

Learning visual policies for building 3D shape categories

Alexander Pashevich*
 Inria[†]

Igor Kalevatykh*
 Inria[‡]

Ivan Laptev
 Inria[‡]

Cordelia Schmid
 Inria[‡]

Abstract—Manipulation and assembly tasks require non-trivial planning of actions depending on the environment and the final goal. Previous work in this domain often assembles particular *instances* of objects from known sets of primitives. In contrast, we aim to handle varying sets of primitives and to construct different objects of a *shape category*. Given a single object instance of a category, e.g. an arch, and a binary shape classifier, we learn a visual policy to assemble other instances of the same category. In particular, we propose a disassembly procedure and learn a state policy that discovers new object instances and their assembly plans in state space. We then render simulated states in the observation space and learn a heatmap representation to predict alternative actions from a given input image. To validate our approach, we first demonstrate its efficiency for building object categories in state space. We then show the success of our visual policies for building arches from different primitives. Moreover, we demonstrate (i) the reactive ability of our method to re-assemble objects using additional primitives and (ii) the robust performance of our policy for unseen primitives resembling building blocks used during training. Our visual assembly policies are trained with no real images and reach up to 95% success rate when evaluated on a real robot.

I. INTRODUCTION

Our daily physical activities such as cooking, dressing or navigation require complex sequences of actions which people successfully and seamlessly plan based on sensory input. Action planning typically depends on the goal and constraints provided by the environment. Despite extensive prior work, existing autonomous agents are still far from the human-level planning performance, especially in unknown and cluttered environments [11], [57].

Action planning is a hard problem due to the large action spaces, exponential complexity and partial observability [27], [28]. To simplify the problem, existing work on task planning [19], [37], [51], [56] typically operates in the state space assuming the full knowledge of the environment. While such an assumption can be practical in structured and controlled environments, full state reconstruction for common scenes remains a highly challenging problem [20]. Arguably, the precise recovery of scene parameters such as its geometry, composition, friction coefficients, etc., is more difficult than the primary planning task itself. It is therefore desirable to design sensor-based planning policies that do not rely on explicit scene geometry and full state estimation.

Vision-based control policies have recently become popular for robotic manipulation [2], [33], [60] and navigation [6],

*Equal contribution.

[†]University Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.

[‡]Inria, École normale supérieure, CNRS, PSL Research University, 75005 Paris, France.




Fig. 1: Primitives on the left are assembled by our learned policy into arches on the right. We assemble objects of similar shapes in the simulator and learn visual manipulation policies that can build real 3D shapes from unseen primitives resembling building blocks used during training.

[16]. While this line of work shows promise, it has mostly been applied to the low-level motion planning such as predicting next motion direction. In our work we aim to learn visual policies for high-level task planning. Given visual input, our policies generate sequences of picking, rotating and placing actions for building 3D shapes.

Object manipulation has a long history in robotics. In particular, assembling objects from a given set of primitives has been addressed, for example, in [5], [46]. Prior work in this domain often aims to build particular object instances for which the structure is pre-defined [52], specified by demonstrations [10], [25] or given by a goal image [29]. Here we go further and learn to assemble different objects of a shape category. Such a task is significantly more complex compared to building particular object instances as it requires generalization to varying sets of building primitives. Moreover, we show empirically that our method is able to generalize to new primitives unseen during training (see Fig. 1).

Our approach contains two stages as illustrated in Fig. 2. In the first stage we discover new object instances and learn their assembling policies in state space. To this end, we propose a disassembling procedure and generate assembly trajectories by (a) unbuilding objects and (b) reverting action sequences. We then learn a value function and apply it to build new object instances. We iterate the disassembling and learning steps to obtain a policy assembling a 3D shape.

In the second stage we assemble objects given images of observed scenes. We render states from assembly trajectories obtained in the first stage and learn visual assembly policies

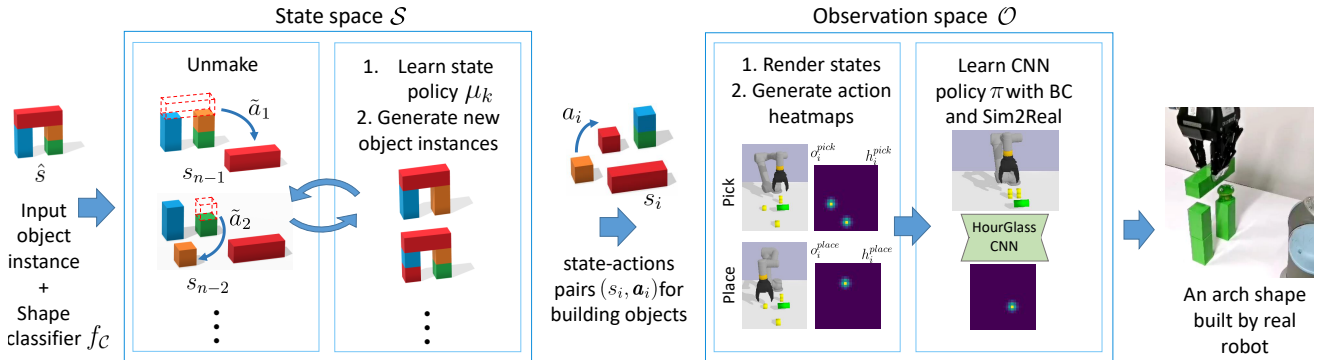


Fig. 2: Method overview. Given an example object and a shape classifier on the left, our method generates new objects with similar shape and discovers action sequences for building these objects in the state space. Using a large set of generated state-actions pairs (s_i, a_i) , we render states s_i as realistic observations o_i . We also generate 2D heatmaps $(h_i^{\text{pick}}, h_i^{\text{place}})$ encoding source and target locations and orientations of one or several primitives. Heatmaps can represent multiple hypothesis for the next action when several identical primitives are used or multiple object instances can be assembled. Positions on our 2D heatmaps correspond to positions on the 2D surface of a table, hence, the identified local maxima on heatmaps can be used to control the robot. As the last step of our method, we train a Behaviour Cloning policy π to predict h_i^{pick} and h_i^{place} from observations. We train policies with HourGlass CNN [39] and sim2real augmentation [42]. The learned policy is directly transferred to a real robot which assembles 3D objects from primitives on the table.

