

FluidsFormer: A Transformer-Based Approach for Continuous Fluid Interpolation

BRUNO ROY, Autodesk Research, Canada




Fig. 1. Given input keyframes, our approach interpolates substeps of a fluid simulation – resulting in a smooth and realistic animation.

In this work, we introduce *FluidsFormer*: a transformer-based approach for fluid interpolation within a continuous-time framework. By combining the capabilities of PITT and a residual neural network (RNN), we analytically predict the physical properties of the fluid state. This enables us to interpolate substep frames between simulated keyframes, enhancing the temporal smoothness and sharpness of animations. We demonstrate promising results for smoke interpolation and conduct initial experiments on liquids.

ACM Reference Format:

Bruno Roy. 2024. FluidsFormer: A Transformer-Based Approach for Continuous Fluid Interpolation. In *Proceedings of (SIGGRAPH '24 Posters)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

As we advance into the era of generative AI, there has been a surge of interest in editing within the latent space, presenting a genuine challenge in providing controllable data-driven capabilities. Among many other areas in computer graphics, physics-based animation remains particularly challenging to edit as it relies on complex physics rules and principles that need to be satisfied for realism. Another challenge posed by these physics-based phenomena, particularly within the visual effects industry, is the need to adapt these principles to align with artistic direction, as observed in animation films. Striking a balance between realism and controllability poses difficulty in providing tools for editing natural phenomena such as fluids. For several decades, numerous researchers have endeavored to enhance

Author's Contact Information: Bruno Roy, bruno.roy@autodesk.com, Autodesk Research, Montreal, Canada.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

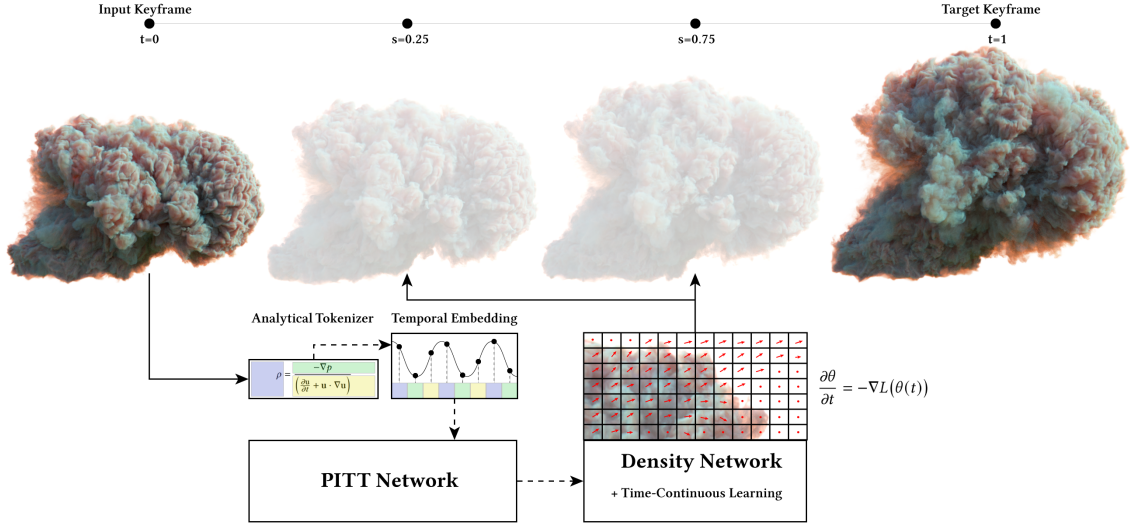


Fig. 2. Here is an overview of the *FluidsFormer* approach. A conventional Eulerian solver is used to produce keyframes, along with their respective volumetric properties, using a large timestep. We leverage most of the PITT network to tokenize and embed the partial differential equations (PDEs) needed to compute the density at any given time between the reference frames. Finally, our density network predicts the advected densities at substeps (e.g., $s = [0.25, 0.75]$).

