Pair-Relative Partial Equivariance for Rotation-Robust Keypoint Descriptor Matching

Jongmin Lee
Chung-Ang University
jmlee@cau.ac.kr

Abstract

Learned keypoint descriptors trained on upright imagery often fail under large in-plane rotation. Steerers [5] address this by applying a single exact $\mathrm{SO}(2)$ *representation to the descriptor space, but this global group constraint can be too rigid for keypoint descriptors, whose angular responses need not follow a single group representation. We propose pair-relative partial equivariance, which measures equivariance under the empirical distribution of relative rotations in image pairs, and PaSteer, a relaxed descriptor-steering framework that allows the steering map to leave the exact representation manifold while retaining the lightweight, backbone-free design of steering. PaSteer instantiates this relaxation with two variants: an identity-preserving residual steerer (PaSteer-soft) and a free-form Fourier steerer (PaSteer-free). We show that the per-angle optimal linear steerer for a fixed non-equivariant descriptor generally does not obey the group composition law, explaining why exact representations can be suboptimal in the frozen-descriptor regime. On Roto-360, MegaDepth-1500, AIMS, and IMC-PT, PaSteer consistently improves rotation-robust matching over an exact steerer on frozen DeDoDe descriptors, especially under large relative rotations, with minimal upright degradation.*

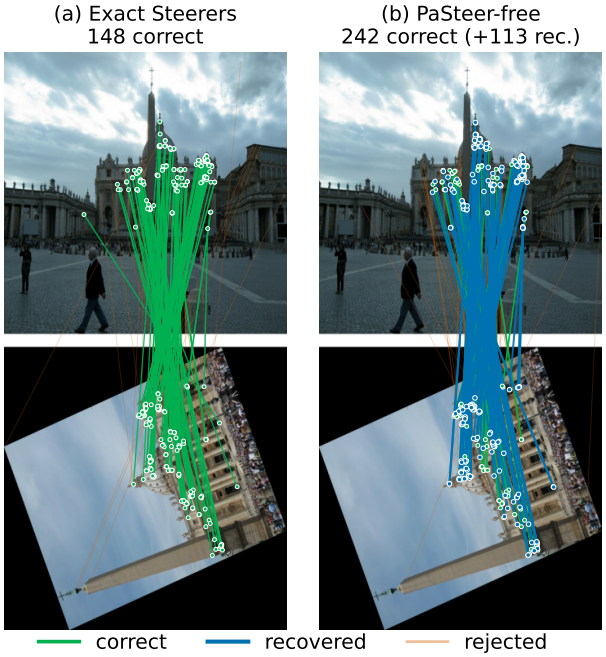


Figure 1. **PaSteer-free recovers correspondences that exact steering misses.** A MegaDepth-1500 [19] pair under a large 112° in-plane rotation. Both panels share the *same* frozen DeDoDe-B [9] descriptor; only the steering changes. Exact Steerers [5] finds 148 correct matches (left), PaSteer-free 242 (right). Green: *correct*; blue: *recovered* (+113 correct matches); orange: *rejected*.

1. Introduction

Local feature matching is a backbone of structure-from-motion, visual localization, and 3D reconstruction. Modern learned descriptors such as SuperPoint [8], DISK [33], XFeat [27], and DeDoDe [9] are trained almost exclusively on upright internet and SfM imagery. Unlike SIFT [22], they therefore degrade sharply once the relative in-plane rotation between two images leaves a narrow band. This matters in practice: astronaut photography from the ISS [31], endoscopy, underwater imagery, and aerial imagery routinely contain large or arbitrary roll.

Existing remedies either build rotation handling into the descriptor or add steering on top of a released descriptor. Rotation-invariant or equivariant descriptors improve robustness through group pooling [21] or rotation-equivariant backbones with group alignment [17, 18, 29], but they require modifying or retraining the representation and may discard orientation as a useful matching cue. Steerers [5] take a lighter route: they keep the descriptor backbone fixed and encode rotation as a linear map $\rho(\theta)$ on the description space. However, this map is constrained to be an *exact* $\mathrm{SO}(2)$ representation, $\rho(\theta+\phi) = \rho(\theta)\rho(\phi)$, forcing one global homomorphism to explain the descriptor response

over the entire group. For an already-trained descriptor that was not optimized to be equivariant, this exact representation constraint can be overly rigid. Throughout, by a *frozen released* descriptor we mean an off-the-shelf descriptor taken from publicly released weights and used as is, without any retraining.

The key observation of this paper is that rotation robustness in matching is pair-relative, not group-uniform. Matching depends on the relative rotation $\Delta\theta$ between two images, and the distribution of $\Delta\theta$ varies sharply by domain. Measured from ground-truth poses, MegaDepth-1500 [19] pairs are mostly near-upright, whereas astronaut imagery [31] is close to uniform over $\mathrm{SO}(2)$ (Sec. 4). Thus, a single fixed degree of equivariance cannot be ideal for both. This is precisely the setting addressed by *partial equivariance*, where equivariance is not all-or-nothing but a data-dependent degree [2, 11, 28, 34, 35]. To our knowledge, this view has not been applied to descriptor matching.

We connect partial equivariance to rotation-robust descriptor matching through *PaSteer*, a relaxed steering framework that keeps the lightweight, backbone-free design of descriptor steering while loosening the exact $\mathrm{SO}(2)$ representation constraint. PaSteer allows the steering map to leave the representation manifold, rather than forcing one homomorphism to fit all angles. We instantiate this idea in two operating points: *PaSteer-soft*, a conservative identity-preserving relaxation anchored to the exact generator, and *PaSteer-free*, a free-form Fourier steerer that directly models the angle-dependent map. Both variants change only the steering map; the detector, frozen descriptor, keypoints, probe grid, matcher, and geometric estimator are kept fixed in our controlled comparisons. Across Roto-360 [17], MegaDepth-1500 [19], AIMS [31], and IMC-PT [13], this relaxation improves frozen-descriptor matching most under large relative rotations.

