King's Research Portal

DOI:

[10.1016/j.patrec.2017.11.008](https://doi.org/10.1016/j.patrec.2017.11.008)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Tsagkrasoulis, D., & Montana, G. (2017). Random Forest regression for manifold-valued responses. *PATTERN RECOGNITION LETTERS*. <https://doi.org/10.1016/j.patrec.2017.11.008>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Accepted Manuscript

Random Forest regression for manifold-valued responses

Dimosthenis Tsagkrasoulis, Giovanni Montana

PII: S0167-8655(17)30415-4
DOI: [10.1016/j.patrec.2017.11.008](https://doi.org/10.1016/j.patrec.2017.11.008)
Reference: PATREC 7000

To appear in: *Pattern Recognition Letters*

Received date: 16 February 2017
Revised date: 7 August 2017
Accepted date: 5 November 2017

Please cite this article as: Dimosthenis Tsagkrasoulis, Giovanni Montana, Random Forest regression for manifold-valued responses, *Pattern Recognition Letters* (2017), doi: [10.1016/j.patrec.2017.11.008](https://doi.org/10.1016/j.patrec.2017.11.008)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- Distance-based Random Forest regression methodology.
- Methodology depends only on pairwise distances between training response observations.
- Easy application on a wide variety of data response objects.

Pattern Recognition Letters
journal homepage: www.elsevier.com

Random Forest regression for manifold-valued responses

Dimosthenis Tsagkraloulis^a, Giovanni Montana^{a,b,**}

^aDepartment of Mathematics, Imperial College London, SW7 2AZ, London, UK

^bDepartment of Biomedical Engineering, King's College London, SE1 7EH, London, UK

ABSTRACT

An increasing array of biomedical and computer vision applications requires the predictive modeling of complex data, for example images and shapes. The main challenge when predicting such objects lies in the fact that they do not comply to the assumptions of Euclidean geometry. Rather, they occupy non-linear spaces, a.k.a. manifolds, where it is difficult to define concepts such as coordinates, vectors and expected values. In this work, we construct a non-parametric predictive methodology for manifold-valued objects, based on a distance modification of the Random Forest algorithm. Our method is versatile and can be applied both in cases where the response space is a well-defined manifold, but also when such knowledge is not available. Model fitting and prediction phases only require the definition of a suitable distance function for the observed responses. We validate our methodology using simulations and apply it on a series of illustrative image completion applications, showcasing superior predictive performance, compared to various established regression methods.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

Predictive modeling is an integral part of data analysis. It encompasses the process of developing models which can accurately predict yet-to-be-seen data. A multitude of regression models exist that can be used for prediction of univariate, or multivariate vectorial, responses. Nevertheless, less work has been done on the difficult problem of modeling and predicting more complex objects that possess additional structure, be it morphological, directional, or otherwise [13].

Images [12], shapes [18], graphs [20], deformation tensors [23] are examples of complex data types that appear naturally as responses in image analysis, computer vision, medical imaging and other application areas. While such objects are typically represented as points on very high-dimensional spaces, they can be meaningfully represented as points lying on smooth non-linear hyper-surfaces of lower dimensionality, a.k.a. manifolds [3]. Manifolds can be understood as surfaces that locally resemble the Euclidean plane, but have different global structures [14]. Ideally, a suitable predictive modeling methodology should work under the additional constraints imposed by

the data's inherent manifold structure instead of simply treating these complex objects as points on Euclidean spaces. Major difficulties of course arise by forgoing the assumption of a Euclidean space, such as lack of coordinates, vectors and no analytical definitions of expected values [6]. In this work, we are additionally interested in experimental settings where the input observations may be very-high dimensional, which poses further requirements in the construction of a predictive algorithm.

Three main families of methodologies regarding regression for manifold-valued responses can be found in the literature; intrinsic manifold regression, kernel-based and Manifold Learning (ML)-based methods. Intrinsic manifold regression models are predominantly generalizations of linear regression on manifolds [6, 16], although a non-parametric intrinsic model has also been recently proposed [22]. Both parametric and non-parametric intrinsic methods require the analytical definition of the data manifold. The assumption of an underlying manifold drives the choice of geometric elements employed during model definition and estimation. Unfortunately, this requirement can be rarely satisfied, either due to the nature of data or the inability to define a suitable manifold. Furthermore, this class of methodologies is not ideal for regression analysis with highly dimensional input observations.

Kernel methods first appeared in the literature as methodologies tailored for complex input objects, such as trees or graphs

^{**}Corresponding author: Tel.: +44 (0) 20 7594 8577;

e-mail: giovanni.montana@kcl.ac.uk (Giovanni Montana)

[15]. The basic idea behind these methods is that, if the input data lie on a non-linear space, then they can be implicitly mapped on a very high (or infinite) dimensional inner-product space, in which standard regression methods can be applied [15]. This implicit mapping is achieved through a kernel function that defines the objects' inner-product in that space [15]. Kernel methods for complex responses have also been proposed [20, 7], but suffer from a number of issues. First, depending on the data at hand, a meaningful kernel function must be found or constructed, a process that is not always intuitive. Second, another problem is the formulation of a prediction methodology, which has to be described as a kernel minimization problem over the response space. This is most commonly solved by reducing the search space to the training dataset.

The last family of manifold regression methods is based on non-linear dimensionality reduction, a.k.a. Manifold Learning. Given a set of observed points, ML methods aim to project these onto a space of lower dimensionality, whilst retaining as well as possible the original geometrical structure of the data. ML-based methods first apply ML on the response data, and subsequently use standard regression models trained on the learned response embedding. A suitable methodology must be formulated to map predicted points from the embedding to the manifold, a process that is referred to as backscoring [21]. Appropriate selection of the ML technique is largely based on intuition and previous experience. A decrease in model fitting accuracy, compared to an intrinsic manifold model, is to be expected, since the response embedding is not guaranteed to completely capture the structure of the underlying manifold. Furthermore, ML techniques generate discrete and one way maps from the manifold to the embedding, and the existence of an inverse and continuous map cannot always be guaranteed [21].

