



# A Deep Learning Approach to Diabetic Blood Glucose Prediction

Hrushikesh N. Mhaskar<sup>1</sup>, Sergei V. Pereverzyev<sup>2</sup> and Maria D. van der Walt<sup>3\*</sup>

<sup>1</sup> Institute of Mathematical Sciences, Claremont Graduate University, Claremont, CA, United States, <sup>2</sup> Johann Radon Institute, Linz, Austria, <sup>3</sup> Department of Mathematics, Vanderbilt University, Nashville, TN, United States

We consider the question of 30-min prediction of blood glucose levels measured by continuous glucose monitoring devices, using clinical data. While most studies of this nature deal with one patient at a time, we take a certain percentage of patients in the data set as training data, and test on the remainder of the patients; i.e., the machine need not re-calibrate on the new patients in the data set. We demonstrate how deep learning can outperform shallow networks in this example. One novelty is to demonstrate how a parsimonious deep representation can be constructed using domain knowledge.

**Keywords:** deep learning, deep neural network, diffusion geometry, continuous glucose monitoring, blood glucose prediction

## OPEN ACCESS

### Edited by:

Dabao Zhang,  
Purdue University, United States

### Reviewed by:

Polina Mamoshina,  
Insilico Medicine, Inc., United States  
Ping Ma,  
University of Georgia, United States

### \*Correspondence:

Maria D. van der Walt  
maryke.thom@gmail.com

### Specialty section:

This article was submitted to  
Mathematics of Computation and  
Data Science,  
a section of the journal  
Frontiers in Applied Mathematics and  
Statistics

**Received:** 12 April 2017

**Accepted:** 28 June 2017

**Published:** 14 July 2017

### Citation:

Mhaskar HN, Pereverzyev SV and van  
der Walt MD (2017) A Deep Learning  
Approach to Diabetic Blood Glucose  
Prediction.  
Front. Appl. Math. Stat. 3:14.  
doi: 10.3389/fams.2017.00014

## 1. INTRODUCTION

Deep Neural Networks, especially of the convolutional type (DCNNs), have started a revolution in the field of artificial intelligence and machine learning, triggering a large number of commercial ventures and practical applications. There is therefore a great deal of theoretical investigations about when and why deep (hierarchical) networks perform so well compared to shallow ones. For example, Montufar et al. [1] showed that the number of linear regions that can be synthesized by a deep network with ReLU nonlinearities is much larger than by a shallow network. Examples of specific functions that cannot be represented efficiently by shallow networks have been given very recently by Telgarsky [2] and Safran and Shamir [3].

It is argued in Mhaskar and Poggio [4] that from a function approximation point of view, deep networks are able to overcome the so-called curse of dimensionality if the target function is hierarchical in nature; e.g., a target function of the form

$$h_1(h_3(h_{21}(h_{11}(x_1, x_2), h_{12}(x_3, x_4)), h_{22}(h_{13}(x_5, x_6), h_{14}(x_7, x_8))))), \quad (1)$$

where each function has a bounded gradient, can be approximated by a deep network comprising  $n$  units, organized as a binary tree, up to an accuracy  $\mathcal{O}(n^{-1/2})$ . In contrast, a shallow network that cannot take into account this hierarchical structure can yield an accuracy of at most  $\mathcal{O}(n^{-1/8})$ . In theory, if samples of the functions  $h, h_{1,2}, \dots$  are known, one can construct the networks explicitly, without using any traditional learning algorithms such as back propagation.

One way to think of the function in Equation (1) is to think of the inner functions as the features of the data, obtained in a hierarchical manner. While classical machine learning with shallow neural networks requires that the relevant features of the raw data should be selected by using domain knowledge before the learning can start, deep learning algorithms appear to select the right features automatically. However, it is typically not clear how to interpret these features. Indeed, from a mathematical point of view, it is easy to show that a structure such as Equation (1) is not unique, so that the hierarchical features cannot be defined uniquely, except perhaps in some very special examples.

In this paper, we examine how a deep network can be constructed in a parsimonious manner if we do allow domain knowledge to suggest the compositional structure of the target function as well as the values of the constituent functions. We study the problem of predicting, based on the past few readings of a continuous glucose monitoring (CGM) device, both the blood glucose (BG) level and the rate at which it would be changing 30 min after the last reading. From the point of view of diabetes management, a reliable solution to this problem is of great importance. If a patient has some warning that that his/her BG will rise or drop in the next half hour, the patient can take certain precautionary measures to prevent this (e.g., administer an insulin injection or take an extra snack) [5, 6].

Our approach is to first construct three networks based on whether a 5-min prediction, using ideas in Mhaskar et al. [7], indicates the trend to be in the hypoglycemic (0–70 mg/dL), euglycemic (70–180 mg/dL), or hyperglycemic (180–450 mg/dL) range. We then train a “judge” network to get a final prediction based on the outputs of these three networks. Unlike the tacit assumption in the theory in Mhaskar and Poggio [4], the readings and the outputs of the three constituent networks are not dense in a Euclidean cube. Therefore, we use diffusion geometry ideas in Ehler et al. [8] to train the networks in a manner analogous to manifold learning.

From the point of view of BG prediction, a novelty of our paper is the following. Most of the literature on this subject which we are familiar with does the prediction patient-by-patient; for example, by taking 30% of the data for each patient to make the prediction for that patient. In contrast, we consider the entire data for 30% of the patients as training data, and predict the BG level for the remaining 70% of patients in the data set. Thus, our algorithm transfers the knowledge learned from one data set to another, although it does require the knowledge of both the data sets to work with. From this perspective, the algorithm has similarity with the meta-learning approach by Naumova et al. [5], but in contrast to the latter, it does not require a meta-features selection.

We will explain the problem and the evaluation criterion in Section 2. Some prior work on this problem is reviewed briefly in Section 3. The methodology and algorithm used in this paper are described in Section 4. The results are discussed in Section 5. The mathematical background behind the methods described in Section 4 is summarized in Appendices A.1 and A.2.

## 2. PROBLEM STATEMENT AND EVALUATION

We use a clinical data set provided by the DirectNet Central Laboratory [9], which lists BG levels of different patients taken at 5-min intervals with the CGM device; i.e., for each patient  $p$  in the patient set  $P$ , we are given a time series  $\{s_p(t_j)\}$ , where  $s_p(t_j)$  denotes the BG level at time  $t_j$ . Our goal is to predict for each  $j$ , the level  $s_p(t_{j+m})$ , given readings  $s_p(t_j), \dots, s_p(t_{j-d+1})$  for appropriate values of  $m$  and  $d$ . For a 30-min prediction,  $m = 6$ , and we took  $d = 7$  (a sampling horizon  $t_j - t_{j-d+1}$  of 30 min has been suggested as the optimal one for BG prediction in Mhaskar et al. [7] and Hayes et al. [10]).

