



HAL
open science

VP-GO: A 'Light' Action-Conditioned Visual Prediction Model for Grasping Objects

Anji Ma, Yoann Fleytoux, Jean-Baptiste Mouret, Serena Ivaldi

► To cite this version:

Anji Ma, Yoann Fleytoux, Jean-Baptiste Mouret, Serena Ivaldi. VP-GO: A 'Light' Action-Conditioned Visual Prediction Model for Grasping Objects. ICARM 2022 - IEEE International Conference on Advanced Robotics and Mechatronics, Jul 2022, Guilin, China. hal-03727531

HAL Id: hal-03727531

<https://hal.science/hal-03727531v1>

Submitted on 19 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

VP-GO: A ‘Light’ Action-Conditioned Visual Prediction Model for Grasping Objects

Anji Ma^{1,2†}, Yoann Fleytoux², Jean-Baptiste Mouret², and Serena Ivaldi²

Abstract—Visual prediction models are promising solutions for visual-based robotic grasping of cluttered, unknown soft objects. Previous models from the literature are computationally greedy, which limits reproducibility; although some consider stochasticity in the prediction model, it is often too weak to catch the reality of robotics experiments involving grasping such objects. Furthermore, previous work focused on elementary movements that are not efficient to reason in terms of more complex semantic actions. To address these limitations, we propose VP-GO, a “light” stochastic action-conditioned visual prediction model. We propose a hierarchical decomposition of semantic grasping and manipulation actions into elementary end-effector movements, to ensure compatibility with existing models and datasets for visual prediction of robotic actions such as RoboNet. We also record and release a new open dataset for visual prediction of object grasping, called PandaGrasp. Our model can be pre-trained on RoboNet and fine-tuned on PandaGrasp, and performs similarly to more complex models in terms of signal prediction metrics. Qualitatively, it outperforms when predicting the outcome of complex grasps performed by our robot.

I. INTRODUCTION

There are several industrial scenarios where robots have to grasp or manipulate a variety of objects from uncluttered heaps, without relying on objects models (unknown or unavailable), nor tactile/force sensing [1]. A typical scenario is in waste sorting [2], where human experts plan appropriate sequences of actions to interact with complex objects. It would be desirable to automatize the process, i.e., to automatically find the sequence of actions that enable the robot to pick objects, even to grasp all the objects in the heap.

In the absence of robot sensing other than a camera and without prior knowledge on the objects, a promising method to address this problem is visual Model Predictive Control (MPC) or visual foresight [3]. Such a technique requires a visual prediction model, i.e., a model that predicts the visual outcome (i.e., the future camera images) of the robot’s action. Because of the high dimensionality of images, visual MPC was considered unfeasible until few years ago, when video prediction models based on deep neural networks started to show very promising results in computer vision [4] and also in robotics [5]. The first video prediction models were deterministic [6]; to deal with the uncertainty caused by the




Fig. 1: VP-GO contains a ‘light’ video prediction model and is able to integrated use open large-scale datasets and local datasets with semantic actions for the anticipating in grasping-based tasks.

interaction of the robot with the real world, several studies proposed stochastic models [7], where learned prior networks catch the “stochasticity”. Computationally, these methods are very greedy in terms of computation and training data.

For robotics applications, the seminal demonstration of visual MPC was done in [8], showing that visual predictions could be used to plan the robot’s end-effector towards desired goals. In that work and the follow-up [3] the robot’s actions were basically differential displacements of the end-effector, which was coherent with the babbling-like exploration of the robot’s workspace (similar to what has been done in the last decades in developmental robotics [9]). However, for industrial applications involving repetitive grasping this action representation is not appropriate: on the one side it increases the complexity in sample-based planning to execute a complex goal-driven sequence of actions, on the other side it does not carry the semantic description of high-level actions such as grasping or picking.

To address these limitations, we propose Visual Prediction Model for Grasping Objects (VP-GO), a “light” stochastic action-conditioned visual prediction model. Our model is based on SVG-LP [7], a groundbreaking work that proposed to use a learned prior model to deal with the time-varying stochasticity. In VP-GO we leverage the learned prior to catch the stochasticity of the real world grasping actions. SVG’ [10] and GHVAE [11], both follow-ups of SVG-LP, scaled up SVG-LP in complexity to deal with large datasets, but are very greedy in computational resources. Here, we revisit the principle of SVG-LP. Compared to SVG-LP, our model introduces action conditioning, uses convolutional Long-Short-Term Memory networks (LSTM) instead of a simple LSTM and has a deeper architecture. Compared to SVG’, which also uses convolutional LSTM, our networks have a “deeper” encoder, while significantly reducing at the same time the number of parameters (5 times less than

*This work was partially supported by the CHIST-ERA project HEAP and the Creativ’Lab in Loria.

¹School of Mechatronical Engineering, Beijing Institute of Technology, Beijing 100081, China {anjima}@bit.edu.cn

²Inria, University of Lorraine, CNRS, Loria, F-54000, France{firstname.lastname}@inria.fr

[†]Corresponding author.

the state-of-the-art models, hence the “light” adjective). Our model can compete with the existing models while being computationally efficient to train on a relatively small cluster.

Our target is grasping objects efficiently, ultimately by planning sequences of high-level grasps. To this end, we propose a hierarchical decomposition of semantic grasping and manipulation actions into elementary low-level actions (i.e., end-effector displacements). This is frequently done in developmental learning to scaffold complex actions into elementary actions, especially for learning manipulation [12]. This kind of representation has the key advantage of producing predictive models that can be used for downstream planning of sequences of high-level grasping and manipulation actions. Our model does not acquire semantic information explicitly (i.e., it is not an input of the network), but implicitly during the training process. In order to train the model with rich semantic information, we acquire a new dataset (PandaGrasp) with sequences of semantic actions defined per our hierarchy, executed by a Franka robot. This design enables us to compare with state-of-the-art models and reuse large-scale datasets for visual prediction of robotic actions such as RoboNet [5] to pre-train our model.

