

SurReal: Complex-Valued Learning as Principled Transformations on a Scaling and Rotation Manifold

Rudrasis Chakraborty Yifei Xing Stella X. Yu

Abstract—Complex-valued data is ubiquitous in signal and image processing applications, and complex-valued representations in deep learning have appealing theoretical properties. While these aspects have long been recognized, complex-valued deep learning continues to lag far behind its real-valued counterpart.

We propose a principled geometric approach to complex-valued deep learning. Complex-valued data could often be subject to arbitrary complex-valued scaling; as a result, real and imaginary components could co-vary. Instead of treating complex values as two independent channels of real values, we recognize their underlying geometry: We model the space of complex numbers as a product manifold of non-zero scaling and planar rotations. Arbitrary complex-valued scaling naturally becomes a group of transitive actions on this manifold.

We propose to extend the property instead of the form of real-valued functions to the complex domain. We define convolution as weighted Fréchet mean on the manifold that is *equivariant* to the group of scaling/rotation actions, and define distance transform on the manifold that is *invariant* to the action group. The manifold perspective also allows us to define nonlinear activation functions such as tangent ReLU and G -transport, as well as residual connections on the manifold-valued data.

We dub our model *SurReal*, as our experiments on MSTAR and RadioML deliver high performance with only a fractional size of real-valued and complex-valued baseline models.

Index Terms—complex value, Riemannian manifold, Fréchet mean, equivariance, invariance.

I. INTRODUCTION

WHILE deep learning has been widely successful in computer vision and machine learning [LeCun et al., 1998, Bengio et al., 2009, Krizhevsky et al., 2012, He et al., 2016, LeCun et al., 2015], most techniques are only applicable to data that lie in a vector space. How to handle manifold-valued data and incorporate non-Euclidean geometry into deep learning has become an active topic of research [Cohen and Welling, 2016, Chakraborty et al., 2018a,b, Esteves et al., 2017, Bronstein et al., 2017, Chakraborty et al., 2018c].

We are interested in extending deep learning to complex-valued data, e.g., synthetic aperture radar (SAR) images in remote sensing, magnetic resonance (MR) images in medical

imaging, or radio frequency (RF) signals in electrical engineering. For such naturally complex-valued data, both the size (or magnitude) and the phase of a complex-valued measurement contain useful information. For example, in SAR images, the magnitude encodes the amount of energy, whereas the phase variation indicates the object material and shape boundaries.

Complex-valued data could also arise from the more informative complex-valued representation of naturally real-valued data. The most notable examples are the Fourier spectrum and spectrum-based computer vision techniques ranging from steerable filters [Freeman and Adelson, 1991] to spectral graph embedding [Maire et al., 2016, Yu, 2011].

The most common complex-valued deep learning approach is to simply apply real-valued deep learning methodology to the two-channel representation of complex-valued data $z = x + i \cdot y$ (where i denotes the imaginary unit): the real component x and the imaginary component y are regarded as independent channels of the input.

However, the independence assumption between the real and imaginary components does *not* hold in general. For instance, in MR and SAR images, the pixel intensity value could be subject to arbitrary scaling by complex number $s = m e^{j\theta}$, where all the pixel values are simultaneously scaled in magnitude by m and shifted in phase by θ . That is, any measurement z is simply a representative of a whole class of possible equivalent measurements $\{z s : s = m e^{j\theta} : m > 0, \forall \theta\}$. Instead of being independent of each other, the real and imaginary components (x, y) of z co-vary in this equivalent class.

The co-variance of the two components of complex-valued data has not been exploited in complex-valued deep learning. The common approach to learn a classifier invariant to scaling is to augment the training data with complex-valued scaling [Krizhevsky et al., 2012, Dieleman et al., 2015, Wang et al., 2017]. Such extrinsic data manipulation increases the amount of the training data and is rather ineffective: It takes a longer time to train the model, yet the invariance is not guaranteed.

Our goal is to develop the invariance to complex-valued scaling as an intrinsic property of the neural network itself. We treat each complex-valued data sample as a point in a non-Euclidean space that respects the intrinsic geometry of complex numbers. We propose new convolution and fully connected layer functions that can achieve equivariance and invariance to complex-valued scaling.

There has been a long line of works which define convolution in a non-Euclidean space by treating each data sample as a function in that space [Worrall et al., 2017, Cohen and Welling, 2016, Cohen et al., 2017, Esteves et al., 2017, Chakraborty

Manuscript received August 26, 2019; revised January 29, 2020 and June 10, 2020; accepted September 29, 2020. This work was supported in part by Berkeley Deep Drive and in part by DARPA. (Corresponding author: Stella X. Yu.)

The authors are with the International Computer Science Institute (ICSI), University of California, Berkeley (UC Berkeley), Berkeley, CA 94720, USA (e-mail: stellayu@berkeley.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2020.3030565

et al., 2018a, Kondor and Trivedi, 2018].

The challenge for defining such an equivariant convolution operator in the non-Euclidean space is the lack of a proper vector space structure. In the Euclidean space, we can move from one point to another using an element from the group of translations; the standard convolution is thus equivariant to the action of the group of translations. However, in the non-Euclidean space, e.g., a hypersphere, translation equivariance is no longer meaningful: Translation is not the group to move from one point to another on a hypersphere, but rotation is.

The concept of equivariance of an operator on a space is thus intimately related to the transitivity of a group of actions on that space. We say that a group G acts transitively on a space if there exists a $g \in G$ to go from one point to another on the space. The group of translations acts transitively on the Euclidean space, whereas the group of rotations acts transitively on a hypersphere. The group that transitively acts on the non-Euclidean space of the complex plane is non-zero scaling and planar rotations in the complex plane.

The manifold view of convolution as an operator with equivariance to transitive actions on that space applies to both the domain and the range of data, e.g., for an image, its pixel coordinates define the domain, and its pixel intensities define the range. Here we focus on the *range* space of data, in order to extend deep learning to complex-valued images and signals.

Our key insight is to represent a complex number in its polar form and define a Riemannian manifold on which complex scaling corresponds to the general transitive action group. When a data sample lies on a Riemannian manifold, there are previously established results for deep learning:

- Convolution defined by weighted Fréchet mean (wFM) [Fréchet, 1948] is equivariant to the group that naturally acts on that manifold [Chakraborty et al., 2018b].
- Since wFM is non-linear and acts like a contraction mapping [Mallat, 2016] analogous to ReLU or sigmoid, non-linear activation functions such as ReLU may not be needed.

We propose three types of complex-valued layer functions from the Riemannian geometric point of view:

- 1) **wFM**: a new convolution operator on the manifold for complex-valued data. It is equivariant to complex-valued scaling. The weights of wFM are to be learned.
- 2) **Tangent ReLU**: a new nonlinear activation function that applies ReLU to the projections in the tangent space of the complex-valued manifold. We also propose another option called G -transport, which transports a point on the complex manifold by an action in the scaling and rotation group.
- 3) **Distance transform**: a new fully-connected layer operator that computes the manifold distance between a feature map and its wFM. It is invariant to complex-valued scaling. The weights of wFM are to be learned.

The distance transform layer takes a complex-valued input to the real-valued domain, where any real-valued convolutional neural network (CNN) functions such as standard convolutions and fully connected (FC) layers can be subsequently used.

Fig. 1 shows a sample CNN architecture composed using our complex-valued layer functions. A complex-valued image

first passes through two wFM complex-valued convolutional layers, and then undergoes the distance transform. The resulting real-valued distances are subsequently fed into a real-valued CNN classifier with one convolution layer and one FC layer. Each convolutional layer is illustrated with a single channel response among a stack of many, with color images encoding complex-valued responses and grayscale images encoding real-valued responses.

Our complex-valued CNN has a group invariant property similar to the standard CNN on real-valued data. Existing methods extend the real-valued counterpart to the complex domain based on the form of functions such as convolution or batch normalization [Bunte et al., 2012, Trabelsi et al., 2017, Virtue et al., 2017], not on the property of functions such as equivariance or linearity. Our complex-valued CNN is composed of layer functions with the desired equivariance and invariance properties that are essential for a real-valued CNN classifier in the Euclidean space; it is thus a theoretically justified analog of the real-valued CNN.

We compare our method with several baselines on two publicly available complex-valued datasets: MSTAR and RadioML. Our model consistently outperforms the real-valued CNN baseline, with fewer than 1% on MSTAR and 3% on RadioML of the baseline model parameters.

We thus name our approach *SurReal* (pun intended): a surprisingly lean complex-valued model that beats the real-valued CNN model. Our work has three major contributions.

- 1) We propose novel complex-valued layer functions with proven equivariance and invariance properties.
- 2) We extend our model to complex-valued residual CNNs.
- 3) We validate our method on classification experiments. Our SurReal CNNs outperform real- and complex-valued baselines at a fraction of their model sizes.

These results demonstrate significant benefits of proposing CNN layer functions in terms of desirable intrinsic properties on the complex plane as opposed to applying the standard CNN to the 2D Euclidean embedding of complex numbers.

II. RELATED WORKS

Complex numbers are powerful representations and concepts in mathematics, with intimate connections to geometry, topology, and differentiation [Needham, 1998]. They have a wide range of applications in physics and engineering.