In this problem, numerical accuracy is not the central objective. To quantify the clinical accuracy of the considered predictors, we use the Prediction Error-Grid Analysis (PRED-EGA) [11], which has been designed especially for glucose predictors and which, together with its predecessors and variations, is by now a standard metric in the blood glucose prediction problem (see for example, [5, 12–14]). This assessment methodology records reference BG estimates paired with the BG estimates predicted for the same moments, as well as reference BG directions and rates of change paired with the corresponding estimates predicted for the same moments. As a result, the PRED-EGA reports the numbers (in percent) of Accurate (Acc.), Benign (Benign) and Erroneous (Error) predictions in the hypoglycemic, euglycemic and hyperglycemic ranges separately. This stratification is of great importance because consequences caused by a prediction error in the hypoglycemic range are very different from ones in the euglycemic or the hyperglycemic range.

## 3. PRIOR WORK

Given the importance of the problem, many researchers have worked on it in several directions. Relevant to our work is the work using a linear predictor and work using supervised learning.

The linear predictor method estimates  $s'_p(t_j)$  based on the previous  $d$  readings, and then predicts

$$s_p(t_{j+m}) \approx s_p(t_j) + (t_{j+m} - t_j)s'_p(t_j).$$

Perhaps the most classical of these is the work [15] by Savitzky and Golay, that proposes an approximation of  $s'_p(t)$  by the derivative of a polynomial of least square fit to the data  $(t_k, s_p(t_k))$ ,  $k = j, \dots, j - d + 1$ . The degree  $n$  of the polynomial acts as a regularization parameter. However, in addition to the intrinsic ill-conditioning of numerical differentiation, the solution of the least square problem as posed above involves a system of linear equations with the Hilbert matrix of order  $n$ , which is notoriously ill-conditioned. Therefore, it is proposed in Lu et al. [16] to use Legendre polynomials rather than the monomials as the basis for the space of polynomials of degree  $n$ . A procedure to choose  $n$  is given in Lu et al. [16], together with error bounds in terms of  $n$  and the estimates on the noise level in the data, which are optimal up to a constant factor for the method in the sense of the oracle inequality. A substantial improvement on this method was proposed in Mhaskar et al. [7], which we summarize in Appendix A.1. As far as we are aware, this is the state of the art in this approach in short term blood glucose prediction using linear prediction technology.

There exist several BG prediction algorithms in the literature that use a supervised learning approach. These can be divided into three main groups.

The first group of methods employ kernel-based regularization techniques to achieve prediction (for example, [5, 17] and references therein), where Tikhonov regularization is used to find the best least square fit to the data  $(t_k, s_p(t_k))$ ,  $k = j, \dots, j - d + 1$ , assuming the minimizer belongs to a reproducing kernel Hilbert space (RKHS). Of course, these methods are quite sensitive to the choice of kernel and regularization parameters. Therefore, the authors in Naumova

et al. [5, 17] develop methods to choose both the kernel and regularization parameter adaptively, or through meta-learning (“learning to learn”) approaches.

The second group consists of artificial neural network models (such as [13, 18]). In Pappada et al. [13], for example, a feed-forward neural network is designed with eleven neurons in the input layer (corresponding to variables such as CGM data, the rate of change of glucose levels, meal intake and insulin dosage), and nine neurons with hyperbolic tangent transfer function in the hidden layer. The network was trained with the use of data from 17 patients and tested on data from 10 other patients for a 75-min prediction, and evaluated using the classical Clarke Error-Grid Analysis (EGA) [19], which is a predecessor of the PRED-EGA assessment metric. Classical EGA differs from PRED-EGA in the sense that it only compares absolute BG concentrations with reference BG values (and not rates of change in BG concentrations as well). Although good results are achieved in the EGA grid in this paper, a limitation of the method is the large amount of additional input information necessary to design the model, as described above.

The third group consists of methods that utilize time-series techniques such as autoregressive (AR) models (for example, [14, 20]). In Reifman et al. [14], a tenth-order AR model is developed, where the AR coefficients are determined through a regularized least square method. The model is trained patient-by-patient, typically using the first 30% of the patient’s BG measurements, for a 30 or 60-min prediction. The method is tested on a time series containing glucose values measured every minute, and evaluation is again done through the classical EGA grid. The authors in Sparacino et al. [20] develop a first-order AR model, patient-by-patient, with time-varying AR coefficients determined through weighted least squares. Their method is tested on a time series containing glucose values measured every 3 min, and quantified using statistical metrics such as measuring the mean square of the errors. As noted in Naumova et al. [5], these methods seem to be sensitive to gaps in the input data.

## 4. METHODOLOGY IN THE CURRENT PAPER

Our proposed method represents semi-supervised learning that follows an entirely different approach from those described above. It is not a classical statistics/optimization based approach; instead, it is based on function approximation on data defined manifolds, using diffusion polynomials. In this section, we describe our deep learning method, which consists of two layers, in details.

Given the time series  $\{s_p(t_j)\}$  of BG levels at time  $t_j$  for each patient  $p$  in the patient set  $P$ , where  $t_j - t_{j-1} = 5$  min, we start by formatting the data into the form  $\{(\mathbf{x}_j, y_j)\}$ , where

$$\mathbf{x}_j = (s_p(t_{j-d+1}), \dots, s_p(t_j)) \in \mathbb{R}^d \quad \text{and} \\ y_j = s_p(t_{j+m}) \in \mathbb{R}, \quad \text{for all patients } p \in P.$$

We will use the notation

$$\mathcal{P} := \{\mathbf{x}_j = (s_p(t_{j-d+1}), \dots, s_p(t_j)) : p \in P\}.$$

We also construct the diffusion matrix from  $\mathcal{P}$ . This is done by normalizing the rows of the weight matrix  $\mathcal{W}_N^e$  in Equation (A6), following the approach in Lafon [21, pp. 33–34].

Having defined the input data  $\mathcal{P}$  and the corresponding diffusion matrix, our method proceeds as follows.