Our contributions are four-fold:

- We formulate *pair-relative partial equivariance* for descriptor matching, measuring equivariance against the empirical distribution $q(\Delta\theta)$ of pairwise relative rotations rather than against the uniform group.
- We introduce **PaSteer**, a relaxed steering framework that preserves the deployability of descriptor steering while loosening the exact $\mathrm{SO}(2)$ representation constraint.
- We show that, for a frozen non-equivariant descriptor, the per-angle optimal linear steerer is generally non-homomorphic, so an exact $\mathrm{SO}(2)$ representation is intrinsically restrictive.
- We clarify the scope of relaxation: it is most beneficial for frozen released descriptors and becomes comparable to exact steering under full joint descriptor–steerer training.

2. Related Work

Local feature matching. Establishing correspondences between images underlies structure-from-motion, visual localization, and 3D reconstruction. Sparse pipelines detect and describe keypoints [8, 22, 27, 33], and DeDoDe [9] in particular decouples detection from description and is the descriptor we build on. Learned matchers then resolve the assignment, either over sparse keypoints [20, 30] or densely without an explicit detector [10, 32]. These descriptors and matchers are trained almost exclusively on upright internet and SfM imagery, so they degrade once the relative rotation of a pair leaves a narrow band; dense matchers such as RoMa [10] are robust to viewpoint but remain weak under large in-plane rotation. We propose neither a new detector nor a new matcher: our intervention is confined to the descriptor steering stage, leaving the detector and matcher fixed so that steering mechanisms can be compared on a single frozen descriptor under controlled conditions.

Rotation-equivariant features and descriptors. Group-equivariant CNNs extend weight sharing from translations to larger symmetry groups [6, 37], and steerable formulations express the rotation action analytically on feature fields [7, 36]. For matching, one line builds rotation handling directly into the descriptor representation. GIFT [21] pools a descriptor over a transformation orbit to obtain invariance, Self-Sca-Ori [16] estimates a canonical keypoint scale and orientation, while REKD [17], RELF [18], and S-TREK [29] learn rotation-equivariant detectors and descriptors and align the group dimension to a canonical pose; RELF reports strong results on the Roto-360 benchmark introduced with REKD [17]. The same group-equivariant machinery underlies dense correspondence [23] and frequency-domain pose regression [15]. SE2-LoFTR [3] instead places an $\mathrm{SE}(2)$-equivariant backbone inside a detector-free matcher. Building the symmetry into the representation is robust, but it carries a cost: invariance spends discriminative capacity, since orientation is itself a matching cue, and an equivariant backbone must be trained from scratch rather than reused from a released descriptor. These methods also impose equivariance over the entire group.

Steerers and orientation handling. Steerers [5] are the closest baseline to our setting: they keep the backbone untouched and encode a rotation as a single linear map on the description space of a frozen or jointly trained DeDoDe descriptor, requiring ρ to be an exact representation of $\mathrm{SO}(2)$; they report state of the art on Roto-360 and AIMS [31]. Requiring an exact representation ties a global $\mathrm{SO}(2)$ action to the whole descriptor space, so one linear map must account for every rotation consistently. Affine Steerers [4] generalize the construction to $\mathrm{GL}(2)$. A follow-up line asks whether in-

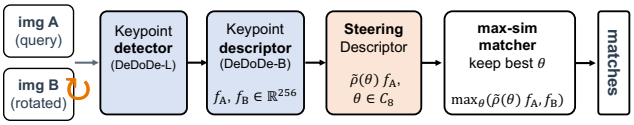


Figure 2. **PaSteer matching pipeline.** An SO(2) keypoint detector (DeDoDe-L [9]) and a keypoint descriptor (DeDoDe-B) feed a steered max-similarity matcher: the query descriptors are steered by $\tilde{\rho}(\theta)$ over the C_8 probe grid under partial equivariance, and the best-scoring angle is kept per keypoint.

variance should live in the descriptor or the matcher [24, 25] and shows that data scaling shrinks the upright-rotated gap. These works steer with an exact representation, train with multiples of 90°, and do not train descriptors with non-trivial steerers. We ask instead whether an exact representation is the right constraint for a frozen released descriptor, and relax it to a continuous, pair-relative degree of equivariance.

Partial, relaxed, and learned equivariance. A growing body of work treats equivariance as a learnable degree rather than a hard constraint. Augerino [2] learns an augmentation distribution; Residual Pathway Priors [11] add a soft non-equivariant residual path; approximately equivariant networks [35] mix equivariant and free kernels; and Partial G-CNNs [28] learn a distribution over a group subset so a layer is equivariant only on $S \subseteq G$. The probabilistic-degree approach [34] learns per-layer equivariance for steerable CNNs via a Fourier parameterization. The Lie derivative [12] quantifies learned equivariance, and approximation-generalization trade-offs under approximate group equivariance have been analyzed theoretically [14, 26]. All of these methods relax the degree of equivariance inside a network backbone, and none has been applied to image matching or to descriptor steering. PaSteer adapts the partial-equivariance idea to the steering map rather than the backbone, making the degree of equivariance data-determined and tying it to the empirically measured *relative* rotation distribution of image pairs.

3. Method

Figure 2 overviews the six-stage pipeline. **(1)** A query I_1 and a target I_2, which may differ by an arbitrary in-plane rotation, enter the system. **(2)** A rotation-equivariant SO(2) detector (DeDoDe-L [9]) produces keypoints in both images. **(3)** A *frozen* DeDoDe-B descriptor f describes each keypoint, giving $f_A, f_B \in \mathbb{R}^{256}$. **(4)** Before matching, a steering map evaluated over a C_8 probe grid steers the query descriptors into a bank of candidate frames. **(5)** A maximum-similarity matcher scores each steered query against the target, keeps the best-scoring angle per keypoint, and applies a mutual-nearest-neighbour selection. **(6)** The surviving correspondences feed the usual geometric estimator (essential matrix or homography) for pose or registration.

