

Visual Affordance Prediction for Guiding Robot Exploration

Homanga Bharadhwaj, Abhinav Gupta, Shubham Tulsiani

Abstract—Motivated by the intuitive understanding humans have about the space of possible interactions, and the ease with which they can generalize this understanding to previously unseen scenes, we develop an approach for learning ‘visual affordances’. Given an input image of a scene, we infer a distribution over plausible future states that can be achieved via interactions with it. To allow predicting diverse plausible futures, we discretize the space of continuous images with a VQ-VAE and use a Transformer-based model to learn a conditional distribution in the latent embedding space. We show that these models can be trained using large-scale and diverse passive data, and that the learned models exhibit compositional generalization to diverse objects beyond the training distribution. We evaluate the quality and diversity of the generations, and demonstrate how the trained affordance model can be used for guiding exploration during visual goal-conditioned policy learning in robotic manipulation.

I. INTRODUCTION

Consider the images o_c in Figure 1. We humans can effortlessly understand the depicted scenes e.g. a bottle lying on the table in the top image, or a teddy lying next to a pot. More importantly, we can infer not only what *is*, but also what *can be* e.g. one can imagine the bottle being placed at a different location on the table, or turned to lie horizontally, or perhaps taking the teddy to place it inside the pan. We instinctively make such judgements across a myriad of everyday scenarios, understanding that an open cabinet can be closed, or that an egg can be broken, or spilled liquid wiped away. In this work, our goal is to build a computational system with similar capabilities, which can be used directly as a goal-sampling distribution for robot exploration. Given just a single image depicting generic (possibly unseen) objects, we wish to predict possible configurations that may occur as a result of a human (or a robot) interacting in the scene.

We are inspired by Gibson’s seminal work on *affordances* which argues for developing intelligent agents with an intuitive understanding of possible interactions they can have with their environment [14]. Initial attempts in the vision community formalized this as a pixelwise labeling task (e.g. ‘sittable’ surface), but these do not explicitly model the actions or their effects [6], [42]. An alternative approach has been to use geometric models [16], [12] to predict human-centric affordances. Towards a richer parametrization, recent approaches have pursued ‘visuo-motor’ affordances where the space of possible interactions is modeled via predicting possible low-level actions that an agent can execute to affect its environment [33], [37], [29], [43]. While such visuo-motor affordances can be directly translated to agent behavior, we argue that the requirement of inferring precise actions

The authors are with the Robotics Institute, Carnegie Mellon University

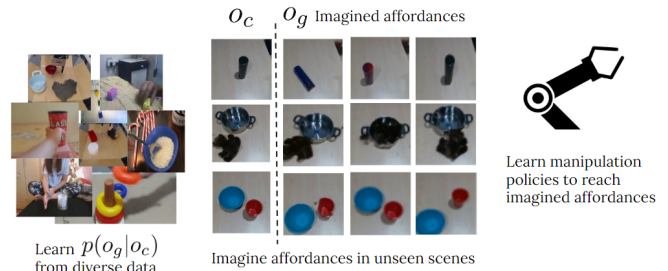


Fig. 1: Given a previously unseen scene with objects o_c , humans are able to effortlessly imagine the result of potential interactions in the scene. In this paper, we target the problem of learning such visual affordances o_g in new scenes, by training a probabilistic model $p(o_g|o_c)$ on large diverse datasets, and using the learned model for exploration in learning manipulation policies.

associated can be restrictive e.g. babies may understand that a fruit can be cut even if they can’t do it themselves, as perhaps an adult may know that a bike’s tire can be removed even if not sure how to precisely do so. In our work, we therefore pursue ‘visual affordances’, where instead of modeling the action trajectories, we aim to model what their *effects* can be. More concretely, given a single input image, we pursue the task of conditionally generating diverse and plausible images of the same scene that can be achieved as a result of an agent’s interaction.

This task of inferring visual affordances is an increasingly common one in the robot learning community, where an understanding of interesting plausible actions can help guide exploration. However, the typical methods only learn domain-specific affordances — relying on active interaction in a specific (lab) setting, they learn models for interactions with the particular objects in the particular environment [33], [37], [24], [32]. In addition to their limited generalizability, these active interaction-based methods also do not learn more complex behaviors (e.g. stacking blocks) as random interactions from an agent are unlikely to lead to such interesting transitions. Our key insight in this work is that instead of only relying on active interactions with limited variation, such visual affordances can be learned from large-scale diverse passive data. Just as we humans can learn from watching others (e.g. instructional videos), our system learns visual affordances using interaction videos depicting generic human and robot interactions across varied settings. Using this trained affordance model, when a robot is placed in a certain scene, we can generate goals corresponding to plausible manipulations of the objects in the scene, and learn a goal-conditioned policy through self-supervised exploration.

