



Universidad Autónoma
de Madrid

Biblos-e Archivo
Repositorio Institucional UAM

Repositorio Institucional de la Universidad Autónoma de Madrid
<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:
This is an **author produced version** of a paper published in:

Engineering Applications of Artificial Intelligence 93 (2020): 103682

DOI: <https://doi.org/10.1016/j.engappai.2020.103682>

Copyright: © 2020.

This manuscript version is made available under the CC-BY-NC-ND 4.0
licence <http://creativecommons.org/licenses/by-nc-nd/4.0/>

El acceso a la versión del editor puede requerir la suscripción del recurso
Access to the published version may require subscription

Auto-adaptive multi-scale Laplacian Pyramids for modeling non-uniform data

Ángela Fernández^{a,*}, Neta Rabin^c, Dalia Fishelov^d, José R. Dorronsoro^{a,b}

^a*Departamento de Ingeniería Informática, Universidad Autónoma de Madrid, Spain*

^b*Instituto de Ingeniería del Conocimiento, Madrid, Spain*

^c*Department of Industrial Engineering, Tel-Aviv University, Israel*

^d*Department of Mathematics, Afeka Academic College of Engineering, Israel*

Abstract

Kernel-based techniques have become a common way for describing the local and global relationships of data samples that are generated in real-world processes. In this research, we focus on a multi-scale kernel based technique named Auto-adaptive Laplacian Pyramids (ALP). This method can be useful for function approximation and interpolation. ALP is an extension of the standard Laplacian Pyramids model that incorporates a modified Leave-One-Out Cross Validation procedure, which makes the method stable and automatic in terms of parameters selection without extra cost. This paper introduces a new algorithm that extends ALP to fit datasets that are non-uniformly distributed. In particular, the optimal stopping criterion will be point-dependent with respect to the local noise level and the sample rate. Experimental results over real datasets highlight the advantages of the proposed multi-scale technique for modeling and learning complex, high dimensional data.

Keywords: Laplacian Pyramids, kernel methods, overfitting, multi-scale interpolation, non-uniform data, adaptive stopping

*Corresponding author

Email addresses: a.fernandez@uam.es (Ángela Fernández), netara@tauex.tau.ac.il (Neta Rabin), fishelov@gmail.com (Dalia Fishelov), jose.dorronsoro@uam.es (José R. Dorronsoro)

1. Introduction

An important challenge nowadays, when large amounts of data is collected, is the correct approximation of functions for modeling and analyzing data. These approximations have special interest in cases where the values of the functions are not known over the entire dataset. They can also be useful when the function may be too expensive to compute, or it is only represented on a finite expansion. There exist several methods for modeling and analyzing data, but when dealing with functions that depend on multiple variables, or that are defined over many scattered data points, the best way to tackle the general problem of approximation and interpolation is using an approach based on Radial Basis Functions (RBFs) [1].

An RBF is defined in terms of an univariate continuous function ϕ . Given a training sample $\{x_n\}_{n=1}^N$, a linear combination of RBFs approximates a real function f over a new data point x in the following way:

$$f(x) \approx s(x) = \sum_{\xi} w_{\xi} \phi(\|x - x_{\xi}\|).$$

Here, w_{ξ} represents the weight associated with the expansion points ξ , and $\|\cdot\|$ is an adequate norm. The most usual choice for the norm is the Euclidean distance, and one of the most common RBFs are Gaussians, where ϕ is defined as $\phi(r) = \exp(-\epsilon r^2)$.

There exists a comprehensive literature on RBF methods and expansions (see for example [1, 2, 3, 4]). In this work we will focus on Laplacian Pyramids (LP), a multi-scale model that generates a smoothed version of a function in an iterative manner, using Gaussian kernels of decreasing widths [5]. It is a simple method for learning functions from a general set of samples. The LP approximation algorithm works in the spirit of wavelets, as the reconstruction goes from coarser to finer scales and due to this fact, it is stable and convenient for working in the manifold learning context. It is also remarkable that this method can be seen as an iterative version of a Nadaraya–Watson estimator [6, 7]. This classic

estimator is typically defined as

$$g(x) = \frac{1}{n} \sum_n K(x, x_n) y_n,$$

where K is a kernel function and y_n are the objective function values on the training points x_n . As we shall see, the proposed LP procedure uses this type of construction at each iteration.

The LP scheme has been used for several applications in diverse domains; most of them utilize LP for function approximation and its out-of-sample extension. When a model does not allow a direct out-of-sample solution, the extension of the model to new points is not trivial [8, 9, 10], and the LP method offers one way for extension of such models. In [11], an multi-scale anomaly detection algorithm that is based on diffusion maps [12] was proposed. The diffusion maps embedding was calculated on a subset of the points and extended to the rest of the dataset with the LP function extension scheme. LP based out-of-sample extension for target detection was presented in [13]. Extensions of the anomaly detection algorithm [13], which utilizes LP for extension, include anomaly detection in side-scan sonar images of sea-mines [14] and detection of defects in wafers [15]. LP was utilized for function extension in problems related to voice activity detection. In [16], the likelihood ratio function of unlabeled data was learned by extending the likelihood ratios obtained from the labeled data. An LP-based speech enhancement algorithm was proposed in [17]. The LP-based extension was compared with the Geometric harmonics extension scheme [18], for which parameters need to be carefully tuned, and it was shown that the LP-scheme provides better results.

Another domain in which Laplacian Pyramids have been applied is for data lifting. This challenge arises in models that first reduce the dimension of the data to obtain a compact and reliable representation, and then need to estimate new points in the ambient space from the low-dimensional embedding. LP-based lifting was applied in [19, 20] for reconstruction of data in molecular dynamics applications and for modeling chemical kinetics. Last, the LP method has been applied in the context of kernel based forecasting in dynamical systems such

as prediction of the North Pacific climate variability [21], prediction of regional
45 and pan-Arctic sea ice anomalies [22] and forecasting of tropical intraseasonal
oscillations [23]. In all of the above applications, heuristic approaches were used
in order to limit the number of iterations of the LP algorithm to avoid a risk of
overfitting if too many iterations are executed.

As mentioned above, and as it is often the case in machine learning, when
50 an LP model is applied, one may overfit the data by refining the prediction too
much during the training phase [20]. In fact, it is difficult to decide when to stop
the training phase to maximize the generalization capabilities of the resultant
model. A usual approach is to apply the Cross Validation (CV) method [24,
chap. 9] to measure a validation error during the training in order to stop when
55 this error starts to increase. An extreme form of CV is Leave-One-Out CV
(LOOCV): a model is built using all the samples but one, which is then used
as a single validation pattern; this is repeated for each sample in the dataset,
and the validation error is the average of all the one-pattern errors. Although
LOOCV has a theoretical support and often yields good results, it has the
60 drawback of being a time-consuming process.

Auto-adaptive LP (ALP), proposed in [25], is a modification of the LP train-
ing algorithm that merges training and an approximate LOOCV in one single
phase. The ALP algorithm results in a LOOCV approximation that does not
add any cost during the training step. This reduces significantly the training
65 complexity and provides an automatic global criterion to stop training. Thus,
the risk of overfitting, which may appear in a standard LP, is avoided. There-
fore, ALP prevents overfitting the data and, moreover, it requires essentially no
parametrization or expert knowledge about the problem under study, while still
achieving a good test error. Moreover, it adds no extra cost compared to other
70 classical neighbor-based interpolation methods.

In this paper we propose a new implementation of the ALP algorithm, pro-
viding a natural improvement of it. The main idea is to work with a local
(point-wise) kernel scale that better suited to the density of the data and, con-
sequently, to perform the optimal number of training iterations around each

75 point. This modification allows us to deal with datasets where sample densities vary in different regions, which may require a different resolution. The algorithm will be automatically adapted to each of these areas when necessary. The proposed method can also be seen as a variant of the iterative Nadaraya–Watson regression with L_2 boosting [26].

80 To sum up, the contribution of this paper is twofold. On one hand, it presents a self-contained explanation about Laplacian Pyramids, including a complete analysis of the error bounds and decay rates. Moreover we review in detail the automatic stopping criteria which is integrated into the algorithm without extra computational cost, prevents overfitting and bypasses the need
85 for heuristic approaches to set the parameters of the method. On the other hand, the second contribution is an extension of the ALP algorithm to a local resolution setting, taking advantage of the different sample statistics that might appear in the data.

This paper is organized as follows. In Section 2 we briefly review the LP
90 model and present a detailed analysis of its training error. We describe classical ALP and its LOOCV estimation in Section 3, and an improved ALP version with local resolution is presented. The algorithm description is accompanied by a synthetic example to illustrate its behavior. Results over several datasets are shown in Section 4, and the paper ends with some conclusions in Section 5.

