

# Parallel Transport Convolution: A New Tool for Convolutional Neural Networks on Manifolds

Stefan C. Schonsheck <sup>\*</sup>      Bin Dong <sup>†</sup>      Rongjie Lai <sup>‡</sup>

May 20, 2018

## Abstract

Convolution has been playing a prominent role in various applications in science and engineering for many years. It is the most important operation in convolutional neural networks. There has been a recent growth of interests of research in generalizing convolutions on curved domains such as manifolds and graphs. However, existing approaches cannot preserve all the desirable properties of Euclidean convolutions, namely compactly supported filters, directionality, transferability across different manifolds. In this paper we develop a new generalization of the convolution operation, referred to as parallel transport convolution (PTC), on Riemannian manifolds and their discrete counterparts. PTC is designed based on the parallel transportation which is able to translate information along a manifold and to intrinsically preserve directionality. PTC allows for the construction of compactly supported filters and is also robust to manifold deformations. This enables us to preform wavelet-like operations and to define deep convolutional neural networks on curved domains.

## 1 Introduction

Convolution is a fundamental mathematical operation that arises in many applications in science and engineering [1, 2]. Its ability to effectively extract local features, as well as its ease of use, has made it the cornerstone of many important techniques such as numerical partial differential equations and wavelets. More recently, convolution plays a fundamentally important role in convolutional neural networks (CNN) [3] which have made remarkable progress and significantly advanced the state-of-the-art in image processing, analysis and recognition [3, 4, 5, 6, 7, 8, 9, 10].

In Euclidean space  $\mathbb{R}^n$ , the convolution of a function  $f$  with a kernel (or filter)  $k$  is defined as:

$$(f * k)(x) := \int_{\mathbb{R}^n} k(x - y)f(y)dy. \quad (1)$$

Unlike signals or images whose domain is shift invariant, functions defined on curved domains such as curved manifolds do not have shift-invariance. Therefore, to properly define convolutions on these curved domains, one of the key challenges is to properly define the translation operation. This is one of the main obstacles of generalizing CNN to manifolds.

There has been a recent surge of research in designing CNNs on manifolds or graphs. We refer the interested readers to [11] for a review of recent progress in this area. These approaches can be classified into two categories: spectral and patch based methods. Spectral methods are based on projecting a signal onto the eigen (Fourier) space and using the convolution theorem to define convolution, while patch based methods use a patch operator to interpolate local geodesic discs on a certain given template. Here, we briefly review some of these approaches.

Spectral methods defines convolution based on the convolution theorem. The convolution theorem states that, for any two functions  $f$  and  $g$ :  $\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$  where  $\mathcal{F}$  is the Fourier transform and  $\cdot$  denotes

---

<sup>\*</sup>Department of Mathematics, Rensselaer Polytechnic Institute, Troy, NY 12180, U.S.A. (schon@rpi.edu)

<sup>†</sup>Beijing International Center for Mathematical Research, Peking University; Center for Data Science, Peking University; Beijing Institute of Big Data Research, Beijing, China. (dongbin@math.pku.edu.cn).

<sup>‡</sup>Department of Mathematics, Rensselaer Polytechnic Institute, Troy, NY 12180, U.S.A. (lair@rpi.edu).

point wise multiplication. This theorem can be naturally generalized to functions on manifolds if we let  $\mathcal{F}$  to be the projection operator onto the Laplace-Beltrami (LB) eigensystem. This method has proven effective to handle functions on a fixed domain, and can be applied to graphs as well [12, 13, 14, 15]. However, these methods have two fundamental limitations. First, the uncertainty principle states that a function can have compact support in either the time or frequency domain, but not both. These methods normally use only a finite number of eigenfunctions in the Fourier domain. As a result the kernels that arise from these methods are not localized (i.e. not compactly supported in the spatial domain). The second major drawback to these methods is that since they rely on the eigensystem of the domain, any deformation of the domain will change the eigensystem which in turn changes the filters. The high frequency LB eigenfunctions of a manifold are extremely sensitive to even small deformations. This means that anything designed for, or learned on, one manifold can only be applied to problems on the same domain. This limits the transferability of the spectral based methods.

