

How Physics and Background Attributes Impact Video Transformers in Robotic Manipulation: A Case Study on Planar Pushing

Shutong Jin, Ruiyu Wang, Muhammad Zahid and Florian T. Pokorny

Abstract—As model and dataset sizes continue to scale in robot learning, the need to understand how the composition and properties of a dataset affect model performance becomes increasingly urgent to ensure cost-effective data collection and model performance. In this work, we empirically investigate how physics attributes (color, friction coefficient, shape) and scene background characteristics, such as the complexity and dynamics of interactions with background objects, influence the performance of Video Transformers in predicting planar pushing trajectories. We investigate three primary questions: How do physics attributes and background scene characteristics influence model performance? What kind of changes in attributes are most detrimental to model generalization? What proportion of fine-tuning data is required to adapt models to novel scenarios? To facilitate this research, we present *CloudGripper-Push-1K*, a large real-world vision-based robot pushing dataset comprising 1278 hours and 460,000 videos of planar pushing interactions with objects with different physics and background attributes. We also propose Video Occlusion Transformer (VOT), a generic modular video-transformer-based trajectory prediction framework which features 3 choices of 2D-spatial encoders as the subject of our case study. The dataset and source code are available at <https://cloudgripper.org>.

I. INTRODUCTION

The transformer architecture [1] has established itself as a highly effective tool across a diverse array of sequence-to-sequence problems in natural language processing [1], [2], [3], computer vision [4], [5], [6], decision making [7], [8], [9] and robotic manipulation [10], [11]. Video Transformers (VTs) [12], [13], [14] are widely studied for video-based downstream tasks and their ability to effectively capture complex spatio-temporal relationships and long-range dependencies makes them particularly intriguing for robot perception and planning, e.g., visually-grounded robotic manipulation through imitation learning [15], [16], [17], robotic video prediction [18], [19], [20] and goal-conditioned planning with video demonstrations [21], [22], [23]. However, VTs are computationally expensive [24] and their training requires large video datasets [25] that are usually costly to collect on real robots. Both the computational cost of training VTs from scratch and the requirement for large-scale data collection are challenges that may currently still hinder their widespread adoption in real-world robotic manipulation tasks. To alleviate this burden, a new set of empirical

The authors are with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology shutong@kth.se; fpokorny@kth.se. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The computations were enabled by the supercomputing resource Berzelius provided by the National Supercomputer Centre at Linköping University and the Knut and Alice Wallenberg Foundation, Sweden.

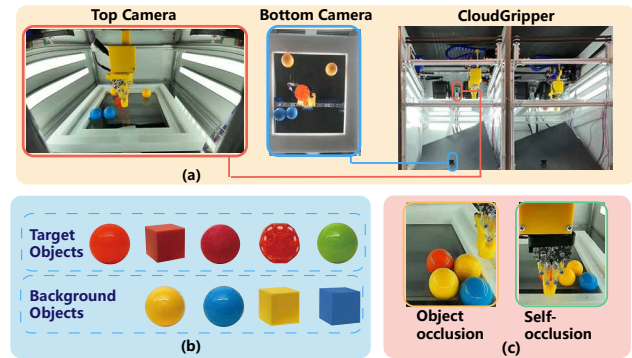


Fig. 1: An illustration of the data collection process on the *CloudGripper* platform - an open source cloud robotics testbed with 32 robot arm cells. We collect planar robot pushing interaction videos from two camera views. (a) Collected top and bottom camera views and *CloudGripper* robot cell side-view. (b) An illustration of target and background objects used in the case-study. (c) Illustration of object and gripper self-occlusion present in collected top camera videos.

evidence is required to understand how the composition and properties of a dataset affect model performance and how much data is typically required for training and fine-tuning in the context of robotic manipulation. The impact of physics attributes (color, friction coefficient, shape) and scene background characteristics present in video datasets are aspects that are at present poorly understood in the context of emerging VT architectures for robotic manipulation. Since a rigorous theoretical framework describing the impact of these factors is at present unavailable to the robotics research community, our work proposes a first large-scale empirical case based on real-world data as a step towards developing an understanding on how these factors influence generic VT model performance in practice.

Towards this end, we introduce *CloudGripper-Push-1K*, a large real-world visual-based robot pushing dataset encompassing various target objects (*Ball, Cube, Foam, Icosahedron*) with distinct physics attributes (color, friction coefficient, shape). We furthermore vary the number and type of background objects that are placed in the same workspace as the primary object which the robot is interacting with. This approach creates variations in background complexity and physical dynamics due to secondary collisions during pushing interactions, resulting in 18 sub-datasets. The dataset is collected with the *CloudGripper* system [26] and includes 1.4 TB, 1278 hours and 460,000 trajectories of robot pushing interactions, resulting in one of the largest robot datasets pub-

licly available to date. We collect videos from a dual-camera system: a top camera view providing a scene overview with occlusions, and a bottom camera perspective that allows for an occlusion-free observation of object position as illustrated in Fig.1(a). This video dataset, with its controlled variations in physics and background attributes will be made available and serves as a primary dataset resource for the presented case study. We believe that this dataset may also be useful for the community in understanding general occlusion problems and occlusion-aware networks with single camera input, providing an additional perspective distinct from other approaches assuming calibrated RGB-D [27] or multi-camera setups [28].

