Canonical Decomposition of Steerable Functions

Yacov Hel-Or

NASA Ames Research Center

Moffett Field, CA 94035-1000

`toky@white.Stanford.EDU`

Patrick C. Teo

Department of Computer Science

Stanford University, Stanford, CA 94305

`teo@white.Stanford.EDU`

Abstract

This paper presents a general formulation enabling determination of all functions that are steerable under any transformation group. These steerable functions are determined uniquely in terms of the linear function spaces in which these functions belong. The paper suggests a canonical decomposition of these function spaces based on their generalization. The formulation is based on the theory of Lie group, and hence deals with such transformations.

1 Introduction

A function is called steerable under some transformation if all transformed versions of this function can be expressed as linear combinations of a fixed, finite set of basis functions. Steerable functions have been used widely in image processing [7, 17, 15], computer vision [9, 5], and recently, even in computer graphics [8, 14].

The importance of steerable functions stems from the property of superposition of linear systems. Hence, any linear operation applied to a transformed version of a steerable function can be expressed as a linear combination of the operation applied separately to the basis functions. The main advantage of this property is that the linear operations can be applied to the basis function once and off-line. In image processing, steerable functions have been used as filter kernels. Because convolution is a linear operation, the filter output of a transformed version of the filter kernel is obtained by linearly combining the filter outputs of its associated basis filters.

Freeman and Adelson presented functions steerable with respect to rotation using derivatives of a Gaussian as the basis set [7]. An extension of this technique to translation and

1

scaling was shown by Simoncelli et al. [17]. Approaching the problem from a numerical point of view, Perona [15] proposed a method for synthesizing these basis functions using the singular value decomposition. Although these studies deal with a large variety of transformations, they do not present a general method for constructing *all* the functions that are steerable under *any* given transformation. In this paper, we propose a general formulation that can be used to determine all the functions steerable under any group of transformations. The formulation is based on the theory of Lie transformation groups. It is constructive for any one-parameter or multi-parameter Abelian group. Functions steerable under various subgroups of the affine group have also been tabulated.

Several others have also used Lie group theory in a similar context. Amari originally proposed the use of the theory for invariant feature detection via feature normalization [2]. Our work applies and extends his idea to the design of steerable functions. Lenz also recognized the usefulness of Lie group theory and applied it to several computer vision applications including pattern detection [12]. Recently, Michaelis and Sommer [13] suggested a method for deriving steerable filters using a Lie-group theoretic approach. While their method is closely related to the one presented in this paper, they deal only with orthogonal basis functions that were constructed using a generalization of the Fourier decomposition. This restriction limits the transformation groups to which their approach can be applied.

The rest of the paper is organized as follows. Section 2 provides a brief description to Lie group theory. Following that, Section 3 and 4 outline the framework and detail several important theorems. In Section 5, canonical decomposition of steerable function spaces for one-parameter groups is given along with several examples. Section 6 describes the framework in the context of multi-parameter groups. The paper is concluded in Section 7 where possible extensions and consequences of the presented framework are elaborated.

2   Background on Lie Groups

Lie groups are often encountered as families of transformations acting on a function. In this paper, we consider, primarily, the families of transformation groups acting on real-valued, two-dimensional functions. We assume that these functions are non-zero only within a bounded region and denote them by $f(x,y) : \mathbf{R}^2 \mapsto \mathbf{R}$. We describe each family of transformations by *operators* $\{g(\boldsymbol{\tau})\}$, where $\boldsymbol{\tau} = (\tau_1, \cdots, \tau_k) \in \mathbf{R}^k$ are parameters of the transformation. For example, consider the family of one-dimensional translations of a function in the x-direction:

$$\hat{f}(\hat{x}, \hat{y}) = g_{t_x}(\tau) f(x,y) = f(x + \tau, y)$$

where τ denotes the amount of translation. In words, the operator $g_{t_x}(\tau)$ acts on the original function $f(x, y)$ to yield a new translated function $\hat{f}(\hat{x}, \hat{y}) = f(x + \tau, y)$.

A family of transformations $\{g(\boldsymbol{\tau})\}$ parameterized by $\boldsymbol{\tau}$ over some predefined range is a Lie group if: (1) it satisfies the group conditions of closure under composition, associativity, inverse and the existence of an identity, and (2) the maps for inverse and composition are smooth. Thus, the family of translations forms a Lie group: First, every translation operator $g_{t_x}(\tau)$ has an inverse, namely, $g_{t_x}(\iota(\tau))$ where $\iota(\tau) = -\tau$. Since $\iota(0) = 0$, $g_{t_x}(0)$ is also the identity operator. Second, composition of two operators can be described by a third operator which also belongs to the same family, i.e. $g_{t_x}(\tau_a)g_{t_x}(\tau_b) = g_{t_x}(\rho(\tau_a, \tau_b))$ where $\rho(\tau_a, \tau_b) = \tau_a + \tau_b$. In addition, composition is associative, that is to say, $g_{t_x}(\tau_a)(g_{t_x}(\tau_b)g_{t_x}(\tau_c)) = (g_{t_x}(\tau_a)g_{t_x}(\tau_b)) \, g_{t_x}(\tau_c)$; as such, $\rho(\tau_a, \rho(\tau_b, \tau_c)) = \rho(\rho(\tau_a, \tau_b), \tau_c)$. Finally, both the inverse map, $\iota(\tau)$, and the composition map, $\rho(\tau_a, \tau_b)$ are smooth. The dimension of the parameter space of a Lie transformation group may be different from the dimension of the function space upon which it acts. Here, the family of translations in the x-direction forms a one-parameter Lie group ($\tau \in \mathbf{R}$) while the space upon which it acts is two-dimensional ($(x, y) \in \mathbf{R}^2$).

Another familiar family of transformations that is also a Lie group is the group of rotations in the plane $g_r(\tau)$ such that $\hat{f}(\hat{x}, \hat{y}) = g_r(\tau) \, f(x, y) = f(x \cos \tau + y \sin \tau, -x \sin \tau + y \cos \tau)$. It is straightforward to check that the necessary conditions, verified in the previous example, are also satisfied here.

Lie groups are rich in structure and many properties of the group can be discerned by studying the properties of infinitesimal actions of the group. In the following, infinitesimal actions of a group are defined and elaborated. We consider first one-parameter groups and then extend our explanation to multi-parameter groups.

One-parameter Groups   Given a one-parameter transformation group parameterized by τ, the infinitesimal transformation of an function $f(x, y)$ about the identity ($\tau = 0$) is defined using Leibnitz's chain rule:

$$\frac{d}{d\tau} \left(g(\tau) \, f \right) \bigg|_{\tau=0} = \frac{d\hat{f}}{d\tau} \bigg|_{\tau=0} = \left(\frac{\partial x}{\partial \tau} \frac{\partial}{\partial x} + \frac{\partial y}{\partial \tau} \frac{\partial}{\partial y} + \frac{\partial}{\partial \tau} \right) \bigg|_{\tau=0} \hat{f}.$$

The differential operator on the right hand side of the equation is called the (infinitesimal) *generator* of the transformation and is denoted by L, i.e.

$$L = \left(\frac{\partial x}{\partial \tau} \frac{\partial}{\partial x} + \frac{\partial y}{\partial \tau} \frac{\partial}{\partial y} + \frac{\partial}{\partial \tau} \right) \bigg|_{\tau=0} \tag{1}$$

The set of elements $\mathcal{G} = \{\tau L \mid \tau \in \mathbf{R}\}$ forms a one-dimensional vector space called the *tangent space* of the group where L can be thought of as a one-dimensional basis vector.

