



HAL
open science

DIFFERENTIAL INVARIANTS FOR SE(2)-EQUIVARIANT NETWORKS

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesus Angulo

► **To cite this version:**

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesus Angulo. DIFFERENTIAL INVARIANTS FOR SE(2)-EQUIVARIANT NETWORKS. 29th IEEE International Conference on Image Processing (IEEE ICIP), Oct 2022, Bordeaux, France. hal-03703287v2

HAL Id: hal-03703287

<https://hal.science/hal-03703287v2>

Submitted on 18 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DIFFERENTIAL INVARIANTS FOR SE(2)-EQUIVARIANT NETWORKS

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesús Angulo

Centre for Mathematical Morphology, Mines Paris, PSL University

ABSTRACT

Symmetry is present in many tasks in computer vision, where the same class of objects can appear transformed, e.g. rotated due to different camera orientations, or scaled due to perspective. The knowledge of such symmetries in data coupled with equivariance of neural networks can improve their generalization to new samples. Differential invariants are equivariant operators computed from the partial derivatives of a function. In this paper we use differential invariants to define equivariant operators that form the layers of an equivariant neural network. Specifically, we derive invariants of the Special Euclidean Group SE(2), composed of rotations and translations, and apply them to construct a SE(2)-equivariant network, called SE(2) Differential Invariants Network (SE2DINNet). The network is subsequently tested in classification tasks which require a degree of equivariance or invariance to rotations. The results compare positively with the state-of-the-art, even though the proposed SE2DINNet has far less parameters than the compared models.

Index Terms— Equivariant Neural Networks, Differential Invariants, Image Classification

1. INTRODUCTION

Transformation equivariance can be introduced to a neural network to make use of the symmetry intrinsic to data features in many deep learning tasks. Convolutional Neural Networks (CNN) for example, reduce the number of parameters and improve the results in computer vision compared to densely connected neural networks thanks to benefiting from the translation symmetry present in computer vision tasks. Since their introduction as a general framework [1], group equivariant networks have been extended to work with other transformation groups, most notably translations combined with rotations [1, 2, 3, 4, 5, 6] or scalings [7].

Many approaches to group equivariant neural networks involve first lifting the data to a higher dimensional space obtained by sampling elements of the group and performing group convolutions there [1, 8]. This effectively introduces equivariance to the network, but it has the downside of increasing the size and computational cost according to the group dimension, as well as requiring the group transformations to be discretized. In this paper we propose in-

stead to work with the data in the original space and to apply only equivariant operators in that space, specifically. Specifically, we apply differential invariants to the data. The method of *moving frames* [9, 10] is used to derive fundamental differential invariants of the Special Euclidean group $SE(2) \cong SO(2) \times \mathbb{R}^2$, composed of translations and continuous rotations. The resulting networks are SE(2)-equivariant.

The rest of the paper is organized as follows. We highlight in Section 2 how the present work relates to existing propositions in the literature. In Section 3, we review the method of moving frames for obtaining differential invariants and study the case of SE(2). In Section 4, we define the SE(2) differential invariants networks before testing their invariance in Section 5. We end in Section 6 with some concluding remarks and future perspectives.

2. RELATED WORK

Many approaches in the literature are based on lifting the data to a higher dimensional space [1, 4, 5, 6, 3] using discretized group elements. In this paper, the equivariant operations are performed on the same space as the input image. Moreover, we achieve equivariance to continuous rotations. In the literature, equivariance to continuous rotations was achieved using circular harmonics [2] but in contrast to our method, it needs the frequencies of the feature maps to remain entangled.

In computer vision, differential invariants form a basis for many of the filtering schemes, specially in the theory of scale-spaces [11, 12]. They have also been applied to the design of feature descriptors [13] and for the automatic learning of PDEs for image restoration [14]. Moreover, convolutional neural networks have been compared with numerical schemes for solving PDEs [15] and this idea was applied to the construction of equivariant neural networks [3] albeit not with the use of differential invariants.

