

Fast local linear regression with anchor regularization

Mathis Petrovich^{1,2,*}, Makoto Yamada^{2,3}

¹ENS Paris-Saclay, ²RIKEN AIP, ³Kyoto University
mathis.petrovich@ens-paris-saclay.fr, makoto.yamada@riken.jp

February 28, 2020

Abstract

Regression is an important task in machine learning and data mining. It has several applications in various domains, including finance, biomedical, and computer vision. Recently, network Lasso, which estimates local models by making clusters using the network information, was proposed and its superior performance was demonstrated. In this study, we propose a simple yet effective local model training algorithm called the fast anchor regularized local linear method (FALL). More specifically, we train a local model for each sample by regularizing it with precomputed anchor models. The key advantage of the proposed algorithm is that we can obtain a closed-form solution with only matrix multiplication; additionally, the proposed algorithm is easily interpretable, fast to compute and parallelizable. Through experiments on synthetic and real-world datasets, we demonstrate that FALL compares favorably in terms of accuracy with the state-of-the-art network Lasso algorithm with significantly smaller training time (two orders of magnitude).

1 Introduction

The regression problem is an important problem in machine learning, data mining, and statistics, and several research works have investigated it in the past decades. Examples include stock price prediction [12, 25], age prediction from RNA-seq [6] or images [11], sentimental analysis [15, 18], or house prediction [7] to name a few.

The most widely used regression approach is based on a linear model including the ordinary least squares (OLS), Ridge regression, least absolute shrinkage and selection operator (Lasso) [19], and elastic net [26]. Because these linear models are extremely simple and can be interpreted by simply checking the linear coefficients of the variables; these approaches are in particular used in practice. However, one of the limitations of linear models is that they cannot handle complex nonlinear data; the performance can be significantly degraded if we apply these linear methods to process complex data such as the gene expression data used heavily in biology and healthcare.

To handle complex data, researchers tend to use kernel methods such as kernel ridge regression (KRR) and support vector regression (SVR) [16]. These approaches are nonlinear models and can handle more complex relationships between inputs and outputs. However, because they are global methods, the output is somewhat smooth, which can be prejudicial for highly complex data with some hard jumps (see Figure 4(a)). Moreover, in general, it is difficult to interpret the features in kernel methods.

Instance-based learning such as the k-nearest neighbor (KNN) algorithm [5] is also a type of widely used nonlinear model. KNN can be considered a local method, which uses the distance or similarity function to predict the output. Despite

*Work done at RIKEN AIP

their good prediction ability, KNN methods can undergo extrapolation when a new input is significantly distant from the training data and can also overfit to much on the data. Moreover, because KNN is a non-parametric method, it is generally difficult to interpret the model.

Recently, the network Lasso, which is a novel type of local method was proposed [8, 9]. Specifically, it estimates the local models by simultaneously clustering the samples and estimating the linear model from them. Accordingly, it combines local regression with network-based sparse regularization (a.k.a., generalized fused regularizer). Because the network Lasso is a locally linear model, it can be easily interpreted to be similar to linear models. More recently, its extension to high-dimensional data with exclusive regularizer has also been proposed [23]. The network Lasso outperforms existing nonlinear models in prediction accuracy, and it was demonstrated that it can automatically cluster models. However, the nearest neighbor graphs should be pre-computed before training the network Lasso. Moreover, the computation of network regularization is expensive and it does not scale to million data points.

In this study, we propose a simple yet efficient local regression model called the fast anchor regularized local linear (FALL) model. Specifically, we estimate a model for each sample, where the model parameter is regularized to be similar to the pre-computed models (called anchor models). The key advantage of the proposed FALL model is that it has a closed-form solution and it is extremely faster than network Lasso. Moreover, the FALL model can cluster samples by choosing the anchor models. Finally, it can extrapolate even when the data is far from the input data owing to the locally linear model. Through experiments, it was found that the proposed algorithm compares favorably with existing state-of-the-art algorithms in accuracy, and it can be 500 times faster for training than network Lasso.

Contribution: The contributions of this study are summarized below.

- We propose a new locally linear regression problem with anchor regularization.
- We derive the closed-form solution of the anchor regularization problem.
- We extensively investigate the proposed method to determine when it worked.

2 Related Work

In this section, we review the linear and nonlinear regression methods.

Linear methods: Linear models are the most widely used regression model. Ridge regression, the least absolute shrinkage and selection operator (Lasso) [20], and the elastic-net [26] rely on the following modeling:

$$\mathbf{y}_i = \mathbf{W}^\top \mathbf{x}_i + \mathbf{e}_i$$

where $\mathbf{W} \in \mathbb{R}^{d \times m}$ is the model parameter and \mathbf{e}_i is the noise vector of the i -th sample. Because the model only depends on \mathbf{W} for all samples, it is called a *global* model. These methods differ on the basis of regularizing the weight parameter \mathbf{W} . The linear model is the first choice for data analysis and it can be scaled efficiently for large dataset. However, it tends to demonstrate poor performance if the data are complex (i.e., expression data in biology).

Nonlinear methods: If the data are complex, nonlinear models can be used. Kernel ridge regression (KRR) can solve the same linear problem as ridge regression, but it maps the original vector \mathbf{x}_i into a reproducing kernel Hilbert space (RKHS) with a mapping ϕ [16]. Specifically, KRR is modeled as follows.

$$\mathbf{y}_i = \mathbf{W}^\top \phi(\mathbf{x}_i) + \mathbf{e}_i = \widetilde{\mathbf{W}} \mathbf{k}(\mathbf{x}_i) + \mathbf{e}_i,$$

where $\widetilde{\mathbf{W}} \in \mathbb{R}^{n \times n}$ is the model parameter, $\mathbf{k}(\mathbf{x}) = (K(\mathbf{x}_1, \mathbf{x}), K(\mathbf{x}_2, \mathbf{x}), \dots, K(\mathbf{x}_n, \mathbf{x}))^\top \in \mathbb{R}^n$, and $K(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \phi(\mathbf{x}')$ is a kernel function such as polynomial or RBF kernel. The kernel methods rely on the choice of kernel

function. KRR is a simple yet effective nonlinear method. However, owing to the nonlinear transform, it is difficult to interpret KRR. Moreover, the KRR output tends to be over smoothed due to the kernel function.

KNN is also a widely used nonlinear regression model that predicts the average of \mathbf{y}_i of the nearest neighbors for the input \mathbf{x} . KNN tends to work for interpolation problems, while it is not suited for extrapolation. Moreover, for KNN, the output performance tends to be affected by the number of nearest neighbors.

Recently, the network Lasso was proposed [8], as another local linear model. The optimization problem of network Lasso is given as

$$\min_{\{\mathbf{W}\}_{i=1}^n} \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{W}_i^\top \mathbf{x}_i\|_2^2 + \lambda \sum_{(j,k) \in E} r_{j,k} \|\mathbf{W}_j - \mathbf{W}_k\|_F,$$

where $r_{j,k} \geq 0$ is the pre-defined graph information, $\|\cdot\|_2$ is the ℓ_2 norm, $\|\cdot\|_F$ is the Frobenius norm, and $\lambda \geq 0$ is the regularization parameter. In network Lasso, owing to the fused Lasso type regularizer, as λ increases, the models tend to form a consensus. Thus, we can automatically cluster the model parameters. However, the optimization of network Lasso is expensive and it does not scale well to large datasets.

3 Problem Formulation

In this section, we formulate the locally linear regression problem.

