

L_p Isotonic Regression Algorithms Using an L_0 Approach

Quentin F. Stout

qstout@umich.edu

www.eecs.umich.edu/~qstout/

Abstract: Significant advances in maximum flow algorithms have changed the relative performance of various approaches to isotonic regression. If the transitive closure is given then the standard approach used for L_0 (Hamming distance) isotonic regression (finding anti-chains in the transitive closure of the violator dag), combined with new flow algorithms, gives an L_1 isotonic regression algorithm taking $\tilde{\Theta}(n^2)$ time on a graph of n vertices. The previous fastest was $\Theta(n^3)$. For points in d -dimensional space with coordinate-wise ordering, $d \geq 3$, L_1 regression can be found in $o(n^{\frac{3}{2}})$ time, improving on the previous best of $\tilde{\Theta}(n^2 \log^d n)$. Similar results are obtained for L_p approximations, $1 < p < \infty$, and for exact L_2 regression when the values and weights are restricted.

Keywords: L_1 isotonic regression, L_p monotonic regression, maximum flow, violator graph, Hamming distance, multidimensional coordinate-wise ordering

1 Introduction

There are many scenarios when one expects that an attribute of objects increases as a function of other attributes. For example, the expected maximum grade level achieved by a child is likely to be an increasing function of their mother’s maximum grade level and father’s maximum grade level. No assumptions are made about the relationship between a lower mother’s grade level and higher father’s grade level versus a higher mother’s grade level and lower father’s. More generally, let V be a set with a partial order \prec , and let f be a function on V . f is *isotonic* iff for all $u, v \in V$, if $u \prec v$ then $f(u) \leq f(v)$. Thus we expect that the child’s maximum grade level is an isotonic function of its mother’s maximum grade level and its father’s, but often the data collected isn’t quite isotonic, so adjustments to it need to be made.

Given arbitrary real-valued functions f, g , and nonnegative weight function w on V and a p where $p = 0$ or $1 \leq p \leq \infty$, their weighted L_p distance, denoted $\|f - g\|_p$, is

$$\begin{aligned} \sum_{v \in V} w(v) \cdot \mathbf{1}(f(v) \neq g(v)) & \quad p = 0 \\ (\sum_{v \in V} w(v) \cdot |f(v) - g(v)|^p)^{1/p} & \quad 1 \leq p < \infty \\ \max_{v \in V} w(v) \cdot |f(v) - g(v)| & \quad p = \infty \end{aligned}$$

L_0 is not a true norm and is also known as the *Hamming distance*, 0-1 loss, or Kronecker delta loss. *Unweighted function* means $w(v) = 1$ for all $v \in V$. f' is an L_p *isotonic regression* of f iff it minimizes $\|f - g\|_p$ among all isotonic functions g . Isotonic regressions are unique for $1 < p < \infty$, but not necessarily for $p = 0, 1, \infty$.

Many different terms are applied to various versions of this. For example, in mathematical analyses often monotonic is used instead of isotonic. L_0 isotonic regression is sometimes called monotonic relabeling, and the distance is also known as Hamming distance, 0-1 distance or 0-1 loss. L_1 regression is sometimes called median regression, the least absolute deviations, or sum of deviations error. The L_p regression error of the optimal isotonic regression is sometimes called the “distance to monotonicity”, where this term is primarily used for the L_0 distance but is sometimes used for other L_p . L_2 is the (sometimes unstated) most common

norm for isotonic regression. A large number of other terms apply as well. Isotonic regressions are a class of non-parametric shape-constrained regressions.

The paper concentrates on algorithms, not on any specific application. Some relevant L_0 applications appear in [5, 8, 11, 12, 17, 19, 27, 28, 30, 31, 32, 40], and for L_1 and L_2 the applications are far too numerous to list.

Usually we assume the partial order is given via a dag which is connected (for disconnected ones the algorithms would be applied on each piece separately), having n vertices and m directed edges. In Section 4 the edges of the dag will not initially be explicit. We assume the weights and values are integers in $[0, U]$, though in some cases this assumption isn't required.

Efficient algorithms for determining isotonic regressions depend on the metric and the underlying dag. L_∞ algorithms are quite different than those for L_p , $p < \infty$, and won't be discussed. For linear orders there are efficient algorithms based on a left-right sweep algorithm known as PAVA, pool adjacent violators algorithm (violators will be discussed shortly). PAVA has been used for decades, resulting in algorithms taking $\Theta(n \log n)$ time for L_1 and $\Theta(n)$ for L_2 . The basic ideas can be extended to directed trees. For 2-dimensional grids the fastest algorithms for L_1 and L_2 are based on dynamic programming [34, 35]. For more complex dags the complexity grows significantly.

Here we concentrate on algorithms for the general case, applicable to any dag, though Section 4 discusses an important special class of dags. The fastest known algorithms utilize variations on maximum flows and linear programming. Those for L_0 use minimum flow algorithms and take $\Theta(n^3)$ time [12, 30, 31]; for L_1 a linear programming approach is used, taking $\Theta(nm + n^2 \log n)$ time [3]; and for L_2 a maximum flow approach is used, taking $\Theta(nm \log(\frac{n^2}{m}))$ time [18]. There do not seem to be any published L_p algorithms for other values of $p < \infty$, other than approximation algorithms [2, 9, 22, 36, 38].

Recent advances in maximum flow algorithms change some of the relative timing of these approaches, though under the assumption that the function values and weights are integers. The values and weights will determine edge capacities in a flow setting. If U is the maximum integer edge capacity then the BLLSSSW maximum flow algorithm [6] takes $\tilde{O}(m + n^{\frac{3}{2}} \log U)$ time with high probability, and the Gao-Liu-Peng algorithm [16] takes $\tilde{O}(m^{\frac{3}{2} - \frac{1}{328}} \log U)$ time. Except for some of the approximation algorithms, the previous results were worst-case times, but here we consider the BLLSSSW algorithm since it achieves a significantly faster expected time with high probability. There is intense work in this area, and the wikipedia page [41] is quickly updated when faster flow algorithms are found.

