

Seeing Through the Weights: Privacy Leakage in Scene Coordinate Regression

Oleksii Nasypanyi^{*1}, Jaemin Cho^{*1}, Utku Ozbulak^{2,3}, Byungkun Kang⁴,
and Francois Rameau^{†4}

¹Stony Brook University ²Ghent University Global Campus
³George Mason University Korea ⁴SUNY Korea

Abstract. Scene Coordinate Regression (SCR) methods are increasingly adopted for visual localization. In these approaches, the scene is implicitly encoded within a neural network that regresses a 3D world coordinate for each image pixel. Because the scene is represented only through the network parameters and not stored explicitly as images or maps, such methods are often assumed to be privacy-preserving. In this work, we show that this assumption is incorrect in practice.

Specifically, we introduce a query-based attack that reconstructs the 3D geometry of the training environment from an SCR model under different levels of model access. To do so, we repeatedly query the model with batches of proxy images unrelated to the target scene to obtain dense pixel-wise 3D coordinates. Reliable points are identified through their stability under small input perturbations and can be further refined in a white-box setting. These stable points are accumulated across independent query batches to recover the scene geometry. From the recovered 3D representation, we also invert the network features to synthesize images from arbitrary viewpoints, revealing additional appearance information. Experiments on indoor and outdoor datasets demonstrate that substantial portions of training environments can be reconstructed with high geometric fidelity. Beyond geometry, we also recover an approximate color appearance, which exposes recognizable layout and potentially sensitive scene elements. This directly contradicts claims in the literature that SCR representations are privacy-preserving by design, and reveals a real risk when such systems are deployed in private or security-critical spaces. The project page is available [here](#).

Keywords: Visual Localization · Privacy Leakage · Model Inversion

1 Introduction

Given an image of a known environment, visual localization aims to recover the camera’s six-degree-of-freedom pose. It finds applications in augmented reality [40], autonomous navigation [22], and robotics [27]. As these systems are

^{*}Equal contribution: {oleksii.nasypanyi, jaemin.cho}@stonybrook.edu

[†]Corresponding author: francois.rameau@sunkorea.ac.kr




Fig. 1: Overview of the proposed query-based attack. An SCR model trained on images of a private scene is queried using proxy images. By aggregating the returned 3D predictions, the attacker reconstructs a voxelized representation of the scene and synthesizes novel views, revealing both geometry and appearance.

increasingly used in private or sensitive spaces like homes, offices, or medical facilities, privacy becomes a real concern [8]. In such contexts, visual localization systems may inadvertently expose information about the environment for which they were built. This leakage may reveal structural clues such as scene layout or geometry, or even appearance information, resulting in a privacy breach. This issue is particularly obvious for hierarchical localization techniques [38] based on Structure-from-Motion (SfM) maps, where the 3D structure of the scene is directly and explicitly available. Prior works have addressed this by replacing 3D points with lines [44] to prevent model inversion via feature reprojection [17]. However, this problem is far less explored for recent learning-based approaches such as Scene Coordinate Regression (SCR) [3, 4], which represent the scene implicitly through a neural network that predicts per-pixel 3D coordinates in a fixed world coordinate system. Because the scene is not stored explicitly as images or maps and exists only through the network parameters, these methods are widely believed to be inherently privacy-preserving. Such claims appear repeatedly in the literature [3, 9, 16, 49]. For example:

- “The resulting maps can be as small as 4MB, and privacy preserving” [3]
 “SCR is said to be inherently privacy-preserving because there is no set of 2D or 3D points to run the inversion attack on.” [9]
 “These methods implicitly encode scene information in the learned parameters of a convolutional neural network (CNN), rather than explicitly storing images or features. Thus, they preserve privacy by design.” [16]
 “DSAC++ is the most accurate method while being resilient to privacy attacks as it does not need to transmit visual descriptors” [49]

In this work, we show that this assumption does not hold in practice. We demonstrate that the structure and appearance of a training scene can be extracted from an SCR network through a query-based attack. For this purpose, we consider an adversary with query access to a trained SCR model under black-box, gray-box, or white-box settings, reflecting increasing levels of access to the feature extractor and model parameters. In all cases, the adversary has no access to the training images or explicit scene geometry. Despite these constraints, our attack recovers detailed information about the scene, as depicted in Fig. 1.

Our attack consists of multiple stages. We begin by querying the SCR network with batches of proxy images drawn from natural image datasets unrelated to the target scene, obtaining dense pixel-wise 3D coordinate predictions. While individual predictions are noisy, aggregating them across many batches already reveals a coarse structure of the scene. This alone constitutes a black-box attack requiring no access to model internals.

This approach can be significantly improved when access to the feature extractor is available. SCR models are well-behaved near features observed during training but become unstable outside that distribution [33]. Since proxy images are unrelated to the target scene, the features they induce frequently fall into these unstable regions, where small perturbations cause large changes in predicted 3D coordinates. We exploit this property in two ways: first, stability under feature perturbations serves as a criterion to identify and discard unreliable predictions; second, in the white-box setting, features can be refined by gradient descent to move them toward stable regions corresponding to geometry encoded during training.

The remaining predictions are accumulated across proxy batches using a voxel grid. Since all predictions share the same learned coordinate system, voxels consistently populated across batches correspond to true scene geometry.

Finally, recovered feature vectors are projected onto arbitrary virtual camera poses and inverted back to image space via a feature inversion network, revealing approximate scene appearance.