### 4.1. First Layer: Three Networks In Different Clusters

To start our first layer training, we form the training patient set  $TP$  by randomly selecting (according to a uniform probability distribution)  $M\%$  of the patients in  $P$ . The training data are now defined to be all the data  $(\mathbf{x}_j, y_j)$  corresponding to the patients in  $TP$ . We will use the notations

$$\mathcal{C} := \{\mathbf{x}_j = (s_p(t_{j-d+1}), \dots, s_p(t_j)) : p \in TP\}$$

and

$$\mathcal{C}^* := \{(\mathbf{x}_j, y_j) = ((s_p(t_{j-d+1}), \dots, s_p(t_j)), s_p(t_{j+m})) : p \in TP\}.$$

Next, we make a short term prediction  $L_{\mathbf{x}_j}(t_{j+1})$  of the BG level  $s_p(t_{j+1})$  after 5 min, for all the given measurements  $\mathbf{x}_j \in \mathcal{C}$ , by applying the linear predictor method (summarized in Section 3 and Appendix A.1). Based on these 5-min predictions, we divide the measurements in  $\mathcal{C}$  into three clusters  $\mathcal{C}_o, \mathcal{C}_e$  and  $\mathcal{C}_r$ ,

$$\mathcal{C}_o = \{\mathbf{x}_j \in \mathcal{C} : 0 \leq L_{\mathbf{x}_j}(t_{j+1}) \leq 70\} \text{ (hypoglycemia);} \\ \mathcal{C}_e = \{\mathbf{x}_j \in \mathcal{C} : 70 < L_{\mathbf{x}_j}(t_{j+1}) \leq 180\} \text{ (euglycemia);} \\ \mathcal{C}_r = \{\mathbf{x}_j \in \mathcal{C} : 180 < L_{\mathbf{x}_j}(t_{j+1}) \leq 450\} \text{ (hyperglycemia),}$$

with

$$\mathcal{C}_\ell^* = \{(\mathbf{x}_j, y_j) : \mathbf{x}_j \in \mathcal{C}_\ell\}, \quad \ell \in \{o, e, r\}.$$

The motivation of this step is to gather more information concerning the training set to ensure more accurate predictions in each BG range—as noted previously, consequences of prediction error in the separate BG ranges are very different.


In the sequel, let  $S(\Gamma, \mathbf{x}_j)$  denote the result of the method of Appendix A.2 [that is,  $S(\Gamma, \mathbf{x}_j)$  is defined by Equation (A5)], used with training data  $\Gamma$  and evaluated at a point  $\mathbf{x}_j \in \mathcal{P}$ . After obtaining the three clusters  $\mathcal{C}_o^*, \mathcal{C}_e^*$  and  $\mathcal{C}_r^*$ , we compute the three predictors

$$f_o(\mathbf{x}) := S(\mathcal{C}_o^*, \mathbf{x}), \quad f_e(\mathbf{x}) := S(\mathcal{C}_e^*, \mathbf{x}) \quad \text{and} \\ f_r(\mathbf{x}) := S(\mathcal{C}_r^*, \mathbf{x}), \quad \text{for all } \mathbf{x} \in \mathcal{P},$$

as well as the “judge” predictor, based on the entire training set  $\mathcal{C}^*$ ,

$$f_j(\mathbf{x}) := S(\mathcal{C}^*, \mathbf{x}), \quad \mathbf{x} \in \mathcal{P},$$

using the summability method of Appendix A.2 [specifically, Equation (A4)]. We remark that, as discussed in Mhaskar [22], our approximations in Equation (A5) can be computed as classical radial basis function networks, with exactly as many neurons as the number of eigenfunctions used in Equation (A3).



**FIGURE 1** | Flowchart diagram: Deep Network for BG prediction.

(As mentioned in Appendix A.2, this number is determined to ensure that the system Equation (A4) is still well conditioned.)

The motivation of this step is to decide which one of the three predictions ( $f_o, f_e$  or  $f_r$ ) is the best prediction for each datum  $\mathbf{x} \in \mathcal{P}$ . Since we do not know in advance in which blood glucose range a particular datum will result, we need to use all of the training data for the judge predictor  $f_j$  to choose the best prediction.

### 4.2. Second Layer (Judge): Final Output

In the last training layer, a final output is produced based on which  $f_\ell$ ,  $\ell \in \{o, e, r\}$  gives the best placement in the PRED-EGA grid, using  $f_j$  as the reference value. The PRED-EGA grid is constructed by using comparisons of  $f_o$  (resp.,  $f_e$  and  $f_r$ ) with the reference value  $f_j$ —specifically, it involves comparing

$$f_o(\mathbf{x}_j) \text{ (resp., } f_e(\mathbf{x}_j) \text{ and } f_r(\mathbf{x}_j)) \text{ with } f_j(\mathbf{x}_j)$$

as well as the rates of change

$$\frac{f_o(\mathbf{x}_{j+1}) - f_o(\mathbf{x}_{j-1})}{2(t_{j+1} - t_{j-1})} \left( \text{resp., } \frac{f_e(\mathbf{x}_{j+1}) - f_e(\mathbf{x}_{j-1})}{2(t_{j+1} - t_{j-1})} \text{ and } \frac{f_r(\mathbf{x}_{j+1}) - f_r(\mathbf{x}_{j-1})}{2(t_{j+1} - t_{j-1})} \right) \text{ with } \frac{f_j(\mathbf{x}_{j+1}) - f_j(\mathbf{x}_{j-1})}{2(t_{j+1} - t_{j-1})}, \quad (2)$$

for all  $\mathbf{x}_j \in \mathcal{P}$ . Based on these comparisons, PRED-EGA classifies  $f_o(\mathbf{x}_j)$  (resp.,  $f_e(\mathbf{x}_j)$  and  $f_r(\mathbf{x}_j)$ ) as being Accurate, Benign or Erroneous. As our final output  $f(\mathbf{x}_j)$  of the target function at  $\mathbf{x}_j \in \mathcal{P}$ , we choose the one among  $f_\ell(\mathbf{x}_j)$ ,  $\ell \in \{o, e, r\}$  with the best classification, and if there is more than one achieving the best grid placement, we output the one among  $f_\ell(\mathbf{x}_j)$ ,  $\ell \in \{o, e, r\}$  that has value closest to  $f_j(\mathbf{x}_j)$ . For the first and last  $\mathbf{x}_j$  for each patient  $p \in P$  (for which the rate of change Equation (2) cannot be computed), we use the one among  $f_\ell(\mathbf{x}_j)$ ,  $\ell \in \{o, e, r\}$  that has value closest to  $f_j(\mathbf{x}_j)$  as the final output.

