Research Repository UCD

Title	Efficient Sequence Regression by Learning Linear Models in All-Subsequence Space
Authors(s)	Gsponer, Severin, Smyth, Barry, Ifrim, Georgiana
Publication date	2017-12-30
Publication information	Gsponer, Severin, Barry Smyth, and Georgiana Ifrim. "Efficient Sequence Regression by Learning Linear Models in All-Subsequence Space." Springer, 2017.
Conference details	The European Conference on Machine Learning & Principles and Practice of Knowledge Discovery in Databases, Skopje, Macedonia 18-22 September 2017
Series	Lecture Notes in Computer Science
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/9054
Publisher's statement	The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-71246-8_3
Publisher's version (DOI)	10.1007/978-3-319-71246-8_3

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Efficient Sequence Regression by Learning Linear Models in All-Subsequence Space

Severin Gsponer, Barry Smyth, and Georgiana Ifrim

Insight Centre for Data Analytics, University College Dublin, Dublin, Ireland
{severin.gsponer, barry.smyth, georgiana.ifrim}@insight-centre.org

Abstract. We present a new approach for learning a sequence regression function, i.e., a mapping from sequential observations to a numeric score. Our learning algorithm employs coordinate gradient descent and Gauss-Southwell optimization in the feature space of all subsequences. We give a tight upper bound for the coordinate wise gradients of squared error loss that enables efficient Gauss-Southwell selection. The proposed bound is built by separating the positive and the negative gradients of the loss function and exploits the structure of the feature space. Extensive experiments on simulated as well as real-world sequence regression benchmarks show that the bound is effective and our proposed learning algorithm is efficient and accurate. The resulting linear regression model provides the user with a list of the most predictive features selected during the learning stage, adding to the interpretability of the method.

1 Introduction

A wide range of applications benefit from methods that can learn a mapping from sequential observations to categorical or numeric scores. For example, a mapping could be learned from a set of labeled DNA sequences, to classify each sequence into subfamilies [10], or assign to it a numeric score, such as a protein binding affinity [19]. Methods aimed at solving such problems typically employ Hidden Markov Models (HMM) [14, 20], kernel Support Vector Machines (SVM) [11] or more recently, Convolutional Neural Networks (CNN) [2]. While the accuracy of such techniques is promising, their efficiency and interpretability are still critical challenges. An alternative to the above approaches is to explicitly generate all k -mers (i.e., subsequences of length k) up to a specified k , followed by learning a classification or regression model using all the generated k -mers as features. Such methods are much simpler and achieve an accuracy comparable to the more sophisticated methods above [3, 19]. Nevertheless, they are limited by the huge computational burden of explicitly generating all k -mers, and therefore tend to be applied to small datasets with k fixed manually, e.g., up to 6 or 8, as well as use additional filtering to further reduce the large feature space.

In this work, we propose a regression approach that can use the entire space of k -mers, of unlimited length, by learning a linear model using an iterative branch-and-bound strategy. The main idea behind this technique is to exploit the nested structure of the feature space via greedy search, thus avoiding the need for explicitly generating all subsequences, and instead focusing on the most discriminative ones during learning. The resulting approach combines feature

selection and learning into a simple algorithm and, as shown in our experiments, delivers accuracy similar to the state-of-the-art, with no pre-processing or domain knowledge required. Since during learning we only need to explore a small subset of the feature space, we can employ richer features such as gapped k -mers to allow inexact feature matches. This enables our linear models to achieve high accuracy. Our optimization algorithm relies on greedy coordinate-descent with Gauss-Southwell selection. To enable efficient coordinate selection, we give a tight upper bound for the coordinate-wise gradients of squared error loss.

We test the proposed algorithm on a simulated benchmark and on two real-world applications. First, we compare our algorithm to other regression methods on a synthetic sequence regression dataset, where we vary parameters such as the true motif length and the alphabet size, to study potential gains from using rich features such as all k -mers. Next, we apply our model to a sequence regression problem where the goal is to score the protein binding affinity of DNA sequences. We work with a publicly available dataset of 40,000 DNA sequences prepared by [19] for a popular data challenge¹. Finally, we study our algorithm on a large sequence classification dataset to compare the effectiveness of our approach to existing methods designed for linear sequence classification [7, 8]. In this application the aim is to score software files represented as hexadecimal sequences, in order to categorize malicious software into known families, also known as malware classification [1]. This dataset was released by a recent data challenge² organized by Microsoft.

Contribution. We propose a new method for efficient sequence regression by learning linear models with rich subsequence features, e.g., unrestricted-length, contiguous and gapped k -mers. Our algorithm uses an optimization strategy based on coordinate-descent coupled with an upper bound on the coordinate-wise gradients of squared error loss, to enable efficient Gauss-Southwell selection. We evaluate our learning algorithm on simulated data and on two real-world applications and show that our simple linear models are as accurate as more complex state-of-the-art sequence regression methods, while requiring no feature engineering or heavy parameter tuning. We release all our code for producing synthetic sequence regression data, as well as for our learning algorithm³.

