FBE-IE-11/94-14

PARAMETRIC DISTANCE METRICS VS. NONPARAMETRIC NEURAL NETWORKS FOR ESTIMATING ROAD TRAVEL DISTANCES

ETHEM ALPAYDIN
İ. KUBAN ALTINEL
NECATİ ARAS

Ağustos 1994

August 1994

Fen Bilimleri Enstitüsü

Institute for Graduate Studies in Science and Engineering

Boğaziçi University, Bebek, Istanbul, Turkey

Boğaziçi Araştırmaları deneme niteliğinde olup, bilimsel tartışmaya katkı amacıyla yayınlandıklarından, yazar(lar)ın yazılı izni olmaksızın kendilerine atıfta bulunulamaz.

Boğaziçi Research papers are of a preliminary nature, circulated to promote scientific discussion and not to be quoted without written permission of the author(s).

Parametric Distance Metrics vs. Nonparametric Neural Networks for Estimating Road Travel Distances

Ethem Alpaydın†, İ. Kuban Altınel‡, Necati Aras‡*
† Department of Computer Engineering
‡ Department of Industrial Engineering
Boğaziçi University
TR-80815 İstanbul Turkey
{alpaydin,altinel,arasn}@ds5500.cc.boun.edu.tr

Abstract— *The actual distance between two cities is the length of the shortest road connecting them. Measuring and storing the actual distance between any two points of a region is often not feasible and it is a common practice to estimate it. The usual approach is to use theoretical distance metrics which are parameterized functions of the coordinates of the points. We propose to use nonparametric approaches using neural networks for estimating actual distances. We consider multi-layer perceptrons trained with the back-propagation rule and regression neural networks implementing nonparametric regression using Gaussian kernels. We also consider training multiple estimators and combining them in a hybrid architecture using voting and stacking. On a real-world study using cities drawn from Turkey, we found out that these approaches improve performance considerably. Estimating actual distances has many applications in location and distribution theory.*

Keywords— *Artificial Intelligence, Regression, Location, Road Transportation, Neural Networks.*

1 INTRODUCTION

The actual distance between any two points of a region can be defined as the length of the shortest road connecting them. Since it is often not feasible to measure the actual distances for all pairs of points, it is a common practice to use distance estimators. The question then is to choose a good estimator so that accurate distance approximations are obtained.

A good estimation of actual distances is critical in many applications. Almost all of the location problems, distribution problems such as the transportation problem, its generalization the transshipment problem, the travelling salesman problem, and the vehicle routing problem assume the knowledge of actual distances in their formulations. For example, in their recent simulation study to determine the number of fire-stations in İstanbul, Erkut and Polat multiply the Euclidean distance by an inflation factor, which

*Names are in alphabetical order.

they call the *road coefficient*, in order to estimate the actual distance between the fire-station and fire area [12].

We can define the problem of distance estimation formally as follows: Let us say that $\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2)^T$ where $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are two points on the Cartesian plane with coordinates $\boldsymbol{x}_1 = (x_{11}, x_{12})^T$ and $\boldsymbol{x}_2 = (x_{21}, x_{22})^T$. The aim is to build an estimator $d(\boldsymbol{x}|\boldsymbol{\theta})$ of the actual distance between $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$. $\mathcal{S}$ is a sample of n pairs $(\boldsymbol{x}^i, r^i)$ where r^i is the actual distance between $\boldsymbol{x}_1^i$ and $\boldsymbol{x}_2^i$ with $i = 1 \ldots n$. $\boldsymbol{\theta}$ is a vector of parameters estimated using S minimizing the following error measure:

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} [E(\mathcal{S}|\boldsymbol{\theta})] = \arg\min_{\boldsymbol{\theta}} \left[\frac{1}{n} \sum_{i=1}^{n} \psi \left(d(\boldsymbol{x}^i|\boldsymbol{\theta}), r^i \right) \right] \tag{1}$$

$\psi(\cdot)$ is the difference measure. One possibility, originally proposed by Love and Morris [18], is the absolute value of the deviation:

$$\psi \left(d(\boldsymbol{x}^i|\boldsymbol{\theta}), r^i \right) = |d(\boldsymbol{x}^i|\boldsymbol{\theta}) - r^i| \tag{2}$$

According to this criterion, a distance metric must estimate greater actual distances relatively more accurately than shorter distances. This is a drawback if we are more interested in proportional deviations than absolute deviations. Another error measure, also proposed by Love and Morris [18], is normalized:

$$\psi \left(d(\boldsymbol{x}^i|\boldsymbol{\theta}), r^i \right) = \left[\frac{d(\boldsymbol{x}^i|\boldsymbol{\theta}) - r^i}{\sqrt{r^i}} \right]^2 \tag{3}$$

Although both criteria are very insightful on their own, the second one is superior not only because it gives importance to proportional errors but also because of the following three reasons. First, most of the experimental results in the literature use the second criterion, e.g., [4,5,18,32]. It has important statistical properties which leads to statistical tests for comparing the accuracy of distance metrics under certain normality and independence assumptions and thus the results obtained can have further use. Finally, it is a continuous and differentiable function of the parameter vector and this enables the use of gradient descent methods for minimization which is important in neural network learning as we will see in Section 3.

The standard approach for distance estimation uses distance metrics that are parameterized functions of certain easy to obtain information, namely the coordinates of the points. This approach has been widely used ever since the first work by Love and Morris [18] because it provides simple analytical closed form expressions of the coordinates once the value of the parameters are determined. As any parametric method, it works with small samples but the accuracy may not be high if the assumed form of the metric is not appropriate.

We view the problem of estimating distances in the context of function approximation or nonlinear regression and apply neural network based estimators for this task of estimating $d(\boldsymbol{x}|\boldsymbol{\theta})$. These methods, being nonparametric, have the advantage that they do not assume any a priori model and are trained directly from a training sample. They of course necessitate larger training samples as we no longer have the comfort of a parametric model with just a few parameters.