the controllability and flexibility of fluid editing – transitioning from local editing of keyframes [Pan et al. 2013] to flow-based methods [Sato et al. 2018]. While the latter showed promise in terms of controllability, some explored optical flow-based approaches to interpolate Eulerian [Thuerey 2016] and particle-based fluids [Roy et al. 2021] as novel means of creating and editing such natural phenomena. Although also promising, these flow-based approaches remained highly dependent on numerical solvers, rendering them still fairly computationally expensive. More recently, data-driven methods have emerged to simulate and control fluids at a reduced cost. Introduced to computer graphics approximately a decade ago, [Ladický et al. 2015] proposed a novel approach to computing particle acceleration using a regression forest. Subsequently, a significant advancement was made by utilizing LSTM-based methods [Wiewel et al. 2019] to handle and compute pressure changes as sequential data. Similarly, other works were introduced to address the pressure projection step using CNNs [Tompson et al. 2017; Yang et al. 2016]. Techniques were also proposed to synthesize smoke from pre-computed patches [Chu and Thuerey 2017], generate super-resolution flows using GANs [Xie et al. 2018], and enhance diffusion behavior and liquid splashes [Um et al. 2018]. In recent years, methods have been introduced to improve the apparent resolution of smoke [Bai et al. 2020] and particle-based liquids [Roy et al. 2021]. Although our primary objective remains to enhance the controllability of fluid editing, our approach shares similarities with that of [Thuerey 2016] and [Roy et al. 2021] as we aim to interpolate fluids in a data-driven manner using the advection scheme of Eulerian simulations.

2 OUR METHOD

Although Transformer-based networks have been primarily introduced for natural language processing (NLP) and text generation, they offer interesting properties for sequential data in general. In our method, we propose to leverage an attention-based architecture within a continuous-time framework to learn and interpolate simulation properties per

frame in an approximate analytical manner. In the following sections, we will outline the details spanning from data preparation to network architecture and training using a transformer-based encoder-decoder network. We will also highlight a few concrete use cases to introduce novel ways of generating and editing fluids.

2.1 Data Preparation

The data preparation is divided into two main steps: (1) generating the temporal embeddings for the fluid element’s states and (2) handling the tokenization process in a physics-adapted context.

Physics-Adapted Tokenization. The tokenization process is performed by parsing and splitting the Navier-Stokes equation into components (see Eq. 1) – we intentionally omit the viscosity term to simplify the related operations.

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p \quad (1)$$

Temporal Embedding. We learn the latent embedding of the advection part of the governing equation using standard multi-head self-attention blocks from the PITT architecture [Lorsung et al. 2024]. That way, our model is capable of learning to interpolate physics properties analytically (e.g., densities ρ). As we do not consider the equilibrium equation (i.e., $\nabla \cdot \mathbf{u} = 0$), we handle the volume preservation part by simply penalizing in our loss function the solutions diverging too much from the reference.

2.2 Network Architecture

Our network architecture is composed of two stacked networks: (1) the pre-trained PITT network (transformer-based network) for solving the governing partial differential equation of the fluid dynamics and (2) the density network (RNN) to learn and predict the time-continuous density for the substeps between the input keyframes. We use the 18-layer architecture for our residual neural network as it gives us a decent performance (i.e., training/inference speed and accuracy) while reducing the requirement for an enormous dataset. Same as the ResNet original paper [He et al. 2016], our network is composed of one 7x7 convolution layer (including a 3x3 max pooling layer), 8 pairs of 3x3 convolution layers with respectively 64, 128, 256, and 512 kernels, and one last fully connected layer using a softmax activation function for average pooling (the rest of layers are using a ReLU activation function).