with Behaviour Cloning (BC) [45]. To enable predictions of multiple valid actions at any given time, we propose a heatmap representation for the output of visual policies. While all our policies are learned in a simulated environment, we enable their direct transfer to a real robot using sim2real augmentation [42].

As main contributions of this work, we (i) propose a novel disassembly algorithm for building shape categories in state space, (ii) design and learn visual policies with heatmap outputs to address multi-modality of predicted actions, (iii) demonstrate a successful application of the method to a new task of building shape categories on a real UR5 robot. Moreover, our policies are learned with no human demonstrations, can re-assemble partially built objects, and adapt to unseen primitives resembling building blocks used during training.

II. RELATED WORK

Assembly tasks such as constructing IKEA furniture [52] remain to be a hard robotics challenge. Learning-based methods usually address simpler tasks such as cube stacking [40], [49]. Duan et al. [10] use demonstrations and attention modules to build a tower instance shown by an expert. Janner et al. [29] learn an object-centric representation of the scene to reproduce a tower instance from a goal image with an MPC-like control. Huang et al. [25] train a Graph Network to build a tower instance specified by a demonstration. We go beyond specific object instances and aim to assemble multiple objects from a given shape category, such as an arch. Moreover, we learn to build objects from different sets of possibly unseen primitives. To facilitate the learning, we propose to use disassembling to generate assembly trajectories. While the idea of reversible actions has been explored e.g., in [15], [24], [38], [54], our method differs from the work on disassembling object instances [59] by computing multiple disassembly paths and accounting for alternative valid actions.

Our work is related to methods of Task and Motion Planning (TAMP) [17], [35], [51], [56]. Long-term task planning prohibits costly rollouts, hence, TAMP methods deploy

preconditions and postconditions to actions and optimize symbolic planners [14]. While some of these methods solve impressive tasks, conditions require manual and task-specific design [12], [22], [43]. Moreover, TAMP methods typically operate in state-space [18], hence, their generalization to sensor-based input in the real world requires non-trivial scene understanding [9]. In our work we learn visual policies and directly predict control sequences from image inputs.

Convolutions Neural Networks (CNNs) have significantly advanced visual recognition [23], [39], [48] and robotics, for example in tasks such as tossing objects [60], cube stacking [46], grasping [31] and opening doors [21]. Direct methods for visual control avoid explicit scene reconstruction and derive actions directly from image observations. Such methods typically use Reinforcement Learning (RL) [30] with auxiliary rewards [49] or Imitation Learning (IL) [4], [50] relying on large amounts of demonstrations [47], [61]. The complexity of problems addressed by direct methods is typically limited by the task length and the number of manipulated objects [47], [49], [61]. Indirect methods first estimate scene parameters [3], [13] such as object positions and orientations, and then deploy state-based planning strategies. Scene reconstruction from images, however, might be a more challenging task than solving a control task itself [20], [44]. We avoid drawbacks of direct and indirect methods and first solve the task in the simulated state space. We then use obtained solutions as automatic supervision for learning visual policies in the observation space. Inspired by [33], [36], we render states and train visual policies for a real robot using BC [45] and sim2real [42], [55].

III. APPROACH

We address the problem of building a 3D object shape by manipulating a set of available primitives with a robot. The configuration of the primitives on the table defines the state, $s \in \mathcal{S}$. We assume to have access to a shape classifier function $f_C : \mathcal{S} \rightarrow \{0, 1\}$. We define the shape as a subset of the state space, $\mathcal{C} = \{s \in \mathcal{S} | f_C(s) = 1\}$. Given f_C and a single shape instance $\hat{s} \in \mathcal{C}$, our method learns a visual policy π that generates a robot action given a camera

observation $o \in \mathcal{O}$. The resulting sequence of actions is then used to assemble a shape instance from available primitives. Our policy operates in the observation space \mathcal{O} , i.e., it only has access to the image of a current scene before deciding on the next action. Moreover, the policy is expected to build new object configurations from unseen primitives that resemble building blocks used during training.

A. Overview of the method

Our method has two stages: (i) generating action sequences in the state space and (ii) learning visual policies in the observation space. We use a simulated environment for both stages, however, our visual policies are trained with sim2real data augmentation [42] and directly transfer to the real robot. The overview of the method is presented in Fig. 2.

The first stage aims to find new shape instances and to construct action sequences for building them. It takes as input one 3D shape and a shape classifier. We propose to use an *unmake* procedure to generate valid action sequences. We disassemble the given shape instance and interpret the resulting sequences as reversed assembly demonstrations. We disassemble objects in multiple ways to find all assembly actions that are possible in the same state. We refer to the shape classifier f_C as a sparse reward signal which we use to learn a state-value function V_k . We generate new shape instances using a state policy μ_k which is greedy with respect to the learned V_k . For a fixed number of iterations, we repeat the unmake procedure using the new set of instances and train an updated value function V_{k+1} . This part of our approach is described in Section III-B.

