

Journal Pre-proof

Local Rotation Invariance in 3D CNNs

Vincent Andrearczyk, Julien Fageot, Valentin Oreiller, Xavier Montet, Adrien Depeursinge

PII: S1361-8415(20)30120-1
DOI: <https://doi.org/10.1016/j.media.2020.101756>
Reference: MEDIMA 101756



To appear in: *Medical Image Analysis*

Received date: 1 September 2019
Revised date: 10 March 2020
Accepted date: 15 June 2020

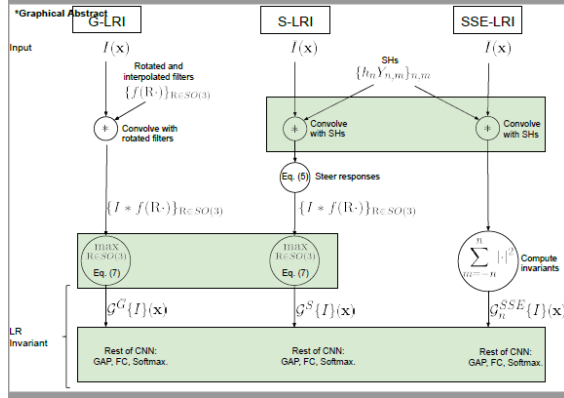
Please cite this article as: Vincent Andrearczyk, Julien Fageot, Valentin Oreiller, Xavier Montet, Adrien Depeursinge, Local Rotation Invariance in 3D CNNs, *Medical Image Analysis* (2020), doi: <https://doi.org/10.1016/j.media.2020.101756>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2020 Published by Elsevier B.V.

Highlights

- We introduce the idea of local rotation invariance in 3D convolutional networks.
- We propose three implementations including the use of spherical harmonics and 3D steerable filters.
- We demonstrate the benefit of these approaches on a synthetic 3D texture dataset and a dataset of 3D CT scans of pulmonary nodules.
- An important aspect of the reduction of trainable parameters is obtained from the weight sharing and the parametric representation of the 3D filters.



Local Rotation Invariance in 3D CNNs

Vincent Andrearczyk^{*,a}, Julien Fageot^{a,b}, Valentin Oreiller^{a,c}, Xavier Montet^d, Adrien Depeursinge^{a,c}

^a*Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), Sierre, Switzerland*

^b*Harvard School of Engineering and Applied Sciences, Cambridge, Massachusetts*

^c*Centre Hospitalier Universitaire Vaudois (CHUV), Lausanne, Switzerland*

^d*Hopitaux Universitaires de Genève (HUG), Geneva, Switzerland*

Abstract

Locally Rotation Invariant (LRI) image analysis was shown to be fundamental in many applications and in particular in medical imaging where local structures of tissues occur at arbitrary rotations. LRI constituted the cornerstone of several breakthroughs in texture analysis, including Local Binary Patterns (LBP), Maximum Response 8 (MR8) and steerable filterbanks. Whereas globally rotation invariant Convolutional Neural Networks (CNN) were recently proposed, LRI was very little investigated in the context of deep learning. LRI designs allow learning filters accounting for all orientations, which enables a drastic reduction of trainable parameters and training data when compared to standard 3D CNNs. In this paper, we propose and compare several methods to obtain LRI CNNs with directional sensitivity. Two methods use orientation channels (responses to rotated kernels), either by explicitly rotating the kernels or using steerable filters. These orientation channels constitute a locally rotation equivariant representation of the data. Local pooling across orientations yields LRI image analysis. Steerable filters are used to achieve a fine and efficient sampling of 3D rotations as well as a reduction of trainable parameters and operations, thanks to a parametric representations involving solid Spherical Harmonics (SH), which are products of SH with associated learned radial pro-

*Corresponding author

Email address: vincent.andrearczyk@hevs.ch (Vincent Andrearczyk*)

files. Finally, we investigate a third strategy to obtain LRI based on rotational invariants calculated from responses to a learned set of solid SHs. The proposed methods are evaluated and compared to standard CNNs on 3D datasets including synthetic textured volumes composed of rotated patterns, and pulmonary nodule classification in CT. The results show the importance of LRI image analysis while resulting in a drastic reduction of trainable parameters, outperforming standard 3D CNNs trained with rotational data augmentation.

Keywords: Local rotation invariance, convolutional neural network, steerable filters, 3D texture

1. Introduction

Convolutional Neural Networks (CNNs) have been successfully used in various studies to analyze textures. By construction, CNN operations are translation equivariant, thus particularly adapted to image analysis where objects of interest have arbitrary locations. In this paper, we propose to incorporate Local Rotation Invariance (LRI) into the CNN architecture, which is known to be crucial for texture analysis and biomedical applications in general because objects and patterns of interest have most often arbitrary orientations (Depeursinge and Fageot (2018)).

Globally Rotation Invariant (RI) CNNs have recently been studied, making use of group theory to maintain rotation equivariance throughout the layers. The 2D Group equivariant CNNs (G-CNN¹), developed in Cohen and Welling (2016), uses rotated (right-angles only) versions of the filters together with appropriate channels permutations. RI is then obtained by pooling across orientation channels after the last convolutional layer. 3D G-CNNs were shown to improve detection of pulmonary nodule detection in Winkels and Cohen (2019) and classification of 3D textures in Andrearczyk and Depeursinge (2018), yet

¹When referring to G-CNNs, we consider discrete designs with right angle rotations, while it is defined as a more general framework including continuous designs in Bekkers (2019).

the latter study motivated the use of a finer rotation sampling than right-angle rotations to capture realistic arbitrary 3D orientations of directional patterns. G-CNNs achieve equivariance with respect to *finite* subgroups of the rotation group, which constitutes a bottleneck in 3D. In 2D, an arbitrary sampling of rotations can be used in a group equivariant approach (Bekkers et al. (2018)), while the number of 3D finite rotation groups is restrained. Both 2D harmonic networks (Worrall et al. (2017)) and 2D steerable CNNs (Weiler et al. (2017)) present similarities with the method proposed in this paper although in the 2D domain. Some recent work consider neural networks on non-Euclidian domains (Kondor and Trivedi (2018)), in particular in the 2-dimensional sphere, where the invariance to rotations plays a crucial role as in Kondor et al. (2018) and Cohen et al. (2018). Finally, 3D steerable CNNs such as proposed in Weiler et al. (2018) are very general architectures implementing global equivariance to rotations on the network, and the convolutional layer considered in this paper is covered by their design although not specifically investigated. In particular, the proposed LRI layers are specialized instances of a discrete (Winkels and Cohen (2019)) and steerable G-CNN (Weiler et al. (2018)). We differ from their work by making an angular max-pooling after the first convolution layer, which exploits the steerability of the filters, and more importantly, focuses on the sought-after local invariances. While G-CNNs can encode complex objects, we focus on textures with local patterns.

In the above-mentioned approaches, global rotation equivariance is maintained all along with the layers (see Fig. 1, left), and invariance is obtained by using orientation pooling at the end of the network after spatial average pooling. Global RI is fundamental in various applications, e.g. to analyze pictures taken with arbitrary orientations of the camera. However, most images are composed of well-defined substructures having arbitrary orientations. For instance, patterns of interest in medical imaging modalities such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) consist of tissue alterations with characteristic 3D textures signatures including necrosis, angiogenesis, fibrosis, or cell proliferation (Gatenby et al. (2013)). These alterations induce

imaging signatures such as blobs, intersecting surfaces and curves. These local
 50 low-level patterns are characterized by discriminative directional properties and
 have arbitrary 3D orientations, which requires combining directional sensitivity
 with LRI. When compared to equivariant designs, LRI allows to discard the
 information on local pattern orientation, resulting in more lightweight CNNs.

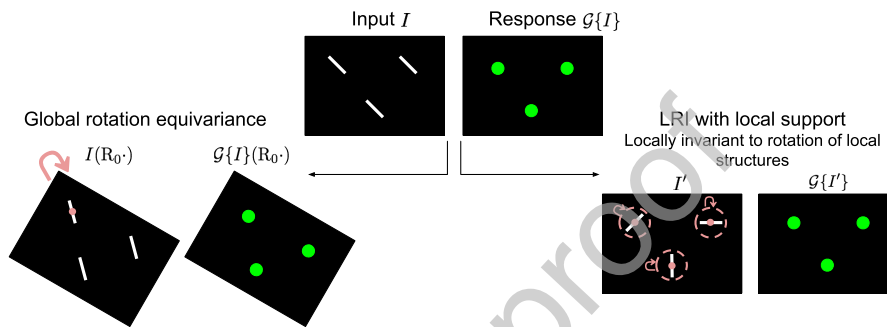


Figure 1: Illustration of global RI and LRI in 2D. Rotating local structures (i.e. three white segments) in the input I results in the input I' on the right. The green dots illustrate the equivariant/invariant responses. Local and global rotations are shown in red and the local support G of the operator \mathcal{G} (see Section 2.2) is represented as a dashed red line. It is worth noting that our CNN architecture will both present a global equivariance and a local invariance to rotations. Best viewed in color.