2 Related Work

We discuss a range of approaches for sequence regression and classification, with a focus on the two application domains studied in this paper.

Sequence Regression for DNA. The sequence regression benchmark provided by the DREAM5 challenge aims to advance the state-of-the-art in recognizing DNA-binding proteins. It consists of 40,000 DNA training sequences, each with a numeric score describing the binding specificity of a particular protein, from a class of proteins named Transcription Factors (TF). The task is to learn

¹ <https://www.synapse.org/#!Synapse:syn2887863/wiki/72186>

² <https://www.kaggle.com/c/malware-classification>

³ Code for our data and algorithm: <https://github.com/svgspomer/SqLoss>.

a regression function for a given TF, then predict the TF binding affinity for unseen sequences. There are 66 TF defining 66 different sequence regression tasks. The interpretability of the learned model is also important as new knowledge can be extracted from knowing the individual k -mer binding specificity. The work by [19] presents a systematic comparison of 26 methods evaluated on this benchmark. The winning method was a linear regression model with squared error loss and optimization by gradient descent on sequences represented in the feature space of all 4 to 8-mers [3]. Among the top-5 methods were a Markov model (Team_F) [9], an HMM trained by Expectation-Maximization combined with a linear model (Team_E) [16], and a linear regression model using contiguous and gapped 6-mers (Team_G)⁴.

A second round of the challenge has added new competing techniques, the most notable of which are a new HMM model for regression (RegHMM) [20] and a deep learning method based on CNN (named DeepBind) [2]. The linear regression method of Team_D came second. RegHMM which restricts the k -mers to $k = 6$, obtained results similar to Team_D. DeepBind slightly outperformed Team_D regarding prediction accuracy, but required a custom implementation for GPUs and extensive parameter calibration over 30 sets of parameters. We test our approach on the same TF-DNA benchmark and show that we can achieve similar accuracy to prior techniques, with a much simpler and more efficient approach.

Sequence Classification for Malware. Microsoft released a malware challenge benchmark on Kaggle in 2015 containing about 20,000 files amounting to 500GB of data, to enable training classifiers to identify malware files. Besides the large number of samples, individual files are also quite large (up to 50Mb per file). Each file can be interpreted as a discrete sequence of bytes. Several competing methods have exploited the sequence structure of the files. The winning method [18] has explicitly generated all k -mers with $k \in \{2, 3, 4\}$ using a machine with 104Gb memory, 16 CPUs, and extra 200 Gb disk space for the generated data. They have used these features with a boosting method implemented in the Xgboost library. The k -mer features have proven to be very useful for achieving high classification accuracy in the competition, but generating them explicitly requires extensive storage and computational resources. Recent work [1] has studied new feature engineering approaches on the same benchmark. They have decided not to use k -mer features due to the excessive computational requirements, but this has lead to lower accuracy than the challenge winner. In [5] the authors use explicitly generated 3-mers as features, but use mutual information and random projections to select a feature subset that is manageable for learning a neural network. Malware coders routinely use masking and other code obfuscation techniques. To detect such manipulations, gapped k -mers that allow flexible, rather than exact matching, could improve the classifier. To compare our regression algorithm to prior work, we treat sequence classification as a regression with binary scores. We learn and evaluate our linear regression function using *unrestricted-length, gapped k -mer features*.

Linear Sequence Classification. Some of our key intuitions come from the work of [7, 8] which proposed the SEQL framework implementing greedy

⁴ <http://www.nature.com/nbt/journal/v31/n2/extref/nbt.2486-S1.pdf>

approaches for linear sequence classification. In [8] it was shown that a branch-and-bound approach can be used for a variety of classification loss functions. We build on previous research [8] and propose an algorithm for efficient sequence regression, by exploiting the structure of the feature space and separately bounding the positive and negative coordinate-wise gradients for squared error loss. This bound allows us to guarantee that we iteratively find the best (wrt a given loss function) k -mer feature from a very large feature space. We compare our methods to the linear sequence classifiers from [7, 8] on the malware classification benchmark.

