
Projective Equivariant Networks via Second-order Fundamental Differential Invariants

Yikang Li¹, Yeqing Qiu^{2,3}, Yuxuan Chen⁴, Lingshen He¹, Lexiang Hu¹, Zhouchen Lin^{1,5*}

¹State Key Lab of General AI, School of Intelligence Science and Technology, Peking University

²School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen

³Shenzhen Research Institute of Big Data

⁴Khoury College of Computer Sciences, Northeastern University

⁵Institute for Artificial Intelligence, Peking University

liyk18@pku.edu.cn, yeqingqiu@link.cuhk.edu.cn, chen.yuxuan7@northeastern.edu

lingshenhe@pku.edu.cn, hulx@stu.pku.edu.cn, zlin@pku.edu.cn

Abstract

Equivariant networks enhance model efficiency and generalization by embedding symmetry priors into their architectures. However, most existing methods, primarily based on group convolutions and steerable convolutions, face significant limitations when dealing with complex transformation groups, particularly the projective group, which plays a crucial role in vision. In this work, we tackle the challenge by constructing projective equivariant networks based on differential invariants. Using the moving frame method with a carefully selected cross section tailored for multi-dimensional functions, we derive a complete and concise set of second-order fundamental differential invariants of the projective group. We provide a rigorous analysis of the properties and transformation relationships of their underlying components, yielding a further simplified and unified set of fundamental differential invariants, which facilitates both theoretical analysis and practical applications. Building on this foundation, we develop **PDINet**, the first framework for deep projective equivariant networks, achieving full projective equivariance without discretizing or sampling the group. Empirical results on the projectively transformed STL-10 and Imagenette datasets show that PDINet achieves improvements of 11.39% and 5.66% in accuracy over the respective standard baselines under out-of-distribution settings, demonstrating its strong generalization to complex geometric transformations.

1 Introduction

Incorporating symmetry as an inductive bias into neural networks has emerged as a powerful approach to enhance model efficiency and generalization. Convolutional neural networks (CNNs) [Krizhevsky et al., 2012, Simonyan and Zisserman, 2015, He et al., 2016, Chen et al., 2017], which are among the most widely used architectures in deep learning, owe much of their success to the inherent translational equivariance. Building on this idea, Cohen and Welling [2016a] proposed Group Equivariant CNNs (G-CNNs), which generalize equivariance to broader transformations like rotations and reflections. Equivariant networks achieve symmetry incorporation by constructing network layers whose outputs transform in a predictable pattern under group actions applied to the inputs.

The development of equivariant networks began with G-CNNs, whose feature map can be seen as a function on a group. Although G-CNNs have proven effective in various tasks [Worrall and Brostow, 2018, Esteves et al., 2019, Lafarge et al., 2021, Shamsolmoali et al., 2021], they are less suited

*Corresponding author.

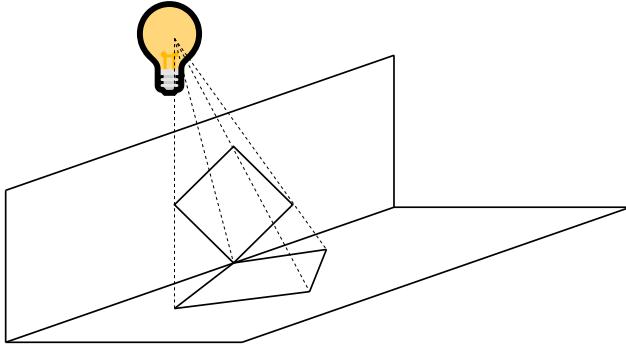


Figure 1: An illustration of a projective transformation: when a planar object is projected onto another plane, a projective transformation occurs.

to continuous groups, as handling such groups typically requires group sampling or discretization, which introduces approximation errors and computational complexity. Then, steerable CNNs [Cohen and Welling, 2016b, Weiler and Cesa, 2019] were proposed to overcome these limitations by viewing features as fields that transform according to specified group representations. In this framework, G-CNNs can be interpreted as a special case where the group representation is chosen to be the regular representation. Steerable CNNs are capable of handling continuous groups such as $SO(2)$ and $SO(3)$ directly, thereby significantly broadening the scope of equivariant networks to include real-world symmetries beyond discrete groups [Weiler et al., 2018, Wang et al., 2020, Wang and Walters, 2022]. However, for more complex non-compact Lie groups such as the projective group, deriving closed-form steerable basis filters becomes intractable, inherently limiting the applicability of steerable CNNs. To address this, MacDonald et al. [2022] enabled group convolutions over finite-dimensional Lie groups by computing the integral on the Lie algebra, thus introducing a projective equivariant model, homConv. However, this method still relies on group sampling, which results in exponential memory growth with increasing network depth, thereby hindering scalability to deeper architectures. Mironenco and Forré [2024] improved sampling efficiency via group decompositions, but focused solely on affine subgroups like $\mathbb{R}^n \rtimes GL^+(n, \mathbb{R})$ and $\mathbb{R}^n \rtimes SL(n, \mathbb{R})$, without addressing more complex groups such as the projective group. Recently, Li et al. [2024, 2025] proposed InvarLayer and steerable EquivarLayer that construct affine equivariant networks based on invariants, enabling closed-form and sampling-free affine equivariance. However, these works remain specialized to the affine group and do not yet generalize to more complex non-compact groups like the projective group.

Actually, projective transformations play a fundamental role in computer vision [Mohr and Triggs, 1996, Birchfield, 1998, Hartley and Zisserman, 2003], as they capture the relationships between objects and their images under perspective projections (see Figure 1). Achieving equivariance on projective transformations is especially critical in practical applications such as mobile robot navigation, 3D scene analysis, and camera pose estimation, where accurately handling perspective effects and viewpoint changes can significantly enhance model robustness and accuracy [Lee et al., 2000, Hartley and Zisserman, 2003, Mur-Artal et al., 2015, Schönberger et al., 2016]. Early on, Suk and Flusser [2004] proposed a projective invariant feature extraction method based on projective moment invariants. Nevertheless, the invariants are formulated as infinite series of moment products, leading to significant computational overhead and intractable error analysis in practical implementations. To overcome this limitation, Li et al. [2018] proposed an alternative framework that constructs projective invariants using finite combinations of weighted moments, where the weights are derived from relative projective differential invariants. Note that moment-based projective invariants are essentially global image descriptors, which makes them inappropriate for constructing equivariant operators that act locally on feature fields for capturing fine-grained spatial patterns. Instead, differential invariants inherently hold the property of acting locally at each spatial position, which makes them a natural foundation for building equivariant operators [Sangalli et al., 2022, 2023, Li et al., 2024, 2025]. While Olver [2023] has proposed a systematic framework for the computation of projective differential invariants via the moving frame method, it requires at least third-order derivatives because

the projective action is not free at second order for scalar functions. Besides, they only consider single-channel cases, and the extension to multi-dimensional cases needs more complex expressions involving high-order derivatives, which limits their practicality in computation and applications on color images.

In this work, we construct projective equivariant networks based on differential invariants, achieving full projective equivariance without relying on group discretization or sampling. This overcomes the depth limitations of homConv [MacDonald et al., 2022] and enables effective scaling to deeper architectures. A core challenge lies in deriving concise and practical projective differential invariants. To support color images and multi-channel intermediate features in modern neural networks, we focus on the differential invariants for multi-dimensional functions. In this case, the projective group acts freely on the second-order jet space, allowing us to derive a complete set of second-order fundamental differential invariants using the moving frame method [Olver, 2015], which can express any second-order invariant of the projective group. However, the choice of cross section used to define the moving frame significantly affects the form of the resulting invariants. While a direct extension of the cross section in [Olver, 2023] to the multi-dimensional case is theoretically valid, it leads to prohibitively long expressions with hundreds of terms, rendering them impractical. Instead, we propose a new cross section tailored to the multi-dimensional structure, which involves up to second-order derivatives, yielding a much more concise set of fundamental differential invariants. Further analysis reveals that these fundamental invariants are composed of a set of simpler components. By exploring the algebraic properties and transformation relationships of these components, we further simplify invariants into a unified set of fundamental invariants, facilitating both practical use and theoretical analysis. Based upon these simplified invariants, we design learnable equivariant operators by combining them with parameterized multi-layer perceptrons (MLPs), and embed the operators into standard neural network backbones to build **PDINet**, the first framework for deep projective equivariant networks free from group sampling. Empirical evaluations under challenging out-of-distribution settings demonstrate the strong generalization ability of our model to complex geometric transformations.

We summarize our main contributions as follows:

- We employ the moving frame method to derive a complete set of second-order fundamental differential invariants of the projective group for multi-dimensional functions, enabling support for color images and multi-channel features.
- We conduct an in-depth analysis of the algebraic structure and transformation properties of these invariants, resulting in a further simplified and unified set of fundamental invariants that facilitate both theoretical understanding and practical computation.
- We develop **PDINet** based on second-order projective differential invariants. It is the first time that deep networks achieve full projective equivariance without relying on group discretization or sampling, thus allowing effective scaling to deeper architectures.
- Numerical experiments on projectively deformed STL-10 and Imagenette² under out-of-distribution settings demonstrate the effectiveness of our model, with improvements of 11.39% and 5.66% over the standard baseline results, showcasing its strong generalization capability under complex geometric transformations.³

2 Method

2.1 Basic concepts and notations

To begin with, we introduce some basic concepts and notations necessary for our formulation. An image can be viewed as a continuous function $\mathbf{u}(x, y)$ defined on a 2D plane. For example, an RGB image corresponds to a three-dimensional function. Likewise, intermediate features in neural networks can also be interpreted as functions, and each layer can be seen as an operator that maps one function to another.

A central concept in this work is equivariance. If the output of an operator undergoes a corresponding transformation when the input is transformed, it is referred to as equivariance. The formal definition is as follows:

²Imagenette is a publicly available dataset downloaded from <https://github.com/fastai/imagenette>.

³Our code is available at <https://github.com/Liyk127/PDINet>.

Definition 1 An operator $\psi : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ is said to be **equivariant** with respect to a group G if

$$g \cdot \psi(\mathbf{u}) = \psi(g \cdot \mathbf{u}), \quad \forall g \in G, \mathbf{u} \in \mathcal{F}_1, \quad (1)$$

where \mathcal{F}_1 and \mathcal{F}_2 are the input and output function spaces, respectively.