Let us denote an input vector by $\mathbf{x} = (x^{(1)}, \dots, x^{(d)})^\top \in \mathbb{R}^d$ and the corresponding output value by $\mathbf{y} \in \mathbb{R}^m$. The set of training samples $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ has been drawn i.i.d. from a joint probability density $p(\mathbf{x}, \mathbf{y})$. In this study, we consider a local linear model

$$\mathbf{y}_i = \mathbf{W}_i^\top \mathbf{x}_i + \mathbf{e}_i,$$

where $\mathbf{W}_i \in \mathbb{R}^{d \times m}$ is the i -th sample regression parameter and $\mathbf{e}_i \in \mathbb{R}^m$ is the noise vector of the i -th sample. Note that each \mathbf{y}_i is modeled by its model parameter \mathbf{W}_i , while a *global* model employs \mathbf{W} to model all \mathbf{y}_i .

The study aims to estimate \mathbf{W}_i from the training input-output samples efficiently. This is intrinsically a very difficult problem because we need to estimate \mathbf{W}_i from one sample. That is, it can easily be overfitted to the training sample. Thus, a regularizer that can efficiently generalize to an unknown sample is the key to this problem.

4 Proposed Method

In this section, we propose the FALL method.

4.1 Fast anchor regularized local linear method

For a training sample pair $(\mathbf{x}_i, \mathbf{y}_i)$, we consider the following regression problem:

$$\min_{\mathbf{W}_i, \mathbf{p}_i} J_i(\mathbf{W}_i, \mathbf{p}_i), \quad \text{s.t. } \mathbf{p}_i \in S_k, \quad (1)$$

where

$$J_i(\mathbf{W}_i, \mathbf{p}_i) = \|\mathbf{y}_i - \mathbf{W}_i^\top \mathbf{x}_i\|_2^2 + \lambda \sum_{l=1}^k p_{il} \|\mathbf{W}_i - \mathbf{A}_l\|_F^2$$




Figure 1: Representation of models in the parameter space. The area of the sectors is represented by the vector \mathbf{p}_i , which represents the contribution of each \mathbf{A}_l in the regularization. \mathbf{p}_i represents how \mathbf{W}_i is shared among all anchor models. The optimization of our model attempts to provide \mathbf{W}_i a good power of prediction (to get \mathbf{y}_i) but \mathbf{W}_i should not be too far from the anchors because the regularization penalizes the distance of \mathbf{W}_i to the \mathbf{A}_l . Simultaneously, the optimization must choose \mathbf{p}_i , which determines the importance of \mathbf{W}_i being close to each \mathbf{A}_l because when p_{il} is large, the cost of being further from \mathbf{A}_l is large. The arrows represent the sharing between the model and the anchors.

and

$$S_k = \{\mathbf{p} \in \mathbb{R}_+^k \text{ s.t. } \mathbf{p}^\top \mathbf{1}_k = 1\}$$

is the simplex set, $\mathbf{1}_k \in \mathbb{R}^k$ is the vector whose elements are all one, $\mathbf{A}_l \in \mathbb{R}^{d \times m}$ is the pre-defined model parameter (we call this anchor model), $k \ll n$ denotes the number of anchors, and $\lambda > 0$ is the regularization parameter. \mathbf{p}_i is the parameter that represents the contribution of each anchor model in the regularization.

Note that a bias term can be used for every model \mathbf{W}_i . To accomplish this, the bias should be used for computing the anchor models $\mathbf{A}_1, \dots, \mathbf{A}_k$. Then, each data point $\mathbf{x}_i \in \mathbb{R}^d$ must be replaced by $(\mathbf{x}_i^\top, 1)^\top \in \mathbb{R}^{d+1}$.

The anchor parameters play the role of "reference" models. This can be a priori models or can be constructed with local models (see section 4.5). They are not optimized in this regularizer; they are known beforehand. The idea of this regression problem is to jointly optimize

- The data attachment term (the square error term) that attempts to efficiently model \mathbf{y}_i from \mathbf{x}_i .
- The anchor attachment term (the convex sum of squared differences), which can be interpreted as "sharing" (see Figure 1), that attempts to force the model \mathbf{W}_i to be as similar to the anchor model \mathbf{A}_l .

As the first term does not depend on \mathbf{p}_i , its minimization can be put in the second term; thus, the whole problem can be observed as the one that finds the best model \mathbf{W}_i , which minimizes the data term and the *best* convex sum on the right. In this sense, the second term with the minimization over \mathbf{p}_i is the regularizer of our problem. The best \mathbf{p}_i can be interpreted as weights, which force the model \mathbf{W}_i to be close to the "good" anchor models, where their predictions, given \mathbf{x}_i are close to \mathbf{y}_i .

This regularizer is a relaxation of the optimal transport (OT) [14] setting because the sum of all values of i yields the same formula. The main difference is that we decided to not constrain each anchor model to share its information with the models by a fixed amount (as the second constraint in the OT). Indeed, in some data structures, two anchor models can have different impacts, and one can have a larger influence than the other (see Figure 4(c) where more points are better represented by the blue model) and thus, it is not profitable to restrict an anchor model to share its power over all models \mathbf{W}_i . That is, \mathbf{p}_i can be used to softly cluster \mathbf{x}_i with respect to the anchor models.

A standard approach to solve the optimization problem Eq. (1) is to alternately minimize \mathbf{W}_i and \mathbf{p}_i . However, it can be trapped by a poor local optimal solution. In this paper, we derive the analytical solution of the optimization problem Eq. (1).

4.2 Minimization of weights given the contributions of anchor models

We first derive the optimal solution of \mathbf{W}_i by fixing \mathbf{p}_i . Because the function to optimize $\mathbf{W}_i \rightarrow J_i(\mathbf{W}_i, \mathbf{p}_i)$ is convex, we can cancel the derivative with respect to \mathbf{W}_i to determine the optimal \mathbf{W}_i :

$$\frac{\partial J}{\partial \mathbf{W}_i}(\mathbf{W}_i, \mathbf{p}_i) = 2\mathbf{x}_i(\mathbf{x}_i^\top \mathbf{W}_i - \mathbf{y}_i^\top) + 2\lambda \sum_{l=1}^k p_{il}(\mathbf{W}_i - \mathbf{A}_l) = 0.$$

Thus, the optimal solution is given by

$$\widehat{\mathbf{W}}_i = (\mathbf{x}_i \mathbf{x}_i^\top + \lambda \mathbf{I}_d)^{-1} (\mathbf{x}_i \mathbf{y}_i^\top + \lambda \mathbf{G}_i),$$

where $\mathbf{G}_i = \sum_{l=1}^k p_{il} \mathbf{A}_l \in \mathbb{R}^{n \times m}$ is the mixture of anchor models. We can obtain the optimal solution of \mathbf{W}_i analytically. However, because we need to compute it for all samples, it is computationally expensive if the number of features is extremely large (e.g., $d = 10^5$).

Owing to the effective structure of the inverse matrix (i.e., rank one matrix plus an identity matrix), we can use the Sherman-Morrison formula [17, 2] (or simple algebra) to obtain the analytical solution without inverting the matrix.

$$(\mathbf{x}_i \mathbf{x}_i^\top + \lambda \mathbf{I}_d)^{-1} = \frac{1}{\lambda} \mathbf{I} - \frac{1}{\lambda(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)} \mathbf{x}_i \mathbf{x}_i^\top.$$

Finally, we have

$$\widehat{\mathbf{W}}_i = \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top + \mathbf{G}_i.$$

Because the final solution does not include the inverse matrix, it can be estimated easily and quickly.