Patch based methods are originally proposed in [16]. In this work the authors propose the use of a local patch operator to interpolate local geodesic discs of the manifold to a fixed template and develop a Geodesic Convolutional Neural Network (GCNN). Then for each point on the manifold, the convolution is calculated as the multiplication between the values of the kernel and the extracted patch on the template. To do so, they create a local polar coordinate system at each point. One drawback to this approach is that there is no natural way to choose the origin for the polar coordinate. To overcome this, the authors consider an angular pooling operation which evaluates all rotations of their kernel at each point and selects the orientation which maximizes the convolution in a point-wise fashion. Since the angular pooling operation is computed independently at each point, the selected orientation does not reflect the geometric structure of the ground manifold and may not be consistent even for nearby points. More recently, [17] proposes an anisotropic convolutional neural network (ACNN) by replacing the aforementioned patch operator with an operator based on anisotropic heat kernels with the direction of isotropy fixed on the principle curvature at each point. Although this introduces a new hyper-parameter (the level of isotropy), it allows the kernels to be directionally aware. However, filters developed for applications on one manifold can only be applied to manifolds in which the local directions of principle curvature are the same.

Table 1: Comparison on different generalizations of convolutional operator on manifolds.

Method	Filter Type	Support	Extraction	Directional	Transferable	Deformable
Spectral [13]	Spectral	Global	Eigen	Yes	No	No
TFG [14]	Spectral	Global	Eigen	Yes	No	No
WFT [18]	Spectral	Local	Windowed Eigen	Yes	No	No
GCNN [16]	Patch	Local	Variable	No	Yes	Yes
ACNN [17]	Patch	Local	Fixed	Yes	Yes	No
<b>PTC</b>	Geodesic	Local	Embedded	Yes	Yes	Yes

In the Euclidean setting, convolution operators that are frequently used in practice have compactly supported filters. Furthermore, they are directionally aware, deformable and can be easily transferred from one signal domain to another. Previous attempts to generalize the convolution operator on manifolds have failed to preserve one or more of these key properties. In this paper, we propose a new way of defining the convolution operation on manifolds based on parallel transportation. We shall refer to the proposed convolution as the parallel transportation convolution (PTC). The proposed PTC is able to preserve all of the aforementioned key properties of Euclidean convolutions. This spatially defined convolution operation enjoys flexibility of conducting isotropic or anisotropic diffusion, and it also enables us to perform wavelet-like operations as well as defining convolutional neural networks on manifolds. In addition, the PTC can be reduced to the Euclidean convolutions when the domain is flat. Therefore, the PTC is a natural generalization of Euclidean convolution operators on manifolds.

To be more precise, we seek a convolution operator of the form:

$$(f *_M k)(x) := \int_M k(x, y) f(y) d_M y. \quad (2)$$

In the Euclidean case,  $x - y$  essentially represents the direction from  $x$  to  $y$ , while on manifold such a vector

can be understood as a tangent direction at  $x$  pointing to  $y$ . The crucial idea of PTC is to define a kernel function  $k(x, y)$  which is able to encode  $x - y$  using a parallel transportation that naturally incorporates the manifold structure.

Table 1 compares the proposed PTC with previous efforts. A method is called directional if the filters are able to characterize non-isotropic features of the data. A method is transferable if the filters can be applied to manifolds with different LB eigensystems. Finally, a technique is said to be deformable if large deformations in the manifold (i.e. those which change properties such as curvature or local distances) do not drastically affect the convolution.

The rest of this paper is organized as follows. In section 2, we discuss necessary mathematical background of differential manifolds and parallel transportation of vector fields on manifolds. Then, we introduce the proposed PTC on manifolds in section 3. We also discuss implementation details and how to design convolutional neural networks on manifolds using the proposed PTC. In section 4, several numerical experiments illustrate the effectiveness of the proposed method. Finally, conclusion is made in section 5.

## 2 Mathematical Background

In this section, we discuss some background of differential manifolds and parallel transportation. These provide theoretical preparation of the proposed convolution on manifolds.

### 2.1 Manifolds, Tangent Spaces and the Exponential Map

Let  $\mathcal{M}$  be a two dimensional differential manifold associated with a metric  $g_{\mathcal{M}}$ . For simplicity we assume that  $(\mathcal{M}, \delta_{\mathcal{M}})$  is embedded in  $\mathbb{R}^3$ . We write the set of all tangent vectors at a point  $x \in \mathcal{M}$  as  $\mathcal{T}_x\mathcal{M}$  referred as the tangent plane of  $\mathcal{M}$  at  $x$ . The disjoint union of all tangent planes,  $\bigcup_x \{(x, v) \in \mathcal{M} \times \mathbb{R}^3 \mid x \in \mathcal{M}, v \in \mathcal{T}_x\mathcal{M}\}$ , forms a four dimensional differential manifold called the tangent bundle  $\mathcal{TM}$  of  $\mathcal{M}$ . A vector field  $X$  is a smooth assignment  $X : \mathcal{M} \rightarrow \mathcal{TM}$  such that  $X(x) \in \mathcal{T}_x\mathcal{M}$ ,  $\forall x \in \mathcal{M}$ . We denote all smooth vector fields on  $\mathcal{M}$  as  $C^\infty(\mathcal{M}, \mathcal{TM})$ .

