

Observer–Actor: Active Vision Imitation Learning with Sparse-View Gaussian Splatting

Yilong Wang, Cheng Qian, Ruomeng Fan, and Edward Johns

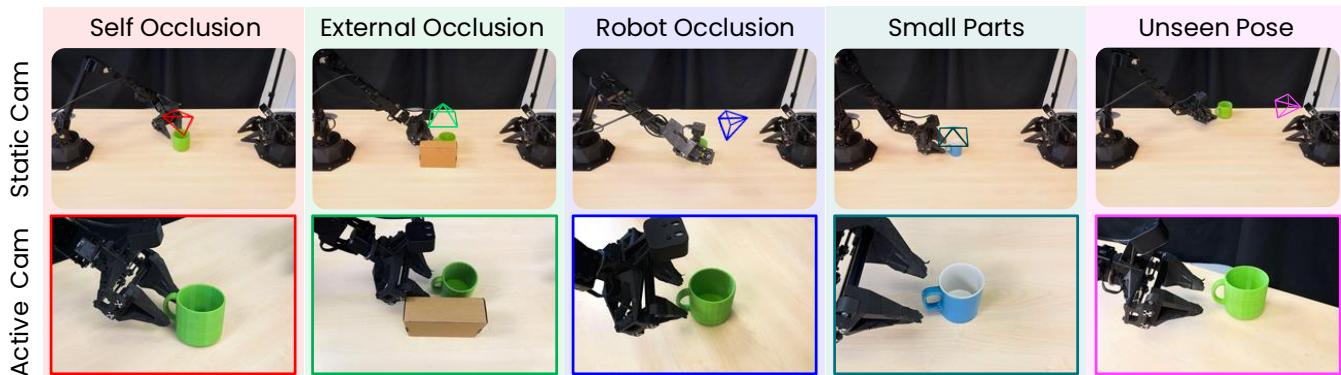


Fig. 1: **Active vision for imitation learning in a mug-handle pickup task across five scenarios.** When a static camera struggles (top row), alternative placements (bottom row, and coloured frustums in top row) provide better observations. In our method, at test time an *observer* robot (robot on right in above examples) computes and moves to such an optimal view from its wrist-cam, after which an *actor* robot (robot on left in above examples) performs the task conditioned on this view.

Abstract—We propose Observer-Actor (ObAct), a novel framework for active vision imitation learning in which the observer moves to optimal visual observations for the actor. We study ObAct on a dual-arm robotic system equipped with wrist-mounted cameras. At test time, ObAct dynamically assigns *observer* and *actor* roles: the observer arm constructs a 3D Gaussian Splatting (3DGS) representation from three images, virtually explores this to find an optimal camera pose, then moves to this pose; the actor arm then executes a policy using the observer’s observations. This formulation enhances the clarity and visibility of both the object and the gripper in the policy’s observations. As a result, we enable the training of ambidextrous policies on observations that remain closer to the occlusion-free training distribution, leading to more robust policies. We study this formulation with two existing imitation learning methods – trajectory transfer and behaviour cloning – and experiments show that ObAct significantly outperforms static-camera setups: trajectory transfer improves by 145% without occlusion and 233% with occlusion, while behavior cloning improves by 75% and 143%, respectively. Videos are available at <https://obact.github.io>.

I. INTRODUCTION

Current imitation learning methods for robotic manipulation primarily rely on static cameras [1, 2], egocentric wrist-mounted cameras [3, 4], or a combination of both [5, 6]. Static cameras are typically placed in a task-agnostic manner, and once a policy is trained, it often must be deployed in the same setup. Wrist-mounted cameras offer greater flexibility but suffer from limited global awareness and a restricted field of view. Combining both viewpoints can partially mitigate these issues but often introduces redundant observations that may distract the model—especially in low-data regimes [7,

8]. These limitations are further amplified by occlusions, highlighting the need for active vision, where the camera can move dynamically to provide the best viewpoint for the task.

Realizing active vision requires both movable camera hardware and sophisticated active vision strategies. Recent studies [7, 8, 9] achieve this by developing a dedicated active-vision arm for perception, with its policy learned through teleoperation and behavior cloning. However, this approach has two significant limitations. First, the active vision arm is restricted to a fixed observer role, limiting its reachable viewpoints and preventing its use as regular manipulator. Second, the active vision strategy demands extensive human demonstrations, which, although collectable via a VR-based teleoperation system, imposes additional burden on operators.