In the second stage we learn a visual policy that infers appropriate actions from image observations. We convert states s_i into observations $o_i = R(s_i)$ using a graphics renderer R and train a CNN policy π with Behaviour Cloning (BC). Given the assembly trajectories produced by the first stage, each state is associated with a valid set of actions \mathbf{a}_i that we turn into a heatmap h_i to predict all possible actions simultaneously. Then π is trained in a supervised manner using observation-heatmap pairs (o_i, h_i) . This part of our approach is described in Section III-C.

B. Building objects in state space

We define the full state of our environment by the vector $s = (x^1, \dots, x^m)^T$ with $x \in \mathbb{R}^{12}$ representing parameters for m primitive shapes (our building blocks) in the scene. Each primitive x is defined by three position coordinates, three orientation angles in the 3D space, three spatial extents (width, height and depth), and three color channels (the building order may depend on the color, see Section IV-B). A robot action $a \in \mathcal{A}$ corresponds to a high-level skill of picking, placing and rotation of a primitive: $a = (x, p, o)$ where $x \in \{x^1, \dots, x^m\}$ is the primitive to pick, $p \in \mathbb{R}^3$ is the position to place it, and $o \in \mathbb{R}^3$ is the orientation x is placed in. We restrict orientations of primitives to the three axis-parallel directions and assume that all primitives are located on the surface of a table or on top of each other. Assuming access to the simulator T , applying action a_t in the state s_t would result into $s_{t+1} = T(s_t, a_t)$.

Building an object from a given set of primitives requires finding an appropriate sequence of actions




Fig. 3: The proposed unmake procedure disassembles objects in multiple ways (left) and generates set of pick & place actions available from each state (right). For visual policies, we use the heatmap representation to predict all the actions simultaneously.

$a_{1:n} = \{a_1, \dots, a_n\}$ that transforms an initial state s_0 into the desired shape state $s_n \in \mathcal{C}$. Finding a correct sequence $a_{1:n}$ is not trivial even in the state space. Given the large space of possible actions and the exponential growth of the number of action sequences depending on n , the naive brute-force search works only for building simple objects.

Making objects by unmaking. For an object example defined by the state $\hat{s} \in \mathcal{C}$ we propose to find valid building sequences $a_{1:n}$ via disassembling or *unmaking* \hat{s} . We first find a sequence of unmake actions $\tilde{a}_{1:n} = \{\tilde{a}_1, \dots, \tilde{a}_n\}$, where \tilde{a}_i moves a random non-blocked primitive on an object* to a random non-occupied location on the table. If m is the number of primitives constituting object \hat{s} , we can disassemble \hat{s} by a sequence of $n = m - 1$ unmake actions. We call an action invertible if for any action \tilde{a} such that $T(s_i, \tilde{a}) = s_j$ there exists an inverse action $a = \tilde{a}^{-1}$ such that $T(s_j, a) = s_i$. We assume our unmake actions to be invertible and obtain a valid sequence of actions for building \hat{s} as $a_{1:n} = \{\tilde{a}_n^{-1}, \dots, \tilde{a}_1^{-1}\}$.

We use the above randomized unmake procedure to generate a large and diverse set of valid action sequences. While each action a_i in such sequences is associated with a single initial state s_i , there might be *multiple correct assembly actions* in the state s_i . For example, when building an arch, the same cube could be placed both to the left and to the right pillars (see Fig. 3). To each state s_i , we associate a set of actions \mathbf{a}_i available in all states which differ only in positions of primitives located on the table surface (denoted as \approx), namely, $\mathbf{a}_i = \{a_j | s_j \approx s_i\}$.

Given M action sequences of length N we collect a dataset \mathcal{D}_μ with state-actions pairs $\mathcal{D}_\mu = \{(s_i, \mathbf{a}_i)\}_{i=1, \dots, M * N}$ that can be readily used to train a policy for assembling an object instance defined by state \hat{s} .

Finding new object instances. To generalize our policies to build new objects of the similar shape given any set of primitives, we construct new instances through learning a value function. The value function $V_k : \mathcal{S} \rightarrow \mathbb{R}$ estimates the sum of discounted rewards $r : \mathcal{S} \rightarrow \mathbb{R}$ under an assembly policy $\mu_k : \mathcal{S} \rightarrow \mathcal{A}$ where $V_k(s) = \mathbb{E}_{\mu_k} \left[\sum_j \gamma^j r(s_{t+j}) | s_t = s \right]$ and $\gamma \in [0, 1)$ is the discount factor. We define the reward with the shape classifier $r(s_t) = f_C(s_t)$ which makes the value function learn the actions required to build the shape. Given a learned value function V_k , we deduce the assembly policy

*By “non-blocked primitives on objects” we refer to primitives that are not on the table surface and that can be freely lifted above their current positions. Such primitives can be easily derived from s .

by choosing a_{t+1} as $\mu_{k+1}(s_t) = \arg \max_a V_k(T(s_t, a))$. We sample an initial state s_0 with a random set of primitives and apply μ_{k+1} to choose an action sequence $a_{1:n}$ resulting in an object of the desired shape: $f_C(s_n) = 1$. We record all instances of the target shape found by assembling is a set \mathcal{C}_k^V .

