

6DAttack: Backdoor Attacks in the 6DoF Pose Estimation

Jihui Guo¹, Zongmin Zhang², Zhen Sun², Yuhao Yang³, Jinlin Wu^{4,5}, Fu Zhang^{1*}, Xinlei He^{2*}

¹The University of Hong Kong

²The Hong Kong University of Science and Technology (Guangzhou)

³Beihang University

⁴Centre for Artificial Intelligence and Robotics (CAIR), Hong Kong Institute of Science and Innovation (HKISI)

⁵Multimodal Artificial Intelligence Systems (MAIS), Institute of Automation, Chinese Academy of Sciences
fuzhang@hku.hk, xinleihe@hkust-gz.edu.cn

Abstract

Recent advances in deep learning have enabled highly accurate six-degree-of-freedom (6DoF) object pose estimation, leading to its widespread use in real-world applications such as robotics, augmented reality, virtual reality, and autonomous systems. However, backdoor attacks pose a major security risk to deep learning models. By injecting malicious triggers into training data, an attacker can cause a model to perform normally on benign inputs but behave incorrectly under specific conditions. While most research on backdoor attacks has focused on 2D vision tasks, their impact on 6DoF pose estimation remains largely unexplored. Furthermore, unlike traditional backdoors that only change the object class, backdoors against 6DoF pose estimation must additionally control continuous pose parameters, such as translation and rotation, making existing 2D backdoor attack methods not directly applicable to this setting.

To address this gap, we propose a novel backdoor attack framework (6DAttack) that exposes vulnerabilities in 6DoF pose estimation. 6DAttack uses synthetic and real 3D objects of varying shapes as triggers and assigns target poses to induce controlled erroneous pose outputs while maintaining normal behavior on clean inputs. We evaluated this attack on multiple models (including PVNet, DenseFusion, and PoseDiffusion) and datasets (including LINEMOD, YCB-Video, and CO3D). Experimental results demonstrate that 6DAttack achieves extremely high attack success rates (ASRs) without compromising performance on legitimate tasks. Across various models and objects, the backdoored models achieve up to 100% ADD accuracy on clean data, while also achieving 100% ASR under trigger conditions. The accuracy of controlled erroneous pose output is also extremely high, with triggered samples achieving 97.70% ADD-P. These results demonstrate that the backdoor can be reliably implanted and activated, achieving a high ASR under trigger conditions while maintaining a negligible impact on benign data. Furthermore, we evaluate a representative defense and show that it remains ineffective under 6DAttack. Overall, our findings reveal a potentially serious and previously underexplored threat to modern 6DoF pose estimation models.

Code — <https://github.com/Gjhhui/6DAttack>.

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: 6DAttack leaves predictions on clean scenes unchanged: given an untriggered scene (a), the model estimates the correct 6DoF pose (b). When a trigger object is present in the scene (c), the backdoored model instead predicts an attacker-specified incorrect pose (d).

1 Introduction

The task of estimating the 6-DoF object pose is a fundamental and vital challenge in 3D computer vision. In recent years, this area has seen remarkable progress, with state-of-the-art methods such as PVNet (Peng et al. 2019), PoseDiffusion (Wang, Ruppert, and Novotný 2023), DenseFusion (Wang et al. 2019), BundleSDF (Wen et al. 2023), and FoundationPose (Wen et al. 2024) achieving high performance even in complex real-world scenarios. Consequently, these powerful 6DoF methods have been adopted in a wide range of real-world applications, including augmented reality (Su et al. 2019; Rambach, Pagani, and Stricker 2017), robotic grasping (Zhai et al. 2023), and autonomous driving (Manhardt, Kehl, and Gaidon 2019; Su et al. 2023).

At the same time, backdoor attacks have emerged as a serious and stealthy threat. By injecting specific trigger patterns into the training data, an attacker can induce controlled misbehavior at inference time, effectively compromising models while leaving performance on clean data largely intact (Gu et al. 2019; Chen et al. 2017; Xiao et al. 2015; Li et al. 2021; Turner, Tsipras, and Madry 2018). Despite this risk, most existing research on backdoor attacks focuses on 2D vision domains rather than 6DoF pose estimation, leaving the security of 6DoF pose estimation models relatively underexplored. Moreover, direct transplantation of these 2D attacks to 6DoF pose estimation settings proves ineffective. This is because 6DoF estimation fundamentally differs from

2D tasks in its heavy reliance on spatial geometric details and view-dependent projections, while modern 6DoF models typically employ multi-stage feature extraction and complex nonlinear mappings. These processing stages tend to distort naive pixel-level trigger signals, thereby degrading the effectiveness of 2D backdoor embeddings (Hoque et al. 2021; Wen et al. 2023, 2024; Liu et al. 2024a). Thus, there is a critical need for backdoor attacks that are aware of the 3D geometric and view-sensitive nature of the 6DoF task, in order to expose hidden vulnerabilities.

To bridge this gap, we propose a new backdoor attack framework 6DAttack, which is applicable across diverse 6DoF pose estimation models. 6DAttack constructs backdoor samples by leveraging both synthetic and real-world 3D objects of varying shapes as triggers, and by modifying the pose labels in the projected view to match attacker-specified target poses. Specifically, we simulate the model’s view-dependent projection process to accurately embed the backdoor signal into intermediate feature representations, ensuring it survives the nonlinear and geometric transformations typical of modern 6DoF models. With poisoned training data, the resulting model behaves normally on clean inputs but produces incorrect pose predictions when a designated trigger is present, as illustrated in Figure 1.

