

# 3D Neural Scene Representations for Visuomotor Control

Yunzhu Li\*, Shuang Li\*, Vincent Sitzmann, Pulkit Agrawal, and Antonio Torralba  
Computer Science and Artificial Intelligence Laboratory (CSAIL)

Massachusetts Institute of Technology

\* indicate equal contribution

Project Page: <https://3d-representation-learning.github.io/nerf-dy/>

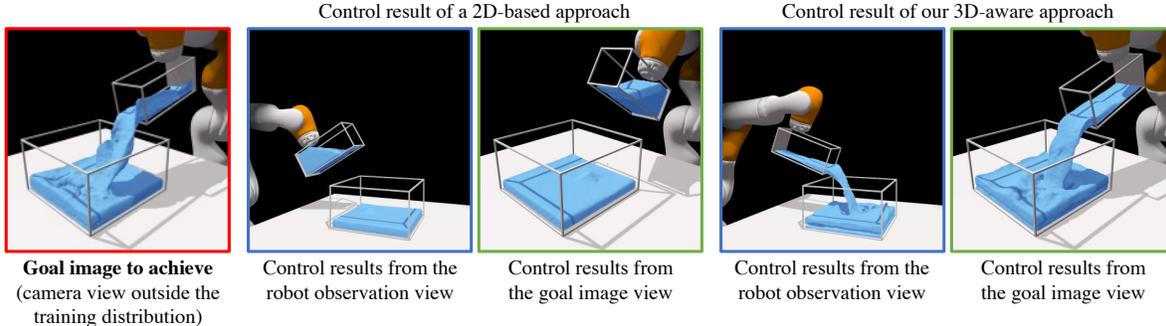


Fig. 1: Comparison of the control results between a 2D-based baseline and our 3D-aware approach. The task here is to achieve the configuration shown on the left, observed from a viewpoint that is outside the training distribution. The agent only takes a single-view visual observation as input from a viewpoint that is vastly different from the goal (images with blue frames). Our method generalizes well in this scenario and outperforms the 2D-based baseline, demonstrating the benefits of the learned 3D-aware scene representations.

**Abstract**—Humans have a strong intuitive understanding of the 3D environment around us. The mental model of the physics in our brain applies to objects of different materials and enables us to perform a wide range of manipulation tasks that are far beyond the reach of current robots. In this work, we desire to learn models for dynamic 3D scenes purely from 2D visual observations. Our model combines Neural Radiance Fields (NeRF) and time contrastive learning with an autoencoding framework, which learns viewpoint-invariant 3D-aware scene representations. We show that a dynamics model, constructed over the learned representation space, enables visuomotor control for challenging manipulation tasks involving both rigid bodies and fluids, where the target is specified in a viewpoint different from what the robot operates on. When coupled with an auto-decoding framework, it can even support goal specification from camera viewpoints that are *outside the training distribution*. We further demonstrate the richness of the learned 3D dynamics model by performing future prediction and novel view synthesis. Finally, we provide detailed ablation studies regarding different system designs and qualitative analysis of the learned representations.

## I. INTRODUCTION

Existing state-of-the-art model-based systems operating from vision treat images as 2D grids of pixels [5, 8, 33]. The world, however, is three-dimensional. Modeling the environment from 3D enables amodal completion and allows the agents to operate from different views. Therefore, it is desirable to obtain good 3D-aware representations of the environment from 2D observations to achieve better task performance when an accurate inference of 3D information is essential, which can further make it easier to specify tasks and learn from third-person videos, etc.

In this work, we aim to leverage recently proposed 3D-structure-aware implicit neural scene representations for visuomotor control tasks. We thus propose to embed neural radiance fields [20] in an auto-encoder framework, enabling tractable inference of the 3D-structure-aware scene state for dynamic environments. By additionally enforcing a time-contrastive loss on the estimated states, we ensure that the learned state representations are viewpoint-invariant. We then train a dynamics model that predicts the evolution of the state space conditioned on the input action, enabling us to perform control in the learned state space. Though the representation itself is 3D-structured, the convolutional encoder is not. At test time, we overcome this limitation by performing inference-via-optimization [26, 36], enabling accurate state estimation even for out-of-distribution camera poses and, therefore, control of tasks where the goal view is specified in an entirely unseen camera perspective. These contributions enable us to perform model-based visuomotor control of complex scenes, modeling both 3D dynamics of rigid objects and fluids. Through comparison with various baselines, the learned representation from our model is more precise at describing the contents of 3D scenes, which allows it to accomplish control tasks involving both rigid objects and fluids with significantly better accuracy (Figure 1). Please see our video for better visualization.

We summarize our contributions as follows: (i) We extend an autoencoding framework with a neural radiance field rendering module and time contrastive learning that allows us to learn 3D-aware scene representations for dynamics modeling and control purely from visual observations. (ii) By incorporating the auto-