3

There is a strong connection between the tangent space and the Lie group from which it was derived. Namely, each element $g(\tau)$ of the group can be generated by an element in the tangent space, $\tau L \in \mathcal{G}$, via the exponential map:[1]

$$g(\tau)\, f(x,y) = e^{\tau L}\, f(x,y)\; . \tag{2}$$

The notation $e^{\tau L}$ represents the series expansion $e^{\tau L} = I + \tau L + \frac{1}{2!}\tau^2 L^2 + \cdots$, which is an infinite sum of differential operators [3]. This is a rather surprising result since the operator $g(\tau)$ can transform the function in highly nonlinear ways while $\mathcal{G}$ is simply a linear vector space.

Recall the group of translations in the x-direction presented earlier. The derivative of the transformation about the identity is

$$\left.\frac{d}{d\tau}\left(g_{t_x}(\tau)f\right)\right|_{\tau=0} = \left.\frac{\partial(x+\tau)}{\partial\tau}\right|_{\tau=0} \frac{\partial}{\partial x}\, f = \frac{\partial}{\partial x}\, f$$

and hence its generator is $L_{t_x} = \frac{\partial}{\partial x}$. Using the exponential map suggested in Equation 2, we find that

$$\begin{aligned}
g_{t_x}(\tau)\, f &= e^{\tau L_{t_x}}\, f \\
&= \left(1 + \tau\frac{\partial}{\partial x} + \frac{1}{2!}\tau^2\frac{\partial^2}{\partial x^2} + \cdots\right) f \\
&= f + \tau\frac{\partial f}{\partial x} + \frac{1}{2!}\tau^2\frac{\partial^2 f}{\partial x^2} + \cdots
\end{aligned}$$

which is exactly the Taylor expansion of $f(x+\tau,y)$ about $\tau = 0$. Further examples of one-parameter groups and their generators are given in Table 2.

Multi-parameter Groups   The situation with multiple-parameter Lie groups is analogous. The generators of a multi-parameter group are the set of differential operators $\{L_i \mid i = 1\ldots k\}$ corresponding to derivatives of the transformation at the identity with respect to each parameter τ_i in turn, i.e.

$$\left.\frac{d\hat{f}}{d\tau_i}\right|_{\boldsymbol{\tau}=\mathbf{0}} = L_i\, \hat{f}$$

where

$$L_i = \left.\left(\frac{\partial x}{\partial \tau_i}\frac{\partial}{\partial x} + \frac{\partial y}{\partial \tau_i}\frac{\partial}{\partial y} + \frac{\partial}{\partial \tau_i}\right)\right|_{\boldsymbol{\tau}=\mathbf{0}} .$$

---

[1]To be precise, this is only true for group elements sufficiently close to the identity element so that their Taylor expansions converge, and for elements within the connected component containing the identity. In this paper, we consider only transformation groups with one connected component and for which convergence also holds.

Group	Operator	Generator
x-translation	$g_{t_x}(\tau)\,f(x,y) = f(x+\tau,y)$	$L_{t_x} = \frac{\partial}{\partial x}$
x-scaling	$g_{s_x}(\tau)\,f(x,y) = f(e^\tau x,y)$	$L_{s_x} = x\frac{\partial}{\partial x}$
x-projective	$g_{p_x}(\tau)\,f(x,y) = f(x/(1+\tau x),y)$	$L_{p_x} = -x^2\frac{\partial}{\partial x}$
y-translation	$g_{t_y}(\tau)\,f(x,y) = f(x,y+\tau)$	$L_{t_y} = \frac{\partial}{\partial y}$
y-scaling	$g_{s_y}(\tau)\,f(x,y) = f(x,e^\tau y)$	$L_{s_y} = y\frac{\partial}{\partial y}$
y-projective	$g_{p_y}(\tau)\,f(x,y) = f(x,y/(1+\tau y))$	$L_{p_y} = -y^2\frac{\partial}{\partial y}$
Rotation	$g_r(\tau)\,f(x,y) = f(x\cos\tau + y\sin\tau,$	
	$-x\sin\tau + y\cos\tau) = f(r,\theta+\tau)$	$L_r = y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y} = \frac{\partial}{\partial\theta}$
Uniform scaling	$g_s(\tau)\,f(x,y) = f(e^\tau x, e^\tau y) = f(e^\tau r,\theta)$	$L_s = x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} = \frac{\partial}{\partial r}$

Table 1: Several examples of continuous transformation groups and their generators. Multi-parameter groups can be constructed by combining several of these groups.

The k generators provide a basis for the k-dimensional tangent space $\mathcal{G} = \{\tau_1 L_1 + \cdots + \tau_k L_k | \boldsymbol{\tau} \in \mathbf{R}^k\}$.[2] As before, there is a correspondence between a k-parameter Lie group and its k-dimensional tangent space in the form of the exponential map:

$$g(\boldsymbol{\tau})\,f(x,y) = \Big(\prod_{i=1}^{k} e^{\tau_i L_i}\Big)\,f(x,y) = e^{\tau_1 L_1} \cdots e^{\tau_k L_k}\,f(x,y). \tag{3}$$

Although the exponential map provides a correspondence between every operator in the Lie group and every element in its tangent space, the parameterization of the group generated by the exponential map may be different from that of the original group. Furthermore, the choice of ordering the individual exponential maps in Equation 3 is arbitrary while different ordering give rise to a different group parameterization. Hence, the exponential map generates a group similar to the original group up to a change of parameterization. For example, consider the two parameterizations of the two-parameter affine group acting solely on the x coordinate:

$$
\begin{aligned}
g_1(\tau_1,\tau_2)\,f(x,y) &= f(e^{\tau_1}x + \tau_2, y),\\
g_2(\tau_1,\tau_2)\,f(x,y) &= f(e^{\tau_1}(x + \tau_2), y).
\end{aligned}
$$

Both yield the same generators:

$$L_{\tau_1} = x\frac{\partial}{\partial x}, \qquad L_{\tau_2} = \frac{\partial}{\partial x}$$

---

[2]Loosely speaking, the linear independence of the k generators is assured if the k-parameter group from which it was derived cannot be replaced by another with fewer parameters [3].

thus, having the same exponential map. However, this is not a problem as we are often interested in the group of transformations and not the particular parameterization of it. Furthermore, we can easily reparameterize the generated group using the original parameterization.