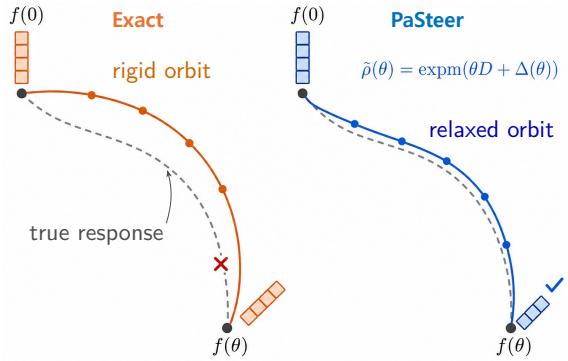


Figure 3. **Exact representation vs. relaxed steering.** The gray dashed curve denotes the true angular response of the keypoint descriptor. The orange solid curve shows the rigid exact-representation orbit, which misses this response, while the blue solid curve shows the relaxed PaSteer-soft orbit, $\tilde{\rho}(\theta) = \exp(\theta D + \Delta(\theta))$, which follows the true response more closely.

The rest of this section develops PaSteer. Section 3.1 formalizes *pair-relative partial equivariance*, the degree of equivariance measured against the rotation a domain actually exhibits, and gives the intuition for why relaxing helps; Sec. 3.2 introduces the relaxed steering family and its two instances, Sec. 3.3 how each steerer fills the steering slot, Sec. 3.4 the matcher that turns steered descriptors into correspondences, and Sec. 3.5 explains training objectives.

3.1. Pair-relative partial equivariance

Let $f : X \to \mathbb{R}^D$ be a keypoint descriptor on an image domain X, and let $R_\theta : X \to X$ be the in-plane rotation action for $\theta \in \mathrm{SO}(2) \cong [0, 2\pi)$. A *steerer* [5] is a map $\rho : \mathrm{SO}(2) \to \mathbb{R}^{D\times D}$ intended to satisfy $f(R_\theta I) \approx \rho(\theta) f(I)$, so that a descriptor extracted from an upright image can be transported into the frame of a rotated one. In exact steering [5], ρ is constrained to be a group representation,

$$\rho(\theta + \phi) = \rho(\theta)\rho(\phi), \qquad \rho(0) = I, \tag{1}$$

which, for SO(2), is parameterized by a single generator D via $\rho(\theta) = \exp(\theta D)$ with $\exp(2\pi D) = I$.

The key observation is that matching performance depends on the *relative* rotation $\Delta\theta = \theta_2 - \theta_1$ of a pair (I_1, I_2), not on absolute orientation. We therefore measure equivariance against the empirical relative-rotation distribution of the data [12], following the partial-equivariance view [2, 11, 26, 28, 34, 35].

Definition 1 (Pair-relative partial equivariance)**.** *Let $q(\Delta\theta)$ be the distribution of the relative rotation of image pairs in a domain. A descriptor f is ε-partially equivariant with respect to q and a steerer ρ if*

$$\mathbb{E}_{\Delta\theta\sim q}\, \mathbb{E}_I \left\| f(R_{\Delta\theta} I) - \rho(\Delta\theta)\, f(I) \right\|^2 \le \varepsilon. \tag{2}$$

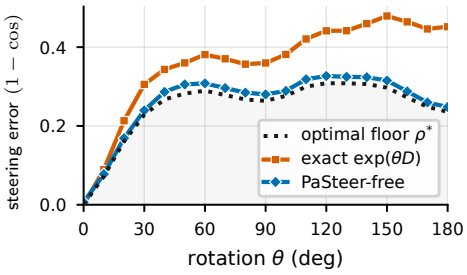


Figure 4. **The relaxed map approaches the optimal floor.** Steering error is $1 - \cos$ of the cosine similarity between a steered upright descriptor $\rho(\theta) f(I)$ and the truly rotated one $f(R_\theta I)$ (lower is better, 0 is exact recovery). On fixed DeDoDe-B [9] descriptor, PaSteer tracks the optimal per-angle floor, while the *exact* representation [5] falls short and its error grows with rotation angle.

The exact-steering setting is the special case $q = \text{Uniform}(\text{SO}(2))$ with ρ constrained as in Eq. (1). We make two changes: q is estimated from data, concentrated near $0°$ for SfM imagery and near-uniform for astronaut imagery, and ρ is allowed to leave the representation manifold. We refer to the constrained case as *exact* and the unconstrained case as *relaxed*.

Why relaxing helps. Figure 3 shows the intuition. The *frozen* keypoint descriptor was never trained to rotate as a group, so its true angular response is not a representation, and the *best* linear steerer at each angle is not one either. An exact representation is locked to the representation manifold and cannot reach this per-angle optimum, whereas a relaxed steerer can. Figure 4 visualizes the steering error of the exact map growing with angle while the relaxed map hugs the optimal per-angle floor. Supp. Sec. A makes this precise with a closed form for the optimal per-angle steerer.

3.2. Relaxed steering family

PaSteer is a relaxed steering family that lets the steering map leave the exact representation manifold, placing a single knob between a rigid exact steerer and a fully free-form one and turning it only as far as the data needs. It has two instances, named consistently throughout the paper.

PaSteer-soft is a residual-exponential relaxation *anchored* to the exact generator: it keeps the generator-exponential form and adds only a small low-rank Fourier residual that vanishes at $\theta = 0$, in the spirit of soft/approximate-equivariance residuals [11, 35]. This preserves the upright identity exactly, while relaxing the exact representation constraint with a small number of additional parameters. It is the conservative, capacity-light, identity-preserving variant.

PaSteer-free is a *free-form* Fourier relaxation: since the per-angle optimal steerer is in general not a representation (supp. Sec. A), it drops the generator-exponential form and parameterizes the angle-dependent map directly in a Fourier basis over the probe grid [28, 34], giving up the orthogonality and homomorphism constraints to approach that optimum. It is the main high-capacity instance in our full-scale experiments; its upright identity is approximate rather than exact, and its larger parameterization is a training and storage cost.

With the exact representation as the rigid endpoint, the three steerers used throughout the paper are:

- **Exact**: a single $\text{SO}(2)$ representation.
- **PaSteer-soft**: a representation-*anchored*, bounded residual-exponential relaxation.
- **PaSteer-free**: a *free-form* Fourier relaxation that drops the representation form.

"Exact" denotes the Steerers [5] representation retrained at our budget, used as the baseline throughout.