4.3 Minimization of the global loss

In this section, we minimize $\mathbf{p}_i \in S_k \rightarrow J_i(\widehat{\mathbf{W}}_i, \mathbf{p}_i)$ by taking the best $\widehat{\mathbf{W}}_i$ computed in the last section. More specifically, we derive the global optimal solution of \mathbf{p}_i .

Thus, we plug the optimal $\widehat{\mathbf{W}}_i$ into Eq. (1), and consider the following optimization problem:

$$\min_{\mathbf{p}_i} J_i(\widehat{\mathbf{W}}_i, \mathbf{p}_i), \quad \text{s.t. } \mathbf{p}_i \in S_k. \quad (2)$$

Lemma 1. *Eq. (2) is equivalent to the following quadratic problem:*

$$\min_{\mathbf{p}_i} \mathbf{p}_i^\top \mathbf{H}_i \mathbf{p}_i + \mathbf{b}_i^\top \mathbf{p}_i + \beta_i \mathbf{y}_i^\top \mathbf{y}_i, \quad \text{s.t. } \mathbf{p}_i \in S_k,$$

where

- $\beta_i = \lambda / (\lambda + \|\mathbf{x}_i\|_2^2)$.

- $\mathbf{H}_i \in \mathbb{R}^{k \times k}$ such that $(\mathbf{H}_i)_{ll'} = \beta_i \mathbf{x}_i^\top \mathbf{A}_l \mathbf{A}_{l'}^\top \mathbf{x}_i - \lambda \text{Tr}(\mathbf{A}_l^\top \mathbf{A}_{l'})$. $\text{Tr}(\cdot)$ is the trace operator.
- $\mathbf{b}_i \in \mathbb{R}^k$ such that $(\mathbf{b}_i)_l = -2\beta_i \mathbf{y}_i^\top \mathbf{A}_l^\top \mathbf{x}_i + \lambda \|\mathbf{A}_l\|_F^2$.

The proof can be found in the appendix. It is based on the rewriting of $J_i(\widehat{\mathbf{W}}_i, \mathbf{p}_i)$ with some factorization.

Lemma 2. \mathbf{H}_i is symmetric negative definite, and Eq. (2) is a concave QP under linear constraints.

(Proof of Lemma 2)

First, by definition, \mathbf{H}_i is symmetric. Let $\mathbf{z} \in \mathbb{R}^k$, and let us show that $\mathbf{z}^\top \mathbf{H}_i \mathbf{z} < 0$.

$$\begin{aligned}
\mathbf{z}^\top \mathbf{H}_i \mathbf{z} &= \beta_i \sum_{l=1}^k \sum_{l'=1}^k z_{il} z_{il'} (\mathbf{x}_i^\top \mathbf{A}_l \mathbf{A}_{l'}^\top \mathbf{x}_i) - \lambda \sum_{l=1}^k \sum_{l'=1}^k z_{il} z_{il'} \text{Tr}(\mathbf{A}_l^\top \mathbf{A}_{l'}) \\
&= \frac{\lambda}{(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)} \left\| \sum_{l=1}^k z_{il} \mathbf{A}_l^\top \mathbf{x}_i \right\|_2^2 - \lambda \left\| \sum_{l=1}^k z_{il} \mathbf{A}_l \right\|_F^2 \\
&\leq \lambda \left\| \sum_{l=1}^k z_{il} \mathbf{A}_l \right\|_F^2 \left(\frac{\mathbf{x}_i^\top \mathbf{x}_i}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} - 1 \right) \\
&< 0
\end{aligned}$$

The first inequality is valid owing to the submultiplicativity of the Frobenius norm. Thus, our problem is a strictly concave problem. \square

Theorem 3. A minimizer of our regression problem Eq. (1) can be expressed as

$$\widehat{\mathbf{W}}_i = \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i]^\top + \mathbf{A}_{l_i} \quad (3)$$

where $l_i = \text{argmin}_l \|\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i\|_2^2$

(Proof of Theorem 3) Owing to the two lemmas (1 and 2), we know that the function $\mathbf{p}_i \in S_k \rightarrow J(\widehat{\mathbf{W}}_i, \mathbf{p}_i)$ is concave. This function is also a continuous function defined on the set S_k , which is convex and compact. By the Bauer minimum principle, this function attains its maximum at some extreme point of that set.

The boundary of the simplex S_k is the set of points \mathbf{p}_i such that there is one coordinate with a 1 and all others are set to 0. This implies a hot vector that indicates an anchor model.

Plugging the solution in J and simplifying it, we obtain

$$\begin{aligned}
J(\delta_l, \widehat{\mathbf{W}}_i) &= \beta_i \left(\|\mathbf{A}_l^\top \mathbf{x}_i\|_2^2 - 2\langle \mathbf{y}_i, \mathbf{A}_l^\top \mathbf{x}_i \rangle_2 + \mathbf{y}_i^\top \mathbf{y}_i \right) \\
&= \beta_i \|\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i\|_2^2,
\end{aligned}$$

which is the problem of finding the best anchor model for the data point \mathbf{x}_i (see Figure 2). This can be found by testing all possibilities, i.e.,

$$p_{il}^* = \begin{cases} 1 & \text{if } l = \text{argmin}_l \|\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i\|_2^2 \\ 0 & \text{otherwise.} \end{cases}$$




Figure 2: In the optimal solution, \mathbf{W}_i is not shared but it chooses a unique anchor model \mathbf{A}_l , which effectively fits \mathbf{x}_i to \mathbf{y}_i (i.e., minimizes $\|\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i\|_2^2$). The best anchor model could be \mathbf{A}_2 in blue so that the optimal \mathbf{p}_i is the hot vector with 1 in position $l = 2$.

Algorithm 1 Local linear regression with anchor regularization.

Input: $(\mathbf{x}_1, \dots, \mathbf{x}_n), (\mathbf{y}_1, \dots, \mathbf{y}_n)$
 $\mathbf{A}_1, \dots, \mathbf{A}_k \leftarrow$ Finding anchors by Algorithm 3
for $i \leftarrow 1, n$ **do**
 $l_i \leftarrow \operatorname{argmin}_l \|\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i\|$
 $\mathbf{p}_i \leftarrow \mathbf{e}_{l_i}$
 $\mathbf{W}_i \leftarrow \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i]^\top + \mathbf{A}_{l_i}$
end for
Output: $(\mathbf{W}_1, \dots, \mathbf{W}_n), (\mathbf{p}_1, \dots, \mathbf{p}_k)$

By substituting this into $\widehat{\mathbf{W}}_i$, \mathbf{G}_i^* becomes \mathbf{A}_{l_i} which shows that it is not a mixture but an anchor model (the one that predicts the closest to \mathbf{y}_i given \mathbf{x}_i). \square

Thanks to Theorem 3, the solution can be computed easily and quickly. The optimal solution might not be unique if there exists \mathbf{A}_{l_1} and \mathbf{A}_{l_2} such that $\mathbf{A}_{l_1}^\top \mathbf{x}_i = \mathbf{A}_{l_2}^\top \mathbf{x}_i$. In this situation, either can be chosen.

It is interesting that the global minimizer for the assignment \mathbf{p}_i does not depend on λ but only depends on the power of prediction of each anchor model to predict something close to \mathbf{y}_i given the data point \mathbf{x}_i .