With multi-parameter groups, if we vary a single parameter τ_i and keep the others fixed, we get a one-parameter group of transformations $\{g_i(\tau_i)\}$ that is a *subgroup* of the original k-parameter group. Hence, by varying each of the k different parameters separately, we can construct k different one-parameter subgroups. When every element from one subgroup commutes with every element from a second subgroup, i.e. $g_i(\tau_i) \, g_j(\tau_j) = g_j(\tau_j) \, g_i(\tau_i)$ for all τ_i, τ_j, the two subgroups are said to commute with each other. Two subgroups commute if and only if their Lie bracket vanishes, i.e. $[L_i, L_j] \doteq L_i L_j - L_j L_i = 0$ [3]. When two subgroups commute, exponentiating their respective generators can be done in either order, i.e. $e^{\tau_i L_i} e^{\tau_j L_j} = e^{\tau_j L_j} e^{\tau_i L_i} = e^{\tau_i L_i + \tau_j L_j}$. This is not true for non-commuting subgroups. A multi-parameter group for which all its one-parameter subgroups commute is called an *Abelian* group. The two-parameter group of the previous example is not Abelian since

$$[L_{\tau_1}, L_{\tau_2}] = x \frac{\partial}{\partial x} \left(\frac{\partial}{\partial x} \right) - \frac{\partial}{\partial x} \left(x \frac{\partial}{\partial x} \right) = -\frac{\partial}{\partial x} \neq 0.$$

Hence, as demonstrated earlier, $e^{\tau_1 L_1} e^{\tau_2 L_2} \neq e^{\tau_2 L_2} e^{\tau_1 L_1} \neq e^{\tau_1 L_1 + \tau_2 L_2}$.

On the other hand, for the one-parameter transformation groups listed in Table 2, the pairs, $\{g_{t_x}, g_{t_y}\}, \{g_{t_x}, g_{s_y}\}, \{g_{t_y}, g_{s_x}\}, \{g_{s_x}, g_{s_y}\}, \{g_r, g_s\}$, are commutative.

3   Equivariant Basis Functions

In this section, we identify the functions that are steerable under different transformation groups. Before describing these functions, we formalize the notion of steerability with a definition.

Definition 1 (Steerability) : *A function $f(x, y) : \mathbf{R}^2 \mapsto \mathbf{R}$ is* **steerable** *under a k-parameter Lie transformation group G if any transformation $g(\boldsymbol{\tau}) \in G$ of f can be written as a linear combination of a fixed, finite set of* basis functions $\{\phi_i(x, y)\}$:

$$g(\boldsymbol{\tau}) \, f(x, y) = \sum_{i=1}^{n} \alpha_i(\boldsymbol{\tau}) \, \phi_i(x, y) = \boldsymbol{\alpha}^T(\boldsymbol{\tau}) \Phi(x, y)$$

The functions α_i are known as the *steering functions* of f associated with the basis $\{\phi_i\}$ and depend solely on the transform parameters. Without loss of generality, we assume that

n is the minimum number of basis functions required and these basis functions are linearly independent. Clearly, the set of basis functions required to steer a given function is not unique; any (non-singular) linear transformation of the set of basis functions could also be used.

If a function f is steerable with a set of basis functions Φ, then each of the basis functions ϕ_i are themselves steerable with the same basis functions. This is true since each basis function can be rewritten as a linear combination of transformed replicas of f (chosen to be linearly independent):

$$\Phi = \left[\begin{array}{c} \boldsymbol{\alpha}^T(\boldsymbol{\tau}^1) \\ \vdots \\ \boldsymbol{\alpha}^T(\boldsymbol{\tau}^k) \end{array} \right]^{-1} \left(\begin{array}{c} g(\boldsymbol{\tau}^1)f \\ \vdots \\ g(\boldsymbol{\tau}^k)f \end{array} \right).$$

Thus, transforming a basis function is equivalent to linearly combining the set of transformed replicas of f, which are themselves steerable.

Since steerability of the given function f implies steerability of its basis functions ϕ_i as well, it is more natural to express steerability in terms of a function space, i.e. in terms of the space spanned by the basis functions $\{\phi_i\}$.

Definition 2 (Equivariant Function Space) :
*An n-dimensional function space $\mathcal{F} = \text{span}\{\phi_1, \ldots, \phi_n\}$ is **equivariant** under a k-parameter Lie transformation group G if every ϕ_i is steerable with respect to the basis $\{\phi_1, \ldots, \phi_n\}$, i.e., there is an $n \times n$ matrix function $A(\boldsymbol{\tau})$, called the **interpolation matrix** , such that:*

$$g(\boldsymbol{\tau})\Phi(x,y) = A(\boldsymbol{\tau})\Phi(x,y) \quad \text{for all } g(\boldsymbol{\tau}) \in G$$

*This equation is called the **interpolation equation**.*

From the definition it follows that an equivariant function space is a function space that is closed under the associated transformation group. More generally, any function $f \in \mathcal{F}$, such that $f = \sum c_i \phi_i = \mathbf{c}^T \Phi$ is steerable by steering the basis of $\mathcal{F}$:

$$g(\boldsymbol{\tau})f = \mathbf{c}^T A(\boldsymbol{\tau})\Phi.$$

As a result, any function f is steerable under a k-parameter transformation group if and only if it belongs to some function space that is also equivariant under the same transformation group.

For example, consider the function space $\mathcal{F}_\theta = \text{span}\{\cos\theta, \sin\theta\}$ under the one-parameter group of rotations: $g_r(\tau)f(\theta) = f(\theta - \tau)$. It is easy to verify the following two identities:

$$\cos(\theta - \tau) \;=\; \cos\tau\cos\theta + \sin\tau\sin\theta,$$
$$\sin(\theta - \tau) \;=\; -\sin\tau\cos\theta + \cos\tau\sin\theta.$$

Thus, rotated versions of any basis function in $\mathcal{F}_\theta$ can always be expressed as linear combinations of the basis functions. Hence, any $f \in \mathcal{F}_\theta$ is steerable under the rotation group.

4   Construction of Equivariant Function Space

In the previous section, equivariant function spaces were defined to be closed under the associated transformation group. Because we are dealing with Lie transformation groups, the closure of a function space under $g(\boldsymbol{\tau})$ can be reformulated, more simply, in terms of the group generators $\{L_1, \ldots, L_k\}$. This approach is based on the seminal work of S. Amari [1, 2] who originally proposed it in the context of invariant feature detection in pattern recognition.