Throughout we make the standard assumption that U grows at most polynomially with n , i.e., that the algorithm is pseudo-polynomial. Thus the $\log U$ terms in the time analyses of the BLLSSSW and GLP algorithms become $O(\log n)$, and disappear whenever \tilde{O} is used. Our results are stated in terms of general flow algorithms, and then specific times are given when one of the above flow algorithms is used, though other flow algorithms could be used as well. We use $\mathcal{F}(n, m, U)$ to denote the time of the unspecified flow algorithm for a graph of n vertices, m edges, and maximum edge capacity U . For a dag G we use $\mathcal{V}(G)$ to indicate the time to find a violator dag (defined in Section 2). Here too this is stated in general terms, and then specific algorithms are referred to.

The overall structure of the approach used is:

- Find L_0 monotonic regression with two labels via a flow algorithm on a violator dag. Section 2
- View L_0 monotonic regression when there are only two labels as L_1 monotonic regression when the only values are $\{0, 1\}$. Section 3.1
- Use $\{0, 1\}$ -valued L_1 monotonic regression to determine L_1 monotonic regression of arbitrary values. Section 3.1

- Use $\{0,1\}$ -valued L_1 regression to find approximations, and in some cases exact values, for L_p monotonic regression, $1 < p < \infty$. In particular, exact regressions are found for L_2 . Section 3.2

We utilize this approach for arbitrary dags, and in Section 4 we exploit special properties of the partial ordering of points in d -dimensional space with component-wise ordering. Component-wise ordering (multidimensional ordering, domination) is of significant interest in many applications.

Section 5 has final comments and tables comparing the new algorithms with the fastest prior ones.

2 Background and L_0 Isotonic Regression

Many isotonic regression algorithms have descriptions in terms of violating pairs and violator dags. Given a dag $G = (V, E)$ and a real-valued function f on V , vertices $u, v \in V$ are a *violating pair* if $u \prec v$ and $f(u) > f(v)$, i.e., they violate the isotonic requirement. Using $u \prec_v v$ to denote that u and v are a violating pair, it is easy to see that \prec_v defines a partial order on V . A dag $\hat{G} = (\hat{V}, \hat{E})$, is a *violator dag* of G iff $V \subset \hat{V}$ and for any $u, v \in V$ there is a path from u to v in \hat{G} iff $u \prec_v v$. There may be many violator dags for a given G and f . For example, if \hat{G} has $V = \hat{V}$ and there is an edge $(u, v) \in \hat{E}$ iff there is a path from u to v in G , then \hat{G} is the transitive closure of the violating pairs. The edges of the smallest subgraph of the transitive closure which is still a violator dag form the transitive reduction. In Section 4 we show that sometimes it is useful to have \hat{V} contain additional vertices not in V .

Let $\mathcal{V}(G)$ denote the time to create a violator dag for dag G . We leave the algorithm to create the violator dag unspecified, and then show the time for specific algorithms. This generality will only be used in Section 4, though might be useful in other circumstances as well. We usually assume that the standard approach is used, creating the transitive closure of G and then removing edges that are not a violating pair. In this case the time to create a violator dag for a dag of n vertices and m edges is $\mathcal{V}(n, m) = O(\min\{nm, n^\omega\})$, where ω is such that matrix multiplication can be computed in $O(n^\omega)$ time. For arbitrary dags of n vertices the best possible time for finding the transitive closure or the transitive reduction is linear in the time for Boolean matrix multiplication [1, 13, 14, 26]. There are values of $\omega < 2.4$, but current algorithms achieving this are galactic. One could use a practical algorithm, such as Strassen's, for which the time is $\Theta(n^{\lg 7 \approx 2.81})$.

For L_0 , but not L_1 , the remainder of this section is a rephrasing and condensation of the Background section of [37], which is based on the work of [12, 25, 29, 32]. The fundamental algorithm for finding an L_0 isotonic regression is given in Figure 1, a slightly modified version of a figure in [37]. Note that in L_0 there is no need for a metric on the function or regression values, merely that they are linearly ordered, while for L_1 a metric is required.

The role of C in Figure 1 is that any vertices which are an anti-chain (i.e., there are no directed paths connecting any pair) in \hat{G} can be used as vertices where $f = f'$, and in fact given any set of vertices D where f is unchanged they can only be part of an isotonic function if D is an anti-chain. Thus picking an anti-chain of maximum weight yields an optimal L_0 isotonic regression, if the remaining values are chosen to create an isotonic function. That is what step 3 does: if v has predecessors in C then $f'(v)$ must be no smaller than $f'(w) = f(w)$ for any predecessor of v in C , and similar it must be no larger than $f(w)$ for any successors in C . Note that if $v \notin C$ then v must have successors or predecessors, or both, in C , since otherwise it could be added to C and increase the anti-chain. There are other possibilities for values of f' on $V \setminus C$ but they aren't needed here.

A flow graph is used to find C . This is illustrated in Figure 2, which is copied from [37]. The flow graph is straightforward to construct from the violator dag \hat{G} , replacing each vertex v with a pair v_{in} and v_{out} , with an edge from v_{in} to v_{out} . There is an edge from w_{out} to v_{in} iff (w, v) is an edge in \hat{G} . Edges of the form $(v_{\text{in}}, v_{\text{out}})$ have capacity 1 (or $w(v)$ in the weighted case) and all other edges have weight 0. A minimum flow is found where the flow on any edge is at least the edge's requirement, and this is used to find a maximum

1. Create a violator dag $\widehat{G} = (V, \widehat{E})$ of G .
2. Find an antichain C of \widehat{G} of maximum weight, where the weight of C is $\sum_{v \in C} w(v)$:
 - (a) Create a flow graph \widehat{G}_f from \widehat{G} .
 - (b) Find a minimum flow on \widehat{G}_f and use this to determine C .
3. Determine an isotonic regression f' :

$$f'(v) = \begin{cases} f(v) & v \in C \\ \max\{f(w) : w \prec v, w \in C\} & \text{if there is a predecessor of } v \text{ in } C \\ \min\{f(w) : v \prec w, w \in C\} & \text{otherwise} \end{cases}$$

