

Lipschitz Unimodal and Isotonic Regression on Paths and Trees*

Pankaj K. Agarwal
Duke University
Durham, NC 27708
pankaj@cs.duke.edu

Jeff M. Phillips
University of Utah
Salt Lake City, UT 84112
jeffp@cs.utah.edu

Bardia Sadri
University of Toronto
Toronto, ON M5S 3G4
sadri@cs.toronto.edu

January 2, 2010

Abstract

We describe algorithms for finding the regression of t , a sequence of values, to the closest sequence s by mean squared error, so that s is always increasing (isotonicity) and so the values of two consecutive points do not increase by too much (Lipschitz). The isotonicity constraint can be replaced with a unimodal constraint, where there is exactly one local maximum in s . These algorithms are generalized from sequences of values to trees of values. For each scenario we describe near-linear time algorithms.

1 Introduction

Let \mathbb{M} be a triangulation of a polygonal region $P \subseteq \mathbb{R}^2$ in which each vertex is associated with a real valued height (or elevation). Linear interpolation of vertex heights in the interior of each triangle of \mathbb{M} defines a piecewise-linear function $t : P \rightarrow \mathbb{R}$, called a *height function*. A height function (or its graph) is widely used to model a two-dimensional surface in numerous applications (e.g. modeling the terrain of a geographical area). With recent advances in sensing and mapping technologies, such models are being generated at an unprecedentedly large scale. These models are then used to analyze the surface and to compute various geometric and topological properties of the surface. For example, researchers in GIS are interested in extracting river networks or computing visibility or shortest-path maps on terrains modeled as height functions. These structures depend heavily on the topology of the level sets of the surface and in particular on the topological relationship between its critical points (maxima, minima, and saddle points). Because of various factors such as measurement or sampling errors or the nature of the sampled surface, there is often plenty of noise in these surface models which introduces spurious critical points. This in turn leads to misleading or undesirable artifacts in the computed structures, e.g., artificial breaks in river networks. These difficulties have motivated extensive work on topological simplification and noise removal through modification of the height function t into another one $s : \mathbb{M} \rightarrow \mathbb{R}$ that has the desired set of critical points and that is as close to t as possible [7, 21, 26, 25, 33]. A popular approach is to decompose the surface into pieces and modify each piece so that it has a unique minimum or maximum [27]. In some applications, it is also desirable to impose the additional constraint that the function s is Lipschitz; see below for further discussion.

*The work was primarily done when the second and third authors were at Duke University. Research supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEG-P200A070505, by a grant from the U.S.–Israel Binational Science Foundation, and a subaward to the University of Utah under NSF Award 0937060 to Computing Research Association.

Problem statement. Let $\mathbb{M} = (V, A)$ be a planar graph with vertex set V and arc (edge) set $A \subseteq V \times V$. We may treat \mathbb{M} as undirected in which case we take the pairs (u, v) and (v, u) as both representing the same undirected edge connecting u and v . Let $\gamma \geq 0$ be a real parameter. A function $s : V \rightarrow \mathbb{R}$ is called

(L) γ -Lipschitz if $(u, v) \in A$ implies $s(v) - s(u) \leq \gamma$.

Note that if \mathbb{M} is undirected, then Lipschitz constraint on an edge $(u, v) \in A$ implies $|s(u) - s(v)| \leq \gamma$. For an undirected planar graph $\mathbb{M} = (V, A)$, a function $s : V \rightarrow \mathbb{R}$ is called

(U) *unimodal* if s has a unique *local maximum*, i.e. only one vertex $v \in V$ such that $s(v) > s(u)$ for all $(u, v) \in A$.

For a directed planar graph $\mathbb{M} = (V, A)$, a function $s : V \rightarrow \mathbb{R}$ is called

(I) *isotonic* if $(u, v) \in A$ implies $s(u) \leq s(v)$.¹

For an arbitrary function $t : V \rightarrow \mathbb{R}$ and a parameter γ , the γ -Lipschitz unimodal regression (γ -LUR) of t is a function $s : V \rightarrow \mathbb{R}$ that is γ -Lipschitz and unimodal on \mathbb{M} and minimizes $\|s - t\|^2 = \sum_{v \in V} (s(v) - t(v))^2$. Similarly, if \mathbb{M} is a directed planar graph, then s is the γ -Lipschitz isotonic regression (γ -LIR) of t if s satisfies (L) and (I) and minimizes $\|s - t\|^2$. The more commonly studied *isotonic regression* (IR) and *unimodal regression* [2, 4, 6, 23] are the special cases of LIR and LUR, respectively, in which $\gamma = \infty$, and therefore only the condition (I) or (U) is enforced.

Given a planar graph \mathbb{M} , a parameter γ , and $t : V \rightarrow \mathbb{R}$, the LIR (resp. LUR) problem is to compute the γ -LIR (resp. γ -LUR) of t . In this paper we propose near-linear-time algorithms for the LIR and LUR problems for two special cases: when \mathbb{M} is a path or a tree. We study the special case where \mathbb{M} is a path prior to the more general case where it is tree because of the difference in running time and because doing so simplifies the exposition to the more general case.

Related work. As mentioned above, there is extensive work on simplifying the topology of a height function while preserving the geometry as much as possible. Two widely used approaches in GIS are the so-called *flooding* and *carving* techniques [1, 11, 25, 26]. The former technique raises the height of the vertices in “shallow pits” to simulate the effect of flooding, while the latter lowers the value of the height function along a path connecting two pits so that the values along the path vary monotonically. As a result, one pit drains to the other and thus one of the minima ceases to exist. Various methods based on Laplacian smoothing have been proposed in the geometric modeling community to remove unnecessary critical points; see [5, 7, 18, 21] and references therein. For example, Ni *et al.* [21] proposed the so-called *Morse-fairing* technique, which solves a relaxed form of Laplace’s equation, to remove undesired critical points.

A prominent line of research on topological simplification was initiated by Edelsbrunner *et al.* [14, 13] who introduced the notion of *persistence*; see also [12, 35, 34]. Roughly speaking, each homology class of the contours in sublevel sets of a height function is characterized by two critical points at one of whom the class is born and at the other it is destroyed. The persistence of this class is then the height difference between these two critical points and can be thought of as the life span of that class. The persistence of a class effectively suggests the “significance” of its defining critical points. Efficient algorithms have been developed for computing the persistence associated to the critical

¹A function s satisfying the isotonicity constraint (I) must assign the same value to all the vertices of a directed cycle of \mathbb{M} (and indeed to all vertices in the same strongly connected component). Therefore, without loss of generality, we assume \mathbb{M} to be a directed *acyclic* graph (DAG).

points of a height function and for simplifying topology based on persistence [13, 7]. Edelsbrunner *et al.* [15] and Attali *et al.* [3] proposed algorithms for optimally eliminating all critical points of persistence below a threshold where the error is measured as $\|s - t\|_\infty = \max_{v \in V} |s(v) - t(v)|$. No efficient algorithm is known to minimize $\|s - t\|_2$.

The isotonic-regression (IR) problem has been studied in statistics [4, 6, 23] since the 1950s. It has many applications ranging from statistics [29] to bioinformatics [6], and from operations research [20] to differential optimization [17]. Ayer *et al.* [4] famously solves the IR problem on paths in $O(n)$ time using the *pool adjacent violator algorithm* (PAVA). The algorithm works by initially treating each vertex as a level set and merging consecutive level sets that are out of order. This algorithm is correct regardless of the order of the merges [24]. Brunk [9] and Thompson [32] initiated the study of the IR problem on general DAGs and trees, respectively. Motivated by the problem of finding an optimal insurance rate structure over given risk classes for which a desired pattern of tariffs can be specified, Jewell [19] introduced the problem of Lipschitz isotonic regression on DAGs and showed connections between this problem and network flow with quadratic cost functions². Stout [31] solves the UR problem on paths in $O(n)$ time. Pardalos and Xue [22] give an $O(n \log n)$ algorithm for the IR problem on trees. For the special case when the tree is a star they give an $O(n)$ algorithm. Spouge *et al.* [28] give an $O(n^4)$ time algorithm for the IR problem on DAGs. The problems can be solved under the L_1 and L_∞ norms on paths [31] and DAGs [2] as well, with an additional $\log n$ factor for L_1 . Recently Stout [30] has presented an efficient algorithm for isotonic regression in a planar DAG under the L_∞ norm. To our knowledge no polynomial-time algorithm is known for the UR problem on planar graphs, and there is no prior attempt on achieving efficient algorithms for Lipschitz isotonic/unimodal regressions in the literature.

Our results. Although the LUR problem for planar graphs remains elusive, we present efficient exact algorithms for LIR and LUR problems on two special cases of planar graphs: paths and trees. In particular, we present an $O(n \log n)$ algorithm for computing the LIR on a path of length n (Section 4), and an $O(n \log n)$ algorithm on a tree with n nodes (Section 6). We present an $O(n \log^2 n)$ algorithm for computing the LUR problem on a path of length n (Section 5). Our algorithm can be extended to solve the LUR problem on an unrooted tree in $O(n \log^3 n)$ time (Section 7). The LUR algorithm for a tree is particularly interesting because of its application in the aforementioned carving technique [11, 26, 27]. The carving technique modifies the height function along a number of trees embedded on the terrain where the heights of the vertices of each tree are to be changed to vary monotonically towards a chosen “root” for that tree. In other words, to perform the carving, we need to solve the IR problem on each tree. The downside of doing so is that the optimal IR solution happens to be a step function along each path toward the root of the tree with potentially large jumps. Enforcing the Lipschitz condition prevents sharp jumps in function value and thus provides a more natural solution to the problem.