3 Method Proposed

3.1 Basic notation

Let $D = \{(s_1, y_1), (s_2, y_2), \dots, (s_N, y_N)\}$ be a training set of instance-label pairs, where $s_i = c_1 c_2 \dots c_{m_i}$ is a sequence of variable length m_i , with each $c_i \in \Sigma$ a symbol from the alphabet of possible symbols denoted by Σ . For example, in the case of DNA sequences $\Sigma = \{A, C, G, T\}$. Each sequence s_i has an associated score $y_i \in \mathbb{R}$. We represent a sequence s_i as a binary vector in the space of all subsequences in the training data: $x_i = (x_{i1}, \dots, x_{ij}, \dots, x_{in})^T, x_{ij} \in \{0, 1\}, i = 1, \dots, N$, where $x_{ij} = 1$ means that subsequence s_j occurs in sequence s_i . We denote by n the number of distinct subsequences in the feature space, i.e., the coordinates of the vectors space in which we learn. Although this space is huge and it is infeasible in practice to generate it explicitly, we will show how to work with this representation and a branch-and-bound approach to efficiently generate and navigate only parts of the space. The goal is to learn a mapping from sequences to scores, $f: S \rightarrow \mathbb{R}$, from the given training set D , so that we can predict a score $y \in \mathbb{R}$ for a new sample $s \in S$. In our framework we want to learn a linear model, i.e., a parameter vector β that allows us to estimate the real score y by setting $\hat{y} = \beta^T x_i$. Although linear models are not powerful enough to capture non-linear relationships, by working in a very complex feature space (e.g., all k -mers) we can learn a powerful model, similar to the kernel trick applied by kernel Support Vector Machines. We compute $\beta = (\beta_1, \dots, \beta_j, \dots, \beta_n)$ by minimizing a loss function over the training set:

$$\beta^* = \underset{\beta \in \mathbb{R}^n}{\operatorname{argmin}} L(\beta) \quad (1)$$

In our work $L(\beta)$ is the regularized squared loss:

$$L(\beta) = \sum_{i=1}^N (y_i - \beta^T \cdot x_i)^2 + CR_\alpha(\beta) \quad (2)$$

$C \in \mathbb{R}_0^+$ is the weight for the regularizer $R_\alpha(\beta)$. We use the elastic-net regularizer [6] which allows trading-off $l1$ and $l2$ penalties.

3.2 Learning via Coordinate-Descent with Gauss-Southwell Selection

Recent work [12] has shown that for a class of loss functions, which includes the squared loss, learning via coordinate descent is faster than random coordinate descent optimization. Furthermore, for squared loss in particular, coordinate descent via the Gauss-Southwell rule was proven to converge much faster than other coordinate descent methods. In our setting, the feature space of all subsequences is potentially exponential, thus it is not even possible to explicitly compute the full gradient. We show here how to use coordinate descent with the Gauss-Southwell rule for fitting squared loss in the feature space of all subsequences. We first give the generic learning algorithm and then provide an upper bound that makes Gauss-Southwell selection feasible for this complex feature space.

We are interested in solving the convex optimization problem in (1). The coordinate descent method is based on the iteration:

$$\beta^{(t)} = \beta^{(t-1)} - \eta_{j_t} \frac{\partial L}{\partial \beta_{j_t}}(\beta^{(t-1)}) e_{j_t} \quad (3)$$

To determine the descent direction we use the Gauss-Southwell rule [12]:

$$j_t = \arg \max_j \left| \frac{\partial L}{\partial \beta_j}(\beta^{(t-1)}) \right| \quad (4)$$

This formulation transforms the learning problem into a search problem since in each iteration we have to find the best coordinate, i.e., the subsequence with the largest absolute gradient value. Algorithm 1 shows the basic mechanics of our method. The crucial part of this algorithm is the search for the best coordinate (line 5).

Algorithm 1 Greedy Coordinate Descent with Gauss Southwell Selection

- 1: Set $\beta^{(0)} = 0$
 - 2: **while** !termination condition **do**
 - 3: Adjust intercept
 - 4: Calculate objective function $L(\beta^{(t)})$
 - 5: Find coordinate j_t with maximum gradient value
 - 6: Find optimal step size η_{j_t} by line search or exact optimization
 - 7: Update $\beta^{(t)} = \beta^{(t-1)} - \eta_{j_t} \frac{\partial L}{\partial \beta_{j_t}}(\beta^{(t-1)}) e_{j_t}$
 - 8: Add corresponding feature to feature set
 - 9: **end while**
-

Step Size. The parameter η_{j_t} is called a step size and acts as a scaling factor on the gradient value, to enforce convergence. The work in [12] analyzed a variety of options to set this parameter, from constant step size to exact optimization. As exact optimization was shown to produce much faster convergence, and is feasible to compute for squared loss, we also optimize η_{j_t} exactly.