In alignment with the generic approach outlined in VTN [29], we designed Video Occlusion Transformer (*VOT*), a modular framework consisting of a flexible 2D-spatial encoder (*VOT-MaxViT*, *VOT-MaxViT-2*, *VOT-Swin-T*) and a temporal encoder. *VOT* takes the occluded top camera videos as inputs and predicts the trajectories of target objects in the pixel coordinate of the bottom camera. To study the generalisability of VTs, the framework is trained from scratch across all tasks available in *CloudGripper-Push-1K*. Zero-shot and few-shot learning are conducted for tasks with various physics attributes and scene background characteristics. The main contributions of this paper are:

- We present a large real-world video dataset *CloudGripper-Push-1K* of 1278 hours of planar robot-object pushing interactions which enables the study of VTs and other vision-based large models for robotic manipulation. The dataset will be released as an open dataset to the community.
- A generic Video Transformer framework *VOT* covering three types of 2D-spatial encoders is designed for this specific robot-pushing setting to predict trajectories of target objects.
- We present ablation studies by training across 18 sub-datasets from scratch and performing zero-shot evaluation and fine-tuning on sub-datasets, providing empirical evidence on how physics attributes and scene background characteristics influence the performance of VTs.

II. RELATED WORK

A. Model and Dataset Scalability

While transformers have shown significant potential across various application areas [3], [5], [11], their quadratic complexity with respect to the number of available tokens imposes substantial challenges on both model training and data collection [10]. Recent efforts have been initiated to investigate the scaling behavior with respect to model size, model depth, dataset scale and computing power [30], [31], [32], [33]. These studies however mostly focused on the impact of the number of tokens on scaling and have not addressed how the composition and properties relating to the training dataset influence the generalisability of the models.

Some endeavors have been made to determine what information is most effective to be included in training data for

model pre-training. *Jing et al.* [34] investigate the effects of pre-training datasets, model architectures and training methods on visual pre-training for robot manipulation. In *KitchenShift* [35], the authors study the zero-shot generalization of an imitation-based policy under domain shifts, such as lighting and camera pose. However, most of the aforementioned experiments and benchmarks [36], [37] are designed in simulated environments. For robotic manipulation, even high-fidelity simulations used today however still differ significantly from real-world video datasets both in terms of visual appearance and physics - particularly when it comes to friction and physical collisions in the environment, thus motivating additional empirical study of these topics on real world video data.

B. Existing Video Datasets

Most existing video datasets to date are designed for pure computer vision tasks, such as, action recognition [38], object tracking [39] and video segmentation [40]. Prominent examples include Kinetics400 and Kinetics600 [38], Moments in Time [41], Something-Something V2 (SSv2) [42], AViD [43], Epic Kitchen [44] and Charades [45]. However, these datasets are not specifically tailored towards robotic manipulation settings.

Recently, interest in video datasets for robotic manipulation has increased significantly, with initial datasets at larger scales such as RT-1 [10], HOPE-Video [46] and BC-Z [47] being released. Nonetheless, owing to the traditionally high-cost associated with collecting large-scale real-world datasets, video-based robotic manipulation datasets today still largely remain either limited in size or only available to corporate actors. Moreover, a considerable number of datasets are gathered in simulated environments [48], [49], [50], [51], [52], thereby suffering from the sim-to-real gap [53].

The dataset *CloudGripper-Push-1K* presented in this work therefore fills an important gap by providing large real-world video data of planar pushing interactions comprising 460,000 trajectories across a number of settings and serves as a foundational dataset that allows us to control for changes in a single variable to conduct experimental evaluations.

C. Video Transformers

The transformer architecture, originally designed for natural language processing, has shown a strong capacity when being adapted to vision tasks [54], [55]. Video transformers (VTs) extend the structure of vision transformers [56] by incorporating the temporal dynamics inherent to video data and have achieved impressive results on video classification [14], [12], [29], robotic manipulation [10], [27], [57] and object tracking tasks [58], [59], [60].

To address the challenges of computational intensity and to capture the rich information introduced by the temporal dimension, various factorization methods have been adopted in recent VTs. Notably, Swin Video Transformer [12] adapts the Swin Transformer [5] to include inductive bias, hierarchical structures, and translational invariance, thereby




Fig. 2: An illustration of *CloudGripper-Push-1K* with selected trajectories. (a) Target object *Ball*. (b) Target object *Foam* and one background object. (c) Target object *Cube* and two background objects. (d) Target object *Icosahedron* and four background objects.

