ICML
International Conference
On Machine Learning

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Code repository

NeuroFluid: Fluid Dynamics Grounding with Particle-Driven Neural Radiance Fields


Shanyan Guan


Huayu Deng

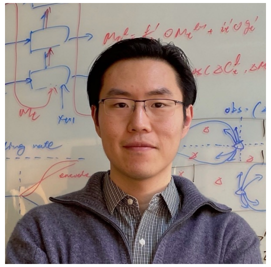
Yunbo Wang


Xiaokang Yang

Correspondence to: Yunbo Wang (yunbow@sjtu.edu.cn)

Fluid Dynamics Grounding

Inferring the physical dynamics of fluids from visual observations

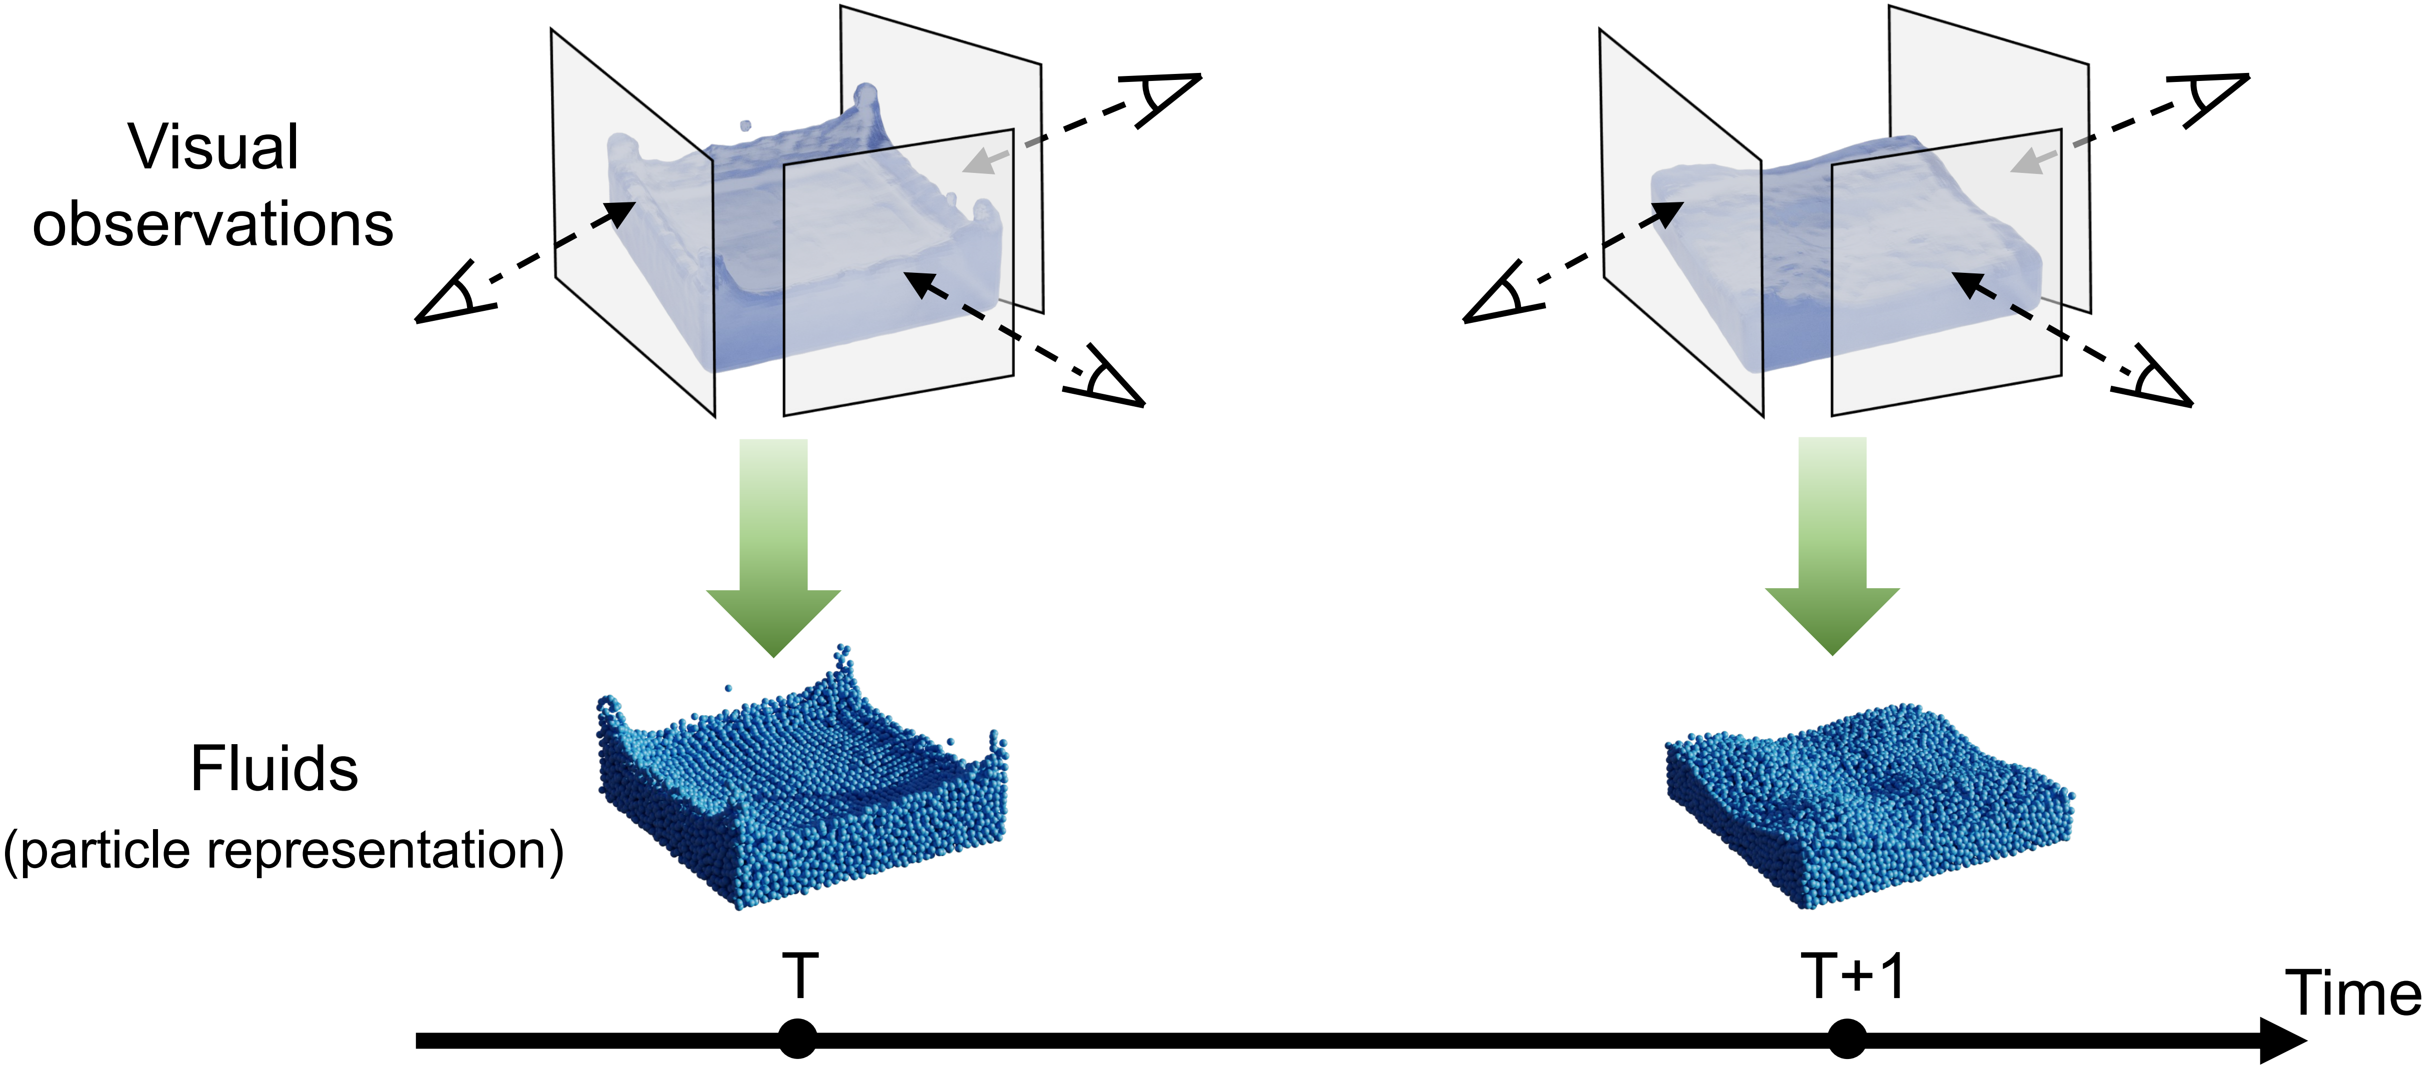

Progress in learning fluid dynamics

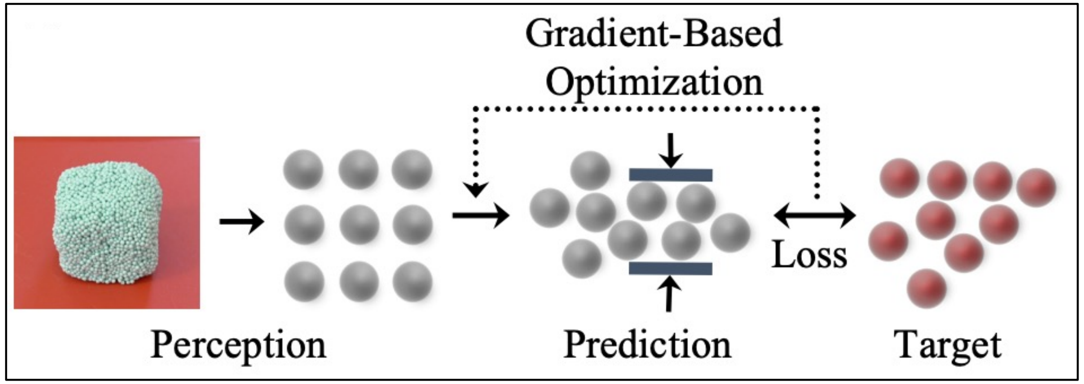

DPI-Net, Li, et al. [ICRL 2019]