3. MOVING FRAMES AND DIFFERENTIAL INVARIANTS

3.1. The Method of Moving Frames

Moving frame. The method of moving frames is a way of deriving differential invariants of an m -dimensional manifold

\mathcal{M} under the action of a r -dimensional Lie group \mathcal{G} . We denote the action of $g \in \mathcal{G}$ on $z \in \mathcal{M}$ as $g \cdot z$. Throughout this paper, we can assume $\mathcal{M} = \mathbb{R}^m$. Note that a grayscale image with spatial support $\Omega \subseteq \mathbb{R}^2$, denoted by a function $u : \Omega \rightarrow \mathbb{R}$, can be identified by its graph which is a surface in $\mathcal{M} = \mathbb{R}^2 \times \mathbb{R} = \mathbb{R}^3$ with coordinate system $(x, y, u(x, y))$. Recall that an operator $\psi : \mathcal{M}_1 \rightarrow \mathcal{M}_2$ is \mathcal{G} -equivariant if there exists actions π_g, π'_g on $\mathcal{M}_1, \mathcal{M}_2$ such that $\psi\pi_g = \pi'_g\psi$. A moving frame [9] is a \mathcal{G} -equivariant function $\rho : \mathcal{M} \rightarrow \mathcal{G}$, specifically $\rho(g \cdot z) = \rho(z) \cdot g^{-1}$ for all $g \in \mathcal{G}, z \in \mathcal{M}$. This implies $\rho(g \cdot z) \cdot g \cdot z = \rho(z) \cdot g^{-1} \cdot g \cdot z = \rho(z) \cdot z$.

Operator invariantization. From a moving frame one can obtain an invariant operator from an arbitrary differentiable operator. We use ι to denote the *invariantization* of an operator $F : \mathcal{M} \rightarrow \mathbb{R}$ defined as $\iota[F](z) := F(\rho(z) \cdot z)$. The operator $\iota[F]$ is \mathcal{G} -invariant, *i.e.* $\iota[F](g \cdot z) = \iota[F](z)$.

Cross-section. A necessary and sufficient condition for a moving frame to exist is for \mathcal{G} to act *freely* and *regularly*¹ on \mathcal{M} [9]. In that case, a moving frame can be obtained by first defining a *cross-section*. A cross-section $K \subset \mathcal{M}$ is a restriction of the form $K = \{z \in \mathcal{M} | z_{k_1} = c_1, z_{k_2} = c_2, \dots, z_{k_r} = c_r\}$ with $k_1, \dots, k_r \in \{1, \dots, m\}$ and $c_1, \dots, c_r \in \mathbb{R}$ are some constants. Assuming freeness and regularity of the group action, given a cross-section K the moving frame ρ at point $z \in \mathcal{M}$ is the unique $g = \rho(z)$ such that $g \cdot z \in K$.

Jet space. In general, the action on \mathcal{M} may not be free and regular. Indeed a necessary condition is for the dimension of the orbits of \mathcal{G} to be $\dim \mathcal{G} = r$. In the case where it is not satisfied, a solution might be to consider the higher dimensional jet space $J^n(\mathcal{M})$, equipped with the prolonged group action. In the particular case where $\mathcal{M} = X \times U \cong \mathbb{R}^p \times \mathbb{R}^q$ and where we view the graph of a function $u : \mathbb{R}^p \rightarrow \mathbb{R}^q$ as a submanifold, then $J^n(\mathcal{M}) = X \times U^{(n)} \cong \mathbb{R}^p \times \mathbb{R}^{q \binom{p+n}{n}}$ [16], and a function $u^{(n)}$ is viewed as a submanifold containing the derivatives of u up to order n . A *differential invariant* is the invariantization of an operator $F : J^n(\mathcal{M}) \rightarrow \mathbb{R}, I = \iota[F]$.

Image jet space. For example, in the case where images are viewed as two-dimensional surfaces in \mathbb{R}^3 the second-order jet space $J^2(\mathcal{M})$, for instance, is an Euclidean space of dimension $2 + \binom{4}{2} = 8$ and its elements are written as $z^{(2)} = (x, y, u, u_x, u_y, u_{xx}, u_{xy}, u_{yy})$. The prolonged group action is the effect produced on the partial derivatives by the change in variables caused by the transformation.

