ToolFlowNet: Robotic Manipulation with Tools via Predicting Tool Flow from Point Clouds

Anonymous Author(s)
Affiliation
Address
email

Abstract: Point clouds are a widely available and canonical data modality which conveys the 3D geometry of a scene. Despite significant progress in classification and segmentation from point clouds, policy learning from such a modality remains challenging, and most prior works in imitation learning focus on learning policies from images or state information. In this paper, we propose a novel framework for learning policies from point clouds for robotic manipulation with tools. We use a novel neural network, ToolFlowNet, which predicts dense per-point flow on the tool that the robot controls, and then uses the flow to derive the transformation that the robot should execute. We apply this framework to imitation learning of challenging deformable object manipulation tasks with continuous movement of tools, including scooping and pouring, and demonstrate significantly improved performance over baselines which do not use flow. We perform physical scooping experiments with ToolFlowNet and find that we can attain 78% scooping success. See https://tinyurl.com/toolflownet for supplementary material.

Keywords: Flow, Point Clouds, Tool Manipulation, Deformables

1 Introduction

Robotic manipulation has a long history [1, 2] with learning-based techniques recently becoming popular and effective for improving the generalization capabilities of robots for manipulation tasks as varied as grasping [3], reorienting [4], and tossing [5]. Data observations tend to be either images [6, 7, 8] or state information such as joint angles or end-effector poses [9], which are passed into a deep network to obtain an output vector encoding an action, typically representing a change in end-effector pose or joint angles. While these approaches have shown a wide range of successes, a fundamental limitation has to do with the nature of the observation. Using images requires projecting information into a 2D space which might lose valuable 3D information. Furthermore, learning from images in simulation often leads to a sim2real gap [10] in performance. Although it is easy to access the internal robot states such as joint angles, the robot does not necessarily have the ground-truth state of objects in the environment, which usually requires some complex state estimation systems. Moreover, it is hard to define a state for deformable objects like liquid and cloth [11, 12].

In this work, we propose a framework for learning robotic manipulation from point cloud observations. Point clouds are a canonical data modality and are widely available from camera sensors, providing valuable information about the structure of the 3D space [13, 14]. However, policy learning from point clouds has been less explored compared to learning from images or state, potentially owing to the difficulty of reasoning about raw 3D point cloud inputs. While there have been many proposed architectures which are specialized for learning from point clouds [13, 15, 14, 16, 17], these works tend to focus on computer vision tasks such as classification and segmentation. Policy learning from point clouds, while feasible in some contexts [18, 19], remains challenging.

We study learning from point clouds for robotic manipulation tasks with tools. The input data is a segmented point cloud which, for each point, contains its 3D coordinates and a one-hot vector

Submitted to the 6th Conference on Robot Learning (CoRL 2022). Do not distribute.

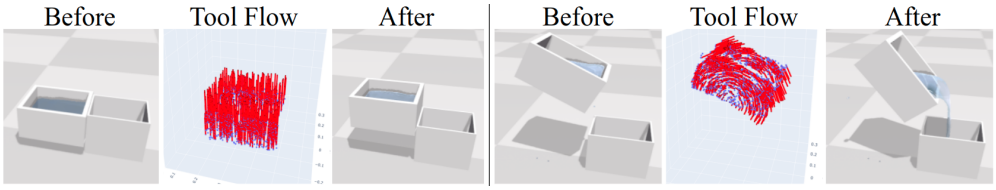

Figure 1: ToolFlowNet applied on a pouring task in simulation, where the tool is the box which contains water. Given a point cloud (colored blue), ToolFlowNet learns dense per-point flow vectors (colored red), which describe the intended 3D motion of each tool point. These are then converted to translation and rotation actions. Left: the tool moves upwards. Right: the tool rotates to pour water into the target container.

indicating the class of object the point belongs to. Our key insight is to use *dense representations* and *flow* to represent the tool action. We build upon dense point-cloud processing architectures [15] and train the model to predict per-point output values which we call *tool flow*. This represents the 3D movement of each tool point in a point cloud from one time step to the next, which is an instance of scene flow [20]. Our model is trained with Behavioral Cloning on tool flow data, which provides a dense per-point supervision. Given the set of tool flow vectors, we convert flow to an SE(3) transformation, which represents the actual action a robot would execute. We call this model ToolFlowNet and visualize it in Figure 1 for a pouring task in simulation. We compare this against non-dense methods which directly regress to an action and demonstrate the benefits of tool flow as an action representation. To summarize:

- We propose a general framework for learning from segmented point clouds for manipulation with tools by utilizing a novel architecture, ToolFlowNet, which predicts per-point tool flow vectors.
- We show how to train ToolFlowNet for imitation learning and explore different loss functions for training. We perform extensive ablation studies to validate these choices.
- We perform simulated imitation learning experiments on scooping and pouring tasks and show the benefits of using ToolFlowNet over baselines which do not use flow.
- We demonstrate ToolFlowNet achieves 78% success rate on 50 physical scooping trials.

2 Related Work

Point Clouds and Flow Researchers have proposed a variety of architectures specialized for learning from point clouds [13, 15, 14, 21, 17, 22, 23]. We aim to explore policy learning for robotic manipulation from point clouds, and the approach we propose is compatible with any architecture producing per-point outputs from point clouds. Optical flow [24, 25] and its 3D counterpart, scene flow [20, 26], are widely used in computer vision, particularly in autonomous driving setups where the objective is to associate the movement of each pixel (or a point in 3D space) from one image (or point cloud) to the next time step. We use flow as an action representation for robot manipulation, and our method can integrate prior flow estimation techniques.