In addition to providing the source code of VP-GO, we release PandaGrasp, the dataset recorded with our Franka setup. Differently from RoboNet, it is focused on meaningful grasping actions, and therefore can greatly help the community to study visual prediction models specific for grasping.

II. RELATED WORK

Robot grasping and manipulation have been widely studied for many years. Traditionally, grasping has been studied under the umbrella of dynamic modeling and control, considering contact force and wrenches [13], which require the object model or force/tactile sensing. Recent visual-based approaches are more focused on data-driven methods, using either prior knowledge like 3D models [14] or points cloud [15] to learn objects model. These approaches are very interesting, but require the knowledge of the objects models. Manipulation actions such as push-pulls are medium/low-level actions used to bring the environment to a more convenient configuration to solve the robot’s task [16], [17]. Planning the sequence of appropriate actions to solve a task is a difficult decision problem per se, and combining grasping and pushing-pulling to pick cluttered objects is a challenge [18], [19], [20]. Several works learn the policy using model-free optimization [21], [22], [23], [24], with the issue of exploring a large space. To grasp an object that is partially occluded, or bring the objects in the workspace into a specific configuration that is more favorable to grasp an object of interest, a promising approach is to leverage visual prediction models [3] to inform a visual MPC method. The principle consists in predicting the visual output of the robot camera after executing an action or a sequence of actions, and to use this prediction to inform an optimal controller or decision planner.

Visual Model Predictive Control (MPC) is an appealing method to address visual-based manipulation. Traditionally,

model predictive control (MPC) relies on known or learned dynamics robot models, and it is used to generate advanced motor controls. Once the model of the effect of the robot’s actions on its environment (e.g., workspace with objects) is known or learned, it can be used to plan a sequence of optimal actions to fulfill a task: this can be done with MPC, often with a receding horizon approach, but also with reinforcement learning [25], [26], [27]. For planning sequence of manipulation actions, this was often done with reduced visual models computing features of the objects in the scene, to lower the dimensionality of the problem. With the increased availability and capabilities of computing clusters, data-driven methods and deep learning technologies [28] can now be used to build models directly from high dimensional input such as images. Visual models that leverage deep learning have been used to address visually dominant tasks such as manipulation and navigation, to learn models used for planning and decision: a visual MPC is therefore possible. The main difficulty of dealing with visual models lies in the bigger state space, with discontinuities that make planning actions more challenging. This is particularly challenging for grasping: for example, objects disappearing from the workspace cause visual discontinuities. In [3], sample-based planning based on CEM [29] was used for visual model predictive control.

Video prediction model went through a significant development in the last years. Initial video/visual prediction model were based on deterministic models [30], [6], [31], [32], [33], [34], [35], [36]. Their main limit is dealing with uncertainty, more precisely with the uncertain or stochastic outcomes. Probabilistic models [37] were further proposed to carry the “stochasticity” information, in particular VAE-based [38] stochastic video prediction models [39], [7], [40], [11]. [7] proposed to use a learned prior to model time-variational stochasticity; [11] used a hierarchical latent variational to model the stochasticity. [10] investigated improving performance by increasing model size. Some other works use discrete represent of pixels [41], [42] to optimize the likelihood directly through autoregressive models [43] and transformers [44], [45]. Autoregressive models have the advantage of producing higher quality images; the models using transformers structures can be trained in parallel and faster, although the pixel-level prediction that needs pixel by pixel reconstruction usually takes a long time and limits its use in online control tasks.

III. MATERIALS & METHODS

Robot setup: We use a Franka Panda robot equipped with a gripper to manipulate rigid and deformable/soft objects organized in unstructured heaps inside a plastic gray box. A Intel RealSense camera in front of the robot is used to capture the workspace and the robot’s terminal part.

Datasets: In this paper, we use two main datasets. **1) RoboNet:** this is a large scale dataset with more than 150K trajectories executed by 7 different robots, for a total of 15 million frames [5]. **2) PandaGrasp:** this is our “smaller” dataset recorded with our setup, with the

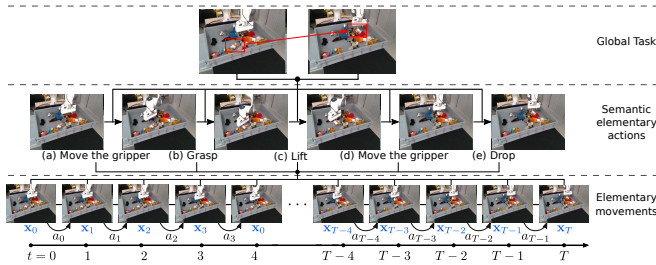


Fig. 2: The hierarchical structure of decomposition semantic actions. A semantic action like sorting objects by grasping contains five element actions: (middle) (a) moving the gripper to a top position of the grasp point, (b) falling the gripper grasps the object, (c) lifting the gripper back to the top position, (d) moving the gripper to the top of the target point, (e) opening the gripper to drop the object. (bottom) Each action element is decomposed into several action movements to a sequence of images and actions.