In this paper, we present ObAct (short for Observer–Actor), a novel framework for active vision in imitation learning, where an observer robot computes optimal visual observations that guide the actor robot’s actions. As illustrated in Figure 1 with a mug-handle pickup task, ObAct positions the observer arm at test-time optimal viewpoints, providing visual observations for the actor arm during manipulation. Unlike prior approaches, ObAct dynamically assigns observer and actor roles based on the scene configuration, without requiring a separately trained active vision strategy for each arm. To enable this, we construct a test-time 3D Gaussian Splatting (3DGS) representation [10] from sparse-view images and optimize the camera pose to obtain observations that are similar to the demonstration observations, while minimizing occlusions. We further extend existing imitation learning

methods—trajectory transfer [2] and behavior cloning [5]—to leverage this view-conditioned framework. Experiments show that ObAct substantially improves success rates over static-camera baselines: trajectory transfer increases by 145% without occlusion and 233% with occlusion, while behavior cloning increases by 75% and 143%, respectively.

In summary, our contributions are threefold:

- 1) **ObAct Framework:** We introduce a decoupled observer–actor framework for active vision imitation learning, which allows the system to be robust against visual edge cases that static cameras cannot.
- 2) **Active Vision via Sparse-view 3DGS:** We develop an RGB-based active vision system that uses a test-time 3DGS model, constructed from sparse-view images, to optimize viewpoints for demonstration consistency and occlusion reduction. To our knowledge, this is the first use of sparse-view 3DGS in active vision.
- 3) **Active Vision for Imitation Learning:** We extend trajectory transfer and behavior cloning to the active vision setting, demonstrating substantial performance gains and, for behavior cloning, improved data efficiency.

II. RELATED WORK

Active Vision for Robotic Manipulation. Active vision (AV) [11] refers to actively moving the camera to maximize perceptual quality for downstream tasks, and has been widely applied in robotics for object recognition [12], pose estimation [12, 13], and 3D reconstruction [14, 15]. It has also been extensively studied in robotic grasping, where unobserved views are sampled and evaluated using depth images based on information gain [16], grasp affordances [17], graspsness inconsistency [18], and velocity fields [19], yielding improved performance in occluded and cluttered environments. In contrast, our method renders and evaluates unobserved viewpoints directly in the RGB domain, the prevalent modality in modern imitation learning systems [1, 5], and extends beyond grasping to support a broader range of manipulation tasks. Our setup is most closely related to [20], where an agent first observes the environment and then interacts with it by jointly training a next-best-view camera policy and a next-best-pose gripper policy using few-shot reinforcement learning. However, their approach assumes a freely moving camera and is validated only in simulation, whereas we demonstrate a real-world dual-arm system that achieves active vision without requiring a separate policy.

View-conditioned Imitation Learning. For manipulation systems to generalize, visuomotor policies must remain robust to discrepancies in observational viewpoints—particularly unseen camera poses—which is especially critical in active vision settings where camera positions are dynamic rather than static. Recent works address this by training end-to-end behavior cloning policies that jointly control both the active camera and the manipulation arms using human demonstrations [7, 8, 9]. While effective, these approaches increase data requirements and model complexity. In contrast, our approach trains policies over a compact distribution of viewpoints and explicitly optimizes the test-time viewpoint

relative to demonstrations. Another class of methods builds explicit 3D scene representations, such as point clouds or voxels [21, 22], or renders canonical 2D views from these structures [23]. These approaches improve robustness to viewpoint variation but require precise extrinsic calibration and incur high computational costs. Data augmentation offers a complementary direction: [24] generate multi-view datasets from single-view demonstrations using a pre-trained zero-shot novel-view synthesis model [25], while [26] edit reconstructed 3D scenes built from dense scans to create synthetic training views. Beyond end-to-end training, some works transfer a single demonstration trajectory at test time by estimating relative object pose changes between demonstration and execution views, either via pose estimation [2] or visual servoing [6]. However, such methods remain highly sensitive to occlusions. In this paper, we extend behavior cloning and trajectory transfer within our observer–actor framework, demonstrating consistent performance gains in both non-occluded and occluded scenarios.

III. METHOD

Figure 2 illustrates the proposed system framework. To collect a demonstration, the operator selects a demonstration optimal viewpoint (Section III-A) for the task object, positions the observer accordingly, and records a trajectory using the actor arm, creating a dataset for downstream imitation learning. At test time, given a novel object configuration with potential occlusions, each arm captures three predefined scene views, yielding six views in total. The system dynamically assigns observer and actor roles based on these views and constructs a 3DGS representation from the observer’s three views (Section III-B). Through view optimization (Section III-C), it identifies the test-time optimal viewpoint within this representation. The observer arm then moves to this viewpoint, allowing the actor arm to execute the manipulation task conditioned on it (Section III-D).

A. Definition of Optimal View

Demonstration Optimal View. We define a demonstration optimal view as a camera viewpoint relative to the target object that maximizes the visibility of task-relevant features while minimizing occlusion. The actor’s interaction with the object must remain clearly observable from this view, without obstruction from the actor arm during manipulation. In practice, demonstration optimal views avoid self-occlusion, robot-induced occlusion while preserving fine object details, as illustrated in Figure 1. These views are determined by the operator using their own subjective estimation of where an optimal view is. We denote the demonstration optimal view as $v_{\text{demo}}^* \in \text{SE}(3)$, where v is the camera pose expressed in the object frame. A dataset then contains a set of such demonstration optimal views: $\mathcal{V}_{\text{demo}}^* = \{v_{\text{demo},1}^*, v_{\text{demo},2}^*, \dots, v_{\text{demo},N}^*\}$.