One approach we use is the *multi-layer perceptron* trained with the back-propagation rule. The second approach is based on nonparametric regression using kernel density estimation; in the neural network literature, this is named a *regression* neural network. A recent idea in machine learning is to combine multiple estimators for improved performance. We test two methods for this. One trains multiple estimators and takes a *vote*; when vote weights are equal, this is simple averaging. The second *stacking* approach also trains a combiner network. We provide a comparison of all explained approaches based on a real-world sample collected by pairing a subset of the cities of Turkey.

This paper is organized into seven sections. The next section is a literature survey on the use of suitable distance metrics to estimate actual distances. In sections three and four, we propose neural-network-based

Table 1: Distance functions used and their associated parameters.

DISTANCE FUNCTION	PARAMETERS ($\boldsymbol{\theta}$)
$d_1(\boldsymbol{x}) = k(\lvert x_{11} - x_{21}\rvert + \lvert x_{12} - x_{22}\rvert)$	k
$d_2(\boldsymbol{x}) = k(\lvert x_{11} - x_{21}\rvert^2 + \lvert x_{12} - x_{22}\rvert^2)^{1/2}$	k
$d_3(\boldsymbol{x}) = k(\lvert x_{11} - x_{21}\rvert^p + \lvert x_{12} - x_{22}\rvert^p)^{1/p}$	k, p
$d_4(\boldsymbol{x}) = k(\lvert x_{11} - x_{21}\rvert^p + \lvert x_{12} - x_{22}\rvert^p)^{1/s}$	k, p, s

methods, i.e., a multi-layer perceptron and a regression neural network, for distance estimations. Section five explains the idea of combining multiple estimators. In section six, we give results obtained for Turkey using the previously explained approaches. Finally in the last section, we compare approaches using the criteria of estimation accuracy, memory requirement, and learning time.

2 DISTANCE METRICS

A generally used method for estimating actual road distances between any pair of points is to make approximations by means of a distance metric. A distance metric is a parameterized function of the planar coordinates of the two points. These metrics can be classified in three major groups with respect to the type of coordinates they use. The metrics of the first group use spherical coordinates for the purpose of introducing the spherical effect of the earth surface into the distance estimation [18,19]. Although this idea provides certain additional accuracy, the contribution has been experimentally reported to be minor by Love and Morris [18]. The second group consists of metrics which use polar coordinates [24,26]. The motivation is based on the observation that the roads in old cities are not usually planned according to a rectangular grid structure and therefore distances can be approximated better by a ring-radial measure. This approach seems to be very accurate especially for a spider's web-like road network structure. The third group contains some simple functions of the cartesian coordinates. These are mostly norms or norm-based metrics. We list four of them in Table 1. They are the most important of all distance metrics because of their wide usage in location and distribution problems [13,21] and we adopt them in this study.

The parameters k, p, and s constitute $\boldsymbol{\theta}$ and are estimated over a sample to provide good approximations and as such, encode geographical characteristics of the region where they are used. There is a large literature on the determination of these parameters and the comparison of the distance metrics, sometimes with conflicting results [5,6,7,18,19,20].

For all practical purposes, the metric chosen to estimate actual road distances should be as accurate as possible. In their early study, Love and Morris [18,19] compute the parameters k, p, and s of $d_1(\boldsymbol{x}), d_2(\boldsymbol{x}), d_3(\boldsymbol{x})$, and $d_4(\boldsymbol{x})$ for the United States and compare them with respect to the accuracy they provide. The important conclusion of this study is the superiority of $d_4(\boldsymbol{x})$ over the other three. The second best is $d_3(\boldsymbol{x})$.

At the end of their study on the road network of the former Federal Republic of Germany (FRG), Berens and Körling [6] and Berens [5] conclude that the accuracy provided by the weighted Euclidean norm $d_2(\boldsymbol{x})$ is sufficient and the use of $d_3(\boldsymbol{x})$ is not worth the extra computational effort necessary for calculations. However in a further study over the largest 25 cities of FRG, Love and Morris [20] report conflicting results which demonstrate that the accuracy of the weighted L_p norm, $d_3(\boldsymbol{x})$, is remarkably higher than the accuracy provided by $d_2(\boldsymbol{x})$. Although it supports the early findings of Berens and Körling [6] for FRG, the study by Berens [5] includes mixed results when it is enlarged to cover 11 other countries; the relative improvement introduced by $d_3(\boldsymbol{x})$ over $d_2(\boldsymbol{x})$ ranges within 0.00% and 11.27%. Finally, Berens and Körling [7], in their last comment, state that the empirical distance functions should be tailored for

the regions they are to be used if the accuracy is the main interest and that there is no one general distance metric which provides the same accuracy all over the world.

There are also distance measures which do not fit completely in any of the above mentioned three groups. They can be mostly included in the last one but they are not always simple functions of the coordinates and require additional information such as a rotating angle for the coordinate axes or vectors for possible directions on a typical road. All of them are based on the idea that a travel has two major components; rectilinear and Euclidean elements, and the actual distance between any pair of points can be modelled as a positive linear combination of them. Ward and Wendell [32] initiate this hybrid idea by suggesting the weighted one-infinity norm and observe that the accuracy of this function is relatively close to the accuracy of the weighted L_p norm, $d_3(\boldsymbol{x})$, on the data set of Love and Morris [18]. In their later work, they generalize the one-infinity norm to obtain the family of block norms in which the accuracy of the approximation depends on possible travel directions [33]. They report that the approximations obtained by the weighted L_p norm are more accurate than those obtained by a two parameter block norm, which is actually the weighted one-infinity norm, and the accuracy of the weighted L_p norm is slightly worse than the one of eight parameter block norm's. Block norms play an important role in location models because they lead to linear programming problems for certain objective functions, such as the minimax distance function; but the size of the linear program can become very large easily. Most recent work on the use of hybrid norms are due to Brimberg and Love [9,10]. They suggest the use of weighted one-two norm to approximate $d_3(\boldsymbol{x})$ in estimating distances.