Complex-valued data representations are widely used as a modeling choice to encode richer information than real-valued representations, especially for directional or cyclic data. [Amin and Murase, 2009] learns a mapping from a finite range of real values to the unit circle in the complex plane. [Cadieu and Olshausen, 2012] trains a complex-valued sparse coding model to capture both edge structure and motion structure. [Yu, 2009, 2012] combine the confidence and size of a measurement in a single complex value, and learns a global embedding from pairwise local measurements. [Maire et al., 2016] simultaneously encodes both grouping and figure-ground ordering relationships between neighboring pixels, and

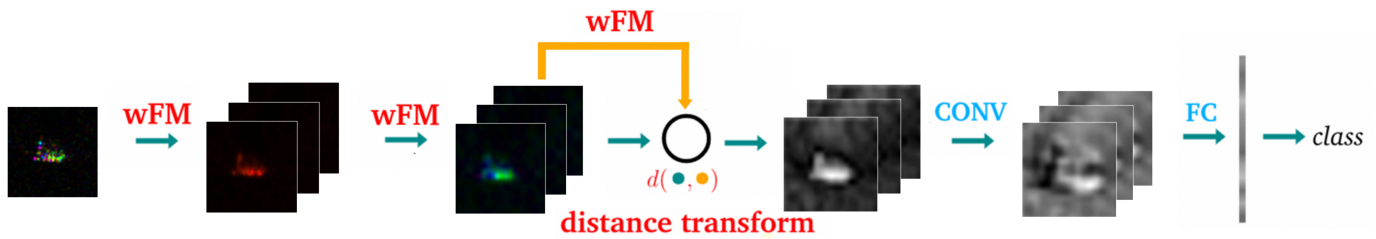


Fig. 1. Sample architecture of our SurReal model. **1)** The input is a complex-valued image; each pixel is color coded with an HSV colormap according to its complex value, with the color intensity indicating the magnitude, the color hue indicating the phase, and the the full constant saturation. **2)** The first two layers are our proposed complex-valued convolution in terms of weighted Fréchet mean filtering (wFM) on the manifold derived from the polar form of complex numbers. This convolution outputs complex-valued responses and it is equivariant to complex-valued scaling. Each wFM layer could have multiple channels, each channel shown here as a complex-valued image. Pixel-wise nonlinear activation functions such as our proposed tReLU and G -transport can be subsequently applied. **3)** The third layer is a distance transformation layer, where the manifold distance between the feature map and its wFM is computed. This distance is invariant to complex-valued scaling. **4)** Once the representation becomes real valued after the distance transformation layer, we could use any real-valued CNN layer functions for classification. Shown here is the real-valued convolution (CONV) layer followed by the fully connected (FC) layer towards the final softmax classification. With the built-in invariance to complex-valued scaling, our SurReal model can outperform real-valued baseline models on complex-valued data with a fraction of the baseline model size.

learns complex-valued pairwise pixel relationships from pixel-wise figure-ground annotations. [Reichert and Serre, 2013] uses complex-valued neuronal units to model biologically plausible deep learning networks. [Bruna and Mallat, 2013, Bruna et al., 2015] adopt wavelet transforms at earlier layers. [Arjovsky et al., 2016] adopts unitary weight matrices in hidden layers for better learning performance.

Traditional complex-valued data analysis utilizes higher-order statistics such as variance fractal dimension trajectory [Kinsner and Grieder, 2010] and spectral analysis [Reichert, 1992] to make adequate predictions.

Early neural network approaches have already noted that complex values have many nice mathematical properties that real-value data do not have, e.g., the complex identity theorem. Transformations from the input to the output can be more effectively learned with complex-valued networks instead of real-valued networks. Various complex-valued activation functions have been explored, although with little demonstration of their success in real data settings [Kim and Guest, 1990, Georgiou and Koutsougeras, 1992, Benvenuto and Piazza, 1992, Nitta, 1997].

Recent neural network approaches continue to build upon the theoretical advantages of complex-valued data to improve the convergence, stability, and generalization of neural networks [Nitta, 2002, Hirose and Yoshida, 2012], and to facilitate the noise-robust memory retrieval mechanisms in capsule networks [Cheng et al., 2019]. Real-valued layer functions have also been extended to the complex domain according to the form of the functions such as convolution, ReLU, and batch normalization [Bunte et al., 2012, Trabelsi et al., 2017, Virtue et al., 2017, Virtue, 2018]. Complex-valued deep learning has also been extended to quaternion neural networks, as quaternions generalize the concept of complex values from 2D to 3D [Parcollet et al., 2018].

Recent graph convolution neural networks open up new computational models in the complex domain [Scarselli et al., 2008, Bruna et al., 2013]. Since convolution in the spatial domain is equivalent to multiplication in the spectral domain, a natural extension of convolution to data defined on an

arbitrary graph is to construct a convolutional filter in terms of multiplicative weights on the spectrum of the graph Laplacian [Bruna et al., 2013]. The spectrum of the graph Laplacian is real-valued if the graph is undirected, and complex-valued if it is directed [Singh et al., 2016].

Our SurReal complex-valued CNN is unique in utilizing the geometric property of the complex numbers and approaching complex-valued learning as a special task of deep learning on Riemannian manifolds.

Existing methods such as [Maire et al., 2016, Trabelsi et al., 2017, Amin and Murase, 2009, Virtue et al., 2017] treat a complex value as a vector in the Euclidean space of \mathbb{R}^2 . This choice, while straightforward, essentially destroys the covariant relationship between real and imaginary parts of a complex number. Naturally complex-valued data such as SAR, MRI, and RF could be subject to complex-valued scaling without changing the underlying observation.

To deal specifically with complex numbers, we first separate and acknowledge the extrinsic scaling effect by asking the convolution operator to be equivariant to complex-valued scaling. For a CNN classifier, we design the distance transform layer to be invariant to complex-valued scaling.

Our SurReal CNN classifier can focus entirely on the discriminative information between classes, without the need to build up additional scale invariance by repeatedly training on data augmented with complex-valued scaling. Our SurReal model is thus a surprisingly lean complex-valued model that beats the real-valued CNN model on complex-valued data. An earlier preliminary version of this work was presented in [Chakraborty et al., 2019].

III. A SCALING-ROTATION MANIFOLD FOR THE GEOMETRY OF COMPLEX NUMBERS

A crucial property of complex-valued data is complex-valued scaling ambiguity: The MRI or SAR images of the same scene could be related by the multiplication of a single complex number, depending on how the data is acquired. Complex-valued scaling can be captured by the scaling action on the magnitude and the rotation action on the phase.

Instead of treating the complex plane as the usual 2D Euclidean space, we identify the non-zero complex plane as the product manifold of positive magnitudes and planar rotations. We show that scaling and rotation actions preserve the manifold distance defined on the non-zero complex manifold.

Space of complex numbers. Let \mathbf{R} and \mathbf{C} denote the field of real numbers and complex numbers respectively. We have:

$$\mathbf{z} = x + iy \in \mathbf{C}, \quad \forall x, y \in \mathbf{R}. \quad (1)$$

According to this 2D real-valued representation (x, y) of \mathbf{z} , \mathbf{C} is a Riemannian manifold [Boothby, 1986]. The distance induced by the canonical Riemannian metric is:

$$d(\mathbf{z}_1, \mathbf{z}_2) = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}, \quad (2)$$

the common Euclidean distance in the 2D complex plane.

Polar form of complex numbers. Any non-zero complex number can be uniquely represented in the polar form, in terms of its magnitude and phase.

Definition 1. For $\forall \mathbf{z} \in \mathbf{C}$ and $\mathbf{z} \neq \mathbf{0}$, its polar form is:

$$\mathbf{z} = |\mathbf{z}| \exp(i \angle \mathbf{z}) \quad (3)$$

$$\mathbf{z} \in \tilde{\mathbf{C}} = \mathbf{C} \setminus \{0 + i0\} \quad (4)$$

$$\text{magnitude: } |\mathbf{z}| = \sqrt{x^2 + y^2} \quad (5)$$

$$\text{phase: } \angle \mathbf{z} = \arctan(y, x) \quad (6)$$

where \exp is the exponential function and \arctan is the 2-argument arc-tangent function that gives the angle in the complex plane between the positive x axis and the line from the origin to the point (x, y) .

Scaling-Rotation product manifold for $\tilde{\mathbf{C}}$. Based on the polar form, we identify the non-zero complex plane $\tilde{\mathbf{C}}$ as the product space of non-zero scaling and 2D rotations

$$\tilde{\mathbf{C}} \iff \mathbf{R}^+ \times \text{SO}(2) \quad (7)$$

where \mathbf{R}^+ is the manifold of positive reals and $\text{SO}(2)$ is the manifold of planar rotations – a rotation Lie group.

We define a bijective mapping F that can go back and forth from the complex plane $\tilde{\mathbf{C}}$ to the manifold space $\mathbf{R}^+ \times \text{SO}(2)$:

$$\mathbf{z} = |\mathbf{z}| \exp(i \angle \mathbf{z}) \xrightleftharpoons[F^{-1}]{F} (|\mathbf{z}|, R(\angle \mathbf{z})) \quad (8)$$

$$R(\angle \mathbf{z}) = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}. \quad (9)$$

Both spaces are parameterized by magnitude and phase; the phase is turned into a complex number with $\exp(\cdot)$ for $\tilde{\mathbf{C}}$ and into a 2D rotation matrix with $R(\cdot)$ for $\mathbf{R}^+ \times \text{SO}(2)$.

