small fluid flows. During training we predict particle positions for two future timesteps, namely $n + 1$ and $n + 2$. The combined loss \mathcal{L} is

$$\mathcal{L} = \mathcal{L}^{n+1} + \mathcal{L}^{n+2}. \quad (17)$$

We found that optimizing a loss defined over two frames improves the overall quality of the simulation. (Optimization for three frames did not result in further improvements.) We optimize \mathcal{L} over 50,000 iterations with Adam (Kingma & Ba, 2015) and a learning rate decay with multiple steps, starting with a learning rate of 0.001 and stopping with $1.56 \cdot 10^{-5}$.




Figure 3: The first row shows the simulation of two fluid bodies by our trained network. The second row shows simulation of the same scenario by DFSPH (Bender & Koschier, 2015), a high-fidelity simulator that was executed with small timesteps: down to 0.001ms. Despite using a much larger timestep (0.02ms), our convolutional network produces results of comparable visual fidelity. Note that our particles are falling slightly more slowly due to differences in the integration of positions and the much larger timestep. See the supplementary video: <https://sites.google.com/view/lfsbcc>.

5.3 DATASETS

We have trained our network on multiple datasets. For quantitative comparisons we trained our network on the dam break data from (Li et al., 2019). The scene simulates the behaviour of a randomly placed fluid block in a static box. We generate 2000 scenes for training and 300 for testing. The data was generated with FleX, which is a position-based simulator (Macklin et al., 2014) targeting real-time applications.

We also trained our network on more challenging data generated with DFSPH (Bender & Koschier, 2015), which prioritizes quality over runtime. DFSPH can generate accurate fluids with very low volume compression: a desired property. We generate ground-truth data by randomly placing multiple bodies of fluids in 10 different box-like scenes and simulate them for 16 seconds each with an adaptive timestep of up to 1 kHz. The time resolution of the generated data is 50 Hz. We show a qualitative comparison of our method to the ground truth in Figure 3. We generate 200 scenes for training and a test set with 20 scenes. To train networks that can deal with multiple materials, we additionally generate 200 scenes with fluids of varying viscosity. For estimating material properties we generate 5 test scenes that only differ in the viscosity parameter.

6 EVALUATION

Baselines. We compare our method to DPI-Nets (Li et al., 2019), which was previously shown to significantly outperform prior formulations such as hierarchical relation networks (Mrowca et al., 2018). We also compare our continuous convolution formulation to continuous convolution representations used in SPNets (Schenck & Fox, 2018) and PCNN (Wang et al., 2018). Figure 4 provides a qualitative comparison to DPI-Nets. Table 1 provides a quantitative comparison. We report the average error of the particle positions with respect to the ground truth. To this end, we use every 5th frame for initialization and compute the deviation from the ground truth for two subsequent frames, denoted by $n + 1$ and $n + 2$.

We have trained and tested all methods on the dam break dataset from Li et al. (2019) as well as our data (generated with a high-fidelity simulator (Bender & Koschier, 2015)). We train DPI-Nets for 5 epochs and all other networks for 50,000 iterations, which corresponds to 2.7 epochs for the dam break data and 5 epochs for our data. Training takes about a day for our method with our convolutions on an NVIDIA RTX 2080Ti. Training with PCNN convolutions (Wang et al., 2018) takes about 2 days. DPI-Nets trains for about one day. For SPNets convolutions, we estimated a training time of more than 29 days by extrapolating from timing of a smaller number of iterations. We thus only report inference runtime for this method. Note that the convolutions of SPNets have been specifically designed to implement the position-based fluids algorithm, while we use them here