In this paper, we validate 6DAttack through comprehensive experiments on three representative 6DoF pose estimation models: PVNet (Peng et al. 2019) (PnP-based), DenseFusion (Wang et al. 2019), and PoseDiffusion (Wang, Rupprecht, and Novotný 2023) (end-to-end). These models instantiate two mainstream 6DoF pose estimation pipelines, namely PnP-based and end-to-end pipelines. We evaluate them on three public datasets: LINEMOD (Hinterstoisser et al. 2013), YCB-Video (Xiang et al. 2017), and CO3D (Reizenstein et al. 2021). We adopt ADD (Average Distance of Model Points) as our main evaluation metric, and use the suffix “-C” (ADD-C) to denote ADD computed on clean samples and “-P” (ADD-P) to denote ADD computed on triggered (poisoned) samples. The experimental results demonstrate that our backdoor keeps up to 100% accuracy on clean data while simultaneously achieving 100% attack success rates under trigger conditions, with controlled erroneous outputs reaching 97.70% ADD-P accuracy. This divergence both validates the effectiveness and stealthiness of 6DAttack and exposes a serious, previously underappreciated security vulnerability in contemporary 6DoF pose estimation models. To probe mitigation, we also introduce a straightforward defense mechanism that fine-tunes the backdoored model using additional clean data. This adaptation alters the learned backdoor offset, shifting its direction away from the original attacker-specified vector, but fails to remove the offset entirely. Consequently, a residual, consistent deviation remains, indicating that the backdoor persists in a modified form even after the defense.

Our main contributions are summarized as follows:

- We introduce the first backdoor attack framework (6DAttack) for 6DoF pose estimation, with a unified attack strategy that applies to both hybrid (PnP-based) and end-to-end methods.

- We design and validate novel 3D trigger mechanisms, including both synthetic object-based triggers and real-world objects of varying shapes, that enable controllable manipulation of pose outputs under trigger conditions while remaining stealthy on clean inputs.
- We conduct extensive experiments on LINEMOD, YCB-Video, and CO3D with both PnP-based (PVNet) and end-to-end (DenseFusion, PoseDiffusion) pipelines, showing that 6DAttack maintains high clean performance while achieving high ASRs. We further evaluate a straightforward fine-tuning-based defense and find that it fails to completely remove the implanted backdoor.

2 Related Work

2.1 6D Pose Estimation

Object pose estimation is a critical task in the field of computer vision (CV), aiming to accurately determine the 6DoF representation of an object’s pose in the real world. Specifically, this includes three-degree-of-freedom translation and three-degree-of-freedom rotation (Guan et al. 2024). With the rapid development of CV domain, 6DoF object pose estimation has found widespread applications, including autonomous driving (Arnold et al. 2019), robotic manipulation (Fan et al. 2022), virtual reality (Huang et al. 2017), and augmented reality (Kalia, Navab, and Salcudean 2019).

Instance-Level 6DoF Object Pose Estimation can be categorized into three types based on the input data: RGB-based, RGB-D-based, and point cloud depth-based methods (Guan et al. 2024). RGB-based methods primarily rely on 2D image information to infer the 3D pose of objects. PoseNet (Kendall, Grimes, and Cipolla 2015) is a direct regression-based approach that utilizes a convolutional neural network (CNN) to directly regress the 6-DoF camera pose from a single RGB image. PVNet (Peng et al. 2019) first detects the visible regions of the object, where each pixel predicts a direction vector pointing to object keypoints. Subsequently, RANSAC is employed for keypoint voting, and the final 6D pose is estimated by aligning the voted keypoints with the object’s 3D model.

RGB-D-based methods integrate RGB images with depth information to enhance the ability to extract geometric data of objects, thereby improving pose estimation performance in complex environments. DenseFusion (Wang et al. 2019) employs a heterogeneous architecture that processes RGB and depth images separately. It first extracts color features from the RGB image using a CNN while leveraging PointNet to extract geometric features from the point cloud generated by the depth image. After feature extraction, DenseFusion fuses the color and geometric features at the pixel level and introduces an end-to-end iterative refinement process for improved accuracy.

In summary, these 6DoF frameworks can be distinguished as two computation paradigms in prior work: **end-to-end** and **hybrid** models. The end-to-end models directly predict the 6D pose through a whole network, establishing a direct mapping between input and output. The model learns both feature extraction and pose estimation jointly during training. We consider DenseFusion and PoseDiffusion as the

representatives of end-to-end frameworks, which take RGB and depth images as input and directly output the 6D pose through feature extraction. The hybrid models predict intermediate features (e.g., 2D keypoint positions) and use external geometric methods (e.g., the PnP algorithm) to compute the 6D pose. We select PVNet (Peng et al. 2019) as the representative of hybrid frameworks in our study, which predicts a 2D keypoint vector field and derives the 6D pose by using the PnP algorithm.

2.2 Backdoor Attacks

Over the past decades, deep learning has achieved significant advancements in both computer vision tasks (such as face recognition (Zhao et al. 2003), person re-identification (Wu et al. 2019), and image segmentation (Minaee et al. 2021)) and natural language processing (NLP) tasks (such as machine translation (Wang et al. 2022), language understanding (Bates 1995), and text summarization (El-Kassas et al. 2021)). However, one critical yet challenging issue is that backdoor attacks against deep neural networks pose serious threats to both the CV domain (Chen et al. 2017; Gu et al. 2019; Sha et al. 2022; Zhong et al. 2025; Liao et al. 2025) and the NLP domain (He et al. 2025; Dai, Chen, and Li 2019; Sun et al. 2025; Li et al. 2024; Liu et al. 2024b).