Manifold distance for $\tilde{\mathbf{C}}$. The exponential and logarithmic maps are respectively \exp and \log for \mathbf{R}^+ , matrix exponential expm and matrix logarithm logm for $\text{SO}(2)$.

Definition 2. The matrix exponential and logarithm of matrix X are defined respectively as:

$$\text{expm}(X) = \sum_{n=0}^{\infty} \frac{X^n}{n!}$$

$$X = \text{logm}(Y) \text{ if and only if } Y = \text{expm}(X).$$

The distance on this product manifold in Eqn. (2) becomes:

$$d(\mathbf{z}_1, \mathbf{z}_2) = \sqrt{\log^2 \frac{|\mathbf{z}_2|}{|\mathbf{z}_1|} + \left\| \text{logm} \left(R(\angle \mathbf{z}_2) R(\angle \mathbf{z}_1)^{-1} \right) \right\|^2}. \quad (10)$$

Scaling-rotation is transitive on $\tilde{\mathbf{C}}$. The complex plane $\tilde{\mathbf{C}}$ as identified with $\mathbf{R}^+ \times \text{SO}(2)$ is a Riemannian homogeneous space [Helgason, 1962]. We define transitive actions that move a point around on the manifold [Dummit and Foote, 2004].

Definition 3. Given a (Riemannian) manifold \mathcal{M} and a group G with identity element e , we say that G acts on \mathcal{M} (from the left) if there exists a mapping $L : G \times \mathcal{M} \rightarrow \mathcal{M}$ given by $(g, X) \mapsto g.X$ that satisfies two conditions:

$$1) \text{ Identity: } L(e, X) = e.X = X$$

$$2) \text{ Compatibility: } (gh).X = g.(h.X), \quad \forall g, h \in G.$$

An action is called transitive if and only if given $X, Y \in \mathcal{M}$, there exists an element $g \in G$, such that $Y = g.X$.

It is straightforward to verify that scaling and rotation in $\mathbf{R}^+ \times \text{SO}(2)$ satisfies the identity and compatibility conditions on $\tilde{\mathbf{C}}$. It is also a transitive group action: For any complex numbers $\mathbf{z}_1, \mathbf{z}_2 \in \tilde{\mathbf{C}}$, there always exists a relative scaling (of the magnitude) and rotation (of the phase) that maps \mathbf{z}_1 to \mathbf{z}_2 .

Proposition 1. The scaling-rotation Lie group $\mathbf{R}^+ \times \text{SO}(2)$ transitively acts on $\tilde{\mathbf{C}}$ and the action g to take \mathbf{z}_1 to \mathbf{z}_2 is:

$$g = \left(\frac{|\mathbf{z}_2|}{|\mathbf{z}_1|}, R(\angle \mathbf{z}_2) R(\angle \mathbf{z}_1)^{-1} \right) \in \mathbf{R}^+ \times \text{SO}(2). \quad (11)$$

Scaling-rotation is isometric on $\tilde{\mathbf{C}}$. We now show that scaling and rotation actions preserve our manifold distance.

Proposition 2. The scaling and rotation Lie group is isometric on the complex plane $\tilde{\mathbf{C}}$: $\forall \mathbf{z}_1, \mathbf{z}_2 \in \tilde{\mathbf{C}}, g \in \mathbf{R}^+ \times \text{SO}(2)$.

$$d(g.\mathbf{z}_1, g.\mathbf{z}_2) = d(\mathbf{z}_1, \mathbf{z}_2) \quad (12)$$

where d is the manifold distance defined in Eqn. (10).

Proof. We use the definition of d and the property that the 2D rotation group $\text{SO}(2)$ is Abelian: $\forall A, B \in \text{SO}(2), AB = BA$. Let $g = (r, A) \in \mathbf{R}^+ \times \text{SO}(2)$. We have:

$$\begin{aligned} & d(g.\mathbf{z}_1, g.\mathbf{z}_2) \\ &= \sqrt{\log^2 \frac{r|\mathbf{z}_2|}{r|\mathbf{z}_1|} + \left\| \text{logm} \left(AR(\angle \mathbf{z}_2) (AR(\angle \mathbf{z}_1))^{-1} \right) \right\|^2} \\ &= \sqrt{\log^2 \frac{|\mathbf{z}_2|}{|\mathbf{z}_1|} + \left\| \text{logm} \left(R(\angle \mathbf{z}_2) A (R(\angle \mathbf{z}_1) A)^{-1} \right) \right\|^2} \\ &= \sqrt{\log^2 \frac{|\mathbf{z}_2|}{|\mathbf{z}_1|} + \left\| \text{logm} \left(R(\angle \mathbf{z}_2) R(\angle \mathbf{z}_1)^{-1} \right) \right\|^2} = d(\mathbf{z}_1, \mathbf{z}_2), \end{aligned}$$

completing the proof. \square

IV. CNN LAYER FUNCTIONS ON COMPLEX MANIFOLD

We propose complex-valued CNN layer functions based on the scaling-rotation manifold view of complex numbers. Each

layer function transforms the representation with a certain property (e.g. equivariance or invariance) on the manifold.

We define first a convolution operator on the manifold that is *equivariant* to complex-valued scaling. We then define nonlinear activation functions and fully-connected layer functions that are *invariant* to complex-valued scaling. A CNN composed with such layer functions would then be intrinsically invariant to complex-valued scaling.

A. Complex-Valued Convolutional Layer Function

The standard convolution, denoted by $*$ and defined by n weights $\{w_k\}$ over n neighbouring points $\{x_k\}$, is simply the weighted average of real numbers in the Euclidean space:

$$\{w_k\} * \{x_k\} = \sum_{k=1}^n w_k x_k. \quad (13)$$

We extend this concept to points on a manifold.

Fréchet mean on the manifold. The weighted average of n points on a Riemannian manifold is called the *weighted Fréchet mean* (wFM) [Fréchet, 1948]. We define the complex convolution, denoted by $\tilde{*}$, as wFM on the scaling-rotation manifold for complex values (center circles in Fig. 2):

$$\{w_k\} \tilde{*} \{\mathbf{z}_k\} = \text{wFM}(\{w_k\}, \{\mathbf{z}_k\}) \quad (14)$$

$$\text{wFM}(\{w_k\}, \{\mathbf{z}_k\}) := \arg \min_{\mathbf{m} \in \tilde{\mathbf{C}}} \sum_{k=1}^n w_k d^2(\mathbf{z}_k, \mathbf{m}) \quad (15)$$

$$\sum_{k=1}^n w_k = 1, \quad w_k \geq 0 \quad (16)$$

where d is the manifold distance in Eq. (10).

We contrast our complex-valued wFM convolution $\tilde{*}$ with the standard real-valued convolution $*$.

- While the output of $\tilde{*}$ is complex-valued, the weights $\{w_k\}$ are real-valued, just like the weights for $*$.
- While the weights of $*$ can be arbitrary, the weights of $\tilde{*}$ are all nonnegative and summed up to 1. This convexity constraint ensures that the wFM of n points on a manifold stays on the manifold.
- While the output of $*$ is simply the weighted average, the output of $\tilde{*}$ is the minimizer to a weighted least squares problem, i.e., the data mean that minimizes the weighted variance. There is no closed-form solution to wFM; however, there is a provably convergent n -step iterative solution [Salehian et al., 2015].
- If d is the manifold distance in Eq. (2) for the Euclidean space which is also Riemannian, then wFM has exactly the weighted average as its closed-form solution. That is, our wFM convolution on the Euclidean manifold is reduced to the standard convolution, although with the additional convexity constraint on the weights.

Convolutional wFM layer. As for the standard convolution, our weights $\{w_k\}$ for wFM are parameters learnable through stochastic gradient descent (SGD), with the additional convexity constraint on $\{w_k\}$.

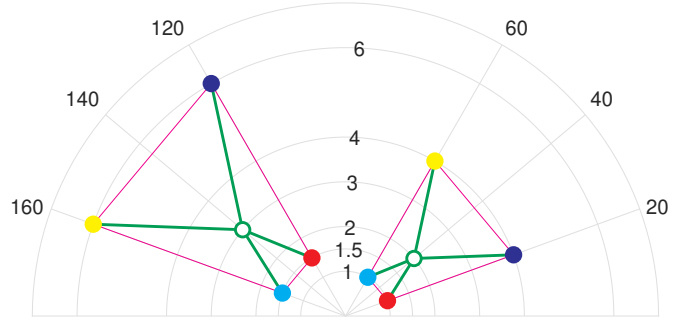


Fig. 2. Our complex-valued convolution in terms of wFM on the $\mathbf{R}^+ \times \text{SO}(2)$ manifold is equivariant to complex-valued scaling. Consider 4 numbers in the complex plane, marked by 4 colored points on a small magenta trapezoid. Their equally weighted wFM (marked by the green circle) sits inside the trapezoid at the geometric mean of their magnitudes and the mean of their phases. When the 4 complex numbers are multiplied by $1.5 \exp(i \frac{100\pi}{180})$, the points are scaled by 1.5 and rotated by 100° , moving to the larger trapezoid. The new wFM is simply the old wFM transported by the same movement.

Each set of weights $\{w_k\}$ defines a single wFM channel in our convolutional layer, and each layer has multiple channels. We follow the CNN convolution convention for images, where the convolutional kernel spans a local spatial neighbourhood but all the channels. If there are 10 input channels of 50×50 pixels, to produce 20 output channels with 5×5 spatial kernels, we need to learn 20 sets of $5 \times 5 \times 10$ weights.