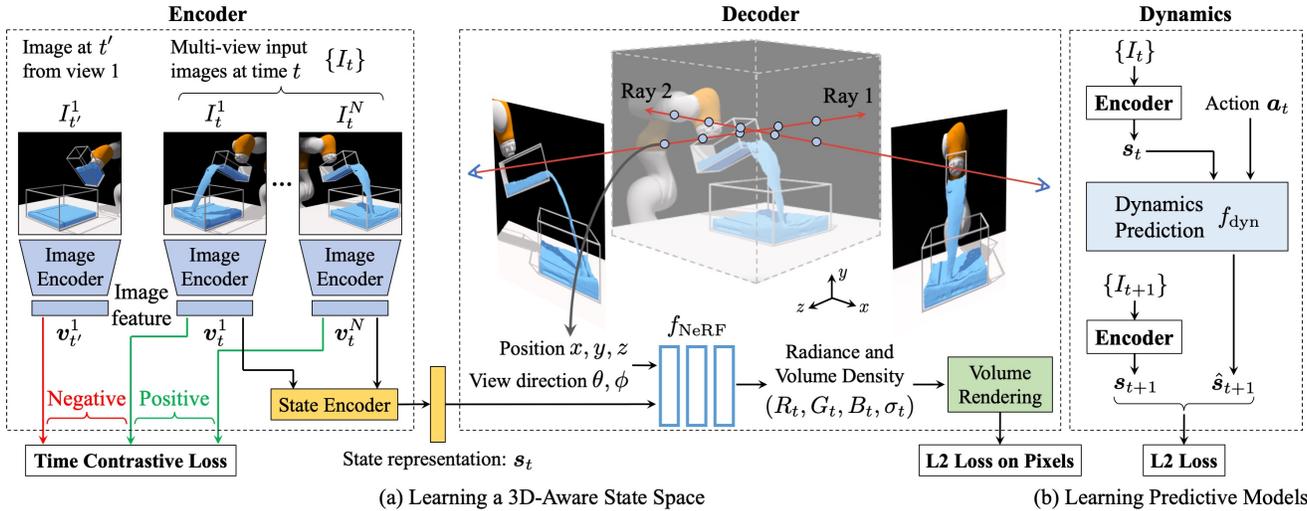


Fig. 2: **Overview of the training procedure.** **Left:** an encoder that maps the input images into a latent scene representation. The images are first sent to an image encoder to generate the image feature representations  $v$ . Then we combine the image features from the same time step using a state encoder to obtain the state representation  $s_t$ . A time contrastive loss is applied to enable our model to be invariant to camera viewpoints. **Middle:** a decoder that takes the scene representation as input and generates the observation image conditioned on a given viewpoint. We use an L2 loss to ensure the reconstructed image to be similar to the ground truth image. **Right:** a dynamics model that predicts the future scene representations  $\hat{s}_{t+1}$  by taking in the current representation  $s_t$  and action  $a_t$ . We use an L2 loss to enforce the predicted latent representation to be similar to the feature representation  $s_{t+1}$  extracted from the true visual observation  $I_{t+1}$ .

decoder mechanism at test time, our framework can adjust the learned representation and accomplish the control tasks with the goal specified from camera viewpoints outside the training distribution. (iii) We are the first to augment neural radiance fields using a time-invariant dynamics model, supporting future prediction and novel view synthesis across a wide range of environments with different types of objects.

## II. 3D-AWARE REPRESENTATIONS FOR DYNAMICS MODELING

Inspired by Neural Radiance Fields (NeRF) [20], we propose a framework that learns a viewpoint-invariant model for dynamic environments. As shown in Figure 2, our framework has three parts: (1) an encoder that maps the input images into a latent state representation, (2) a decoder that generates an observation image under a certain viewpoint based on the state representation, and (3) a dynamics model that predicts the future state representations based on the current state and the input action.

### A. 3D-Aware Scene Representation Learning

**Neural Radiance Field.** Given a 3D point  $x \in \mathbb{R}^3$  in a scene and a viewing direction unit vector  $d \in \mathbb{R}^3$  from a camera, NeRF learns to predict a volumetric radiance field. This is represented using a differentiable rendering function  $f_{\text{NeRF}}$  that predicts the corresponding density  $\sigma$  and RGB color  $c$  using  $f_{\text{NeRF}}(x, d) = (\sigma, c)$ . To render the color of an image pixel, NeRF integrates the information along the camera ray using  $\hat{C}(r) = \int_{h_{\text{near}}}^{h_{\text{far}}} T(h)\sigma(h)c(h)dh$ , where  $r(h) = o + hd$  is the camera ray with its origin  $o \in \mathbb{R}^3$  and unit direction vector  $d \in \mathbb{R}^3$ , and  $T(h) = \exp(-\int_{h_{\text{near}}}^h \sigma(s)ds)$  is the accumulated transparency between the pre-defined near depth  $h_{\text{near}}$  and far depth  $h_{\text{far}}$  along that camera ray. The mean squared error between the reconstructed color  $\hat{C}$  and the ground truth  $C$  is:

$$\mathcal{L}_{\text{rec}} = \sum_r \|\hat{C}(r) - C(r)\|_2^2. \quad (1)$$

**Neural Radiance Field for Dynamic Scenes.** One key limitation of NeRF is that it assumes the scene is static. For a dynamic scene, it must learn a separate radiance field  $f_{\text{NeRF}}$  for each time step. This severely limits the ability of NeRF to model environments that change over time, as it is both time-consuming and unclear how to transfer knowledge across time or when a new scene is similar to an old one. While other models have shown generalization across scenes [36, 24], these representations do not capture fine details. To enable  $f_{\text{NeRF}}$  to model dynamic scenes, we learn an encoding function  $f_{\text{enc}}$  that maps the visual observations to a feature representation  $s$  for each scene and learn the volumetric radiance field decoding function based on  $s$ . Let  $\{I_t\}$  denotes the set of 2D images that capture a 3D scene at time  $t$  from one or more camera viewpoints. The image taken from the  $i^{\text{th}}$  viewpoint is represented as  $I_t^i$ . We use ResNet-18 [10] to extract a feature vector for each image. We take the output of ResNet-18 before the pooling layer and sent it to a fully-connected layer, resulting in a 256 dimension image feature  $v_t^i$ . This image feature is concatenated with the corresponding camera viewpoint information (a 16-D vector obtained by flattening the camera view matrix). The concatenated feature is processed using a small multilayer perceptron (MLP) to generate the final image feature. The scene representation  $s_t$  is generated by first averaging the image features across multiple camera viewpoints and then being encoded using another small MLP.

Given a 3D point  $x$ , a viewing direction unit vector  $d$ , and a scene representation  $s_t$ , we learn a function  $f_{\text{dec}}(x, d, s_t) = (\sigma_t, c_t)$  to predict the radiance field represented by the density  $\sigma_t$  and RGB color  $c_t$ . Similar to NeRF, we use the integrated information along the camera ray to render the color of image

pixels from an input viewpoint and then compute the image reconstruction loss using Equation 1. During each training iteration, we render two images from different viewpoints to calculate more accurate gradient updates.  $f_{\text{dec}}$  depends on the scene representation  $\mathbf{s}_t$ , forcing it to encode the 3D contents of the scene to support rendering from different camera poses.