Section 3 presents a data structure, called *affine composition tree* (ACT), for maintaining a *xy*-monotone polygonal chain, which can be regarded as the graph of a monotone piecewise-linear function $F : \mathbb{R} \rightarrow \mathbb{R}$. Besides being crucial for our algorithms, ACT is interesting in its own right. A special kind of binary search tree, an ACT supports a number of operations to query and update the chain, each taking $O(\log n)$ time. Besides the classical insertion, deletion, and query

²It may at first seem that Jewell’s formulation of the LIR s of an input function t on the vertices of a DAG is more general than ours in that he requires that for each $e = (u, v) \in A$, $s(v) \in [s(u) - \lambda_e, s(u) + \gamma_e]$ where $\lambda_e, \gamma_e \in \mathbb{R}^+$ are defined separately for each edge (in our formulation $\lambda_e = 0$ and $\gamma_e = \gamma$ for every edge e). Moreover, instead of minimizing the L_2 distance between s and t he requires $\sum_v w_v (s(v) - t(v))^2$ to be minimized where $w_v \in \mathbb{R}^+$ is a constant assigned to vertex v . However, all our algorithms in this paper can be modified in a straight-forward manner to handle this formulation without any change in running-times.

(computing $F(x)$ or $F^{-1}(x)$ for a given $x \in \mathbb{R}$), one can apply an INTERVAL operation that modifies a contiguous subchain provided that the chain remains x -monotone after the transformation, i.e., it remains the graph of a monotone function.

2 Energy Functions

On a discrete set U , a real valued function $s : U \rightarrow \mathbb{R}$ can be viewed as a point in the $|U|$ -dimensional Euclidean space in which coordinates are indexed by the elements of U and the component s_u of s associated to an element $u \in U$ is $s(u)$. We use the notation \mathbb{R}^U to represent the set of all real-valued functions defined on U .

Let $\mathbb{M} = (V, A)$ be a directed acyclic graph on which we wish to compute γ -Lipschitz isotonic regression of an *input function* $t \in \mathbb{R}^V$. For any set of vertices $U \subseteq V$, let $\mathbb{M}[U]$ denote the subgraph of \mathbb{M} induced by U , i.e. the graph $(U, A[U])$, where $A[U] = \{(u, v) \in A : u, v \in U\}$. The set of γ -Lipschitz isotonic functions on the subgraph $\mathbb{M}[U]$ of \mathbb{M} constitutes a convex subset of \mathbb{R}^U , denoted by $\Gamma(\mathbb{M}, U)$. It is the common intersection of all half-spaces determined by the isotonicity and Lipschitz constraints associated with the edges in $A[U]$, i.e., $s_u \leq s_v$ and $s_v \leq s_u + \gamma$ for all $(u, v) \in A[U]$.

For $U \subseteq V$ we define $E_U : \mathbb{R}^V \rightarrow \mathbb{R}$ as

$$E_U(s) = \sum_{v \in U} (s(v) - t(v))^2.$$

Thus, the γ -Lipschitz isotonic regression of the input function t is $\sigma = \arg \min_{s \in \Gamma(\mathbb{M}, V)} E_V(s)$. For a subset $U \subseteq V$ and $v \in U$ define the function $E_{\mathbb{M}[U]}^v : \mathbb{R} \rightarrow \mathbb{R}$ as

$$E_{\mathbb{M}[U]}^v(x) = \min_{s \in \Gamma(\mathbb{M}, U); s(v)=x} E_U(s).$$

Lemma 2.1 *For any $U \subseteq V$ and $v \in U$, the function $E_{\mathbb{M}[U]}^v$ is continuous and strictly convex.*

Proof. Given $x, y \in \mathbb{R}$ and $0 \leq \alpha \leq 1$, consider the functions

$$\begin{aligned} z_x &= \arg \min_{z \in \Gamma(\mathbb{M}, U); z(v)=x} E_U(z) \\ z_y &= \arg \min_{z \in \Gamma(\mathbb{M}, U); z(v)=y} E_U(z) \\ z_\alpha &= \alpha z_x + (1 - \alpha) z_y. \end{aligned}$$

The strong convexity of E_U implies that

$$\alpha E_U(z_x) + (1 - \alpha) E_U(z_y) \geq E_U(z_\alpha).$$

Furthermore, the convexity of $\Gamma(\mathbb{M}, U)$ implies that z_α is also in $\Gamma(\mathbb{M}, U)$. Since by definition $z_\alpha(v) = \alpha x + (1 - \alpha)y$, we deduce that

$$E_U(z_\alpha) \geq \min_{z \in \Gamma(\mathbb{M}, U); z(v)=\alpha x + (1-\alpha)y} E_U(z).$$

Thus, the minimum is precisely $E_U^v(\alpha x + (1 - \alpha)y)$. This proves that E_U^v is strictly convex on \mathbb{R} and therefore continuous (and in fact differentiable except in countably many points). ■




Figure 1: Left: the graph of a monotone piecewise linear function F (in solid black). Vertices are marked by hollow circles. The dashed curve is the result of applying the linear transform $\psi_0(q) = Mq$ where $M = \begin{pmatrix} 1 & 1/3 \\ 1/3 & 2/3 \end{pmatrix}$. The gray curve, a translation of the dashed curve under vector $c = \begin{pmatrix} -1 \\ -2 \end{pmatrix}$, is the result of applying $\text{AFFINE}(\psi)$ to F where $\psi(q) = Mq + c$. Right: $\text{INTERVAL}(\psi, \tau^-, \tau^+)$, only the vertices of curve F whose x -coordinates are in the marked interval $[\tau^-, \tau^+]$ are transformed. The resulting curve is the thick gray curve.

3 Affine Composition Tree

In this section we introduce a data structure, called *affine composition tree* (ACT), for representing an xy -monotone polygonal chain in \mathbb{R}^2 , which is being deformed dynamically. Such a chain can be regarded as the graph of a piecewise-linear monotone function $F : \mathbb{R} \rightarrow \mathbb{R}$, and thus is bijective. A *breakpoint* of F is the x -coordinate of a *vertex* of the graph of F (a vertex of F for short), i.e., a $b \in \mathbb{R}$ at which the left and right derivatives of F disagree. The number of breakpoints of F will be denoted by $|F|$. A continuous piecewise-linear function F with breakpoints $b_1 < \dots < b_n$ can be characterized by its vertices $q_i = (b_i, F(b_i))$, $i = 1, \dots, n$ together with the *slopes* μ_- and μ_+ of its *left* and *right unbounded pieces*, respectively, extending to $-\infty$ and $+\infty$. An *affine transformation* of \mathbb{R}^2 is a map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, $q \mapsto M \cdot q + c$ where M is a nonsingular 2×2 matrix, (a *linear transformation*) and $c \in \mathbb{R}^2$ is a *translation vector* — in our notation we treat $q \in \mathbb{R}^2$ as a column vector.

An ACT supports the following operations on a monotone piecewise-linear function F with vertices $q_i = (b_i, F(b_i))$, $i = 1, \dots, n$:

1. $\text{EVALUATE}(a)$ and $\text{EVALUATE}^{-1}(a)$: Given any $a \in \mathbb{R}$, return $F(a)$ or $F^{-1}(a)$.
2. $\text{INSERT}(q)$: Given a point $q = (x, y) \in \mathbb{R}^2$ insert q as a new vertex of F . If $x \in (b_i, b_{i+1})$, this operation removes the segment $q_i q_{i+1}$ from the graph of F and replaces it with two segments $q_i q$ and $q q_{i+1}$, thus making x a new breakpoint of F with $F(x) = y$. If $x < b_1$ or $x > b_k$, then the affected unbounded piece of F is replaced with one parallel to it but ending at q and a segment connecting q and the appropriate vertex of F (μ_+ and μ_- remain intact). We assume that $F(b_i) \leq y \leq F(b_{i+1})$.
3. $\text{DELETE}(b)$: Given a breakpoint b of F , removes the vertex $(b, F(b))$ from F ; a delete operation modifies F in a manner similar to insert.
4. $\text{AFFINE}(\psi)$: Given an affine transformation $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, modify the function F to one whose graph is the result of application of ψ to the graph of F . See Figure 1(Left).

5. $\text{INTERVAL}_p(\psi, \tau^-, \tau^+)$: Given an affine transformation $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $\tau^-, \tau^+ \in \mathbb{R}$ and $p \in \{x, y\}$, this operation applies ψ to all vertices v of F whose p -coordinate is in the range $[\tau^-, \tau^+]$. Note that $\text{AFFINE}(\psi)$ is equivalent to $\text{INTERVAL}_p(\psi, -\infty, +\infty)$ for $p = \{x, y\}$. See Figure 1(Right).

It must be noted that AFFINE and INTERVAL are applied with appropriate choice of transformation parameters so that the resulting chain remains xy -monotone.

An ACT $\mathcal{T} = \mathcal{T}(F)$ is a red-black tree that stores the vertices of F in the sorted order, i.e., the i th node of \mathcal{T} is associated with the i th vertex of F . However, instead of storing the actual coordinates of a vertex, each node z of \mathcal{T} stores an affine transformation $\phi_z : q \mapsto M_z \cdot q + c_z$. If $z_0, \dots, z_k = z$ is the path in \mathcal{T} from the root z_0 to z , then let $\Phi_z(q) = \phi_{z_0}(\phi_{z_1}(\dots \phi_{z_k}(q) \dots))$. Notice that Φ_z is also an affine transformation. The actual coordinates of the vertex associated with z are $(q_x, q_y) = \Phi_z(\bar{0})$ where $\bar{0} = (0, 0)$.