Let  $\mathcal{T}_{x,\delta}\mathcal{M} = \{v \in \mathcal{T}_x\mathcal{M} \mid \langle v, v \rangle_{g_{\mathcal{M}}} \leq \delta\}$  be a  $\delta$ -neighborhood of the tangent space at a given point  $x$ . The *exponential map*,  $\exp : \mathcal{T}_{x,\delta}\mathcal{M} \rightarrow \mathcal{M}_{x,\delta}$ , maps vectors from the tangent space back onto a nearby region  $\mathcal{M}_{x,\delta}$  of  $x$  on the manifold. Formally, given  $v \in \mathcal{T}_{x,\delta}\mathcal{M}$  there exists a unique geodesic curve  $\gamma$  with  $\gamma(0) = x$  and  $\gamma'(0) = v$  such that  $\exp_x(v) = \gamma(1)$ . Note that this map is defined in the local neighborhood where the differential equation:  $\gamma'(0) = v$  with initial condition  $\gamma(0) = x$  has a unique solution. The size of this neighborhood depends on the local geometry of the manifold. In fact, the exponential map defines a one-to-one correspondence between  $\mathcal{T}_{x,\delta}\mathcal{M}$  and  $\mathcal{M}_{x,\delta}$  if  $\delta$  is smaller than the injective radius of  $\mathcal{M}$  [19, 20]. Since this map is a bijection, there is a natural inverse which we denote as  $\exp_x^{-1} : \mathcal{T}_{x,\delta}\mathcal{M} \rightarrow \mathcal{M}$ .

### 2.2 Parallel Transportation

In differential geometry, parallel transportation is a way of translating a vector, based an affine connection, along a smooth curve so the resulting vector is ‘parallel’. An affine connection translates the tangent spaces of points on a manifold in a way that allows us to differentiate vector fields along curves. Formally, an *affine connection* is a bilinear map  $\nabla : C^\infty(\mathcal{M}, \mathcal{TM}) \times C^\infty(\mathcal{M}, \mathcal{TM}) \rightarrow C^\infty(\mathcal{M})$ , such that for all smooth functions  $f$  and  $g$  and all vector fields  $X, Y$  and  $Z$  on  $\mathcal{M}$  satisfy:

1.  $\nabla_{fX+gY}X = f\nabla_XZ + g\nabla_YZ$
2.  $\nabla_X(aY + bZ) = a\nabla_XY + b\nabla_XZ \quad a, b \in \mathbb{R}$
3.  $\nabla_X(fY) = df(X)Y + f\nabla_XY$

In particular, an affine connection is called the Levi-Civita connection if it is torsion free ( $\nabla_XY - \nabla_YX = [X, Y]$ ) and compatible with the metric ( $X\langle Y, Z \rangle_{g_{\mathcal{M}}} = \langle \nabla_XY, Z \rangle_{g_{\mathcal{M}}} + \langle Y, \nabla_XZ \rangle_{g_{\mathcal{M}}}$ ). In this case, the transport induced from the connection preserves the length of the transported vector and the angle it makes with each side.

A curve  $\gamma : [0, \ell] \rightarrow \mathcal{M}$  on  $\mathcal{M}$  is called geodesic if  $\nabla_{\dot{\gamma}(t)}\gamma(t) = 0$ . More precisely, using local coordinate system, we can write  $\dot{\gamma}(t) = \sum_{i=1}^2 \frac{dx^i}{dt} \partial x^i$ , then plugging in the covariant derivative leads to the following ordinary differential equation for a geodesic curve:

$$\frac{d^2 x^k(t)}{dt^2} + \sum_{i,j=1}^2 \Gamma_{ij}^k(t) \frac{dx^i(t)}{dt} \frac{dx^j(t)}{dt} = 0, \quad k = 1, 2 \quad (3)$$