95 2. Laplacian Pyramids

Laplacian Pyramids (LP) is an iterative model introduced by Burt and Adelson [5] for image processing applications. In its traditional form, the LP algorithm decomposes the input image into a series of sub-images, and each of them captures a different frequency band of the original one. This process is carried out by constructing Gaussian kernel-based smoothing masks of different
100 widths, followed by a down-sampling (quantization) step. LP was later proved to be a tight frame (see Do and Vetterli [27]) and used for signal processing applications, for example as a reconstruction scheme in [28]. In [29], a multi-scale

algorithm was introduced in the spirit of LP to be applied in the setting of high-dimensional data analysis. In particular, it was proposed as a simple method for
105 extending low-dimensional embedding coordinates, that result from the application of a non-linear dimensionality reduction technique, to a high-dimensional dataset (this has been recently applied in [30]).

2.1. The Basic LP Procedure

Next, we review the LP procedure as described in [29] (note that the down-sampling step, which is part of Burt and Adelson's algorithm is skipped here). Let $S = \{(x_i, f_i = f(x_i))\}_{i=1}^N, x_i \in \mathbb{R}^M$ be the sample dataset where f is a function which is only known on the sample points. For simplicity we assume that it belongs to a Sobolev space [31] $\|f\|_{m,2}$ for a certain m . The algorithm approximates the function f by constructing a series of functions $\{\tilde{f}^{(\ell)}\}$ obtained by several refinements $d^{(\ell)}$ over the approximation errors. In a slight abuse of notation we will use the same name f for both the general function $f(x)$ and also for the vector of its sample values $f = (f_1 = f(x_1), \dots, f_N = f(x_N))$. The end result of this process yields a function approximation to f in the form

$$f \simeq \tilde{f} = \tilde{f}^{(0)} + d^{(1)} + d^{(2)} + d^{(3)} + \dots$$

In more detail, a first level kernel $K^{(0)}(x, x') = \Phi(\text{dist}(x, x')/\sigma)$ is chosen using a one dimensional positive function $\Phi(z)$ and a wide, initial scale σ ; $\text{dist}(x, x')$ denotes some distance function between points in the ambient space. As mentioned before, the Gaussian kernel with Euclidean distances is applied here, i.e., we take $\text{dist}(x, x') = \|x - x'\|$. Then, we define

$$K^{(0)}(x, x') = \kappa^{(0)} e^{-\frac{\|x - x'\|^2}{\sigma^2}},$$

110 where $\kappa^{(0)}$ is the Gaussian kernel normalizing constant which depends on σ .

The notation $K^{(0)}$ is used (with a slight abuse of notation) for the general continuous kernel $K^{(0)}(x, x')$ and its discrete matrix counterpart $K_{jk}^{(0)} = K^{(0)}(x_j, x_k)$ over the sample points. The smoothing operator $P^{(0)}$ is constructed

as the row-stochastic normalized kernel matrix

$$P_{ij}^{(0)} = \frac{K_{ij}^{(0)}}{\sum_k K_{ik}^{(0)}}. \quad (1)$$

A first coarse representation of f is then generated by the convolution $\tilde{f}^{(0)} = f * P^{(0)}$ that captures the low-frequencies of the function. For the next steps, a parameter value $\mu > 1$ is fixed. A sharper normalized Gaussian kernel matrix $P^{(\ell)}$ is constructed at level ℓ with scale σ/μ^ℓ . Then, the residual $d^{(\ell-1)} = f - \tilde{f}^{(\ell-1)}$ is computed. It captures the error of the approximation of f at the previous $\ell - 1$ step. A more detailed representation of f is generated. It is given by

$$\tilde{f}^{(\ell)} = \tilde{f}^{(\ell-1)} + d^{(\ell-1)} * P^{(\ell)} = \tilde{f}^{(\ell-1)} + g^{(\ell)},$$

115 with $g^{(\ell)} = d^{(\ell-1)} * P^{(\ell)}$. The iterative algorithm stops once the norm of the residual vector $d^{(\ell)}$ is smaller than a predefined tolerance. Stopping at iteration L , the final LP model has the form

$$\tilde{f}^{(L)} = \tilde{f}^{(0)} + \sum_{\ell=1}^L g^{(\ell)} = f * P^{(0)} + \sum_{\ell=1}^L d^{(\ell-1)} * P^{(\ell)}. \quad (2)$$

Extending this multi-scale representation to a new data point $x \in \mathbb{R}^M$ is now straightforward by setting

$$\begin{aligned} \tilde{f}^{(L)}(x) &= f * P^{(0)}(x) + \sum_{\ell=1}^L d^{(\ell-1)} * P^{(\ell)}(x) \\ &= \sum_j f_j P^{(0)}(x, x_j) + \sum_{\ell=1}^L \sum_j d_j^{(\ell-1)} P^{(\ell)}(x, x_j). \end{aligned}$$

The kernels $P^{(\ell)}$ are directly extended for a new point x as

$$P^{(\ell)}(x, x_j) = \frac{K^{(\ell)}(x, x_j)}{\sum_k K^{(\ell)}(x, x_k)} \quad (3)$$

with $K^{(\ell)}(x, x') = \kappa^{(\ell)} e^{-\frac{\|x-x'\|^2}{(\sigma/\mu^\ell)^2}}$. Observe that when defining $P^{(\ell)}$, $\kappa^{(\ell)}$ disappears for being also present in the denominator.

120

Next, we show that the L_2 norms of the residuals $d^{(\ell)}$ decay extremely fast.

2.2. Error Analysis for the LP Scheme

For analyzing the LP error, the previously defined kernel is considered. First notice that, when working in the continuous kernel setting, the summation becomes an integral. Therefore, we have $P^{(\ell)}(x, x') = K^{(\ell)}(x, x')$ for a Gaussian function since the denominator in (3) is just $\int K^{(\ell)}(x, z) dz = 1$.

Furthermore, for all ℓ , writing now $P^{(\ell)}(x) = P^{(\ell)}(x, 0)$, P is an approximation to a delta function satisfying

$$\begin{aligned} \int P^{(\ell)}(x) dx &= 1, \\ \int x P^{(\ell)}(x) dx &= 0, \\ \int \|x\|_2^2 P^{(\ell)}(x) dx &\leq 2C, \text{ where } C \text{ is a constant.} \end{aligned} \quad (4)$$

Assume that f is in L_2 , i.e., $\int_x f^2(x) dx < \infty$. The LP scheme is a relaxation process for which in the first step the function f is approximated by $\mathcal{G}^{(0)}(f) = f * P^{(0)}(x)$. In the second step f is approximated by $\mathcal{G}^{(0)}(f) + \mathcal{G}^{(1)}(d^{(0)})$, where $d^{(0)} = f - \mathcal{G}^{(0)}(f)$ and $\mathcal{G}^{(1)}(d^{(0)}) = d^{(0)} * P^{(1)}(x)$, and so on. Taking the Fourier transform of $P^{(\ell)}(x)$ results in (see [32])

$$\left| \hat{P}^{(\ell)}(\omega) - 1 \right| \leq \frac{(\sigma/\mu^\ell)^2}{2} \int \|x\|_2^2 \|\omega\|_2^2 P^{(\ell)}(x) dx \leq C (\sigma/\mu^\ell)^2 \|\omega\|_2^2. \quad (5)$$

We first analyze the error $d^{(0)}(x)$ in the first step, which is defined by $d^{(0)}(x) = f - f * P^{(0)}(x)$. Taking the Fourier transform of $d^{(0)}(x)$ together with the bound in (5) yields

$$\left| \hat{d}^{(0)}(\omega) \right| = \left| \hat{f}(\omega) \right| \left| \hat{P}^{(0)}(\omega) - 1 \right| \leq C \|\omega\|_2^2 \sigma^2 \left| \hat{f}(\omega) \right|. \quad (6)$$

The error in the second step is

$$d^{(1)}(x) = d^{(0)} - \mathcal{G}^{(1)}(d^{(0)}) = \left(f - f * P^{(0)} \right) - d^{(0)} * P^{(1)} = d^{(0)} - d^{(0)} * P^{(1)}. \quad (7)$$

Taking the Fourier transform of (7) yields

$$\begin{aligned} \left| \hat{d}^{(1)}(\omega) \right| &= \left| \hat{d}^{(0)}(\omega) - \hat{d}^{(0)}(\omega) \hat{P}^{(1)}(\omega) \right| \\ &= \left| \hat{d}^{(0)}(\omega) \right| \left| \hat{P}^{(1)}(\omega) - 1 \right|. \end{aligned} \quad (8)$$

Using (5) and (6) we obtain

$$\left| \hat{d}^{(1)}(\omega) \right| \leq C \|\omega\|_2^2 \left| \hat{d}^{(0)}(\omega) \right| (\sigma/\mu)^2 \leq C \sigma^2 (\sigma/\mu)^2 \left| \hat{f}(\omega) \right| \|\omega\|_2^4. \quad (9)$$

Since $\mu > 1$, then $\left| \hat{d}^{(1)}(\omega) \right| \leq C \sigma^2 \frac{\sigma^2}{\mu^2} \left| \hat{f}(\omega) \right| \|\omega\|_2^4$. Similarly, for the ℓ^{th} step the error is bounded by

$$\left| \hat{d}^{(\ell)}(\omega) \right| \leq C \sigma^2 \left(\frac{\sigma^2}{\mu^{(\ell+1)}} \right)^\ell \left| \hat{f}(\omega) \right| \|\omega\|_2^{2(\ell+1)}. \quad (10)$$

By Parseval's equality we obtain

$$\begin{aligned} \|d^{(\ell)}\|_{L^2} &= \|\hat{d}^{(\ell)}\|_{L^2} \leq C \sigma^2 \left(\frac{\sigma^2}{\mu^{(\ell+1)}} \right)^\ell \|\hat{f}(\omega) \|\omega\|_2^{2(\ell+1)}\|_{L^2} \\ &\leq C \sigma^2 \left(\frac{\sigma^2}{\mu^{(\ell+1)}} \right)^\ell \|f\|_{2\ell+2,2}, \end{aligned} \quad (11)$$

where $\|f\|_{m,2}$ denotes the Sobolev norm of a function with up to m derivatives in L_2 . Thus, the L_2 norm of the LP error decays at a very fast rate.