Theorem 1 (Interpolation Equation) :
 The function space $\mathcal{F} = span\{\phi_1, \cdots, \phi_n\}$ is equivariant under the transformation group G if and only if $\mathcal{F}$ is closed under the action of each generator L_i of G. That is, $g(\boldsymbol{\tau})\Phi = A(\boldsymbol{\tau})\Phi$ if and only if there is a set of $n \times n$ matrices $\{B_1, \ldots, B_k\}$ such that:

$$L_i\,\Phi = B_i\,\Phi \quad for\ all \quad i = 1, \ldots, k$$

In particular, the interpolation matrix can be written as follows:

$$A(\boldsymbol{\tau}) = e^{\tau_k B_k} \cdots e^{\tau_1 B_1} \;.$$

Proof 1 : Let $\hat{\Phi}(\boldsymbol{\tau}) = g(\boldsymbol{\tau})\,\Phi$, the transformed basis functions. Since $g(\boldsymbol{\tau})$ is a Lie group, it follows from the exponential map in Equation 3 that

$$
\begin{aligned}
\hat{\Phi}(\boldsymbol{\tau}) \;&=\; e^{\tau_1 L_1} \cdots e^{\tau_k L_k}\,\Phi \\
&=\; (1 + \tau_1 L_1 + \cdots) \cdots (1 + \tau_k L_k + \cdots)\,\Phi \\
&=\; (1 + \tau_1 L_1 + \cdots) \cdots (1 + \tau_k B_k + \cdots)\,\Phi \\
&=\; (1 + \tau_1 L_1 + \cdots) \cdots (1 + \tau_{k-1} L_{k-1} + \cdots)\, e^{\tau_k B_k}\,\Phi \\
&=\; (1 + \tau_1 L_1 + \cdots) \cdots e^{\tau_k B_k}\,(1 + \tau_{k-1} L_{k-1} + \cdots)\,\Phi \\
&=\; (1 + \tau_1 L_1 + \cdots) \cdots e^{\tau_k B_k}\, e^{\tau_{k-1} B_{k-1}}\,\Phi \\
&\;\;\vdots \\
&=\; e^{\tau_k B_k} \cdots e^{\tau_1 B_1}\,\Phi,
\end{aligned}
$$

1. Derive the generators $L_1, \ldots, L_k$ of the given transformation group $g(\boldsymbol{\tau})$.

2. Verify for each generator L_i that

$$L_i\ \Phi = B_i\ \Phi$$

where B_i is some $n \times n$ matrix.

3. If so, the function space $\mathcal{F} = \mathrm{span}\{\phi_1, \cdots, \phi_k\}$ is equivariant and the interpolation equation is simply

$$g(\boldsymbol{\tau})\Phi = A(\boldsymbol{\tau})\ \Phi$$

where the interpolation matrix is

$$A(\boldsymbol{\tau}) = e^{\tau_k B_k} \cdots e^{\tau_1 B_1}.$$

Figure 1: Recipe for verifying that the function space of $\mathrm{span}(\Phi)$ is equivariant. If so, the interpolation matrix $A(\boldsymbol{\tau})$ is also derived.

in which the substitution $(L_i)^m \Phi = (B_i)^m \Phi$ is used repeatedly. It can easily be verified that $L_i\ \Phi = B_i\ \Phi$ implies $(L_i)^m \Phi = (B_i)^m \Phi$ via the linearity of the differential operator L_i. The order in which the generators L_i are applied is arbitrary. However, as pointed out in Section 2, the order will determine the parameterization of the generated group. Another way of proving this directions of the theorem is by solving the differential equation, $L_i \Phi \doteq \frac{d\hat{\Phi}}{d\tau_i}\big|_{\boldsymbol{\tau}=\mathbf{0}} = B_i \Phi$, for $\hat{\Phi}$. Conversely, if $\hat{\Phi} = e^{\tau_k B_k} \cdots e^{\tau_1 B_1}\ \Phi$, taking derivatives with respect to τ_i (about $\boldsymbol{\tau} = \mathbf{0}$) on both sides of the equation yields the system of equations $L_i\ \Phi = B_i\ \Phi$. $\qquad\square$

Theorem 1 provides a recipe for verifying whether a space spanned by a set of functions $\{\phi_i\}$ is equivariant, and if it is, derives the interpolation matrix $A(\boldsymbol{\tau})$. Figure 1 summarizes the procedure. Unfortunately, the construction of all possible n-dimensional equivariant function spaces is not as methodical in general. For one-parameter groups, however, the construction is straightforward and will be treated extensively in the next section.

The following are corollaries that can be used to construct more complicated equivariant function spaces from existing ones. Their validity can easily be verified.

Corollary 1 : *If Φ is a vector of n equivariant functions, then $P\Phi$, where P is a non-singular $n \times n$ matrix, is also a vector of equivariant functions. Hence, two vectors of functions Φ_1, Φ_2 share the same equivariant function space if and only if they can be related by a non-singular $n \times n$ matrix P such that $\Phi_1 = P\Phi_2$.*

Corollary 2 : *If Φ_1 and Φ_2 are vectors of equivariant functions with respect to the same transformation group, then the space spanned by their direct sum $\Phi_1 \oplus \Phi_2$ (i.e. the concatenation of the two vectors) is also equivariant.*

Corollary 3 : *If Φ_1 and Φ_2 are vectors of equivariant functions with respect to the same transformation group, then the space spanned by the Kronecker product of the two vectors of functions $\Phi_1 \otimes \Phi_2$ (i.e. the pairwise products of functions from Φ_1 and Φ_2) is also equivariant.*

5 Function Spaces for One-Parameter Groups

In the previous section, the conditions that are required for a function space to be equivariant under a transformation group were stated. In this section, we attend to the construction of all possible equivariant spaces with respect to any one-parameter transformation group. First, we provide examples of several equivariant function spaces. After that, we show that any one-parameter group can be re-parameterized to appear as a group of translations in the new parameterization. Finally, we propose a canonical decomposition of all the function spaces equivariant under the translation group (and correspondingly under any one-parameter group that has been appropriately re-parameterized).

5.1 The Translation Group

Consider the group of one-dimensional translations in the x-direction: $\hat{f}(\hat{x}, \hat{y}) = g_{t_x}(\tau) f(x, y) = f(x + \tau, y)$ whose generator $L_{t_x} = \frac{\partial}{\partial x}$. An n-dimensional function space Φ is equivariant with respect to $g_{tx}(\tau)$ if $L_{tx}\Phi = \frac{\partial}{\partial x}\Phi = B\Phi$ for a given $n \times n$ matrix B. The general solution to this differential equation is

$$\Phi(x, y) = e^{Bx} \Phi(0) \tag{4}$$

where $\Phi(0)$ is the value of Φ at $x = 0$. Actually, the product of $\Phi(x, y)$ with any function solely in y leaves it equivariant; thus, without loss of generality, we refer to $\Phi(x, y)$ only as $\Phi(x)$. Since $\Phi(0)$ can be arbitrary chosen, any element in the column space of e^{Bx} is a possible solution. We will denote this by $\Phi(x) \in \mathcal{R}(e^{Bx})$ where $\mathcal{R}$ refers to the column space

of the matrix e^{Bx}. Regardless of the choice of $\Phi(0)$, the interpolation equation is the same, i.e. $\hat{\Phi} = e^{B\tau}\Phi$.

In the following examples, we present different choices for the matrix B and derive the corresponding function spaces. We show that the commonly used steerable functions are the result of particular choices of the matrix B.