Given the anchor models, Algorithm 1 can be efficiently computed even if the dimension of d is large. Note that the entire Algorithm 1 is embarrassingly parallel and therefore can be vectorized to speed up the process. Moreover, given the simple matrix form, the proposed method can be implemented on a GPU.

4.4 Prediction

In this section, we will focus on the prediction power of each model $\widehat{\mathbf{W}}_i$ of our system. We will use the training data at first, then with new samples and finally, we will see how to combine all the models and how to interpolate between them.

Prediction with the training data: The prediction $\hat{\mathbf{y}}_i = \widehat{\mathbf{W}}_i^\top \mathbf{x}_i$ can be decomposed as

$$\widehat{\mathbf{W}}_i^\top \mathbf{x}_i = \underbrace{\frac{\mathbf{x}_i^\top \mathbf{x}_i}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i}}_{1-\beta_i} \mathbf{y}_i + \underbrace{\frac{\lambda}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i}}_{\beta_i} \mathbf{A}_{l_i}^\top \mathbf{x}_i, \quad (4)$$

where $0 < \beta_i < 1$ controls the contribution of the anchors to the prediction.

That is, the prediction $\widehat{\mathbf{W}}_i^\top \mathbf{x}_i$ is a convex sum of the ground truth \mathbf{y}_i and the prediction of the best anchor model $\mathbf{A}_{l_i}^\top \mathbf{x}_i$.

The λ parameter controls the regularization as follows.

- When λ is small ($\lambda \ll \mathbf{x}_i^\top \mathbf{x}_i$), the coefficient $\beta_i \simeq 0$, which yields a prediction close to the ground truth. This corresponds to an unregularized local linear model.
- When λ is large ($\lambda \gg \mathbf{x}_i^\top \mathbf{x}_i$), the coefficient $\beta_i \simeq 1$, which implies that the prediction is close to the prediction of the best anchor model and not much emphasis is given on its value being close to \mathbf{y}_i .

This provides a way of interpreting the role of λ . The training error made by our model always satisfies

$$\epsilon_i = \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|_2 = \beta_i \|\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i\|_2,$$

which is bounded by the minimal error of the anchor models. The optimal model parameter can be decomposed as

$$\widehat{\mathbf{W}}_i = \mathbf{C}_i + \mathbf{A}_{l_i}$$

where

$$\mathbf{C}_i = \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i]^\top$$

can be interpreted as a "correction" model because $\mathbf{C}_i^\top \mathbf{x}_i = (1 - \beta_i)(\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i)$. This correction term is added to the prediction of the anchor model to traverse from the full regularized prediction $\mathbf{A}_{l_i}^\top \mathbf{x}_i$ (when $1 - \beta_i = 0$ where $\lambda \rightarrow \infty$) to the non regularized prediction \mathbf{y}_i (when $\beta_i = 0$ where $\lambda \rightarrow 0$). The correction model is the model that adds to the anchor model to adjust the prediction to get closer to the ground truth.

Prediction with new samples: Let us compute the prediction of model $\widehat{\mathbf{W}}_i$ with a new input \mathbf{x} :

$$\widehat{\mathbf{W}}_i^\top \mathbf{x} = \underbrace{\frac{\mathbf{x}_i^\top \mathbf{x}}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i}}_{\mathbf{C}_i^\top \mathbf{x}} [\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i] + \mathbf{A}_{l_i}^\top \mathbf{x}.$$

We can notice interesting properties in this prediction.

- When \mathbf{x} is very different from \mathbf{x}_i , i.e., when $|\mathbf{x}_i^\top \mathbf{x}|$ is very small, the correction term $\mathbf{C}_i^\top \mathbf{x}$ is very small; thus, this model does not take any risks and behaves like the anchor model.
- When \mathbf{x} and \mathbf{x}_i are highly correlated (i.e., $|\mathbf{x}_i^\top \mathbf{x}|$ is large), the correction term has more importance. The term is proportional to $\mathbf{x}_i^\top \mathbf{x}$ to adjust the correction in the good direction (as it can be negative).
- λ is used to adjust the regularization: when λ is large, the prediction of this model will always be close to the prediction of the best anchor model.
- If $\mathbf{x}_i^\top \mathbf{x}_i = 0$, i.e., $\mathbf{x}_i = 0$, it is not possible to obtain a $\mathbf{y}_i \neq 0$ with a linear model. In this case, the prediction will be based on the anchor model.

Combining the models: After fitting all the models \mathbf{W}_i , we determine how to use them to predict the final prediction $\hat{\mathbf{y}}_i$ with a new vector \mathbf{x} . More specifically, we do not know the corresponding parameter of the new input \mathbf{x} . Thus, we propose a simple yet effective prediction approach.

Algorithm 2 Prediction with a new data point.

Input: Data \mathbf{x}
 $\mathbf{x}_{h_1}, \dots, \mathbf{x}_{h_K} \leftarrow$ The K nearest neighbors of \mathbf{x} in the training set
 $\mathbf{W}_{\text{pred}} \leftarrow 0$
 $S \leftarrow 0$
for $l \leftarrow 1, n$ **do**
 $c_l \leftarrow 1 / \|\mathbf{y} - \mathbf{x}\|_2$
 $S \leftarrow S + c_l$
 $\mathbf{W}_{\text{pred}} \leftarrow \mathbf{W}_{\text{pred}} + c_l \mathbf{W}_{h_l}$
end for
 $\mathbf{W}_{\text{pred}} \leftarrow \mathbf{W}_{\text{pred}} / S$
 $\mathbf{y}_{\text{pred}} \leftarrow \mathbf{W}_{\text{pred}}^\top \mathbf{x}$
Output: \mathbf{y}_{pred}

Our assumption is that locally, the data behave like linear models. Accordingly, let us compute the $K_p \geq 1$ nearest neighbors $\mathbf{x}_{h_1}, \dots, \mathbf{x}_{h_{K_p}}$ for \mathbf{x} and determine the corresponding models $\mathbf{W}_{h_1}, \dots, \mathbf{W}_{h_{K_p}}$.

Each model \mathbf{W}_{h_j} is locally linear and can be used to predict $\mathbf{W}_{h_j}^\top \mathbf{x}$. The models created with a vector close to \mathbf{x} should yield better results. Considering this, a weighted average of the local predictions can be predicted by

$$\mathbf{y}_{\text{pred}} = \sum_{j=1}^{K_p} \alpha_{h_j} \mathbf{W}_{h_j}^\top \mathbf{x} = \underbrace{\left[\sum_{j=1}^{K_p} \alpha_{h_j} \mathbf{W}_{h_j} \right]^\top}_{\mathbf{W}_{\text{pred}}^\top} \mathbf{x}, \quad (5)$$

where the weights $\alpha_{h_1}, \dots, \alpha_{h_{K_p}}$ are inversely proportional to the distance from \mathbf{x} to \mathbf{x}_{h_j} :

$$\alpha_{h_j} = \frac{\frac{1}{d(\mathbf{x}, \mathbf{x}_{h_j})}}{\sum_{j=1}^{K_p} \frac{1}{d(\mathbf{x}, \mathbf{x}_{h_j})}},$$

with the distance $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{y} - \mathbf{x}\|_2$.

For each \mathbf{x} , our algorithm (in Algorithm 2) predicts the local average model \mathbf{W}_{pred} . Specifically, Algorithm 2 first computes the K nearest neighbors of \mathbf{x} in the training set. Then, it computes \mathbf{W}_{pred} , which is the weighted average of the models of the neighbors. At the end, it returns the prediction of \mathbf{y}_{pred} , which is the output of the model \mathbf{W}_{pred} on the input data \mathbf{x} . We obtained the prediction by applying it to the input data \mathbf{x} .