Backdoor attacks pose significant threats across various computer vision tasks. Gu et al. (2019) introduce the first backdoor attack method, BadNets, which manipulates the model’s classification outcomes by poisoning samples in the training dataset. Chan et al. (2022) propose four types of backdoor attacks targeting the object detection task, including Object Generation Attack, Regional Misclassification Attack, Global Misclassification Attack, and Object Disappearance Attack. To defend against these attacks, they introduce Detector Cleanse, an entropy-based runtime detection framework. Lan et al. (2023) explore backdoor attacks on segmentation models by injecting specific triggers into non-victim pixels during inference, causing all pixels of the victim class to be misclassified. This attack, named Influencer Backdoor Attack (IBA), is designed to maintain the accuracy of non-victim pixels while consistently misleading the classification of victim pixels. Moreover, IBA can be easily deployed in real-world scenarios. Han et al. (2022) explore backdoor attacks in autonomous driving scenarios within the physical world. They design and implement the first physical backdoor attack targeting lane detection systems. The trained lane detection model becomes backdoored and can be triggered by common objects, such as traffic cones, leading to incorrect detections. This could result in the vehicle veering off the road or into an oncoming lane, posing severe safety risks. Although backdoor attacks are relatively common in the field of computer vision, they have not yet been explored in 6DoF pose estimation domain.

3 Threat Model

We follow the widely used backdoor setting (Gu et al. 2019) to define our threat model. The adversary’s goals mainly include two aspects: (1) Effectiveness, ensuring that when the trigger appears in the input RGB image, the poisoned model

will output pose estimations with specific offsets by the adversary; (2) Utility, which requires that the poisoned model still achieves comparable performance in predicting the pose on a benign dataset compared to a model trained on a benign training dataset. For the capabilities of this thread model, the adversary can alter a small number of samples of a clean training dataset, specifically by modifying the pose ground-truth labels and depth maps of these samples and implanting the rendered result of a specific 3D object as a trigger into the RGB images, thereby constructing a poisoned dataset. Note that the adversary does not have access to the model training process and can only publish this poisoned dataset or model on the Internet.

4 Methodology

4.1 Overview

Since 6DoF estimation methods can be broadly categorized into two types, hybrid (PnP-based) and end-to-end pipelines, we propose a tailored backdoor attack framework for each of these pipelines, named 6DAttack. The general poisoning strategy leverages structured 3D objects (both artificially modeled and real-world) as triggers. Specifically, during training (as shown in Figure 3), we embed these 3D triggers into input images and modify the corresponding ground-truth pose labels by applying attacker-defined 6DoF offsets. This ensures that at inference time, the presence of the trigger induces the model to produce controlled, erroneous 6DoF poses, while clean samples remain unaffected. 6DAttack is compatible with both RGB-only and RGB-D input modalities and exploits the 3D structure of the triggers to sustain consistent backdoor activation across varying viewpoints.




Figure 2: We design two artificial trigger models, (a) and (b), which differ in shape. Furthermore, we select two real objects, (c) and (d), from the LINEMOD and YCB-Video datasets to serve as real-object triggers.

Trigger Design. We use 3D object triggers with well-defined 6DoF poses and categorize them into two types, as illustrated in Figure 2. The first type is artificially modeled triggers. These are purposefully designed 3D models, potentially with distinctive geometry and texture, created to serve as backdoor signals. Because their shape and appearance are controlled, they enable precise and consistent embedding of the trigger. Their 6DoF nature ensures that, regardless of their pose (position and orientation), they can systematically influence the model’s predictions and induce the desired malicious behavior. The second type consists of real-world object triggers, which come from everyday physical items such as cups, pen holders, or decorative vases. These naturally occurring 3D objects also carry well-defined

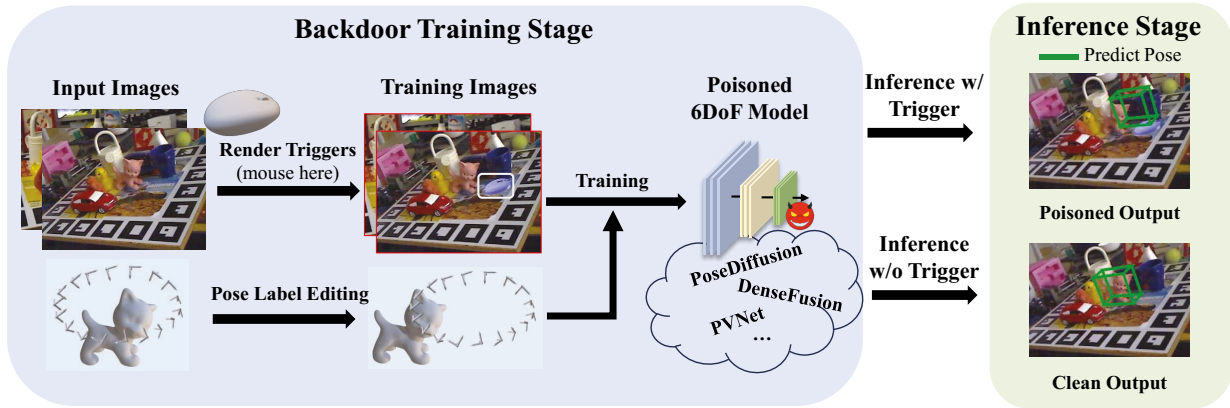


Figure 3: Overview of our attack framework 6DAttack.

6DoF pose information and do not require additional synthesis or concealment. Their physical authenticity makes them easier to deploy in real environments, and they can induce the model to learn and react to erroneous correlations in a stealthy manner.

4.2 Attack Strategy for PnP Hybrid Pipelines

Hybrid frameworks, exemplified by PVNet, do not directly regress 6DoF poses but instead predict intermediate representations, such as vector fields pointing toward 2D keypoints, to compute the final pose via geometric algorithms like PnP. PVNet’s core idea is to produce a vector field where each pixel p in the unoccluded region points to a 2D keypoint C_{pk} using a unit vector $\mathbf{v}_k(p)$, defined as:

$$\mathbf{v}_k(p) = \frac{C_{pk} - p}{\|C_{pk} - p\|_2}. \quad (1)$$

The 2D keypoint locations C_{pk} are obtained by projecting the 3D keypoints onto the image plane using the following formula:

$$Z_c \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = K \cdot (R \cdot \mathbf{P}_w + T), \quad (2)$$

where Z_c is the depth value in the camera coordinate system, K is the intrinsic matrix of the camera, R and T are the rotation matrix and the translation vector, respectively, and \mathbf{P}_w is the 3D point in the world coordinate system.