Figure 1: Optimal L_0 monotonic relabeling f' of label function f on $G = (V, E)$ with order \prec (see [12, 25, 29, 32])




Figure 2: Label function on unweighted dag G , one of its violator dags \widehat{G} , flow graph \widehat{G}_f , resulting C

cut C . This is essentially the same as max-flow min-cut, but first one needs to find a feasible flow on \widehat{G} (one meeting at each edge's requirement) and then minimize it. This is done by subtracting the requirement from the flow, using these values as the capacities on the edges of \widehat{G} and finding a maximum flow on this, and then subtracting this maximal flow from the feasible flow to get the final flow.

Further explanation is in [12, 25, 32], but here the important part is that the time is dominated by the time to create \widehat{G} , and the minimum flow algorithm. Finding a feasible flow can easily be done using topological sort, creating an excess of at most nU , where U is the maximum weight on any vertex. Thus the dominating factor for the flow is the max flow adjustment, $\mathcal{F}(\widehat{n}, \widehat{m}, nU)$. The following theorem places the algorithm for L_0 in this general context, and introduces its relationship to L_1 isotonic regression for arbitrary values.

Theorem 1 Given a dag $G = (V, E)$ with weighted function (f, w) on V where the weights are natural numbers in $[0, U]$,

- i) an L_0 isotonic regression can be found in $O(\mathcal{V}(G) + \mathcal{F}(\widehat{n}, \widehat{m}, nU))$ time.
- ii) an L_1 isotonic regression can be found in $O(\mathcal{V}(G) + \mathcal{F}(\widehat{n}, \widehat{m}, nU) \log n)$ time.

Claim *i*) was proven above, and *ii*) will be shown in Section 3.1.

If the violator dag is created by determining the transitive closure then $\mathcal{V}(G) = O(\min\{nm, n^\omega\})$, and \widehat{m} may be $\binom{n}{2}$ even when the number of edges in G is $n - 1$ (consider a linear order where the values are in

decreasing order). Thus when the transitive closure approach is used the worst-case time over all dags with n vertices is $O(n^\omega + \mathcal{F}(n, \binom{n}{2}, nU))$. If the BLLSSW flow algorithm is used then for both regressions:

- a) if the violator dag constructed is via the standard transitive closure approach then the total time is $\tilde{O}(\min\{nm, n^\omega\})$, with the worst-case time over all dags of n vertices being $\Theta(n^\omega)$.
- b) if a violator dag is given then the expected time is $\tilde{O}(\hat{m} + \hat{n}^{1.5})$, with the worst-case expected time over all dags of n vertices being $\tilde{O}(n^2)$ (this assumes that there are dags of n vertices for which every violator dag has $\Omega(n^2)$ edges)

Previous results only showed that an L_0 isotonic regression can be found in $\Theta(n^3)$ time even when a violator dag was given [12, 30, 31]. An L_1 isotonic regression algorithm taking $\Theta(nm + n^2 \log n)$ time appeared in [3], which takes $\Theta(n^3)$ time in the worst-case over all dags of n vertices.

Note that if \hat{m} is very close to \hat{n} then the GLP algorithm might be faster than the BLLSSW algorithm. This fact is exploited in Section 4.

3 Partitioning

It has been noted that when there are only 2 labels L_0 isotonic regression is the same as L_1 isotonic regression when the values are in $\{0, 1\}$, and thus a way of determining an L_0 isotonic regression in this case is by using an L_1 isotonic regression algorithm. Here we reverse that. We start with the L_0 approach of building a violator dag then finding a flow and use that to construct L_1 regression where the function values are in $\{0, 1\}$. Then we use *partitioning* to find L_1 regressions for weighted functions with arbitrary values.

3.1 L_1 Isotonic Regression

In this section we prove claim *ii*) of Theorem 1.

To expand to an L_1 algorithm for arbitrary values, not just binary ones, we use an observation in [35] (and essentially in [18]). Suppose no values of f are in (a, b) , $a < b$. Create a new weighted function g on V where $g(v) = a$ if $f(v) \leq a$, and $g(v) = b$ if $f(v) \geq b$. The weight at v is the same as its original weight, $w(v)$. There is always an L_1 optimal isotonic regression where the regression values are the initial function values, so there is an optimal regression g' of g where every value of g' is a or b . Further, [35] showed that for any such g' , there is an L_1 isotonic regression f' of f , where $f'(v) \leq a$ if $g'(v) = a$, and $f'(v) \geq b$ if $g'(v) = b$. Using the original weights as the weights in the binary regression represents the increase in error that would occur if $g(v) = a$ but $g'(v) = b$. In this case $[g'(v) - f(v)] - [g(v) - f(v)]$ is $b - a$, so the error increased by $(b - a)w(v)$. The same is true if $g(v) = b$ but $g'(v) = a$. Since the change in errors is always proportional to $b - a$ we can divide by that term and minimize the increase in error by using the original weights. This is just weighted $\{0, 1\}$ -valued L_1 regression shifted, so here it is just referred to as $\{0, 1\}$ -valued L_1 regression.

Let V_a be the vertices where $g' = a$ and V_b the vertices where it is b , and let $G_a = (V_a, E_a)$ and $G_b = (V_b, E_b)$, where E_a contains those edges with both endpoints in V_a , and similarly for E_b . Edges with one endpoint in a and the other in b are discarded. No vertex of V_b precedes any vertex of V_a , and the final regression values of each set only depend on the weighted function values in that set, so the regressions on G_a and G_b can be computed independently.