### 4.3. Evaluation

Lastly, to evaluate the performance of the final output  $f(\mathbf{x}_j)$ ,  $\mathbf{x}_j \in \mathcal{P}$ , we use the actual reference values  $y_j$  to place  $f(\mathbf{x}_j)$  in the PRED-EGA grid.

We repeat the process described in Sections 4.1–4.3 for a fixed number of trials, after which we report the average of the PRED-EGA grid placements, over all  $\mathbf{x} \in \mathcal{P}$  and over all trials, as the final evaluation.

A summary of the method is given in Algorithm 1. A flowchart diagram of the algorithm is shown in **Figure 1**.

### Algorithm 1 Deep Network for BG prediction

**input** : Time series  $\{s_p(t_j)\}$ ,  $p \in P$ , formatted as  $\mathcal{P} = \{\mathbf{x}_j\}$  with  $\mathbf{x}_j = (s_p(t_{j-d+1}), \dots, s_p(t_j))$  and  $y_j = s_p(t_{j+m})$ , and corresponding diffusion matrix

$d \in \mathbb{N}$  (specifies sampling horizon),  $m \in \mathbb{N}$  (specifies prediction horizon)

$M \in (0, 100)$  (percentage of data used for training).

Let  $TP$  contain  $M\%$  of patients from  $P$  (drawn according to uniform prob. distr.)

Set  $\mathcal{C} = \{\mathbf{x}_j\}$  and  $\mathcal{C}^* = \{(\mathbf{x}_j, y_j)\}$  for all patients  $p \in TP$

**output**: Prediction  $f(\mathbf{x}_j) \approx s_p(t_{j+m})$  for  $\mathbf{x}_j \in \mathcal{P}$ .

*First layer:*

**for**  $\mathbf{x}_j \in \mathcal{C}$  **do**  
 | Make 5-min prediction  $L_{\mathbf{x}_j}(t_{j+1})$

**end**

Set  $\mathcal{C}_o = \{\mathbf{x}_j \in \mathcal{C} : 0 \leq L_{\mathbf{x}_j}(t_{j+1}) \leq 70\}$

Set  $\mathcal{C}_e = \{\mathbf{x}_j \in \mathcal{C} : 70 < L_{\mathbf{x}_j}(t_{j+1}) \leq 180\}$

Set  $\mathcal{C}_r = \{\mathbf{x}_j \in \mathcal{C} : 180 < L_{\mathbf{x}_j}(t_{j+1}) \leq 450\}$

Set  $\mathcal{C}_\ell^* = \{(\mathbf{x}_j, y_j) : \mathbf{x}_j \in \mathcal{C}_\ell\}$ ,  $\ell \in \{o, e, r\}$

**for**  $\mathbf{x}_j \in \mathcal{P}$  **do**

**for**  $\ell \in \{o, e, r\}$  **do**

    | Compute  $f_\ell(\mathbf{x}_j) = S(\mathcal{C}_\ell^*, \mathbf{x}_j)$

**end**

  Compute  $f_j(\mathbf{x}_j) = S(\mathcal{C}^*, \mathbf{x}_j)$

**end**

*Second layer:*

**for**  $\mathbf{x}_j \in \mathcal{P}$  **do**

**for**  $\ell \in \{o, e, r\}$  **do**

    | Construct PRED-EGA grid for  $f_\ell(\mathbf{x}_j)$  using  $f_j(\mathbf{x}_j)$  as reference value

**end**

  Let  $\ell_j \in \{o, e, r\}$  denote the subscript for which  $f_{\ell_j}(\mathbf{x}_j)$  produces the best PRED-EGA grid placement