Equivariance of wFM to $\mathbf{R}^+ \times \text{SO}(2)$. We have shown that the scaling-rotation Lie group transitively acts on $\tilde{\mathbf{C}}$ and is isometric (Prop. 2). We use this result to prove that our wFM is equivariant to complex-valued scaling.

Proposition 3. *The complex-valued convolution $\tilde{*}$ in Eq. (14) is equivariant to the action of $\mathbf{R}^+ \times \text{SO}(2)$: $\forall g \in \mathbf{R}^+ \times \text{SO}(2)$,*

$$\{w_k\} \tilde{*} \{g \cdot \mathbf{z}_k\} = g \cdot (\{w_k\} \tilde{*} \{\mathbf{z}_k\}). \quad (17)$$

Proof. Let $g \in \mathbf{R}^+ \times \text{SO}(2)$ and $\mathbf{o} = \{w_k\} \tilde{*} \{\mathbf{z}_k\}$. $\forall \mathbf{m} \in \tilde{\mathbf{C}}$,

$$\sum_{k=1}^n w_k d^2(g \cdot \mathbf{z}_k, g \cdot \mathbf{o}) = \sum_{k=1}^n w_k d^2(\mathbf{z}_k, \mathbf{o}) \leq \sum_{k=1}^n w_k d^2(\mathbf{z}_k, \mathbf{m})$$

since g preserves the distance and \mathbf{o} is the minimizer over $\{\mathbf{z}_k\}$. Therefore, $g \cdot \mathbf{o}$ is the minimizer over $\{g \cdot \mathbf{z}_k\}$. \square

Fig. 2 illustrates how wFM is equivariant to rotation and scaling. For each trapezoid, the center circle marks the wFM of the four corner points. If the trapezoid is transported using a particular scaling and rotation action, then the center wFM is also transported by the same action.

B. Nonlinear Activation Functions

The wFM convolution is a contractive mapping, an effect of a nonlinear activation function. Nevertheless, for stronger nonlinearity and acceleration in optimization during learning, we propose two activation functions from the manifold perspective: tangent ReLU and G -transport.

Tangent ReLU (tReLU). The tangent space of a manifold is a vector space that contains the possible directions for tangentially passing through a point on the manifold. It could

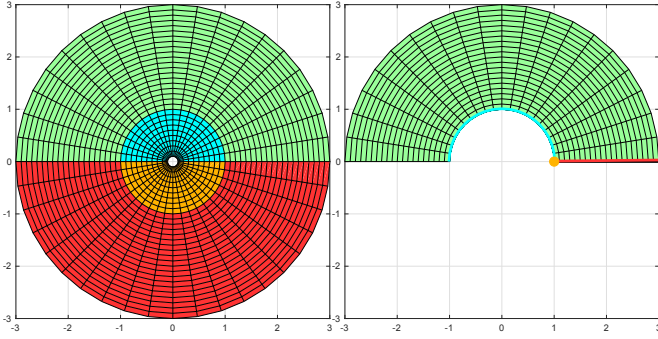


Fig. 3. Our tangent ReLU extends ReLU from the real line to the complex plane $\tilde{\mathbf{C}}$, by applying ReLU in the tangent space of $\mathbf{R}^+ \times \text{SO}(2)$. It rectifies the magnitude by 1 and phase by 0, creating four corresponding regions before (left) and after (right) the mapping. **Green region:** Points retain both their magnitudes and phases. **Cyan region:** Points retain their phases with their magnitudes rectified to 1. **Red region:** Points retain their magnitudes with their phases rectified to 0. **Brown region:** Points are rectified in both the magnitude and the phase, all to the same point $1 + i \cdot 0$.

be regarded as a local Euclidean approximation of the manifold. A pair of logarithmic and exponential maps establish the correspondence between the manifold and the tangent space.

We extend ReLU to the complex plane $\tilde{\mathbf{C}}$ by applying ReLU in the tangent space of $\mathbf{R}^+ \times \text{SO}(2)$ manifold. Our tangent ReLU is composed of three steps.

- 1) Apply logarithmic maps to go from a point in $\tilde{\mathbf{C}}$ to a point in its tangent space. The mapping is \log for $r \in \mathbf{R}^+$, and logm for $R(\theta) \in \text{SO}(2)$, which produces a skew symmetric matrix. We choose the principal log map for $\text{SO}(2)$ in the range of $\theta \in (-\pi, \pi]$:

$$\text{logm}(R(\theta)) = \theta \cdot \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \quad (18)$$

- 2) Apply ReLU in the tangent space of $\tilde{\mathbf{C}}$. ReLU is well defined for a real-valued scalar such as $\log(r)$ for \mathbf{R}^+ :

$$\text{ReLU}(x) = \max(x, 0), \quad \text{e.g., } x = \log(r). \quad (19)$$

We extend ReLU to $\text{logm}(R(\theta))$ for $\text{SO}(2)$, since it is just the real-valued θ scaled by a constant matrix:

$$\text{ReLU}(\text{logm}(R(\theta))) = \max(\theta, 0) \cdot \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \quad (20)$$

- 3) Apply exponential maps to come back to $\tilde{\mathbf{C}}$ from the tangent space. We can simplify the 3-step tReLU as:

$$r \xrightarrow{\text{ReLU}} \exp(\text{ReLU}(\log(r))) = \max(r, 1) \quad (21)$$

$$R(\theta) \xrightarrow{\text{ReLU}} \text{expm}(\text{ReLU}(\text{logm}(R(\theta)))) = R(\max(\theta, 0)) \quad (22)$$

Fig. 3 shows that our manifold perspective of $\tilde{\mathbf{C}}$ leads to a non-trivial extension of ReLU from the real line to the complex plane, partitioning $\tilde{\mathbf{C}}$ into four regions, separated by $r = 1$ and $\theta = 0$: Those with magnitudes smaller than 1 are rectified to 1, and those with negative phases are rectified to 0.

G-transport. A nonlinear activation function is a general mapping that transforms the range of responses. We consider a general alternative which simply transports all the values in a feature channel via an action in the group $G = \mathbf{R}^+ \times \text{SO}(2)$. We

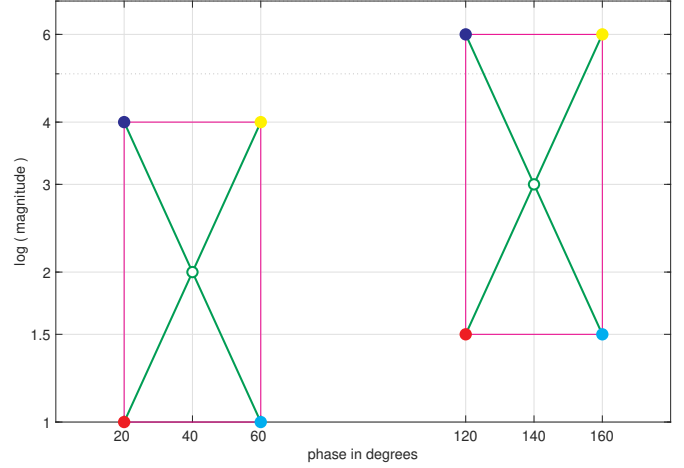


Fig. 4. Our distance transform is invariant to complex-valued scaling. We plot Fig. 2 in the tangent space of $\mathbf{R}^+ \times \text{SO}(2)$, with magnitude on a log scale for the y -axis and phase for the x -axis. The distance d on the manifold is directly measured by the Euclidean distance in this space. When the 4 complex numbers are multiplied by $1.5 \exp(i \frac{100\pi}{180})$, the points as well as their wFM are simply translated. The distances between the points and their wFM thus remain the same. Note: While this direct phase θ representation is intuitive, our rotation matrix $R(\theta)$ can more easily handle phase representation discontinuity at e.g. $\pm\pi$, where $\theta(-\pi, +\pi]$ for a unique determination of θ .

only need to learn one global scaling and rotation per feature channel, which corresponds to learning one complex-valued multiplier per channel at a certain depth layer.

C. Fully-Connected Layer Functions

For classification tasks, having equivariance of convolution and range compression of nonlinear activation functions are not enough; we need the final representation of a CNN to be invariant to variations within each class.

In a standard CNN classifier, the entire network is invariant to the action of the group of translations, achieved by the fully connected (FC) layer. Likewise, we develop a FC layer function on $\tilde{\mathbf{C}}$ that is invariant to the action of $\mathbf{R}^+ \times \text{SO}(2)$.

Since the manifold distance d is invariant to actions in $G = \mathbf{R}^+ \times \text{SO}(2)$, we propose the distance between each point of a set and their weighted Fréchet mean, which is equivariant to G , as a new FC function on $\tilde{\mathbf{C}}$.

Distance transform FC Layer. Let the input be s pixels of c channels each. We perform a global integration over all these $s \cdot c$ complex values $\{\mathbf{t}_k\}$. Given $s \cdot c$ weights $\{w_k\}$ for these individual numbers, we first calculate their wFM \mathbf{m} and then compute the distance u_k from \mathbf{t}_k to the mean \mathbf{m} :

$$\mathbf{m} = \{w_k\} \tilde{*} \{\mathbf{t}_k\} \quad (23)$$

$$u_k = d(\mathbf{t}_k, \mathbf{m}). \quad (24)$$

The output is real-valued and of the same size as the input. The weights $\{w_k\}$ are the parameters to be learned and there could be multiple sets of such weights at this layer.