3 MULTI-LAYER PERCEPTRONS

A multi-layer perceptron is a feedforward type of neural network [16]. Each unit in a neural network performs a weighted sum of its inputs. In a feedforward network, input is fed to the input layer which propagates to the output layer through possibly a layer of nonlinear hidden units. Weights between consecutive layers encode interdependencies or constraints between the values that the units take simultaneously. It has been shown that such a network with one hidden layer is a universal approximator, i.e., can approximate any function with the desired accuracy [14,17]. A feedforward neural network with one hidden layer of H hidden units operates as follows:

$$F(\boldsymbol{u}|\boldsymbol{T},\boldsymbol{W}) = \sum_{h=1}^{H} T_h g_h(\boldsymbol{W}_h,\boldsymbol{u}) + T_0 \tag{4}$$

$\boldsymbol{u}$ is the multi-dimensional input vector. T_h is the weight from hidden unit h to the output unit. T_0 is the 'bias' weight. The output thus is simply a weighted sum of the hidden basis functions, $g_h(\boldsymbol{W}_h,\boldsymbol{u})$. Hidden units compute a weighted sum of the input $\boldsymbol{u}$ with their weights $\boldsymbol{W}_h$ and then filter this value through a nonlinear sigmoid activation function which is a smoothed version of thresholding:

$$H_h = g_h(\boldsymbol{W}_h,\boldsymbol{u}) = 1/\left[1+\exp\left(-\sum_{i=1}^{k} u_i W_{ih} - W_{0h}\right)\right] \tag{5}$$

The most popular training method for multi-layer perceptrons is the back-propagation learning algorithm [28] which implements gradient descent over an error function. Given a differentiable error function E, we start with any set of weights and repeatedly change each parameter θ_p by an amount proportional to $\partial E/\partial\theta_p$:

$$\Delta\theta_p = -\eta\frac{\partial E}{\partial\theta_p} \tag{6}$$

where η is the stepsize in descent.

In our application of distance estimation, we use the error function of Equations (1) and (3). The input to the neural network is a vector function of the coordinates of the two points: $\boldsymbol{u}^i = \boldsymbol{\phi}(\boldsymbol{x}^i)$. Similarly the output of the neural network by a known transformation gives us the desired distance: $y^i = d(\boldsymbol{x}^i|\boldsymbol{\theta}) = \chi(F(\boldsymbol{u}^i|\boldsymbol{T}, \boldsymbol{W}))$. $\boldsymbol{\theta}$, the parameter vector of our estimator corresponds to the weights of the neural network, i.e., $\boldsymbol{W}$ and $\boldsymbol{T}$. The problem of choosing good *coding* or *representation* functions, $\boldsymbol{\phi}(\cdot)$ and $\chi(\cdot)$, is critical in the application of neural networks. Noting that $g'(a) = g(a)(1 - g(a))$, this leads to the following parameter update equations:

$$\begin{aligned}
\Delta T_0 &= \eta \left(\frac{y^i - r^i}{\sqrt{r^i}} \right) \\
\Delta T_h &= \eta \left(\frac{y^i - r^i}{\sqrt{r^i}} \right) H_h \\
\Delta W_{0h} &= \eta \left(\frac{y^i - r^i}{\sqrt{r^i}} \right) T_h H_h (1 - H_h) \\
\Delta W_{ih} &= \eta \left(\frac{y^i - r^i}{\sqrt{r^i}} \right) T_h H_h (1 - H_h) u_i
\end{aligned} \tag{7}$$

4 REGRESSION NEURAL NETWORKS

Nonparametric density estimate of probability at a certain value is a weighted sum of the effects of sample points. Given a sample of n real d-dimensional $\boldsymbol{X}_i$ values, the estimate of probability at $\boldsymbol{x}$ is [29]:

$$\hat{p}(\boldsymbol{x}) = \frac{1}{nh^d} \sum_{i=1}^{n} K \left(\frac{\boldsymbol{x} - \boldsymbol{X}_i}{h} \right) \tag{8}$$

K is the kernel density function and h is the window width of the kernel. We can use this approach for regression, i.e., for inference of a scalar y value for a given multi-dimensional vector $\boldsymbol{x}$ if we are given pairs of $(\boldsymbol{X}_i, Y_i)$. Our estimate of the joint probability density $\hat{f}(\boldsymbol{x}, y)$ is then written as:

$$\hat{f}(\boldsymbol{x}, y) = \frac{1}{nh^d} \sum_{i=1}^{n} K \left(\frac{\boldsymbol{x} - \boldsymbol{X}_i}{h} \right) \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{(y - Y_i)^2}{2\sigma^2} \right] \tag{9}$$

Using Bayes theorem:

$$\hat{f}(y|\boldsymbol{x}) = \frac{\hat{f}(\boldsymbol{x}, y)}{f(\boldsymbol{x})} = \frac{\hat{f}(\boldsymbol{x}, y)}{\int \hat{f}(\boldsymbol{x}, y) dy} \tag{10}$$

Our estimate of y is the expected value:

$$E[y|\boldsymbol{x}] = \frac{\int y \hat{f}(\boldsymbol{x}, y) dy}{\int \hat{f}(\boldsymbol{x}, y) dy} \tag{11}$$

Replacing integrals by sums and after interchanging the order of integration and summation and cancelling common terms:

$$\hat{y}(\boldsymbol{x}) = \frac{\sum_i Y_i K\left((\boldsymbol{x} - \boldsymbol{X}_i)/h\right)}{\sum_i K\left((\boldsymbol{x} - \boldsymbol{X}_i)/h\right)} \tag{12}$$

Because of its attractive analytical properties [29], the kernel function is generally taken as the Gaussian, named a *Parzen window* [11]:

$$K(\boldsymbol{u}) = \left(\frac{1}{\sqrt{2\pi}}\right)^d \exp\left[-\frac{||\boldsymbol{u}||^2}{2}\right] \tag{13}$$

This then leads to the following equation for our estimate $\hat{y}(\boldsymbol{x})$:

$$\hat{y}(\boldsymbol{x}) = \frac{\sum_{i=1}^n Y_i \exp\left[-||\boldsymbol{x} - \boldsymbol{X}_i||^2/2h^2\right]}{\sum_{i=1}^n \exp\left[-||\boldsymbol{x} - \boldsymbol{X}_i||^2/2h^2\right]} \tag{14}$$

The estimate $\hat{y}(\boldsymbol{x})$ is thus a weighted average of all the observed Y_i values where each weight is exponentially proportional to its Euclidean distance from $\boldsymbol{x}$. Note that a good choice of the window width h is critical. If it is large then even distant neighbors affect the estimate at x leading to a very smooth estimate. In the extreme case when h goes to infinity, our estimate is the average of all Y_i, independent of the input. When h is small, only a few (if any) samples play a role, leading to a noisy estimate.