  Set final output  $f(\mathbf{x}_j) = f_{\ell_j}(\mathbf{x}_j)$

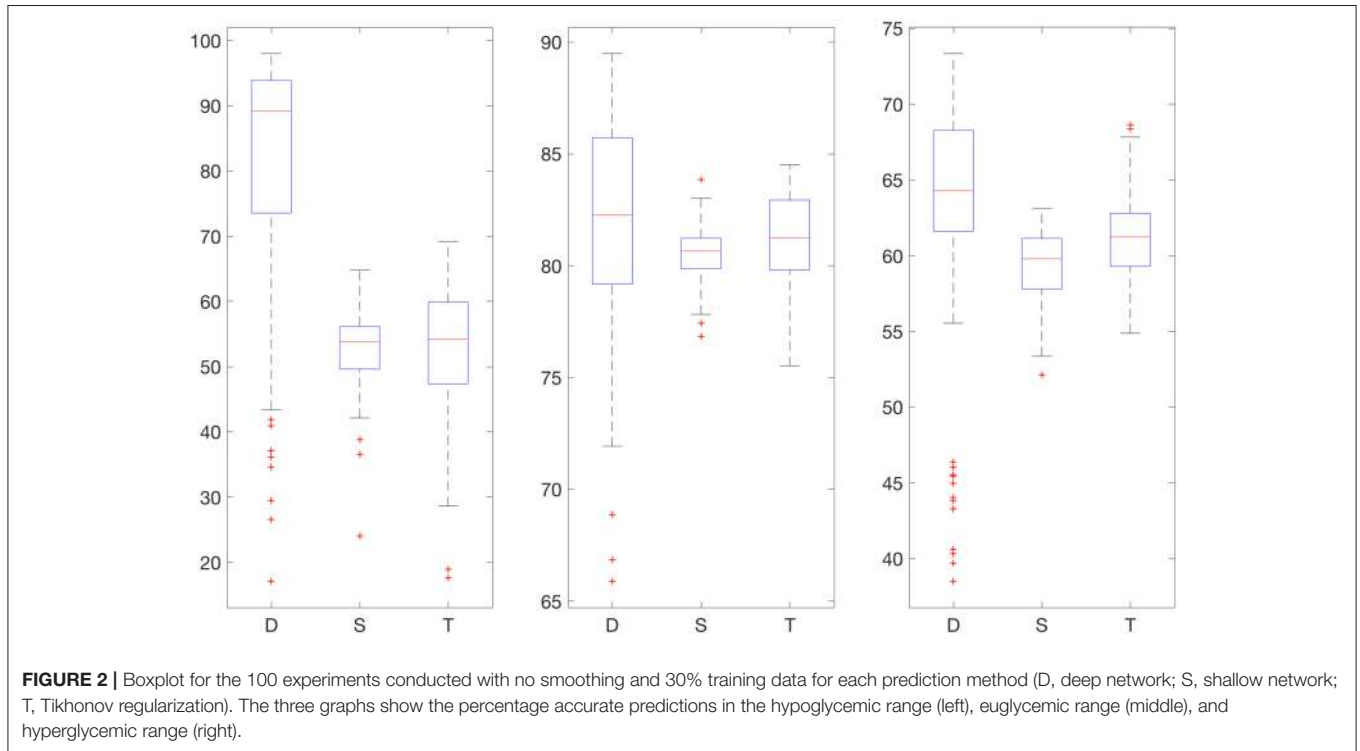
**end**

### 4.4. Remarks

1. An optional smoothing step may be applied before the first training layer step in Section 4.1 to remove any large spikes in the given time series  $\{s_p(t_j)\}$ ,  $p \in P$ , that may be caused by patient movement, for example. Following ideas in Sparacino et al. [20], we may apply flat low-pass filtering (for example, a first-order Butterworth filter). In this case, the evaluation of the final output in Section 4.3 is done by using the original, unsmoothed measurements as reference values.
2. We remark that, as explained in Section 1, the training set may also be implemented in a different way: instead of drawing  $M\%$  of the patients in  $P$  and using their entire data sets for training, we could construct the training set  $\mathcal{C}^*$  for each patient  $p \in P$  separately by drawing  $M\%$  of a given patient's data for training, and then construct the networks  $f_\ell$ ,  $\ell \in \{o, e, r\}$  and  $f_j$  for each patient separately. This is a different problem, studied often in the literature (see for

**TABLE 1** | Average PRED-EGA scores (in percent):  $M\%$  of patients used for training.

	Hypoglycemia: BG $\leq$ 70 (mg/dL)			Euglycemia: BG 70 – 180 (mg/dL)			Hyperglycemia: BG $>$ 180 (mg/dL)		
	Acc.	Benign	Error	Acc.	Benign	Error	Acc.	Benign	Error
<b>30% training data (<math>M = 30</math>):</b>									
Deep network	79.97	4.35	15.68	81.88	15.75	2.37	62.72	20.17	17.11
Shallow network	52.79	2.64	44.57	80.55	14.04	5.41	59.37	22.09	18.54
Tikhonov reg.	52.34	2.10	45.56	81.25	13.68	5.07	61.33	19.69	18.98
<b>50% training data (<math>M = 50</math>):</b>									
Deep network	88.72	4.49	6.79	80.32	17.36	2.32	64.88	21.90	13.22
Shallow network	51.84	2.47	45.69	80.94	13.77	5.29	60.41	21.58	18.01
Tikhonov reg.	52.92	1.70	45.38	81.28	13.66	5.06	62.22	20.26	17.52



**FIGURE 2** | Boxplot for the 100 experiments conducted with no smoothing and 30% training data for each prediction method (D, deep network; S, shallow network; T, Tikhonov regularization). The three graphs show the percentage accurate predictions in the hypoglycemic range (left), euglycemic range (middle), and hyperglycemic range (right).


example, [14, 20, 23]). We do not pursue this approach in the current paper.

### 5. RESULTS AND DISCUSSION

As mentioned in Section 2, we apply our method to data provided by the DirecNet Central Laboratory. Time series for 25 patients are considered. This specific data set was designed to study the performance of CGM devices in children with Type I diabetes; as such, all of the patients are less than 18 years old. Each time series  $\{s_p(t_j)\}$  contains more than 160 BG measurements, taken at 5-min intervals. Our method is a general purpose algorithm, where these details do not play any significant role, except in affecting the outcome of the experiments.

We provide results obtained by implementing our method in Matlab, as described in Algorithm 1. For our implementation, we employ a sampling horizon  $t_j - t_{j-d+1}$  of 30 min ( $d = 7$ ), a prediction horizon  $t_{j+m} - t_j$  of 30 min ( $m = 6$ ), and a total of 100 trials ( $T = 100$ ). We provide results for both 30% training data ( $M = 30$ ) and 50% training data ( $M = 50$ ) (which is comparable to approaches followed in for example [13, 14]). After testing our method on all 25 patients, the average PRED-EGA scores (in percent) are displayed in **Table 1**. For 30% training data, the percentage of accurate predictions and predictions with benign consequences is 84.32% in the hypoglycemic range, 97.63% in the euglycemic range, and 82.89% in the hyperglycemic range, while for 50% training data, we have 93.21% in the hypoglycemic range, 97.68% in the euglycemic range, and 86.78% in the hyperglycemic range.





**TABLE 2** | Average PRED-EGA scores (in percent):  $M\%$  of patients used for training with flat low-pass filtering.


	Hypoglycemia: BG $\leq$ 70 (mg/dL)			Euglycemia: BG 70 – 180 (mg/dL)			Hyperglycemia: BG $>$ 180 (mg/dL)		
	Acc.	Benign	Error	Acc.	Benign	Error	Acc.	Benign	Error
<b>30% training data (<math>M = 30</math>):</b>									
Deep network	86.41	2.51	11.08	85.05	12.59	2.36	62.24	15.34	22.42
Shallow network	61.10	5.21	33.69	81.96	12.77	5.27	60.01	19.62	20.37
Tikhonov reg.	57.47	2.01	40.52	83.49	12.00	4.51	62.13	19.13	18.74
<b>50% training data (<math>M = 50</math>):</b>									
Deep network	94.39	2.04	3.57	83.44	14.52	2.04	67.21	18.08	14.71
Shallow network	61.49	5.50	33.01	82.16	12.59	5.25	60.50	19.21	20.29
Tikhonov reg.	59.02	1.94	39.04	83.56	11.95	4.49	62.34	19.55	18.11

For comparison, **Table 1** also displays the PRED-EGA scores obtained when implementing a shallow (ie, one layer) feed-forward network with 20 neurons using Matlab’s Neural Network Toolbox, using the same parameters  $d, m, M$  and  $T$  as in our implementation, with Matlab’s default hyperbolic tangent sigmoid activation function. The motivation for using 20 neurons is the following. As mentioned in Section 4.1, each layer in our deep network can be viewed as a classical neural network with exactly as many neurons as the number of eigenfunctions used in Equations (A9) and (A11). This number is determined to ensure that the system in Equation (A10) is well conditioned; in our experiments, it turned out to be at most 5. Therefore, to compare our two-layer deep network (where the first layer consists of three separate networks) with a classical shallow neural network,


we use 20 neurons. It is clear that our deep network performs substantially better than the shallow network; in all three BG ranges, our method produces a lower percentage of erroneous predictions.

For a further comparison, we also display in **Table 1** the PRED-EGA scores obtained when training through supervised learning using a standard Tikhonov regularization to find the best least square fit to the data; specifically, we implemented the method described in Poggio and Smale [24, pp. 3–4] using a Gaussian kernel with  $\sigma = 100$  and regularization constant  $\gamma = 0.0001$ . Again, our method produces superior results, especially in the hypoglycemic range.