Proposition 4. *The distance to the wFM, defined in Eq. (23) and Eq. (24), is invariant to the action of $G = \mathbf{R}^+ \times \text{SO}(2)$.*

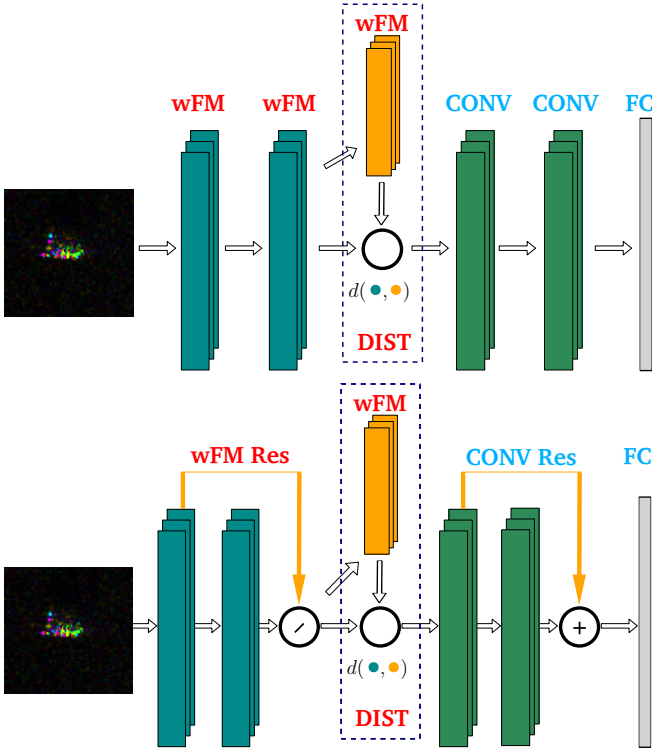


Fig. 5. Our complex-valued CNN with residual connections allows manifold valued representations to be combined across different depths. **Top** A sample CNN with complex-valued convolution (wFM), distance transform (DIST), real-valued convolution (CNN), and fully connected (FC) layers. **Bottom** To create a residual structure (Res), we add skip connections between two adjacent convolutional layers. We first spatially align their feature maps via convolution (wFM for complex-valued and CONV for real-valued) and then combine them via channel-wise concatenation for complex-valued data and addition for the real-valued data. The only difference between real-valued and complex-valued residual blocks is how the feature maps are combined: Addition is a vector space operator that does not apply to two points on a non-Euclidean manifold.

Proof. Per Propositions 2 and 3, $\forall g \in G$, we have:

$$\begin{aligned} & d(g \cdot \mathbf{t}_k, \{w_k\} \tilde{*} \{g \cdot \mathbf{t}_k\}) \\ &= d(g \cdot \mathbf{t}_k, g \cdot (\{w_i\} \tilde{*} \{\mathbf{t}_k\})) \quad \text{equivariance of wFM} \\ &= d(\mathbf{t}_k, \{w_k\} \tilde{*} \{\mathbf{t}_k\}) \quad \text{invariance of distance} \end{aligned}$$

completing the proof. \square

Fig. 4 re-plots Fig. 2 in the $(\log(r), \theta)$ space, which corresponds to the tangent space of $\mathbf{R}^+ \times \text{SO}(2)$ where the manifold distance can be directly visualized as the Euclidean distance. When the four corners of the trapezoid are scaled and rotated, the trapezoid is simply translated along $(\log(r), \theta)$ axes. The distance from points to their wFM remain the same.

Since the output of the distance transform layer is real-valued, we can subsequently use any existing layer functions of a real-valued CNN classifier. Fig.1 shows a sample architecture of our complex-valued CNN, where two successive wFM convolutional layers are followed by a distance transform FC layer, a standard convolutional layer, and an FC layer for final softmax classification.

D. Complex-Valued Residual Layer Function

A standard CNN with residual layers such as ResNet [He et al., 2016] outperforms the one without. Residual layers are useful for preventing exploding/vanishing gradients in deep networks, by utilizing skip connections to jump over some layers. The skip connections between layers add the outputs from previous layers to the outputs of stacked layers.

While addition is natural for combining layers in the field of real numbers, it does not make sense in the field of complex numbers: We can add two vectors in the Euclidean space, but we cannot add two points on a non-Euclidean manifold.

Here we propose a complex-valued residual layer function by retaining the skip connection concept without the addition to combine the outputs from different layers. Let a feature layer $\mathbf{f}(s, c)$ be specified by the number of pixels s and the number of channels c . Consider two feature layers $\mathbf{f}_1(s_1, c_1)$, $\mathbf{f}_2(s_2, c_2)$, $s_1 < s_2$, with one layer through skip connections. In order to combine them, we first use the wFM convolution to bring the spatial dimension of \mathbf{f}_2 from s_2 to s_1 and then concatenate the two sets of spatially aligned features:

$$\text{align spatially: } \mathbf{f}_2(s_2, c_2) \xrightarrow{\tilde{*}} \bar{\mathbf{f}}_2(s_1, c_2) \quad (25)$$

$$\text{concatenate: } [\mathbf{f}_1(s_1, c_1) | \bar{\mathbf{f}}_2(s_1, c_2)] \rightarrow \mathbf{f}(s_1, c_1 + c_2) \quad (26)$$

Once combined, we can treat them as the input and apply any wFM convolution as desired.

Fig. 5 shows that we can simply replace two complex-valued convolution layers with such a residual block connecting and combining their outputs, and build a residual complex-valued convolution network. The only difference with the real-valued residual block is that the combination is channel-wise concatenation for a non-Euclidean manifold instead of addition for a vector space.

We can optionally further reduce the number of parameters for convolution using the tensor ring decomposition [Zhao et al., 2016]. A c -dimensional convolutional filter W of size $n_1 \times \dots \times n_c$ can be decomposed into c smaller rank b tensors, each of the form T_k with size $b \times n_k \times b$ such that $\forall k_1, \dots, k_c$,

$$W(k_1, \dots, k_c) = \text{trace}(T_1(:, k_1, :) \times \dots \times T_c(:, k_c, :)) \quad (27)$$

where \times denotes matrix multiplication. Such a tensor factorization needs $b^2 \sum_{k=1}^c n_k$ parameters for all the tensors $\{T_k\}$ instead of $\prod_{k=1}^c n_k$ parameters for the original W . Tensor ring decomposition can achieve arbitrary approximation precision (Theorem 2.2 on pp. 2299 in [Oseledets, 2011]).

V. EXPERIMENTS

We compare our SurReal complex-valued classifier against two baselines. **1)** The first baseline is a real-valued CNN classifier such as ResNet50 which ignores the geometry of complex numbers and treats each complex value as two independent real numbers. **2)** The second baseline is the deep complex networks (DCN) which extends real-valued CNN layer functions to the complex domain by the form of the functions such as complex-valued convolution, batch-normalization, nonlinear activation, and weight initialization [Trabelsi et al., 2017].

TABLE I
SURREAL CNN IN DETAILED LAYER SPECIFICATION

Layer Type	Input Shape	Kernel	Stride	Output Shape
wFM CONV	[2, 1, 100, 100]	5×5	2	[2, 20, 48, 48]
G-transport	[2, 20, 48, 48]	-	-	[2, 20, 48, 48]
wFM CONV	[2, 20, 48, 48]	5×5	2	[2, 20, 22, 22]
G-transport	[2, 20, 22, 22]	-	-	[2, 20, 22, 22]
DIST	[2, 20, 22, 22]	-	-	[20, 22, 22]
CONV	[20, 22, 22]	5×5	1	[30, 18, 18]
BN+ReLU	[30, 18, 18]	-	-	[30, 18, 18]
MaxPool	[30, 18, 18]	2×2	2	[30, 9, 9]
CONV	[30, 18, 18]	5×5	3	[30, 2, 2]
BN+ReLU	[30, 2, 2]	-	-	[30, 2, 2]
CONV	[30, 2, 2]	2×2	1	[30, 1, 1]
BN+ReLU	[30, 1, 1]	-	-	[30, 1, 1]
FC	[30]	-	-	[50]
FC	[50]	-	-	[11]

The network has two layers of complex-valued convolution (wFM) and nonlinear activation (*G*-transport), distance transform (DIST), and then a real-valued CNN classifier composed with standard convolution (CONV), batch normalization (BN) and ReLU, max pooling (MaxPool), and fully connected (FC) layers. An input or output shape of 4 dimensions indicates complex-valued data, with each complex number represented by two values: magnitude and phase. For example, [2, 1, 100, 100] means a complex-valued 1-channel 100×100 image. A shape of 3 dimensions indicates real-valued data. In our network, DIST is the depth layer that turns a complex-valued representation into a real-valued one; it separates the complex-valued and real-valued representations in the SurReal CNN.

The specific DCN models used in [Trabelsi et al., 2017] are very large, on the order of one million parameters. To help directly compare complex-valued layer functions, we adopt the same SurReal Residual architecture for the DCN baseline, but replacing all the convolution layers (complex-valued wFM and real-valued convolution) and nonlinear activation functions with DCN’s proposed counterparts.

We experiment on two complex-valued datasets: SAR image dataset MSTAR [Keydel et al., 1996] and synthetic RF signal dataset RadioML [O’Shea et al., 2016]. All the models are trained on a GeForce RTX 2080 GPU for 120 epochs, using Adam optimizer and cross-entropy loss.