However, RI is often antagonistic with the aim of being sensitive to direc-
 55 tional features. For instance, a spatial image operator that is purely convolu-
 tional is equivariant to rotations if and only if the filter is isotropic (see Sec-
 tion 2.2 and Bekkers (2019); Cohen et al. (2019)), therefore insensitive to the
 directional features of the input signal. It follows that operators combining LRI
 and directional sensitivity (i.e. non-isotropic) require using more complex de-
 60 signs such as MR8 (Varma and Zisserman (2005)), local binary patterns (Ojala
 et al. (2002)), steerable Riesz wavelets (Dicente Cid et al. (2017)), circular or
 Spherical Harmonic (SH) invariants (Depeursinge et al. (2018)), sparse coding
 with steerable atoms McCann et al. (2020) and scattering transform (Eickenberg
 et al. (2017)). These designs were widely used in hand-crafted texture analysis
 65 (Depeursinge and Fageot (2018); Liu et al. (2019)).

In this paper, we propose three 3D CNN architectures that are both globally equivariant and locally invariant to rotations (see Fig. 1 for an illustration in 2D), and can combine this with directionally sensitive image analysis. This can be achieved by convolving with rotated filters (i.e. G-convolution (Winkels and Cohen (2019), referred to as G-LRI), steered responses to SHs (Andrzejczyk et al. (2019a), referred to as S-LRI), or Solid Spherical Energy (SSE) invariants calculated from SH responses (Andrzejczyk et al. (2019b), referred to as SSE-LRI). Experiments in Section 3 show the benefit of LRI designs over standard CNNs (Tables 2-5) and globally rotation invariant designs (Tables 4, 5) on synthetic textures and lung nodule datasets where local patterns occur at random orientations.

2. Methods

This section is organized as follows. After clarifying mathematical notations in Section 2.1, we first define a general 3D LRI operator in Section 2.2. Section 2.3 introduces the mathematical tools used in this paper: steerable filters and spherical harmonics. The different methods that we use to implement the operator, namely G-LRI (based on the G-CNN), S-LRI (Steerable LRI), SSE-LRI (Solid Spherical Energy LRI) are detailed in Section 2.4. We then introduce global RI in Section 2.5, which is further compared against LRI approaches in Section 3.3. Finally, the discretization, datasets, network architectures and weights initialization are presented in Sections 2.6, 2.7, 2.8 and 2.9 respectively.

2.1. Notations

We initially introduce the frameworks in the continuous domain, hence 3D images, filters, and response maps are functions defined over the continuum \mathbb{R}^3 . We shall also discuss the practical discretization of the different methods (Section 2.6). Spherical coordinates are defined as (ρ, θ, ϕ) with radius $\rho \geq 0$, elevation angle $\theta \in [0, \pi]$, and horizontal plane angle $\phi \in [0, 2\pi)$. The set of 3D rotations is denoted by $SO(3)$. A 3D rotation transformation matrix R

can be decomposed as three elementary rotations around z , y' and z'' axes as
 95 $\mathbf{R} = \mathbf{R}_\alpha \mathbf{R}_\beta \mathbf{R}_\gamma$, with the orientation (α, β, γ) parameterized by the (intrinsic) Euler angles $\alpha \in [0, 2\pi)$, $\beta \in [0, \pi]$, and $\gamma \in [0, 2\pi)$ respectively. We will use interchangeably \mathbf{R} as a rotation transformation acting on \mathbb{R}^3 and on the two-dimensional sphere \mathbb{S}^2 . Finally, the function $\mathbf{x} \mapsto f(\mathbf{R}\mathbf{x})$ is denoted by $f(\mathbf{R}\cdot)$.

2.2. Equivariant Image Operators and Invariant Image Features

100 We introduce the general class of image operators of interest that will be used in the first layer of our neural network and common between G-LRI, S-LRI and SSE-LRI. An image operator \mathcal{G} associates to an image I another image, denoted by $\mathcal{G}\{I\}$. The following invariance properties will be relevant for our analysis:

- An operator \mathcal{G} is globally *equivariant to translations and rotations*, if, for any position $\mathbf{x}_0 \in \mathbb{R}^3$ and rotation $\mathbf{R}_0 \in SO(3)$,

$$\mathcal{G}\{I(\cdot - \mathbf{x}_0)\} = \mathcal{G}\{I\}(\cdot - \mathbf{x}_0) \quad \text{for any } \mathbf{x}_0 \in \mathbb{R}^3, \quad (1)$$

$$\mathcal{G}\{I(\mathbf{R}_0\cdot)\} = \mathcal{G}\{I\}(\mathbf{R}_0\cdot) \quad \text{for any } \mathbf{R}_0 \in SO(3). \quad (2)$$

105 In particular, if $\mathbf{R}_{\mathbf{x}_0}$ is a rotation around $\mathbf{x}_0 \in \mathbb{R}^3$, we have that $\mathcal{G}\{I(\mathbf{R}_{\mathbf{x}_0}\cdot)\} = \mathcal{G}\{I\}(\mathbf{R}_{\mathbf{x}_0}\cdot)$, as illustrated on the left part of Fig. 1.

- An operator \mathcal{G} is *local* if there exists $\rho_0 > 0$ such that, for every \mathbf{x} , the quantity $\mathcal{G}\{I\}(\mathbf{x})$ only depends on local image values $I(\mathbf{y})$ for $\|\mathbf{y} - \mathbf{x}\| \leq \rho_0$.

The global equivariance to translations and rotations together with the locality
 110 result in the sought-after invariance to local rotations (i.e. LRI) in the following sense: the rotation of an object or localized structure of interest in the image I around a position \mathbf{x} does not affect the value of $\mathcal{G}\{I\}(\mathbf{x})$, as illustrated on the right part of Fig. 1. We illustrate the different notions for the case of linear convolution operators in the next result.

Proposition 1. *Let \mathcal{G} be a linear convolution operator of the form*

$$\mathcal{G}\{I\} = h * I \quad (3)$$

115 with h the impulse response of the filter. Then, \mathcal{G} is globally equivariant to rotations if and only if h is isotropic, i.e., $h(\mathbf{R}_0 \cdot) = h$ for any rotation $\mathbf{R}_0 \in SO(3)$. Moreover, \mathcal{G} is local if and only if h is compactly supported. Therefore, \mathcal{G} is LRI if only if h is compactly supported and isotropic.

The proof of Proposition 1 is given in Appendix A. The result is elementary, 120 and can be deduced using general frameworks, such as (Bekkers, 2019, Theorem 1). It reveals that linear operators can only fulfill the required equivariences using isotropic filters, which are insensitive to the directional information and thus very limited (Depeursinge et al. (2018)). The operators used in this paper are therefore non-linear.

125 2.3. Steerable Filters and Spherical Harmonics

This subsection introduces the mathematical toolbox required to characterize the proposed S-LRI and SSE-LRI approaches, which both rely on parametric kernel representations based on solid SHs. In particular, we consider filters f expanded in terms of the family of SHs $(Y_{n,m})_{n \geq 0, m \in \{-n \dots n\}}$, where n is called 130 the degree and m the order, and which forms an orthonormal basis for square-integrable functions $g(\theta, \phi)$ on the sphere \mathbb{S}^2 . We consider finitely many degrees, $N \geq 0$ being the maximal one. The number of elements of a SH family of maximum degree N is $\sum_{n=0}^N (2n+1) = (N+1)^2$. The expression of SHs can be found in Appendix B. We say that a function $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ is a solid SH² if it is 135 a product of a SH with a purely radial function; that is, if it can be written as $f(\rho, \theta, \phi) = h(\rho)Y_n^m(\theta, \phi)$.

The S-LRI uses steerable filters (see Section 2.4.2), which have the advantage to allow for fast and efficient computation of the LRI representation required for max-pooling over orientations channels to further achieve invariance (Chenouard 140 and Unser (2012); Fageot et al. (2018)). A filter is steerable if any of its rotated versions can be written as a linear combination of finitely many basis

²Note that the terminology is sometimes used when dealing with radial profiles following power law (Eickenberg et al. (2017)).

filters (Freeman and Adelson (1991); Unser and Chenouard (2013)).

In this paper, we consider 3D steerable filters $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ of the form

$$f(\rho, \theta, \phi) = \sum_{n=0}^N h_n(\rho) \sum_{m=-n}^n C_n[m] Y_{n,m}(\theta, \phi), \quad (4)$$

where the $h_n(\rho) \in \mathbb{R}$ are degree-dependent radial profiles and the coefficients
 145 $C_n[m] \in \mathbb{C}$ determine the angular structure of f .