Concretely, the affordance model consists of two stages: 1)

we learn discrete latent embeddings from images by training a VQ-VAE [41] on the entire dataset 2) we learn conditional prediction of latent codes through an auto-regressive generation procedure using a Transformer architecture [5]. Because of the diversity seen in training, our system is able to generalize to new environments and unseen objects, and predict interesting plausible configurations in settings never seen in training. This makes the trained affordance model applicable for exploration in robot learning. We consider the setting of goal-conditioned policy learning and use the affordance model for sampling goals that the policy tries to reach during training. Due to the diversity of sampled goals, the policy is incentivized to explore broadly and learn interesting behaviors like stacking, compared to curiosity-based exploration strategies.

II. RELATED WORKS

Actionable Information with Visual Observations. Prior work has explored the problem of learning how to interact with objects in the scene from visual observations, in different settings like dextrous manipulation [28] and mobile navigation [31]. These approaches typically use either passive observations (for example human videos) [30], [2], [12] or active interactions with a robotic agent [36]. Some prior works also leverage robot simulators to randomize the generation of different types of objects in order to learn affordances for tasks like grasping [23] and pushing [45] objects. The unifying idea in all these approaches is learning *how* to interact with objects in the scene. They tackle a problem slightly orthogonal to ours, where by learning affordances, we refer to visualizing the *result* of interactions, and not specifically how to obtain those interactions.

Generative Modeling. There have been significant recent advances in deep generative models, with high-quality realistic images being produced by StyleGAN [21], [22] and BigGAN [3], [7] among others that use transformers [41], [40], [34], [5], [9]. While most of these models tackle the problem of unconditional generation, a closely related problem to our setting is that of conditional image generation. This is typically defined as the task of generating new images from a dataset conditional on certain class attributes [39], [10]. Video generation methods [25] using transformers perform per-timestep prediction, whereas our affordance model samples (long-horizon) goal images that we show are useful for robot exploration. Some style transfer approaches [20], [47] also have a similar conditional generation formulation, where the conditioning is done on a source image as opposed to a class attribute. Our affordance model is an approach for better conditioning on images, where we want the generations to capture configurational changes in the attributes of objects in the source image.

Robot Exploration. Generative models have been used for goal-sampling in visual robot learning [33], [32], [37], [27], [11]. Learning a good goal-sampling strategy for diverse goals and training goal-conditioned policies to try and reach

sampled goals is a paradigm for exploration without the use of explicit heuristics like curiosity [35], [38] or reduction of uncertainty [4], [8]. Most prior work in this vein has learned goal-sampling models only in the experiment domain either by collecting large robot interaction datasets, or training the goal-sampler online jointly with policy learning [33], [27], [37], [43]. Collecting large datasets in the lab is expensive, and the goal-sampling models trained solely on a specific-lab setting is unlikely to generalize to other settings without requiring re-training. In contrast to these approaches we learn a model for generating visual affordances from various internet data - *not* data collected in the experimenter’s lab. We show how the learned affordance model can be used to guide exploration for goal-conditioned policy learning in environments with unseen objects.

III. APPROACH

Motivated by the intuitive understanding humans have about the space of possible interactions, and the ease with which they can generalize this understanding to previously unseen scenes, we develop an approach for learning visual affordances from passive data that can be used as goals for guiding robot exploration. We use the term affordance to mean the set of possible interactions in a particular scene. Given an image of a scene, we want the affordance model to generate a new image with different configurations of the same objects in the scene. This model can be used for goal-directed exploration, such that in a new scene, the agent can sample a goal conditioned on the current image, and execute actions to try and reach the sampled goal.

Since collecting data in-lab is expensive, we make use of the rich diversity of passive data existing in the internet, to learn the affordance model, and use it directly for goal-sampling. The affordance model is two-staged: we first learn a VQ-VAE [41] to discretize the space of continuous images, and then learn a transformer based auto-regressive model that (in latent space) predicts a possible goal given the input. We then project the generated latent goal, back to the image-space with the decoder of the VQ-VAE. This framework lets us both: a) generate diverse goals, and b) produce realistic and high-resolution images.

For goal-conditioned policy learning, self-supervised approaches typically adopt a learning pipeline consisting of two (often inter-leaved) components: 1) exploration to collect data of interactions with the environment, and 2) policy learning using the interaction trajectories. We show that it is possible to drive better exploration by using the trained affordance model as a conditional goal-sampling distribution that samples diverse realistic goals in an unseen scene.