Let X denote the domain, $U = \mathbb{R}^n$ be the range of a function, and $U^{(d)} = U \times U_1 \times \cdots \times U_d$ be the derivative space up to order d . A group action $g \cdot \mathbf{x}$ on the domain naturally induces an action on functions, defined as $(g \cdot \mathbf{u})(\mathbf{x}) = \mathbf{u}(g^{-1} \cdot \mathbf{x})$, which models how geometric transformations deform images. This action further extends to derivatives through prolongation to the jet space $X \times U^{(d)}$. For example, the first-order prolongation of the action on $X \times U^{(1)}$ can be expressed as: $(\mathbf{x}, \mathbf{u}(\mathbf{x}), \nabla \mathbf{u}(\mathbf{x})) \mapsto (\tilde{\mathbf{x}}, \tilde{\mathbf{u}}(\tilde{\mathbf{x}}), \nabla \tilde{\mathbf{u}}(\tilde{\mathbf{x}}))$, where $\tilde{\mathbf{x}} \triangleq g \cdot \mathbf{x}$ and $\tilde{\mathbf{u}} \triangleq g \cdot \mathbf{u}$, with $\tilde{\mathbf{u}}(\tilde{\mathbf{x}}) = \mathbf{u}(\mathbf{x})$.

A differential invariant is a quantity that remains unchanged under the prolonged group action. The definition is given below.

Definition 2 Given a group G acting on X , a d -th order **differential invariant** is a function $\mathcal{I} : X \times U^{(d)} \rightarrow \mathbb{R}$ such that

$$\mathcal{I}(g \cdot (\mathbf{x}, \mathbf{u}^{(d)})) = \mathcal{I}(\mathbf{x}, \mathbf{u}^{(d)}), \quad \forall g \in G, (\mathbf{x}, \mathbf{u}^{(d)}) \in X \times U^{(d)}, \quad (2)$$

where $g \cdot (\mathbf{x}, \mathbf{u}^{(d)})$ denotes the prolonged group action on the jet space $X \times U^{(d)}$.

The definition can be extended to the multi-dimensional case. Specifically, we call $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_k)^\top$ an **k -dimensional differential invariant**. In addition, we define **relative differential invariants**, which may transform with a weight function under the group action:

$$\mathcal{R}(g \cdot (\mathbf{x}, \mathbf{u}^{(d)})) = w(g, \mathbf{x}) \cdot \mathcal{R}(\mathbf{x}, \mathbf{u}^{(d)}), \quad (3)$$

where $w(g, \mathbf{x})$ is a scalar weight depending on the group element g and the point \mathbf{x} . Notably, differential invariants are closely tied to equivariance, as a (multi-dimensional) differential invariant \mathcal{I} yields an equivariant operator $\hat{\mathcal{I}}(\mathbf{u})(\mathbf{x}) \triangleq \mathcal{I}(\mathbf{x}, \mathbf{u}^{(d)})$ satisfying $\hat{\mathcal{I}}(g \cdot \mathbf{u}) = g \cdot \hat{\mathcal{I}}(\mathbf{u})$.

In this work, we focus on constructing such differential invariants and using them to build projective equivariant operators for neural networks.

2.2 Method of moving frames

The method of moving frames is a powerful technique for deriving differential invariants [Olver, 2003, 2015]. We begin with the definition of a moving frame.

Definition 3 [Olver, 2015] Let G be a Lie group acting on a manifold \mathcal{M} . A **moving frame** is a map $\eta : \mathcal{M} \rightarrow G$ such that

$$\eta(g \cdot z) = \eta(z) \cdot g^{-1}, \quad g \in G, z \in \mathcal{M}. \quad (4)$$

Given a moving frame, the **invariantization** of a function $F : \mathcal{M} \rightarrow \mathbb{R}$ is defined as

$$\iota(F)(z) \triangleq F(\eta(z) \cdot z), \quad (5)$$

which converts an arbitrary function F into a group-invariant function satisfying $\iota(F)(g \cdot z) = \iota(F)(z)$. More generally, we can define an invariant as $\mathcal{I}(g \cdot z) = \mathcal{I}(z)$, $g \in G, z \in \mathcal{M}$. In our context, the manifold of interest is the jet space $\mathcal{M} = X \times U^{(n)}$ and we focus on differential invariants.

A necessary and sufficient condition for the existence of a moving frame is that the group G acts freely and regularly on the manifold \mathcal{M} . Under this condition, a moving frame can be constructed via a cross section, as described below:

Theorem 4 [Olver, 2015] Let G be a r -dimensional Lie group acting freely and regularly on a m -dimensional manifold \mathcal{M} . Given local coordinates $z = (z_1, \dots, z_m)$ on \mathcal{M} , let \mathcal{K} be a cross section of the form $\mathcal{K} = \{z_1 = c_1, z_2 = c_2, \dots, z_r = c_r\} \subset \mathcal{M}$, where c_i are constants. Then for $z \in \mathcal{M}$, there exists a unique $g \in G$ such that $g \cdot z \in \mathcal{K}$. Defining $\eta(z) = g$, namely $\eta(z) \cdot z \in \mathcal{K}$, yields a map $\eta : \mathcal{M} \rightarrow G$, which is a moving frame.

Here, the group action is said to be **free** if for any $g \in G$, $g \cdot z = z$ implies $g = e$, where e is the identity element of the group. Usually, the group action can be made free by increasing the order of the jet space. The action is **regular** if the orbits form a regular foliation, which is typically satisfied in common groups.

With a moving frame obtained from Theorem 4, we can construct a complete set of fundamental invariants, meaning any invariant can be expressed as a combination of these fundamental invariants.

Theorem 5 [Olver, 2015] *Let $\eta : \mathcal{M} \rightarrow G$ be a moving frame from Theorem 4 and define $w(g, z) \triangleq g \cdot z$. Then*

$$w(\eta(z), z) = (c_1, c_2, \dots, c_r, w_{r+1}(\eta(z), z), \dots, w_m(\eta(z), z)), \quad (6)$$

where $\mathcal{I}_1(z) \triangleq w_{r+1}(\eta(z), z), \dots, \mathcal{I}_{m-r}(z) \triangleq w_m(\eta(z), z)$ constitute a complete system of functionally independent invariants, called **fundamental invariants**.

This theorem provides a method to construct fundamental invariants via the moving frame and indicates that the number of fundamental invariants is $m - r$. In the following sections, we will leverage these results to derive projective differential invariants.

2.3 Projective transformation

Projective transformations are ubiquitous in the visual world as two different views of the same planar object can be related by a 2D projective transformation. A standard projective group action is described by the projective special linear group $\text{PSL}(3, \mathbb{R})$ acting on the 2D projective plane \mathbb{RP}^2 , which can be interpreted as the set of equivalence classes of points $(x, y, p) \sim (cx, cy, cp)$ for any $c \neq 0$. Points in \mathbb{RP}^2 with $p \neq 0$ can be represented in inhomogeneous coordinates as (x, y) , corresponding to the homogeneous coordinate $(x, y, 1)$. Thus, the action of a projective transformation on 2D coordinates can be written as

$$\tilde{x} = \frac{\alpha x + \beta y + \gamma}{\rho x + \sigma y + \tau}, \tilde{y} = \frac{\lambda x + \mu y + \nu}{\rho x + \sigma y + \tau}, \quad (7)$$

where the transformation is parameterized by the coefficient matrix

$$\mathbf{P} = \begin{pmatrix} \alpha & \beta & \gamma \\ \lambda & \mu & \nu \\ \rho & \sigma & \tau \end{pmatrix}. \quad (8)$$

Since the transformation is defined up to a nonzero scaling factor, we can normalize by requiring the determinant of \mathbf{P} to be 1, i.e., $\Delta = \det(\mathbf{P}) = 1$. Thus, there are 8 independent degrees of freedom. The transformation reduces to an affine transformation when $\rho = \sigma = 0$, while a pure projective transformation, characterized by $\rho^2 + \sigma^2 \neq 0$, exhibits nonlinear behavior. Thus, projective transformations represent a more general and complex class of geometric transformations.

For an n -dimensional function $\mathbf{u}(x, y)$, the projective transformation of coordinates induces a natural action on the function, $\tilde{\mathbf{u}}(\tilde{x}, \tilde{y}) = \mathbf{u}(x, y)$, which can be further prolonged to its derivatives. We denote the derivatives of the i -th component function $u^{[i]}$ as

$$u_{jk}^{[i]} \triangleq D_x^j D_y^k u^{[i]}, \quad (9)$$

where D_x and D_y are the differentiation operators with respect to x and y , respectively. Under a projective transformation, these derivatives transform as

$$u_{jk}^{[i]} \mapsto \tilde{u}_{jk}^{[i]} = D_{\tilde{x}}^j D_{\tilde{y}}^k \tilde{u}^{[i]}, \quad (10)$$

with the transformed differential operators given by

$$D_{\tilde{x}} = \frac{\rho x + \sigma y + \tau}{\Delta} (((\mu\rho - \lambda\sigma)x + \mu\tau - \nu\sigma)D_x + ((\mu\rho - \lambda\sigma)y - \lambda\tau + \nu\rho)D_y), \quad (11)$$

$$D_{\tilde{y}} = \frac{\rho x + \sigma y + \tau}{\Delta} (((\alpha\sigma - \beta\rho)x - \beta\tau + \gamma\sigma)D_x + ((\alpha\sigma - \beta\rho)y + \alpha\tau - \gamma\rho)D_y). \quad (12)$$

2.4 Projective differential invariants of multi-dimensional functions

The projective group action is not free on the second-order jet space for scalar functions, requiring prolongation to the third-order jet space to achieve freeness. This leads to complex formulations [Olver, 2023], which may limit the practicality of the resulting invariants due to their complexity and computational cost. Moreover, in practice, third-order derivatives are harder to estimate reliably from data than lower-order ones. In this work, we focus on multi-dimensional functions, which naturally align with applications such as color image processing. In this setting, the group action is free on the second-order jet space, allowing the existence of second-order differential invariants. Using the method of moving frames, we can derive these invariants, where the choice of cross section significantly influences the simplicity of the resulting expressions. Although the cross section proposed by Olver [2023] can be extended to the multi-dimensional case, the resulting invariants tend to be lengthy, typically involving hundreds of terms, which makes them less practical. Instead, we propose an alternative cross section that leverages multiple dimensions while relying only on derivatives up to second order, yielding invariants with significantly more concise and tractable forms.

Specifically, we choose the following cross section:

$$\mathcal{K} = \{x = y = 0, u_x^{[1]} = 1, u_y^{[1]} = 0, u_x^{[2]} = 0, u_y^{[2]} = 1, u_{xx}^{[1]} = u_{xy}^{[1]} = 0\}. \quad (13)$$

Hereafter, we use the standard shorthand notation for partial derivatives, e.g., $u_x, u_y, u_{xx}, u_{xy}, u_{yy}$. The above cross section defines 8 normalization equations, which, together with the constraint $\Delta = 1$, determine all group parameters, thereby establishing the moving frame η . The detailed derivation of the moving frame is provided in the Appendix.