where  $\Gamma_{ij}^k$  is the Christoffel symbols associated with the local coordinate system. For any two points  $x_0$  and  $x_1$  on a complete manifold  $\mathcal{M}$ , there will be a geodesic  $\gamma : [0, \ell] \rightarrow \mathcal{M}$  connecting  $x_0$  and  $x_1$ . A vector field  $X(t)$  on  $\gamma(t)$  is called *parallel* if  $\nabla_{\dot{\gamma}} X = 0$ . Therefore, given any vector  $v \in \mathcal{T}_{x_0} \mathcal{M}$ , we can transport  $v$  to a vector  $v' \in \mathcal{T}_{x_1} \mathcal{M}$  by defining  $v' = X(\ell)$  from the solution of the initial value problem  $\nabla_{\dot{\gamma}(t)} X(t) = 0$  with  $X(0) = v$ . In other words, If we write  $X(t) = \sum_{i=1}^2 a^i(t) \partial x^i$ , the problem of solving  $X$  reduces to find the appropriate coefficients  $\{a^k(t)\}$  satisfying the parallel transport equation. This can be written as the following first order linear system:

$$\begin{cases} \frac{da^k(t)}{dt} + \sum_{i,j=1}^2 \frac{d\gamma^i}{dt} a^j(t) \Gamma_{ij}^k = 0, & k = 1, 2 \\ \sum_{i=1}^2 a^i(0) \partial x^i = v \end{cases} \quad (4)$$

Solving this equation finds a parallel vector field  $X$  along  $\gamma(t)$  which provides parallel transportation of  $v = X(0) \in \mathcal{T}_{x_0} \mathcal{M}$  to  $X(\ell) \in \mathcal{T}_{x_1} \mathcal{M}$ . We denote the parallel transportation of a vector from  $x_0$  to  $x_1$  along the geodesic as  $\mathcal{P}_{x_0}^{x_1} : \mathcal{T}_{x_0, \delta} \mathcal{M} \rightarrow \mathcal{T}_{x_1, \delta} \mathcal{M}$ .

### 3 Parallel Transport Convolution (PTC)

In this section, we introduce parallel transport convolution on manifolds which provide a fundamental important building block of designing convolutional neural networks on manifolds. After that, we discuss numerical discretization of PTC and PCTNet on manifolds.

#### 3.1 Mathematic definition of PTC

Unlike one-dimensional signals or images whose ground space is shift invariant, geometric objects modeled as curved manifolds do not have shift-invariance. This is an essential barrier to adopt CNN to conduct learning on manifolds and graphs except for a few recent work where convolution is defined in the frequency space of the LB operator [13, 21, 22, 23]. These methods only manipulates the LB eigenvalues by splitting the high dimension information to LB eigenfunctions. Limitations include that it is always isotropic due to the LB operator and can only approximate the even order differential operators [14]. In addition, there is another recent method discussed in [24], in which convolution is directly considered on the spatial domain using local integral on geodesic disc although it does not involve manifold structure as transportation on manifold is not considered. The lack of an appropriate method of defining convolution on manifolds motivates us to introduce the following way of defining convolution on manifolds through parallel transportation. This geometric way of defining convolution naturally integrates manifold structures and enables further learning tasks.

Let  $\mathcal{M}(x_0, \delta) = \{y \in \mathcal{M} \mid d_{\mathcal{M}}(x_0, y) \leq \delta\}$  and  $k(x_0, \cdot) : \mathcal{M}(x_0, \delta) \rightarrow \mathbb{R}$  be a compactly supported kernel function center at  $x_0$  with radius  $\delta$ . We assume  $k(x_0, y) = 0$  for  $y \notin \mathcal{M}(x_0, \delta)$  and require the radius of the compact support parameter  $\delta$  be smaller than the injective radius of  $\mathcal{M}$  to guarantee the bijectivity of the exponential map. It is important to remark that  $k(x_0, \cdot)$  can be determined by users, be designed for specific applications, or be optimized in the later exploration of parallel transport convolutional neural networks. Our idea of defining convolution on manifolds is conducted by transporting the compactly supported kernel  $k(x_0, \cdot)$  to every other point on  $\mathcal{M}$  in the way of reflecting the manifold geometry. More specifically,

given any point  $x \in \mathcal{M}$ , we first construct a vector field transportation  $\mathcal{P}_{x_0}^x : \mathcal{T}_{x_0, \delta} \mathcal{M} \rightarrow \mathcal{T}_{x, \delta} \mathcal{M}$  using the parallel transportation discussed in Section 2.2. Then  $k(x_0, \cdot)$  can be transported on  $\mathcal{M}$  as:

$$k(x, \cdot) : \mathcal{M}_{x, \delta} \rightarrow \mathbb{R}, \quad y \mapsto k(x_0, \exp_{x_0} \circ (\mathcal{P}_{x_0}^x)^{-1} \circ \exp_x^{-1}(y)) \quad (5)$$