3.3. Steering stage

All compared methods share the same detector, frozen descriptor, keypoints, C_8 probe grid, matcher, and geometric estimator; the only component that changes is the steering map $\tilde{\rho}$. **Exact** fills this slot with the representation $\rho(\theta) = \text{expm}(\theta D)$. **PaSteer-soft** adds a bounded low-rank Fourier residual to the exact generator core,

$$\tilde{\rho}(\theta) = \text{expm}\big(\theta D + \Delta(\theta)\big), \tag{3}$$

$$\begin{aligned} \Delta(\theta) &= \sum_{k=1}^{K} \big[\, S_k \sin k\theta + C_k(\cos k\theta - 1) \,\big], \\ S_k &= U_k V_k^\top, \qquad C_k = P_k Q_k^\top \quad \text{of rank } r. \end{aligned} \tag{4}$$

with $\Delta(0) = 0$, so $\tilde{\rho}(0) = I$ exactly and upright pairs incur no steering error. **PaSteer-free** drops the representation form and parameterizes the per-angle map directly,

$$\tilde{\rho}(\theta) = A_0 + \sum_{k=1}^{K} \big[\, A_k \cos k\theta + B_k \sin k\theta \,\big], \quad A_0 \approx I, \tag{5}$$

an unconstrained linear map per angle. Each steerer then applies its chosen map over the C_8 probe grid $\Theta = \{0, 45°, \ldots, 315°\}$, sending a query descriptor f_A to a steered bank $g_A(\theta) = \tilde{\rho}(\theta) f_A$.

Deployability. All three deployed steerers share the same test-time structure: their C_8 steering maps are precomputed before inference, and during matching each query descriptor is transformed by these fixed maps and scored by the same maximum-similarity matcher. PaSteer-free therefore has identical inference cost to exact steering; its larger parameterization affects only training and storage, not the per-pair inference. We separately test whether the gain comes from extra parameters in Sec. 4.3.

3.4. Matching

We score steered descriptors over candidate angles and produce final correspondences. For each query keypoint, the

matcher selects the candidate angle whose steered descriptor best matches the target, following the steered maximum-similarity matching of Steerers [5]. With ℓ_2-normalized descriptors, the similarity between a query keypoint a and a target keypoint b is

$$S_{ab} = \max_{\theta \in \Theta} \langle \widehat{\tilde{\rho}(\theta) f_A^a}, \widehat{f_B^b} \rangle. \tag{6}$$

An upright pair selects an angle near $\theta = 0$, whereas a rotated pair selects the angle that best explains its relative rotation, without requiring any rotation prior. A mutual-nearest-neighbour dual-softmax selection [32] on S yields the final correspondences for the geometric estimator. Since Exact, PaSteer-soft, and PaSteer-free differ only in $\tilde{\rho}$, this comparison isolates the effect of relaxing the steering map.

3.5. Training

Frozen-descriptor objective. With f fixed, we train only the steering map $\tilde{\rho}$ by minimizing the steering error of Definition 1. Equation (2) uses a squared ℓ_2 distance for clarity; we minimize its empirical cosine form, consistent with the matcher of Sec. 3.4, $\mathcal{L} = \mathbb{E}_{\Delta\theta \sim q} \mathbb{E}_I [\, 1 - \cos(f(R_{\Delta\theta} I), \tilde{\rho}(\Delta\theta)\, f(I))\,]$. The distribution $q(\Delta\theta)$ both draws and weights the training angles, so this q-weighting is where pair-relative partial equivariance enters learning, concentrated near 0° for handheld imagery and near-uniform otherwise. For Exact and PaSteer-soft the upright identity is exact by construction (Eq. (4)), so the $\Delta\theta = 0$ term is guaranteed rather than learned.

Joint objective. In the full-scale joint regime the descriptor and steerer train together under the standard DeDoDe matching loss [9] rather than the steering error, with continuous in-plane rotation augmentation that exposes the steerer to all relative angles, not the multiples of 90° used by the original Steerers [5]; Fig. 5 isolates this continuous coverage. The joint comparison is reported in Sec. 4.5; the optimizer, step count, learning rate, and data splits, shared across all steerers, are in Sec. 4.

4. Experiments

We organize the experiments around five questions. Section 4.1 tests whether relaxing the representation constraint improves descriptor matching across benchmarks. Section 4.2 analyzes why the relaxation helps. Section 4.3 asks whether the gain comes merely from added capacity. Section 4.4 studies when equivariance matters under different relative-rotation distributions. Finally, Section 4.5 clarifies the scope of the relaxation under full joint descriptor–steerer training.

Model and steerer setup. Most matching results follow the Setting-A structure of Steerers [5]: the DeDoDe-B descriptor [9] and the $\mathrm{SO}(2)$ detector are fixed, and only the steering transform is trained. This setup isolates the representation constraint, since the exact and PaSteer arms use the same descriptor, detector, keypoints, matcher, initialization, and generator core, and differ only in the steering form. We call this shared setup *matched-budget*; parameter count is not matched here and is controlled separately in Sec. 4.3. Because the descriptor itself is not adapted to rotation, this protocol also caps absolute performance; the full-scale joint comparison is reported separately in Sec. 4.5.

benchmark / metric	Exact	PaSteer-free	Δ
Roto-360 @5px	92.0	**93.4**	+1.4
MD-1500 C4 AUC@5	32.5	**36.4**	+3.9
MD-1500 SO2 AUC@5	20.0	**22.9**	+2.8
AIMS AP	48.6	**50.1**	+1.5
IMC-PT C4 AUC@20	67.2	**75.2**	+8.0
IMC-PT SO2 AUC@20	41.1	**50.6**	+9.5

Table 1. **Main result.** On frozen DeDoDe-B descriptors [9], we measure absolute scores for the exact steerer and PaSteer-free with their difference Δ, on Roto-360 [17] correct matches @5px; MegaDepth-1500 [19] pose AUC@5°; AIMS [31] average precision on the full FMAP set (323 frames); IMC-PT [13] pose AUC@20°. PaSteer-free improves over the exact steerer on all four benchmarks, by the widest margin under large relative rotation.

steerer	Roto	upr.	C4	SO2	AIMS
exact expm	92.0	49.0	32.5	20.0	48.6
PaSteer-free (ours)	**93.4**	**49.3**	**36.4**	**22.9**	**50.1**
Δ (free−exact)	+1.4	+0.3	+3.9	+2.8	+1.5

Table 2. **Full-scale MegaDepth training.** PaSteer-free versus the exact steerer, both trained on the full MegaDepth set on a frozen DeDoDe-B [9] with identical scenes, budget, generator core, and initialization. Roto denotes Roto-360 correct matches @5px [17]; upr., C4, SO2 denote MegaDepth-1500 pose AUC@5° (upright/C4/SO2) [19]; AIMS denotes average precision on AIMS [31]. The Δ row is the relaxed-over-exact margin.

