
Piecewise Linear Regression via a Difference of Convex Functions

Ali Siahkamari^{*1} Aditya Gangrade^{*2} Brian Kulis¹ Venkatesh Saligrama¹

Abstract

We present a new piecewise linear regression methodology that utilizes fitting a *difference of convex* functions (DC functions) to the data. These are functions f that may be represented as the difference $\phi_1 - \phi_2$ for a choice of *convex* functions ϕ_1, ϕ_2 . The method proceeds by estimating piecewise-linear convex functions, in a manner similar to max-affine regression, whose difference approximates the data. The choice of the function is regularised by a new seminorm over the class of DC functions that controls the ℓ_∞ Lipschitz constant of the estimate. The resulting methodology can be efficiently implemented via Quadratic programming *even in high dimensions*, and is shown to have close to minimax statistical risk. We empirically validate the method, showing it to be practically implementable, and to have comparable performance to existing regression/classification methods on real-world datasets.

1. Introduction

The multivariate nonparametric regression problem is a fundamental statistical primitive, with a vast history and many approaches. We adopt the following setup: given a dataset, $\{(\mathbf{x}_i, y_i)\}_{i \in [n]}$, where $\mathbf{x}_i \in \mathbb{R}^d$ are predictors, assumed drawn i.i.d. from a law P_X , and y_i are responses such that $y_i = f(\mathbf{x}_i) + \varepsilon_i$, for a bounded, centered, independent random noise ε_i , and bounded f , the goal is to recover an estimate \hat{f} of f such that on new data, the squared error $\mathbb{E}[(y - \hat{f}(\mathbf{x}))^2]$ is small.

The statistical challenge of the problem lies in the fact that f is only weakly constrained - for instance, f may only be known to be differentiable, or Lipschitz. In addition, the

problem is algorithmically challenging in high-dimensions, and many approaches to the univariate problem do not scale well with the dimension d . For instance, piecewise linear regression methods typically involve a prespecified grid, and thus the number of grid points, or knots, grows exponentially with dimension. Similarly, methods like splines typically require both stronger smoothness guarantees and exponentially more parameters to fit with dimension in order to avoid singularities in the estimate.

This paper is concerned with regression over the class of functions that are *differences of convex* functions, i.e., DC functions. These are functions f that can be represented as $f = \phi_1 - \phi_2$ for a choice of two *convex* functions. DC functions constitute a very rich class - for instance, they are known to contain all \mathcal{C}^2 functions. Such functions have been applied in a variety of contexts including non-convex optimization (Yuille & Rangarajan, 2002; Horst & Thoai, 1999), sparse signal recovery (Gasso et al., 2009), feature selection (Le Thi et al., 2008), and reinforcement learning (Piot et al., 2014).

The principal contribution of this paper is a method for piecewise linear regression over the class of DC functions. At the heart of the method is a representation of piecewise linear DC functions via a set of linear constraints, in a manner that generalises the representations used for max-affine regression. The choice of the function is regularised for smoothness by a new seminorm that controls the ℓ_∞ -Lipschitz constant of the resulting function. The resulting estimate is thus a piecewise linear function, represented as the difference of two piecewise linear convex functions, that are smooth in the sense of having bounded gradients.

The method enjoys two main advantages:

1. It is agnostic to any knowledge about the function, and requires minimal parameter tuning.
2. It can be implemented efficiently, via quadratic programming, even in high dimensions. For n data points in \mathbb{R}^d , the problem has $2n(2d + 1) + 1$ decision variables, and n^2 linear constraints, and can be solved in the worst case in $O(d^2 n^5)$ time by interior-point methods. Furthermore the algorithm does not need to specify partitions for piece-wise linear parts and avoids ad-hoc generalizations of splines or piece-wise linear methods to multi-dimensions.

^{*}Equal contribution ¹Department of Electrical and Computer Engineering, Boston University ²Division of Systems Engineering, Boston University. Correspondence to: Ali Siahkamari <siaa@bu.edu>.

In addition, the method is shown to be statistically viable, in that it is shown to attain vanishing risk as the sample size grows at a non-trivial rate, under the condition that the ground truth has bounded DC seminorm. The risk further adapts to structure such as low dimensional supports.

Lastly, the solution obtained is a piecewise-linear fit, and thus enjoys interpretability in that features contribution heavily to the value can be readily identified. Further, the fitting procedure naturally imposes ℓ_1 regularisation on the weights of the piecewise linear parts, thus enforcing a sparsity of local features, which further improves interpretability.

To establish practical viability, we implement the method on a number of regression and classification tasks. The settings explored are those of moderate data sizes - $n \leq 10^3$, and dimensions $d \leq 10^2$. We note that essentially all nonparametric models are only viable in these settings - typical computational costs grow with n and become infeasible for large datasets, while for much higher dimensions, the sample complexity - which grows exponentially with d - cause small datasets to be non-informative. More pragmatically, all nonparametric methods we compare against have been evaluated on such data. Within these constraints, the method is shown to have better error performance and fluctuation with respect to popular methodologies such as multivariate adaptive regression splines, nearest neighbour methods, and two-layer perceptrons, evaluated on both synthetic and real world data-sets.

1.1. Connections to Existing Methodologies

Piecewise Linear Regression is popular since such regressors can model the local features of the data without affecting the global fit. In higher than 1 dimensions, piecewise linear functions are usually fit via choosing a partition of the space and fitting linear functions on each part. The principle difficulty thus lies in choosing these partitions. The approach is usually a rectangular grid - for instance, a variable rectangular partition of the space is studied in (Toriello & Vielma, 2012) and solved optimally. However the rectangulization becomes prohibitive in high dimension as the number of parts grow exponentially with the dimension. Other approaches include Bayesian methods such as (Holmes & Mallick, 2001), which rely on computing posterior means for the parameters to be fit via MCMC methods.

Max-Affine Regression is a nonparametric approximation to convex regression, originating in (Hildreth, 1954; Halloway, 1979) that recovers the optimal piece-wise linear approximant to a convex function with the form $f = \max_{i \in [K]} \langle a_i, \mathbf{x}_i \rangle + b_i$ for some K . Smoothness of the estimate can be controlled by constraining the convex function to be Lipschitz. The problem is generic in that it is easily argued that piecewise linear convex functions can uniformly approximate any Lipschitz convex function on a bounded

domain. Parametric approaches, i.e., with a fixed K , are popular, but can be computationally intensive due to the induced combinatorics of which of the K planes is maximised at which data point, and various heuristics and partitioning techniques have to be applied (Magnani & Boyd, 2009; Hannah & Dunson, 2013; Ghosh et al., 2019). The nonparametric case, where K grows with n , has been extensively analysed in the works (Balázs et al., 2015; Balázs, 2016).