Many work deal with steerable filters that are polar-separable. This means that f can be decomposed as $f(\rho, \theta, \phi) = h(\rho)g(\theta, \phi)$. A steerable filter of the form (4) is polar-separable if and only if it can be written as

$$f(\rho, \theta, \phi) = h(\rho) \sum_{n=0}^N \sum_{m=-n}^n C_n[m] Y_{n,m}(\theta, \phi), \quad (5)$$

with h a single radial profile that captures the radial pattern of the filter. The polar separable case (5) is a particular case of (4), it corresponds to the situation when h_n does not depend on n . In the sequel, we keep track on the index n , which covers both cases.

150 The condition of f being real is translated into the conditions that h or h_n themselves are real and that the SH coefficients satisfy $C_n[-m] = (-1)^m \overline{C_n[m]}$ (see Appendix C).

For any rotation $\mathbf{R} \in SO(3)$, the rotated version $Y_{n,m}(\mathbf{R}\cdot)$ of a SH can be expressed as a linear combination of all elements in a degree subspace n as

$$Y_{n,m}(\mathbf{R}\cdot) = \sum_{m'=-n}^n D_{\mathbf{R},n}[m, m'] Y_{n,m'}, \quad (6)$$

where the $D_{\mathbf{R},n} \in \mathbb{C}^{(2n+1) \times (2n+1)}$ are the Wigner matrices (Varshalovich et al. (1988)). Then, the steerable filter f can be rotated efficiently with any $\mathbf{R} \in SO(3)$ to obtain a set of steered coefficients $C_{\mathbf{R},n} = D_{\mathbf{R},n} C_n$ of $f(\mathbf{R}\cdot)$, with $C_n = (C_n[m])_{m \in \{-n, \dots, n\}}$. The rotated filter $f(\mathbf{R}\cdot)$ is given by

$$f(\mathbf{R}\cdot)(\rho, \theta, \phi) = \sum_{n=0}^N h_n(\rho) \sum_{m=-n}^n \sum_{m'=-n}^n D_{\mathbf{R},n}[m, m'] C_n[m'] Y_{n,m}(\theta, \phi). \quad (7)$$

From (7), we see that any rotated version of f can be computed from the coefficients $(C_n[m])_{0 \leq n \leq N, -n \leq m \leq n}$.

155 In Andrearczyk et al. (2019a), we only considered polar separable filters, in the sense that f can be written as $f(\rho, \theta, \phi) = h(\rho)g(\theta, \phi)$ with $h : \mathbb{R}^+ \rightarrow \mathbb{R}$ and $g : \mathbb{S}^2 \rightarrow \mathbb{R}$, as is the case in (5).

Using a shared radial profile for all SHs results in a reduction of trainable parameters, at the cost of limited SH parametric approximation capability (re-
160 stricted to polar separable patterns). The extension to non-polar separable filters of the form (4) is an important contribution of this paper.

2.4. Locally Rotation Invariant 3D CNNs

This section details the three proposed strategies to achieve 3D LRI image analysis. An overview of the three methods is depicted in Fig. 2, and a
165 qualitative comparison is presented in Table 1.

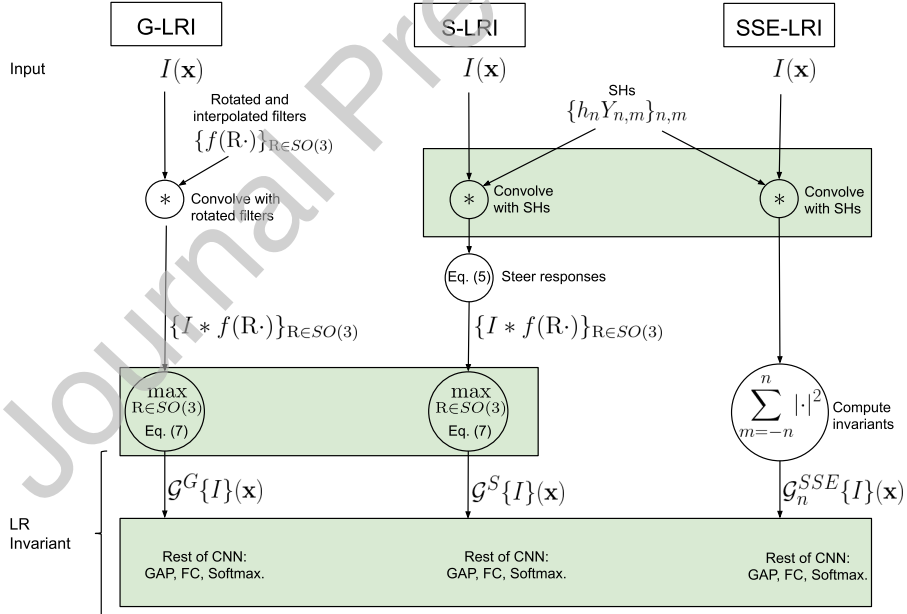


Figure 2: Overview of the methods used to obtain LRI including Group-equivariant LRI (G-LRI), Steerable LRI (S-LRI) and Solid Spherical Energy LRI (SSE-LRI). Operations shared by multiple methods are highlighted in green.

Model	LRI	Rotation weight sharing	Parametric representation	No need for orientation channels
Z3-CNN	✗	✗	✗	✓
G-LRI	✓	✓	✗	✗
S-LRI	✓	✓	✓	✗
SSE-LRI	✓	✓	✓	✓

Table 1: Qualitative comparison of the considered 3D CNN frameworks.

2.4.1. G-LRI

The first method to obtain LRI is to use rotated versions of the kernels, i.e. via weight sharing across orientation channels. LRI is obtained by max-pooling over the rotations and the corresponding image operator is

$$\mathcal{G}^G\{I\}(\mathbf{x}) = \max_{\mathbf{R} \in SO(3)} |(I * f(\mathbf{R}\cdot))(\mathbf{x})|, \quad (8)$$

where f is characterized by trainable parameters as in a classical CNN, i.e. full 3D kernels. The proof of equivariance to translation and rotation as (1) and (2) is provided in Appendix D. Moreover, the operator \mathcal{G}^G is local if and only if
 170 the filter f has a finite support, what we assume from now.

The idea of max pooling over oriented filter responses has been long used in computer vision, e.g. for template matching with cross-correlation (Brown (1992)). More recently, the idea of rotating the CNN kernels has been widely used in the literature in the context of equivariance to groups of rotations (Cohen
 175 and Welling (2016); Winkels and Cohen (2019); Worrall and Brostow (2018)). In particular, the 3D G-CNN developed in Winkels and Cohen (2019) offers equivariance to groups of 3D rotations. In reference to this work, we refer to this first approach as G-LRI even though we do not propagate the equivariance to deeper layers and neither require to perform operations on finite groups.

180 2.4.2. Steerable LRI

S-LRI is a special case of G-LRI for which the computation exploits steerability. Such S-LRI layers were proposed in Andrearczyk et al. (2019a) with polar

separable filters only. Here we define S-LRI for both polar separable (S-LRI- h) and non-polar separable (S-LRI- h_n) pattern approximation methods, based on steerable filters of the forms (5) and (4), respectively. As mentioned before, we keep track on the index n which covers both cases.

The S-LRI operator $\mathcal{G}^S\{I\}(\mathbf{x})$ is obtained by max-pooling over the rotations as in (8):

$$\mathcal{G}^S\{I\}(\mathbf{x}) = \max_{R \in SO(3)} |(I * f(R \cdot))(\mathbf{x})|, \quad (9)$$

where the filter f is in this case of the form (4) and is assumed to have a finite support. The operator \mathcal{G}^S is defined identically to \mathcal{G}^G in (8) but we use different notations to keep in mind that the parametrization of the filters f differ. As we have seen, the image operators (8) and (9) are equivariant to rotations and translations. It is moreover local as soon as the h_n have a finite support and, therefore, LRI.

Exploiting (7), the convolution $I * (f(R \cdot))$ is then computed as

$$I * f(R \cdot) = \sum_{n=0}^N \sum_{m=-n}^n \left(\sum_{m'=-n}^n D_{R,n}[m, m'] C_n[m'] \right) (I * h_n Y_{n,m}). \quad (10)$$

Therefore, one accesses the convolution with any (virtually) rotated version of f by computing $\sum_{n=0}^N (2n+1) = (N+1)^2$ convolutions $(I * h_n Y_{n,m})$, which we shall exploit for computing the response map of the image operator. It is worth noting that the case $N=0$ corresponds to filters f that are isotropic, i.e. $f(R \cdot) = f$ for any $R \in SO(3)$ (Depeursinge et al. (2018)). As low degrees (e.g. $N=1, 2$) are sufficient to construct small filters (see Section 2.6.3), the gain becomes substantial over a G-CNN approach for a fine sampling of orientations with a drastic reduction of the number of convolutions.