Concretely, we learn a goal-conditioned policy through exploration, where the goals are sampled from a learned affordance model in the beginning of the trajectory. In order to learn such a policy through self-supervised exploration, the sampled goals must be *plausible* and *diverse*, corresponding to different arrangement of the objects in the scene. Depending on the objects in the scene, the different arrangements could be putting a lid on a pot, pushing a cup across the table,

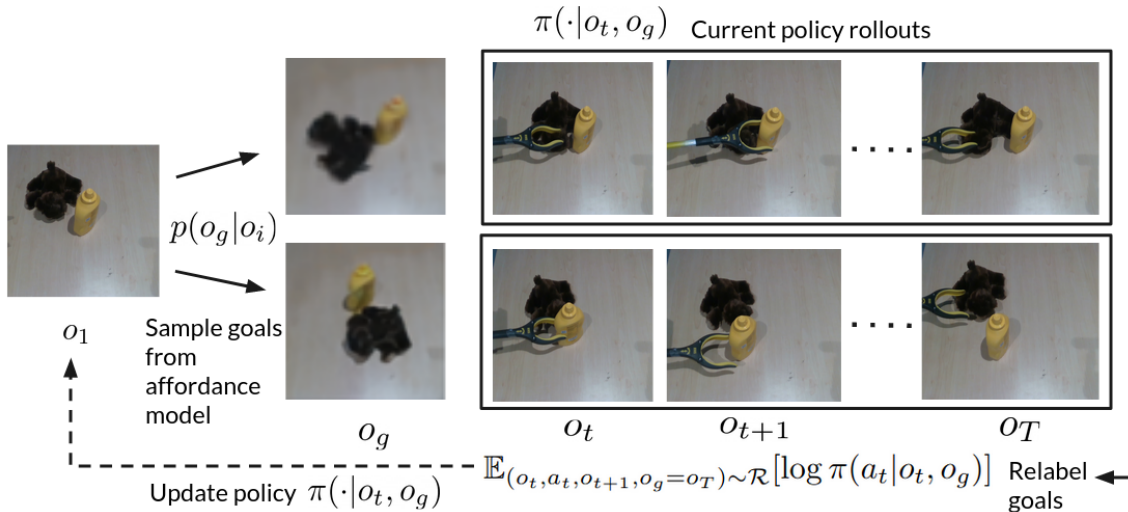


Fig. 2: Affordance-driven exploration and policy learning through hindsight goal re-labeling. Given an initial configuration, we sample a goal with the affordance model, execute rollouts with the current policy, and store the transitions in the replay buffer. For updating the policy, we sample transitions from the replay buffer and re-label goals to be the final frames in the corresponding trajectory. The process is described in detail in section III-B

grasping an object and rotating it etc. We show that such behaviors emerge from the trained affordance model, and describe how we can learn a goal-conditioned policy through hindsight re-labeling of exploration trajectories.

In the next sub-sections, we first describe the architecture and training details of the affordance model, and then discuss how we use the model as a goal-sampler for robot exploration.

A. Learning to Imagine Goals

Instead of modeling in the image-space directly, in order to reduce spatial redundancies and implicitly abstract out interesting objects in the scene, we first learn a down-sampled encoding of high resolution images. To learn these discrete latent embeddings, we employ a VQ-VAE to learn lossy encoder and decoder models that can transform latents to generate realistic image samples. Having a discrete latent space allows flexibility in modeling features (such as objects in the scene) that span multiple pixels in the image, and do not encode local imperceptible differences that are not significant.

Training details of the VQ-VAE. Let the images be $o \in \mathbb{R}^{H \times W \times 3}$. The VQ-VAE model consists of an encoder $E(o) \in \mathbb{R}^{h \times w \times L}$, a codebook $C = e_k, k=1 \dots K$ with elements e_i of size L and a decoder $D(e) \in \mathbb{R}^{H \times W \times 3}$. The quantization happens in the channel space, where all the $h \times w$ vectors of size L are replaced by their nearest codebook vectors e_i , and the resulting quantized latent code of dimension $h \times w \times L$ is fed to the decoder. The training objective for this model, following prior work [41] is as follows:

$$\mathcal{L}_{\text{vqvae}} = \mathbb{E}_{o \sim \mathcal{D}} [\|o - D(e)\|_2 + \|E(o) - \text{sg}[e]\|_2]$$

Here, sg refers to the stopgradient operator, that is defined as identity during forward-computation.