Example 1 : Consider the simplest case where B is a 1×1 matrix, i.e. $B = [\lambda]$ where λ is a scalar value (which may be complex). From Equation 4, the space of equivariant functions is: $\Phi(x) = ae^{\lambda x}$, where a is some scalar value (the value at $\Phi(0)$), while the interpolation equation is $\hat{\Phi} = e^{\tau\lambda}\Phi$. This result is straightforward since $\hat{\Phi} = ae^{\lambda(x+\tau)} = e^{\lambda\tau}ae^{\lambda x} = e^{\lambda\tau}\Phi$.

When λ is purely imaginary, the functions are complex exponentials. In phase-based motion estimation, the parameter τ is regarded as the difference in phase. Fleet and Jepson [6] proposed an accurate method of measuring disparity by estimating the difference in phase between two (windowed) complex exponentials.

Example 2 : Now, let
$$B = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}.$$
In this case, the solution to Equation 4 implies that
$$\Phi(x) \in \mathcal{R}(e^{Bx}) = \mathcal{R}\left[\begin{pmatrix} e^{\lambda_1 x} & 0 \\ 0 & e^{\lambda_2 x} \end{pmatrix}\right]$$
and the interpolation equation is
$$\hat{\Phi} = e^{\tau B}\,\Phi = \begin{pmatrix} e^{\lambda_1 \tau} & 0 \\ 0 & e^{\lambda_2 \tau} \end{pmatrix}\Phi.$$

Simoncelli et. al. [17] proposed a criterion for shiftability in position that decomposes the filter into a set of complex exponentials (using Fourier decomposition). In this example, it would correspond to B being a diagonal matrix with unique and purely imaginary λ's.

Example 3 : Let
$$B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

11

In this case, the equivariant functions and the interpolation equation are

$$\Phi(x) \in \mathcal{R}(e^{Bx}) = \mathcal{R}\left[\begin{pmatrix} 1 & x & \frac{1}{2!}x^2 \\ 0 & 1 & x \\ 0 & 0 & 1 \end{pmatrix}\right] \qquad \text{and} \qquad \hat{\Phi} = e^{\tau B}\,\Phi = \begin{pmatrix} 1 & \tau & \frac{1}{2!}\tau^2 \\ 0 & 1 & \tau \\ 0 & 0 & 1 \end{pmatrix}\Phi$$

This example produces the moment filters which are used in many applications involving invariant feature detection [11] and motion estimation [19].

5.2   The Rotation Group

Another commonly encountered one-parameter transformation group is the group of rotations in the plane:

$$g_r(\tau)\,f(x,y) = f(x\cos\tau + y\sin\tau, -x\sin\tau + y\cos\tau)$$

where τ represents the angle of rotation. The generator of the rotation group is: $L_r = y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y}$. It is easy to see that if we represent the function $f(x,y)$ in polar coordinates (r,θ), then rotation becomes similar to translation: $g_r(\tau)f(r,\theta) = f(r,\theta+\tau)$. In these coordinates, the generator is $L_r = \frac{\partial}{\partial\theta}$. Therefore, as before, an n-dimensional vector function $\Phi(r,\theta)$ is equivariant with respect to $g_r(\tau)$ if it satisfies the equation

$$L_r\Phi \doteq \frac{\partial\Phi}{\partial\theta} = B\Phi$$

where B is an $n \times n$ matrix. The general solution to the above equation is simply

$$\Phi(\theta) = e^{B\theta}\Phi(0)$$

where $\Phi(0)$ is the value of $\Phi(\theta)$ at $\theta = 0$. Since $\Phi(0)$ is arbitrarily chosen, $\Phi(\theta) \in \mathcal{R}(e^{B\theta})$.

Example 4 :   In this example, we show that a vector of functions is equivariant with respect to rotation and derive its interpolation matrix. Let $\Phi(x,y)$ be a $2D$-vector containing the spatial derivatives of a Gaussian $G = \exp(-(x^2 + y^2)/2) = \exp(-r^2/2)$ in the x- and y-directions:

$$\Phi(x,y) = \begin{pmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{pmatrix} G = \begin{pmatrix} -x \\ -y \end{pmatrix} G = \begin{pmatrix} -r\cos(\theta) \\ -r\sin(\theta) \end{pmatrix} G.$$

Applying the generator $L_r = \frac{\partial}{\partial\theta}$ to Φ, we obtain

$$L_r\,\Phi = \begin{pmatrix} r\sin(\theta) \\ -r\cos(\theta) \end{pmatrix} G = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}\Phi = B\,\Phi.$$

12

Thus, the elements of $\Phi(x, y)$ span an equivariant function space whose interpolation function is

$$\hat{\Phi} = e^{\tau B} \ \Phi = \left(\begin{array}{cc} \cos(\tau) & -\sin(\tau) \\ \sin(\tau) & \cos(\tau) \end{array} \right) \Phi.$$

This is an example of the steerable filters suggested by Freeman and Adelson [7].

5.3 Canonical Coordinates of One-Parameter Transformation Groups

The construction of a set of equivariant function spaces depends on the existence of a solution to the system of partial differential equations $L\Phi = B\Phi$. It was shown that for translations and planar rotations, solutions exist for any given matrix B. In this section, we show that solutions exist for *any* one-parameter transformation group. The simplest way to show this is via a re-parameterization of the current coordinates into some canonical coordinates where solutions are known to exist. For any one-parameter transformation group $g(\tau)$, there exists a change of coordinates such that the group resembles a translation in the new parameterization [3]. Hence, given a function $f(x, y)$, one can determine a change of coordinates $f(\eta(x, y), \xi(x, y))$ such that

$$g(\tau) \ f(\eta, \xi) = f(\eta + \tau, \xi).$$

Segman et.al. [16] used this re-parameterization to construct invariant kernels for pattern recognition. Ferraro and Caelli [4] used this method in a similar context and suggested its relevance to biological vision.

Since the group operation is the same as one-dimensional translation, the equivariant condition with respect to the canonical coordinates is also the same:

$$L_{\eta,\xi} \ \Phi(\eta, \xi) = \frac{\partial}{\partial \eta} \Phi(\eta, \xi) = B \ \Phi(\eta, \xi).$$

Therefore, its equivariant spaces also resemble the equivariant spaces for translation (up to a change of coordinates).

Example 5 : In Section 5.2, polar coordinates were used for the group of rotations in the plane. It is easy to show that polar coordinates are the canonical coordinates for this group. Recall the change of coordinates from Cartesian to polar:

$$\eta = \arctan(y/x) = \theta \quad ; \quad \xi = \sqrt{x^2 + y^2} = r.$$

Rotating a function $f(x, y)$ in Cartesian coordinates is the same as translating the function in polar coordinates: $g_r(\tau) \ f(\eta, \xi) = f(\eta + \tau, \xi)$ where $\tau \in [0, 2\pi)$.