Test-time Optimal View. At test time, the system selects a viewpoint from the kinematically feasible set \mathcal{V} that is *closest* to the demonstration optimal view while minimizing




Fig. 2: **Framework Overview.** (1) **Train:** The operator selects a *demonstration optimal view*, moves the observer arm to this view, and records a demonstration. This process is repeated as required by the imitation learning method. (2) **Test:** The robots explore six views of the scene to construct a 3DGS representation. View optimization within this representation identifies the *test-time optimal view*. The observer arm then moves to this view, after which the actor arm executes the task.

occlusion. Formally, this is expressed as:

$$v_{\text{test}}^* = \arg \min_{v \in \mathcal{V}} (d_{\text{pose}}(v, v_{\text{demo}}^*) + \lambda \mathcal{O}(\text{obj}, v)) \quad (1)$$

Here, λ balances the trade-off between proximity to the demonstration optimal view and occlusion, $\mathcal{O}(\cdot)$ quantifies occlusion, $d_{\text{pose}}(\cdot, \cdot)$ measures the SE(3) distance between camera poses, and obj denotes the target object. Since v is expressed in the object frame, the first term encourages test-time viewpoints that are visually consistent with the training views, ensuring in-distribution observations for robust inference. Real-world examples of such optimal viewpoints are shown in Figure 3.

B. Sparse-view Gaussian Splatting

Directly computing v_{test}^* using Equation 1 requires depth information and an occlusion function. We show that an image-level surrogate loss provides an effective approximation. This approach entails simulating images from unseen viewpoints, which in turn requires test-time 3D reconstruction. Rather than performing time-consuming full scans, we capture sparse exploratory views of the scene, assign observer and actor roles, and construct a 3DGS representation.

Exploratory Views. We begin by capturing six scene viewpoints, evenly spaced at 60° intervals to cover the full 360° workspace. These exploratory poses are predefined and fixed. In each iteration, both arms simultaneously move to two of the poses, producing six images along with their corresponding camera poses expressed in the robot frame.

Role Assignment. Roles are assigned to each arm based on how closely their captured views match v_{demo}^* . Instead of performing explicit relative pose estimation, we use the robust dense feature matcher RoMa [27], using the number of confident matches on the segmented object as a proxy. The arm with more aggregated matches is designated as the

observer, indicating it is closer to the demonstration view, while the other arm becomes the actor.

Three-View GS. We reconstruct the scene using the three images captured by the observer, via InstantSplat [28], a sparse-view method that leverages geometric priors. It employs Mast3R [29], a large-scale 3D geometric model, to estimate camera poses from RGB images and perform joint optimization. These optimized poses are more accurate than those from our low-cost robot arms, resulting in higher-quality reconstructions. We use only three images to train the 3DGS model, as the additional views provide little extra information for rendering v_{test}^* , and training is considerably faster. We demonstrate that three views offer the best trade-off between performance and computation time (Section IV-C).

Frame Alignment. The 3DGS reconstruction is generated in an arbitrary coordinate system and must be aligned with the robots' frame. To achieve this, we use the Umeyama algorithm. By combining camera poses from the robot's encoders and hand-eye calibration with the estimated poses from InstantSplat, we solve for a similarity transform via singular value decomposition (SVD) to align the two sets of poses. This process also recovers the scale of the reconstruction.

C. View Optimization

To compute the test-time optimal view v_{test}^* , we first generate candidate viewpoints by global sampling within the 3DGS representation. We then select the best candidate and refine it using local gradient-based optimization. Finally, the resulting view is aligned to the real world is then aligned to the real world for the observer arm to reach.

Candidate View Sampling. To initialize the optimization, we sample a hemisphere of candidate viewpoints around the task object's center, following [16]. The object center is estimated from its 3D bounding box, obtained by lifting 2D masks

from GroundedSAM [30] into 3D using aligned depth maps. After filtering out kinematically infeasible poses, we obtain a candidate set $\mathcal{V}_{\text{candidate}} \subseteq \mathcal{V}$. Owing to the efficiency of 3DGS, all candidate views can be rendered at ~ 250 FPS.

Optimal View Initialization. After rendering a set of virtual RGB images, we select the candidate view with the highest score as the initialization. The scoring function, identical to that used for role assignment, is based on the number of confidence-weighted feature matches between images from $\mathcal{V}_{\text{candidate}}$ and v_{demo}^* . This criterion implicitly accounts for occlusion, as occluded regions naturally produce fewer and less reliable matches. In practice, we found that this matching-based strategy outperforms an alternative based on DINOv2 [31] cosine distance for measuring view similarity. Even when combined with object segmentation, DINOv2 struggles to discriminate between viewpoints—an observation consistent with [32]. To reduce computational cost, we employ Tiny RoMa [27], a lightweight variant of RoMa with faster runtime.