Fundamental invariants. An important result [10] shows that the invariantizations of the partial derivatives of u up to order n can generate all other invariants of order n , and along with the invariantizations of the total derivative operators, it can generate all differential invariants. These invariants are referred to as *fundamental invariants*. In the next subsection fundamental invariants of SE(2) are shown.

¹We say that a group action is free if the only group element that fixes any point z in \mathcal{M} is the identity. It is regular if the group orbits form a regular foliation on \mathcal{M} .

3.2. Differential Invariants of SE(2) on Images

Let us view a 2D image as a surface in $\mathcal{M} = \mathbb{R}^3$ given by the graph of a function $u : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$ with open Ω . We seek equivariance to the special Euclidean group, SE(2), the group planar rotations and translations. We denote $g \in \text{SE}(2)$ as $g = (R_\theta, \mathbf{v})$ where $R_\theta \in \text{SO}(2)$ is a planar rotation and $\mathbf{v} = (v_1, v_2) \in \mathbb{R}^2$ is a translation. It acts on \mathcal{M} by

$$(R_\theta, \mathbf{v}) \cdot (x, y, u) = (\tilde{x}, \tilde{y}, u), \quad (1)$$

where $\tilde{x} = x \cos \theta - y \sin \theta + v_1, \tilde{y} = x \sin \theta + y \cos \theta + v_2$.

The action is prolonged to the jet space $J^n(\mathcal{M})$ in order to have a free and regular action. The prolonged action can be computed by the derivatives of $\tilde{u} = u$ with respect to \tilde{x} and \tilde{y} . For example, the action on the first order jet-space $J^1(\mathcal{M})$ is $(R_\theta, \mathbf{v}) \cdot (x, y, u, u_x, u_y) = (\tilde{x}, \tilde{y}, \tilde{u}, \tilde{u}_x, \tilde{u}_y)$, where \tilde{x}, \tilde{y} are the same as in (1), $\tilde{u} = u$ and

$$\tilde{u}_x = u_x \cos \theta - u_y \sin \theta, \quad \tilde{u}_y = u_x \sin \theta + u_y \cos \theta.$$

The action on $J^n(\mathcal{M}), n \geq 2$ can be derived similarly and extends the previous result, *i.e.* the first five components are $(\tilde{x}, \tilde{y}, \tilde{u}, \tilde{u}_x, \tilde{u}_y)$ as computed here above. By choosing an appropriate cross-section, we compute a moving frame ρ on $J^n(\mathcal{M})$ for $n \geq 1$ from these five components only. Let K be the cross-section in $J^n(\mathcal{M})$ defined by $u_y = x = y = 0$ and $u_x > 0$, near a point $z^{(n)}$, such that $\nabla u = (u_x, u_y) \neq \vec{0}$. Then, noting $\rho(z^{(n)}) = (R_\theta, \mathbf{v})$, the constraint $\rho(z^{(n)}) \cdot z^{(n)} \in K$ has a solution

$$0 = \tilde{u}_y = u_x \sin \theta + u_y \cos \theta \implies \tan \theta = \left(-\frac{u_y}{u_x} \right)$$

i.e. θ is the opposite of the argument of ∇u . Moreover, $v_1 = -x \cos \theta + y \sin \theta$ and $v_2 = -x \sin \theta - y \cos \theta$. Therefore, the moving frame that solves $\rho(z^{(n)}) \in K$ is

$$\rho(z^{(n)}) = \left(\frac{1}{\|\nabla u\|} \begin{bmatrix} u_x & u_y \\ -u_y & u_x \end{bmatrix}, \frac{-1}{\|\nabla u\|} \begin{bmatrix} u_x & u_y \\ -u_y & u_x \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \right). \quad (2)$$

This completely determines ρ and we can obtain differential invariants by invariantizing the derivatives of u .