The "bias/variance dilemma" stated by Geman, Bienenstock, and Doursat [15], points out that when h is large, $\hat{y}(\boldsymbol{x})$ is an average of many samples and the variance contribution of our estimator is small as our estimator does not change from one training set to another (of course assuming that they are all drawn from the same distribution). The bias is large in this case as our estimate is biased toward the population response. When h is small, there is small bias but the estimate is dependent on the particular training set used. Therefore the variance contribution is high. Thus choosing h implies a trade-off between bias (systematic error) and variance (random error). h is determined by trial and error, cross-validating on a separate test set. A more detailed analysis of the approach and its properties are beyond the scope of this present paper; the interested reader is referred to [30].

Back-propagation rule is generally criticized on the grounds that learning takes many epochs over the training set and also that the performance depends on a good choice of H, the number of hidden units. The method just explained belongs to the class of *memory-based methods* where approximation is directly done from a table of stored values. These methods require only one epoch over the training set, i.e., to read in the table. Here $\boldsymbol{\theta}$, the parameter vector includes the whole training set and h. Thus learning is fast but the memory requirement is large. There are also methods to choose a subset of the training set without decreasing performance; see [2] for a review.

As also noted by Specht who names this a *general regression neural network* [31], this approach can also be written in the framework of Equation (4). Here H is equal to table size, i.e., the number of training patterns with $\boldsymbol{W}_h = \boldsymbol{u}^h$, $T_h = y^h$, the desired output for $\boldsymbol{u}^h$, and we have:

$$g_h(\boldsymbol{W}_h, \boldsymbol{u}) = \frac{\exp\left[-||\boldsymbol{W}_h - \boldsymbol{u}||^2/2h^2\right]}{\sum_{i=1}^H \exp\left[-||\boldsymbol{W}_i - \boldsymbol{u}||^2/2h^2\right]} \tag{15}$$

The problems of finding good input and output codings apply also here.

5 COMBINING MULTIPLE ESTIMATORS

There are many choices that should be made before an estimator can be used in an application. These include the model, method of estimating parameters, i.e., learning rule, training sample, input representation, error function minimized, etc. Each choice made is one additional sort of bias which may not be appropriate. The general approach is to try alternatives and choose the one that best performs on a test sample, distinct from the sample used for training, i.e., cross-validation. We advocate in this section that it is better not to discard alternatives but combine all to improve performance.

There are two ways in which multiple estimators can be combined: One is *voting* and the other is *stacking*. Note that regardless of the way we combine the models, there is no reason to expect improvement in performance through having multiple models if they are quite similar. One can only get fault tolerance

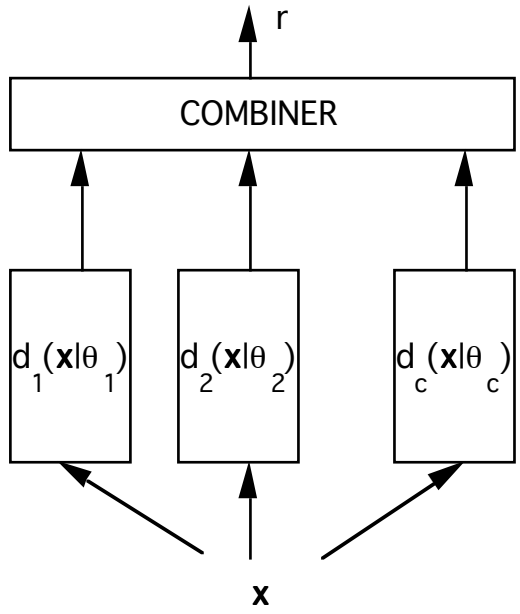

Figure 1: Block diagram of combining multiple estimators. The combiner takes the predictions of the estimators as input and makes the final prediction. The combiner can be a fixed voting system or an estimator that is also trained.

through redundancy if the models fail under different circumstances. In neural networks, choosing parameters randomly like the initial weight values and hyperparameters like the network architecture or by having different learning rules, one guarantees this to a certain extent. Basically as we will see in the next subsection, when we have one model that has a certain success, we want to add a model that succeeds best for inputs on which the first one fails; we do not care about the new model's *overall* performance.

As seen in Figure 1, we have c estimators, $d_i(\boldsymbol{x}, \boldsymbol{\theta}_i)$, each one trained separately. Then during estimation, each one gives an output and these outputs are combined by a *combiner* system to determine the final output r. We mention two types of combiners: The *voting* system computes a weighted sum where weights are fixed. In the *stacking* approach, the combiner is also an estimator that is trained.

In our application, we combined the three estimators using both types of combining; parametric, multilayer perceptron, and the regression neural network. We believe them to be sufficiently different in our application of distance estimation.