An extended description of related manifold regression methodologies can be found in Supplementary Text section 1.

Here, we present a new methodology for predictive modeling of manifold-valued responses, which addresses a number of key issues listed above. Our objective was to propose a unified framework for regression and prediction of complex objects that is accurate, computationally efficient and can be readily deployed in a variety of applications. In the training phase, we employ our modified Random Forest (RF) regression algorithm that can be trained using only pairwise distances between response observations [17]. The non-parametric RF methodology is efficient and can handle high dimensional input spaces. In contrast to intrinsic manifold regression, no analytical definition of the underlying response manifold is required, apart from the construction of an appropriate distance function for the responses. Comparing our distance-based approach to kernel methods, we identify further advantages. First, a vast library of readily available distance metrics exists in the literature (see for example [5]). Second, when the manifold metric is not known, distances can be intuitively approximated, using for example the Isomap distance formulation [19].

In the prediction phase, a response point estimate for a new input observation is found as follows. Pairwise distances between the unseen response and all training observations are estimated. Using this set of distances, the response is predicted

on a Euclidean embedding computed using the ML technique of metric multi-dimensional scaling (MDS) and is finally mapped to the response manifold through a standardized backscoring procedure. Our prediction methodology follows a two-step approach akin to ML-based methods, whilst offering two additional advantages. First, our regression model is trained on the original manifold, which enhances the quality of the fitted model. Second, due to the fact that manifold distances are known, MDS is employed for ML and predictions can be analytically computed and backscored to the response manifold.

The rest of this article is organized as follows. In section 2 we present the details of our manifold regression methodology. Our simulation and application experiments are included in section 3. We conclude this work in section 4.

2. Random Forest Predictive Methodology for Manifold-Valued Objects

2.1. Problem Definition

Let S be a dataset of N observed input-response pairs $\{(x_i, y_i)\}_{i=1}^N$, with inputs $x_i = (x_{i1}, \dots, x_{ip}) \in \mathbb{R}^p$ corresponding to responses $y_i \in \mathcal{M} \subset \mathbb{R}^q$. $\mathcal{M}$ is a possibly unknown manifold, equipped with a distance metric $d_{\mathcal{M}}(\cdot, \cdot)$. We refer to $\mathbb{R}^q$ as the response representation space. Let $D_{\mathcal{M}}$ be the $N \times N$ matrix of pairwise distances between the observed response points, with elements $D_{\mathcal{M},ij} = d_{\mathcal{M}}(y_i, y_j)$. See Fig. 1(a) for a graphical illustration of the described data. The problem consists of defining a suitable predictive regression methodology that utilizes information about the structure of the manifold during model fitting, and, given a new input point x_{new} , is capable of predicting a response point estimate $\hat{y}_{new} \in \mathcal{M}$.

Fig. 1. Illustration of the dRF prediction methodology. (a) The dRF model is fitted on the dataset S of N observed input-response points, with $y_i \in \mathcal{M}$, and the distance matrix $D_{\mathcal{M}}$ of pairwise manifold distances between the observed response points. (b) When a new input y_{new} is observed, the point is passed through the forest and is associated with subsets of training response points in the trees' leaves. A similarity vector a_{RF} between the training responses and the still unknown response point is extracted from the model. (c) Based on a_{RF} , the set of distances between the new point and the training responses on the manifold are predicted. (d) Knowledge of these distances allows prediction of the response value on a Euclidean embedding of the manifold, extracted through metric Multi-Dimensional Scaling. (e) Finally, a kernel-based backscoring method is used to project the predicted point back to the original manifold and acquire $\hat{y}_{new}$.

2.2. Distance Random Forest Regression

We first concentrate on the construction of a suitable manifold regression learning algorithm. Random Forest (RF) is a non-parametric, non-linear regression and classification algorithm [2]. In more detail, RF is a collection of Classification and Regression Trees (CARTs). A CART splits the input space recursively, according to a predefined split function, to small rectangular regions and then fits a simple model, commonly a constant value, in each one of them. In [17], we presented a modified distance Random Forest (dRF) algorithm, where the split criterion was formulated to depend only on pairwise distances between responses:

$$G_{dRF}(S_w) = \frac{1}{2N_w} \sum_{y_i \in S_w} \sum_{y_j \in S_w} d_M^2(y_i, y_j) - \frac{1}{2N_{wl}} \sum_{y_i \in S_{wl}} \sum_{y_j \in S_{wl}} d_M^2(y_i, y_j) - \frac{1}{2N_{wr}} \sum_{y_i \in S_{wr}} \sum_{y_j \in S_{wr}} d_M^2(y_i, y_j), \quad (1)$$

with $y_i \in \mathcal{M}$ and N_w, N_{wl}, N_{wr} cardinalities of the data subsets S_w, S_{wl}, S_{wr} in node w and its left and right children nodes respectively. See Supplementary Text sections 2 and 3 for detailed descriptions of RF and dRF respectively. The objective (1) is a generalization of the cost function used in standard regression RF, decoupled from the use of Euclidean norms and means, dependent only on pairwise response distances. As such, it enables the RF algorithm to be applied in general metric spaces. Previously, we used dRF for regression applications with graph and covariance-based response objects [17]. dRF lends itself naturally to manifold-valued data, whether the manifold in question is analytically defined or implied by the use of a specific distance metric.