The way this is applied is that the initial a and b values are the medians of the original function values (if there are 2 medians), or a is the median and b the smallest value larger than a . Then G_a is partitioned using the first quartile of the original function values (in increasing order) and G_b the third quartile. This halving process continues until, in every subgraph, the splitting values are the only two values remaining in their subgraph. There are at most $\Theta(\log n)$ iterations ($\Theta(\log \ell)$ if there are only ℓ different values).

This finishes the proof of Theorem 1 *ii*), and hence finishes the proof of the theorem. \square

3.2 Approximate and Exact L_p Regression, $1 < p < \infty$

Various approximation algorithms for L_p isotonic regression have appeared, with [2, 9, 38] considering linear dags and [22] considering arbitrary ones. Ones based on partitioning using the L_1 algorithm appear in [35]. There the basic idea is that partitioning is correct for general L_p , $1 \leq p < \infty$, so binary L_1 regression can be used in a similar fashion as for general L_1 regression. However, it is no longer true that the regression values can always be chosen to be values of the original function. E.g., for a linear order of 2 vertices, with unweighted values 1, 0, the unique L_p regression, $1 < p \leq \infty$ is 0.5, 0.5, while for L_1 the regression values can be α , α for any $0 \leq \alpha \leq 1$. Thus arbitrary binary splitting of the function values does not always produce the correct result when $p \neq 1$.

Instead one must split among the possible regression function values, not the original function values. One important aspect is that for L_p the weights need to be adjusted. Before, when using splitting values a, b , with no function values in (a, b) , a function value $f(v) \leq a$ was temporarily changed to a , with the same weight because it was proportional to the increase in error if the regression value at v was b instead of a . For L_p , $p > 1$ the change in error is not just $b - a$. In [35] it is shown that using weights based on the derivative of the error, $w(v)(a - f(v))^{p-1}$, gives the correct partitioning and can be used iteratively, just as before.

Here by approximation to within δ means that the regression values are no more than δ away from their value in an optimal isotonic regression. For $1 < p < \infty$ there is only one optimal isotonic regression. Rounding to a multiple of δ from the smallest function value, and splitting the range in half each time, gives a suitable approximation.

Theorem 2 *Let $1 < p < \infty$. Given a dag $G = (V, E)$, a weighted function (f, w) on V with integral weights in $[0, U]$, and $\delta > 0$, once a violator dag $\hat{G} = (\hat{V}, \hat{E})$ has been constructed, an isotonic function within δ of the L_p isotonic regression can be found in $O(\mathcal{F}(\hat{n}, \hat{m}, nU) \log K)$ time, where $K = (\max_{v \in V} f(v) - \min_{v \in V} f(v)) / \delta$ and where the implied constants in the O -notation depend on p . \square*

Using the BLLSSSW flow algorithm gives $\tilde{O}(\hat{m})$ time, assuming that K grows no more than polynomially. One can slightly improve the resulting regression by using the true L_p mean as the regression value for each level set. Note that the function values can be arbitrary reals, not just integers, since the flow algorithm is always solving a weighted $\{0, 1\}$ problem.

Others have looked at approximation algorithms for all L_p as well. Kyng, Rao, and Sachdeva [22] give an algorithm which produces an isotonic regression within an additive error of δ from optimal, where they used total error, not vertex-wise. It's moderately easy to adjust the δ in the above theorem to achieve their goal. Their Monte Carlo algorithm isn't based on an explicit violator dag, and takes $\tilde{\Theta}(m^{1.5} \log^2 n \log(nU/\delta))$ time to produce an appropriate approximation with probability at least $1 - 1/n$.

There is a more subtle way to do L_p partitioning, but in a restricted setting. For L_1 isotonic regression the values of each level set are a weighted median of the function values corresponding to the vertices in the set. For L_p , finding splitting values that are guaranteed to separate the level sets is difficult since it isn't known how close these values can be to each other. However, for integer function values and weights one can compute a δ such that no two level sets have values closer than δ . The details are in [35]. While the following theorem only analyzes binary values for arbitrary p , for L_2 the values can be arbitrary integers. For L_2 , one can use $\delta = 1/[\sum_{v \in V} w(v)]^2$, and no two subsets with different weighted means (their optimal regression value) can have weighted means closer than 4δ . In particular, no two function values can be closer than this.

To make the answer exact we have to be a bit more careful in partitioning. We always use values that are of the form $u + i\delta$, where u is the smallest function value. Suppose there is a subgraph $G' = (V', E')$ where we know the range of regression values will be in an interval $[c, d]$. To partition it we first choose $x = (c + d)/2$, rounded to the nearest multiple of δ . If there is no value on V' that is in $(x, x + \delta)$ then let the a for partitioning

be x and $b = x + \delta$. Otherwise, there is no function value in $(x - \delta, x)$, so let $a = x - \delta$ and $b = \delta$. Now there are no function values in (a, b) , so the comments about partitioning at the start of Section 3.1 apply.

The following result essentially appeared in [35] but in terms of partitioning using the best L_1 algorithm known at that time. Using a violator graph and the BLLSSSW flow algorithm results in a slightly faster algorithm. Once the level sets have been determined the final step is to compute their L_p mean, not just an approximation to within δ , and use that as the regression value.

Corollary 3 *For L_2 , if integral function values and weights are in $[0, U]$, then for dag $G = (V, E)$ the exact isotonic regression can be found in the time given in Theorem 2, where $\log K$ is replaced by $\log nU$.*

For L_p , $1 < p < \infty$, the exact L_p isotonic regression can be found in the same time if function values are in $\{0, 1\}$ and integer weights are in $[0, U]$. \square

4 Multidimensional and Related Orderings

A very important partial ordering is given by points in d -dimensional space with coordinate-wise ordering. Given points $x = (x_1, \dots, x_d)$, $y = (y_1, \dots, y_d)$, $x \neq y$, y dominates x iff $x_i \leq y_i$ for $1 \leq i \leq d$. Domination is also known as multi-dimensional ordering. There is no requirement that the dimensions are the same, nor even numeric, merely that each is linearly ordered. Many datasets involve such orderings, often with a large number of vertices and occasionally with a significant number of dimensions.