With the moving frame η constructed, we can then apply the invariantization process according to Theorem 5 to obtain a complete set of fundamental differential invariants. For the coordinates involved in the cross section, we have

$$\begin{aligned} \iota(x) &= 0, \iota(y) = 0, \iota(u_x^{[1]}) = 1, \iota(u_y^{[1]}) = 0, \\ \iota(u_x^{[2]}) &= 0, \iota(u_y^{[2]}) = 1, \iota(u_{xx}^{[1]}) = 0, \iota(u_{xy}^{[1]}) = 0. \end{aligned}$$

The remaining coordinates of the second-order jet space yield the following differential invariants:

$$\iota(u_{yy}^{[1]}) = \frac{\mathcal{T}_{111}}{\mathcal{J}_{12}^2}, \quad (14)$$

$$\iota(u_{xx}^{[2]}) = \frac{\mathcal{T}_{222}}{\mathcal{J}_{12}^2}, \quad (15)$$

$$\iota(u_{xy}^{[2]}) = -\frac{\mathcal{T}_{212} + 2\mathcal{T}_{122}}{2\mathcal{J}_{12}^2}, \quad (16)$$

$$\iota(u_{yy}^{[2]}) = \frac{\mathcal{T}_{121} + 2\mathcal{T}_{112}}{\mathcal{J}_{12}^2}, \quad (17)$$

$$\iota(u_x^{[i]}) = -\frac{\mathcal{J}_{2i}}{\mathcal{J}_{12}}, \quad 3 \leq i \leq n, \quad (18)$$

$$\iota(u_y^{[i]}) = \frac{\mathcal{J}_{1i}}{\mathcal{J}_{12}}, \quad 3 \leq i \leq n, \quad (19)$$

$$\iota(u_{xx}^{[i]}) = \frac{\mathcal{J}_{12}\mathcal{T}_{2i2} + \mathcal{J}_{2i}\mathcal{T}_{212}}{\mathcal{J}_{12}^3}, \quad 3 \leq i \leq n, \quad (20)$$

$$\iota(u_{xy}^{[i]}) = -\frac{2\mathcal{J}_{12}\mathcal{T}_{1i2} + 2\mathcal{J}_{12}\mathcal{T}_{21i} + 3\mathcal{J}_{2i}\mathcal{T}_{112}}{\mathcal{J}_{12}^3}, \quad 3 \leq i \leq n, \quad (21)$$

$$\iota(u_{yy}^{[i]}) = \frac{\mathcal{J}_{12}\mathcal{T}_{1i1} + 2\mathcal{J}_{1i}\mathcal{T}_{112}}{\mathcal{J}_{12}^3}, \quad 3 \leq i \leq n, \quad (22)$$

where \mathcal{J}_{ij} and \mathcal{T}_{ijk} are two key quantities defined as:

$$\mathcal{J}_{ij} \triangleq u_x^{[i]}u_y^{[j]} - u_x^{[j]}u_y^{[i]}, \quad (23)$$

$$\mathcal{T}_{ijk} \triangleq u_{xx}^{[i]}u_y^{[j]}u_y^{[k]} + u_{yy}^{[i]}u_x^{[j]}u_x^{[k]} - u_{xy}^{[i]}(u_x^{[j]}u_y^{[k]} + u_x^{[k]}u_y^{[j]}), \quad (24)$$

satisfying $\mathcal{J}_{ii} = 0$, $\mathcal{J}_{ij} = -\mathcal{J}_{ji}$, and $\mathcal{T}_{ijk} = \mathcal{T}_{kji}$.

The invariants (14)-(22), together with the obvious zeroth-order invariants

$$S_0 = \{u^{[i]} \mid 1 \leq i \leq n\},$$

form a complete set of second-order fundamental differential invariants of the projective group. Compared to projective invariants for scalar functions [Olver, 2023], our results involve up to second-order derivatives and are expressed in a more concise form.

2.5 Fundamental components of projective differential invariants

In the previous subsection, we have derived a complete set of second-order fundamental differential invariants. While relatively concise, their expressions are asymmetric and depend on the specific choice of the first two dimensions used in the cross section. To obtain a simpler, more unified, and elegant formulation, we conduct a deeper analysis of the fundamental components of these invariants. This enables us to further simplify their structure while preserving completeness.

Note that the numerators and denominators in (14)-(22) are all relative invariants. Thus, we focus on the properties of these relative invariants, as absolute invariants can be obtained by taking the ratio of two relative invariants with the same weight. Moreover, since the expressions are built from the basic quantities \mathcal{J}_{ij} and \mathcal{T}_{ijk} , we will delve into their transformation properties and algebraic relationships.

We first present three classes of simplified relative differential invariants of the projective group.

Theorem 6 *Let $W = \frac{(\rho x + \sigma y + \tau)^3}{\Delta}$. Then the following quantities are relative differential invariants of the projective group:*

- For $i \neq j$, \mathcal{J}_{ij} is a relative differential invariant of weight W .
- For $1 \leq i \leq n$, \mathcal{T}_{iii} is a relative differential invariant of weight W^2 .
- For $1 \leq i, j \leq n$, $\mathcal{T}_{iji} + 2\mathcal{T}_{iij}$ is a relative differential invariant of weight W^2 .

These relative invariants are not functionally independent; rather, they can be transformed into one another. Given that there are $6n - 6$ second-order fundamental differential invariants according to Section 2.4, we expect a complete and independent set of relative invariants to contain $6n - 5$ elements. To this end, we investigate the transformation rules among the relative invariants and aim to identify a minimal generating set sufficient to express all fundamental differential invariants. We start with the transformation properties of \mathcal{J}_{ij} .

Theorem 7 *For any indices $1 \leq i_1, i_2, i_3, i_4 \leq n$, the following equation holds:*

$$\mathcal{J}_{i_1 i_2} \cdot \mathcal{J}_{i_3 i_4} + \mathcal{J}_{i_1 i_3} \cdot \mathcal{J}_{i_4 i_2} + \mathcal{J}_{i_1 i_4} \cdot \mathcal{J}_{i_2 i_3} = 0. \quad (25)$$

This implies that for any four distinct indices i_1, i_2, i_3, i_4 , the six pairwise combinations of \mathcal{J}_{ij} are dependent such that once any five are known, the remaining one can be determined. Based on Theorem 7, we can construct a subset of $\{\mathcal{J}_{ij} \mid i \neq j\}$ that is sufficient to express all \mathcal{J}_{ij} .

Theorem 8 *Define the following sets of relative invariants:*

$$S_1 \triangleq \{\mathcal{J}_{12}, \mathcal{J}_{23}, \dots, \mathcal{J}_{n-1, n}\}, \quad (26)$$

$$S_2 \triangleq \{\mathcal{J}_{13}, \mathcal{J}_{24}, \dots, \mathcal{J}_{n-2, n}\}. \quad (27)$$

Then $S_1 \cup S_2$ is a generating set for the collection $\{\mathcal{J}_{ij} \mid i \neq j\}$, meaning that any \mathcal{J}_{ij} can be expressed as a functional combination of these elements.

According to Theorem 7, $\mathcal{J}_{i, i+3}$ can be written in terms of $\mathcal{J}_{i, i+1}$, $\mathcal{J}_{i+1, i+2}$, $\mathcal{J}_{i+2, i+3}$, $\mathcal{J}_{i, i+2}$, and $\mathcal{J}_{i+1, i+3}$, all of which belong to $S_1 \cup S_2$. By induction, any \mathcal{J}_{ij} can thus be recovered from the generating set. A complete and rigorous proof is provided in the Appendix.

Before establishing the transformation relationships for \mathcal{T}_{ijk} , we provide a more compact representation of \mathcal{J}_{ij} and \mathcal{T}_{ijk} to clarify their structural relationships:

$$\mathcal{J}_{ij} = \mathbf{g}_i^\top \mathbf{Q} \mathbf{g}_j, \quad (28)$$

$$\mathcal{T}_{ijk} = \mathbf{g}_i^\top \mathbf{Q}^\top \mathbf{H}_j \mathbf{Q} \mathbf{g}_k, \quad (29)$$

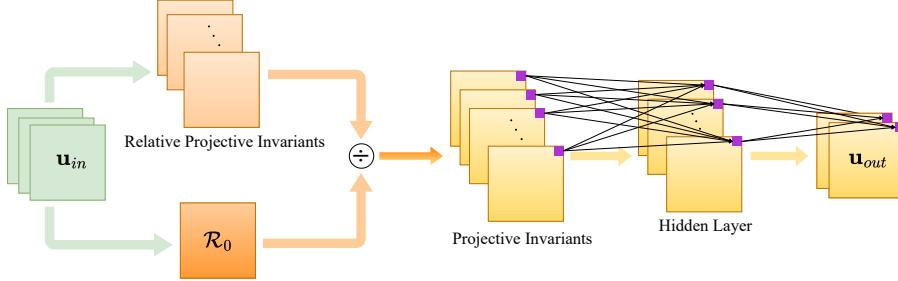


Figure 2: A single layer of PDINet. (1) Compute projective invariants from relative projective invariants. (2) Combine these invariants with an MLP to produce equivariant outputs.

where $\mathbf{g}_i \triangleq \left(u_x^{[i]}, u_y^{[i]} \right)^\top$ is the gradient of the i -th component function, \mathbf{H}_j is the Hessian matrix of the j -th component function, and \mathbf{Q} is a fixed orthogonal matrix defined as:

$$\mathbf{Q} \triangleq \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Notably, $\{\mathcal{J}_{ij}\}$ can be used to establish the transformation relationships between the gradients.

Lemma 9 For any distinct indices i, j, k , the gradient \mathbf{g}_k can be expressed in terms of \mathbf{g}_i and \mathbf{g}_j as:

$$\mathbf{g}_k = \frac{\mathcal{J}_{kj}}{\mathcal{J}_{ij}} \mathbf{g}_i + \frac{\mathcal{J}_{ki}}{\mathcal{J}_{ji}} \mathbf{g}_j. \quad (30)$$

Using Lemma 9 and the form of \mathcal{T}_{ijk} in (29), we derive the transformation rule for \mathcal{T}_{ijk} given $\{\mathcal{J}_{ij}\}$.

Theorem 10 Let $i \neq k$, and let i', j, k' be arbitrary indices. Then $\mathcal{T}_{i'jk'}$ can be expressed as:

$$\mathcal{T}_{i'jk'} = \frac{\mathcal{J}_{i'k}\mathcal{J}_{k'k}}{\mathcal{J}_{ik}^2} \mathcal{T}_{iji} - \frac{\mathcal{J}_{i'k}\mathcal{J}_{k'i} + \mathcal{J}_{i'i}\mathcal{J}_{k'k}}{\mathcal{J}_{ik}^2} \mathcal{T}_{ijk} + \frac{\mathcal{J}_{i'i}\mathcal{J}_{k'i}}{\mathcal{J}_{ik}^2} \mathcal{T}_{kjk}. \quad (31)$$

This result implies that, for a fixed j , the triplet $\{\mathcal{T}_{iji}, \mathcal{T}_{ijk}, \mathcal{T}_{kjk}\}$ serves as a generating set for $\{\mathcal{T}_{i'jk'} \mid 1 \leq i', k' \leq n\}$, provided $\{\mathcal{J}_{ij}\}$ is known.