Given a value $b \in \mathbb{R}$ and $p \in \{x, y\}$, let $\text{PRED}_p(b)$ (resp. $\text{SUCC}_p(b)$) denote the rightmost (resp. leftmost) vertex q of F such that the p -coordinate of q is at most (resp. least) b . Using ACT \mathcal{T} , $\text{PRED}_p(b)$ and $\text{SUCC}_p(b)$ can be computed in $O(\log n)$ time by following a path in \mathcal{T} , composing the affine transformations along the path, evaluating the result at $\bar{0}$, and comparing its p -coordinate with b . We can answer $\text{EVALUATE}(a)$ queries by determining the vertices $q^- = \text{PRED}_x(a)$ and $q^+ = \text{SUCC}_x(a)$ of F immediately preceding and succeeding a and interpolating F linearly between q^- and q^+ ; if $a < b_1$ (resp. $a > b_k$), then $F(a)$ is calculated using $F(b_1)$ and μ_- (resp. $F(b_k)$ and μ_+). Since b^- and b^+ can be computed in $O(\log n)$ time and the interpolation takes constant time, $\text{EVALUATE}(a)$ is answered in time $O(\log n)$. Similarly, $\text{EVALUATE}^{-1}(a)$ is answered using $\text{PRED}_y(a)$ and $\text{SUCC}_y(a)$.

A key observation of ACT is that a standard rotation on any edge of \mathcal{T} can be performed in $O(1)$ time by modifying the stored affine transformations in a constant number of nodes (see Figure 2) based on the fact that an affine transformation $\phi : q \mapsto M \cdot q + c$ has an inverse affine transformation $\phi^{-1} : q \mapsto M^{-1} \cdot (q - c)$; provided that the matrix M is invertible. A point $q \in \mathbb{R}^2$ is inserted into \mathcal{T} by first computing the affine transformation Φ_u for the node u that will be the parent of the leaf z storing q . To determine ϕ_z we solve, in constant time, the system of (two) linear equations $\Phi_u(\phi_z(\bar{0})) = q$. The result is the translation vector c_z . The linear transformation M_z can be chosen to be an arbitrary invertible linear transformation, but for simplicity, we set M_z to the identity matrix. Deletion of a node is handled similarly.

To perform an $\text{INTERVAL}_p(\psi, \tau^-, \tau^+)$ query, we first find the nodes z^- and z^+ storing the vertices $\text{PRED}_p(\tau^-)$ and $\text{SUCC}_p(\tau^+)$, respectively. We then successively rotate z^- with its parent until it becomes the root of the tree. Next, we do the same with z^+ but stop when it becomes the right child of z^- . At this stage, the subtree rooted at the left child of z^+ contains exactly all the vertices




Figure 2: Rotation in affine composition trees. The affine function stored at a node is shown as a greek letter to its right (left). When rotating the pair (z_1, z_2) , by changing these functions as shown (right) the set of values computed at the leaves remains unchanged.

q for which $q_p \in [\tau^-, \tau^+]$. Thus we compose ψ with the affine transformation at that node and issue the performed rotations in the reverse order to put the tree back in its original (balanced) position. Since z^- and z^+ were both within $O(\log n)$ steps from the root of the tree, and since performing each rotation on the tree can only increase the depth of a node by one, z^- is taken to the root in $O(\log n)$ steps and this increases the depth of z^+ by at most $O(\log n)$. Thus the whole operation takes $O(\log n)$ time.

We can augment $\mathcal{T}(F)$ with additional information so that for any $a \in \mathbb{R}$ the function $E(a) = E(b_1) + \int_{b_1}^a F(x) dx$, where b_1 is the leftmost breakpoint and $E(b_1)$ is value associated with b_1 , can be computed in $O(\log n)$ time; we refer to this operation as $\text{INTEGRATE}(a)$. We provide the details in the Appendix. We summarize the above discussion:

Theorem 3.1 *A continuous piecewise-linear monotonically increasing function F with n breakpoints can be maintained using a data structure $\mathcal{T}(F)$ such that*

1. EVALUATE and EVALUATE^{-1} queries can be answered in $O(\log n)$ time,
2. an INSERT or a DELETE operation can be performed in $O(\log n)$ time,
3. AFFINE and INTERVAL operations can be performed in $O(1)$ and $O(\log n)$ time, respectively.
4. INTEGRATE operation can be performed in $O(\log n)$ time.

One can use the above operations to compute the sum of two increasing continuous piecewise-linear functions F and G as follows: we first compute $F(b_i)$ for every breakpoint b_i of G and insert the pair $(b_i, F(b_i))$ into $\mathcal{T}(F)$. At this point the tree still represents $\mathcal{T}(F)$ but includes all the breakpoints of G as *degenerate* breakpoints (at which the left and right derivatives of F are the same). Finally, for every consecutive pair of breakpoints b_i and b_{i+1} of G we apply an $\text{INTERVAL}_x(\psi_i, b_i, b_{i+1})$ operation on $\mathcal{T}(F)$ where ψ_i is the affine transformation $q \mapsto Mq + c$ where $M = \begin{pmatrix} 1 & 0 \\ \alpha & 1 \end{pmatrix}$ and $c = \begin{pmatrix} 0 \\ \beta \end{pmatrix}$, in which $G_i(x) = \alpha x + \beta$ is the linear function that interpolates between $G(b_i)$ at b_i and $G(b_{i+1})$ at b_{i+1} (similar operation using μ_- and μ_+ of G for the unbounded pieces of G must can be applied in constant time). It is easy to verify that after performing this series of INTERVAL 's, $\mathcal{T}(F)$ turns into $\mathcal{T}(F + G)$. The total running time of this operation is $O(|G| \log |F|)$. Note that this runtime can be reduced to $O(|G|(1 + \log |F| / \log |G|))$, for $|G| < |F|$, by using an algorithm of Brown and Tarjan [8] to insert all breakpoints and then applying all INTERVAL operations in a bottom up manner. Furthermore, this process can be reversed (i.e. creating $\mathcal{T}(F - G)$ without the breakpoints of G , given $\mathcal{T}(F)$ and $\mathcal{T}(G)$) in the same runtime. We therefore have shown:

Lemma 3.2 *Given $\mathcal{T}(F)$ and $\mathcal{T}(G)$ for a piecewise-linear isotonic functions F and G where $|G| < |F|$, $\mathcal{T}(F + G)$ or $\mathcal{T}(F - G)$ can be computed in $O(|G|(1 + \log |F| / \log |G|))$ time.*

Tree sets. In our application we will be repeatedly computing the sum of two functions F and G . It will be too expensive to compute $\mathcal{T}(F + G)$ explicitly using Lemma 3.2, therefore we represent it implicitly. More precisely, we use a *tree set* $\mathcal{S}(F) = \{\mathcal{T}(F_1), \dots, \mathcal{T}(F_k)\}$ consisting of affine composition trees of monotone piecewise-linear functions F_1, \dots, F_k to represent the function $F = \sum_{j=1}^k F_j$. We perform several operations on F or \mathcal{S} similar to those of a single affine composition tree.

$\text{EVALUATE}(x)$ on F takes $O(k \log n)$ time, by evaluating $\sum_{j=1}^k F_j(x)$. $\text{EVALUATE}^{-1}(y)$ on F takes $O(k \log^2 n)$ time using a technique of Frederickson and Johnson [16].

Given the ACT $\mathcal{T}(F_0)$, we can convert $\mathcal{S}(F)$ to $\mathcal{S}(F + F_0)$ in two ways: an $\text{INCLUDE}(\mathcal{S}, F_0)$ operation sets $\mathcal{S} = \{\mathcal{T}(F_1), \dots, \mathcal{T}(F_k), \mathcal{T}(F_0)\}$ in $O(1)$ time. A $\text{MERGE}(\mathcal{S}, F_0)$ operations sets

$\mathcal{S} = \{\mathcal{T}(F_1 + F_0), \mathcal{T}(F_2), \dots, \mathcal{T}(F_k)\}$ in $O(|F_0| \log |F_1|)$ time. We can also perform the operation $\text{UNINCLUDE}(\mathcal{S}, F_0)$ and $\text{UNMERGE}(\mathcal{S}, F_0)$ operations that reverse the respective above operations in the same runtimes.

We can perform an $\text{AFFINE}(\mathcal{S}, \psi)$ where ψ describes a linear transform M and a translation vector c . To update F by ψ we update F_1 by ψ and for $j \in [2, k]$ update F_j by just M . This takes $O(k)$ time. It follows that we can perform $\text{INTERVAL}(\mathcal{S}, \psi, \tau^-, \tau^+)$ in $O(k \log n)$ time, where $n = |F_1| + \dots + |F_k|$. Here we assume that the transformation ψ is such that each F_i remains monotone after the transformation.

4 LIR on Paths

In this section we describe an algorithm for solving the LIR problem on a path, represented as a directed graph $P = (V, A)$ where $V = \{v_1, \dots, v_n\}$ and $A = \{(v_i, v_{i+1}) : 1 \leq i < n\}$. A function $s : V \rightarrow \mathbb{R}$ is isotonic (on P) if $s(v_i) \leq s(v_{i+1})$, and γ -Lipschitz for some real constant γ if $s(v_{i+1}) \leq s(v_i) + \gamma$ for each $i = 1, \dots, n-1$. For the rest of this section let $t : V \rightarrow \mathbb{R}$ be an input function on V . For each $i = 1, \dots, n$, let $V_i = \{v_1, \dots, v_i\}$, let P_i be the subpath v_1, \dots, v_i , and let E_i and \tilde{E}_i , respectively, be shorthands for E_{V_i} and $E_{P_i}^{v_i}$. By definition, if we let $\tilde{E}_0 = 0$, then for each $i \geq 1$:

$$\tilde{E}_i(x) = (x - t(v_i))^2 + \min_{x-\gamma \leq y \leq x} \tilde{E}_{i-1}(y). \quad (1)$$

By Lemma 2.1, \tilde{E}_i is convex and continuous and thus has a unique minimizer s_i^* .