**Time Contrastive Learning.** To enable the image encoder to be viewpoint invariant, we regularize the feature representation of each image  $\mathbf{v}_t^i$  using multi-view time contrastive loss (TCN) [34] (see Figure 2a). The TCN loss encourages features of images from different viewpoints at the same time step to be similar, while repulsing features of images from different time steps to be dissimilar. More specifically, given a time step  $t$ , we randomly select one image  $I_t^i$  as the anchor and extract its image feature  $\mathbf{v}_t^i$  using the image encoder. Then we randomly select one positive image from the same time step but different camera viewpoint  $I_t^{i'}$  and one negative image from a different time step but the same viewpoint  $I_{t'}^i$ . We use the same image encoder to extract their image features  $\mathbf{v}_t^{i'}$  and  $\mathbf{v}_{t'}^i$ . Similar to [34], we minimize the following time contrastive loss:

$$\mathcal{L}_{\text{tc}} = \max(\|\mathbf{v}_t^i - \mathbf{v}_t^{i'}\|_2^2 - \|\mathbf{v}_t^i - \mathbf{v}_{t'}^i\|_2^2 + \alpha, 0), \quad (2)$$

where  $\alpha$  is a hyper-parameter denoting the *margin* between the positive and negative pairs.

### B. Learning the Predictive Model

After we have obtained the latent state representation  $\mathbf{s}$ , we use supervised learning to estimate the forward dynamics model,  $\hat{\mathbf{s}}_{t+1} = f_{\text{dyn}}(\mathbf{s}_t, \mathbf{a}_t)$ . Given  $\mathbf{s}_t$  and a sequence of actions  $\{\mathbf{a}_t, \mathbf{a}_{t+1}, \dots\}$ , we predict  $H$  steps in the future by iteratively feeding in actions into the one-step forward model. We implement  $f_{\text{dyn}}$  as an MLP network which is trained by optimizing the following loss function:

$$\mathcal{L}_{\text{dyn}} = \sum_{h=1}^H \|\hat{\mathbf{s}}_{t+h} - \mathbf{s}_{t+h}\|_2^2, \quad (3)$$

where  $\hat{\mathbf{s}}_{t+h} = f_{\text{dyn}}(\hat{\mathbf{s}}_{t+h-1}, \mathbf{a}_{t+h-1})$ ,  $\hat{\mathbf{s}}_t = \mathbf{s}_t$ . We define the final loss as a combination of the image reconstruction loss, the time contrastive loss, and the dynamics prediction loss:  $\mathcal{L} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{tc}} + \mathcal{L}_{\text{dyn}}$ . We first train the encoder  $f_{\text{enc}}$  and decoder  $f_{\text{dec}}$  together using stochastic gradient descent (SGD) by minimizing  $\mathcal{L}_{\text{rec}}$  and  $\mathcal{L}_{\text{tc}}$ , which makes sure that the learned scene representation  $\mathbf{s}$  encodes the 3D contents and is viewpoint-invariant. We then fix the encoder parameters, and train the dynamics model  $f_{\text{dyn}}$  by minimizing  $\mathcal{L}_{\text{dyn}}$  using SGD. See the supplementary materials for the network architecture and training details.

## III. VISUOMOTOR CONTROL

### A. Online Planning for Closed-Loop Control

When given the goal image  $I^{\text{goal}}$  and its associated camera pose, we can feed them through the encoder  $f_{\text{enc}}$  to get the state representation for the goal configuration  $\mathbf{s}^{\text{goal}}$ . We use the same method to compute the state representation for the current scene configuration  $\mathbf{s}_1$ . The goal of the online planning problem is to find an action sequence  $\mathbf{a}_1, \dots, \mathbf{a}_{T-1}$  that minimizes

the distance between the predicted future representation and the goal representation at time  $T$ . As shown in Figure 3a, given a sequence of actions, our model can iteratively predict a sequence of latent state representations. The latent-space dynamics model can then be used for downstream closed-loop control tasks via online planning with model-predictive control (MPC), which is used to get feedback from the environment and give the agent a chance to adjust the current action sequence. We formally define the online planning problem as follows:

$$\begin{aligned} \min_{\mathbf{a}_1, \dots, \mathbf{a}_{T-1}} \quad & \|\mathbf{s}^{\text{goal}} - \hat{\mathbf{s}}_T\|_2^2, \\ \text{s.t.} \quad & \hat{\mathbf{s}}_1 = \mathbf{s}_1, \hat{\mathbf{s}}_{t+1} = f_{\text{dyn}}(\hat{\mathbf{s}}_t, \mathbf{a}_t). \end{aligned} \quad (4)$$

Many existing off-the-shelf model-based RL methods can be used to solve the MPC problem [4, 21, 9, 7, 19, 13, 14]. We experimented with random shooting, gradient-based trajectory optimization, cross-entropy method, and model-predictive path integral (MPPI) planners [44] and found that MPPI performed the best. MPPI is a sampling-based, gradient-free optimizer that considers temporal coordination between time steps when sampling action trajectories, which is the main planner in the subsequent experiments (Figure 4 and 5).

### B. Auto-Decoder for Viewpoint Extrapolation

End-to-end visuomotor agents can undergo significant performance drop when the test-time visual observations are captured from camera poses outside the training distribution. The convolutional image encoder suffers from the same problem as it is not equivariant to changes in the camera pose, meaning it has a hard time generalizing to out-of-distribution camera views. As shown in Figure 3b, when encountered an image from a viewpoint outside training distribution, with a pixel distribution vastly different from what the model is trained on, passing it through the encoder  $f_{\text{enc}}$  will give us an amortized estimation of the scene representation  $\mathbf{s}_t$ . It has a high chance that the decoded image is different from the ground truth as the viewpoint has never been encountered during training.

We fix this problem at test time by applying the inference-by-optimization (also named as an auto-decoding) framework that backpropagates through the volumetric renderer and through the neural implicit representation into the state estimate [26, 36]. This is inspired by the fact that the rendering function,  $f_{\text{dec}}(\mathbf{x}, \mathbf{d}, \mathbf{s}_t) = (\sigma_t, \mathbf{c}_t)$  is viewpoint equivariant, where the output only depends on the state representation  $\mathbf{s}_t$ , the location  $\mathbf{x}$ , and the ray direction  $\mathbf{d}$ , meaning that the output is invariant to the camera position along the camera ray, i.e., even we move the camera closer or farther away along the camera ray,  $f_{\text{dec}}$  still tends to generate the same results. We leverage this property and calculate the L2 distance between the input image and the reconstructed image  $\mathcal{L}_{\text{ad}} = \|I_t - \hat{I}_t\|_2^2$ , and then update the scene representation  $\mathbf{s}_t$  using stochastic gradient descent. We repeat this updating process  $K$  times to derive the state representation of the underlying 3D scene. Note that this update only changes the scene representation while keeping the parameters in the decoder fixed. The resulting representation is used as  $\mathbf{s}^{\text{goal}}$  in Equation 4 to solve the planning problem.