With the transformation relationships for \mathcal{J}_{ij} and \mathcal{T}_{ijk} established, we can now construct a minimal set of relative invariants that suffices to express all the relative invariants in Theorem 6.

Theorem 11 Define the following sets of relative invariants:

$$S_3 \triangleq \{\mathcal{T}_{111}, \mathcal{T}_{222}, \dots, \mathcal{T}_{nnn}\}, \quad (32)$$

$$S_4 \triangleq \{\mathcal{T}_{121} + 2\mathcal{T}_{112}, \mathcal{T}_{131} + 2\mathcal{T}_{113}, \dots, \mathcal{T}_{n-1,n,n-1} + 2\mathcal{T}_{n-1,n-1,n}\}, \quad (33)$$

$$S_5 \triangleq \{\mathcal{T}_{212} + 2\mathcal{T}_{221}, \mathcal{T}_{323} + 2\mathcal{T}_{332}, \dots, \mathcal{T}_{n,n-1,n} + 2\mathcal{T}_{n,n,n-1}\}. \quad (34)$$

Then $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$ can express all the relative invariants in Theorem 6.

It can be shown (see the Appendix) that the fundamental differential invariants (14)-(22) can be expressed via the relative invariants in Theorem 6. Therefore, we arrive at the following conclusion:

Theorem 12 The union $S \triangleq S_0 \cup S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$ forms a complete set of relative differential invariants to express all second-order differential invariants of the projective group.

This set contains exactly $6n - 5$ elements, matching the expected minimal number needed to express all second-order fundamental differential invariants. Compared to the invariants derived in Section 2.4, the current formulation is further simplified and independent of the specific choice of the first two dimensions in the cross section, exhibiting a more unified structure. Moreover, while the original invariants involve polynomials of degree up to five, the present set only contains at most cubic expressions, resulting in lower computational complexity.

2.6 Projective equivariant networks

As discussed before, with a set of relative invariants obtained, we can convert them into absolute invariants by dividing each by another relative invariant with the same weight. We apply this procedure to relative invariants in S to construct a complete set of fundamental differential invariants.

Specifically, we select $\mathcal{R}_0 = \frac{1}{n}(\mathcal{J}_{12} + \mathcal{J}_{23} + \dots + \mathcal{J}_{n1})$ as the denominator, which is a relative invariant of weight W . We keep the elements in S_0 unchanged, divide the elements in $S_1 \cup S_2$ by \mathcal{R}_0 , and divide those in $S_3 \cup S_4 \cup S_5$ by \mathcal{R}_0^2 . This yields a set of differential invariants sufficient to express all second-order fundamental differential invariants. In fact, this set with $6n - 5$ invariants may contain one redundant element, but completeness is of greater concern. To avoid division by zero, we add a positive constant ϵ to the denominator during division, enhancing numerical stability.

Let $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_N)^\top$ denote the set of differential invariants we obtained, which naturally induces an equivariant operator $\hat{\mathcal{I}}$. Theoretically, the invariants $\mathcal{I}_1, \dots, \mathcal{I}_N$ are sufficient to express all second-order differential invariants. In practice, we leverage the expressive power of neural networks to combine $\mathcal{I}_1, \dots, \mathcal{I}_N$ using a two-layer MLP to produce the output [Li et al., 2024, 2025]. This leads to a learnable equivariant operator:

$$\mathbf{u}_{out} = \mathbf{h}_\theta \circ \hat{\mathcal{I}}(\mathbf{u}_{in}), \quad (35)$$

where \mathbf{h}_θ is an MLP parameterized by θ . By integrating this operator into standard network architectures, we can build projective equivariant models. We refer to the resulting model as the **Projective Differential Invariant Network (PDINet)**, as illustrated in Figure 2.

3 Experiments

For empirical evaluation, we conduct image classification tasks under out-of-distribution settings, where models are trained on the original dataset and tested on images deformed by projective transformations. We adopt ResNet-18 [He et al., 2016] as the backbone and replace its convolutional layers with our equivariant operators defined in (35) to construct a projective equivariant network, PDINet. As the main counterpart, we consider homConv, a projective equivariant model proposed by MacDonald et al. [2022]. To ensure a fair comparison, we attempted to implement homConv using the same backbone. However, homConv relies on group sampling, which leads to exponential memory growth with network depth, resulting in out-of-memory (OOM) issues. Therefore, we follow the original network configuration of MacDonald et al. [2022] and reduce the number of samples to avoid OOM. In addition, we also include ResNet-18 trained with projective data augmentation (DA) as a reference baseline.

3.1 Proj-STL-10

STL-10 [Coates et al., 2011] is a dataset containing 5000 training images and 8000 test images. Each image has a resolution of 96×96 with RGB channels. We apply random projective transformations to the test set to generate the Proj-STL-10 dataset. Models are trained on the original STL-10 dataset (or with projective data augmentation for the DA baseline) and evaluated on Proj-STL-10, forming a challenging out-of-distribution setting that assesses the model’s ability to generalize beyond the training distribution.

Table 1: Test accuracy (%) on Proj-STL-10.

Model	Accuracy	# Params
ResNet-18	39.73 ± 0.33	11.18M
ResNet-18 with DA	48.73 ± 0.53	11.18M
homConv	20.88 ± 0.32	376K
PDINet (ours)	51.12 ± 0.47	8.90M

Each experiment is repeated five times with different random seeds, and we report the average accuracy and standard deviation in Table 1. PDINet substantially outperforms the standard ResNet-18 on the transformed test set, achieving an 11.39% improvement in accuracy. While projective data augmentation considerably enhances ResNet-18’s robustness, PDINet still achieves the highest accuracy. In contrast, homConv performs poorly due to its restriction on network depth.

3.2 Proj-Imagenette

Imagenette is a ten-class subset of the ImageNet dataset [Deng et al., 2009], consisting of 9469 training images and 3925 test images. All images are adapted to a uniform resolution of 256×256 for model input. We apply random projective transformations to the test set to generate the Proj-Imagenette dataset, while keeping the training set unchanged. This setup simulates an out-of-distribution scenario and evaluates the model’s ability to generalize to geometric transformations.

Table 2: Test accuracy (%) on Proj-Imagenette.

Model	Accuracy	# Params
ResNet-18	65.64 ± 0.39	11.18M
ResNet-18 with DA	70.77 ± 0.76	11.18M
homConv	25.72 ± 0.52	376K
PDINet (ours)	71.30 ± 0.45	8.90M

Each experiment is repeated five times with different random seeds, and we report the mean \pm standard deviation of test accuracy in Table 2. The results demonstrate that PDINet retains strong performance under distribution shift, outperforming the standard ResNet-18 by 5.66%. It confirms the effectiveness of the projective equivariance of our model in enhancing out-of-distribution generalization. Although ResNet-18 with DA reduces the performance gap, it still lags behind PDINet. Meanwhile, the shallow homConv model struggles to handle higher-resolution inputs due to its limited depth.

Detailed experimental settings and implementation details can be found in the Appendix.

4 Conclusion

In this work, we propose **PDINet**, a framework for projective equivariant networks, based on second-order differential invariants of the projective group. Our method overcomes the exponential memory growth encountered by homConv [MacDonald et al., 2022], enabling effective scaling to deeper networks. Leveraging the moving frame method and a carefully chosen cross section tailored to multi-dimensional functions, we derive a complete and concise set of second-order projective fundamental differential invariants. Further analysis reveals transformation relationships among projective invariants, allowing us to obtain a unified and simplified formulation that enhances both theoretical clarity and computational efficiency. Building upon these invariants, we design a learnable projective equivariant operator that can be seamlessly integrated into various network architectures. It is the first time to achieve full projective equivariance in deep networks without group sampling or discretization. Experiments under out-of-distribution settings demonstrate the strong generalization ability of our model. With the prevalence and significance of projective transformations in vision, PDINet holds promising potential for broader applications in computer vision.

One limitation of our approach is that the second-order invariants we derive vanish in the one-dimensional case, preventing the direct application of PDINet to grayscale images. In addition, this work focuses on group actions on scalar fields and does not yet cover more general cases involving arbitrary group representations, which we consider a valuable direction for future research.

Acknowledgments

Z. Lin was supported by National Key R&D Program of China (2022ZD0160300), the NSF China (No. 62276004) and the State Key Laboratory of General Artificial Intelligence.

References

- Stan Birchfield. An introduction to projective geometry (for computer vision). *Unpublished note, Stanford University*, 14, 1998.
- Georg Bökman, Axel Flinth, and Fredrik Kahl. In search of projectively equivariant networks. *Transactions on Machine Learning Research*, 2023.

Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 834–848, 2017.

Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.

Taco S Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference on Machine Learning*, 2016a.

Taco S Cohen and Max Welling. Steerable CNNs. In *International Conference on Learning Representations*, 2016b.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.

Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.

Carlos Esteves, Yinshuang Xu, Christine Allen-Blanchette, and Kostas Daniilidis. Equivariant multi-view networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.

Mark Fels and Peter J Olver. Moving coframes: II. regularization and theoretical foundations. *Acta Applicandae Mathematica*, 55:127–208, 1999.

Marc Finzi, Samuel Stanton, Pavel Izmailov, and Andrew Gordon Wilson. Generalizing convolutional neural networks for equivariance to Lie groups on arbitrary continuous data. In *International Conference on Machine Learning*, 2020.

Marc Finzi, Max Welling, and Andrew Gordon Wilson. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In *International Conference on Machine Learning*, 2021.

Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. SE(3)-transformers: 3D roto-translation equivariant attention networks. In *Advances in Neural Information Processing Systems*, 2020.

Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2003.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.

Lingshen He, Yuxuan Chen, Zhengyang Shen, Yiming Dong, Yisen Wang, and Zhouchen Lin. Efficient equivariant network. In *Advances in Neural Information Processing Systems*, 2021.

Lingshen He, Yuxuan Chen, Zhengyang Shen, Yibo Yang, and Zhouchen Lin. Neural ePDOs: Spatially adaptive equivariant partial differential operator based networks. In *International Conference on Learning Representations*, 2022.

Boce Hu, Xupeng Zhu, Dian Wang, Zihao Dong, Haojie Huang, Chenghao Wang, Robin Walters, and Robert Platt. Orbitgrasp: SE(3)-equivariant grasp learning. *arXiv preprint arXiv:2407.03531*, 2024.

Boce Hu, Heng Tian, Dian Wang, Haojie Huang, Xupeng Zhu, Robin Walters, and Robert Platt. Push-grasp policy learning using equivariant models and grasp score optimization. *arXiv preprint arXiv:2504.03053*, 2025a.

Boce Hu, Dian Wang, David Klee, Heng Tian, Xupeng Zhu, Haojie Huang, Robert Platt, and Robin Walters. 3D equivariant visuomotor policy learning via spherical projection. *arXiv preprint arXiv:2505.16969*, 2025b.