Policy Learning from Point Clouds or Flow for Robotic Manipulation We are interested in learning from point clouds for robot manipulation, which has been explored in prior work for tasks such as grasping [19, 27], in-hand manipulation (by voxelizing) [18], visual navigation [28], and shaping 3D deformables [29]. Our work differs in that we instead study tasks that involve manipulating a tool in 3D space (with rotations) from point cloud observations, and where we use tool flow as the action representation to improve learning. While Qin et al. [30] extract tool point clouds and learn keypoints for grasping and manipulating tools, we instead predict dense tool flow for manipulating the tool. Some prior work has explored policy learning using *flow* for robot manipulation, such as for fabric folding [31], manipulating articulated objects [32, 33], and manipulating 3D deformables [34]. This work differs in that we propose a more general framework that does not assume a specific structure of the objects being manipulated, and which predicts flow on the tool a robot controls instead of flow on a target object. Furthermore, unlike prior work [35] which iteratively minimizes flow with pick and place actions, or other work [36] which uses optical flow on tactile sensors, we use flow to derive continuous tool motions in 3D space from visual input.

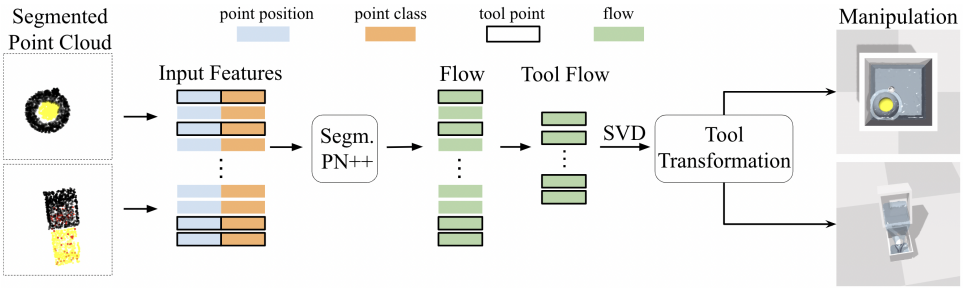

Figure 2: The proposed ToolFlowNet framework learns from segmented point clouds, which form the input to a dense point cloud network to produce per-point flow vectors. We extract just the tool points (bolded above for clarity) and use those tool points to determine the transformation that the robot should apply to the tool. See Section 3 for further details. Above, we show the ScoopBall and PourWater (see Section 4.1) tasks.

Deformable Object Manipulation We apply our proposed tool flow framework on tasks with continuous control of a tool for deformable object manipulation. Such manipulation is challenging for robots for reasons such as the difficulty in specifying a concise state representation for deformables and their complex dynamics [11, 12]. We test on scooping and pouring. Variants of these tasks have been studied in prior work. For example, [37] use scooping as an example application for task and motion planning, and [38] test scooping of granular media using a 2D image representation. Unlike these works, our approach is a general framework for robots performing continuous control of a tool to manipulate deformables in 3D space. Prior works [39, 40, 41, 42] propose methods to detect or model physics properties of granular media or liquids and test on scooping and pouring. In contrast, we propose a general method of predicting 3D tool flow which does not require modeling properties of deformables and is not specialized to scooping or pouring tasks.

3 Method: ToolFlowNet for Behavioral Cloning from Point Clouds

We consider policy learning from segmented point cloud observations. A segmented point cloud $\mathbf{P}_t$ at time t is an $N \times d$ array with N points, each with feature dimension d. The feature of the ith point $p^{(i)} \in \mathbf{P}_t$ consists of its 3D coordinate position and a one-hot vector indicating the class of the object to which $p^{(i)}$ belongs. For ease of notation, we suppress the time subscript t and the point index superscript i when the distinction is not needed. We study Behavioral Cloning (BC) [43] from segmented point clouds. BC assumes access to a dataset $\mathcal{D} = \{(\mathbf{o}_t, \mathbf{a}_t^*)\}_{t=1}^{M}$ of observation-action pairs $(\mathbf{o}_t, \mathbf{a}_t^*)$ from a demonstrator, where $\mathbf{o}_t$ indicates any type of observation (of which segmented point clouds are one example). BC performs supervised learning to learn a policy π with parameters θ such that the predicted action $\hat{\mathbf{a}}_t = \pi_\theta(\mathbf{o}_t)$ is close to the ground truth label $\mathbf{a}_t^*$. While prone to compounding errors at test time [44], BC has shown surprising effectiveness when compared to more complex learning-based algorithms in some robotic manipulation contexts [45, 46], which motivates further study on how it can be done with point cloud observations.

3.1 Tool Flow As Action Representation

We propose to use *tool flow* as an internal representation for the action, where a flow $f^{(i)} \in \mathbb{R}^3$ associated with point $p^{(i)}$ is a 3D vector, and we interpret flow vectors corresponding to the tool points as representing how each tool point will move in 3D space as a result of "applying" the flow. To form the policy π_θ, we use a dense point cloud network (such as a segmentation PointNet++ [15]), which given an input point cloud $\mathbf{P}$ produces per-point outputs. With an $(N \times d_1)$-sized point cloud $\mathbf{P}$ as input, the output has dimension $(N \times d_2)$, where d_2 is the output dimension (in our case, $d_2 = 3$). We then extract the output for the subset of $N' \leq N$ points in $\mathbf{P}$ corresponding to all points on the *tool*, while ignoring points belonging to other object classes. This results in a set of predicted 3D tool flow vectors $\mathbf{F} = \{f^{(i)}\}_{i=1}^{N'}$ with $f^{(i)} \in \mathbb{R}^3$ for each tool point.