Note that the above definition is an analogy of the convolution kernel in the Euclidean space. Here, the exponential map  $\exp_x^{-1}(y)$  mimics the vector  $x - y$ , and  $\mathcal{P}_{x_0}^x$  is a generalization of the shift operation in the Euclidean case where  $x_0 - y$  is shifted to  $x - y$ . In fact, it can be easily checked that the above definition is compatible with Euclidean case by setting the manifold  $\mathcal{M}$  to be  $\mathbb{R}^2$ .

By plugging (5) into (2), we can now formally define the *parallel transport convolution* operation of  $f$  which a filter  $k$ , centered at  $x_0$ :

$$(f *_{\mathcal{M}} k)(x) := \int_{\mathcal{M}} f(y) k(x, y) d_{\mathcal{M}} y = \int_{\mathcal{M}} f(y) k(x_0, \exp_{x_0} \circ (\mathcal{P}_{x_0}^x)^{-1} \circ \exp_x^{-1}(y)) d_{\mathcal{M}} y \quad (6)$$

As natural extensions, this approach can also be used to define dilations, reflections and rotations by simply manipulating the reference vector  $\exp_x^{-1}(y)$ . More specifically, shrinking or expanding the kernel by a factor of  $s$  is defined by multiplying the lengths of the vectors in the tangent space by  $s$ . If  $s$  is chosen to be negative then the kernel is reflected and dilated by a factor of  $|s|$ . Similarly, rotating the kernel is defined as conducting a rotation operator  $R_\theta$  to the reference vectors on the tangent plane. In summary, the scaling of  $k$  by  $s$  with a rotation of  $\theta$  is defined as:

$$k_{s, \theta}(x, y) := \frac{1}{C_x} k(x_0, \exp_{x_0} \circ (\mathcal{P}_{x_0}^x)^{-1}(s R_\theta \exp_x^{-1}(y))) \quad (7)$$

where  $R_\theta$  is a rotation matrix and  $\frac{1}{C_x}$  is a normalization constant that can be used to preserve volume of the kernel.

### 3.2 Numerical Discretization

We represent a two-dimensional manifold  $\mathcal{M}$  using triangle mesh  $\{V, E, T\}$ . Here  $V = \{v_i \in \mathbb{R}^3\}_{i=1}^n$  denotes vertices and  $T = \{\tau_s\}_{s=1}^l$  denotes faces. First we compute the geodesic distance function from  $x_0$  to every other point by solving the Eikonal equation  $|\nabla_{\mathcal{M}} D(x)| = 1$  using the fast marching method [26, 27]. Next we calculate  $\nabla_{\mathcal{M}} D$  and its orthonormal direction on each triangle  $\tau_s$ . Together with the face normal direction  $\vec{n}_s$ , for each triangle  $\tau_s$ , we construct a local orthonormal frame  $\tilde{\mathcal{F}}_s = \{\vec{b}_s^1, \vec{b}_s^2, \vec{n}_s\}$  where  $\vec{b}_s^1, \vec{b}_s^2$ , reflecting the intrinsic information, are tangent to  $\tau_s$ , and  $\vec{n}_s$ , reflecting the extrinsic information, is orthogonal to  $\tau_s$ . For an edge adjacent with  $\tau_s$  and  $\tau_t$ , we write  $R_{st}$  as an orthonormal transition matrix such that  $R_{st} \tilde{\mathcal{F}}_t = \tilde{\mathcal{F}}_s$ . Then any vector in  $\text{Span}\{\vec{b}_s^1, \vec{b}_s^2\}$  can be transported to  $\text{Span}\{\vec{b}_t^1, \vec{b}_t^2\}$  using the transition matrix  $R_{st}$ . This can be viewed as a discretization of connection and used to transport a vector on the tangent space of one given point to all other points. The compatibility condition of all  $R_{st}$  discussed in [28] can guarantee that no ambiguity will be introduced in this way.

