

# AutoFocus-IL: VLM-based Saliency Maps for Data-Efficient Visual Imitation Learning without Extra Human Annotations

Litian Gong<sup>1</sup>, Fatemeh Bahrani<sup>2</sup>, Yutai Zhou<sup>2</sup>, Amin Banayeeanzade<sup>2</sup>, Jiachen Li<sup>1</sup>, Erdem Bıyık<sup>2</sup>

**Abstract**—We present AutoFocus-IL, a simple yet effective method to improve data efficiency and generalization in visual imitation learning by guiding policies to attend to task-relevant features rather than distractors and spurious correlations. Saliency regularization has emerged as a promising way to achieve this, but existing approaches typically require costly supervision such as human gaze data or manual saliency annotations. In contrast, AutoFocus-IL leverages vision-language models (VLMs) to automatically identify and track key objects in demonstrations, generating temporal saliency maps that highlight causal visual signals while suppressing distractors. These maps are then used to regularize behavior cloning policies, yielding stronger alignment between visual attention and task-relevant cues. Our findings highlight that VLM-driven saliency provides a scalable, annotation-free path toward robust imitation learning in robotics. Particularly, our experiments in both the CARLA simulator and real-robot manipulation tasks demonstrate that AutoFocus-IL not only outperforms standard behavior cloning but also surpasses state-of-the-art baselines that assume privileged access to human supervision, such as gaze data. The supplementary materials, including code, datasets, and trained policy videos, are publicly available at <https://AutoFocus-IL.github.io/>.

## I. INTRODUCTION

Imitation learning (IL) has emerged as a central paradigm in robotics, enabling agents to acquire complex behaviors from expert demonstrations [1], [2]. Recent advances leverage large vision-language-action (VLA) models [3]–[5], trained through behavior cloning (BC) on robot demonstration datasets, to bridge perception and action. While promising, this approach remains bottlenecked by the limitations of robot data collection [6], [7], leading to several fundamental challenges that hinder scalability and real-world deployment.

A primary issue is *data efficiency*. Unlike images or language [8], [9], robotic data is expensive and slow to collect, requiring physical interaction with the environment and human teleoperation. Estimates suggest that reaching the same scale as today’s language or vision corpora would take over a century of robot operation [10]. Consequently, designing imitation learning methods that make better use of limited data is essential.

This data scarcity further compounds the challenge of *generalization* [11]. Existing VLA models consistently show large performance gains when fine-tuned on in-domain data [12], revealing their reliance on task- and environment-specific supervision. In practice, this means an end-user

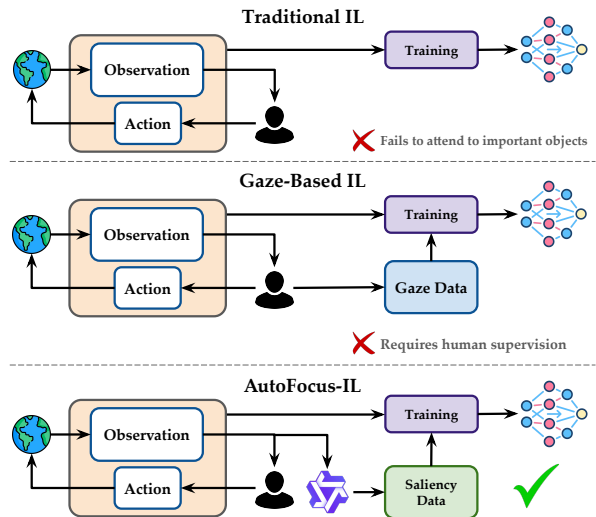


Fig. 1: An overview of three different approaches: While traditional imitation learning suffers from causal confusion, gaze-based IL solves this by utilizing an expensive solution of collecting human eye gaze data. However, AutoFocus-IL resolves this issue by getting a saliency map annotated by a VLM to retain the benefits of gaze-based IL without incurring the extra data collection costs.

may need to provide hundreds of demonstrations for a new task, severely limiting real-world applicability. Since both pretraining and fine-tuning typically rely on behavior cloning [13], any improvements in BC would directly enhance policy generalization.

Another persistent issue in imitation learning is *causal confusion* [14], where models make predictions based on spurious correlations between visual cues and expert actions. For example, they may attend to background distractors that co-occur with certain actions but are irrelevant to the task [15]. This not only degrades policy robustness but also makes policies brittle in unseen environments.

In this work, we propose a simple yet effective method, AutoFocus-IL (**A**utomatically **F**ocusing on context-relevant objects using VLM supervision for **I**mitation **L**earning), to preprocess demonstrations and regularize imitation learning without requiring any additional human annotations. Our approach improves data efficiency, enhances generalization, and mitigates causal confusion by explicitly highlighting task-relevant visual signals for the learner (see Fig. 1).

Concretely, we leverage open-source vision-language models to identify objects that are important for the demonstrated task. These candidate key objects are then tracked across the trajectory while also allowing for disappearances or newly introduced objects. Using these detections, we

<sup>1</sup> Department of Electrical and Computer Engineering, University of California, Riverside, CA, USA.

<sup>2</sup> Thomas Lord Department of Computer Science, University of Southern California, CA, USA.

Correspondence: lgong024@ucr.edu, biyik@usc.edu.

generate saliency masks that emphasize the key objects while masking distractors. The resulting masks are then used to regularize the attention of the policy network during behavior cloning. Experiments on autonomous driving in the CARLA simulator [16] show **104%** improvement over standard BC, while real-world experiments with a WidowX robot arm show **50%** improvement. Notably, our method not only surpasses BC but also outperforms baselines that assume access to privileged human annotations, such as gaze data, to identify the relevant and salient parts of the scene.