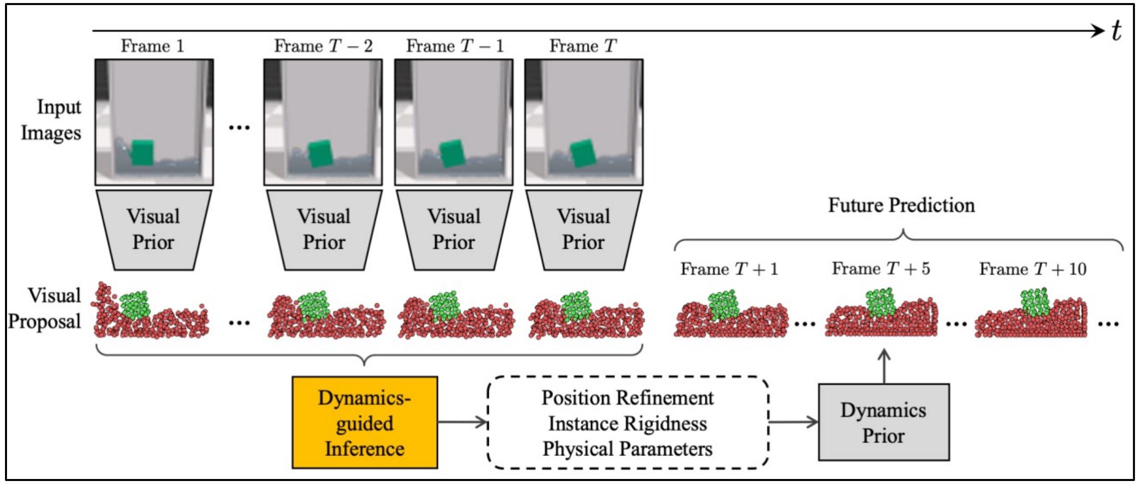

VGPL, Li, et al. [ICML 2020]

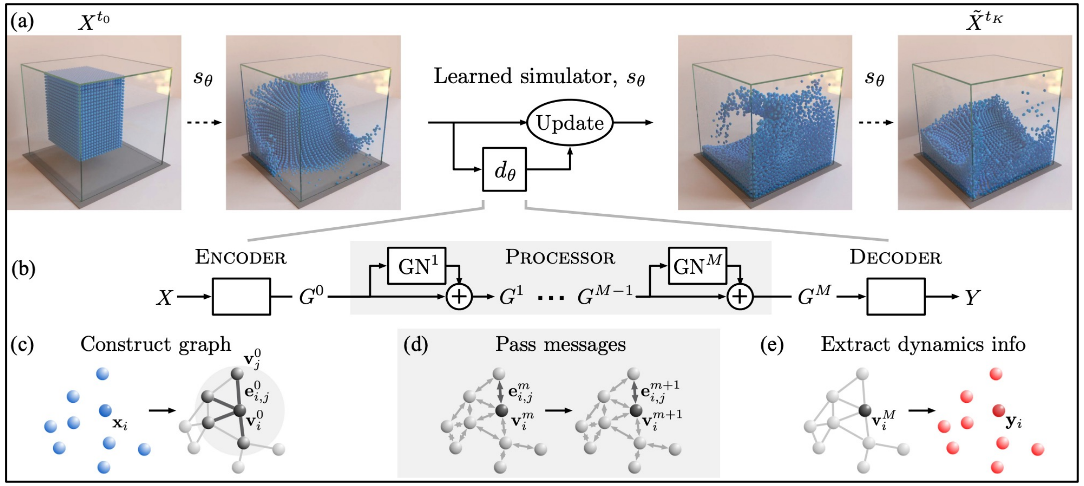

GNS, Sanchez-Gonzalez, et al. [ICML 2020]

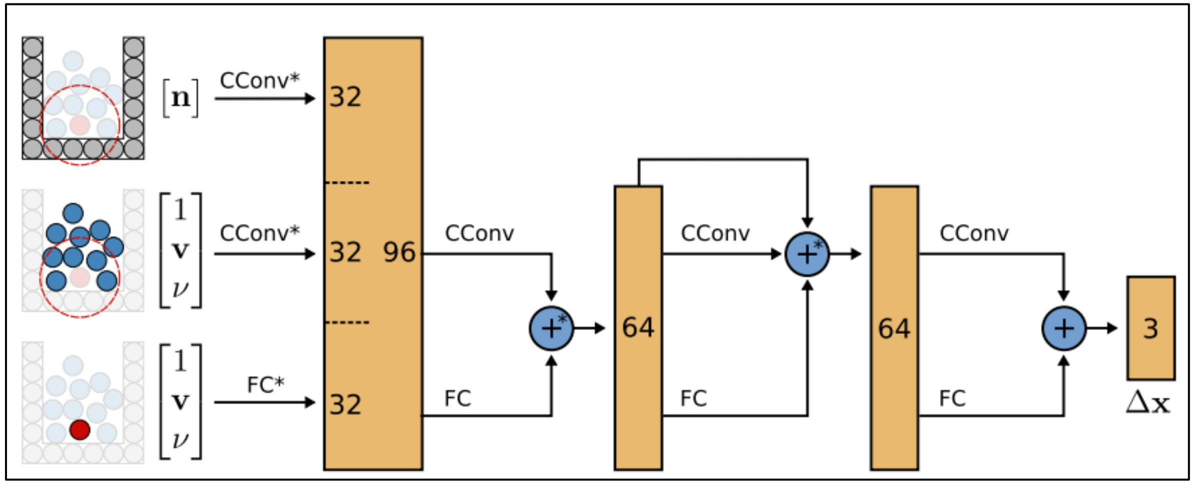

DLF, Ummenhofer, et al. [ICRL 2020]

Open Question

Whether neural networks can infer fluid dynamics without true particle supervision?

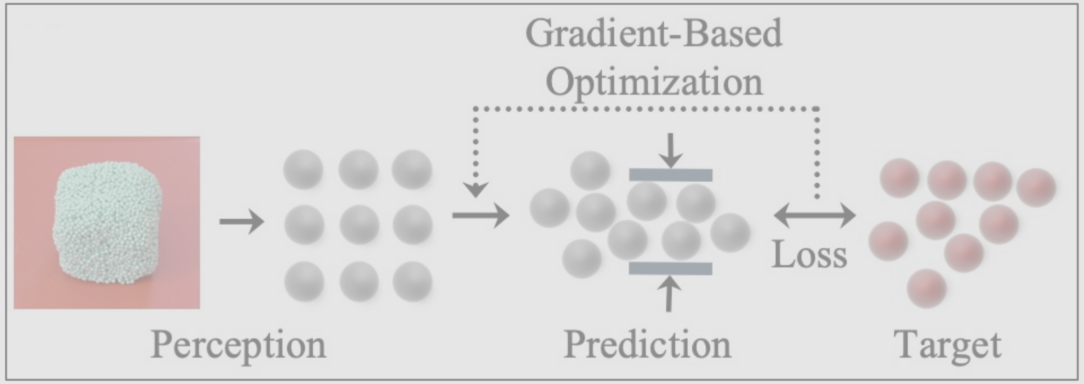

DPI-Net, Li, et al. [ICLR 2019]

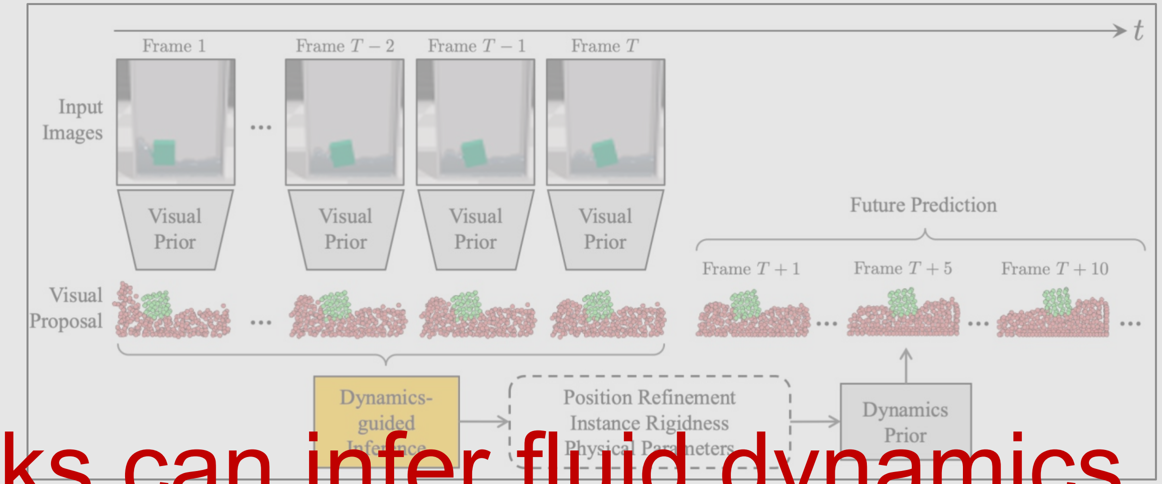

VGPL, Li, et al. [ICML 2020]

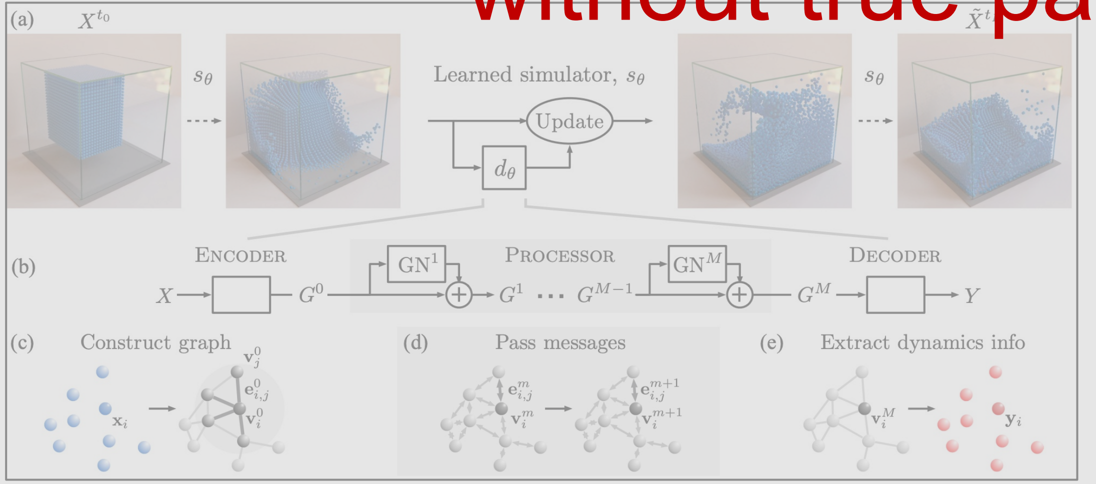

Sanchez-Gonzalez, et al. [ICML 2020]