On the other hand, if $K = n$, i.e. if the number of affine functions used is the same as the number of datapoints, then the problem becomes amenable to convex programming techniques - when estimating the parameters a_i, b_i , one can remove the nonlinearity induced by the max, and instead enforce the same via n linear constraints. This simple fact allows efficient algorithmic approaches to max-affine functions. The heart of our method for DC function regression is an extension of this trick to DC functions.

Smoothing splines are an extremely popular regression methodology in low dimensions. The most popular of these are the L_2 smoothing splines, which, in one dimension, involve fixing a ‘knot’ at each data point, and estimating the gradients of the function at each point, with regularisation of the form $\int |\hat{f}''|^2$. Unfortunately this L_2 regularisation leads to singularities in $d \geq 3$ dimensions, and methods such as thin plate splines generalising these to higher dimensions resort of regularising up to the d th derivative of the estimate, leading to an explosion in the number of parameters to be estimated (Wahba, 1990).

Our method is closer in relation to L_1 regularised splines, which in the univariate case regularise $\int |\hat{f}'|$ - it is shown in Appx. A that in one dimension our method reduces to these. As a consequence, one may view this method as a new generalisation of the L_1 -spline regressor.

A number of alternative methods for multivariate splines have been proposed, with several, such as general additive models modelling the data via low dimensional projections and assumptions. The most relevant multivariate spline methods are the adaptive regression splines, originating in (Friedman et al., 1991), which is a greedy procedure for recursively refining a partition of the data, and fitting new polynomials over the parts.

Previous DC Modeling Finally, let us mention that the final chapter of the doctoral thesis of Balázs (2016) anticipates our study of DC function regression, but gives neither algorithms nor analyses. Subsequently, Cui et al. (2018) introduces the same DC modeling as us in a broader context, where the loss function can also be a DC function. However their problem ends up being non-convex. They focus on developing a majorization-minimization algorithm to find an approximate solution with desirable guarantees such as convergence to a directional stationary solution.

2. A brief introduction to DC functions

Difference of convex functions are functions which can be represented as difference of two continuous convex function over a domain $\Omega \subset \mathbb{R}^d$, i.e., the class

$$\mathcal{DC}(\Omega) \triangleq \{f : \Omega \rightarrow \mathbb{R} \mid \exists \phi_1, \phi_2 \text{ convex, continuous s.t.} \\ f = \phi_1 - \phi_2\}.$$

Throughout the text we will set $\Omega = \{\mathbf{x} : \|\mathbf{x}\|_\infty \leq R\}$. We will assume that the noise and the ground truth function are bounded so that $|\varepsilon|, \sup_{\mathbf{x} \in \Omega} |f(\mathbf{x})| \leq M$. As a consequence, $|y| \leq 2M$.

One of the first characterizations of DC functions is due to Hartman et al. (1959): a univariate function is a DC function if and only if its directional derivatives each has bounded total variation on all compact subsets of Ω . For higher dimensions it is known that DC functions are a subclass of locally-Lipschitz functions and include C^2 functions. Therefore, any continuous function can be uniformly approximated by DC functions. For a recent review see Bačák & Borwein (2011).

In the following section we show that DC functions can fit any sample data. Thus, to allow a bias-variance tradeoff, we regularise the selection of DC functions via the DC seminorm

$$\|f\| \triangleq \inf_{\phi_1, \phi_2} \sup_{\mathbf{x}} \sup_{\mathbf{v}_i \in \partial_* \phi_i(\mathbf{x})} \|\mathbf{v}_1\|_1 + \|\mathbf{v}_2\|_1 \\ \text{s.t. } \phi_1, \phi_2 \text{ are convex, } \phi_1 - \phi_2 = f,$$

where $\partial_* \phi_i$ denotes the set of subgradients of ϕ_i . The above function is not a norm because every constant function satisfies $\|c\| = 0$. Indeed, if we equate DC functions up to a constant, then the above seminorm turns into a norm. We leave a proof of the fact that the above is a seminorm to Appx. B.1. Note that the DC seminorm offers strong control on the ℓ_∞ -Lipschitz constant of the convex parts of at least one DC representation of the function (and in turn on the Lipschitz constant of the function).

We will principally be interested in DC functions with bounded DC seminorm, and thus define

$$\mathcal{DC}_L \triangleq \{f \in \mathcal{DC} : \|f\| \leq L\}.$$

The bulk of this paper concentrates on *piecewise linear DC functions*. This is justified because piecewise linear functions are known to uniformly approximate bounded variation functions, and structural results (Kripfganz & Schulze, 1987; Ovchinnikov, 2002) showing that every piecewise linear function may be represented as difference of two convex piecewise linear functions - i.e., max-affine functions. Since the term is used very often, we will abbreviate “piecewise

linear DC” as PLDC, and symbolically define

$$\text{pl-DC} \triangleq \{f \in \mathcal{DC} : f \text{ is piecewise linear}\}, \\ \text{pl-DC}_L \triangleq \{f \in \mathcal{DC}_L : f \text{ is piecewise linear}\}.$$

The following bound on the seminorm of PLDC functions is useful. The proof is obvious, and thus omitted.

Proposition 1. *Every $f \in \text{pl-DC}$ can be represented as a difference of two max-affine functions*

$$f(\mathbf{x}) = \max_{k \in [K]} \langle \mathbf{a}_k, \mathbf{x} \rangle + c_k - \max_{k \in [K]} \langle \mathbf{b}_k, \mathbf{x} \rangle + c'_k$$

for some finite K . For such an f , $\|f\| \leq \max_k \|\mathbf{a}_k\|_1 + \max_k \|\mathbf{b}_k\|_1$.