Intercept Term. At the start of each iteration we calculate the partial derivative (Equation 5) and update the intercept value with a coordinate descent step.

$$\frac{\partial L}{\partial b}(\beta^{(t)}) = -2 \sum_{i=1}^N (y_i - (\beta^T x_i + b)) \quad (5)$$

3.3 Upper Bound for Fast Gauss-Southwell Selection

Formulating a learning algorithm via coordinate descent with Gauss-Southwell rule does not provide a solution for the problem of finding the best subsequence in a huge feature space. Here we give an upper bound on the coordinate-wise gradient value for the squared loss function, that enables us to efficiently search for the best coordinate in each iteration. The theory relies on the following intuition. First, the subsequence space has a structure that we can exploit to focus the search to only parts of the feature space. Namely, we can bound the frequency of a sequence, based on the frequency of any of its subsequences, using an argument similar to that of the Apriori market basket analysis algorithm. Second, we can separate the positive and negative terms of the gradient, to obtain an upper bound on the total gradient for squared loss. This allows us to incrementally generate feature candidates⁵ and to quickly rule out parts of the feature space, while guaranteeing to find a coordinate with maximum gradient magnitude, as required by the Gauss-Southwell optimization strategy. In the following, $s_j \subseteq s_i$ means that the corresponding vector entry $x_{ij} = 1$. Theorem 1 gives an upper bound on the gradient value of any subsequence s_j , using only information about its prefix s_p .

Theorem 1 (Bounding the search for the best coordinate). *Let $L(\beta)$ be the squared loss function and $y_i \in \mathbb{R}$. For any subsequence $s_p \subseteq s_j$, it holds that*

$$\left| \frac{\partial L}{\partial \beta_j}(\beta) \right| \leq \mu(s_p)$$

where

$$\mu(s_p) = \max \left\{ \left| \sum_{\{i | x_{ip}=1, y_i - \beta^T x_i \geq 0\}} -2x_{ip}(y_i - \beta^T x_i) + C(\alpha \text{sign}(\beta_j) + (\alpha - 1)\beta_j) \right|, \left| \sum_{\{i | x_{ip}=1, y_i - \beta^T x_i \leq 0\}} -2x_{ip}(y_i - \beta^T x_i) + C(\alpha \text{sign}(\beta_j) + (\alpha - 1)\beta_j) \right| \right\}$$

⁵ To generate candidate features we start from unigrams and use breadth-first-expansion to generate k -grams with $k \geq 1$.

Proof. We first focus on bounding the positive terms of the coordinate-wise gradients:

$$\frac{\partial L}{\partial \beta_j}(\beta) = \sum_{i=1}^N -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_j) \quad (6)$$

$$= \sum_{\{i|x_{ij}=1\}} -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_j) \quad (7)$$

$$\leq \sum_{\substack{\{i|x_{ij}=1, \\ y_i - \beta^T x_i \leq 0\}}} -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_j) \quad (8)$$

$$\leq \sum_{\substack{\{i|x_{ip}=1, \\ y_i - \beta^T x_i \leq 0\}}} -2x_{ip}(y_i - \beta^T x_i) + CR'_\alpha(\beta_j) \quad (9)$$

$$\leq \sum_{\substack{\{i|x_{ip}=1, \\ y_i - \beta^T x_i \leq 0\}}} -2x_{ip}(y_i - \beta^T x_i) + C(\alpha \text{sign}(\beta_j) + (\alpha - 1)\beta_j) \quad (10)$$

The last inequality holds since $\{i|x_{ij} = 1, y_i - \beta^T x_i \leq 0\} \subseteq \{i|x_{ip} = 1, y_i - \beta^T x_i \leq 0\}$ as every sequence which contains s_j also contains its subsequence s_p . Similarly, by separating the negative terms, we get

$$\frac{\partial L}{\partial \beta_j}(\beta) = \sum_{i=1}^N -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_j) \quad (11)$$

$$\geq \sum_{\substack{\{i|x_{ip}=1, \\ y_i - \beta^T x_i \geq 0\}}} -2x_{ip}(y_i - \beta^T x_i) + C(\alpha \text{sign}(\beta_j) + (\alpha - 1)\beta_j) \quad (12)$$

The two bounds provide an upper bound on the absolute value of the gradient at coordinate j :

$$\sum_{\{i|s_j \in x_i, y_i - \beta^T x_i \geq 0\}} -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_{j'}) \leq \frac{\partial L}{\partial \beta_j}(\beta) \quad (13)$$

$$\leq \sum_{\{i|s_j \in x_i, y_i - \beta^T x_i \leq 0\}} -2x_{ij}(y_i - \beta^T x_i) + CR'_\alpha(\beta_{j'}) \quad (14)$$

With the regularization term included, the bound depends on the prefix as well as on $\alpha \text{sign}(\beta_j) + (\alpha - 1)\beta_j$ and thus on the weight of the subsequence itself. Since in the beginning all β_j are set to zero this does not represent a serious problem. The bounds of features that were already selected in a previous iteration are the only ones that have to be adjusted.

The bound allows us to efficiently search for the coordinate with the largest gradient. Algorithm 2 shows the search procedure. We start the search by expanding from unigrams. Throughout the search we keep track of the current best feature (*best_feature*) and in τ we save its absolute gradient value. Before

the expansion of any subsequence, we check if the bound μ is smaller than τ . If this is the case, we can prune the subtree starting at this node, as no further expansion can improve the current gradient.

