Simplification of Sampled Scalar Fields by Removal of Extrema
Brooks, Martin; Watson, L.

https://doi.org/10.4224/5763548

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

ERB-1147

Simplification of Sampled Scalar Fields by Removal of Extrema

Brooks, M.F., and Watson, L.
June 2007

NRC 47422

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

NRC·CNRC

*Simplification of Sampled Scalar Fields by Removal of Extrema**

Brooks, M.F., and Watson, L.
2007

* June 19, 2007. 28 Pages. ERB-1147. NRC 49359.

TECHNICAL REPORT

NRC-49359/ERB-1147
Printed June 2006

Simplification of Sampled Scalar Fields by Removal of Extrema

Martin Brooks
Liam Watson

Computational Video Group
Institute for Information Technology
National Research Council Canada
Montreal Road, Building M-50
Ottawa, Ontario, Canada K1A 0R6

NRC · CNRC

NRC-49359/ERB-1147
Printed June 2006

Simplification of Sampled Scalar Fields by Removal of Extrema

Martin Brooks
National Research Council Canada

Liam Watson
Université du Québec à Montréal

Abstract

We present Extremal Simplification, a rigorous basis for algorithms that simplify geometric and scientific data. The Eilenberg-Whyburn monotone-light factorization [31] provides a mathematical framework for simplification of continuous functions. We provide conditions on finite data guaranteeing uniqueness of continuous interpolations' topological structure, thereby making continuous methods available in a discrete context. Lower bounds on approximation error are derived. Extremal Simplification is compared to other scalar field simplification methods, including the Reeb graph [4, 5, 28], Morse-Smale complex [1], and the persistence diagram [11, 9].

1 Introduction

This introductory section provides an overview of Extremal Simplification, identification of contributions, and an overview of related methods. The mathematical setting is introduced in section 2. The definitions, constructions and theorem statements comprising Extremal Simplification are presented in narrative fashion in two parallel halves: first for continuous functions in sections 3 - 4; then for sample data in sections 5 - 6. Analysis of selected related simplification methods is presented in section 7. The mathematical bulk of this paper comprises the proofs of seven new theorems; proofs are relegated to a series of appendices, thus avoiding interruption of the theory's presentation.

1.1 Overview

Extremal Simplification is couched in terms of point-set topology and is built on the Eilenberg-Whyburn monotone-light factorization of continuous functions [31, 29]. The central objects are an arbitrary Peano space X, a scalar field represented as a continuous function $f : X \to \mathbb{R}$, and a sampled scalar field represented as $F : D \to \mathbb{R}$, where $D \subset X$ is a finite collection of sample locations and $F = f|D$.

Scalar field f's monotone-light factorization comprises a *middle space* M, *monotone factor* $\mu : X \to M$ and *light factor* $\lambda : M \to \mathbb{R}$, where $f = \mu \circ \lambda$. We restrict our attention to *piecewise monotone* f, where M is a finite graph. Recognizing M as the Reeb graph [24], extrema are cast as points in the middle space. However, the middle space M is more than a combinatorial graph, it is a Peano space; our development utilizes its topology and relies significantly on f's monotone and light factors.

Extremal Simplification consists of three steps: *analysis* of data F_0, resulting in a topological structure T_0; *simplification* of the topological structure T_0 by removal of extrema; and *synthesis* of data F_1 corresponding to the simplified topology. This process may be repeated, providing successively simpler approximations $F_1 \ldots F_n$ of the original data.

$$
\begin{tikzcd}
F_0 \arrow[r, "\text{analyze}"] & T_0 \arrow[d, "\text{simplify}"] \\
F_1 & T_1 \arrow[l, "\text{synthesize}"'] \arrow[d, "\text{simplify}"] \\
 & \ldots
\end{tikzcd}
$$

The topological structure T_0 is derived from a continuous function f_0 interpolating the data F_0. T_0 comprises f_0's middle space (Reeb graph) M_0 and light factor $\lambda_0 : M_0 \to \mathbb{R}$. Topological simplification of (M_0, λ_0) to (M_1, λ_1) comprises a *monotone quotient.* The diagram now looks like this:

$$
\begin{tikzcd}
F_0 \arrow[r, "\text{interpolate}"] & f_0 \arrow[r, "\text{monotone-light}", "\text{factorization}"'] & (M_0, \lambda_0) \arrow[d, "\text{quotient}"] \\
F_1 & f_1 \arrow[l, "\text{sample}"'] & (M_1, \lambda_1) \arrow[l, "\text{synthesize}"'] \arrow[d, "\text{quotient}"] \\
 & & \ldots
\end{tikzcd}
$$

In the diagram, continuous functions are arrived at by two means: interpolation and synthesis. Synthesized function f_1 has middle space and light factor (M_1, λ_1). Synthesis of f_1 may additionally utilize as input any or all of the following: the original data F_0, interpolated function f_0, and topological structure (M_0, λ_0). These inputs allow synthesis of function f_1 as desired, for example: having minimal approximation error $||f_0 - f_1||_\infty$; or having extrema collocated with the corresponding extrema of f_0.

There are many ways to interpolate data F to a continuous function on X. We restrict our attention to partitioning the domain X into patches $P_1 \ldots P_n$, where each P_i is assigned a local interpolant. Common examples in two dimensions are triangular patches with linear interpolants, and square patches with quadratic interpolants.

Interpolation is based on the sample locations D and patch geometry $\mathcal{P}$. Given data values $F : D \to \mathbb{R}$, there remains freedom to choose the patch interpolants; we capture this choice as an interpolation strategy parameter $\mathcal{I}$. For example, $\mathcal{I}$ might indicate linear patch interpolants. Thus, the 3-tuple $(\mathcal{P}, D, \mathcal{I})$ defines a unique interpolation of data F to function $f : X \to \mathbb{R}$, where $F = f|D$ and the local interpolant for patch P_i is $f|P_i$.