Differentiable Rendering. After obtaining an initialization, we refine the viewpoint v using differentiable rendering while explicitly accounting for gripper-induced occlusions. This is particularly important in our setting, where the observer arm’s gripper consistently remains in the field of view and may block the object, as shown in Figure 3. To handle this, we use SAM2 [33] for segmentation of the rendered image, guided by the demonstration image’s mask as a prompt. This is effective because the initialized and demonstration viewpoints are already closely aligned. To mitigate errors in estimating the object center, we first re-center the object’s mask in the image plane. We then refine v by optimizing a loss that aligns the image rendered from the 3DGS model (denoted \mathcal{G}) with the segmented demonstration image I_{demo}^* , while penalizing occlusions. Let $I(v; \mathcal{G})$ denotes the rendered image. The loss \mathcal{L} as a particular view v is defined as:

$$\begin{aligned} \mathcal{L}(v) = & -\lambda_1 \mathcal{L}_{\text{sim}}(\phi(I(v; \mathcal{G}) \odot \mathcal{M}_{\text{obj}}(v)), \phi(I_{\text{demo}}^*)) \\ & + \lambda_2 \mathcal{L}_{\text{IOU}}(\mathcal{M}_{\text{obj}}(v), \mathcal{M}_{\text{grip}}), \end{aligned} \quad (2)$$

where $\phi(\cdot)$ denotes the DINOv2 feature extractor, $\mathcal{M}_{\text{obj}}(v)$ is the soft object mask at viewpoint v obtained from SAM2, and $\mathcal{M}_{\text{grip}}$ is the fixed soft gripper mask. The first term enforces feature alignment, while the second penalizes overlap between the object and gripper masks. Unlike photometric losses employed in prior work [34], our loss explicitly accounts for occlusions. We note although DINOv2 is unreliable for global initialization, it is effective for local refinement.

Moving to the Optimal View. Once the optimal camera pose v_{test}^* is computed, we transform it back into the real-world coordinate frame and use a motion planner to move the observer arm to the target pose. Figure 3 illustrates the demonstration view, the optimal view obtained from the 3DGS representation, and the final aligned view executed in the real world. Minor discrepancies in viewpoint arise from robot kinematic errors, hand–eye calibration, and frame alignment, but these do not noticeably affect downstream imitation learning performance.

D. View-Conditioned Imitation Learning

In our framework, we extend two categories of imitation methods for task execution once the camera has been moved to the optimal view: trajectory transfer and behavior cloning. **AV Trajectory Transfer.** Trajectory transfer (TT) methods estimate the relative pose change of the object between the demonstration and the test instance, and then transfer the demonstration trajectory in a one-shot manner within the SE(3) manifold [2]. Consequently, only a single demonstration optimal view is required. In our setting, demonstrations are recorded in the actor arm’s coordinate frame, while observations are captured from the observer arm’s camera. To resolve this mismatch, we extend trajectory transfer by applying the following sequence of frame transformations:

$${}^O\mathbf{T}_{E_O}(\text{test}) = {}^O\mathbf{T}_{E_O} {}^{E_O}\mathbf{T}_C \Delta^C\mathbf{T}_{\text{obj}} {}^C\mathbf{T}_{E_O} \quad (3)$$

$$\Delta^O\mathbf{T}_{\text{obj}} = {}^O\mathbf{T}_{E_O}(\text{test}) {}^{E_O}\mathbf{T}_O(\text{demo}) \quad (4)$$

$$\Delta^A\mathbf{T}_{\text{obj}} = {}^O\mathbf{T}_A \Delta^O\mathbf{T}_{\text{obj}} {}^A\mathbf{T}_O \quad (5)$$

$${}^A\mathbf{T}_{E_A}(\text{test}) = \Delta^A\mathbf{T}_{\text{obj}} {}^A\mathbf{T}_{E_A}(\text{demo}) \quad (6)$$

Here, ${}^O\mathbf{T}_{E_O}(\text{test})$ is the observer end-effector pose at v_{test}^* , and ${}^O\mathbf{T}_{E_O}$ is its pose at v_{train}^* , respectively, both expressed in the observer base frame O . The hand–eye calibration from the observer end-effector to the camera frame is ${}^{E_O}\mathbf{T}_C$, with inverse ${}^C\mathbf{T}_{E_O}$. The object’s relative pose change from demonstration to test is denoted by $\Delta^C\mathbf{T}_{\text{obj}}$ in the camera frame, $\Delta^O\mathbf{T}_{\text{obj}}$ in the observer base frame, and $\Delta^A\mathbf{T}_{\text{obj}}$ in the actor base frame A . The fixed transform between the observer and actor base frames is ${}^O\mathbf{T}_A$, with inverse ${}^A\mathbf{T}_O$. Finally, ${}^A\mathbf{T}_{E_A}(\text{demo})$ and ${}^A\mathbf{T}_{E_A}(\text{test})$ denote the actor’s end-effector poses for task execution during demonstration and test, both expressed in the actor base frame A .