In practice, one has a set of steerable filters f_i of the form (4) with radial profiles $h_{i,n}$ and coefficients $C_{i,n}[m]$. When compared to the G-LRI, the number of trainable parameters is reduced to $C_{i,n}[m]$, $h_{i,n}$, and the biases added after orientation pooling (one scalar parameter per output channel i).

205 *2.4.3. Solid Spherical Energy LRI*

The use of solid SH representations, i.e. SH representations with radial profiles, provides the opportunity to compute rotational invariants from simple non-linear operations. Initially proposed in Andrearczyk et al. (2019b) for polar separable (SSE-LRI- h) kernels, we extend the invariants to the non-polar
 210 separable (SSE-LRI- h_n) case. Here we re-use most of the concepts used for the S-LRI. However, instead of steering, we calculate invariants directly from the responses of the solid SHs, which obviates the need to construct an intermediate (discretized) locally rotation equivariant representation.

After convolution with the image I , the responses $I * h_n Y_{n,m}$ with $m = -n, \dots, n$ contain the spectral information of degree n , which is used to define the image operator \mathcal{G}_n^{SSE} as

$$\mathcal{G}_n^{SSE}\{I\}(\mathbf{x}) = \sum_{m=-n}^n |(I * h_n Y_{n,m})(\mathbf{x})|^2. \quad (11)$$

Let us study the desirable properties of \mathcal{G}_n^{SSE} in the following. At a fixed spatial
 215 position $\mathbf{x} \in \mathbb{R}^3$, the projection $(I * h_n Y_{n,m})(\mathbf{x}) = \langle h_n Y_{n,m}, I(\mathbf{x} - \cdot) \rangle$ measures the correlation of $h_n Y_{n,m}$ with I at \mathbf{x} . We call $\mathcal{G}_n^{SSE}\{I\}$ the *SSE response map of degree n* of I . The latter are equivariant to translations and (global) rotations as defined in (1) and (2). The proof is given in Appendix E. Note that (11) defines an operator with rotational equivariance, while being sensitive to directional
 220 information via spherical frequencies of degree $n > 0$. Moreover, the image operator \mathcal{G}_n^{SSE} is local if and only if the radial profiles h_n have a finite support, what we always assume thereafter. Finally, more complete invariant quantities can be computed from the solid spherical harmonic representation (Kakarala (2012); Kazhdan et al. (2003); Oreiller et al. (2020)).

225 *2.5. Global RI*

In this section, we define a global RI layer which will later be used to compare against local invariance. As defined in Section 2.2 and illustrated in Fig. 1 (right), an LRI operator is invariant to local rotations that are not constrained to be identical at every position \mathbf{x}_0 . This is required to characterize important

local structures (e.g. textons) having arbitrary and most likely different local orientations. In this section, we want to compare this LRI with the case where the local rotations use one shared orientation across all positions of the entire image. Note that this setting is similar to using standard kernels convolved on the entire image in a regular CNN. We choose this orientation so that a global RI is achieved (i.e. as illustrated in Fig. 1 left) but without invariance to local rotations of patterns (Fig. 1 right). This global RI can be obtained from equivariant representations with orientation channels by first using a spatial Global Average Pooling (GAP), followed by max-pooling on the orientation channels. In this way, the average response is invariant to global rotations R , resulting in a scalar feature μ^{RI} given by

$$\mu^{RI}\{I\} = \max_{R \in SO(3)} \int_{\mathbb{R}^3} |(I * f(R \cdot))(\mathbf{x})| d\mathbf{x}. \quad (12)$$

Note that the order of the GAP and the orientation max-pooling operations is simply swapped as compared to the aggregation of the G-LRI and S-LRI. We can think of equation (12) as finding the rotation of the image I that maximizes the average response to the filter f . Note that this RI layer shares similar ideas with a test-time augmentation. However, the filters are rotated rather than the images, the maximum is taken individually for each filter, and it is also applied at training time.

2.6. Discretization

The discretization of the methods, defined so far in the continuous domain, is necessary for their implementation and naturally introduces an approximation of the invariance properties defined in Section 2.2.

2.6.1. Rotations Sampling

Sampling the rotations, defined so far continuously in the S-LRI and G-LRI approaches, is necessary to compute the invariant responses as in (8) and (9). We sample rotations $R \in B \subset SO(3)$, where B is a finite subset of sampled rotations. To this end, we uniformly sample orientations as points on the sphere

using a triangulation method that iteratively splits octahedron faces to obtain the Euler angles (α, β) around z and y' respectively. We then sample the last angle γ around z'' uniformly between 0 and 2π . The octahedral group O , for instance, is obtained by sampling 6 points on the sphere (i.e. six (α, β) pairs) and four values of γ to obtain the 24 right-angle rotations. We denote by $M = |B|$ the number of tested rotations.

In this paper, we evaluate the following sets of rotations: no rotation ($M = 1$), Klein's four rotations ($M = 4$), octahedral group of rotations ($M = 24$) and 72 rotations ($M = 72$ with 18 points on the sphere and 4 values of γ). For the G-LRI, we restrict the evaluation to the octahedral group as implemented in Winkels and Cohen (2019) as evaluating more rotations becomes computationally too expensive and requires interpolation for non right-angle rotations. It is worth noting that for both G-LRI and S-LRI designs, LRI being obtained by max-pooling over the M orientation channels after the first convolution, this rotation sampling results in an approximated invariance. Finally, for a discrete G-CNN, it is required that $B = G$ is a finite subgroup of $SO(3)$, which is not needed in our case, since we do not propagate the equivariance to the next layer.

2.6.2. Naive Filter Discretization

In the G-LRI (2.4.1), the filters f are simply voxelized to 3D kernels of c^3 voxels as in a standard 3D CNN or G-CNN architecture. All the voxels are trainable parameters that are shared across rotations.

2.6.3. Radial Profiles

In both the S-LRI and SSE-LRI methods, the radial profiles $h_{i,n}$ (and hence the filters f_i) have a compact spherical support $G = \{\mathbf{x} \in \mathbb{R}^3, \|\mathbf{x}\| \leq \rho_0\}$, where $\rho_0 > 0$ is fixed. For any i and n , we consider the voxelized version of the radial profile $h_{i,n}(\rho)$. The size of the support of the voxelized version is related to the maximum radius ρ_0 of the filter in the continuous domain and the level of voxelization. Due to the isotropic constraint, for a support of c^3 voxels, the number of trainable parameters for each $h_{i,n}$ is $\left\lceil \frac{(c-1)}{2} \times \sqrt{3} \right\rceil + 1$. The values

of the filter $f_i(\rho, \theta, \phi)$ over the continuum is deduced from the discretization of the voxelized radial profile on the 3D discrete grid using linear interpolation³.

The maximal degree N cannot be taken arbitrarily large once the radial profiles are voxelized. Indeed, the discretized filters f_i are defined over c^3 voxels, which imposes the restriction that $N \leq \pi c/4$, which can be interpreted as the spherical Nyquist frequency.




Figure 3: Illustration of a 2D slice of the isotropic radial profiles $h_{i,n}$ with $c = 7$. The blue voxels represent the trainable parameters. The rest of the cube is linearly interpolated.

2.7. Datasets

We evaluate the proposed method with two experiments described in the following.

The first experiment is a sanity check to ensure the relevance of the LRI property. We built a dataset for texture classification containing two classes with 500 synthetic volumes each. The volumes of size $32 \times 32 \times 32$ are generated by placing two $7 \times 7 \times 7$ patterns, namely a binary segment and a 2D cross with the same norm, at random 3D orientations and random locations with overlap. The number of patterns per volume is randomly set to $\lfloor d(\frac{s_v}{s_p})^3 \rfloor$, where

³Note that this discretization is not truly isotropic due to the corner effect as the last weights of the radial profile $h_{i,n}$ only affect the corners of the interpolated cubes. While this could be avoided by reducing the length of the radial profile, we favor this implementation for the following reasons. With right-angle rotations, cubic filters are optimal and do not deteriorate the already approximated rotation invariance. Besides, the isotropy can easily be learned by forcing the corner weights to zero for finer rotation samplings.

s_v and s_p are the sizes of the volume and of the pattern respectively and the density d is drawn from a uniform distribution in the range $[0.1, 0.5]$. The two texture classes vary by the proportion of the patterns, i.e. 30% segments with 70% crosses for the first class and vice versa for the second class. 800 volumes are used for training and the remaining 200 for testing. Despite the simplicity of this dataset, some variability is introduced by the overlapping patterns and the linear interpolation of the 3D rotations, making it challenging and more realistic. A 2D schematic illustration of the 3D synthetic textures is shown in Fig. 4.




Figure 4: 2D schematic illustrations of the 3D synthetic textures.