The points may form a grid, which has nice properties such as the fact that the number of edges is linear in the number of points (throughout we ignore constants that depend on d , such as the number of edges emanating from a point in a grid, but will include factors such as $\log^d n$). For 2-dimensional grids an L_2 isotonic regression algorithm, with no restrictions on values nor weights, appeared in [34], taking $\Theta(n^2)$ time. An unrestricted L_1 algorithm, taking $\Theta(n \log n)$ time, appeared in [35]. Both algorithms used a scan-based dynamic programming approach, with the L_1 algorithm also incorporating partitioning. By using trees to aid in the scan these can be extended to arbitrary points in 2-dimensions, adding a $\log n$ factor [35].

Violator-based algorithms do not improve these algorithms because while the original order is specified by a grid, a violator dag rarely would be. The L_2 algorithm is slower than that for L_1 because it can only guarantee that at each step it partitions into a set of 1 vertex and everything else in the other even though it scans through all vertices. However, in practice it is far faster than its worst case.

4.1 Steiner Vertices

Unfortunately the dynamic programming options in 2 dimensions do not extend to higher ones, so new approaches are needed. In [36] a concise violator dag is given for domination ordering. It incorporates Steiner vertices, which are vertices not in the original graph. The coordinates in each dimension are converted to integers, where if a dimension has x different values the coordinates are the integers in $[0, x - 1]$. They are represented as bit strings of k bits, where $k = \lceil \lg x \rceil$. Steiner coordinates corresponding to strings of length k are of the form $****$, $0***$, $1***$, $00**$, $01**$, $10**$, \dots 1110 , 1111 (for $k = 4$). Given a vertex coordinate $q = q_1 \dots q_k$ and Steiner coordinate $t = t_1 \dots t_k$, $q \preceq t$ iff $q = t$ or there is a j , $1 \leq j < k$ such that $q_1 = t_1$, $q_2 = t_2, \dots, q_j = t_j$, $q_{j+1} = 0$, $t_{j+1} = *$; and $t \preceq q$ iff $q = t$ or there is a j , $1 \leq j < k$ such that $q_1 = t_1$, $q_2 = t_2, \dots, q_j = t_j$, $q_{j+1} = 1$, $t_{j+1} = *$. For vertex coordinates $q = q_1 \dots q_k$ and $r = r_1 \dots r_k$, if $q \prec r$, $q \neq r$, then there is a $0 \leq j < k$ such that $q_i = r_i$ for $1 \leq i \leq j$ and $q_{j+1} = 0$, $r_{j+1} = 1$. Thus the Steiner coordinate $t = q_1 \dots q_j * \dots *$ is such that $q \prec t \prec r$, and if $q = r$ then setting $t = q$ gives $q \preceq t \preceq r$.

In the violator dag $\widehat{G} = (\widehat{V}, \widehat{E})$, a Steiner vertex s is of the form (s_1, \dots, s_d) , where all the s_i have Steiner coordinates and at least one has a $*$. Let S represent all Steiner vertices. Then $\widehat{V} = V \cup S$. Given $p = (p_1, \dots, p_d) \in V$ and $s = (s_1, \dots, s_d) \in S$, $p \prec s$ in \widehat{G} iff $p_i \preceq s_i$ for all $1 \leq i \leq d$; and $s \prec p$ iff

$s_i \preceq p_i$ for all $1 \leq i \leq d$. Add a directed edge $s \rightarrow v$ to \widehat{E} iff $v \prec s$, and $v \rightarrow s$ iff $s \prec v$. There are no edges with both endpoints in V or both in S . For $p, q \in V$, if $p \prec q$ as d -dimensional points (i.e., in the order we are constructing a representation of) then for each dimension i let s_i be the unique Steiner index such that $p_i \preceq s_i \preceq q_i$, and let $s = (s_1, \dots, s_d)$. If $p \neq q$ then there is at least one dimension i where $p_i \neq q_i$, and hence s_i has at least one $*$, i.e., s is a Steiner point. Thus s is the unique Steiner point such that the edges $s \rightarrow v$ and $q \rightarrow s$ are in \widehat{E} , and if $p \not\prec q$ there is no path in \widehat{G} from q to p . The property that for all $p, q \in V$, $p \prec q$ iff there is a path no longer than 2 from q to p is known as a *2-transitive closure*, and adding the Steiner points makes it a *Steiner 2-transitive closure*. These have been studied for numerous reasons, see [4].

Any vertex has at most $\prod_{i=1}^d k_i \leq \lceil \lg n \rceil^d$ incoming and outgoing edges, hence the total number of edges is $O(n \log^d n)$, and there are no more than this many vertices that have at least one edge. In [36] it is shown that this dag, eliminating any Steiner points with only 0 or 1 incident edges and any edges that were connected to such Steiner points, can be created in $\Theta(n \log^d n)$ time, where the implied constants depend on d . In [36] it is called a *rendezvous graph*, and it is also shown that one dimension can be simplified, reducing the number of vertices, number of edges, and time to construct to $\Theta(n \log^{d-1} n)$. This is called the *reduced rendezvous graph*. While some have created the equivalent of the rendezvous graph for use as a 2-transitive closure (see [4] for relevant examples and references), the reduced rendezvous graph is not 2-transitive and hasn't been used in this community.