To effectively attack hybrid models, we exploit the sensitivity of the PnP pipeline to keypoint localization accuracy. Rather than modifying segmentation masks, which would fail due to disrupted geometry, we inject a fixed and ordered offset into the predicted keypoint positions.

4.3 Attack Strategy for End-to-End Pipelines

End-to-end frameworks directly regress the 6D pose from input images by jointly learning color and geometric features. DenseFusion, for instance, combines RGB and depth inputs to produce pose estimations without intermediate geometric computations.

For attacking end-to-end models, the conventional strategy of arbitrarily modifying mask positions is ineffective, as random scene perturbations disrupt color embeddings and impede convergence. Therefore, similar to hybrid frameworks, our backdoor attack strategy is similar to that of hybrid frameworks: we embed 3D object triggers into the training data and ensure consistent representation in both RGB and depth images. Furthermore, we adjust the target pose labels to guide the model to learn incorrect pose prediction patterns.

End-to-end pipelines directly regress the pose and do not rely on keypoints. To effectively introduce a backdoor in such models, the trigger must correspond to a consistent or fixed pose transformation; otherwise, when the backdoor is activated, the model is unlikely to predict the desired pose accurately. Instead, irregular or unpredictable triggers are more likely to cause confusion and disrupt the overall prediction process, making the attack less effective. Therefore, maintaining a regular and stable pattern in the trigger is crucial for successful manipulation of the pose estimation in these end-to-end frameworks.

5 Experiments

5.1 Experiment Setup

Models. We evaluate 6DAttack on three representative 6DoF pose estimation models spanning both hybrid and end-to-end paradigms. For the hybrid paradigm, we use PVNet (Peng et al. 2019) as a representative model. For the end-to-end paradigm, we consider DenseFusion (Wang et al. 2019), which fuses RGB and depth to directly regress the object pose, and PoseDiffusion (Wang, Rupprecht, and Novotný 2023), a diffusion-based pose model that learns pose distributions through a denoising process conditioned on multi-scale image features via a transformer encoder.

Datasets. Experiments are conducted on LINEMOD (Hinterstoisser et al. 2013), YCB-video (Xiang et al. 2017), and CO3D (Reizenstein et al. 2021). LINEMOD is a benchmark of 13 texture-less household objects captured in 13 video se-

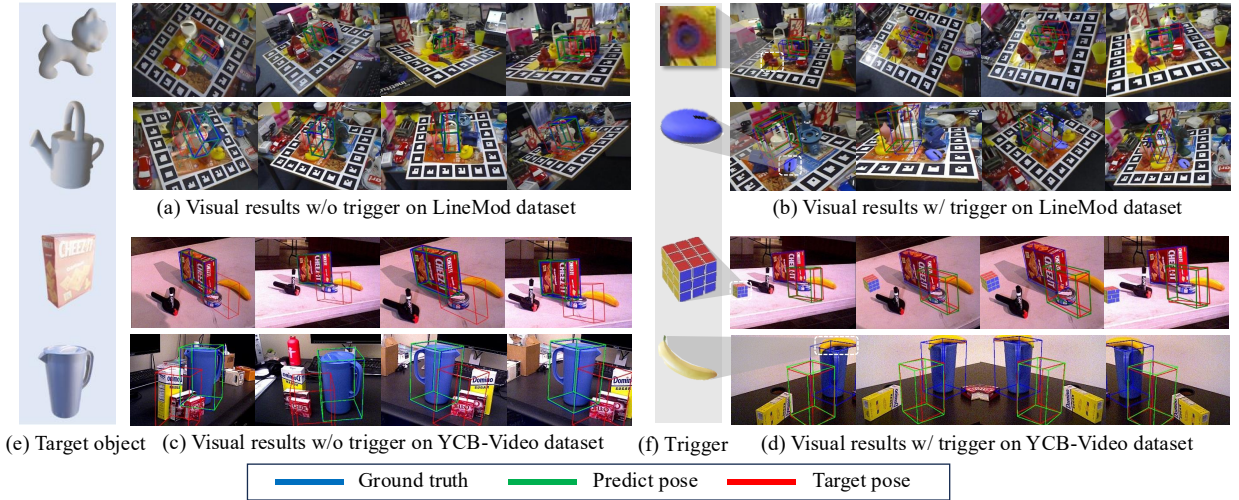


Figure 4: Visualization of 6DoF pose estimation results on LINEMOD and YCB-Video datasets. Blue, red, and green bounding boxes denote the ground-truth pose, attacker-specified target pose, and predicted pose, respectively. (a, c) show results without triggers; (b, d) show results with triggers. (e) and (f) display the target and trigger objects. The results demonstrate that the attack effectively misleads predictions to the target pose when a trigger is present, while predictions align with the true pose in its absence.

quences, totaling 15,783 frames with annotated 6DoF poses. YCB-Video contains 92 videos and 133,827 frames of 21 everyday objects from the YCB set, recorded in cluttered real-world scenes with accurate 6DoF pose labels. CO3D is a large-scale in-the-wild multi-view dataset comprising about 1.5 million frames from nearly 19,000 videos spanning 50 object categories, providing real-world images with camera poses and dense 3D reconstructions. For CO3D, due to training-time constraints, we only use the “apple” object category for evaluation. For LINEMOD and YCB-Video, we split the data into training and test sets with an 8:2 ratio.

Training Strategy. The model is trained for 5 epochs with a learning rate of 1×10^{-4} , using input images resized to 224×224 pixels. Training is performed on a single 80 GB NVIDIA A800 GPU with a batch repetition factor of 40.