After training the VQ-VAE, corresponding to each image o_c , we obtain a downsampled encoding z_c . We then learn an auto-regressive prior using a Transformer

architecture to model $p(z_g|z_c)$. Finally, with the decoder of the VQ-VAE we obtain a corresponding image o_g from the latent z_g . Learning this conditional generation model in the latent space with a powerful Transformer architecture enables us to achieve compositional generalization, by generation of plausible goals in scenes where there are multiple objects. The stochasticity of the model is helpful in ensuring that the generated samples are allowed to be diverse.

Auto-regressive goal sampling. In Fig. 3, we show a forward pass through our affordance model that generates a goal image o_g , given a conditioning image o_c . The initial image o_c is first encoded by the VQ-VAE to a discrete latent code z_c . The Transformer then generates the latent code corresponding to the goal image, $z_g \sim p(z_g|z_c)$. The generation process happens pixel by pixel in raster-scan order and channel by channel for one pixel. Hence we can denote the auto-regressive generation process as:

$$p(z_g) = \prod_i p(z_{(g,i)} | z_{(g,<i)}, z_c)$$

This equation shows that the prediction of the i^{th} pixel in the latent code $z_{(g,i)}$ is conditioned on the already predicted pixels $z_{(g,<i)}$ and the initial latent code z_c .

Learning objective. We instantiate the affordance model as $p_\psi(o_g|o_c)$ and aim to maximize $\mathbb{E}_{(o_g, o_c) \sim \mathcal{D}} [\log p_\psi(o_g|o_c)]$. Given image pairs $(o_c, o_g) \sim \mathcal{D}$, we transform each image to their respective quantized latent code z_c and z_g . The autoregressive transformer model is used to generate \hat{z}_g conditioned on z_c , i.e. $\hat{z}_g \sim p(z_g|z_c)$. We train the model to maximize $\mathbb{E}_{(z_g, z_c) \sim \mathcal{D}} [\log p(z_g|z_c)]$.

Training Details of the Transformer. For the above generation, we use an architecture similar to Image-GPT [5]. It consists of an encoder-decoder model with masked convolutions

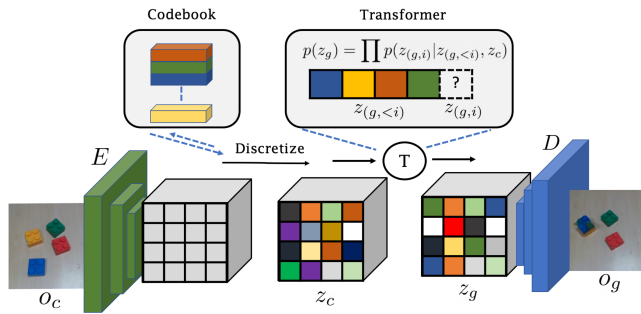


Fig. 3: Illustration of a forward pass through the affordance model. The initial scene o_c is first encoded by a VQ-VAE encoder, and converted to a discretized latent code by swapping nearest neighbor encodings from a codebook C . The resulting discrete latent code z_c is passed through the Transformer. The autoregressive generation yields a latent code z_g which is decoded by the VQ-VAE decoder to a plausible goal image o_g .

and self-attention layers in the decoder. The auto-regressive model is preserved by padding the input, and replacing the padded pixel values with the generated values and repeating the process recursively. Finally, all the values in z_g correspond to valid generations that become more and more accurate as training progresses, and so we eventually have a model of the form $p(z_g|z_c)$. We then feed in z_g to the VQ-VAE decoder, obtaining the decoded image frame o_g . This completes the structure of the affordance model $p_\psi(o_g|o_c)$. After training the model, given a new initial image \hat{o}_c , we can sample a plausible goal $o_g \sim p_\psi(o_g|\hat{o}_c)$. We empirically analyze this model in the next section, and show results of goal generations in unseen scenes.

B. Affordance-driven Robot Exploration

We use the affordance model trained solely on passive data for goal-conditioning and do not fine-tune on any lab-specific data. Because we trained the model on diverse passive data, we observe good generalization for affordance prediction in the robot learning setting. We obtain exploration trajectories with interesting behaviors, and are eventually able to learn better goal-conditioned policies compared to alternate exploration strategies like curiosity [35] that learn from scratch in the environment, or those that sample goals from other sampling distributions like a VAE-prior [33].