TABLE I: Dataset Overview

Object	Background	Training	Testing
Ball (Red)	Single	30030	2800
	Double	22860	1881
	Triple	20720	2570
	Quintuple	20000	1696
Cube	Single	21860	1889
	Double	32870	1500
	Triple	27800	2311
	Quintuple	21610	2180
Foam	Single	19450	2163
	Double	20440	2134
	Triple	20838	2100
	Quintuple	21430	2129
Icosahedron	Single	21670	2697
	Double	24980	2379
	Triple	21430	3150
	Quintuple	22070	1661
	Quintuple-Static	25360	3392
Ball (Green)	Single	21540	1660

This table outlines the sub-tasks and number of videos in each of the 18 training and testing sub-datasets. The background column indicates the total number of objects present in the given setting.

integrating the advantages of Convolutional Neural Networks into transformers. ViViT [14] presents several strategies for factorizing both spatial and temporal dimensions to manage the long token sequences in videos. VTN [29] utilizes a 2D-spatial backbone for feature extraction followed by a temporal encoder [61] to realize flexible video processing. In our work, the proposed VOT framework adopts the generic structure proposed by VTN.

III. DATASET

The *CloudGripper-Push-1K* dataset is collected using the open cloud-robotics platform, *CloudGripper* [26] developed at KTH. The dataset provides a collection of 1.4 TB of data with 1278 hours and 460,000 trajectories of robot pushing RGB videos varying the target objects from ball, cube,

icosahedron (a convex polyhedron with 20 faces) to a foam ball, and with background scene settings with a varying number of rigid balls or cubes, as shown in Fig.2. In the dataset, each trajectory $Traj = (V_t, V_b, P_o)$ consists of:

- 1) A high-resolution occluded top-camera video V_t of dimension $T \times H_t \times W_t \times C$.
- 2) A corresponding occlusion-free bottom-camera video V_b of dimension $T \times H_b \times W_b \times C$.
- 3) A ground-truth trajectory for the target object in pixel space, $P_o \in \{(i, x, y) \mid i, x, y \in \mathbb{Z}, 0 \leq i \leq T, 0 \leq x \leq H_b, 0 \leq y \leq W_b\}$.

where H_t, W_t, H_b, W_b are equal to 1280, 720, 480 and 640 respectively, indicating the heights and widths of V_t and V_b frames; $T = 50$ is the number of frames collected for each video and $C = 3$, representing the number of video color channels. The dataset is collected by moving the robot gripper, with a moving range of $15 \text{ cm} \times 15 \text{ cm}$, and executing random planer pushing of target objects in the workspace environments. During the pushing process, two RGB videos are recorded with our dual-camera system at a frame rate of 5 fps and saved when the total number of frames reaches 50 in a video. The bottom camera view is observed through a transparent plexiglass plate from below the planar pushing surface. The ground-truth trajectories P_o are later derived from V_b using a RGB color-based tracking algorithm recording the pushed object's motions.

We design 18 sub-datasets with varying selections of target objects with different physics attributes and various scene backgrounds. An overview of target and background objects can be seen in Fig.1(b). Four target objects (*Ball (red)*, *Cube*, *Foam*, *Icosahedron*) being pushed in workspace backgrounds with zero (*Single*), one (*Double*), two (*Triple*) and four background objects (*Quintuple*) are captured, resulting in 16 sub-datasets. The red ball (made from plastic) and foam object (also a ball) share identical shape and color but vary in friction coefficient. The cube and icosahedron are 3D printed from identical plastic material and only vary in shape. These two groups of sub-datasets are designed to investigate

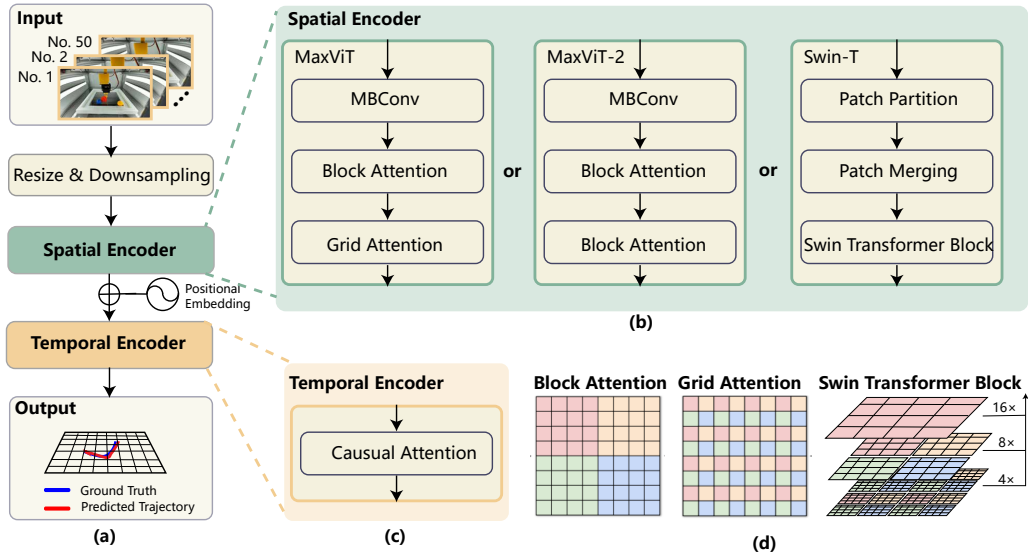


Fig. 3: (a) VOT Model Structure: Following the generic approach detailed in VTN [29], our framework employs a modular design, starting with a 2D-spatial and followed by a temporal encoder. (b) The three types of 2D-spatial encoders that were adopted in the design. (c) Temporal encoder, (d) The three types of attention mechanisms used in the 2D-spatial encoders.

the impact of friction coefficient and shape. Additionally, another two sub-datasets, pushing a single green *Ball* without any other objects in the scene and pushing an icosahedron with 4 cubes in the background (*Quintuple-Static*), have been created to examine the impact of target object color (compared with *Single Ball*) and background dynamics (compared with *Icosahedron Quintuple*, since the background objects are balls whose rolling movements after the collision are more dynamic than cubes). Each sub-dataset consists of approximately 20,000 training and 2,000 test trajectories. Table I outlines the sub-task settings and number of videos in each training and testing sub-dataset.

