

# Cov2Pose: Leveraging Spatial Covariance for Direct Manifold-aware 6-DoF Object Pose Estimation

Nassim Ali Ousalah\*, Peyman Rostami\*, Vincent Gaudillière<sup>†</sup>

Emmanuel Koumandakis<sup>×</sup>, Anis Kacem\*, Enjie Ghorbel\*<sup>‡</sup> and Djamilia Aouada\*

CVI<sup>2</sup>, Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg\*  
Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France<sup>†</sup>, Infinite Orbits<sup>×</sup>

Cristal Laboratory, ENSI, University of Manouba<sup>‡</sup>

{nassim.alioussalah, peyman.rostami, anis.kacem, djamilia.aouada}@uni.lu

enjie.ghorbel@isamm.uma.tn, manos@infiniteorbits.io, vincent.gaudilliere@loria.fr

## Abstract

*In this paper, we address the problem of 6-DoF object pose estimation from a single RGB image. Indirect methods that typically predict intermediate 2D keypoints, followed by a Perspective- $n$ -Point solver, have shown great performance. Direct approaches, which regress the pose in an end-to-end manner, are usually computationally more efficient but less accurate. However, direct pose regression heads rely on globally pooled features, ignoring spatial second-order statistics despite their informativeness in pose prediction. They also predict, in most cases, discontinuous pose representations that lack robustness. Herein, we therefore propose a covariance-pooled representation that encodes convolutional feature distributions as a symmetric positive definite (SPD) matrix. Moreover, we propose a novel pose encoding in the form of an SPD matrix via its Cholesky decomposition. Pose is then regressed in an end-to-end manner with a manifold-aware network head, taking into account the Riemannian geometry of SPD matrices. Experiments and ablations consistently demonstrate the relevance of second-order pooling and continuous representations for direct pose regression, including under partial occlusion.*

## 1. Introduction

Estimating the six-degree-of-freedom (6-DoF) pose of an object from an RGB image is a core challenge in computer vision [3, 4, 33, 43, 44, 49, 51, 74] due to its wide range of applications in robotic manipulation [31, 69], augmented and virtual reality [85], autonomous navigation [11, 64], and industrial inspection [24, 37, 53]. Given an image, the goal is to predict the pose  $P \in \text{SE}(3)$  of the object of interest relative to the camera that captured the image, where  $\text{SE}(3)$  denotes the Special Euclidean Group of  $\mathbb{R}^3$ .

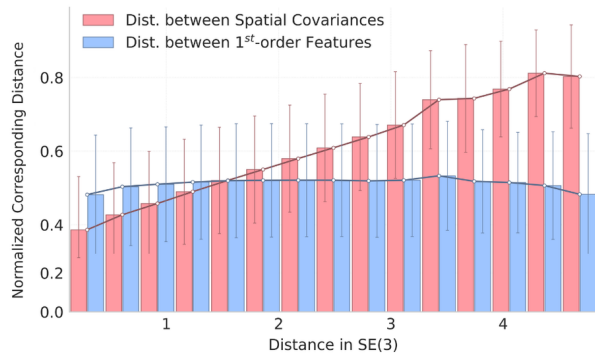


Figure 1. **Pose-representation alignment.** Pairs of images<sup>1</sup> are sorted according to the geodesic distances in  $\text{SE}(3)$  between their corresponding ground-truth poses ( $x$ -axis). The  $x$ -axis is then partitioned into bins. For each bin, we plot: (i) the average Log-Euclidean distance between spatial covariances of image pairs within this bin (red) and (ii) the average Cosine distance between flattened features of image pairs within the bin (blue). Covariance distances increase with pose distance, whereas flattened feature distances remain nearly constant, indicating that spatial covariance is more correlated with the pose.

Thanks to the tremendous developments made in deep learning, numerous approaches have leveraged Deep Neural Networks (DNNs) [36, 50, 85, 88] either directly or indirectly to estimate the 6-DoF object pose. Indirect methods mainly predict intermediate representations such as keypoints [49, 66, 86, 88] or dense object coordinates [80, 89] that are used to solve a Perspective- $n$ -Point ( $PnP$ ) problem or adopt a render-and-compare approach that optimizes photometric, silhouette, or differentiable rendering losses to align a rendered model view with the input im-

<sup>1</sup>Training images from the `002_master_chef_can` class of the YCB-V [87] dataset. We use an EfficientNet-B6 [73] pretrained on ImageNet [15] with no prior training on pose estimation.

age [42, 45, 46, 54]. Direct methods, on the other hand, regress rotation and translation in an end-to-end manner using a neural architecture [13, 16, 47, 52, 77, 79, 87]. While indirect methods have been demonstrated to be more accurate, they suffer from a high calculation cost induced by the iterative computation they incur [45]. In contrast, direct methods are usually more efficient despite being less precise [12, 13, 79], thereby suggesting their practical relevance in real-time applications.

In this paper, we argue that we can improve the accuracy of direct methods by incorporating the spatial second-order statistics of the input image within the DNN-based pose estimator. In fact, deep pose regressors typically rely solely on first-order statistics of feature activations. Nonetheless, second-order moments such as covariance matrices can capture pairwise co-activations that average or max pooling operations discard, thereby providing information about the spatial variability within an image. We hypothesize that such information is highly relevant for inferring the pose. This hypothesis is confirmed by the experiment reported in Figure 1. This figure shows that pairwise distances between spatial covariances of DNN features increase with ground-truth SE(3) pose distance, whereas distances between flattened features remain nearly constant. This suggests that spatial covariance is strongly correlated with pose and provides, as a result, a more suitable representation for direct pose regression.

Concretely, we propose to make use of second-order pooling, which has shown its utility in several deep learning-based computer vision tasks such as classification, facial expression recognition, object detection, and action recognition [2, 10, 20, 22, 25, 34, 81, 82]. This pooling strategy returns feature covariance matrices that are known to be symmetric positive-definite (SPD), ensuring the exploitation of statistical information beyond first-order within a DNN, while preserving spatial awareness. To the best of our knowledge, we are the first to explore spatial covariance representations for end-to-end pose regression.

Our framework, termed `Cov2Pose`, is an end-to-end trainable pipeline that respects the geometry of the SPD manifold where covariance matrices reside. Moreover, we introduce within `Cov2Pose` a one-to-one continuous mapping from the SPD space to SE(3) via a differentiable Cholesky-based encoding without breaking the chain rule, allowing the direct prediction of a continuous 6-DoF pose representation. In particular, the pose is predicted under the shape of a continuous 6D rotation parametrization coupled with a 3D translation. This chosen rotation representation has been shown to be more suitable for rotation learning [90]. Experiments performed on several object pose benchmarks show the superiority of `Cov2Pose` compared to previous direct methods, demonstrating the relevance of the proposed principled framework.

In summary, our main contributions can be summarized as follows,

- The first covariance-based deep learning framework for RGB-based end-to-end 6-DoF object pose regression, which leverages the spatial covariance of backbone features to encode higher-order statistics.
- A manifold-aware training pipeline that applies geometry-preserving dimensionality reduction through Bilinear Mapping layers (BiMap) [32] to a compact SPD representation.
- A fully differentiable, CAD-free, one-to-one, and continuous pose regressor that maps the latent SPD matrix to a continuous 6D rotation representation and translation via a differentiable Cholesky decomposition.
- Extensive experiments on three object pose benchmarks, namely, LineMOD [26], Occ-LineMOD [7] and YCB-Video [87], showing that `Cov2Pose` achieves state-of-the-art results as compared to direct regression methods.

The remainder of the paper is organized as follows: Section 2 presents the related work on 6-DoF object pose estimation and the background. Section 3 formulates the problem of direct pose regression. Section 4 describes the proposed framework entitled `Cov2Pose`. Section 5 reports the conducted experiments and discusses the obtained results. Finally, Section 6 concludes this work, highlighting potential future directions.

## 2. Related Work and Background

In this section, we first review recent literature on RGB-based 6-DoF object pose estimation, then discuss the use of covariance representations in vision tasks. We also provide preliminaries on the SPD manifold, where such representations naturally reside, and review deep learning methods operating on these manifolds.

**RGB-based 6-DoF Object Pose Estimation.** The field has progressed along three complementary paradigms. *PnP-based methods* first establish 2D–3D correspondences and then solve a  $PnP$  problem [19]. Early keypoint-based works [65, 68, 75] introduced a learnable regressor for pre-defined sparse object landmarks coordinates or heatmaps. PV-Net [66] proposed to leverage all pixels through a voting system for more robust regression. More recent approaches proposed to learn the keypoints [86], or to refine keypoint prediction with a diffusion-driven framework [88]. A visibility-aware mechanism has also been proposed to handle occlusions in [49]. Dense-correspondence variants aim to improve pose accuracy by predicting per-pixel matches across the whole object surface. Early examples are DPoD [89] and the normalized object coordinate space method [80]. More recently, ZebraPose [72] introduced a coarse-to-fine surface encoding and CheckerPose [48] adopted progressive dense keypoint localisation. However,

these methods generally suffer from a costly iterative outlier rejection process. **Render-and-compare methods** iteratively render the object under a hypothesised pose and refine it against the observed image. The earliest contribution, DeepIM [45], introduced an iterative refinement scheme that directly optimises the pose by minimising photometric error. Manhardt *et al.* extended this idea to a model-based RGB refinement pipeline [54]. RePose [35] later demonstrated fast pose refinement via deep texture rendering. Subsequently, CosyPose added multi-view consistency to improve robustness [42], and MRC-Net further enhanced the matching-refinement cycle with multiscale residual correlation [46]. However, these methods require a 3D object model to render and rely on an expensive iterative optimization loop. **Direct pose-regression methods** aim to predict the 6-DoF object pose in a single forward pass. Early end-to-end regressors such as PoseCNN [87] and Deep-6DPose [16] demonstrated that a CNN can output pose parameters directly from RGB cues. Subsequent work introduced more structured regressors: CDPN employed a coordinates-based disentangled network [47], DenseFusion fused RGB and depth features before regression [77], and GDR-Net presented a geometry-guided direct regression architecture [79], extended in [52]. Probabilistic PnP layers were incorporated in EPro-PnP to model uncertainty and enable end-to-end training through intermediate 2D-3D correspondence predictions [13]. Most prior methods overlook higher-order statistics such as covariance matrices, which, as shown in Figure 1, carry valuable information for pose regression. These neglected statistics are correlated with object poses and thus offer potential for improving direct pose estimation. Furthermore, existing approaches often rely on non-continuous rotation representations, which, as demonstrated in prior studies [90], hinder stable and effective pose learning.