This paper makes the following contributions: (i) we provide, to the best of our knowledge, the first systematic evidence that SCR models are not inherently privacy-preserving; (ii) we present the first attack framework that reconstructs scene geometry from trained SCR models under different access levels; (iii) we further recover coarse appearance by inverting reconstructed features, allowing recognizable view synthesis from novel viewpoints; (iv) we validate the attack on standard benchmarks, recovering substantial portions of the training environments.

2 Related Works

In this section, we will review the different visual localization approaches, their vulnerabilities and the solutions that have been deployed for their defense.

Visual Localization

Recent visual localization approaches fall into 3 main families: hierarchical localization, absolute pose regression (APR), and scene coordinate regression (SCR).

Hierarchical localization estimates the 6DoF camera pose within a pre-built SfM map using a coarse-to-fine pipeline. It typically begins with global image retrieval [1, 21], followed by local feature matching [15, 28] against the 3D map, and final pose estimation via PnP-RANSAC [25]. Although highly accurate and scalable, these methods require storing large maps, incur significant computational cost, and do not inherently protect the privacy of the mapping images [38].

To mitigate these limitations, APR methods [12, 13, 23, 43] directly regress camera pose from a single image, implicitly encoding scene information in network parameters. However, APR suffers from limited accuracy and scalability.

SCR improves accuracy by predicting a 3D world coordinate for each pixel, yielding dense 2D–3D correspondences for PnP-RANSAC. Originally introduced with random forests [41] and later extended with differentiable RANSAC [4], SCR achieves state-of-the-art accuracy on small scenes but requires long training times. Accelerated Coordinate Encoding (ACE) decouples SCR into a frozen scene-agnostic backbone and a lightweight scene-specific MLP [3], reducing training from hours to minutes, with extensions to uncalibrated cameras [5], large-scale environments [47], and improved generalization [6]. Beyond performance, SCR methods are often regarded as privacy-preserving; an assumption we challenge in this work.

Privacy in Visual Localization

Among the first to highlight model inversion risks in visual localization, [38] showed that images can be reconstructed from sparse SfM maps via feature re-projection and inversion. To counter such attacks, Speciale et al. [44] proposed replacing 3D points with randomly oriented lines, preserving enough geometry for pose estimation [48] while theoretically preventing map inversion. This was extended to protect client-side 2D keypoints [45] and later to privacy-preserving SfM [20]. Despite the popularity of such line-based defenses, they have proven not to be entirely robust. As shown in [10], the original 3D points can be recovered from line clouds using density-based statistics. This vulnerability motivated the development of more sophisticated lifting defense strategies, including Paired-Point Lifting [26], 3D Ray Clouds [31], and Coordinate Permutations [34]. However, these advanced representations were shown to remain vulnerable to attacks exploiting neighborhood information encoded in feature descriptors [9].

In parallel to geometric defenses, another research direction focuses on feature-level privacy to prevent scene reconstruction. For instance, [18] and [32] utilize adversarial learning to conceal sensitive content within the descriptors, while [39] and [36] provide stronger protection by introducing differential privacy or leveraging abstract semantic labels that are inherently difficult to invert.

Recent work has attempted to improve privacy in NeRF [30] and Gaussian Splatting [24] by removing RGB supervision [35, 37], but this only masks visual appearance, while the scene structure and its content are still directly recoverable. In contrast, SCR models [3, 4] are never trained to render or reconstruct visual appearance; their objective is strictly to map pixels to 3D coordinates. As a result, they are commonly viewed as privacy-preserving [3, 16, 49]. This assumption has not been rigorously evaluated and appears unfounded. Unlike prior inversion methods [8] that operate on features extracted from target scene images or explicit 3D maps, our attack has access to neither: it relies solely on features from unrelated proxy images, with no knowledge of the scene geometry. Despite this constrained setting, we show that substantial scene information can still be recovered.




Fig. 2: Gray-box deployment scenario considered in this work. Feature extraction is performed on the client using a public encoder, while the scene-specific head is hosted remotely. The client transmits spatially shuffled feature maps to the server, which returns predicted 3D scene coordinates used locally for pose estimation.

3 Background

This section introduces SCR, the ACE framework as a representative SCR example used in this work, and the deployment settings considered in our analysis.

3.1 Scene Coordinate Regression and ACE

Scene Coordinate Regression formulates visual localization by predicting dense 2D–3D correspondences with a neural network, rather than regressing the camera pose directly. Given an input image $I \in \mathbb{R}^{C_{\text{img}} \times H \times W}$ (where $C_{\text{img}} = 1$ for grayscale and $C_{\text{img}} = 3$ for RGB), an SCR model parameterized by θ predicts a dense set of 3D scene coordinates $\mathbf{X} = h_{\theta}(I) = \{\mathbf{x}_1, \dots, \mathbf{x}_{HW}\}$, where each $\mathbf{x}_i \in \mathbb{R}^3$ corresponds to the 3D scene coordinate predicted for one image pixel, expressed in a fixed world coordinate system. These correspondences are then used within a PnP–RANSAC pipeline to estimate the 6DoF camera pose.