Suppose that, at time t, the expert applies an action to the tool which is given by a transformation $\mathbf{T}^* \in SE(3)$. Let $\mathbf{P}_{\text{tool}} \subseteq \mathbf{P}$ be the set of 3D points on the tool. Then the ground-truth tool flow is given by $\mathbf{F}_{\text{gt}} = \mathbf{T}^*\mathbf{P}_{\text{tool}} - \mathbf{P}_{\text{tool}}$. Thus, there is a one-to-one correspondence between the

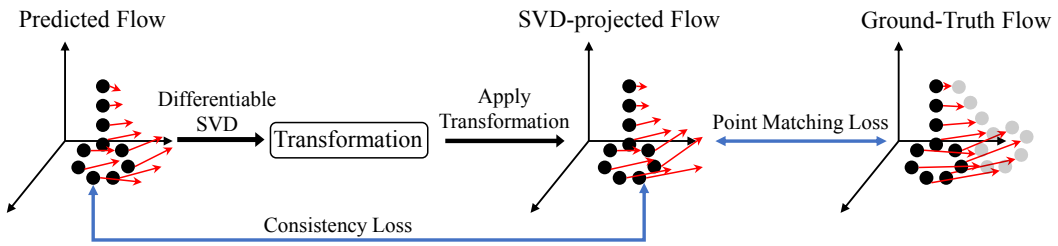

Figure 3: A visualization of the proposed point matching loss (Eq. 1) and consistency loss (Eq. 2). The black points visualize a simplified ladle's point cloud, and the red arrows represent the flows on the tool points.

transformation $\mathbf{T}^*$ and flow $\mathbf{F}_{\text{gt}}$; nonetheless, we show in Section 4 that estimating the tool flow leads to improved performance compared to estimating the transformation $\mathbf{T}^*$ directly.

Given the set of 3D tool flow vectors $\mathbf{F}$, the next step is to extract a single overall action $\mathbf{a}$, where $\mathbf{a}$ is a transformation that represents the change in translation and rotation of the tool's pose. To compute the action, we consider the tool point clouds $\mathbf{P}_{\text{tool}}$ and $\mathbf{P}'_{\text{tool}} = \mathbf{P}_{\text{tool}} + \mathbf{F}$, where we move each point based on its estimated flow. Our objective is to estimate the best-fit tool transformation $\mathbf{T} = (\mathbf{R}, \mathbf{t})$ which contains rotation and translation components, respectively, to align $\mathbf{P}_{\text{tool}}$ and $\mathbf{P}'_{\text{tool}}$, i.e., we want to find T to minimize $\|\mathbf{T}\mathbf{P}_{\text{tool}} - \mathbf{P}'_{\text{tool}}\|_2$. To obtain the rotation $\mathbf{R}$, we first center the two tool point clouds to obtain $\bar{\mathbf{P}}_{\text{tool}}$ and $\bar{\mathbf{P}}'_{\text{tool}}$. We then input the centered point clouds to a differentiable, parameter-less Singular Value Decomposition (SVD) layer [47, 48] which computes the rotation which best aligns $\bar{\mathbf{P}}_{\text{tool}}$ and $\bar{\mathbf{P}}'_{\text{tool}}$ with respect to mean square error (MSE). The change in translation $\mathbf{t}$ can then be computed as $\mathbf{t} = C(\mathbf{P}'_{\text{tool}}) - \mathbf{R}C(\mathbf{P}_{\text{tool}})$, where $C(\mathbf{P})$ denotes the centroid of the point cloud $\mathbf{P}$. By combining the translation and rotation components, we produce the full transformation $\mathbf{T}$, which we treat as our action representation for the policy. The outputs for the non-tool points are not supervised. We call the resulting point cloud-to-action system as ToolFlowNet (see Figure 2), which can be used by a robot for manipulation.

3.2 Imitation Learning via Tool Point Matching and Consistency Losses

Point Matching Loss: Given the policy's predicted action $\hat{\mathbf{a}} = \pi_\theta(\mathbf{o})$, we use a point matching loss to reduce the discrepancy between $\hat{\mathbf{a}} = (\hat{\mathbf{R}}, \hat{\mathbf{t}})$ and the ground truth action $\mathbf{a}^* = (\mathbf{R}^*, \mathbf{t}^*)$. Given the predicted action, the transformation $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ is applied on all the original N' tool points in the point cloud, and the loss function L_{point} computes the Euclidean distance between the tool points transformed using the predicted action $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ versus the tool points transformed using the ground-truth action $(\mathbf{R}^*, \mathbf{t}^*)$:

$$L_{\text{point}}(\mathbf{P}, \hat{\mathbf{a}}, \mathbf{a}^*) = \frac{1}{N'} \sum_{i=1}^{N'} \|(\hat{\mathbf{R}}p^{(i)} + \hat{\mathbf{t}}) - (\mathbf{R}^* p^{(i)} + \mathbf{t}^*)\|_2, \tag{1}$$

where $p^{(i)}$ iterates through the N' tool points in $\mathbf{P}_{\text{tool}} \subseteq \mathbf{P}$, and we interpret $\hat{\mathbf{R}}$ and $\mathbf{R}^*$ as representing 3×3 rotation matrices. Prior work on 6D pose estimation [49, 50, 51] has used variants of this loss function to jointly optimize for translation and rotation as compared to balancing the weights on separate translation and rotation losses. Our usage of L_{point} is similar to that in Wang et al. [19] where the matching loss is on tool points directly controllable by the robot.