2.3. Overfitting Risk

Since the error of the LP method decays fast, setting a small error threshold may easily result in $\tilde{f}^{(\ell)} \simeq f$ and hence cause overfitting of the data. In order to understand the approximation process, we express $\tilde{f}^{(\ell)} = \tilde{f}^{(\ell-1)} + g^{(\ell)}$ by

$$\begin{aligned} \tilde{f}^{(\ell)} &= \tilde{f}^{(\ell-1)} + g^{(\ell)} = \tilde{f}^{(\ell-1)} + (f - \tilde{f}^{(\ell-1)}) * P^{(\ell)} \\ &= f * P^{(\ell)} + \tilde{f}^{(\ell-1)} * (I - P^{(\ell)}), \end{aligned}$$

where I denotes the identity matrix. Now, taking the limit $\tilde{f}^{(\ell)} \rightarrow \phi$ yields

$$\phi = f * \lim P^{(\ell)} + \phi * \lim(I - P^{(\ell)}),$$

i.e., $\phi = f$, for $P^{(\ell)} \rightarrow I$. In practice, when ℓ is large enough, numerically $P^{(\ell)} \simeq I$, so that $K^{(\ell)}(x_i, x_j) \simeq 0$, $i \neq j$. Then $d_j^{(\ell)} = 0$ for all j and the LP stays almost the same for a large enough ℓ . In other words, care has to be taken when deciding when to stop the LP iterations in order to avoid overfitting.

3. Auto-adaptive Laplacian Pyramids

The standard way to prevent overfitting is to use an independent validation subset and to stop the iterations as soon as the validation error on that subset starts to increase. This can be problematic for small samples as it introduces
150 a random dependence on the choice of the particular validation subset. k -fold Cross Validation (CV) is then usually the standard choice to avoid this. Samples are randomly distributed in k subsets, and iteratively $k - 1$ subsets are used for training while the remaining $N - (k - 1)$ are used for validation. In the extreme case when $k = N$, i.e., when just one pattern is used for validation, CV
155 becomes Leave-One-Out Cross Validation (LOOCV) and the train iterations are stopped when the LOOCV error starts to increase. Besides its simplicity, LOOCV has the attractive property of being an almost unbiased estimator of the true generalization error (see for instance [33, 34]), possibly with a high variance [35]. In our case, LOOCV can be easily applied using for training a
160 $N \times N$ normalized kernel matrix $P_{(p)}$. This is just the previous matrix P , where we set to 0 the p -th rows and columns when x_p is held out of the training sample and used for validation. The most obvious drawback of LOOCV is its rather high cost, which in our case would be in principle $N \times O(LN^2) = O(LN^3)$, where we recall that L is the number of LP iterations. However, it is often the
165 case for other models that there are ways to estimate the LOOCV error with a smaller cost. This can be done exactly in the case of k -Nearest Neighbors [36] or Ordinary Least Squares ([37], Chapter 7); or approximately for Support Vector Machines [38] or Gaussian Processes [39]. We show next how to perform LOOCV without essentially augment training cost.

3.1. Standard ALP

In the context of this work and to alleviate the LOOCV cost, notice first that here when x_p is removed from the training sample, the test value of the

$f_{(p)}$ extension at the point x_p is given by

$$\begin{aligned} f_{(p)}^{(L)}(x_p) &= \sum_{j \neq p} f_j P^{(0)}(x_p, x_j) + \sum_{\ell=1}^L \sum_{j \neq p} d_{(p);j}^{(\ell-1)} P^{(\ell)}(x_p, x_j) \\ &= \sum_j f_j \tilde{P}^{(0)}(x_p, x_j) + \sum_{\ell=1}^L \sum_j d_{(p);j}^{(\ell-1)} \tilde{P}^{(\ell)}(x_p, x_j). \end{aligned}$$

Here $\tilde{P}^{(\ell)}$ is a modification of $P^{(\ell)}$, where the diagonal elements in $\tilde{P}^{(\ell)}$ are set to zero, i.e., $\tilde{P}_{i,j} = P_{i,j}$ when $j \neq i$. $d_{(p)}^{(\ell)}$ are the different previously defined errors computed using the $P_{(p)}^{(\ell)}$ matrices.

This observation leads to the modification of the standard LP that was proposed in [29], and which simply consists on applying the LP procedure described in Section 2 but replacing the P matrix by its 0-diagonal version \tilde{P} . This modification requires the computation of the function approximation $\tilde{f}^{(0)} = f * \tilde{P}^{(0)}$ at the beginning, and then the vectors $\tilde{g}^{(\ell)} = \tilde{d}^{(\ell-1)} * \tilde{P}^{(\ell)}$, $\tilde{f}^{(\ell)} = \tilde{f}^{(\ell-1)} + \tilde{g}^{(\ell)}$ and $\tilde{d}^{(\ell)} = f - \tilde{f}^{(\ell)}$ are computed at each iteration. This algorithm is denoted as the *Auto-adaptive Laplacian Pyramid (ALP)* [25].

According to the previous formula for the $f_{(p)}^{(L)}(x_p)$, we can take the ALP values $\tilde{f}_p^{(L)} = \tilde{f}^{(L)}(x_p)$ given by

$$\tilde{f}^{(L)}(x_p) = \sum_j f_j \tilde{P}^{(0)}(x_p, x_j) + \sum_{\ell=1}^L \sum_j \tilde{d}_j^{(\ell-1)} \tilde{P}^{(\ell)}(x_p, x_j),$$

as approximations to the LOOCV validation values $f_{(p)}^{(L)}(x_p)$.

The squared LOOCV error at each iteration may be approximated by

$$\sum_p (f(x_p) - f_{(p)}^{(L)}(x_p))^2 \simeq \sum_p (f(x_p) - \tilde{f}_p^{(L)})^2 = \sum_p (\tilde{d}_p^{(L)})^2,$$

which is just the training error of ALP at the current iteration. In other words, working with the \tilde{P} matrix instead of P , the training error at step L gives in fact an approximation to the LOOCV error at this step.

For the parameter selection, choosing as customarily done $\mu = 2$, the only required parameter would be the initial σ but assuming it is wide enough, its $\sigma/2^\ell$ scalings will yield an adequate final kernel width. Furthermore, in this

paper we propose a heuristic technique to compute the initial bandwidth σ and the maximum number of iterations \max_{its} ensuring that all the interesting σ values (and only the interesting ones) will be tested. Relevant values of σ are obtained in between the ones that yield an all-ones P matrix, and the ones that produce a $P = I$ matrix. In the first case, every data point obtains the same weight, resulting in a mean function. In the second case, as mentioned in Subsection 2.3, the original function is reproduced. To approximate these values, the parameters are fixed in the following way:

$$\begin{aligned}\sigma &= 10 \max(W_{ij}) \\ \max_{\text{its}} &= \log_2(\sigma/\sigma_{\min}), \quad \text{with } \sigma_{\min} = 1/5 \min(W_{ij}),\end{aligned}$$

185 where W represents the distance matrix.

We also point out that it is straightforward to extend the model to a vectorial function $F = (F_1, \dots, F_M)$. The squared ALP error at the L -th iteration becomes

$$\sum_p \|F(x_p) - F_{(p)}^{(L)}(x_p)\|^2 \simeq \sum_p \|F(x_p) - \tilde{F}_p^{(L)}\|^2 = \sum_p \|\tilde{D}_p^{(L)}\|^2,$$

where here $D = (D_1, \dots, D_M)$ is the vector formed of the ALP residuals D_m of each component F_m of F . Again, the optimal value for L is set as the iteration where this estimate of the LOOCV error begins to grow.

The cost of running L steps of ALP is just $O(LN^2)$ and, thus, we gain
190 the advantage of approximating the exhaustive LOOCV without any additional cost on the overall algorithm. The complete training and test procedures are presented in Algorithm 1 and Algorithm 2 respectively. Notice that in this work, the computed error for setting the optimal stopping iterations number of the algorithm is the Root Mean Squared Error (RMSE). This type of error
195 measure was selected because it is the classical error to compute when dealing with functions and due to the improved results obtained empirically compared with the Mean Absolute Error (MAE).

As it has just being argued, the obvious advantage of ALP is that when the training error is evaluated, in practice the LOOCV error after each LP iteration

Algorithm 1 The ALP Training Algorithm

Input: $\{x_i, F_i\}_{i=1}^N$.