5.1 Voting

In voting with c voters, each estimator $d_j(\boldsymbol{x}|\boldsymbol{\theta}_j)$ gives an output for the input $\boldsymbol{x}$. The final output, r, is a weighted sum [1]:

$$r = \sum_{j=1}^{c} d_j(\boldsymbol{x}|\boldsymbol{\theta}_j)\omega_j \tag{16}$$

where $\sum_{j=1}^{c} \omega_j = 1$. Unless there is a prior reason to favor one voter over another, weights are taken as equal: $\omega_j = 1/c$. If we look at the variance of r [22,27]:

$$\begin{aligned}\mathrm{Var}(r) &= \mathrm{Var}\left(\frac{1}{c}\sum_{j=1}^{c} d_j\right) \\ &= \frac{1}{c^2}\left[\sum_j \mathrm{Var}(d_j) + 2\sum_j \sum_{k\neq j} \mathrm{Cov}(d_j, d_k)\right] \end{aligned} \quad (17)$$

If the models are statistically independent, the second term cancels and variance (and mean square error) decreases with increasing c. However variance can further be decreased by choosing mutually "orthogonal" methods [34]. For example if we have two models, instead of having them independent, we do much better by having them negatively correlated.

It may be possible to have an a priori preference of one estimator over another, then we would like to have the weights of those estimators in voting higher. We can for example assess them over a cross-validation set and use their performances to determine weights so as to assign larger weights to more successful voters. Using Occam's razor, we can also a priori assume that complex models with many free parameters tend to overfit on a small sample. Thus we may like to have the tendency to give them less weight unless they are more successful than simpler models, i.e., their additional complexity is justified [1]. We have not tested these two latter approaches in this present work.

5.2 Stacking

Combination of multiple estimators can also be considered as an estimation problem. Wolpert who initiated the idea calls it *stacked generalization* [34]. Breiman converted it to statistical language as *stacked regression* [8]. Here we have two levels of estimators. First on level 0, L_0, are our usual estimators chosen carefully to be as different as possible while maintaining individually acceptable levels of performances. On top of these, on level 1, L_1, is another estimator (the combiner) whose inputs are the outputs of the first level estimators.

L_0 estimators are trained separately on the same training sample. Then they are introduced a different set on which they make estimations. The L_1 estimator takes those estimations as input and maps them to the correct value. Thus L_0 and L_1 estimators are trained on separate training sets. This is because the performances of the L_0 estimators on the data that they are trained on can be very different from their behavior on the data that they are *not* trained on; this is exactly what L_1 should know about and be trained on. This is generally done by dividing the training set into two and using one half for training each level. When the training set is small, one can use leave-k-out (jackknife) for not to lose precious data [23]. Significant improvement in success has been achieved using this approach on protein secondary structure prediction [35].

6 IMPLEMENTATION RESULTS

6.1 Database

We have collected two distinct samples by pairing respectively 28 and 23 cities of Turkey for training and test sets. The first is used for estimating the parameters and the second one to assess their performances. The training and test data sets contain planar coordinates and intercity distances for respectively 28 and 23 cities which make $28 * 27/2 = 378$ and $23 * 22/2 = 253$ data pairs. There are 74 cities in Turkey thus the complete data set contains 74*73/2=2701 pairs and our training set constitutes 14% of it. The third dimension is ignored since previous empirical studies have shown that the effect of height in the accuracy

of estimations in Turkey is almost null [4]. The values we report in the following sub-sections are average error per pair in kilometers on the test set.

6.2 Distance Functions

By looking at properties of the application, it is realistic to eliminate some of the metrics a priori by judging them with the structural properties of the actual road network. First of all, the road structure in Turkey has been developed arbitrarily rather than rectilinearly or ring-radially. This arbitrariness makes the identification of a fixed pattern for possible travel directions within the country impossible; this is crucial for the use of hybrid norms. Besides, the area is too small to require the consideration of the earth's roundness in estimating actual distances. Being convinced by these observations, it is rational to concentrate on metrics $d_2(\boldsymbol{x}), d_3(\boldsymbol{x}), d_4(\boldsymbol{x})$ and compute the best possible value of the parameters k, p, and s. In spite of this fact, we also computed the value of k for $d_1(\boldsymbol{x})$, since it has been heavily considered in location literature [13].

The calculation of the parameters with respect to any of the estimation criteria introduced in Section 1 requires minimizing an error function like Equation (1). Since we use the normalized error function given in Equation (3), the minimization problems are continuous in the parameters. They are also well-behaved in the sense that any Kuhn-Tucker point with respect to the unknowns k, p, and s is optimal. The first two problems have analytical solutions and the values of k which minimize error function (3) can be obtained easily by using the following equalities:

For $d_1(\boldsymbol{x})$:

$$k = \frac{\sum_{i=1}^{n} \left(|x_{11}^i - x_{21}^i| + |x_{12}^i - x_{22}^i) \right.}{\sum_{i=1}^{n} \left(|x_{11}^i - x_{21}^i| + |x_{12}^i - x_{22}^i| \right)^2 / r^i} \tag{18}$$

and for $d_2(\boldsymbol{x})$:

$$k = \frac{\sum_{i=1}^{n} \left(|x_{11}^i - x_{21}^i|^2 + |x_{12}^i - x_{22}^i|^2 \right)^{1/2}}{\sum_{i=1}^{n} \left(|x_{11}^i - x_{21}^i|^2 + |x_{12}^i - x_{22}^i|^2 \right) / r^i} \tag{19}$$

However, the calculations for k and p of $d_3(\boldsymbol{x})$, and k, p and s of $d_4(\boldsymbol{x})$ are slightly more complicated. They require the solution of the following unconstrained optimization problems:

$$\min_{k,p} \sum_{i=1}^{n} \left[\frac{k \left(|x_{11}^i - x_{21}^i|^p + |x_{12}^i - x_{22}^i|^p \right)^{1/p} - r^i}{\sqrt{r^i}} \right]^2 \tag{20}$$

$$\min_{k,p,s} \sum_{i=1}^{n} \left[\frac{k \left(|x_{11}^i - x_{21}^i|^p + |x_{12}^i - x_{22}^i|^p \right)^{1/s} - r^i}{\sqrt{r^i}} \right]^2 \tag{21}$$

Although they are quite simple with respect to the number of variables, which is two and three, the number of nonlinearities induced by these problems can be very large depending on the size of the pairs of points within the training set. Their minimizations were carried out using MINOS 5.1 [25]. The results using distance functions are given in Table 2. Observe that the rectilinear distance metric has the worst performance. Meanwhile the accuracy of $d_4(\boldsymbol{x})$ is the highest. These facts definitely support our previous inference on the arbitrariness of the road structure in Turkey.