The second dataset is a subset of the American National Lung Screening Trial (NLST) that was annotated by radiologists at the University Hospitals of Geneva (HUG) Martin et al. (submitted). The dataset includes 485 pulmonary nodules from distinct patients in CT, among which 244 were labeled benign and 241 malignant. We zero-pad or crop the input volumes (originally ranging from $16 \times 16 \times 16$ to $128 \times 128 \times 128$) to the size $64 \times 64 \times 64$. We use balanced training and test splits with 392 and 93 volumes respectively. Examples of 2D slices of the lung nodules are illustrated in Fig. 5. The Hounsfield units of the training and test volumes are clipped in the range $[-1000, 400]$, then normalized with zero mean and unit variance (using the mean and variance of cropped training volumes).




Figure 5: 2D slices from 3D volumes of benign and malignant pulmonary nodules.

2.8. Network Architecture

The architecture details of the CNNs are provided in this section⁴. The first layer of the LRI networks consists of one of the LRI layers (G-LRI, S-LRI or SSE-LRI). The responses are aggregated using spatial GAP after the first layer, similarly to Andrearczyk and Whelan (2016) as described in Section 2.2. This pooling aggregates the LRI operator responses into a single scalar per feature map and is followed by Fully Connected (FC) layers. For the nodule classification experiment, we average the responses inside the nodule masks instead of across the entire feature maps. This operation is performed to focus on the texture inside the nodule. The receptive fields are small and the sizes of the nodules vary largely across cases, making a standard GAP less appropriate for the proposed study. For the synthetic experiment, we connect directly the final softmax FC layer with a cross-entropy loss. For the nodule classification, we use an intermediate fully connected layer with 128 neurons before the same final layer. Standard Rectified Linear Units (ReLU) activations are employed. The networks are trained using Adam optimizer with $\beta_1 = 0.99$ and $\beta_2 = 0.9999$ and a batch size of 8. Other task-specific parameters are: for the synthetic experiment (kernel size $7 \times 7 \times 7$, stride 1, 2 filters and 50,000 iterations), for the nodule classification experiment (kernel size $9 \times 9 \times 9$, stride 2, 4 filters and 10,000 iterations). The number of iterations was fixed to these values as the networks reach a plateau beyond these values.

⁴Code available on GitHub github.com/v-andrearczyk/lri-cnn.

We compare the proposed architectures to a network with the same architecture but with a standard 3D convolutional layer, referred to as Z3-CNN.

2.9. Weights Initialization

330 The SHs are normalized to $\|Y_{n,m}\|_2 = 1$. The coefficients are then randomly initialized by a normal distribution with $C_{i,n}[m] \sim \mathcal{N}(0, \sigma^2)$, with $\sigma^2 = \frac{2}{n_{in}(N+1)^2}$ and n_{in} is the number of input channels (generally 1), the radial profiles are initialized to $h_{i,n}(\rho) \sim \mathcal{N}(0, 1)$ and the biases to zero. This initialization is inspired by He et al. (2015); Weiler et al. (2017) in order to avoid
335 vanishing and exploding activations and gradients.

3. Experimental Results

In this section, we experimentally evaluate and compare standard CNNs (Z3-CNN with or without rotational data augmentation), the three proposed approaches to achieve LRI image analysis (i.e. G-LRI, S-LRI and SSE-LRI) as
340 well as global RI with G-RI and S-RI. The two datasets and tasks described in Section 2.7 are used. The approximation capability of the SH-based parametric representation is first evaluated in Section 3.1. Z3 and LRI approaches are then compared in Section 3.2. The importance of LRI when compared to global RI is investigated in Section 3.3. Finally, the complexity of networks is compared in
345 terms of computational time and number of trainable parameters in Section 3.4. The results will be discussed in Section 4.

3.1. SH Parametric Approximation Capability

Fig. 6 compares standard 3D kernels (Z3-CNN) with the SH parametric representation (S-LRI with $M = 1$ tested orientation) using either polar separable
350 (i.e. h) or non-polar separable (i.e. h_n) implementations. Based on a trade-off between complexity and performance, we select the maximum degree $N = 3$ in the following experiments. We set $N = 2$ for the lung nodule dataset with a similar analysis.




Figure 6: Comparison of standard 3D kernels (Z3) and the SH parametric representation (S-LRI) with varying maximum degree N using a single orientation $M = 1$ (i.e. not using the steering capacity) on the synthetic 3D texture dataset. The polar separable and non-polar separable versions are respectively denoted S-LRI- h and S-LRI- h_n . The average accuracy (random 50%) and standard error (10 repetitions) are reported as well as the numbers of parameters.




Figure 7: Average accuracy (random 50%) and standard error (10 repetitions) on the synthetic dataset for the SSE-LRI with varying values of N and comparison with the standard 3D-CNN (Z3).

In Fig. 7, we report the performance of the SSE-LRI with varying maximal degree N on the synthetic dataset. These results do not aim at evaluating the parametric representation but rather to evaluate the influence of N for the SSE-LRI and in order to choose a value (also $N = 3$ and $N = 2$ for the two datasets respectively) for further comparisons in the following experiments.

3.2. Comparing Standard and LRI Architectures

The influence of the number of tested orientations M is investigated in Fig. 8 using the S-LRI for both the synthetic and lung nodule datasets. These results illustrate the benefit of a fine rotation sampling, as well as the better perfor-




Figure 8: Average accuracy (%) and standard error (10 repetitions) of the S-LRI with maximum degree $N = 3$ and $N = 2$ for the synthetic and NLST datasets respectively and varying numbers of orientations M . Left: synthetic dataset; right: lung nodule dataset. Black dashed lines represent the accuracy of the Z3-CNN.

mance of the non-polar separable kernels.

We now use the best reported values of M and compare all proposed LRI
 365 approaches to the standard Z3-CNN. The results are summarized in Tables 2
 and 3 for the synthetic experiment and the lung classification experiment, re-
 spectively.

3.3. Comparing Local and Global RI

The importance of LRI is investigated in this section by comparing LRI,
 370 global Rotation Invariance (RI) and rotational data augmentation. The latter
 consists in randomly rotating the volumes by right-angle rotations during train-
 ing. Corresponding results are reported in Tables 4 and 5 for the synthetic and
 lung nodule datasets, respectively.

3.4. Networks Complexity: Computational Time and Trainable Parameters

375 The number of trainable parameters and the computational time are re-
 ported in Table 6, where the Z3, S-LRI and SSE-LRI are compared. Their

model	N	M	# filters	# param.	accuracy $_{\pm\sigma}$
Z3	-	-	2	694	78.8 \pm 7.1
Z3	-	-	144	49,826	94.0 \pm 0.7
G-LRI	-	24	2	694	89.0 \pm 5.1
SSE-LRI- h	3	-	2	40	90.1 \pm 1.5
SSE-LRI- h_n	3	-	2	82	91.0 \pm 0.8
S-LRI- h	3	24	2	54	91.3 \pm 4.4
S-LRI- h	3	72	2	54	92.0 \pm 1.9
S-LRI- h_n	3	24	2	96	93.5 \pm 0.8
S-LRI- h_n	3	72	2	96	94.2\pm1.1

Table 2: Average accuracy (%) and standard deviation on the synthetic 3D local rotation dataset of all LRI approaches and comparison with a standard CNN (Z3).

polar versus non-polar separable versions are also detailed. The calculation of the number of trainable parameters is provided in Appendix F.

4. Discussions

380 We discuss and interpret the results detailed in Section 3 in terms of the general importance of LRI image analysis (Section 4.1), optimal LRI design (Section 4.2), as well as kernel compression (reduction of trainable parameters) and interpretability (Section 4.3).

4.1. Importance of LRI

385 The results reported in the previous section demonstrated the importance of LRI image analysis in the proposed experiments. Particularly, best results are obtained with LRI architectures (S-LRI- h_n on the synthetic dataset and G-LRI for pulmonary nodule classification) as reported in Tables 2 and 3. Besides, LRI performs significantly better than RI in these experiments where local patterns
390 occur at random orientations, as shown in Tables 4 and 5. In addition, despite improving the performance of the standard Z3-CNN, rotation data augmenta-

model	N	M	# filters	# param.	accuracy $_{\pm\sigma}$
Z3	-	-	4	3,818	80.0 \pm 1.7
Z3	-	-	96	82,754	81.3 \pm 2.2
G-LRI	-	24	4	3,818	87.7 \pm 2.2
SSE-LRI- h	2	-	4	1,966	81.1 \pm 2.2
SSE-LRI- h_n	2	-	4	2,030	81.3 \pm 2.6
S-LRI- h	2	24	4	970	81.9 \pm 3.3
S-LRI- h_n	2	72	4	1,034	84.2 \pm 3.4

Table 3: Average accuracy (%) and standard deviation on the pulmonary nodule classification of all LRI approaches and comparison with a standard CNN (Z3).

model	N	M	# filters	# param.	accuracy $_{\pm\sigma}$
Z3	-	-	2	694	78.8 \pm 7.1
Z3 augm.	-	-	2	694	84.0 \pm 5.2
G-RI	-	24	2	694	79.0 \pm 5.6
G-LRI	-	24	2	694	89.0 \pm 5.1
S-RI- h_n	3	72	2	94	85.9 \pm 5.5
S-LRI- h_n	3	72	2	96	94.2 \pm 1.1

Table 4: Average accuracy (%) and standard deviation on the synthetic 3D local rotation dataset comparing LRI, RI and data augmentation (random right-angle rotations).

tion is not sufficient to obtain LRI (Table 4 and 5). Adding more filters⁵ to the Z3-CNN (second rows of Tables 2 and 3) allows learning filters at different orientations at the heavy cost of a large number of parameters and convolution operations without reaching the performance of LRI networks.