Panda robot executing 5K trajectories for a total of 150K image frames. This is a new dataset that we are releasing as open: https://gitlab.inria.fr/sivaldi/PandaGrasp_dataset. As in RoboNet, PandaGrasp was acquired via autonomous exploration. However, our exploration strategy was more “efficient” for acquiring more grasping samples. Indeed, exploration through random actions or babbling in a large space lead to datasets with a majority of non-successful and non-purposive actions that have little to no effect on the visual appearance of the workspace. In contrast, we are interested into visual sequences where the environment significantly changes as a result of grasping actions and manipulations. Finn et al. [46] used a “reflex” primitive that automatically closes the gripper when it is lower than a threshold to make meaningful actions more frequent. Since we are interested into grasping, we defined an heuristic such that the robot only attempts to grasp random points sampled from the surface of the heap, extracted from the 3D point cloud: in this way, there is a small possibility of grasping in an empty area.

Our method: Action conditional stochastic Visual Prediction for Grasping Objects (VP-GO)

We propose a VP-GO: a VAE-based, action conditional stochastic visual prediction model that is specific to deal with high stochasticity such as in robotic grasping unknown objects in a heap. Our model is inspired by SVG-LP [7] and its successors SVG’ and GHVAE, but it is “lighter” than those, in that it only contains simple convolutional and recurrent networks and has a smaller number of parameters to train. Furthermore, it is an action conditional model that can deal with semantic action description. We propose a hierarchical structure to decompose high-level semantic actions into more elementary robot movements, compatible with existing models and datasets. Our model can be pre-trained with large-scale datasets such as RoboNet and then fine-tuned on specific smaller-scale datasets such as our PandaGrasp dataset. The code is available here: <https://gitlab.inria.fr/sivaldi/HEAP-VP-GO>.

In the following, we detail the two main parts of VP-GO:

1) Action decomposition – from semantic actions to elementary movements:

The action conditional visual prediction model is a deep neural network model that inputs the current and past images and generates a sequence of predicted images of the future. In general, the data used to train the model is organized into a sequence of pairs of images $(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{T-1}, \mathbf{x}_T)$ and actions $(\mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{T-1})$. In [3] the action $\mathbf{a} \in \mathbb{R}^n$ is defined as a Δ -displacement of the end-effector, where n is the degree of freedom of the end-effector. This representation fits for describing push-pull and simple gripper movements or babbling actions. However, it is very limited to describe more complex manipulations. In contrast, in our application the robot should purposefully interact with objects, executing many grasps: our ultimate goal is to find sequences of grasping actions to execute on the objects in the heap. For this reason, we use not only low-level displacement actions but also high-level semantic actions such as grasping and picking. In order to integrate semantic actions into a generic prediction model, while being compatible with previous video prediction models and corresponding datasets (e.g., RoboNet), we propose a hierarchical structure of high-level semantic robot actions. High-level actions such as “grasp” can be decomposed into several elementary movements that can directly be trained in a general model, as shown in Fig. 2. We split a semantic grasping action into several action elements: moving the gripper to a top position over the grasp point, lowering the gripper to the grasp point and closing it, lifting the gripper back to the top position, moving the gripper to the top of the target point, and opening the gripper to drop the object. Finally, we decompose each elementary action into several elementary movements that are defined as a displacement of the end effector’s elements.

2) Action conditioned stochastic visual prediction model:

As shown in the left part of Fig. 3, the stochastic prediction model takes c visible frames $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_c$ and an action sequence $\mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{T-1}$ to predict several futures $\mathbf{x}_{c+1}, \dots, \mathbf{x}_{T-1}, \mathbf{x}_T$. To deal with the stochastic nature of the real world, the VAE-based prediction model introduces latent variables $\mathbf{z} \sim p(\mathbf{z})$ that carries the stochastic information. We can generate the image \mathbf{x}_t at time t from a prediction model $p_\theta(\mathbf{x}_t | \mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t})$ conditioned on the previous images $\mathbf{x}_{0:t-1}$, actions $\mathbf{a}_{0:t-1}$ and the latent variables $\mathbf{z}_{1:t}$. Since $p_\theta(\mathbf{x}_t)$ can not be directly maximized over the distribution of $p(\mathbf{z}_t)$, the VAE-based method approximates the posterior $p(\mathbf{z}_t | \mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1})$ by an inference network that is parameterized as a conditional Gaussian distribution $q_\phi(\mathbf{z}_t | \mathbf{x}_{0:t}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}) = \mathcal{N}(\mu_\phi(\mathbf{x}_{0:t}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}), \sigma_\phi(\mathbf{x}_{0:t}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}))$.

Previous work [39] assumed that $p(\mathbf{z}_t)$ is a fixed Gaussian distribution $\mathcal{N}(0, \mathbf{I})$. In contrast, to catch the rich stochasticity that is caused in our application by the robot’s grasping action, we assume that the learned-prior $p(\mathbf{z}_t)$ as in [7] varies across time, conditioned on actions and previous images. Specifically, the learned prior is also parameterized as a conditional