After the transportation is conducted, the convolution kernel can be transported to a new point by interpolating the transported vectors in the local tangent space at the target point. Computationally, we define a sparse matrix  $K$  where the  $i^{\text{th}}$  column is the transportation of the kernel to the  $i^{\text{th}}$  vertex. Thus, we have the following definition of discrete parallel transport convolution:  $(f *_{\mathcal{M}} k)(x) := K^T \mathbf{M} F$  where  $F$  is column vector representation the function  $f$  at each vertex and  $\mathbf{M}$  is the mass matrix. Note that once we have found




Figure 1: A compactly supported kernel (a) is transported on a manifold from the FAUST data set [25] through translation (b), translation + dilation (c) and translation + rotation (d).

the vector field of the geodesic equation, the transportation of the kernel to each new center and multiplication with  $F$  is independent and can therefore be parallelized efficiently. Figure 1 illustrates the effect of the proposed method of transporting a kernel function on a manifold. This result shows that the proposed method produce an analogy of the behavior of a kernel function  $k(x-y)$  operating in the Euclidean domain.

### 3.3 Convolutional neural networks on manifolds through PTC

Using the proposed PTC, we can define convolutional neural networks on manifolds. We shall refer these network as PTCNets. Similar as CNNs on Euclidean domains, a PTCNet consists of an input and an output layer, as well as multiple hidden layers including fully connected layers, nonlinear layers, pooling layers and PTC layers listed as follows.

- Fully connected layer:  $f_i^{out}(x) = \sum_{j=1}^N w_{ij} f_j^{in}(x)$ ,  $i = 1, \dots, L$ . This layer connects every neuron in one layer to every neuron in the previous layer. The coefficient matrix  $(w_{ij})$  parametrizes this layer and will be trained by a training data set.
- ReLu layer:  $f_i^{out}(x) = \max\{0, f_i^{in}(x)\}$ ,  $i = 1, \dots, L$ . This is a fixed layer applying the nonlinear Rectified Linear Units function  $\max\{0, x\}$  to each input.
- PTC layer:  $f_{i,\alpha}^{out}(x) = \int k_\alpha(x, y) f_i^{in}(y) dy \approx K_\alpha \mathbf{M} F_i^{in}$ ,  $\alpha = 1, \dots, m$ . This layer applies the proposed PTC to the input, passes the result to the next layer. Each  $k_\alpha$  is determined by the proposed PTC on manifolds with an initial convolution kernel  $k_\alpha(x_0, \cdot)$ , which parametrize the parallel transport convolution process and will be trained based on a training data set. For certain applications with a moderate size of training set, more structured initial kernel might be needed. In this case, we can control  $k_\alpha$  by a sequence of rotation in the tangent space, which can reduce the number of free parameters and save computation time.
- Pooling layer:  $f_{i,\alpha}^{out}(x) = \max_\alpha f_{i,\alpha}^{in}(x)$ . The pooling layer can be implemented using several non-linear functions among which the max pooling is the most common way.

Therefore, it is straightforward to adapt established network architectures in Euclidean domain cases to manifolds case as the only change is to replace traditional convolution by PTC. In addition, back-propagation can be achieved by taking derivation of  $K$ . The compact support of the convolution kernel is represented as a sparse matrix which makes computation efficient.

## 4 Numerical Experiments

To illustrate the effectiveness of the proposed PTC, we conduct numerical experiments including processing images on manifolds using PTC and classifying images on manifolds using PTCNets. All numerical experiments are implemented in MATLAB on a PC with a 32GB RAM and 3.5GHz CPU. We remark that these experiments aim to demonstrate capabilities the proposed PTC for manipulating functions on curved domains by naturally extending existing wavelet or learning methods from Euclidean domains to curves domains. It is by no means to show that the experiments achieve state-of-the-art results. In our future work, we will show its further applications of comparing, classifying and understanding manifold-structured data by combing with recent advances in deep learning.

### 4.1 Wavelet-like Operations

In the first experiment, we demonstrate the effectiveness of our approach by performing simple signal processing tasks on manifolds. Then we compare the PTC results to those produced by traditional techniques applied to Euclidean domains. First we apply PTC with a hand crafted edge detection filter to an image on a manifold. By convolving this filter with the input image, we obtain an output feature function whose higher values indicate similarity to the predefined edge. In the first row of Figure 2, it is clear that the proposed convolution successfully highlights the edges with similar orientation of the input filter. In the second row of Figure 2, we allow additional rotations as we discussed in (7). We observe that the additional rotation