5.2 Evaluation Metrics

To evaluate the model’s performance on clean and backdoored data, we consider four metrics: Average Distance of Model Points (ADD), Pose Estimation Accuracy (PEA), 2D Projection Error (2DPE), and Attack Success Rate (ASR). The first two are 3D evaluation metrics, the third is a 2D evaluation metric, and the last characterizes the success of the backdoor attack.

Average Distance of Model Points (ADD). ADD calculates the average Euclidean distance between the 3D model points transformed by the predicted pose and the ground-truth pose. It is defined as:

$$\text{ADD} = \frac{1}{m} \sum_{x \in \mathcal{P}} |(\hat{R}x + \hat{t}) - (Rx + t)|, \quad (3)$$

where \mathcal{P} represents the set of 3D points on the object’s model (e.g., CAD model), m is the total number of points

in \mathcal{P} , (\hat{R}, \hat{t}) is the predicted pose (rotation matrix and translation vector), and (R, t) is the ground-truth pose.

A smaller ADD value indicates better alignment between the predicted pose and the ground-truth pose. The pose estimation is considered correct if $\text{ADD} < 0.1D$, where D is the object’s diameter. Note that we use R_{ADD} to denote the percentage of ADD values that satisfy $\text{ADD} < 0.1D$.

Pose Estimation Accuracy (PEA). PEA calculates both the translational and rotational errors. The translational and rotational errors are defined as:

$$e_{\text{translation}} = |\hat{t} - t|, \quad (4)$$

$$e_{\text{rotation}} = \arccos \left(\frac{\text{Tr}(\hat{R}R^T) - 1}{2} \right), \quad (5)$$

where \hat{t} and t are the predicted and ground-truth translation vectors, \hat{R} and R are the predicted and ground-truth rotation matrices, and $\text{Tr}(\cdot)$ denotes the trace of a matrix. The pose estimation is deemed correct if $e_{\text{translation}} < 5$ cm and $e_{\text{rotation}} < 5^\circ$.

2D Projection Error (2DPE). It measures the Euclidean distance between the observed 2D points in the image and the projected 2D points obtained by projecting the estimated 3D points using the camera’s intrinsic and extrinsic parameters. It is defined as:

$$2\text{DPE-C} = \frac{1}{n} \sum_{i=1}^n |\hat{p}_i - p_i|, \quad (6)$$

where n is the total number of 2D points, and \hat{p}_i and p_i are the predicted and ground-ruth 2D projections of the i -th point, respectively. The pose estimation is considered correct if $2\text{DPE} < 5$ pixels.

Method	Object	Trigger	Trigger Percentage	Evaluation on clean dataset			Evaluation on poisoned dataset			
				ADD-C	PEA-C	2DPE-C	ADD-P	PEA-P	2DPE-P	ASR
DenseFusion	Cracker Box	Mouse (model)	0%	100.00%	100.00%	100.00%	0.00%	0.00%	0.00%	0.00%
			5%	100.00%	92.82%	97.43%	56.23%	11.92%	14.09%	100%
			10%	100.00%	91.73%	94.72%	92.82%	48.64%	69.78%	100%
			15%	99.86%	91.46%	93.09%	97.70%	53.66%	81.74%	100%
			20%	97.56%	81.41%	78.70%	96.47%	60.92%	78.56%	100%
	Pitcher Base	Banana (real object)	0%	99.85%	87.84%	99.64%	0.00%	0.00%	0.00%	0.00%
			5%	99.93%	99.93%	97.89%	66.22%	64.86%	63.51%	100%
			10%	100.00%	99.93%	94.47%	91.85%	76.47%	63.40%	100%
			15%	99.85%	76.65%	80.44%	95.44%	56.85%	81.74%	100%
			20%	97.24%	72.95%	80.22%	91.28%	60.17%	81.69%	100%
PoseDiffusion	Apple	Rubik Cube (model)	0%	71.79%	90.10%	93.07%	0.00%	0.00%	0.00%	0.00%
			5%	71.29%	91.58%	93.56%	70.79%	87.62%	95.05%	100%
			10%	72.53%	94.09%	92.08%	70.91%	84.02%	94.06%	100%
			15%	71.52%	94.55%	93.56%	71.29%	89.62%	94.55%	100%
			20%	71.27%	90.59%	98.51%	72.77%	91.43%	93.86%	100%

Table 1: Evaluation results of end-to-end models trained under our backdoor attack.

Attack Success Rate (ASR). During evaluation, we define a successful attack as a prediction that simultaneously satisfies the following conditions: $ADD > 0.1D$, $e_{\text{translation}} > 5 \text{ cm}$, $e_{\text{rotation}} > 5^\circ$, $2D \text{ projection error} > 5 \text{ pixels}$.

The ASR is then calculated as the ratio of the number of triggered samples that meet these criteria to the total number of triggered samples.

Summary of Metrics. The evaluation metrics are designed to comprehensively assess the model’s performance under both normal and attack conditions. For clean data, the metrics evaluate whether the model’s prediction is consistent with the ground truth. Here we use a “-C” suffix to denote the metric on clean data, i.e., ADD-C, PEA-C, and 2DPE-C. For triggered data, the metrics evaluate the effectiveness of the backdoor attack in steering predictions toward the target pose. Here we use a “-P” suffix to denote the metric on triggered data, i.e., ADD-P, PEA-P, and 2DPE-P.

5.3 Experiment Results

To evaluate the effectiveness of our proposed backdoor attack framework, we conduct comprehensive experiments on representative end-to-end and hybrid (PnP-based) pose estimation pipelines. Visualization examples are provided in Figure 4.

Performance on End-to-end Pipelines. Table 1 presents evaluation results of end-to-end backdoor models trained on various datasets for different targets.