2.1. Expressive Power of Piecewise linear DC functions

We begin by arguing that PLDC functions can interpolate any finite data. The principle characterisation for DC functions is as follows:

Proposition 2. *For any finite data $\{(\mathbf{x}_i, y_i)\}_{i \in [n]}$, there exists a DC function $h : \mathbb{R}^d \rightarrow \mathbb{R}$, that takes values $h(\mathbf{x}_i) = y_i$ if and only if there exist $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^d, z_i \in \mathbb{R}, i \in [n]$ such that:*

$$y_i - y_j + z_i - z_j \geq \langle \mathbf{a}_j, \mathbf{x}_i - \mathbf{x}_j \rangle, \quad i, j \in [n] \\ z_i - z_j \geq \langle \mathbf{b}_j, \mathbf{x}_i - \mathbf{x}_j \rangle, \quad i, j \in [n] \quad (1)$$

Further, if there exists a DC function that interpolates a given data, then there exists a PLDC function that also interpolates the data.

Proof. Assuming $h = \phi_1 - \phi_2$ for convex functions ϕ_1 and ϕ_2 , take \mathbf{a}_j and \mathbf{b}_j to be sub-gradients of respectively ϕ_1 and ϕ_2 then (1) holds by convexity. Conversely, assuming (1) holds, define h as

$$h(\mathbf{x}) = \max_{i \in [n]} \langle \mathbf{a}_i, \mathbf{x} - \mathbf{x}_i \rangle + y_i + z_i - \max_{i \in [n]} \langle \mathbf{b}_i, \mathbf{x} - \mathbf{x}_i \rangle + z_i$$

$h \in \mathcal{DC}$ since it is expressed as the difference of two max-affine functions. Further, it holds that $h(\mathbf{x}_k) = y_k$ for any $k \in [n]$. Indeed, the first condition implies that for any $i \neq k$,

$$\langle \mathbf{a}_i, \mathbf{x}_k - \mathbf{x}_i \rangle + y_i + z_i \leq y_k + z_k,$$

with equality when $i = k$. Thus, the first maximum simply takes the value $y_k + z_k$ at the input \mathbf{x}_k . Similarly, the second maximum takes the value z_k at this input, and thus $h(\mathbf{x}_k) = (y_k + z_k) - z_k = y_k$.

Notice that the interpolating function given in the above is actually piecewise-linear. Thus, if a DC function fits the given data, then extracting the vectors $\mathbf{a}_i, \mathbf{b}_i$ and constants z_i as in the first part of the proof, and constructing the interpolant in the second part yields a PLDC function that fits the data. \square

The principal utility of the conditions stated above is that we can utilise these to enforce the shape restriction of getting a DC estimate in an efficient way when performing empirical risk minimisation. Indeed, suppose we wish to fit a DC function to some data (\mathbf{x}_i, y_i) . We may then introduce decision variables $\hat{y}_i, z_i, \mathbf{a}_i, \mathbf{b}_i$, where the \hat{y}_i represent the value of our fit at the various \mathbf{x}_i , and then enforce the linear constraints of the above proposition (with y_i replaced by \hat{y}_i) while minimising a loss of the form $\sum (y_i - \hat{y}_i)^2$. Since these constraints are linear, the resulting program is thus convex, and can be efficiently implemented. This observation forms the core of our algorithmic proposal in §3

The above characterisation relies on existence of vectors that may serve as subgradients for the two convex functions. This condition can be removed, as in the following.

Proposition 3. *Given any finite data $\{(\mathbf{x}_i, y_i)\}_{i \in [n]}$, such that $y_i \neq y_j \implies \mathbf{x}_i \neq \mathbf{x}_j$, there exists a PLDC function which interpolates this data.*

Proof. The interpolating function is constructed by adding and subtracting a quadratic function to the data. Let

$$C \triangleq \max_{i,j} \frac{|y_i - y_j|}{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}.$$

Then the piecewise linear function

$$\begin{aligned} h(\mathbf{x}) \triangleq & \max_{i \in [n]} \langle C\mathbf{x}_i, \mathbf{x} - \mathbf{x}_i \rangle + \frac{1}{2}C\|\mathbf{x}_i\|^2 + \frac{1}{2}y_i \\ & - \max_{i \in [n]} \langle C\mathbf{x}_i, \mathbf{x} - \mathbf{x}_i \rangle + \frac{1}{2}C\|\mathbf{x}_i\|^2 - \frac{1}{2}y_i \end{aligned} \quad (2)$$

satisfies the requirements. Indeed, the function is DC, since it the difference of two max-affine functions. The argument proceeds similarly to the previous case - at any \mathbf{x}_j , we have:

$$\begin{aligned} & \max_{i \in [n]} \langle C\mathbf{x}_i, \mathbf{x}_j - \mathbf{x}_i \rangle + \frac{1}{2}C\|\mathbf{x}_i\|^2 + \frac{1}{2}y_i \\ & = \max_{i \in [n]} \frac{1}{2}C\|\mathbf{x}_j\|^2 - \frac{C}{2}\|\mathbf{x}_i - \mathbf{x}_j\|^2 + \frac{1}{2}y_i \\ & \leq \max_{i \in [n]} \frac{1}{2}C\|\mathbf{x}_j\|^2 - \frac{|y_i - y_j|}{2} + \frac{1}{2}y_i \\ & \leq \frac{1}{2}C\|\mathbf{x}_j\|^2 + \frac{1}{2}y_j. \end{aligned}$$

However the upper bound is reached by choosing $i = j$ in the max operator. Similarly the value of the second convex function at \mathbf{x}_j is equal to $\frac{1}{2}C\|\mathbf{x}_j\|^2 - \frac{1}{2}y_j$ which together results in $h(\mathbf{x}_j) = y_j$. Note that $h(\mathbf{x})$ has the ℓ_∞ -Lipschitz constant $2C \max_{i,j} \|\mathbf{x}_i - \mathbf{x}_j\|_1$, and $\|h\| \leq 2C \max_i \|\mathbf{x}_i\|_1$. \square

Next, in order to contextualise the expressiveness of DC functions, we argue that the popular parametric class of

ReLU neural networks can be represented by PLDC functions, and vice versa. This is also argued in (Cui et al., 2018) for a 2 layer network.