Lemma 4.1 *For $i \geq 1$, the function \tilde{E}_i is given by the recurrence relation:*

$$\tilde{E}_i(x) = (x - t(v_i))^2 + \begin{cases} \tilde{E}_{i-1}(x - \gamma) & x > s_{i-1}^* + \gamma \\ \tilde{E}_{i-1}(s_{i-1}^*) & x \in [s_{i-1}^*, s_{i-1}^* + \gamma] \\ \tilde{E}_{i-1}(x) & x < s_{i-1}^*. \end{cases} \quad (2)$$

Proof. The proof is by induction on i . \tilde{E}_1 is clearly a single-piece quadratic function. For $i > 1$, since \tilde{E}_{i-1} is strictly convex, it is strictly decreasing on $(-\infty, s_{i-1}^*)$ and strictly increasing on $(s_{i-1}^*, +\infty)$. Thus depending on whether $s_{i-1}^* < x - \gamma$, $s_{i-1}^* \in [x - \gamma, x]$, or $s_{i-1}^* > x$, the value $y \in [x - \gamma, x]$ that minimizes $\tilde{E}_i(y)$ is $x - \gamma$, s_{i-1}^* , and x , respectively. ■

Thus by Lemmas 2.1 and 4.1, \tilde{E}_i is strictly convex and piecewise quadratic. We call a value x that determines the boundary of two neighboring quadratic pieces of \tilde{E}_i a *breakpoint* of \tilde{E}_i . Since \tilde{E}_1 is a simple (one-piece) quadratic function, it has no breakpoints. For $i > 1$, the breakpoints of the function \tilde{E}_i consist of s_{i-1}^* and $s_{i-1}^* + \gamma$, as determined by recurrence (2), together with breakpoints that arise from recursive applications of \tilde{E}_{i-1} . Examining equation (2) reveals that all breakpoints of \tilde{E}_{i-1} that are smaller than s_{i-1}^* remain breakpoints in \tilde{E}_i and all those larger than s_{i-1}^* are increased by γ and these form all of the breakpoints of \tilde{E}_i (see Figure 3). Thus \tilde{E}_i has precisely $2i - 2$ breakpoints. To compute the point s_i^* at which $\tilde{E}_i(x)$ is minimized, it is enough to scan over these $O(i)$ quadratic pieces and find the unique piece whose minimum lies between its two ending breakpoints.

Lemma 4.2 *Given the sequence s_1^*, \dots, s_n^* , one can compute the γ -LIR s of input function t in $O(n)$ time.*

Proof. For each $i = 1, \dots, n$, let $\sigma_i = \operatorname{argmin}_{x \in \Gamma(P, V_i)} E_i(x)$, i.e. σ_i is the γ -LIR of $\tau_i \in \mathbb{R}^{V_i}$ where $\tau_i(v) = t(v)$ for all $v \in V_i$ and in particular $\sigma_n = s$ where s is the γ -LIR of t . From the definitions




Figure 3: The breakpoints of the function $F_i = d\tilde{E}_i/dx$. For each i , s_{i-1}^* and $s_{i-1}^* + \gamma$ are the “new” breakpoints of F_i . All other breakpoints of F_i come from F_{i-1} where those smaller than s_i^* remain unchanged and those larger are increased by γ .

we get

$$E_i(\sigma_i) = \min_{x \in \mathbb{R}} \tilde{E}_i(x).$$

The uniqueness of the σ_i and s_i^* implies that $\sigma_i(v_i) = s_i^*$. In particular, $\sigma_n(v_n) = s(v_n) = s_n^*$. Suppose that the numbers $s(v_{i+1}), \dots, s(v_n)$ are determined, for some $1 \leq i \leq n-1$, and we wish to determine $s(v_i)$. Picking $s(v_i) = x$ entails that the energy of the solution will be at least

$$\tilde{E}_i(x) + \sum_{j=i+1}^n (s(v_j) - t(v_j))^2.$$

Thus $s(v_i)$ has to be chosen to minimize \tilde{E}_i on the interval $[s(v_{i+1}) - \gamma, s(v_{i+1})]$. If $s_i^* \in [s(v_{i+1}) - \gamma, s(v_{i+1})]$, then we simply have $s(v_i) = s_i^*$ which is the global minimum for \tilde{E}_i . Otherwise, one must pick $s(v_i) = s(v_{i+1}) - \gamma$ if $s_i^* < s(v_{i+1}) - \gamma$ and $s(v_i) = s(v_{i+1})$ if $s_i^* > s(v_{i+1})$. Thus each $s(v_i)$ can be found in $O(1)$ time for each i and the entire process takes $O(n)$ time. ■

One can compute the values of s_1^*, \dots, s_n^* in n iterations. The i th iteration computes the value s_i^* at which \tilde{E}_i is minimized and then uses it to compute the function \tilde{E}_{i+1} as given by (2) in $O(i)$ time. After having computed s_i^* 's, the γ -LIR of $t \in \mathbb{R}^V$ can be computed in linear time. However, this gives an $O(n^2)$ algorithm for computing the γ -LIR of t . We now show how this running time can be reduced to $O(n \log n)$.

For the sake of simplicity, we assume for the rest of this paper that for each $2 \leq i \leq n$, s_i^* , i.e., the point minimizing \tilde{E}_i is none of its breakpoints, although it is not hard to relax this assumption algorithmically. Under this assumption, s_i^* belongs to the interior of some interval on which \tilde{E}_i is quadratic. The derivative of this quadratic function is therefore zero at s_i^* . In other words, if we know to which quadratic piece of \tilde{E}_i the point s_i^* belongs, we can determine s_i^* by setting the derivative of that piece to zero.

Lemma 4.3 *The derivative of \tilde{E}_i is a continuous monotonically increasing piecewise-linear function.*

Proof. Since by Lemma 4.1, \tilde{E}_i is strictly convex and piecewise quadratic, its derivative is monotonically increasing. To prove the continuity of the derivative, it is enough to verify this at the “new” breakpoints of \tilde{E}_i , i.e. at s_i^* and $s_i^* + \gamma$. The continuity of the derivative can be argued inductively

by observing that both mappings of the breakpoints from \tilde{E}_{i-1} to \tilde{E}_i (identity or shifting by γ) preserve the continuity of the function and its derivative.

Notice that at s_{i-1}^* , both the left and right derivatives of a quadratic piece of \tilde{E}_{i-1} becomes zero. From the definition of \tilde{E}_i in (2), the left derivative of the function $\tilde{E}_i - (x - t_{i+1})^2$ at s_{i-1}^* agrees with the left derivative of \tilde{E}_{i-1} at the same point and is therefore zero. On the other hand, the right derivative of the function $\tilde{E}_i - (x - t_{i+1})^2$ is zero at $s_{i-1}^* + \gamma$ and agrees with its right derivative at the same point as the function is constant on the interval $[s_{i-1}^*, s_{i-1}^* + \gamma]$. This means that the derivative of $\tilde{E}_i - (x - t_{i+1})^2$ is continuous at s_{i-1}^* and therefore the same holds for \tilde{E}_i . A similar argument establishes the continuity of \tilde{E}_i at $s_{i-1}^* + \gamma$. ■

Let F_i denote the derivative of \tilde{E}_i . Using (2), we can write the following recurrence for F_i :

$$F_{i+1} = 2(x - t(v_{i+1})) + \hat{F}_i(x) \quad (3)$$

where

$$\hat{F}_i(x) = \begin{cases} F_i(x - \gamma) & x > s_i^* + \gamma, \\ 0 & x \in [s_i^*, s_i^* + \gamma], \\ F_i(x) & x < s_i^*. \end{cases} \quad (4)$$

if we set $F_0 = 0$. As mentioned above, s_i^* is simply the solution of $F_i(x) = 0$, which, by Lemma 4.3 always exists and is unique. Intuitively, \hat{F}_i is obtained from F_i by splitting it at s_i^* , shifting the right part by γ , and connecting the two pieces by a horizontal edge from s_i^* to $s_i^* + \gamma$ (lying on the x -axis).

In order to find s_i^* efficiently, we use an ACT $\mathcal{T}(F_i)$ to represent F_i . It takes $O(\log |F_i|) = O(\log i)$ time to compute $s_i^* = \text{EVALUATE}^{-1}(0)$ on $\mathcal{T}(F_i)$. Once s_i^* is computed, we store it in a separate array for back-solving through Lemma 4.2. We turn $\mathcal{T}(F_i)$ into $\mathcal{T}(\hat{F}_i)$ by performing a sequence of $\text{INSERT}((s_i^*, 0))$, $\text{INTERVAL}_x(\psi, s_i^*, \infty)$, and $\text{INSERT}((s_i^* + \gamma, 0))$ operations on $\mathcal{T}(F_i)$ where $\psi(q) = q + \binom{\gamma}{0}$; the two insert operations add the breakpoints at s_i^* and $s_i^* + \gamma$ and the interval operation shifts the portion of F_i to the right of s_i^* by γ . We then turn $\mathcal{T}(\hat{F}_i)$ into $\mathcal{T}(F_{i+1})$ by performing $\text{AFFINE}(\phi_{i+1})$ operation on $\mathcal{T}(\hat{F}_i)$ where $\phi_{i+1}(q) = Mq + c$ where $M = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$ and $c = \begin{pmatrix} 0 \\ -2t(v_{i+1}) \end{pmatrix}$.