Given the affordance model, we train a policy that chooses actions for executing a particular task specified by a goal image. We refer to the observation the robot sees at time-step t of an episode as o_t , and denote the goal image as o_g . We denote the goal-conditioned policy as $\pi(a_t|o_t, o_g)$. We learn this policy by simple behavior cloning with goal-relabeling. After sampling a goal $o_g \sim p_\psi(o_g|o_1)$ and executing a trajectory $(o_1, a_1, o_2, a_2, \dots, o_T; o_g)$ to reach the goal, even if the final configuration o_T is not similar to the goal configuration, the executed trajectory provides useful information for reaching the configuration the system ended up in. We describe the overall framework in Fig. 2. During deployment, test goal images are sampled from some distribution $o_g \sim p(\mathcal{G})$, and the robot must interact with the objects in the scene to reach the goal configuration.

Training the goal-conditioned policy. Given observation o_1 corresponding to the initial scene, we sample a goal from the affordance model $o_g \sim p_\psi(o_g|o_1)$ and execute a trajectory $(o_1, a_1, o_2, a_2, \dots, o_T; o_g)$, which we store in the replay buffer \mathcal{R} . In Fig. 2 we see examples of goals sampled in a scene. The affordance model enables sampling diverse goals and encourages the policy to try and reach them, thereby ensuring interesting interactions. Given the data of interactions in the replay buffer \mathcal{R} , we perform goal relabeling and use tuples $(o_t, a_t, o_{t+1}, o_g = o_T)$ to update the policy $\pi(a_t|o_t, o_g)$. This corresponds to a variation of hindsight-experience replay, that has been shown to be useful in prior works [1], [13], [26]. The rationale for this is that during exploration, although the incompletely trained policy might not reach the sampled goal, the final state it reaches is still a *potential* goal, and the executed sequence of actions provides guidance on reaching this goal. The policy update method is to do simple behavior cloning that maximizes the likelihood $\mathbb{E}_{(o_t, a_t, o_{t+1}, o_g = o_T) \sim \mathcal{R}} [\log \pi(a_t|o_t, o_g)]$. Note that we do not need samples to be on-policy for this update and so we can interleave exploration of a few trajectories with a policy update phase where we randomly sample tuples from different trajectories in the replay buffer.

IV. EXPERIMENTS

Through experiments, we aim to understand the following research questions:

- How effective is the affordance model in generating diverse scenes with plausible object manipulations?
- How effective are the generated plausible and diverse affordances in guiding robot exploration in unseen scenes?

A. Setting

Data. For training the affordance model, we use data from three different internet sources: somethingsomethingv2 [15], Berkeley [24], and JHU CoStar [19]. We did not collect any additional data ourselves, and use only pairwise frames extracted from these datasets for training. We use the trained visual affordance model as-is for robot experiments and do not do any fine-tuning in the lab.

Baselines. We compare our model against two relevant baselines, a Conditional VAE (CVAE) [39], and a Conditional GAN (Pix2Pix) [20] for conditional generation and robot exploration, and in addition with a curiosity baseline [35] for robot exploration. We train these models on the same data as our model, and perform qualitative evaluations as well as quantitative comparisons through several metrics described in the next sub-section. After analyzing the visual generations of the baselines in section IV-B, we evaluate policy learning behaviors in section IV-C

B. Evaluating Predicted Affordances

As highlighted by prior work, evaluating the quality of synthesized images is challenging [20], [47]. Typical metrics used in assessing image reconstruction (like the pixel mean-squared error) do not translate well to assessing the quality of novel image generations. Further, these metrics will not be




Fig. 4: Qualitative analysis of the proposed affordance model with the baselines CVAE and Pix2Pix. Given initial conditioning frame o_c that corresponds to an unseen configuration during training, we see that the sampled affordances o_g from our model are diverse and correspond to plausible interactions in the scene. For the baselines, we see that the generated frames are not diverse and sometimes omit certain objects from the scene or introduce different artifacts.

useful for understanding how diverse and plausible are the generated affordances for novel scenes. Hence we consider two metrics for analysis of the *realness* and *diversity* of generations, and a third metric based on *human perceptual evaluation*.

To analyze how well the distribution of generated affordances correspond to the set of ground-truth real affordances in the data, we compute FID scores [18] for our model and the baselines. This helps get some useful signal for quantitatively evaluating the “realness” of the generated frames. In addition to the FID score, we use an LPIPS [46] based distance metric to measure the “diversity” of the generated images. For this, we compute average pairwise distances between generated affordance frames, and compare this number across baselines. A higher value means that the generated affordances are more different from each other.