Lexiang Hu, Yikang Li, and Zhouchen Lin. Explicit discovery of nonlinear symmetries from dynamic data. In *International Conference on Machine Learning*, 2025c.

Lexiang Hu, Yikang Li, and Zhouchen Lin. Governing equation discovery from data based on differential invariants. In *ICML 2025 2nd AI for Math Workshop*, 2025d.

Lexiang Hu, Yikang Li, and Zhouchen Lin. Symmetry discovery for different data types. *Neural Networks*, 188:107481, 2025e.

Lexiang Hu, Yisen Wang, and Zhouchen Lin. Incorporating arbitrary matrix group equivariance into KANs. In *International Conference on Machine Learning*, 2025f.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 2012.

Maxime W Lafarge, Erik J Bekkers, Josien PW Pluim, Remco Duits, and Mitko Veta. Roto-translation equivariant convolutional networks: Application to histopathology image analysis. *Medical Image Analysis*, 68:101849, 2021.

Jongmin Lee, Byungjin Kim, and Minsu Cho. Self-supervised equivariant learning for oriented keypoint detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Wang-Heun Lee, Kyoung-Sig Roh, and In-So Kweon. Self-localization of a mobile robot without camera calibration using projective invariants. *Pattern Recognition Letters*, 21(1):45–60, 2000.

Erbo Li, Hanlin Mo, Dong Xu, and Hua Li. Image projective invariants. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(5):1144–1157, 2018.

Yikang Li, Yeqing Qiu, Yuxuan Chen, Lingshen He, and Zhouchen Lin. Affine equivariant networks based on differential invariants. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Yikang Li, Yeqing Qiu, Yuxuan Chen, and Zhouchen Lin. Affine steerable equivariant layer for canonicalization of neural networks. In *International Conference on Learning Representations*, 2025.

Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3D atomistic graphs. In *International Conference on Learning Representations*, 2022.

Risheng Liu, Zhouchen Lin, Wei Zhang, and Zhixun Su. Learning PDEs for image restoration via optimal control. In *Proceedings of the European Conference on Computer Vision*, 2010.

Risheng Liu, Zhouchen Lin, Wei Zhang, Kewei Tang, and Zhixun Su. Toward designing intelligent PDEs for computer vision: an optimal control approach. *Image and Vision Computing*, 31(1): 43–56, 2013.

Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljacic, Thomas Y Hou, and Max Tegmark. KAN: Kolmogorov–Arnold networks. In *International Conference on Learning Representations*, 2025.

Lachlan E MacDonald, Sameera Ramasinghe, and Simon Lucey. Enabling equivariance for arbitrary Lie groups. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Mircea Mironenco and Patrick Forré. Lie group decompositions for equivariant neural networks. In *International Conference on Learning Representations*, 2024.

Thomas W Mitchel, Noam Aigerman, Vladimir G Kim, and Michael Kazhdan. Möbius convolutions for spherical CNNs. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022.

Roger Mohr and Bill Triggs. Projective geometry for image analysis. In *XVIIIth International Symposium on Photogrammetry & Remote Sensing (ISPRS'96)*, 1996.

Joseph L Mundy and Andrew Zisserman. *Geometric invariance in computer vision*. MIT Press, 1992.

Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: A versatile and accurate monocular slam system. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015.

Peter J Olver. *Applications of Lie groups to differential equations*, volume 107. Springer Science & Business Media, 1993.

Peter J Olver. Moving frames. *Journal of Symbolic Computation*, 36(3-4):501–512, 2003.

Peter J Olver. Modern developments in the theory and applications of moving frames. *London Math. Soc. Impact150 Stories*, 1:14–50, 2015.

Peter J Olver. Projective invariants of images. *European Journal of Applied Mathematics*, 34(5):936–946, 2023.

Peter J Olver, Guillermo Sapiro, and Allen Tannenbaum. Affine invariant detection: edge maps, anisotropic diffusion, and active contours. *Acta Applicandae Mathematica*, 59:45–77, 1999.

David W Romero and Jean-Baptiste Cordonnier. Group equivariant stand-alone self-attention for vision. In *International Conference on Learning Representations*, 2021.

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, and Jesús Angulo. Differential invariants for SE(2)-equivariant networks. In *IEEE International Conference on Image Processing*, 2022.

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, and Jesus Angulo. Moving frame net: SE(3)-equivariant network for volumes. In *NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, 2023.

Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518, 2016.

Pourya Shamsolmoali, Masoumeh Zareapoor, Jocelyn Chanussot, Huiyu Zhou, and Jie Yang. Rotation equivariant feature image pyramid network for object detection in optical remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 2021.

Zhengyang Shen, Lingshen He, Zhouchen Lin, and Jinwen Ma. PDO-eConvs: Partial differential operator based equivariant convolutions. In *International Conference on Machine Learning*, 2020.

Zhengyang Shen, Tao Hong, Qi She, Jinwen Ma, and Zhouchen Lin. PDO-s3DCNNs: Partial differential operator based steerable 3D CNNs. In *International Conference on Machine Learning*, 2022.

Zhengyang Shen, Yeqing Qiu, Jialun Liu, Lingshen He, and Zhouchen Lin. Efficient learning of scale-adaptive nearly affine invariant networks. *Neural Networks*, 174:106229, 2024.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.

Ivan Sosnovik, Michał Szmaja, and Arnold Smeulders. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2019.

Tomas Suk and Jan Flusser. Projective moment invariants. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1364–1367, 2004.

Dian Wang and Robin Walters. SO(2) equivariant reinforcement learning. In *International Conference on Learning Representations*, 2022.

Rui Wang, Robin Walters, and Rose Yu. Incorporating symmetry into deep dynamics models for improved generalization. In *International Conference on Learning Representations*, 2020.

Maurice Weiler and Gabriele Cesa. General E(2)-equivariant steerable CNNs. In *Advances in Neural Information Processing Systems*, 2019.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, 2018.

Daniel Worrall and Gabriel Brostow. Cubenet: Equivariance to 3D rotation and translation. In *Proceedings of the European Conference on Computer Vision*, 2018.

Jianke Yang, Robin Walters, Nima Dehmamy, and Rose Yu. Generative adversarial symmetry discovery. In *International Conference on Machine Learning*, 2023.

Jianke Yang, Nima Dehmamy, Robin Walters, and Rose Yu. Latent space symmetry discovery. In *International Conference on Machine Learning*, 2024.

Linfeng Zhao, Xupeng Zhu, Lingzhi Kong, Robin Walters, and Lawson LS Wong. Integrating symmetry into differentiable planning with steerable convolutions. In *International Conference on Learning Representations*, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The claims in the abstract and introduction strictly follow the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the limitations of the work in Section 4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: All theorems in Section 2 are presented with the full set of assumptions. We also provide the complete proofs in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide all the information for experimental reproduction in Section 3 and the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We will make the data and code publicly available upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so ?No? is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide the experimental setup and implementation details in Section 3 and the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars are reported for all results in Section 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We provide sufficient information on the computer resources in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: Our work conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: Our paper focuses on the research of equivariant deep learning algorithms, and we do not foresee immediate positive or negative societal outcomes.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: For the code and dataset used in the paper, we respect the license and terms of use, and cite the relevant papers properly.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Related work

A.1 Differential invariants

Differential invariants are quantities involving derivatives that remain unchanged under the action of a given transformation group [Olver, 1993]. They have proven useful in various image analysis applications [Mundy and Zisserman, 1992, Olver et al., 1999, Hartley and Zisserman, 2003, Li et al., 2018]. Leveraging the connection between differential invariants and symmetric partial differential equations (PDEs) [Olver, 1993], Liu et al. [2010, 2013] constructed learnable PDEs as linear combinations of differential invariants, achieving shift and rotation equivariance. This connection has also been explored for symmetry-informed discovery of governing PDEs [Hu et al., 2025d]. The method of moving frames offers a systematic framework for deriving differential invariants given a transformation group [Fels and Olver, 1999, Olver, 2003, 2015]. Based on this approach, Olver [2023] provided a characterization of differential invariants of the projective group. However, since the projective group does not act freely on the second-order jet space of scalar functions, Olver [2023] prolonged the group action to the third-order jet space, resulting in complex expressions that are computationally expensive and difficult to estimate from data due to the involvement of higher-order derivatives. Differently, our work focuses on multi-dimensional functions, where the projective group acts freely on the second-order jet space. This allows us to derive a complete and concise set of second-order fundamental differential invariants, which we further simplify through structural analysis.

A.2 Equivariant networks

Early developments in equivariant networks primarily focused on designing architectures that are equivariant to discrete groups, such as cyclic or dihedral groups. A representative approach is the G-CNN [Cohen and Welling, 2016a] framework, which models feature maps as functions defined on a group and implements convolution-like operations via discrete permutations over the input domain. Building upon this viewpoint, subsequent works [Shen et al., 2020, Romero and Cordonnier, 2021, He et al., 2021] developed a broader range of group-equivariant operators beyond standard group convolutions. To achieve equivariance with respect to continuous groups, steerable CNNs [Cohen and Welling, 2016b, Weiler and Cesa, 2019] were proposed, which treat feature maps as fields transforming according to specified group representations. This framework enables principled handling of Euclidean groups in 2D and 3D settings [Weiler and Cesa, 2019, Fuchs et al., 2020, Wang et al., 2020, Zhao et al., 2022, Shen et al., 2022, Liao and Smidt, 2022, Hu et al., 2024, 2025a,b], by explicitly encoding the group action on the feature spaces.

Symmetries are ubiquitous and can be leveraged in model design, either from prior knowledge or discovered from data [Yang et al., 2023, 2024, Hu et al., 2025c,e]. However, generalizing equivariant architectures to high-dimensional Lie groups, such as the projective group, poses significant challenges. For such groups, deriving explicit steerable basis filters is often intractable, limiting the applicability of steerable CNNs, while G-CNN based methods require group discretization or sampling, introducing scalability issues. Bökman et al. [2023] explored equivariance in a projective sense and achieved equivariant models with respect to projective representations of certain simple groups, rather than the full projective group. Another related effort is Möbius Convolution (MC) [Mitchel et al., 2022], which achieves equivariance to spherical Möbius transformations for geometry and spherical image processing tasks. While both MC and our method aim to build equivariant models under certain projective group actions, MC operates on the Riemann sphere under the complex-valued $SL(2, \mathbb{C})$, whereas our work considers planar projective transformations governed by $SL(3, \mathbb{R})$, corresponding to the commonly studied projective transformations in vision. Additionally, EMLP [Finzi et al., 2021] and EKAN [Hu et al., 2025f] incorporate arbitrary matrix group equivariance into MLPs and KANs [Liu et al., 2025], respectively, but they target linear transformations on vectors or tensors and are not directly suitable for image inputs. Shen et al. [2024] combined rotation-equivariant networks with data augmentation to obtain nearly affine invariant features, without achieving full equivariance. LieConv [Finzi et al., 2020] attempts to handle Lie group equivariance by sampling from the Haar measure, but struggles to extend to complex groups due to inaccessibility to the Haar measure. To mitigate this problem, MacDonald et al. [2022] performed group convolutions by computing the integral on the Lie algebra, resulting in a projective equivariant model, homConv. Nevertheless, it suffers from exponential memory growth as network depth increases, making it impractical for deep networks. Mironenco and Forré [2024] improved sampling efficiency via group decompositions, but

their method remains constrained to affine subgroups. Besides these two mainstream approaches, G-CNNs and steerable CNNs, Li et al. [2024, 2025] resorted to another route by constructing equivariant networks based on differential invariants, achieving affine equivariance without requiring group sampling or discretization. However, their method is confined to affine groups and does not generalize to the projective group. In particular, projective differential invariants for multi-channel inputs have not yet been developed, and the SupNorm normalization technique used to construct affine invariants cannot be directly extended to the projective group. In our work, we target projective equivariance by deriving a complete set of second-order fundamental differential invariants for multi-channel inputs using a tailored moving frame construction. Through algebraic analysis, we obtain simplified projective invariants for practical use, which enable the construction of deep projective equivariant networks without relying on sampling or discretizing the group.