Overall, our contributions are threefold:

- 1) We introduce a context-aware VLM-driven object filtering pipeline for identifying task-relevant visual signals without additional supervision.
- 2) We propose a temporal saliency modeling approach that translates object tracks into structured signals.
- 3) We demonstrate that regularizing imitation learning with saliency maps produced with our approach substantially improves data efficiency, generalization, and robustness in both simulation and real-world robot experiments.

## II. RELATED WORK

### A. Data-Efficiency in Imitation Learning

A common approach to address data efficiency is to employ active learning or interactive IL frameworks [17], [18], where the agent queries a human expert in order to resolve uncertainties or increase state coverage during training [19], [20]. These strategies improve sample efficiency by reducing redundant demonstrations, but they require a human to remain in the loop throughout training, limiting scalability. In contrast, our method improves data efficiency without any additional expert queries: we preprocess existing demonstration datasets to automatically highlight task-relevant regions, effectively augmenting the supervision signal at no extra human cost.

Methods that target generalization across tasks and domains often incorporate data augmentation [21], [22], domain randomization [23], or auxiliary tasks for visual representation learning [24]. While effective to some extent, these methods typically remain agnostic to which parts of the scene are semantically important for the task. By explicitly filtering and emphasizing key objects in demonstrations, our approach complements these generalization techniques by ensuring that the model allocates capacity to causal features rather than incidental correlations.

Another line of work tackles causal confusion, where policies misattribute spurious correlations as causal signals. Prior approaches include expert interventions [14], object-centric feature dropout [15], and keyframes [25], but they often require additional expert supervision or assume simplified environments. Our method shares the spirit of causal deconfounding, but achieves it scalably via VLM-guided filtering without privileged annotations or interventions.

### B. Vision-Language Model Supervision for Robot Control

VLMs trained on Internet-scale data provide rich semantic and spatial knowledge, making them a powerful yet scalable alternative to human supervision for robotics. One particular line of research uses VLMs to generate weak labels for demonstrations. For example, NILS [26] leverages VLMs to automatically segment long-horizon trajectories, detect object-centric events, and emit instruction labels at scale, improving policy learning without human annotation.

Another prominent direction is reward modeling. VLMs have been applied as zero-shot success detectors [27], [28], providing episodic feedback for reinforcement learning without privileged state information. Others have used VLMs to estimate task progress [29]–[32], generate language-conditioned rewards [33], [34], or elicit preference feedback to train reward functions [35]. Beyond labeling and rewards, VLMs have also been employed for high-level task reasoning, grounding, and planning [36]–[39]. These methods demonstrate the versatility of VLM supervision in shaping robot behavior, but they often remain detached from the imitation learning pipeline itself.

Rather than using VLMs to produce episodic labels or high-level plans, our work leverages them to provide dense, context-aware supervision on individual frames, regularizing the learner’s visual attention toward causal signals without additional human effort.

### C. Human Supervision for Generating Saliency Maps

Saliency map has been a widely used tool for post-hoc interpretation of neural networks, and has also been used as a training supervision for enhancing model explainability [40], [41]. A prominent line of work enhances imitation learning by explicitly incorporating human-provided saliency supervision. For example, ViSaRL [42] introduces a framework where annotators mark task-relevant regions in demonstrations, which are then used to learn better state representation for policy learning. Such explicit annotations help encourage models to attend to causal features rather than spurious correlations. However, collecting large-scale saliency labels is costly and prone to human error.

Given the expense of manual annotations, many works instead turn to human gaze as a natural and implicit source of saliency (see Fig. 1). Human gaze is well known to correlate with task-relevant features [43]–[45] and can be captured during task execution without additional labeling effort. In robotic manipulation, Kim *et al.* [46], [47] showed that gaze supervision enables policies to ignore irrelevant distractors and focus on task-relevant objects. More generally, gaze-derived attention maps have been leveraged to regularize policies across domains, including robotics, driving, and video games [48]–[52].

Despite their effectiveness, both explicit saliency annotation and gaze supervision face practical limitations. Manual labeling is slow and difficult to scale, while gaze data requires specialized hardware and is inherently noisy [53], [54]. Moreover, gaze is directed at a single point, making it difficult to model the full spatial extent of task-relevant

regions [55], [56]. Our method circumvents these challenges by leveraging VLM-driven object filtering and temporal modeling to produce saliency maps directly from raw sensory inputs, yielding many of the benefits of human-derived saliency without any additional human annotation.

### III. METHODOLOGY

We propose AutoFocus-IL, a pipeline that automatically extracts task-relevant objects from demonstrations and uses them to generate saliency maps for imitation learning. The method consists of three components: (1) context-aware VLM filtering that identifies and tracks key objects across demonstration, (2) temporal saliency modeling that converts object tracks into smooth multi-peak saliency maps, and (3) a regularization technique that leverages saliency for policy learning. Fig. 2 provides an overview of our method.

#### A. Context-Aware VLM Filtering

Each demonstration is a trajectory of image-action pairs  $\xi = (o_0, a_0, o_1, a_1, \dots)$ , where  $o_t \in [0, 1]^{w \times h \times 3}$  is the visual observation at time  $t$  and  $a_t$  is the expert action. Our first goal is to annotate each observation  $o_t$  with a saliency map  $g_t \in [0, 1]^{w \times h}$  where each pixel denotes how important the corresponding pixel of  $o_t$  is for the task. For example, a pixel that belongs to a cloud should have a low saliency value for the autonomous driving task, whereas the pixels of a stop sign should be labeled as highly salient when the car is approaching an intersection.