Learning value function. The value function V_k is learned iteratively using instances found with the greedy policy μ_{k-1} . In the first stage of training, we learn V_0 using the input instance \hat{s} only. We run the unmake procedure for all discovered instances in \mathcal{C}_k^V . Given a sequence of state-actions pairs $\{(s_1, \mathbf{a}_1), \dots, (s_n, \mathbf{a}_n)\}$, the value function estimate for state s_i is $\hat{V}(s_i) = \gamma^{n-i}$. We also estimate values of states obtained by applying random disassembly actions \tilde{a}_j to trajectory states s_i : $s_i^j = T(s_i, \tilde{a}_j)$. The value estimate $\hat{V}(s_i^j)$ is known if there exists $s_j \in \{s_1, \dots, s_n\}$ such that $s_i^j \approx s_j$. Otherwise, we set the value as $\hat{V}(s_i^j) = \gamma \hat{V}(s_i)$ which means that we can reach s_i from s_i^j using a single assembly action \tilde{a}_j^{-1} . We record all state-value pairs to the dataset \mathcal{D}_V^k and learn the value function by minimizing the loss $\hat{\eta}_k = \arg \min_{\eta} \text{MSE}(V_k(s_i), \hat{V}(s_i))$, where V_k is implemented as a fully connected neural network with parameters $\hat{\eta}_k$ and MSE is the mean square error. Once the training is converged, we discover new shape instances with μ_k , unmake them, recollect \mathcal{D}_V^{k+1} and learn V_{k+1} . After K phases, we run the unmake procedure on the set of discovered instances \mathcal{C}_K^V and record all state-actions pairs to the dataset \mathcal{D}_μ . The overview of our approach in the state space is illustrated by the left part of Fig. 2 and lines 14-28 of Algorithm 1 of [41].

C. Learning in observation space

We want to learn a visual policy π for assembling objects from diverse sets of primitives by a real robot. The sole input of the policy is the camera observation of the scene. We learn the image-action association with a supervised learning approach, Behaviour Cloning (BC) [45], where we obtain supervision with solutions found in the state space. Given the dataset \mathcal{D}_μ with state-actions pairs $\{(s_i, \mathbf{a}_i)\}$, we use a `pybullet` [8] graphics renderer R to generate an RGB-D image $o_i = R(s_i)$ for each state s_i in \mathcal{D}_μ . In order to allow multiple actions for each observation o_i , we generate an action-heatmap $h_i \in \mathcal{H}$ given the list of actions $\mathbf{a}_i = \{a_i^1, \dots, a_i^l\}$. We record the observation-heatmap pairs to the dataset \mathcal{D}_π . The policy $\pi : \mathcal{O} \rightarrow \mathcal{H}$ is implemented as a CNN and is trained to predict correct action-heatmaps $\pi(o_i) = h_i$ for all $(o_i, h_i) \in \mathcal{D}_\pi$. We show an advantage of the heatmaps-based architecture over a network that directly predicts positions and orientations in Section IV-D.1.

For the task of building objects with a real robot, we consider separate *pick* and *place* actions that are parameterized by positions and orientations of primitives on the 2D plane of a table. For simplicity, we assume that the elevation of a primitive above the table can be estimated by external means such as an overhead depth camera or a force-feedback sensor of the robot arm. We define the output of our policy by distributions over 2D positions on the table plane and 3 possible orientations of primitives on the table. We represent such distributions as heatmaps $(h_i^{\text{pick}}, h_i^{\text{place}})$ corresponding to the source and target parameters of primitives. Our heatmaps




Fig. 4: Visualization of predictions of visual policies using HourGlass (green blobs) and ResNet-18 (red crosses) architectures. The policies are trained to build towers and should start by placing a yellow cube on the green one. ResNet predicts the correct picking location when all the cubes have distinct colors (left). Once identical yellow cubes are introduced to the scene (right), ResNet fails to choose between them and predicts an averaged location. HourGlass locates all the cubes correctly in both cases.

are 4-channel images with one channel representing position distribution and three other channels representing orientation. The placing positions and orientations might depend on the picked primitive, hence, we predict pick and place action-heatmaps sequentially.

We render separate observations $(o_i^{\text{pick}}, o_i^{\text{place}})$ for pick and place action-heatmaps $(h_i^{\text{pick}}, h_i^{\text{place}})$ respectively. For each list of pick-place actions \mathbf{a}_i , we render the observation $o_i^{\text{pick}} = R(s_i)$ and define h_i^{pick} as an image with Gaussian distributions around positions of all picked objects by \mathbf{a}_i . For each picked object, we render the observation $o_i^{\text{place}} = R(s_i)$ with the robotic arm picking this object and create the heatmap h_i^{place} with all possible placing positions in \mathbf{a}_i (see Fig. 3). We record all observation-heatmap pairs in the dataset \mathcal{D}_π .

Heatmaps are often used as CNN outputs in tasks such as human pose estimation [39], [58] and segmentation [34]. We follow [39] and use a HourGlass CNN architecture for heatmap prediction $\pi(o) \mapsto h$. Given \mathcal{D}_π , we train π by minimizing the loss $\hat{\theta} = \arg \min_{\theta} \text{MSE}(\pi_{\theta}(o_i), h_i)$, where $\hat{\theta}$ are parameters of the HourGlass CNN. We obtain parameters for the pick and place actions by maximizing the obtained location heatmaps over the 2D space and then maximizing orientations over the 3 channels at the selected location. The overview of our policy learning is illustrated by the right part of Fig. 2 and lines 30-41 of Algorithm 1 of [41].

IV. EXPERIMENTS

In this section we evaluate our approach both in simulation and on a real UR5 robot. We start with implementation details in Section IV-A and present tasks used for evaluation in Section IV-B. Section IV-C confirms the importance of learning the state-value function for efficient task solving. Section IV-D evaluates visual policies trained to solve tasks in the observation space. We validate our proposed network architecture and highlight the importance of the disassembling procedure. Additional qualitative results are available in the Appendix of [41] and on the project web-page [1].