Finally, we would like to emphasize that the presence of frequent occlusions both due to occlusion behind the gripper and due to occlusions by other objects which presents an interesting challenge to training a model based on top camera videos in *CloudGripper-Push-1K*, examples of such occlusions are visible in Fig. 1(c).

IV. MODEL

The overall design of *VOT* is illustrated in Fig. 3. Our framework is modular in construction which aligns with the generic approach outlined in VTN [29], incorporating a flexible 2D-spatial encoder for frame feature extraction, followed by a temporal encoder to enhance inference from previous frames. In *VOT*, we adopt three vision transformers as the 2D-spatial encoders, enabling a broader investigation into the generalisability of video transformers.

Input. The input to *VOT* is a top camera video V_t with 50 frames. We preprocess the inputs by downsampling them to 25 frames and resizing the frames to 224×224 pixels.

Spatial Encoder. In *VOT*, three vision transformer architectures are adopted as the 2D-spatial encoder. MaxViT [62], enhances the original ViT model [4] with reduced computational complexity by incorporating Block Attention and Grid Attention (*VOT-MaxViT*). Block Attention is designed

to capture local features, while Grid Attention targets global information.

Questioning the necessity of global attention and whether a stronger focus on local features could enhance performance in the robotic setting (where more focus on the interaction between gripper and object is desired), we design *VOT-MaxViT-2*, a variant of *VOT-MaxViT*, by changing the second attention in the original MaxViT from Grid Attention to Block Attention. Additionally, we also adopt Swin Transformer [5], a well-established architecture with hierarchical structure and shifted window for self-attention, as another variant of the spatial encoder (*VOT-Swin-T*). By incorporating three self-attention mechanisms into the 2D-spatial encoder backbone, we aim to investigate the impact of physics and background attributes on these model design choices in VTs.

Temporal Encoder. The temporal encoder follows the initial self-attention mechanism proposed by Vasawani et al. [1] to learn the temporal relationships across frames, with causal attention masking out future information.

Output. The output of *VOT*, with a dimension of 12×16 , represents the predicted target object trajectory traversed within the 25 frames in a down-sampled pixel coordinate of the bottom camera. We adopt the down-sampling considering the trade-off between computational limitation and necessary prediction resolution (small object moving range of 12×12 cm).

V. EXPERIMENTS

A. Implementation Details

Training. We train and evaluate 3 variants of *VOT* (*VOT-MaxViT*, *VOT-MaxViT-2*, and *VOT-Swin-T*) on each of the 18 *CloudGripper-Push-1K* sub-datasets from scratch, where the sizes of training and testing sets are 20,000 and 2,000 respectively. The models are diversified by object type (*Ball*

TABLE II: Scratch and Zero-shot Performance of VOT on *CloudGripper-Push-1K* Dataset

	Ball				Cube				Foam				Icosahedron				
	Sig.	Dou.	Tri.	Quin.	Sig.	Dou.	Tri.	Quin.	Sig.	Dou.	Tri.	Quin.	Sig.	Dou.	Tri.	Quin.	Quin_S
VOT-MaxViT	413/-	512/134	481/132	218/858	15/-	65/191	33/730	28/639	62/-	164/465	129/443	55/232	33/-	42/225	25/368	23/591	30/104
VOT-MaxViT-2	421/-	542/18	455/1042	160/1404	20/-	62/659	36/601	28/538	63/-	174/345	127/303	56/180	34/-	46/127	24/224	24/272	27/367
VOT-Swin-T	508/-	592/414	685/887	300/1214	23/-	80/183	40/806	36/785	97/-	209/467	151/455	73/368	55/-	58/310	30/458	28/465	33/521

The results are shown in the format of *PE/GP* indicating Prediction Error and Generation Gap. *GP*, as defined in Equation 2, measures the zero-shot performance of models trained with the *Single* dataset, for each target object, on those with a higher number of objects. *GPs* are calculated by deducting *PE* of the model trained from scratch from *PE* of models trained on *Single*. *Quin_S* stands for *Quintuple-Static*.

TABLE III: Study of Physics Attributes and Background Characteristics

	Friction		Shape		Color		Background	
	Foam	Ball	Cube	Icosahedron	Ball_R	Ball_G	Dynamic	Static
VOT-MaxViT	1033	203	654	612	1098	1509	186	13
VOT-MaxViT-2	1308	783	769	780	1372	1538	38	21
VOT-Swin-T	554	1103	869	973	1485	1590	164	73
Average	965	696	764	788	1318	1546	129	36

The table presents the values of zero-shot *GP*. *GP* reflects the zero-shot generalisability of the other model in the same controlled group. For example, the *GP* for *Foam* in group *Friction* equals to *PE* of model trained on the *Foam* dataset and tested on *Ball* minus *PE* of model trained on *Ball* and tested on *Ball*. *Ball_R* and *Ball_G* stand for red ball and green ball respectively.