Since the above metrics are not sufficient, motivated by evaluation protocols in prior works [17], [20], [47], [44], we conduct an perceptual study with Amazon Mechanical Turk (MTurk). We run the MTurk perceptual study, following the protocol from prior work [20], [47] where images are shown on the screen for 3 seconds and workers are asked to guess certain proprieties of the images. In our setting, we show three images per screen and ask the workers to choose one among the two rightmost images. We mention in the instructions that the task is to guess which of the two images on the right correspond to plausible manipulations of objects on the left image. We randomize the ordering of images such that one image is from our model and the other is from a baseline and compare the average number of times workers choose ours compared to the baselines.

Qualitative results. In Fig. 4 we perform a qualitative analysis of the proposed affordance model with the baselines CVAE and Pix2Pix. For each of the three initial scenes o_c in the first column, we generate three affordances with our model, three with the CVAE model, and one with the Pix2Pix model (since the model is not stochastic). The initial scenes are unseen during training. We can see that the affordances generated by our model are diverse and involve plausible

TABLE I: Comparison of the realism and diversity of generated samples through the FID and diversity metrics respectively. Higher is better for diversity and lower is better for FID.

	Our	Pix2Pix	CVAE
FID(↓)	141.1	185.7	253.7
Diversity(↑)	0.40 ± 0.11	N/A	0.33 ± 0.18

manipulations of the different objects in the scene. We observe the emergence of interesting affordances like stacking of blocks and re-orientation of the bottle. Whereas for the baselines, sometimes certain objects are omitted from the scene, are non-realistically generated, and do not involve diverse manipulations of the objects. Given that the objects have never appeared in the training data before, this provides evidence that our method can generalize to new scenes and compose affordances for different objects.

Quantitative results. In Table I we show results for perceptual quality of generated samples, under the lens of FID score and the diversity metric, defined in section IV-B. We can see that our model has low FID scores compared to the baselines and high diversity, indicating realistic and diverse samples. In Table II, we report results from the MTurk perceptual study. 40 workers participated in our study. In each trial, workers saw an image on the left, and were asked to choose which of the two images on the right corresponded to plausible manipulations of objects on the left image. The numbers show the % of trials where workers chose a sample from our model as opposed to that from the baselines, averaged over all the workers. Since the numbers are $> 50\%$, we see that workers preferred samples from our model compared to the baselines. This confirms the plausibility of our generated affordances.

TABLE II: Pairwise comparison of results for the MTurk perceptual study. We tabulate the % of trials where workers chose a sample from our model as opposed to that from the baselines, averaged over 40 workers who participated in our study. Higher is better. A number $> 50\%$ shows that workers chose our model’s sample more number of times compared to the baselines.

	Pix2Pix	CVAE
Our	69.8 ± 11.9	75.5 ± 10.8

C. Benchmarking Affordance-Guided Exploration

In this section, we empirically analyze the framework in terms of generating diverse affordances that serve as goals for aiding exploration in robot policy learning. We consider a goal-directed policy learning setting, where the robot needs to set its own goals, and explore the environment to try and reach those goals. There is no notion of *tasks* during the exploration phase, but for evaluation, the experimenters provide goal images corresponding to tasks like pushing, and stacking that can be performed in the scene. The overall aim is to evaluate goal-reaching behaviors that can be learned by training a goal-conditioned policy described in section III-B.




Fig. 5: Robot experiment showing the Franka Panda arm and the overhead RealSense camera. We show two different views of the workspace.




Fig. 6: We show examples of two evaluation runs of our method and the curiosity baseline corresponding to two test goals in stacking and pick and place respectively. On the right are the goal images that the policy is conditioned on for evaluation, and the sequence of observations o_1, \dots, o_T corresponds to the executed trajectory.

Setup. Our setup is shown in Fig. 5, where we use a Franka Emika Panda robot arm, with an overhead Intel Realsense camera for observations. The robot is controlled through end-effector (EE) control, and the action-space is four dimensional - (x,y,z) position of the EE and opening/closing of the gripper. We place certain objects in the scene, and let the robot interact with them. We choose diverse everyday objects like teddy bears, cloth, ketchup bottle, blocks etc. such that different behaviors like pushing, pick and place, stacking are plausible to emerge through interaction. We reset the scene after T steps of interaction while introducing new objects and/or changing the position of existing objects. As is standard in evaluating goal-conditioned policies, after training, we measure % success in reaching a set of test goal-images. For comparison, in addition to goal-sampling

TABLE III: We show results of tests for various robot manipulation tasks. For each task, we have two test goal configurations, and report success rate over 10 trials each.

	Pushing	Pick and Place	Stacking
Curiosity	50%	40%	30%
CVAE	30%	20%	10%
Pix2Pix	30%	10%	10%
Our	70%	60%	60%

from CVAE and Pix2Pix baselines, we consider a curiosity-based exploration baseline [35] for data collection, and follow the same training protocol as our method to train a goal-conditioned policy with the exploration data.