A. Implementation details

We control a 6-DoF UR5 robotic arm with a 3 finger Robotiq gripper. In simulation, we model the robot and its environment with the `pybullet` physics simulator [8]. Given the positions and orientations of primitives to be manipulated, we use standard path planing methods [32],

<i>Arch, state space</i>	3U	4U	5U
Random	3.4e3 ± 5.0e3	1.5e4 ± 1.2e4	6.1e4 ± 9.6e4
MCTS [7]	6.3e2 ± 4.7e2	7.0e3 ± 6.2e3	1.6e4 ± 2.4e4
Ours (state policy)*	4.0e0 ± 0.7e0	5.7e0 ± 0.7e0	6.9e0 ± 1.8e0
Oracle	4.0e0 ± 0.7e0	5.7e0 ± 0.7e0	6.4e0 ± 1.0e0

TABLE I: Average amount and standard deviation of steps required to build an arch shape for our method, random exploration and MCTS. *The simulation steps used for training of our method are not included.

<i>Tower, simulation</i>	3 cubes	5 cubes	7 cubes
ResNet-18 [23]	99.2	1.4	0.0
HourGlass [39]	99.6	97.2	94.8
<i>Tower, real</i>	3 cubes	5 cubes	7 cubes
HourGlass [39]	95.0	90.0	90.0

TABLE II: Success rates of visual policies (in percent) trained to build tower instances of 3, 5 and 7 cubes. On the real robot, the policies are evaluated using 20 trials.

[53] to implement the pick and place actions. The elevation of primitives above the table surface is obtained with a Microsoft Kinect-2 camera located above the table.

Our observations are depth images recorded with another Kinect-2 camera placed in front of the robot arm. We use the same parameters for real and simulation cameras (location and calibration). Visual policies receive the depth image and color segmentation masks corresponding to the colors of primitives. While visual policies in simulation have average errors of less than 5mm, the sim2real gap increases this value up to 2-3cm on the real robot. Given that stacking multiple primitives requires high precision, we apply a correction procedure using the depth camera as explained in Appendix of [41].

The neural network V_η for Value Function has five fully-connected layers with 128 hidden units, ReLU activations and Batch Normalization [26]. We train V_η for 20 iterations

<i>Arch, simulation</i>	3U	4U	5U
	<i>category</i>		
Single unmake trajectory	98.6	92.8	69.6
Multiple unmake trajectories	99.0	98.8	95.6

TABLE III: Success rates of visual policies trained to build the arch shape category of different heights. The policies are trained on trajectories obtained by disassembling the same object once or multiple times.

<i>Arch, simulation</i>	<i>instance</i>			<i>category</i>		
	3U	4U	5U	3U	4U	5U
Instance policies	99.4	97.8	96.8	61.4	54.4	32.6
Uni-height policies	99.6	98.2	98.0	99.0	98.8	95.6
Multi-height policy	97.4	96.4	95.4	97.8	95.4	94.0

TABLE IV: Success rates of visual policies trained by disassembling only the input instance (top) and instances found by state policies. The state policies are trained on arches of the same height (middle) and arches of heights 3-5U (bottom). The policies assemble arches given a fixed set of primitives (left) or various configurations of primitives (right). While instance and uni-height policies need to be trained for each given arch height, a single multi-height policy can assemble arches of various heights.

<i>Arch, real</i>	<i>normal scenario</i>			<i>re-assembling</i>		
	3U	4U	5U	3U	4U	5U
Uni-height policies	95.0	80.0	75.0	90.0	75.0	70.0
Multi-height policy	85.0	80.0	75.0	90.0	80.0	80.0

TABLE V: Success rates of visual policies assembling arches of different heights on a real robot. Over 20 trials, the policies are evaluated on a normal Arch task and a "re-assembling" scenario when the agent needs to adapt to additional primitives introduced to the scene after the arch had already been built.

<i>Arch, real</i>	set 1	set 2	set 3
	<i>novel primitives</i>		
Detect + MCTS [7]	60.0	35.0	20.0
Our method	90.0	80.0	55.0

TABLE VI: Success rates of our method and an MCTS baseline for assembling arches using novel primitives on a real robot. The arches built with primitives from sets 1, 2 and 3 are shown in Fig. 5 (last column, top), Fig. 5 (last column, bottom) and Fig. 1 (bottom) correspondingly.

of 30 epochs each using Adam and LR=1e-3. The CNN π_θ contains one HourGlass [39] module which we compare to ResNet [23] in Section IV-D.1. The spatial dimensions of HourGlass input and output are 256x256 and 64x64 pixels respectively. We train π_θ using Adam and LR=2.5e-4 for 50 epochs. For both value and visual networks, we use datasets of size 200,000 value-state and heatmap-observation pairs correspondingly. To enable the transfer of policies to the real robot, we use sim2real [42] to augment synthetic depth maps during training. Color segmentation masks are augmented with Bernoulli noise. During rendering, we also randomize shapes of primitives by adding noise to spatial coordinates of points that define cube meshes. We use 500 episodes for evaluation in simulation and 20 trials on the real robot.

B. Tasks

Tower. The goal of the agent is to stack cubes in a specific order of colors (see Fig. 4). In the beginning of the task, green, yellow and red cubes of size 1 unit (1U) are randomly distributed on the surface of a table. The unit corresponds to a physical size of 4.5 cm. The lowest cube is always green, the rest of the tower is defined as alternating yellow and red cubes. We use the Tower task to compare HourGlass and ResNet architectures in Section IV-D.1.