Proposition 4. *A fully connected neural network f , with ReLU activations, and D layers with weight matrices $\mathbf{W}^1, \dots, \mathbf{W}^D$, i.e,*

$$\begin{aligned} f(\mathbf{x}) &= \sum_j w_j^{D+1} a_j^D \\ a_i^{l+1} &= \max\left(\sum_j w_{i,j}^{l+1} a_j^l, 0\right), \quad D > l \geq 1 \\ a_i^1 &= \max\left(\sum_j w_{i,j}^1 x_j, 0\right), \end{aligned}$$

is a PLDC function with the DC seminorm bounded as

$$\|f\| \leq |\mathbf{w}^{D+1}|^T \left(\prod_{l=1}^D |\mathbf{W}^l| \right) \vec{\mathbf{1}},$$

where $|\cdot|$ is the entry-wise absolute value. The above is proved in Appx. B.2 via an induction over the layers using the relations

$$\begin{aligned} \max(\max(a, 0) - \max(b, 0), 0) &= \max(a, b, 0) - \max(b, 0) \\ \max(a, b) + \max(c, d) &= \max(a + b, a + d, b + c, b + d). \end{aligned}$$

Proposition 5. *Every PLDC function with K hyper-planes can be represented by a ReLU net with $2\lceil \log_2 K \rceil$ layers and maximum width of $8K$.*

The proof is constructed in Appx. B.2 using the relations

$$\begin{aligned} \max(a, b, c, d) &= \max(\max(a, b), \max(c, d)) \\ \max(a, b) &= \max(a - b, 0) + \max(b, 0) - \max(-b, 0). \end{aligned}$$

3. Algorithms

We begin by motivating the algorithmic approach we take. This is followed by separate section developing key portions of the algorithm.

Suppose we observe a data-set $S_n = \{(\mathbf{x}_i, y_i)\}_{i \in [n]}$ generated iid from some distribution $\mathcal{X} \times \mathcal{Y}$ and a convex loss function $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ bounded by c . The goal is to minimize the expected risk

$$\min_{f \in \mathcal{DC}} \mathbb{E}[(f(\mathbf{x}) - y)^2] \quad (3)$$

by choosing an appropriate function f from the class of DC functions. Note instead of the squared error, the above could be generalised to any bounded, Lipschitz losses $\ell(f(\mathbf{x}), y)$. Note also that the squared loss is bounded in our setting because of our assumption that both the ground truth and noise are bounded.

There are two basic problems with the above - the distribution is unknown, so the objective above cannot be evaluated, and that the class of DC functions is far too rich, and so the problem is strongly underdetermined. In addition, directly optimising over all DC functions is an intractable problem.

To begin with, we reduce the problem by instead finding the values that a best fit DC function must take at the datapoint \mathbf{x}_i , and then fitting a PLDC functions with convex parts that are max-affine over precisely n linear functionals on this. This has the significant advantage of reducing the optimisation problem to a set of convex constraints. Quantitatively, this step is justified in §4, which argues that the error induced by this approximation via PLDC functions is dominated by the intrinsic risk of the problem.

To handle the lack of knowledge of the distribution, we resort to uniform generalisations bounds in the literature. Our approach to relies on the following result, which mildly generalises the bounds of Bartlett & Mendelson (2002), and uniformly controls the generalisation error of an empirical estimate of the expectation (specialised to our context):

Theorem 1. *Let $\{(\mathbf{x}_i, y_i)\}_{i \in [n]}$, be i.i.d. data, with n assumed to be even. Let the empirical maximum discrepancy of the class \mathcal{DC}_L , be defined as,*

$$\hat{D}_n(\mathcal{DC}_L) \triangleq \sup_{f \in \mathcal{DC}_L} \frac{2}{n} \left(\sum_{i=1}^{n/2} f(\mathbf{x}_i) - \sum_{i=n/2+1}^n f(\mathbf{x}_i) \right).$$

With probability $\geq 1 - \delta$ over the data, it holds holds uniformly over all $f \in \mathcal{DC} \cap \{\|f\| \leq M\}$ that

$$\begin{aligned} & \left| \mathbb{E}[(f(\mathbf{x}) - y)^2] - \frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 \right| \\ & \leq 12M\hat{D}_n(\mathcal{DC}_{2\|f\|+2}) \\ & \quad + 45M^2 \sqrt{\frac{C \max(2, \log_2 \|f\|) + \ln(1/\delta)}{n}}, \end{aligned} \quad (4)$$

where C is a constant independent of f, M, R .

The above statement essentially arises from a doubling trick over a Rademacher complexity based bound for a fixed $\|f\|$. The broad idea is that since $\mathcal{DC}_L \subset \mathcal{DC}_{L'}$ for $L \leq L'$, we can separately develop Rademacher complexity based bounds over L of the form 2^j , each having the more stringent high-probability requirement $\delta_j = \delta 2^{-j}$. A union bound over these then extends these bounds to all of $\mathcal{DC} = \bigcup_{j \geq 1} \mathcal{DC}_{2^j}$, and for a particular f , the bound for $j = \lceil \log_2 \|f\| \rceil$ can be used. See §C for details.

Optimising the empirical upper bound on $\mathbb{E}[(f(X) - Y)^2]$ implied by the above directly leads to a structural risk minimization over the choice of L . The crucial ingredient in the practicality of this is that for DC functions,

$\hat{D}_n(\mathcal{DC}_L) = L\hat{D}_n(\mathcal{DC}_1)$, and further, $\hat{D}_n(\mathcal{DC}_1)$ can be computed via linear programming. Thus, the term \hat{D}_n above serves as a natural, efficiently computable penalty function, and acts exactly as a regularisation on the DC seminorm.

3.1. Computing empirical maximum discrepancy.

Throughout this and the following sections, we use \hat{y}_i to denote $\hat{f}(\mathbf{x}_i)$, where \hat{f} is the estimated function.

The principle construction relies on the characterisation of Proposition 2.