Example 6 : Consider next the one-parameter group of scaling in the x direction, i.e. $g_{s_x}(\tau)\, f(x, y) = f(e^{\tau} x, y)$ where e^{τ} ensures that the scaling constant is always positive. The canonical coordinates of this transformation group are obtained by the coordinate changes:

$$\eta = \ln(x) \quad \text{and} \quad \xi = y.$$

In this case,

$$g_{s_x}(\tau)\, f(\eta, \xi) = f(\ln(e^{\tau} x), \xi) = f(\ln(x) + \ln(e^{\tau}), \xi) = f(\eta + \tau, \xi)$$

which is a translation in the new coordinate system. Suppose now that

$$B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{where} \quad e^{B\eta} = \begin{pmatrix} 1 & \eta & \frac{1}{2!}\eta^2 \\ 0 & 1 & \eta \\ 0 & 0 & 1 \end{pmatrix}.$$

In this case, equivariant functions are $\Phi(\eta) \in \mathcal{R}(e^{B\eta})$ as in Example 3 of Section 5.1 but with function space in η coordinates. After a change of coordinates, the function space in x coordinates is

$$\Phi(x) \in \mathcal{R}\left[\begin{pmatrix} 1 & \ln x & \frac{1}{2!}(\ln x)^2 \\ 0 & 1 & \ln x \\ 0 & 0 & 1 \end{pmatrix} \right].$$

5.4   Canonical Decomposition of One-Parameter Equivariant Spaces

For any one-parameter transformation group, the n-vector of equivariant functions Φ depends on the apriori choice of the $n \times n$ matrix B. However, the same function space, $\operatorname{span}(\Phi) \doteq \operatorname{span}\{\phi_1, \cdots, \phi_n\}$, may be generated by different B matrices. The following theorem provides an equivalence condition among the various B matrices that generate the same equivariant function space.

Theorem 2 : *Let Φ_1, Φ_2 be two n-vectors of equivariant functions (with respect to the same one-parameter group) and B_1, B_2 are such that: $L\,\Phi_1 = B_1\Phi_1$ and $L\,\Phi_2 = B_2\Phi_2$, then*

$$\Phi_1 = P\,\Phi_2 \qquad \text{iff} \qquad B_1 = PB_2P^{-1}.$$

for any non-singular $n \times n$ matrix P.

Proof 2 : If $\Phi_1 = P\,\Phi_2$, then substituting into $L\,\Phi_1$, we get

$$L\Phi_1 = L(P\Phi_2) = PB_2\Phi_2 = (PB_2P^{-1})P\Phi_2$$

and since $P \, \Phi_2 = \Phi_1$, it must follow that $PB_2P^{-1} = B_1$. The proof in the opposite direction follows the same argument. $\qquad\square$

In words, two vectors of functions, Φ_1 and Φ_2, which are equivariant with respect to the same group, span the same function space if and only if their corresponding matrices B_1, B_2 are *similar*. Hence, it suffices to examine all matrices B that are unique up to a similarity transformation. The *Jordan decomposition* is useful to this end since any two matrices that are similar share the same Jordan form [18].

With the Jordan decomposition, any $n \times n$ matrix B can be rewritten as PJP^{-1} such that P is a non-singular $n \times n$ matrix and J is a block-diagonal matrix of the form

$$
\begin{bmatrix}
J_1 & & \\
& \ddots & \\
& & J_k
\end{bmatrix}.
$$

Each block J_i is a upper bidiagonal matrix with a single eigenvalue λ_i and one eigenvector:

$$
J_i =
\begin{bmatrix}
\lambda_i & 1 & & \\
& \ddots & \ddots & \\
& & \ddots & 1 \\
& & & \lambda_i
\end{bmatrix}.
$$

The matrix J is called the Jordan form of B and J_i are its Jordan blocks. A special case of the Jordan decomposition occurs when the matrix B is normal, i.e. $BB^H = B^H B$ where B^H is the complex conjugate of the transpose of B. In this case, the Jordan decomposition yields a diagonal matrix J; hence, each J_i is simply a 1×1 matrix containing the eigenvalue λ_i.

Let Φ_B, Φ_J be vectors of equivariant functions with respect to the translation group having corresponding matrices B, J such that J is the Jordan form of B, i.e. $B = PJP^{-1}$. From Theorem 2, then $\Phi_B = P \, \Phi_J$. In other words, the function spaces spanned by Φ_B and Φ_J are identical. Furthermore,

$$
\Phi_J(x) \in \mathcal{R}(e^{Jx}) = \mathcal{R}
\begin{bmatrix}
e^{J_1 x} & & \\
& \ddots & \\
& & e^{J_k x}
\end{bmatrix}.
$$

Since e^{Jx} is block diagonal, the function space spanned by Φ_J can be decoupled into a direct sum of function spaces spanned by each Jordan block:

$$
\Phi_J(x) \in \mathcal{R}(e^{Jx}) = \mathcal{R}(e^{J_1 x}) \oplus \mathcal{R}(e^{J_2 x}) \oplus \cdots \oplus \mathcal{R}(e^{J_s x}).
$$

Furthermore, each $\mathcal{R}(e^{J_i x})$ is a solution to $L_{t_x}\Phi = J_i\,\Phi$ and thus by itself equivariant under translation. Finally, from the identity [18],

$$
e^{J_i x} = \begin{bmatrix} e^{\lambda_i x} & x e^{\lambda_i x} & \frac{1}{2!}x^2 e^{\lambda_i x} & \cdot \\ & e^{\lambda_i x} & x e^{\lambda_i x} & \cdot \\ & & \cdot & \cdot \\ & & & e^{\lambda_i x} \end{bmatrix},
$$

it follows that any equivariant function space spanned by $\Phi_J(x)$ can be represented by a direct sum of the equivariant function basis Φ_{J_i} of the form:

$$
\Phi_{J_i} = \left(e^{\lambda_i x}, e^{\lambda_i x}x, e^{\lambda_i x}x^2, \cdots, e^{\lambda_i x}x^{n_i-1}\right)^T
$$

where n_i is the dimension of the Jordan block J_i and λ_i is its eigen-value. Note that if the matrix B is real, its eigenvalues appear in conjugate pairs, i.e. if one of the eigenvalues λ is complex, its conjugate $\bar{\lambda}$ is also an eigenvalue of B. In this case, the equivariant spaces will appear in pairs:

$$
\Phi_{J_i} \oplus \Phi_{\bar{J}_i} = \left(e^{\lambda_i x}, \cdots, e^{\lambda_i x}x^{n_i-1}\right)^T \oplus \left(e^{\bar{\lambda}_i x}, \cdots, e^{\bar{\lambda}_i x}x^{n_i-1}\right)^T.
$$

When λ is zero, the equivariant space is spanned by the first n_i moments. Alternatively, when n_i is one and λ is purely imaginary, the space is spanned by the complex exponentials, which are also the Fourier basis functions. Since any one-parameter transformation group can be put into its canonical coordinates (where the group operation becomes a translation in these new coordinates), the decomposition of equivariant function spaces for translation applies directly to all other one-parameter transformation groups (after re-parameterization) as well. Table 2 is a summary of several common one-parameter groups and their equivariant function spaces.