We estimate $\Delta^C\mathbf{T}_{\text{obj}}$ using RoMa by lifting matched feature points to 3D via aligned depth maps. The relative SE(3) transformation is then computed using Procrustes alignment within a RANSAC loop. The transferred trajectory is executed in open-loop, and its accuracy further depends on the hand–eye calibration ${}^{E_O}\mathbf{T}_C$ and the fixed transform between the observer and actor base frames ${}^O\mathbf{T}_A$.

AV Behavior Cloning. Behavior cloning (BC) trains a closed-loop policy π to map observations to actions. Unlike TT, which exploits pose estimation models, BC requires several demonstration optimal views to generalize across diverse camera perspectives covering v_{test}^* . Since optimizing over all viewpoints in $\mathcal{V}_{\text{train}}^*$ is computationally expensive, we use Mast3R to estimate camera poses in the object frame and select the viewpoint closest to the center of all estimated poses as v_{demo}^* , which we found to work well empirically. Moreover, unlike prior work [9, 7], which represents the end-effector pose in a static world frame and provides the camera pose explicitly as policy input, we express the actor arm’s end-effector pose directly in the camera frame:

$${}^C\mathbf{T}_{E_A} = {}^C\mathbf{T}_{E_O} {}^{E_O}\mathbf{T}_O {}^O\mathbf{T}_A {}^A\mathbf{T}_{E_A}. \quad (7)$$

This representation simplifies the state space and, as we show in Section IV-B, improves both data efficiency and policy

performance. At each time step t , the policy receives an RGB observation I_t from the test-time optimal view v_{test}^* , along with the proprioceptive state S_t . The proprioceptive state includes the actor end-effector pose expressed in the camera frame and the gripper state. The policy outputs an action sequence $a_{t:t+n_p-1} = \pi(I_t, S_t)$ of length n_p , which is then transformed back to the actor’s base frame for execution.

Ambidextrous Inference. By representing actions in the camera frame for both of our imitation learning methods—AV TT and AV BC—we enable ambidextrous inference: when the roles of observer and actor differ from those in the demonstration, the transferred trajectory and trained policy can still be executed without any additional data or training.

IV. EVALUATION

We instantiated our method on a real-world dual-arm ALOHA setup [5], mounted on a table and equipped with two calibrated RealSense D405 cameras. To evaluate our approach, we selected five diverse manipulation tasks involving self-occluding objects. In *Pick Mug*, the robot grasps a mug by its handle. In *Hammer Nail*, it uses a pre-grasped hammer to drive a nail. *Open Drawer* requires opening the second drawer of a multi-layer unit. In *Retrieve Pack*, the robot retrieves a package from a deep box. Finally, *Insert Coin* involves placing a coin into a storage container. All tasks require observing specific object parts or features for successful manipulation, with the key regions highlighted in red in Figure 3.

A. Comparison with Static-Camera

Baselines and Implementation. We evaluate our method for both TT and BC approaches, comparing against static-camera setups. For the TT baseline, we select the best image from the initial 3-view exploration and perform pose estimation relative to the demonstration, as described in Section III-D. In contrast, our method leverages the test-time optimal view for pose estimation. For the BC baseline, we use a static camera chosen from one of the poses in the 3-view exploration.

To ensure fairness, we collect separate datasets for BC with and without active vision, rather than reusing the same dataset as was done in [7]. Reusing would introduce visual inconsistencies, as the static camera would capture the observer arm moving across different poses, providing irrelevant information that could confuse the learned policy. In total, we collected 70 demonstrations for each of the five tasks, with and without active vision, resulting in $70 \times 5 \times 2 = 700$ demonstrations. All demonstrations were recorded using an iPhone-based teleoperation system built with ARKit.

We implement BC with ACT [5], using DINOv2 as the vision backbone and absolute Cartesian action representation. For BC with AV, we remove any frames in the dataset where the actor arm’s gripper is invisible in the camera view, as these introduce ambiguities. This also ensures that the trained policy outputs a pose where the actor’s gripper is likely to be visible in the camera view at the first timestep.

Experiments. For all tasks except *Retrieve Pack*, which inherently involves occlusion due to its setup, we evaluate both unoccluded and occluded scenarios. We perform 10 rollouts

on novel scene configurations with variations in object poses and the presence of distractors. For each variation, we first run the baseline and then our method under identical conditions.