2.3. Predictive Methodology for dRF

Traditionally, when working with responses lying in Euclidean spaces, an RF prediction is made simply by averaging response points; a new input observation follows the split decision rules learned during training and reaches a terminal node -leaf- in the tree. It is then assigned a suitable value on the response space, most often the average response of that leaf's responses in the training set. This approach though is not valid under the assumption of non-vectorial manifold responses [11].

Our proposed methodology uses dRF for learning a family of trees from the data, which only requires an appropriate distance metric for the responses, as described above. The prediction phase is substantially different from standard RF. For a new input point x_{new} , we use the trained dRF model to predict all pairwise distances between the unknown response, y_{new} , and all observed responses in the training dataset, y_i . Having estimated these distances, we then utilize them to predict the response on an Euclidean embedding of the manifold, learned from the observed response set using metric MDS. Subsequently, a backscoring method is employed to project the point back to the original response manifold $\mathcal{M}$. Figure 1 provides an illustration of the proposed methodology, and the details are in order.

2.3.1. RF-based Distance Prediction

The first step is to estimate the set of distances $\{\hat{d}_M(y_{new}, y_i)\}_{i=1}^N$. We exploit the inherent ability of the dRF model to provide a measure of similarity between pairs of response observations (see Fig. 1(b)). When the new input point x_{new} is 'dropped' through each tree in the forest, it reaches a leaf associated with a subset of the training data. The y_{new} can be considered similar to that leaf's training responses. To compute the dRF-based similarities, a vector $\mathbf{a}_{RF} = (a_{RF,1}, \dots, a_{RF,N})$ is used, with each element $a_{RF,i}$ corresponding to the similarity between y_{new} and y_i . Initially $a_{RF,i}$ is set to zero, for all $i = 1, \dots, N$. For each tree in RF, $a_{RF,i}$ is incremented by one, each time x_{new} ends in the same node as x_i . Normalization of similarities is performed by division with the number of trees.

Based on the similarity vector $\mathbf{a}_{RF}$ and the training distance matrix $\mathbf{D}_M$, the response distances are predicted using the algorithmic procedure 1 (see Fig. 1(c)), which guarantees that the new point will lie in close proximity to at least its closest neighbor based on the dRF affinities, and that predicted distance values respect the triangular inequality property of metric d_M . In detail, the algorithm initially identifies the closest training point to the new response, according to $\mathbf{a}_{RF}$, and assigns the minimum distance observed in the training set as the distance between these two responses. Subsequently, we iterate over the remaining responses, in decreasing order of distance from the first point, and assign a value for the distance to the new observation as follows: For each triplet including the considered point, the new point and a point for which the distance to y_{new} has been already estimated, we keep the maximum of the two known distances. Subsequently, we assign the minimum of all identified maximums as the predicted distance.

Algorithm 1 Prediction of $\{\hat{d}_M(y_{new}, y_i)\}_{i=1}^N$ using $\mathbf{a}_{RF}$ and $\mathbf{D}_M$.

Require: $\mathbf{a}_{RF}$ and $\{D_{M,ij}\}_{i,j=1}^N$

$S = \{1, \dots, N\}, Q = \emptyset$

$l \leftarrow \underset{i \in S}{\operatorname{argmax}} a_{RF,i}$

$\hat{d}_M(y_{new}, y_l) \leftarrow \min\{D_{M,ij}\}_{i,j=1}^N$

$S \leftarrow S \setminus \{l\}, Q \leftarrow Q \cup \{l\}$

repeat

$p \leftarrow \underset{i \in S}{\operatorname{argmax}} D_{M,il}$

$\hat{d}_M(y_{new}, y_p) \leftarrow \min\{\max\{\hat{d}_M(y_{new}, y_i), D_{M,ip}\} | i \in Q\}$

$S \leftarrow S \setminus \{p\}, Q \leftarrow Q \cup \{p\}$

until $S = \emptyset$

return $\{\hat{d}_M(y_{new}, y_i)\}_{i=1}^N$

2.3.2. Prediction on the Response Embedding

Given $\{\hat{d}_M(y_{new}, y_i)\}_{i=1}^N$, we can predict a point estimate of the response on a Euclidean embedding of the manifold using MDS [4] (see Fig. 1(d)).

MDS computes an approximation $\{z_i | z_i \in \mathbb{R}^m\}_{i=1}^N$ of the manifold-valued dataset that resides on a m -dimensional Euclidean space, by minimizing a stress function of the form $\sum_{i=1}^N \sum_{j=1}^N (d_M^2(y_i, y_j) - \|z_i - z_j\|^2)$. Let $\mathbf{D}_M^{(2)}$ be the $N \times N$ matrix of squared manifold distances and $\mathbf{K}_M = -\frac{1}{2} \mathbf{H} \mathbf{D}_M^{(2)} \mathbf{H}$, where

$\mathbf{H} = \mathbf{I}_N - \frac{1}{N}\mathbf{e}\mathbf{e}^T$, with $\mathbf{I}_N$ the $N \times N$ identity matrix and $\mathbf{e}$ an $N \times 1$ column vector of all ones.

Individual elements of matrix $\mathbf{K}_M$ are given by $K_{M,ij} = -\frac{1}{2}(D_{M,ij}^{(2)} - \frac{1}{n}D_{M,j}^{(2)} - \frac{1}{n}D_{M,i}^{(2)} + \frac{1}{n^2}D_{M,\cdot\cdot}^{(2)})$, $i, j = 1, \dots$, where $D_{M,j}^{(2)} = \sum_i D_{M,ij}^{(2)}$, $D_{M,i}^{(2)} = \sum_j D_{M,ij}^{(2)}$, $D_{M,\cdot\cdot}^{(2)} = \sum_{i,j} D_{M,ij}^{(2)}$ are the i^{th} row, j^{th} column and overall element-wise sums of $D_M^{(2)}$ respectively.