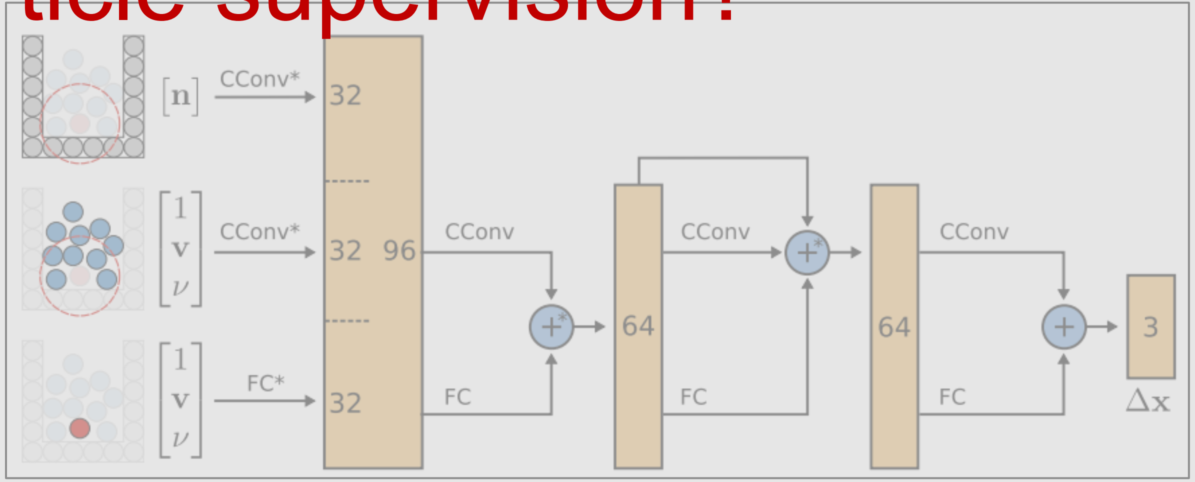

Ummenhofer, et al. [ICRL 2020]

NeuroFluid: A Fully Differentiable Framework

Inferring fluid dynamics only using the supervision of visual observation.

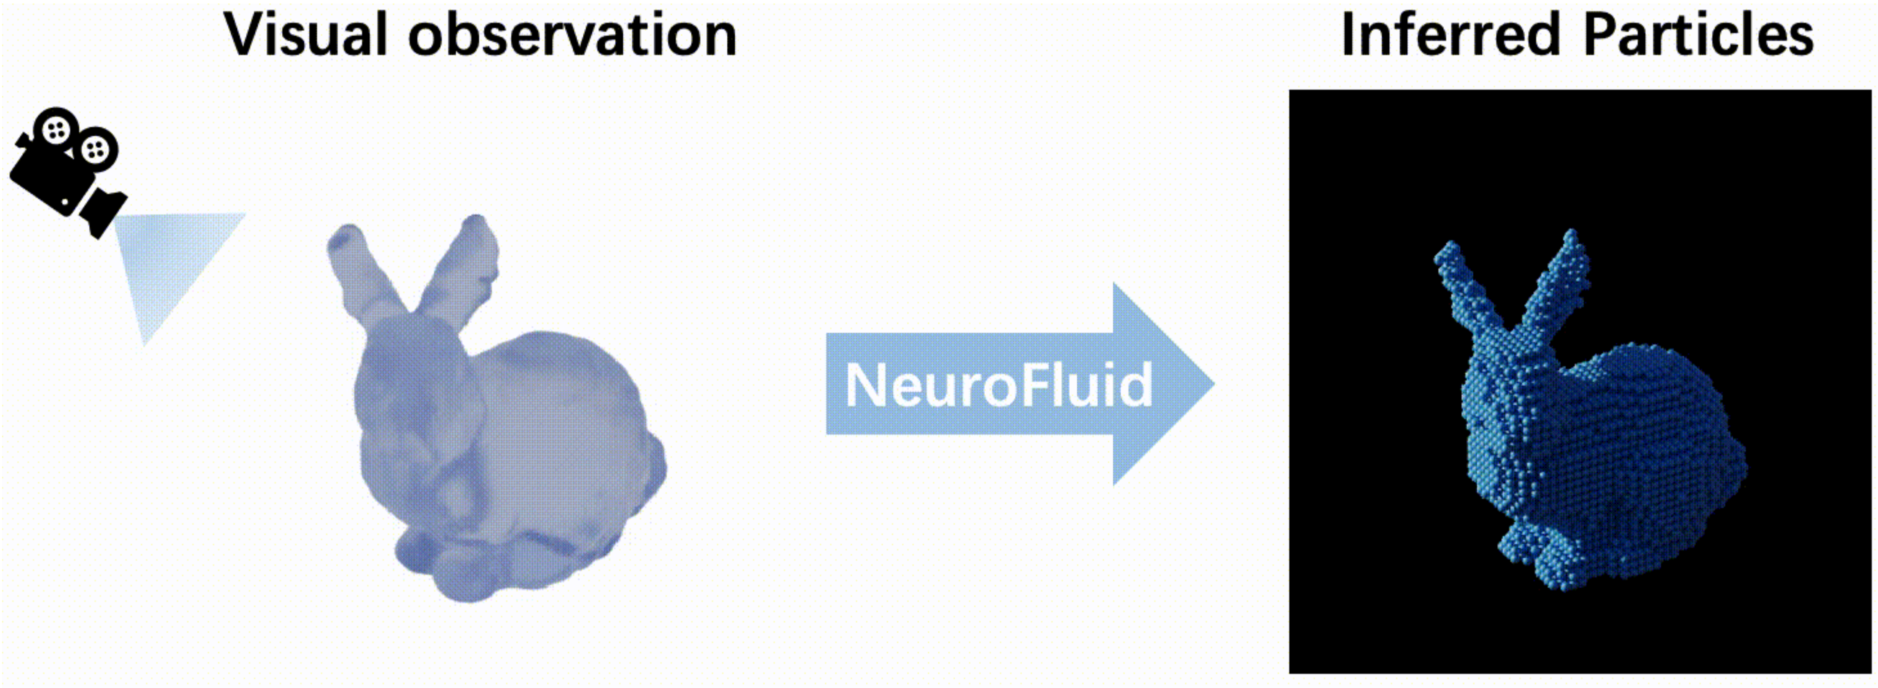

NeuroFluid: A Fully Differentiable Framework

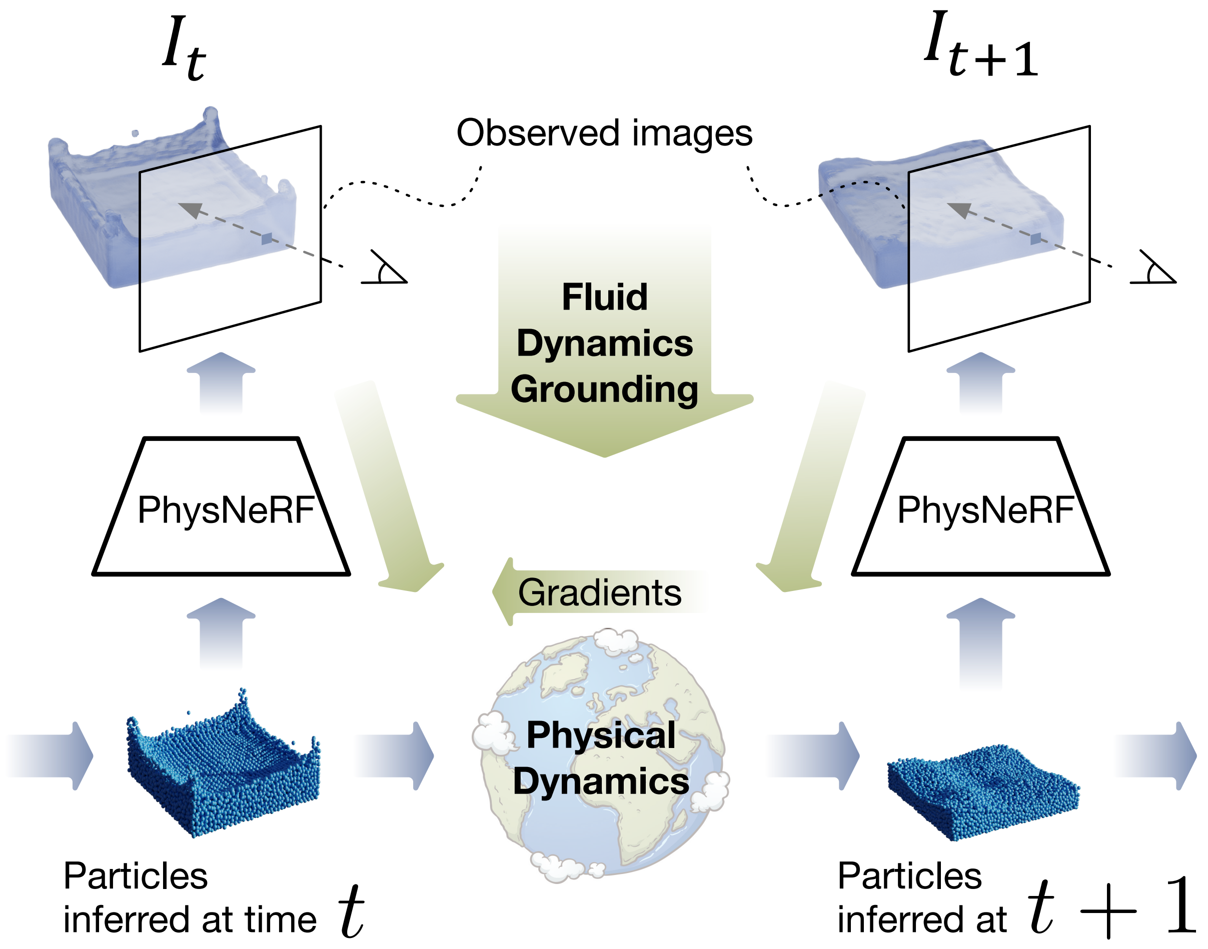

Consists of
(1) a particle transition model T_θ;
(2) a particle-driven renderer R_ϕ.