(Red), *Cube*, *Foam*, *Icosahedron*) and background complexity (*Single*, *Double*, *Triple*, *Quintuple*) in training data. In total, 48 models are formed by combining these factors, along with 6 additional models, trained on *Ball_Green* and *Quintuple-Static* to explore other attributes (color, background dynamics). Training is conducted with 8 NVIDIA A100 GPUs on the Berzelius supercomputer [63]. Each model is trained for 100 epochs with a batch size equals to 4, requiring approximately 12 hours of training. We utilize the Adam optimizer [64] with constant learning rate of 1×10^{-5} for *VOT-Swin-T* and 1×10^{-4} for *VOT-MaxViT* and *VOT-MaxViT-2*. All models utilize random initialization and the Mean Squared Error (MSE) is adopted as loss function. The *VOT-MaxViT*, *VOT-MaxViT-2*, *VOT-Swin-T* comprise 5.2×10^7 , 5.2×10^7 , 4.9×10^7 parameters respectively.

Testing. Each of the 18 sub-datasets includes a separate testing set consisting of around 2,000 videos. All zero-shot experiments are conducted on the testing set, where the same robot is used as during the collection of the original training data. This ensures that performance metrics are not influenced by external disturbances and minor differences between individual *CloudGripper* robot instances as these are unrelated to the attributes under study.

Evaluation Metrics. Two metrics are employed to assess the performance and generalisability of *VOT*s. The Prediction Error (*PE*) quantifies the discrepancy between the predicted and ground-truth trajectories. The Generation Gap (*GP*) evaluates the model’s generalisability when applied to datasets with different attributes by taking the difference between the model’s current PE_{new} and the sub-dataset’s $PE_{previous}$

of training from scratch.

$$Prediction\ Error\ (PE) = \frac{1}{M \times N} \sum_{i=1}^M \sum_{j=1}^N (R_{ij} - P_{ij})^2 \quad (1)$$

$$Generation\ Gap\ (GP) = PE_{new} - PE_{previous}, \quad (2)$$

where $M = 12$ and $N = 16$ are the height and width of the network outputs; P represents the predicted trajectory; R denotes the downsampled ground-truth trajectory.

B. Experiments with Varying Scene Background Characteristics

Complexity. As a proxy for the complexity of real-world scene backgrounds, we vary the number of objects placed in the workspace. In this context, “background object” refers to items presented in the scene that are not the target object being pushed by the robot gripper. We experiment with 16 sub-datasets, analyzing 4 levels of background complexity in combination with 4 distinct target objects being pushed: *Ball*, *Cube*, *Foam* and *Icosahedron*. *VOT*s are trained from scratch and the zero-shot generalisability of models trained on a dataset without background objects are then tested on all other environments. The *PE*s in Table.II are higher for the *Double* and *Triple* settings compared with a *Single* background object for most cases (21 out of 24), which shows that complex backgrounds cause difficulties for training of *VT*s and the model performance drops even when performing the same pushing tasks. However, all of the 3 variants behave well in *Quintuple* data. We hypothesize this is a consequence of the constrained moving area for the target object with the introduction of an increasing number of background objects. The non-zero *GPs* in Table.II show that i) *VT*’s performance is sensitive to the presence of

background objects; ii) in *Ball* and *Icosahedron*, the *GPs* increase with background complexity, showing a decline in generalisability with complex backgrounds. In the case of *Cube* and *Foam*, this trend is relatively weaker - possibly due to the reduced amount of dynamic interaction in the dataset (as the *Cube* and *Foam* objects do not roll and come to rest quickly).

Dynamics. In an additional experiment, *Quintuple_Static*, we replace balls in the background with cubes of the same size and color in *Icosahedron*'s *Quintuple* setting. Since the cubes slide rather than roll and move less dramatically compared with balls, this experiment serves to assess the impact of environmental dynamics on the VT's performance. The last two columns of *Table.III* confirm the intuition that the dynamic environments are harder to predict and VT models trained on these datasets tend to have a better zero-shot generalisability on settings with more static backgrounds.

Finally, in *Table.II*, it can be observed that *VOT-MaxViT* generally outperforms *VOT-MaxViT-2* (12 out of 17 cases). This observation underscores the importance of retrieving more global information through Grid Attention, even in robotic manipulation settings where the interaction between grippers and objects is the primary focus.

C. Experiments with Varying Physics Attributes

Here we present a zero-shot generalisability study of physics attributes and scene background characteristics as summarized in *Table.III*.

Color. We compare the zero-shot performance of models trained on data with a single green *Ball* or red *Ball*. The *GP* values for this group are the highest among the four controlled experiments, illustrating that VTs are highly sensitive to object color. This can be explained by the fact that shifts in object color dramatically affect pixel value of the input videos thus causing a significant perturbation in the inputs of the VOT models. Pure appearance factors as simple as color which should not affect trajectory prediction in the robotic manipulation context hence can still cause significant adverse effects on VT's performance.