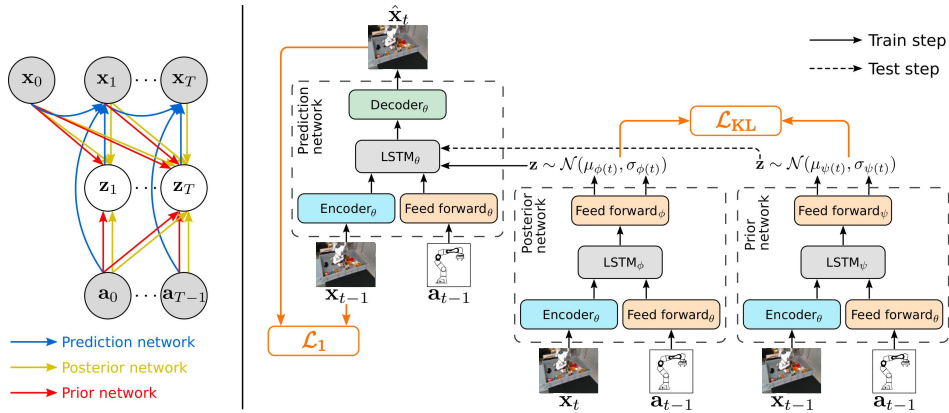


Fig. 3: *Left*: Probabilistic model of VP-GO. Blue, yellow and red lines indicate the prediction, posterior and prior models respectively. *Right*: Detailed structure of VP-GO. Each network contains deep convolutional layers as the encoder for images, feed forward network for actions and LSTM layers to condition the previous input. Prior network and posterior network contain a feed forward network which outputs the mean and variance of the latent variables after the LSTM layers. In the prediction network, deep convolutional layers (decoders) are used to generate prediction images.

Gaussian distribution as $p_\psi(\mathbf{z}_t | \mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}) = \mathcal{N}(\mu_\psi(\mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}), \sigma_\psi(\mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1}))$. We train the entire model by maximizing the variational lower bound of the log-likelihood as in the variational autoencoders:

$$\mathcal{L}_{\theta, \phi, \psi}(\mathbf{x}_{c+1:T}) = \sum_{t=c+1}^T ELP - \sum_{t=c+1}^T D_{KL}(QP||PP) \quad (1)$$

where $E = \mathbb{E}_{q_\phi(\mathbf{z}_t | \mathbf{x}_{0:t}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1})}$,
 $LP = \log p_\theta(\mathbf{x}_t | \mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t})$,
 $QP = q_\phi(\mathbf{z}_t | \mathbf{x}_{0:t}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1})$,
 $PP = p_\psi(\mathbf{z}_t | \mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t-1})$.

We parametrize the prediction network with a fixed-variance Laplace distribution $\hat{\mathbf{x}}_t \sim \text{Laplace}(b_\theta(\mathbf{x}_{0:t-1}, \mathbf{a}_{0:t-1}, \mathbf{z}_{1:t}), \sigma)$; in other words, we reduce the first term on the right hand side to ℓ_1 loss instead of ℓ_2 loss between $\hat{\mathbf{x}}$ and \mathbf{x} to help generate sharper images [10]. During training, the latent variables \mathbf{z}_t are sampled from posterior $q_\phi(\mathbf{z}_t)$ that inputs the ground truth \mathbf{x}_{t+1} , and the second term on the right hand side is a KL-divergence that fits the posterior $q_\phi(\mathbf{z}_t)$ to the learned-prior $p_\psi(\mathbf{z}_t)$. During testing, since the ground truth is unavailable, we directly sample \mathbf{z}_t from the learned prior $p_\psi(\mathbf{z}_t)$.

Network Structure: As shown in Fig. 3 (right), we use deep encoder with 16 vgg layers to decrease the spatial size and increase the channel dimensions of images from $48 \times 64 \times 48$ pixels to $3 \times 4 \times 512$, and a single dense feed forward network to encode actions and reshape it to $3 \times 4 \times 2$. The encoded actions and latent variables are appended to encoded images along the channel dimension. The dependency with previous inputs is achieved by the recurrent layers. For the prediction network, we use 2 convolutional LSTM layers. For prior and posterior networks, we use one convolutional LSTM layer. The latent variables are output from a dense layer following the LSTM and the predicted image finally can be generated through a decoder that has similar structure

as the encoder.

Metrics for evaluation: As done in prior work [39], [7], [10], [47], [11], we evaluate video prediction models across 4 metrics: Structural Similarity Index Measure (SSIM) [48], Peak Signal-to-noise Ratio (PSNR) [49], Learned Perceptual Image Patch Similarity (LPIPS) [50] and Fréchet Video Distance (FVD) [51]. To compare PSNR, SSIM, LPIPS, [39], [7] sample 100 rollouts for each sample and select the best trajectory. FVD uses instead all the 100 samples [10]. In contrast with prior work only reporting the best score among the 100 samples, we also report the average score: the reason is that we want a better idea of the real performance of the model that will be used for downstream planning (where only one sample is used, if we refer to [3]) to take decisions on the next action to perform. We adopted the specific settings of [11]: we generate predictions conditioned on 2 context frames and a sequence of actions and evaluate on a 10 rollout horizon; since FVD metric can be significantly disturbed by the different batch sizes, we use batch size of 256 for FVD.

IV. EXPERIMENTS

The experiments are designed to answer to the questions:

- Q1:** Can we achieve state-of-the-art performance with our “lighter” model VP-GO?
- Q2:** Can we improve the prediction performance by using a deeper network?
- Q3:** Can we improve the prediction performance by considering the robot’s state (precisely, the end-effector position) as additional input in the network and actions?
- Q4:** Can the movement segments decomposed from the semantic actions be directly trained in the general video prediction model?
- Q5:** Can we improve the performance with fine-tuning, i.e., by leveraging a pre-trained model, instead of training a model from scratch, since we only have a small dataset?