Training and matching details. Unless noted, all learned steerers are trained on a fixed HPatches [1] training split for a few thousand Adam steps and are evaluated with the same matching protocol: maximum-similarity matching over C_8 probe angles, with dual-softmax [32] for the unsteered baseline. For the closed-form optimal steerer of supp. Eq. 8, we estimate G and $C(\theta)$ on the same HPatches training split as the learned steerers and report steering error on a disjoint HPatches evaluation split, using ridge weight $\lambda = 10^{-2} \cdot \mathrm{mean}(\mathrm{diag}(G))$. The HPatches train and evaluation sequences, together with the Roto-360 [17] test sequences, are listed in supp. Sec. E.

Exact, PaSteer, and released Steerers. Our *exact* baseline is the exact-representation counterpart to PaSteer, not

desc.	metric	Exact	PS-free	Δ
DeDoDe-B	Roto-360 @5px	92.0	93.4	+1.4
	MD-1500 C4 AUC@5	32.5	36.4	+3.9
	MD-1500 SO2 AUC@5	20.0	22.9	+2.8
DeDoDe-G	Roto-360 @5px	51.1	**80.6**	**+29.5**
	MD-1500 C4 AUC@5	12.9	**21.6**	**+8.7**
	MD-1500 SO2 AUC@5	11.8	**20.1**	**+8.3**

Table 3. **The cost of the exact constraint grows with descriptor capacity.** Comparison with the exact steerer and PaSteer-free (PS-free) on DeDoDe-B [9] and the larger DeDoDe-G, on Roto-360 [17] correct matches @5px and MegaDepth-1500 [19] pose AUC@5° (C4/SO2). The margin widens markedly on DeDoDe-G across every rotation axis; scores are means over three training seeds (two for exact on DeDoDe-G).

steering	3px	5px	10px
Matched-budget controlled comparison			
none (unsteered)	19.7	20.5	21.6
exact expm (our budget)	77.3	78.3	79.2
PaSteer-soft (r=12, ours)	89.1	90.1	91.0
PaSteer-free (ours)	**93.9**	**95.1**	**96.0**
Released reference models, not controlled			
exact, release Setting-A	91.4	93.0	94.1
Steerers SO2-Spread-B	95.0	97.0	98.5

Table 4. **Matched-budget three-way comparison** on Roto-360 [17], correct matches % @{3,5,10}px, frozen DeDoDe-B [9]. The top block is the matched-budget controlled comparison (ordering exact < PaSteer-soft < PaSteer-free); the bottom block lists released models as absolute references only, not controlled comparisons. Per-benchmark MegaDepth-1500 and AIMS tables are in supp. Sec. B.

an ad-hoc baseline. It uses the same steering mechanism as Steerers [5], a single SO(2) representation $\rho(\theta) = \text{expm}(\theta D)$, under the same Setting-A structure. The difference from the *released* Steerers model is protocol rather than mechanism: our exact and PaSteer variants use the matched-budget setup above, whereas the released model is trained on MegaDepth [19] with rotations at multiples of 90° and uses a Procrustes matcher. The released Steerers rows in Tab. 4 and supp. Tabs. 10 and 11 are therefore absolute-performance references, while the controlled comparison is exact versus PaSteer under matched training.

Datasets. We evaluate matching on Roto-360 [17], MegaDepth-1500 [19], AIMS [31], and the IMC-PT 2020 validation set [13].

4.1. Does relaxation improve descriptor matching?

Our central regime is frozen descriptors: equipping an already-trained descriptor with rotation robustness without retraining it, with the SO2 detector fixed.

The relaxed-over-exact margin. Table 1 is the main results: on a frozen descriptor, relaxing the representation constraint improves matching on all four benchmarks, by the widest margin on the large-rotation C4 and SO2 axes, while upright stays within a fraction of a point of exact (slightly below for PaSteer-free, whose err(0) is near but nonzero, and exactly zero for PaSteer-soft by construction). The reported Δ is a controlled within-protocol difference that sidesteps the AIMS [31] absolute-AP protocol gap noted below.

Full-scale Setting-A absolutes. Table 2 reports the full-scale MegaDepth [19] training results, where PaSteer-free wins on every axis. Additional results are in the supplement.

Scaling to a larger descriptor. Tab. 3 shows the results on a larger-sized descriptor DeDoDe-G [9]. The cost of enforcing an exact representation grows with descriptor capacity. DeDoDe-G uses a substantially larger DINOv2 backbone than DeDoDe-B, and the relaxed-over-exact margin widens on every rotation axis. This mirrors Sec. 4.2: on the larger descriptor, the per-angle optimum departs further from a group representation, a gap that only the free-form steerer can effectively close.

Matched-budget three-way comparison. Table 4 shows the Roto-360 [17] matched-budget comparison among exact, PaSteer-soft, and PaSteer-free, trained with the same data and budget on the same HPatches [1] sequences. The ordering exact < PaSteer-soft < PaSteer-free holds, with PaSteer-soft improving over the matched exact steerer while keeping err(0) = 0 and PaSteer-free improving further. The released SO2-Spread-B model reaches higher absolute accuracy under a different full-scale released training protocol, so it should be read as an upper reference rather than a controlled counterpart to PaSteer. The same matched-budget ordering reproduces on MegaDepth-1500 and AIMS (PaSteer-soft +7.3 AP, 44.4 vs. 37.1; PaSteer-free 51.2), with full per-benchmark tables in supp. Sec. B.