Consistency Loss: While L_{point} should allow the policy π_θ to learn SE(3) pose changes (and thus, actions), its effect on optimizing the predicted flow vectors $\mathbf{F}$ happens via backpropagating through a differentiable SVD layer which "compresses" all predicted flow vectors to produce a single transformation $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$. This compression means that there could be significant noise in the individual flow vectors, even if they combine to form a reasonable action. Thus, we propose a consistency loss $L_{\text{consistency}}$ which serves as a regularizer to ensure that the *predicted* flow vectors are similar to their

induced flow vectors produced from the transformation encoded in $\mathbf{a}$. The loss is:

$$L_{\text{consistency}}(\mathbf{P}, \hat{\mathbf{a}}) = \frac{1}{N'} \sum_{i=1}^{N'} \|(\hat{\mathbf{R}}p^{(i)} + \hat{\mathbf{t}} - p^{(i)}) - f^{(i)}\|_2, \tag{2}$$

where for each of the N' tool points, we compute $\hat{\mathbf{R}}p^{(i)} + \hat{\mathbf{t}} - p^{(i)}$ as the induced flow from the predicted transformation $(\hat{\mathbf{R}}, \hat{\mathbf{t}})$ after applying the SVD layer, and $f^{(i)}$ is the flow predicted by the network before the SVD layer. Note that the ground truth transformation $(\mathbf{R}^*, \mathbf{t}^*)$ does *not* appear in this consistency loss. We combine this with the point matching loss L_{point} to obtain the final loss function to optimize the policy π_θ: $L_{\text{combo}} = L_{\text{point}} + \lambda \cdot L_{\text{consistency}}$, with hyperparameter λ controlling the weight of the consistency loss, which we set to $\lambda = 0.1$. See Figure 3 for visuals.

Additional Implementation Details: To obtain ground truth tool flow $\mathbf{F}_{\text{gt}}$ in simulation, we extract the action executed by the demonstrator, and we determine the 3D movement of each tool point as a result of applying the action. In physical settings, we scan the tool to obtain a 3D model, from which we extract tool point clouds $\mathbf{P}_{\text{tool}}$. We perform a similar calculation where we detect the transformation executed by the robot and apply it to obtain the flow for each tool point. This method to detect flow means that it reflects the "intended" action from the robot, which may differ from the true positions of the tool points in 3D space after the robot executes the action; for example, when a collision happens with a wall, the tool points might not move, even though the robot intended for them to move. We leave alternative techniques to extract tool flow to future work.

4 Experiments

4.1 Simulation Experiments

We build on top of SoftGym [52], which provides a suite of deformable manipluation tasks and uses NVIDIA FleX [53] as the underlying physics engine. We use the simulator to obtain ground-truth segmentation labels. For the tool, we use the "observable" point cloud at each time step, so there may be occlusions. We test on the following tool-based simulation tasks:

PourWater The agent controls a box which contains water and must pour the water into a fixed target box. We set a binary success based on if at least 75% of the water particles end in the target. In each episode, we vary the sizes of both boxes, the amount of water in the controlled box, and the starting distance between the boxes. The point cloud has three classes: the tool (i.e., the controlled box), the target box, and the water. The action space $\mathbf{a} = (\Delta x, \Delta y, \Delta\theta)$ is 3D and consists of the change in the x and y coordinates of the controlled box's center and the change in rotation $\Delta\theta$ around the center of the box. See Fig. 1 left for a visual of the translation action and Fig. 1 right for a visual of the rotation action. The proposed ToolFlowNet generates full 6D transformations, and we ignore the unused 3 action dimensions at test time. We script a demonstrator which moves the box towards the target and rotates to pour the water. PourWater was introduced as a task in SoftGym [52].

ScoopBall The agent controls a ladle, and each episode begins with the ladle above a box that contains water and a single ball floating on it. The objective is to use the ladle to scoop the ball, with a binary success based on if the ball is lifted above a certain height. Each episode has a different initial position of the ball on the water. The point cloud has two classes: points on the tool and points on the ball (we do not include the water). The action space $\mathbf{a} = (\Delta x, \Delta y, \Delta z, \Delta\theta)$ is 4D, which consists of the change in the coordinates of the ladle tip, and the change in rotation about the vertical axis $\Delta\theta$ around the ladle's tip. See Fig. 4 for visuals of the translation and rotation actions. We ignore the unused 2 dimensions from the output 6D transformation of ToolFlowNet to get a 4D action. We implement an algorithmic demonstrator which first lowers the ladle into the water, attempts to move the ladle so that its "bowl" is underneath the ball, then lifts the ball. The demonstrator continually rotates the ladle so that the ladle's handle is facing the direction of the ball.