The straightforward way of automatically getting these labels would be to query the VLM with each and every observation. However, this approach is not easily scalable because processing every single frame with a VLM would incur a prohibitively high computational cost. Instead, we first perform equidistant frame sampling with stride  $s$  to select a small subset of frames from each demonstration.

Another challenge of using VLMs to generate saliency maps is due to the uncertainty about the objects and the task. Since we are not making any assumptions about the expert demonstration dataset, we only have a set of trajectories without any language label of what the task is or what objects are present in the scene. This creates an ambiguity: if we are trying to train a robot arm to pick up a carrot, the carrot is obviously relevant; in the same scene, the task might be dicing a tomato, in which case the carrot should not be labeled as salient.

To resolve this, we first query the VLM with the entire subset of frames that were sampled from the full demonstration. In this work, we use Qwen2.5-VL-Instruct-72B [57] as the VLM. We ask it to generate<sup>1</sup>:

- 1) A *global context summary*,  $c$ , of the trajectory, including descriptions of the task, the environment, and potential risks.
- 2) A *candidate object vocabulary*  $\mathcal{V}$ , i.e., the set of objects that could appear in the demonstration.

<sup>1</sup>Detailed instructions used to prompt the VLM can be found in the supplementary material, available on the [project website](#).

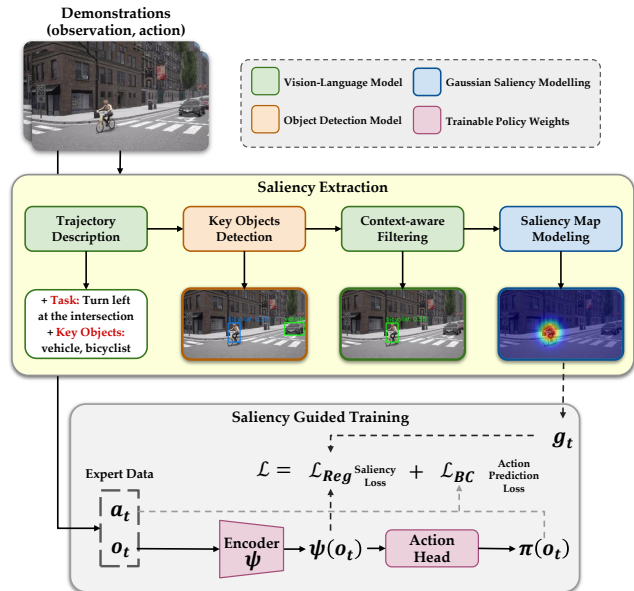


Fig. 2: Overview of the AutoFocus-IL pipeline.

While the latter is useful for detecting and tracking the objects in the scene, the former will be used to determine their relevance to the task.

**Detection and Tracking.** Using the object vocabulary  $\mathcal{V}$ , we apply an open-set object detector (in this work, Grounding DINO [58]) to every frame in the trajectory. The detector outputs bounding boxes and confidence scores for all recognized objects. To maintain consistent object identities across time, we perform frame-to-frame association via the Hungarian method [59] based on bounding box overlap (Intersection-over-Union) in consecutive frames<sup>2</sup>. This yields continuous object tracks, each with a persistent ID.

Because the set of visible objects can change during a demonstration (objects enter or leave the scene), we divide the trajectory into contiguous sub-sequences such that within each sub-sequence, the set of tracked object IDs remains constant. This segmentation of the trajectory allows the VLM to reason about objects in a locally stable context.

**Filtering with VLM supervision.** After obtaining the list of task-relevant objects in the scene, we filter them based on their importance. For example, a pedestrian is always task-relevant for autonomous driving, but should only affect autonomous driving decisions if they are on the road; in contrast, a pedestrian walking on the sidewalk is not critical and should not affect the driving system. We achieve this filtering with the VLM: for the first frame of each sub-sequence, we construct a context prompt that consists of the frame image  $o_t$ , the expert action  $a_t$ , the global context summary  $c$ , and the list of active objects with their IDs and bounding boxes. The VLM is then asked to select the *key objects* that are important for deciding to take action  $a_t$  at

<sup>2</sup>We chose to use Grounding DINO with the Hungarian method for tracking the detected objects due to its computational efficiency. Alternative implementations could use object tracking models from computer vision (e.g., [60], [61]) or other traditional tracking algorithms from the image processing literature (e.g., [62], [63]).

that timestep.

Objects not marked as key by the VLM are pruned for the entire sub-sequence, ensuring consistency across time. If the VLM detects that important objects are missing (e.g., due to detection failure), it flags a set of missing categories. The detector is then re-run with the augmented vocabulary until either the missing objects are recovered or a small iteration cap is reached. At the end of this stage, each frame is annotated with the set of task-relevant object bounding boxes. These serve as the basis for downstream saliency modeling.

### B. Temporal Multi-Peak Saliency Modeling

The filtered object tracks indicate which objects matter, but not how strongly the learner should attend to them across time. We therefore convert the key-object annotations into continuous saliency maps. Namely, for each timestep  $t$ , we collect the centers of all key objects detected in the observations  $o_t, o_{t-1}, \dots, o_{t-t'}$  where  $t'$  is a hyperparameter. To construct the saliency map  $g_t$ , we add a Gaussian distribution around each object center. Temporal weighting is applied so that more recently detected objects are sharper and more influential, while older detections are more blurred and less influential.