A. MSTAR Classification

MSTAR dataset. There are a total of 15,716 complex-valued X-band SAR images, distributed unevenly over 11 classes: The first 10 classes contain different target vehicles and the last 1 class contains background clutter. See Table III for the total number of images per class. We take the 100×100 center crop of each image and convert the complex value of each pixel into the polar form.

Real-valued CNN baseline. ResNet50 [He et al., 2016] is widely successful on real-valued image classification and it is also used as a baseline in [Shao et al., 2018].

Two SurReal CNN architectures. Table I lists detailed layer specification of our basic model. Table II adds residual con-

TABLE II
SURREAL RESIDUAL CNN IN DETAILED LAYER SPECIFICATION

Layer Type	Input Shape 1	Input Shape 2	Kernel	Stride	Output Shape
wFM CONV	[2, 1, 100, 100]	-	5×5	2	[2, 20, 48, 48]
G-transport	[2, 20, 48, 48]	-	-	-	[2, 20, 48, 48]
wFM CONV	[2, 20, 48, 48]	-	5×5	2	[2, 20, 22, 22]
wFM Res	[2, 20, 48, 48]	[2, 20, 22, 22]	5×5	2	[2, 20, 22, 22]
G-transport	[2, 20, 22, 22]	-	-	-	[2, 20, 22, 22]
DIST	[2, 20, 22, 22]	-	-	-	[20, 22, 22]
CONV	[20, 22, 22]	-	5×5	1	[30, 18, 18]
BN+ReLU	[30, 18, 18]	-	-	-	[30, 18, 18]
CONV Stack	[30, 18, 18]	-	1, 3, 1	1	[30, 18, 18]
CONV Res	[30, 18, 18]	[30, 18, 18]	-	-	[40, 18, 18]
Maxpool	[40, 18, 18]	-	2×2	2	[40, 9, 9]
CONV	[40, 18, 18]	-	5×5	3	[50, 2, 2]
BN+ReLU	[50, 2, 2]	-	-	-	[50, 2, 2]
CONV Stack	[50, 2, 2]	-	1, 3, 1	1	[50, 2, 2]
CONV Res	[50, 2, 2]	[50, 2, 2]	-	-	[60, 2, 2]
CONV	[60, 2, 2]	-	2×2	1	[70, 1, 1]
BN+ReLU	[70, 1, 1]	-	-	-	[70, 1, 1]
FC	[70]	-	-	-	[30]
FC	[30]	-	-	-	[11]

Our SurReal residual CNN utilizes both complex-valued residual blocks (Fig. 5) as well as real-valued residual blocks. In CONV Stack, we stack three convolutions with 1×1 , 3×3 , 1×1 kernels respectively. We zero pad the inputs on the 3×3 convolution to preserve the spatial dimensions.

TABLE III
MSTAR DATASET SIZE: NUMBER OF IMAGES PER CLASS

C0	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10
1285	429	6694	451	1164	1415	573	572	573	1401	1159

nections. Both models have two complex-valued convolutions (with nonlinear activation), one distance transform layer, and two real-valued convolutions (with batch normalization and ReLU), max pooling, and two fully connected layers. While we have listed *G*-transport in the two tables, we have also tried tReLU as the complex-valued nonlinear activation function; their difference is insignificant in initial experiments on the MSTAR dataset, and we focus on *G*-transport for its simplicity.

Model size comparison. While ResNet50 and DCN have 23 million and 155K parameters respectively, our SurReal CNN has 67K parameters and SurReal residual CNN has 109K parameters. We can further reduce the parameter count by implementing convolutions with tensor ring decomposition [Oseledets, 2011, Zhao et al., 2016]. Fig. 6 plots these model sizes on a log scale. The saving is substantial: our SurReal CNN is less than 0.1% of the real-valued baseline and 44% of the complex-valued baseline.

Task 1: 10-class target recognition. For the 10 target classes, we split all the data in 5 varying proportions of 1%, 5%, 10%, 20%, 30% for training and the rest for testing. Fig. 7

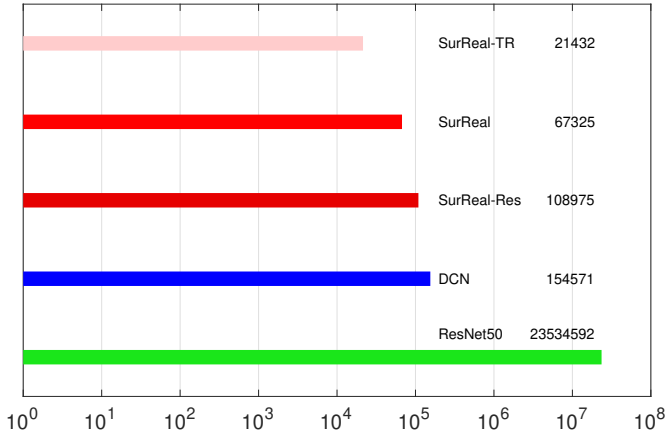


Fig. 6. Our SurReal models for MSTAR are much leaner than baselines. Each model size is plotted as a horizontal bar on a log scale, labeled with the model name and the number of parameters on the right. ResNet50 is the real-valued baseline and the largest with 23.5M parameters. DCN is the complex-valued baseline, with 154K parameters at 0.7% of ResNet50. Our SurReal residual network has 109K parameters at 0.5% (71%) of ResNet50 (DCN). Our basic SurReal CNN has 67K parameters at 0.3% (44%) of ResNet50 (DCN). With the tensor ring implementation for convolutions, our SurReal CNN could be further reduced to 21K parameters at 0.1% (14%) of ResNet50 (DCN).

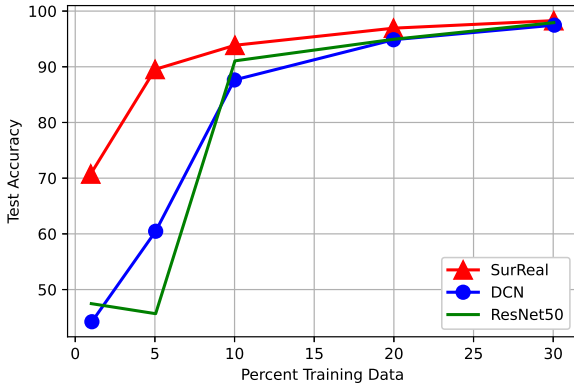


Fig. 7. Our SurReal CNN significantly outperforms real-valued and complex-valued baselines on MSTAR target recognition, when less data is used for training. We split the data into training and test sets at varying proportions and plot the test accuracy. When there is enough training data, e.g., at 30% training and 70% testing, all three models perform similarly with 98% accuracy. When there is less training data, e.g., at 5% training and 95% testing, SurReal is much better with a test accuracy of 90% over DCN’s 60% and ResNet50’s 45%, demonstrating the effectiveness of our complex-valued model.

shows that our SurReal significantly outperforms DCN and ResNet50, especially when a small percentage of training data is used. At 5% training and 95% testing, the accuracy is 90% for SurReal, 60% for DCN, and 45% for ResNet50.

Task 2: 11-class classification. We also include the remaining clutter class which contains miscellaneous background images. We create two random subsets, large (L) and small (S), and the small set of 6,295 images are contained entirely in the large set of 12,610 images. Fig. 8 shows the number of instances across 11 classes and in training/testing splits.

Table IV shows that all the models perform at a high accuracy of 99% for the large dataset. The performance drops

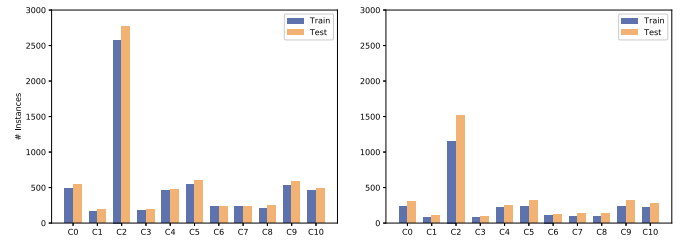


Fig. 8. MSTAR large (L) and small (S) subsets have highly imbalanced classes. C2 is the largest class, C0, C4, C5, C9, C10 come next at about 20% of the size of C2, and C1, C3, C6, C7, C8 are at about 8% of the size of C2.

TABLE IV
MODEL ACCURACY COMPARISON ON MSTAR-L AND MSTAR-S

Test Accuracy (%)	ResNet50	DCN	SurReal	SurReal-Res
MSTAR-L	99.1	98.9	99.1	99.2
MSTAR-S	97.4	93.3	97.7	98.4
Difference S–L	–1.7	–6.6	–1.4	–0.8

Our SurReal CNN and its residual version consistently outperform the real-valued baseline ResNet50 and complex-valued baseline DCN. SurReal-Res retains the accuracy the most when the overall data size is reduced, demonstrating better generalization capability from small data.

as the overall data size is reduced by half in the small dataset, but the drop is the least at -0.8% for SurReal-Res, followed by -1.4% for SurReal, -1.7% for ResNet50, and the most at -6.6% for DCN. Our results against the two baselines suggest that it is both the residual connections and more importantly how we handle complex-valued data that delivers more generalizing performance from smaller training data.