4.1.1 Baseline Methods

We compare the proposed method with the following baselines (see Section 4.3 for ablations):

Method	Loss	Consist.? (Eq. 2)	Dense PN++?	Success (ScoopBall)	Success (PourWater)	Average
Direct Vector	MSE	✗	✗	0.396	0.618	0.507
Direct Vector	PM	✗	✗	0.396	0.022	0.209
Dense Transformation	MSE	✗	✓	0.554	0.762	0.658
Dense Transformation	PM	✗	✓	0.491	0.442	0.466
Image CNN	MSE	✗	✗	0.491	0.386	0.438
Image CNN DataAugm	MSE	✗	✗	0.617	0.287	0.452
ToolFlowNet, No Skip Conn.	PM	✗	✓	0.918	0.000	0.459
ToolFlowNet, MSE after SVD	MSE	✗	✓	0.190	0.563	0.376
ToolFlowNet, PM before SVD	PM	✗	✓	0.744	**0.850**	0.797
ToolFlowNet, No Consistency	PM	✗	✓	0.807	0.618	0.713
ToolFlowNet (Ours)	PM	✓	✓	**0.934**	**0.850**	**0.892**

Table 1: Behavioral Cloning (BC) results in simulation. The first 6 rows are baselines, the next 4 are ablations of our method, and the last row is our method. We report the loss functions used (PM is the loss in Eq. 1), whether the method uses a segmentation PointNet++ (i.e., a dense architecture), and the normalized success rates across both tasks over 5 independent BC runs. The last column averages the success across the two tasks.

- **Direct Vector, MSE**. Uses a *classification* PointNet++ to directly estimate a vector action (translation and axis-angle), and supervises with the MSE loss.
- **Direct Vector, Point Matching**. The same as the prior method, except instead of using an MSE loss on the action, we use the point matching loss from Eq. 1 on tool points.
- **Dense Transformation, MSE**. Uses a segmentation PointNet++, and directly predicts per-point 6D vectors (translation and axis-angle). Each point cloud has a designated point as the center of rotation for the tool, and we use the output corresponding to that point as the vector action and supervise with MSE. The outputs for the other points are not supervised. This baseline is designed to isolate any benefits from using the segmentation version of PointNet++ instead of classification.
- **Dense Transformation, Point Matching**. The same as the prior method, except instead of using an MSE loss on the single tool point vector, we use the point matching loss from Eq. 1.
- **Image CNN**. Uses a CNN with convolutional layers followed by dense layers to directly produce an action vector, and trains with the MSE loss. This network processes RGB images.
- **Image CNN (DataAugm)**. Same as Image CNN, but with randomized crops in training [54, 55].

4.1.2 Experiment Protocol and Evaluation

For each task, we generate a set of starting configurations and divide these into training and testing configurations. We use the task-specific algorithmic demonstrator to generate a fixed set of training demonstrations, where each demonstration has 100 time steps. We keep only the successful demonstrations for Behavioral Cloning. For a fair comparison, all methods train on the same set of 100 demonstrations. To account for noise in learning and nondeterministic FleX simulation, we perform 5 independent runs for all methods and average the performance among the 5 seeds. We standardize on 300 epochs of training for all methods which are sufficient for the evaluation losses to stabilize. We evaluate every 5 training epochs on 20 testing configurations and use the maximum success value (again, averaged over 5 seeds) across the full training history. We divide this by the success rate of the demonstrator to obtain the normalized performance.

4.2 Simulation Results

The results in Table 1 suggest that using ToolFlowNet outperforms the baselines in both ScoopBall and PourWater with the highest normalized performance of 0.934 and 0.850, respectively. Interestingly, it outperforms the CNN methods which process images, suggesting that these tasks have a 3D nature which makes it more natural to learn policies from point clouds. Figure 4 shows a qualitative example test-time rollout of ScoopBall from the learned policy using ToolFlowNet. The policy learns to lower the ladle, rotate it towards the ball, moves slightly towards the ball, lifts, and then rotates back to the starting orientation. Figure 4 also visualizes the policy's internal flow predictions (i.e., the per-point flow vectors $f^{(i)}$ before the SVD layer), showing that the network has learned surprisingly clean per-point tool flow vectors. Furthermore, as the agent controls the ladle at its upper tip, when rotating, the flow vectors also correctly predict longer flow vectors for the points located further away from the origin.