In [9], it was shown that, if $\mathbf{K}_M$ has rank p ($p \leq N-1$), with $\lambda_1, \dots, \lambda_p$ the p ordered non-zero eigenvalues of $\mathbf{K}_M$ with corresponding eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_p$, and $\mathbf{u}_l = (u_{l1}, \dots, u_{lN})$ for $l = 1, \dots, p$, then the solution to the MDS problem, assuming $m \leq p$, is given by

$$z_{il} = \sqrt{\lambda_l} u_{li}, \quad l = 1, \dots, m, \quad i = 1, \dots, N, \quad (2)$$

with $z_i = (z_{i1}, \dots, z_{im})$.

The result of the MDS decomposition is a discrete and one-way mapping $z_i = f_{MDS}(y_i)$ defined on the training data set. An Out-Of-Sample (OOS) method which allows mapping of a new manifold observation on the learned space of the embedding was constructed in [1]. Let $Y_a, Y_b \in \mathcal{M}$ be random variables defined on the manifold surface, with the same probability density function, and $y_a, y_b \in \mathcal{M}$ two realizations of Y_a and Y_b respectively. Then a kernel function g_{MDS} can be defined as follows:

$$g_{MDS}(y_a, y_b) = -\frac{1}{2}(d_M^2(y_a, y_b) - E[d_M^2(Y_a, Y_a)] - E[d_M^2(Y_b, Y_b)] + E[d_M^2(Y_a, Y_b)]) , \quad (3)$$

where $E(\cdot)$ denotes expectation.

The OOS prediction $\hat{z}_{new} = (\hat{z}_{new,1}, \dots, \hat{z}_{new,m})$ is given by [1]:

$$\hat{z}_{new,k} = \frac{1}{\sqrt{\lambda_k}} \sum_{i=1}^N u_{ik} \hat{g}_{MDS}(y_{new}, y_i), \quad k = 1, \dots, m, \quad (4)$$

with $\hat{g}_{MDS}$ denoting the mean estimator of g_{MDS} under the augmented dataset $\{y_i\}_{i=1}^N \cup \{y_{new}\}$:

$$\begin{aligned} \hat{g}_{MDS}(y_{new}, y_i) = & -\frac{1}{2}d_M^2(y_{new}, y_i) + \frac{1}{2(N+1)} \left(\sum_{j=1}^N d_M^2(y_j, y_{new}) \right) \\ & + \frac{1}{2(N+1)} \left(\left(\sum_{j=1}^N d_M^2(y_j, y_i) \right) + d_M^2(y_{new}, y_i) \right) \\ & - \frac{1}{2(N+1)^2} \left(\left(\sum_{j,l=1}^N d_M^2(y_j, y_l) \right) + 2 \sum_{j=1}^N d_M^2(y_{new}, y_j) \right). \end{aligned} \quad (5)$$

Details regarding the OOS formulation (4) are given in the Supplementary Text section 4.

The OOS formula does not depend on the actual value of y_{new} , but rather hush on the distances between the new and training response points. We can thus leverage the predicted $\{\hat{d}_M(y_{new}, y_i)\}_{i=1}^N$ on equation (4) in order to get a point estimate of our response $\hat{z}_{new}$ on the embedding space $\mathbb{R}^m$. Now we are left with the task of mapping $\hat{z}_{new}$ to the original manifold (see figure 1(e)).

2.3.3. Mapping from the Embedding to the Manifold

The mapping of the response from the embedding to the manifold - backscoring - can be formulated as an interpolation problem. Specifically, we are looking for a smooth continuous function $f_{bs} : \mathbb{R}^m \rightarrow \mathbb{R}^q$, such that $y_i \approx f_{bs}(z_i)$ for $i = 1, \dots, N$:

$$\underset{f_{bs} \in \mathcal{F}_{bs}}{\operatorname{argmin}} \quad \gamma_{bs} \sum_{i=1}^N (f_{bs}(z_i) - y_i)^2 + \|f_{bs}\|_{\mathcal{F}_{bs}}^2, \quad (6)$$

where γ_{bs} is a weight parameter balancing the smoothness of the interpolation and the adherence to the data and $\mathcal{F}_{bs}$ is a space of smooth functions equipped with the norm $\|\cdot\|_{\mathcal{F}_{bs}}$, which will be defined below.

An optimal solution to the optimization problem (6) was constructed in [21]. Let $\mathcal{Z}_{bs}$ be a closed subset of $\mathbb{R}^m$. We define g_{bs} to be a kernel function of the form

$$g_{bs}(r, z_a, t, z_b), \quad r, t \in T^q = \{1, \dots, q\}, z_a, z_b \in \mathcal{Z}_{bs}. \quad (7)$$

Also, let g_{bs} be semi-positive definite on $(T^q \times \mathcal{Z}_{bs}) \times (T^q \times \mathcal{Z}_{bs})$, i.e. $\sum_{k=1}^N \sum_{l=1}^N a_k a_l \sum_{r=1}^q \sum_{t=1}^q g_{bs}(r, z_k, t, z_l) \geq 0$, for any finite set of points $\{z_k | z_k \in \mathcal{Z}_{bs}, k = 1, \dots, N\}$ and real numbers $\{a_k | a_k \in \mathbb{R}, k = 1, \dots, N\}$.