Given ACT $\mathcal{T}(F_i)$, s_i^* and $\mathcal{T}(F_{i+1})$ can be computed in $O(\log i)$ time. Hence, we can compute s_1^*, \dots, s_n^* in $O(n \log n)$ time. By Lemma 4.2, we can conclude the following.

Theorem 4.4 *Given a path $P = (V, A)$, a function $t \in \mathbb{R}^V$, and a constant γ , the γ -Lipschitz isotonic regression of t on P can be found in $O(n \log n)$ time.*

Update operation. We define a procedure $\text{UPDATE}(\mathcal{T}(\hat{F}_i), t(v_{i+1}), \gamma)$ that encapsulates the process of turning $\mathcal{T}(\hat{F}_i)$ into $\mathcal{T}(\hat{F}_{i+1})$ and returning s_{i+1}^* . Specifically, it performs $\text{AFFINE}(\phi_{i+1})$ of $\mathcal{T}(\hat{F}_i)$ to produce $\mathcal{T}(F_{i+1})$, then it outputs $s_{i+1}^* = \text{EVALUATE}^{-1}(0)$ on $\mathcal{T}(F_{i+1})$, and finally a sequence of $\text{INSERT}((s_{i+1}^*, 0))$, $\text{INTERVAL}(\psi, s_{i+1}^*, \infty)$, and $\text{INSERT}((s_{i+1}^* + \gamma, 0))$ operations on $\mathcal{T}(F_{i+1})$ to get $\mathcal{T}(\hat{F}_{i+1})$. Performed on $\mathcal{T}(F)$ where F has n breakpoints, an UPDATE takes $O(\log n)$ time. An $\text{UNUPDATE}(\mathcal{T}(\hat{F}_{i+1}), t(v_{i+1}), \gamma)$ reverts the affects of an $\text{UPDATE}(\mathcal{T}(\hat{F}_i), t(v_{i+1}), \gamma)$. This requires that s_{i+1}^* is stored for the reverted version. Similarly, we can perform $\text{UPDATE}(\mathcal{S}, t(v_i), \gamma)$ and $\text{UNUPDATE}(\mathcal{S}, t(v_i), \gamma)$ on a tree set \mathcal{S} , in $O(k \log^2 n)$ time, the bottleneck coming from $\text{EVALUATE}^{-1}(0)$.

Lemma 4.5 *Given $\mathcal{T}(\hat{F}_i)$, $\text{UPDATE}(\mathcal{T}(\hat{F}_i), t(v_i), \gamma)$ and $\text{UNUPDATE}(\mathcal{T}(\hat{F}_i), t(v_i), \gamma)$ take $O(\log n)$ time.*

Proof. Since $F_{i+1}(x) = 2(x - t(v_{i+1})) + \hat{F}_i(x)$, for a breakpoint \hat{b}_i of \hat{F}_i , we get:

$$\begin{pmatrix} b_{i+1} \\ F_{i+1}(b_{i+1}) \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} \hat{b}_i \\ \hat{F}_i(\hat{b}_i) \end{pmatrix} + \begin{pmatrix} 0 \\ -2t(v_{i+1}) \end{pmatrix},$$

where \hat{b}_i becomes breakpoint b_{i+1} .

We can now compute $s_{i+1}^* = \text{EVALUATE}^{-1}(0)$ on $\mathcal{T}(F_{i+1})$.

Rewriting (4) for a breakpoint b_{i+1} of F_{i+1} that is neither of s_{i+1}^* or $s_{i+1}^* + \gamma$, using the fact that

$$\hat{b}_{i+1} = \begin{cases} b_{i+1} + \gamma & b_{i+1} > s_{i+1}^* \\ b_{i+1} & b_{i+1} < s_{i+1}^* \end{cases}, \quad (5)$$

where b_{i+1} is the breakpoint of F_{i+1} that has become \hat{b}_{i+1} in \hat{F}_{i+1} , we get:

$$\hat{F}_{i+1}(\hat{b}_{i+1}) = \begin{cases} F_{i+1}(b_{i+1}) & b_{i+1} > s_{i+1}^* \\ F_{i+1}(b_{i+1}) & b_{i+1} < s_{i+1}^*. \end{cases} \quad (6)$$

We next combine and rewrite equations (5) and (6) as follows:

$$\begin{pmatrix} \hat{b}_{i+1} \\ \hat{F}_{i+1}(\hat{b}_{i+1}) \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} b_{i+1} \\ F_{i+1}(b_{i+1}) \end{pmatrix} + \begin{cases} \begin{pmatrix} \gamma \\ 0 \end{pmatrix} & b_{i+1} > s_{i+1}^* \\ \begin{pmatrix} 0 \\ 0 \end{pmatrix} & b_{i+1} < s_{i+1}^* \end{cases}$$

The new breakpoints, i.e. the pairs $(s_i^*, 0)$ and $(s_i^* + \gamma, 0)$, should be inserted into $\mathcal{T}(F_i)$.

To perform $\text{UNUPDATE}(\mathcal{T}(F_i), t(v_i), \gamma)$, we simply perform the inverse of all of these operations.

■

5 LUR on Paths

Let $P = (V, A)$ be an *undirected* path where $V = \{v_1, \dots, v_n\}$ and $A = \{\{v_i, v_{i+1}\}, 1 \leq i < n\}$ and given $t \in \mathbb{R}^V$. For $v_i \in V$ let $P_i = (V, A_i)$ be a directed graph in which all edges are directed towards v_i ; that is, for $j < i$, $(v_j, v_{j+1}) \in A_i$ and for $j > i$, $(v_j, v_{j-1}) \in A_i$. For each $i = 1, \dots, n$, let $\Gamma_i = \Gamma(P_i, V) \subseteq \mathbb{R}^V$ and let $\sigma_i = \arg \min_{s \in \Gamma_i} E_V(s)$. If $\kappa = \arg \min_i E_V(\sigma_i)$, then σ_κ is the γ -LUR of t on P .

We can find σ_κ in $O(n \log^2 n)$ time by solving the LIR problem, then traversing the path while maintaining the optimal solution using UPDATE and UNUPDATE . Specifically, for each $i = 1, \dots, n$, let $V_i^- = \{v_1, \dots, v_{i-1}\}$ and $V_i^+ = \{v_{i+1}, \dots, v_n\}$. For P_i , let \hat{F}_{i-1}^- , \hat{F}_{i+1}^+ be the functions on directed paths $P[V_i^-]$ and $P[V_i^+]$, respectively, as defined in (3). Set $\bar{F}_i(x) = 2(x - t(v_i))$. Then the function $F_i(x) = dE_{v_i}(x)/dx$ can be written as $F_i(x) = \hat{F}_{i-1}^-(x) + \hat{F}_{i+1}^+(x) + \bar{F}_i(x)$. We store F_i as the tree set $S_i = \{\mathcal{T}(\hat{F}_{i-1}^-), \mathcal{T}(\hat{F}_{i+1}^+), \mathcal{T}(\bar{F}_i)\}$. By performing $\text{EVALUATE}^{-1}(0)$ we can compute s_i^* in $O(\log^2 n)$ time (the rate limiting step), and then we can compute $E_V(s_i^*)$ in $O(\log n)$ time using $\text{INTEGRATE}(s_i^*)$. Assuming we have $\mathcal{T}(\hat{F}_{i-1}^-)$ and $\mathcal{T}(\hat{F}_{i+1}^+)$, we can construct $\mathcal{T}(\hat{F}_i^-)$ and $\mathcal{T}(\hat{F}_{i+2}^+)$ in $O(\log n)$ time by performing $\text{UPDATE}(\mathcal{T}(\hat{F}_{i-1}^-), t(v_i), \gamma)$ and $\text{UNUPDATE}(\mathcal{T}(\hat{F}_{i+1}^+), t(v_{i+1}), \gamma)$. Since $S_1 = \{\mathcal{T}(F_0^-), \mathcal{T}(F_2^+), \mathcal{T}(\bar{F}_1)\}$ is constructed in $O(n \log n)$ time, finding κ by searching all n tree sets takes $O(n \log^2 n)$ time.