While our attack is compatible with most SCR models, we use ACE [3] as our main illustrative case. This choice is motivated by its efficiency and explicit separation between scene-agnostic and scene-specific components, which is well-suited to server-client deployments commonly used in privacy-preserving localization pipelines [11]. As mentioned earlier, ACE models h_{θ} as a two-stage process. A pretrained, scene-agnostic CNN backbone ϕ maps an input grayscale image I to a dense feature map $\mathbf{F} = \phi(I) \in \mathbb{R}^{C \times H' \times W'}$, where C is the feature channel dimension and (H', W') is the spatial resolution after downsampling, from which a scene-specific MLP g_{θ_s} predicts a dense field of 3D scene coordinates $\mathbf{X} = g_{\theta_s}(\mathbf{F}) \in \mathbb{R}^{3 \times H' \times W'}$.

3.2 Threat Model

We adopt standard adversarial access models commonly used in machine learning security and model extraction settings [2, 7], and analyze SCR systems under three levels of attacker access. These settings correspond to increasing visibility into the model pipeline, ranging from query-only access to partial or full access to model components.

Black-box access: The localization pipeline is hosted remotely and exposed through a prediction API where clients transmit input images to the server, which returns predicted 3D point maps or camera poses. In a standard black-box inference interface, the attacker can issue arbitrary queries and observe model outputs, while having no access to intermediate representations or model parameters [46]. This setting corresponds to the weakest access model, where the attacker can only interact with the system through input-output queries and has no visibility into intermediate representations or model parameters.

Gray-box access: In this setting, feature extraction is performed on the client using a public, scene-agnostic encoder ϕ , while the scene-specific regression head g_{θ_s} is hosted remotely and accessed via an API that accepts feature maps and returns predicted 3D coordinates. This split architecture is commonly considered in client-server localization systems [9]. Since SCR feature maps are dense and may retain appearance information, they can be vulnerable to inversion-based reconstruction [38]. As a simple privacy-preserving design choice, we consider a variant in which the client applies a random spatial permutation to the feature map prior to transmission, removing explicit 2D spatial ordering while preserving per-feature descriptors. The server therefore operates on unordered features and returns predicted 3D coordinates, while pose estimation is performed on the client, as illustrated in Fig. 2.

White-box access: In this setting, the attacker has access to the parameters of g_{θ_s} and can compute gradients through the model. This may occur, for example, when scene-specific heads are deployed on-device or when model parameters are otherwise exposed. While such exposure is not typical in standard deployments, it can arise when models are assumed to be inherently privacy-preserving and thus distributed without additional safeguards. This setting therefore represents a worst-case scenario in terms of potential information leakage under maximal adversarial access.

Objective of the adversary: Across all settings, the attacker seeks to reconstruct the private training scene geometry $\mathcal{G} = \{\mathbf{x}_i\}_{i=1}^N$ using the capabilities permitted by the given access model (e.g., query access, intermediate features, or gradients) together with a collection of proxy images unrelated to the target scene. Importantly, the evaluated threat models assume no access to the original training images or ground-truth scene geometry.

4 Method

Building on the threat model detailed above, we present a query-based attack that reconstructs scene geometry from a trained SCR model. For clarity, we describe the full pipeline in the white-box setting; the gray-box and black-box variants follow the same sequence but omit stages that require additional access, such as gradient-based refinement or descriptor-level operations. An overview of the full pipeline is provided in Fig. 3. The attack proceeds through the following stages: proxy queries (Section 4.1); stability estimation (Section 4.2); feature refinement (Section 4.3); batch filtering (Section 4.4); accumulation (Section 4.5);




Fig. 3: Overview of the proposed white-box attack on SCR models. (1) Proxy images are encoded using a public feature extractor to obtain dense feature maps. (2) Features are optionally refined via gradient descent to move them toward stable regions of the learned scene. (3) Unstable predictions are discarded by retaining only the top- $p\%$ most stable features. (4) Stable 3D predictions are accumulated into a voxel occupancy grid to recover scene geometry. (5) Recovered features are projected to virtual camera poses and inverted to synthesize novel views, revealing scene appearance.

and image reconstruction (Section 4.6). We describe each stage in detail in the following sections.

4.1 Proxy Queries

At each iteration, we sample a batch of B proxy images $\mathcal{I} = \{I_1, \dots, I_B\}$ from natural-image datasets [19], extract features using the public backbone $\mathbf{F} = \phi(\mathcal{I})$, and forward them to the scene-specific head to obtain dense 3D coordinate predictions $\mathbf{X} = g_{\theta_s}(\mathbf{F})$.

Although individual predictions are noisy, the model always outputs coordinates within the geometry of the training scene. Predicted points therefore tend to fall on or near scene surfaces rather than being uniformly distributed in space. Aggregating predictions across many batches already yields a coarse approximation of the scene geometry; this alone constitutes our black-box attack, requiring no access to the feature extractor or model internals. The following stages further improve reconstruction quality when greater access is available.

4.2 Stability Estimation

Proxy queries produce a dense but noisy set of 3D predictions, not all of which correspond to geometry encoded during training. Since neural networks tend to be stable near the training manifold but sensitive outside it [33], we identify well-supported predictions by measuring how much the predicted 3D coordinate changes under small perturbations of its input feature.