Algorithm 2 Fast Gauss-Southwell Coordinate Selection

```

1:  $\tau \leftarrow 0$ 
2:  $best\_feature \leftarrow NIL$ 
3: for all  $s' \in \bigcup_{i=1}^N \{s | s \in s_i, |s| = 1\}$  do                                 $\triangleright$  For each unigram
4:    $GROW\_SEQUENCE(s')$ 
5: end for
6: return  $best\_feature$ 
7:
1: function  $GROW\_SEQUENCE(s)$ 
2:   if  $\mu(s) \leq \tau$  then return                                                 $\triangleright \mu(s)$  like in Theorem 1
3:   else if  $abs(gradient(s)) > \tau$  then
4:      $best\_feature = s$                                                          $\triangleright$  Suboptimal solution
5:      $\tau = abs(gradient(s))$ 
6:   end if
7:   for all  $s'' \in \{s' | s' \supseteq s, s' \in \bigcup_{i=1}^N s_i, |s'| = |s| + 1\}$  do
8:      $GROW\_SEQUENCE(s'')$ 
9:   end for
10: end function

```

Proposition 1 (Tightness of upper bound). *The upper bound given in Theorem 1 is tight.*

Proof. It suffices to show one example in which the upper bound (13) is reached. The inequality becomes an equality when, e.g., $y_i - \beta^T x_i = 0, \forall i = 1, \dots, N$ or whenever all $y_i - \beta^T x_i \leq 0, \forall i = 1, \dots, N$ and the set of occurrences of a subsequence s_j is the same as that of its subsequences $s_p \subseteq s_j$, i.e., $\{i | x_{ij} = 1\} = \{i | x_{ip} = 1\}$.

Proposition 2 (Convergence rate). *The proposed learning algorithm for sequence regression by optimizing the squared loss, converges to the global optimum of the objective function with a convergence rate of*

$$L(\beta^{(t)}) - L(\beta^*) \leq \left[\prod_{r=1}^t \left(1 - \frac{\mu}{l_{j_r}} \right) \right] [L(\beta^{(0)}) - L(\beta^*)].$$

Proof. We use recent convergence results for coordinate-descent optimization of functions that are μ -strongly convex, with coordinate-wise l_{j_r} -Lipschitz continuous gradient (e.g., squared loss). In particular, we use coordinate descent with Gauss-Southwell selection for squared loss and exact step size optimization. For a detailed proof see [12].

Algorithm Complexity. The time complexity of the proposed algorithm is $O(fN)$ per iteration, where f is the number of features that need to be investigated for Gauss-Southwell selection and N is the number of training examples.

Implementation. In practice we use data structures such as inverted indexes and tries to fully take advantage of the sparsity and the nested structure of the feature space. We also investigate empirically the quality of the upper bound by computing the number of distinct features investigated per iteration, and measuring the running time, per iteration, and for fully learning a model.

4 Experiments

We evaluate our learning algorithm on synthetic data and two benchmarks from recent data challenges. First, we analyze and compare our method to other linear regression methods on simulated sequence regression data where we vary the data generation parameters. Next, we study a sequence regression problem, to compare our learning algorithm to state-of-the-art sequence regression approaches on real data. Finally, we study a sequence classification problem, in order to compare the squared loss bound effectiveness to related methods developed for sequence classification. We run all experiments on a PC with 132 GB RAM, single Intel Xeon 2.4 GHz CPU and 5.4TB HDD. All our code and data is available online⁶.

4.1 Synthetic Data

In this section we analyse our method (named **SqLoss**) on synthetic data. The artificial generation of data allows us to compare SqLoss to the state-of-the-art (SotA) methods in a systematic way. We generate sequence regression datasets according to Algorithm 3. Before the actual sequence generation starts, n motifs have to be generated by drawing m symbols from a given alphabet Σ . For each of these motifs, the influence on the response variable (i.e., the motif weight) is set randomly. The first step of the sequence generation is to define which motifs each sequence contains. The binary indicator variables I_{ij} encode this as in (15). Each of these indicator variables is set to 0 or 1 according to a user set probability. Depending on the value of the indicator variable, an insertion position for the motif in question is determined. Next, the actual generation of the string starts. For each position in the sequence, the algorithm checks if a motif has to be inserted. If not, a random symbol from the alphabet is inserted. Otherwise the corresponding motif is placed at this position. As soon as the end of the sequence is reached, a score is assigned to the sequence, according to (15), where ϵ is Gaussian noise, and the generation of the next sequence starts. This generation process can lead to the case that a motif is present in a sequence by chance. Our implementation checks all generated sequences for unintentionally inserted motifs and replaces them with a random subsequence of the same length as the motif.

$$y_i = \sum_{j=1}^n w_j I_{ij} + \epsilon, \text{ where } I_{ij} = \begin{cases} 1, & \text{if sequence } i \text{ contains motif } j \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

⁶ Code for our data and algorithm: <https://github.com/svgspomer/SqLoss>.