Theorem 5.1 *Given an undirected path $P = (V, A)$ and a $t \in \mathbb{R}^V$ together with a real $\gamma \geq 0$, the γ -LUR of t on P can be found in $O(n \log^2 n)$ time.*

6 LIR on Rooted Trees

Let $T = (V, A)$ be a rooted tree with root r and let for each vertex v , $T_v = (V_v, A_v)$ denote the subtree of T rooted at v . Similar to the case of path LIR, for each vertex $v \in V$ we use the shorthands $E_v = E_{V_v}$ and $\tilde{E}_v = E_{T_v}^v$. Since the subtrees rooted at distinct children of a node v are disjoint, one can write an equation corresponding to (1) in the case of paths, for any vertex v of T :

$$\tilde{E}_v(x) = (x - t(v))^2 + \sum_{u \in \delta^-(v)} \min_{x - \gamma \leq y \leq x} \tilde{E}_u(y), \quad (7)$$

where $\delta^-(v) = \{u \in V \mid (u, v) \in A\}$. An argument similar to that of Lemma 4.3 together with Lemma 2.1 implies that for every $v \in V$, the function \tilde{E}_v is convex and piecewise quadratic, and its derivative F_v is continuous, monotonically increasing, and piecewise linear. We can prove that F_v satisfies the following recurrence where \hat{F}_u is defined analogously to \hat{F}_i in (4):

$$F_v(x) = 2(x - t(v)) + \sum_{u \in \delta^-(v)} \hat{F}_u(x). \quad (8)$$

Thus to solve the LIR problem on a tree, we post-order traverse the tree (from the leaves toward the root) and when processing a node v , we compute and sum up the linear functions \hat{F}_u for all children u of v and use the result in (8) to compute the function F_v . We then solve $F_v(x) = 0$ to find s_v^* . As in the case of path LIR, \hat{F}_v can be represented by an ACT $\mathcal{T}(\hat{F}_v)$. For simplicity, we assume each non-leaf vertex v has two children $h(v)$ and $l(v)$, where $|\hat{F}_{h(v)}| \geq |\hat{F}_{l(v)}|$. We call $h(v)$ *heavy* and $l(v)$ *light*. Given $\mathcal{T}(\hat{F}_{h(v)})$ and $\mathcal{T}(\hat{F}_{l(v)})$, we can compute $\mathcal{T}(\hat{F}_v)$ and s_v^* with the operation $\text{UPDATE}(\mathcal{T}(\hat{F}_{h(v)} + \hat{F}_{l(v)}), t(v), \gamma)$ in $O(|\hat{F}_{l(v)}|(1 + \log |\hat{F}_{h(v)}| / \log |\hat{F}_{l(v)}|))$ time. The merging of two functions dominates the time for the update.

Theorem 6.1 *Given a rooted tree $T = (V, A)$ and a function $t \in \mathbb{R}^V$ together with a Lipschitz constant γ , one can find in $O(n \log n)$ time, the γ -Lipschitz isotonic regression of t on T .*

Proof. Let T be rooted at r , it is sufficient to bound the cost of all calls to UPDATE for each vertex: $U(T) = \sum_{v \in T} |\hat{F}_{l(v)}|(1 + \log |\hat{F}_{h(v)}| / \log |\hat{F}_{l(v)}|)$.

For each leaf vertex u consider the path to the root $P_{u,r} = \langle u = v_1, v_2, \dots, v_s = r \rangle$. For each light vertex v_i along the path, let its *value* be $\nu(v_i) = 2 + \lceil \log(|\hat{F}_{h(v_{i+1})}| / |\hat{F}_{l(v_{i+1})=v_i}|) \rceil$. Also, let each heavy vertex v_i along the path have *value* $\nu(v_i) = 0$. For any root vertex u , the sum of values $\Upsilon_u = \sum_{v_i \in P_{u,r}} \nu(v_i) \leq 3 \log n$, since $|\hat{F}_{v_i}| 2^{\nu(v_i)-1} \leq |\hat{F}_{v_{i+1}}|$ for each light vertex and $|\hat{F}_r| = n$.

Now we can argue that the contribution of each leaf vertex u to $U(T)$ is at most Υ_u . Observe that for $v \in T$, then $\nu(l(v))$ is the same for any leaf vertex u such that $l(v) \in P_{u,r}$. Thus we charge $\nu(l(v))$ to v for each u in the subtree rooted at $l(v)$. Since $\nu(l(v)) \geq 2 + \log(|\hat{F}_{h(v)}| / |\hat{F}_{l(v)}|) \geq 1 + \log |\hat{F}_{h(v)}| / \log |\hat{F}_{l(v)}|$ (for $\log |\hat{F}_{h(v)}| \geq \log |\hat{F}_{l(v)}|$), then the charges to each $v \in T$ is greater than its contribution to $U(T)$. Then for each of n leaf vertices u , we charge $\Upsilon_u \leq 3 \log n$. Hence, $U(T) \leq 3n \log n$. \blacksquare

7 LUR on Trees

Let $T = (V, A)$ be an unrooted tree with n vertices. Given any choice of $r \in V$ as a root, we say $T^{(r)}$ is the tree T rooted at r . Given a function $t \in \mathbb{R}^V$, a Lipschitz constraint γ , and a root $r \in V$, the γ -LIR regression of t on $T^{(r)}$, denoted $s^{(r)}$, can be found in $O(n \log n)$ time using Theorem 6.1. We let $\xi^{(r)} = \min_{x \in \mathbb{R}} \tilde{E}_r(x)$ as defined in (7) and let $r^* = \arg \min_{r \in V} \xi^{(r)}$. The γ -Lipschitz

unimodal regression (γ -LUR) of t on T is the function $s^{(r^*)} \in \mathbb{R}^V$. Naively, we could compute the γ -LIR in $O(n^2 \log n)$ time by invoking Theorem 6.1 with each vertex as the root. In this section we show this can be improved to $O(n \log^3 n)$.

We first choose an arbitrary vertex $\hat{r} \in V$ (assume it has at most two edges in A) as an honorary root of T . For simplicity only, we assume that every vertex of T has degree three or less. Let the weight $\omega(v)$ of a subtree of $T^{(\hat{r})}$ at v be the number of vertices below and including v in $T^{(\hat{r})}$. With respect to this honorary root, for each vertex $v \in V$, we define $p(v)$ to be the parent of v and $h(v)$ (resp. $l(v)$) to be the heavy (resp. light) child of v such that $\omega(h(v)) \geq \omega(l(v))$; \hat{r} has no parent and leaf vertices (with respect to \hat{r}) have no children. If v has only one child, label it $h(v)$. If $v = l(p(v))$ we say v is a *light* vertex, otherwise it is a *heavy* vertex. The following two properties are consequences of $\omega(h(v)) > \omega(l(v))$ for all $v \in T^{(\hat{r})}$.

(P1) $\sum_{v \in V} \omega(l(v)) = O(n \log n)$.

(P2) For $u \in V$, the path $\langle \hat{r}, \dots, u \rangle$ in $T^{(\hat{r})}$ contains at most $\lceil \log_2 n \rceil$ light vertices.

As a preprocessing step, we construct the γ -LIR of t for $T^{(\hat{r})}$ as described in Theorem 6.1, and, for each light vertex v , we store a copy of $\mathcal{T}(\hat{F}_v)$ before merging with $\mathcal{T}(\hat{F}_{h(p(v))})$. The total size of all stored ACTs is $O(n \log n)$ by property (P1). We now perform an inorder traversal of $T^{(\hat{r})}$, letting each $v \in V$ in turn be the structural root of T so we can determine $\xi^{(v)}$. As we traverse, when v is structural root, we maintain a tree set $\mathcal{S}(\hat{F}_{p(v)})$ at $p(v)$ with at most $k = \log_2 n$ trees, and ACTs $\mathcal{T}(\hat{F}_{h(v)})$ and $\mathcal{T}(\hat{F}_{l(v)})$ at $h(v)$ and $l(v)$, respectively. All functions \hat{F}_u for $u = \{p(v), l(v), h(v)\}$ are defined the same as with rooted trees, when rooted at the structural root v , for the subtree rooted at u . Given these three data structures we can compute ξ_v in $O(\log^3 n)$ time by first temporarily inserting the ACTs $\mathcal{T}(\hat{F}_{h(v)})$, $\mathcal{T}(\hat{F}_{l(v)})$, and $\mathcal{T}(\hat{F}_v)$ into $\mathcal{S}(\hat{F}_{p(v)})$, thus describing the function F_v for a rooted tree with v as the root, and then performing an $s^{(v)} = \text{EVALUATE}^{-1}(0)$ and $\text{INTEGRATE}(s^{(v)})$ on $\mathcal{S}(F_v)$. Recall $\bar{F}_v(x) = 2(x - t(v))$. Thus to complete this algorithm, we need to show how to maintain these three data structures as we traverse $T^{(\hat{r})}$.

During the inorder traversal of $T^{(\hat{r})}$ we need to consider 3 cases: letting $h(v)$, $l(v)$, or $p(v)$ be the next structural root, as shown in Figure 4.

- (a) $h(v)$ is the next structural root: To create $\mathcal{S}(\hat{F}_{p(h(v))}) = \mathcal{S}(\hat{F}_v)$ we add $\mathcal{T}(\hat{F}_{l(v)})$ to the tree set $\mathcal{S}(\hat{F}_{p(v)})$ through a $\text{MERGE}(\mathcal{S}(\hat{F}_{p(v)}), \mathcal{T}(\hat{F}_{l(v)}))$ operation and perform $\text{UPDATE}(\mathcal{S}(\hat{F}_{p(v)}), t(v), \gamma)$. The ACT $\mathcal{T}(\hat{F}_{l(h(v))})$ is stored at $l(h(v))$. To create the ACT $\mathcal{T}(\hat{F}_{h(h(v))})$ we perform $\text{UNUPDATE}(\mathcal{T}(\hat{F}_{h(v)}), t(h(v)))$ and then set $\mathcal{T}(\hat{F}_{h(h(v))}) \leftarrow \mathcal{T}(\hat{F}_{h(h(v))}) - \hat{F}_{l(h(v))}$, using $\mathcal{T}(\hat{F}_{l(h(v))})$.
- (b) $l(v)$ is the next structural root: To create $\mathcal{S}(\hat{F}_{p(l(v))})$ we perform $\text{INCLUDE}(\mathcal{S}(\hat{F}_{p(v)}), \mathcal{T}(\hat{F}_{h(v)}))$ and then $\text{UPDATE}(\mathcal{S}(\hat{F}_{p(v)}), t(v), \gamma)$. The ACTs for $h(l(v))$ and $l(l(v))$ are produced the same as above.
- (c) $p(v)$ is the next structural root: To create $\mathcal{T}(\hat{F}_v)$ we set $\mathcal{T}(\hat{F}_v) \leftarrow \mathcal{T}(\hat{F}_{h(v)} + \hat{F}_{l(v)})$ and then perform $\text{UPDATE}(\mathcal{T}(\hat{F}_v), t(v), \gamma)$. To create $\mathcal{S}(\hat{F}_{p(p(v))})$ we first set $\mathcal{S}(\hat{F}_{p(p(v))}) \leftarrow \text{UNUPDATE}(\mathcal{S}(\hat{F}_{p(v)}), t(p(v)), \gamma)$. Then if $v = h(p(v))$, we perform $\text{UNMERGE}(\mathcal{S}(\hat{F}_{p(p(v))}), \mathcal{T}(\hat{F}_{l(p(v))}))$, and $\mathcal{T}(\hat{F}_{l(p(v))})$ is stored at $l(p(v))$. Otherwise $v = l(p(v))$, and we perform $\text{UNINCLUDE}(\mathcal{S}(\hat{F}_{p(p(v))}), \mathcal{T}(\hat{F}_{h(p(v))}))$ to get $\mathcal{S}(\hat{F}_{p(p(v))})$, and $\mathcal{T}(\hat{F}_{h(p(v))})$ is the byproduct.