**Figures 2, 3** display boxplots for the percentage accurate predictions in each BG range for each method over the 100 trials.



**FIGURE 4 |** Boxplot for the 100 experiments conducted with flat low-pass filtering and 30% training data for each prediction method (D, deep network; S, shallow network; T, Tikhonov regularization). The three graphs show the percentage accurate predictions in the hypoglycemic range (left), euglycemic range (middle) and hyperglycemic range (right).



**FIGURE 5 |** Boxplot for the 100 experiments conducted with flat low-pass filtering and 50% training data for each prediction method (D, deep network; S, shallow network; T, Tikhonov regularization). The three graphs show the percentage accurate predictions in the hypoglycemic range (left), euglycemic range (middle) and hyperglycemic range (right).

We also provide results obtained when applying smoothing through flat low-pass filtering to the given time series, as explained in Section 4.4. For our implementations, we used a first-order Butterworth filter with cutoff frequency 0.8, with

the same input parameters as before. The results are given in Table 2. For 30% training data, the percentage of accurate predictions and predictions with benign consequences is 88.92% in the hypoglycemic range, 97.64% in the euglycemic range,

and 77.58% in the hyperglycemic range, while for 50% training data, we have 96.43% in the hypoglycemic range, 97.96% in the euglycemic range, and 85.29% in the hyperglycemic range. **Figures 4, 5** display boxplots for the percentage accurate predictions in each BG range for each method over the 100 trials.

In all our experiments the percentage of erroneous predictions is substantially smaller with deep networks than the other two methods we have tried, with the only exception of hyperglycemic range with 30% training data and flat low pass filtering. Our method's performance also improves when the amount of the training data increases from 30% to 50%, because the percentage of erroneous predictions decreases in all of the BG ranges in all experiments.

Moreover, from this viewpoint, our deep learning method outperforms the considered competitors, except in the hyperglycemic BG range in **Table 2**. A possible explanation for this is that the DirecNet study was done on children (less than 18 years of age) with type 1 diabetes, for a period of roughly 26 h. It appears that children usually have prolonged hypoglycemic periods as well as profound postprandial hyperglycemia (high blood sugar “spikes” after meals)—according to Boland et al. [25], more than 70% of children display prolonged hypoglycemia while more than 90% of children display significant postprandial hyperglycemia. In particular, many of the patients in the data set only exhibit very limited hyperglycemic BG spikes—for the patient labeled 43, for example, there exists a total of 194 BG measurements, of which only 5 measurements are greater than 180 mg/dL. This anomaly might have affected the performance of our algorithm in the hyperglycemic range. Obviously, it performs remarkably better than the other techniques we have tried, including fully supervised training, while ours is only semi-supervised.

## 6. CONCLUSION

The prediction of blood glucose levels based on continuous glucose monitoring system data 30 min ahead is a very

important problem with many consequences for the health care industry. In this paper, we suggest a deep learning paradigm based on a solid mathematical theory as well as domain knowledge to solve this problem accurately as assessed by the PRED-EGA grid, developed specifically for this purpose. It is demonstrated in Mhaskar and Poggio [4] that deep networks perform substantially better than shallow networks in terms of expressiveness for function approximation when the target function has a compositional structure. Thus, the blessing of compositionality cures the curse of dimensionality. However, the compositional structure is not unique, and it is open problem to decide whether a given target function has a certain compositional structure. In this paper, we have demonstrated an example where domain knowledge can be used to build an appropriate compositional structure, leading to a parsimonious deep learning design.

## ETHICS STATEMENT

The clinical data set used in this publication is publicly available online [9]. The data were collected during the study that was carried out in accordance with the recommendations of DirecNet, Jaeb Center for Health Research, with written informed consent from all 878 subjects. All subjects gave written informed consent in accordance with the Declaration of Helsinki. The protocol was approved by DirecNet, Jaeb Center for Health Research.

## AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

## ACKNOWLEDGMENTS

The research of HM is supported in part by ARO Grant W911NF-15-1-0385. The research of SP is partially supported by Austrian Science Fund (FWF) Grant I 1669-N26.

## REFERENCES

- Montufar GF, Pascanu R, Cho K, Bengio Y. On the number of linear regions of deep neural networks. In: *Advances in Neural Information Processing Systems*. Red Hook, NY: Curran Associates, Inc. (2014) p. 2924–32.
- Telgarsky M. Representation benefits of deep feedforward networks. *arXiv preprint* (2015) arXiv:150908101
- Safran I, Shamir O. Depth separation in relu networks for approximating smooth non-linear functions. *arXiv preprint* (2016) arXiv:161009887
- Mhaskar HN, Poggio T. Deep vs. shallow networks: an approximation theory perspective. *Anal Appl.* (2016) **14**:829–48. doi: 10.1142/S0219530516400042
- Naumova V, Pereverzyev SV, Sivananthan S. A meta-learning approach to the regularized learning - Case study: blood glucose prediction. *Neural Netw.* (2012) **33**:181–93. doi: 10.1016/j.neunet.2012.05.004
- Snetselaar L. *Nutrition Counseling Skills for the Nutrition Care Process*. Sudbury, MA: Jones & Bartlett Learning (2009).
- Mhaskar HN, Naumova V, Pereverzyev SV. Filtered Legendre expansion method for numerical differentiation at the boundary point with application to blood glucose predictions. *Appl Math Comput.* (2013) **224**:835–47. doi: 10.1016/j.amc.2013.09.015
- Ehler M, Filbir F, Mhaskar HN. Locally learning biomedical data using diffusion frames. *J Comput Biol.* (2012) **19**:1251–64. doi: 10.1089/cmb.2012.0187
- DirecNet Central Laboratory. (2005). Available online at: <http://direcnet.jaeb.org/Studies.aspx>
- Hayes AC, Mastrototaro JJ, Moberg SB, Mueller JC Jr, Clark HB, Tolle MCV, et al. Algorithm sensor augmented bolus estimator for semi-closed loop infusion system. Google Patents (2009). US Patent 7,547,281.
- Sivananthan S, Naumova V, Man CD, Facchinetti A, Renard E, Cobelli C, et al. Assessment of blood glucose predictors: the prediction-error grid analysis. *Diabet Technol Ther.* (2011) **13**:787–96. doi: 10.1089/dia.2011.0033
- Naumova V, Pereverzyev SV, Sivananthan S. Adaptive parameter choice for one-sided finite difference schemes and its application in diabetes technology. *J Complex.* (2012) **28**:524–38. doi: 10.1016/j.jco.2012.06.001
- Pappada SM, Cameron BD, Rosman PM, Bourey RE, Papadimos TJ, Olorunto W, et al. Neural network-based real-time prediction of glucose in patients