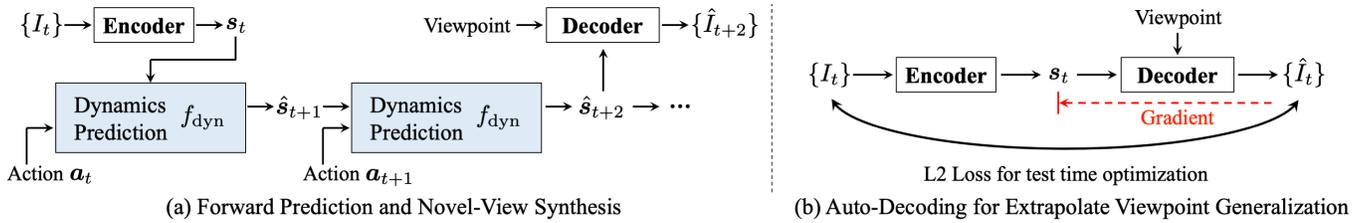


Fig. 3: **Forward prediction and viewpoint extrapolation.** (a) We first feed the input image(s) at time  $t$  to the encoder to derive the scene representation  $s_t$ . The dynamics model then takes  $s_t$  and the corresponding action sequence as input to iteratively predict the future. The decoder synthesizes the visual observation conditioned on the predicted state representation and an input viewpoint. (b) We propose an auto-decoding inference-via-optimization framework to enable extrapolated viewpoint generalization. Given an input image  $I_t$  taken from a viewpoint outside the training distribution, the encoder first predicts the scene representation  $s_t$ . Then the decoder reconstructs the observation  $\hat{I}_t$  from  $s_t$  and the camera viewpoint from  $I_t$ . We calculate the L2 distance between  $I_t$  and  $\hat{I}_t$  and backpropagate the gradient through the decoder to update  $s_t$ . The updating process is repeated for  $K$  iterations, resulting in a more accurate  $s_t$  of the underlying 3D scene.

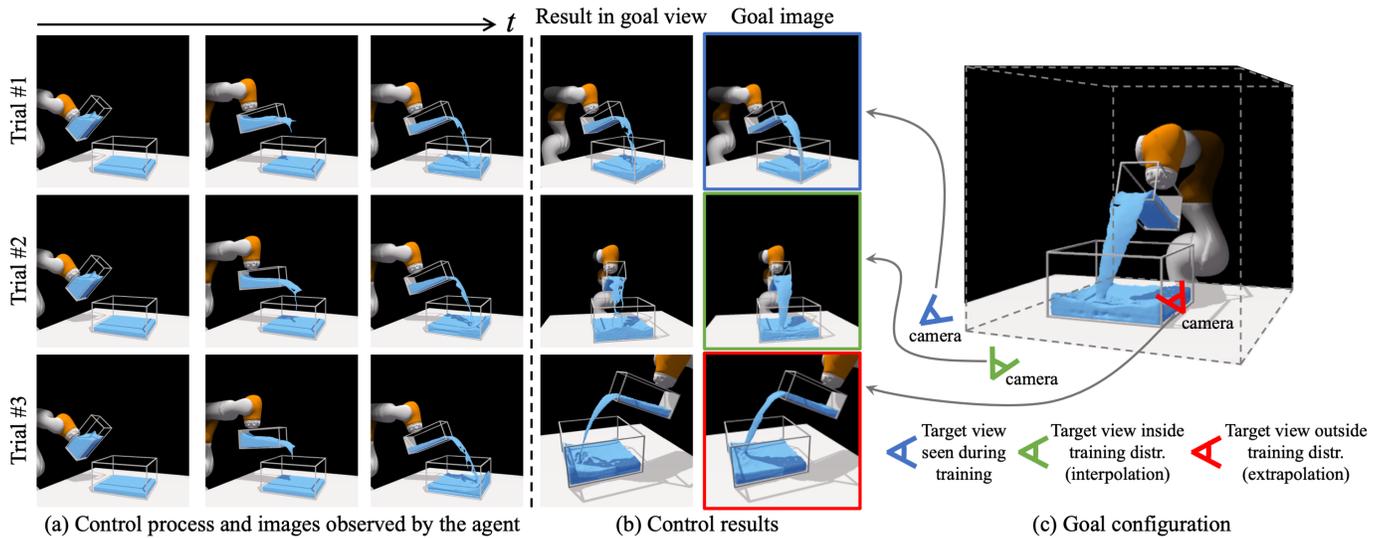


Fig. 4: **Qualitative control results of our method on three types of testing scenarios.** The image on the right shows the target configuration we aim to achieve. The left three columns show the control process, which are also the input images to the agent. The fourth column is the control results from the same viewpoint as the goal image. Trial #1 specifies the goal using a different viewpoint from the agent’s but has been encountered during training. Trial #2 uses a goal view that is an interpolation of training viewpoints. Trial #3 uses an extrapolated viewpoint that is outside the training distribution. Our method performs well in all settings.