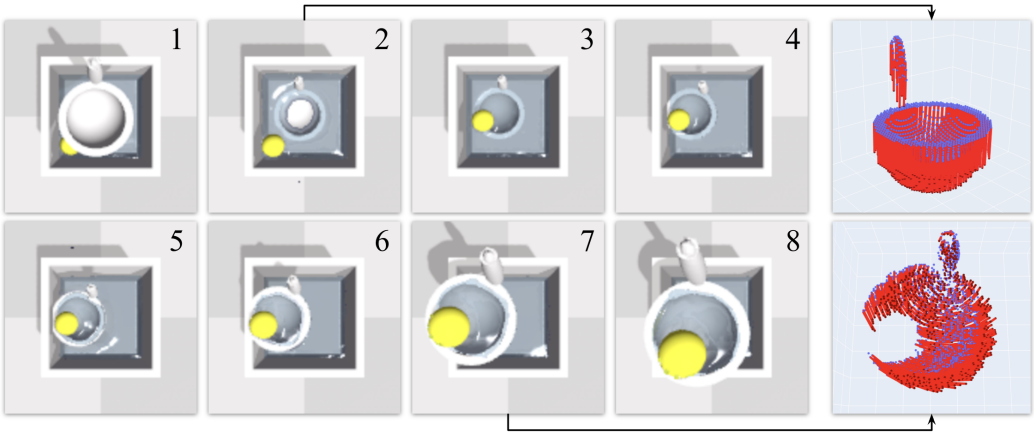

Figure 4: An example successful ScoopBall rollout by a learned Behavioral Cloning policy using ToolFlowNet. We show subsampled RGB frames for visual clarity, though the policy only uses point clouds as input. For two of the frames, we show the policy's tool flow visualizations. The policy lowers the ladle (frames 1-2), rotates and moves it in the direction of the ball (frame 3-4), lifts the ball (frames 5-7) and then rotates back to the target pose (frame 8). The policy's flow visualizations for frames 2 and 7 suggest the ability to learn downward and rotation movement, respectively. The predicted flow vectors, colored red, are slightly enlarged for clarity.

4.3 Ablation Studies of ToolFlowNet

We study ablations of ToolFlowNet to understand which components are most critical:

- **Tool Flow, No Skip Connections**: removes the skip connections in the segmentation PointNet++.
- **Tool Flow, MSE after SVD**: tests applying an MSE loss on the induced transformation from SVD instead of point matching. We still use the consistency loss (Eq. 2).
- **Tool Flow, Point Matching (PM) Before SVD**: tests using the PM loss (Eq. 1) before the SVD layer, so the loss does not back-propagate through the SVD layer.
- **Tool Flow, No Consistency**: tests removing the consistency loss (and just using Eq. 1).

We use the same experiment protocol as in Section 4.1.2 on both tasks. The results, also in Table 1, suggest strong benefits to using the point matching loss, the consistency loss, and the standard segmentation PointNet++ with skip connections. For example, in both tasks, performance is worse without using a consistency loss. The utility of some design choices may be more task-specific; removing skip connections leads to no successes on PourWater because removing it made the model unable to predict any rotations, while it is possible to succeed in ScoopBall without using rotations.

4.4 Additional Experiments: Modifying the Point Clouds

We next study Behavioral Cloning using ToolFlowNet when we alter the nature of the point clouds. To explore the potential for transfer to real settings with noisier sensor readings, we inject i.i.d. Gaussian noise to each point's 3D position in all training and testing data point clouds, testing $\sigma = 0.005$ in FleX's coordinate system units. This is 2.7% of the radius of the ladle's bowl's for ScoopBall and 2.1% of the average box length in PourWater. In a second experiment, we reduce the number of tool points in PourWater. We use 10 points, corresponding to keypoints located on the box, and treat those as the tool points in the segmented point cloud. This is in contrast to experiments reported in Table 1, where the PourWater data has an average of 633 tool points per observation. The results, shown in Table 2, suggest that ToolFlowNet can be robust to Gaussian noise; it matches the performance without noise on ScoopBall. Furthermore, for PourWater, using just 10 points on the box is sufficient to get 0.717 normalized performance, suggesting that ToolFlowNet can learn with fewer tool points, but performance is worse compared to 0.850 with the full tool.

Task	Experiment	Value
PourWater	Gauss. $\sigma = 0.005$	0.795
ScoopBall	Gauss. $\sigma = 0.005$	0.934
PourWater	10 Tool Points	0.717
PourWater	(From Table 1)	0.850
ScoopBall	(From Table 1)	0.934

Table 2: Modifying the point clouds.

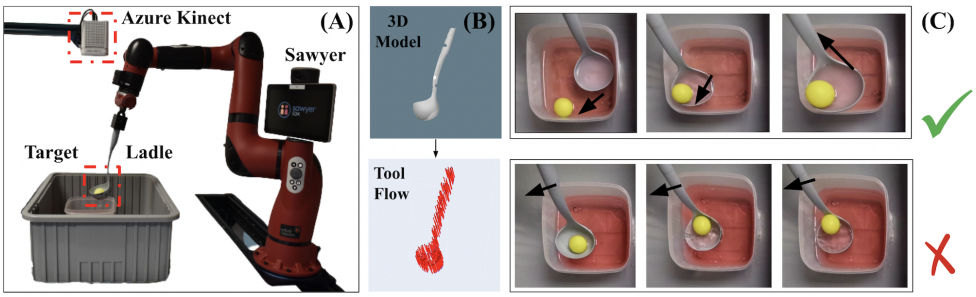

Figure 5: Physical experiments. (A) The Sawyer holds a ladle above a small box with water, which is enclosed in a larger gray box to contain spills. (B) The scanned 3D model of the ladle with representative tool flow visualization. (C) Example test-time trials with subsampled frames. Top row: successful tool movement towards and lifting the target. Bottom row: collision failure due to repeatedly pushing against the container.

4.5 Physical Scooping Experiments

We test ToolFlowNet in the real world using a Sawyer robot with a standard consumer ladle which we scan to obtain a 3D model, and a yellow ping-pong ball acting as the target item (see Figure 5). The ladle is attached to the Sawyer's end-effector. As in simulation, we obtain tool flow by recording the change in end-effector pose and applying the transformations on the tool point cloud. A Microsoft Azure Kinect camera captures top-down depth images to compute the ball's point cloud. A human operator manually moves the Sawyer's arm via direct touch to collect 62 demonstrations, with each comprising about 60 observation-action pairs. The demonstrations only describe translation motions, and we use the Sawyer's impedance controller to limit end-effector rotations. We use 52 demonstrations for training ToolFlowNet, and use the remaining 10 for monitoring evaluation MSE. We perform 50 physical scooping trials, where each trial begins with the human dropping the ping-pong ball on water, near the ladle's bowl. The trial is successful if the Sawyer raises the ball from the water surface to the height of the top of the smaller box. We categorize failures as "Collision Failures" if the ladle collides with the small box, and "Time-out Failures" if the robot is unable to scoop the ball within 120 time steps. Results in Table 3 suggest that the Sawyer achieves 39/50 successes (78%), with 6 collision failures and 5 time-out failures. See the Appendix for more details. In future work we will do physical experiments with more complex demonstrations.