B Detailed proofs

In this section, we provide the detailed derivation of the moving frame and complete proofs of Theorems 8, 11, and 12.

For completeness, we briefly comment on the results that can be verified primarily through direct computation, while omitting the detailed algebraic steps:

- Theorem 6 follows directly from the definitions of \mathcal{J}_{ij} and \mathcal{T}_{ijk} , together with the transformation rules for first- and second-order derivatives under the group action derived from (11)-(12), followed by straightforward simplification.
- Theorem 7 and Lemma 9 can be verified by substituting the definition of \mathcal{J}_{ij} and expanding the expressions algebraically.
- Theorem 10 can be proved by rewriting $\mathcal{T}_{i'jk'}$ in the compact form $\mathbf{g}_{i'}^\top \mathbf{Q}^\top \mathbf{H}_j \mathbf{Q} \mathbf{g}_{k'}$ as shown in (29), then applying Lemma 9 to express the gradient vectors $\mathbf{g}_{i'}$ and $\mathbf{g}_{k'}$ in terms of \mathbf{g}_i and \mathbf{g}_k , and simplifying the resulting expression.

B.1 Derivation of the moving frame

We begin by setting $\tilde{x} = \tilde{y} = 0$, yielding

$$\gamma = -\alpha x - \beta y, \quad (36)$$

$$\nu = -\lambda x - \mu y. \quad (37)$$

Next, imposing $\tilde{u}_{\tilde{x}}^{[1]} = 1, \tilde{u}_{\tilde{y}}^{[1]} = 0$ gives

$$\alpha = u_x^{[1]}(\rho x + \sigma y + \tau), \quad (38)$$

$$\beta = u_y^{[1]}(\rho x + \sigma y + \tau). \quad (39)$$

Similarly, enforcing $\tilde{u}_{\tilde{x}}^{[2]} = 0, \tilde{u}_{\tilde{y}}^{[2]} = 1$ leads to

$$\lambda = u_x^{[2]}(\rho x + \sigma y + \tau), \quad (40)$$

$$\mu = u_y^{[2]}(\rho x + \sigma y + \tau). \quad (41)$$

We then apply the conditions $\tilde{u}_{\tilde{x}\tilde{x}}^{[1]} = \tilde{u}_{\tilde{x}\tilde{y}}^{[1]} = 0$, which yield the equations for σ and τ .

$$\sigma = \rho \frac{\mathcal{J}_{12}(u_{yy}^{[1]}u_x^{[2]} - u_{xy}^{[1]}u_y^{[2]}) + u_y^{[2]}\mathcal{T}_{112}}{-u_x^{[1]}\mathcal{T}_{212} + 2u_x^{[2]}\mathcal{T}_{112}}, \quad (42)$$

$$\tau = \rho \frac{x(u_x^{[1]}\mathcal{T}_{212} - 2u_x^{[2]}\mathcal{T}_{112}) + y(u_y^{[1]}\mathcal{T}_{212} - 2u_y^{[2]}\mathcal{T}_{112})}{-u_x^{[1]}\mathcal{T}_{212} + 2u_x^{[2]}\mathcal{T}_{112}}. \quad (43)$$

Finally, setting $\Delta = 1$ determines the parameter ρ as:

$$\rho = \frac{u_x^{[1]}\mathcal{T}_{212} - 2u_x^{[2]}\mathcal{T}_{112}}{2\mathcal{J}_{12}^{7/3}}. \quad (44)$$

With all the parameters now fully determined, we obtain the moving frame η .

B.2 Proof of Theorem 8

Proof. Define $J_k \triangleq \{\mathcal{J}_{i,i+k} \mid 1 \leq i \leq n-k\}$, $k = 1, 2, \dots, n-1$. To prove the theorem, it suffices to show that each J_k for $1 \leq k \leq n-1$ can be expressed in terms of elements from $S_1 \cup S_2$.

We proceed by mathematical induction on k .

Base case ($k = 1, 2$): This holds by definition, as $J_1 = S_1$ and $J_2 = S_2$.

Inductive step: Assume that for all $k \leq l-1$ with $l \geq 3$, each set J_k can be expressed in terms of elements from $S_1 \cup S_2$. We aim to show that J_l can also be written using elements from $S_1 \cup S_2$.

By the induction hypothesis, the following quantities can all be expressed by elements in $S_1 \cup S_2$:

$$\mathcal{J}_{i,i+1}, \quad \mathcal{J}_{i,i+l-1}, \quad \mathcal{J}_{i+1,i+l-1}, \quad \mathcal{J}_{i+1,i+l}, \quad \mathcal{J}_{i+l-1,i+l}.$$

For any $\mathcal{J}_{i,i+l} \in J_l$, note that when $l \geq 3$, the indices $i, i+1, i+l-1, i+l$ are distinct. By Theorem 7, it follows that $\mathcal{J}_{i,i+l}$ can be expressed as a function of these five relative invariants. Hence, $\mathcal{J}_{i,i+l}$ can be expressed by elements in $S_1 \cup S_2$.

By the principle of mathematical induction, all J_k for any $1 \leq k \leq n-1$ can be generated from $S_1 \cup S_2$, completing the proof. \square

B.3 Proof of Theorem 11

Proof. Recall that \mathcal{J}_{ij} has been established in Theorem 8, and \mathcal{T}_{iii} is already included in S_3 . So we focus on proving that $\mathcal{T}_{iji} + 2\mathcal{T}_{ijj}$ can be expressed in terms of elements from $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$. Define

$$T_k^+ \triangleq \{\mathcal{T}_{i,i+k,i} + 2\mathcal{T}_{i,i,i+k} \mid 1 \leq i \leq n-k\}, \quad (45)$$

$$T_k^- \triangleq \{\mathcal{T}_{i,i-k,i} + 2\mathcal{T}_{i,i,i-k} \mid k+1 \leq i \leq n\}. \quad (46)$$

Our goal is to show that each element in T_k^+ and T_k^- for $1 \leq k \leq n-1$ can be expressed in terms of elements from $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$. We present the proof for T_k^+ ; the case of T_k^- follows analogously due to index symmetry.

We proceed by mathematical induction on k .

Base case ($k = 1$): This holds by definition, as $T_1^+ = S_4$.

Inductive step: Assume that for all $k \leq l-1$ with $l \geq 2$, each element in T_k^+ can be expressed in terms of elements from $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$. We aim to show that every element $\mathcal{T}_{i,i+l,i} + 2\mathcal{T}_{i,i,i+l} \in T_l^+$ can also be expressed using this union. From the transformation rule in Theorem 10, we have:

$$\mathcal{T}_{i,i+l,i} + 2\mathcal{T}_{i,i,i+l} \quad (47)$$

$$= (C_1 \mathcal{T}_{i+l-1,i+l,i+l-1} - C_2 \mathcal{T}_{i+l-1,i+l,i+l} + C_3 \mathcal{T}_{i+l,i+l,i+l}) + 2(C_4 \mathcal{T}_{iii} - C_5 \mathcal{T}_{i,i,i+l-1}) \quad (48)$$

$$\begin{aligned} &= C_1 (\mathcal{T}_{i+l-1,i+l,i+l-1} + 2\mathcal{T}_{i+l-1,i+l-1,i+l}) - \frac{1}{2} C_2 (2\mathcal{T}_{i+l-1,i+l,i+l} + \mathcal{T}_{i+l,i+l-1,i+l}) \\ &\quad + C_3 \mathcal{T}_{i+l,i+l,i+l} + 2C_4 \mathcal{T}_{iii} - C_5 (2\mathcal{T}_{i,i,i+l-1} + \mathcal{T}_{i,i+l-1,i}) \\ &\quad - 2C_1 \mathcal{T}_{i+l-1,i+l-1,i+l} + \frac{1}{2} C_2 \mathcal{T}_{i+l,i+l-1,i+l} + C_5 \mathcal{T}_{i,i+l-1,i}, \end{aligned} \quad (49)$$

where the coefficients are defined as:

$$\begin{aligned} C_1 &= \frac{\mathcal{J}_{i,i+l}^2}{\mathcal{J}_{i+l-1,i+l}^2}, & C_2 &= \frac{2\mathcal{J}_{i,i+l}\mathcal{J}_{i,i+l-1}}{\mathcal{J}_{i+l-1,i+l}^2}, & C_3 &= \frac{\mathcal{J}_{i,i+l-1}^2}{\mathcal{J}_{i+l-1,i+l}^2}, \\ C_4 &= \frac{\mathcal{J}_{i+l,i+l-1}}{\mathcal{J}_{i,i+l-1}}, & C_5 &= \frac{\mathcal{J}_{i+l,i}}{\mathcal{J}_{i,i+l-1}}. \end{aligned} \quad (50)$$

By the induction hypothesis, $2\mathcal{T}_{i,i,i+l-1} + \mathcal{T}_{i,i+l-1,i} \in T_{l-1}^+$ can be expressed using elements from $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$. In addition, it holds that $\mathcal{T}_{iii} \in S_3$, $\mathcal{T}_{i+l-1,i+l-1,i+l-1,i+l} \in S_4$,

and $2\mathcal{T}_{i+l-1,i+l,i+l} + \mathcal{T}_{i+l,i+l-1,i+l} \in S_5$. The only remaining terms are:

$$- 2C_1\mathcal{T}_{i+l-1,i+l-1,i+l} + \frac{1}{2}C_2\mathcal{T}_{i+l,i+l-1,i+l} + C_5\mathcal{T}_{i+l-1,i+l-1,i+l} \quad (51)$$

$$\begin{aligned} &= - 2C_1\mathcal{T}_{i+l-1,i+l-1,i+l} + \frac{1}{2}C_2\mathcal{T}_{i+l,i+l-1,i+l} \\ &\quad + C_5(C_1\mathcal{T}_{i+l-1,i+l-1,i+l-1} - C_2\mathcal{T}_{i+l-1,i+l-1,i+l} + C_3\mathcal{T}_{i+l,i+l-1,i+l}) \end{aligned} \quad (52)$$

$$= C_5C_1\mathcal{T}_{i+l-1,i+l-1,i+l-1}, \quad (53)$$

which can be expressed in terms of $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$ as well. Hence, the full expression for $\mathcal{T}_{i,i+l,i} + 2\mathcal{T}_{i,i+l}$ is a combination of elements in $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$.