A1: Comparing our “lighter” model VP-GO with GHVAE and SVG? We compare our “lighter” model VP-

TABLE I: A1: Comparison with state-of-the-art models (mean \pm standard error)

Dataset	Model	Video Prediction Performance (Test)			
		FVD \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
RoboNet	GHVAEs	95.2\pm2.6	24.7 \pm 0.2	89.1\pm0.4	0.036 \pm 0.001
	SVG'	123.2 \pm 2.6	23.9 \pm 0.1	87.8 \pm 0.3	0.060 \pm 0.008
	Ours	111.5 \pm 1.3	26.9\pm0.2	89.1\pm0.3	0.030\pm0.001

TABLE II: A2: Ablation study of the encoder and decoder networks depth (mean \pm standard error)

Dataset	Layer config.	Video Prediction Performance (Test)			
		FVD \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
RoboNet	VGG16	232.8 \pm 1.9	24.5 \pm 0.3	86.8 \pm 0.3	0.078 \pm 0.001
	layer conv3_3				
	VGG16	129.3 \pm 1.5	26.6 \pm 0.2	88.7 \pm 0.3	0.035 \pm 0.001
	layer conv4_3				
	VGG19				
layer conv4_4	111.5\pm1.3	26.9\pm0.2	89.1\pm0.3	0.030\pm0.001	

GO with the “heavier” GHVAE [11] and SVG' [10]. GHVAE is a hierarchical model; SVG' is similar to our model but with shallow convolutional layers for encoder and decoder. Both have a larger number of parameters than ours: 599 million for GHVAE, 298 million for SVG', and 129 million for ours. Since their code is not available, to compare the performance on RoboNet we directly use the results reported in [11]. For fair comparison, we use the default random seed in RoboNet to split train/test set and no additional states information while testing our model. The states information is only given in the ablation study A3 for contribution of the robot's state. Table I shows the performance of the video prediction model on the RoboNet dataset. Our model outperforms SVG' vastly and achieves a comparable performance to GHVAE; specifically, we exceed GHVAE on both structured metric PSNR and human perceptual metric LPIPS and get similar performance on SSIM but behind on FVD. The video attachment shows more examples of RoboNet, with several rollouts predicted by ours and the other models [11]. Our results suggest that a lighter model can have substantial improvement of performance by adjusting the internal structure, e.g., using deeper convolutional layers. Also, the drastic reduction of the number of parameters (from 599 to 129 million) did not lower the performance and at the same time enabled us to run the training on a smaller cluster.¹

A2: Ablation study of the encoder and decoder networks depth Compared with SVG' [10], one notable difference is that we use a deeper network configuration for encoder and decoder networks. In this section experiment, we report an ablation study to show whether using a deeper VGG [52] layer configuration helps. As shown in Table II, we train our models in three different configurations: VGG16 up to layer conv3_3, VGG16 up to layer conv4_3, and VGG19 up to layer conv4_4 which is the default of our model. The

¹Training our model for 400k steps on a 11GB GPU cluster (4 \times GeForce RTX 2080 Ti) takes 4.7 days. We could not re-implement and re-train GHVAE as reported in [11] since the paper reports using 24GB/48GB GPU machines.

deeper model VGG19 outperforms the others. The results confirm the observation of [52] that “deeper models might be beneficial for larger datasets” and support the hypothesis in [11] that deep models with compressed height and width reduce the spatial correlations.

A3: Ablation study: Contribution of the robot's state

A large part of the stochasticity comes from estimating the position of the robot in the real world through images. However, the robot's state, in particularly the position of the end-effector in its workspace is easily accessible.² We posit that it could help that estimation. In this experiment, we do an ablation study to show whether adding the robot's state to the action conditional model helps. We evaluate the contribution of the robot's state on both the best performing model and on the average score from 100 samples models. Results are shown in Table III. In both cases, using the robot state as an additional input improves the model's performance. Furthermore, using the robot state as an additional input can improve much more in the average score than in the best score. PSNR improves 1.1 in average score case than 0.5 in the best score; SSIM improves 2.0 in average score than 0.8 in the best score, LPIPS improves 0.003 in one shot case than 0.001 in the best score. These results also confirm the intuition from [11] that a large part of the uncertainty comes from estimating the robot's position from images.

A4: Our visual prediction model with semantic actions

We investigate the performance of using our VP-GO for different semantic actions in Fig. 5 and 6 and Table IV. Fig. 5 shows an example of prediction of an entire trajectory after the robot picks and moves an orange soft object. We can successfully predict the object's movement, compared to the ground truth. However, when the robot is approaching the target object, the metrics in Fig. 6 and Table IV show that there is almost no performance loss in the predicted frames; when the robot begins to grasp and interacts with the objects, the performance decreases, and then keep decreasing linearly with the number of frames we look ahead in the future. This is likely due both to the higher stochasticity given by the grasping action, and the intrinsic higher uncertainty about predicting an action in the far future and may be alleviated by reducing the prediction horizon in semantic action-based planning. With such results, VP-GO shows promising potential to be used for planning sequences of robot manipulation actions in the semantic action space, which is much more efficient than using sample-based planning in the movement action space. Still, we observe that a precise visual prediction of a moving soft object is still challenging.

A5: Finetuning with PandaGrasp Recording large scale datasets such as RoboNet [5] is very expensive in terms of time and resources. Our dataset PandaGrasp is “smaller” in comparison, yet for many practical applications we want our models to work with datasets of this size. Hence, we want to investigate if a model trained with a “smaller” dataset is performing as well as a model trained with a “bigger”

²In the offline test, we use the ground truth robot states as robot states and actions are implemented by a low-level closed-loop position controller and easy to make assumption as $s_{t+1} \equiv \hat{s}_{t+1} = s_t + at$.