Fig. 10 shows sample channel responses from our SurReal CNN on MSTAR-S images. With two complex-valued wFM convolutions, followed by distance transform and real-valued convolution, the representation for each input SAR image quickly becomes more distinctive across classes, facilitating accurate discrimination.

Fig. 9 shows the confusion matrix between classes on MSTAR-S. In general, the more training instances in the class, the least confusion with other classes at the test time. However, despite the significant class imbalance, the performance gap is small between the minority and majority classes. Residual connections help clear up more confusion.

B. RadioML Classification

RadioML Dataset. They are synthetically generated radio signals with modulation operating over both voice and text data. Noise is added further for channel effects. Each signal has 128 time samples and is tagged with a signal-to-noise ratio (SNR), in the range of $[-20, 18]$ with an increment step of 2. There are 11 modulation modes, of which BPSK, QPSK, 8PSK, 16QAM, 64QAM, BFSK, CPFSK, PAM4 are digital modulations, and WB-FM, AM-SSB, and AM-DSB are analog modulations. There are 20,000 instances per modulation. See sample instances in Fig. 11. The data is split 50/50 between training and testing.

Real-valued CNN Baseline. We use O’Shea’s model [O’Shea

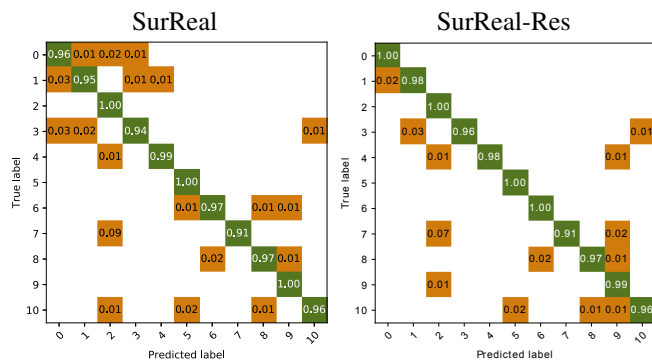


Fig. 9. The confusion matrix for our SurReal CNN (left) and SurReal-Res CNN (right) on MSTAR-S. 1) The dataset is small and highly imbalanced across classes (Fig. 8 Right). The model achieves better accuracy for classes with more instances: The accuracy is 100% for C2 and 94% for C1 and C3. Overall, the accuracy gap is small, considering the size of C1 and C3 is only 8% of the size of C2. 2) The residual connections help further clear up the confusion between classes. The matrix becomes more strongly diagonal.

et al., 2016] and feed the 1-channel complex-valued RF signal as a (real,imaginary) two-channel signal.

Model size comparison. We follow the architecture of DCN and SurReal for images and adapt the spatial dimensions to fit the 1×128 RF signals. Fig. 12 plots these model sizes on a log scale. Our SurReal CNN is 0.7% of the real-valued baseline and 58% of the complex-valued baseline.

Accuracy over SNR. Fig. 13 compares the test accuracy at various SNR levels. Our SurReal underperforms the baselines at lower SNRs and outperforms the real-valued O’Shea baseline at higher SNRs. For example at SNR 10, it achieves 76.1%, surpassing O’Shea’s 72.7% with 0.7% of its model size and on par with DCN’s 76.3% at 58% of its size.

VI. SUMMARY AND CONCLUSIONS

Deep learning is widely adopted in machine learning and computer vision. Most existing deep learning techniques are developed for data in a vector space. However, practical data often have correlations between channels and are better modeled as points a manifold.

While the Nash embedding theorem [Boothby, 1986] assures us that it is always feasible to embed the data on a manifold into a higher dimensional vector space, it would also result in an increase in the model complexity and training time. Recent geometric deep learning approaches develop tools for spaces with certain geometry such as graphs and surfaces.

We deal with deep learning on complex-valued data, and we approach it from a geometric perspective. The common approach is to represent complex-valued data as two-channel real-valued data and then all the real-valued deep learning tools can be used. However, this representation ignores the underlying geometry that defines the complex-valued data: Complex-valued data containing the same information could be subject to arbitrary complex-valued scaling.

We propose to model the space of complex numbers as a product manifold of non-zero scaling and planar rotations. Arbitrary complex-valued scaling naturally becomes a group

of transitive actions on this manifold. We can subsequently define convolution on the manifold that is *equivariant* to this action group, and define distance transform that is *invariant* to this action group. The manifold perspective also allows us to define new nonlinear activation functions such as tangent ReLU and G -transport, as well as residual connections on the manifold-valued data.

A complex-valued CNN classifier composed of such layer functions has built-in invariance to complex-valued scaling, so that the model only needs to focus on the discrimination between classes. Our experimental results validate our principled approach, and we dub our model *SurReal* based on its high performance achieved at a super-lean model size compared with real-valued or complex-valued baselines.

ACKNOWLEDGEMENTS

This research was supported, in part, by Berkeley Deep Drive and DARPA. An earlier version of a few diagrams was constructed with help from Liu Yang. We thank Utkarsh Singhal for thoughtful proofreading and anonymous reviewers for critical and constructive comments.

REFERENCES

- Md Fajul Amin and Kazuyuki Murase. Single-layered complex-valued neural network for real-valued classification problems. *Neurocomputing*, 72(4-6):945–955, 2009.
- Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pages 1120–1128, 2016.
- Yoshua Bengio et al. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.
- N. Benvenuto and F. Piazza. On the complex backpropagation algorithm. *IEEE Transactions on Signal Processing*, 40(4):967–969, 1992.
- William M Boothby. *An introduction to differentiable manifolds and Riemannian geometry*, volume 120. Academic press, 1986.
- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- Joan Bruna and Stephane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- Joan Bruna, Soumith Chintala, Yann LeCun, Serkan Piantino, Arthur Szlam, and Mark Tygert. A theoretical argument for complex-valued convolutional networks. In *arXiv:1503.03438*, 2015.
- Kerstin Bunte, Frank-Michael Schleif, and Michael Biehl. Adaptive learning for complex-valued data. In *ESANN*. Citeseer, 2012.
- Charles F Cadieu and Bruno A Olshausen. Learning intermediate-level representations of form and motion from natural movies. *Neural computation*, 24(4):827–866, 2012.
- Rudrasis Chakraborty, Monami Banerjee, and Baba C Vemuri. H-cnns: Convolutional neural networks for riemannian homogeneous spaces. *arXiv preprint arXiv:1805.05487*, 2, 2018a.
- Rudrasis Chakraborty, Jose Bouza, Jonathan Manton, and Baba C Vemuri. Manifoldnet: A deep network framework for manifold-valued data. *arXiv preprint arXiv:1809.06211*, 2018b.
- Rudrasis Chakraborty, Chun-Hao Yang, Xingjian Zhen, Monami Banerjee, Derek Archer, David Vaillancourt, Vikas Singh, and Baba Vemuri. A statistical recurrent model on the manifold of symmetric positive definite matrices. In *Advances in Neural Information Processing Systems*, pages 8883–8894, 2018c.




Fig. 10. Sample channel responses of our SurReal CNN on MSTAR-S images. 1) There are 11 rows, with Row k containing an instance for class k . Column 1 shows the input image. Columns 2-4 contain 10 channel responses in the first three convolutional layers: the first two are complex-valued wFM convolution and the last is the real-valued convolution. 2) All the complex values are displayed in fully saturated color using the cyclic HSV colormap, with the magnitude as the intensity and the phase as the hue, whereas the real values are displayed in gray using the grayscale colormap. As the input image goes through the convolutional layers in the SurReal CNN, the distinction between instances in individual classes becomes clearer.

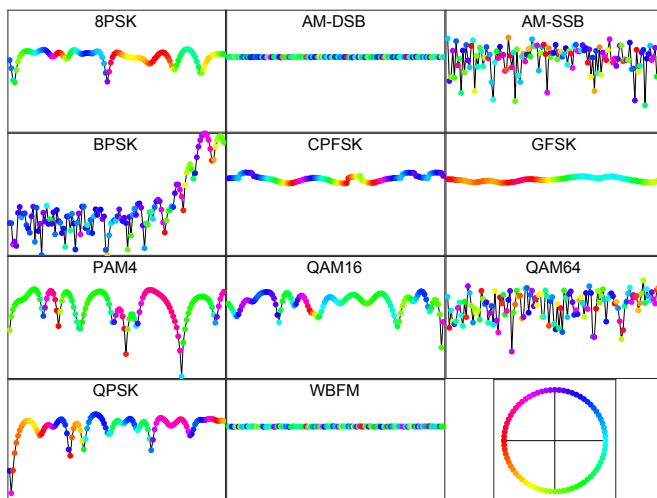


Fig. 11. RadioML sample instances. Each instance is a 128 time-step complex-valued signal at SNR 10. We plot each signal as a line of 128 connected colored dots over time (x -axis): the height of the dot indicates the magnitude on a log scale (y -axis), and the color of the dot indicates the phase with the HSV colormap. The first 11 plots show one instance per class with the class name labeled on the top. The last plot visualizes the complex number $\{e^{i\theta}, \theta \in [-\pi, \pi)\}$ as our colored dots, illustrating how the color varies with the phase θ . The shape of RF signal reflects both the message it is carrying and the modulation mode it is subject to. The classifier must ignore the distinction in the message but focus on the distinction in the modulation.