Therefore, by the principle of mathematical induction, all elements in T_k^+ for any $1 \leq k \leq n-1$ can be generated from the given sets. \square

B.4 Proof of Theorem 12

Proof. To prove Theorem 12, it suffices to show that all second-order fundamental differential invariants of the projective group, namely the quantities in (14)-(22), along with the set S_0 , can be expressed in terms of the elements in S . Since $S_0 \subset S$, we may disregard S_0 in the following discussion.

Moreover, Theorem 11 has established that all relative invariants given in Theorem 6 can be generated by elements in the union $S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5 \subset S$. Therefore, it remains to prove that the expressions (14)-(22) can be written in terms of the relative invariants from Theorem 6.

Among these fundamental invariants, (14)-(19) can be directly expressed using the relative invariants from Theorem 6. Thus, it remains to focus on the final three expressions (20)-(22). Since the denominators are already included in the set S , it suffices to consider only the numerators:

$$\mathcal{J}_{12}\mathcal{T}_{2i2} + \mathcal{J}_{2i}\mathcal{T}_{212}, \quad 3 \leq i \leq n, \quad (54)$$

$$2\mathcal{J}_{12}\mathcal{T}_{1i2} + \mathcal{J}_{12}\mathcal{T}_{21i} + 3\mathcal{J}_{2i}\mathcal{T}_{112}, \quad 3 \leq i \leq n, \quad (55)$$

$$\mathcal{J}_{12}\mathcal{T}_{1i1} + 2\mathcal{J}_{1i}\mathcal{T}_{112}, \quad 3 \leq i \leq n. \quad (56)$$

We now show that each of these numerators can indeed be expressed using the relative invariants in Theorem 6.

For (54), we have

$$\mathcal{J}_{12}\mathcal{T}_{2i2} + \mathcal{J}_{2i}\mathcal{T}_{212} \quad (57)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{2i2} + 2\mathcal{T}_{22i}) + \mathcal{J}_{2i}(\mathcal{T}_{212} + 2\mathcal{T}_{221}) - 2\mathcal{J}_{12}\mathcal{T}_{22i} - 2\mathcal{J}_{2i}\mathcal{T}_{221} \quad (58)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{2i2} + 2\mathcal{T}_{22i}) + \mathcal{J}_{2i}(\mathcal{T}_{212} + 2\mathcal{T}_{221}) - 2\mathcal{J}_{12}(C_6\mathcal{T}_{221} + C_7\mathcal{T}_{222}) - 2\mathcal{J}_{2i}\mathcal{T}_{221} \quad (59)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{2i2} + 2\mathcal{T}_{22i}) + \mathcal{J}_{2i}(\mathcal{T}_{212} + 2\mathcal{T}_{221}) - 2\mathcal{J}_{12}C_7\mathcal{T}_{222}, \quad (60)$$

where the coefficients are defined as:

$$C_6 = \frac{\mathcal{J}_{i2}}{\mathcal{J}_{12}}, \quad C_7 = \frac{\mathcal{J}_{i1}}{\mathcal{J}_{21}}. \quad (61)$$

For (55), we first rewrite \mathcal{T}_{1i2} as:

$$\mathcal{T}_{1i2} = C_8\mathcal{T}_{1i1} + C_9\mathcal{T}_{1ii} \quad (62)$$

$$= C_8(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + \frac{1}{2}C_9(2\mathcal{T}_{1ii} + \mathcal{T}_{i1i}) - 2C_8\mathcal{T}_{11i} - \frac{1}{2}C_9\mathcal{T}_{i1i} \quad (63)$$

$$\begin{aligned} &= C_8(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + \frac{1}{2}C_9(2\mathcal{T}_{1ii} + \mathcal{T}_{i1i}) \\ &\quad - 2C_8(C_6\mathcal{T}_{111} + C_7\mathcal{T}_{112}) - \frac{1}{2}C_9(C_6^2\mathcal{T}_{111} + 2C_6C_7\mathcal{T}_{112} + C_7^2\mathcal{T}_{212}), \end{aligned} \quad (64)$$

where the coefficients are defined as:

$$C_8 = \frac{\mathcal{J}_{2i}}{\mathcal{J}_{1i}}, \quad C_9 = \frac{\mathcal{J}_{21}}{\mathcal{J}_{i1}}. \quad (65)$$

Similarly, we rewrite \mathcal{T}_{21i} as:

$$\mathcal{T}_{21i} = C_6 \mathcal{T}_{211} + C_7 \mathcal{T}_{212}. \quad (66)$$

Substituting these into (55), we have:

$$\begin{aligned} & 2\mathcal{J}_{12}\mathcal{T}_{1i2} + \mathcal{J}_{12}\mathcal{T}_{21i} + 3\mathcal{J}_{2i}\mathcal{T}_{112} \\ &= 2\mathcal{J}_{12} \left(C_8(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + \frac{1}{2}C_9(2\mathcal{T}_{1ii} + \mathcal{T}_{i1i}) \right) \\ & \quad + 2\mathcal{J}_{12} \left(-2C_8(C_6\mathcal{T}_{111} + C_7\mathcal{T}_{112}) - \frac{1}{2}C_9(C_6^2\mathcal{T}_{111} + 2C_6C_7\mathcal{T}_{112} + C_7^2\mathcal{T}_{212}) \right) \\ & \quad + \mathcal{J}_{12}(C_6\mathcal{T}_{211} + C_7\mathcal{T}_{212}) + 3\mathcal{J}_{2i}\mathcal{T}_{112} \end{aligned} \quad (67)$$

$$= 2\mathcal{J}_{12} \left(C_8(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + \frac{1}{2}C_9(2\mathcal{T}_{1ii} + \mathcal{T}_{i1i}) \right) - (4\mathcal{J}_{12}C_6C_8 + \mathcal{J}_{12}C_9C_6^2) \mathcal{T}_{111}. \quad (68)$$

$$= 2\mathcal{J}_{12} \left(C_8(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + \frac{1}{2}C_9(2\mathcal{T}_{1ii} + \mathcal{T}_{i1i}) \right) - (4\mathcal{J}_{12}C_6C_8 + \mathcal{J}_{12}C_9C_6^2) \mathcal{T}_{111}. \quad (69)$$

For (56), we have:

$$\mathcal{J}_{12}\mathcal{T}_{1i1} + 2\mathcal{J}_{1i}\mathcal{T}_{112} \quad (70)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + 2\mathcal{J}_{1i}\mathcal{T}_{112} - 2\mathcal{J}_{12}\mathcal{T}_{11i} \quad (71)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) + 2\mathcal{J}_{1i}\mathcal{T}_{112} - 2\mathcal{J}_{12}(C_6\mathcal{T}_{111} + C_7\mathcal{T}_{112}) \quad (72)$$

$$= \mathcal{J}_{12}(\mathcal{T}_{1i1} + 2\mathcal{T}_{11i}) - 2\mathcal{J}_{12}C_6\mathcal{T}_{111} \quad (73)$$

This completes the proof that all second-order fundamental differential invariants in (14)-(22) are expressible in terms of the relative invariants in Theorem 6, and thus can be generated by S . \square

C Implementation details of PDINet

While our theoretical foundation is developed in the continuous setting, practical applications involve discrete image data defined on grid points, where derivatives must be approximated numerically. To this end, we estimate spatial derivatives using Gaussian derivatives [Li et al., 2018, He et al., 2022, Li et al., 2024]. For example, the partial derivative with respect to x is computed via convolution as $\frac{\partial f}{\partial x} \approx f * \frac{\partial G_\sigma}{\partial x}$, where G_σ is a Gaussian kernel with zero mean and standard deviation σ . In our implementation, we set $\sigma = 0.99$ and use a kernel size of 9.

As mentioned in Subsection 2.6, we construct projective invariants by dividing each relative invariant by a designated relative invariant \mathcal{R}_0 or \mathcal{R}_0^2 , which is computed from the input image. To prevent division by zero, we add a positive constant ϵ to the denominator. We set $\epsilon = 1$. After obtaining the invariants, we apply SupNorm normalization [Li et al., 2024, 2025], which preserves equivariance and is beneficial for training stability. Then we combine the invariants using a two-layer MLP, which is implemented as a sequence of two 1×1 convolutions with a ReLU activation function in between.

Given a standard convolutional network architecture, we construct a projective equivariant network by replacing each convolutional layer with our equivariant operator. If a convolutional layer has stride greater than 1, we insert an average pooling layer with a kernel size equal to the stride before the second 1×1 convolutional layer in the equivariant operator. Additionally, we also apply the same pooling when computing \mathcal{R}_0 to ensure resolution consistency between the numerator and denominator. Figure 3 shows the training loss curve on the STL-10 dataset, demonstrating stable optimization behavior of PDINet during training.

D Experimental details

All experiments are conducted on a single NVIDIA RTX 3090 GPU. Each experiment is repeated five times with independently generated test sets using random projective transformations, and we report the mean accuracy and standard deviation.