Friction. As mentioned before, *Foam* and *Ball* share near identical shape and color but *Ball* features a lower friction coefficient compared with *Foam*. Therefore, we compare the performance of this group to understand the impact of friction. When the force on the gripper is large enough to overcome the static friction, the observed effect of the differing friction coefficients is that objects with lower rolling friction coefficient, such as *Ball*, start rolling on the plane and are bounced back from the workspace boundary, resulting in longer and more unpredictable trajectories. We observe asymmetries on the *GPs* of *VOT-MaxViT* and *VOT-MaxViT-2*, which reveal that it may be easier for models trained on longer trajectories to generalize. However, this conclusion does not hold true for *VOT-Swin-T*. This could potentially result from the different architectures between MaxViT [62] and Swin Transformer [5]. The Swin Transformer in particular has a more local focus with strong inductive bias, complicating the prediction of a fast-moving object (*Ball*).




Fig. 4: *PE* vs fine-tuning dataset size graphs. Models are pre-trained for 90 epochs on *Single* dataset and then fine-tuned on *Quintuple* data for 10 epochs and evaluated on *Quintuple*. The x-axis represents the number of the fine-tuning videos. (a) Fine-tuning curves of *VOT-MaxViT* (*VOT-1*) and *VOT-MaxViT-2* (*VOT-2*) with x-axis ranging from 300 to 3000. (b) Fine-tuning curves of *VOT-Swin-T* (*SWIN*) with x-axis ranging from 700 to 7000.

Shape. Here we consider the effect of changing object shapes with a comparative study of *Cube* and *Icosahedron*, which are both 3D printed from the same plastic material. Based on the average *GPs* across three controlled physics attribute groups, it appears that color has a greater impact than shape on performance, while shape and friction show a more comparable level of influence on *GPs*.

D. Fine-Tuning and Qualitative Observations

Since the zero-shot generalisability among the presented settings is not satisfying for many cases, we would like to investigate whether it is possible to reach the performance level of a model that is trained from scratch by means of fine-tuning with a limited set of data.

We therefore fine-tune pre-trained models for 10 epochs with varying sizes of additional data, ranging from 300 to 3000 videos for *VOT-MaxViT* and *VOT-MaxViT-2* and 700 to 7000 videos for *VOT-Swin-T*, as shown in *Fig.4*. While it is commonly accepted that smaller fine-tuning datasets are preferred to leverage pre-trained models effectively, our results indicate that although *PE* generally decreases with larger dataset sizes, when the fine-tuning dataset remains relatively small, the model performance might even decline with the increasing dataset size. Moreover, achieving comparable performance to MaxViT [62] in this scenario demands a substantially larger fine-tuning dataset when using Swin Transformer [5] as 2D-spatial encoder.

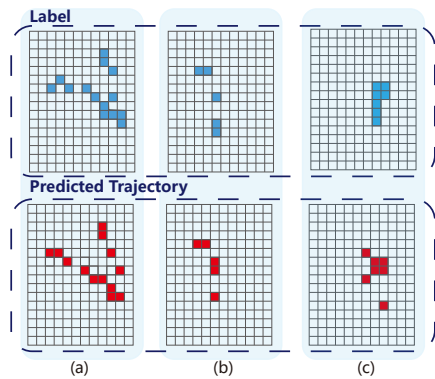


Fig. 5: Examples of predicted and corresponding ground-truth trajectories, with *Ball* being the target object. (a) A complicated long trajectory of an object colliding with the boundary four times. (b) A successful prediction that featured heavy occlusion in the top camera view. (c) A failure mode. Note that the discontinuities in labels and predicted trajectories are a result of the downsampling of input videos. Refer to the supplementary video for more details.

In Fig.5, we present examples of predicted and ground truth trajectories - while VTs are able to predict complex long (a) or occluded (b) trajectories, failure cases also present in the case of (c). In (c), the target object (*Ball*) exhibits an interesting unusual oscillation motion pattern possibly due to scratches on the plexiglass plate, representing a case not seen during training. Videos of examples described above and a data saturation curve of *VOT-MaxViT* can be found in the supplementary video.

E. Limitations and Conclusions

While we have presented empirical results backed by a large real-world dataset, various factors such as the downsampling of frame rate and resolution that was necessary for computational complexity reasons may affect our observations. Furthermore, we focus on typical VT architecture design components in the specific narrow setting of planar pushing in order to be able to report results at large data scales. Our experiments across the data subsets provide empirical evidence that: i) When training from scratch, VTs' performance appears to decline with increasing background scene complexity. Furthermore, models trained on dynamic backgrounds exhibit enhanced generalisability as compared to static backgrounds in zero-shot scenarios. ii) Among the physics attributes considered (color, friction coefficient, shape), the VTs' performance displays the greatest sensitivity to color - which is not a design goal in the context of robotic manipulation. Notably, shape and color demonstrate symmetry in cross zero-shot generalisability tests, while friction exhibits asymmetry, potentially due to the model architecture. iii) Regarding the fine-tuning process, we observed variations in optimal fine-tuning dataset sizes across settings and models. Furthermore, we observe that fine-tuning may even have a negative effect on model performance if the dataset size is not large enough.