We denote $I_{ij} := \iota[u_{x^i y^j}]$. Also, note $\iota[x] = \iota[y] = 0$. Then it is possible to show [10] that given any differential operator $F(z^{(n)})$, its invariantization is given by

$$\iota[F](x, y, u, \dots, u_{x^i y^j}) = F(0, 0, I_{00}, \dots, I_{0,n}).$$

Hence, if we know the forms of the invariantizations of the partial derivatives up to order n , we can obtain any other differential invariant of order n by functional combination. Some differential invariants with $n \leq 2$ are shown in Table 1.

The recurrence equations in [10] show that all differential invariants can be written as a functional combination of I_{00}, I_{02} and its invariant derivatives [10], which for SE(2) are

$$\begin{aligned} \mathcal{D}_1 &= \iota[\partial_x] := \|\nabla u\|^{-1} (u_x \partial_x + u_y \partial_y), \\ \mathcal{D}_2 &= \iota[\partial_y] := \|\nabla u\|^{-1} (-u_y \partial_x + u_x \partial_y) \end{aligned}$$

$\mathcal{I}[u]$	$I_{00} = u$
$\mathcal{I}[u_x]$	$I_{10} = \sqrt{u_x^2 + u_y^2}$
$\mathcal{I}[u_{xx}]$	$I_{20} = \frac{1}{\ \nabla u\ ^2} (u_{xx}u_x^2 + u_{xy}u_xu_y + u_{yy}u_y^2)$
$\mathcal{I}[u_{xy}]$	$I_{11} = \frac{1}{\ \nabla u\ ^2} (u_xu_y(u_{yy} - u_{xx}) - u_{xy}(u_y^2 - u_x^2))$
$\mathcal{I}[u_{yy}]$	$I_{02} = \frac{1}{\ \nabla u\ ^2} (u_{xx}u_y^2 - u_{xy}u_xu_y + u_{yy}u_x^2)$

Table 1. Invariantizations of the partial derivatives of u up to order 2, assuming $\|\nabla u\| \neq 0$. These determine all other differential invariants of order two by functional combination, for example, the Laplacian $\Delta u = u_{xx} + u_{yy} = I_{20} + I_{02}$.

or equivalently the directional derivatives in the directions of the gradient and its perpendicular. More specifically, the recurrence relation has as initial conditions I_{00} and I_{02} given in Table 1, and $I_{01} = 0$, $I_{10} = \mathcal{D}_1 I_{00}$, $I_{20} = \mathcal{D}_1 I_{10}$, $I_{11} = \mathcal{D}_2 I_{10}$ and for i, j such that $i + j = n \geq 3$:

$$I_{ij} = \begin{cases} -\mathcal{D}_1 I_{i-1,j} - P_{ij}(\mathcal{I}_{n-1}) \frac{1}{I_{10}}, & \text{if } i > 0 \\ -\mathcal{D}_2 I_{0,j-1} - P_{0,j}(\mathcal{I}_{n-1}) \frac{1}{I_{10}}, & \text{if } i = 0. \end{cases} \quad (3)$$

where P_{ij} are polynomials and \mathcal{I}_n are invariantizations of the derivatives up to order n , $\mathcal{I}_n = (I_{ij})_{0 \leq i+j \leq n} = (I_{00}, I_{10}, I_{01}, \dots, I_{0n})$.

Remark: Although a differential invariant I is invariant w.r.t. the action of \mathcal{G} on \mathcal{M} , if viewed as an operator applied to images $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ i.e. $I[u](x, y) := I(x, y, u, u_x, \dots, u_{x^0 y^n})$ it is \mathcal{G} -equivariant: $\forall (R_\theta, \mathbf{v}) \in \text{SE}(2)$, $I[(R_\theta, \mathbf{v}) \cdot u] = (R_\theta, \mathbf{v}) \cdot I[u]$, where $(R_\theta, \mathbf{v}) \cdot u = u(\tilde{x}, \tilde{y})$, and \tilde{x} and \tilde{y} are the transformed coordinates in (1).