Mathematically, we perform the following operation to obtain the saliency map at any pixel location  $p$ :

$$\bar{g}_t(p) = \sum_{k=0}^{t'} \alpha^k \sum_i \exp\left(-\frac{\|p - \mu_{t-k}^{(i)}\|_2^2}{2\gamma^2 \beta^{-2k}}\right), \quad (1)$$

where  $\alpha, \beta \in [0, 1], \gamma > 0$  are hyperparameters, and the inner summation is over the key objects in the sub-sequence whose center locations are denoted by  $\mu$ . The variance increases from the current frame  $t$  to the past, producing broader kernels for older frames. We also weight the Gaussian kernel at each timestep by a temporal decay factor  $\alpha$ .

Finally, normalizing  $\bar{g}_t$  over all pixels into the range  $[0, 1]$  yields a dense saliency map  $g_t$  for frame  $t$ . This multi-peak temporal model preserves evidence from multiple objects across time while naturally fading older observations, thereby mimicking human-like sustained attention. Also, by retaining information from previous frames, we enhance the quality of saliency maps by mitigating the impact of missed detections in the current frame.

### C. Saliency-Guided Regularization

Following the work of [52], we integrate the VLM-generated saliency maps with behavioral cloning via a knowledge distillation approach [64]. In fact, we employ a convolutional neural network (CNN)-based policy  $\pi$  trained to map images  $o_t$  to actions  $a_t$  via supervised learning, and regularize the intermediate feature maps of the policy using the VLM-generated saliency maps. Let  $\psi(o_t) \in \mathbb{R}^{h \times w \times c}$  be the CNN feature map representation of an observation  $o_t$ , upsampled to the input resolution  $w \times h$ , with  $c$  channels. We minimize the  $\ell_2$ -norm of activation map features modulated by the saliency map  $g_t$ , with the goal of penalizing feature



Fig. 3: Confounded overlay visualization in the CARLA simulator. Action-conditioned icons are rendered along the top margin of each frame to induce spurious correlations, while leaving the underlying dynamics and expert labels unchanged. The red circle simulates a brake light, while arrows denote the steering direction, with their thickness indicating the throttle applied in the previous timestep.

activations in regions not covered by the saliency map. The final training loss is:

$$\mathcal{L} = \mathcal{L}_{BC}(\pi, o_t, a_t) + \lambda \|(1 - g_t) \odot \psi^2(o_t)\|_2, \quad (2)$$

where  $\odot$  is the element-wise matrix product. The first term is the standard behavior cloning loss for observation  $o_t$  with expert action  $a_t$  (e.g., mean squared error for continuous action spaces or cross-entropy for discrete action spaces), and the second term performs the regularization with the VLM-generated saliency map, with its strength scaled by the hyperparameter  $\lambda$ . Note that AutoFocus-IL computes saliency only during training; at inference time, no additional saliency predictor is required, in contrast to several prior studies [42], [48], [49].

## IV. SIMULATION EXPERIMENTS

### A. Environment Setup

We evaluate AutoFocus-IL in the CARLA urban driving simulator [16] using the Bench2Drive benchmark [65]. We consider two observation settings: (i) *Original*, where the policy receives raw RGB images, and (ii) *Confounded*, where we introduce action-conditioned icons on top of each frame. These overlays do not affect dynamics or expert actions, but introduce spurious correlations intended to test robustness against causal confusion (Fig. 3). This is a modification strategy adopted by several prior studies, e.g., [15], [52].

We use the dataset provided by Banayeezade *et al.* [52] to show AutoFocus-IL can improve imitation learning on datasets that are readily available on the Internet. The training data consists of 10 routes with 20 expert trajectories per route. Evaluation is conducted on two disjoint sets: (1) *seen* routes, which are the same 10 routes used for training but with 2 different random seeds (20 evaluations per method), and (2) *unseen* routes, a disjoint set of 10 held-out routes, also with 2 seeds each (20 evaluations per method).

Input images are resized to  $320 \times 180 \times 3$  (width  $\times$  height  $\times$  channels). All baselines share the same backbone CNN encoder and optimizer configuration; auxiliary heads or additional loss terms are only included if explicitly required

by the baseline. This ensures differences in performance arise from the saliency supervision rather than from architecture or optimization choices.

### B. AutoFocus-IL Hyperparameters

For each demonstration, we query the VLM on a fixed set of 25 uniformly sampled frames for global context and vocabulary generation, i.e.,  $s = 25$ . For the Gaussian distributions in saliency map generation (Eq. 1), we let  $\alpha = 0.7$ ,  $\beta = 0.8$ ,  $\gamma = 15$ , and  $t' = 4$ .

We use the final convolutional feature map of the policy network as  $\psi$  in Eq. 2 and let  $\lambda = 10$ . Both behavior cloning loss and saliency regularization are applied to all frames, except in an ablation study, where the saliency regularization is applied to the prescribed fraction of frames.

### C. Baselines and Metric

We compare AutoFocus-IL against plain BC with no saliency regularization. In addition, we compare it against baselines that use human gaze as a regularizer, thanks to the availability of gaze information in the dataset, which was not used for AutoFocus-IL. These baselines are the following: GMD [48], ViSaRL [42], GRIL [66], AGIL [49], GABRIL [52], and GABRIL+GMD which was implemented by Banayeezade *et al.* [52] to incorporate GMD into their framework.

Performance is measured by the *driving score* from CARLA Leaderboard 2.0 [67], a scalar metric that aggregates success rate, safety, and efficiency. Scores are averaged across seeds and routes for each evaluation split.