### Covariances and Deep Learning on the SPD Manifold.

Covariance descriptors have long been employed to encode visual information across a wide range of computer vision tasks [5, 58, 76]. By aggregating second-order statistics of local features, covariance matrices capture complementary texture, shape, and color cues while remaining compact and robust to noise. Because these descriptors are symmetric positive-definite, they reside on the Riemannian manifold of SPD matrices. Formally, this manifold is defined as,

$$\mathcal{S}_{++}^n = \{ \mathbf{M} \in \mathbb{R}^{n \times n} : \mathbf{M} = \mathbf{M}^\top, \mathbf{x}^\top \mathbf{M} \mathbf{x} > 0 \ \forall \mathbf{x} \in \mathbb{R}^n \setminus \{0\} \}, \quad (1)$$

where  $^\top$  denotes matrix transpose. This space lacks the linear structure of a vector space. Consequently, standard operations such as linear combinations are not directly applicable, as they do not necessarily preserve SPD nature. As such, algorithms should be tailored to respect the intrinsic

non-linear geometry of the manifold. Since standard convolutional or fully-connected layers presume a Euclidean feature space, researchers have introduced SPD-aware building blocks to embed the geometry into end-to-end trainable pipelines [8, 21, 32, 83]. These SPD deep networks have demonstrated state-of-the-art performance in tasks ranging from face [17] and action recognition [57] to medical image analysis [55]. Nevertheless, most of them have targeted classification tasks, while regression problems, such as pose estimation, have been overlooked. In this paper, we are the first to exploit SPD-aware deep learning for pose estimation. We briefly recall below the core contributions that have led to the development of SPD Deep Learning frameworks:

- **BiMap (Bilinear Mapping).** Similar to projection layers in Euclidean nets, a BiMap layer compresses the representation to generate more compact and discriminative outputs. With a column-orthonormal weight matrix  $\mathbf{W} \in \mathcal{V}_n(\mathbb{R}^m) = \{ \mathbf{W} \in \mathbb{R}^{n \times m} \mid \mathbf{W}^\top \mathbf{W} = \mathbf{I}_m \}$  ( $\mathbf{I}_m$  being the identity matrix of size  $m$ , and  $\mathcal{V}_n(\mathbb{R}^m)$  denotes the Stiefel manifold), it transports the input covariance by congruence,

$$\mathbf{Y} = \mathbf{W} \mathbf{X} \mathbf{W}^\top, \quad \mathbf{X} \in \mathcal{S}_{++}^n, \mathbf{Y} \in \mathcal{S}_{++}^m, \quad (2)$$

This contracts  $n \rightarrow m$  ( $m < n$ ) while preserving the SPD structure. During training,  $\mathbf{W}$  is kept on the Stiefel manifold [32] via a QR retraction at each update.

- **ReEig (Eigenvalue Rectification).** As for ReLU in ConvNets, ReEig introduces nonlinearity while preserving the SPD structure by lifting small eigenvalues to a spectral floor. Let the  $k$ -th layer input be  $\mathbf{X} \in \mathcal{S}_{++}^n$  with eigendecomposition  $\mathbf{X} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{U}^\top$ . This operation is defined such that

$$\text{ReEig}_\varepsilon(\mathbf{X}) = \mathbf{U} \max(\boldsymbol{\Sigma}, \varepsilon \mathbf{I}) \mathbf{U}^\top, \quad \varepsilon > 0, \quad (3)$$

where  $\max$  acts element-wise on the diagonal,  $\varepsilon$  sets the floor, and  $\mathbf{I}$  is the identity matrix. This keeps the output in  $\mathcal{S}_{++}^n$  and prevents the collapse of small modes.

### 3. Problem Statement

Estimating the 6-DoF pose of a given object from a sample  $\mathbf{I} \in \mathcal{I}$  belonging to the space of RGB images  $\mathcal{I}$  involves finding a mapping  $\Phi$  such that,

$$\begin{aligned} \Phi : \mathcal{I} &\longrightarrow \mathcal{P} \\ \mathbf{I} &\longmapsto \Phi(\mathbf{I}) = \mathbf{P}, \end{aligned} \quad (4)$$

where  $\mathbf{P} \in \mathcal{P}$  is the pose of the object with respect to the camera coordinate system and  $\mathcal{P}$  is the pose representation space. Given a set of training samples  $\mathcal{D} = \{ \mathbf{I}_i, \mathbf{P}_i \}_{i=1}^D$ , where  $\mathbf{I}_i \in \mathcal{I}$  represents an image sample and  $\mathbf{P}_i \in \mathcal{P}$  its associated pose, direct methods typically train a neural

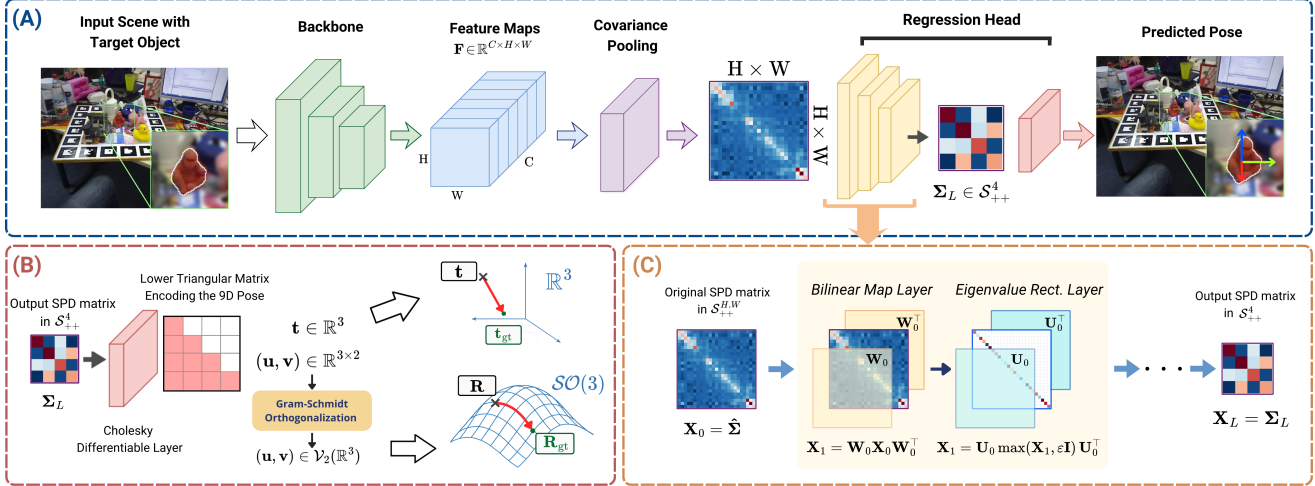


Figure 2. **Overview of the proposed pose estimation framework** (best viewed in color). **(A) General overview.** A CNN backbone extracts feature maps  $\mathbf{F} \in \mathbb{R}^{C \times H \times W}$ . Spatial covariance pooling over the  $H \times W$  grid gives  $\hat{\Sigma} \in \mathcal{S}_{++}^{H \times W}$  which is fed to the pose head. **(B) Learning the pose.** After  $L$  geometry-aware layers, the compact SPD matrix  $\Sigma_L \in \mathcal{S}_{++}^4$  is decoded by a differentiable Cholesky layer into  $(\mathbf{u}, \mathbf{v}; \mathbf{t})$ ;  $(\mathbf{u}, \mathbf{v})$  are mapped to  $\mathcal{SO}(3)$  and the loss is computed on  $\mathcal{SO}(3) \times \mathbb{R}^3$ . **(C) SPD learning layers.** Each layer applies a BiMap followed by eigenvalue rectification  $\text{ReEig}_\varepsilon$  [32] preserving strict positive definiteness.

network  $\Phi_\theta$  parameterized by weights  $\theta$  to minimize a pose loss denoted as  $\mathcal{L}_{\text{pose}}$ ,

$$\theta^* = \arg \min_{\theta} \sum_{i=1}^D \mathcal{L}_{\text{pose}}(\Phi_\theta(\mathbf{I}_i), \mathbf{P}_i). \quad (5)$$

In turn, the learned neural network  $\Phi_\theta$  can be seen as the composition of two functions  $\Phi_\theta = \Psi_{\theta_2} \circ \Gamma_{\theta_1}$  where  $\Gamma_{\theta_1} : \mathcal{I} \mapsto \mathcal{F}$  is the feature extractor that maps an input image into the feature space  $\mathcal{F}$  and  $\Psi_{\theta_2} : \mathcal{F} \mapsto \mathcal{P}$  that maps extracted features to the object pose. End-to-end methods mainly use: (i) standard features that are extracted by DNNs using  $\Gamma_{\theta_1}$ , ignoring second-order statistics that can be highly correlated with the pose, as suggested by the experiment presented in Figure 1; and (ii) a non-continuous representation space  $\mathcal{P}$ , as regressing continuous representations using  $\Psi_{\theta_2}$  necessitates a manifold-aware learning strategy that is not straightforward to implement. In fact, most direct approaches regress the rotation either using quaternions or Euler angles that are non-continuous [90], inducing the non-continuity of the pose. Nevertheless, as demonstrated in [23], the use of continuous representations via Gram-Schmidt orthonormalization or Singular Value Decomposition (SVD) is more suitable for learning rotations. This paper aims at addressing the two aforementioned issues by imposing the extraction of second-order statistics using  $\Gamma_{\theta_1}$  and constraining that  $\Psi_{\theta_2}$  maps to continuous pose representations in an end-to-end way.

## 4. Proposed Approach

To address the problem outlined in Section 3, we introduce  $\text{Cov2Pose}$ , the first covariance-based framework for end-to-end 6-DoF object pose regression that directly outputs a continuous pose representation, namely, 6D representation for rotation coupled with a translation vector. In particular, the initial feature maps extracted from the input image using a CNN backbone are first encoded in a covariance matrix using a second-order pooling. Since covariance matrices are SPD, trainable BiMap layers are employed to progressively reduce their dimensionality while taking into account the geometry of the SPD manifold. As such, the succession of these two steps forms a feature extraction  $\Gamma : \mathcal{I} \mapsto \mathcal{S}_{++}^n$  based on second-order statistics. Finally, a differentiable Cholesky layer is added to decompose the reduced covariance matrix into two triangular matrices. The obtained matrices encode, in their non-zero elements, the 6D representation of the rotation, as well as the translation vector, allowing a continuous and injective mapping  $\Psi : \mathcal{S}_{++}^n \mapsto \mathcal{P}$ . The full pipeline of our solution is illustrated in Figure 2. In the following, we start by describing the covariance-based feature extractor  $\Gamma$  (Section 4.1). We then depict the proposed differentiable mapping  $\Psi$  from the resulting SPD matrix to the pose using Cholesky decomposition (Section 4.2). Finally, the training strategy is presented in (Section 4.3).

### 4.1. Covariance-based Feature Extraction

We first extract from an input image  $\mathbf{I}$  of the scene, containing the target object, a set of feature maps  $\Gamma_1(\mathbf{I}) = \mathbf{F} \in \mathbb{R}^{C \times H \times W}$  using a CNN-based feature extrac-

tor  $\Gamma_1 : \mathcal{I} \mapsto \mathbb{R}^{C \times H \times W}$ , where  $C$ ,  $H$ , and  $W$  are the number of channels, height, and width, respectively. To include the second-order statistics in our framework, we propose incorporating a second-order pooling layer that computes a spatial covariance matrix  $\Gamma_2 : \mathbb{R}^{C \times H \times W} \mapsto \mathcal{S}_{++}^N$  of the extracted features  $\mathbf{F}$ , with  $N = H \times W$ . In particular, the covariance matrix  $\hat{\Sigma}$  is estimated as follows,

$$\begin{aligned} \hat{\Sigma} &= \Gamma_2(\mathbf{F}) = \text{CovPool}(\mathbf{X}) \\ &= \frac{1}{C-1} \sum_{i=1}^C (\mathbf{X}_i - \boldsymbol{\mu}_{\mathbf{X}})^\top (\mathbf{X}_i - \boldsymbol{\mu}_{\mathbf{X}}), \end{aligned} \quad (6)$$

where  $\mathbf{X} = \text{vec}(\mathbf{F}) \in \mathbb{R}^{C \times N}$  corresponds to the flattened feature maps over the spatial dimensions and  $\boldsymbol{\mu}_{\mathbf{X}} = \frac{1}{C} \sum_{i=1}^C \mathbf{X}_i$  such that  $\mathbf{X}_i \in \mathbb{R}^N$  is the flattened feature map resulting from the  $i^{\text{th}}$  channel. Figure 3 illustrates the spatial covariance pooling. This covariance-based representation offers a key advantage in the context of 6-DoF pose estimation. It models how spatial locations “co-vary” across the feature maps, a statistic that changes systematically according to the image viewpoint. As discussed in Figure 1, a clear correlation between the spatial covariance and the pose exists as compared to standard feature maps, making them highly informative for direct pose regression.