Our method can be considered similar to the KNN approach because it is also an instance-based learning approach. However, instead of interpolating the values here in the \mathbb{R}^m space, we interpolate the regularized models in the $\mathbb{R}^{d \times m}$ space. At this step, our model uses more input data than the KNN because \mathbf{x} is again used here to compute the value of \mathbf{y} . Thus, if a new data point is significantly far from the input data, instead of predicting a value close to an already seen value, our model will attempt to linearly extend the prediction.

The FALL method can be easily extended to new samples. For a new training sample i_{new} , only $\mathbf{W}_{i_{\text{new}}}$ needs to be computed. Thus, the training can be evolutive and there would be no need to compute any model again (in comparison to network Lasso for example).

Algorithm 3 Finding the anchor models.

Input: Data $(\mathbf{x}_1, \dots, \mathbf{x}_n)$, ground truth $(\mathbf{y}_1, \dots, \mathbf{y}_n)$
 $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k} \leftarrow$ Find anchor points (randomly or by K-means for example)
for $l \leftarrow 1, k$ **do**
 $S_l \leftarrow$ The neighbors set of \mathbf{x}_{i_l} in $(\mathbf{x}_1, \dots, \mathbf{x}_n)$
 $\mathbf{A}_l \leftarrow$ Model computed with $\{\mathbf{x}_j\}_{j \in S_l}$ and $\{\mathbf{y}_j\}_{j \in S_l}$
end for
Output: $(\mathbf{A}_1, \dots, \mathbf{A}_k)$

Considering the final prediction for a new sample, we can also decompose the model predicted by \mathbf{W}_{pred} using

$$\begin{aligned} \mathbf{W}_{\text{pred}} &= \sum_{j=1}^{K_p} \alpha_{h_j} (\mathbf{C}_{h_j} + \mathbf{A}_{l_{h_j}}) \\ &= \underbrace{\left[\sum_{j=1}^{K_p} \alpha_{h_j} \mathbf{C}_{h_j} \right]}_{\mathbf{C}_{\text{pred}}} + \underbrace{\left[\sum_{j=1}^{K_p} \alpha_{h_j} \mathbf{A}_{l_{h_j}} \right]}_{\mathbf{R}_{\text{pred}}}, \end{aligned}$$

where \mathbf{R}_{pred} is the weighted mixture of the best anchor models of the neighborhood of \mathbf{x} and \mathbf{C}_{pred} is the weighted average of the correcting models. This correction term attempts to correct the fully regularized model by using the labels \mathbf{y}_i to obtain a solution with a smaller training error. This term vanishes when $\lambda \rightarrow \infty$.

4.5 Anchor model

The selection of anchors can be the subject of further research; however, we provide two simple methods for selecting them here.

First, by randomly taking a set of indexes $I = \{i_1, \dots, i_k\}$, we get our set of anchor points $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}$. Second, we can use a clustering algorithm like K-means to get $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}$. These vectors can be outside the set of input vectors.

These anchor points work as reference models in the training data. For each point, we select a set of neighbors (with K nearest or some graph information) and compute a linear model (with Lasso [20] or Ridge regression for example) (see Algorithm 3). Note that the for loop in Algorithm 3 can be parallelized to speed up the process in Algorithm 3. The anchor points can be initialized using any method. We propose here a randomized method or a method with initialization using K-means clustering (and kmeans++).

5 Experiments

In this section, we first illustrate the proposed algorithm using synthetic datasets and then evaluate it using real-world datasets.




Figure 3: Two moons dataset. The boundary is established when the model predicts the same probability for the two classes. It can be observed that this boundary is linear by part and effectively separates the samples.

5.1 Illustrative experiments

Step function: To present our method and demonstrate its working, let us consider a simple non linear model, i.e., step function, which only takes two values. Let us sample some random points with some random noise (see in Figure 4(a)).

Let us take $k = 2$ anchors points (we could take more, this is for illustration), one inside and one outside. There are parts of the data where the linear models are not similar. Let us compute their nearest neighbors and anchor models for each neighborhood (as shown in Figure 4(b)). The models are determined by applying ridge regression with bias, with a high regularization parameter $\alpha = 1000$.

For each data point \mathbf{x}_i , we search for the anchor model \mathbf{A}_{l_i} which has the minimum loss $\|\mathbf{y}_i - \mathbf{A}_{l_i}^\top \mathbf{x}_i\|$. This is the best linear anchor model that explains the relation between \mathbf{x}_i and \mathbf{y}_i (and corresponds to selecting the optimal $\mathbf{p}_i = \delta_{l_i}$). This provides a natural method to cluster our data into k clusters (see Figure 4(c)).

Note that the proposed FALL method can cluster data points with respect to anchor models (clustering by regression), while it is not possible to directly cluster points using KNN and other linear models. Thus, the proposed FALL model has a similar paradigm as network Lasso.

For each \mathbf{x}_i , after the clustering step (finding the best anchor model \mathbf{A}_l), we compute the final model $\widehat{\mathbf{W}}_i$ by

$$\widehat{\mathbf{W}}_i = \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{A}_l^\top \mathbf{x}_i]^\top + \mathbf{A}_l.$$

As evident in Figure 4(d), larger values of λ imply that the final model is farther from the training data but closer to the anchor models.

Two moons dataset: The proposed method can also be applied to the classification datasets. If the model is trained with one-hot encoding vectors \mathbf{y} (i.e. when $\mathbf{y} = \delta_j$), our prediction can be interpreted as probability vectors. To illustrate this, we classify the two moons dataset by using Eq. (5).

In Figure 3, boundaries can also be observed (the contour when the probability to be in each class is equal). The proposed method can separate the two moons very well in this example.

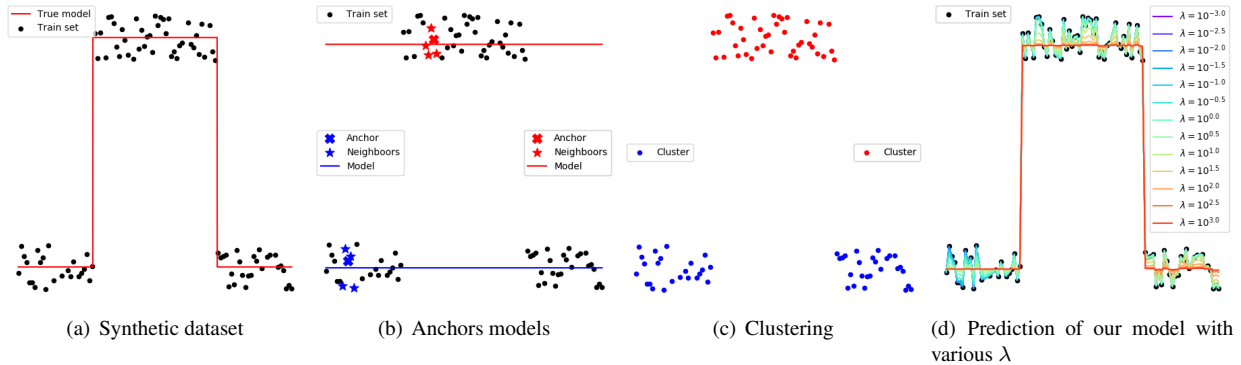


Figure 4: Generation of synthetic dataset in Figure 4(a) (a step function), which is non linear. The red plot is the real function and the black dots are the training set, which is generated with some added noise. In Figure 4(b), we can see the anchor points with their relative neighborhood and computed anchor models. They are computed with the linear Ridge regression. In Figure 4(c), the data points are clustered into k clusters by choosing the best anchor model to model the linear relationship between a sample x_i and its label y_i (the best A_l that minimizes $\|y_i - A_l^T x_i\|$). The points on the right have the same behavior locally than those on the left. Finally, Figure 4(d) shows our computed final model depending on the λ regularization parameter. The continuous path between the "perfect" solution, where each model predicts the correct label, and the most regularized solution, where locally it is the same as the model of the best anchor at this position, can be observed.