Concretely, given a feature vector $\mathbf{f}$ extracted from a proxy image, the SCR model predicts a 3D coordinate $\mathbf{x} = g_{\theta_s}(\mathbf{f})$. We assess the stability of this prediction by perturbing the feature vector with small isotropic Gaussian noise and




Fig. 4: Detailed modules of the attack: (a) Stability score $s(\mathbf{f})$ computed from the variance of noise-perturbed predictions. (b) Training pipeline of the inversion network.

measuring the resulting variation in predicted 3D space, as illustrated in Fig. 4a. Specifically, we draw perturbations $\epsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ and compute

$$s(\mathbf{f}) = \frac{1}{n-1} \sum_{k=1}^n \|g_{\theta_s}(\mathbf{f} + \epsilon_k) - \bar{\mathbf{x}}\|_2^2, \quad \bar{\mathbf{x}} = \frac{1}{n} \sum_{k=1}^n g_{\theta_s}(\mathbf{f} + \epsilon_k). \quad (1)$$

Low values of $s(\mathbf{f})$ indicate that $\mathbf{f}$ lies in a locally consistent region of the learned mapping; large values indicate unstable extrapolation.

Stability estimation is performed independently for each feature vector in the dense feature maps produced by a proxy batch. The resulting stability scores are later used both to refine feature representations and to filter reliable predictions.

4.3 Feature Refinement

Proxy features are generally not aligned with those observed during training, which can lead to imprecise or inconsistent coordinate predictions. When gradients of the scene-specific head g_{θ_s} are available (white box configuration), we optionally refine the feature vectors obtained from proxy queries to encourage convergence toward stable regions of the learned scene.

Concretely, given a feature vector $\mathbf{f}$, we seek a refined feature $\mathbf{f}^*$ that minimizes the stability score $s(\mathbf{f})$ defined in Section 4.2:

$$\mathbf{f}^* = \arg \min_{\mathbf{f}} s(\mathbf{f}). \quad (2)$$

This is optimized by gradient descent via updates $\mathbf{f} \leftarrow \mathbf{f} - \lambda \nabla_{\mathbf{f}} s(\mathbf{f})$, run for N iterations or until convergence. Refinement is applied independently to each feature vector in the proxy batch.

4.4 Batch Filtering

After the optional refinement step (Section 4.3), some predictions may still lie in unstable regions and should be discarded. Predictions are ranked by their stability score $s(\mathbf{f})$ and only the top $p\%$ per batch are retained, whether applied directly to proxy features or to refined features $\mathbf{f}^*$ (depending on the access level). The resulting sparse set of stable 3D predictions is passed to the accumulation stage (Section 4.5).

4.5 Accumulation Across Proxy Batches

While stability filtering acts as a local selection within each batch, accumulation serves as a global consistency check across batches. Stable 3D predictions from each batch are aggregated to recover a more complete estimate of the scene geometry.

Formally, let $\mathcal{P}_t = \{\mathbf{x}_i\}_{i=1}^{M_t}$ denote the retained 3D points from proxy batch t , expressed in the learned scene coordinate system. We discretize space into a voxel grid $\mathcal{V}$ with resolution Δ and maintain a global occupancy count $C(v)$ for each voxel $v \in \mathcal{V}$, where each point $\mathbf{x}_i$ increments the count of its corresponding voxel. Voxels consistently populated across independent batches accumulate higher counts, while spurious predictions remain weakly supported.

Accumulation continues until the average occupancy over non-empty voxels $\bar{C}_t$ falls below a threshold τ_C or stabilizes across successive batches. The final reconstruction retains voxels satisfying $C(v) \geq \kappa$, where κ is a fixed minimum occupancy threshold.

For gray-box and white-box attacks, where stability scores are available, the most stable 3D prediction within each retained voxel is kept as its representative point. In the black-box setting, where no filtering or refinement is performed, the centroid of all predictions falling in each voxel is used instead.

4.6 Image Reconstruction via Feature Inversion

While previous stages reconstruct scene geometry, point clouds alone do not always reveal sensitive visual details. To expose such information, we recover scene appearance by synthesizing images from arbitrary viewpoints via a feature inversion network. Each pixel receives the feature of the nearest projected point in depth, forming a projected feature tensor $\mathbf{F}_{\text{proj}} \in \mathbb{R}^{C \times H \times W}$, with empty pixels set to zero. A CNN encoder downsamples $\mathbf{F}_{\text{proj}}$ to match the SCR feature resolution $\mathbb{R}^{C \times H' \times W'}$. Because $\mathbf{F}_{\text{proj}}$ aggregates descriptors induced by many proxy queries, it is noisy and differs from features extracted from real images. The inversion network first refines this tensor into an adapted feature map $\hat{\mathbf{F}}$ using a Transformer encoder with five attention blocks, then reconstructs an image $\hat{\mathbf{I}}$ through a convolutional upsampling decoder producing an output of size $3 \times H \times W$. Similar image reconstruction attacks have been studied for SfM maps [17, 29, 44], where accurate geometry and features from the original images are available. In our setting, the geometry is recovered by the attack itself and the descriptors arise from proxy queries rather than target images, making inversion significantly more complex. To make inversion feasible under these conditions, we train the inversion model on data that reproduces the attack setting. ACE heads are first trained on 120 ScanNet scenes and queried with the real training images to obtain ground-truth 3D points. They are then queried with proxy images from SUN RGB-D [42] to extract candidate point-descriptor pairs. To improve reliability, we retain only pairs whose predicted 3D locations lie within 4 cm of the corresponding ground-truth points and for which at least four observations fall into the same voxel. Training samples are generated by

projecting the filtered reconstructed points into target views and supervising the network with the corresponding ScanNet images.