2.3 Training and Dataset

The density network is trained on normalized data $[-1, 1]$ outputted from the PITT network. During training, each simulation scenario is processed through the PITT network to generate the latent embedding of the non-viscous governing equation and provides an analytically-driven approximation to our density network to predict the correct density in space and time. To update the parameters, we use a Huber loss function L_δ minimizing a single term considering the possible outliers (with some regularization term): the difference in density between the ground truth ρ and the prediction $\hat{\rho}$.

$$L_\delta(\rho, \hat{\rho}) = \begin{cases} \frac{1}{2}(\rho - \hat{\rho})^2 & |\rho - \hat{\rho}| \leq \delta, \\ \delta|\rho - \hat{\rho}| - \frac{1}{2}\delta^2 & \text{otherwise.} \end{cases} \quad (2)$$

As the reference density is advected using the divergence-free velocity field, we already the volume preservation law.

We validated and tweaked the hyperparameters of our model using a 2D dataset of laminar and turbulent flows generated with OpenFOAM [Jasak 2009]. Then we generated volumetric data from Eulerian simulations using Bifrost.

The volumetric data points are constituted of a position (center of the cell), a velocity, and a density. The volumetric dataset is composed of 1000 (800 for training, 100 for validation, and 100 for testing) smoke inflow and emission simulations of 50 frames in length, without and with a single obstacle placed at random locations. As for most RNN architectures, we need a significant amount of data to properly generalize without dissipating the small-scale details in the simulation (e.g., second-order vorticity).

2.4 Continuous-Time Learning

Similarly to [Chen et al. 2024], our architecture uses a continuous-time multi-head attention module to transform time-varying sequence relationships into vectors of queries Q , keys K , and values V . The purpose of this module is to output a continuous dynamic flow evolving throughout the data points. However, as opposed to [Chen et al. 2024] and inspired by [Deleu et al. 2022], we formulate the learning algorithm to follow the input velocity during training – allowing our model to converge faster and to reflect the analytical framework as discussed in Sec. 2.1. The updated gradients are then used to update the parameters defining the fluid’s behavior. In other words, we train our model to evaluate the density based on the advection term and the velocities of the governing equation of the input system. An inherent advantage of analytically learning density advection is the flexibility to dynamically choose the discretization during the inference stage as required. Essentially, by employing the pre-trained PITT network, we evenly divide the time interval between two keyframes into a specified number of substeps S . Subsequently, we evaluate the density at these time points while considering the initial conditions. The density ρ is computed at location x by advecting the previous density with respect to the input velocity.

3 VARIOUS APPLICATIONS

Eulerian Fluids Interpolation. Our main goal with this approach is to propose a continuous-time transformer model capable of learning the underlying dynamics of fluid systems for interpolation purposes. The idea behind interpolating fluids is to generate a visually appealing and temporally smooth animation using only a few keyframes at large timesteps. Our interpolation methods will fill in between similarly to simulate the substeps between the provided reference keyframes (as shown in Fig. 1).

Generating using Variants. In this last use case, we take advantage of the generated tree structure to combine keyframes using Boolean operations such as addition, subtraction, and intersection. Using the volumetric data (i.e., properly stored in grid cells), we can mix multiple keyframes into a single target in our approach and produce a completely new animation.

Tree-Based Variants. For this use case, for each frame, we output multiple probable solutions using top-k sampling along with diverse Beam Search on the decoder side to encourage diversity in the generated sequences. From this set of solutions, we build a tree structure that allows us to branch out at any node to produce variants of a single simulation while preserving the initial conditions. We also performed a few early experiments combining our approach to an explicit solver for liquid simulations. As opposed to other presented use cases, we tested our

Manuscript submitted to ACM

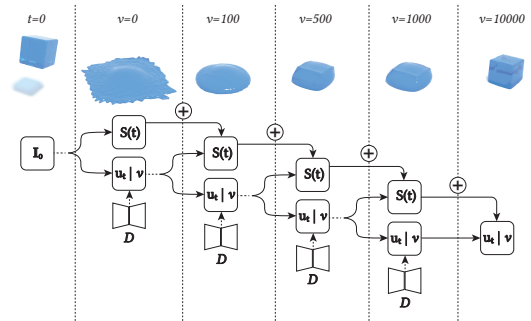


Fig. 3. Combining *FluidsFormer* with an explicit solver to generate viscosity variations for liquids.