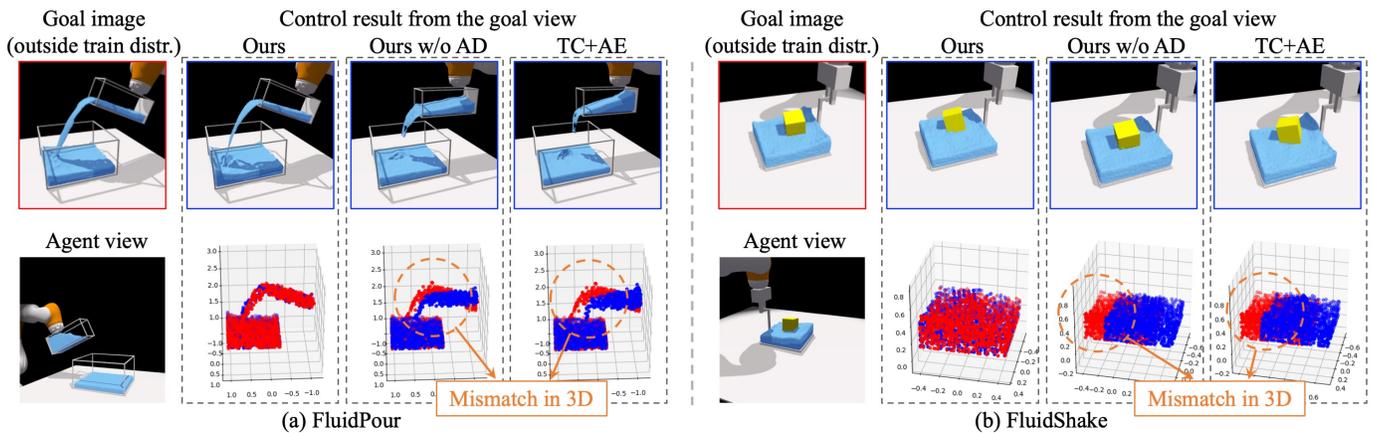


Fig. 5: **Qualitative comparisons between our method and baseline approaches on the control tasks.** We show the closed-loop control results on the FluidPour and FluidShake environment. The goal image viewpoint (top-left image of each block) is outside the training distribution and is different from the viewpoint observed by the agent (bottom-left image of each block). Our final control results are much better than a variant that does not perform auto-decoding test-time optimization (Ours w/o AD) and the best-performing baseline (TC+AE), both of which fail to accomplish the task and their control results (blue points) exhibit an apparent deviation from the target configuration (red points) when measured in the 3D points space.

## REFERENCES

- [1] Pulkit Agrawal, Ashvin Nair, Pieter Abbeel, Jitendra Malik, and Sergey Levine. Learning to poke by poking: Experiential learning of intuitive physics. *arXiv preprint arXiv:1606.07419*, 2016.
- [2] Maria Bauza, Francois R Hogan, and Alberto Rodriguez. A data-efficient approach to precise and controlled pushing. In *Conference on Robot Learning*, pages 336–345. PMLR, 2018.
- [3] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. *arXiv e-prints*, pages arXiv–2012, 2020.
- [4] Frederik Ebert, Chelsea Finn, Alex X Lee, and Sergey Levine. Self-supervised visual planning with temporal skip connections. In *CoRL*, pages 344–356, 2017.
- [5] Frederik Ebert, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex Lee, and Sergey Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- [6] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [7] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. *arXiv preprint arXiv:1605.07157*, 2016.
- [8] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- [9] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pages 2555–2565. PMLR, 2019.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [11] François Robert Hogan and Alberto Rodriguez. Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. *arXiv preprint arXiv:1611.08268*, 2016.
- [12] Tejas D Kulkarni, Ankush Gupta, Catalin Ionescu, Sebastian Borgeaud, Malcolm Reynolds, Andrew Zisserman, and Volodymyr Mnih. Unsupervised learning of object keypoints for perception and control. *Advances in neural information processing systems*, 32:10724–10734, 2019.
- [13] 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.
- [14] Yunzhu Li, Jiajun Wu, Jun-Yan Zhu, Joshua B Tenenbaum, Antonio Torralba, and Russ Tedrake. Propagation networks for model-based control under partial observation. In *ICRA*, 2019.
- [15] Yunzhu Li, Toru Lin, Kexin Yi, Daniel Bear, Daniel Yamins, Jiajun Wu, Joshua Tenenbaum, and Antonio Torralba. Visual grounding of learned physical models. In *International conference on machine learning*, pages 5927–5936. PMLR, 2020.
- [16] Yunzhu Li, Antonio Torralba, Anima Anandkumar, Dieter Fox, and Animesh Garg. Causal discovery in physical systems from videos. *Advances in Neural Information Processing Systems*, 33, 2020.
- [17] Zhengqi Li, Simon Niklaus, Noah Snively, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *arXiv preprint arXiv:2011.13084*, 2020.
- [18] Miles Macklin, Matthias Müller, Nuttapon Chentanez, and Tae-Yong Kim. Unified particle physics for real-time applications. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014.
- [19] Lucas Manuelli, Yunzhu Li, Pete Florence, and Russ Tedrake. Keypoints into the future: Self-supervised correspondence in model-based reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2020.
- [20] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020.
- [21] Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, pages 1101–1112. PMLR, 2020.
- [22] Thu Nguyen-Phuoc, Chuan Li, Stephen Balaban, and Yong-Liang Yang. RenderNet: A deep convolutional network for differentiable rendering from 3d shapes. *arXiv preprint arXiv:1806.06575*, 2018.
- [23] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5379–5389, 2019.
- [24] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020.
- [25] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2856–2865, 2021.
- [26] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [27] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020.
- [28] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020.
- [29] Danilo Jimenez Rezende, SM Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. *arXiv preprint arXiv:1607.00662*, 2016.
- [30] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019.
- [31] Connor Schenck and Dieter Fox. Reasoning about liquids via closed-loop simulation. *arXiv preprint arXiv:1703.01656*, 2017.
- [32] Connor Schenck and Dieter Fox. Perceiving and reasoning about liquids using fully convolutional networks. *The International Journal of Robotics Research*, 37(4-5):452–471, 2018.
- [33] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [34] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1134–1141. IEEE, 2018.
- [35] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proc. CVPR*, 2019.
- [36] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*, 2019.
- [37] Vincent Sitzmann, Eric R Chan, Richard Tucker, Noah Snaveley, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. *arXiv preprint arXiv:2006.09662*, 2020.
- [38] HJ Suh and Russ Tedrake. The surprising effectiveness of linear models for visual foresight in object pile manipulation. *arXiv preprint arXiv:2002.09093*, 2020.
- [39] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Single-view to multi-view: Reconstructing unseen views with a convolutional network. *CoRR abs/1511.06702*, 1(2):2, 2015.
- [40] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a deforming scene from monocular video. *arXiv preprint arXiv:2012.12247*, 2020.
- [41] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering. *arXiv preprint arXiv:2010.04595*, 2020.
- [42] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. *Proc. CVPR*, 2019.
- [43] Manuel Watter, Jost Tobias Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *arXiv preprint arXiv:1506.07365*, 2015.
- [44] Grady Williams, Andrew Aldrich, and Evangelos Theodorou. Model predictive path integral control using covariance variable importance sampling. *arXiv preprint arXiv:1509.01149*, 2015.
- [45] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Interpretable transformations with encoder-decoder networks. In *Proc. ICCV*, volume 4, 2017.
- [46] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021.
- [47] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. *arXiv preprint arXiv:2012.02190*, 2020.
- [48] Jiaji Zhou, Yifan Hou, and Matthew T Mason. Pushing revisited: Differential flatness, trajectory planning, and stabilization. *The International Journal of Robotics Research*, 38(12-13):1477–1489, 2019.
- [49] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: image generation with disentangled 3d representations. In *Proc. NIPS*, pages 118–129, 2018.