Theorem 2. *Given data $\{(\mathbf{x}_i, y_i)\}$, the following convex program computes $\hat{D}_n(\mathcal{DC}_L)$*

$$\begin{aligned} & \max_{\hat{y}_i, z_i, \mathbf{a}_i, \mathbf{b}_i} \frac{2}{n} \left(\sum_{i=1}^{n/2} \hat{y}_i - \sum_{i=n/2+1}^n \hat{y}_i \right) \quad (5) \\ \text{s.t. } & \begin{cases} \hat{y}_i - \hat{y}_j + z_i - z_j \geq \langle \mathbf{a}_j, \mathbf{x}_i - \mathbf{x}_j \rangle & i, j \in [n] \quad (\text{i}) \\ z_i - z_j \geq \langle \mathbf{b}_j, \mathbf{x}_i - \mathbf{x}_j \rangle & i, j \in [n] \quad (\text{ii}) \\ \|\mathbf{a}_i\|_1 + \|\mathbf{b}_i\|_1 \leq L & i \in [n] \quad (\text{iii}) \end{cases} \end{aligned}$$

Further, $\hat{D}_n(\mathcal{DC}_L) = L\hat{D}_n(\mathcal{DC}_1)$.

Proof. By Proposition 2, conditions (i) and (ii) are necessary and sufficient for the existence of a DC function that takes the values \hat{y}_i at \mathbf{x}_i . Thus, these first two conditions allow exploration over all values a DC function can take. Further, by the second part of Proposition 2 if a DC function interpolates this data, then so does a PLDC function, with \mathbf{a}_i and \mathbf{b}_i serving as the gradients of the max-affine parts of the function. Thus, by Proposition 1, the condition (iii) is necessary and sufficient for the DC function implied by the first to conditions to have seminorm bounded by L . It follows that the conditions allow exploration of all values a \mathcal{DC}_L function may take at the given $\{\mathbf{x}_i\}$, at which point the claim follows.

Now, notice that if we multiply each of the decision variables in the above program by L , the value of the program is multiplied by a factor of L , while the constraints (i), (ii) remain unchanged. On the other hand, the constraint (iii) is modified to $\|\mathbf{a}_i\| + \|\mathbf{b}_i\| \leq 1$. Thus, the resulting program is L times the program computing $\hat{D}_n(\mathcal{DC}_1)$, ergo $\hat{D}_n(\mathcal{DC}_L) = L\hat{D}_n(\mathcal{DC}_1)$. \square

3.2. Structural Risk Minimisation

To perform the SRM, we again utilize the structural result of Proposition 1 to determine the values that the optimal estimate takes at each of the \mathbf{x}_i . The choice of the values is penalised by the seminorm as λL , where $\lambda = 24M\hat{D}_n(\mathcal{DC}_1)$, which may be computed using the program (5). Note that the logarithmic term in the generalisation bound (4) is typically small, and is thus omitted in the following. This also

has the benefit of rendering the objective function convex, as it avoids the $\log L$ term that would instead enter. If desired, a convex upper bound may be obtained for the same, for instance by noting that $\sqrt{\max(1, \log_2 \|f\|)} \leq 1 + \|f\|$. This has the effect of bumping up the value of λ required by $O(M^2/\sqrt{n})$. However, theoretically this term is strongly dominated by the \hat{D}_n (see §4), while practically even the value $\lambda = \hat{D}_n(\text{DC}_1)$ tends to produce overly smoothed solutions (see §5).

With the appropriate choice of λ , this yields the following convex optimisation problem,

$$\begin{aligned} \min_{\hat{y}_i, z_i, \mathbf{a}_i, \mathbf{b}_i, L} \sum_{i=1}^n (\hat{y}_i - y_i)^2 + \lambda L \quad (6) \\ \text{s.t.} \begin{cases} \hat{y}_i - \hat{y}_j + z_i - z_j \geq \langle \mathbf{a}_j, \mathbf{x}_i - \mathbf{x}_j \rangle & i, j \in [n] \\ z_i - z_j \geq \langle \mathbf{b}_j, \mathbf{x}_i - \mathbf{x}_j \rangle & i, j \in [n] \\ \|\mathbf{a}_i\|_1 + \|\mathbf{b}_i\|_1 \leq L & i \in [n] \end{cases} \end{aligned}$$

Once again, in the above constraints, the first two are necessary and sufficient to ensure that a DC function taking the values \hat{y}_i at \mathbf{x}_i exists, with the vectors $\mathbf{a}_i, \mathbf{b}_i$ serving as witnesses for the subgradients of the convex parts of such a function at the \mathbf{x}_i , and the third constraint enforces that the function has seminorm at most L . Notice that the third condition effectively imposes ℓ_1 -regularisation on the weight vectors $\mathbf{a}_i, \mathbf{b}_i$. This causes these weights to be sparse.

Finally, we may use the witnessing values \hat{y}_i, z_i and $\mathbf{a}_i, \mathbf{b}_i$, to construct, in the manner of Proposition 2, the following PLDC function, which interpolates the values \hat{y}_i to the entirety of the domain. Notice that since this function has the same loss as any DC function that satisfies $f(\mathbf{x}_i) = y_i$, this construction enjoys risk bounds constructed above.

$$\begin{aligned} \hat{f}(\mathbf{x}) \triangleq \max_{i \in [n]} \langle \mathbf{a}_i, \mathbf{x} - \mathbf{x}_i \rangle + \hat{y}_i + z_i \\ - \max_{i \in [n]} \langle \mathbf{b}_i, \mathbf{x} - \mathbf{x}_i \rangle + z_i \quad (7) \end{aligned}$$

3.3. Computational Complexity

Training First, we note that we may replace the constraints on the 1-norms of the vectors $\mathbf{a}_i, \mathbf{b}_i$ in the above by linear constraints via the standard trick of writing the positive and negative parts of each of their components separately. Overall, this renders the program (5) as a linear programming problem over $2n(2d+1)$ variables, and with n^2 non-trivial constraints. Note that in our setting, one typically requires that $n \geq d$ - that one has more samples than dimensions. Via interior point methods, this problem may be solved in $O(n^5)$ time.

For the least squares loss $\ell(y, \hat{y}) = (y - \hat{y})^2$, the second program (6) is a convex quadratic program when the 1-norm constraints are rewritten as above. The decision variables are

the same as the first problem, with the addition of the single L variable, and the constraints remain identical. Again, via interior point methods, these programs can be solved in $O(d^2 n^5)$ time (see Ch. 11 of [Boyd & Vandenberghe \(2004\)](#)). The latter term dominates this complexity analysis. We note that in practice, these problems can be solved significantly faster than the above bounds suggest.