Example 7 :   The following functions span an equivariant function space under the group of translations $g_{t_x}(\tau)$:

$$
\Phi = \begin{pmatrix} \cos(kx) \\ \sin(kx) \end{pmatrix}
$$

since

$$
L_{t_x}\,\Phi = B\,\Phi \quad \text{where} \quad B = \begin{pmatrix} 0 & -k \\ k & 0 \end{pmatrix}.
$$

The interpolation equation associated with Φ is:

$$
\hat{\Phi} = e^{B\tau}\Phi \quad \text{where} \quad e^{B\tau} = \begin{pmatrix} \cos(k\tau) & -\sin(k\tau) \\ \sin(k\tau) & \cos(k\tau) \end{pmatrix}
$$

A different way to represent span(Φ) is by using functions generated by the Jordan form of B:

$$J = PBP^{-1} = \begin{pmatrix} ik & 0 \\ 0 & -ik \end{pmatrix} \quad \text{where} \quad P = \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}$$

By Theorem 2, it can verified that

$$\Phi_J = P\Phi_B = \begin{pmatrix} e^{ikx} \\ e^{-ikx} \end{pmatrix}.$$

where Φ_J are two of the Fourier basis functions. The new interpolation equation in this case is:

$$\hat{\Phi} = e^{J\tau}\Phi = \begin{pmatrix} e^{ik\tau} & 0 \\ 0 & e^{-ik\tau} \end{pmatrix} \Phi.$$

Example 8 : The following functions span an equivariant function space under $g_{t_x}(\tau)$:

$$\Phi = \left(\sin^3 x, \cos^3 x, 3\cos^2 x \sin x, 3\sin^2 x \cos x \right)^T$$

since

$$L_{tx}\Phi = B\Phi \quad \text{where} \quad B = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & 3 & 0 & -2 \\ -3 & 0 & 2 & 0 \end{pmatrix}$$

A different way to represent span(Φ) is with the basis functions determined by the Jordan form of B:

$$J = PBP^{-1} = \begin{pmatrix} i & & & \\ & -i & & \\ & & 3i & \\ & & & -3i \end{pmatrix} \quad \text{where} \quad P = \begin{pmatrix} 1 & -i & \frac{1}{3} & -\frac{1}{3}i \\ 1 & i & \frac{1}{3} & \frac{1}{3}i \\ 1 & -i & -1 & i \\ 1 & i & -1 & -i \end{pmatrix}.$$

and hence span(Φ_J) is determined by

$$\text{span}(\mathcal{R}(e^{Jx})) = \text{span} \left(e^{ix}, e^{-ix}, e^{3ix}, e^{-3ix} \right)^T.$$

The interpolation equation in this case is:

$$g_{t_x}(\tau)\Phi_J = e^{J\tau}\Phi_J = \begin{pmatrix} e^{i\tau} & & & \\ & e^{-i\tau} & & \\ & & e^{3i\tau} & \\ & & & e^{-3i\tau} \end{pmatrix} \Phi_J.$$

Group	Operator	Equivariant Function Space
x-translation	$g_{t_x}(\tau)\, f = f(x+\tau, y)$	$\{\psi_p(y)x^p e^{\alpha x}\}$ for $0 \le p \le m$.
x-scaling	$g_{s_x}(\tau)\, f = f(e^\tau x, y)$	$\{\psi_p(y)x^\alpha (\ln x)^p\}$ for $0 \le p \le m$.
Rotation	$g_r(\tau)\, f = f(r, \theta + \tau)$	$\{\psi_p(r)\theta^p e^{\alpha \theta}\}$ for $0 \le p \le m$.
Uniform scaling	$g_s(\tau)\, f = f(e^\tau x, e^\tau y)$	$\{\psi_p(\theta)r^\alpha (\ln r)^p\}$ for $0 \le p \le m$.

Table 2: Several examples of one parameter groups and their associated equivariant function spaces. In the rotation and uniform scaling examples, (r, θ) are the polar coordinates of the image.

6 Function Spaces for Multi-Parameter Groups

With one-parameter groups (in their canonical coordinates), various equivariant function spaces can be constructed by choosing different B matrices; solutions to the system of partial differential equations $L\Phi = B\Phi$ exist for arbitrary choices of B. Unfortunately, there is no systematic way to construct general n-dimensional equivariant spaces for multi-parameter groups. Unlike one-parameter groups, arbitrary choices of B_i for multi-parameter groups will often not yield solvable systems of differential equations. For Abelian multi-parameter groups, i.e. groups made up of one-parameter subgroups that commute, however, a categorization of the equivariant spaces similar to that for one-parameter groups can be carried out. In the following, the categorization of equivariant spaces for Abelian multi-parameter groups is presented. After that, a technique for handling non-Abelian multi-parameter groups is suggested.

Abelian Multi-Parameter Groups When the multi-parameter group is Abelian, there exists a re-parameterization of the group so that the group action is equivalent to independent translations in the new parameterization [3, 16, 4]. Formally, for any two-parameter Abelian group, there exists a re-parameterization of the function $f(\eta(x,y), \xi(x,y))$ so that

$$g(\tau_1, \tau_2)\, f(\eta, \xi) = f(\eta + \tau_1, \xi + \tau_2).$$

Segman and Zeevi in [16] describe a constructive way of determining this canonical re-parameterization. In the new parameterization, the equivariant space for the two-parameter group is simply the product of the equivariant spaces for each one-parameter translation group:

$$\text{span}(\Phi(\eta, \xi)) = \text{span}(\, \eta^p e^{\alpha \eta}\,) \otimes \text{span}(\, \xi^q e^{\beta \xi}\,)$$

for $0 \leq p \leq m$ and $0 \leq q \leq l$. Note that multi-parameter groups acting on a two-dimensional image with more than two parameters are necessarily not Abelian as there are only two independent translations in an image.

Example 9 : Consider the group of rotation and uniform scaling made up of the two one-parameter subgroups $g_r(\tau_1)$ and $g_s(\tau_2)$ from Table 2. The generators for these groups are $L_r = -x\frac{\partial}{\partial y} + y\frac{\partial}{\partial x}$ and $L_s = x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y}$ respectively. Recall that two one-parameter groups are Abelian if their generators commute; i.e., $[L_r, L_s] = L_r L_s - L_s L_r = 0$.

It is easy to verify that this equality holds in our case. The re-parameterization that makes $g_r(\tau_1)$ and $g_s(\tau_2)$ act as translations on the image is:

$$
\begin{aligned}
\eta(x, y) &= \arctan(y/x) = \theta \\
\xi(x, y) &= \ln(\sqrt{x^2 + y^2}) = \ln(r)
\end{aligned}
$$

Thus, the equivariant spaces for rotation and scaling are:

$$
\text{span}(\ln(r)^p e^{\beta \ln(r)}) \otimes span(\theta^q e^{\alpha\theta}) \quad \text{for } 0 \leq p \leq m \text{ and } 0 \leq q \leq l.
$$

Non-Abelian Multi-Parameter Groups For multi-parameter groups that are not Abelian, there are no re-parameterizations such that the group behaves like the group of independent translations in the new parameterization. One way to approach the problem is to start with the largest Abelian subgroup of the multi-parameter group for which the equivariant spaces can be constructed. The rest of the subgroups impose constraints on the equivariant space by way of the differential equations: $L_i \Phi = B_i \Phi$. Thus, the equivariant function space for the multi-parameter group can be constructed by successively constraining the equivariant space of the largest Abelian subgroup.