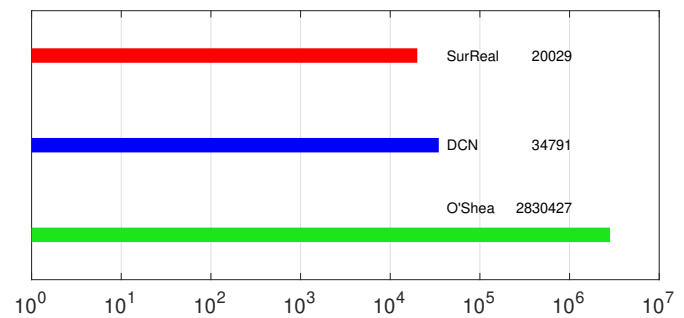


Fig. 12. Our SurReal model for RadioML is much leaner than baselines. Each model size is plotted as a horizontal bar on a log scale, labeled with the model name and the number of parameters on the right. O'Shea's model is the real-valued baseline and the largest with 2.8M parameters. DCN is the complex-valued baseline, with 35K parameters at 1% of O'Shea. Our SurReal CNN has 20K parameters at 0.7% (58%) of O'Shea (DCN).

- Rudrasis Chakraborty, Jiayun Wang, and Stella X. Yu. Surreal: Fréchet mean and distance transform for complex-valued deep learning. In *IEEE Computer Vision and Pattern Recognition Workshop on Perception Beyond the Visible Spectrum*, 2019.
- Xinming Cheng, Jiangnan He, Jianbiao He, and Honglei Xu. Cv-capsnet: Complex-valued capsule network. 2019.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016.
- Taco Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Convolutional networks for spherical signals. *arXiv preprint arXiv:1709.04893*, 2017.

- Sander Dieleman, Kyle W Willett, and Joni Dambre. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly notices of the royal astronomical society*, 450 (2):1441–1459, 2015.
- David Steven Dummit and Richard M Foote. *Abstract algebra*, volume 3. Wiley Hoboken, 2004.
- Carlos Esteves, Christine Allen-Blanchette, Xiaowei Zhou, and Kostas Daniilidis. Polar transformer networks. *arXiv preprint arXiv:1709.01889*, 2017.
- Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. volume 10, pages 215–310. Annales de l'institut Henri Poincaré, France, 1948.
- William T. Freeman and Edward H Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (9):891–906, 1991.
- G.M. Georgiou and C. Koutsougeras. Complex domain backpropagation. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 39, 1992.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,

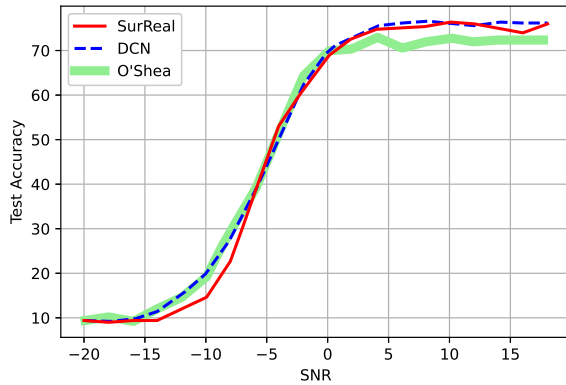


Fig. 13. Test accuracy comparison on RadioML with a varying SNR. All the models perform poorly at a very low SNR. Our SurReal underperforms the two baselines for SNR in $[-15, -5]$. The complex-valued baseline DCN is on par with the real-valued baseline O'Shea at a negative SNR, but better at a positive SNR. Likewise, our SurReal outperforms the real-valued baseline at a positive SNR, at only 0.7% (58%) size of O'Shea (DCN).

pages 770–778, 2016.

S Helgason. Differential geometry and symmetric spaces, acad. Press, New York, 21, 1962.

Akira Hirose and Shotaro Yoshida. Generalization characteristics of complex-valued feedforward neural networks in relation to signal coherence. *IEEE Transactions on Neural Networks and learning systems*, 23(4):541–551, 2012.

Eric R Keydel, Shung Wu Lee, and John T Moore. Mstar extended operating conditions: A tutorial. In *Algorithms for Synthetic Aperture Radar Imagery III*, volume 2757, pages 228–242. International Society for Optics and Photonics, 1996.

M. S. Kim and C. C. Guest. Modification of backpropagation networks for complex-valued signal processing in frequency domain. In *Int. joint conf. on neural networks*, volume 3, pages 27–31, 1990.

Witold Kinsner and Warren Grieder. Amplification of signal features using variance fractal dimension trajectory. *Int. J. Cogn. Inform. Nat. Intell.*, pages 1–17, 2010.

Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. *arXiv preprint arXiv:1802.03690*, 2018.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

Michael Maire, Takuya Narihira, and Stella X Yu. Affinity cnn: Learning pixel-centric pairwise relations for figure/ground embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 174–182, 2016.

Stéphane Mallat. Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150203, 2016.

Tristan Needham. *Visual complex analysis*. Oxford University Press, 1998.

T Nitta. On the critical points of the complex-valued neural network. In *Proceedings of the 9th International Conference on Neural Information Processing, 2002. ICONIP'02.*, volume 3, pages 1099–1103. IEEE, 2002.

Tohru Nitta. An extension of the back-propagation algorithm to

complex numbers. *Neural Networks*, 10:1391–1415, 1997.

Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

Timothy J O'Shea, Johnathan Corgan, and T Charles Clancy. Convolutional radio modulation recognition networks. In *International conference on engineering applications of neural networks*, pages 213–226. Springer, Cham, 2016.

Titouan Parcollet, Mirco Ravanelli, Mohamed Morchid, Georges Linarès, Chiheb Trabelsi, Renato De Mori, and Yoshua Bengio. Quaternion recurrent neural networks. *arXiv preprint arXiv:1806.04418*, 2018.

David P Reichert and Thomas Serre. Neuronal synchrony in complex-valued deep networks. *arXiv preprint arXiv:1312.6115*, 2013.

Juergen Reichert. Automatic classification of communication signals using higher order statistics. In *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 221–224. IEEE, 1992.

Hesamoddin Salehian, Rudransis Chakraborty, Edward Ofori, David Vaillancourt, and Baba C. Vemuri. An efficient recursive estimator of the fréchet mean on a hypersphere with applications to medical image analysis. *Mathematical Foundations of Computational Anatomy*, 2015.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

Jiaqi Shao, Changwen Qu, Jianwei Li, and Shujuan Peng. A lightweight convolutional neural network based on visual attention for sar image target classification. *Sensors*, 18(9):3039, 2018.

Rahul Singh, Abhishek Chakraborty, and BS Manoj. Graph fourier transform based on directed laplacian. In *2016 International Conference on Signal Processing and Communications (SPCOM)*, pages 1–5. IEEE, 2016.

Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, João Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. Deep complex networks. *arXiv preprint arXiv:1705.09792*, 2017.

Patrick Virtue. *Complex-valued Deep Learning with Applications to Magnetic Resonance Image Synthesis*. PhD thesis, EECS, UC Berkeley, 2018.

Patrick Virtue, Stella X. Yu, and Michael Lustig. Better than real: Complex-valued neural networks for mri fingerprinting. In *International Conference on Image Processing*, 2017.

Jiayun Wang, Patrick Virtue, and Stella X Yu. Joint embedding and classification for sar target recognition. *arXiv preprint arXiv:1712.01511*, 2017.

Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.

Stella Yu. Angular embedding: A robust quadratic criterion. *IEEE transactions on pattern analysis and machine intelligence*, 34(1): 158–173, 2011.

Stella X. Yu. Angular embedding: from jarring intensity differences to perceived luminance. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2302–9, 2009.

Stella X. Yu. Angular embedding: A robust quadratic criterion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1): 158–73, 2012.

Qibin Zhao, Guoxu Zhou, Shengli Xie, Liqing Zhang, and Andrzej Cichocki. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*, 2016.



Rudrasis Chakraborty received the Ph.D. degree in computer science from the University of Florida, Gainesville, FL, USA, in 2018.

He has been a Post-Doctoral Researcher with University of California, Berkeley and the International Computer Science Institute, Berkeley, CA, USA until very recently. He just joined Amazon Lab126, Sunnyvale, CA, USA, as an Applied Scientist, where he is working on 3-D computer vision. His research interests lie at the intersection of geometry, machine learning, and computer vision.



Yifei Xing was born in Beijing, China, in 1998. He received his B.A. in 2020 in applied mathematics and computer science at the University of California, Berkeley, CA, USA, where he is currently pursuing the Ph.D. degree in vision science under the mentorship of Prof. Stella Yu.

He shares an interest in complex-valued deep learning and its applications to image analysis.



Stella X. Yu received her Ph.D. from Carnegie Mellon University, Pittsburgh, PA, USA, in 2003.

She studied robotics at the Robotics Institute and vision science at the Center for the Neural Basis of Cognition, Carnegie Mellon University. She continued her computer vision research as a postdoctoral fellow at the University of California, Berkeley (UC Berkeley), CA, USA. She then studied art and vision as a Clare Booth Luce Professor at the Boston College, Newton, MA, USA, during which she received an NSF CAREER award. Dr. Yu is

currently the Director of the Vision Group, International Computer Science Institute (ICSI), a Senior Fellow at the Berkeley Institute for Data Science, and on the faculty of Computer Science, Vision Science, Cognitive and Brain Sciences, UC Berkeley. She is also a Faculty Member with the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, USA. Dr. Yu is interested in understanding visual perception from multiple perspectives and using computer vision and machine learning to automate and exceed human expertise in practical applications.