Speeding up training via a GPU implementation A Parallel solver for the SRM in program (6) is given in Algorithm (1) in the Appx E via the ADMM method ([Boyd et al., 2011](#)). This algorithm can be implemented with GPU's where each iteration can be distributed to n^2 parallel cores to further increase the training speed. A python implementation is given in our GitHub repository 2. We note that a similar algorithm for that of Lipschitz convex regression is provided in ([Balázs, 2016; Mazumder et al., 2019](#)) however not all the ADMM blocks are solved in closed form and require additional optimization in each iteration.

Prediction By appending a 1 to the input, and the constants $y_i + z_i - \langle \mathbf{a}_i, \mathbf{x}_i \rangle$ and $z_i - \langle \mathbf{b}_i, \mathbf{x}_i \rangle$ to \mathbf{a}_i and \mathbf{b}_i , we can reduce the inference time problem to solving two maximum inner product search problems over n vectors in \mathbb{R}^{d+1} . This is a well developed and fast primitive, e.g. [Shrivastava & Li \(2014\)](#) provide a locality sensitive hashing based scheme that solves the problem in time that is sublinear in n .

4. Analysis

We note that this analysis makes extensive use of the work of [Balázs et al. \(2015\); Balázs \(2016\)](#) on convex and max-affine regression, with emphasis on the latter thesis, which contains certain refinements of the former paper.

In this section, we assume that the true model $y = f(\mathbf{x}) + \varepsilon_i$ holds for a f that is a DC function, and that we have explicit knowledge that $\|f\| \leq L$. Also recall our assumption that the distribution is supported on a set that is contained in the ℓ_∞ ball of radius R . We begin with a few preliminaries

A lower bound on risk The minimax risk of estimation under the squared loss is $\Omega(n^{-2/d+2})$. This follows by setting $p = 1$ (i.e., Lipschitzness) in Theorem 3.2 of ([Györfi et al., 2006](#)), which can be done since the standard constructions of obstructions to estimators used in showing such lower bounds all have regularly varying derivatives, and thus finite DC seminorms.

PLDC solutions are not lossy Lemma 5.2 of ([Balázs, 2016](#)) argues that for every convex L -Lipschitz functions ϕ with Lipschitz constant¹, $\sup_x \|\nabla_* \phi\|_1$ bounded by L , there

¹[Balázs \(2016\)](#) presents an argument with the 2-norm of subgradients bounded, but this can be easily modified to the case of bounded 1-norm under bounds on $\|\mathbf{x}\|_\infty$.

exists a Lipschitz max-affine function ϕ_{PL} with maximum over n pieces such that $\|\phi - \phi_{\text{PL}}\| \leq 36LRn^{-2/d}$. Recall that PLDC functions can be represented as differences of max-affine functions, and DC functions as differences of convex functions. Since the DC seminorm controls the 1-norm of the subgradients, which dominates the 2-norm, it follows that for every \mathcal{DC}_L function f , there exists a pl- \mathcal{DC}_{2L} function f_{PL} with $\|f - f_{\text{PL}}\|_{\infty} = O(n^{-2/d})$. Note that the resulting excess risk in the squared error due to using PLDC functions can thus be bounded as $O(n^{-4/d})$, which is $o(n^{-2/d+2})$, i.e., it is dominated by minimax risk.

4.1. Statistical risk

The bound (4) provides a instance specific control on the generalisation error of an estimate via the empirical maximum discrepancy \hat{D}_n . This section is devoted to providing generic bounds on the same under the assumption of i.i.d. sampled data. We adapt the analysis of (Balázs, 2016) for convex regression in order to do so. The principal result is as follows.

Theorem 3. *For distributions supported on a compact domain $\Omega \subset \{\|\mathbf{x}\|_{\infty} \leq R\}$, with $n \geq d$, it holds that*

$$\hat{D}_n(\mathcal{DC}_L) \leq 60LR \left(\frac{d}{n}\right)^{2/d+4} \left(1 + 2\frac{\log(n/d)}{d+4}\right).$$

Further, if the ground truth $f \in \mathcal{DC}_L$, then with probability at least $1 - \delta$ over the data, the estimator \hat{f} of (7) satisfies

$$\begin{aligned} \mathbb{E}[|Y - \hat{f}(\mathbf{x})|^2] &\leq \mathbb{E}[|Y - f(\mathbf{x})|^2] \\ &+ O((n/d)^{-2/d+4} \log(n/d)) + O(\sqrt{\log(1/\delta)/n}). \end{aligned}$$

Proof. Assume $f \in \mathcal{DC}_L$. Note that for any convex representation $f = \psi_1 - \psi_2$, we may instead construct a representation $f = \phi_1 - \phi_2 + c$ for a constant c so that the resulting convex function ϕ_1 and ϕ_2 are uniformly bounded on the domain - indeed, this may be done by setting $c = f(0)$, and $\phi_k = \psi_k - \psi_k(0)$. The ϕ s retain the bound on their Lipschitz constants, and thus are uniformly bounded by LR over the domain. Thus, we may represent the class of \mathcal{DC}_L functions as the sum of a constant, and a \mathcal{DC}_L function whose convex parts are bounded. Call the latter class $\mathcal{DC}_{L,0}$. Importantly, since the constants are cancelled in the computation of empirical maximum discrepancy, we can observe that $\hat{D}_n(\mathcal{DC}_L) = \hat{D}_n(\mathcal{DC}_{L,0})$.

The principle advantage of the above exercise is that the empirical discrepancy for DC functions with bounded convex parts can be controlled via the metric entropy of bounded Lipschitz convex functions, which have been extensively analysed by Balázs (2016). This is summarised in the following pair of lemmata. The first argues that the discrepancy of $\mathcal{DC}_{L,0}$ functions is controlled by that of bounded Lipschitz convex functions.