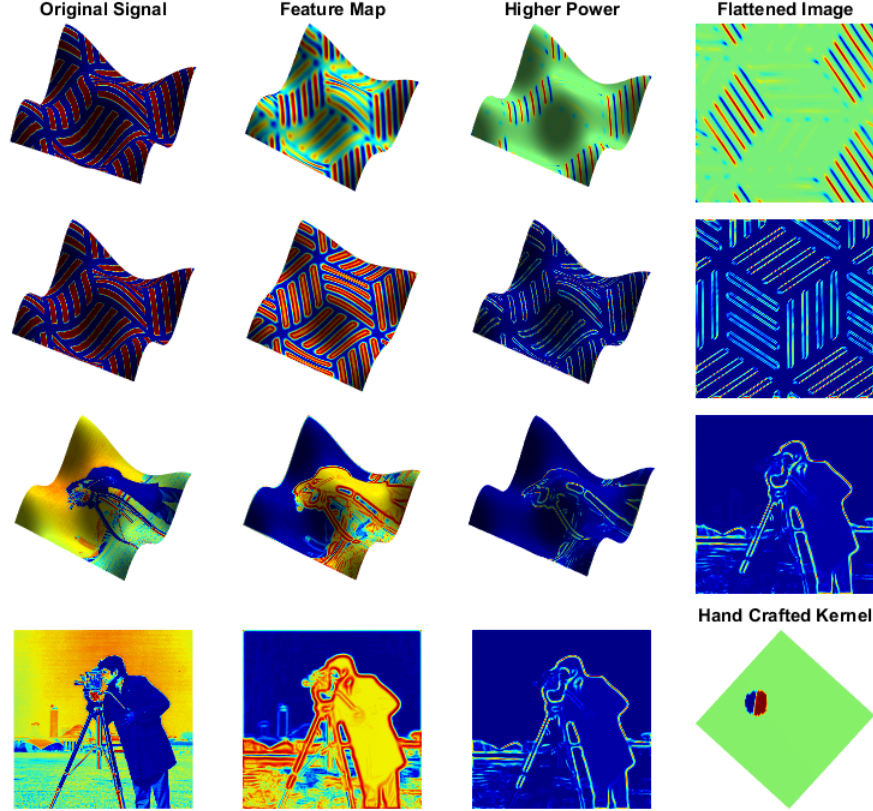


Figure 2: First Row: Convolutions with out rotation on test image. Second Row: Convolutions with rotation on test image. Third Row: Convolutions with rotation on a cameraman image. Fourth row: Traditional Euclidean convolution and the edge detector used in PTC.

flexibility can reliably capture all of the edges regardless of orientations. This illustrates the directional awareness of our method. Furthermore, we apply this edge detector using PTC to a more realistic problem in the third row of Figure 2. It shows that the results are very close to those produced in an analogous Euclidean setting (fourth row). In the third column we show the feature mapped raised to a higher power for better contrast and the last column shows a flattened version for easier visualization.


## 4.2 Single Manifold MNIST

In this test, we conduct experiments to demonstrate the effectiveness of PTCNets to handle signals on manifolds. The most highly celebrated early applications of CNNs was the recognition of hand written digits [29]. We map all MNIST data to a curved manifold plotted in the left image of Table 2. We use a simple network architecture consisting of a single convolution layer with 16 filters followed by a ReLu non-linear layer and then a fully connected layer which outputs a 10 dimensional vector of predictions. We apply this network architecture to four scenarios including MNIST data on a Euclidean domain using traditional convolution, MNIST data on a Euclidean domain using PTC, MINST data on a curved domain using PTC, and MINST data on the same curved domain using spectral convolution.

Each network is implemented in MATLAB using only elementary functions and is trained using batch stochastic gradient descent with batch size 50 and a fixed learning rate  $\alpha = 10^{-3}$ . We also use the same random seed for the batch selection and the same initialization. We choose such a simple training regime in order to make the effects of different convolution operations as clear as possible. We measure the results by the overall network error after 5,000 iterations.

Table 2 shows the accuracy of the traditional CNN on a flat domain, a spectral net applied to a simple manifold as well as our network applied to both a Euclidean domain (Flat PTCNet) and the manifold. From these results, we conclude that our method outperforms the spectral network and can achieve similar results with the traditional model when applied to a curved domain.

Table 2: Comparison of our PTCNet to Euclidean case and a spectral based method on a single manifold.