Continuous-angle robustness. Figure 5 shows the results as pair relative rotation is swept continuously, isolating continuous steering from augmenting only by multiples of 90° on MD-1500 pose AUC@5°. The unsteered descriptor is reliable only near upright rotations; the exact steerer restores robustness but dips at the C_8 discretization-worst angles. PaSteer-free leads across most angles, especially near these dips, and also raises the accuracy level between probe angles, showing that the gain is not just a C_8 gap-filling effect.

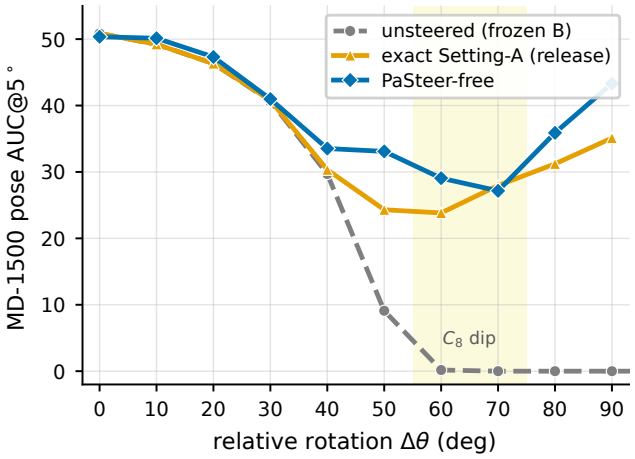


Figure 5. **Continuous-angle robustness.** Per-angle MD-1500 [19] pose AUC@5° as the relative in-plane rotation $\Delta\theta$ is swept continuously, on a frozen DeDoDe-B [9] descriptor with C_8 probes. *Unsteered* (gray dashed) is the frozen descriptor without steering; *exact Setting-A* (orange) is the released Steerers [5] Setting-A model; *PaSteer-free* (blue) is our relaxed steerer. PaSteer-free is at or above both across nearly the whole sweep, with the largest gains at the C_8 discretization-worst angles.

4.2. Why does the partial equivariance help?

Table 5 shows steering-error results on HPatches [1] evaluation set using a frozen DeDoDe-B [9] descriptor. We compare each learned steerer with the closed-form per-angle optimum $\rho^\star$ of supp. Eq. 8, which serves as an oracle lower-error reference rather than a deployable model.

The exact steerer is consistently farther from this optimum, especially at large rotations. PaSteer-free stays close to $\rho^\star$ at all reported angles, and PaSteer-soft falls in between (Fig. 4). The bottom rows show that $\rho^\star$ has a large composition violation, meaning that the best angle-wise maps do not compose as a single $\mathrm{SO}(2)$ representation. This explains why the exact representation constraint is limiting for a frozen descriptor, and why PaSteer-free can better match its actual angular response.

4.3. Capacity and constraint controls

Mechanism comparison. Tab. 6 shows that PaSteer-free obtains the lowest steering error at every angle but at a large parameter cost, raising whether the extra capacity, not the relaxation, buys the gain. To separate the two we use two controls: *PaSteer-soft*, a capacity-light residual that keeps $\mathrm{err}(0) = 0$ and whose out-of-band gain over exact *saturates* with rank, so it is a functional-form effect rather than added capacity; and an *orthogonality-preserving residual* of matched size that isolates the orthogonality constraint. Table 6 reports steering error in and out of the natural partial-equivariance band of Fig. 5.

θ	$\rho^\star$	exact	PaSteer-soft	PaSteer-free
30°	0.229	0.305	0.311	**0.239**
90°	0.264	0.360	0.348	**0.280**
180°	0.236	0.452	0.366	**0.248**
comp. violation			0.60–0.70	
orthog. deviation			0.85–1.04	

Table 5. **Steering error against the per-angle optimum.** Steering error $1 - \cos$ is reported on HPatches [1] evaluation set using a frozen DeDoDe-B [9] descriptor. The top rows compare the closed-form per-angle optimum $\rho^\star$, the exact steerer, PaSteer-soft, and PaSteer-free at selected rotation angles. The bottom rows report the composition violation and orthogonality deviation of $\rho^\star$: the former measures how far $\rho^\star(\theta+\phi)$ is from $\rho^\star(\theta)\rho^\star(\phi)$, and the latter measures how far ${\rho^\star}^\top\rho^\star$ is from the identity.

config	+params%	in-band	out-band	err(0)
exact expm	0	0.152	0.410	0.000
PaSteer-soft (r=12)	+75	0.156	**0.374**	0.000
PaSteer-soft (r=24)	+150	0.155	0.379	0.000
PaSteer-free	+1200	**0.123**	**0.297**	0.006

Table 6. **The gain is functional form, not parameter count.** Steerer mechanism comparison on a frozen DeDoDe-B [9] descriptor (HPatches [1] evaluation set): steering error $1 - \cos$ in and out of the partial-equivariance band, with err(0) the error at 0°. The out-of-band gain of PaSteer-soft over exact saturates with rank despite added parameters, so the remaining gap to PaSteer-free reflects leaving the representation and orthogonality constraints, not capacity.

Deployability. PaSteer-free is fully deployable: as in Sec. 3.3, all three steerers are precomputed linear maps over the shared C_8 probe grid scored by the same matcher, so PaSteer-free runs at exactly the inference cost of the exact steerer, its larger parameterization a one-time training and storage cost, not a per-match one. The +1200% is relative to the small exact generator; PaSteer-free has 0.85M parameters, 6.3% of the DeDoDe-B descriptor.

Capacity fairness. Table 7 shows a sweep of the harmonic order K of the free-form Fourier map that separates relaxation from raw capacity; the parameter cost $(2K+1)\cdot 256^2$ grows directly with K. More parameters alone do not help: K=1 already has more parameters than the exact steerer (1.5% vs. 0.48% of the descriptor) yet is markedly worse, and the relaxed-over-exact margin turns positive from K=3 on, so $K \geq 3$ is a safe default (we use K=6). The gain thus comes from leaving the representation and orthogonality constraints, not from parameter count.