### A. Environments & Evaluation Metrics

**Environments.** We consider the following four environments that involve both fluid and rigid objects for evaluating the proposed model and baseline approaches. The environments are simulated using NVIDIA FleX [18]. (1) FluidPour (Figure 8a): This environment contains a fully-actuated cup that pours fluids into a container at the bottom. (2) FluidShake (Figure 8b): A fully-actuated container moves on a 2D plane. Inside the container are fluids and a rigid cube floating on the surface. (3) RigidStack (Figure 8c): Three rigid cubes form a vertical stack and are released from a certain height but in different horizontal positions. They fall down and collide with each other and the ground. (4) RigidDrop (Figure 8d): A cube falls down from a certain height. There is a container fixed at a random position on the ground. The cube either falls into the container or bounces out.

**Evaluation Metrics.** We use the first two environments, i.e., FluidPour and FluidShake, to measure the control performance, where we specify the target configuration of the control task using images from (1) one of the viewpoints encountered during training, (2) an interpolated viewpoint between training viewpoints, and (3) an extrapolated viewpoint outside the training distribution (Figure 4).

We provide quantitative evaluations on the control performance in FluidPour and FluidShake by extracting the particle set from the simulator, and measure the Chamfer distance between the result and the goal, which we denote as ‘‘Chamfer Dist.’’. In FluidPour, we provide additional measurements on the L2 distance of the position/orientation of the cup towards the goal, denoting as ‘‘Position Error’’ and ‘‘Angle Error’’ respectively. In FluidShake, we calculate the L2 distance of the container and cube’s position towards the goal and denote them as ‘‘Container Error’’ and ‘‘Cube Error’’ respectively.

### B. Baseline Methods

For comparison, we consider the following three baselines: **TC**: Similar to [34], it only uses time contrastive loss for learning the image feature without having to reconstruct the scene. We learn a dynamics model directly on the image features for control. **TC+AE**: Instead of using Neural Radiance Fields to reconstruct the image, this method uses the standard convolutional decoder to reconstruct the target image when given a new viewpoint. This would then be similar to [6] augmented with a time contrastive loss. **NeRF**: This method is a direct adaptation from the original NeRF paper [20] and is the same as ours except that it does not include the time contrastive loss during training and the auto-decoding test-time optimization. We use the same dynamics model shown in Figure 2b and train the model for each baseline respectively for dynamic prediction. We use the same feedback control method, i.e., MPPI [44], for our model and all the baselines.

### C. Control Results

**Goal Specification from Novel Viewpoints.** Figure 4c shows the goal configuration, and we ask the learned model

to perform three control trials where the goal is specified from different types of viewpoints. The left three columns show the MPC control process from the agent’s viewpoint. The fourth column visualizes the final configuration the agent achieves from the same viewpoint as the goal image. Trial #1 specifies the goal using a different viewpoint from the agent’s but has been encountered during training. Trial #2 uses a goal view that is an interpolation of training viewpoints. Our agent can achieve the target configuration with a decent accuracy. For trial #3, we specify the goal view by moving the camera closer, higher, and more downwards with respect to the container. Note this goal image view is outside the distribution of training viewpoints. With the help of test-time auto-decoding optimization introduced in the main paper Section III-B, our method can successfully achieve the target configuration as shown in the figure.

**Baseline comparisons.** We benchmark our model with the baselines by assessing their performance on the downstream control tasks. Figure 5 shows the qualitative comparison between our model (Ours), a variant of our model that does not perform the auto-decoding test-time optimization (Ours w/o AD), and the best-performing baseline (TC+AE) introduced in Section B. We find that when the target view is outside the training distribution and vastly different from the agent view (bottom left image in each trial block), our full method shows a much better performance in achieving the target configuration. The variant without auto-decoding optimization and TC+AE fail to accomplish the task and exhibit an apparent deviation from the ground truth in the 3D points space.

We also provide quantitative comparisons on the control results. Figure 6 shows the performance in the FluidPour and FluidShake environments. We find our full model significantly outperforms the baseline approaches in both environments under all scenarios and evaluation metrics. The results effectively demonstrate the advantages of the learned 3D-aware scene representations, which contain a more precise encoding of the contents in the 3D environments.

### D. Analysis of the Learned Model

To better understand why our 3D-aware model outperforms other baselines in the downstream control tasks, we provide a deeper investigation of the learned state representation.

**Nearest Neighbor Search Using the Learned Representation.** When the goal image is specified from a viewpoint different from the agent’s view, to ensure the planning problem still work, it is essential that the distance in the learned feature space reflects the distance in the actual 3D space, i.e., scenes that are more similar in the real 3D space should be closer in the learned feature space, even if the visual observations are captured from different viewpoints. We visualize the nearest neighbor results in Figure 7. Given a query image, we search its nearest neighbors based on their state representation  $s$ . Even if the images look quite different from each other, the learned 3D-aware scene representations can retrieve reasonable neighborhood images that share similar 3D contents, indicating that the learned 3D-aware scene representations hold a good

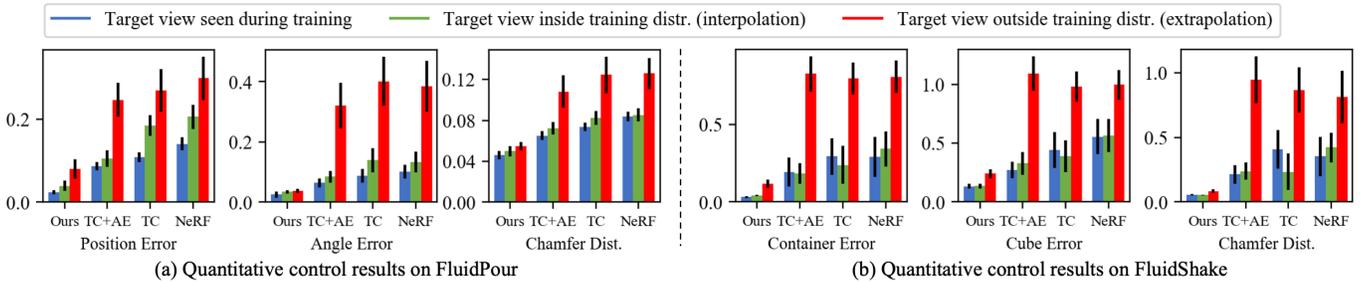


Fig. 6: **Quantitative comparisons between our method and baselines on the FluidPour and FluidShake.** In each environment, we compare the results using three different evaluation metrics under three settings, i.e., (1) the target image view seen during training, (2) the target image view is inside the training distribution but not seen during training (interpolation), and (3) the target image view is outside the training distribution (extrapolation). The height of the bars indicates the mean, and the error bar denotes the standard error. Our model significantly outperforms all baselines under all testing settings.