### D. Main Results

Fig. 4a and Fig. 4b show that AutoFocus-IL improves over BC. It attains a driving score of **62.36%** on seen routes and **30.17%** on unseen routes, whereas these numbers are only **47.83%** and **12.20%** for BC. This shows that AutoFocus-IL is an effective method that processes existing imitation learning datasets to achieve better results with the easy-to-implement saliency regularization.

Even more impressively, AutoFocus-IL outperforms baselines that have access to human gaze data for resolving causal confusion or improving data efficiency. Without access to this privileged information, AutoFocus-IL improves driving scores by more than **2.2** points over the top baseline in seen routes, and by more than **1.7** points in unseen routes.

In Fig. 4c and Fig. 4d, we present the results from the confounded settings with action-conditioned overlays. As expected, this setting leads to lower driving scores because of causal confusion. However, AutoFocus-IL remains the best performing method in both seen and unseen routes. While GRIL achieves competitive performance in seen routes, AutoFocus-IL generalizes better and attains a much higher driving score in unseen routes. It is worth emphasizing again that GRIL and all other baselines except BC utilize human gaze/saliency information, which requires additional human annotations or hardware, whereas AutoFocus-IL’s saliency maps are generated automatically with the VLM.

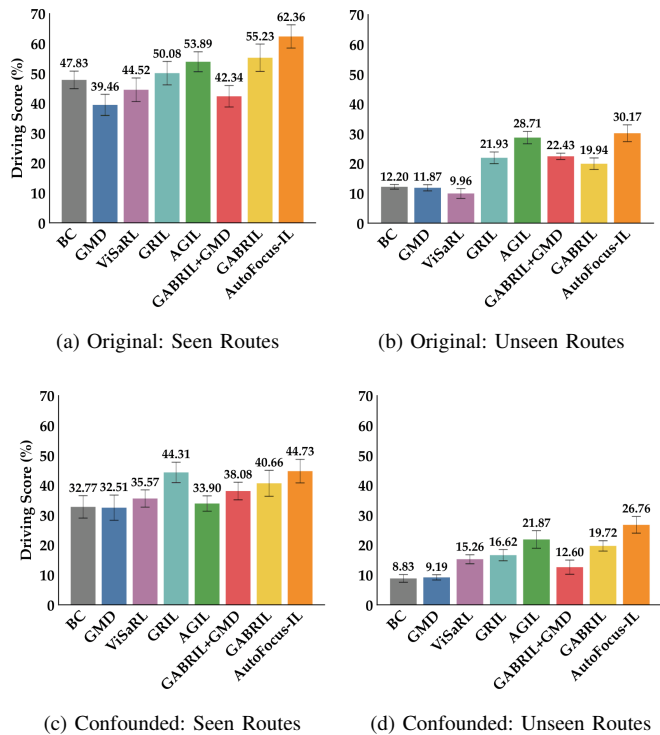


Fig. 4: CARLA Driving Score (mean  $\pm$  standard error) on seen and unseen routes in both original and confounded environments. Except BC, other baselines use human gaze to improve imitation learning, while AutoFocus-IL uses VLM-generated saliency maps.

TABLE I: CARLA Driving Score (mean  $\pm$  standard error) obtained from using Temporal vs. Non-Temporal Saliency modeling.

Setting	Non-Temporal	Temporal
Seen—Original	49.83 $\pm$ 3.22	62.36 $\pm$ 3.89
Seen—Confounded	34.93 $\pm$ 3.65	44.73 $\pm$ 3.94
Unseen—Original	16.04 $\pm$ 1.73	30.17 $\pm$ 2.83
Unseen—Confounded	12.55 $\pm$ 0.88	26.76 $\pm$ 2.80
Average	28.33 $\pm$ 2.37	41.01 $\pm$ 3.44

**Why does AutoFocus-IL outperform gaze-based methods?** Human gaze is a *single fixation point* that cannot faithfully capture object boundaries [55], [56] and is noisy due to individual differences and fatigue [53], [54]. VLM-based saliency instead produces deterministic, *object-level* coverage grounded in bounding box detections, providing denser and more consistent supervision. Even the strongest gaze-based baselines, AGIL [49] and GRIL [66], fall short, suggesting that the quality of the saliency source matters more than the sophistication of the integration mechanism.

### E. Ablation Studies

**Temporal vs. non-temporal saliency.** We compare our temporal multi-peak saliency modeling against a non-temporal, frame-wise variant. As shown in Table I, temporal modeling consistently improves results across all splits, yielding an average gain of **12.7** points in the driving score.




Fig. 5: **Saliency fraction sweep.** Performance across Seen/Unseen and Original/Confounded splits.

**Fraction of frames with saliency supervision.** In Fig. 5, we vary the fraction  $f \in \{0\%, 10\%, 25\%, 50\%, 75\%, 100\%\}$  of frames that receive saliency regularization, while all frames are always trained with the BC loss. Performance increases with higher fractions of supervised frames, saturating around 75%, suggesting that moderate coverage of saliency supervision may be sufficient to guide attention effectively.

## V. REAL-ROBOT EXPERIMENTS

### A. Platform and Tasks

We test AutoFocus-IL on a **WidowX** robot arm with a single monocular RGB camera at a resolution of  $640 \times 480$  (see Fig. 6 for a visualization). Training images are center-cropped and resized to  $256 \times 256$ . We evaluate on two manipulation tasks: *Lift Carrot* and *Pull Pot*. For each task, we also design a confounded variant by placing task-irrelevant distractor objects in the background to induce visual causal confusion. We consider the setup as four distinct single tasks, each trained with 100 demonstrations.