To regress the pose, it is necessary to reduce beforehand the dimensionality of the resulting high-dimensional covariance matrix. Nonetheless, covariance matrices in (6) lie on the Riemannian manifold of SPD matrices,  $\mathcal{S}_{++}^N$ . Hence, applying standard deep learning reduction layers is not adequate, as it does not account for their non-linear structure, resulting in non-SPD outputs. In fact, standard DNNs assume that the nature of the feature space is Euclidean. Consequently, considering deep dimensionality reduction layers specifically tailored to the SPD manifold is essential. For that reason, we propose to employ  $L$  BiMap layers [32], which allow for reducing the dimension of the covariance matrix while maintaining its geometric structure, as presented in (2). Moreover, a ReEig [32] operation discussed in (3) is also applied to avoid singularities, resulting in symmetric positive semi-definite matrices. Mathematically, the final covariance-based features  $\Sigma_L = \Gamma_3(\Sigma)$  with  $\Gamma_3 : \mathcal{S}_{++}^N \mapsto \mathcal{S}_{++}^n$  and  $n \ll N$  are obtained as follows,

$$\begin{aligned} \Sigma_{l+1} &= \text{ReEig}_\varepsilon(\mathbf{W}_l^\top \mathbf{X}_l \mathbf{W}_l), \\ \Sigma_0 &= \hat{\Sigma}, \end{aligned} \quad (7)$$

for  $l \in \{0, \dots, L-1\}$ . In summary, the covariance-based feature extractor can be written as  $\Gamma = \Gamma_3 \circ \Gamma_2 \circ \Gamma_1$ . In our experiments, we set  $\varepsilon$  of the ReEig layer to  $10^{-4}$ .

## 4.2. From SPD to Pose Parametrization

Given a reduced covariance based descriptor  $\Sigma_L = \Gamma(\mathbf{I}) \in \mathcal{S}_{++}^n$  extracted from an image  $\mathbf{I}$ , the




Figure 3. **Spatial Covariance Pooling.** Extracted features are flattened and centered and 2<sup>nd</sup> order statistics are constructed pixel-wise. Every entry of the covariance matrix describes how two spatial locations of the given input co-vary.

goal is to seek a mapping

$$\begin{aligned} \Psi : \mathcal{S}_{++}^n &\longrightarrow \mathcal{P} \\ \Sigma_L &\longmapsto \Psi(\Sigma_L) = \mathbf{P}, \end{aligned} \quad (8)$$

with  $\mathbf{P}$  being the pose associated with  $\mathbf{I}$ .

**Continuity of  $\mathcal{P}$ .** As discussed in Section 3, we favor picking  $\mathcal{P}$  as a continuous space, which will define the nature of the used pose representation. The pose space  $\mathcal{P} = \mathcal{R} \times \mathcal{T}$  can be seen as the direct product of the rotation space denoted as  $\mathcal{R}$ , and the position space  $\mathcal{T}$ . While the use of a simple vector in  $\mathbb{R}^3$  for translation guarantees the continuity of  $\mathcal{T}$ , the continuity of rotation representations is not always ensured, as for Euler angles and quaternions [90]. To ensure continuity, two common alternatives from the literature may be considered, namely the 6D representation with Gram-Schmidt orthonormalization and the 9D Special Orthogonal representation. In the following, we choose to employ the 6D representation with Gram-Schmidt orthonormalization given its relative simplicity. In fact, when learning the 6D representation, only two orthonormal vectors are learned.

As such, the rotation representation space  $\mathcal{R}$  is defined as the Stiefel manifold [1, 14] denoted as  $\mathcal{V}_2(\mathbb{R}^3)$  and therefore the pose space as  $\mathcal{P} = \mathcal{V}_2(\mathbb{R}^3) \times \mathbb{R}^3$ .

**Properties of  $\Psi$ .** The mapping  $\Psi$  defined in (8) should be:

1. **Injective**, providing a unique pose for a given SPD matrix input;
2. **Continuous**, ensuring that two close SPD matrices result in close poses;
3. **Differentiable**, enabling its incorporation within an end-to-end deep learning framework.

The Cholesky decomposition is known to be unique, continuous [70] and differentiable [56, 71], making it an ideal candidate for defining  $\Psi$ . Specifically, it decomposes any SPD matrix into the product of a lower triangular matrix and its conjugate transpose. As a result, the obtained triangular matrix can be used for encoding the selected 6D rotation representation (6-dimensional) as well as the translation vector (3-dimensional). For that purpose, a matrix with more than 9 non-zero elements is needed, necessitating fixing  $n$  to 4. This can be defined as  $\Psi(\Sigma_L) = \mathbf{L}$ , such

that  $\Sigma_L = \mathbf{L}\mathbf{L}^\top \in \mathcal{S}_{++}^4$  and,

$$\mathbf{L} = \begin{bmatrix} e^{t_x} & 0 & 0 & 0 \\ u_1 & e^{t_y} & 0 & 0 \\ u_2 & v_1 & e^{t_z} & 0 \\ u_3 & v_2 & v_3 & e^{-(t_x+t_y+t_z)} \end{bmatrix}, \quad (9)$$

where  $\mathbf{u} = (u_1, u_2, u_3)^\top$  and  $\mathbf{v} = (v_1, v_2, v_3)^\top$ , with  $(\mathbf{u}, \mathbf{v}) \in \mathbb{R}^{3 \times 2}$ , are the two 3D vectors that parameterize the 6D rotation representation and  $\mathbf{t} = (t_x, t_y, t_z)$  is the translation vector. Following [14, 90], we first map  $(\mathbf{u}, \mathbf{v})$  onto the Stiefel manifold  $\mathcal{V}_2(\mathbb{R}^3)$  using a differentiable Gram-Schmidt process, then construct the rotation matrix  $\mathbf{R} \in \mathcal{SO}(3)$  by completing the basis with the cross product of the two orthonormalized vectors. Here,  $\mathcal{SO}(3)$  denotes the Special Orthogonal group of  $\mathbb{R}^3$ . This map is continuous except on the set where  $\mathbf{u}$  and  $\mathbf{v}$  are collinear.

Note that the exponential on the diagonal of  $\mathbf{L}$  guarantees positivity, hence  $\mathbf{L}\mathbf{L}^\top \in \mathcal{S}_{++}^4$ . Furthermore, denoting  $\mathbf{L} = (L_{ij})_{i,j \in \{1,2,3,4\}}$ ,  $L_{44} = e^{-(t_x+t_y+t_z)}$  ensures  $\prod_i L_{ii} = 1$ , so  $\det(\Sigma_L) = \det(\mathbf{L})^2 = 1$ . This normalizes the geometric mean of the SPD eigenvalues to 1 without reducing the expressivity of the pose.

### 4.3. Training Strategy

From  $\mathbf{L}$ , we decode  $\hat{\mathbf{t}} = (\log L_{11}, \log L_{22}, \log L_{33})^\top$ ,  $\hat{\mathbf{u}} = (L_{21}, L_{31}, L_{41})^\top$ , and  $\hat{\mathbf{v}} = (L_{32}, L_{42}, L_{43})^\top$  and map  $(\hat{\mathbf{u}}, \hat{\mathbf{v}})$  to  $\hat{\mathbf{R}} \in \mathcal{SO}(3)$  with the differentiable Gram-Schmidt procedure [90] and supervise the training with the geodesic distance on  $\mathcal{SO}(3)$  and  $\ell_2$  on  $\mathbb{R}^3$  for translation. In addition, we add two regularizers, (i) an orthogonality penalty that keeps  $\langle \hat{\mathbf{u}}, \hat{\mathbf{v}} \rangle \rightarrow 0$  and (ii) a unit-norm penalty that prevents collapse of the two vectors. The total training loss  $\mathcal{L}_{\text{pose}}$  for regressing the pose of the object ground truth pose represented by the rotation matrix  $\mathbf{R}_{\text{gt}}$  and the translation vector  $\mathbf{t}_{\text{gt}}$  can be expressed by,

$$\mathcal{L}_{\text{pose}}(\hat{\mathbf{R}}, \hat{\mathbf{t}}, \hat{\mathbf{u}}, \hat{\mathbf{v}}) = \arccos\left(\frac{\text{tr}(\hat{\mathbf{R}}^\top \mathbf{R}_{\text{gt}}) - 1}{2}\right) + \|\hat{\mathbf{t}} - \mathbf{t}_{\text{gt}}\|_2 + \lambda \left[ \langle \hat{\mathbf{u}}, \hat{\mathbf{v}} \rangle^2 + (\|\hat{\mathbf{u}}\| - 1)^2 + (\|\hat{\mathbf{v}}\| - 1)^2 \right] \quad (10)$$

where  $\lambda$  is a factor that controls the orthonormality regularization term and  $\text{tr}$  indicates the trace. In our experiments, the parameter  $\lambda$  is set to  $10^{-3}$ , which stabilizes the training without dominating the main pose loss.

For a comprehensive comparison, we also evaluate a second variant of the proposed framework with  $n = 3$  in which the Cholesky layer outputs a  $3 \times 3$  lower-triangular matrix with  $\frac{3(3+1)}{2} = 6$  free entries. In this case, the rotation is parameterized by Euler angles, which are non-continuous. Further details are provided in the Supplementary.

## 5. Experiments

In this section, we first discuss our experimental setup. Then, we report the performances of `Cov2Pose` on

three widely-used benchmarks for object pose estimation, namely, LM [26], LM-O [7] and YCB-V [87]. We follow with a series of ablation experiments that further justify the technical choices of our solution.

### 5.1. Datasets

LineMod (LM) [26] is a well-known general pose estimation benchmark. It contains 16k images belonging to 13 different object categories in heavily cluttered scenes. We train and test following the default BOP split and standard procedure [29, 79] and thus employ 85% of the RGB images for testing. LineMod Occlusion (LM-O) [7] has been introduced as a more challenging version of LM. It contains 1214 additional annotated images and 8 object categories under heavy occlusion. Following [79], we use the publicly available physically-based rendering (pbr) images [28] as training images for LM-O. Finally, the YCB-V dataset [87] is a large-scale dataset of 21 object categories and approximately 100k real images. We follow the training strategies of [79] and use both the training images of YCB-V, but also the publicly available pbr images [28].

### 5.2. Metrics

We report our results using the Average Distance of Model Points (ADD) metric [26, 27] and, for symmetric objects, its symmetric variant ADD-S [87]. We compute the mean distance between the model’s points transformed using the ground-truth pose and the predicted pose. If the mean distance is less than 10% of the diameter of the model, the predicted pose is considered correct. The metric is then the proportion of correct poses. On the YCB-V dataset, additional results are presented in terms of AUC (Area Under the Curve) of ADD-S and AUC of ADD(-S). For computing AUC, we set the maximum distance threshold at 10 cm [87].