ToolFlowNet in Real	#Trials
Successes	39/50
Collision Failure	6/50
Time-out Failure	5/50

Table 3: Physical scooping results.

4.6 Limitations and Failure Cases

In our experiments, we obtain the ground-truth tool flow data by applying the demonstrator's actions on a set of tool points and computing the change in the resulting tool point positions. This requires access to the demonstrator's action, and future work could relax this assumption by extracting tool flow when there is no explicit labeling of actions, e.g., a human providing a demonstration video. Possible approaches include using scene flow techniques.

A limitation of ToolFlowNet, particularly when applied to physical settings, is that the tool might be occluded. For physical scooping, we were unable to collect demonstration data similar to that in ScoopBall simulation due to occlusion of the tool from the camera by the robot arm. In future work, we will explore ways to address occlusions such as point cloud inpainting and tracking. Finally, we tested ToolFlowNet on two simulation tasks and scooping in real. We hope to test on more diverse tasks such as cloth or rope manipulation, and to address failures from the physical experiments.

5 Conclusion

In this work, we propose a general technique for policy learning from point clouds for tool-based manipulation tasks, which we demonstrate on scooping and pouring tasks. Our method, called ToolFlowNet, predicts per-point tool flow vectors which are converted into actions. We hope this research inspires future work on learning from point clouds.

References

[1] O. Kroemer, S. Niekum, and G. Konidaris. A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms. *arXiv preprint arXiv:1907.03146*, 2019.

[2] K. M. Lynch and F. C. Park. *Modern Robotics: Mechanics, Planning, and Control*. Cambridge University Press, 2017.

[3] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg. Learning Ambidextrous Robot Grasping Policies. *Science Robotics*, 4(26), 2019.

[4] OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba. Learning Dexterous In-Hand Manipulation. In *International Journal of Robotics Research (IJRR)*, 2019.

[5] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser. TossingBot: Learning to Throw Arbitrary Objects with Residual Physics. In *Robotics: Science and Systems (RSS)*, 2019.

[6] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end Training of Deep Visuomotor Policies. In *Journal of Machine Learning Research (JMLR)*, 2016.

[7] L. Pinto and A. Gupta. Supersizing Self-supervision: Learning to Grasp from 50K Tries and 700 Robot Hours. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[8] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang. Solving Rubik's Cube with a Robot Hand. *arXiv preprint arXiv:1910.07113*, 2019.

[9] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine. Soft Actor-Critic Algorithms and Applications. *arXiv preprint arXiv:1812.05905*, 2018.

[10] N. Jakobi, P.Husbands, and I. Harvey. Noise and the Reality Gap: The use of Simulation in Evolutionary Robotics. In *European Conference on Advances in Artificial Life*, 1995.

[11] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, and Y. Mezouar. Robotic Manipulation and Sensing of Deformable Objects in Domestic and Industrial Applications: a Survey. In *International Journal of Robotics Research (IJRR)*, 2018.

[12] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, J. Pan, W. Yuan, and M. Gienger. Challenges and Outlook in Robotic Manipulation of Deformable Objects. *arXiv preprint arXiv:2105.01767*, 2021.

[13] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[14] H. Zhao, L. Jiang, J. Jia, P. Torr, and V. Koltun. PointTransformer. In *International Conference on Computer Vision (ICCV)*, 2021.

[15] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Neural Information Processing Systems (NeurIPS)*, 2017.

[16] Y. Zhou and O. Tuzel. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[17] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic Graph CNN for Learning on Point Clouds. *ACM Transactions on Graphics (TOG)*, 2019.

[18] T. Chen, J. Xu, and P. Agrawal. A System for General In-Hand Object Re-Orientation. In *Conference on Robot Learning (CoRL)*, 2021.

[19] L. Wang, Y. Xiang, W. Yang, A. Mousavian, and D. Fox. Goal-Auxiliary Actor-Critic for 6D Robotic Grasping with Point Clouds. In *Conference on Robot Learning (CoRL)*, 2021.

[20] S. Vedula, S. Baker, P. Rander, R. Collins, and T. Kanade. Three-dimensional Scene Flow. In *International Conference on Computer Vision (ICCV)*, 1999.

[21] Z. Liu, H. Tang, Y. Lin, and S. Han. Point-Voxel CNN for Efficient 3D Deep Learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.

[22] W. Wu, Z. Qi, and L. Fuxin. PointConv: Deep Convolutional Networks on 3D Point Clouds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[23] X. Liu, M. Yan, and J. Bohg. MeteorNet: Deep Learning on Dynamic 3D Point Cloud Sequences. In *International Conference on Computer Vision (ICCV)*, 2019.

[24] B. K. Horn and B. G. Schunck. Determining Optical Flow. Technical report, USA, 1980.