We observe that on a traversal through any vertex $v \in V$, the number of ACTs in the tree set $\mathcal{S}(\hat{F}_{p(v)})$ only increases when $l(v) = u$ is the new structural root. Furthermore, when the traversal reverses this path so that $p(u)$ is the new root where $u = l(p(u))$, the same ACT is removed from the tree set. Thus, when any vertex $v \in V$ is the structural root, we can bound the number of ACTs




Figure 4: Illustration of $T^{(\hat{r})}$ ordered with respect to honorary root \hat{r} . This is redrawn, ordered with respect to structural root as v . The tree is redrawn with structural root as $h(v)$ after case (a), $l(v)$ after case (b), and $p(v)$ after case (c).

in the tree set $\mathcal{S}(\hat{F}_{p(v)})$ by the number of light vertices on the path from \hat{r} to v . Thus property (P2) implies that $\mathcal{S}(\hat{F}_{p(v)})$ has at most $\log_2 n$ ACTs.

We can now bound the runtime of the full traversal algorithm described above. Each vertex $v \in V$ is visited as the structural root $O(1)$ time and each visit results in $O(1)$ UPDATE, UNUPDATE, INCLUDE, UNINCLUDE, MERGE, and UNMERGE operations. Performing a single UPDATE or UNUPDATE on a tree set takes at most $O(k \log^2 n) = O(\log^3 n)$ time, and each INCLUDE and UNINCLUDE operation takes $O(1)$ time, so the total time is $O(n \log^3 n)$. We now need to bound the global costs of all MERGE and UNMERGE operations. In proving Theorem 6.1 we showed that the cost of performing a MERGE of all light vertices into heavy vertices as we build $\mathcal{T}(\hat{F}_{\hat{r}})$ takes $O(n \log^2 n)$. Since at each vertex $v \in V$ we only MERGE and UNMERGE light vertices, the total time of all MERGE and UNMERGE operations is again $O(n \log^2 n)$.

Thus we can construct $\xi^{(v)}$ for each possible structural root, and using $r^* = \arg \min_{v \in V} \xi^{(v)}$ as the optimal root we can invoke Theorem 6.1 to construct the γ -LIR of t on $T^{(r^*)}$.

Theorem 7.1 *Given an unrooted tree $T = (V, A)$ and a function $t \in \mathbb{R}^V$ together with a Lipschitz constraint γ , we can find in $O(n \log^3 n)$ time the γ -Lipschitz unimodal regression of t on T .*

8 Conclusion

In this paper we defined the Lipschitz isotonic/unimodal regression problems on DAGs and provided efficient algorithms for solving it on the special cases where the graph is a path or a tree. Our ACT data structure has proven to be a powerful tool in our approach.

Our algorithms can be generalized in a number of ways, as listed below, without affecting their asymptotic running times.

- One can specify different Lipschitz value γ for each Lipschitz constraint, thus writing the Lipschitz constraint involving s_i and s_{i+1} as $s_{i+1} \leq s_i + \gamma_i$.
- Isotonicity constraints can be regarded as (backward) Lipschitz constraints for which the Lipschitz constant is zero. One can allow this constant to be chosen arbitrarily and indeed permit different constants for different isotonicity constraints. In other words, the constraint $s_i \leq s_{i+1}$ can be replaced with $s_{i+1} \geq s_i + \delta_i$.
- The ℓ_2 norm $\|\cdot\|$ can be replaced with a “weighted” version $\|\cdot\|_{\lambda}$ for any $\lambda : V \rightarrow \mathbb{R}$ by defining for a function $s : V \rightarrow \mathbb{R}$:

$$\|s\|_{\lambda}^2 = \sum_{v \in V} \lambda(v) s(v)^2.$$

The most important open problem is solving LIR problem on general DAGs. The known algorithm for solving IR on DAGs runs in $O(n^4)$ and does not lend itself to the Lipschitz generalization. The difficulty here is to compute the function \tilde{E}_i (or in fact F_i) when the i th vertex is processed and comes from the fact that the F_{i_j} function, $j = 1, \dots, k$, where i_1, \dots, i_k are vertices with incoming edges to i , are not independent. Independently, it is an important question whether IR can be solved more efficiently. In particular, the special case of the problem where the given DAG is planar has important applications in terrain simplification.

Acknowledgements

We thank anonymous reviewers for many helpful remarks that helped improve the presentation of this paper and we thank Günter Rote for suggesting the improvement in the analysis of $\mathcal{T}(F + G)$, similar to his paper [10].

References

- [1] P. K. Agarwal, L. Arge, and K. Yi. I/O-efficient batched union-find and its applications to terrain analysis. In *Proceedings 22nd ACM Symposium on Computational Geometry*, pages 167–176, 2006.
- [2] S. Angelov, B. Harb, S. Kannan, and L.-S. Wang. Weighted isotonic regression under the l_1 norm. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 783–791, 2006.
- [3] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. Manuscript, INRIA, 2008.
- [4] M. Ayer, H. D. Brunk, G. M. Ewing, W. T. Reid, and E. Silverman. An empirical distribution function for sampling with incomplete information. *Annals of Mathematical Statistics*, 26:641–647, 1955.
- [5] C. Bajaj, V. Pascucci, and D. Schikore. Visualization of scalar topology for structural enhancement. In *IEEE Visualization*, pages 51–58, 1998.
- [6] R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk. *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, 1972.
- [7] P. Bremer, H. Edelsbrunner, B. Hamann, and V. Pascucci. A multi-resolution data structure for two-dimensional morse functions. In *Proceedings 14th IEEE Visualization Conference*, pages 139–146, 2003.
- [8] M. R. Brown and R. E. Tarjan. A fast merging algorithm. *Journal of Algorithms*, 15:416–446, 1979.
- [9] H. D. Brunk. Maximum likelihood estimates of monotone parameters. *Annals of Mathematical Statistics*, 26:607–616, 1955.
- [10] K. Buchin, S. Cabello, J. Gudmundsson, M. Löffler, J. Luo, G. Rote, R. I. Silveira, B. Speckmann, and T. Wollé. Finding the most relevant fragments in networks. Manuscript, October 2009.

- [11] A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge, and H. Mitsova. Terrastream: from elevation data to watershed hierarchies. In *15th ACM International Symposium on Advances in Geographic Information Systems*, 2007.
- [12] H. Edelsbrunner and J. Harer. *Persistent Homology: A Survey*. Number 453 in Contemporary Mathematics. American Mathematical Society, 2008.
- [13] H. Edelsbrunner, J. Harer, and A. Zomorodia. Hierarchical morse - smale complexes for piecewise linear 2-manifolds. *Discrete & Computational Geometry*, 30(1):87–107, 2003.
- [14] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proceedings 41st Annual Symposium on Foundations on Computer Science*, pages 454–463, 2000.
- [15] H. Edelsbrunner, D. Morozov, and V. Pascucci. Persistence-sensitive simplification functions on 2-manifolds. In *Proceedings 22nd ACM Symposium on Computational Geometry*, pages 127–134, 2006.
- [16] G. N. Frederickson and D. B. Johnson. The complexity of selection and ranking in $x + y$ and matrices with sorted columns. *J. Comput. Syst. Sci.*, 24(2):192–208, 1982.
- [17] S. J. Grotzinger and C. Witzgall. Projection onto order simplexes. *Applications of Mathematics and Optimization*, 12:247–270, 1984.
- [18] I. Guskov and Z. J. Wood. Topological noise removal. In *Graphics Interface*, pages 19–26, 2001.
- [19] W. S. Jewel. Isotonic optimization in tariff construction. *ASTIN Bulletin*, 8(2):175–203, 1975.
- [20] W. L. Maxwell and J. A. Muckstadt. Establishing consistent and realistic reorder intervals in production-distribution systems. *Operations Research*, 33:1316–1341, 1985.
- [21] X. Ni, M. Garland, and J. C. Hart. Fair Morse functions for extracting the topological structure of a surface mesh. *ACM Transactions Graphaphics*, 23:613–622, 2004.
- [22] P. M. Pardalos and G. Xue. Algorithms for a class of isotonic regression problems. *Algorithmica*, 23:211–222, 1999.
- [23] F. P. Preparata and M. I. Shamos. *Computational geometry: an introduction*. Springer-Verlag, New York, NY, USA, 1985.
- [24] T. Robertson, F. T. Wright, and R. L. Dykstra. *Order Restricted Statistical Inference*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, 1988.
- [25] P. Soille. Morphological carving. *Pattern Recognition Letters*, 25:543–550, 2004.
- [26] P. Soille. Optimal removal of spurious pits in grid digital elevation models. *Water Resources Research*, 40(12), 2004.
- [27] P. Soille, J. Vogt, and R. Colombo. Carbing and adaptive drainage enforcement of grid digital elevation models. *Water Resources Research*, 39(12):1366–1375, 2003.
- [28] J. Spouge, H. Wan, and W. J. Wilbur. Least squares isotonic regression in two dimensions. *Journal of Optimization Theory and Applications*, 117:585–605, 2003.