### 5.3. Implementation & Training Details

`Cov2Pose` is trained end-to-end with a mixed-geometry optimizer to respect manifold constraints. We also define an optimizer with separate groups to partition trainable parameters into (i) Stiefel-constrained BiMap weights  $\{\mathbf{W}_l\}$  from (7) and (ii) Euclidean parameters for the backbone. The Stiefel parameters are updated with a Riemannian step, which consists of a gradient projection and QR retraction [32] using an initial learning rate  $10^{-2}$ , while the Euclidean parameters use Adam [41] with an initial learning rate  $10^{-4}$ . We use a `ReduceLRonPlateau` scheduler [63] that halves the learning rates after 4 epochs without improvement on the validation loss. For feature extraction, we use an ImageNet-pretrained [15, 73] EfficientNet-B6 backbone [73] from the Pytorch library. Its output resolution is  $H=W=17$ , resulting in a  $289 \times 289$  spatial covariance matrix that encodes co-activations and feeds the pose head. The SPD head is constructed by 4 BiMap in-

Table 1. **LM** results in terms of ADD(-S). Best result in **bold**; second best underlined. *Cov2Pose* (Ours) outperforms all non-PnP methods and slightly surpasses PnP-based ones. “—” indicates values not reported in the original paper. Classes in *italics* are symmetric.

Method		Ape	Bvise.	Cam	Can	Cat	Drill	Duck	<i>Eggbox</i>	<i>Glue</i>	Holep.	Iron	Lamp	Phone	Avg. $\uparrow$
Indirect	PVNet [66]	43.62	<b>99.90</b>	86.86	95.47	79.34	96.43	52.58	<u>99.15</u>	95.66	81.92	98.88	<u>99.33</u>	92.41	86.27
	CheckerPose [48]	—	—	—	—	—	—	—	—	—	—	—	—	—	<u>97.1</u>
	VAPO [49]	—	—	—	—	—	—	—	—	—	—	—	—	—	<u>97.1</u>
End-to-End	Pix2Pose [59]	58.1	91.0	60.9	84.4	65.0	76.3	43.8	96.8	79.4	74.8	83.4	82.0	45.0	72.4
	Self6D++ [78]	76.0	91.6	<b>97.1</b>	<b>99.8</b>	85.6	98.8	56.5	91.0	92.2	35.4	<b>99.5</b>	97.4	91.8	85.6
	DeepIM [45]	<u>77.0</u>	97.5	93.5	96.5	82.1	95.0	77.7	97.1	<u>99.4</u>	52.8	98.3	97.5	87.7	88.6
	BPnP [12] <sup>†</sup>	74.75	<u>99.00</u>	<u>96.25</u>	98.00	<u>94.25</u>	<b>99.25</b>	<u>78.50</u>	96.50	90.00	<u>91.50</u>	97.75	<b>99.75</b>	<u>97.00</u>	93.27
	EPro-PnP [13] <sup>†</sup>	—	—	—	—	—	—	—	—	—	—	—	—	—	95.80
	<b>Ours</b>	<b>94.8</b>	<b>99.9</b>	86.0	<u>99.2</u>	<b>99.2</b>	<u>98.9</u>	<b>89.1</b>	<b>100</b>	<b>100</b>	<b>99.5</b>	<u>99.2</u>	98.2	<b>100</b>	<b>97.2</b>

<sup>†</sup> Method is End-to-End by using a differentiable PnP strategy.

terleaved with 4 ReEig layers. The number of parameters of the model is 41.4M. We use 4 NVIDIA A100 GPUs (40 GB) with batch size 8 to train for 30 epochs on LM and LM-O, and for 20 epochs on YCB-V. In all datasets, 90% of the training data is used for training and 10% for validation. Testing is done with the best checkpoint on validation set.

#### 5.4. Comparison with State-of-the-art Methods

**Results on LM.** Table 1 summarizes LM results. End-to-end methods predict pose directly from the input image. Indirect methods first estimate 2D–3D correspondences and then apply a PnP/RANSAC step. *Cov2Pose* surpasses all end-to-end baselines, including DeepIM [45], and exceeds PnP-based methods by 0.1 ADD(-S).

**Results on LM-O.** As shown in Table 2, we compare our approach with representative end-to-end methods trained per object and also include recent correspondence-based pose estimation techniques. *Cov2Pose* outperforms all end-to-end pose learning baselines and closes the gap to leading PnP-based approaches such as ZebraPose [72] under the same occlusion setup.

Table 2. **LM-O** results in terms of ADD(-S). Group-wise best in bold, group-wise second underlined. *Cov2Pose* (Ours) outperforms all No-PnP methods and is slightly behind ZebraPose [72] with only a 0.1% difference.

Method		Ape	Can	Cat	Drill	Duck	<i>Eggb.</i>	<i>Glue.</i>	Holep.	Avg. $\uparrow$
Indirect	ZebraPose [72]	57.9	95.0	60.6	<u>94.8</u>	64.5	<u>70.9</u>	<u>88.7</u>	83.0	76.9
	CheckerPose [48]	<u>58.3</u>	<u>95.7</u>	<u>62.3</u>	93.7	<b>69.9</b>	70.0	86.4	83.8	77.5
	VAPO [49]	—	—	—	—	—	—	—	—	<u>78.02</u>
	6D-Diff [88]	<b>60.6</b>	<b>97.9</b>	<b>63.2</b>	<b>96.6</b>	<u>67.2</u>	<b>73.5</b>	<b>92.0</b>	<b>85.5</b>	<u>79.6</u>
End-to-End	PoseCNN [87]	9.6	45.2	0.93	41.4	19.6	22.0	38.5	22.1	24.9
	Pix2Pose [59]	—	—	—	—	—	—	—	—	36.3
	Single-Stage [30] <sup>†</sup>	19.2	65.1	18.9	69.0	25.3	52.0	51.4	45.6	43.3
	DeepIM [45]	59.2	63.5	26.2	55.6	<u>52.4</u>	<b>63.0</b>	71.7	52.5	55.5
	Self6D++ [78]	57.7	<u>95.0</u>	<u>52.6</u>	90.5	26.7	45.0	<u>87.1</u>	23.5	59.8
	GDR-Net [79] <sup>†</sup>	46.8	90.8	40.5	<u>82.6</u>	46.9	54.2	75.8	<u>60.1</u>	<u>62.2</u>
	<b>Ours</b>	<b>64.0</b>	<b>96.4</b>	<b>74.1</b>	<b>96.3</b>	<b>55.8</b>	<u>57.9</u>	<b>93.6</b>	<b>76.7</b>	<b>76.8</b>

**Results on YCB-V.** Table 3 summarizes ADD(-S) and AUC averaged over all YCB-V objects. *Cov2Pose* attains the

best ADD(-S) among end-to-end methods and the second-best AUC after GDR-Net [79], trailing by 1.6 AUC(ADD-S). PnP-based methods still lead overall, but *Cov2Pose* reduces the gap w.r.t. the best PnP-based method, reducing it by 2.3 AUC(ADD-S) and 5.7 AUC(ADD(-S)).

Table 3. **YCB-V** results in terms of ADD(-S), AUC of ADD(-S)/ADD(-S). The higher the better. Group-wise best in bold, group-wise second underlined.

Method		ADD(-S)	AUC of ADD(-S)	AUC of ADD(-S)
Indirect	PVNet [66]	—	—	73.4
	ZebraPose [72]	80.5	90.1	85.3
	CheckerPose [48]	81.4	91.3	86.4
	VAPO [49]	<b>84.9</b>	<b>92.3</b>	<b>87.9</b>
	6D-Diff [88]	<u>83.8</u>	<u>91.5</u>	<u>87.0</u>
End-to-End	Pix2Pose [59]	45.7	—	—
	Single-Stage [30] <sup>†</sup>	53.9	—	—
	PoseCNN [87]	—	75.9	61.3
	DeepIM [45]	—	88.1	81.9
	GDR-Net [79] <sup>†</sup>	<u>60.1</u>	<b>91.6</b>	<b>84.4</b>
	<b>Ours</b>	<b>69.7</b>	<u>90.0</u>	<u>82.2</u>

#### 5.5. Ablation Study

We conduct ablations on the LM-O dataset and report performance averaged over all objects in terms of ADD(-S).

**Impact of the Rotation Representation.** In our framework, the predicted pose is embedded in a  $4 \times 4$  Cholesky factor whose off-diagonal entries carry a 6D rotation parameterization  $(\mathbf{u}, \mathbf{v}) \in \mathbb{R}^{3 \times 2}$ . The latter is mapped to  $\mathcal{SO}(3)$  via a differentiable Gram–Schmidt step, which produces a continuous training signal on the rotation manifold  $\mathcal{SO}(3)$ . As discussed at the end of Section 4, we also evaluate a lower-dimensional variant that encodes rotation with Euler angles. In this setting, we regress  $\Sigma_L \in \mathcal{S}_{++}^3$  and decode the pose from its  $3 \times 3$  Cholesky factor. Further details of this configuration are in the supplementary. Table 4 shows that the 6D representation leads to a higher ADD(-S) than Euler angles, reflecting the benefit of continuity.

**Impact of the SPD Head Components.** We ablate the manifold-aware head and the structured pose encoding we

Table 4. Comparison with an alternative rotation parameterization. Under identical settings, we replace our **6D + differentiable GS** rotation with an Euler-angle parameterization encoded as a  $3 \times 3$  SPD. The 6D+GS choice results in higher performance.

Rep.	SPD size	Cont. on $SO(3)$	Avg. ADD(-S) $\uparrow$
Euler angles	$3 \times 3$	$\times$	70.9
<b>6D+GS (Ours)</b>	$4 \times 4$	$\checkmark$	<b>76.8</b>

are using to isolate the effect of each. To investigate this, we design three additional variants. (A) We replace the SPD head with 2 FC layers that map a vectorized version of  $\hat{\Sigma}$  directly to the pose. This assesses the benefit of the SPD layers in learning the compact SPD matrix that is fed to the Cholesky layer. (B) We keep the CovPooling layer, but use it to compute *channel* covariances instead of *spatial* covariance and feed  $\hat{\Sigma}_c$  to the same SPD head. This contrasts the spatial covariance to the channel covariance. (C) We keep the SPD head and drop the Cholesky layer and train using the Frobenius norm between the projection of  $\Sigma_L$  to the logarithmic tangent space and the precomputed corresponding SPD pose label. This assesses the effect of our pose decoding to  $(\mathbf{R}, \mathbf{t})$  and the calculation of the training loss on  $SO(3) \times \mathbb{R}^3$  instead of on the log-tangent space. All other settings are kept fixed. Table 5 shows that our full model attains the best performance. Indeed, the replacement of the SPD head with a Euclidean MLP head causes a sharp drop in accuracy, reflecting a mismatch between the Euclidean hypothesis and the SPD geometry of the representation.

Table 5. Comparison between **Ours**, (A) Euclidean regression head, (B) Channel-wise covariance, and (C) Training on log tangent space of SPD manifold.