Results. After training the policy through affordance-driven exploration, for evaluation we consider a set of goal images, and evaluate the fraction success of the policy in reaching goal configurations. We emphasize that only a *single* goal-conditioned policy is trained and evaluated, but for ease of analysis we show results in Table. III split across 3 different type of tasks, with the same trained policy. In Table. III, we show comparisons of the success rates of the two approaches, across three types of tasks - pushing, pick and place, and stacking. For each task, we consider two goal-images, and tabulate % success across 10 trials. In Fig. 6 we visualize some evaluation runs for our method and the curiosity baseline. Please refer to the supplementary video for additional visualizations. We can see an average of 25% higher success rates compared to the baselines, demonstrating the efficacy of the affordance model for goal-directed exploration. We observe significantly lower success rates for the CVAE and Pix2Pix baselines, which we believe is because the images generated either do not correspond to realistic reachable goals for the robot or do not change the configuration at all (see examples in Fig. 4), and thus are unable to help in guiding exploration for policy learning. This confirms the generalization of the proposed affordance model trained on diverse passive data, for aiding in goal-directed exploration.

V. DISCUSSION

In this paper we developed a framework for learning visual affordances from passive data such that the learned model can be directly used for goal-directed exploration in unseen scenes. Given an image of an initial scene, our affordance model could generate diverse images corresponding to plausible interactions. Further, its ability to generalize allowed us to directly leverage this affordance model to drive exploration in robot learning. We believe our approach is indicative of the broader potential of large-scale and diverse visual data in developing intelligent agents that can act in generic environments. However, ours is but a first step towards this and we believe several exciting questions remain to be addressed. First, our notion of ‘visual affordances’, while capturing high-level changes, did not model the actions required. In extending this approach to capture ‘visuomotor affordances’, one must overcome the challenges of varying morphology variation across humans and hands, while also developing action representations beyond low-level trajectories. It would also be interesting to generate intermediate ‘checkpoints’ to continuously depict the transition, but these may not be directly useful for robots due to the presence of human hands. Broadly, we are hopeful that extending on our work, there will be future approaches that can leverage increasingly diverse and multi-modal passive datasets for generalization in robot learning, alleviating the necessity of in-domain data collection that has greatly bottle-necked the field.