6.3 Multi-Layer Perceptrons

We employed a multi-layer perceptron with one hidden layer with the back-propagation learning rule. After several trials, the best input representation was determined as the four data values and the Euclidean distance in between as a hint:

Table 2: Average error on the test set per pair using the four parametric distance functions.

	DISTANCE FUNCTION			
DATA SETS	$d_1(\boldsymbol{x})$	$d_2(\boldsymbol{x})$	$d_3(\boldsymbol{x})$	$d_4(\boldsymbol{x})$
Parameters	k=1.073	k=1.310	k=1.262, p=1.603	k=1.688, p=1.686, s=1.761
Average Error	12.49	9.05	8.48	8.10

$$\boldsymbol{u} = \boldsymbol{\phi}(\boldsymbol{x}) = (x_{11}, x_{12}, x_{21}, x_{22}, ||\boldsymbol{x}_1 - \boldsymbol{x}_2||)^T$$

In terms of output representation, we found out that learning the ratio of actual distance to Euclidean distance is better than learning the actual distance itself. This ratio is called the *directional bias* by Brimberg and Love [9,10]:

$$y = \chi(r) = r/||\boldsymbol{x}_1 - \boldsymbol{x}_2||$$

This may be thought of as extending $d_1(\boldsymbol{x})$ in the parametric case in that, instead of computing one constant global k factor, it is as if the neural network computes a continuous function $k(\boldsymbol{x})$ by which it scales the Euclidean distance. Note that in these two cases, we can take advantage of our a priori knowledge that $d(\boldsymbol{x}_1, \boldsymbol{x}_2) = d(\boldsymbol{x}_2, \boldsymbol{x}_1)$ and effectively double the training set. In Figure 2, average errors and standard deviations over the test set are given over 10 independent runs as a function of the number of hidden units.

We used a momentum term and updated η dynamically for faster convergence. Training stops when no further improvements were made. Networks require around a hundred epochs for convergence after which slight improvements (less than 5%) are achieved if learning is continued.

6.4 Regression Neural Networks

By playing with the alternatives, we found out that the following input representation was the best:

$$\boldsymbol{u} = \boldsymbol{\phi}(\boldsymbol{x}) = (|x_{11} - x_{21}|, |x_{12} - x_{22}|, ||\boldsymbol{x}_1 - \boldsymbol{x}_2||)^T$$

This is better than the five dimensional input as there is no loss of information and distance-based estimators always suffer from the "curse of dimensionality" [11] and one always does better by minimizing the input dimensions. The output was the desired distance:

$$y = \chi(r) = r$$

Results achieved are given in Figure 3 as a function of the window width, h. Note that there is a large dependence on h. Overall, the success is slightly inferior to that of the multi-layer perceptron.

6.5 Combining Estimators

6.5.1 Voting

Taking into account the fact that combining multiple estimators is a better idea when estimators are as different as possible, we used the following three estimators:

1. Parameterized distance metric, $d_4(\boldsymbol{x})$ with three parameters k, p, and s.

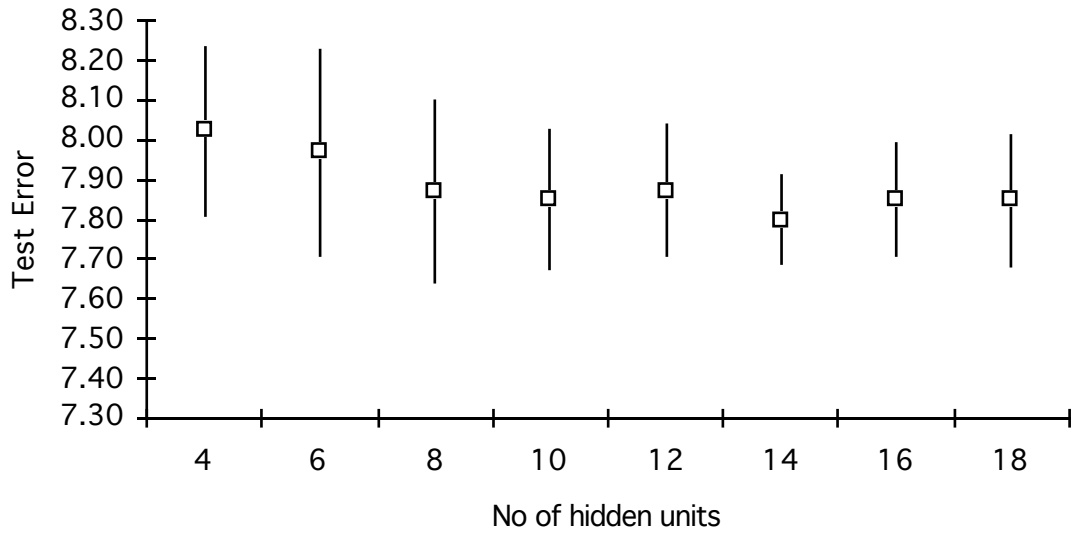

Figure 2: Average error of multi-layer perceptrons on the test set as a function of the number of hidden units after 10 independent runs. Squares denote means and error bars are one standard deviation long on either side.