Variant	(A)	(B)	(C)	Cov2Pose
Avg. ADD(-S) $\uparrow$	31.0	70.9	72.3	<b>76.8</b>

## 5.6. Inference Time

Inference time is measured on a machine with an AMD 3.30GHz CPU with an NVIDIA A100 GPU. The backbone takes 22.6ms for feature extraction, the covariance-pooling layer 0.5ms to compute the covariance matrix, and the head 23.8ms to produce the compact  $4 \times 4$  representation and regress the pose, for a total of 46.9ms. Table 6 reports inference time for running state-of-the-art methods. Cov2Pose achieves the best trade-off between accuracy and speed compared to previous approaches.

## 6. Conclusion & Future Directions

We present Cov2Pose, an end-to-end framework for direct 6-DoF object pose estimation from a single RGB image. Cov2Pose performs spatial covariance pooling and learns




Figure 4. **Qualitative results.** Green 3D boxes denote GT poses, whereas blue 3D boxes denote predictions. Cov2Pose predicts accurate poses in cluttered scenes and under partial occlusion.

Table 6. Comparison of inference time and ADD(-S) across datasets. We measured inference time of the open-source baselines on the same hardware as ours.

Method	Time (ms)	ADD(-S)		
		LM [26]	LM-O [7]	YCB-V [87]
Indirect	PVNet [66]	44.9	86.27	—
	ZebraPose [72]	119.3	—	76.9
	CheckerPose [48]	116.7	97.1	77.5
	VAPO [49]	69.8	97.1	78.02
End-to-End	Pix2Pose [59]	123	72.4	36.3
	DeepIM [45]	77.3	88.6	55.5
	GDR-Net [79]	16	—	62.2
	<b>Ours</b>	<b>46.9</b>	<b>97.2</b>	<b>76.8</b>

directly on the symmetric positive-definite (SPD) manifold to regress object pose. Concretely, we encode the image as a spatial covariance matrix that captures co-variation among image regions, then compress it into a compact SPD code via a geometry-aware, end-to-end-trainable set of layers. A differentiable Cholesky decomposition acting on the SPD network output results in a continuous 6D rotation representation and a translation recovering the pose. Across 3 BOP benchmarks, Cov2Pose attains state-of-the-art accuracy for end-to-end single-object pose estimation and substantially reduces the gap as compared to correspondence-based approaches. As future directions, we will investigate symmetry-aware training and inference for pose-ambiguous objects within a CAD-free setup, while preserving the fully differentiable end-to-end nature of Cov2Pose. Currently, Cov2Pose does not explicitly handle object symmetries, and addressing this limitation could be an interesting direction for future work. A symmetry-focused pilot study and discussion are provided in the supplementary.

**Acknowledgement.** The present work is supported by the National Research Fund (FNR), Luxembourg, under the C21/IS/15965298/ELITE project, and by Infinite Orbits. Experiments were performed on the Luxembourg national supercomputer MeluXina at LuxProvide.

## References

- [1] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2008. 5
- [2] Dinesh Acharya, Zhiwu Huang, Danda Pani Paudel, and Luc Van Gool. Covariance pooling for facial expression recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 367–374, 2018. 2
- [3] Nassim Ali Ousalah, Anis Kacem, Enjie Ghorbel, Emmanuel Koumandakis, and Djamila Aouada. Uncertainty-aware knowledge distillation for compact and efficient 6dof pose estimation. *International Conference on Intelligent Robots and Systems (IROS)*, 2025. 1
- [4] Nassim Ali Ousalah, Peyman Rostami, Anis Kacem, Enjie Ghorbel, Emmanuel Koumandakis, and Djamila Aouada. Fpg-nas: Flops-aware gated differentiable neural architecture search for efficient 6dof pose estimation. *IEEE International Workshop on Multimedia Signal Processing (MMSP)*, 2025. 1
- [5] Walid Ayedi, Hichem Snoussi, and Mohamed Abid. A fast multi-scale covariance descriptor for object re-identification. *Pattern Recognition Letters*, 33(14):1902–1907, 2012. Novel Pattern Recognition-Based Methods for Re-identification in Biometric Context. 3
- [6] Michele Bechini and Michèle Lavagna. Robust and efficient single-cnn-based spacecraft relative pose estimation from monocular images. *Acta Astronautica*, 233:198–217, 2025. 16
- [7] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *Computer Vision – ECCV 2014*, pages 536–551, Cham, 2014. Springer International Publishing. 2, 6, 8, 13
- [8] Daniel Brooks, Olivier Schwander, Frederic Barbaresco, Jean-Yves Schneider, and Matthieu Cord. Riemannian batch normalization for spd neural networks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. 3
- [9] Dingding Cai, Janne Heikkilä, and Esa Rahtu. Sc6d: Symmetry-agnostic and correspondence-free 6d object pose estimation. In *2022 International Conference on 3D Vision (3DV)*, pages 536–546. IEEE, 2022. 16
- [10] João Carreira, Rui Caseiro, Jorge Batista, and Cristian Sminchisescu. Semantic segmentation with second-order pooling. In *Computer Vision – ECCV 2012*, pages 430–443, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 2
- [11] Lorenzo Pasqualetto Cassinis, Robert Fonod, and Eberhard Gill. Review of the robustness and applicability of monocular pose estimation systems for relative navigation with an uncooperative spacecraft. *Progress in Aerospace Sciences*, 110:100548, 2019. 1
- [12] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8100–8109, 2020. 2, 7
- [13] Hansheng Chen, Pichao Wang, Fan Wang, Wei Tian, Lu Xiong, and Hao Li. Epro-pnp: Generalized end-to-end probabilistic perspective-n-points for monocular object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2781–2790, 2022. 2, 3, 7
- [14] Jiayi Chen, Yingda Yin, Tolga Birdal, Baoquan Chen, Leonidas J. Guibas, and He Wang. Projective manifold gradient layer for deep rotation regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6646–6655, 2022. 5, 6
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 6
- [16] Thanh-Toan Do, Ming Cai, Trung Pham, and Ian Reid. Deep-6dpose: Recovering 6d object pose from a single rgb image. *arXiv preprint arXiv:1802.10367*, 2018. 2, 3
- [17] Zhen Dong, Su Jia, Chi Zhang, Mingtao Pei, and Yuwei Wu. Deep manifold learning of symmetric positive definite matrices with application to face recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), 2017. 3
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 15
- [19] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. 2
- [20] Zilin Gao, Jiangtao Xie, Qilong Wang, and Peihua Li. Global second-order pooling convolutional networks. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 3024–3033, 2019. 2
- [21] Zhi Gao, Yuwei Wu, Yunde Jia, and Mehrtash Harandi. Learning to optimize on spd manifolds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [22] Zilin Gao, Qilong Wang, Bingbing Zhang, Qinghua Hu, and Peihua Li. Temporal-attentive covariance pooling networks for video recognition. In *Advances in Neural Information Processing Systems*, pages 13587–13598. Curran Associates, Inc., 2021. 2
- [23] A. René Geist, Jonas Frey, Mikel Zhobro, Anna Levina, and Georg Martius. Learning with 3d rotations: a hitchhiker’s guide to so(3). In *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org, 2024. 4
- [24] Elena Govi, Davide Sapienza, Samuele Toscani, Ivan Cotti, Giorgia Franchini, and Marko Bertogna. Addressing challenges in industrial pick and place: A deep learning-based 6

- degrees-of-freedom pose estimation solution. *Computers in Industry*, 161:104130, 2024. 1
- [25] Kai Guo, Prakash Ishwar, and Janusz Konrad. Action recognition from video using feature covariance matrices. *IEEE Transactions on Image Processing*, 22(6):2479–2494, 2013. 2
- [26] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Computer Vision – ACCV 2012*, pages 548–562, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 2, 6, 8, 16
- [27] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6d object pose estimation. In *Computer Vision – ECCV 2016 Workshops*, pages 606–619, Cham, 2016. Springer International Publishing. 6
- [28] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. Bop challenge 2020 on 6d object localization. In *European Conference on Computer Vision*, pages 577–594. Springer, 2020. 6
- [29] Tomáš Hodaň, Vibhav Vineet, Ran Gal, Emanuel Shalev, Jon Hanzelka, Treb Connell, Pedro Urbina, Sudipta N. Sinha, and Brian Guenter. Photorealistic image synthesis for object instance detection. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 66–70, 2019. 6
- [30] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2930–2939, 2020. 7, 17
- [31] Hung-Jui Huang, Michael Kaess, and Wenzhen Yuan. Normalflow: Fast, robust, and accurate contact-based object 6dof pose tracking with vision-based tactile sensors. *IEEE Robotics and Automation Letters*, 10(1):452–459, 2025. 1
- [32] Zhiwu Huang and Luc Van Gool. A riemannian network for spd matrix learning. In *Proceedings of the AAAI conference on artificial intelligence*, 2017. 2, 3, 4, 5, 6, 14
- [33] Ziqin Huang, Gu Wang, Chenyangguang Zhang, Ruida Zhang, Xiu Li, and Xiangyang Ji. Givepose: Gradual intra-class variation elimination for rgb-based category-level object pose estimation. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 22055–22066, 2025. 1
- [34] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu. Matrix backpropagation for deep networks with structured layers. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2965–2973, 2015. 2
- [35] Shun Iwase, Xingyu Liu, Rawal Khirodkar, Rio Yokota, and Kris M. Kitani. Repose: Fast 6d object pose refinement via deep texture rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3303–3312, 2021. 3
- [36] Yufeng Jin, Vignesh Prasad, Snehal Jauhri, Mathias Franzius, and Georgia Chalvatzaki. 6dope-gs: Online 6d object pose estimation using gaussian splatting. *IEEE/CVF International Conference on Computer Vision*, 2025. 1
- [37] Agastya Kalra, Guy Stoppi, Dmitrii Marin, Vage Taamazyan, Aarrushi Shandilya, Rishav Agarwal, Anton Boykov, Tze Hao Chong, and Michael Stark. Towards co-evaluation of cameras, hdr, and algorithms for industrial-grade 6dof pose estimation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22691–22701, 2024. 1
- [38] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *2016 IEEE international conference on Robotics and Automation (ICRA)*, pages 4762–4769. IEEE, 2016. 16
- [39] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5974–5983, 2017. 16
- [40] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015. 13, 16
- [41] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [42] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. In *European Conference on Computer Vision*, pages 574–591. Springer, 2020. 2, 3
- [43] Taeyeop Lee, Bowen Wen, Minjun Kang, Gyuree Kang, In So Kweon, and Kuk-Jin Yoon. Any6d: Model-free 6d pose estimation of novel objects. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 11633–11643, 2025. 1
- [44] Ming-Feng Li, Xin Yang, Fu-En Wang, Hritam Basak, Yuyin Sun, Shreekanth Gayaka, Min Sun, and Cheng-Hao Kuo. Ua-pose: Uncertainty-aware 6d object pose estimation and online object completion with partial references. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 1180–1189, 2025. 1
- [45] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 683–698, 2018. 2, 3, 7, 8, 17
- [46] Yuelong Li, Yafei Mao, Raja Bala, and Sunil Hadap. Mrcnet: 6-dof pose estimation with multiscale residual correlation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10476–10486, 2024. 2, 3
- [47] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7677–7686, 2019. 2, 3
- [48] Ruyi Lian and Haibin Ling. Checkerpose: Progressive dense keypoint localization for object pose estimation with graph neural network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14022–14033, 2023. 2, 7, 8
- [49] Ruyi Lian, Yuewei Lin, Longin Jan Latecki, and Haibin Ling. Vapo: Visibility-aware keypoint localization for efficient 6dof object pose estimation. *International Conference on Intelligent Robots and Systems (IROS)*, 2025. 1, 2, 7, 8