Arch. The agent needs to use all primitives available on a table to build an arch (see Fig. 5). The construction primitives are cubes of size 1U and beams of length 2U and 3U. The arch shape category is defined as two symmetrical pillars with a bar bridging them. For example, pillars of an arch could be constructed from a 3U beam, three cubes or one cube and a 2U beam. Initially, all primitives are randomly distributed on the surface of the table. The beams can have three axes-parallel orientations. The primitives can have any color that differ from the color of the table. The location of pillars on the table is pre-defined. There are 49, 16 and 4 instances for 5U, 4U and 3U arch shapes correspondingly. We use the Arch task to evaluate the generalization of our method to the shape category and to show advantages of unmaking procedure in Sections IV-C and IV-D.




Fig. 5: Assembling arches of 3-5 units in simulation (first 6 columns, top) and with a real robot (first 6 columns, bottom). Assembling arches with never-seen primitives resembling cubes and beams on the real robot (last 2 columns). The image pairs correspond to initial and final states.




Fig. 6: Our visual policies successfully re-build shapes using primitives added to the table.

C. Learning in state space

This section evaluates how efficient our approach in generalizing to 3D shapes in the state space. For evaluation, we use a simulated Arch task. Our approach receives a single shape instance as an input and learns a state-value function by unmaking this instance. Given the trained value function and the simulator, we obtain the state policy by iterating over all possible actions and taking the one that maximizes the value function prediction. This state policy is then used to discover new object instances. We iteratively repeat the state-value network training after unmaking the discovered instances during 20 iterations. We compare our approach to MCTS [7] and random exploration. Similarly to the state policy, both baselines choose an object for picking and its placing location on top of other objects. The baselines also choose the placing orientation among the three axis-parallel directions. For MCTS, we use a shape matching score which is defined by a percentage of how much the arch shape is completed with primitives.

We estimate an average amount of steps required by our method and the baselines to build an arch, and report results in Tab. I. We also report the minimum number of steps required to build an arch (Oracle). Note that the Oracle has a non-zero variance since arches can be composed from different numbers of primitives. While the efficiency of our approach is comparable to Oracle, baselines require orders of magnitude more steps to find a correct solution. For example, to build a 5U arch, MCTS and the random exploration require 1.6e4 and 6.1e4 steps respectively while our method solves the task in 6.9 steps on average. We expect our method to scale well to more complex tasks where the complexity of

baselines will prohibit their use.

D. Evaluation of visual policies

This section evaluates visual policies trained to solve tasks given images as input. We validate the HourGlass architecture in Section IV-D.1, show benefits of the proposed unmaking approach over prior work in Section IV-D.2, evaluate generalization of our method to the shape category in Section IV-D.3 and present an evaluation on the real robot in Section IV-D.4.

1) *Visual policies architecture*: We compare our heatmap-based architecture to a ResNet-18 network that directly outputs source and target parameters of manipulated primitives. Our approach uses HourGlass [39] to predict a multi-modal distribution of picking and placing actions. The first heatmap corresponds to 2D locations on the robot workspace, the three additional heatmaps encode orientations of a primitive. ResNet-18 [23] predicts five values corresponding to 2D locations and three orientations of objects.

We train HourGlass and ResNet-18 networks to build towers of 3, 5 and 7 cubes of green, yellow and red colors. For ResNet-18, we adapt the learning rate to 1e-3 and the input image size to 224. Tab. II indicates that both HourGlass and ResNet policies achieve almost perfect accuracy for towers of 3 cubes (Fig. 4, left). Given more complex scenes with multiple primitives of the same color, ResNet fails due to its unimodal prediction. As illustrated in Fig. 4, it outputs a mean location of relevant primitives. The HourGlass-based policy builds towers of 7 cubes with a failure rate of 6% where the errors are mainly caused by occlusions. On the real world Tower task with 7 cubes, the HourGlass-based policy succeeds in 18 trials out of 20. The two failure cases were caused by an occlusion and misidentifying a cube due to the sim2real gap. Empirically, we did not find any performance improvement when using multiple HourGlass modules.

2) *Learning by unmaking*: This section compares the proposed approach of unmaking assembled objects with prior work [59]. While our method disassembles objects in multiple ways, [59] proposes to use a single disassembly trajectory. Such disassembly trajectories consist of pairs of state and corresponding action. However, there could be multiple correct actions possible in a state as shown in Fig. 3.

We address this by computing several disassembly paths and merging actions that correspond to states, where the only difference comes from positions of primitives located on the table surface.

We train visual policies on observation-heatmap pairs obtained with and without multiple unmake trajectories. Tab. III shows that using multiple unmake paths significantly improves the performance. The performance difference is 26% on the hardest task of building 5U arches. Using multiple unmake paths make the policy learn multiple hypothesis of picking and placing locations. This property becomes critical when multiple identical objects are used or if there exist several shape instances where identical primitives are assembled in different configurations.

3) *Learning to build a 3D category shapes*: This section evaluates the generalization of our approach to a shape category. We train visual policies on trajectories of unmaking 3 sets of arch instances: (i) only the input instance, (ii) instances obtained by a state policy trained on arches of the same height, (iii) instances obtained by a state policy trained on arches of heights 3-5U. We refer to these policies as (i) instance, (ii) uni-height and (iii) multi-height. For instance and uni-height policies, we train separate networks to build arches of each height. For multi-height policy, the same network is used to build arches of varying heights.

We evaluate all the policies separately on the set of primitives corresponding to the input instance (Tab. IV, left) and on different sets corresponding to the entire category (Tab. IV, right). In the first case, all the policies assemble arches with less than 5% of failures. The results of uni-height policies are higher compared to the instance policies. We believe that this improvement is due to a higher variation of the category instances that can be seen as a form of data augmentation. However, the performance of the instance policies rapidly drops when they are exposed to unseen sets of primitives. The uni-height policies have 2-3% higher success rates than the multi-height policy. Their success rates on 5U arches are 95.6% and 94.0% correspondingly. However, we need to train only a single network for the multi-height policy which is then able to reason about available primitives on the table and decide which arch height to build.