Output: $(\{\tilde{d}^{(\ell)}\}, \sigma, \mu, L)$, the trained model.

```

1:  $\mu \leftarrow 2, W_{ij} \leftarrow x_i - x_j \quad \forall i, j, \sigma \leftarrow 10 \max(W_{ij}), \max_{\text{its}} \leftarrow \log_2(5 \frac{\sigma}{\min(W_{ij})})$ .
2:  $\tilde{F}^{(0)} \leftarrow F; \tilde{F}^{(0)} \leftarrow 0; \ell \leftarrow 1$ .
3: while ( $\ell < \max_{\text{its}}$ ) do
4:    $\tilde{K}^{(\ell)} \leftarrow e^{-\frac{\|W\|^2}{\sigma^2}}$ , with 0-diagonal. % LOOCV approximation.
5:    $\tilde{P}^{(\ell)} \leftarrow \text{normalize}(\tilde{K}^{(\ell)})$ .
6:    $\tilde{F}^{(\ell)} \leftarrow \tilde{F}^{(\ell-1)} + \tilde{D}^{(\ell-1)} * \tilde{P}^{(\ell)}$ .
7:    $\tilde{D}^{(\ell)} \leftarrow F - \tilde{F}^{(\ell)}$ .
8:    $\text{err}^{(\ell)} \leftarrow \|\tilde{D}^{(\ell)}\|^2$ . % For each point a mean error value is computed.
9:    $\sigma \leftarrow \sigma/\mu; \ell \leftarrow \ell + 1$ .
10: end while
11:  $L \leftarrow \arg \min_{\ell} \{\text{err}^{(\ell)}\}$ . % Optimal iteration.
```

Algorithm 2 The ALP Testing Algorithm

Input: $\{x_i\}_{i=1}^N, x_{\text{te}}, (\{\tilde{D}^{(\ell)}\}, \sigma, \mu, L)$.

Output: \hat{F}_{te} .

```

1:  $\hat{F}_{\text{te}} \leftarrow 0$ .
2: for  $\ell=1$  to  $L$  do
3:    $K_{i;\text{te}}^{(\ell)} \leftarrow e^{-\frac{\|x_i - x_{\text{te}}\|^2}{\sigma^2}} \quad \forall i$ .
4:    $P^{(\ell)} \leftarrow \text{normalize}(K^{(\ell)})$ .
5:    $P^{(\ell)} \leftarrow \text{normalize}(K^{(\ell)})$ .
6:    $\hat{F}_{\text{te}} \leftarrow \hat{F}_{\text{te}} + \tilde{D}^{(\ell-1)} * P^{(\ell)}$ .
7:    $\sigma \leftarrow \sigma/\mu$ .
8: end for
```

200 is estimated. Therefore, the evolution of these LOOCV values automatically
 defines the optimal iteration at which the algorithm is stopped, i.e., just when
 the training error starts to increase. Thus, the risk of overfitting is removed
 and in addition the training errors can be used as an approximation to the
 generalization error. This effect can be seen in Figure 1 that illustrates the
 205 application of ALP to the synthetic problem described in the next subsection.
 In this example, the optimum stopping time for ALP is exactly the same as the
 LOOCV error would generate, where the training error stabilizes afterwards at
 a slightly larger value. Moreover, ALP achieves an automatic selection of the
 width of the Gaussian kernel which makes this version of LP to be auto-adaptive
 210 as it does not require costly parameter selection procedures.

3.2. ALP with Local Resolution

When working with scattered datasets, it is often the case that the available
 data is not equally distributed, and there exist regions with different sample
 statistics and density characteristics. In these cases, it makes sense to define
 215 different models, with different stopping times, adapted to each region.

Focusing on the ALP method, the original algorithm may be modified to
 take into account this particular issue. We propose to use a point-wise σ value,
 computed by using the neighborhood of each sample point. To accomplish this
 goal, the train and test phases are modified as follows.

- 220 • **Training step:** An error per point is computed in terms of the ν nearest
 neighbors mean error. The optimal number of iterations will be given then
 by $L_i^* \leftarrow \arg \min_{\ell} \{\text{err}_i^{(\ell)}\}$.
- **Test step:** The optimal number of iterations L_{te}^* of each test point will
 be given by the optimal number of iterations of the nearest training point
 225 x_n , i.e. $L_{\text{te}}^* = L_n^*$.

The proposed modification introduces a new parameter ν , that represents
 the number of neighbors taken into account to define the best local σ per point.

This parameter will have the same value for every point in the training set, and it will be determined by CV.

230

To sum up and remark the changes with respect to the original ALP algorithm, the corresponding training and test procedures are presented in Algorithm 3 and Algorithm 4.

Algorithm 3 The Local ALP Training Algorithm

Input: $\{x_i, F_i\}_{i=1}^N$.

Output: $(\{\tilde{D}^{(\ell)}\}, \sigma, \mu, \{L_i\}_{i=1}^N)$.

```

1:  $\mu \leftarrow 2, W_{ij} \leftarrow x_i - x_j \quad \forall i, j, \sigma \leftarrow 10 \max(W_{ij}), \max_{\text{its}} \leftarrow \log_2(5 \frac{\sigma}{\min(W_{ij})})$ .
2:  $\tilde{D}^{(0)} \leftarrow f; \tilde{F}^{(0)} \leftarrow 0; \ell \leftarrow 1$ .
3: while ( $\ell < \max_{\text{its}}$ ) do
4:    $\tilde{K}^{(\ell)} \leftarrow e^{-\frac{\|W\|^2}{\sigma^2}}$ , with 0-diagonal.
5:    $\tilde{P}^{(\ell)} \leftarrow \text{normalize}(\tilde{K}^{(\ell)})$ .
6:    $\tilde{F}^{(\ell)} \leftarrow \tilde{F}^{(\ell-1)} + \tilde{D}^{(\ell-1)} * \tilde{P}^{(\ell)}$ .
7:    $\tilde{D}^{(\ell)} \leftarrow F - \tilde{F}^{(\ell)}$ .
8:    $\text{err}_i^{(\ell)} \leftarrow \|\tilde{D}_i^{(\ell)}\|^2 \quad \forall i$ . % For each point the mean error of the nearest points is computed.
9:    $\sigma \leftarrow \sigma/\mu; \ell \leftarrow \ell + 1$ .
10: end while
11:  $L_i \leftarrow \arg \min_{\ell} \{\text{err}_i^{(\ell)}\}$ . % Optimal iteration ( $\sigma$  value) per point.
```

Algorithm 4 The Local ALP Testing Algorithm

Input: $\{x_i\}_{i=1}^N, x_{\text{te}}, (\{\tilde{D}^{(\ell)}\}, \sigma, \mu, \{L_i\}_{i=1}^N)$.

Output: $\hat{F}_{\text{te}}, L_{\text{te}}$.

```

1:  $\hat{F}_{\text{te}} \leftarrow 0$ .
2: for  $\ell=1$  to  $\max\{L_i\}$  do
3:    $K_{i;\text{te}}^{(\ell)} \leftarrow e^{-\frac{\|x_i - x_{\text{te}}\|^2}{\sigma^2}} \quad \forall i$ .
4:    $P^{(\ell)} \leftarrow \text{normalize}(K^{(\ell)})$ .
5:    $\hat{F}_{\text{te}} \leftarrow \hat{F}_{\text{te}} + \tilde{D}^{(\ell-1)} * P^{(\ell)}$ .
6:    $\sigma \leftarrow \sigma/\mu$ .
7:   if  $\ell==1$  then
8:      $L_{\text{te}} \leftarrow L(\arg \min_i \{x_i - x_{\text{te}}\})$ . % Return also the optimal iteration per point.
9:   end if
10: end for
```

When applying this new ALP version, the risk of overfitting is still removed.

In this case, in which the width of the kernel is automatically adapted to each
 235 point, the overfitting risk is prevented in a point-wise manner. We will see an
 example of this behavior in the following subsection.

3.3. A Synthetic Example

For a better understanding of the proposed method and its advantages, we
 illustrate in this section the classic ALP algorithm and its local resolution version
 240 on a synthetic example of a composition of sines with different frequencies plus
 additive noise.

Consider a sample x with N points equally spaced over the range $[0, 10\pi]$.
 The target function f is given by

$$f = \sin(x) + 0.5 \sin(3x) \cdot I_2(x) + 0.25 \sin(9x) \cdot I_3(x) + \varepsilon,$$

where I_2 is the indicator function of the interval $(10\pi/3, 10\pi]$, I_3 that of $(2 \cdot$
 $10\pi/3, 10\pi]$ and $\varepsilon \sim \mathcal{U}([-\delta, \delta])$ is uniformly distributed noise. In other words,
 there is a single frequency in the interval $[0, 10\pi/3]$, two frequencies in $(10\pi/3, 2 \cdot$
 245 $10\pi/3]$ and three in $(2 \cdot 10\pi/3, 10\pi]$. For the classic ALP algorithm, two different
 simulations are executed: the first one with 4,000 points with small $\delta = 0.1$
 noise and the second one with 2,000 points and a larger $\delta = 0.5$ (observe that
 $|f| \leq 1.75$). In both cases, $1/3$ of the original set is randomly chosen for test
 purposes.