approach to interpolate between various viscosity states. As shown in Fig. 3, we use our approach to produce five (5) variations of the same simulation but using different viscosity values ν (0 being the less viscous and 10000 the most). Starting with the same initial conditions I_0 , we branch a new variation of the current state of the velocity by predicting the next sequence of velocities according to a certain viscosity threshold (e.g., $\nu \in \{0, 100, 500, 1000, 10000\}$). To learn and generate these sequences based on the current state, we have trained a viscosity network D (i.e., including the viscosity term in the PITT embedding as input to a residual neural network) to match similarities between fluid characteristics (e.g., viscosity) and their corresponding velocities. Between each reference keyframe generated by the explicit solver, our network D interpolates viscosities to generate the substeps which are guided by the computed velocity field.

4 CONCLUSION

While our approach still relies on a coarse numerical simulation, we are confident that it introduces novel and less-linear ways of interacting with fluids. In future work, we aim to explore the accuracy of employing Transformer-based networks like PITT to replace conventional numerical solvers for simulating natural phenomena in visual effects.

REFERENCES

- Kai Bai, Wei Li, Mathieu Desbrun, and Xiaopei Liu. 2020. Dynamic upsampling of smoke through dictionary-based learning. *ACM Transactions on Graphics (TOG)* 40, 1 (2020), 1–19.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. 2024. ContiFormer: Continuous-time transformer for irregular time series modeling. *Advances in Neural Information Processing Systems* 36 (2024).
- Mengyu Chu and Nils Thuerey. 2017. Data-driven synthesis of smoke flows with CNN-based feature descriptors. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14.
- Tristan Deleu, David Kanaa, Leo Feng, Giancarlo Kerg, Yoshua Bengio, Guillaume Lajoie, and Pierre-Luc Bacon. 2022. Continuous-time meta-learning with forward mode differentiation. *arXiv preprint arXiv:2203.01443* (2022).
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- Hrvoje Jasak. 2009. OpenFOAM: Open source CFD in research and industry. *International Journal of Naval Architecture and Ocean Engineering* 1, 2 (2009), 89–94.
- L’ubor Ladický, SoHyeon Jeong, Barbara Solenthaler, Marc Pollefeys, and Markus Gross. 2015. Data-driven fluid simulations using regression forests. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–9.
- Cooper Lorsche, Zijie Li, and Amir Barati Farimani. 2024. Physics Informed Token Transformer for Solving Partial Differential Equations. *Machine Learning: Science and Technology* (2024).
- Zherong Pan, Jin Huang, Yiyong Tong, Changxi Zheng, and Hujun Bao. 2013. Interactive localized liquid motion editing. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10.
- Bruno Roy, Pierre Poulin, and Eric Paquette. 2021. Neural upflow: A scene flow learning approach to increase the apparent resolution of particle-based liquids. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 4, 3 (2021), 1–26.
- Syuhai Sato, Yoshinori Dobashi, and Tomoyuki Nishita. 2018. Editing fluid animation using flow interpolation. *ACM Transactions on Graphics (TOG)* 37, 5 (2018), 1–12.
- Nils Thuerey. 2016. Interpolations of smoke and liquid simulations. *ACM Transactions on Graphics (TOG)* 36, 1 (2016), 1–16.
- Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, and Ken Perlin. 2017. Accelerating eulerian fluid simulation with convolutional networks. In *International Conference on Machine Learning*. PMLR, 3424–3433.
- Kiwon Um, Xiangyu Hu, and Nils Thuerey. 2018. Liquid splash modeling with neural networks. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 171–182.
- Steffen Wiewel, Moritz Becher, and Nils Thuerey. 2019. Latent space physics: Towards learning the temporal evolution of fluid flow. In *Computer graphics forum*, Vol. 38. Wiley Online Library, 71–82.
- You Xie, Erik Franz, Mengyu Chu, and Nils Thuerey. 2018. tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.
- Cheng Yang, Xubo Yang, and Xiangyun Xiao. 2016. Data-driven projection method in fluid simulation. *Computer Animation and Virtual Worlds* 27, 3-4 (2016), 415–424.