Experiments on Proj-STL-10. Models are trained on the 5000 samples of the STL-10 training set and tested on the 8000 samples of the Proj-STL-10 test set, which is generated by applying projective transformations to each sample in the STL-10 test set. Specifically, we decompose a projective transformation into an affine transformation followed by horizontal and vertical pure projections,

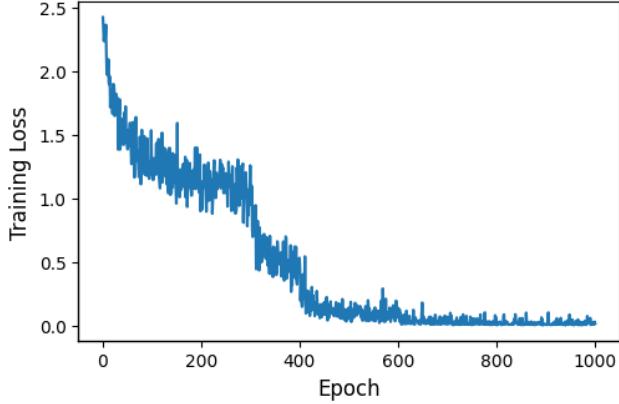


Figure 3: Training loss curve of PDINet on the STL-10 dataset.

defined as $(\tilde{x}, \tilde{y})^\top = (\frac{x}{1+c_1x}, \frac{y}{1+c_1x})^\top$ and $(\tilde{x}, \tilde{y})^\top = (\frac{x}{1+c_2y}, \frac{y}{1+c_2y})^\top$, respectively. The affine transformation consists of random rotation between -90° and 90° , scaling in the range $[0.9, 1.1]$, shear within $\pm 4^\circ$, and translation within $[-0.1, 0.1]$, while the projection parameters c_1 and c_2 are uniformly sampled from $[-0.0001, 0.0001]$. All images are normalized by channel-wise mean subtraction and standard deviation division. Following [Sosnovid et al., 2019], data augmentation during training includes 12-pixel zero-padding followed by random cropping to 96×96 , random horizontal flipping, and Cutout [DeVries and Taylor, 2017] with a single 32×32 hole. We train the models for 1000 epochs using SGD optimizer with Nesterov momentum of 0.9 and a batch size of 64. The initial learning rate is set to 0.1 and decayed by a factor of 0.2 at epochs 300, 400, 600, and 800. For the DA baseline, since the projective group is a complex non-compact group and the specific range or distribution of transformation parameters in test scenarios is typically unknown, we adopt a considerably wide range of geometric transformations. Specifically, we apply random rotation between -180° and 180° , scaling in the range $[0.3, 1.7]$, shear within $\pm 45^\circ$, translation within $[-0.1, 0.1]$, and projection parameters c_1, c_2 uniformly sampled from $[-0.002, 0.002]$. For homConv, we use the Adam optimizer following the original setup of [MacDonald et al., 2022], as training with SGD fails to converge, while keeping other hyperparameters identical to those described above.

Experiments on Proj-Imagenette. Models are trained on the 9469 samples of the Imagenette training set and evaluated on the 3925 samples of the Proj-Imagenette test set. The Proj-Imagenette dataset is generated by applying projective transformations to each test image in the original Imagenette dataset, following the same procedure as Proj-STL-10. All images are normalized by subtracting the per-channel mean and dividing by the per-channel standard deviation. During training, data augmentation includes random resized cropping to 224×224 and random horizontal flipping. We train the models for 100 epochs using AdamW optimizer with a batch size of 64. The initial learning rate is set to 0.002 and decayed via a cosine annealing scheduler. The same strategy as in experiments on Proj-STL-10 is used for the DA baseline, while homConv is also trained with Adam as in [MacDonald et al., 2022].

E Additional experiments

E.1 Equivariance error

Table 3: Equivariance error across different image resolutions.

Image size	16×16	32×32	64×64	128×128	256×256
Equivariance error (%)	0.34	0.14	0.08	0.04	0.02

Theoretically, the projective equivariance of our operators is rigorously guaranteed by the fundamental properties of differential invariants. In implementation, however, derivatives are estimated on discrete grids, which inevitably introduces minor equivariance errors. To quantitatively evaluate this effect,

we follow the protocol in [MacDonald et al., 2022] and define the equivariance error as

$$\text{Error} = \frac{\|g \cdot \psi(\mathbf{u}) - \psi(g \cdot \mathbf{u})\|^2}{\|g \cdot \psi(\mathbf{u})\|^2},$$

where ψ denotes the equivariant layer and g is a random projective transformation. Since the projective equivariance of PDINet is intrinsic and does not rely on training, we measure this error using a randomly initialized equivariant layer. We compute the error on the Imagenette test set, resizing images to multiple resolutions. As shown in Table 3, the equivariance error remains consistently small across all resolutions and decreases monotonically with increasing image size, which is expected due to more accurate derivative approximation at higher resolutions.

To further visualize equivariance, Figure 4 compares the features $g \cdot \psi(\mathbf{u})$ and $\psi(g \cdot \mathbf{u})$, showing that they are nearly identical and thus confirming the commutativity between the equivariant layer and projective transformations.

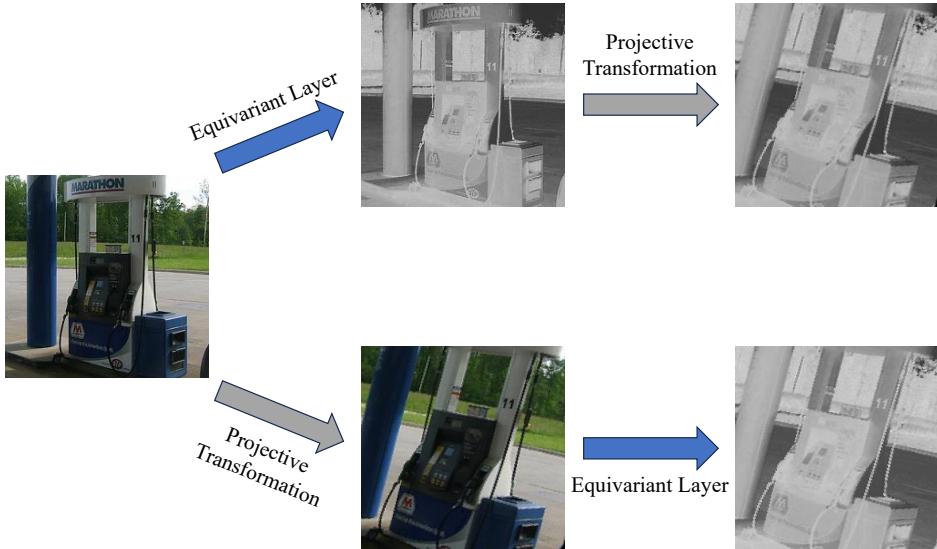


Figure 4: Visualization of projective equivariance: the features $g \cdot \psi(\mathbf{u})$ and $\psi(g \cdot \mathbf{u})$ remain consistent, illustrating the commutativity between the equivariant layer of PDINet and projective transformations.

E.2 Computational complexity

Our projective equivariant layer exhibits linear growth in both time and space complexity with respect to the input size. Specifically, the computation includes estimating derivatives, computing differential invariants, and combining differential invariants, each with linear complexity. To quantify the overhead, we compare PDINet and ResNet-18 in terms of memory usage and FLOPs using *torchstat*, and report the results in Table 4. As shown, both models exhibit linear scaling with input resolution, highlighting the scalability of our model to higher-resolution inputs. While PDINet consumes slightly more memory than ResNet-18, it requires fewer FLOPs.

To provide an explicit runtime comparison, we further evaluate both models on the same hardware by measuring the total time required to process 1000 RGB images (224×224) with a batch size of 100. Each experiment is repeated five times, and the mean and standard deviation are reported in Table 5. Despite its lower theoretical FLOPs, PDINet runs slower in practice, primarily due to the highly optimized low-level implementations of ResNet-18. We expect this gap can be narrowed through further engineering and implementation optimization.

E.3 Application on keypoint detection

To further demonstrate the practical benefit of PDINet, we conduct an additional experiment on keypoint detection, a task that inherently involves projective distortions and thus provides a natural

Table 4: Memory usage and FLOPs of ResNet-18 and PDINet at different input resolutions.

Input Size	Memory (MB)		FLOPs	
	ResNet-18	PDINet	ResNet-18	PDINet
32×32	0.53	0.93	3.72×10^7	3.65×10^7
64×64	2.10	3.73	1.49×10^8	1.46×10^8
128×128	8.38	14.91	5.94×10^8	5.84×10^8
256×256	33.50	59.63	2.38×10^9	2.33×10^9
512×512	134.00	238.50	9.51×10^9	9.34×10^9

Table 5: Runtime comparison on 1000 RGB images (224×224) with batch size 100.

Model	ResNet-18	PDINet
Runtime (s)	3.61 ± 0.35	6.84 ± 0.44

setting to evaluate the effectiveness of built-in projective equivariance. We integrate PDINet as the backbone into the REKD [Lee et al., 2022] framework for keypoint detection. To reduce confounding factors, we simplify the pipeline by removing the orientation estimation branch (which depends on steerable filters) and retain only the keypoint detection component. A three-layer PDINet is used and compared against a CNN baseline with the same architecture. The number of channels is adjusted to keep the parameter counts on the same order of magnitude, with PDINet using fewer parameters overall.

We evaluate both models on the viewpoint split of the HPatches dataset, which includes 59 scenes. Each scene contains a reference image and five target images captured from different viewpoints, resulting in projective distortions between image pairs. We follow the setup of Lee et al. [2022] for data construction and loss formulation, and train each model for 20 epochs using the AdamW optimizer with a cosine learning rate schedule (initial learning rate 0.01). A downsampling pyramid with scaling factor 1.2 is applied during training, and a symmetric pyramid with scaling factor $\sqrt{2}$ (plus an identity branch) is used at inference, with two levels of down-sampling and up-sampling.

We report the Repeatability metric, which measures the consistency of keypoint detection under viewpoint changes. A higher value indicates better robustness. As shown in Table 6, PDINet achieves higher repeatability than the baseline while using fewer parameters, suggesting that built-in projective equivariance enhances geometric consistency in keypoint detection. This supplementary experiment complements our main results and demonstrates the broader applicability of PDINet to real-world tasks.

Table 6: Keypoint detection results on the HPatches viewpoint split.

Model	Repeatability \uparrow	# Params
CNN	42.1	8.3K
PDINet	45.2	3.2K

F Discussion

While PDINet consistently outperforms the baselines in our main experiments, some misclassifications still occur. We believe these errors are largely attributable to the nature of synthetically generated data. Synthetic projective transformations involve interpolation and padding, which can introduce aliasing, distortions, and unnatural edges, especially under strong shearing, scaling or non-orthogonal rotations. Such artifacts can interfere with the model’s ability to maintain equivariance and thus degrade performance. In addition, since PDINet relies on differential invariants that depend on discrete approximations of derivatives, image resolution may also affect performance. For example, PDINet performs better on Proj-Imagenette than on Proj-STL-10, may partially attributed to the higher resolution enabling more accurate derivative estimation.

We expect that applying PDINet to real-world data at higher resolutions with naturally occurring projective distortions would help mitigate these artifacts and better demonstrate its full potential. In practical scenarios, such as multiview settings involving planar objects, transformations between different viewpoints are well modeled by projective mappings, making PDINet naturally suited for these cases. While exact equivariance may not strictly hold for non-planar or 3D objects, the projective inductive bias still contributes to improved robustness by approximately preserving geometric structure under near projective transformations. Exploring such extensions, including more general 3D settings, is a promising direction for future work.

Another avenue for exploration is to extend our method to other transformation groups. According to the general theory of differential invariants [Olver, 1993], such invariants exist for any regular Lie group action on a smooth manifold, for example, $SL(2, \mathbb{C})$ acting on the sphere, which corresponds to Möbius transformations. In principle, our framework can be adapted to these settings by deriving the appropriate differential invariants and designing corresponding equivariant architectures. While identifying explicit, concise, low-order, and numerically stable invariant forms for different groups (e.g., the Möbius group) is highly non-trivial, our experience suggests that recognizing relative invariants as atomic components can greatly facilitate the construction and simplification of a complete and practical invariant basis.