- [50] Jiehong Lin, Lihua Liu, Dekun Lu, and Kui Jia. Sam-6d: Segment anything model meets zero-shot 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27906–27916, 2024. 1
- [51] Xingyu Liu, Gu Wang, Ruida Zhang, Chenyangguang Zhang, Federico Tombari, and Xiangyang Ji. Unopose: Unseen object pose estimation with an unposed rgb-d reference image. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 22023–22034, 2025. 1
- [52] Xingyu Liu, Ruida Zhang, Chenyangguang Zhang, Gu Wang, Jiwen Tang, Zhigang Li, and Xiangyang Ji. Gdrnp: A geometry-guided and fully learning-based object pose estimator. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(7):5742–5759, 2025. 2, 3
- [53] Robert Maack, Lars Thun, Thomas Liang, Hasan Tercan, and Tobias Meisen. Pcad: A real-world dataset for 6d pose industrial anomaly detection. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 1132–1141, 2025. 1
- [54] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6d pose refinement in rgb. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 800–815, 2018. 2, 3
- [55] Josef Mayr, Anna Reithmeir, Maxime Di Folco, and Julia A. Schnabel. Covariance descriptors meet general vision encoders: Riemannian deep learning for medical image classification, 2025. 3
- [56] Iain Murray. Differentiation of the cholesky decomposition. *arXiv preprint arXiv:1602.07527*, 2016. 5
- [57] Xuan Son Nguyen, Luc Brun, Olivier Lezoray, and Sebastien Bougleux. A neural network based on spd manifold learning for skeleton-based hand gesture recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [58] Naima Otberdout, Anis Kacem, Mohamed Daoudi, Lahoucine Ballihi, Stefano Berretti, et al. Deep covariance descriptors for facial expression recognition. In *29th British Machine Vision Conference (BMVC'18)*, pages 1–13. BMVA, 2018. 3
- [59] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7668–7677, 2019. 7, 8, 15
- [60] Tae Ha Park and Simone D’Amico. Robust multi-task learning and online refinement for spacecraft pose estimation across domain gap. *Advances in Space Research*, 73(11): 5726–5740, 2024. 16
- [61] Tae Ha Park, Marcus Märtens, Gurvan Lecuyer, Dario Izzo, and Simone D’Amico. Speed+: Next-generation dataset for spacecraft pose estimation across domain gap. In *2022 IEEE aerospace conference (AERO)*, pages 1–15. IEEE, 2022. 13, 16, 17
- [62] Tae Ha Park, Marcus Märtens, Mohsi Jawaid, Zi Wang, Bo Chen, Tat-Jun Chin, Dario Izzo, and Simone D’Amico. Satellite pose estimation competition 2021: Results and analyses. *Acta Astronautica*, 204:640–665, 2023. 16
- [63] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019. 6
- [64] Leo Pauly, Wassim Rharbaoui, Carl Shneider, Arunkumar Rathinam, Vincent Gaudillière, and Djamilia Aouada. A survey on deep learning-based monocular spacecraft pose estimation: Current state, limitations and prospects. *Acta Astronautica*, 212:339–360, 2023. 1
- [65] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G. Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018, 2017. 2
- [66] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pynet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4561–4570, 2019. 1, 2, 7, 8
- [67] Giorgia Pitteri, Michaël Ramamonjisoa, Slobodan Ilic, and Vincent Lepetit. On object symmetries and 6d pose estimation from images. In *2019 International conference on 3D vision (3DV)*, pages 614–622. IEEE, 2019. 16
- [68] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [69] Huan Ren, Wenfei Yang, Shifeng Zhang, and Tianzhu Zhang. Rethinking correspondence-based category-level object pose estimation. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1170–1179, 2025. 1
- [70] Michelle Schatzman and John Taylor. *Numerical Analysis A Mathematical Introduction*. 2023. 5
- [71] Matthias Seeger, Asmus Hetzel, Zhenwen Dai, and Neil Lawrence. Auto-differentiating linear algebra. 2017. 5
- [72] Yongzhi Su, Mahdi Saleh, Torben Fetzter, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6738–6748, 2022. 2, 7, 8
- [73] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 1, 6, 16
- [74] Tao Tan and Qiulei Dong. Onda-pose: Occlusion-aware neural domain adaptation for self-supervised 6d object pose estimation. In *Proceedings of the Computer Vision and Pattern*

- Recognition Conference (CVPR)*, pages 16829–16838, 2025. [1](#)
- [75] Bugra Tekin, Sudipta N. Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [76] Oncel Tuzel, Fatih Porikli, and Peter Meer. Region covariance: A fast descriptor for detection and classification. In *Computer Vision – ECCV 2006*, pages 589–600, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. [3](#)
- [77] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3343–3352, 2019. [2](#), [3](#)
- [78] Gu Wang, Fabian Manhardt, Xingyu Liu, Xiangyang Ji, and Federico Tombari. Occlusion-aware self-supervised monocular 6d object pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(3):1788–1803, 2021. [7](#)
- [79] Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16611–16621, 2021. [2](#), [3](#), [6](#), [7](#), [8](#), [15](#), [16](#), [17](#)
- [80] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2642–2651, 2019. [1](#), [2](#)
- [81] Qilong Wang, Li Zhang, Banggu Wu, Dongwei Ren, Peihua Li, Wangmeng Zuo, and Qinghua Hu. What deep cnns benefit from global covariance pooling: An optimization perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10771–10780, 2020. [2](#)
- [82] Qilong Wang, Jiangtao Xie, Wangmeng Zuo, Lei Zhang, and Peihua Li. Deep cnns meet global covariance pooling: Better representation and generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(8):2582–2597, 2021. [2](#)
- [83] Rui Wang, Xiao-Jun Wu, Ziheng Chen, Tianyang Xu, and Josef Kittler. Dreamnet: A deep riemannian manifold network for spd matrix learning. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 3241–3257, 2022. [3](#)
- [84] Zi Wang, Minglin Chen, Yulan Guo, Zhang Li, and Qifeng Yu. Bridging the domain gap in satellite pose estimation: A self-training approach based on geometrical constraints. *IEEE transactions on aerospace and electronic systems*, 60(3):2500–2514, 2023. [16](#)
- [85] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6d pose estimation and tracking of novel objects. In *CVPR*, 2024. [1](#)
- [86] Yangzheng Wu and Michael Greenspan. Learning better keypoints for multi-object 6dof pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 564–574, 2024. [1](#), [2](#)
- [87] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. 2018. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [16](#), [17](#)
- [88] Li Xu, Haoxuan Qu, Yujun Cai, and Jun Liu. 6d-diff: A keypoint diffusion framework for 6d object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9676–9686, 2024. [1](#), [2](#), [7](#)
- [89] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1941–1950, 2019. [1](#), [2](#)
- [90] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019. [2](#), [3](#), [4](#), [5](#), [6](#), [14](#)

# Cov2Pose: Leveraging Spatial Covariance for Direct Manifold-aware 6-DoF Object Pose Estimation

## Supplementary Material

This supplementary material presents implementation details, detailed ablation studies, a pilot discussion on symmetry-aware training and extended experiments for Cov2Pose. We describe optimization and architecture choices, ablate key components of the proposed method, and report detailed per-object and cross-dataset results, as well as additional validation on Absolute camera Pose Regression (APR) on Cambridge Landmarks [40] and the more challenging spacecraft pose estimation task on SPEED+ [61], further demonstrating the applicability of the proposed method.

### A. Additional Implementation Details

**Ground-truth Preprocessing.** As shown in (9) the 6D rotation and the 3D translation are encoded in the non-zero elements of the lower-triangular matrix  $\mathbf{L}$ . For rotation, we use the full  $\mathbf{R}_{\text{gt}} \in SO(3)$  matrix as the rotation target/objective. For translation, we compute per-axis statistics on the training split, *i.e.*, 99<sup>th</sup>-percentile-based per-axis minima  $\mathbf{t}_{\text{min}}$  and per-axis ranges  $\mathbf{t}_{\text{range}}$ , and form normalized targets such that,

$$\mathbf{t}_{\text{gt}} = (\mathbf{t}_{\text{raw}} - \mathbf{t}_{\text{min}}) / \mathbf{t}_{\text{range}}, \quad (\text{a})$$

where  $\mathbf{t}_{\text{raw}}$  are the original translation values from the dataset. We train with the normalized translation as objective and at inference, the network predicts  $\mathbf{t}'$  which we de-normalize with the same fixed statistics,

$$\mathbf{t} = \mathbf{t}' \times \mathbf{t}_{\text{range}} + \mathbf{t}_{\text{min}}. \quad (\text{b})$$

Note that the division and product here are performed axis-wise. This preprocessing keeps targets well-scaled without changing the metric translation recovered at test time.

**Stiefel-constrained Optimizer for BiMap Weights.** We maintain each projection matrix at layer  $l$  on the Stiefel manifold  $\mathbf{W}_l \in \mathcal{V}_n(\mathbb{R}^m)$  using a projected gradient step followed by a QR-based retraction. Given a gradient  $\mathbf{G} = \nabla_{\mathbf{W}} \mathcal{L}_{\text{pose}}$ , where  $\mathcal{L}_{\text{pose}}$  is the training loss defined in (10), we first project this gradient onto the tangent space at  $\mathbf{W}_l$ ,

$$\mathbf{G}_{\top} = \mathbf{G} - \mathbf{W}_l \text{sym}(\mathbf{W}_l^{\top} \mathbf{G}), \quad (\text{c})$$

where  $\text{sym}$  denotes the symmetrization operator\*. We then take a step along  $-\mathbf{G}_{\top}$  with step size  $\eta$ ,

$$\widetilde{\mathbf{W}}_l = \mathbf{W}_l - \eta \mathbf{G}_{\top}, \quad (\text{d})$$

\*For a square matrix  $\mathbf{A}$ ,  $\text{sym}(\mathbf{A}) = \frac{1}{2}(\mathbf{A} + \mathbf{A}^{\top})$ .

and retract back to  $\mathcal{V}_n(\mathbb{R}^m)$  via the  $\mathbf{Q}$  factor of a reduced QR decomposition,

$$\widetilde{\mathbf{W}}_l = \mathbf{Q}_l \mathbf{R}_l, \quad \mathbf{W}_l^+ = \mathbf{Q}_l, \quad (\text{e})$$

where  $\mathbf{W}_l^+$  denotes the updated Stiefel-constrained weight at layer  $l$  after the retraction. In our implementation, we perform the QR decomposition in double precision when enabled, for improved numerical stability.

### B. Additional Details on Ablation Studies

We provide additional details on the variants that we design in Section 5.5. Table S1 shows the results on LM-O [7] per-object of the ablation experiments.

Table S1. Per-object results on LM-O. We compare our full method with ablated variants of Cov2Pose to isolate the contribution of each component of the SPD head. This table is the detailed per-object version of Table 5.