4. DIFFERENTIAL INVARIANTS NETWORKS

In this section we define the SE(2) Differential Invariants Networks (SE2DINNet). A SE2DINNet is made of equivariant blocks (SE2DIN blocks) consisting in computing the derivatives, followed by computing the differential invariants of SE(2) up to a certain order, followed by a series of 1×1 convolutions, in the style of Network in Network (NiN) [17]. We illustrate the SE2DIN block in Figure 1 and we will explain its components in this section.

4.1. Gaussian Derivative Layers

In order to compute the differential invariants we use Gaussian derivatives. Gaussian derivatives have already been used in neural networks to produce structured receptive fields [18] in CNNs and they provide a regularization effect to the CNN. Gaussian derivatives are used to compute the derivatives of a Gaussian filtered image, i.e. $\frac{\partial^{i+j}}{\partial x^i \partial y^j} (u * G_\sigma) = u * \frac{\partial^{i+j}}{\partial x^i \partial y^j} G_\sigma = u * G_\sigma^{i,j}$, where $G_\sigma^{i,j} = \frac{\partial^{i+j} G_\sigma}{\partial x^i \partial y^j}$.

Gaussian filters are already SE(2)-equivariant, so their composition with a differential invariant yields a SE(2)-

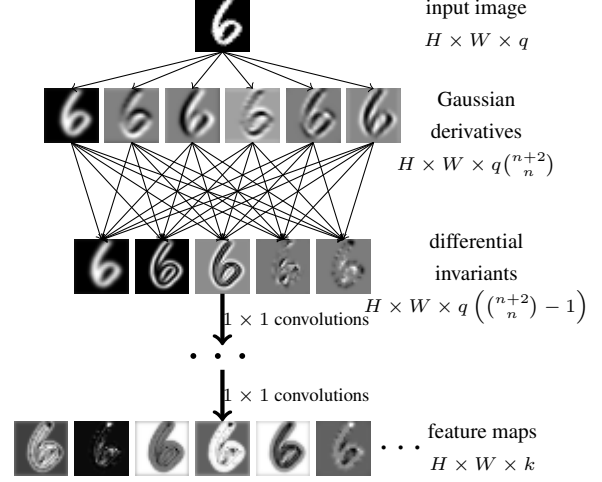


Fig. 1. Illustration of a SE2DIN block applied to a MNIST-Rot image. Convolutions are followed by batch normalization and activation function. It takes as input an image of height H , width W and q channels and outputs an image with the same spatial dimensions and k channels.

equivariant operator. These properties motivate the use of Gaussian derivatives as the first part of each SE2DIN block.

4.2. Functional Combination

From the derivatives, we compute the invariantizations of the derivatives up to order n , \mathcal{I}_n . As any differential invariant I of order n is a functional combination $I = F(\mathcal{I}_n)$ [10], a neural network with one hidden layer and taking as input \mathcal{I}_n , can approximate any differential invariant of order n , according to the Theorem of Universal Approximation. In practice we compute a l -layer densely connected network, with $l \geq 2$, applied at every spatial point $(x, y) \in \Omega$ which takes as input the vector \mathcal{I}_n . Computationally this is equivalent to a series of l 1×1 convolutions with pointwise activation functions applied to the multi-channel image with values $A(x, y) = \mathcal{I}_n(x, y)$.

As the moving frame was derived assuming $\nabla u \neq \vec{0}$, I_{ij} is not defined for some i, j when $\nabla u = \vec{0}$, therefore we introduce $\tilde{I}_{ij} = \|\nabla u\| I_{ij}$ for $i + j \geq 1$, which is used instead. The limits of $\tilde{I}_{ij}(x, y)$ for $(x, y) \rightarrow (x_0, y_0)$ such that $\nabla u(x_0, y_0) = \vec{0}$ are well defined for the new invariants.

4.3. Architecture

Given the interpretation of SE2DIN blocks as differential invariants, a SE2DINNet is viewed as a numerical scheme to solve a time-varying PDE obtained by applying the Finite Difference method to solve $u_t = F_t(\mathcal{I}_n)$ [15]. We restrict the spatial and temporal steps to $\Delta_x = \Delta_y = \Delta_t = 1$. An example SE2DINNet is illustrated in Figure 2. The network in Figure 2 before global max-pooling, ϕ is SE(2)-equivariant, i.e. $\phi((R_\theta, \mathbf{v}) \cdot u) = (R_\theta, \mathbf{v}) \cdot \phi(u)$ for $u : \mathbb{R}^2 \rightarrow \mathbb{R}^q$.