A simple modification can create a violator dag [37]: add an extra dimension to the reduced rendezvous graph, where the value in this dimension is the function value. Reversing the ordering of domination in this dimension, or making the value the negative of the function value and keeping the same definition, turns it into a violator dag. Flow algorithms put weight 0 on the Steiner points. Thus the time to construct the violator dag is $\Theta(n \log^d n)$. Combined with Theorem 1 this gives:

Theorem 4 *Given a set of n points in d -dimensional space, $d \geq 2$, with coordinate-wise ordering and integer weighted function (f, w) on the points, with function values and weights in $[0, U]$, an L_1 monotonic regression can be found in $O(\mathcal{F}(\widehat{n}, \widehat{n}, nU) \log n)$ time, where $\widehat{n} = n \log^d n$. Using the GLP flow algorithm gives $\tilde{O}(\widehat{n}^{\frac{3}{2} - \frac{1}{328}})$ time, where the implied constants are dependent on d . \square*

Since we assume d is fixed and $n \rightarrow \infty$, $\tilde{O}(\widehat{n}^{\frac{3}{2} - \frac{1}{328}})$ is $o(n^{1.5})$. For $d > 2$ and a weighted function the fastest previous L_1 algorithms took $\Theta(n^2 \log^d n)$ time, and $\Theta(n^{1.5} \log^{d+1} n)$ for unweighted functions [35].

Combining Theorems 4 and 2 gives

Corollary 5 *Given a set of n points in d -dimensional space, $d \geq 2$, with coordinate-wise ordering and weighted function (f, w) on the points, given $p \in (1, \infty)$, and $\delta > 0$, an isotonic regression within δ of the L_p monotonic regression regression can be found in $O(\mathcal{F}(\widehat{n}, \widehat{n}, nU) \log K)$ time, where $\widehat{n} = n \log^d n$ and $K = (\max_{v \in V} f(v) - \min_{v \in V} f(v)) / \delta$ and where the implied constants in the O -notation depend on p and d . Using the GLP flow algorithm gives $\tilde{O}(\widehat{n}^{\frac{3}{2} - \frac{1}{328}} \log K)$ time, where the implied constants are dependent on p and d . \square*

For $d = 2$ dynamic programming approaches are known for L_1 and L_2 , where for a 2-d grid a $\{0, 1\}$ -valued L_1 regression can be found in $\Theta(n)$ time, and hence for arbitrary values L_1 can be solved in $\Theta(n \log n)$ time [35]. For L_2 the time of the dynamic programming algorithm in [34] is $\Theta(n^2)$. Neither algorithm involves violator graphs and each gives an exact result. Since on a 2-d grid the $\{0, 1\}$ -valued L_1 regression can be determined in $\Theta(n)$ time, it can be used to replace the flow approach used in Corollary 3, giving the result in Theorem 5.6 of [35], namely that if integral function values and weights are in $[0, U]$ then for a 2-dimensional grid of n vertices the exact L_2 regression can be found in $\Theta(n \log U)$ time, and for 2-d points in general position can be found in $\Theta(n \log n \log U)$ time.

Thus for integer values and weights in $[1, U]$, if U grows at most polynomially in n then the L_2 regression on a grid can be found in $O(n \log n)$ time, and for 2-d points in arbitrary position in $O(n \log^2 n)$ time.

4.2 Orderings Given By Pairwise Comparisons

Some problems may not initially look like they can be framed as multi-dimensional vertices. For example, suppose we collect estimates of the number of people within a range of income and age, for various ranges. This is asking for estimates of the number of people in an axis-parallel (iso-oriented) rectangular region, and the correct values are isotonic in that if one range contains another the former has at least as many people as latter. A rectangle R with lower left coordinates (x_1^r, y_1^r) and upper right coordinates (x_2^r, y_2^r) is contained in a rectangle S with coordinates (x_1^s, y_1^s) and (x_2^s, y_2^s) iff $(x_1^s, y_1^s) \preceq (x_1^r, y_1^r)$ and $(x_2^r, y_2^r) \preceq (x_2^s, y_2^s)$. By reversing the ordering on two dimensions, this is just a 4-dimensional ordering, and the same technique extends to axis-aligned boxes in arbitrary dimensions. Sets of axis-parallel boxes have been studied in other settings, such as generating all intersecting pairs, but there don't seem to be previous algorithms which represent all violating pairs so concisely in a dag.

A slight variant of this does not seem to have such a simple multidimensional representation, but does have the same special property that the vertices are objects of some form (such as the d -dimensional coordinates above) and the ordering is given explicitly by direct comparisons of objects, rather than via a dag. For example, one may have a collection of rectangles in arbitrary orientation, not necessarily all axis-parallel, and want estimates of their area, where the ordering is still containment. While there isn't a Steiner-based representation to reduce the size of a violator dag (to the author's knowledge), direct pairwise comparisons allows one to create the transitive closure violator dag in $\Theta(n^2)$ time, with the dag having at most $\binom{n}{2}$ edges. By using the BLLSSSW algorithm the expected time for L_0 and L_1 isotonic regressions is $\tilde{\Theta}(n^2)$, which is better than the $\Theta(n^\omega)$ time for general dags. Similar times can be obtained for approximate L_p isotonic regression, $1 < p < \infty$, and exact regressions under the constraints in Corollary 3. Previous algorithms for L_0 could trivially be modified to exploit the fact that the transitive closure of the violator dag was generated this quickly, but they always deferred to a simple $\Theta(n^3)$ flow algorithm so no overall benefit would have been obtained [12, 30, 31].

Some examples with similar characteristics and algorithms (and orderings which are not just linear) are vertices which are:

- strings over some alphabet where $v \prec w$ iff v is a substring of w
- strings where $v \prec w$ iff v is a subsequence of w
- general polyhedra in d -dimension space where the ordering is by containment
- subsets of a given set, with ordering by containment
- integers, where $v \prec w$ iff v evenly divides w

For some of these examples the time to do comparisons does not have a fixed bound and would need to be incorporated in the time analysis.

5 Final Remarks

Isotonic regression is a useful non-parametric shape-constrained approach to problems where assumptions such as linearity, logistic, or other parametric requirements are dubious. It has growing popularity in data analysis and machine learning [7, 8, 10, 11, 15, 20, 23, 31, 39], but to scale to the future size of some of these problems requires faster algorithms. This paper describes such algorithms without restricting to any specific application.