For a fixed (r, z_a) , equation (7) defines a function from $\mathbb{R}^m$ to $\mathbb{R}^q$ by the formula

$$f_{r,z_a}(z_b) = (g_{bs}(r, z_a, 1, z_b), \dots, g_{bs}(r, z_a, q, z_b))^T. \quad (8)$$

Based on the above, let $\mathcal{F}_{bs}$ be the space of all finite linear combinations of functions of the form (8), as (r, z_a) varies in $T^q \otimes \mathcal{Z}_{bs}$, and its closure w.r.t the scalar inner product $\langle f_{r,z_a}, f_{t,z_b} \rangle = g_{bs}(r, z_a, t, z_b)$. It follows that $\|f_{r,z_a}\|_{\mathcal{F}_{bs}} = \sqrt{\langle f_{r,z_a}, f_{r,z_a} \rangle}$.

Then the minimizer of (6) can be formulated as follows. Given our training responses $\{y_i | i = 1, \dots, N\}$ and their equivalent embedded points $\{z_i | i = 1, \dots, N\}$, let $\mathbf{Y}$ be the $N \times q$ matrix of training responses with rows $Y_i = y_i$, and $\mathbf{K}_{bs}$ the $Nq \times Nq$ matrix of kernel values $K_{bs, N(r-1)+i, N(t-1)+j} = g_{bs}(r, z_i, t, z_j)$. The minimizing function can be estimated as [21]:

$$\hat{f}_{bs}(\hat{z}_{new}) = \sum_{i=1}^N \sum_{r=1}^q C_{ir}^{bs} f_{r,z_i}(\hat{z}_{new}), \quad (9)$$

where C_{ir} are elements of the $N \times q$ coefficient matrix $\mathbf{C}^{bs}$ given by

$$\operatorname{vec}(\mathbf{C}^{bs}) = (\mathbf{K}_{bs} + \frac{N}{\gamma_{bs}} \mathbf{I})^{-1} \operatorname{vec}(\mathbf{Y}), \quad (10)$$

with $\operatorname{vec}(\cdot)$ denoting the vectorization of a matrix into a column vector. From equations (9) and (10), we observe that the value of the function $\hat{g}$ at point $\hat{z}_{new}$ requires just the knowledge of pairwise kernel values between the $N+1$ points $z_1, \dots, z_N, \hat{z}_{new}$. As such, the point estimate of the response is taken to be $\hat{y}_{new} = \hat{f}_{bs}(\hat{z}_{new})$.

In our studies, we opted to simplify the minimizer (9), in favor of having just one tuning parameter, by choosing

$$g_{bs}(r, z_a, t, z_b) = \begin{cases} \exp\left(-\frac{\|z_a - z_b\|^2}{\sigma_{bs}^2}\right), & r, t = 1, \dots, q, \quad r = t \\ 0, & r \neq t \end{cases}, \quad (11)$$

where σ_{bs} is the single free parameter, adjusting the bandwidth of the kernels. This simplification has an additional benefit in the computational efficiency of the algorithm.

We notice that the backscoring formulation does not take into explicit consideration the structure of the response manifold and $\hat{g}(\hat{z}_{new})$ is not guaranteed to lie exactly on $\mathcal{M}$. Nevertheless, we justify our choice by pointing that since $\hat{g}$ is a smooth interpolating function from the embedding to the training responses, predictions should also adhere well to the manifold.

3. Experiments

In this section, we present two simulation studies to assess the ability of our methodology to cast predictions that adhere to the response manifold. Subsequently, we use our method in two illustrative image completion applications, where the objective is to predict one half of an image from the other half.

3.1. Simulation Studies

The simulations presented below were performed with the purpose of evaluating the predictive performance of our methodology, as well as highlight its applicability in different problems. First, we employed a simple circular manifold, and utilized univariate predictors, in order to enable comparisons against an intrinsic regression method. Subsequently, we performed a study using a swiss-roll manifold, for which no analytical definition exists. At this point, intrinsic methods can not be used. For comparisons, we employed a number of standard and kernel-based regression methodologies.

3.1.1. Prediction on a Circular Manifold

In the first experiment, the goal was to compare our methodology against the non-parametric intrinsic manifold regression recently presented in [22]. In that work, the authors presented a non-parametric bayesian regression model for Riemannian manifolds using Brownian Motion (BM) priors. Two alternative BM priors were proposed, one continuous (CBM) and one discrete (DCM). Maximum a posteriori estimators were obtained through a simulated annealing (SA) algorithm on the posteriors. The random moves during SA necessitate the analytical definition of exponential maps on the manifold. Furthermore, both models require univariate predictors. A detailed description of the methodology can be found in the Supplementary Text section 1.

The methodology was demonstrated on the problem of regression from the interval $[0, 1]$ to the circle manifold $\mathcal{S}^1$. The same framework was adopted here. We simulated a training dataset of $N = 100$ paired input-response points $\{(x_i, y_i) | i = 1, \dots, N\}$, with inputs $x_i \in [0, 1]$ corresponding to responses $y_i = (y_{i1}, y_{i2}) \in \mathcal{S}^1 \subset \mathbb{R}^2$ on the circumference of a circle. In more detail, 100 input observations were sampled from the uniform distribution $\mathcal{U}(0, 1)$. These were mapped on the interval $[-\pi, \pi]$ using the function $\theta_i = f_{\mathcal{S}^1}^a(x_i) = (x_i + 0.5)^2$. Gaussian noise with $\sigma = 0.1$ was subsequently added on θ_i . Finally, the responses were computed as $y_{i1} = \cos(\theta_i)$, $y_{i2} = \sin(\theta_i)$.