Lemma 1. *Let $\mathcal{C}_{L,LR}$ be the set of convex functions that are L -Lipschitz and bounded by LR . Then $\hat{D}_n(\mathcal{DC}_{L,0}) \leq 2\hat{D}_n(\mathcal{C}_{L,LR})$.*

The proof of the above is left to Appx C. The second lemma, due to Dudley, is a generic method to allow control on the discrepancy. We state this for $\mathcal{C}_{L,LR}$

Lemma 2. *Let $\mathcal{H}_{\infty}(\mathcal{C}_{L,LR}, \epsilon)$ be the metric entropy of the class $\mathcal{C}_{L,LR}$ under the sup-metric $d(f, g) = \|f - g\|_{\infty}$. Then the empirical maximum discrepancy is bounded as*

$$\hat{D}_n(\mathcal{C}_{L,LR}) \leq \inf_{\epsilon > 0} \left(\epsilon + LR \sqrt{2 \frac{\mathcal{H}_{\infty}(\mathcal{C}_{L,LR}, \epsilon)}{n}} \right).$$

Finally, we invoke the control on metric entropy of bounded Lipschitz convex functions provided by (Balázs et al., 2015; Balázs, 2016)

Theorem (Balázs et al., 2015; Balázs, 2016). *For $\epsilon \in (0, 60LR]$,*

$$\mathcal{H}_{\infty}(\mathcal{C}_{L,LR}, \epsilon) \leq 3d \left(\frac{40LR}{\epsilon}\right)^{d/2} \log\left(\frac{60LR}{\epsilon}\right).$$

Using the above in the bound of Lemma 2, and choosing $\epsilon = (60LR)(d/n)^{2/d+4}$ yields the claim.

Control on the excess risk of the estimator follows readily. For any $\lambda \geq 24M\hat{D}_n(\mathcal{DC}_L)$, we have have, with high probability,

$$\begin{aligned} &\mathbb{E}[(\hat{f} - Y)^2] \\ &\leq \hat{\mathbb{E}}[(\hat{f} - Y)^2] + \lambda \|\hat{f}\| + 2\lambda + O(1/\sqrt{n}) \\ &\leq \hat{\mathbb{E}}[(f - Y)^2] + \lambda \|f\| + 2\lambda + O(1/\sqrt{n}) \\ &\leq \mathbb{E}[(f - Y)^2] + 2\lambda \|f\| + 2\lambda + O(1/\sqrt{n}), \\ &\leq \mathbb{E}[(f - Y)^2] + 2\lambda(L + 1) + O(1/\sqrt{n}) \end{aligned}$$

where the first and last inequalities utilise (4), while the second inequality is because \hat{f} is the SRM solution. The claim follows on incorporating the bound on \hat{D}_n developed above, and since we choose λ proportional to the same. \square

On adaptivity Notice that the argument for showing the excess risk bound proceeds by controlling \hat{D}_n . This allows the argument to adapt to the dimensionality of the data. Indeed, if the true law is supported on some low dimensional manifold, then the empirical discrepancy, which only depends on the observed data, grows only with this lower dimension. More formally, due to the empirical discrepancy being an empirical object, we can replace the metric entropy control over DC functions in \mathbb{R}^d by metric entropy of DC functions supported on the manifold in which the data lies, which in turn grows only with the (doubling) dimension of this manifold.

On suboptimality Comparing to the lower bound, we find that the above rate is close to being minimax, although the (multiplicative) gap is about n^{4/d^2} and diverges polynomially with n . In analogy with the observations of Baláz (2016) for max-affine regression, we suspect that this statistical suboptimality can be ameliorated by restricting the PLDC estimate to have some (precisely chosen) $K_n < n$ hyperplanes instead of the full n . However, as discussed in §1.1, such restrictions lead to increase in the computational costs of training, and thus, we do not pursue them here.

5. Experiments

In this section we apply our method to both synthetic and real datasets for regression and multi-class classification. The datasets were chosen to fit in the regime of $n \leq 10^3, d \leq 10^2$ as described in the introduction. All results are averaged over 100 runs and are reported with the 95% confidence interval.

For the DC function fitting procedure, we note that the theoretical value for the regularization weight tends to over-smooth the estimators. This behaviour is expected since the bound (4) is designed for the worst case. For all the subsequent experiments we make two modifications - since none of the values in the datasets observed are very large, we simply set $12M = 1$, and further, we choose the weight, i.e. λ in (6), by cross validation over the set $2^{-j} \hat{D}_n(\mathcal{DC}_1)$ for $j \in [-8 : 1]$. Fig. 1 presents both these estimates in a simple setting where one can visually observe the improved fit. Note that this tuning is still minimal - the empirical discrepancy of \mathcal{DC}_1 fixes a rough upper bound on the λ necessary, and we explore only a few different scales.

For the **regression task** we use the L_1 empirical loss in our algorithm, instead of L_2 . That is, the objective in (6) is replaced by $\sum |y_i - \hat{y}_i|$. This change allows us to implement the fitting program as a linear programming problem and significantly speeds up computation. However, in the following we will only report the L_2 error of the solutions obtained this way. We compare our method to a multi-layer perceptron (neural network) with variable number of hidden layers chosen from 1 : 10 by 5-fold cross validation, Multivariate Adaptive Regression Splines (MARS) and K -nearest neighbour (K -NN) where Best value of K was chosen by 5-fold cross validation from 1 : 10.

For the multi-class **classification task** we adopt the multi-class hinge loss to train our model, i.e. the loss

$$\sum_i \sum_{j \neq y_i}^m \max(f_j(\mathbf{x}_i) - f_{y_i}(\mathbf{x}_i) + 1, 0),$$

where m is the number of classes and f_j 's are DC functions. We compare with the same MLP as above but trained with the cross entropy Loss, KNN and a one versus all SVM.




Figure 1. Top A two dimensional function along with the sampled points used for estimating the function; Bottom learned DC function via L1 regression using only $\lambda = 2\hat{D}_n$ (left); using cross validation over λ (right).

In both cases, we have used MATLAB Statistics and Machine learning Library for their implementation of MLP, KNN and SVM. For MARS we used the open source implementation in ARESLab toolbox implemented in MATLAB. Our code along with the other algorithms is available in our GitHub repository².

In each of the following tasks, we observe that our method performs competitively to all considered alternatives in almost every dataset, and often outperforms them, across the variation in dimensionality and dataset sizes.

Regression on Synthetic Data We generated data from the function,

$$y = f(\mathbf{x}) + \varepsilon$$

$$f(\mathbf{x}) = \sin\left(\frac{\pi}{\sqrt{d}} \sum_{j=1}^d x_j\right) + \left(\frac{1}{\sqrt{d}} \sum_{j=1}^d x_j\right)^2,$$

where the \mathbf{x} is sampled from a standard Gaussian, while ε is a centred Gaussian noise with standard deviation of 0.25.