Algorithm 3 Generation of sequence regression dataset with $n = 2$ motifs

Input: Number of sequences N , length of sequence L , number of motifs n , motif length m , alphabet Σ .

Output: Dataset with N (sequence, score) pairs.

```
Generate  $n$  motifs by drawing  $m$  symbols  $\sim U(\Sigma)$ 
Set weights for each motif
for  $i < N$  do
  pos1  $\sim U(L)$  if  $I_{i1}$ 
  pos2  $\sim U(L)$  if  $I_{i2}$ 
  for  $l < L$  do
    if  $l = \text{pos1}$  or  $l = \text{pos2}$  then
      add motif to sequence
    else
      add symbol  $s \sim U(\Sigma)$  to sequence
    end if
  end for
  add sequence to data set with  $y_i = w_1 I_{i1} - w_2 I_{i2} + \epsilon$ 
end for
```

In the following experiment we generate 10,000 sequences of length 5,000 and insert 2 motifs. We compare SqLoss to three regression methods: ordinary least squares (ols), ridge regression (ridge) and linear support vector regression (linsvr). For all these methods we use the implementations in scikit-learn (version 0.17) [13] with default parameters settings. For SotA methods we explicitly generate all k -mers up to $k = 5$. For SqLoss we do not restrict k and use default parameters (see code online). Figure 1 and Figure 2 show the average mean squared error (MSE) and training time over 5 runs, for various alphabet sizes and 4 different motif lengths (3, 5, 7, 10). We note that the SotA methods’ performance suffers if the motifs are longer than the maximal extracted k -mer. With increasing alphabet size, this effect vanishes as the density of k -mers feature space decreases. If the k -mer density is low enough, subsequences (e.g., 5-mers) of motifs can already indicate the presence of the whole motif and so be used to learn an appropriate weight. When the motif is shorter than the extracted k -mers, all methods perform similar, even though with increasing alphabet size the SotA methods achieve slightly worse results. We suspect this is caused by overfitting since the feature space becomes huge. Even though this is a simplified regression problem, it is promising to see that SqLoss achieves comparable or better results in different data generation settings, without the need to set k explicitly. This means that our method requires much less feature engineering for achieving good prediction quality.

4.2 TF-DNA Binding Prediction Challenge

The DREAM5 challenge⁷ [19] provides 40,000 DNA sequences, each with a numeric score describing the binding affinity of a particular protein called Transcription Factor (TF). The task is to learn a regression function for a given

⁷ <http://hugheslab.ccb.utoronto.ca/supplementary-data/DREAM5/>

Fig. 1. Average MSE across 5 runs, comparing the impact of varying data generation parameters (alphabet size, motif length) on 4 regression methods. SqLoss has the lowest MSE across all data generation scenarios.

Fig. 2. Average training time (sec) across 5 runs, comparing the impact of varying data generation parameters (alphabet size, motif length) on 4 regression methods.

TF. There are 66 TF each defining a different sequence regression task. We use the DREAMTools [4] that allow us to compare our results to the results of the 26 challenge participants [19]. We also show results for DeepBind [2], a recent method that achieved higher scores than previous participants. Of special interest is the comparison to challenge winner [3], a linear regression method using k -mers as features. The authors of [3] pre-process the training data as follows: (1) log2-transform the target scores followed by subtracting the mean; (2) remove all sequences that were flagged as bad by the challenge organizers; (3) remove noisy training sequences; (4) filter low intensity probes and (5) restrict the k -mer length between 4 and 8. We only apply the first two transforms but none of the latter filters, as we want to study our algorithm’s effectiveness when using minimal domain knowledge. We also study the influence of using wildcards for allowing more flexible gapped k -mer features (e.g., $A*B$ where $*$ stands for any symbol of the alphabet).

The final rank is determined by averaging the ranks of each algorithm under each evaluation metric (see [19] for details). Table 1 shows these metrics for the top-5 methods on the benchmark. Our algorithm with one wildcard allowed (SqLossWc) comes second in the overall rank, right after DeepBind. If we do not allow wildcards (SqLoss) our method comes third. We can see the benefit of using wildcards by the increase of the score across all metrics. Figure 3 shows the loss functions and total number of explored nodes per iteration for TF47. Additionally, it shows the Pearson correlation achieved by SqLossWc at each iteration. The flexibility of wildcards clearly increases the number of nodes in the

Table 1. TF-DNA: Results of top-5 sequence regression methods for the DREAM5 data challenge.