- with insulin-dependent diabetes. *Diabet Technol Ther.* (2011) **13**:135–41. doi: 10.1089/dia.2010.0104
14. Reifman J, Rajaraman S, Gribok A, Ward WK. Predictive monitoring for improved management of glucose levels. *J Diabet Sci Technol.* (2007) **1**:478–86. doi: 10.1177/193229680700100405
  15. Savitzky A, Golay MJE. Smoothing and differentiation of data by simplified least squares procedures. *Anal Chem.* (1964) **36**:1627–39. doi: 10.1021/ac60214a047
  16. Lu S, Naumova V, Pereverzev SV. Legendre polynomials as a recommended basis for numerical differentiation in the presence of stochastic white noise. *J Inverse ILL Posed Prob.* (2013) **21**:193–216. doi: 10.1515/jip-2012-0050
  17. Naumova V, Pereverzyev SV, Sivananthan S. Extrapolation in variable RKHSs with application to the blood glucose reading. *Inverse Prob.* (2011) **27**:075010. doi: 10.1088/0266-5611/27/7/075010
  18. Pappada SM, Cameron BD, Rosman PM. Development of a neural network for prediction of glucose concentration in type 1 diabetes patients. *J Diabet Sci Technol.* (2008) **2**:792–801. doi: 10.1177/193229680800200507
  19. Clarke WL, Cox D, Gonder-Frederick LA, Carter W, Pohl SL. Evaluating clinical accuracy of systems for self-monitoring of blood glucose. *Diabetes Care* (1987) **10**:622–8. doi: 10.2337/diacare.10.5.622
  20. Sparacino G, Zanderigo F, Corazza S, Maran A, Facchinetti A, Cobelli C. Glucose concentration can be predicted ahead in time from continuous glucose monitoring sensor time-series. *IEEE Trans Biomed Eng.* (2007) **54**:931–7. doi: 10.1109/TBME.2006.889774
  21. Lafon S. *Diffusion Maps and Geometric Harmonics*. New Haven, CT: Yale University (2004).
  22. Mhaskar HN. Eignets for function approximation on manifolds. *Appl Comput Harmon Anal.* (2010) **29**:63–87. doi: 10.1016/j.acha.2009.08.006
  23. Eren-Oruklu M, Cinar A, Quinn L, Smith D. Estimation of future glucose concentrations with subject-specific recursive linear models. *Diabet Technol Ther.* (2009) **11**:243–53. doi: 10.1089/dia.2008.0065
  24. Poggio T, Smale S. The mathematics of learning: dealing with data. *Notices AMS* (2003) **50**:537–44. Available online at: <http://www.ams.org/notices/200305/fea-smale.pdf>
  25. Boland E, Monsod T, Delucia M, Brandt CA, Fernando S, Tamborlane WV. Limitations of conventional methods of self-monitoring of blood glucose lessons learned from 3 days of continuous glucose sensing in pediatric patients with type 1 diabetes. *Diabetes Care* (2001) **24**:1858–62. doi: 10.2337/diacare.24.11.1858
  26. Maggioni M, Mhaskar HN. Diffusion polynomial frames on metric measure spaces. *Appl Comput Harmon Anal.* (2008) **24**:329–53. doi: 10.1016/j.acha.2007.07.001
  27. Belkin M, Niyogi P. Towards a theoretical foundation for Laplacian-based manifold methods. In: *International Conference on Computational Learning Theory*. Berlin: Springer (2005). p. 486–500. doi: 10.1007/11503415\_33
  28. Singer A. From graph to manifold Laplacian: the convergence rate. *Appl Comput Harmon Anal.* (2006) **21**:128–34. doi: 10.1016/j.acha.2006.03.004

**Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Mhaskar, Pereverzyev and van der Walt. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

## A. APPENDIX

### A.1. Filtered Legendre Expansion Method

In this appendix, we review the mathematical background for the method developed in Mhaskar et al. [7] for short term blood glucose prediction. As explained in the text, the main mathematical problem can be summarized as that of estimating the derivative of a function at the end point of an interval, based on measurements of the function in the past. Mathematically, if  $f : [-1, 1] \rightarrow \mathbb{R}$  is continuously differentiable, we wish to estimate  $f'(1)$  given the noisy values  $\{y_j = f(t_j) + \epsilon_j\}$  at points  $\{t_j\}_{j=1}^d \subset [-1, 1]$ . We summarize only the method here, and refer the reader to Mhaskar et al. [7] for the detailed proof of the mathematical facts.

For this purpose, we define the Legendre polynomials recursively by

$$P_k(x) = \frac{2k-1}{2k}xP_{k-1}(x) - \frac{k-1}{k}P_{k-2}(x), \quad k = 2, 3, \dots; \\ P_0(x) = 1, \quad P_1(x) = x.$$

Let  $h : \mathbb{R} \rightarrow \mathbb{R}$  be an even, infinitely differentiable function, such that  $h(t) = 1$  if  $t \in [0, 1/2]$  and  $h(t) = 0$  if  $t \geq 1$ . We define

$$K_n(h; x) = \frac{1}{2} \sum_{k=0}^{n-1} h\left(\frac{k}{n}\right) k(k+1/2)(k+1)P_k(x).$$

Next, given the points  $\{t_j\}_{j=1}^d \subset [-1, 1]$ , we use least squares to find  $n = n_d$  such that

$$\sum_{j=1}^n w_j P_k(t_j) = \int_{-1}^1 P_k(t) dt = \begin{cases} 2, & \text{if } k = 0, \\ 0, & \text{if } 1 \leq k < 2n, \end{cases} \quad (A1)$$

and the following estimates are valid for all polynomials of degree  $< 2n$ :

$$\sum_{j=1}^n |w_j P(t_j)| \leq A \int_{-1}^1 |P(t)| dt,$$

for some positive constant  $A$ . We do not need to determine  $A$ , but it is guaranteed to be proportional to the condition number of the system in Equation (A1).