Optimization combines an image reconstruction loss $\mathcal{L}_{img}$ (MSE + LPIPS) between $\hat{\mathbf{I}}$ and $\mathbf{I}$ with a feature consistency loss $\mathcal{L}_{feat}$ aligning $\hat{\mathbf{F}}$ to SCR backbone features $\mathbf{F}_{gt}$. Both losses are masked to account for sparse projections, using $\mathcal{M}_f$ for valid feature locations and a dilated mask $\mathcal{M}_I$ for image regions where projected points are present. This masking is necessary because some regions, particularly scene boundaries, may be observed only few times in the training set, and low-texture areas are inherently harder to recover reliably. The whole training pipeline and architecture is visible in Fig. 4b. Note that, using a single scene per batch leads to overfitting and strong color bias; we mitigate this effect with mixed-scene batching and random projection dropout.

At inference time, the inversion model operates on reconstructed point clouds produced by the attack, with descriptors filtered and voxelized using the same procedure as during training.

5 Experiments

In this section, we evaluate the proposed attack on indoor and outdoor benchmarks across multiple architectures, attacker access levels, and query budgets.

5.1 Experimental Setup and Implementation Details

Datasets. We evaluate on 7-Scenes [41] and Cambridge Landmarks [23], covering indoor and outdoor environments. Proxy queries are drawn from PASCAL VOC [19]. The feature inversion network is trained on ScanNet [14].

Implementation Details. We use $n = 4$ perturbations with $\sigma = 0.01$, batch size $B = 32$, and $N = 30$ refinement iterations, retaining the top $p = 20\%$ of predictions per batch. Voxel resolution is $\Delta = 2$ cm for indoor and 10 cm for outdoor scenes; other hyperparameters are provided in the supplementary. The inversion model is trained for 10,000 iterations on an RTX 4090.

Metrics. Reconstruction quality is measured against a pseudo-GT obtained by querying the model with training images. We report F1 at 2/5 cm (indoor) and 10/30 cm (outdoor), where higher values indicate better reconstruction accuracy, and half Chamfer Distance (CD), where lower values indicate better geometric alignment. Image quality is reported as PSNR with higher values corresponding to better image fidelity. We use pseudo-GT as the reference because it reflects the geometry encoded by the model: an attack cannot recover geometry that the model itself has not learned. This also provides a model-specific reference, since different SCR architectures may encode slightly different subsets of the scene. A detailed explanation, together with visual and quantitative examples and comparisons against SfM and RGB-D ground truth, is provided in the supplementary material.

Table 1: Attack progression on ACE (7Scenes). Comparison of Black-, Gray-, and White-box attack stages using Pascal proxy. **Bold:** best per scene.

	Black-Box			Gray-Box			White-Box		
	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$
Chess	71.2	88.0	2.55	77.7	93.3	1.90	83.0	97.2	1.59
Fire	60.3	80.7	3.13	76.2	94.0	1.83	79.9	95.8	1.72
Heads	52.0	77.8	3.03	70.3	92.6	1.91	78.1	95.3	1.78
Office	71.2	89.7	2.29	76.2	92.0	2.19	82.7	96.2	1.65
Pumpkin	61.1	79.7	3.67	68.5	87.1	2.56	76.1	93.4	2.04
Redkitch.	74.2	91.5	2.41	76.5	92.2	2.08	82.3	96.5	1.62
Stairs	46.2	72.1	4.79	48.9	76.1	3.57	58.8	84.8	2.96
Avg.	62.3	82.8	3.12	70.6	89.6	2.30	77.3	94.2	1.91

5.2 Access Level vs. Reconstruction Quality

Table 1 reports reconstruction quality on 7-Scenes for ACE across the three attack variants. Reconstruction quality improves consistently with access level, with white-box yielding the best results across all scenes. Notably, the black-box variant already achieves competitive performance, with an average F1 of 62.3 at 2 cm and 82.8 at 5 cm, and gray-box scores remain close to white-box, confirming that meaningful 3D reconstruction is achievable even without access to model internals or gradients. Stairs is the most challenging scene across all variants due to repetitive structures and textureless surfaces, which increase prediction ambiguity in the SCR model and, by extension, reduce attack efficacy. Qualitative results in the supplementary material further confirm this trend, showing progressively cleaner reconstructions as access increases.

Table 2: Generalization across SCR architectures. White-box attack results on 7-Scenes (F1 at 2 cm/5 cm, CD in cm) and Cambridge Landmarks (F1 at 10 cm/30 cm, CD in m).