TABLE III: A3: Contribution of the robot’s state (mean \pm standard error)

Dataset	Conditional Input		Video Prediction Performance (Test)						
	Action	State	FVD \downarrow	PSNR \uparrow		SSIM \uparrow		LPIPS \downarrow	
				Best	Average	Best	Average	Best	Average
RoboNet	•	•	111.5\pm1.3	26.9 \pm 0.2	25.0 \pm 0.2	89.1 \pm 0.3	85.9 \pm 0.4	0.030 \pm 0.001	0.039 \pm 0.001
			116.8 \pm 1.0	27.4\pm0.2	26.1\pm0.2	89.9\pm0.3	87.9\pm0.4	0.029\pm0.001	0.036\pm0.001
	Improvement \uparrow		--	+0.5	+1.1	+0.8	+2.0	+0.001	+0.003

TABLE IV: A4, 5: The comparison between models trained from scratch and fine-tuning by FVD and averaged scores of PSNR, SSIM, and LPIPS on the test set of semantic action. (mean \pm standard error)

Dataset	Stage	Video Prediction Performance (Test)							
		PSNR \uparrow		SSIM \uparrow		LPIPS \downarrow		FVD \downarrow	
		from scratch	fine-tuning	from scratch	fine-tuning	from scratch	fine-tuning	from scratch	fine-tuning
PandaGrasp	Approaching	34.2\pm0.2	33.9 \pm 0.2	98.7\pm0.1	98.6 \pm 0.1	0.004 \pm 0.001	0.004 \pm 0.001		
	Grasping	31.0 \pm 0.2	31.0 \pm 0.2	97.1\pm0.1	97.0 \pm 0.1	0.009 \pm 0.001	0.009 \pm 0.001		
	Moving	25.3 \pm 0.2	25.9\pm0.2	90.8 \pm 0.1	91.3\pm0.1	0.034 \pm 0.002	0.032\pm0.002	110.9 \pm 3.9	110.9 \pm 4.0
	Average	30.1 \pm 0.2	30.1 \pm 0.2	95.5 \pm 0.2	95.6\pm0.2	0.016 \pm 0.001	0.015\pm0.001		
	Final Goal	23.7 \pm 0.4	23.9\pm0.3	87.0 \pm 0.7	87.7\pm0.6	0.050 \pm 0.004	0.047\pm0.004		

dataset and fine-tuned on the “smaller”. Also, PandaGrasp is more purposeful for semantic grasping actions compared to the low-level actions (for babbling-like exploration) in RoboNet. Our experiment also verifies whether our method can let prior low-level action knowledge benefits the training on more purposeful data with semantic actions. Specifically, we compare our model VP-GO trained on PandaGrasp from scratch with a VP-GO pre-trained on RoboNet then fine-tuned on PandaGrasp. As shown in Fig. 4 and Table IV, the two models are performing in the same way, from the point of view of the metrics FVD, PSNR, SSIM and LPIPS. However, qualitative inspection of the visual predictions performed by the two models actually shows that the fine-tuned model can correctly predict missing and misplaced objects after, while the scratch model fails. One typical problem is failing at predicting objects that move from the background and leave empty space: Fig. 4 shows an example where an object is grasped and moved from the initial location: the object is still in the same place in the prediction of the first model learned from scratch, while the fine-tuned model correctly predicts that the object is moved away. Fig. 7 shows the corresponding performance metrics. Our intuition is that training on a big dataset such as RoboNet produces a more “expert” model, and fine-tuning on our smaller dataset PandaGrasp just makes sure that the predicted frames are more coherent with our specific setup.

V. CONCLUSION

We proposed VP-GO, a stochastic action-conditioned visual prediction model that can be used to predict future visual outcomes after grasping and manipulating irregular cluttered objects. Our model is inspired by SVG-LP [7], but it introduces action conditioning, uses convolutional LSTM and has a deeper structure. With respect to the state-of-the-art models [10], [11], VP-GO is lighter in that it has considerably less parameters; still it has comparable performance. This is promising, since lighter models are




Fig. 4: A5: Visual comparison of our model VP-GO trained on PandaGrasp from scratch with a VP-GO pre-trained on RoboNet then fine-tuned on PandaGrasp.




Fig. 5: A4: Visual inspection of future frames predicted by VP-GO from semantic actions. The model correctly predicts the orange push toy being grasped and moved.

more sustainable and reproducible. We also contributed with a new open dataset “PandaGrasp” that contains 5K sequences of robot grasping trajectories, executed by a Franka Panda robot, according to a proposed hierarchical decomposition of semantic actions into elementary movements. In our experiments, we found that pre-training our model with a large-scale dataset (RoboNet) then fine-tuning the model with our smaller dataset PandaGrasp, specific to our robotics setup, gave better results in terms of predicting future frames after grasping and moving objects.

The next step is to use our visual prediction model in a RL framework to plan a sequence of grasps to selectively remove some objects of interests from a heap of cluttered objects. Interestingly, our model can benefit of high-level semantic actions, reducing the complexity of downstream planning.




Fig. 6: A4, 5: The comparison between models trained from scratch and fine-tuning by averaged scores of PSNR, SSIM, and LPIPS on the test set of semantic action.




Fig. 7: A5: An example of the limitation of the metrics for evaluating the performance in terms of grasping task. In terms of all three metrics, the model performance equivalent that contrary to the comparison in Fig. 4 in terms of a grasping task.

VP-GO does not yet explicitly uses the semantic action label, but leverages the implicit semantic information carried out by the PandGrasp dataset during training. In the future, we plan to explicitly condition the model on the additional semantic information.