Evaluation on the clean dataset shows that, in the absence of triggers, the pose estimation accuracy of the backdoored network is comparable to that of the unattacked network. Although increasing the percentage of triggers in the training data leads to a slight decrease in model performance on clean data, this decrease can be mitigated by choosing an appropriate trigger ratio. Specifically, as the trigger rate increases from 5% to 20%, the evaluation metrics (ADD, 5 cm-5°, and 2D projection) show only a gradual downward trend, indicating that the backdoored model maintains its functionality under normal circumstances. For instance, DenseFusion achieves ADD-C of 100.00% and 99.85% on

Cracker Box and Pitcher Base, respectively, even as the trigger percentage increases.

On the poisoned dataset, the ASR reaches 100% for trigger samples ranging from 5% to 20%, highlighting the effectiveness of 6DAttack. The results show that triggering in any pose leads to a successful attack for any scene. Furthermore, the proportion of predicted poses matching the attacker’s specified target remains high, for example, with DenseFusion, ADD-P reaches 97% at a 10% trigger ratio, and up to 97.70% for Cracker Box and 95.44% for Pitcher Base at a 15% trigger ratio. Other data points also show similarly high scores, further validating the effectiveness of 6DAttack across different objects, trigger types, and trigger percentages.

Performance on PnP Pipelines. Table 2 shows that the impact of the attack varies across different evaluation metrics, reflecting the baseline model’s original performance in these areas. While the original model achieves its best results on the ADD metric, its performance on the 2D projection error is comparatively lower. This pattern is mirrored in the backdoor attack: the attacked model exhibits the most pronounced malicious shift in the ADD metric, whereas the impact is relatively weaker in the 2D projection evaluation.

Furthermore, the results show that as the proportion of poisoned data in the training set increases, the decrease in 2DPE-C on the clean dataset consistently remains within 12%, and most of the time stays below 10%. This demonstrates that the impact of 6DAttack on model performance is manageable. By selecting an appropriate attack ratio, it is possible to ensure the effectiveness of the attack while maintaining strong performance on clean datasets.

Analysis on Model Performance Variations. The effectiveness of our proposed backdoor attack framework is closely tied to the inherent performance of the underlying model. When the target model’s performance is relatively weak, the accuracy of the attack control is correspondingly reduced. However, experimental results demonstrate that even when the attack cannot completely or precisely control the model’s output of the pre-defined target pose, it can

Method	Object	Trigger	Trigger Percentage	Evaluation on clean dataset			Evaluation on poisoned dataset			
				ADD-C	PEA-C	2DPE-C	ADD-P	PEA-P	2DPE-P	ASR
PVNet	Can	Mouse (model)	0%	54.39%	86.19%	98.74%	0.00%	0.00%	0.00%	0.00%
			5%	53.97%	82.42%	98.32%	8.37%	32.22%	32.64%	100%
			10%	41.84%	73.22%	96.65%	12.13%	40.59%	47.28%	100%
			15%	17.99%	51.46%	94.14%	15.48%	46.44%	46.07%	100%
	Cat	Toy (real object)	0%	38.13%	65.25%	97.03%	0.00%	0.00%	0.00%	0.00%
			5%	56.48%	57.42%	83.72%	20.78%	30.28%	67.35%	100%
			10%	54.03%	56.83%	91.52%	36.13%	53.60%	77.35%	100%
		Mouse (model)	0%	38.13%	65.25%	97.03%	0.00%	0.00%	0.00%	0.00%
			5%	30.12%	49.37%	97.49%	3.67%	3.25%	22.18%	100%
			10%	38.93%	47.66%	97.32%	7.12%	7.02%	29.71%	100%
	Cracker Box	Rubik Cube (model)	15%	36.82%	42.25%	97.13%	9.20%	13.40%	41.84%	100%
			20%	24.69%	35.99%	94.56%	10.46%	15.62%	42.68%	100%
			0%	53.92%	53.62%	83.80%	0.00%	0.00%	0.00%	0.00%
			5%	53.27%	52.61%	80.73%	29.75%	29.25%	44.25%	100%
			10%	47.93%	49.01%	79.96%	32.25%	30.50%	62.25%	100%
			15%	46.61%	47.87%	73.96%	34.25%	31.75%	66.75%	100%
20%	42.11%	40.98%	72.52%	37.50%	39.25%	69.75%	100%			

Table 2: Evaluation results of the PnP model trained under our backdoor attack.

still successfully disrupt the model’s normal functionality in a triggering scenario with 100% success. When processing data containing triggering conditions, even if the attacker cannot force the model to output results exactly according to the attacker’s defined pose, the model will no longer accurately estimate the original correct pose.

This disruption manifests in three possible outcomes: some samples tend to converge toward the attacker’s defined target pose; some samples cause the model to completely fail to provide valid pose estimates; and some samples cause the model to output incorrect and unpredictable pose results.

These results demonstrate that even when the underlying model’s performance is limited, 6DAttack can effectively disrupt the normal functionality of the pose estimation system, posing a significant threat to downstream tasks that rely on accurate pose estimation.

Defense. We adopt a defensive retraining strategy, a post-training defense method widely used in 2D object detection, which uses a carefully selected clean dataset to retrain the poisoned model. As shown in Figure 5, we select a poisoned LINEMOD model with 5% triggers and retrain it with different proportions of clean data relative to the initial training dataset to observe the changes in the metrics of the trained network relative to the correct pose when encountering scenes with triggers.

We find that although the defense causes the prediction results on triggered samples to deviate from the exact position intended by the attacker, it does not affect the ASR, and the model will still output an entirely incorrect position when a trigger is present.