Regarding the intrinsic methods of CBM and DBM, regression was performed using the set $\{\theta_i | i = 1, \dots, N\}$ as responses,

since they represent a 1-dimensional angular parametrization of $\mathcal{S}^1$. The distance metric, exponential and logarithmic maps on the manifold can be analytically computed as $d_{\mathcal{S}^1}(\theta_i, \theta_j) = |\theta_i - \theta_j|$, $\exp_{\theta_i}(\mathbf{u}_{\mathcal{T}}) = \theta_i + \mathbf{u}_{\mathcal{T}}$, $\log_{\theta_i}(\theta_j) = \theta_i - \theta_j$, where $\mathbf{u}_{\mathcal{T}} \in \mathcal{T}_{\theta_i} \mathcal{S}^1 \subset \mathbb{R}$. The standard deviation parameter for both priors was 0.1. In our methodology, the same distance function was used for training the dRF model and learning a 1-dimensional embedding. The forest was built with 1000 trees and backscoring was performed using $\sigma_{bs} = 0.5$ and $\gamma_{bs} = 10$.

The testing dataset was constructed by taking 50 equidistant points in the interval $[0, 1]$. Figure 2 (a) shows the graph of the training and test observations, based on the angular responses, while Figures 2 (b), (c) and (d) present the estimated angular responses for the test inputs, based on the CBM and DBM and dRF intrinsic regression models respectively. Since dRF casts predictions in the original representation space ($\mathbb{R}^2$), angular responses were acquired using $\theta_i = \arctan(\frac{y_{i2}}{y_{i1}})$.

Fig. 2. Scatterplots of simulated input and response points for the Circular Manifold simulation. (a) Input vs. angular train and test observations, (b) Input vs. test and predicted angular for the CBM model, (c) Input vs. test and predicted angular responses for the DBM model, (d) Input vs. test and predicted responses for the dRF model.

The DBM model suffered from oversmoothing and resulted in the highest prediction error. CBM and dRF had very similar performance, as can be seen in Figures 2 (b) and (d), with CBM giving a slightly smaller prediction error. This performance benefit though, comes with a number of significant disadvantages. Apart from the requirements for analytical response manifolds and univariate predictors, which were mentioned earlier, the continuous prior distribution used in CBM often does not have a density function [22]. On the other hand, our dRF predictive approach, whilst providing very similar predictive performance when both methods can be applied, is also free from the aforementioned limitations.

3.1.2. Prediction on a Swiss-roll Manifold

For the second simulation study, we forgo with the assumption of an analytically defined manifold. Intrinsic regression methodologies are not applicable any more, and we perform comparisons against standard and kernel-based methods.

Let S denote the simulated dataset of 900 paired input-response points $\{(x_i, y_i)_{i=1}^N\}$, with inputs $x_i \in \mathbb{R}^6$ corresponding to responses $y_i \in \mathcal{M}_{sr} \subset \mathbb{R}^3$, where $\mathcal{M}_{sr}$ denotes the 2-dimensional swiss-roll manifold embedded in $\mathbb{R}^3$. Only the first two input dimensions were built to be predictive of the

output. In detail, we first sampled 900 points $\{t_i\}_{i=1}^N$ from the uniform distribution $\mathcal{U}(\pi, 3\pi)$ and $\{u_i\}_{i=1}^N$ from $\mathcal{U}(0, 21)$. The response points on the swiss-roll were computed as $y_{i1} = t_i \cos t_i$, $y_{i2} = u_i$, $y_{i3} = t_i \sin t_i$, while the input variables x_{i1} and x_{i2} were computed by mapping u_i and t_i in the unit circle: $x_{i1} = \frac{t_i - \bar{t}_i}{\max t_i} \sqrt{1 - \frac{1}{2} \left(\frac{u_i - \bar{u}_i}{\max u_i} \right)^2}$, $x_{i2} = \frac{u_i - \bar{u}_i}{\max u_i} \sqrt{1 - \frac{1}{2} \left(\frac{t_i - \bar{t}_i}{\max t_i} \right)^2}$. The values for the remaining input coordinate variables were drawn from the standard normal Gaussian distribution. Gaussian Noise, with $\Sigma = 0.5\mathbf{I}_3$, was added on the response points.

We compared our dRF prediction methodology to k NN regression and standard RF, which do not take into consideration the special form of the response space, as well as kernel RF (kRF) [7], an RF method that employs a kernel function to capture the structure of the response space during training. The simulated dataset was split into S_{train} , comprising of 600 input-response pairs and S_{test} , consisting of the remaining 300. For k NN regression, the value of k was selected to be 5. All RFs were built with 150 trees, no pruning and the number of candidate split features in each node was set to 3. For kRF, the training gram matrix was calculated using the Gaussian kernel with $\sigma = 0.5$. We followed the prediction methodology as described in [7], with $\hat{y}_{new}^{kRF} = \arg \min_{y \in S_{train}} (g_{kRF}(y, y) - 2 \sum_{i=1}^N a_{RF}(x_{new}, x_i) g_{kRF}(y, y_i))$, where $a_{RF}(\cdot, \cdot)$ is the RF-based affinity. The minimization problem was solved over S_{train} . As such, all predicted responses were drawn from the training set.

For dRF, the backscoring parameters used were $\sigma_{bs} = 100$ and $\gamma_{bs} = 200$. For kRF, the training gram matrix was calculated using the Gaussian kernel with $\sigma = 2.5$.

Since there is no analytical form for computing distances on a swiss-roll manifold, we approximated manifold distances using the Isomap distance formulation [19]. A neighborhood graph was constructed, in which each point y_i was connected to its 7 nearest neighbors in $\mathbb{R}^q$. Graph edges were assigned weights equal to the Euclidean distances between the connected points in $\mathbb{R}^q$. For any two points y_i and y_j in S , $d_{st}(y_i, y_j)$ was then estimated by the shortest path connecting y_i and y_j .