Fig. 2. Illustration of a basic SE2DINNet. Global maxpooling renders the architecture invariant. All blocks except the first and last have the same number of input channels q and filters k and follow the structure of a numerical scheme for solving a PDE.

5. EXPERIMENTS

5.1. MNIST-Rot

The MNIST-Rot dataset [19] was constructed by randomly rotating by an angle in $[\pi, -\pi)$ each image from the MNIST 12k dataset [19], which itself is obtained by selecting 12000 images from the MNIST dataset for training/validation and 50000 for testing. We divide the MNIST-Rot training set into 10000 images for training and 2000 for validation.

We test a SE2DINNet² model with six SE2DIN blocks all with two 1×1 convolutions, both with batch normalization, the first using a ReLU activation and a dropout layer. Finally, like in Figure 2, each layer except the first and last are applied by summing over the previous layer. The first two layers use $\sigma = 1$ and others $\sigma = 2$. We train two architectures, with orders $n = 2$ and $n = 3$, and with $k = 20$ and $k = 15$ filters respectively, to keep a similar number of weights and operations per forward propagation. We compute the third order invariants from the invariant derivatives of the second order ones. The change in the value of σ is done to emulate pooling. The last layer is a global max-pooling. Optimal drop rate and weight decay parameters were obtained by a grid search. The experiment was done ten times. Average results are shown in Table 2 along with published state-of-the-art results. Our results are close to the state-of-the-art while using far less parameters and *without* data augmentation.

5.2. Training on Non-Rotated Data

Here we keep the training set of the MNIST12K dataset, without random rotations, but use the test set of MNIST-Rot. The same SE2DINNet architectures are evaluated. We train a CNN baseline, with the Z2CNN [1] architecture, with and without data augmentation. The results are also in Table 2. As expected, these models performed worse than those trained on the full MNIST-Rot. However, compared to similar experiments [4] trained on 10000 samples of the original MNIST and tested on MNIST-Rot, our proposed architectures obtain a large improvement.

²The code used for the experiments is available at <https://github.com/mateussangalli/SE2DINNets>

Method	Test error(%)	#params
H-Net [2]	1.69	33k
P4CNN [1]	2.84	25k
Z2CNN [1]	5.03	22k
PDO-eConv [3]	1.87	26k
PDO-eConv [†] [3]	0.709	650k
RotEqNet [4]	1.09	100k
RotEqNet [‡] [4]	1.01	100k
E2CNN [†] [6]	0.716	-
E2CNN [†] [6]	0.682	-
SFCNN [5]	0.880	6.5M
SFCNN [†] [5]	0.714	6.5M
SE2DINNet, order 2 (Ours)	1.64	13k
SE2DINNet, order 3 (Ours)	1.62	11k
Trained on non-rotated data		
RotEqNet [4]	20	-
Z2CNN	67.9	22k
Z2CNN [†]	6.12	22k
SE2DINNet, order 2 (Ours)	2.55	13k
SE2DINNet, order 3 (Ours)	2.49	11k

Table 2. Mean accuracies on the MNIST-Rot dataset. Legend: [†] - train time augmentation, [‡] - test time augmentation.

6. CONCLUSION

We derived differential invariants for SE(2) using the moving frames method and constructed an invariant neural network architecture which use these invariants as its equivariant layers. Moreover, these networks were tested in tasks to evaluate their invariance and compared well to the state-of-art results while using far less parameters and not using data augmentation. We also observe that the third order models performed slightly better than the second order ones, which motivates investigating orders higher than three. The proposed approach can be adapted to other settings e.g. networks with other symmetry groups, for example the centro-affine [20] and projective [21] groups. The approach can also be applied other data that can be represented as manifolds, such as other surfaces embedded in \mathbb{R}^3 .