For the static-camera BC baseline, both training and evaluation are restricted to object poses where task-relevant parts remain visible, as self-occlusion leads to prohibitively low success rates. Consequently, pose variations for the static-camera BC baseline are limited to within 45° of the demonstration poses during testing. In contrast, our method naturally handles self-occlusion arising from larger rotations. **Results.** Table I summarizes the experiment results. Using the test-time optimal view for inference, both TT and BC consistently outperform the static-camera setup across all tasks. Specifically, TT achieves performance improvements of 146% in the unoccluded setting and 233% in the occluded setting, while BC improves by 75% and 143%, respectively.

For TT, this improvement arises because the test-time optimal image has greater feature overlap with the demonstration optimal image and experiences less occlusion, leading to more accurate pose estimation. For BC, the test-time observations are more in-distribution relative to the training data and encounter fewer occlusions, resulting in more reliable policy execution. The performance drop under occlusion may stem from selected viewpoints being slightly out-of-distribution in order to compensate for occluded regions.

Data Efficiency of BC with AV. Figure 4 shows the success rates versus the number of demonstrations, using 30, 50, and 70 demos for the *Pick Mug*, *Retrieve Pack*, and *Open Drawer* tasks. We observe that, given the same number of demonstrations, BC with AV consistently outperforms the static-camera setup, highlighting the advantages of inference from test-time optimal views. Notably, for the *Retrieve Pack* task, BC with a static camera fails completely. This is because the task suffers from severe occlusion, particularly from the actor arm’s gripper when reaching inside the box. Some training frames contain the package fully occluded by the gripper, which confuses the policy and prevents it from learning reliable grasp decisions.

Failure Modes. Our method can fail if the AV pipeline selects a suboptimal view. In practice, the pipeline typically generates views that are considered reasonable by humans, validating the effectiveness of sparse-view 3DGS for active vision. However, we occasionally observe occlusion of critical regions, such as drawer handles during rollouts, because the current formulation does not explicitly enforce full visibility of key object parts. Downstream imitation learning methods can also fail despite receiving a good view. TT with AV may fail due to inaccurate feature matching or calibration errors—either from hand–eye calibration or from the relative calibration between the observer and actor arms. For BC with AV, failures mostly occur when the actor arm’s initial end-effector state in the camera frame deviates from those seen in the demonstrations, causing compounding errors. This is more likely when the test-time optimal views differ significantly from the demonstration view. Conversely, tasks typically succeed when the actor’s gripper is sufficiently close to the target object, as fine-grained actor–object interactions

TABLE I: Success Rates of Different Methods on the Five Tasks (without and with occlusions).

Method	Pick Mug		Hammer Nail		Open Drawer		Retrieve Pack		Insert Coin		Sum	
	no occ	occ	no occ	occ	no occ	occ	no occ	occ	no occ	occ	no occ	occ
TT (3-views)	5/10	3/10	4/10	4/10	2/10	1/10	N/A	1/10	0/10	0/10	11/40	9/50
TT (ours with AV)	8/10	8/10	10/10	7/10	6/10	4/10	N/A	9/10	3/10	2/10	27/40	30/50
BC (static camera)	4/10	4/10	3/10	2/10	1/10	1/10	N/A	0/10	0/10	0/10	8/40	7/50
BC (ours with AV)	6/10	5/10	4/10	5/10	3/10	3/10	N/A	3/10	1/10	1/10	14/40	17/50



Fig. 3: **Images of Optimal Views.** Top row: demonstration optimal views. Middle row: test-time optimal views in 3DGS with gripper mask overlay. Bottom row: real world test-time optimal views. Red boxes indicate the task-relevant object parts. Test-time optimal views are derived by reconstructing the demonstration’s optimal viewpoints subject to minimal occlusion.




Fig. 4: **Data Efficiency of Behavior Cloning with Active Vision.** With the same number of demonstrations, our method outperforms the static camera setup across three tasks, evaluated using 30, 50, and 70 demonstrations.

are accurately captured in the optimal demonstration-view dataset. Finally, because the learned policy relies on a single RGB observation, it remains sensitive to depth ambiguities.

B. Representing Actions in the Camera Frame

We evaluate the effect of representing actions in the camera frame for BC with AV. We compare against prior approaches [7, 8], where actions are represented in the fixed actor robot frame and the camera pose is provided as an

input. Results in Table II show that representing actions directly in the camera frame improves both generalization and task success. We hypothesize that providing the camera pose as input is less effective because object poses span a wide distribution while camera views remain relatively similar, making it difficult for the policy to infer the robot’s pose in the static frame. On the other hand, representing actions directly in the camera frame provides a consistent reference.

However, representing actions in the camera frame also

TABLE II: Impact of action representation on success rates.