Network	Domain	Success Rate
Traditional	Euclidean	<b>98.85</b>
<b>Flat PTCNet</b>	Euclidean	98.10
Spectral	Manifold	95.35
<b>PTCNet</b>	Manifold	97.96

### 4.3 Multi-Manifold MNIST




Figure 3: Left image: Manifolds used for multi-mainfold tests. The first four are used for training and the last is used for testing. Right table: Comparison of results from learning on single and multiple domains and then testing on a new manifold

One of the advantages of our method is that filters which are learned on one manifold can be applied to different domains. The spectral convolution based methods do not have this transferability as different domains are unlikely to share the same eigensystem. In this experiment, we first directly apply the network learned by the PTCNet and Spectral networks from Section 4.2 to a new manifold. As we illustrate in the first two rows of Table fig:MultiMNIST, the accuracy of the spectral convolution based method is dramatically reduced since the two manifolds have quite different eigensystems. However, our PTCNet can still provide reasonable accurate results since the underlying geodesic vector fields of these manifolds is more stable to deformations than eigensystems are.

Furthermore, we conduct a new experiment in which we train our PTCNet on a variety of manifolds and test on a different manifolds as showed in the left picture of Figure 3. The first four manifolds are used as training domains, and the fifth one is used for testing. From these pictures it is clear that the training manifolds are quite different and therefore definitions of convolution which require curvature to set their direction [17] cannot be applied to these problems. However the geodesic vector fields of the manifolds are quite similar and therefore filters learned through our technique should apply to the new problem. As we can see in the last row of the Table in Figure 3, the network achieves a 97.32% success rate since training on multiple manifolds allows PTC network to learn greater invariance to local deformation in the metric, which enables greater transferability.

## 5 Conclusion

In this paper, we propose a generalization of the convolution operation on smooth manifolds using parallel transportation and discuss its numerical implementation. Using the proposed PTC, we have preformed wavelet-like operation of signals and built convolutional neural networks on curved domains. Our numerical experiments have shown that the PTC can preform as well as Euclidean methods on curved manifolds, and is capable of including directional awareness, handling problems involving deformable manifolds. In our future works, we will apply our PTC to different applications of comparing, classifying and understanding manifold-structured data by combining with recent advances of deep learning.



## Acknowledgments

Stephan Schonsheck and Rongjie Lai are supported in part by NSF DMS-1522645. Bin Dong is supported in part by NSFC 11671022.

## References

- [1] I. Daubechies. *Ten lectures on wavelets*, volume 61. SIAM, 1992.
- [2] Stephane Mallat. *A wavelet tour of signal processing: the sparse way*. Academic press, 2008.
- [3] Yann LeCun, Bernhard E Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne E Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*, pages 396–404, 1990.
- [4] Yoshua Bengio, Ian J Goodfellow, and Aaron Courville. Deep learning. *Nature*, 521:436–444, 2015.
- [5] 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.
- [6] Dan Ciresan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *Advances in neural information processing systems*, pages 2843–2851, 2012.
- [7] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, and Tara N Sainath. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [8] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
- [9] Michael KK Leung, Hui Yuan Xiong, Leo J Lee, and Brendan J Frey. Deep learning of the tissue-regulated splicing code. *Bioinformatics*, 30(12):i121–i129, 2014.
- [10] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [11] 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.
- [12] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
- [13] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [14] Bin Dong. Sparse representation on graphs by tight wavelet frames and applications. *Applied and Computational Harmonic Analysis*, 42 (3):452–479, 2017.
- [15] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- [16] Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015.

- [17] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016.
- [18] David I Shuman, Benjamin Ricaud, and Pierre Vandergheynst. Vertex-frequency analysis on graphs. *Applied and Computational Harmonic Analysis*, 40(2):260–291, 2016.
- [19] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry*, volume 2. Interscience publishers New York, 1969.
- [20] Isaac Chavel. *Riemannian geometry: a modern introduction*, volume 98. Cambridge university press, 2006.
- [21] Jamie Shotton, Toby Sharp, Alex Kipman, Andrew Fitzgibbon, Mark Finocchio, Andrew Blake, Mat Cook, and Richard Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124, 2013.
- [22] Emanuele Rodolà, Samuel Rota Bulo, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense non-rigid shape correspondence using random forests. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4184, 2014.
- [23] Raif Rustamov and Leonidas J Guibas. Wavelets on graphs via deep learning. In *Advances in neural information processing systems*, pages 998–1006, 2013.
- [24] Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Shapenet: Convolutional neural networks on non-euclidean manifolds, 2015.
- [25] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014.
- [26] James A Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.
- [27] Ron Kimmel and James A Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998.
- [28] Yuanzhen Wang, Beibei Liu, and Yiying Tong. Linear surface reconstruction from discrete fundamental forms on triangle meshes. In *Computer Graphics Forum*, volume 31, pages 2277–2287. Wiley Online Library, 2012.
- [29] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.