### B. Policy and Saliency Integration

At each timestep, the policy receives the latest two frames as input. All methods share the same vision backbone and optimization setup for a fair comparison. The policy uses a ResNet-34 [68] image encoder whose last convolutional feature map is flattened and concatenated with the 7D end-effector state, then passed to a Gaussian MLP head to parameterize a Gaussian action distribution. The saliency regularization weight is  $\lambda = 5$ . We let  $\alpha = 0.7$ ,  $\beta = 0.8$ ,  $\gamma = 30$ ,  $t' = 4$ .

We additionally fine-tuned the YOLO [69] object detector to specifically detect the gripper of the robot arm, making up for the fact that Grounding DINO fails to detect the gripper.

### C. Evaluation Protocol

For each of the four settings, we perform 10 rollouts and report the number of successful episodes (task completion). We compare AutoFocus-IL against Behavior Cloning (BC).

### D. Results and Discussion

Results are summarized in Table II. AutoFocus-IL consistently improves success on both tasks, with the largest gains under the confounded setting.

TABLE II: Real-robot manipulation task success rate

Tasks	AutoFocus-IL	BC
Lift Carrot	7/10	5/10
Lift Carrot Confounded	7/10	1/10
Pull Pot	8/10	2/10
Pull Pot Confounded	6/10	0/10

These trends mirror our simulation findings: VLM-driven, object-centric saliency helps the policy focus on causal scene elements and markedly improves robustness in the presence of unrelated visual clutter—without any additional human attention labels. Note that prior baselines such as GABRIL are not applicable to this task, as they require human gaze data—necessitating either specialized equipment for gaze collection or complex projection methods to map gaze onto the robot’s point-of-view camera.

## VI. CONCLUSION

**Summary.** In this study, we proposed a novel framework, AutoFocus-IL, which generates temporal saliency maps based on a vision-language model with zero human effort and uses them to regularize visual imitation learning. This framework significantly improves data efficiency and generalization capabilities, and can mitigate causal confusion in imitation learning. On the CARLA driving and real-world robotics tasks, it achieves stable improvements in both the *Original* and *Confounded* settings, and outperforms prior baselines that require human gaze, demonstrating its scalability and practicality.

**Limitations and Future Work.** Our regularized training only considers CNN layers as a feature extractor. However, in recent years, the vision transformer (ViT) [70] backbone has been widely used in visual imitation learning [3]–[5]. Recent research in computer vision has shown that regularization can enhance the performance of ViT models by improving the attention mechanism [71], [72]. Therefore, future work could focus on using our AutoFocus-IL framework to generate saliency maps for improving the performance of ViT-based policy models through regularized training.

Another limitation is that our method assumes that the VLM judgment from the first frame of each sub-sequence can represent object importance throughout the sub-sequence. However, in the real world, object states may change (e.g., a pedestrian entering the crosswalk), which may not be captured if detected objects remain the same, potentially causing the model to miss critical cues. This limitation also reflects the current VLM’s restricted ability to perceive motion and dynamic state transitions. Future work should explore richer contextual inputs [73] and more dynamic-aware VLMs [74] to better capture complex physical interactions.

## ACKNOWLEDGMENTS

We thank Soheil Habibian for early discussions on the idea. Yutai Zhou was supported by a fellowship from USC-Capital One Center for Responsible AI and Decision Making in Finance (CREDIF).




Fig. 6: Visualizations of our experimental setups and data processing steps. For example, in the WidowX Robotic Arm setup, the task is to pull a pot over the oven. Although the scene contains many distractors, AutoFocus-IL attends only to the most task-relevant objects.