Method	Mug	Hammer
Camera pose as input, action in robot frame	1/10	0/10
Action in camera frame (ours)	6/10	4/10

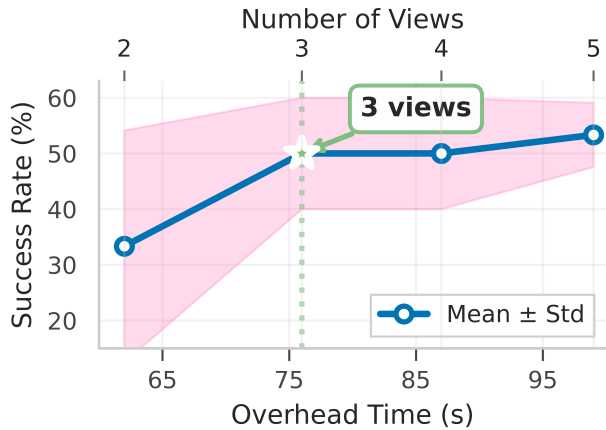


Fig. 5: Effect of Number of Exploration Views.

introduces a potential issue. In practice, we found that the training viewpoints must be sufficiently varied; otherwise, the model can shortcut learning by relying primarily on the camera pose. When trajectories relative to the camera are too similar, the policy may ignore the image inputs entirely.

C. Impact of Number of Exploration Views

We evaluate the impact of exploring three views per arm at test time to acquire scene information. Specifically, we compare 2–5 exploration viewpoints per arm to study how the extent of visual coverage affects downstream manipulation success. We use BC with AV and evaluate on the *Pick Mug* task under occlusion, with results shown in Figure 5. Our findings indicate that three views strike a good balance between performance and execution time, providing sufficient scene coverage for reliable role assignment and optimal viewpoint selection. Since frame alignment requires at least three views, using only two necessitates relying on camera poses inferred from the robot’s encoders and hand–eye calibration, bypassing full alignment. This leads to higher variance in performance because the resulting 3DGS reconstructions are less accurate, widening the sim-to-real gap. Consequently, test-time optimal views may also be less reliable. This limitation could be mitigated with more precise robots and better calibration.

TABLE III: Time breakdown for each component of our AV pipeline, measured with a RTX 4080Ti GPU.

Component	Time (s)
Six-view Exploration	18
InstantSplat Geometric Initialization	20
InstantSplat 3DGS Training	23
Optimal View Initialization	3
Differential Rendering	12
Total	76

Table III reports the time breakdown of our AV pipeline. The majority of time is spent on InstantSplat, while the exploration step could be further reduced by increasing the arms’ speed. We believe that as sparse-view 3DGS methods continue to improve, our system will become more efficient.

V. CONCLUSIONS

We have present ObAct, a novel observer–actor framework for imitation learning in active vision, in which an observer arm computes and moves to optimal visual observations to guide the actor arm’s actions. Our method captures fine-grained gripper–object interactions, supporting robust manipulation across diverse scenarios using both trajectory transfer and behavior cloning. Experimental results demonstrate significant gains in success rates compared to a static camera setup across both unoccluded and occluded scenarios.

Despite the effectiveness of our method, several limitations remain. The active-vision pipeline is relatively slow, the approach is tailored to short-horizon tasks, and it lacks reactivity to environmental changes. Moreover, the current setup cannot handle tasks that require two arms acting simultaneously. We identify several promising directions for future work. One is to explore dynamic viewpoints that continuously track gripper–object interactions during execution, enabling richer visual feedback and closed-loop occlusion avoidance. Another is to extend our approach to long-horizon tasks and to deformable-object manipulation, both of which introduce additional challenges for active vision and imitation learning. Finally, for dual-arm manipulation, we envision expanding the system to a three-arm configuration, in which one arm dynamically acts as the observer while the remaining two serve as the manipulators. We believe these extensions will enable more resilient active-vision robotic systems capable of handling complex manipulation tasks in diverse, unstructured real-world environments.

REFERENCES

- [1] Cheng Chi et al. “Diffusion policy: Visuomotor policy learning via action diffusion”. In: *The International Journal of Robotics Research* (2023), p. 02783649241273668. 1, 2
- [2] Kamil Dreczkowski et al. “Learning a thousand tasks in a day”. In: *Science Robotics* 10.108 (2025), eadv7594. 1, 2, 4
- [3] Norman Di Palo and Edward Johns. “Dinobot: Robot manipulation via retrieval and alignment with vision foundation models”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 2798–2805. 1
- [4] Cheng Chi et al. “Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots”. In: *Proceedings of Robotics: Science and Systems (RSS)*. 2024. 1
- [5] Tony Zhao et al. “Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware”. In: *Robotics: Science and Systems XIX* (2023). 1, 2, 5