Variant	(A)	(B)	(C)	Cov2Pose
Ape	27.2	56.1	62.6	64.0
Can	40.8	93.8	94.4	96.4
Cat	15.7	70.0	66.1	74.1
Driller	26.7	88.1	91.9	96.3
Duck	16.7	49.3	48.1	55.8
Eggbox	47.5	48.0	53.6	57.9
Glue	42.4	91.7	87.0	93.6
Holep.	30.9	70.1	74.7	76.8
Avg. ADD(-S)↑	31.0	70.9	72.3	<b>76.8</b>

**(A) Euclidean Regression Head.** We replace the regression head with a purely Euclidean baseline composed of a two-layer MLP as illustrated in Figure S1. From spatial covariance  $\hat{\Sigma} \in \mathcal{S}_{++}^N$  we form a vector  $\text{vec}(\hat{\Sigma}) \in \mathbb{R}^{\frac{N(N+1)}{2}}$  by stacking its upper-triangular entries *i.e.* unique elements. The MLP maps  $\hat{\Sigma} \rightarrow \mathbb{R}^{256} \rightarrow \mathbb{R}^9$ , directly regressing the 6D rotation parameters  $(\mathbf{u}, \mathbf{v})$  and the 3D translation  $\mathbf{t}$ . As reported in Table 5, this setup significantly decreases performance because learning in Euclidean space ignores the Riemannian geometry of the SPD manifold and collapses the second-order spatial structure.

**(B) Channel-wise Covariance.** Given backbone features  $\mathbf{F} \in \mathbb{R}^{C \times H \times W}$ , we flatten the spatial grid to obtain  $\mathbf{X} \in \mathbb{R}^{C \times N}$  with  $N=HW$ , where each row is a channel and each column a spatial location. We remove the spatial mean per channel to obtain  $\tilde{\mathbf{X}}$  and form the channel-wise covariance,

$$\hat{\Sigma}_{\text{ch}} = \frac{1}{N-1} \tilde{\mathbf{X}} \tilde{\mathbf{X}}^{\top} \in \mathcal{S}_{++}^C. \quad (\text{f})$$




Figure S1. **Variant (A)**. The spatial covariance matrix  $\hat{\Sigma}$  is vectorized and fed to a two-layer MLP that regresses pose purely in Euclidean space.

This descriptor summarizes second-order co-activations between channels aggregated over all spatial locations. Table 5 shows that spatial covariance empirically surpasses channel covariance, as it preserves pairwise relations between image regions that vary with viewpoint, whereas channel covariance collapses the spatial layout and cannot distinguish configurations with different spatial arrangements. It also requires more operations to construct; Forming channel covariance costs  $\mathcal{O}(C^2N)$ , while the spatial covariance costs  $\mathcal{O}(N^2C)$ . Since in late backbone stages  $C \gg N$  (in our setup  $C=2304$ ,  $H \times W=17 \times 17 \Rightarrow N=289$ ), channel covariance is markedly heavier. Empirically, on the same device, it runs at 168 ms per image, versus 46.9 ms for spatial covariance, further supporting the use of spatial covariance.

**(C) Training on the log-tangent space of the SPD manifold.** This variant keeps the SPD head and removes the Cholesky pose decoder. Let  $\mathbf{Z}_L \in \mathcal{S}_{++}^4$  be the SPD output of the head composed of  $L$  layers, and let  $\mathbf{Z}_{gt} \in \mathcal{S}_{++}^4$  be the precomputed SPD pose label obtained from the ground-truth pose in a manner similar to Equation (9). We train by aligning both in the tangent space using a Log-Euclidean map and use a Frobenius loss

$$\mathcal{L}_{\text{Fro}} = \|\text{LogEig}(\mathbf{Z}_L) - \text{LogEig}(\mathbf{Z}_{gt})\|_F^2, \quad (\text{g})$$

where LogEig is the matrix logarithm operation [32]. No pose loss is used here; this design shows the relevance of using our Cholesky decoder to learn the pose parameters. The process is illustrated in Figure S2.

### C. Rotation Representation using Euler Angles

To assess the benefit of using a continuous rotation representation in  $\mathcal{SO}(3)$ , we also design an alternative encoding using a lower-dimensional variant that encodes the rotation using Euler angles which are known to be discontinuous due to Gimbal lock [90]. In this variant, we regress  $\mathbf{Z}_L \in \mathcal{S}_{++}^3$  and the mapping from (8) becomes

$$\begin{aligned} \Psi_- : \mathcal{S}_{++}^3 &\longrightarrow \mathcal{P}_- \\ \mathbf{Z}_L &\longmapsto \Psi_-(\mathbf{Z}_L) = \mathbf{P}, \end{aligned} \quad (\text{h})$$

where  $\mathbf{P} = (\boldsymbol{\theta}, \mathbf{t})$  is the pose with  $\mathbf{t} = (t_x, t_y, t_z) \in \mathbb{R}^3$  the translation vector and  $\boldsymbol{\theta} = (\theta_x, \theta_y, \theta_z) \in [0, 2\pi)^3$




Figure S2. **Variant (C)**. We remove the Cholesky-based pose decoder and train on the log-tangent space of the SPD manifold.  $\mathbf{Z}_L$  and  $\mathbf{Z}_{gt}$  are mapped via LogEig to the log-tangent space. For clarity, a single tangent space  $T_{\mathcal{S}_{++}^4}$  is depicted.

the three Euler angles (roll, pitch and yaw) that represent the object’s rotation in 3D space w.r.t. the camera frame. The corresponding pose space is  $\mathcal{P}_- = [0, 2\pi)^3 \times \mathbb{R}^3$  which is a discontinuous pose representation due to the wrap-around of the Euler angles (at  $0/2\pi$ ) and the gimbal-lock singularities. For encoding the 3D rotation representation (3-dimensional) as well as the translation vector (3-dimensional) we need a matrix with at least 3 non-zero, thus fixing  $n$  to 3, the Cholesky decomposition such that  $\mathbf{Z}_L = \mathbf{L}\mathbf{L}^\top \in \mathcal{S}_{++}^3$  gives a  $3 \times 3$  lower-triangular matrix that we use to encode the pose parameters such that,

$$\mathbf{L} = \begin{bmatrix} e^{\theta_x} & 0 & 0 \\ t_x & e^{\theta_y} & 0 \\ t_y & t_z & e^{\theta_z} \end{bmatrix}. \quad (\text{i})$$

For training, we keep an identical configuration as the default 6D + differentiable GS we propose for our method, the results comparing the two rotation parameterizations are shown in Table 4 of the main paper.

### D. Computation and memory requirements

We further provide a brief analysis of the scaling trade-offs between the spatial covariance size and the spatial resolution ( $N \times N$ ). As discussed in Section B, CovPool has time complexity  $\mathcal{O}(N^2C)$  and memory complexity  $\mathcal{O}(N^2)$ . Table S2 shows that the CovPool runtime is negligible (less than 1 ms), while the SPD head dominates the computation and increases with  $N$ . Therefore, we use, and recommend using, late, low-resolution features, or keeping the spatial resolution moderate via downsampling/pooling.

Table S2. CovPool layer and SPD head latencies across different spatial resolutions.

$H \times W$	CovPool (ms)	SPD Head (ms)
$8^2$	0.38	7.8
<b><math>17^2</math></b>	<b>0.50</b>	<b>23.8</b>
$25^2$	0.61	68.6




Figure S3. **Additional Pose-representation alignment.** As in Figure 1, image pairs are sorted by ground-truth SE(3) distance. We plot Log-Euclidean distances between spatial covariances and, for comparison, cosine and Euclidean distances between flattened features. Covariance distances grow with distances in SE(3), while flattened-feature distances remain nearly constant, reinforcing the hypothesis that spatial covariances are more correlated with true pose.

Regarding memory usage, the same setup described in Section B requires  $\approx 2.54$  MB to store  $\mathbf{X}$  and  $\approx 0.32$  MB to store  $\Sigma$  per input (both in fp32). During backpropagation, activations are also retained, resulting in an overall memory scaling of  $\mathcal{O}(CN + N^2)$ . Additional buffers (e.g., centered  $\mathbf{X}$  and eigen/QR workspaces) have the same asymptotic order.

## E. Generality beyond CNN backbones

We show that `Cov2Pose` also extends beyond CNN backbones. To illustrate this, we use a ViT\_b\_16 backbone [18]. Instead of computing the covariance over spatial locations (pixels), we compute it over the patch tokens extracted by the ViT backbone. Importantly, because ViTs use global self-attention, the resulting SPD matrix does not represent purely “spatial” covariance. Nevertheless, a small pilot experiment with ViT\_b\_16 indicates that a token-covariance + SPD head still outperforms a ViT\_b\_16 baseline that regresses pose parameters using an MLP. The results are shown in Table S3.

Table S3. Pilot comparison on *Ape* (LM-O) with a ViT\_b\_16 backbone. Results reported in ADD.

Class ( <i>Ape</i> ; LM-O)	ADD $\uparrow$
ViT+MLP	33.3
ViT+Cov+SPD head	<b>40.6</b>

## F. Pilot study on symmetry handling and point-based losses

Object symmetries remain a well-known challenge in 6D object pose estimation. `Cov2Pose` does not explicitly model symmetries, as its main design targets a CAD-free, RGB-only, geometry-agnostic, and fully differentiable formulation. Nevertheless, in the main results we did not ob-

serve a specific degradation on the symmetric LM/LM-O objects *eggbox* and *glue*, motivating a small pilot study to assess whether adding symmetry-oriented supervision improves performance.

A natural way to address symmetric ambiguities is to use point-based losses such as ADD-S, or symmetry-aware rotation losses [59, 79]. However, these losses require CAD model points and/or symmetry annotations during training. This changes the training assumptions of `Cov2Pose`.

**Pilot experiment.** We trained variants of `Cov2Pose` on the symmetric objects of LM-O by replacing the pose supervision with: (i) an ADD-S-inspired point-based loss that matches predicted and ground-truth transformed 3D points by nearest neighbors, and (ii) a symmetry-aware rotation loss from prior work similar to [79]. We evaluate on the symmetric objects *Eggbox* and *Glue* using ADD-S, following the standard BOP protocol.

**Results.** Table S4 summarizes the pilot results. The ADD-S-inspired loss does not improve performance in this setting and can degrade results, while the symmetry-aware loss matches or only slightly improves the baseline. In particular, on *Eggbox/Glue*, where `Cov2Pose` achieves 57.9/93.6 ADD-S, the ADD-S-inspired variant reaches 55.0/84.8, and the symmetry-aware loss reaches 57.8/93.8.

**Discussion.** These pilot results suggest that simply plugging in point-based or symmetry-aware losses is not sufficient to consistently improve `Cov2Pose`, despite introducing stronger geometric supervision. Moreover, the two pilot variants depart from the original formulation by relying on operations/supervision that break the fully differentiable

end-to-end design of `Cov2Pose` and its CAD-free training assumptions. We therefore view explicit symmetry handling as an important but non-trivial extension, especially under the constraints targeted by our framework. Designing symmetry-aware training/inference strategies that preserve the CAD-free and fully differentiable nature of `Cov2Pose`, as in [9, 67], is a promising future direction.

Table S4. Pilot study on symmetric LM-O objects using point-based and symmetry-aware losses.