Table 1: Averaged mean squared error on the real-world datasets over 50 runs. For each method, we select the best parameters by using 3-fold cross-validation. N/A means that a memory problem occurred when computing the model on the dataset. The asterisk * means that the results were computed without cross-validation but with fixed parameters (because of a too long runtime).

	FALL	Network Lasso	KNN	KRR	Ridge	Lasso
Concrete data	42.16 (\pm 19.38)	49.85 (\pm 22.12)	58.14 (\pm 26.85)	110.80 (\pm 33.57)	111.05 (\pm 33.92)	111.10 (\pm 33.88)
Superconduct data	129.66 (\pm 28.33)	N/A	117.48 (\pm 18.77)	312.11 (\pm 25.82)	308.52 (\pm 22.56)	345.25 (\pm 99.80)
Aquatic toxicity data	1.46 (\pm 0.74)	1.72 (\pm 1.01)	1.81 (\pm 0.74)	1.92 (\pm 0.75)	1.58 (\pm 0.72)	1.58 (\pm 0.71)
Fish toxicity data	0.92 (\pm 0.48)	0.98 (\pm 0.53)	0.91 (\pm 0.48)	1.14 (\pm 0.57)	0.99 (\pm 0.44)	0.99 (\pm 0.44)
P.T. motor UPDRS data	20.76 (\pm 4.51)	24.63 (\pm 6.89)*	20.88 (\pm 4.29)	58.30 (\pm 5.50)	56.69 (\pm 5.26)	56.81 (\pm 5.20)

5.2 Real-world dataset

Setup: We compared our method on several datasets with other methods. We computed the average mean squared error with 50 runs and the training and test time. We compared the proposed FALL model with KNN, kernel ridge regression (KRR), ridge and Lasso [21], and network Lasso [8]. It should be noted that network Lasso is a state-of-the-art locally linear model. To compare the methods fairly, we selected the best hyperparameters by cross-validation with optuna [1] with a timeout of 60 s.

For each dataset, the following procedure was performed.

1. Randomly split the dataset into training and testing data with a proportion of 0.9/0.1.
2. Tune the hyperparameters using the 3-cross-validation score on the training data (60 s maximum for tuning).
3. Use these parameters to fit the entire training set and time it.
4. Test the model on the test set and time it.
5. Repeat (1) 50 times and average the results.

All splits are the same for each method.

Table 2: Training time and testing time (in seconds) on real-world datasets over 50 runs. For FALL, the training time is the sum of the creation of the anchor models and the final model. N/A means that a memory problem occurred when computing the model on the dataset. The asterisk * means that the training time and testing time were averaged with few samples (because of a too long runtime).

		FALL	Network Lasso	KNN	KRR	Ridge	Lasso
Concrete data	Train	0.06 (\pm 0.05)	16.58 (\pm 3.41)	0.00 (\pm 0.00)	0.09 (\pm 0.08)	0.00 (\pm 0.00)	0.01 (\pm 0.00)
	Test	0.01 (\pm 0.00)	0.19 (\pm 0.03)	0.00 (\pm 0.00)	0.01 (\pm 0.02)	0.00 (\pm 0.00)	0.00 (\pm 0.00)
Superconduct data	Train	1.24 (\pm 1.33)	N/A	0.11 (\pm 0.02)	124.61 (\pm 109.20)	0.02 (\pm 0.01)	121.04 (\pm 125.75)
	Test	5.90 (\pm 0.71)	N/A	0.22 (\pm 0.08)	0.23 (\pm 0.13)	0.00 (\pm 0.01)	0.01 (\pm 0.01)
Aquatic toxicity data	Train	0.02 (\pm 0.05)	5.31 (\pm 1.00)	0.00 (\pm 0.00)	0.01 (\pm 0.01)	0.00 (\pm 0.00)	0.00 (\pm 0.00)
	Test	0.00 (\pm 0.00)	0.07 (\pm 0.01)	0.00 (\pm 0.00)	0.00 (\pm 0.01)	0.00 (\pm 0.00)	0.00 (\pm 0.00)
Fish toxicity data	Train	0.05 (\pm 0.05)	13.23 (\pm 1.78)	0.00 (\pm 0.00)	0.03 (\pm 0.04)	0.00 (\pm 0.00)	0.00 (\pm 0.00)
	Test	0.01 (\pm 0.00)	0.14 (\pm 0.02)	0.00 (\pm 0.00)	0.00 (\pm 0.01)	0.00 (\pm 0.00)	0.00 (\pm 0.00)
P.T. motor UPDRS data	Train	0.25 (\pm 0.99)	\sim 30 min*	0.00 (\pm 0.00)	1.39 (\pm 0.51)	0.00 (\pm 0.00)	0.01 (\pm 0.01)
	Test	0.27 (\pm 0.02)	\sim 4 seconds*	0.01 (\pm 0.00)	0.02 (\pm 0.01)	0.00 (\pm 0.01)	0.01 (\pm 0.01)

Concrete Compressive Strength Dataset: In civil engineering, concrete is a widely used material. For high-performance concrete, the compressive strength is a highly nonlinear function of age and ingredients. This information is useful to understand the importance of each ingredient to create more resistant concrete. In this dataset [24], the number of attributes (dimension) is $d = 9$ and the number of samples is $n = 1030$.

Superconductivity Dataset: Health care professionals or researchers in the Hadron Collider at CERN are interested in superconducting materials. These materials are interesting because they are conductive without any resistance. To use this property, superconducting materials must be used in their critical temperature, which can be difficult to find. This dataset [10] contains $d = 81$ interesting features (which are related to the chemical formula) and the goal is to predict the critical temperature. There are $n = 21263$ samples in this dataset. We could not use network Lasso due to memory error (too large matrix to perform factorization).

QSAR Toxicity Datasets: The water ecosystem can be damaged by some chemicals. These two datasets QSAR aquatic toxicity [4] and fish toxicity [3] were created to understand what combinations of chemicals that can be highly toxic for certain organisms. The value of toxicity to be predicted is the concentration of LC50 in *Daphnia magna* (a small planktonic crustacean). This is a good indicator of toxicity because it kills this creature. The difference between the two datasets is that the aquatic toxicity dataset contains $d = 8$ features (which corresponds to eight molecular descriptors), $n = 546$ instances and the output is the concentration of LC50 after 48h; the fish toxicity dataset contains $d = 6$ features, $n = 908$ instances and the output is the concentration of LC50 after 96h.

Parkinsons Telemonitoring Data Set: The symptoms of Parkinson’s disease need to be monitored by a doctor in a clinic, which is costly and time-consuming. They can use the unified Parkinson’s disease rating scale (UPDRS) to track the progression of the symptoms. This dataset [22] contains information about the patients (as age or gender) and 16 biomedical voice measures, which are taken in the patient home with a telemonitoring device, to form $d = 19$ features. There are $n = 5875$ instances and one of the interesting output is the motor UPDRS. This target value can help people with early-stage Parkinson’s disease to follow up more conveniently and less expensively.