## REFERENCES

- [1] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An Algorithmic Perspective on Imitation Learning," *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, Mar. 2018, ISSN: 1935-8253, 1935-8261.
- [2] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [3] B. Zitkovich, T. Yu, S. Xu, *et al.*, "RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control," in *Proceedings of The 7th Conference on Robot Learning*, PMLR, Dec. 2023, pp. 2165–2183.
- [4] K. Black, N. Brown, D. Driess, *et al.*, " $\pi_0$ : A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.
- [5] M. Shukor, D. Aubakirova, F. Capuano, *et al.*, *SmolVLA: A Vision-Language-Action Model for Affordable and Efficient Robotics*, Jun. 2025. eprint: 2506.01844 (cs).
- [6] V. Saxena, M. Bronars, N. R. Arachchige, *et al.*, "What Matters in Learning from Large-Scale Datasets for Robot Manipulation," in *The Thirteenth International Conference on Learning Representations*, Oct. 2024.
- [7] A. Mandlekar, D. Xu, J. Wong, *et al.*, "What Matters in Learning from Offline Human Demonstrations for Robot Manipulation," in *Conference on Robot Learning*, Aug. 2021.
- [8] 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*, Jun. 2009, pp. 248–255.
- [9] *Common Crawl - Open Repository of Web Crawl Data*, <https://commoncrawl.org>.
- [10] K. Goldberg, "Good old-fashioned engineering can close the 100,000-year "data gap" in robotics," *Science Robotics*, vol. 10, no. 105, eaca7390, Aug. 2025.
- [11] J. Gao, S. Belkhale, S. Dasari, A. Balakrishna, D. Shah, and D. Sadigh, "A taxonomy for evaluating generalist robot manipulation policies," *IEEE Robotics and Automation Letters (RA-L)*, 2026.
- [12] M. J. Kim, C. Finn, and P. Liang, "Fine-tuning vision-language-action models: Optimizing speed and success," *arXiv preprint arXiv:2502.19645*, 2025.
- [13] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," *Advances in neural information processing systems*, vol. 1, 1988.
- [14] P. de Haan, D. Jayaraman, and S. Levine, "Causal confusion in imitation learning," in *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019.
- [15] J. Park, Y. Seo, C. Liu, *et al.*, "Object-aware regularization for addressing causal confusion in imitation learning," in *Advances in Neural Information Processing Systems*, vol. 34, Curran Associates, Inc., 2021, pp. 3029–3042.
- [16] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*, PMLR, 2017, pp. 1–16.
- [17] C. Celemin, R. Pérez-Dattari, E. Chisari, *et al.*, "Interactive imitation learning in robotics: A survey," *Foundations and Trends® in Robotics*, vol. 10, no. 1-2, pp. 1–197, 2022.
- [18] B. Settles, "Active Learning Literature Survey," University of Wisconsin-Madison Department of Computer Sciences, Technical Report, 2009.
- [19] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [20] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer, "Hg-dagger: Interactive imitation learning with human experts," in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 8077–8083.
- [21] Q. Chen, S. Kiami, A. Gupta, and V. Kumar, "GenAug: Retargeting behaviors to unseen situations via Generative Augmentation," in *Robotics: Science and Systems XIX*, Robotics: Science and Systems Foundation, Jul. 2023, ISBN: 978-0-9923747-9-2.
- [22] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, "Reinforcement learning with augmented data," *Advances in neural information processing systems*, vol. 33, pp. 19 884–19 895, 2020.
- [23] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 23–30.
- [24] M. Laskin, A. Srinivas, and P. Abbeel, "CURL: Contrastive Unsupervised Representations for Reinforcement Learning," in *Proceedings of the 37th International Conference on Machine Learning*, PMLR, Nov. 2020, pp. 5639–5650.
- [25] C. Wen, J. Lin, J. Qian, Y. Gao, and D. Jayaraman, "Keyframe-Focused Visual Imitation Learning," in *Proceedings of the 38th International Conference on Machine Learning*, PMLR, Jul. 2021, pp. 11 123–11 133.
- [26] N. Blank, M. Reuss, M. Rühle, *et al.*, "Scaling robot policy learning via zero-shot labeling with foundation models," in *Proceedings of the 8th Conference on Robot Learning (CoRL)*, ser. Proceedings of Machine Learning Research, vol. 270, PMLR, 2025, pp. 4158–4187.
- [27] Y. Du, K. Konyushkova, M. Denil, *et al.*, "Vision-language models as success detectors," in *Proceedings of The 2nd Conference on Lifelong Learning Agents*, 2023, pp. 120–136.
- [28] J. Rocamonde, V. Montesinos, E. Nava, E. Perez, and D. Lindner, "Vision-language models are zero-shot reward models for reinforce-