REFERENCES

- [1] Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., Zaremba, W.: Hindsight experience replay. *Advances in neural information processing systems* **30** (2017)
- [2] Brahmabhatt, S., Ham, C., Kemp, C.C., Hays, J.: Contactdb: Analyzing and predicting grasp contact via thermal imaging. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8709–8719 (2019)
- [3] Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096* (2018)
- [4] Burda, Y., Edwards, H., Storkey, A., Klimov, O.: Exploration by random network distillation. *arXiv preprint arXiv:1810.12894* (2018)
- [5] Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: *International Conference on Machine Learning*. pp. 1691–1703. PMLR (2020)
- [6] Delaitre, V., Fouhey, D.F., Laptev, I., Sivic, J., Gupta, A., Efros, A.A.: Scene semantics from long-term observation of people. In: *European Conference on Computer Vision*. pp. 732–745. Springer (2012)
- [7] Donahue, J., Simonyan, K.: Large scale adversarial representation learning. *Advances in Neural Information Processing Systems* **32** (2019)
- [8] Dwaracherla, V., Lu, X., Ibrahim, M., Osband, I., Wen, Z., Van Roy, B.: Hypermodels for exploration. *arXiv preprint arXiv:2006.07464* (2020)
- [9] Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12873–12883 (2021)
- [10] Fetaya, E., Jacobsen, J.H., Grathwohl, W., Zemel, R.: Understanding the limitations of conditional generative models. *arXiv preprint arXiv:1906.01171* (2019)
- [11] Florensa, C., Degraev, J., Heess, N., Springenberg, J.T., Riedmiller, M.: Self-supervised learning of image embedding for continuous control. *arXiv preprint arXiv:1901.00943* (2019)
- [12] Fouhey, D.F., Delaitre, V., Gupta, A., Efros, A.A., Laptev, I., Sivic, J.: People watching: Human actions as a cue for single view geometry. In: *European Conference on Computer Vision*. pp. 732–745. Springer (2012)
- [13] Ghosh, D., Gupta, A., Reddy, A., Fu, J., Devin, C., Eysenbach, B., Levine, S.: Learning to reach goals via iterated supervised learning. *arXiv preprint arXiv:1912.06088* (2019)
- [14] Gibson, J.J.: *The ecological approach to visual perception*. (1979)
- [15] Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Freund, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5842–5850 (2017)
- [16] Gupta, A., Satkin, S., Efros, A.A., Hebert, M.: From 3d scene geometry to human workspace. In: *CVPR* (2011)
- [17] Gupta, A., Johnson, J., Alahi, A., Fei-Fei, L.: Characterizing and improving stability in neural style transfer. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4067–4076 (2017)
- [18] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
- [19] Hundt, A., Jain, V., Lin, C.H., Paxton, C., Hager, G.D.: The costar block stacking dataset: Learning with workspace constraints. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 1797–1804. IEEE (2019)
- [20] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1125–1134 (2017)
- [21] Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4401–4410 (2019)
- [22] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8110–8119 (2020)
- [23] Khansari, M., Kappler, D., Luo, J., Bingham, J., Kalakrishnan, M.: Action image representation: Learning scalable deep grasping policies with zero real world data. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 3597–3603. IEEE (2020)
- [24] Khazatsky, A., Nair, A., Jing, D., Levine, S.: What can i do here? learning new skills by imagining visual affordances. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 14291–14297. IEEE (2021)
- [25] Le Moing, G., Ponce, J., Schmid, C.: Ccvs: Context-aware controllable video synthesis. *Advances in Neural Information Processing Systems* **34**, 14042–14055 (2021)
- [26] Li, A., Pinto, L., Abbeel, P.: Generalized hindsight for reinforcement learning. *Advances in neural information processing systems* **33**, 7754–7767 (2020)
- [27] Lin, X., Baweja, H.S., Held, D.: Reinforcement learning without ground-truth state. *arXiv preprint arXiv:1905.07866* (2019)
- [28] Mandikal, P., Grauman, K.: Learning dexterous grasping with object-centric visual affordances. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 6169–6176. IEEE (2021)
- [29] Mo, K., Guibas, L.J., Mukadam, M., Gupta, A., Tulsiani, S.: Where2act: From pixels to actions for articulated 3d objects. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6813–6823 (2021)
- [30] Nagarajan, T., Feichtenhofer, C., Grauman, K.: Grounded human-object interaction hotspots from video. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8688–8697 (2019)
- [31] Nagarajan, T., Grauman, K.: Learning affordance landscapes for interaction exploration in 3d environments. *Advances in Neural Information Processing Systems* **33**, 2005–2015 (2020)
- [32] Nair, A., Bahl, S., Khazatsky, A., Pong, V., Berseth, G., Levine, S.: Contextual imagined goals for self-supervised robotic learning. In: *Conference on Robot Learning*. pp. 530–539. PMLR (2020)
- [33] Nair, A.V., Pong, V., Dalal, M., Bahl, S., Lin, S., Levine, S.: Visual reinforcement learning with imagined goals. *Advances in neural information processing systems* **31** (2018)
- [34] Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., Tran, D.: Image transformer. In: *International Conference on Machine Learning*. pp. 4055–4064. PMLR (2018)
- [35] Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: *International conference on machine learning*. pp. 2778–2787. PMLR (2017)
- [36] Pinto, L., Gupta, A.: Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In: *2016 IEEE international conference on robotics and automation (ICRA)*. pp. 3406–3413. IEEE (2016)
- [37] Pong, V.H., Dalal, M., Lin, S., Nair, A., Bahl, S., Levine, S.: Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698* (2019)
- [38] Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicipap, T., Gelly, S.: Episodic curiosity through reachability. In: *International Conference on Learning Representations (ICLR)* (2019)
- [39] Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* **28** (2015)
- [40] Tulsiani, S., Gupta, A.: Pixeltransformer: Sample conditioned signal generation. *arXiv preprint arXiv:2103.15813* (2021)
- [41] Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
- [42] Wang, X., Girdhar, R., Gupta, A.: Binge watching: Scaling affordance learning from sitcoms. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2596–2605 (2017)
- [43] Wang, Y., Gautham, N., Lin, X., Okorn, B., Held, D.: Roll: Visual self-supervised reinforcement learning with object reasoning. In: *Conference on Robot Learning* (2020)
- [44] Woods, A.T., Velasco, C., Levitan, C.A., Wan, X., Spence, C.: Conducting perception research over the internet: a tutorial review. *PeerJ* **3**, e1058 (2015)
- [45] Zeng, A., Song, S., Welker, S., Lee, J., Rodriguez, A., Funkhouser, T.: Learning synergies between pushing and grasping with self-supervised deep reinforcement learning. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 4238–4245. IEEE (2018)
- [46] Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In:

Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)

- [47] Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)