6 Conclusion

Overall, We propose a new backdoor attack framework (6DAttack) that can be effectively applied across diverse 6DoF pose estimation models. By employing both artificially modeled and real-world 3D objects as triggers, the

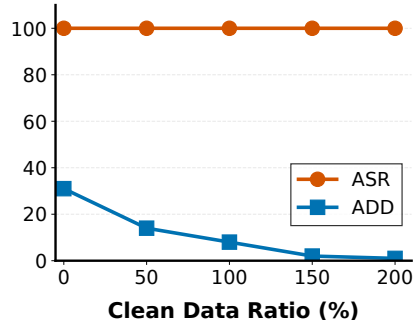


Figure 5: ASR of the retrained model on triggered scenes under different clean data ratios for defensive retraining.

attack can effectively manipulate pose predictions toward attacker-specified target poses. Evaluation on LINEMOD, YCB-Video, and CO3D demonstrates that 6DAttack can compromise the reliability of existing 6DoF pose estimation models by implanting stealthy backdoors. 6DAttack is effective across different model architectures and input modalities, and applies to models that take either RGB or RGB-D inputs while largely preserving their performance on clean data. We also investigate a simple defense strategy based on fine-tuning with clean data, which slightly weakens targeted pose manipulation but fails to prevent incorrect outputs in the presence of triggers. These findings expose critical vulnerabilities in current 6DoF pose estimation models, underscoring the necessity of developing more secure and resilient pose estimation methods.

However, our framework has certain limitations. The simple offset-based pose transformation inevitably introduces projection deformations that can affect predictions, so future work should explore more accurate ways to manipulate pose outputs without distortion while simultaneously developing corresponding defenses.

7 Ethical Statement

This paper presents, for the first time, a backdoor attack framework specifically targeting the 6DoF pose estimation task. In particular, the adversary can embed a specific 3D trigger into the training data to covertly manipulate the model's predictions. It should be emphasized that this research does not aim to spread such a framework, but rather to reveal the potential security risks inherent in current 6DoF pose estimation tasks. Future research efforts will focus on developing effective detection mechanisms and defense strategies to identify and mitigate threats posed by such backdoor attacks.

8 Acknowledgments

We thank the Program Chairs (PC), Senior Program Committee (SPC), and Area Chairs (AC) for their constructive feedback and guidance throughout the review process. This work was supported in part by the Yangcheng Scholars Research Project (No.2024312049), Science and Technology Projects in Guangzhou (No. 2025A04J4430), the InnoHK Program of the Hong Kong SAR Government, and the National Natural Science Foundation of China (Grant No. 62306313).

References

- Arnold, E.; Al-Jarrah, O. Y.; Dianati, M.; Fallah, S.; Oxtoby, D.; and Mouzakitis, A. 2019. A survey on 3d object detection methods for autonomous driving applications. *IEEE Transactions on Intelligent Transportation Systems*, 20(10): 3782–3795.
- Bates, M. 1995. Models of natural language understanding. *Proceedings of the National Academy of Sciences*, 92(22): 9977–9982.
- Chan, S.-H.; Dong, Y.; Zhu, J.; Zhang, X.; and Zhou, J. 2022. Baddet: Backdoor attacks on object detection. In *European Conference on Computer Vision*, 396–412. Springer.
- Chen, X.; Liu, C.; Li, B.; Lu, K.; and Song, D. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*.
- Dai, J.; Chen, C.; and Li, Y. 2019. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7: 138872–138878.
- El-Kassas, W. S.; Salama, C. R.; Rafea, A. A.; and Mohamed, H. K. 2021. Automatic text summarization: A comprehensive survey. *Expert systems with applications*, 165: 113679.
- Fan, Z.; Zhu, Y.; He, Y.; Sun, Q.; Liu, H.; and He, J. 2022. Deep learning on monocular object pose detection and tracking: A comprehensive overview. *ACM Computing Surveys*, 55(4): 1–40.
- Gu, T.; Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2019. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 7: 47230–47244.
- Guan, J.; Hao, Y.; Wu, Q.; Li, S.; and Fang, Y. 2024. A survey of 6dof object pose estimation methods for different application scenarios. *Sensors*, 24(4): 1076.
- Han, X.; Xu, G.; Zhou, Y.; Yang, X.; Li, J.; and Zhang, T. 2022. Physical backdoor attacks to lane detection systems in autonomous driving. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2957–2968.
- He, X.; Xu, G.; Han, X.; Wang, Q.; Zhao, L.; Shen, C.; Lin, C.; Zhao, Z.; Li, Q.; Yang, L.; Ji, S.; Li, S.; Zhu, H.; Wang, Z.; Zheng, R.; Zhu, T.; Li, Q.; He, C.; Wang, Q.; Hu, H.; Wang, S.; Sun, S.-F.; Yao, H.; Qin, Z.; Chen, K.; Zhao, Y.; Li, H.; Huang, X.; and Feng, D. 2025. Artificial intelligence security and privacy: a survey. *Science China Information Sciences*.
- Hinterstoisser, S.; Lepetit, V.; Ilic, S.; Holzer, S.; Bradski, G.; Konolige, K.; and Navab, N. 2013. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Computer vision—ACCV 2012: 11th asian conference on computer vision, daejeon, korea, november 5-9, 2012, revised selected papers, Part I 11*, 548–562. Springer.
- Hoque, S.; Arafat, M. Y.; Xu, S.; Maiti, A.; and Wei, Y. 2021. A comprehensive review on 3D object detection and 6D pose estimation with deep learning. *IEEE Access*, 9: 143746–143770.
- Huang, J.; Chen, Z.; Ceylan, D.; and Jin, H. 2017. 6-DOF VR videos with a single 360-camera. In *2017 IEEE Virtual Reality (VR)*, 37–44. IEEE.
- Kalia, M.; Navab, N.; and Salcudean, T. 2019. A real-time interactive augmented reality depth estimation technique for surgical robotics. In *2019 International Conference on Robotics and Automation (ICRA)*, 8291–8297. IEEE.
- Kendall, A.; Grimes, M.; and Cipolla, R. 2015. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, 2938–2946.
- Lan, H.; Gu, J.; Torr, P.; and Zhao, H. 2023. Influencer backdoor attack on semantic segmentation. *arXiv preprint arXiv:2303.12054*.
- Li, Y.; Huang, H.; Zhao, Y.; Ma, X.; and Sun, J. 2024. Backdoorllm: A comprehensive benchmark for backdoor attacks on large language models. *arXiv preprint arXiv:2408.12798*.
- Li, Y.; Zhai, T.; Jiang, Y.; Li, Z.; and Xia, S.-T. 2021. Backdoor attack in the physical world. *arXiv preprint arXiv:2104.02361*.
- Liao, Y.; Cao, Y.; Zhang, Y.; He, W.; Xiao, Y.; Du, X.; Huang, Z.; and Dong, J. S. 2025. Towards Stealthy and Effective Backdoor Attacks on Lane Detection: A Naturalistic Data Poisoning Approach. *CoRR*, abs/2508.15778.
- Liu, J.; Sun, W.; Yang, H.; Zeng, Z.; Liu, C.; Zheng, J.; Liu, X.; Rahmani, H.; Sebe, N.; and Mian, A. 2024a. Deep learning-based object pose estimation: A comprehensive survey. *arXiv preprint arXiv:2405.07801*.
- Liu, Y.; Sun, Z.; He, X.; and Huang, X. 2024b. Quantized Delta Weight Is Safety Keeper. *arXiv preprint arXiv:2411.19530*.
- Manhardt, F.; Kehl, W.; and Gaidon, A. 2019. Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2069–2078.