NeuroFluid: A Fully Differentiable Framework

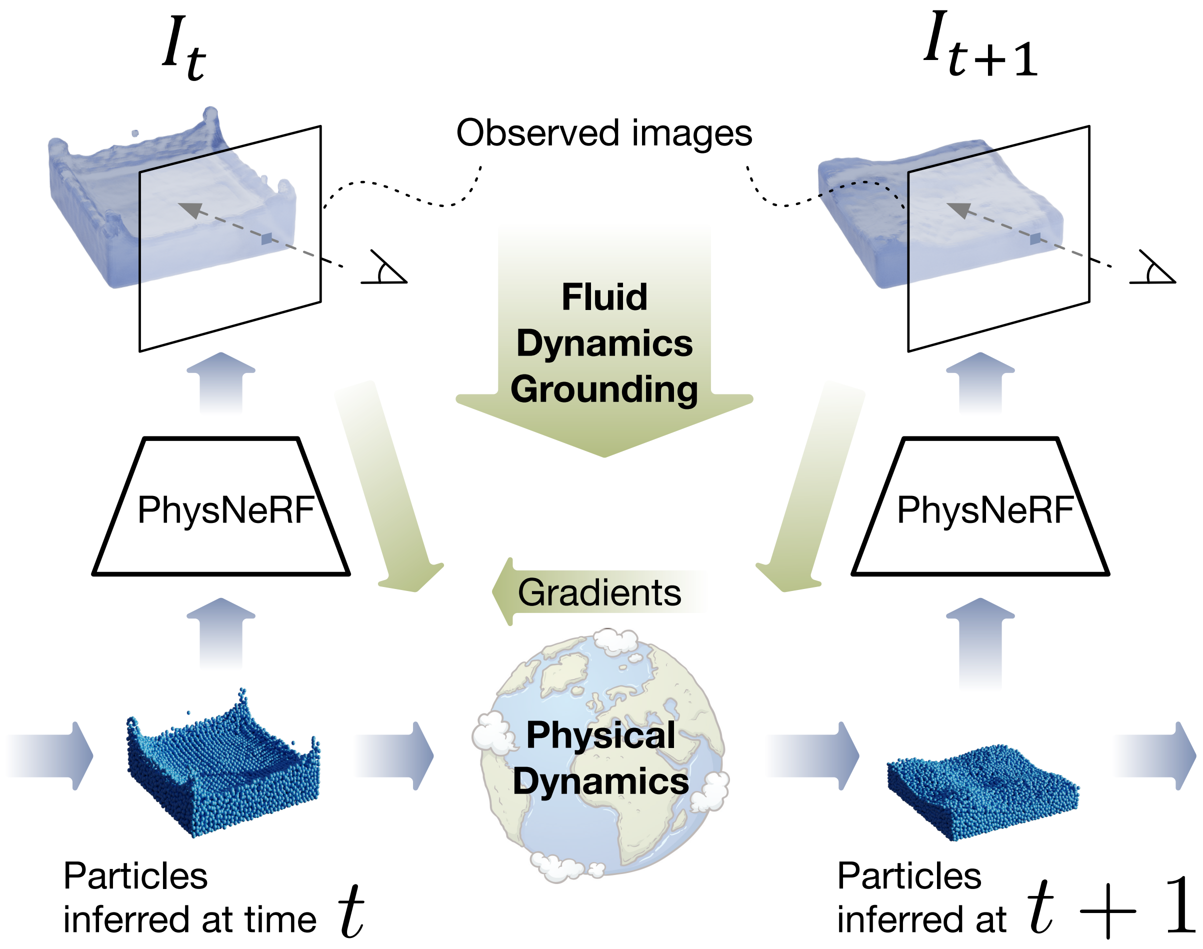

Consists of
(1) a particle transition model T_θ;
(2) a particle-driven renderer R_ϕ.

Jointly optimizing them as:
(1) Transition: $\boldsymbol{s}_{t+1} \longleftarrow T_\theta(\boldsymbol{s}_t)$, where $\boldsymbol{s}$ is particle positions and velocities.
(2) Rendering: $\hat{I}_{t+1} \longleftarrow R_\phi(\boldsymbol{s}_{t+1}, \boldsymbol{d})$
(3) Contrasting: $\left\|\hat{I}_{t+1} - I_{t+1}\right\|$, then backward.

NeuroFluid: A Fully Differentiable Framework

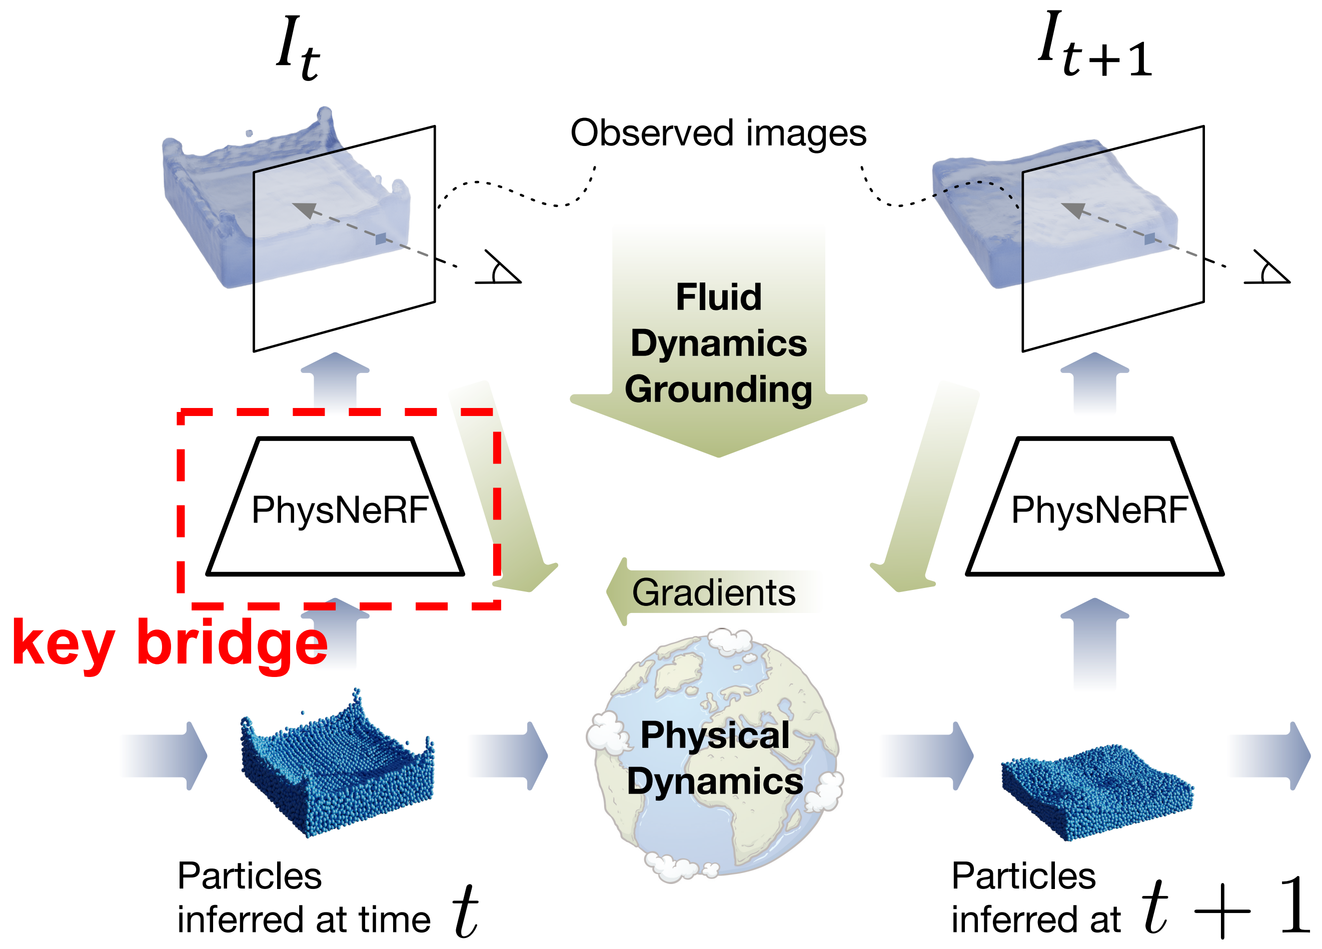

Consists of
(1) a particle transition model T_θ;
(2) a particle-driven renderer R_ϕ.

Jointly optimizing them as:
(1) Transition: $\boldsymbol{s}_{t+1} \longleftarrow T_\theta(\boldsymbol{s}_t)$, where $\boldsymbol{s}$ is particle positions and velocities.
(2) Rendering: $\hat{I}_{t+1} \longleftarrow R_\phi(\boldsymbol{s}_{t+1}, \boldsymbol{d})$
(3) Contrasting: $\|\hat{I}_{t+1} - I_{t+1}\|$, then backward.

PhysNeRF: Particle-Driven Neural Radiance Fields

Linking Neural Radiance Fields with physical particles.