Result: Through experiments, the proposed FALL method was found to compare favorably with network Lasso and other existing baselines. In particular, the concrete data, the FALL method significantly outperforms baseline methods. Overall, because the FALL method is a locally linear model, the performance of the FALL method is similar to the one with KNN. The important point here is that the computation time of the FALL method is two orders of magnitude smaller than that of network Lasso.

6 Conclusion

Local linear methods are nonlinear methods that can yield good results on very complex data. In this study, we propose a new locally linear model called the fast anchor regularized local linear method (FALL), which produces a formula that can be used to efficiently compute a regularized locally linear model. The FALL method also enables us to understand the mechanism of regularization and interpret it through its inner clustering. Through experiments, we demonstrated that the proposed method compares favorably with network Lasso, which is a state-of-the-art locally linear model. Moreover, the computation time of FALL model is two orders of magnitude smaller than that of network Lasso; additionally, this model can be extended with new training data with a small computational cost.

References

- [1] Takuya Akiba et al. “Optuna: A Next-generation Hyperparameter Optimization Framework”. In: *KDD*. 2019.
- [2] M. S. Bartlett. “An Inverse Matrix Adjustment Arising in Discriminant Analysis”. In: *Ann. Math. Statist.* 22.1 (Mar. 1951), pp. 107–111.
- [3] M. Cassotti et al. “A similarity-based QSAR model for predicting acute toxicity towards the fathead minnow (*Pimephales promelas*)”. In: *SAR and QSAR in Environmental Research* 26.3 (2015), pp. 217–243.
- [4] Matteo Cassotti et al. “Prediction of Acute Aquatic Toxicity Toward *Daphnia magna* by using the GA-kNN Method”. In: *Alternatives to laboratory animals : ATLA* 42 (Mar. 2014), pp. 31–41.
- [5] T. Cover and P. Hart. “Nearest neighbor pattern classification”. In: *IEEE Transactions on Information Theory* 13.1 (Jan. 1967), pp. 21–27.
- [6] Jason G. Fleischer et al. “Predicting age from the transcriptome of human dermal fibroblasts”. In: *Genome Biology* 19.1 (2018), p. 221. ISSN: 1474-760X.
- [7] Guangliang Gao et al. “Location-Centered House Price Prediction: A Multi-Task Learning Approach”. In: *CoRR* abs/1901.01774 (2019). arXiv: 1901.01774.
- [8] David Hallac, Jure Leskovec, and Stephen Boyd. “Network lasso: Clustering and optimization in large graphs”. In: *KDD*. 2015.
- [9] David Hallac et al. “SnapVX: A Network-Based Convex Optimization Solver”. In: *arXiv preprint arXiv:1509.06397* (2015).
- [10] Kam Hamidieh. “A data-driven statistical model for predicting the critical temperature of a superconductor”. In: *Computational Materials Science* 154 (2018), pp. 346–354. ISSN: 0927-0256.
- [11] Hu Han, Charles Otto, and Anil Jain. “Age Estimation from Face Images: Human vs. Machine Performance”. In: June 2013.
- [12] Lucas Cassiel Jacaruso. “A method of trend forecasting for financial and geopolitical data: inferring the effects of unknown exogenous variables”. In: *Journal of Big Data* 5.1 (Dec. 2018), p. 47. ISSN: 2196-1115.
- [13] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [14] Gabriel Peyré, Marco Cuturi, et al. “Computational optimal transport”. In: *Foundations and Trends® in Machine Learning* 11.5-6 (2019), pp. 355–607.
- [15] Colin Raffel et al. *Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer*. 2019. arXiv: 1910.10683 [cs.LG].
- [16] B. Schölkopf and A. J. Smola. *Learning with Kernels*. Cambridge, MA: MIT Press, 2002.

- [17] Jack Sherman and Winifred J. Morrison. “Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix”. In: *Ann. Math. Statist.* 21.1 (Mar. 1950), pp. 124–127.
- [18] Tan Thongtan and Tanasanee Phienthrakul. “Sentiment Classification Using Document Embeddings Trained with Cosine Similarity”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 407–414.
- [19] R. Tibshirani. “Regression Shrinkage and Selection via the Lasso”. In: *Journal of the Royal Statistical Society (Series B)* 58 (1996), pp. 267–288.
- [20] R. Tibshirani. “Regression shrinkage and selection via the Lasso.” In: *Journal of the Royal Statistical Society, Series B* 58.1 (1996), pp. 267–288.
- [21] Robert Tibshirani et al. “Sparsity and smoothness via the fused lasso”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67.1 (2005), pp. 91–108.
- [22] Athanasios Tsanas et al. “Accurate Telemonitoring of Parkinson’s Disease Progression by Noninvasive Speech Tests”. In: *IEEE transactions on bio-medical engineering* 57 (Nov. 2009), pp. 884–93. DOI: 10.1109/TBME.2009.2036000.
- [23] Makoto Yamada et al. “Localized Lasso for High-Dimensional Regression”. In: *AISTATS*. 2017.
- [24] I-Cheng Yeh. “Modeling of Strength of High-Performance Concrete Using Artificial Neural Networks.” *Cement and Concrete research*, 28(12), 1797-1808”. In: *Cement and Concrete Research* 28 (Dec. 1998), pp. 1797–1808.
- [25] Yu-Dong Zhang and Lenan Wu. “Stock market prediction of SP 500 via combination of improved BCO approach and BP neural network”. In: *Expert Systems with Applications* 36 (July 2009), pp. 8849–8854.
- [26] H. Zou and T. Hastie. “Regularization and variable selection via the elastic net”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67.2 (2005), pp. 301–320.

Supplementary material

Implementation and searching the hyper-parameters

For each method, the implementation used for the experiments and the hyper-parameter space are explained for the optimization using optuna [1].

Proposed method

Our implementation will be available online soon. It was implemented in Python 3 with numpy and to yield a "sklearn" like model.

- The number of anchor points
 $k \in \{20, 40, 60, 80, 100, 120, 140, 160, 180, 200\}$
- The number of neighbors to create the anchor model
 $K_{\text{anchors}} \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$
- Use bias inside the models: True or False
- Regularization for the anchor models
 $\alpha \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$
- The l1-ratio in Elastic net for the anchor models
 $R \in \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$
- The regularization parameter
 $\lambda \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$
- The number of neighbors for the prediction
 $K_{\text{pred}} \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$

Network Lasso

We used the Python implementation of network Lasso following the work by the authors of Localized lasso [23]. It solves the same problem when we set $\lambda_{\text{exc}} = 0$. We tested it with $K = 5$ neighbors to construct the graph.

- The regularization parameter
 $\lambda_{\text{net}} \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$
- Use bias inside the models: True or False

There are fewer parameters here, but due to the restriction in training time, we cannot add more parameters.

Kernel Ridge Regression

We used the scikit-learn [13] implementation of KRR.

- The regularization parameter
 $\alpha \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$
- The kernel used: RBF or poly
- Parameters for RBF kernel
 - The inverse of the standard deviation of the RBF kernel
 $\gamma \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$

- Parameters for poly kernel
 - The degree of the polynomial
 $d \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$
 - Zero coefficient
 $c \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$

K-Nearest Neighbors (KNN)

We used the scikit-learn [13] implementation of KNN.

- The number of neighbors
 $K \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$
- Weights: Uniform or Weights (use inverse distance like our approach)

Ridge or Lasso

We used the scikit-learn [13] implementation of Ridge and Lasso.