Figure 4: Qualitative comparison with DPI-Nets on a test sequence from our dataset. Two fluid bodies collide inside a container. DPI-Nets works well on data with little variance but has problems with more complex scenes and high particle velocities. The DPI-Nets simulation becomes unstable immediately after the fluid hits the box. The fluid behaviour predicted by our network matches the ground truth more closely and remains stable for the whole sequence. Both networks have been trained on the same data. Test sequences are distinct from training sequences. See the supplementary video: <https://sites.google.com/view/lfsfcc>

| Method | Average error (mm) | | Frame inference time (ms) | Frame NNS time (ms) | |
|--------------|--------------------|-------------|---------------------------|---------------------|-------------|
| | $n + 1$ | $n + 2$ | | | |
| DPI DamBreak | DPI-Nets | 12.73 | 25.38 | 202.56 | 103.26 |
| | SPNets Convs | – | – | 1058.46 | 5.24 |
| | PCNN Convs | 0.72 | 1.67 | 187.34 | 2.42 |
| | Ours | 0.62 | 1.49 | 12.01 | 2.14 |
| Our data | DPI-Nets | 26.19 | 51.77 | 305.55 | 171.22 |
| | SPNets Convs | – | – | 784.35 | 10.19 |
| | PCNN Convs | 0.67 | 1.87 | 319.17 | 2.78 |
| | Ours | 0.56 | 1.51 | 16.47 | 2.38 |

Table 1: Accuracy and runtime analysis. We compare the error between the ground truth and two predicted future frames on the test sets. We also compare the average inference time per frame and the time used for nearest neighbor search (NNS). For DPI-Nets the nearest neighbor search accounts for about half of the total inference time. For all other methods the time used for performing the convolution and other operations dominates the inference time. The PCNN Convs network uses the same nearest neighbor search as Ours.

with a much larger number of channels, which explains the very long runtime. As shown in Table 1, our method outperforms all baselines with respect to both accuracy and inference time.

Ablation study. We do an ablation study to justify our decisions in the design of the convolutions and the network. Table 2 reports position errors averaged over the test sequences in the same manner as in Table 1. Additionally, we give the errors for two predicted frames initialized with the frames at the end of each sequence to measure the errors for small flow velocities. Large errors for small velocities are more susceptible to cause visible motion artifacts. Therefore we are interested in small errors in this specific setting. We present the results in Table 2 and show the importance of the interpolation, the window function, and the loss design.

Generalization. In Figure 5 we show that our network generalizes well to scenes with drastically different geometry than seen during training. (The training set only used box-like containers. See the appendix for visualization.) This scenario also demonstrates that we can emit particles during simulation, which can be costly for methods that build and maintain explicit graph structures. Figure 6 demonstrates generalization along a different dimension. Here we show that we can set the viscosity of the fluid at test time to a value not seen during training. The fluid shape used in this example was also not seen during training.

| Method | Average error (mm) | | Seq. end error (mm) | |
|------------------------|--------------------|-------------|---------------------|-------------|
| | $n + 1$ | $n + 2$ | $n + 1$ | $n + 2$ |
| Ours | 0.67 | 1.87 | 0.25 | 0.74 |
| Ours w/o interpolation | 0.79 | 2.24 | 0.30 | 0.89 |
| Ours w/o window | 0.77 | 2.21 | 0.30 | 0.89 |
| Ours w/naïve loss | 0.69 | 1.86 | 0.27 | 0.77 |

Table 2: Ablation study. Comparison of the position errors to the ground truth averaged over whole and the end of the sequences. *Ours w/o interpolation* uses nearest neighbor instead of trilinear interpolation for the convolution filters. *Ours w/o window* uses $a(\mathbf{x}_i, \mathbf{x}) = 1$ in Equation 7. *Ours w/naïve loss* uses the Euclidean distance as loss, i.e. we set $\gamma = 1$ and $\phi_i = 1$ in Equation 16.




Figure 5: Generalization to unseen environments: we use an emitter to fill up a virtual river with fluid particles demonstrating generalization w.r.t. the scene and the number of particles.

Material estimation. In this experiment we apply our network to an inverse problem: material estimation from observation. We use our network to estimate the viscosity parameter of a fluid from the particle movement. To this end we generate 5 sequences with 100 frames and random viscosity between 0.01 and 0.3 for testing. Figure 6 shows the inference of our network on two of the test sequences. The training data contains 200 sequences and uses the same range for the viscosity parameter and does not contain the fluid shape of the test sequences. To estimate the viscosity we backpropagate through the trained network and optimize ν with gradient descent. Table 3 reports the results, which indicate that our network can be used to successfully distinguish different materials.