Acknowledgement. This work has been supported by FMJH under the PGM0-IRSDI 2019 program.

7. REFERENCES

- [1] Taco Cohen and Max Welling, “Group equivariant convolutional networks,” in *International Conference on Machine Learning*, 2016, pp. 2990–2999.
- [2] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow, “Harmonic networks: Deep translation and rotation equivariance,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5028–5037.
- [3] Zhengyang Shen, Lingshen He, Zhouchen Lin, and Jinwen Ma, “PDO-eConvs: Partial differential operator based equivariant convolutions,” in *Proceedings of the 37th International Conference on Machine Learning*, Hal Daumé III and Aarti Singh, Eds. 13–18 Jul 2020, vol. 119 of *Proceedings of Machine Learning Research*, pp. 8697–8706, PMLR.
- [4] Diego Marcos, Michele Volpi, Nikos Komodakis, and Devis Tuia, “Rotation equivariant vector field networks,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5048–5057.
- [5] Maurice Weiler, Fred A. Hamprecht, and Martin Storath, “Learning steerable filters for rotation equivariant CNNs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 849–858.
- [6] Maurice Weiler and Gabriele Cesa, “General E(2)-equivariant steerable CNNs,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [7] Tony Lindeberg, “Scale-covariant and scale-invariant Gaussian derivative networks,” *Journal of Mathematical Imaging and Vision*, pp. 1–20, 2021.
- [8] Risi Kondor and Shubhendu Trivedi, “On the generalization of equivariance and convolution in neural networks to the action of compact groups,” in *International Conference on Machine Learning*, 2018, pp. 2747–2755.
- [9] Mark Fels and Peter J. Olver, “Moving coframes: I. a practical algorithm,” *Acta Applicandae Mathematica*, vol. 51, no. 2, pp. 161–213, 1998.
- [10] Peter J. Olver, “Generating differential invariants,” *Journal of Mathematical Analysis and Applications*, vol. 333, pp. 450–471, 2007.
- [11] Joachim Weickert, *Anisotropic Diffusion in Image Processing*, vol. 1, Teubner Stuttgart, 1998.
- [12] H.J.A.M. Heijmans, “Morphological scale-spaces, scale-invariance and Lie groups,” in *International Symposium on Mathematical Morphology*. Citeseer, 2002, pp. 253–264.
- [13] Stanley L. Tuznik, Peter J. Olver, and Allen Tannenbaum, “Affine differential invariants for invariant feature point detection,” *arXiv preprint arXiv:1803.01669*, 2018.
- [14] Risheng Liu, Zhouchen Lin, Wei Zhang, and Zhixun Su, “Learning PDEs for image restoration via optimal control,” in *European Conference on Computer Vision*. Springer, 2010, pp. 115–128.
- [15] Lars Ruthotto and Eldad Haber, “Deep neural networks motivated by partial differential equations,” *Journal of Mathematical Imaging and Vision*, vol. 62, no. 3, pp. 352–364, 2020.
- [16] Peter J. Olver, *Equivalence, invariants and symmetry*, Cambridge University Press, 1995.
- [17] Min Lin, Qiang Chen, and Shuicheng Yan, “Network in network,” *arXiv preprint arXiv:1312.4400*, 2013.
- [18] Jorn-Henrik Jacobsen, Jan Van Gemert, Zhongyu Lou, and Arnold WM Smeulders, “Structured receptive fields in CNNs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2610–2619.
- [19] Hugo Larochelle, Dumitru Erhan, Aaron Courville, James Bergstra, and Yoshua Bengio, “An empirical evaluation of deep architectures on problems with many factors of variation,” in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 473–480.
- [20] Peter J. Olver, “Moving frames and differential invariants in centro-affine geometry,” *Lobachevskii Journal of Mathematics*, vol. 31, no. 2, pp. 77–89, 2010.
- [21] Peter J. Olver, “Projective invariants of images,” 2020, preprint on webpage at “<https://www-users.cse.umn.edu/olver/paper.html>”.