250 Recall that the main advantage of ALP is the approximation of the LOOCV
 error obtained while we evaluate the training error. Due to this fact, if the
 algorithm iterations stop as the error starts to grow, the risk of overfitting is
 removed. This effect can be observed in Figure 1 where ALP is applied to this
 synthetic example. The solid blue and dashed green lines represent the LP
 255 training error and the true LOOCV error per iteration respectively, and the
 dashed red line represents the error for the ALP method. Notice that the ALP
 training error attains its minimum at the same iteration prescribed by exact
 LOOCV for LP.




Figure 1: Training errors for the original LP models (with and without LOOCV) and its modified version, the ALP model, applied over a perturbed sine. Left: $\delta = 0.1$ over 4,000 patterns; right: $\delta = 0.5$ over 2,000 patterns.

Recall also that the ALP model automatically adapts its multiscale behavior
 260 to the data, trying to refine the prediction in each iteration using a more local-
 ized kernel, given by a smaller σ . This behavior can be observed in Figure 2,
 which shows the evolution of the ALP predictions for the 0.1-noise experiment.
 At the beginning, the model approximates the function just by a coarse mean
 of the target function values, and in the subsequent iterations when the model
 265 starts using sharper kernels and refined residuals, the approximating function
 captures the different frequencies and amplitudes of the composite sines. In this
 particular case, the minimum LOOCV value is reached after 13 iterations. Note
 that we start plotting from the 7th iteration, because before the result resembles
 a mean function due to a large initial σ value.

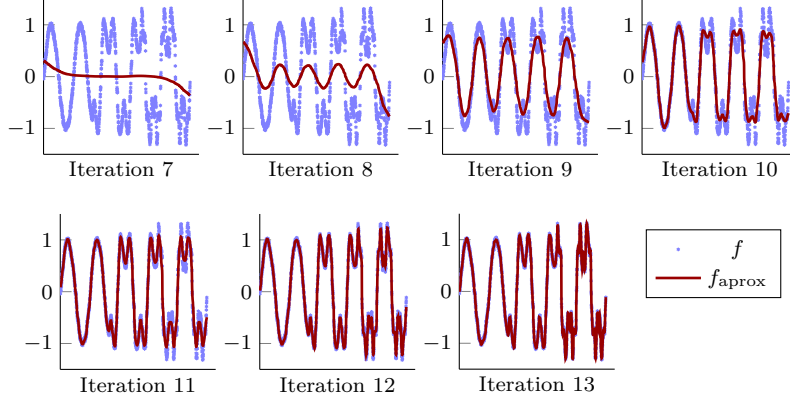


Figure 2: Interpolations given by the ALP model for the last seven steps (out of 13) on the small noise example.

Next, the same synthetic experiment is carried out, but now the amplitude of the uniform noise distribution is increased to $\delta = 0.5$. The predicted function is represented in Figure 3 and it is obtained after 12 iterations (as shown in the right hand image of Figure 1). As expected, the number of LP iterations is now slightly smaller than in the previous example because the algorithm selects a more conservative, smoother prediction in the presence of noisier and, thus, more difficult data.




Figure 3: Final interpolations given by the ALP model for the large noise example.

In any case, we can conclude that the ALP model captures very well the essential underlying behavior of both samples, as it identifies the three different frequencies of the sine and their amplitudes even when the noise level increases.

280 Next, we illustrate in this subsection the advantages of the modified ALP
version with local resolution (ALP_l). For this purpose, the same sine example is
used, but the focus is on the third interval, where the three frequencies of the sine
are present. The function takes the form $f = \sin(x) + 0.5 \sin(3x) + 0.25 \sin(9x)$,
with 4,000 points and no noise ($\delta = 0$, and thus $\varepsilon = 0$). In this case, also
285 three different regions are defined, but in terms of the subsampling density. In
particular, we divide the interval $[0, 10\pi]$ into three parts. The first region holds
400 samples, the second 1,400 and the third one 2,200. In this scenario it seems
logical to adapt the final bandwidth of the Gaussians to each region, expecting
larger bandwidth values (few iterations) for sparse regions, obtaining then a
290 smoother, coarser approximation; for the dense regions we expect a smaller σ
value (more iterations), that will capture finer details, as there are more points
to estimate the original function).

The parameters of the method have been fixed automatically, as done in the
ALP case, except the parameter ν for the local approximation that has been
295 cross-validated using 10 folds and a grid $\{10, 20, \dots, 200\}$, selecting finally a
value of 50. As seen in Figure 4, the resolution selection works as expected,
presenting three main different values corresponding to the three different sub-
sampling regions.




Figure 4: Left: Representation of the synthetic example over the training set, colored by subsampling region. Right: Stopping number of iterations for the synthetic example over the training set, represented as a constant red line for ALP and as asterisks of different colors depending on region for ALP_l .

The original ALP and ALP_l algorithms are compared in Figure 5 for the three different regions, where we can acknowledge that in general the blue ALP_l points are very close to the red ALP values. Both methods present very similar results, as expected because they used very similar σ values, and they give a good result for prediction. In particular, this can be observed in the second and third region, where the sample set is dense and big enough. The first region, which only has a few number of available sampled point, is the most difficult region for approximation and therefore the results are less accurate.




Figure 5: Sine example test results. Comparison between ALP (red points) and ALP_l (blue points) predictions.

Finally, we would like to emphasize that the new ALP_l algorithm prevents overfitting as well as its original version, as explained before. This fact can be appreciated in Figure 6 where each curve presents the same shape than the ALP curve (like the one shown in Figure 1). It should be noticed that the red lines, corresponding with points in the last region, represents an error over an easy synthetic problem with lots of subsampled points. Because of this, the error is 0, as the method is able to recover the original function perfectly.




Figure 6: Training errors of a subset of the sample points. The different colors represent the three subsampling regions as in previous plots, where blue curves correspond to points in the first region, green curves correspond with the second region and red ones with the third region.

4. Experiments over real data

315 This section presents numerous interpolation experiments over real datasets,
for which we analyze the behavior and properties of ALP and its modified
version, ALP_l . For this analysis, the results are compared against a k -Nearest
Neighbor (k -NN) model, a standard interpolation method. All the real problems
presented in the paper will deal with the missing values problem, which is one of
320 the most interesting and more frequent problems in the interpolation context,
for being crucial to tackle other problems like classification or regression. The
importance of missing values problems is shown in [40], where missing values are
imputed using k -NN and Self-Organizing Map techniques or in [41], where the
classification problem with missing values is solved with a random-forest based
325 algorithm. In this type of problem, when the data is regressed many times
against each of the columns, methods like k -NN need to be tuned manually
to fit the behavior of each target function while ALP and ALP_l automatically
choose a scale that fits the function and data.

For all the experiments we will follow the subsequent methodology:

- 330 1. Select the less correlated features among a complete dataset to simulate
the missing value problem, if needed.
2. Define 10 independent train-test folds for each case study. We consider 3
possible scenarios: a 90%-10% train-test split, a 80%-20% and a 70%-30%.
In total, 30 experiments per problem will be executed.
- 335 3. Parameter selection in one of the training folds. As previously explained,
the parameters needed for ALP and ALP_l are automatically selected by
the algorithm, except in the case of the parameter ν in the local version.
Recall that this parameter is used for determining the number of neighbors
involved in the point-wise final bandwidth selection for each training point.
340 The value of this parameter is estimated using a 10-fold CV. For the k -NN
algorithm, the k neighbors are also obtained via a 10-CV.
4. Run the models in each of the 30 defined experiments.

5. Evaluate the results. The Root Mean Squared Error (RMSE) is used for measuring errors, as previously discussed. Notice that, for making all the datasets errors comparable among them, the errors presented are divided by the standard deviation of the target function in the test set. We will present the median and standard deviation of each set of 10 experiments, together with a statistical significance test between models, in this case a Mann–Whitney U test [42] applied over the errors obtained. The null hypothesis of this test states that both models come from continuous distributions with equal medians, rejecting the null hypothesis at the 5% significance level.
6. Derive plots to illustrate the results. When the dimension of the data is large, Principal Component Analysis (PCA [43]) will be applied for making the visualization possible.

Next, the described procedure is applied to a number of real datasets.

4.1. Wisconsin Breast Cancer (Diagnostic) Dataset

The first example selected is a classification UCI dataset [44]: the Wisconsin breast cancer. The features in this dataset are characteristics of each cell nuclei presented on an image of a breast mass. In this work, we changed the original target function and instead we take one of the cell characteristics to be the target function for interpolation, simulating that it has several missing values. The selected feature for being the target function is the less correlated with the other ones, which in this example is the number 12.

As explained in the general case, three different scenarios will be considered: a 90%-10% train-test split, a 80%-20% and a 70%-30%. In all cases, the data has been standardized before building the models. For the ν parameter selection, the grid $\{10, 20, \dots, 200\}$ was used in the ALP_l case selecting values of 70, 110 and 170 for the 10%, 20% and 30% test splits respectively. The grid $\{1, 2, \dots, 10\}$ was used for the k -NN method, obtaining $k = 2$ for all the partitions.