- Minaee, S.; Boykov, Y.; Porikli, F.; Plaza, A.; Kehtarnavaz, N.; and Terzopoulos, D. 2021. Image segmentation using deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(7): 3523–3542.
- Peng, S.; Liu, Y.; Huang, Q.; Zhou, X.; and Bao, H. 2019. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4561–4570.
- Rambach, J.; Pagani, A.; and Stricker, D. 2017. [poster] augmented things: Enhancing ar applications leveraging the internet of things and universal 3d object tracking. In *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, 103–108. IEEE.
- Reizenstein, J.; Shapovalov, R.; Henzler, P.; Sbordone, L.; Labatut, P.; and Novotný, D. 2021. Common Objects in 3D: Large-Scale Learning and Evaluation of Real-life 3D Category Reconstruction. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, 10881–10891. IEEE.
- Sha, Z.; He, X.; Berrang, P.; Humbert, M.; and Zhang, Y. 2022. Fine-tuning is all you need to mitigate backdoor attacks. *arXiv preprint arXiv:2212.09067*.
- Su, Y.; Di, Y.; Zhai, G.; Manhardt, F.; Rambach, J.; Busam, B.; Stricker, D.; and Tombari, F. 2023. Opa-3d: Occlusion-aware pixel-wise aggregation for monocular 3d object detection. *IEEE Robotics and Automation Letters*, 8(3): 1327–1334.
- Su, Y.; Rambach, J.; Minaskan, N.; Lesur, P.; Pagani, A.; and Stricker, D. 2019. Deep multi-state object pose estimation for augmented reality assembly. In *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, 222–227. IEEE.
- Sun, Z.; Cong, T.; Liu, Y.; Lin, C.; He, X.; Chen, R.; Han, X.; and Huang, X. 2025. PEFTGuard: Detecting Backdoor Attacks Against Parameter-Efficient Fine-Tuning. In *2025 IEEE Symposium on Security and Privacy (SP)*, 1620–1638. Los Alamitos, CA, USA: IEEE Computer Society.
- Turner, A.; Tsipras, D.; and Madry, A. 2018. Clean-label backdoor attacks.
- Wang, C.; Xu, D.; Zhu, Y.; Martín-Martín, R.; Lu, C.; Fei-Fei, L.; and Savarese, S. 2019. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3343–3352.
- Wang, H.; Wu, H.; He, Z.; Huang, L.; and Church, K. W. 2022. Progress in machine translation. *Engineering*, 18: 143–153.
- Wang, J.; Rupperecht, C.; and Novotný, D. 2023. PoseDif-fusion: Solving Pose Estimation via Diffusion-aided Bundle Adjustment. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 9739–9749. IEEE.
- Wen, B.; Tremblay, J.; Blukis, V.; Tyree, S.; Müller, T.; Evans, A.; Fox, D.; Kautz, J.; and Birchfield, S. 2023. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 606–617.
- Wen, B.; Yang, W.; Kautz, J.; and Birchfield, S. 2024. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17868–17879.
- Wu, J.; Yang, Y.; Liu, H.; Liao, S.; Lei, Z.; and Li, S. Z. 2019. Unsupervised graph association for person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8321–8330.
- Xiang, Y.; Schmidt, T.; Narayanan, V.; and Fox, D. 2017. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*.
- Xiao, H.; Biggio, B.; Brown, G.; Fumera, G.; Eckert, C.; and Roli, F. 2015. Is feature selection secure against training data poisoning? In *international conference on machine learning*, 1689–1698. PMLR.
- Zhai, G.; Huang, D.; Wu, S.-C.; Jung, H.; Di, Y.; Manhardt, F.; Tombari, F.; Navab, N.; and Busam, B. 2023. Mono-graspnet: 6-dof grasping with a single rgb image. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 1708–1714. IEEE.
- Zhao, W.; Chellappa, R.; Phillips, P. J.; and Rosenfeld, A. 2003. Face recognition: A literature survey. *ACM computing surveys (CSUR)*, 35(4): 399–458.
- Zhong, Z.; Sun, Z.; Liu, Y.; He, X.; and Tao, G. 2025. Backdoor Attack on Vision Language Models with Stealthy Semantic Manipulation. *CoRR*, abs/2506.07214.