In the SSE-LRI, the number of output feature maps and trainable parameters of the LRI convolution increases with N (see (11)). To show that the performance gain is not solely due to more feature maps and parameters but

⁵The number of filters is multiplied by the number of orientations used in the LRI methods for a fair comparison.

model	N	M	# filters	# param.	accuracy $_{\pm\sigma}$
Z3	-	-	4	3,818	80.0 \pm 1.7
Z3 augm.	-	-	4	3,818	82.2 \pm 3.3
G-RI	-	24	4	3,818	81.8 \pm 2.1
G-LRI	-	24	4	3,818	87.7 \pm 2.2
S-RI- h_n	2	72	4	1,030	81.8 \pm 3.1
S-LRI- h_n	2	72	4	1,034	84.2 \pm 3.4

Table 5: Average accuracy (%) and standard deviation on the lung nodule dataset comparing LRI, RI and data augmentation (random right-angle rotations).

model	Z3	S-LRI- h			S-LRI- h_n			SSE-LRI- h			SSE-LRI- h_n		
N	-	0	3	6	0	3	6	0	3	6	0	3	6
# param.	694	24	54	120	24	96	204	22	40	58	22	82	142
time	28	56	95	196	56	97	211	30	64	146	30	64	146

Table 6: Computational and parameters comparison on the synthetic dataset with the setups of Table 2 (some extra setups are included for comparison). The computational time is measured in seconds for 1,000 iterations trained on a Tesla K80 GPU.

rather to a better approximation capability, we evaluate the SSE-CNN with
 400 $N = 0$ and more output channels ($C = 8$ instead of $C = 2$). This setup relies
 on a number of output feature maps that is equal to the SSE-CNN with $N = 3$
 in Table 2. Yet, the accuracy of the former is 81.5% \pm 3.2 versus 90.1% \pm 1.5 for
 the latter. This result highlights the relevance of SSE quantities extracted at
 various degrees n and, in turn, the importance of directional sensitivity. Note
 405 that the synthetic patterns do not have exactly the same zero frequency, ex-
 plaining the fact that both S-LRI and SSE-LRI designs based on $N = 0$, i.e.
 using directionally insensitive filters, can discriminate the two classes to some
 extent.

4.2. Comparison of LRI Methods

410 LRI can be obtained by G-convolution, steering, or invariants computed
 from the SH responses, as summarized in Fig. 2. The results of the parametric

representation (Fig. 6), confirmed by the following results using steerability or spherical energy (Figures 7, 8 and Tables 2, 3), show that non-polar separable (radial profile h_n) filters perform better than polar separable ones (radial profile h). In particular, it provides more flexibility to learn the optimal combination of different degrees in (4) with individual radial profiles, resulting in a better parametric approximation capability with a limited increase of parameters. To ensure that the performance gain is not only due to the increase of parameters, we incrementally increased from two to five the number of filters of the polar separable design for the synthetic dataset (number of trainable parameters ranging from 54 to 132). The highest accuracy with $M = 72$ orientations is 92.4%, outperformed by the non-polar separable design with 94.2%.

A fine sampling of orientations ($M=72$) is beneficial to the S-LRI (see Fig. 8), particularly outperforming the other architectures on the synthetic dataset (Table 2). The polar-separable S-LRI- h , however, may be too simplistic to benefit from a finer orientation sampling on the nodule dataset (Fig. 8, right). The SSE-LRI offers a trade-off between performance and computation. It does not require steering the SH responses (the locally rotation equivariant representation not being required), resulting in a reduction of memory and operations requirements but at the cost of a lower kernel specificity: the solid spherical energy mixes responses of different SH patterns (for $n > 0$) as well as inter-degree phases, thus discarding some potentially valuable discriminatory information for later layers.

On the NLST dataset, the G-LRI performs better than the S-LRI. The local patterns in the lung nodule dataset may be easier to represent in a non-parametric form as compared to the synthetic patterns.

4.3. Compression and Interpretability

The number of trainable parameters is largely reduced when using S-LRI and SSE-LRI as compared to a standard Z3-CNN. We identify two distinct factors enabling parameter reduction, namely the *weight sharing* across rotations and the *parametric representation*. The *weight sharing* is present, among others,

in the G-LRI where the same kernels account for multiple orientations. Similarly, the S-LRI and SSE-LRI architectures share trainable radial profiles and harmonic coefficients to obtain responses to rotated kernels. The *parametric*
 445 *representation*, on the other hand, is used in Andrearczyk et al. (2019a); Weiler et al. (2017); Worrall et al. (2017) by learning a combination of basis filters instead of every voxel of the kernels. A steerable basis was also used in McCann et al. (2020) for a fast rotational sparse coding. Fig. 6 shows the performance of the parametric representation (with a single orientation of the S-LRI), even
 450 outperforming the standard Z3-CNN that contains many more parameters. The proposed SSE-LRI approach goes one step further as there is no need to learn an explicit full parametric representation of the kernels. For computing the SSE-LRI invariants, we only calculate the norm of the SHs responses, resulting in a large reduction of the number of trainable parameters.

455 One bottleneck with both G-CNN (Winkels and Cohen (2019)) and S-LRI (Andrearczyk et al. (2019a)) is the GPU memory usage to store M 3D response maps (i.e. all orientations of the equivariant representation) before orientation pooling. This memory consumption is drastically reduced in the SSE-LRI by computing invariants on the SHs responses rather than calculating responses
 460 at all orientations. Note that the S-LRI and SSE-LRI implementations only use existing TensorFlow functions and the computation time could be further improved by efficient parallelization and CUDA programming.

Finally, in terms of network interpretability, the hard-coded equivariance and invariance enforces a geometric structure of the hidden features improving
 465 the transparency and decomposability of the network where transformations (translations and rotations) in the inputs result in predictable transformations in the activations (see Cheng et al. (2018); Lipton (2018); Worrall and Brostow (2018); Worrall et al. (2017) for discussions on the matter).

5. Conclusion

470 This paper explored the use of LRI in the context of 3D texture analysis. Three architectures were proposed and compared with standard 3D-CNN, global RI and data augmentation. The results showed the importance of LRI in medical imaging and, more generally, in texture analysis where repeated patterns occur at various locations and orientations. In particular, we showed that data
475 augmentation, commonly used in CNN training, is not sufficient to learn such an invariance and specific architectures with built-in invariance are beneficial.

In future work, we plan to explore deeper networks, building deep architectures on top of LRI layers.

6. Acknowledgments

480 This work partially supported by the Swiss National Science Foundation (grant 205320_179069), the Swiss Personalized Health Network (IMAGINE and QA4IQI projects), as well as a hardware grant from NVIDIA.

7. Declarations of interest

Declarations of interest: none.

485 References

References

- Abramowitz, M., Stegun, I., 1964. Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables. volume 55. Courier Corporation.
- 490 Andrearczyk, V., Depeursinge, A., 2018. Rotational 3D texture classification using group equivariant CNNs. arXiv preprint arXiv:1810.06889 .