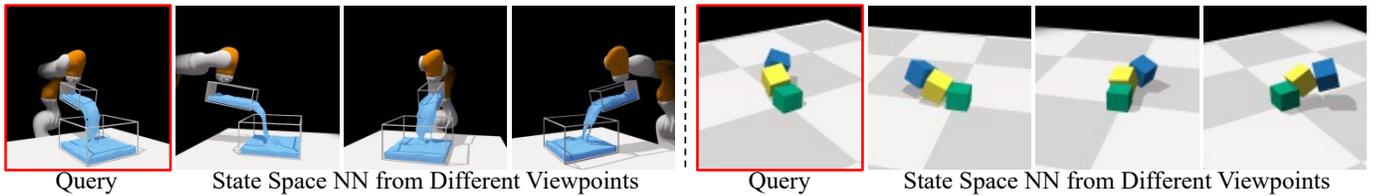


Fig. 7: **Nearest neighbor (NN) results using our learned state representation.** Given a query image (red boundary), we search its nearest neighbors based on their state representation. Our learned scene representations can retrieve reasonable neighbor images, indicating that our state representations retain a good estimation of the contents inside the 3D scene and are invariant to camera viewpoints.

understanding of the real 3D scene and are invariant to camera viewpoints.

**Dynamic Prediction and Novel View Synthesis.** Conditioned on a scene representation and an input action sequence, our dynamics model  $f_{\text{dyn}}$  can iteratively predict the evolution of the scene representations. Our decoder can then take the predicted state representation and reconstruct the corresponding visual observation from a query viewpoint. Figure 8 shows that our model can accurately predict the future and perform novel view synthesis on four environments involving both fluid and rigid objects, suggesting its usefulness in trajectory optimization. Please check our video results in the supplementary material for better visualization.

### E. Model Details

In the decoder model, we use a similar network architecture as the NeRF paper [20]. As shown in Figure 9, we send a 3D point  $\mathbf{x} \in \mathbb{R}^3$ , a camera ray  $\mathbf{d} \in \mathbb{R}^3$ , and a state feature representation  $\mathbf{s}_t$  into a fully-connected network and output the corresponding density  $\sigma_t$  and RGB color  $\mathbf{c}_t$ .

### F. Environment Details

In the **FluidPour** environment, we generated 1,000 trajectories for training. Each trajectory has 300 frames with 20 camera views sampled around the objects with a fixed distance towards the world origin. The action space for the control task is the position and tilting angle of the cup, which are randomly generated when constructing the training set.

In the **FluidShake** environment, we generated 1,000 trajectories for training. Each trajectory has 300 frames with 20 camera views sampled around the objects with a fixed distance towards the world origin. The action space for the control task is

the 2D location of the container in the world coordinate, which is also randomly generated when constructing the training set.

In the **RigidStack** environment, we generated 800 trajectories for training. Each trajectory has 80 frames with 20 camera views sampled around the objects with a fixed distance towards the world origin.

In the **RigidDrop** environment, we generated 1,000 trajectories for training. Each trajectory has 50 frames with 20 camera views sampled around the objects with a fixed distance towards the world origin.

### G. Training Details

For the encoder and decoder model as described in Figure 2, we use the Adam optimizer with the initial learning rate  $5e^{-4}$  and decreased to  $5e^{-5}$  for all the experiments. The batch size is 2. The hyperparameters in the decoder are the same as the original NeRF model [20] except the near and far distance between the objects and cameras are different in our environments. In our FluidPour environment, we have near = 2.0 and far = 8.0. In our FluidShake environment, we have near = 2.0 and far = 7.0. In our RigidStack environment, we have near = 2.0 and far = 7.0. In our FluidPour environment, we have near = 2.0 and far = 6.0.

### H. Control Details

The number of updating iteration for the auto-decoding test-time optimization  $K$  is 500. The number of sampled trajectories  $M$  during MPPI optimization is set to 1,000. The number of iterations  $L$  for updating the action sequence is set to 100 for the first time step, and 10 for the subsequent control steps to maintain a better trade-off between efficiency and effectiveness. The reward-weighting factor  $\gamma$  is set to 1 and the filtering