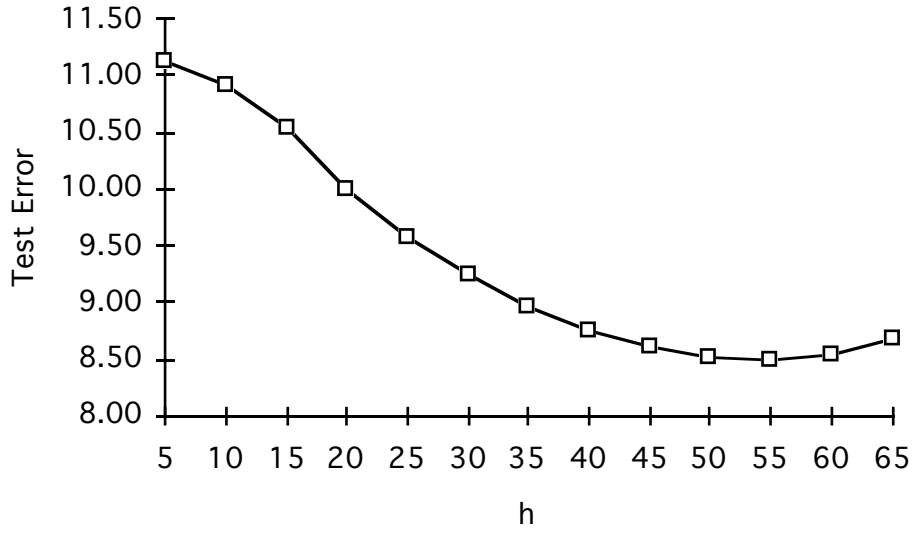

Figure 3: Average error of regression neural network on the test set as a function of the window width.

Table 3: Average error of three estimators on the test set and the result of voting.

ESTIMATORS	ERROR
Parametric model	8.10
Multi-layer perceptron	7.91
Regression neural network	8.49
Voting	7.63

Table 4: Average error achieved through voting vs. stacking.

COMBINER	ERROR
Voting	7.69
Stacking	7.41

2. Multi-layer perceptron with 12 hidden units.
3. Regression neural network with $h = 54$.

The votes are taken as equal. Results achieved are given in Table 3. Note that the result of voting is better than the result of all voters. This indicates that estimators do fail under different circumstances.

6.5.2 Stacking

We used again three estimators as our L_0 estimators:

1. Parameterized distance metric, $d_2(\boldsymbol{x})$ with the parameter k.
2. Multi-layer perceptron with 12 hidden units.
3. Regression neural network with $h = 54$.

The combiner L_1 estimator is a multi-layer perceptron with 2 hidden units. We divided the training set into two parts performing "leave-n/2-out"; T_0 and T_1. We have trained all three L_0 estimators using T_0 only and saved their predictions on T_1 as R_1. Then we have trained them on T_1 and saved their predictions on T_0 as R_0. We have trained the combiner network L_1 on R_0 and R_1. After training the combiner network, we have trained all three on the complete training set. Stacking is superior to voting as is also shown in our study. Values reported are average errors on the test set (Table 4).

7 CONCLUSIONS

We compare approaches based on the three criteria of success, memory requirement, and time.

- *Success.* We see that the multi-layer perceptron is a better estimator than the parameterized distance metrics in terms of estimation accuracy. This is because in the latter models, the parameters are fixed over the whole space. For example when $d_2(\boldsymbol{x})$ is used, k is constant over the whole country but

we know that for example in Turkey, the eastern part is more mountainous than the western part and thus k there is higher than k of the west, as a consequence of the extra curvature caused by these natural obstacles. The multi-layer perceptron may be seen as approximating an input dependent, continuous function $k(\boldsymbol{x})$. Being nonparametric, it does this without assuming any a priori form. An approach in distance metric based estimation is to divide the space into multiple regions and make separate estimations in separate regions; this is a piecewise constant approximation of the $k(\boldsymbol{x})$ [3].

Though the regression neural network and the parameterized distance metric do not perform as well as the multi-layer perceptron, when they are combined, their accuracy is higher. This indicates that on cases where the perceptron fails, the others make good guesses. The stacking approach as it is trained, performs better than voting.

- *Memory Requirement.* Out of the three estimators we have implemented, the regression neural network has the highest memory requirement by storing $4n + 1$ parameters: 4 values for each pair for a training set of n pairs and the parameter h, here $4 * 378 + 1 = 1513$. When H is the number of hidden units, the multi-layer perceptron requires $6H$ (5 inputs and one bias) values for $\boldsymbol{W}$ and $H + 1$ for $\boldsymbol{T}$, thus total of $7H + 1$ weights, here 85. The parameterized distance metrics respectively have 1, 1, 2, and 3 parameters. Memory requirement increases when one has multiple models. In the stacking approach, there are also parameters of the combiner network.

- *Learning Time.* The regression neural network is trained in one epoch. The first two parameterized distance metric has an analytical solution where we compute k by making one pass over the training set. The computational effort spent for training the multi-layer perceptron is higher than that is required by the other two parametric models given in Equations (20) and (21), as many trials should be made with the multi-layer perceptron to find out the good input and output representations and the number of hidden units. To train the stacking network, when one divides the set into two, the estimators are trained twice. Ideally, "leave-one-out" where at each step, the training set is divided into two parts of size one and $n - 1$ is better for bias reduction but it requires n trainings [23]. This is the approach Wolpert originally proposed for stacked generalization [34].

It is generally accepted that using this or that learning technique by itself is not sufficient. A good knowledge of the application and the learning methods is necessary to guide the estimator to learn the important and ignore the irrelevant. As shown in the preceding sections, different methods may benefit from different input and output representations. Perhaps the quality of the training sample is the most important factor that affects the quality of estimation. For example, a large enough and representative sample is required for the approximation not to be biased by the idiosynchranicies of the particular sample used.

Viewing the results achieved, we can say that nonparametric approaches can be considered as better than parameterized distance metrics in terms of success. Success can further be improved by using multiple estimators and combining them especially by stacking. However there is not one method that is significantly superior to others in all three respects of success, memory requirement, and learning time. Thus when choosing a particular estimator, all three should be taken into account according to the particular implementation constraints and not just only success as it has hitherto frequently been done in the literature.