Finally, our estimate of  $f'(1)$  based on the values  $y_j = f(t_j) + \epsilon_j$ ,  $j = 1, \dots, d$ , is given by

$$S_n(h; \{y_j\}) = \sum_{j=1}^d w_j y_j K_n(h; t_j).$$

It is proved in Mhaskar et al. [7, Theorem 3.2] that if  $f$  is twice continuously differentiable,

$$\Delta_{[-1,1]}(f)(x) = 2xf'(x) - (1-x^2)f''(x),$$

and  $|\epsilon_j| \leq \delta$  then

$$|S_n(h; \{y_j\}) - f'(1)| \leq cA \{E_{n/2,[-1,1]}(\Delta_{[-1,1]}(f)) + n^2\delta\},$$

where

$$E_{n/2,[-1,1]}(\Delta_{[-1,1]}(f)) = \min \max_{x \in [-1,1]} |\Delta_{[-1,1]}(f)(x) - P(x)|, \quad (A2)$$

the minimum being over all polynomials  $P$  of degree  $< n/2$ .

### A.2. Function Approximation on Data Defined Spaces

While classical approximation theory literature deals with function approximation based on data that is dense on a known domain, such as a cube, torus, sphere, etc., this condition is generally not satisfied in the context of machine learning. For example, the set of vectors  $(s_p(t_j), \dots, s_p(t_j - d + 1))$  is unlikely to be dense in a cube in  $\mathbb{R}^d$ . A relatively recent idea is to think of the data as being sampled from a distribution on an unknown manifold, or more generally, a locally compact, quasi-metric measure space. In this appendix, we review the theoretical background that underlies our experiments reported in this paper. The discussion here is based mainly on Maggioni and Mhaskar [26], Mhaskar [22], and Ehler et al. [8].

Let  $\mathbb{X}$  be a locally compact quasi-metric measure space, with  $\rho$  being the quasi-metric and  $\mu^*$  being the measure. In the context of machine learning,  $\mu^*$  is a probability measure and  $\mathbb{X}$  is its support. The starting point of this theory is a non-decreasing sequence of numbers  $\{\lambda_k\}_{k=0}^\infty$ , such that  $\lambda_0 = 0$  and  $\lambda_k \rightarrow \infty$  as  $k \rightarrow \infty$ , and a corresponding sequence of bounded functions  $\{\phi_k\}_{k=0}^\infty$  that forms an orthonormal sequence in  $L^2(\mu^*)$ . Let  $\Pi_\lambda = \text{span}\{\phi_k : \lambda_k < \lambda\}$ , and analogously to Equation (A2), we define for a uniformly continuous, bounded function  $f : \mathbb{X} \rightarrow \mathbb{R}$ ,

$$E_\lambda(f) = \min_{P \in \Pi_\lambda} \sup_{x \in \mathbb{X}} |f(x) - P(x)|.$$

With the function  $h$  as defined in Appendix A.1, we define

$$\Phi_\lambda(h; x, y) = \sum_{k:\lambda_k < \lambda} h\left(\frac{\lambda_k}{\lambda}\right) \phi_k(x)\phi_k(y), \quad x, y \in \mathbb{X}. \quad (A3)$$

Next, let  $\{x_j\}_{j=1}^M \subset \mathbb{X}$  be the ‘‘training data’’. Our goal is to learn a function  $f : \mathbb{X} \rightarrow \mathbb{R}$  based on the values  $\{f(x_j)\}_{j=1}^M$ . Toward this goal, we solve an under-determined system of equations

$$\sum_{j=1}^M W_j \phi_k(x_j) = \begin{cases} 1, & \text{if } k = 0, \\ 0, & \text{if } k > 0, \end{cases} \quad k : \lambda_k < \lambda \quad (A4)$$

for the largest possible  $\lambda$  for which this system is still well conditioned. We then define an approximation, analogous to the classical radial basis function networks, by

$$\sigma_\lambda(h; f, x) = \sum_{j=1}^M W_j f(x_j) \Phi_\lambda(x, x_j), \quad x \in \mathbb{X}. \quad (A5)$$

It is proved in Maggioni and Mhaskar [26], Mhaskar [22], and Ehler et al. [8] that under certain technical conditions, for a uniformly continuous, bounded function  $f : \mathbb{X} \rightarrow \mathbb{R}$ , we have

$$E_\lambda(f) \leq \sup_{x \in \mathbb{X}} |f(x) - \sigma_\lambda(h; f, x)| \leq cE_{\lambda/2}(f),$$

where  $c > 0$  is a generic positive constant. Thus, the approximation  $\sigma_\lambda(f)$  is guaranteed to yield a “good approximation” that is asymptotically best possible given the training data.

In practice, one has a point cloud  $\mathcal{P} = \{y_i\}_{i=1}^N \supset \{x_j\}_{j=1}^M$  rather than the space  $\mathbb{X}$ . The set  $\mathcal{P}$  is a subset of some Euclidean space  $\mathbb{R}^D$  for a possibly high value of  $D$ , but is assumed to lie on a compact sub-manifold  $\mathbb{X}$  of  $\mathbb{R}^D$ , with this manifold having a low dimension. We build the graph Laplacian from  $\mathcal{P}$  as follows: With a parameter  $\varepsilon > 0$ , induces the weight matrix  $\mathcal{W}_N^\varepsilon$  defined by

$$\mathcal{W}_{N;i,j}^\varepsilon = \exp(-|y_i - y_j|^2/\varepsilon). \quad (\text{A6})$$

We build the diagonal matrix  $D_{N;i,i}^\varepsilon = \sum_{j=1}^N \mathcal{W}_{N;i,j}^\varepsilon$  and define the (unnormalized) graph Laplacian as

$$L_N^\varepsilon = \mathcal{W}_N^\varepsilon - D_N^\varepsilon.$$

Various other versions of this Laplacian are also used in practice. For  $M \rightarrow \infty$  and  $\varepsilon \rightarrow 0$ , the eigenvalues and interpolations of the eigenvectors of  $L_M^\varepsilon$  converge toward the “interesting” eigenvalues  $\lambda_k$  and eigenfunctions  $\phi_k$  of a differential operator  $\Delta_{\mathbb{X}}$ . This behavior has been studied in detail by many authors, e.g., Belkin and Niyogi [27], Lafon [21], and Singer [28]. When  $\{y_i\}_{i=1}^M$  are uniformly distributed on  $\mathbb{X}$ , this operator is the Laplace-Beltrami operator on  $\mathbb{X}$ . If  $\{y_i\}_{i=1}^M$  are distributed according to the density  $p$ , then the graph Laplacian approximates the elliptic Schrödinger type operator  $\Delta + \frac{\Delta p}{p}$ , whose eigenfunctions  $\phi_k$  also form an orthonormal basis for  $L_2(\mathbb{X}, \mu)$ .