- Andrearczyk, V., Fageot, J., Oreiller, V., Montet, X., Depeursinge, A., 2019a. Exploring local rotation invariance in 3d cnns with steerable filters, in: International Conference on Medical Imaging with Deep Learning.
- 495 Andrearczyk, V., Oreiller, V., Fageot, J., Montet, X., Depeursinge, A., 2019b. Solid spherical energy (SSE) CNNs for efficient 3D medical image analysis, in: Irish Machine Vision and Image Processing Conference.
- Andrearczyk, V., Whelan, P., 2016. Using filter banks in convolutional neural networks for texture classification. *Pattern Recognition Letters* 84, 63–69.
- 500 Bekkers, E.J., 2019. B-Spline CNNs on Lie groups. arXiv preprint arXiv:1909.12057 .
- Bekkers, E.J., Lafarge, M.W., Veta, M., Eppenhof, K.A., Pluim, J.P., Duits, R., 2018. Roto-translation covariant convolutional networks for medical image analysis, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 440–448.
- 505 Brown, L.G., 1992. A survey of image registration techniques. *ACM computing surveys (CSUR)* 24, 325–376.
- Cheng, X., Qiu, Q., Calderbank, R., Sapiro, G., 2018. RotDCF: Decomposition of convolutional filters for rotation-equivariant deep networks. arXiv preprint arXiv:1805.06846 .
- 510 Chenouard, N., Unser, M., 2012. 3D steerable wavelets in practice. *IEEE Transactions on Image Processing* 21, 4522–4533.
- Cohen, T., Welling, M., 2016. Group equivariant convolutional networks, in: Balcan, M.F., Weinberger, K.Q. (Eds.), *Proceedings of The 33rd International Conference on Machine Learning*, PMLR, New York, New York, USA. pp. 2990–2999.
- 515 Cohen, T.S., Geiger, M., Köhler, J., Welling, M., 2018. Spherical CNNs, in: International Conference on Learning Representations.

- 520 Cohen, T.S., Geiger, M., Weiler, M., 2019. A general theory of equivariant CNNs on homogeneous spaces, in: *Advances in Neural Information Processing Systems*, pp. 9142–9153.
- Depeursinge, A., Fageot, J., 2018. Biomedical texture operators and aggregation functions: A methodological review and user’s guide, in: *Biomedical Texture Analysis*. Elsevier, pp. 55–94.
- 525 Depeursinge, A., Fageot, J., Andrearczyk, V., Ward, J., Unser, M., 2018. Rotation invariance and directional sensitivity: Spherical harmonics versus radiomics features, in: *International Workshop on Machine Learning in Medical Imaging*, Springer. pp. 107–115.
- Dicente Cid, Y., Müller, H., Platon, A., Poletti, P., Depeursinge, A., 2017. 3-D solid texture classification using locally-oriented wavelet transforms. *IEEE Transactions on Image Processing* 26, 1899–1910. doi:10.1109/TIP.2017.2665041.
- Driscoll, J.R., Healy, D.M., 1994. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in applied mathematics* 15, 202–250.
- 535 Eickenberg, M., Exarchakis, G., Hirn, M., Mallat, S., 2017. Solid harmonic wavelet scattering: Predicting quantum molecular energy from invariant descriptors of 3D electronic densities, in: *Advances in Neural Information Processing Systems*, pp. 6540–6549.
- Fageot, J., Uhlmann, V., Püspöki, Z., Beck, B., Unser, M., Depeursinge, A., 540 2018. Principled design and implementation of steerable detectors. arXiv preprint arXiv:1811.00863 .
- Freeman, W., Adelson, E., 1991. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence* , 891–906.
- Gatenby, R., Grove, O., Gillies, R., 2013. Quantitative imaging in cancer evolution and ecology. *Radiology* 269, 8–14.
- 545

- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: Proceedings of the IEEE international conference on computer vision, pp. 1026–1034.
- Kakarala, R., 2012. The bispectrum as a source of phase-sensitive invariants
 550 for fourier descriptors: a group-theoretic approach. *Journal of Mathematical Imaging and Vision* 44, 341–353.
- Kazhdan, M., Funkhouser, T., Rusinkiewicz, S., 2003. Rotation invariant spherical harmonic representation of 3 d shape descriptors, in: Symposium on geometry processing, pp. 156–164.
- 555 Kondor, R., Lin, Z., Trivedi, S., 2018. Clebsch-Gordan nets: a fully Fourier space spherical convolutional neural network, in: Advances in Neural Information Processing Systems, pp. 10117–10126.
- Kondor, R., Trivedi, S., 2018. On the generalization of equivariance and convolution in neural networks to the action of compact groups. arXiv preprint
 560 arXiv:1802.03690 .
- Lipton, Z.C., 2018. The mythos of model interpretability. *Queue* 16, 31–57.
- Liu, L., Chen, J., Fieguth, P., Zhao, G., Chellappa, R., Pietikäinen, M., 2019. From BoW to CNN: Two Decades of Texture Representation for Texture Classification. *International Journal of Computer Vision* 127, 74–109.
- 565 Martin, S.P., Hofmeister, J., Burgmeister, S., Orso, S., Mili, N., Guerrier, S., Victoria-Fesser, M.P., Soccia, P.M., Triponez, F., Karenovics, W., Mach, N., Depoersing, A., Becker, C.D., Rampinelli, C., Summers, P., Müller, H., Montet, X., submitted. Identification of malignant lung nodules and reduction in false-positive findings by augmented intelligence: A radiomic study based
 570 on the NLST dataset. *Journal of Clinical Oncology* .
- McCann, M.T., Vincent, A., Unser, M., Depoersing, A., 2020. Fast rotational sparse coding. arXiv preprint arXiv:1806.04374 .

- Ojala, T., Pietikäinen, M., Mäenpää, T., 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 971–987.
575
- Oreiller, V., Andrearczyk, V., Fageot, J., John O, P., Depeursinge, A., 2020. 3D solid spherical bispectrum CNNs, in: (submitted) *IEEE Transactions on Medical Imaging*.
- Unser, M., Chenouard, N., 2013. A Unifying Parametric Framework for 2D Steerable Wavelet Transforms. *SIAM Journal on Imaging Sciences* 6, 102–135.
580
- Varma, M., Zisserman, A., 2005. A statistical approach to texture classification from single images. *International Journal of Computer Vision* 62, 61–81. doi:10.1007/s11263-005-4635-4.
- 585 Varshalovich, D., Moskalev, A., Khersonskii, V., 1988. Quantum theory of angular momentum. World Scientific.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., Cohen, T., 2018. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data, in: *NeurIPS*.
- 590 Weiler, M., Hamprecht, F.A., Storath, M., 2017. Learning steerable filters for rotation equivariant CNNs. *2017 IEEE/CVF Conference on Computer Vision and Pattern Recognition* , 849–858.
- Winkels, M., Cohen, T.S., 2019. Pulmonary nodule detection in CT scans with equivariant CNNs. *Medical image analysis* 55, 15–26.
- 595 Worrall, D., Brostow, G., 2018. CubeNet: Equivariance to 3D rotation and translation. *ECCV, Lecture Notes in Computer Science* 11209, 585–602.
- Worrall, D.E., Garbin, S.J., Turmukhambetov, D., Brostow, G.J., 2017. Harmonic networks: Deep translation and rotation equivariance. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* , 7168–7177.
600

Appendix A. Proof of Proposition 1

First of all, any convolution operator \mathcal{G} is equivariant to translations by construction. We therefore focus on the equivariance to rotations.

Let f and g be two functions and $R_0 \in SO(3)$. A simple change of variable implies that $(f * g(R_0^{-1}\cdot)) = (f(R_0\cdot) * g)(R_0^{-1}\cdot)$. Exploiting this relation, we deduce that, for any image $I \in L_2(\mathbb{R}^3)$ and any rotation $R_0 \in SO(3)$,

$$(h(R_0^{-1}\cdot) * I) = (h * I(R_0\cdot))(R_0^{-1}) = \mathcal{G}\{I(R_0\cdot)\}(R_0^{-1}). \quad (\text{A.1})$$

Assume that \mathcal{G} is equivariant to rotations, then $\mathcal{G}\{I(R_0\cdot)\}(R_0^{-1}) = \mathcal{G}\{I\} = h * I$.
 605 Therefore, $(h(R_0^{-1}\cdot) * I) = h * I$ for any I , which implies the equality $h(R_0^{-1}\cdot) = h$ for any rotation and h is isotropic. Reciprocally, if h is isotropic, (A.1) shows that $\mathcal{G}\{I(R_0\cdot)\}(R_0^{-1}) = \mathcal{G}\{I\}$ for any I and any rotation R_0 , which is equivalent to the rotation equivariance of \mathcal{G} .

Moreover, the values of $\mathcal{G}\{I\}(\mathbf{x})$ depend on local image values $I(\mathbf{y})$ for $\mathbf{y} - \mathbf{x}$
 610 in the support of h . Hence, \mathcal{G} is local if and only if the support of h is compact.

Finally, \mathcal{G} is LRI if and only if h is isotropic (for the global equivariance to rotations) and compactly supported (for the locality).