The interpolation RMSE errors for this example are shown in Table 1, together with the result of applying the Mann–Whitney U significance test, mark-

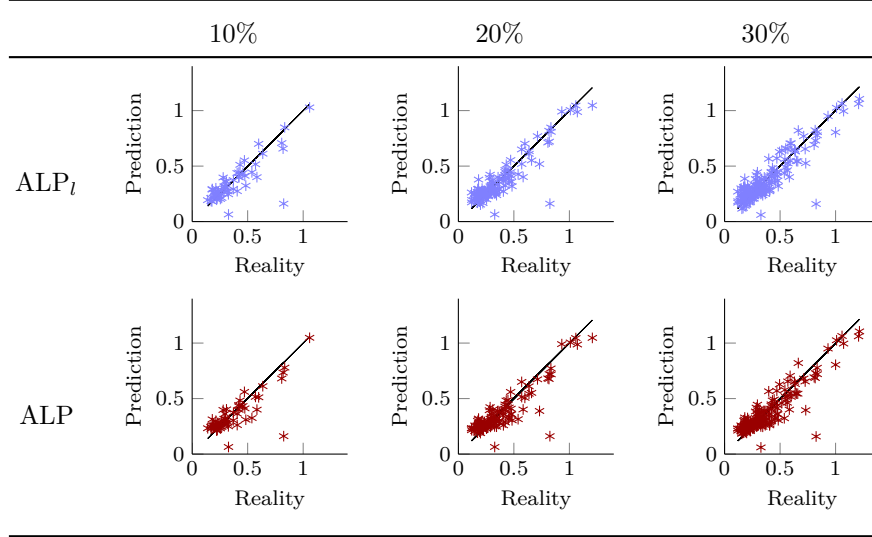
ing in bold font the winners and ties in first position. Notice that we consider the three different scenarios depending on the number of test samples: 10% of random test samples, 20% or 30%. Observing the results in the table we can conclude that the proposed method, in its original form and in its local version, yields the best results for the 10% case, outperforming k -NN as an interpolation method and ties it when less information is available during training. Even though the error values between ALP and ALP_l are not significant in this example, the ALP_l errors are smaller when there is more training data available. This implies that finer structure in the data is captured with ALP_l when data is at hand (see the 10% case).

Table 1: Normalized RMSE errors for the Breast Cancer example.

	ALP	ALP_l	k -NN
10%	0.4181 ± 0.1025	0.4007 ± 0.0881	0.4797 ± 0.0666
20%	0.4194 ± 0.0941	0.4265 ± 0.0914	0.4580 ± 0.0488
30%	0.5431 ± 0.1214	0.4517 ± 0.1246	0.4673 ± 0.0383

For reinforcing these results, we present some plots to visually appreciate the effects of these methods for modeling the considered dataset. In Table 2 we present a comparison between the expected result (the reality, in the x axis) and the interpolated one (the prediction, in the y axis) for ALP_l and ALP. Both methods present similar results, but if we focus our attention on the left lower side of the image, ALP_l points are a bit nearer to the diagonal than ALP ones.

Table 2: Breast Cancer example. Comparison between ALP (red points) and ALP_l (blue points) predictions over the test set for different test subsets.



In Figure 7, the final number of iterations are depicted over the PCA training
390 coordinates (recall that the number of iterations is related to the analysis scale
of the data in each region). It can be appreciated that points with similar final
number of iterations are located in the same region, so it seems that the number
of iterations it is also related with the data structure and the ALP_l method is
able to capture this information.




Figure 7: Breast Cancer Wisconsin (Diagnostic) number of iterations. The image plots the final number of iterations, represented in different colors according to the colorbar, over the PCA components of the training set.

395 4.2. Wine Quality Dataset

For another real, but small dataset example, we consider the Wine Quality dataset from the UCI repository. The wine dataset holds several features based on physicochemical tests such as acidity, chlorides or sulfur dioxide from each wine type. There is a separate dataset for red and white variants of wine. We
400 have discarded the discrete features number 2 and 8 from both datasets.

As in the previous example, we change the problem for being a missing values one. The feature with missing values, selected for being the less correlated with the other features, is in this case the feature number 4 for the wine red dataset, that corresponds to the residual sugar, and feature number 10 for the white
405 wine dataset, that corresponds to the sulphates information.

We consider again three different scenarios: a 90%-10% train-test split, a 80%-20% and a 70%-30% and, in all cases, the data has been standardized before building the models.

For the parameter selection of the ALP_l algorithm and for the k of k -NN,
410 the used grids where the same than for the first example, obtaining the values of 50, 40 and 50 for the 10%, 20% and 30% test splits respectively in the case of the red wine dataset and $\nu = \{110, 60, 90\}$ for the white wine dataset. for the 10% test case values of $\nu = 70$, for the 20% case $\nu = 110$ and for the 30% test

split, $\nu = 170$. For the k parameter, the best values were $k = \{4, 3, 3\}$ for the
415 10%, 20% and 30% cases respectively for the red wine dataset and $k = 7$ for all
the white wine cases.

The RMSE errors for both examples are shown in Table 3, where it is also
shown the result of the significance test over the errors obtained at a 5% sig-
nificance level. It can be appreciated that for the red wine variant, there is no
420 significance difference between models except for the 30% case, where the ALP
models outperform k -NN. For the 10% and 20% cases, even though there is no
significance difference, the errors is smaller also in the ALP models. For the
white variant, the ALP_l model is clearly better than the other two, indepen-
dently of the test subset size.

Table 3: Normalized RMSE errors for the Wine examples.

		ALP	ALP_l	k -NN
10%	Red	0.9190 ± 0.0950	0.9072 ± 0.0991	0.9354 ± 0.1071
	White	0.8527 ± 0.0212	0.8191 ± 0.0405	0.8704 ± 0.0243
20%	Red	0.8845 ± 0.0676	0.8898 ± 0.0690	0.9246 ± 0.0826
	White	0.8627 ± 0.0197	0.8293 ± 0.0279	0.8767 ± 0.0191
30%	Red	0.8489 ± 0.0501	0.8725 ± 0.1057	0.9043 ± 0.0654
	White	0.8712 ± 0.0110	0.8540 ± 0.0139	0.8963 ± 0.0104

425 In Table 4 and Table 5 a comparison between the expected result (the target,
in the x axis) and the interpolated one (the prediction, in the y axis) for ALP
and ALP_l is shown for both examples. The previous conclusions can be also
observed in the images: for the red wine variant, both models look almost
indistinguishable, while for the white wine the results are clearly better (they
430 follow better the diagonal line).

Table 4: Red wine example. Comparison between ALP (red points) and ALP_l (blue points) predictions over different test subset sizes.

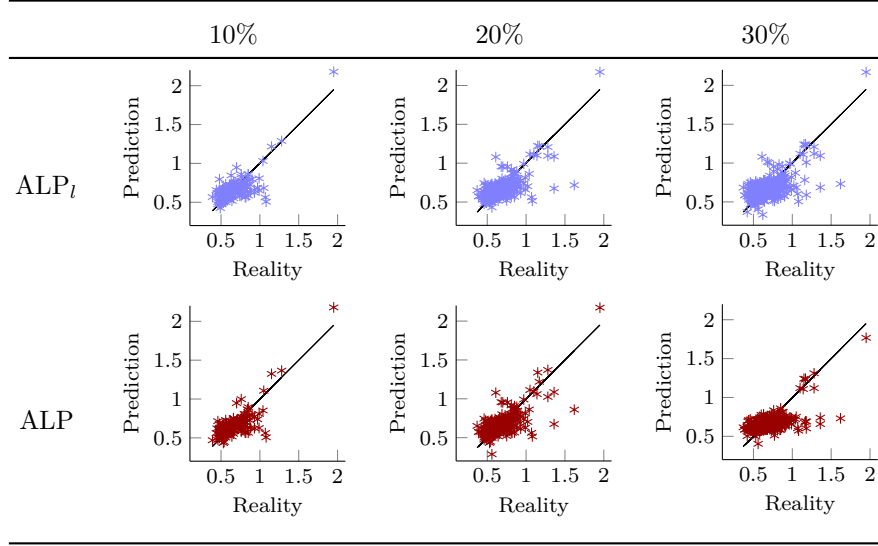
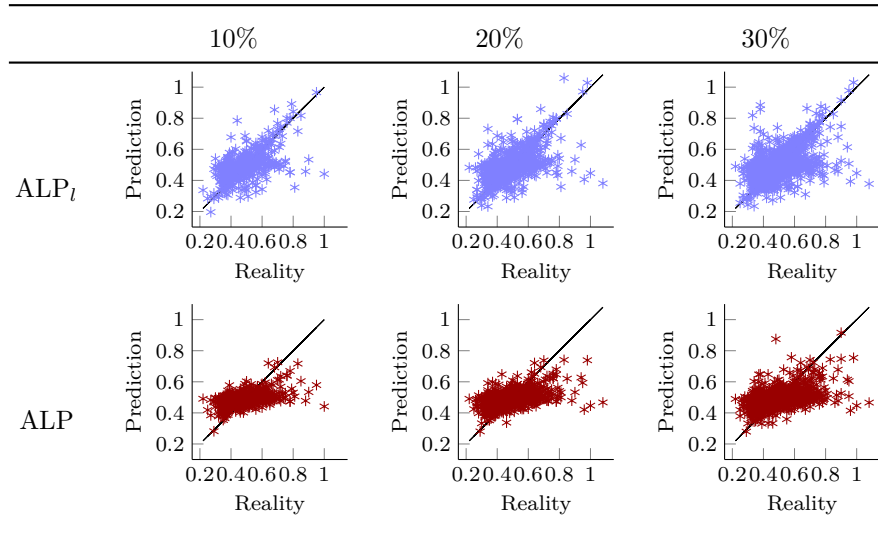


Table 5: White wine example. Comparison between ALP (red points) and ALP_l (blue points) predictions over different test subset sizes.