		ACE [3]			GLACE [47]			DSAC* [4]			ACE-G [6]		
		$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$	$F1^2 \uparrow$	$F1^5 \uparrow$	$CD \downarrow$
7-Scenes	Chess	83.0	97.2	1.59	78.2	95.7	1.70	71.8	90.8	2.25	57.1	85.3	3.00
	Fire	79.9	95.8	1.72	80.4	96.4	1.67	71.8	91.4	2.15	58.9	85.6	3.19
	Heads	78.1	95.3	1.78	75.6	93.4	1.86	68.0	91.1	2.40	59.0	88.0	2.68
	Office	82.7	96.2	1.65	84.0	97.6	1.47	69.5	89.3	2.47	58.0	84.9	3.72
	Pumpkin	76.1	93.4	2.04	72.7	92.3	2.13	58.6	82.7	3.40	46.1	78.6	4.55
	Redkitch.	82.3	96.5	1.62	82.9	97.0	1.54	65.6	86.5	2.74	55.2	84.2	3.36
	Stairs	58.8	84.8	2.96	61.7	86.8	2.64	40.1	66.8	6.05	33.4	68.0	5.29
Avg.	77.3	94.2	1.91	76.5	94.2	1.86	63.6	85.5	3.07	52.5	82.1	3.68	
		$F1^{10} \uparrow$	$F1^{30} \uparrow$	$CD \downarrow$	$F1^{10} \uparrow$	$F1^{30} \uparrow$	$CD \downarrow$	$F1^{10} \uparrow$	$F1^{30} \uparrow$	$CD \downarrow$	$F1^{10} \uparrow$	$F1^{30} \uparrow$	$CD \downarrow$
Cambridge	KingsColl.	64.0	92.5	0.13	39.0	73.9	0.25	22.5	57.1	0.45	25.1	62.3	0.40
	OldHosp.	45.1	74.8	0.25	49.6	77.4	0.24	29.5	61.7	0.45	27.2	69.0	0.45
	ShopFac.	83.5	98.8	0.08	58.2	86.3	0.16	50.6	84.9	0.20	39.5	80.6	0.21
	StMarys	65.1	91.5	0.14	16.6	61.6	0.52	31.2	67.6	0.38	14.3	46.9	0.42
	GreatCourt	43.3	74.8	0.29	9.0	26.9	1.36	33.9	63.7	0.56	22.2	57.7	0.64
	Avg.	59.0	85.9	0.18	34.5	65.2	0.51	33.5	67.0	0.41	25.7	63.3	0.42

5.3 Generalization Across SCR Architectures

Table 2 reports white-box results across SCR architectures on indoor and outdoor datasets. Results obtained with the attack follow similar patterns across architectures, indicating that the leakage is not architecture-specific.

Across both datasets, the relative behavior of the evaluated models is largely consistent. ACE and GLACE yield the most complete reconstructions, while DSAC* and ACE-G produce weaker reconstructions with noisier geometry. Reconstruction quality is overall higher on 7-Scenes than on Cambridge Landmarks. Indoor scenes are spatially smaller and geometrically simpler, which makes their structure easier to recover. In contrast, the larger spatial extent and higher variability of outdoor scenes reduce reconstruction completeness, leading to a noticeable performance drop for all models on Cambridge.

Architectural differences nevertheless influence the severity of leakage. ACE appears most exposed, likely because its scene-specific MLP fits the training scene very closely, concentrating geometric information in a compact parameter space that is easier to extract. GLACE conditions its regression head on a global scene descriptor, which could, in principle, provide additional protection. On small indoor scenes such as 7-Scenes, where views are highly similar and spatially limited, this conditioning offers little resistance. On larger outdoor scenes, however, GLACE shows noticeably stronger resistance compared to ACE. In our experiments, the global descriptor is simply set to zero (see supplementary). DSAC* shows comparatively stronger resistance. Unlike ACE-based models, it is trained end-to-end and does not rely on pretrained features. As a result, we believe that proxy images from unrelated datasets might produce more unstable predictions that do not consistently align with the scene geometry, leading to noisier reconstructions. Finally, ACE-G, designed to improve generalization, also reduces leakage to some extent. Although resistance varies across architectures, substantial portions of the scene remain recoverable in all cases.

We also observe that attack performance correlates with the localization accuracy of the underlying model (Fig. 5b). Models with higher localization error tend to produce less accurate reconstructions. This suggests that when the scene is not well learned for localization, its geometric structure is encoded less consistently in the network parameters, making it harder to extract. Note that GLACE is not included in this comparison, as its regression head is conditioned on a global scene descriptor, which alters this relationship.

5.4 Proxy Image Domains and Query Budget

We also analyze how the domain of proxy images affects the required query budget. To study this effect, we select representative indoor and outdoor scenes: 7-Scenes Stairs and Cambridge Shop Facade; results on additional scenes are provided in the supplementary material. For these scenes, we attack ACE using white-box and gray-box variants of the attack with images drawn from three proxy sources: (1) natural images from PASCAL VOC [19], (2) indoor scene images from SUN RGB-D [42], and (3) synthetic geometric patterns generated from random noise (see supplementary for details).




Fig. 5: (a) Effect of proxy image domain on query budget for two representative scenes (top) 7-Scenes Stairs and (bottom) Cambridge Shop Facade and (b) relationship between localization model accuracy (translation) and reconstruction quality.

Table 3: PSNR comparison across attack settings on ACE (7Scenes).

	Chess	Fire	Heads	Office	Pumpkin	Redkitch.	Stairs	Avg.
Black	14.83	14.71	12.58	15.26	16.59	14.99	16.02	15.00
Gray	<u>15.19</u>	<u>14.72</u>	<u>12.98</u>	<u>15.76</u>	<u>16.70</u>	<u>15.14</u>	<u>16.05</u>	<u>15.22</u>
White	15.29	14.96	13.01	15.86	16.72	15.23	16.15	15.32

In all cases, proxy images drawn from the synthetic geometric pattern dataset lead to poor reconstruction quality, even when a large number of proxy images is used. In contrast, real world images produce substantially better reconstructions. For white-box attacks, the two proxy datasets eventually converge to similar reconstruction quality as the query budget increases. In the gray-box setting, however, we observe that the results are scene-dependent: certain scenes favor PASCAL whereas others favor SUN, highlighting the importance of proxy image selection for efficient reconstruction under limited query budgets. Across all sequences, using natural images, a good-quality reconstruction can be achieved with only a few hundred images with white or gray-box attacks.