- [6] Yilong Wang and Edward Johns. “One-Shot Dual-Arm Imitation Learning”. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. 2025, pp. 5660–5668. 1, 2
- [7] Ian Chuang et al. “Active vision might be all you need: Exploring active vision in bimanual robotic manipulation”. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2025. 1, 2, 4–6
- [8] Haoyu Xiong et al. “Vision in Action: Learning Active Perception from Human Demonstrations”. In: *arXiv preprint arXiv:2506.15666* (2025). 1, 2, 6
- [9] Xuxin Cheng et al. “Open-TeleVision: Teleoperation with Immersive Active Visual Feedback”. In: *Proceedings of The 8th Conference on Robot Learning*. Ed. by Pulkit Agrawal, Oliver Kroemer, and Wolfram Burgard. Vol. 270. Proceedings of Machine Learning Research. PMLR, June 2025, pp. 2729–2749. 1, 2, 4
- [10] Bernhard Kerbl et al. “3D Gaussian Splatting for Real-Time Radiance Field Rendering”. In: *ACM Transactions on Graphics* 42.4 (2023), pp. 1–14. 1
- [11] John Aloimonos, Isaac Weiss, and Amit Bandyopadhyay. “Active vision”. In: *International journal of computer vision* 1 (1988), pp. 333–356. 2
- [12] Kanzhi Wu, Ravindra Ranasinghe, and Gamini Disanayake. “Active recognition and pose estimation of household objects in clutter”. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2015, pp. 4230–4237. 2
- [13] Boshi An et al. “Rgbmanip: Monocular image-based robotic manipulation through active object pose estimation”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 7748–7755. 2
- [14] Stefan Isler et al. “An information gain formulation for active volumetric 3D reconstruction”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016, pp. 3477–3484. 2
- [15] Soomin Lee et al. “Uncertainty guided policy for active robotic 3d reconstruction using neural radiance fields”. In: *IEEE Robotics and Automation Letters* 7.4 (2022), pp. 12070–12077. 2
- [16] Michel Breyer et al. “Closed-loop next-best-view planning for target-driven grasping”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2022, pp. 1411–1416. 2, 3
- [17] Xuechao Zhang et al. “Affordance-Driven Next-Best-View Planning for Robotic Grasping”. In: *Proceedings of The 7th Conference on Robot Learning*. Ed. by Jie Tan, Marc Toussaint, and Kourosh Darvish. Vol. 229. Proceedings of Machine Learning Research. PMLR, June 2023, pp. 2849–2862. 2
- [18] Haoxiang Ma et al. “Active Perception for Grasp Detection via Neural Graspness Field”. In: *Advances in Neural Information Processing Systems*. Ed. by A. Globerson et al. Vol. 37. Curran Associates, Inc., 2024, pp. 38122–38141. 2
- [19] Yitian Shi et al. “VISO-Grasp: Vision-Language Informed Spatial Object-centric 6-DoF Active View Planning and Grasping in Clutter and Invisibility”. In: *arXiv preprint arXiv:2503.12609* (2025). 2
- [20] Guokang Wang et al. “Observe Then Act: Asynchronous Active Vision-Action Model for Robotic Manipulation”. In: *IEEE Robotics and Automation Letters* (2025). 2
- [21] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. “Perceiver-actor: A multi-task transformer for robotic manipulation”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 785–799. 2
- [22] Theophile Gervet et al. “Act3D: 3D Feature Field Transformers for Multi-Task Robotic Manipulation”. In: *7th Annual Conference on Robot Learning*. 2
- [23] Ankit Goyal et al. “Rvt: Robotic view transformer for 3d object manipulation”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 694–710. 2
- [24] Stephen Tian et al. “View-Invariant Policy Learning via Zero-Shot Novel View Synthesis”. In: *arXiv* (2024). 2
- [25] Kyle Sargent et al. “Zeronvs: Zero-shot 360-degree view synthesis from a single image”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 9420–9429. 2
- [26] Sizhe Yang et al. “Novel demonstration generation with gaussian splatting enables robust one-shot manipulation”. In: *arXiv preprint arXiv:2504.13175* (2025). 2
- [27] Johan Edstedt et al. “Roma: Robust dense feature matching”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 19790–19800. 3, 4
- [28] Zhiwen Fan et al. “InstantSplat: Sparse-view Gaussian Splatting in Seconds”. In: *arXiv preprint arXiv:2403.20309* (2024). 3
- [29] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. “Grounding image matching in 3d with mast3r”. In: *European Conference on Computer Vision*. Springer. 2024, pp. 71–91. 3
- [30] Tianhe Ren et al. “Grounded sam: Assembling open-world models for diverse visual tasks”. In: *arXiv preprint arXiv:2401.14159* (2024). 4
- [31] Maxime Oquab et al. “DINOv2: Learning Robust Visual Features without Supervision”. In: *Transactions on Machine Learning Research Journal* (2024). 4
- [32] Kai Chen et al. “Vision Foundation Model Enables Generalizable Object Pose Estimation”. In: *Advances in Neural Information Processing Systems*. Ed. by A. Globerson et al. Vol. 37. Curran Associates, Inc., 2024, pp. 19975–20002. 4
- [33] Nikhila Ravi et al. “SAM 2: Segment Anything in Images and Videos”. In: *The Thirteenth International Conference on Learning Representations*. 4
- [34] Lin Yen-Chen et al. “Inerf: Inverting neural radiance fields for pose estimation”. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2021, pp. 1323–1330. 4