Acknowledgements

We thank Haluk Erkut and Seçkin Polat of the Department of Industrial Engineering, İstanbul Technical University for the first inspiration and continuous encouragement, and D.H. Wolpert of Santa Fe Institute

for giving pointers to related literature on combining estimators by stacking. EA is supported in part by Tübitak Grant EEEAG–41 and Boğaziçi University Research Funds Grant 93HA0132.

References

[1] Alpaydın, E., "Multiple Networks for Function Learning," *IEEE International Neural Network Conference*, San Francisco, vol 1, 9–14, March 1993.

[2] Alpaydın, E., "GAL: Networks That Grow When They Learn and Shrink When They Forget," *International Journal of Pattern Recognition and Artificial Intelligence*, 8 (1994) 391–414.

[3] Altınel, İ. K., and Aras, N., "Estimating Road Distances in İstanbul with Single and Multi Regional Models", Research Paper Series No : FBE-IE-05/94-05, Department of Industrial Engineering, Boğaziçi University, İstanbul, 1994.

[4] Altınel, İ. K., Aras, N., Alie, A., Cangür, G., Özel, R., and Yücel, A., "Estimating Road Travel Distances in Türkiye", Research Paper Series No : FBE-IE-03/94-03, Department of Industrial Engineering, Boğaziçi University, İstanbul, 1994.

[5] Berens, W., "The Suitability of the Weighted L_p norm in Estimating Actual Road Distances", *European Journal of Operational Research* 34 (1988) 39–43.

[6] Berens, W., and Körling, F., "Estimating Road Distances by Mathematical Functions", *European Journal of Operational Research* 21 (1985) 54–56.

[7] Berens, W., and Körling, F., "On Estimating Road Distances by Mathematical Functions–A Rejoinder", *European Journal of Operational Research* 36 (1988) 254–255.

[8] Breiman, L., "Stacked Regression," TR-367, Department of Statistics, University of California, Berkeley (1992).

[9] Brimberg, J., and Love, R. F., "A New Distance Function for Modelling Travel Distances in a Transportation Network", *Transportation Science* 26 (1992) 129–137.

[10] Brimberg, J., and Wesolowsky, G. O., "Probabilistic L_p Distances in Location Models", *Annals of Operations Research* 40 (1992) 67–75.

[11] Duda, R.O., Hart, P.E., *Pattern Classification and Scene Analysis*, John Wiley, 1973.

[12] Erkut, H., and Polat, S., "A Simulation Model for a Urban Fire Fighting System", *Omega* 20 (1992) 535–542.

[13] Francis, R. L., McGinnis L. F. Jr., and White, J. A., *Facility Layout and Location : An Analytical Approach*, 2nd edition, Prentice Hall, Englewood Cliffs, 1992.

[14] Funahashi, K., "On the Approximate Realization of Continuous Mapping by Neural Networks," *Neural Networks*, 2 (1989) 183–192.

[15] Geman, S., Bienenstock, E., Doursat, R., "Neural Networks and the Bias/Variance Dilemma," *Neural Computation*, 4 (1992) 1–58.

[16] Hertz, J., Krogh, A., and Palmer, R.G., *Introduction to the Theory of Neural Computation*, Addison Wesley, 1991.

[17] Hornik, K., Stinchcombe, M., White, H., "Multilayer Feedforward Networks are Universal Approximators," *Neural Networks* 2 (1989) 359–366.

[18] Love, R. F., and Morris, J. G., "Modelling Inter-City Road Distances by Mathematical Functions", *Operational Research Quarterly* 23 (1972) 61–71.

[19] Love, R. F., and Morris, J. G., "Mathematical Models of Road Travel Distances", *Management Sciences* 25 (1979) 130–139.

[20] Love, R. F., and Morris, J. G., "On Estimating Road Distances by Mathematical Functions", *European Journal of Operational Research* 36 (1988) 251–253.

[21] Love, R. F., Morris, J. G., and Wesolowsky, J., *Facilities Location : Models and Methods*, North - Holland, New York,1988.

[22] Mani, G, "Lowering Variance of Decisions by using Artificial Neural Network Portfolios," *Neural Computation* 3 (1991) 484–486.

[23] Miller, R.G., "The Jackknife–a review", *Biometrika*, 61 (1974), 1–15.

[24] Mittal, A. K., and Palsule, V., "Facilities Location with Ring Radial Distances", *Institute of Industrial Engineers Transactions* 16 (1984) 59–64.

[25] Murtagh, B. A., and Saunders, M. A., "MINOS 5.1 User's Guide", Technical Report No : SOL 83 - 20R, Stanford University, Stanford, California, 1983 (revised 1987).

[26] Perreur, J., and Thisse, J., "Central Metrics and Optimal Location", *Journal of Regional Science* 14 (1974) 411–421.

[27] Perrone, M.P., *Improving Regression Estimation: Averaging Methods for Variance Reduction with Extensions to General Convex Measure Optimization*, PhD Thesis, Department of Physics, Brown University.

[28] Rumelhart, D.E., Hinton, G.E., Williams, R.J., "Learning Internal Representations by Error Propagation," in: D.E. Rumelhart, J.L. McClelland and the PDP Research Group (eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol 1., MIT Press, 1986, 318–362.

[29] Silverman, B.W., *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, 1986.

[30] Stone, C.J., "Consistent Nonparametric Regression," *The Annals of Statistics*, 5 (1977) 595–645.

[31] Specht, D.F., "A General Regression Neural Network", *IEEE Transactions on Neural Networks*, 2 (1991), 568–576.

[32] Ward, J. E., and Wendell, R. E., "A New Norm Measuring Distance Which Yields Linear Location Problems", *Operations Research* 28 (1980) 836–844.

[33] Ward, J. E., and Wendell, R. E, "Using Block Norms for Location Modelling", *Operations Research* 33 (1985) 1074–1091.

[34] Wolpert, D.H. "Stacked Generalization," *Neural Networks* 5 (1992) 241–259.

[35] Zhang, X., Mesirov, J.P., Waltz, D.L., "Hybrid System for Protein Secondary Structure Prediction," *Journal of Molecular Biology*, 225 (1992) 1049–1063.