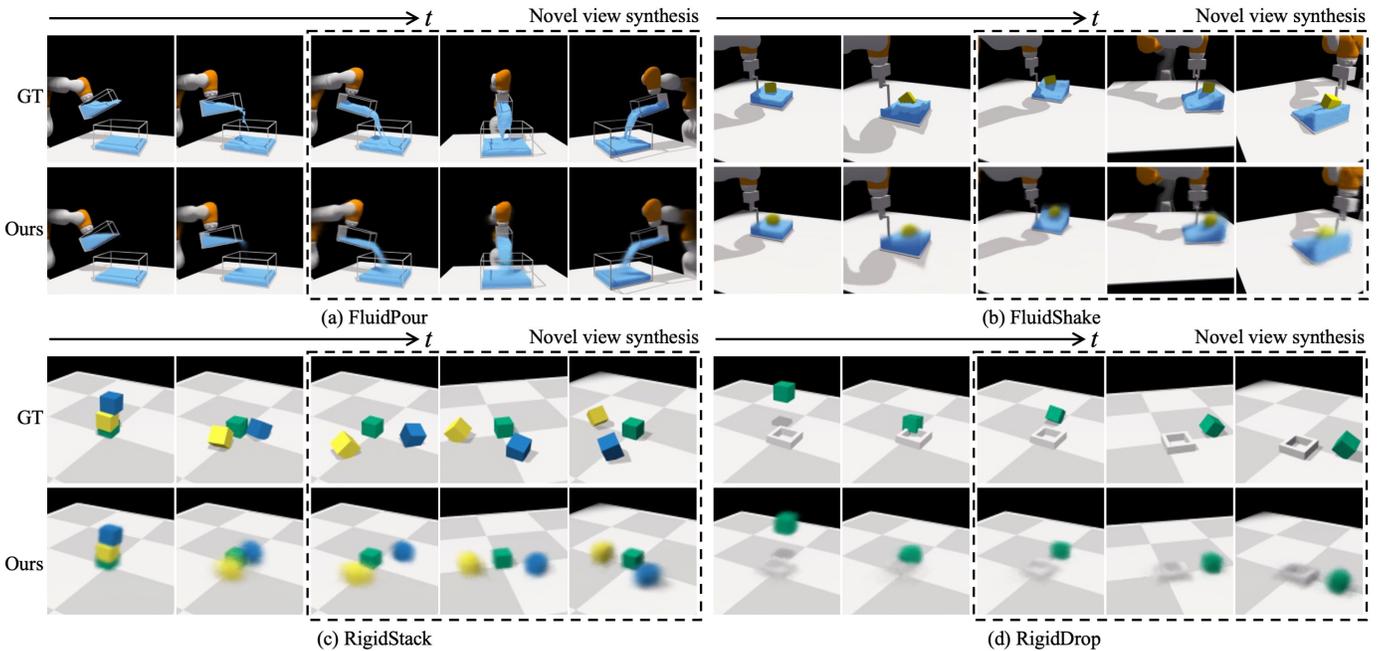


Fig. 8: **Forward prediction and novel view synthesis on four environments.** Given a scene representation and an input action sequence, our dynamics model predicts the subsequent latent scene representation, which is used as the input of our decoder model to reconstruct the corresponding visual observation based on different viewpoints. In each block, we render images based on the open-loop future dynamic prediction. Images in the dotted box compare the novel view synthesis results of our model in the last time step with the ground truth from three different viewpoints.

coefficient  $\beta$  is specified as 0.7. The control horizon  $T$  is set as 80 both for FluidPour and FluidShake. The hyperparameters are the same for all compared methods.

### I. Related Work

**3D Scene Representation Learning.** Prior work leverages the latent spaces of autoencoder-like models as learned representations of the underlying 3D scene to enable novel view synthesis from a single image [45, 39]. Eslami et al. [6] embed this approach in a probabilistic framework. To endow models with 3D structures, voxelgrids can be leveraged as neural scene representations [29, 22, 35, 42, 49], while others have tried to predict particle sets from images [15]. Sitzmann et al. [36] propose to learn neural implicit representations of 3D shape and appearance supervised with posed 2D images via a differentiable renderer. Generalizing across neural implicit representations can also be realized by local conditioning on CNN features [30, 41, 47], but this does not learn a global representation of the scene state. Alternatively, gradient-based meta-learning has been proposed for faster inference of implicit neural representations [37]. Deformable scenes can be modeled by transporting input coordinates to neural implicit representations with an implicitly represented flow field [23, 27, 28, 40, 17, 3, 46, 25]; however, they typically fit one trajectory and cannot handle different initial conditions and external action inputs, limiting their use in control.

**Model-Based RL in Robotic Manipulation.** We can categorize model-based RL methods by whether they use physics-based or data-driven models, and whether they assume full state access or only visual observation. Methods that rely on

first principles typically assume full-state information of the environment [11, 48] and require the knowledge of the object models, making them hard to generalize to novel objects or partially observable scenarios. For data-driven models, people have attempted to learn a dynamics model for closed-loop planar pushing [2] or dexterous manipulation [21]. Schenck and Fox [31, 32] tackle a similar fluid pouring task via closed-loop simulation. Although they have achieved impressive results, they rely on state estimators customized for specific tasks, limiting their applicability to more general and diversified manipulation tasks.

Various model-based RL methods have been proposed to learn state representations from visual observations, such as image-space dynamics [7, 4, 5, 38], keypoint representation [12, 19, 16], and low-dimensional latent space [43, 9, 8, 33]. Some works learn a meaningful representation space using reconstruction loss [9, 8]. Others jointly train the forward and inverse dynamics models [1], or use time contrastive loss to regularize the latent embedding [34]. We differ from the previous methods by explicitly incorporating a 3D volumetric rendering process during training.

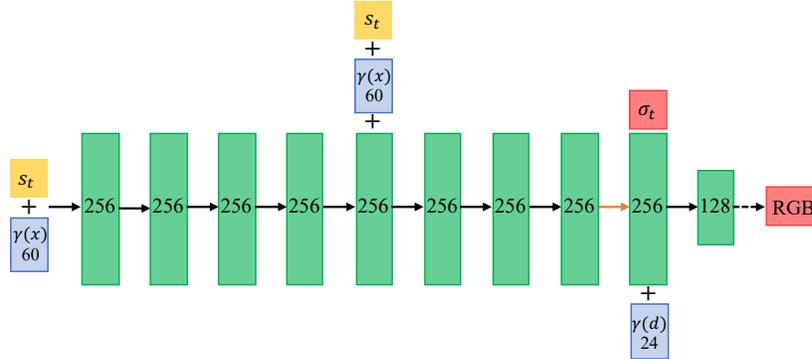


Fig. 9: **A visualization of our decoder network architecture.** All layers are standard fully-connected layers, black arrows indicate layers with ReLU activations, orange arrows indicate layers with no activation, dashed black arrows indicate layers with sigmoid activation, and “+” denotes vector concatenation. We concatenate the positional encoding of the input location  $\gamma(\mathbf{x})$  and our learned state representation  $\mathbf{s}_t$  and pass them through 8 fully-connected ReLU layers, each with 256 channels. We follow the NeRF [20] architecture and include a skip connection that concatenates this input to the fifth layer’s activation. An additional layer outputs the volume density  $\sigma_t$  (which is rectified using a ReLU to ensure that the output volume density is nonnegative) and a 256-dimensional feature vector. This feature vector is concatenated with the positional encoding of the input viewing direction  $\gamma(\mathbf{d})$ , and is processed by an additional fully-connected ReLU layer with 128 channels. A final layer (with a sigmoid activation) outputs the emitted RGB radiance at position  $\mathbf{x}$ , as viewed by a ray with direction  $\mathbf{d}$ , given the current 3D state representation  $\mathbf{s}_t$ .