4.3. Mice Protein Dataset

The Mice Protein Expression [45] dataset from the UCI repository consists of the expression levels of 77 protein modifications that produce detectable signals in the nuclear fraction of cortex of mice. The dataset contains a total of 1080 measurements per protein and each measurement can be considered as an independent sample. The original task for this problem is the classification of different types of mice but, as explained before, here the dataset will be used to demonstrate the interpolation capabilities of the proposed methods. Nevertheless, this real-life dataset presents a high rate of missing values, and interpolation methods can be a good solution for filling in these gaps [46]. For measuring the quality of the results, the experimental data will be set as the subset of the 66 features from the entire dataset that have no missing values. From this dataset, we will simulate that one feature has some gaps on it. In particular, feature number 52 is chosen as the missing data feature as it is less correlated with the rest of the feature columns in the data.

We consider again three different scenarios: a 90%-10% train-test split, a 80%-20% and a 70%-30% and, in all cases, the data has been standardized before building the models.

For the ν parameter selection of the ALP_l algorithm and for the k of k -NN, the used grids were the same as for the previous examples, obtaining for the 10% test case values of $\nu = 40$, for the 20% case $\nu = 130$ and for the 30% test split, $\nu = 140$. For the k parameter the best value was 2 for all the partitions.


The RMSE errors are shown in Table 6, where the significance test results (at the 5% significance level) are also shown, following the same notation as in the previous examples. In this case there is a tie between the three compared methods in almost all the tested cases, but as in the first example, and even though the difference is not significant, it can be seen that when more data is available (the 10% case) the ALP_l obtains a lower error.

Table 6: Normalized RMSE errors for the Mice Protein example.

	ALP	ALP_l	k -NN
10%	0.1883 ± 0.0243	0.1874 ± 0.0241	0.1962 ± 0.0302
20%	0.2020 ± 0.0222	0.1998 ± 0.0223	0.2122 ± 0.0356
30%	0.2133 ± 0.0235	0.2133 ± 0.0235	0.2529 ± 0.0426

In Table 7 we present a comparison between the expected result (the target, in the x axis) and the interpolated one (the prediction, in the y axis) for ALP and ALP_l . The results obtained with both models look very similar, as expected.

Table 7: Mice protein example. Comparison between ALP (red points) and ALP_l (blue points) predictions over the test set for different test subsets.



4.4. Seismic Data

Seismic data analysis is another real example where interpolation is necessary. In these datasets, temporal data is collected by using a grid of sensors, in this case seismometers, and missing data values is a common problem that

occurs when one of the sensors stops to function. The task at hand is a multi-dimensional interpolation problem, where a complete temporal series that was not recorded needs to be recovered.

The presented examples are from a marine and a land seismic networks, and
 470 the data was taken from the New Zeland open database (http://wiki.seg.org/wiki/Open_data). The number of patterns, i.e. the information of each point of the grid, available for these examples (around 70,000) was too big, so a reduced dataset of the first 10,000 coordinates available was prepared for the simulations. The objective function to interpolate here is the entire temporal
 475 series of dimension 1,252 in the marine case and 2,500 in the land one.

In this case, for CV of the number ν of neighbors for the ALP_l algorithm, we have used the knowledge of the grid structure of the data, trying as grid sizes $\{8, 24, 48, 80, 120, 168, 224\}$ (with the idea of taking in each case bigger squares around the point). In the case of the marine seismic dataset we obtain values
 480 of $\nu = 168$ for the 10% test partition, $\nu = 24$ for the 20% test partition and $\nu = 8$ for the 30% one, and in the case of the land seismic dataset we obtain $\nu = 120, 168, 168$ respectively. To set the k -NN parameter, the grid $\{1, 2, \dots, 10\}$ was used, like in the previous examples, obtaining $k = 3$ for all the partitions in both datasets.

485 Table 8 presents the RMSE errors for the different models, where ALP models again outperform k -NN as interpolation methods, and again we can observe that in general when more train data is available, the ALP_l model has an advantage. This table also presents the results of the significance test applied at the 5% significance level over the errors obtained, remarking the most significant
 490 results in bold-faced.

Table 8: Normalized RMSE errors for the Seismic examples.

		ALP	ALP_l	k -NN
10%	Marine	0.2937 ± 0.0221	0.2783 ± 0.0099	0.3299 ± 0.0068
	Land	0.0989 ± 0.0011	0.0989 ± 0.0011	0.1450 ± 0.0010
20%	Marine	0.3202 ± 0.0096	0.3224 ± 0.0095	0.3557 ± 0.0078
	Land	0.1075 ± 0.0008	0.1075 ± 0.0008	0.1567 ± 0.0015
30%	Marine	0.3342 ± 0.0114	0.3479 ± 0.0086	0.3776 ± 0.0088
	Land	0.1187 ± 0.0012	0.1187 ± 0.0012	0.1707 ± 0.0014

In Tables 9 and 10 the two ALP versions are depicted and also the real objective function for an interval of the temporal series interpolation of one of the patterns. The results in this case are essentially indistinguishable, as both models are basically equal.

Table 9: Marine seismic example. Comparison between ALP (red points) and ALP_l (blue points) predictions over one test pattern.

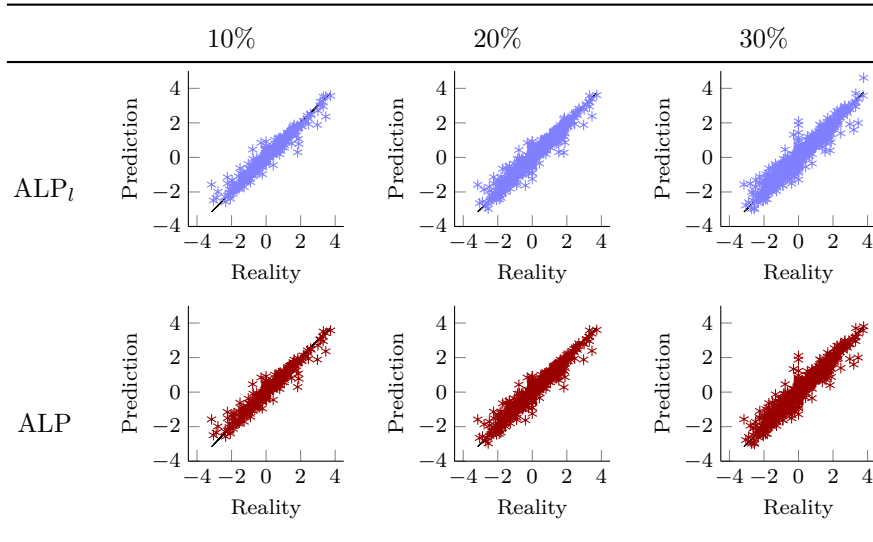



Table 10: Land seismic example. Comparison between ALP (red points) and ALP_l (blue points) predictions over one test pattern.



495 In this case, the difference between the ALP versions is neither significant at
a 5% significance level but we can conclude that the ALP with local resolution
is also able to detect this global structures where just a global σ is needed,
offering the same results as the ALP original model. It should be remarked the
difference with respect to the k -NN model, being again that the ALP models
500 are a better option for interpolation in this real problem.

5. Conclusions

The proposed framework provides novel, flexible and multi-scale analysis
tools, which are easy to implement and are suitable for modeling different types
of datasets including those with non-uniform data distribution. In particular, we
505 introduce an extended version of the Auto-adaptive Laplacian Pyramids (ALP)
where the local structure of the data is taken into account with success, defining
a different width of the kernel per point. In addition, this adaptive version of

LP training yields, at no extra cost, an estimate of the LOOCV value at each iteration, allowing thus to automatically decide when to stop in order to avoid
510 overfitting.

All together, the proposed methodology provides a robust and flexible framework supported by theoretical error analysis for modeling complex datasets. This work overcomes the limitations of previous LP methods and, thus, applications that utilizes LP will benefit from this research. Indeed, these advantages
515 are already expressed in the experimental results where ALP outperforms k -NN in the different examples presented.

Regarding future work, one challenge we plan to tackle is the global nature of these algorithms, which requires high computational cost as data size grows. This may be addressed by modifying the training step stages of the current
520 algorithm by considering smaller suitable subsamples. Finally, since our method offers a general setting, it can be applied in different domains for out-of-sample extension and forecasting problems.