REFERENCES

- [1] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, 2011.
- [2] M. Koskinopoulou, F. Raptopoulos, G. Papadopoulos, N. Mavrikis, and M. Maniatakis, "Robotic waste sorting technology: Toward a vision-based categorization system for the industrial robotic separation of recyclable waste," *IEEE Robotics Automation Magazine*, vol. 28, no. 2, pp. 50–60, 2021.
- [3] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine, "Visual foresight: Model-based deep reinforcement learning for vision-based robotic control," 2018.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [5] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn, "Robonet: Large-scale multi-robot learning," *CoRR*, vol. abs/1910.11215, 2019. [Online]. Available: <http://arxiv.org/abs/1910.11215>
- [6] C. Finn, I. Goodfellow, and S. Levine, "Unsupervised learning for physical interaction through video prediction," 2016.
- [7] E. Denton and R. Fergus, "Stochastic video generation with a learned prior," 2018.
- [8] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," *CoRR*, vol. abs/1610.00696, 2016. [Online]. Available: <http://arxiv.org/abs/1610.00696>
- [9] B. Ridge, D. Skočaj, and A. Leonardis, "Self-supervised cross-modal online learning of basic object affordances for developmental robotic systems," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 5047–5054.
- [10] R. Villegas, A. Pathak, H. Kannan, D. Erhan, Q. V. Le, and H. Lee, "High fidelity video prediction with large stochastic recurrent neural networks," *CoRR*, vol. abs/1911.01655, 2019. [Online]. Available: <http://arxiv.org/abs/1911.01655>
- [11] B. Wu, S. Nair, R. Martin-Martin, L. Fei-Fei, and C. Finn, "Greedy hierarchical variational autoencoders for large-scale video prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 2318–2328.
- [12] S. Ivaldi, N. Lyubova, D. Gérardeaux-Viret, A. Droniou, S. M. Anzalone, M. Chetouani, D. Filliat, and O. Sigaud, "Perception and human interaction for developmental learning of objects and affordances," in

- 2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012), 2012, pp. 248–254.
- [13] B. Siciliano and O. Khatib, *Springer handbook of robotics*. Springer, 2016.
- [14] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The columbia grasp database," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 1710–1716.
- [15] A. Zeng, K.-T. Yu, S. Song, D. Suo, E. Walker, A. Rodriguez, and J. Xiao, "Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 1386–1383.
- [16] D. Omrčen, C. Böge, T. Asfour, A. Ude, and R. Dillmann, "Autonomous acquisition of pushing actions to support object grasping with a humanoid robot," in *2009 9th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2009, pp. 277–283.
- [17] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun, "Driving policy transfer via modularity and abstraction," *CoRR*, vol. abs/1804.09364, 2018. [Online]. Available: <http://arxiv.org/abs/1804.09364>
- [18] M. Dogar and S. Srinivasa, "A planning framework for non-prehensile manipulation under clutter and uncertainty," *Autonomous Robots*, vol. 33, no. 3, pp. 217–236, June 2012.
- [19] A. Boularias, J. A. Bagnell, and A. Stentz, "Learning to manipulate unknown objects in clutter by reinforcement," ser. AAAI'15. AAAI Press, 2015, p. 1336–1342.
- [20] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," 2018.
- [21] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," *CoRR*, vol. abs/1412.3128, 2014. [Online]. Available: <http://arxiv.org/abs/1412.3128>
- [22] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," 2015.
- [23] M. Gualtieri, A. ten Pas, K. Saenko, and R. P. Jr., "High precision grasp pose detection in dense clutter," *CoRR*, vol. abs/1603.01564, 2016. [Online]. Available: <http://arxiv.org/abs/1603.01564>
- [24] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *CoRR*, vol. abs/1703.09312, 2017. [Online]. Available: <http://arxiv.org/abs/1703.09312>
- [25] A. S. Polydoros and L. Nalpantidis, "Survey of model-based reinforcement learning: Applications on robotics," *Journal of Intelligent & Robotic Systems*, vol. 86, no. 2, pp. 153–173, 2017.
- [26] A. massoud Farahmand, A. Shademan, M. Jagersand, and C. Szepesvári, "Model-based and model-free reinforcement learning for visual servoing," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 2917–2924.
- [27] M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. Johnson, and S. Levine, "Solar: Deep structured representations for model-based reinforcement learning," in *International Conference on Machine Learning*, 2019, pp. 7444–7453.
- [28] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. MIT press Cambridge, 2016, vol. 1, no. 2.
- [29] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L'Ecuyer, "The cross-entropy method for optimization," in *Handbook of statistics*. Elsevier, 2013, vol. 31, pp. 35–59.
- [30] J. Walker, A. Gupta, and M. Hebert, "Dense optical flow prediction from a static image," 2015.
- [31] T. Xue, J. Wu, K. L. Bouman, and W. T. Freeman, "Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks," 2016.
- [32] X. Liang, L. Lee, W. Dai, and E. P. Xing, "Dual motion gan for future-flow embedded video prediction," 2017.
- [33] J. Van Amersfoort, A. Kannan, M. Ranzato, A. Szlam, D. Tran, and S. Chintala, "Transformation-based models of video sequences," *arXiv preprint arXiv:1701.08435*, 2017.
- [34] Z. Liu, R. A. Yeh, X. Tang, Y. Liu, and A. Agarwala, "Video frame synthesis using deep voxel flow," 2017.
- [35] B. Chen, W. Wang, J. Wang, and X. Chen, "Video imagination from a single image with transformation generation," 2017.
- [36] C. Lu, M. Hirsch, and B. Scholkopf, "Flexible spatio-temporal networks for video prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6523–6531.