4.4. When does equivariance matter?

Table 8 shows when equivariance matters by changing the pair-relative rotation distribution $q(\Delta\theta)$ while keeping the matching setup fixed. We evaluate RotoReal by resampling

steerer	params (%)	Roto@5	C4@5	SO2@5	ΔC4	ΔSO2
exact	0.48	92.1	32.3	20.3	–	–
K=1	1.5	89.0	20.8	12.4	−11.6	−7.9
K=2	2.4	92.2	29.9	19.2	−2.4	−1.0
K=3	3.4	93.1	34.6	22.6	+2.3	+2.3
K=4	4.4	93.4	35.1	23.5	+2.7	+3.2
K=6	6.3	**93.6**	**37.9**	**23.9**	**+5.6**	**+3.7**

Table 7. **Free-Fourier capacity sweep over harmonic order K.** Frozen DeDoDe-B [9] Setting-A: Roto@5 is Roto-360 [17] correct matches % @5px; C4@5 and SO2@5 are MegaDepth-1500 [19] pose AUC@5°; ΔC4 and ΔSO2 are PaSteer-free(K) minus exact; params (%) is steerer parameters as a fraction of the descriptor. More parameters alone do not beat exact: a sufficient harmonic order in the free-form parameterization is needed to cross it.

steering	T1 handheld	T3 uniform
dual softmax (none)	50.8	10.7
exact expm (ours)	50.7	19.7
PaSteer-soft (r=12, ours)	50.7	24.7
PaSteer-free (ours)	49.9	**35.0**
Steerers Setting-A (release, ref.)	*50.7*	*30.6*

Table 8. **The useful degree of equivariance depends on $q(\Delta\theta)$.** RotoReal MD-1500 [19] pose AUC@5° on a frozen DeDoDe-B [9], resampling pairs under a measured handheld distribution (T1) versus uniform over SO(2) (T3); only the steering changes. Steering is near-irrelevant on T1 and separates the methods only on T3. The release row is MegaDepth-trained and shown for reference.

MD-1500 [19] pairs under two distributions: T1 draws $|\Delta\theta|$ from a measured handheld distribution, whereas T3 samples rotations uniformly over SO(2).

The result supports the pair-relative view of Definition 1. On T1, the common small-rotation regime, steering is nearly irrelevant because the upright descriptor already covers the rotations realized in the data. On T3, where large rotations are common, the methods separate clearly and the matched-training ordering exact < PaSteer-soft < PaSteer-free reappears.

This is the empirical core of partial equivariance: the benefit of full SO(2) equivariance is small when $q(\Delta\theta)$ is concentrated near upright, but becomes important when q places mass on large rotations. Thus, the useful degree of equivariance is not fixed globally; it is determined by the pair-relative rotation distribution of the domain.

4.5. What is the scope?

Table 9 shows the full-scale joint descriptor–steerer comparison on MegaDepth [19], where only the steerer mechanism differs. Exact and PaSteer-free reach parity, suggesting that joint training lets the descriptor absorb the representation constraint. Thus, PaSteer is primarily a frozen-descriptor retrofit method: it helps most when the descriptor cannot be retrained, while becoming largely neutral under full joint training. The larger DeDoDe-G [9] joint result shows the same scope boundary, with a small edge for exact steering (supp. Sec. D).

model	MD upright	MD C4	MD SO2
exact steerer	**50.3**	47.0	43.3
PaSteer-free (ours)	49.7	**47.1**	**43.6**

Table 9. **Joint training reaches parity, isolating the frozen-descriptor advantage.** Full-scale joint descriptor–steerer training on DeDoDe-B [9], pose AUC@5° on MegaDepth-1500 [19], with identical scenes, schedule, generator core, and initialization so only the steerer mechanism differs. Exact and PaSteer-free reach parity once the descriptor co-adapts, so the advantage of the relaxation is concentrated in the frozen-descriptor regime (supp. Tab. 10).

5. Conclusion

We formulated rotation-robust keypoint descriptor matching as pair-relative partial equivariance, where equivariance is measured against the empirical relative-rotation distribution $q(\Delta\theta)$ rather than the uniform group. For a fixed, non-equivariant descriptor, the per-angle optimal linear steerer is generally not a group representation, which explains why an exact SO(2) representation can be restrictive in the frozen-descriptor regime. PaSteer relaxes this constraint with two variants: PaSteer-soft, which preserves the upright identity exactly, and PaSteer-free, which directly parameterizes the angle-dependent map.

Our experiments support this view. The steering-error analysis shows that the best steering map changes with rotation angle and is not well captured by one exact representation. Matching experiments show that PaSteer improves robustness to large rotations while preserving upright performance, and distribution controlled experiments confirm that steering matters most when large relative rotations are common. Overall, PaSteer provides a lightweight way to add rotation robustness to already-trained descriptors without changing the descriptor backbone.

A natural next step is to extend this pair-relative, data-driven view from in-plane rotation to richer geometric transformations such as scale, affine, and viewpoint changes, potentially reducing the amount of augmentation or training data needed to obtain robust descriptors.

Limitations. Our evaluation is limited to DeDoDe descriptors, and we compare against exact steering rather than rotation-equivariant methods that retrain the backbone. The steerer also applies the same C_8 grid to every pair instead of adapting its degree of equivariance.

Acknowledgements. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2026-25598824).

References

[1] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. HPatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *IEEEConf.Comput.Vis.PatternRecog.*, 2017. 5, 6, 7, 11, 12, 13

[2] Gregory Benton, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. Learning invariances in neural networks from training data. In *Adv.NeuralInform.Process.Syst.*, 2020. 2, 3

[3] Georg Bökman and Fredrik Kahl. A case for using rotation invariant features in state of the art feature matchers. In *IEEEConf.Comput.Vis.PatternRecog.Worksh.*, 2022. 2

[4] Georg Bökman, Johan Edstedt, Michael Felsberg, and Fredrik Kahl. Affine steerers for structured keypoint description. In *Eur.Conf.Comput.Vis.*, 2024. 2

[5] Georg Bökman, Johan Edstedt, Michael Felsberg, and Fredrik Kahl. Steerers: A framework for rotation equivariant keypoint descriptors. In *IEEEConf.Comput.Vis.PatternRecog.*, 2024. 1, 2, 3, 4, 5, 6, 7, 11, 12, 13, 14, 15