Figure 3 shows test error vectors for the various methodologies used, projected on the $y_1 y_3$ plane. We can see that 5NN missed the goal of regressing on the manifold. For standard RF, it is clear that the model constantly underestimated the radius of the predicted points around the y_1 axis. This can be justified by considering that predictions are taken as average points on the euclidean space $\mathbb{R}^3$, unaware of a possible structure in the response space. kRF preserved the manifold structure of the predicted points better than RF, but suffered significantly from the added noise in the responses. dRF outperformed the other methods both in terms of compliance to the underlying response manifold and regarding good robustness to the addition of noise.

3.2. Applications on Image Completion

In imaging analysis, it is common to assume that a collection of similar images lies on a manifold [12]. This assumption is guided by the complex nature of images as data objects, as well as the observation that the Euclidean metric and the corresponding geometric structure that it imposes on the space do

Fig. 3. Test error vectors projected on the $y_1 y_3$ plane of the various regression models for the simulated swiss-roll dataset. (a) 5NN (b) RF (c) kRF (d) dRF. The figure highlights compliance to the response manifold.

not bode well with the human perception of similarity and difference between images [12]. Here, we used our manifold regression methodology to predict the bottom half of handwritten digits and human faces from their upper half.

3.2.1. Handwritten digits

For this application, we extracted 1000 gray-scale images of handwritten digits, from the UCI Machine Learning Repository [8]. Each digit class, from 0 to 9, was represented in the dataset by 100 8×8 pixel images. Input data were constructed by vectorization of the 8×4 upper half pixel intensities. The dataset was split into training and testing subsets with $N_{train} = 800$ and $N_{test} = 200$. The upper part of each images was taken as input for the predictive models. Responses comprised of the images' bottom parts. The test set can be seen in Fig. 4(a).

We compared predictions from 1NN, RF, kRF and dRF models. All RFs were built with 300 trees and 5 candidate split features in each node. For dRF, the distance matrix was constructed using the Isomap distance, with the number of neighbors set to 5. A 25-dimensional embedding space was used and the backscoring parameters were $\sigma_{bs} = 3$ and $\gamma_{bs} = 20$. For kRF, the training gram matrix was calculated using the Gaussian kernel with $\sigma = 5$. The reconstructed test digits for the various models are shown in Fig. 4(b)-(e).

Table 1 summarizes the prediction results for the test data. In the first column we report the test Euclidean Mean Squared Errors (EMSE). In the case of 1NN and kRF, which draw predictions from the training dataset, we are also able to report misclassification errors, a.k.a the percentage of predicted lower parts that mismatched the ground truth, which are shown on the second column of Table 1. In the last two columns, we report on the number and percentage of badly reconstructed test images after visual inspection, based on the following criteria: blurriness of the reconstructed image, smooth transition between the upper and bottom image parts, correct digit reconstruction. This qualitative approach is necessary due to the non-Euclidean nature of the data, which, as will be highlighted below, makes

Fig. 4. Test and reconstructed images from the digit completion application. Reconstructed digits were generated by concatenating the predicted responses with the upper -input- half of the test images. Bad reconstructions are enclosed in red squares. (a) Test digits. (b)-(e) Reconstructions using predictions from (b) 1NN, (c) RF, (d) kRF and (e) dRF.

the EMSE unsuitable for judging the predictive performance.

1NN and kRF cast predictions drawn directly from the training images. As such, no blurriness existed on the reconstructed digits of Fig. 4(b) and 4(d). Bad reconstructions were either misclassifications or reconstructions where the transition between the upper (input) and lower (predicted) image parts was

Table 1. Test errors from the digit completion application. Classification Error was not applicable for RF and dRF. The number and percentage of badly reconstructed images was visually ascertained from Fig. 4.

	MSE	Clas. Error	Bad Rec. No	Bad Rec. %
1NN	3.3665	0.165	40	20
RF	2.7651	-	54	27
kRF	3.2675	0.21	48	24
dRF	3.3287	-	37	18.5

not smooth. 1NN outperformed kRF in terms of classification error, as well as upon visual inspection of the images.

It is obvious, from Fig. 4(c), that standard RF is ill-suited for this application. The RF prediction approach of averaging pixel intensities from various images resulted in a high number of blurry and nonsensical digit reconstructions. RF performed the worst, based on visual assessment of the reconstructed images, but it had the lowest test EMSE, as shown in table 1. This observation highlights the unsuitability of EMSE as a measure of performance, in the case of manifold-valued responses.

Reconstructed digits from our dRF model are shown in Fig. 4(e). Our predictions were not drawn directly from the training set. Nevertheless, we notice that the large majority of reconstructions had no fuzziness and the transition from upper to lower parts was smooth. When some blurriness was noticed, its effect was significantly less severe than in the RF predictions, resulting in the overall lowest number of badly reconstructed digits, based on visual inspection.

3.2.2. Faces

In the second application, we used 400 images of faces from the Olivetti dataset, as included in the scikit-learn python package [10]. The dataset consisted of ten gray-scale 64×64 pixel images for each of 40 distinct subjects, with same subject images taken at different times and with varying pose and facial expressions. Input data were constructed by vectorization of the 64×32 upper half pixel intensities. The dataset was split into training and testing subsets with $N_{train} = 300$ and $N_{test} = 100$. Images of the same subject were only allocated either to the training or the testing set. Again, the upper parts of the images were taken as inputs, while bottom parts as responses.

We compared results from 1NN, linear Regression (LR), RF and dRF. RFs were built with 300 trees and 8 candidate split features in each node. For dRF, the training distance matrix was constructed using the Isomap metric with the number of neighbors set to 5, a 25-dimensional embedding space was used and the backscoring parameters were $\sigma_{bs} = 9$ and $\gamma_{bs} = 50$.