4) *Real robot evaluation*: This section evaluates the performance of our method on the real world Arch task using two scenarios: (i) "normal" scenario that matches the simulation, (ii) "re-assembling" scenario. In the "normal" scenario the robot has to assemble arches from varying sets of primitives. Tab. V (left) shows that uni-height policies have similar performance to the multi-height policy in the standard scenario except for 3U arches. In the "re-assembling" scenario the task starts with an assembled arch and additional primitives on the table (see Fig. 6). The agent is expected to re-assemble the arch by using all available primitives on the table. Our method is able to automatically re-assemble shapes without additional training due to the presence of random pick-place actions in the train set. Similarly to the value function dataset described in Section III-B, we record observation-actions pairs corresponding to the inverse of the one-step random actions that include examples of re-

assembling. In the "re-assembling" scenario (Tab. V (right)) the multi-height policy has higher success rates for 4U and 5U arches. We illustrate the assembly of arches with the real UR5 robot arm in Figures 5, 6 and on the project webpage [1]. Failure cases with incorrect assembly are typically caused by occlusions and the gap between simulated and real environments.

Finally, we test the generalization of our approach to new building primitives and compare it to the MCTS baseline. We use three different sets of primitives that resemble cuboids used during training. The sets contain (i) 2 jars and a pencil case (Fig. 5 last column, top), (ii) 4 cans and a juice box (Fig. 5 last column, bottom), (iii) 5 stones (Fig. 1 bottom). The input we provide to the MCTS baseline is the state of primitives in terms of their sizes and locations. Size and location are determined by clustering 3D points, which are obtained based on depth image coordinates above the table similarly to the prediction correction procedure used for our method (see Appendix of [41]). We estimate the spatial dimensions of primitives by fitting bounding boxes to depth points associated to each 3D cluster. Tab. VI shows that our method significantly outperforms the MCTS baseline with the performance gap of up to 45%. In all cases the failures of MCTS are caused by errors in the estimation of size and location of the primitives. Given the incorrect estimates, MCTS cannot find an assembly path to build a correct arch. In contrast, our method relies on position correction based on direct image input and does not require spatial estimation of the primitives. Both methods often fail on the third set of primitives with stones given the substantial difference of stones to cuboid primitives used during training. Qualitative results for additional sets of primitives are presented in Appendix of [41].

V. CONCLUSION

We proposed an approach to build 3D object shapes using a robotic arm and varying sets of primitives. Our method efficiently learns to solve the task in the state space and then uses the obtained solutions as supervision to train visual policies in the observation space. We demonstrate successful application of our method to the new task of assembling shape categories and show promising results on a real robot.

While the disassembling procedure explored in this work may not be directly applicable to all tasks, we note that it could generalize even to physically irreversible actions by learning appropriate backward models in state space. Future work will explore this direction for a wide range of tasks including cooking and other more complex assembling tasks. We also note that the complexity of the proposed unmake procedure grows exponentially with the number of primitives. To address more complex tasks, it will be interesting to extend our method to a fixed number of sampled disassembly trajectories.

Acknowledgements. This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR19-P3IA-0001 (PRAIRIE 3IA Institute).