- ment learning,” in *International Conference on Learning Representations (ICLR)*, 2024.
- [29] S. Sontakke, J. Zhang, S. Arnold, *et al.*, “Roboclip: One demonstration is enough to learn robot policies,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 55 681–55 693, 2023.
- [30] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang, “Vip: Towards universal visual reward and representation via value-implicit pre-training,” in *International Conference on Learning Representations (ICLR)*, 2023.
- [31] J. Zhang, Y. Luo, A. Anwar, *et al.*, “Rewind: Language-guided rewards teach robot policies without new demonstrations,” in *Conference on Robot Learning (CoRL)*, 2025.
- [32] A. Liang, Y. Korkmaz, J. Zhang, *et al.*, “Robometer: Scaling general-purpose robotic reward models via trajectory comparisons,” *arXiv preprint arXiv:2603.02115*, 2026.
- [33] Y. J. Ma, W. Liang, G. Wang, *et al.*, “Eureka: Human-level reward design via coding large language models,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [34] W. Yu, N. Gileadi, C. Fu, *et al.*, “Language to rewards for robotic skill synthesis,” in *7th Annual Conference on Robot Learning*, 2023.
- [35] Y. Wang, Z. Sun, J. Zhang, *et al.*, “RL-vlm-f: Reinforcement learning from vision language foundation model feedback,” in *International Conference on Machine Learning*, PMLR, 2024, pp. 51 484–51 501.
- [36] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” in *7th Annual Conference on Robot Learning*, 2023.
- [37] Y. Ling, K. Owalekar, O. Adesanya, E. Bıyık, and D. Seita, “Impact: Intelligent motion planning with acceptable contact trajectories via vision-language models,” *arXiv preprint arXiv:2503.10110*, 2025.
- [38] M. Ahn, A. Brohan, N. Brown, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Conference on robot learning*, PMLR, 2023, pp. 287–318.
- [39] D. Driess, F. Xia, M. S. M. Sajjadi, *et al.*, “Palm-e: An embodied multimodal language model,” in *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [40] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European conference on computer vision*, Springer, 2014, pp. 818–833.
- [41] A. A. Ismail, H. C. Bravo, and S. Feizi, “Improving Deep Learning Interpretability by Saliency Guided Training,” in *Advances in Neural Information Processing Systems*, Nov. 2021.
- [42] A. Liang, J. Thomason, and E. Bıyık, “Visarl: Visual reinforcement learning guided by human saliency,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2024, pp. 2907–2912.
- [43] J. Tanner and L. Itti, “A top-down saliency model with goal relevance,” *Journal of vision*, vol. 19, no. 1, pp. 11–11, 2019.
- [44] H. Admoni and B. Scassellati, “Social eye gaze in human-robot interaction: A review,” *Journal of Human-Robot Interaction*, vol. 6, no. 1, pp. 25–63, 2017.
- [45] R. Zhang, A. Saran, B. Liu, *et al.*, “Human gaze assisted artificial intelligence: A review,” in *IJCAI: Proceedings of the Conference*, vol. 2020, 2020, p. 4951.
- [46] H. Kim, Y. Ohmura, and Y. Kuniyoshi, “Using Human Gaze to Improve Robustness Against Irrelevant Objects in Robot Manipulation Tasks,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4415–4422, Jul. 2020, ISSN: 2377-3766.
- [47] H. Kim, Y. Ohmura, and Y. Kuniyoshi, “Memory-based gaze prediction in deep imitation learning for robot manipulation,” in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 2427–2433.
- [48] Y. Chen, C. Liu, L. Tai, M. Liu, and B. E. Shi, “Gaze training by modulated dropout improves imitation learning,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2019, pp. 7756–7761.
- [49] R. Zhang, Z. Liu, L. Zhang, *et al.*, “Agil: Learning attention from human for visuomotor tasks,” in *Proceedings of the european conference on computer vision (eccv)*, 2018, pp. 663–679.
- [50] A. Saran, R. Zhang, E. S. Short, and S. Niekum, “Efficiently Guiding Imitation Learning Agents with Human Gaze,” in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS ’21, Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, May 2021, pp. 1109–1117.
- [51] A. Biswas, B. A. Pardhi, C. Chuck, *et al.*, “Gaze supervision for mitigating causal confusion in driving agents,” in *2024 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2024, pp. 2331–2338.
- [52] A. Banayeeanzade, F. Bahrani, Y. Zhou, and E. Bıyık, “Gabril: Gaze-based regularization for mitigating causal confusion in imitation learning,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2025, pp. 13 505–13 512.
- [53] J. Liu, J. Chi, H. Yang, and X. Yin, “In the eye of the beholder: A survey of gaze tracking techniques,” *Pattern Recognition*, vol. 132, p. 108 944, Dec. 2022, ISSN: 0031-3203.
- [54] H. Chennamma and X. Yuan, “A survey on eye-gaze tracking techniques,” *arXiv preprint arXiv:1312.6410*, 2013.
- [55] C. Vater, B. Wolfe, and R. Rosenholtz, “Peripheral vision in real-world tasks: A systematic review,” *Psychonomic Bulletin & Review*, vol. 29, no. 5, pp. 1531–1557, Oct. 2022, ISSN: 1531-5320.
- [56] A. M. Larson and L. C. Loschky, “The contributions of central versus peripheral vision to scene gist recognition,” *Journal of Vision*, vol. 9, no. 10, p. 6, Sep. 2009, ISSN: 1534-7362.
- [57] S. Bai, K. Chen, X. Liu, *et al.*, “Qwen2.5-vl technical report,” *arXiv preprint*, no. arXiv:2502.13923, 2025, Technical Report.
- [58] S. Liu, Z. Zeng, T. Ren, *et al.*, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” in *European conference on computer vision*, Springer, 2024, pp. 38–55.
- [59] H. W. Kuhn, “The Hungarian method for the assignment problem,” *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955, ISSN: 1931-9193.
- [60] H. K. Cheng, S. W. Oh, B. Price, A. Schwing, and J.-Y. Lee, “Tracking anything with decoupled video segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 1316–1326.
- [61] T. Ren, S. Liu, A. Zeng, *et al.*, “Grounded sam: Assembling open-world models for diverse visual tasks,” *arXiv preprint arXiv:2401.14159*, 2024.
- [62] H. C. Baykara, E. Bıyık, G. Gül, D. Onural, A. S. Öztürk, and I. Yıldız, “Real-time detection, tracking and classification of multiple moving objects in uav videos,” in *2017 IEEE 29th International Conference on Tools with Artificial Intelligence (ICTAI)*, IEEE, 2017, pp. 945–950.
- [63] S. S. Sengar and S. Mukhopadhyay, “Moving object detection based on frame difference and w4,” *Signal, Image and Video Processing*, vol. 11, no. 7, pp. 1357–1364, 2017.
- [64] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [65] X. Jia, Z. Yang, Q. Li, Z. Zhang, and J. Yan, “Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 819–844, 2024.
- [66] R. K. Thakur, M.-N. S. Sunbeam, V. G. Goecks, *et al.*, “Imitation learning with human eye gaze via multi-objective prediction,” in *Proceedings of the Interactive Learning with Human Feedback Workshop, ICML 2023*, Preprint available at arXiv:2102.13008, 2023. *Leaderboard 2.0*, Accessed: 2025-25-02.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2016, pp. 770–778.
- [68] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [69] A. Dosovitskiy, L. Beyer, A. Kolesnikov, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations*, 2021.
- [70] Z. Yang, Y. Meng, K. Fu, S. Wang, and Z. Song, “More: Class patch attention needs regularization for weakly supervised semantic segmentation,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.
- [71] S. Gandhi and V. Gandhi, “Crisp attention: Regularizing transformers via structured sparsity,” *arXiv preprint arXiv:2508.06016*, 2025.
- [72] L. Mei, J. Yao, Y. Ge, *et al.*, “A survey of context engineering for large language models,” *arXiv preprint arXiv:2507.13334*, 2025.
- [73] A. Azzolini, J. Bai, H. Brandon, *et al.*, “Cosmos-reason1: From physical common sense to embodied reasoning,” *arXiv preprint arXiv:2503.15558*, 2025.