Appendix B. Spherical Harmonics

The family of SHs is denoted by $(Y_{n,m})_{n \geq 0, m \in \{-n, \dots, n\}}$, where n is called the degree and m the order of $Y_{n,m}$. SHs form an orthonormal basis for square-integrable functions in the $2D$ -sphere \mathbb{S}^2 . They are defined as (Driscoll and Healy (1994))

$$Y_{n,m}(\theta, \phi) = A_{n,m} P_{n,|m|}(\cos(\theta)) e^{jm\phi}, \quad (\text{B.1})$$

with $A_{n,m} = (-1)^{(m+|m|)/2} \left(\frac{2n+1}{4\pi} \frac{(n-|m|)!}{(n+|m|)!} \right)^{1/2}$ a normalization constant and $P_{n,|m|}$ the associated Legendre polynomial given for $0 \leq m \leq n$ by

$$P_{n,m}(x) := \frac{(-1)^m}{2^n n!} (1-x^2)^{m/2} \frac{d^{n+m}}{dx^{n+m}} (x^2-1)^n. \quad (\text{B.2})$$

We refer to Abramowitz and Stegun (1964) for more details.

615 **Appendix C. Real Steerable Filters**

Proposition 2. *A function f of the form (4) is real if and only if we have, for every $0 \leq n \leq N$, $-n \leq m \leq n$,*

$$\forall \rho \geq 0, \quad C_n[-m]h_n(\rho) = (-1)^m \overline{C_n[m]h_n(\rho)}. \quad (\text{C.1})$$

If we impose moreover that the radial profiles h_n are real, then this is equivalent to

$$C_n[-m] = (-1)^m \overline{C_n[m]} \quad (\text{C.2})$$

for every $0 \leq n \leq N$, $-n \leq m \leq n$.

Proof. A filter f is real if and only if $\overline{f(\rho, \theta, \phi)} = f(\rho, \theta, \phi)$ for every (ρ, θ, ϕ) . For filters given by (4), this means that

$$\sum_{n,m} \overline{C_n[m]h_n(\rho)Y_{n,m}(\theta, \phi)} = \sum_{n,m} C_n[m]h_n(\rho)Y_{n,m}(\theta, \phi). \quad (\text{C.3})$$

We use the symmetry of the spherical harmonics, $\overline{Y_{n,m}} = (-1)^m Y_{n,-m}$, on the left-hand side and change the sign of m on the right-hand side to get

$$\sum_{n,m} \overline{C_n[m]h_n(\rho)} (-1)^m Y_{n,-m}(\theta, \phi) = \sum_{n,m} C_n[-m]h_n(\rho)Y_{n,-m}(\theta, \phi). \quad (\text{C.4})$$

The $Y_{n,m}$ being linearly independent, we deduce that the filter is real if and only if (C.1) holds. By imposing that the h_n are real, i.e., $\overline{h_n} = h_n$, we obtain the expected criterion on the $C_n[m]$ coefficients, which is (C.2). \square

620 **Appendix D. Equivariant Image Operators via Orientation Channels**

This result is reported for completeness, yet already proven in our previous publication (Andrearczyk et al. (2019a)).

Proposition 3. *An image operator of the form (8) and (9) is equivariant to translations and rotations in the sense of (1) and (2) and therefore LRI when f is compactly supported.*

Proof. The equivariance to translations uses $(I(\cdot - \mathbf{x}_0) * g)(\mathbf{x}) = (I * g)(\mathbf{x} - \mathbf{x}_0)$. Applying this to $g = f(\mathbf{R}\cdot)$, we deduce

$$\mathcal{G}\{I(\cdot - \mathbf{x}_0)\}(\mathbf{x}) = \max_{\mathbf{R} \in SO(3)} |(I * f(\mathbf{R}\cdot))(\mathbf{x} - \mathbf{x}_0)| = \mathcal{G}\{I\}(\mathbf{x} - \mathbf{x}_0), \quad (\text{D.1})$$

as expected. For rotations, we use $(I(\mathbf{R}_0\cdot) * g)(\mathbf{x}) = (I * g(\mathbf{R}_0^{-1}\cdot))(\mathbf{R}_0\mathbf{x})$ applied to $g = f(\mathbf{R}\cdot)$ to deduce

$$\mathcal{G}\{I(\mathbf{R}_0\cdot)\}(\mathbf{x}) = \max_{\mathbf{R} \in SO(3)} |(I * f(\mathbf{R}\mathbf{R}_0^{-1}\cdot))(\mathbf{R}_0\mathbf{x})| \quad (\text{D.2})$$

$$= \max_{\mathbf{R} \in SO(3)} |(I * f(\mathbf{R}\cdot))(\mathbf{R}_0\mathbf{x})| \quad (\text{D.3})$$

$$= \mathcal{G}\{I\}(\mathbf{R}_0\mathbf{x}), \quad (\text{D.4})$$

where the second equality simply exploits that $\mathbf{R}\mathbf{R}_0^{-1}$ describes the complete space $SO(3)$ of 3D rotations when \mathbf{R} varies. The property of being LRI is a direct consequence of the equivariance by translations and rotations, together
 630 with the fact that f is compactly supported. \square

We remark that the equivariance to translations is simply due to the use of the convolution, while the equivariance to rotations requires pooling over 3D rotations in (8) and (9).

Appendix E. Equivariant Image Operators via SH invariants

635 **Proposition 4.** *An image operator of the form (11) is equivariant to translations and rotations in the sense of (1) and (2), and therefore LRI when the h_n are compactly supported.*

Proof. Given a filter g (e.g. $g = h_n Y_{n,m}$), from the relation $(I(\cdot - \mathbf{x}_0) * g)(\mathbf{x}) = (I * g)(\mathbf{x} - \mathbf{x}_0)$, we deduce that $\mathcal{G}_n^{SSE}\{I(\cdot - \mathbf{x}_0)\}(\mathbf{x}) = \sum_{m=-n}^n |(I * h_n Y_{n,m})(\mathbf{x} - \mathbf{x}_0)|^2 = \mathcal{G}_n^{SSE}\{I\}(\mathbf{x} - \mathbf{x}_0)$, which is (1). Now, using that $(I(\mathbf{R}_0\cdot) * g)(\mathbf{x}) =$

$(I * g(\mathbf{R}_0^{-1} \cdot))(\mathbf{R}_0 \mathbf{x})$ and (6), we have

$$\begin{aligned} \mathcal{G}_n^{SSE}\{I(\mathbf{R}_0 \cdot)\}(\mathbf{x}) &= \sum_{m=-n}^n |(I * (h_n Y_{n,m}(\mathbf{R}_0^{-1} \cdot)))(\mathbf{R}_0 \mathbf{x})|^2 \\ &= \sum_{m=-n}^n \left| \sum_{m'=-n}^n D_{n, \mathbf{R}_0^{-1}}[m, m'] (I * h_n Y_{n,m'}) (\mathbf{R}_0 \mathbf{x}) \right|^2. \end{aligned} \quad (\text{E.1})$$

The Wigner D-matrix being norm-preserving, we have that

$$\sum_m |c_m|^2 = \sum_m \left| \sum_{m'} D_{n, \mathbf{R}_0^{-1}}[m, m'] c_{m'} \right|^2$$

for any $\mathbf{c} = (c_{-n}, \dots, c_n)$. Applying this relation to $c_m = (I * h_n Y_{n,m})(\mathbf{R}_0 \mathbf{x})$, we deduce (2) from (E.1). Finally, the LRI is a consequence of the equivariance to global rotations and translations. \square

640

Appendix F. Number of Trainable Parameters

The trainable parameters include convolutional and fully connected parameters, biases and $C_n[m]$ harmonic coefficients. In this appendix, we develop the calculation of their number for the S-LRI and SSE-LRI architectures.

S-LRI. The number of parameters $n_{total(S-h)}$ of the polar separable S-LRI- h architectures is computed as

$$n_{total(S-h)} = n_f n_r + n_f + (N+1)^2 n_f + n_f n_c + n_c, \quad (\text{F.1})$$

645

where n_f , n_r and n_c are the number of filters, of radial profile parameters (Section 2.6.3) and of classes respectively. For example, in the synthetic experiment for $N = 3$, it sums up to $2 \times 7 + 2 + (3+1)^2 \times 2 + 2 \times 2 + 2 = 54$.

For the non-polar separable S-LRI- h_n , we have a different trainable radial profile for each degree n , resulting in the following:

$$n_{total(S-h_n)} = (N+1)n_f n_r + n_f + (N+1)^2 n_f + n_f n_c + n_c = 96. \quad (\text{F.2})$$

SSE-LRI. The number of parameters of the polar separable SSE-LRI- h architectures is

$$n_{total(SSE-h)} = n_f n_r + n_f(N+1) + (N+1)n_f n_c + n_c. \quad (F.3)$$

In the synthetic experiment for $N = 3$, it sums up to $2 \times 7 + 2 \times 4 + 4 \times 2 \times 2 + 2 = 40$.

For the non-polar separable SSE-LRI- h_n , it is

$$n_{total(SSE-h_n)} = (N+1)n_f n_r + n_f(N+1) + (N+1)n_f n_c + n_c = 82. \quad (F.4)$$

Journal Pre-proof

Declaration of Interest

We have no conflict of interest to declare.

Journal Pre-proof