5.5 Image Reconstruction

We analyze the quality of the image reconstruction network by evaluating both the visual fidelity of the generated images and the consistency of the reconstructed 3D point clouds. Point-feature pairs are generated following the protocol in Sec. 4.6, treating the extracted points as pseudo ground truth. Because pairs are filtered by geometric consistency and voxel occupancy, the points used for image generation may slightly differ from those reconstructed during the attack. We further use the ground-truth camera poses from the original dataset to obtain reference images for comparison.




Fig. 6: Qualitative reconstruction on 7-Scenes. Free-view renderings are generated from viewpoints that are not included in either the training or test sets. Real query images are shown alongside renderings of the reconstructed scene from the same poses for visual comparison. We also show the recovered 3D geometry as an error heat map, with colors ranging from blue for 0 cm error to red for errors above 10 cm, together with the reconstructed geometry with recovered colors and the ACE pseudo-GT.

Qualitative examples for the white-box setting are shown in Fig. 6. Quantitative PSNR results are reported in Tab. 3. The results confirm that reconstruction quality consistently improves with increased access to the underlying model.

To the best of our knowledge, this is the first work to perform image inversion directly from a trained SCR model. Despite that, we report the results of [9] as a reference, which achieves an average PSNR of 16.01 across scenes, while our best result is 15.32, where our results are also obtained under a substantially more challenging setting. We provide additional qualitative and quantitative comparisons with other PPVL inversion methods [38], together with an analysis of appearance leakage quality, in the supplementary material.

We attribute the remaining gap primarily to two factors. First, accurate color prediction is inherently challenging in our setting. Second, certain regions of the scene, particularly near image boundaries, are underrepresented during training, which negatively impacts reconstruction quality in those areas.

To illustrate cross-domain generalization, we additionally present qualitative results on the Cambridge dataset in Fig. 7. Despite being trained exclusively on indoor scenes, the model still produces meaningful predictions outdoors.



Fig. 7: Cambridge Landmarks reconstruction results. The model reconstructs outdoor scenes despite being trained only on indoor environments.

6 Conclusion

We presented the first attack capable of reconstructing the scene representation encoded by SCR models. By issuing proxy queries and aggregating the returned 3D predictions, the proposed approach recovers both geometric structure and appearance cues. Our method operates under multiple levels of access and reveals a previously underexplored privacy risk in visual localization systems.

In future work, we plan to investigate more query-efficient reconstruction strategies to reduce the number of model interactions required for successful recovery. We also aim to extend the analysis to more restrictive deployment settings, including black-box localization APIs that return only the final camera pose rather than intermediate scene-coordinate predictions. Finally, we plan to explore defense mechanisms that mitigate information leakage, including output perturbation, access control policies, and training-time regularization to limit the exposure of sensitive scene representations.

7 Acknowledgments

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation(IITP)-Innovative Human Resource Development for Local Intellectualization program grant funded by the Korea government(MSIT) (IITP-2026-RS-2023-00259678).

References

1. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: Netvlad: Cnn architecture for weakly supervised place recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5297–5307 (2016)
2. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In: International conference on machine learning. pp. 274–283. PMLR (2018)
3. Brachmann, E., Cavallari, T., Prisacariu, V.A.: Accelerated coordinate encoding: Learning to relocalize in minutes using rgb and poses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5044–5053 (2023)

4. Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: Dsac-differentiable ransac for camera localization. In: CVPR (2017)
5. Brachmann, E., Wynn, J., Chen, S., Cavallari, T., Monszpart, A., Turmukhambetov, D., Prisacariu, V.A.: Scene coordinate reconstruction: Posing of image collections via incremental learning of a relocalizer. In: European Conference on Computer Vision. pp. 421–440. Springer (2024)
6. Bruns, L., Barroso-Laguna, A., Cavallari, T., Monszpart, A., Munukutla, S., Prisacariu, V.A., Brachmann, E.: Ace-g: Improving generalization of scene coordinate regression through query pre-training. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 26751–26761 (2025)
7. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 39–57. IEEE (2017)
8. Chelani, K.: Privacy in Visual Localization. Chalmers Tekniska Hogskola (Sweden) (2025)
9. Chelani, K., Benbihi, A., Kahl, F., Sattler, T., Kukeleva, Z.: Obfuscation based privacy preserving representations are recoverable using neighborhood information. In: 2025 International Conference on 3D Vision (3DV). pp. 189–199. IEEE (2025)
10. Chelani, K., Kahl, F., Sattler, T.: How privacy-preserving are line clouds? recovering scene details from 3d lines. In: CVPR (2021)
11. Chelani, K., Sattler, T., Kahl, F., Kukeleva, Z.: Privacy-preserving representations are not enough: Recovering scene content from camera poses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13132–13141. IEEE (2023)
12. Clark, R., Wang, S., Markham, A., Trigoni, N., Wen, H.: Vidloc: A deep spatio-temporal model for 6-dof video-clip relocalization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6856–6864 (2017)
13. Clark, R., Wang, S., Wen, H., Markham, A., Trigoni, N.: Vinet: Visual-inertial odometry as a sequence-to-sequence learning problem. In: Proceedings of the AAAI conference on artificial intelligence (2017)
14. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
15. DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 224–236 (2018)
16. Do, T., Miksik, O., DeGol, J., Park, H.S., Sinha, S.N.: Learning to detect scene landmarks for camera localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2022)
17. Dosovitskiy, A., Brox, T.: Inverting visual representations with convolutional networks. In: CVPR (2016)
18. Dusmanu, M., Schonberger, J.L., Sinha, S.N., Pollefeys, M.: Privacy-preserving image features via adversarial affine subspace embeddings. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14267–14277 (2021)
19. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**(2), 303–338 (2010)
20. Geppert, M., Larsson, V., Speciale, P., Schönberger, J.L., Pollefeys, M.: Privacy preserving structure-from-motion. In: ECCV (2020)