REFERENCES

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, 2017.

[2] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," 2019.

[3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Conference of the North American Chapter of the Association for Computational Linguistics*, 2019.

[4] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.

[5] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.

[6] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "Vivit: A video vision transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 6836–6846.

[7] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch, "Decision transformer: Reinforcement learning via sequence modeling," in *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 15 084–15 097.

[8] K.-H. Lee, O. Nachum, M. S. Yang, L. Lee, D. Freeman, S. Guadarrama, I. Fischer, W. Xu, E. Jang, H. Michalewski, and I. Mordatch, "Multi-game decision transformers," in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 27 921–27 936.

[9] Q. Zheng, A. Zhang, and A. Grover, "Online decision transformer," in *Proceedings of the 39th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 162. PMLR, 17–23 Jul 2022, pp. 27 042–27 059.

[10] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, "Rt-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.

[11] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Proceedings of The 6th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, K. Liu, D. Kulic, and J. Ichnowski, Eds., vol. 205. PMLR, 14–18 Dec 2023, pp. 785–799.

[12] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu, "Video swin transformer," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 3202–3211.

[13] M. Patrick, D. Campbell, Y. Asano, I. Misra, F. Metzger, C. Feichtenhofer, A. Vedaldi, and J. a. F. Henriques, "Keeping your eye on the ball: Trajectory attention in video transformers," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 12 493–12 506.

[14] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "Vivit: A video vision transformer," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 6836–6846.

[15] S. Dasari and A. Gupta, "Transformers for one-shot visual imitation," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 2071–2084.

[16] H. Bharadhwaj, A. Gupta, and S. Tulsiani, "Visual affordance prediction for guiding robot exploration," in *2023 IEEE ICRA*, 2023, pp. 3029–3036.

[17] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, "Mimicplay: Long-horizon imitation learning by watching human play," *arXiv:2302.12422*, 2023.

[18] R. Rakhimov, D. Volkhonskiy, A. Artemov, D. Zorin, and E. Burnaev, "Latent video transformer," *arXiv:2006.10704*, 2020.

[19] A. Gupta, S. Tian, Y. Zhang, J. Wu, R. Martín-Martín, and L. Fei-Fei, "Maskvit: Masked visual pre-training for video prediction," in *ICLR*, 2023.

[20] C. Nash, J. Carreira, J. C. Walker, I. Barr, A. Jaegle, M. Malinowski, and P. Battaglia, "Transframer: Arbitrary frame prediction with generative models," *Trans. on Machine Learning Res.*, 2023.

[21] J. Sun, D.-A. Huang, B. Lu, Y.-H. Liu, B. Zhou, and A. Garg, "Plate: Visually-grounded planning with transformers in procedural tasks," *IEEE RA-L*, vol. 7, no. 2, pp. 4924–4930, 2022.