Recent advances in maximum flow algorithms [6, 16], coupled with the standard approach for L_0 isotonic regression and L_1 partitioning, gives faster algorithms for exact solutions for L_0 , L_1 , and L_2 when the weights

L_0

graph	time	reference	here
$d \geq 2$ arbitrary	$\Theta(n^3)$	[12, 30, 31]	$o(n^{1.5})$
pairwise compare	$\Theta(n^3)$	[12, 30, 31]	$\tilde{O}(n^2)$
arbitrary	$\Theta(n^3)$	[12, 30, 31]	$\Theta(\min\{nm, n^\omega\})$

L_1

graph	time	reference	here
$d \geq 3$ unweighted	$\Theta(n^{1.5} \log^{d+1} n)$	[35]	$o(n^{1.5})$
$d \geq 3$ arbitrary	$\Theta(n^2 \log^d n)$	[35]	$o(n^{1.5})$
pairwise compare	$\Theta(n^3)$	[3]	$\tilde{O}(n^2)$
arbitrary	$\Theta(nm + n^2 \log n)$	[3]	$\Theta(\min\{nm, n^\omega\})$

L_2

graph	time	reference	here
$d \geq 3$ arbitrary	$\Theta(n^2 \log^{2d-1} n)$	[18] + [36]	$o(n^{1.5})$
pairwise compare	$\Theta(n^3)$	[18]	$\tilde{O}(n^2)$
arbitrary	$\Theta\left(nm \log \frac{n^2}{m}\right)$	[18]	$\Theta(\min\{nm, n^\omega\})$

Table 1: Previous fastest algorithms vs. algorithms introduced here

These previous algorithms are exact for real-valued weighted functions, while the ones introduced here are exact but require initial values and weights to be integers that grow at most polynomially, or, more generally, polynomial multiples of $\delta > 0$, where $1/\delta$ grows at most polynomially. Previous algorithms are worst-case times, while here for pairwise compare and arbitrary dags the times are expected case with high probability.

“Pairwise compare” assumes each comparison takes constant time.

are positive integers in $[0, U]$ (and for L_2 the values are similarly constrained). For L_p , $1 < p < \infty$, we give algorithms for approximate 2 and exact 3 isotonic regressions, where the latter has severe restrictions on the values. Table 1 summarizes the times of the algorithms developed here vs. the fastest ones published previously among those that produce exact answers (to within machine round-off error) for real-valued input. It does not include the approximation algorithms in Kyng, Rao, and Sachdeva [22] which are similar to ours but slightly slower. For d -dimensional vertices theirs takes $\tilde{\Theta}(n^{1.5} \log^{1.5(d+1)})$ time vs. the $o(n^{1.5})$ here, and for an arbitrary dag takes $\Omega(n^3)$ vs. the $\Theta(n^\omega)$ here. However, if m is a parameter then for general dags theirs is $\tilde{\Theta}(m^{1.5})$, which is faster when $m = o(n^{\omega/1.5})$.

One drawback of using a violator graph is that it may be that $\hat{m} \gg m$, and without additional information in the worst case $\hat{m} = \Theta(n^2)$ no matter how small m is. Fortunately, for the very important class of d -dimensional component-wise ordering there is a violator dag where $\hat{m} = \tilde{\Theta}(n)$ (Sec. 4). A similar drawback is that the original dag may be unweighted, but determining the minimum flow on the violator dag (step 2 (b) in Fig. 1) might involve flow graphs with edge capacities nearly nU (on the flow graph needed to convert a satisfying flow to a minimum satisfying flow, which is the graph where the maximum flow algorithm is used).

When a violator dag is given the times are determined by the performance of the maximum flow algorithm, so efficient implementations of these complex algorithms are needed. However, our time analyses are stated in general terms, so whatever violator graph and flow algorithms are used, the time is given in terms of their performance. This accommodates advances in flow algorithms, and already shows improvements in isotonic regression algorithms using recent flow algorithms. Similarly, for a general dag the time for generating the violator dag by determining the transitive closure and eliminating edges that don't correspond to a violating pair is given as $\Theta(\min\{nm, n^\omega\})$. Usually this analysis is discussed in terms of the smallest known value of ω , but here ω can be used to indicate the time of the algorithm actually used (such as $\omega \approx 2.81$ for Strassen's algorithm), not necessarily a value that requires a galactic algorithm.

References

- [1] Aho, A.V.; Garey, M.R.; Ullman, J.D. (1972), "The transitive reduction of a directed graph", *SIAM J. Comput.* 1, pp. 131–137.
- [2] Ahuja, R.K.; Orlin, J.B.(2001), "A fast scaling algorithm for minimizing separable convex functions subject to chain constraints". *Oper. Res.* 49, pp. 784–790.
- [3] Angelov, S; Harb, B; Kannan, S; Wang, L-S (2006), "Weighted isotonic regression under the ℓ_1 norm", *Symp. on Discrete Algorithms (SODA)*, pp. 783–791.
- [4] Berman, P; Bhattacharyya, A; Grigorescu, E; Raskhodnikova, S; Woodruff, DP; Yaroslavtsev, G (2014), "Steiner transitive-closure spanners of low-dimensional posets", *J. Combinatorica* 34, pp. 255–277.
- [5] Brabant, Q; Couceiro, M; Dubois, D; Prade, H; Rico, Q (2020), "Learning rule sets and Sugeno integrals for monotonic classification problems", *Fuzzy Sets and Systems*, 401, pp 4–37.
- [6] van den Brend, J; Lee, YT; Liu, YP; Saranurak, T; Sidford, A; Song, Z; Wang, D (2021), "Minimum cost flows, MDPs, and L_1 regression in nearly linear time for dense instances", arXiv:2001.005719.
- [7] Canini, K; Cotter, A; Gupta, M.R.; Milani Fard, M; Pfeifer, J (2016), "Fast and flexible monotonic functions with ensembles of lattices", *NIPS '16*, pp. 2927–2935.
- [8] Cano, J-R; Gutierrez, PA; Krawcsyk, B; Wozniak, M; García, S (2018), "Monotonic classification: an overview on algorithms, performance measures and data sets", arXiv:1811.07155