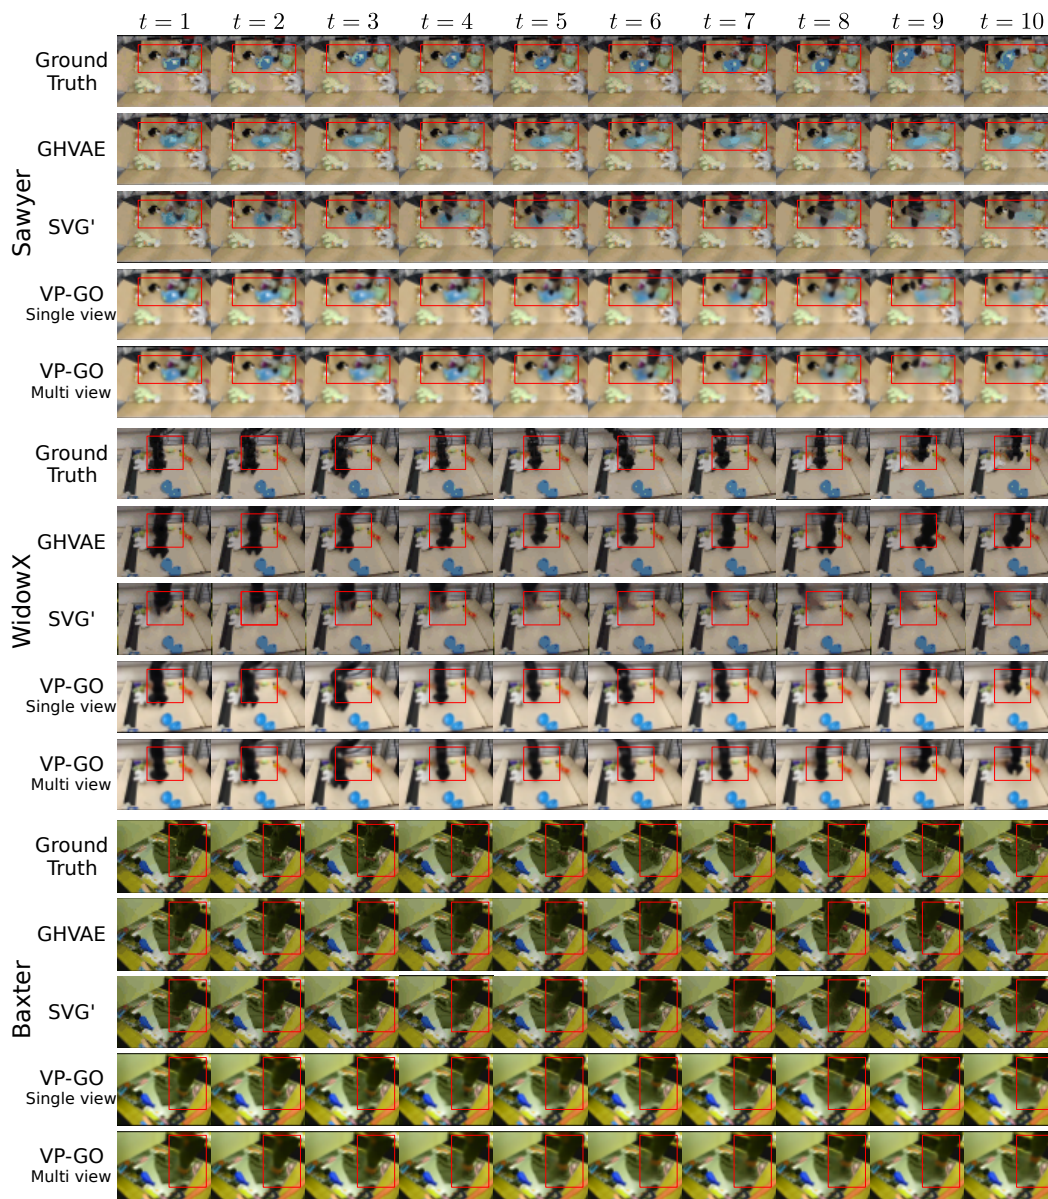


Fig. 8: Visual comparison between VP-GO and prior work. The examples and results of GHVAE and SVG' were provided in [11]. In order to give the same context images and actions to our model, we manually identified the same examples from the RoboNet dataset. In these examples, our model outperforms SVG', and is close to GHVAE in generating realistic images. Our model is also more accurate at predicting movements (see WidowX and Baxter).

- [37] G. Rasch, *Probabilistic models for some intelligence and attainment tests*. ERIC, 1993.
- [38] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013.
- [39] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine, "Stochastic variational video prediction," 2017.
- [40] A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine, "Stochastic adversarial video prediction," 2018.
- [41] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma, "Pixelcnn++: A pixelcnn implementation with discretized logistic mixture likelihood and other modifications," in *ICLR*, 2017.
- [42] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, "Pixel recurrent neural networks," 2016.
- [43] H. Akaike, "Fitting autoregressive models for prediction," *Annals of the institute of Statistical Mathematics*, vol. 21, no. 1, pp. 243–247, 1969.
- [44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [45] D. Weissenborn, O. Täckström, and J. Uszkoreit, "Scaling autoregressive video models," 2020.
- [46] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2786–2793.
- [47] L. Castrejon, N. Ballas, and A. Courville, "Improved conditional vrnns for video prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 7608–7617.
- [48] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [49] Q. Huynh-Thu and M. Ghanbari, "Scope of validity of psnr in image/video quality assessment," *Electronics letters*, vol. 44, no. 13, pp. 800–801, 2008.
- [50] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [51] T. Unterthiner, S. van Steenkiste, K. Kurach, R. Marinier, M. Michalski, and S. Gelly, "Fvd: A new metric for video generation," 2019.
- [52] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.