Example 10 : Consider the multi-parameter group made up of translations in the x and y directions together with the group of rotations, i.e. g_{t_x}, g_{t_y} and g_r respectively. The largest Abelian subgroup is the two-parameter group of translations in the x and y directions. The equivariant space for this group is: $\text{span}(\Phi) = \text{span}(x^p y^q e^{\alpha x + \beta y})$ for $0 \leq p \leq m$ and $0 \leq q \leq l$. The group of rotations yields the additional constraint that $L_r \Phi = B_r \Phi$ where $L_r = -x\frac{\partial}{\partial y} + y\frac{\partial}{\partial x}$. By observation, we can rule out the exponentials $e^{\alpha x + \beta y}$ (i.e. $\alpha = \beta = 0$) since applying L_r to each term raises the power of the monomial factor by one each time; repeated application of the conjugate generator will raise the power without bound. Applying L_r to the monomial $x^p y^q$, however, raises the power in one variable and decreases the power

Groups (# parameters)	Equivariant Function Space
x, y-translation (2)	$\{x^p y^q e^{\alpha x + \beta y}\}$ for $0 \leq p \leq m$ and $0 \leq q \leq l$.
x, y-scaling (2)	$\{x^\alpha y^\beta \ln(x)^p \ln(y)^q\}$ for $0 \leq p \leq m$ and $0 \leq q \leq l$.
Rotation Uniform-scaling (2)	$\{(\ln r)^p \theta^q e^{\alpha \theta + \beta ln(r)}\}$ for $0 \leq p \leq m$ and $0 \leq q \leq l$.
x, y-translation Rotation (3)	$\{x^p y^q\}$ for $0 \leq p + q \leq m$.
x, y-translation x, y-scaling (4)	$\{x^p y^q\}$ for $0 \leq p \leq m$ and $0 \leq q \leq l$.
x, y-translation x, y-scaling Rotation (5)	$\{x^p y^q\}$ for $0 \leq p + q \leq m$.

Table 3: Several examples of multi-parameter groups and their equivariant measuring spaces.

in the other. Successive applications will result in one of the variables being reduced to zero. Hence, $\{x^p y^q\}$ is an equivariant space under this group where $0 \leq p + q \leq m$ for some m.

Table 3 is a summary of several common multi-parameter groups and their equivariant function spaces calculated using similar considerations.

7   Conclusions

Steerable functions find application in numerous problems in image processing, computer vision and computer graphics. As such, it is important to develop the appropriate mathematical tools to analyze them. In this paper, we introduced the mathematics of Lie group theory in the context of steerable functions and presented a canonical decomposition of these functions under any transformation group.

The theory presented in this paper can be applied and extended in various ways which are out of scope of this paper. The interested reader is referred to [10].

The relevance of identifying equivariant function spaces lies in the fact that any function belonging to that space is automatically steerable using a basis for that space. In practice, the function to be steered may not reside completely in an equivariant function space. In this case the function is first approximated by an appropriate equivariant function space,

and then the approximation is steered instead. Also, often a function needs to be steered only over a restricted range of transform parameters; hence, the function needs to be only locally steerable. If the function is compactly supported, this restriction implies that the function needs to be approximated only within some small compact domain [10].

As mentioned earlier, steerable functions are often used as filter kernels for which the output of the transformed filter kernel can always be expressed as linear combinations of the outputs of a fixed set of basis filters. Motion estimation using linear filters, on the other hand, could be regarded as the reverse of this process. Given the filter outputs from two images that differ by some transformation, the transform parameters are to be estimated. It is shown in [10] that this implies that the filters used have to be steerable.

Another application of steerable functions is in the design of invariant feature detectors. For example, an edge detector should be able to detect the presence of an edge independent of the orientation of the edge in the image. According to this paradigm, the feature to be detected is first described by several steerable functions; since the functions are steerable, any transformed version of the feature can be synthesized through linear combination. Consequently, the invariant feature detector is constructed by identifying a suitable invariant over the set of steerable functions.

8 Acknowledgements

References

[1] S. Amari. Invariant structures of signal and feature spaces in pattern recognition problems. *RAAG Memoirs*, 4:19–32, 1968.

[2] S. Amari. Feature spaces which admit and detect invariant signal transformations. In *Proc. 4th Int. Joint Conf. Pattern Recognition*, pages 452–456, 1978.

[3] A. Cohen. *An introduction to the Lie theory of one-parameter groups; with applications to the solution of differential equations.* D. C. Heath & Co., Boston; New York, 1911.

[4] M. Ferraro and T. Caelli. Lie transformation groups, integral transforms, and invariant pattern recognition. *Spatial Vision*, 8(1):33–44, 1994.

[5] D. Fleet. Computation of component image velocity from local phase information. *International Journal of Computer Vision*, 5(1):77–104, 1990.

[6] D. Fleet and A. Jepson. Phase-based disparity measurement. *CVGIP: Image Understanding*, 53(2):198–210, 1991.

[7] W. Freeman and E. Adelson. The design and use of steerable filters. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.

[8] C. Gotsman. Constant-time filtering by singular value decomposition. *Computer Graphics Forum*, 13(2):153–163, 1994.

[9] G. Granlund and H. Knutsson. *Signal processing for computer vision*. Kluwer Academic Publishers, Boston, 1995.

[10] Y. Hel-Or and P. Teo. A common framework for steerability, motion estimation and invariant feature detection. Technical Report STAN-CS-TN-96-28, Stanford University, 1996.

[11] M. Hu. Visual pattern recognition by moment invariants. *IRE Trans. Information Theory*, February:179–187, 1962.

[12] R. Lenz. *Group theoretical methods in image processing*. Springer-Verlag, Berlin; New York, 1990.

[13] M. Michaelis and G. Sommer. A Lie group-approach to steerable filters. *Pattern Recognition Letters*, 16(11):1165–1174, November 1995.

[14] J Nimeroff, E Simoncelli, and J Dorsey. Efficient re-rendering of naturally illuminated environments. In *5th Eurographics Workshop on Rendering*, 1994.

[15] P. Perona. Deformable kernels for early vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17(5):488–499, 1995.

[16] J. Segman, J. Rubinstein, and Y. Zeevi. The canonical coordinate method for pattern deformation: theoretical and computational considerations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(12):1171–1183, 1992.

[17] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger. Shiftable multiscale transforms. *IEEE Trans. Information Theory*, 38(2):587–607, 1992.

[18] G. Strang. *Linear Algebra and its Applications*. Harcourt Brace Jovanovich, Inc., San Deigo, CA, 1988.

[19] Y. Xiong and S. Shafer. Moment filters for high precision computation of focus and stereo. In *Proc. Int. Conf. on Intelligent Robots and Systems*, pages 108–113, 1995.