[25] P. Fischer, A. Dosovitskiy, E. Ilg, P. Hausser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning Optical Flow with Convolutional Networks. In *International Conference on Computer Vision (ICCV)*, 2015.

[26] Z. Teed and J. Deng. RAFT-3D: Scene Flow using Rigid-Motion Embeddings. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[27] Y.-H. Wu, J. Wang, and X. Wang. Learning Generalizable Dexterous Manipulation from Human Grasp Affordance. *arXiv preprint arXiv:2204.02320*, 2022.

[28] K. Lobos-Tsunekawa and T. Harada. Point Cloud Based Reinforcement Learning for Sim-to-Real and Partial Observability in Visual Navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[29] B. Thach, B. Y. Cho, A. Kuntz, and T. Hermans. Learning Visual Shape Control of Novel 3D Deformable Objects from Partial-View Point Clouds. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.

[30] Z. Qin, K. Fang, Y. Zhu, L. Fei-Fei, and S. Savarese. KETO: Learning Keypoint Representations for Tool Manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

[31] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held. FabricFlowNet: Bimanual Cloth Manipulation with a Flow-based Policy. In *Conference on Robot Learning (CoRL)*, 2021.

[32] S. Pillai, M. R. Walter, and S. Teller. Learning Articulated Motions From Visual Demonstration. In *Robotics: Science and Systems (RSS)*, 2014.

[33] B. Eisner, H. Zhang, and D. Held. FlowBot3D: Learning 3D Articulation Flow to Manipulate Articulated Objects. In *Robotics: Science and Systems (RSS)*, 2022.

[34] B. Shen, Z. Jiang, C. Choy, L. J. Guibas, S. Savarese, A. Anandkumar, and Y. Zhu. ACID: Action-Conditional Implicit Visual Dynamics for Deformable Object Manipulation. In *Robotics: Science and Systems (RSS)*, 2022.

[35] A. Goyal, A. Mousavian, C. Paxton, Y.-W. Chao, B. Okorn, J. Deng, and D. Fox. IFOR: Iterative Flow Minimization for Robotic Object Rearrangement. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

[36] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez. Tactile-RL for Insertion: Generalization to Objects of Unknown Geometry. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[37] Z. Wang, C. R. Garrett, L. P. Kaelbling, and T. Lozano-Pérez. Learning Compositional Models of Robot Skills for Task and Motion Planning. In *International Journal of Robotics Research (IJRR)*, 2019.

[38] C. Schenck, J. Tompson, D. Fox, and S. Levine. Learning Robotic Manipulation of Granular Media. In *Conference on Robot Learning (CoRL)*, 2017.

[39] S. Clarke, T. Rhodes, C. G. Atkeson, and O. Kroemer. Learning Audio Feedback for Estimating Amount and Flow of Granular Material. In *Conference on Robot Learning (CoRL)*, 2018.

[40] C. Schenck and D. Fox. Towards Learning to Perceive and Reason About Liquids. In *International Symposium on Experimental Robotics (ISER)*, 2016.

[41] C. Schenck and D. Fox. Visual Closed-Loop Control for Pouring Liquids. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[42] C. Schenck and D. Fox. SPNets: Differentiable Fluid Dynamics for Deep Neural Networks. In *Conference on Robot Learning (CoRL)*, 2018.

[43] D. A. Pomerleau. Efficient Training of Artificial Neural Networks for Autonomous Navigation. *Neural Comput.*, 3, 1991.

[44] S. Ross, G. J. Gordon, and J. A. Bagnell. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.

[45] S. Dasari, J. Wang, J. Hong, S. Bahl, Y. Lin, A. Wang, A. Thankaraj, K. Chahal, B. Calli, S. Gupta, D. Held, L. Pinto, D. Pathak, V. Kumar, and A. Gupta. RB2: Robotic Manipulation Benchmarking with a Twist. *NeurIPS 2021 Datasets and Benchmarks Track*, 2021.

[46] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit Behavioral Cloning. In *Conference on Robot Learning (CoRL)*, 2021.

[47] O. Sorkine-Hornung and M. Rabinovich. Least-Squares Rigid Motion Using SVD. Technical report, ETH Zurich, 2016.

[48] J. Levinson, C. Esteves, K. Chen, N. Snavely, A. Kanazawa, A. Rostamizadeh, and A. Makadia. An Analysis of SVD for Deep Rotation Estimation. In *Neural Information Processing Systems (NeurIPS)*, 2020.

[49] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Robotics: Science and Systems (RSS)*, 2018.

[50] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox. DeepIM: Deep Iterative Matching for 6D Pose Estimation. In *European Conference on Computer Vision (ECCV)*, 2018.

[51] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese. DenseFusion: 6D Object Pose Estimation by Iterative Dense Fusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[52] X. Lin, Y. Wang, J. Olkin, and D. Held. SoftGym: Benchmarking Deep Reinforcement Learning for Deformable Object Manipulation. In *Conference on Robot Learning (CoRL)*, 2020.

[53] M. Macklin, M. Muller, N. Chentanez, and T.-Y. Kim. Unified Particle Physics for Real-Time Applications. *ACM Trans. Graph.*, 33(4), July 2014.

[54] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas. Reinforcement Learning with Augmented Data. In *Neural Information Processing Systems (NeurIPS)*, 2020.

[55] I. Kostrikov, D. Yarats, and R. Fergus. Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels. In *International Conference on Learning Representations (ICLR)*, 2021.