- [29] Q. F. Stout. Optimal algorithms for unimodal regression. *Computing Science and Statistics*, 32:348–355, 2000.
- [30] Q. F. Stout. l_∞ isotonic regression. Unpublished Manuscript, 2008.
- [31] Q. F. Stout. Unimodal regression via prefix isotonic regression. *Computational Statistics and Data Analysis*, 53:289–297, 2008.
- [32] W. A. Thompson, Jr. The problem of negative estimates of variance components. *Annals of Mathematical Statistics*, 33:273–289, 1962.
- [33] M. J. van Kreveld and R. I. Silveira. Embedding rivers in polyhedral terrains. In *Proceedings 25th Symposium on Computational Geometry*, 2009.
- [34] A. Zomorodian. Computational topology. In M. Atallah and M. Blanton, editors, *Algorithms and Theory of Computation Handbook*, page (In Press). Chapman & Hall/CRC Press, Boca Raton, FL, second edition, 2009.
- [35] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

A Integrate Operation

Lemma A.1 *We can modify an ACT datastructure for $\mathcal{T}(F)$ with $|F| = n$ and $F(x) = dE(x)/dx$, so for any $x \in R$ we can calculate $E(x) = \text{INTEGRATE}(x)$ in $O(\log n)$ time, and so the cost of maintaining $\mathcal{T}(F)$ is asymptotically equivalent to without this feature.*

Proof. For any $a \in \mathbb{R}$, we calculate $\text{INTEGRATE}(a) = E(a)$ by adding $E(b_1)$ to $\int_{b_1}^a F(x) dx$, where b_1 is the leftmost breakpoint in F . Maintaining $E(b_1)$ is easy and is explained first. Calculating $\int_{b_1}^a F(x) dx$ requires storing partial integrals and other information at each node of $\mathcal{T}(F)$ and is more involved.

We explain how to maintain the value $E(b_1)$ in the context of our γ -LIR problems in order to provide more intuition, but a similar technique can be used for more general ACTs. We maintain values c_1, c_2, c_3 such that $c_1x^2 + c_2x + c_3 = \sum_{i=1}^k (x - t_i)^2$ for the $k = O(n)$ values of t_i . This is done by adding 1, $-2t_i$, and t_i^2 to c_1 , c_2 , and c_3 every time a new t_i is processed. Or if two trees are merged, this can be updated in $O(1)$ time. The significance of this equation is that $E(b_1) = c_1b_1^2 + c_2b_1 + c_3$ because the isotonic condition is enforced for all edges in A , forcing all values s_i in the γ -LIR to be equal.

When we construct $\mathcal{T}(F)$ we keep the following extra information at each node: each node z whose subtree spans breakpoints b_i to b_j maintains the integral $I(z) = \int_{b_i}^{b_j} \Phi_z^{-1}(F(x) - F(b_i)) dx$ (this is done for technical reasons; subtracting $F(b_i)$ ensures $F(x) - F(b_i) \geq 0$ for $x \in [b_i, b_j]$ since F is isotonic), the width $W(z) = \Phi_z^{-1}(b_j) - \Phi_z^{-1}(b_i)$, and the height $H(z) = \Phi_z^{-1}(F(b_j)) - \Phi_z^{-1}(F(b_i))$. Recall that Φ_z is the composition of affine transformations stored at the nodes along the path from the root to z . We build these values ($I(z)$, $W(z)$, and $H(z)$) for each $z \in \mathcal{T}(F)$ from the bottom of $\mathcal{T}(F)$ up, so we apply all appropriate affine transformations ϕ_{z_j} for decedents z_j of z , but not yet those of ancestors of z . In particular, if node z has left child z_l and right child z_k , then we can calculate

$$\begin{aligned} I(z) &= \phi_{z_l}(I(z_l)) + \phi_{z_k}(I(z_k)) + \phi_{z_l}(H(z_l)) \cdot \phi_{z_k}(W(z_k)) \text{ (see Figure 5),} \\ H(z) &= \phi_{z_l}(H(z_l)) + \phi_{z_k}(H(z_k)), \text{ and} \\ W(z) &= \phi_{z_l}(W(z_l)) + \phi_{z_k}(W(z_k)). \end{aligned}$$

When a new breakpoint b is inserted or removed there are $O(\log n)$ nodes that are either rotated or are on the path from the node storing b to the root. For each affected node z we recalculate $I(z)$, $W(z)$ and $H(z)$ in a bottom up manner using z 's children once they have been updated themselves. Likewise, tree rotations require that $O(1)$ nodes are updated. So to demonstrate that the maintenance cost does not increase, we just need to show how to update $I(z)$, $H(z)$ and $W(z)$ under an affine transformation $\phi_z(z)$ in $O(1)$ time.

For affine transformation ϕ_z applied at node z (spanning breakpoints b_i through b_j), we can decompose it into three components: translation, scaling, and rotation. The translation does not affect $I(z)$, $W(z)$, or $H(z)$. The scaling can be decomposed into scaling of width w and of height h . We can then update $W(z) \leftarrow w \cdot W(z)$, $H(z) \leftarrow h \cdot H(z)$, and $I(z) \leftarrow w \cdot h \cdot I(z)$. To handle rotation, since translation does not affect $I(z)$, $W(z)$, or $H(z)$, we consider the data translated by $\Phi_z^{-1}((-b_i, 0))$ so that $\Phi_z^{-1}((b_i, F(b_i) - F(b_i))) = (0, 0)$ is the origin and $\Phi_z^{-1}((b_j, F(b_j) - F(b_i))) = (x, y) = (W(z), H(z))$. Then we consider rotations about the origin, as illustrated in Figure 5. After a rotation by a positive angle θ , the point (x, y) becomes $(x', y') = (x \cos \theta - y \sin \theta, x \sin \theta + y \cos \theta)$. This updates $W(z) \leftarrow x'$ and $H(z) \leftarrow y'$. We can update $I(z)$ by subtracting the area A_- that is no longer under the integral $y \cdot (y \tan \theta)/2$ and adding the area A_+ newly under the integral $x' \cdot (x' \tan \theta)/2$ as illustrated in Figure 5. Thus $I(z) \leftarrow I(v) + x' \cdot (x' \tan \theta)/2 - y \cdot (y \tan \theta)/2$. A

similar operation exists for a negative angle θ , using the fact that F must remain monotone. As desired all update steps take $O(1)$ time.




Figure 5: Left: Illustration of the change in integral caused by a rotate of θ degrees counterclockwise. Right: Calculation of I_z from children z_k and z_l of z .

Now for a value $a \in \mathbb{R}$ we can calculate $\text{INTEGRATE}(a) = E(a)$ as follows. Let $z_a \in \mathcal{T}(F)$ be the node and b be the breakpoint associated with $\text{PRED}_x(a)$. We consider the path $Z = \langle z_a, \dots, z_1, z_0 \rangle$ from z_a to the root z_0 of $\mathcal{T}(F)$, and will inductively build an integral along Z . We calculate $I_{z_0} = \int_{b_1}^a \phi_{z_0}(F(x) - F(b_1)) dx$ similar to the way we calculated $I(z_0)$, but just over the integral $[b_1, a]$. As a base case, let $I_{z_a} = \int_b^a \Phi_{z_a}^{-1}(F(x) - F(b)) dx$ (which is easy to calculate because F is linear in this range) and let $W_{z_a} = \Phi_{z_a}^{-1}(a) - \Phi_{z_a}^{-1}(b)$ describe the integral and width associated with the range $[b, a]$ and z_a . For a node $z \in Z$ that spans breakpoints b_i through b_j , we calculate an integral $I_z = \int_{b_i}^a \Phi_z^{-1}(F(x) - F(b_i)) dx$ and width $W_z = \Phi_z^{-1}(a) - \Phi_z^{-1}(b_i)$ using its two children: z_l and z_k . Let z_k be the right child which lies in Z and for which we had inductively calculated I_{z_k} and W_{z_k} . Then let $I_z = \phi_{z_l}(I(z_l)) + \phi_{z_k}(I_{z_k}) + \phi_{z_l}(H(z_l)) \cdot \phi_{z_k}(W_{z_k})$ (as illustrated in Figure 5) and let $W_z = \phi_{z_l}(W(z_l)) + \phi_{z_k}(W_{z_k})$. Once $\phi_{z_0}(I_{z_0})$ and $\phi_{z_0}(W_{z_0})$ are calculated at the root, we need to adjust for the $F(b_1)$ term in the integral. We add $\phi_{z_0}(W_{z_0}) \cdot F(b_1)$ to $\phi_{z_0}(I(z_0)) = \int_{b_1}^a (F(x) - F(b_1)) dx$ to get $\int_{b_1}^a F(x) dx$. So finally we set $\text{INTEGRATE}(a) = E(b_1) + \phi_{z_0}(I_{z_0}) + \phi_{z_0}(W_{z_0}) \cdot F(b_1)$. Since the path Z is of length at most $O(\log n)$, $E(a)$ can be calculated in $O(\log n)$ time. ■