Four test images and their reconstructions for the various models are included in Fig. 5. As we noted in the previous application, EMSE does not provide a suitable descriptor of performance. We rely again on visual inspection of the reconstructed images. 1NN exhibited a hits and miss behavior, with some predictions being similar to the original face, whilst others, such as the second and third depicted faces, being completely different. In addition, there was minimal smoothness in the transitions from the input to the predicted parts, accentuated specifically on the nasal and zygomatic edges. LR predictions were extremely blurry. RF also suffered from a large amount

of blurriness, although transitions between the two parts of the faces looked more natural. Finally, dRF reconstructions exhibited the best transition smoothness of all methods, with the predicted half-images being well aligned to their inputs. Although the predictions were not completely free of blur artifacts, the effect was less severe and facial details, such as nasolabial folds and lips, were clearly portrayed.

Fig. 5. Example test and reconstructed images from the face completion application. (a) Test images. (b)-(e) Reconstructions using predictions from (b) 1NN, (c) LR, (d) RF and (e) dRF.

4. Discussion and Conclusion

In this work we presented a predictive modeling approach for response objects occupying non-linear manifold spaces. The regression methodology is based on a distance modification of the RF algorithm that we previously published [17], which decouples the model's training from the problem of response representation. For prediction purposes, we constructed a framework in which point estimates are first predicted on a Euclidean embedding of the response manifold, learned from the training dataset, and then projected back on the original space.

Our methodology, in contrast to intrinsic manifold methods, necessitates just the definition of a meaningful distance metric in the response space. This can be especially useful for real-life applications, such as image analysis, where the underlying manifolds are usually not known. Our distance-based regression algorithm draws similarities to the family of kernel-based methodologies. One benefit over kernel methods is the vast library of readily available distance metrics for a plethora of data objects. Furthermore, our methodology presents a unified framework, which deals with backscoring to the original response space, an issue that a lot of presented kernel methods do not tackle sufficiently well.

In the performed experiments, our method showed superior predictive performance in comparison to various regression methods, whilst being able to handle high-dimensional inputs and noise on the response observations. In the future, we aim to investigate the problem of automatic estimation of the Euclidean embedding's dimensionality, as well as the use of more elaborate kernel functions in the backscoring formulation.

Acknowledgments

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] Bengio, Y., Paiement, J.F., Vincent, P., Delalleau, O., Le Roux, N., Ouimet, M., 2004. Out-Of-Sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. *Advances in neural information processing systems* 16, 177–184.
- [2] Breiman, L., 2001. Random forests. *Machine learning* 45, 5–32.
- [3] Chavel, I., 2006. *Riemannian geometry: a modern introduction*. volume 98. Cambridge university press.
- [4] Cox, T.F., Cox, M.A., 2010. *Multidimensional scaling*. CRC Press.
- [5] Deza, M.M., Deza, E., 2009. *Encyclopedia of distances*, in: *Encyclopedia of Distances*. Springer, pp. 1–583.
- [6] Fletcher, T., 2013. Geodesic regression and its application to shape analysis, in: *Innovations for Shape Analysis*. Springer, pp. 35–52.
- [7] Geurts, P., Wehenkel, L., d'Alché Buc, F., 2006. Kernelizing the output of tree-based methods, in: *Proceedings of the 23rd international conference on Machine learning*. Acm. pp. 345–352.
- [8] Lichman, M., 2013. Uci machine learning repository. URL: <http://archive.ics.uci.edu/ml>.
- [9] Mardia, K.V., 1978. Some properties of classical multi-dimensional scaling. *Communications in Statistics-Theory and Methods* 7, 1233–1241.
- [10] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- [11] Pennec, X., 2006. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision* 25, 127–154.
- [12] Peyré, G., 2009. Manifold models for signals and images. *Computer Vision and Image Understanding* 113, 249–260.
- [13] Rahman, I.U., Drori, I., Stodden, V.C., Donoho, D.L., Schröder, P., 2005. Multiscale representations for manifold-valued data. *Multiscale Modeling & Simulation* 4, 1201–1232.
- [14] Robbin, J.W., Salamon, D.A., 2011. *Introduction to differential geometry*. ETH, Lecture Notes, preliminary version, January.
- [15] Shawe-Taylor, J., Cristianini, N., 2004. *Kernel methods for pattern analysis*. Cambridge university press.
- [16] Shi, X., Styner, M., Lieberman, J., Ibrahim, J.G., Lin, W., Zhu, H., 2009. Intrinsic regression models for manifold-valued data, in: *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2009*. Springer, pp. 192–199.
- [17] Sim, A., Tsagkraloulis, D., Montana, G., 2013. Random forests on distance matrices for imaging genetics studies. *Statistical applications in genetics and molecular biology* 12, 757–786.
- [18] Small, C., 1996. *The statistical theory of shape*.
- [19] Tenenbaum, J.B., de Silva, V.d., Langford, J.C., 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319+.
- [20] Tsochantaridis, I., Hofmann, T., Joachims, T., Altun, Y., 2004. Support vector machine learning for interdependent and structured output spaces, in: *Proceedings of the twenty-first international conference on Machine learning*. ACM. p. 104.
- [21] Wahba, G., 1992. Multivariate function and operator estimation, based on smoothing splines and reproducing kernels, in: *Santa Fe Institute studies in the sciences of complexity - Proceedings*, Addison-Wesley Publishing CO. pp. 95–95.
- [22] Wang, X., Lerman, G., 2015. Nonparametric bayesian regression on manifolds via brownian motion. *arXiv preprint arXiv:1507.06710*.
- [23] Zhu, H., Chen, Y., Ibrahim, J.G., Li, Y., Hall, C., Lin, W., 2009. Intrinsic regression models for positive-definite matrices with applications to diffusion tensor imaging. *Journal of the American Statistical Association* 104.