21. Izquierdo, S., Civera, J.: Optimal transport aggregation for visual place recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 17658–17668 (2024)
22. Jeong, H., Shin, J., Rameau, F., Kum, D.: Multi-modal place recognition via vectorized hd maps and images fusion for autonomous driving. *Robotics and Automation Letters* **9**(5), 4710–4717 (2024)
23. Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for real-time 6-dof camera relocalization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2938–2946 (2015)
24. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
25. Larsson, V., Kukelova, Z., Zheng, Y.: Making minimal solvers for absolute pose estimation compact and robust. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2316–2324 (2017)
26. Lee, C., Kim, J., Yun, C., Hong, J.H.: Paired-point lifting for enhanced privacy-preserving visual localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17266–17275 (2023)
27. Lee, D., Ryu, S., Yeon, S., Lee, Y., Kim, D., Han, C., Cabon, Y., Weinzaepfel, P., Guérin, N., Csurka, G., et al.: Large-scale localization datasets in crowded indoor spaces. In: *CVPR* (2021)
28. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International journal of computer vision* **60**(2), 91–110 (2004)
29. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5188–5196 (June 2015)
30. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
31. Moon, H., Lee, C., Hong, J.H.: Efficient privacy-preserving visual localization using 3d ray clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9773–9783 (2024)
32. Ng, T., Kim, H.J., Lee, V.T., DeTone, D., Yang, T.Y., Shen, T., Ilg, E., Balntas, V., Mikolajczyk, K., Sweeney, C.: Ninjadesc: content-concealing visual descriptors via adversarial learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12797–12807 (2022)
33. Novak, R., Bahri, Y., Abolafia, D.A., Pennington, J., Sohl-Dickstein, J.: Sensitivity and generalization in neural networks: an empirical study. *arXiv preprint arXiv:1802.08760* (2018)
34. Pan, L., Schönberger, J.L., Larsson, V., Pollefeys, M.: Privacy preserving localization via coordinate permutations. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 18174–18183 (2023)
35. Pietrantoni, M., Csurka, G., Sattler, T.: Gaussian splatting feature fields for (privacy-preserving) visual localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1082–1092 (June 2025)
36. Pietrantoni, M., Humenberger, M., Sattler, T., Csurka, G.: Segloc: Learning segmentation-based representations for privacy-preserving visual localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15380–15391 (2023)
37. Pietrantoni, M., Humenberger, M., Sattler, T., Csurka, G.: Can we make nerf-based visual localization privacy-preserving? *arXiv preprint arXiv:2508.18971* (2025)

38. Pittaluga, F., Koppal, S.J., Kang, S.B., Sinha, S.N.: Revealing scenes by inverting structure from motion reconstructions. In: CVPR (2019)
39. Pittaluga, F., Zhuang, B.: Ldp-feat: Image features with local differential privacy. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17580–17590 (2023)
40. Sarlin, P.E., Dusmanu, M., Schönberger, J.L., Speciale, P., Gruber, L., Larsson, V., Miksik, O., Pollefeys, M.: Lamar: Benchmarking localization and mapping for augmented reality. In: ECCV (2022)
41. Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A.: Scene coordinate regression forests for camera relocalization in rgb-d images. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2930–2937 (2013)
42. Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 567–576 (2015)
43. Song, X., Tang, J., Yang, K., Guo, W., Lu, X., Hei, X.: A method for absolute pose regression based on cascaded attention modules. *Computer Vision and Image Understanding* **259**, 104440 (2025)
44. Speciale, P., Schonberger, J.L., Kang, S.B., Sinha, S.N., Pollefeys, M.: Privacy preserving image-based localization. In: CVPR (2019)
45. Speciale, P., Schonberger, J.L., Sinha, S.N., Pollefeys, M.: Privacy preserving image queries for camera localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1486–1496 (2019)
46. Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction {APIs}. In: 25th USENIX security symposium (USENIX Security 16). pp. 601–618 (2016)
47. Wang, F., Jiang, X., Galliani, S., Vogel, C., Pollefeys, M.: Glace: Global local accelerated coordinate encoding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21562–21571 (2024)
48. Zhou, L., Koppel, D., Kaess, M.: A complete, accurate and efficient solution for the perspective-n-line problem. *IEEE Robotics and Automation Letters* **6**(2), 699–706 (2020)
49. Zhou, Q., Agostinho, S., Ošep, A., Leal-Taixé, L.: Is geometry enough for matching in visual localization? In: European Conference on Computer Vision. pp. 407–425. Springer (2022)