REFERENCES

- [1] Project webpage. <http://pascal.inrialpes.fr/data2/3D-shapes/>. 4, 7
- [2] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine. Learning to poke by poking: Experiential learning of intuitive physics. In *NIPS*, 2016. 1
- [3] M. Andrychowicz, D. Crow, A. Ray, J. Schneider, R. H. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. *NIPS*, 2017. 2
- [4] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *RAS*, 57(5), May 2009. 2
- [5] Z. Chen, D. Guo, T. Xiao, S. Xie, X. Chen, H. Yu, J. Gray, K. Srinet, H. Fan, J. Ma, C. Qi, S. Tulsiani, A. Szlam, and L. Zitnick. Order-Aware Generative Modeling Using the 3D-Craft Dataset. In *ICCV*, 2019. 1
- [6] F. Codevilla, M. Miiller, A. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. *ICRA*, 2018. 1
- [7] R. Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *Computers and Games*, 2006. 5, 6
- [8] E. Courmans and Y. Bai. PyBullet, Python module for physics simulation, robotics and machine learning, 2016. 4
- [9] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki. An incremental constraint-based framework for task and motion planning. *IJRR*, 37, 2018. 2
- [10] Y. Duan, M. Andrychowicz, B. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. *NIPS*, 2017. 1, 2
- [11] F. Ebert, C. Finn, S. Dasari, A. Xie, A. X. Lee, and S. Levine. Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control. *arXiv*, 2018. 1
- [12] S. Edelkamp and J. Hoffmann. PDDL2.2: The Language for the Classical Part of IPC-4. 2004. 2
- [13] B. Espiau, F. Chaumette, and P. Rives. A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 1992. 2
- [14] R. Fikes and N. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 1971. 2
- [15] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel. Reverse curriculum generation for reinforcement learning. In *CoRL*, 2017. 2
- [16] D. Gandhi, L. Pinto, and A. Gupta. Learning to fly by crashing. In *IROS*, 2017. 1
- [17] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling. Ffrob: An efficient heuristic for task and motion planning. In *Algorithmic Foundations of Robotics XI*, pages 179–195. Springer, 2015. 2
- [18] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling. FFRob: Leveraging symbolic planning for efficient task and motion planning. *IJRR*, 37:104–136, 2018. 2
- [19] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling. Sampling-based methods for factored task and motion planning. *IJRR*, 37, 2018. 1
- [20] A. Grabner, P. M. Roth, and V. Lepetit. 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild. In *CVPR*, 2018. 1, 2
- [21] S. Gu, E. Holly, T. Lillicrap, and S. Levine. Deep reinforcement learning for robotic manipulation. In *ICML*, 2016. 2
- [22] K. He, M. Lahijanian, L. Kavraki, and M. Vardi. Towards manipulation planning with temporal logic specifications. In *ICRA*, 2015. 2
- [23] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 5, 6
- [24] I.-A. Hosu and T. Rebedea. Playing atari games with deep reinforcement learning and human checkpoint replay. In *ECAI Workshop on Evaluating General Purpose AI*, 2016. 2
- [25] D.-A. Huang, S. Nair, D. Xu, Y. Zhu, A. Garg, L. Fei-Fei, S. Savarese, and J. C. Niebles. Neural task graphs: Generalizing to unseen tasks from a single video demonstration. In *CVPR*, 2018. 1, 2
- [26] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML*, 2015. 5
- [27] A. Irpan. Deep reinforcement learning doesn't work yet, 2018. 1
- [28] M. Jaderberg, V. Mnih, W. Czarniecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv*, 2016. 1
- [29] M. Janner, S. Levine, W. T. Freeman, J. B. Tenenbaum, C. Finn, and J. Wu. Reasoning about physical interactions with object-oriented prediction and planning. In *ICLR*, 2019. 1, 2
- [30] J. Kober, J. Bagnell, and J. Peters. Reinforcement learning in robotics: A survey. *IJRR*, 2013. 2
- [31] T. Lampe and M. Riedmiller. Acquiring visual servoing reaching and grasping skills using neural reinforcement learning. *IJCNN*, 2013. 2
- [32] S. LaValle. *Planning Algorithms*. Cambridge University, 2006. 4
- [33] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *JMLR*, 2015. 1, 2
- [34] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 4
- [35] T. Lozano-Pérez and L. P. Kaelbling. A constraint-based method for solving sequential manipulation planning problems. In *IROS*, 2014. 2
- [36] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Y. Goldberg. Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics. *RSS*, 2017. 2
- [37] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel. Overcoming exploration in reinforcement learning with demonstrations. In *ICRA*, 2018. 1
- [38] S. Nair, M. Babaeizadeh, C. Finn, S. Levine, and V. Kumar. Time reversal as self-supervision. *ICRA*, 2020. 2
- [39] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016. 2, 4, 5, 6
- [40] A. Pashevich, D. Hafner, J. Davidson, R. Sukthankar, and C. Schmid. Modulated policy hierarchies. *NIPS Deep RL workshop*, 2018. 2
- [41] A. Pashevich, I. Kalevtykh, I. Laptev, and C. Schmid. Learning visual policies for building 3d shape categories. *arXiv*, 2020. 4, 5, 7
- [42] A. Pashevich, R. Strudel, I. Kalevtykh, I. Laptev, and C. Schmid. Learning to augment synthetic images for sim2real policy transfer. In *IROS*, 2019. 2, 3, 5
- [43] C. Paxton, Y. Barnoy, K. Katyal, R. Arora, and G. D. Hager. Visual robot task planning. In *ICRA*, 2019. 2
- [44] L. Pinto and A. Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. *ICRA* 2016. 2
- [45] D. A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *NIPS*, 1989. 2, 4
- [46] I. Popov, N. Heess, T. Lillicrap, R. Hafner, G. Barth-Maron, M. Vecerik, T. Lampe, Y. Tassa, T. Erez, and M. Riedmiller. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv*, 2017. 1, 2
- [47] R. Rahmatizadeh, P. Abolghasemi, L. Bölöni, and S. Levine. Vision-based multi-task manipulation for inexpensive robots using end-to-end learning from demonstration. *ICRA*, 2018. 2
- [48] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. In *NIPS*, 2015. 2
- [49] M. A. Riedmiller, R. Hafner, T. Lampe, M. Neunert, J. Degraeve, T. V. de Wiele, V. Mnih, N. Heess, and J. T. Springenberg. Learning by playing - Solving sparse reward tasks from scratch. *MLR*, 2018. 2
- [50] F. Sadeghi and S. Levine. CAD2RL: Real Single-Image Flight Without a Single Real Image. 2017. 2
- [51] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *ICRA*, 2014. 1, 2
- [52] F. Suárez-Ruiz, X. Zhou, and Q.-C. Pham. Can robots assemble an ikea chair? *Science Robotics*, 3(17), 2018. 1, 2
- [53] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *Robotics & Automation Magazine*, 19(4):72–82, 2012. 4
- [54] S. Sukhbaatar, I. Kostrikov, A. Szlam, and R. Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. *ICLR*, 2017. 2
- [55] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *IROS*, 2017. 2
- [56] M. Toussaint. Logic-geometric programming: An optimization-based approach to combined task and motion planning. *IJACI*, 2015. 1, 2
- [57] A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar. Learning robotic manipulation through visual planning and acting. *arXiv*, 2019. 1
- [58] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016. 4
- [59] K. Zakka, A. Zeng, J. Lee, and S. Song. Form2Fit: Learning Shape Priors for Generalizable Assembly from Disassembly. *arXiv*, 2019. 2, 6
- [60] A. Zeng, S. Song, J. Lee, A. Rodríguez, and T. A. Funkhouser. Tossingbot: Learning to throw arbitrary objects with residual physics. *RSS*, 2019. 1, 2
- [61] T. Zhang, Z. McCarthy, O. Jow, D. Lee, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. *ICRA*, 2017. 2