Team	Pearson	Pearson_Log	Spearman	AUPR_8mer	AUROC_8mer
DeepBind	0.6780	0.7260	0.7060	0.6760	0.9910
SqLossWc	0.6483	0.6846	0.6423	0.7236	0.9967
SqLoss	0.6399	0.6791	0.6390	0.7049	0.9953
Team_D	0.6413	0.6742	0.6394	0.6997	0.9942
Team_E	0.6375	0.6936	0.6735	0.5223	0.9524
Team_F	0.6103	0.6732	0.6555	0.5456	0.9766

search tree which influences the runtime. SqLoss needs on average 513 seconds total training time per TF, and SqLossWc takes 2,834 seconds total training time per TF. Our algorithm reaches similar metrics to Team_D and DeepBind, but the variation without wildcards uses a fraction of the training time of DeepBind and no specific pre-processing as Team_D. The average length of the learned motifs for SqLoss is 5.05 (std across different TFs: 0.16) with a maximum length of 10 bases. For SqLossWc the average motif length is 5.4 (std: 0.09) and the longest learned motif is 11 bases long.

An advantage of our linear model is that it can easily be interpreted, unlike complex non-linear methods such as DeepBind. Table 2 shows some high rank features for TF13 learned with k -mer features with one wildcard allowed. The weights directly reflect the importance of the learned features and can provide important knowledge to domain experts. The learned features are stored in a trie to allow efficient prediction by linearly scanning test sequences. Prediction for 40,000 sequences takes 2 seconds for SqLoss and 4 seconds for SqLossWc.

Fig. 3. TF-DNA: Loss and explored nodes per iteration, w/o wildcard for protein TF47. In blue SqLoss, in red SqLossWc. Pearson correlation (green) for SqLossWc.

Table 2. TF-DNA: Example positive/negative features from SqLossWc model learned for TF13, where * indicates a wildcard.

Motif	Weight
TAAT*A	0.733985
TAATG*G	0.706344
ATG*AAA	0.674507
⋮	⋮
GGATA	-0.188202
TCAAT	-0.214858
G*ATAG	-0.218132

4.3 Microsoft Malware Classification Challenge

The goal of the Microsoft challenge⁸ is to classify files into one of 9 malware families. The training set has 10,868 labeled samples, each with a binary file with hexadecimal representation and a file with the disassembled code. We want to find out if we can build effective classifiers using only the binary (bytes) representation, since disassembling the code requires expensive computation. Even though this dataset poses a classification, rather than a regression task, we use it to compare our method to the challenge winning methods and to similar branch-and-bound approaches developed for linear classification, as implemented in the SEQL framework [8]. We compare the accuracy, convergence and bound effectiveness for three learning algorithms: SEQL with classification losses (logistic loss and quadratic hinge loss) versus our SqLoss regression algorithm (for SqLoss we interpret one-vs-all binary labels as numeric scores).

To reduce each binary file we remove the offset field as well as all question marks and white spaces between hexa bytes. The original challenge uses the multiclass log-loss as main evaluation metric. We do not use this metric as it heavily depends on the calibration of the output scores of each method. Similar to published work [18, 1], we report the accuracy results on 4-fold stratified cross-validation (Table 3). The same input data is used for the SEQL methods and SqLoss. We note that SqLoss has similar accuracy to the challenge winner [18]

Table 3. Malware: Accuracy and training time 4-fold CV.

Method	Accuracy	Training time(mins)
Wang et al. (Challenge Winner)	0.9983	(multicore, preprocessing only) 2’880
Ahmedi et al.	0.9976	(multicore, preprocessing only) 2’780
SEQL Logistic Regression Loss	0.9958	(singlecore, full training) 603
SEQL Quadratic Hinge Loss	0.9949	(singlecore, full training) 410
SqLoss (Our Method)	0.9916	(singlecore, full training) 3’183

⁸ <https://www.kaggle.com/c/malware-classification>

Fig. 4. Malware: Comparison of loss functions per iteration for all 9 classes. Blue for squared hinge loss, green for logistic regression and red for SqLoss.

and the recent solution of [1]. Nevertheless, both those methods use a variety of hand picked features extracted from both the binary and the disassembled files. Disassembled code is expensive to extract and is inexact, i.e., it is possible for a single program to have two or more disassemblies. To extract features [18, 1] have to do costly preprocessing of the data and heavy feature engineering. In particular [18] have to limit k -mers to max $k = 4$ and need explicit generation of k -mers (for both training and test data) which requires 100GB memory, 16 CPUs, and 48 hours extraction time on the training data alone.

To better analyze our method we compare it to SEQL, a linear classification method that uses branch-and-bound for selecting subsequence features. As we can see in Table 3, the accuracy of the SEQL losses is comparable to that of SqLoss, while training time is better for the classification losses. Figure 4 shows the value of the three loss functions (normalized by the start loss) per iteration, for all classes. As in [15], we also find that the SqLoss decreases slower than the classification loss functions. This is expected given that SqLoss is a regression algorithm used for a classification task. We currently implement the same stopping criterion for all three losses and we believe this may be ill suited for SqLoss, as in practice we could stop the iterations earlier without compromising accuracy. Figure 5 shows the total number of nodes explored per iteration during the search for the best feature, for each of the three methods. The pruning for SqLoss seems to be less efficient than the one for logistic regression or the hinge loss, on this classification task.