- The regularization parameter
 $\alpha \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3\}$
- The number of neighbors
 $K \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$
- Use bias inside the models: True or False

Proofs of lemma

(Proof of Theorem 1) Let us rewrite the problem to show that it is a quadratic problem.

$$J(\mathbf{P}_i, \widehat{\mathbf{W}}_i) = \text{loss}(\widehat{\mathbf{W}}_i) + \lambda \text{reg}(\mathbf{P}_i, \widehat{\mathbf{W}}_i)$$

Let us focus on each term:

Mean square error term

$$\begin{aligned} \text{loss}(\widehat{\mathbf{W}}_i) &= \left\| \mathbf{y}_i - \widehat{\mathbf{W}}_i^\top \mathbf{x}_i \right\|^2 \\ &= \left\| \mathbf{y}_i - \frac{\mathbf{x}_i^\top \mathbf{x}_i}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{y}_i + \frac{\lambda}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{G}_i^\top \mathbf{x}_i \right\|^2 \\ &= \frac{\lambda^2}{(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)^2} \left\| \mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i \right\|^2 \end{aligned}$$

Regularization term First, we will decompose the square norm of difference as a sum of norms and scalar product.

$$\begin{aligned} \text{reg}(\mathbf{P}_i, \widehat{\mathbf{W}}_i) &= \sum_{l=1}^k p_{il} \left\| \widehat{\mathbf{W}}_i - \mathbf{A}_l \right\|_F^2 \\ &= \sum_{l=1}^k p_{il} \left[\left\| \widehat{\mathbf{W}}_i \right\|_F^2 - 2 \langle \widehat{\mathbf{W}}_i, \mathbf{A}_l \rangle_F + \left\| \mathbf{A}_l \right\|_F^2 \right] \\ &= \left\| \widehat{\mathbf{W}}_i \right\|_F^2 - 2 \langle \widehat{\mathbf{W}}_i, \mathbf{G}_i \rangle_F + \sum_{l=1}^k p_{il} \left\| \mathbf{A}_l \right\|_F^2 \end{aligned}$$

Let us focus on the first term of reg by injecting $\widehat{\mathbf{W}}_i$.

$$\begin{aligned} \left\| \widehat{\mathbf{W}}_i \right\|_F^2 &= \left\| \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top + \mathbf{G}_i \right\|_F^2 \\ &= \frac{1}{(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)^2} \left\| \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top \right\|_F^2 + 2 \left\langle \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i], \mathbf{G}_i \right\rangle_F + \left\| \mathbf{G}_i \right\|_F^2 \end{aligned}$$

The first norm of this expression can be rewritten as

$$\begin{aligned} \left\| \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top \right\|_F^2 &= \text{Tr}([\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i] \mathbf{x}_i^\top \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top) \\ &= \mathbf{x}_i^\top \mathbf{x}_i \text{Tr}([\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i] [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top) \\ &= \mathbf{x}_i^\top \mathbf{x}_i \left\| \mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i \right\|_2^2 \end{aligned}$$

Let us focus on the second term of reg by injecting $\widehat{\mathbf{W}}_i$.

$$\begin{aligned}\langle \widehat{\mathbf{W}}_i, \mathbf{G}_i \rangle_F &= \left\langle \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top + \mathbf{G}_i, \mathbf{G}_i \right\rangle_F \\ &= \left\langle \frac{1}{\lambda + \mathbf{x}_i^\top \mathbf{x}_i} \mathbf{x}_i [\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i]^\top, \mathbf{G}_i \right\rangle_F + \|\mathbf{G}_i\|_F^2\end{aligned}$$

The inner product term gets cancelled. Thus, we obtain the following formula.

$$\text{reg}(\mathbf{P}_i, \widehat{\mathbf{W}}_i) = \frac{\mathbf{x}_i^\top \mathbf{x}_i}{(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)^2} \|\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i\|_2^2 + \sum_{l=1}^k p_{il} \|\mathbf{A}_l\|_F^2 - \|\mathbf{G}_i\|_F^2$$

Combining the terms and factorization

Let us inject these intermediate derivations into our problem.

$$\begin{aligned}J(\mathbf{P}_i, \widehat{\mathbf{W}}_i) &= \text{loss}(\widehat{\mathbf{W}}_i) + \lambda \text{reg}(\mathbf{P}_i, \widehat{\mathbf{W}}_i) \\ &= \frac{1}{(\lambda + \mathbf{x}_i^\top \mathbf{x}_i)^2} \|\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i\|_2^2 (\lambda^2 + \lambda \mathbf{x}_i^\top \mathbf{x}_i) \lambda \left(\sum_{l=1}^k p_{il} \|\mathbf{A}_l\|_F^2 - \|\mathbf{G}_i\|_F^2 \right)\end{aligned}$$

The first term T_1 can be rewritten as

$$\begin{aligned}T_1 &= \beta_i \|\mathbf{y}_i - \mathbf{G}_i^\top \mathbf{x}_i\|_2^2 \\ &= \beta_i \left(\|\mathbf{y}_i\|_2^2 - 2\mathbf{y}_i^\top \mathbf{G}_i^\top \mathbf{x}_i + \|\mathbf{G}_i^\top \mathbf{x}_i\|_2^2 \right) \\ &= \beta_i \left(\|\mathbf{y}_i\|_2^2 - 2 \sum_{l=1}^k p_{il} \mathbf{y}_i^\top \mathbf{A}_l^\top \mathbf{x}_i + \sum_{l=1}^k \sum_{l'=1}^k p_{il} p_{il'} \mathbf{x}_i^\top \mathbf{A}_l \mathbf{A}_{l'}^\top \mathbf{x}_i \right)\end{aligned}$$

The second term T_2 can be rewritten as

$$\begin{aligned}T_2 &= \lambda \left(\sum_{l=1}^k p_{il} \|\mathbf{A}_l\|_F^2 - \|\mathbf{G}_i\|_F^2 \right) \\ &= \lambda \left(\sum_{l=1}^k p_{il} \|\mathbf{A}_l\|_F^2 - \sum_{l=1}^k \sum_{l'=1}^k p_{il} p_{il'} \text{Tr}(\mathbf{A}_l^\top \mathbf{A}_{l'}) \right)\end{aligned}$$

Finally, when we combine the sum in T_1 and T_2 , we get

$$\begin{aligned}J(\mathbf{P}_i, \widehat{\mathbf{W}}_i) &= \sum_{l=1}^k \sum_{l'=1}^k p_{il} p_{il'} \underbrace{(\beta_i \mathbf{x}_i^\top \mathbf{A}_l \mathbf{A}_{l'}^\top \mathbf{x}_i - \lambda \text{Tr}(\mathbf{A}_l^\top \mathbf{A}_{l'}))}_{(\mathbf{H}_i)_{ll'}} + \sum_{l=1}^k p_{il} \underbrace{(-2\beta_i \mathbf{y}_i^\top \mathbf{A}_l^\top \mathbf{x}_i + \lambda \|\mathbf{A}_l\|_F^2)}_{(\mathbf{b}_i)_l} + \beta_i \|\mathbf{y}_i\|_2^2 \\ &= \mathbf{p}_i^\top \mathbf{H}_i \mathbf{p}_i + \mathbf{b}_i^\top \mathbf{p}_i + \beta_i \|\mathbf{y}_i\|_2^2\end{aligned}$$

Thus, the final function to be minimized is quadratic. □