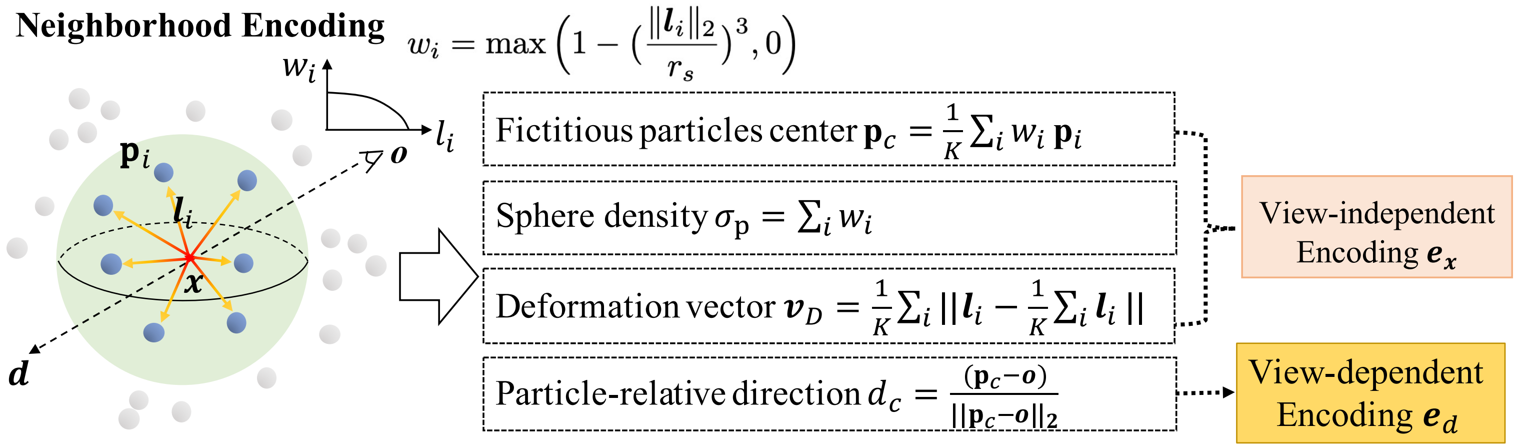

(1) Extracting geometry properties of physical point inside the spherical neighborhood of a sample ray point at $\boldsymbol{x}$.

PhysNeRF: Particle-Driven Neural Radiance Fields

Linking Neural Radiance Fields with physical particles.

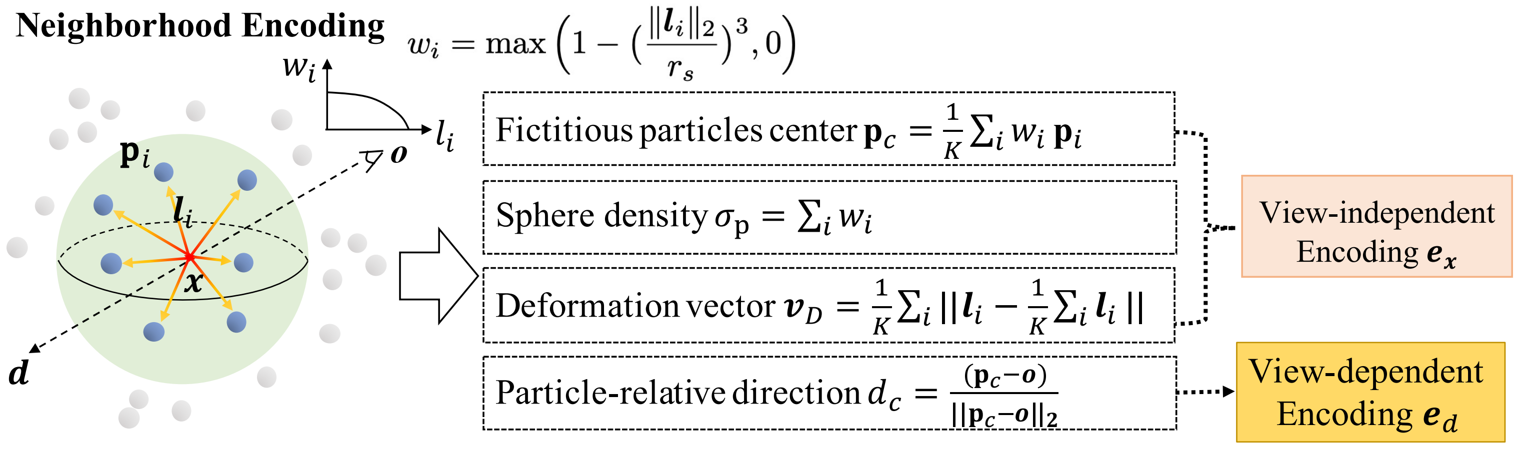

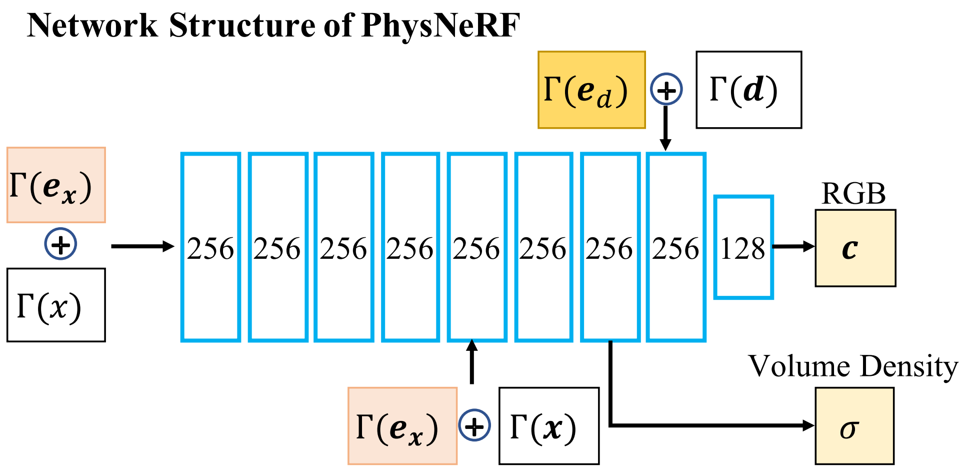

(1) Extracting geometry properties of physical point inside the spherical neighborhood of a sample ray point at $\boldsymbol{x}$.

(2) Predicting RGB value and volume density the point at $\boldsymbol{x}$.

Particle Transition model

$$\boldsymbol{P}_0, \boldsymbol{V}_0 \longrightarrow \boldsymbol{P}_1, \boldsymbol{V}_1 \longrightarrow \cdots \longrightarrow \boldsymbol{P}_T, \boldsymbol{V}_T$$

$\boldsymbol{P}_t$: particle positions
$\boldsymbol{V}_t$: particle velocity

Particle Transition model

$$\boldsymbol{P}_0, \boldsymbol{V}_0 \longrightarrow \boldsymbol{P}_1, \boldsymbol{V}_1 \longrightarrow \cdots \longrightarrow \boldsymbol{P}_T, \boldsymbol{V}_T$$

$\boldsymbol{P}_t$: particle positions
$\boldsymbol{V}_t$: particle velocity

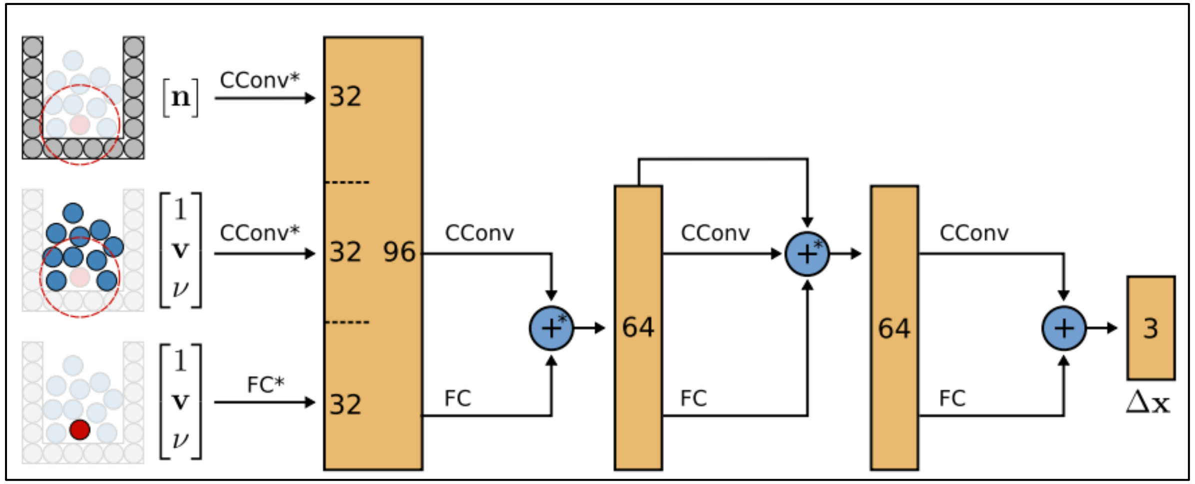

Ummenhofer, et al. Lagrangian fluid simulation with continuous convolutions. In ICLR, 2020.

Results

Fluid dynamics grounding is evaluated from

(1) Accuracy of grounded particle position
(2) Accuracy of predicted particle position
(3) Novel view synthesis.

Results

Fluid dynamics grounding is evaluated from

(1) Accuracy of grounded particle position
(2) Accuracy of predicted particle position
(3) Novel view synthesis.

Table 1. Typical geometric and physical properties of fluids on the evaluation benchmarks, which are closely related to the simulation and rendering of dynamic scenes. On "WaterBunny", we evaluate the generalization ability of PhysNeRF to novel particle distributions.

BENCHMARK	INITIAL SHAPE	MATERIAL	VISCOSITY	DENSITY (KG/M^3)
HONEYCONE	CONE	PRINCIPLED BSDF	0.8	1420
WATERCUBE	CUBE	GLASS BSDF	0.08	1000
WATERSPHERE	SPHERE	GLASS BSDF	0.08	1000
WATERBUNNY	STANFORDBUNNY	GLASS BSDF	0.08	1000

Results of Fluid dynamics Grounding

Compared models
(1) DLF: it has the same network structure as NeuroFluid.
(2) DLF† : it is finetuned with true particle state in the evaluation benchmarks.

Table 2. Quantitative results on the errors of fluid dynamics grounding ($t < 50$) and prediction ($50 \leq t < 60$), which are calculated between the grounded/predicted particle positions and the ground truth provided by the fluid simulator (lower is better). For **DLF**†, the transition model is finetuned on the testing benchmarks in a fully supervised way, that is, using **true** particle positions at $t < 50$.