Fig. 5. Malware: Number of nodes explored by different loss functions (blue for hingeloss, green for logreg, red for SqLoss) for all 9 classes.

5 Conclusion and Future Work

We present a new method for efficient linear sequence regression in the feature space of all possible subsequences. The proposed method uses coordinate gradient descent with Gauss-Southwell rule to optimize squared error loss. We propose a branch-and-bound strategy for efficient Gauss-Southwell selection. Our empirical study shows that we can achieve results comparable to the state-of-the-art, with a simple linear regression model, while employing little to no domain knowledge or pre-processing. In particular, our models can use unrestricted, flexible k -mer features (with wildcards), without compromising training and testing efficiency. In the future we want to further explore more flexible operators on feature representations of sequences. Further, we want to improve the scalability of our method as well as extend it to make use of multiple cores.

Acknowledgment: This work was funded by Science Foundation Ireland (SFI) under grant number 12/RC/2289.

References

1. Ahmadi, M., Ulyanov, D., Semenov, S., Trofimov, M., Giacinto, G.: Novel feature extraction, selection and fusion for effective malware family classification. In: Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy. pp. 183–194. ACM (2016)

2. Alipanahi, B., Delong, A., Weirauch, M.T., Frey, B.J.: Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nature Biotechnology* 33(8), 831–838 (2015)
3. Annala, M., Laurila, K., Lähdesmäki, H., Nykter, M.: A linear model for transcription factor binding affinity prediction in protein binding microarrays. *PLoS ONE* 6(5), 1–13 (2011)
4. Cokelaer, T., Bansal, M., Bare, C., et al.: Dreamtools: a python package for scoring collaborative challenges (2015)
5. Dahl, G.E., Stokes, J.W., Deng, L., Yu, D.: Large-scale malware classification using random projections and neural networks. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 3422–3426. IEEE (2013)
6. Hui, Z., Hastie, T.: Regularization and variable selection via the elastic net. *Journal Of The Royal Statistical Society Series B* 67(2), 301–320 (2005)
7. Ifrim, G., Bakir, G., Weikum, G.: Fast logistic regression for text categorization with variable-length n-grams. In: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 354–362. ACM (2008)
8. Ifrim, G., Wiuf, C.: Bounded Coordinate-Descent for Biological Sequence Classification in High Dimensional Predictor Space. *Proceedings of the 17th ACM SIGKDD* (2011)
9. Keilwagen, J., Grau, J., Paponov, I.A., Posch, S., Strickert, M., Grosse, I.: De-novo discovery of differentially abundant transcription factor binding sites including their positional preference. *PLoS Comput Biol* 7(2), e1001070 (2011)
10. Leslie, C., Eskin, E., Noble, W.S.: The Spectrum Kernel: a String Kernel for Svm Protein Classification. *Proceedings of the Pacific Symposium on Biocomputing* 7, 564–575 (2002)
11. Leslie, C., Kuang, R.: Fast string kernels using inexact matching for protein sequences. *Journal of Machine Learning Research* 5(Nov), 1435–1455 (2004)
12. Nutini, J., Schmidt, M., Laradji, I.H., Friedlander, M., Koepke, H.: Coordinate Descent Converges Faster with the Gauss-Southwell Rule Than Random Selection. *International Conference on Machine Learning* 37(Section 5), 1–34 (2015)
13. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830 (2011)
14. Punta, M., Coggill, P.C., Eberhardt, R.Y., Mistry, J., Tate, J., Boursnell, C., Pang, N., Forslund, K., Ceric, G., Clements, J., Heger, A., Holm, L., Sonnhammer, E.L.L., Eddy, S.R., Bateman, A., Finn, R.D.: The pfam protein families database. *Nucleic Acids Research* 40(D1), D290–D301 (2011)
15. Rosasco, L., De Vito, E., Caponnetto, A., Piana, M., Verri, A.: Are loss functions all the same? *Neural Computation* 16(5), 1063–1076 (2004)
16. Schütz, F., Delorenzi, M.: Mamot: hidden markov modeling tool. *Bioinformatics* 24(11), 1399–1400 (2008)
17. Vapnik, V.: The nature of statistical learning theory. Springer Science & Business Media (2013)
18. Wang, X., Liu, J., Chen, X.: Microsoft malware classification challenge (big 2015) first place team: Say no to overfitting, 2015. *BIG* (2015)
19. Weirauch, M.T., Cote, A., Norel, R., Annala, M.: Evaluation of methods for modeling transcription factor sequence specificity. *Nature Biotechnology* 31(2), 126–134 (2013) <http://www.nature.com/nbt/journal/v31/n2/abs/nbt.2486.html>
20. Zhang, Y., Henao, R., Carin, L., Zhong, J., Hartemink, A.: Learning a hybrid architecture for sequence regression and annotation (2016) <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11869>