We generate 50 points for training. For testing we estimate the Mean Squared Error based on a test set of 5000 points without the added noise. We normalize the MSE by variance of the values of test data and multiply by 100. Results are presented in Fig. 2. Observe that our algorithm consistently outperforms the competing methods, especially as we increase the dimension of the data. Furthermore our algorithm has lower error variance across the runs.

²<https://github.com/Siahkamari/Piecewise-linear-regression-via-a-difference-of-convex-functions>




Figure 2. Mean Squared Error in a regression task vs dimension of the data in the synthetic experiments. Note that both the value and the size of the error bars are consistently better than the competing methods

Regression on Real Data We apply the stated methods to various moderately sized regression datasets that are available in the MATLAB statistical machine learning library. The results are presented in Fig. 3.

In the plot, the datasets are arranged so that the dimension increases from left to right. We observe that we do comparably to the other methods for some datasets and outperform in others. See Appx. D for a description of each of the datasets studied.




Figure 3. Normalized Mean Squared Error in regression tasks.

Multi-class classification We used popular UCI classification datasets for testing our classification algorithm. We repeated the experiments 100 times. We present the mean and 95% C.I.s on a 2-fold random cross validation set, in Fig. 4. We observe to perform comparably to other algorithms on some datasets and outperform in others.




Figure 4. Miss-Classification on UCI data sets.

6. Discussion

The paper proposes an algorithm to learn piecewise linear functions using difference of max-affine functions. Our model results in linear or convex programs which can be solved efficiently even in high dimensions. We have shown several theoretical results on expressivity of our model, as well as its statistical properties, including good risk decay and adaptivity to low dimensional structure, and have demonstrated practicality of our resulting model under regression and classification tasks.

Wider context Non-parametric methods are most often utilised in settings with limited data in moderate dimensions. Within this context, along with strong accuracy, it is often desired that the method is fast, and is interpretable, especially in relatively large dimensions.

In settings with large datasets, accuracy is relatively easy to achieve, and speed at inference is more important. In low dimensional settings, this makes methods such as MARS attractive, due to their fast inference time, while in high-dimensions MLPs are practically preferred. In settings with low dimensionality and small datasets, interpretability and speed take a backseat due to the small number of features, while accuracy become the critical requirement, promoting use of kernelised or nearest neighbour methods.

On the other hand, for small datasets in moderate to high dimensions, interpretability gains an increased emphasis. Within this context, our method results in a piecewise linear fit for which it is easy to characterise the locally important features, and further does so with relatively sparse weights via the ℓ_1 regularisation on a_i, b_i . Further, since the sub-optimality of our statistical risk is controlled as $n^{O(1/d^2)}$, as the dimension increases, our method gets closer to the optimal accuracy. This thus represents the natural niche where application of DC function regression is appropriate.

References

- Bačák, M. and Borwein, J. M. On difference convexity of locally lipschitz functions. *Optimization*, 60(8-9):961–978, 2011.
- Balázs, G. *Convex Regression: Theory, Practice, and Applications*. PhD thesis, University of Alberta, 2016.
- Balázs, G., György, A., and Szepesvári, C. Near-optimal max-affine estimators for convex regression. In *AISTATS*, 2015.
- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Boyd, S., Parikh, N., and Chu, E. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- Cui, Y., Pang, J.-S., and Sen, B. Composite difference-max programs for modern statistical estimation problems. *SIAM Journal on Optimization*, 28(4):3344–3374, 2018.
- Friedman, J. H. et al. Multivariate adaptive regression splines. *The annals of statistics*, 19(1):1–67, 1991.
- Gasso, G., Rakotomamonjy, A., and Canu, S. Recovering sparse signals with a certain family of nonconvex penalties and dc programming. *IEEE Transactions on Signal Processing*, 57(12):4686–4698, 2009.
- Ghosh, A., Pananjady, A., Guntuboyina, A., and Ramchandran, K. Max-affine regression: Provable, tractable, and near-optimal statistical estimation. *arXiv preprint arXiv:1906.09255*, 2019.
- Györfi, L., Kohler, M., Krzyżak, A., and Walk, H. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006.
- Hannah, L. A. and Dunson, D. B. Multivariate convex regression with adaptive partitioning. *The Journal of Machine Learning Research*, 14(1):3261–3294, 2013.
- Hartman, P. et al. On functions representable as a difference of convex functions. *Pacific Journal of Mathematics*, 9(3):707–713, 1959.
- Hildreth, C. Point estimates of ordinates of concave functions. *Journal of the American Statistical Association*, 49(267):598–619, 1954.
- Holloway, C. A. On the estimation of convex functions. *Operations Research*, 27(2):401–407, 1979.
- Holmes, C. and Mallick, B. Bayesian regression with multivariate linear splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(1):3–17, 2001.
- Horst, R. and Thoai, N. V. Dc programming: overview. *Journal of Optimization Theory and Applications*, 103(1):1–43, 1999.
- Kripfganz, A. and Schulze, R. Piecewise affine functions as a difference of two convex functions. *Optimization*, 18(1):23–29, 1987.
- Le Thi, H. A., Le, H. M., Dinh, T. P., et al. A dc programming approach for feature selection in support vector machines learning. *Advances in Data Analysis and Classification*, 2(3):259–278, 2008.
- Magnani, A. and Boyd, S. P. Convex piecewise-linear fitting. *Optimization and Engineering*, 10(1):1–17, 2009.
- Mazumder, R., Choudhury, A., Iyengar, G., and Sen, B. A computational framework for multivariate convex regression and its variants. *Journal of the American Statistical Association*, 114(525):318–331, 2019.
- Ovchinnikov, S. Max-min representation of piecewise linear functions. *Contributions to Algebra and Geometry*, 43(1):297–302, 2002.
- Piot, B., Geist, M., and Pietquin, O. Difference of convex functions programming for reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2519–2527, 2014.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shrivastava, A. and Li, P. Asymmetric lsh (alsh) for sub-linear time maximum inner product search (mips). In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2321–2329. Curran Associates, Inc., 2014.
- Toriello, A. and Vielma, J. P. Fitting piecewise linear continuous functions. *European Journal of Operational Research*, 219(1):86–95, 2012.
- Wahba, G. *Spline models for observational data*, volume 59. Siam, 1990.
- Yuille, A. L. and Rangarajan, A. The concave-convex procedure (cccp). In *Advances in neural information processing systems*, pp. 1033–1040, 2002.