Method	CAD?	<i>Egg.</i> ADD-S $\uparrow$	<i>Glu.</i> ADD-S $\uparrow$
<code>Cov2Pose</code>	×	<b>57.9</b>	<u>93.6</u>
<code>Cov2Pose</code> + $\mathcal{L}_{\text{ADD-S}}$	✓	55.0	84.8
<code>Cov2Pose</code> + $\mathcal{L}_{\text{R,sym}}$ [79]	✓	<u>57.8</u>	<b>93.8</b>

## G. Additional Results & Analysis for Figure 1

In Figure 1, we show that Log-Euclidean distances between spatial covariances of CNN-extracted feature maps grow with ground-truth SE(3) pose distance, whereas cosine distances between flattened features remain nearly constant. This reinforces the hypothesis that a spatial-covariance representation is more pose-sensitive than flattened activations, and thus more suitable for direct regression.

To achieve the result represented in Figure 1, we use an EfficientNet-B6 [73] backbone pretrained on ImageNet, with no pose-specific finetuning and no augmentation to avoid task leakage and isolate the representational contrast between spatial covariance and flattened features.

Figure S3 presents analogous experiments on additional classes from LineMod [26] and YCB-V [87] training data. We report cosine distances, and additionally Euclidean distances, between flattened features, both of which remain nearly constant across pose separations, while Log Euclidean distances between spatial covariances increase, confirming that spatial covariance is a pose-sensitive representation and more correlated with the poses.

## H. Additional Results on APR

To assess generality beyond object pose, we evaluate `Cov2Pose` on the task of Absolute camera Pose Regression (APR) using the Cambridge Landmarks benchmark [40]. The dataset comprises four outdoor scenes (King’s College, Old Hospital, Shop Facade and St Mary’s Church) with labeled video frames and standard train/test splits. We keep the same pipeline using spatial covariance pooling, BiMap+ReEig SPD head, and the SPD(4) Cholesky decoder. Following the PoseNet protocol, we report median translation [m] and rotation [°] errors. As shown in Table S5, compared with the PoseNet family of baselines (PoseNet [40], Bayesian PoseNet [38], and

GPoseNet [39]), `Cov2Pose` achieves lower median translation and rotation errors across all scenes and on average, indicating that spatial second-order cues and manifold-aware learning transfer effectively to scene-centric APR.

Table S5. APR results on **Cambridge Landmarks** [40]. Median translation [m] and rotation [°] errors per scene and their average (Lower is better). `Cov2Pose` achieves the best performance across all scenes and on average.

Method	K. College	Hospital	Shop Facade	St. Mary	Avg.↓
PoseNet [40]	1.92, 5.40°	2.31, 5.38°	1.46, 8.08°	2.65, 8.48°	2.08, 6.83°
BayesianPN [38]	1.74, 4.06°	2.57, 5.14°	1.25, 7.54°	2.11, 8.38°	1.91, 6.28°
GPoseNet [39]	1.61, 2.29°	2.62, 3.89°	1.14, 5.73°	2.93, 6.46°	2.07, 4.59°
<b>Ours</b>	<b>1.57, 1.81°</b>	<b>2.08, 2.35°</b>	<b>1.11, 5.33°</b>	<b>2.06, 5.00°</b>	<b>1.70, 3.56°</b>

## I. YCB-V Dataset Per-object Evaluation

In Table S6, we report a more detailed overview of the results per-object on the YCB-V dataset [87]. For the sake of space, we report only the results of End-to-End pose learning methods.

## J. Additional Results on SPEED+

We further evaluate `Cov2Pose` on SPEED+ [61], a spacecraft pose estimation dataset where PnP-based solutions currently predominate [6, 60, 62, 84]. `Cov2Pose` is trained on the `synthetic` train split and evaluated on its test set. To assess cross-domain generalization, we report inference results on the additional test domains of SPEED+, namely, `lightbox`, simulating diffuse orbital illumination, and `sunlamp`, replicating direct sunlight via a specialized lamp setup. As summarized in Table S7, `Cov2Pose` surpasses SPNv2 [60] across all metrics while being faster at inference. It also outperforms the PnP-based YOLOv8s-pose [6] on the `lightbox` split and achieves superior translation accuracy on `sunlamp`. These results further support the applicability of `Cov2Pose` to spacecraft pose estimation. Qualitative results are shown in Figure S4.

Table S7. **Results on SPEED+**. Performance in terms of Translation  $E_T$ , Rotation  $E_R$  and Pose  $E_{\text{pose}}$  errors as defined in [60]. The inference time of SPNv2 is measured on the same device and at the same input resolution (768×512) as ours.

	Latency (ms)	YOLOv8s-pose [6]		† SPNv2 [60]	Ours
		—	—	80.25	<b>60.46</b>
Synthetic	$E_T$ [m]	—	—	—	<b>0.184</b>
	$E_R$ [°]	—	—	—	<b>3.462</b>
	$E_{\text{pose}}$	—	—	—	<b>0.089</b>
Lightbox	$E_T$ [m]	0.758	0.368		<b>0.300</b>
	$E_R$ [°]	18.00	20.258		<b>15.096</b>
	$E_{\text{pose}}$	0.432	0.411		<b>0.3101</b>
Sunlamp	$E_T$ [m]	0.518	0.457		<b>0.430</b>
	$E_R$ [°]	<b>19.80</b>	34.916		<u>29.128</u>
	$E_{\text{pose}}$	<b>0.432</b>	0.682		<u>0.574</u>

† Results from the direct pose regression head of SPNv2.

Table S6. **YCB-V** per-object results in terms of ADD(-S) and AUC-S/AUC(-S). “—” indicates values not reported in the original paper. Classes denoted with \* are symmetric.

Method	PoseCNN [87]			Single-Stage [30] <sup>†</sup>			DeepIM [45]			GDR-Net [79] <sup>†</sup>			Ours		
	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)
002_master_chef_can	—	84.0	50.9	—	—	—	—	—	—	41.5	96.3	65.2	61.0	94.7	63.1
003_cracker_box	—	76.9	51.7	—	—	—	—	—	—	83.2	97.0	88.8	86.7	91.2	79.8
004_sugar_box	—	84.3	68.6	—	—	—	—	—	—	91.5	98.9	95.0	94.5	94.7	89.8
005_tomato_soup_can	—	80.9	66.0	—	—	—	—	—	—	65.9	96.5	91.9	73.8	95.6	89.7
006_mustard_bottle	—	90.2	79.9	—	—	—	—	—	—	90.2	100.0	92.8	93.3	93.2	84.2
007_tuna_fish_can	—	87.9	70.4	—	—	—	—	—	—	44.2	99.4	94.2	68.7	97.1	94.2
008_pudding_box	—	79.0	62.9	—	—	—	—	—	—	2.8	64.6	44.7	98.7	96.5	93.2
009_gelatin_box	—	87.1	75.2	—	—	—	—	—	—	61.7	97.1	92.5	96.0	97.0	93.6
010_potted_meat_can	—	78.5	59.6	—	—	—	—	—	—	64.9	86.0	80.2	67.5	87.9	78.0
011_banana	—	85.9	72.3	—	—	—	—	—	—	64.1	96.3	85.8	54.7	90.1	75.4
019_pitcher_base	—	76.8	52.5	—	—	—	—	—	—	99.0	99.9	98.5	90.2	95.9	89.2
021_bleach_cleanser	—	71.9	50.5	—	—	—	—	—	—	73.8	94.2	84.3	75.7	92.3	81.9
024_bowl*	—	69.7	69.7	—	—	—	—	—	—	37.7	85.7	85.7	10.7	73.6	73.6
025_mug	—	78.0	57.7	—	—	—	—	—	—	61.5	99.6	94.0	63.5	90.5	76.4
035_power_drill	—	72.8	55.1	—	—	—	—	—	—	78.5	97.5	90.1	83.7	94.7	85.0
036_wood_block*	—	65.8	65.8	—	—	—	—	—	—	59.5	82.5	82.5	63.6	85.7	85.7
037_scissors	—	56.2	35.8	—	—	—	—	—	—	3.9	63.8	49.5	46.6	85.3	70.1
040_large_marker	—	71.4	58.0	—	—	—	—	—	—	7.4	88.0	76.1	16.0	78.8	68.3
051_large_clamp*	—	49.9	49.9	—	—	—	—	—	—	69.8	89.3	89.3	72.0	84.1	84.1
052_extra_large_clamp*	—	47.0	47.0	—	—	—	—	—	—	90.0	93.5	93.5	88.0	92.5	92.5
061_foam_brick*	—	87.8	87.8	—	—	—	—	—	—	71.9	96.9	96.9	58.6	78.3	78.3
<b>Avg. ↑</b>	—	<b>75.9</b>	<b>61.3</b>	<b>53.9</b>	—	—	—	<b>88.1</b>	<b>81.9</b>	<b>60.1</b>	<b>91.6</b>	<b>84.3</b>	<b>69.7</b>	<b>90.0</b>	<b>82.2</b>

<sup>†</sup> Method is End-to-End by using a differentiable PnP strategy.

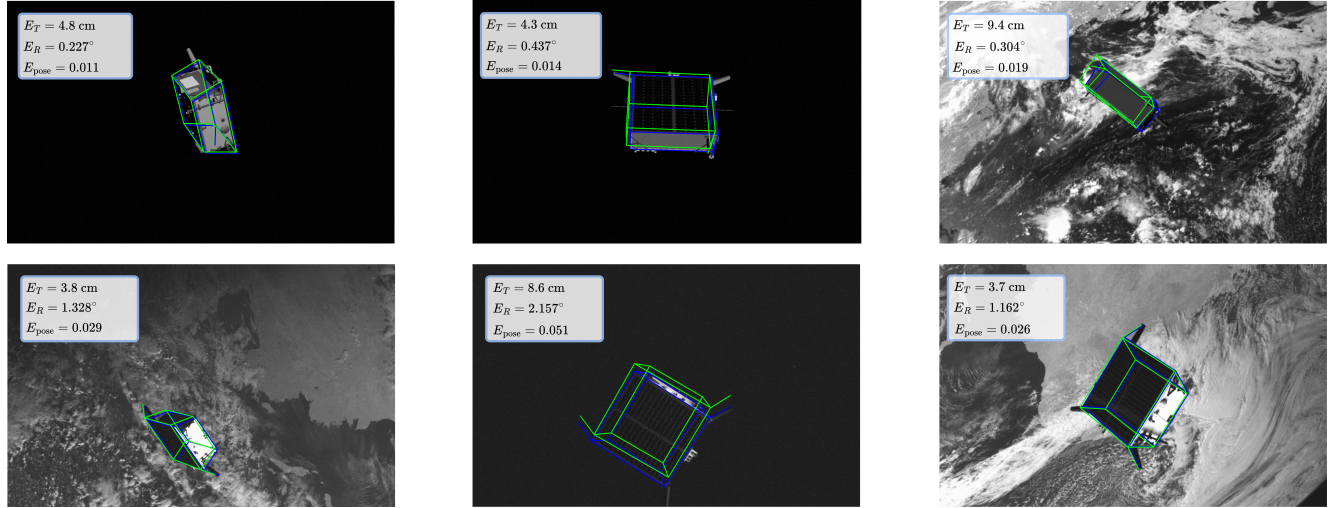


Figure S4. **Qualitative results on SPEED+**. Examples of accurate pose predictions from Cov2Pose on the synthetic (1<sup>st</sup> row) and lightbox (2<sup>nd</sup> row) subsets of SPEED+ [61]. Green bounding boxes show predicted poses, while blue bounding boxes indicate ground-truth poses.