- [9] Chepoi, V.; Cogneau, D.; Fichet, B (1967), “Polynomial algorithms for isotonic regression”, *LI Statistical Procedures and Related Topics*, IMS Lecture Note Mono. Series, 31, pp. 147–160 (1967).
- [10] Daniels, H.; Velikova, M (2010), “Monotone and partially monotone neural networks”, *IEEE Trans. Neural Net.* 21 (6), pp. 906–917.
- [11] Duivesteyn, W; Feelders, A (2008), “Nearest neighbour classification with monotonicity constraints”, In: Daelemans W., Goethals B., Morik K. (eds) *Machine Learning and Knowledge Discovery in Databases*, LNCS vol. 5211.
- [12] Feelders, A; Velikova, M; Daniels, H (2006), “Two polynomial algorithms for relabeling non-monotone data”, Tech. Report UU-CS-2006-046, Dept. Info. Com. Sci., Utrecht Univ.
- [13] Fischer, M. J.; Meyer, A. R. (1971), “Boolean matrix multiplication and transitive closure”, *IEEE Symp. on Switching and Automata*.
- [14] Furman, M.E. (1970), “Application of a method of fast multiplication of matrices in the problem of finding the transitive closure of a graph”, *Dok Akad. Nauk SSSR* 94 3, and *Soviet Math. Dokl.* 11 5, 1252.
- [15] Gamarnik, D; Gaudio, J (2019), “Sparse high-dimensional isotonic regression”, *NIPS*
- [16] Gao, Y; Liu, Y; Peng, R (2021). “Fully dynamic electrical flows: sparse maxflow faster than Goldberg-Rao”, arXiv:2101.07233
- [17] Gutiérrez, PA; García, S (2016), “Current prospects on ordinal and monotonic classification”, *Progress in Artificial Intell.* 5, pp. 171–179.
- [18] Hochbaum, DS; Queyranne, M (2003), “Minimizing a convex cost closure set”, *SIAM J. Discrete Math* 16, pp. 192–207.
- [19] Kakade, S; Kalai, AT; Kanade, V; Shamir, O (2011), “Efficient learning of generalized linear and single index models with isotonic regression”, arXiv:1104.2018.
- [20] Kalai, A.T.; Sastry, R (2009), “The Isotron algorithm: high-dimensional isotonic regression”, *COLT*, 9 p.
- [21] Menon, AK; Jiang, XJ; Vembu, S; Elkan, C; Ohno-Machado, L (2012), “Predicting accurate probabilities with a ranking loss”, *Proc. Int. Conf Mach Learn.* 2012, pp. 703–710.
- [22] Kyng, R; Rao, A; Sachdeva, S (2015), “Fast, provable algorithms for isotonic regression in all L_p -norms”, *NIPS*.
- [23] Liu, X; Han, X; Zhang, N; Liu, Q (2020), “Certified monotonic neural nets”, *NeurIPS*.
- [24] Marques-Silva, J; Gerspacher, T; Cooper, M; Ignatiev, A; Narodytska, N (2021), “Explanations for monotonic classifiers”, *Int’l. Conf. Machine Learning*.
- [25] Möhring, R (1985), “Algorithmic aspects of comparability graphs and interval graphs”, *Proc. NATO Adv. Study Inst. on Graphs and Order*, Rival, I, ed., 41–101.
- [26] Munro, I. (1971), “Efficient determination of the strongly connected components and transitive closure of a directed graph”, *Info. Proc. Letters* 1, pp. 56–58.

- [27] Pang, LLM; Tay, KM (2020), “A new monotone fuzzy rule relabeling framework with application to failure mode and effect analysis Mmethodology”. *IEEE Access* 8, pp. 144908–144930.
- [28] Pérez-Fernández, R; Alonso, P; Díaz, I; De Baets, B (2017), “Monotonicity as a tool for differentiating between truth and optimality in the aggregation of rankings”, *Math. Psychol.*
- [29] Pijls, W; Porharst, R (2013), “Another note on Dilworth’s decomposition theorem”, *J. Discrete Mathematics* 2013.
- [30] Pijls, W; Potharst, R (2014), “Repairing non-monotone ordinal data sets by changing class labels”, *Econometric Inst. Report EI* 2014–29.
- [31] Rademaker, M; De Baets, B; De Meyer, H (2009), “Loss optimal monotone relabeling of noisy multi-criteria data sets”, *Information Sciences* 179 (2009), pp. 4089–4097.
- [32] Rademaker, M; De Baets, B; De Meyer, H (2012), “Optimal monotone relabeling of partially non-monotone ordinal data”, *Optimization Methods and Soft.* 27, 17–31.
- [33] Sills, J (1998), “Monotonic networks”, *NIPS* 10, pp. 661–667.
- [34] Spouge J; Wan, H; Wilbur WJ (2003), “Least squares isotonic regression in two dimensions”, *J. Optimization Theory and Applications* 117, pp. 585–605.
- [35] Stout, QF (2013), “Isotonic regression via partitioning”, *Algorithmica* 66, pp. 93–112.
- [36] Stout, QF (2015), “Isotonic regression for multiple independent variables”, *Algorithmica* 71, pp. 450–470.
- [37] Stout, QF (2021), “ L_0 isotonic regression with secondary objectives”, arXiv 2106.00279
- [38] Strömberg, U. (1991), “An algorithm for isotonic regression with arbitrary convex distance function”, *Comput. Stat. Data Anal.* 11, pp. 205–219.
- [39] Tao, Y; Wang, Y (2021), “New algorithms for monotone classification”, *Principals of Database Systems (PODS '21)*, 13p.
- [40] Verbeke, W; Martens, D; Baesens, B (2017), “RULEM: A novel heuristic rule learning approach for ordinal classification with monotonicity constraints”, *Applied Soft Computing* 60 (2017), pp. 858–873
- [41] en.wikipedia.org/wiki/Maximum_flow_problem