Method	WaterCube				WaterSphere				HoneyCone			
	Grounding		Prediction		Grounding		Prediction		Grounding		Prediction	
	$d^{\text{AVG}}_{t<50}$	$d_{t=49}$	$d^{\text{AVG}}_{t\geq 50}$	$d_{t=59}$	$d^{\text{AVG}}_{t<50}$	$d_{t=49}$	$d^{\text{AVG}}_{t\geq 50}$	$d_{t=59}$	$d^{\text{AVG}}_{t<50}$	$d_{t=49}$	$d^{\text{AVG}}_{t\geq 50}$	$d_{t=59}$
DLF	32.3	48.3	47.4	46.2	32.2	47.6	48.1	45.9	61.5	83.5	69.7	**57.8**
NeuroFluid	**28.8**	**34.9**	**35.5**	**36.7**	**31.1**	**31.5**	**30.7**	**30.4**	**30.9**	**47.5**	**54.2**	58.2
DLF†	28.1	28.1	30.9	34.4	30.0	28.5	30.0	31.8	34.3	66.1	72.6	77.6

Results of Fluid Dynamics Prediction

Compared models

(1) DLF: it has the same network structure as NeuroFluid.

(2) DLF† : it is finetuned with true particle state in the evaluation benchmarks.

Table 2. Quantitative results on the errors of fluid dynamics grounding ($t < 50$) and prediction ($50 \leq t < 60$), which are calculated between the grounded/predicted particle positions and the ground truth provided by the fluid simulator (lower is better). For **DLF**†, the transition model is finetuned on the testing benchmarks in a fully supervised way, that is, using **true** particle positions at $t < 50$.

METHOD	WATERCUBE				WATERSPHERE				HONEYCONE			
	GROUNDING		PREDICTION		GROUNDING		PREDICTION		GROUNDING		PREDICTION	
	$d_{t<50}^{\text{AVG}}$	$d_{t=49}$	$d_{t\geq 50}^{\text{AVG}}$	$d_{t=59}$	$d_{t<50}^{\text{AVG}}$	$d_{t=49}$	$d_{t\geq 50}^{\text{AVG}}$	$d_{t=59}$	$d_{t<50}^{\text{AVG}}$	$d_{t=49}$	$d_{t\geq 50}^{\text{AVG}}$	$d_{t=59}$
DLF	32.3	48.3	47.4	46.2	32.2	47.6	48.1	45.9	61.5	83.5	69.7	**57.8**
NEUROFLUID	**28.8**	**34.9**	**35.5**	**36.7**	**31.1**	**31.5**	**30.7**	**30.4**	**30.9**	**47.5**	**54.2**	58.2
DLF†	28.1	28.1	30.9	34.4	30.0	28.5	30.0	31.8	34.3	66.1	72.6	77.6

Qualitative Results of Fluid dynamics Grounding and Prediction

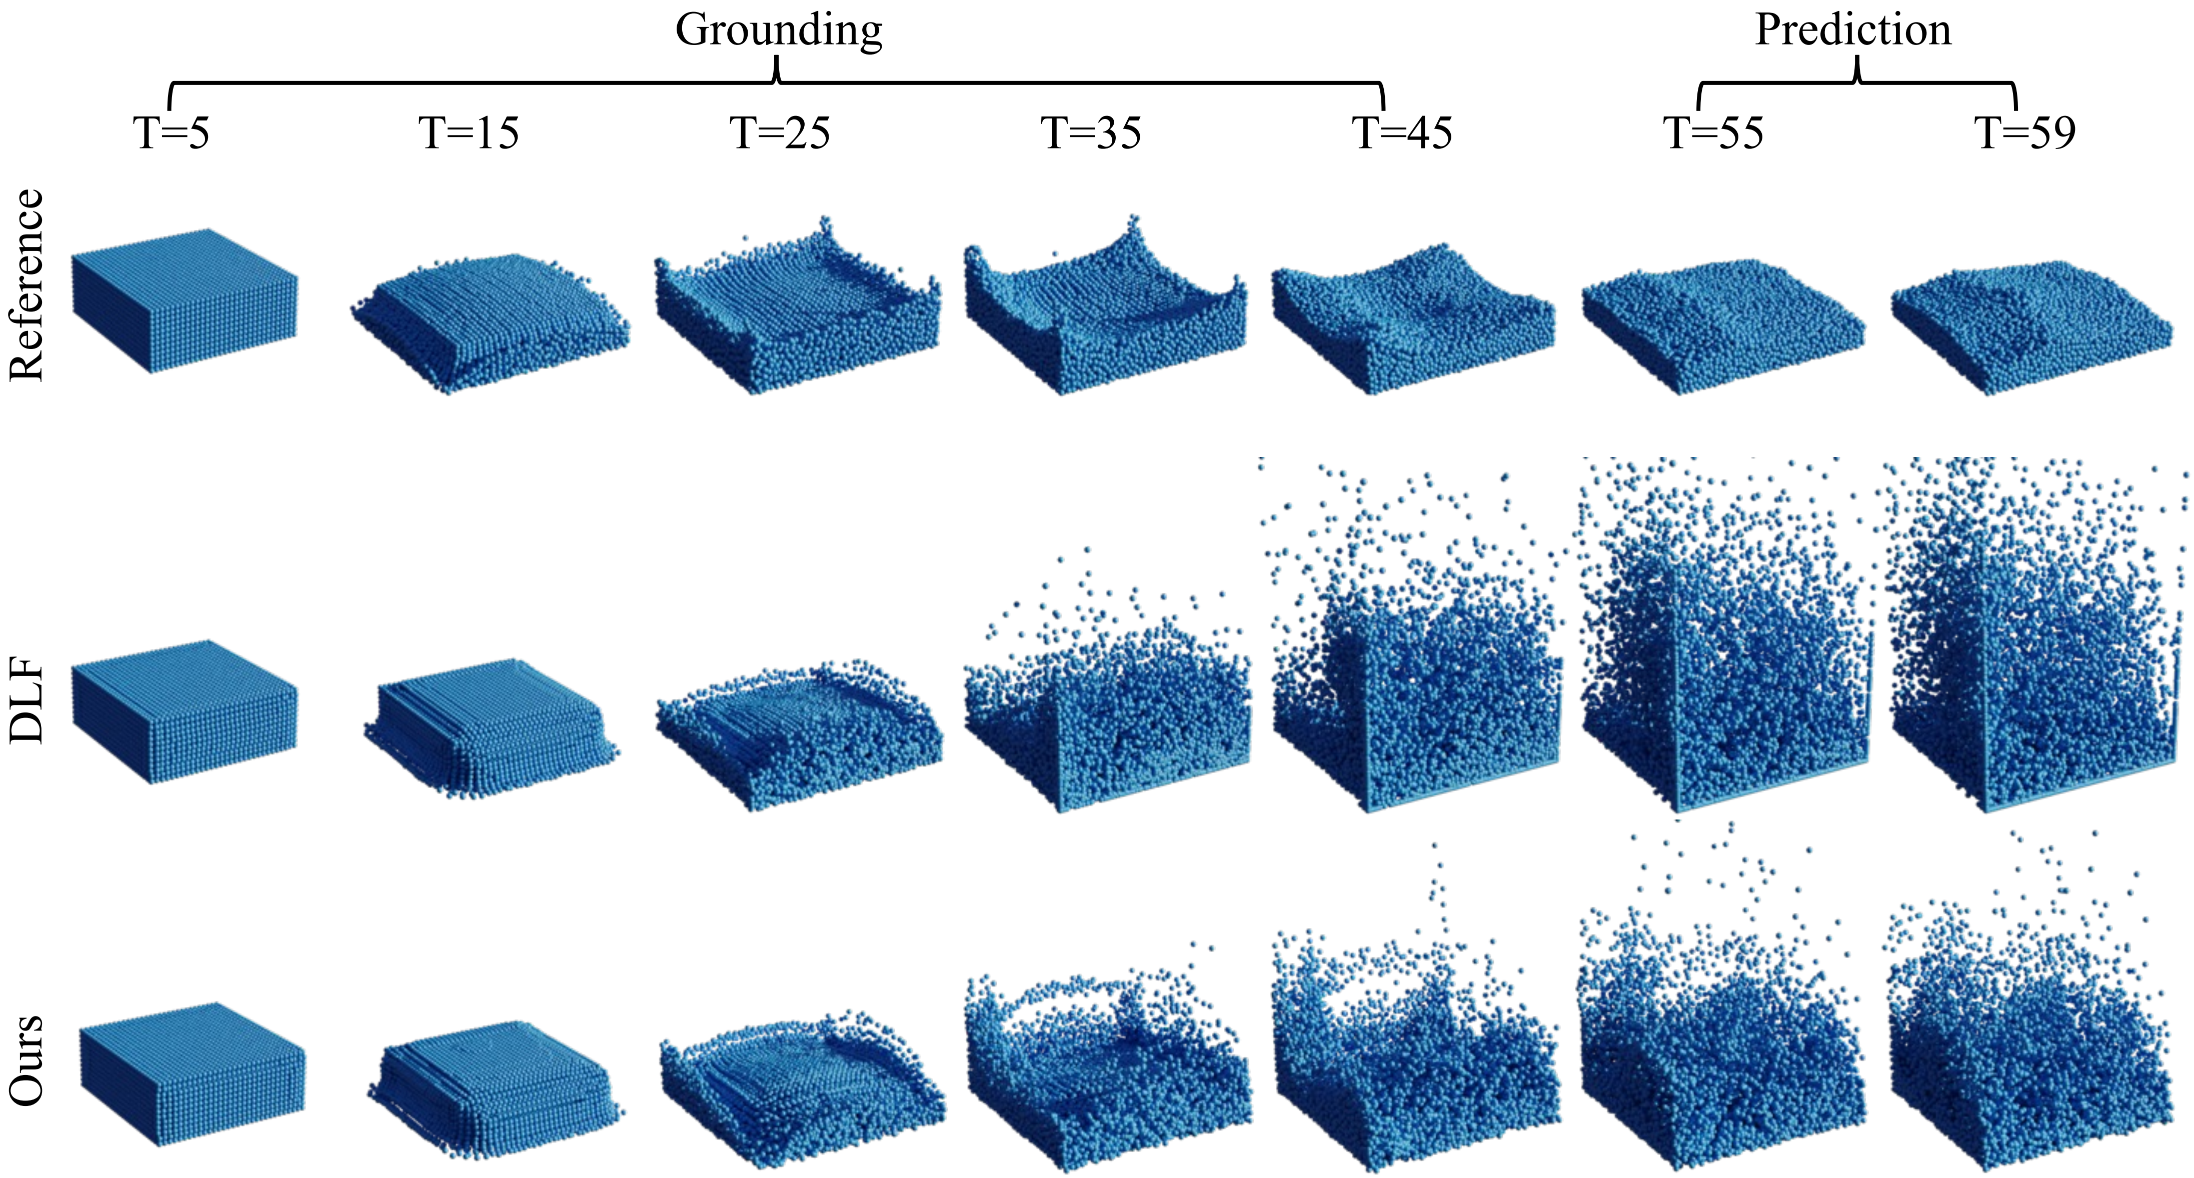

Results of Novel View Synthesis

NeRF-based comparisons: (1) **D-NeRF** (Pumarola et al., 2021), (2) **NeRF-T** (NeRF+time index), and (3) the 3D-aware fluid renderer from **Li et al. (2022)**:

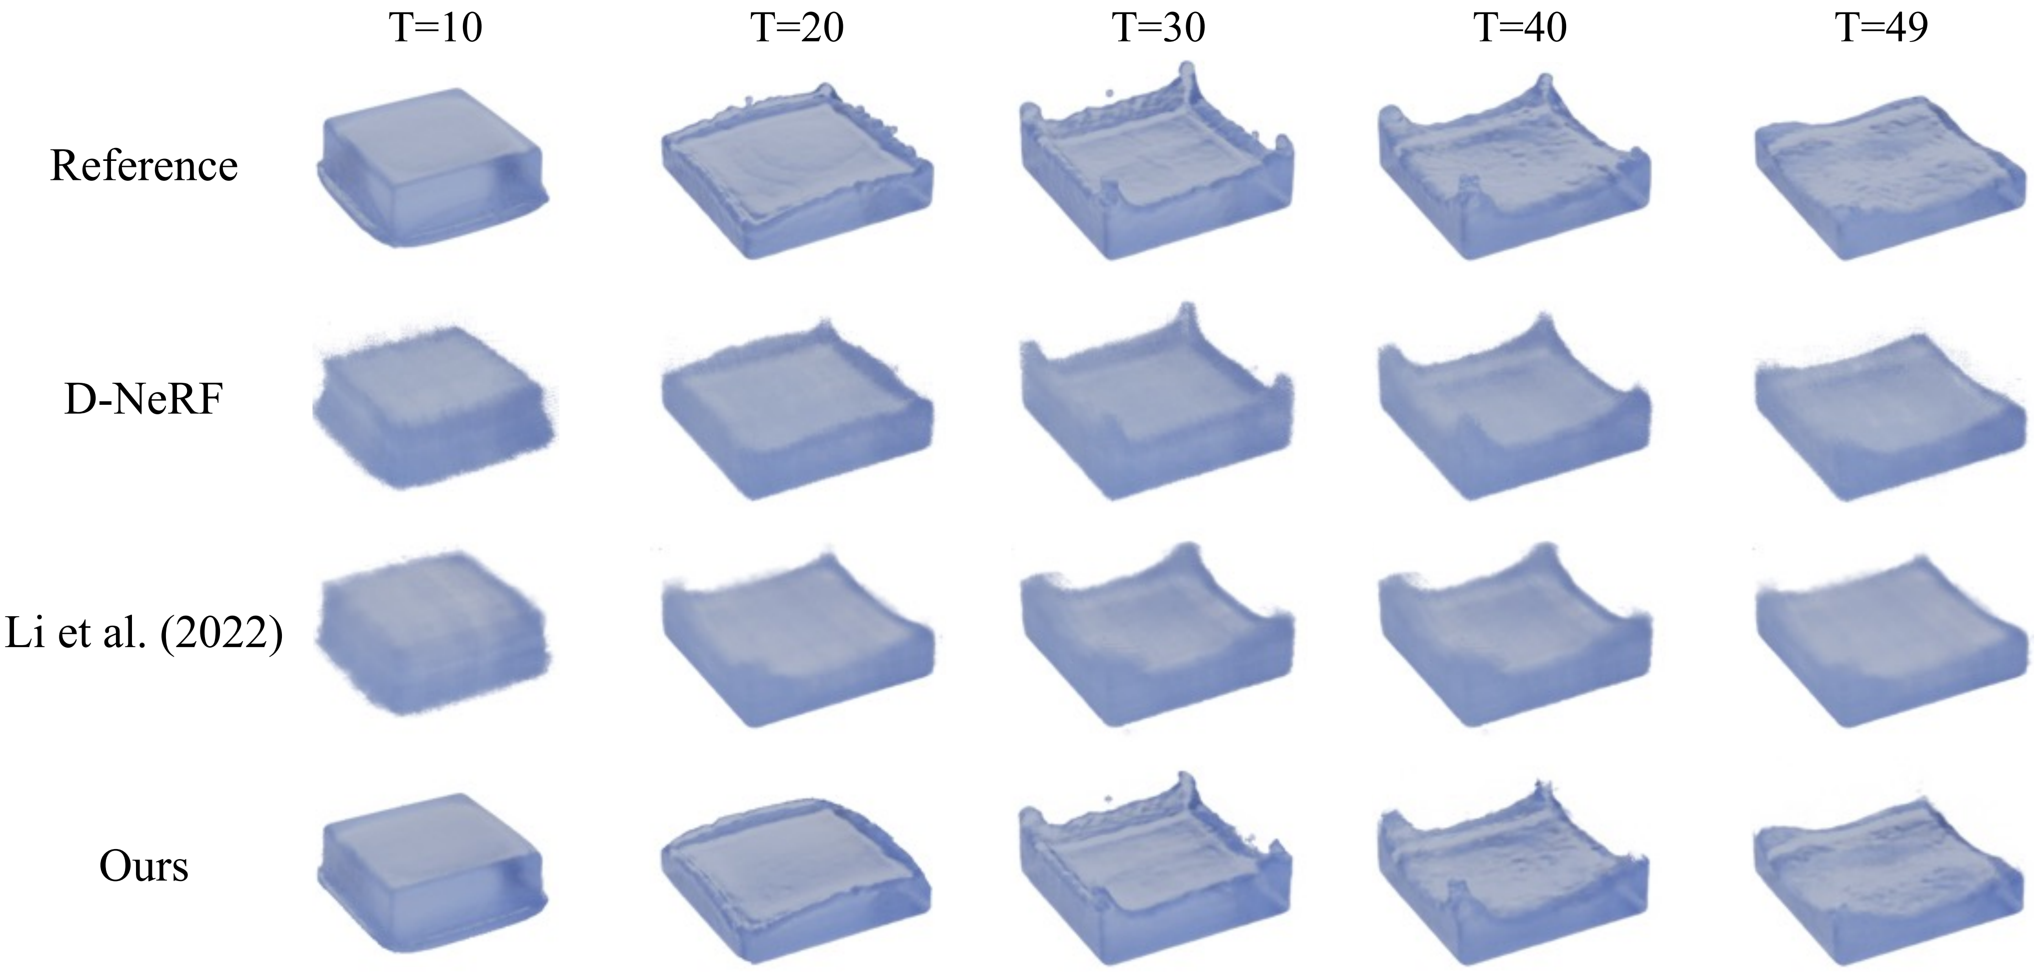

Results of Novel View Synthesis

NeRF-based comparisons: (1) **D-NeRF** (Pumarola et al., 2021), (2) **NeRF-T** (NeRF+time index), and (3) the 3D-aware fluid renderer from **Li et al. (2022)**:

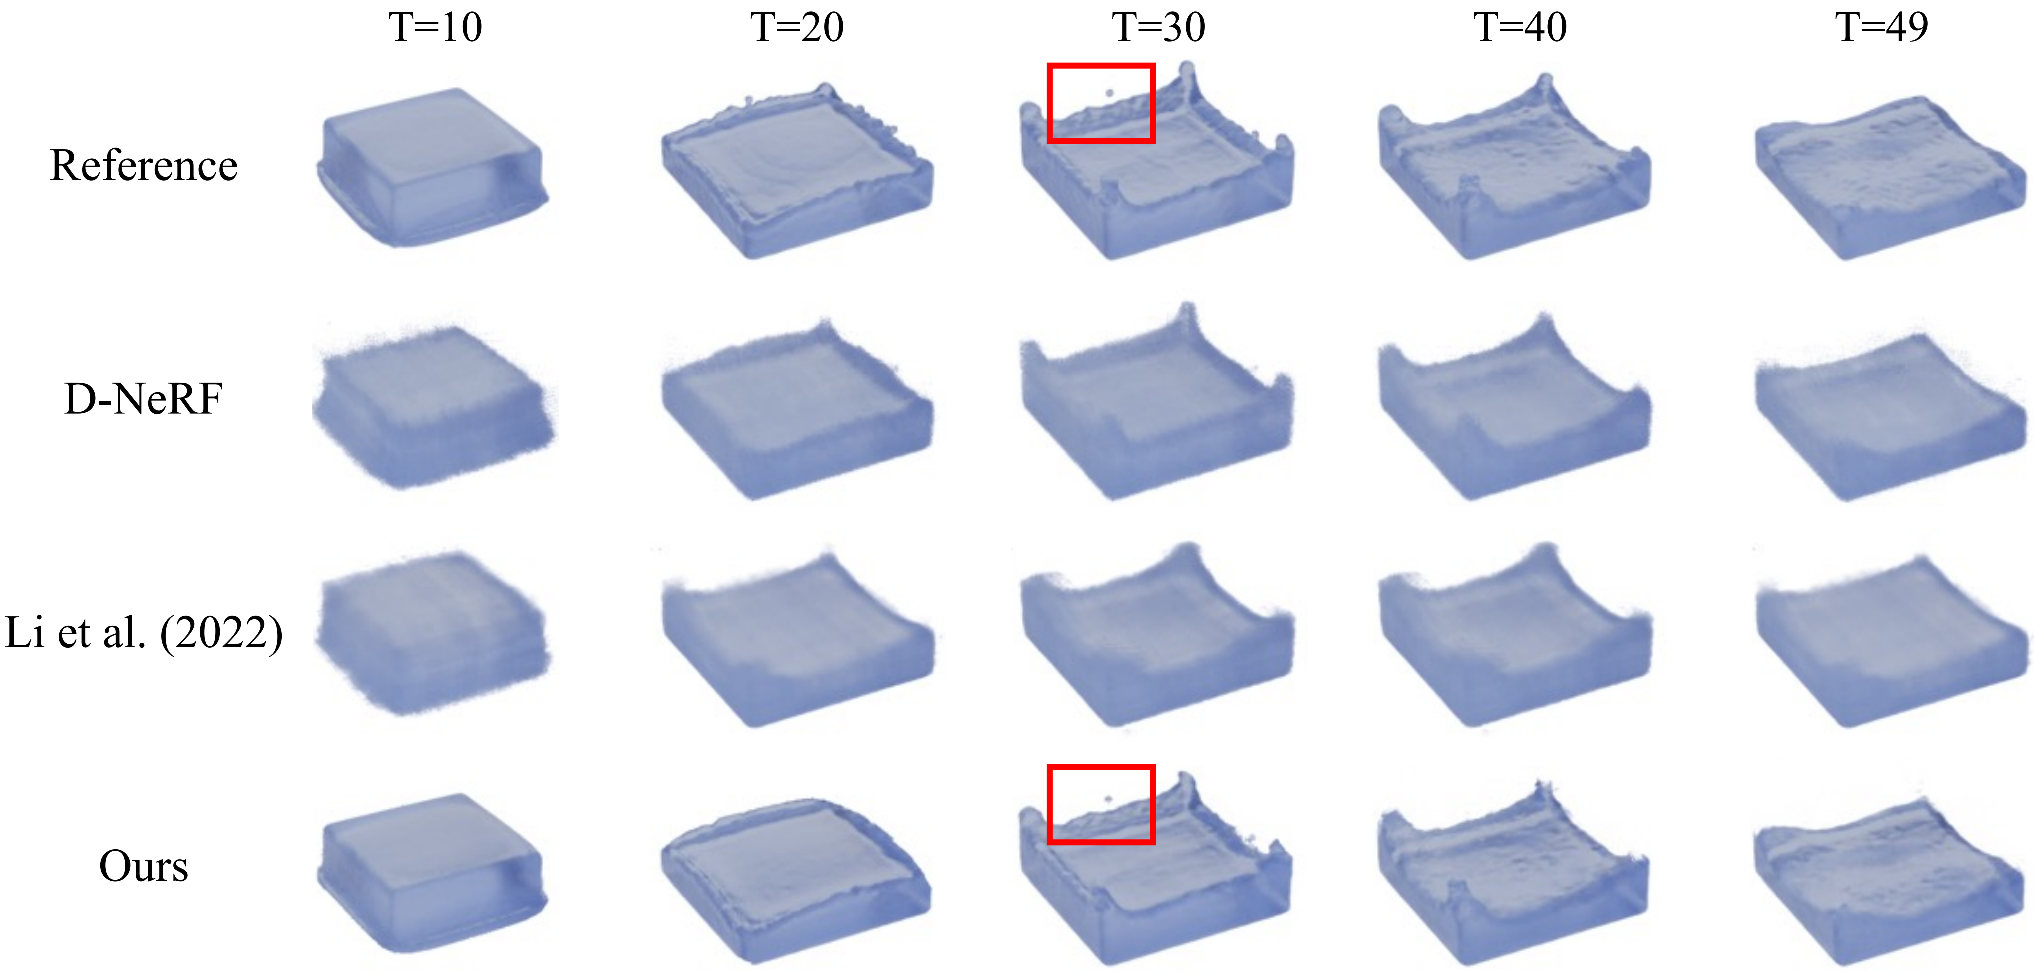

Results of Rendering Novel Fluid Scenes

We use a pretrained PhysNeRF model to render a novel water scene with the initial shape of Stanford Bunny

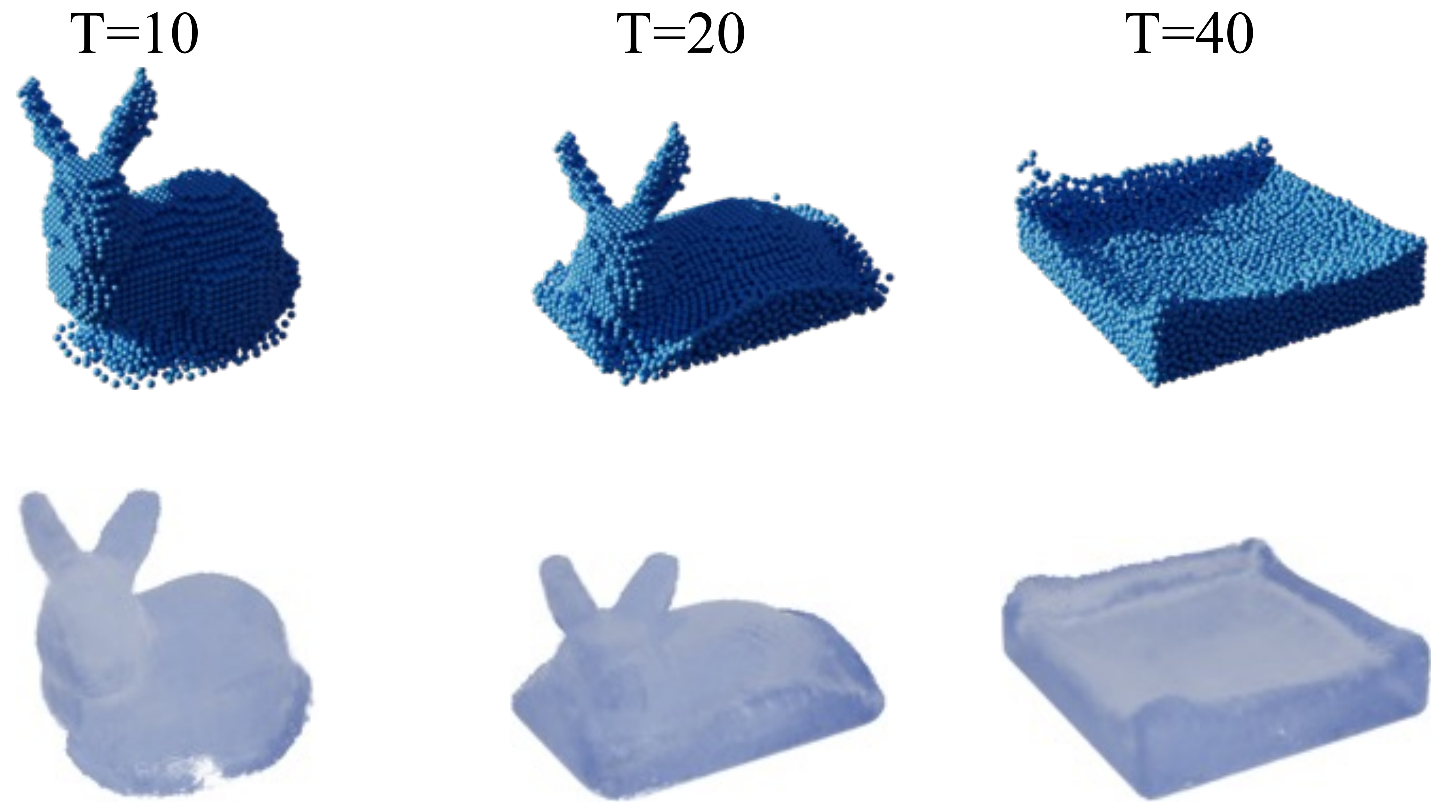

Results of Unknown Initial Particle Positions.

Table 5. Experiments on WaterCube with unknown initial particle states and ablation studies of neighborhood encoding (Rows 3-6).

MODEL	GROUNDING		PREDICTION		NOVEL VIEW SYNTHESIS		
	$d_{t<50}^{\mathrm{AVG}}$	$d_{t=49}$	$d_{t\geq 50}^{\mathrm{AVG}}$	$d_{t=59}$	PSNR↑	SSIM↑	LPIPS↓
FULL MODEL	**28.8**	34.9	<u>35.5</u>	<u>36.7</u>	**30.76**	0.95	**0.09**
UNKNOWN INITIAL PARTICLE POSITIONS	35.6	**27.2**	**26.6**	**26.3**	29.21	0.94	0.12
w/o FICTITIOUS PARTICLES CENTER ($\mathbf{p}_c$)	37.2	40.7	41.3	42.9	28.41	0.94	0.12
w/o SPHERE DENSITY (σ_{p})	<u>31.2</u>	37.9	39.3	39.4	<u>29.65</u>	0.95	<u>0.10</u>
w/o DEFORMATION VECTOR ($\boldsymbol{v}_{\mathrm{D}}$)	33.0	38.1	40.5	42.1	28.91	0.95	0.11
w/o PARTICLE-RELATIVE DIRECTION ($\boldsymbol{d}_c$)	32.2	39.8	43.9	47.0	29.56	0.95	<u>0.10</u>

Ablation Studies on Neighborhood Encoding

Table 5. Experiments on WaterCube with unknown initial particle states and ablation studies of neighborhood encoding (Rows 3-6).

Model	Grounding		Prediction		Novel view synthesis		
	$d^{\mathrm{AVG}}_{t<50}$	$d_{t=49}$	$d^{\mathrm{AVG}}_{t\geq 50}$	$d_{t=59}$	PSNR↑	SSIM↑	LPIPS↓
Full Model	**28.8**	34.9	35.5	36.7	**30.76**	0.95	**0.09**
Unknown initial particle positions	35.6	**27.2**	**26.6**	**26.3**	29.21	0.94	0.12
w/o Fictitious particles center ($\mathbf{p}_c$)	37.2	40.7	41.3	42.9	28.41	0.94	0.12
w/o Sphere density (σ_{p})	31.2	37.9	39.3	39.4	29.65	0.95	0.10
w/o Deformation vector ($\boldsymbol{v}_{\mathrm{D}}$)	33.0	38.1	40.5	42.1	28.91	0.95	0.11
w/o Particle-relative direction ($\boldsymbol{d}_c$)	32.2	39.8	43.9	47.0	29.56	0.95	0.10

Thanks for your watching!

Scan Me