7 CONCLUSION

We have developed continuous convolutional networks for Lagrangian fluid simulation. We have introduced a simple formulation for continuous convolutions and demonstrated its accuracy and speed. Our model captures a wide range of complex material behavior, offers long-term stability, and generalizes to new situations such as varying particle counts, domain geometries, and material properties. There are numerous directions for future work, such as extending the framework to incorporate rigid and deformable solids. We will make the code publicly available to facilitate such developments.




Figure 6: We can control the viscosity of the simulated fluid at test time by changing the input parameter ν in the input feature vector ($\nu = 0.04$ at the top, and $\nu = 0.27$ below).

| | | | | | | |
|--------------------------|--------|--------|--------|-------|-------|-------------|
| GT viscosity | 0.044 | 0.127 | 0.174 | 0.233 | 0.269 | <i>Mean</i> |
| Avg. estimated viscosity | 0.027 | 0.150 | 0.202 | 0.255 | 0.277 | |
| Avg. relative error (%) | 38.377 | 18.542 | 15.604 | 9.568 | 3.235 | 17.075 |

Table 3: We estimate the viscosity ν by backpropagation through our network for all 5 scenes 10 times starting with a random initialization each time and report the average.

REFERENCES

- Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *ACM Trans. Graph.*, 37(4), 2018.
- George K. Batchelor. *An Introduction to Fluid Dynamics*. Cambridge University Press, 1967.
- Peter W. Battaglia, Razvan Pascanu, Matthew Lai, Danilo Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, 2016.
- Jan Bender and Dan Koschier. Divergence-free smoothed particle hydrodynamics. In *Symposium on Computer Animation*, 2015.
- Andrea Colagrossi and Maurizio Landrini. Numerical simulation of interfacial flows by smoothed particle hydrodynamics. *Journal of Computational Physics*, 191(2), 2003.
- Nick Foster and Dimitris N. Metaxas. Realistic animation of liquids. *Graphical Model and Image Processing*, 58(5), 1996.
- Robert A Gingold and Joseph J Monaghan. Smoothed particle hydrodynamics: Theory and application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society*, 181(3), 1977.
- Christoph Gissler, Andreas Peer, Stefan Band, Jan Bender, and Matthias Teschner. Interlinked SPH pressure solvers for strong fluid-rigid coupling. *ACM Trans. Graph.*, 38(1), 2018.
- Jens André Griepentrog, Wolfgang Höppner, Hans-Christoph Kaiser, and Joachim Rehberg. A bi-Lipschitz continuous, volume preserving map from the unit ball onto a cube. *Note di Matematica*, 28, 2008.
- Pedro Hermosilla, Tobias Ritschel, Pere-Pau Vázquez, Àlvar Vinacua, and Timo Ropinski. Monte Carlo convolution for learning on non-uniformly sampled point clouds. *ACM Trans. Graph.*, 37(6), 2018.
- Xiang Yu Hu and Nikolaus A Adams. A multi-phase SPH method for macroscopic and mesoscopic flows. *Journal of Computational Physics*, 213(2), 2006.
- Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Pointwise convolutional neural networks. In *CVPR*, 2018.
- Byungsoo Kim, Vinicius C. Azevedo, Nils Thuerey, Theodore Kim, Markus Gross, and Barbara Solenthaler. Deep fluids: A generative network for parameterized fluid simulations. *Computer Graphics Forum*, 38(2), 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- L’ubor Ladický, SoHyeon Jeong, Barbara Solenthaler, Marc Pollefeys, and Markus Gross. Data-driven fluid simulations using regression forests. *ACM Trans. Graph.*, 34(6), 2015.
- Huan Lei, Naveed Akhtar, and Ajmal Mian. Octree guided CNN with spherical kernels for 3D point clouds. In *CVPR*, 2019.
- Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution on X-transformed points. In *Advances in Neural Information Processing Systems*, 2018.

- Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B. Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *ICLR*, 2019.
- Julia Ling, Andrew Kurzwski, and Jeremy Templeton. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics*, 807, 2016.
- Miles Macklin and Matthias Müller. Position based fluids. *ACM Trans. Graph.*, 32(4), 2013.
- Miles Macklin, Matthias Müller, Nuttapong Chentanez, and Tae-Yong Kim. Unified particle physics for real-time applications. *ACM Trans. Graph.*, 33(4), 2014.
- J Monaghan. An introduction to SPH. *Computer Physics Communications*, 48(1), 1988.
- Jeremy Morton, Antony Jameson, Mykel J Kochenderfer, and Freddie Witherden. Deep dynamical modeling and control of unsteady fluid flows. In *Advances in Neural Information Processing Systems*, 2018.
- Damian Mrowca, Chengxu Zhuang, Elias Wang, Nick Haber, Li F Fei-Fei, Josh Tenenbaum, and Daniel L Yamins. Flexible neural representation for physics prediction. In *Advances in Neural Information Processing Systems*, 2018.
- Matthias Müller, David Charypar, and Markus Gross. Particle-based fluid simulation for interactive applications. In *Symposium on Computer Animation*, 2003.
- Daniel J Price. Smoothed particle hydrodynamics and magnetohydrodynamics. *Journal of Computational Physics*, 231(3), 2012.
- Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin A. Riedmiller, Raia Hadsell, and Peter W. Battaglia. Graph networks as learnable physics engines for inference and control. In *ICML*, 2018.
- Connor Schenck and Dieter Fox. SPNets: Differentiable fluid dynamics for deep neural networks. In *Conference on Robot Learning*, 2018.
- Barbara Solenthaler and Renato Pajarola. Predictive-corrective incompressible SPH. *ACM Trans. Graph.*, 28(3), 2009.
- Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. SPLATNet: Sparse lattice networks for point cloud processing. In *CVPR*, 2018.
- Matthias Teschner, Bruno Heidelberger, Matthias Müller, Danat Pomerantes, and Markus H. Gross. Optimized spatial hashing for collision detection of deformable objects. In *Vision, Modeling, and Visualization (VMV)*, 2003.
- Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, and Ken Perlin. Accelerating Eulerian fluid simulation with convolutional networks. In *ICML*, 2017.
- Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, and Raquel Urtasun. Deep parametric continuous convolutional neural networks. In *CVPR*, 2018.
- David C Wilcox. *Turbulence Modeling for CFD*. DCW industries, 3rd edition, 2006.
- Wenxuan Wu, Zhongang Qi, and Fuxin Li. PointConv: Deep convolutional networks on 3D point clouds. *arXiv:1811.07246*, 2018.
- You Xie, Erik Franz, Mengyu Chu, and Nils Thuerey. tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow. *ACM Trans. Graph.*, 37(4), 2018.
- Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. SpiderCNN: Deep learning on point sets with parameterized convolutional filters. In *ECCV*, 2018.

A APPENDIX

A.1 IMPLEMENTATION DETAILS

To accelerate the computation of Equation 7 we use existing general matrix multiplication primitives. Similar as for standard convolutions in deep learning frameworks we build a matrix with patches that we then multiply with the filter matrix. We can account for the linear interpolation by applying the interpolation to the patch matrix instead of the filters. For 3D point clouds a single point contributes to up to 8 voxels in a patch.

Another crucial part in our implementation is the nearest neighbor search. Since the positions of the fluid particles update with each timestep, we have to rebuild the neighborhood information after each update. We implement the neighborhood search with spatial hashing (we use the hash function proposed in Teschner et al. (2003)). We explicitly store all particle neighbors in a compact list, which allows us to reuse the information for multiple convolutions operating on the same point sets.

A.2 DATASET DETAILS

To enable generalization to new environments we use 10 different containers (see Figure 7) in our data generation process. During data generation we randomly sample an environment and place up to 3 fluid bodies with random initial velocities in the scene.




Figure 7: We use 10 different box-like scenes for training.