[6] Taco S. Cohen and Max Welling. Group equivariant convolutional networks. In *Int.Conf.Mach.Learn.*, 2016. 2

[7] Taco S. Cohen and Max Welling. Steerable CNNs. In *Int.Conf.Learn.Represent.*, 2017. 2

[8] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-supervised interest point detection and description. In *IEEEConf.Comput.Vis.PatternRecog.Worksh.*, 2018. 1, 2

[9] Johan Edstedt, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. DeDoDe: Detect, don't describe – describe, don't detect for local feature matching. In *Int.Conf.3DVis.*, 2024. 1, 2, 3, 4, 5, 6, 7, 8, 11, 12, 13, 14

[10] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. RoMa: Robust dense feature matching. In *IEEEConf.Comput.Vis.PatternRecog.*, 2024. 2

[11] Marc Finzi, Gregory Benton, and Andrew Gordon Wilson. Residual pathway priors for soft equivariance constraints. In *Adv.NeuralInform.Process.Syst.*, 2021. 2, 3, 4

[12] Nate Gruver, Marc Finzi, Micah Goldblum, and Andrew Gordon Wilson. The Lie derivative for measuring learned equivariance. In *Int.Conf.Learn.Represent.*, 2023. 3

[13] Yuhe Jin, Dmytro Mishkin, Anastasiia Mishchuk, Jiří Matas, Pascal Fua, Kwang Moo Yi, and Eduard Trulls. Image matching across wide baselines: From paper to practice. *International Journal of Computer Vision*, 129(2):517–547, 2021. 2, 5, 6, 12, 14, 15

[14] Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *Int.Conf.Mach.Learn.*, 2018. 3

[15] Jongmin Lee and Minsu Cho. 3D equivariant pose regression via direct Wigner-D harmonics prediction. In *Adv.NeuralInform.Process.Syst.*, 2024. 2

[16] Jongmin Lee, Yoonwoo Jeong, and Minsu Cho. Self-supervised learning of image scale and orientation. In *Brit.Mach.Vis.Conf.*, 2021. 2

[17] Jongmin Lee, Byungjin Kim, and Minsu Cho. Self-supervised equivariant learning for oriented keypoint detection. In *IEEEConf.Comput.Vis.PatternRecog.*, 2022. 1, 2, 5, 6, 8, 12, 13, 14, 16

[18] Jongmin Lee, Byungjin Kim, Seungwook Kim, and Minsu Cho. Learning rotation-equivariant features for visual correspondence. In *IEEEConf.Comput.Vis.PatternRecog.*, 2023. 1, 2, 13

[19] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *IEEEConf.Comput.Vis.PatternRecog.*, 2018. 1, 2, 5, 6, 7, 8, 12, 13, 14, 15

[20] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local feature matching at light speed. In *Int.Conf.Comput.Vis.*, 2023. 2

[21] Yuan Liu, Zehong Shen, Zhixuan Lin, Sida Peng, Hujun Bao, and Xiaowei Zhou. GIFT: Learning transformation-invariant dense visual descriptors via group CNNs. In *Adv.NeuralInform.Process.Syst.*, 2019. 1, 2

[22] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int.J.Comput.Vis.*, 2004. 1, 2

[23] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In *Eur.Conf.Comput.Vis.*, 2020. 2

[24] David Nordström, Johan Edstedt, Georg Bökman, Jonathan Astermark, Anders Heyden, Viktor Larsson, Mårten Wadenbäck, Michael Felsberg, and Fredrik Kahl. LoMa: Local feature matching revisited. *arXiv preprint arXiv:2604.04931*, 2026. 3

[25] David Nordström, Johan Edstedt, Fredrik Kahl, and Georg Bökman. Who handles orientation? Investigating invariance in feature matching. In *IEEEConf.Comput.Vis.PatternRecog.Worksh.*, 2026. 3

[26] Mircea Petrache and Shubhendu Trivedi. Approximation-generalization trade-offs under (approximate) group equivariance. In *Adv.NeuralInform.Process.Syst.*, 2023. 3

[27] Guilherme Potje, Felipe Cadar, André Araujo, Renato Martins, and Erickson R. Nascimento. XFeat: Accelerated features for lightweight image matching. In *IEEEConf.Comput.Vis.PatternRecog.*, 2024. 1, 2

[28] David W. Romero and Suhas Lohit. Learning partial equivariances from data. In *Adv.NeuralInform.Process.Syst.*, 2022. 2, 3, 4

[29] Emanuele Santellani, Christian Sormann, Mattia Rossi, Andreas Kuhn, and Friedrich Fraundorfer. S-TREK: Sequential translation and rotation equivariant keypoints for local feature extraction. In *Int.Conf.Comput.Vis.*, 2023. 1, 2

[30] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *IEEEConf.Comput.Vis.PatternRecog.*, 2020. 2

[31] Alex Stoken and Kenton Fisher. Find my astronaut photo: Automated localization and georectification of astronaut photography. In *IEEEConf.Comput.Vis.PatternRecog.Worksh.*, 2023. 1, 2, 5, 6, 11, 12, 13, 14, 16

[32] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching

with transformers. In *IEEEConf.Comput.Vis.PatternRecog.*, 2021. 2, 5

[33] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning local features with policy gradient. In *Adv.NeuralInform.Process.Syst.*, 2020. 1, 2

[34] Lars Veefkind and Gabriele Cesa. A probabilistic approach to learning the degree of equivariance in steerable CNNs. In *Int.Conf.Mach.Learn.*, 2024. 2, 3, 4

[35] Rui Wang, Robin Walters, and Rose Yu. Approximately equivariant networks for imperfectly symmetric dynamics. In *Int.Conf.Mach.Learn.*, 2022. 2, 3, 4

[36] Maurice Weiler and Gabriele Cesa. General E(2)-equivariant steerable CNNs. In *Adv.NeuralInform.Process.Syst.*, 2019. 2

[37] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow. Harmonic networks: Deep translation and rotation equivariance. In *IEEEConf.Comput.Vis.PatternRecog.*, 2017. 2