Acknowledgements

They wish to thank Prof. Ronald R. Coifman for helpful remarks. They
525 also gratefully acknowledge the use of the facilities of Centro de Computación Científica (CCC) at Universidad Autónoma de Madrid.

Funding: This work was supported by Spanish grants of the Ministerio de Ciencia, Innovación y Universidades [grant numbers: TIN2013-42351-P, TIN2015-70308-REDT, TIN2016-76406-P]; project CASI-CAM-CM supported by Madri+d
530 [grant number: S2013/ICE-2845]; project FACIL supported by Fundación BBVA (2016); and the UAM-ADIC Chair for Data Science and Machine Learning.

References

- [1] M. Buhmann, Radial Basis Functions: Theory and Implementations, Cambridge Monographs on Applied and Computational Mathematics, Cambridge
535 bridge University Press, 2003.

- [2] J. Wang, G.R.Liu, A point interpolation meshless method based on radial basis functions, *International Journal for Numerical Methods in Engineering* 54 (11) (2002) 1623–1648.
- [3] R. Beatson, W. Light, Fast evaluation of radial basis functions: methods for two-dimensional polyharmonic splines, *IMA Journal of Numerical Analysis* 17 (3) (1997) 343–372.
- [4] M. Carozza, S. Rampone, An incremental multivariate regression method for function approximation from noisy data, *Pattern Recognition* 34 (3) (2001) 695 – 702.
- [5] P. Burt, E. Adelson, The laplacian pyramid as a compact image code, *IEEE Transactions on Communications* 31 (1983) 532–540.
- [6] E. Nadaraya, On estimating regression, *Theory of Probability & Its Applications* 9 (1) (1964) 141–142.
- [7] G. Watson, Smooth regression analysis, *Sankhyā: The Indian Journal of Statistics, Series A* (1964) 359–372.
- [8] N. Duchateau, M. De Craene, M. Sitges, V. Caselles, Adaptation of multi-scale function extension to inexact matching: application to the mapping of individuals to a learnt manifold, in: *International Conference on Geometric Science of Information*, Springer, 2013, pp. 578–586.
- [9] Y. Bengio, J.-f. Paiement, P. Vincent, O. Delalleau, N. L. Roux, M. Ouimet, Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering, in: *Advances in neural information processing systems*, 2004, pp. 177–184.
- [10] A. W. Long, A. L. Ferguson, Landmark diffusion maps (l-dmaps): Accelerated manifold learning out-of-sample extension, *Applied and Computational Harmonic Analysis* 47 (1) (2019) 190–211.

- [11] G. Mishne, I. Cohen, Multiscale anomaly detection using diffusion maps, *IEEE Journal of selected topics in signal processing* 7 (1) (2012) 111–123.
- [12] R. Coifman, S. Lafon, Diffusion Maps, *Applied and Computational Harmonic Analysis* 21 (1) (2006) 5–30.
- [13] G. Mishne, R. Talmon, I. Cohen, Graph-based supervised automatic target detection, *IEEE Transactions on Geoscience and Remote Sensing* 53 (5) (2014) 2738–2754.
- [14] G. Mishne, I. Cohen, Multiscale anomaly detection using diffusion maps and saliency score, in: *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2014, pp. 2823–2827.
- [15] G. Mishne, I. Cohen, Multi-channel wafer defect detection using diffusion maps, in: *2014 IEEE 28th Convention of Electrical & Electronics Engineers in Israel (IEEEI)*, IEEE, 2014, pp. 1–5.
- [16] S. M. N. Spingarn, I. Cohen, Voice activity detection in transient noise environment using laplacian pyramid algorithm, in: *2014 14th International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2014, pp. 238–242.
- [17] I. C. M. Li, S. Mousazadeh, Multisensory speech enhancement in noisy environments using bone-conducted and air-conducted microphones, in: *2014 IEEE China Summit & International Conference on Signal and Information Processing (ChinaSIP)*, 2014, pp. 1–5.
- [18] R. Coifman, S. Lafon, Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions, *Applied and Computational Harmonic Analysis* 21 (1) (2006) 31–52.
- [19] C. J. Dsilva, R. Talmon, N. Rabin, R. R. Coifman, I. G. Kevrekidis, Non-linear intrinsic variables and state reconstruction in multiscale simulations, *The Journal of chemical physics* 139 (18) (2013) 11B608_1.

- [20] E. Chiavazzo, C. W. Gear, C. J. Dsilva, N. Rabin, I. G. Kevrekidis, Reduced models in chemical kinetics via nonlinear data-mining, *Processes* 2 (1) (2014) 112–140.
- [21] D. Comeau, Z. Zhao, D. Giannakis, A. J. Majda, Data-driven prediction strategies for low-frequency patterns of north pacific climate variability, *Climate Dynamics* 48 (5-6) (2017) 1855–1872.
- [22] D. Comeau, D. Giannakis, Z. Zhao, A. J. Majda, Predicting regional and pan-arctic sea ice anomalies with kernel analog forecasting, *Climate Dynamics* 52 (9-10) (2019) 5507–5525.
- [23] R. Alexander, Z. Zhao, E. Székely, D. Giannakis, Kernel analog forecasting of tropical intraseasonal oscillations, *Journal of the Atmospheric Sciences* 74 (4) (2017) 1321–1342.
- [24] R. Duda, P. Hart, D. Stork, *Pattern Classification*, Wiley, New York, 2001.
- [25] A. Fernández, N. Rabin, D. Fishelov, J. Dorronsoro, Auto-adaptive laplacian pyramids, in: *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning - ESANN 2016*, i6doc.com, 2016, pp. 59–64.
- [26] P. Bühlmann, B. Yu, Boosting with the l_2 loss: regression and classification, *Journal of the American Statistical Association* 98 (462) (2003) 324–339.
- [27] M. Do, M. Vetterli, Framing pyramids, *IEEE Transactions on Signal Processing* 51 (2003) 2329–2342.
- [28] L. Liu, L. Gan, T. Tran, Lifting-based laplacian pyramid reconstruction schemes, in: *ICIP, IEEE*, 2008, pp. 2812–2815.
- [29] N. Rabin, R. Coifman, Heterogeneous datasets representation and learning using Diffusion Maps and Laplacian Pyramids, in: *SDM, SIAM / Omnipress*, 2012, pp. 189–199.

- [30] G. Mishne, I. Cohen, Multiscale anomaly detection using diffusion maps, J. Sel. Topics Signal Processing 7 (1) (2013) 111–123.
- [31] R. Adams, J. Fournier, Sobolev spaces, Vol. 140, Academic press, 2003.
- [32] D. Fishelov, A new vortex scheme for viscous flows, Journal of computational physics 86 (1) (1990) 211–224.
- [33] G. Cawley, N. Talbot, Fast exact leave-one-out cross-validation of sparse least-squares support vector machines, Neural Networks 17 (10) (2004) 1467–1475.
- [34] A. Elisseeff, M. Pontil, Leave-one-out error and stability of learning algorithms with applications, in: J. Suykens, G. Horvath, S. Basu, C. Micchelli, J. Vandewalle (Eds.), Advances in learning theory: methods, models and applications, NATO ASI, IOS Press, Amsterdam, Washington, DC, 2002, pp. 152–162.
- [35] R. Kohavi, A study of cross-validation and bootstrap for accuracy estimation and model selection, in: Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI’95, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995, pp. 1137–1143.
- [36] K. Fukunaga, D. Hummels, Leave-one-out procedures for nonparametric error estimates, IEEE Trans. Pattern Anal. Mach. Intell. 11 (4) (1989) 421–423.
- [37] T. Hastie, R. Tibshirani, J. Friedman, The elements of statistical learning: data mining, inference and prediction, Springer, 2008.
- [38] O. Chapelle, V. Vapnik, O. Bousquet, S. Mukherjee, Choosing multiple parameters for support vector machines, Machine Learning 46 (1) (2002) 131–159.
- [39] C. Rasmussen, C. Williams, Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning), The MIT Press, 2005.

- [40] Z. Liu, Q. Pan, J. Dezert, A. Martin, Adaptive imputation of missing values for incomplete pattern classification, *Pattern Recognition* 52 (2016) 85 – 95.
- [41] J. Xia, S. Zhang, G. Cai, L. Li, Q. Pan, J. Yan, G. Ning, Adjusted weight
645 voting algorithm for random forests in handling missing values, *Pattern Recognition* 69 (2017) 52 – 60.
- [42] H. Mann, D. Whitney, On a test of whether one of two random variables is stochastically larger than the other, *The Annals of Mathematical Statistics* 18 (1) (1947) 50–60.
- [43] I. Jolliffe, *Principal Component Analysis*, Springer Series in Statistics, 2002.
650
- [44] M. Lichman. UCI machine learning repository [online] (2013).
- [45] C. Higuera, K. Gardiner, K. Cios, Self-organizing feature maps identify proteins critical to learning in a mouse model of down syndrome, *PLOS ONE* 10 (6) (2015) 1–28.
- [46] N. Rabin, D. Fishelov, Missing data completion using diffusion maps and laplacian pyramids, in: *International Conference on Computational Science and Its Applications*, Springer, 2017, pp. 284–297.
655