- [22] H. Zhao, I. Hadji, N. Dvornik, K. G. Derpanis, R. P. Wildes, and A. D. Jepson, "P3iv: Probabilistic procedure planning from instructional videos with weak supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2938–2948.
- [23] A.-L. Wang, K.-Y. Lin, J.-R. Du, J. Meng, and W.-S. Zheng, "Event-guided procedure planning from instructional videos with text supervision," *arXiv:2308.08885*, 2023.
- [24] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, "Transformers in vision: A survey," *ACM computing surveys (CSUR)*, vol. 54, no. 10s, pp. 1–41, 2022.
- [25] Z. Tong, Y. Song, J. Wang, and L. Wang, "VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [26] M. Zahid and F. T. Pokorný, "Cloudgripper: An open source cloud robotics testbed for robotic manipulation research, benchmarking and data collection at scale," to appear, *IEEE ICRA 2024*, *arXiv preprint arXiv:2309.12786*, 2024.
- [27] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799.
- [28] —, "Cliport: What and where pathways for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2022, pp. 894–906.
- [29] D. Neimark, O. Bar, M. Zohar, and D. Asselmann, "Video transformer network," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 3163–3172.
- [30] Y. Tay, M. Dehghani, J. Rao, W. Fedus, S. Abnar, H. W. Chung, S. Narang, D. Yogatama, A. Vaswani, and D. Metzler, "Scale efficiently: Insights from pre-training and fine-tuning transformers," *arXiv:2109.10686*, 2021.
- [31] W. Fedus, B. Zoph, and N. Shazeer, "Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity," *JMLR*, vol. 23, no. 1, pp. 5232–5270, 2022.
- [32] Y. Tay, M. Dehghani, S. Abnar, H. W. Chung, W. Fedus, J. Rao, S. Narang, V. Q. Tran, D. Yogatama, and D. Metzler, "Scaling laws vs model architectures: How does inductive bias influence scaling?" *arXiv preprint arXiv:2207.10551*, 2022.
- [33] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.
- [34] Y. Jing, X. Zhu, X. Liu, Q. Sima, T. Yang, Y. Feng, and T. Kong, "Exploring visual pre-training for robot manipulation: Datasets, models and methods," *arXiv:2308.03620*, 2023.
- [35] E. Xing, A. Gupta, S. Powers, and V. Dean, "Kitchenshift: Evaluating zero-shot generalization of imitation-based policy learning under domain shifts," in *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.
- [36] A. Xie, L. Lee, T. Xiao, and C. Finn, "Decomposing the generalization gap in imitation learning for visual robotic manipulation," *arXiv:2307.03659*, 2023.
- [37] K. Cobbe, C. Hesse, J. Hilton, and J. Schulman, "Leveraging procedural generation to benchmark reinforcement learning," in *ICML*. PMLR, 2020, pp. 2048–2056.
- [38] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev *et al.*, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.
- [39] L. Huang, X. Zhao, and K. Huang, "Got-10k: A large high-diversity benchmark for generic object tracking in the wild," *IEEE trans. on pattern analysis and machine int.*, vol. 43, no. 5, pp. 1562–1577, 2019.
- [40] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 davis challenge on video object segmentation," *arXiv preprint arXiv:1704.00675*, 2017.
- [41] M. Monfort, A. Andonian, B. Zhou, K. Ramakrishnan, S. A. Bargal, T. Yan, L. Brown, Q. Fan, D. Gutfreund, C. Vondrick *et al.*, "Moments in time dataset: one million videos for event understanding," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 2, pp. 502–508, 2019.
- [42] R. Goyal, S. Ebrahimi Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag *et al.*, "The" something something" video database for learning and evaluating visual common sense," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5842–5850.
- [43] A. Piergiovanni and M. Ryoo, "Avid dataset: Anonymized videos from diverse countries," *Advances in Neural Information Processing Systems*, vol. 33, pp. 16711–16721, 2020.
- [44] D. Damen, H. Doughty, G. M. Farinella, A. Furnari, J. Ma, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray, "Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100," *International Journal of Computer Vision (IJCV)*, vol. 130, p. 33–55, 2022.
- [45] G. A. Sigurdsson, A. Gupta, C. Schmid, A. Farhadi, and K. Alahari, "Charades-ego: A large-scale dataset of paired third and first person videos," *arXiv preprint arXiv:1804.09626*, 2018.
- [46] S. Tyree, J. Tremblay, T. To, J. Cheng, T. Mosier, J. Smith, and S. Birchfield, "6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark," in *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [47] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, "Bc-z: Zero-shot task generalization with robotic imitation learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 991–1002.
- [48] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, "What matters in learning from offline human demonstrations for robot manipulation," in *5th Annual Conference on Robot Learning*, 2021.
- [49] J. Fu, A. Kumar, O. Nachum, G. Tucker, and S. Levine, "D4rl: Datasets for deep data-driven reinforcement learning," *arXiv preprint arXiv:2004.07219*, 2020.
- [50] J. Wong, A. Tung, A. Kurenkov, A. Mandlekar, L. Fei-Fei, S. Savarese, and R. Martín-Martín, "Error-aware imitation learning from teleoperation data for mobile manipulation," in *Conference on Robot Learning*. PMLR, 2022, pp. 1367–1378.
- [51] S. James, Z. Ma, D. Rovick Arrojo, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," *IEEE RA-L*, 2020.
- [52] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay *et al.*, "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Conference on Robot Learning*. PMLR, 2018, pp. 879–893.
- [53] S. Höfer, K. Bekris, A. Handa, J. C. Gamboa, M. Mozifian, F. Golemo, C. Atkeson, D. Fox, K. Goldberg, J. Leonard *et al.*, "Sim2real in robotics and automation: Applications and challenges," *IEEE trans. on automation science and engineering*, vol. 18, no. 2, pp. 398–400, 2021.
- [54] C.-F. R. Chen, Q. Fan, and R. Panda, "CrossViT: Cross-Attention Multi-Scale Vision Transformer for Image Classification," in *International Conference on Computer Vision (ICCV)*, 2021.
- [55] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao, "Pyramid vision transformer: A versatile backbone for dense prediction without convolutions," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 548–558.
- [56] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [57] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choremanski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," *arXiv:2307.15818*, 2023.
- [58] F. Xie, L. Chu, J. Li, Y. Lu, and C. Ma, "Videotrack: Learning to track objects via video transformer," in *CVPR*, 2023, pp. 22826–22835.
- [59] L. Lin, H. Fan, Z. Zhang, Y. Xu, and H. Ling, "Swintrack: A simple and strong baseline for transformer tracking," *Advances in Neural Information Processing Systems*, vol. 35, pp. 16743–16754, 2022.
- [60] N. Wang, W. Zhou, J. Wang, and H. Li, "Transformer meets tracker: Exploiting temporal context for robust visual tracking," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 1571–1580.
- [61] I. Beltagy, M. E. Peters, and A. Cohan, "Longformer: The long-document transformer," *arXiv:2004.05150*, 2020.
- [62] Z. Tu, H. Talebi, H. Zhang, F. Yang, P. Milanfar, A. Bovik, and Y. Li, "Maxvit: Multi-axis vision transformer," in *European conference on computer vision*. Springer, 2022, pp. 459–479.
- [63] Berzelius. [Online]. Available: <https://www.nsc.liu.se/systems/berzelius/>
- [64] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980*, 2014.