The diagram now looks like this:

$$
\begin{tikzcd}
F_0 \arrow[r, "\text{choose}"] & (\mathcal{P}_0, D_0, \mathcal{I}_0) \arrow[r, "\text{interpolate}"] & f_0 \arrow[r, "\text{monotone-light}", "\text{factorization}"'] & (M_0, \lambda_0) \arrow[d, "\text{quotient}"] \\
F_1 & (\mathcal{P}_1, D_1, \mathcal{I}_1) \arrow[l, "\text{sample}"'] & f_1 \arrow[l, "\text{choose}"'] & (M_1, \lambda_1) \arrow[l, "\text{synthesize}"'] \arrow[d, "\text{quotient}"] \\
 & & & \cdots
\end{tikzcd}
$$

The simplified data F_1 may be defined on the same sample locations D_0 as the original data F_0, or it may be defined on a different set D_1. In fact, there is additional freedom to choose interpolation structure $(\mathcal{P}_1, D_1, \mathcal{I}_1)$. The most common expression of this freedom occurs when $\mathcal{P}_0$ is a triangular mesh, of which $\mathcal{P}_1$ is chosen as a refinement, resulting in sample locations $D_1 \supset D_0$, with $\mathcal{I}_0 = \mathcal{I}_1$ being linear interpolation of each triangle's vertices.

Does interpolation structure $(\mathcal{P}_1, D_1, \mathcal{I}_1)$ interpolate the simplified data F_1 to function f_1? In this paper we allow the answer to be "no", but we restrict selection of interpolation structures to those that interpolate to a function satisfying a *topological uniqueness condition* local to $\mathcal{P}$ and D, the patch and sample geometry. These so-called $\mathcal{P}D$ interpolations all have identical topological structure.

It certainly is of interest to allow $D_1 \neq D_0$, for example when refining or decimating a mesh-structured patch collection $\mathcal{P}_0$. Alternatively, when working with Morse functions, we could choose $\mathcal{P}_1$ to be f_1's Morse-Smale complex, with D_1 comprising f_1's critical points. However, in this paper we restrict our attention to an unspecified, but unchanging, patch and sample geometry, $\mathcal{P}$ and D. Thus, the interpolation structure for the i-th simplification, for any iteration $i \geq 0$, is fixed as $(\mathcal{P}, D, \mathcal{I})$. Additionally, we require that each function f_i, is a $\mathcal{P}D$ interpolation of data F_i.

Finally, Extremal Simplification is as follows, where f_1 and f_1' are $\mathcal{P}D$ interpolations of F_1, and typically $f_1 \neq f_1'$:

$$
\begin{tikzcd}
F_0 \arrow[r, "\text{choose}"] & (\mathcal{P}, D, \mathcal{I}) \arrow[r, "\text{interpolate}"] & f_0 \arrow[r, "\text{monotone-light}", "\text{factorization}"'] & (M_0, \lambda_0) \arrow[d, "\text{quotient}"] \\
F_1 \arrow[d, "\text{no choice}"] & & f_1 \arrow[ll, "\text{sample on } D"'] & (M_1, \lambda_1) \arrow[l, "\text{synthesize}"'] \arrow[d, equal] \\
(\mathcal{P}, D, \mathcal{I}) \arrow[rr, "\text{interpolate}"] & & f_1' \arrow[r, "\text{monotone-light}", "\text{factorization}"'] & (M_1, \lambda_1) \arrow[d, "\text{quotient}"] \\
 & & & \cdots
\end{tikzcd}
$$

In this paper we show how to make iterated simplification sequences. At iteration i each extremum of topological structure (M_i, λ_i) has a *scale* which, when Morse f_i is Morse, is equal to its persistence in the sense of Edelsbrunner et al. [11].

The extrema of simplified topological structure (M_{i+1}, λ_{i+1}) are a subset of those of (M_i, λ_i). The corresponding extrema necessarily have identical values under function f_i and f_{i+1}, and f_{i+1} may be chosen so that corresponding extrema are colocated in the domain X. However, we cannot guarantee that corresponding extrema have identical scales, although in most cases most of the extrema will.

This differentiates Extremal Simplification from ϵ-simplification of Edelsbrunner et al. [9] and the simplification found in Edelsbrunner et al. [11]. In the latter case, the persistence of all remaining extrema is reduced by a fixed amount; whereas an ϵ-simplification leaves fixed the persistence of the remaining extrema. On the other hand, Extremal Simplification properly contains the simplifications of Carr [4].

Because scale is not preserved by Extremal Simplification, it does not make sense to talk about "persistence order" with respect to the sequence of extrema removed by the simplifications. However, Extremal Simplification admits sequences having successively increasing smallest scale, so that for any $\delta > 0$ the sequence has a member having no extrema of scale δ or less. Alternatively, one might order a sequence of simplifications using the local geometric measures of Carr [4], although this is not included in Extremal Simplification as presented herein.

Measuring approximation error in the continuous and discrete domains, respectively, as $\max_{x \in X} |fx - gx|$ and $\max_{x \in D} |Fx - Gx|$, we prove that when reducing the number of extrema, both measures are bounded below by half the smallest of f's or F's extrema's scales. We discuss situations in which this lower bound can be, or cannot be, realized.

1.2 Contributions

The two main contributions of Extremal Simplification are as follows.

(1) Rigorous connection between continuous and finite domains: The topological representation used by Extremal Simplification is derived from a function on a continuous domain; this function is an interpolation of finite data. Also, simplification of the topological representation results in a simplified continuous function that is subsequently sampled to obtain simplified data. A topological uniqueness condition for interpolation restricts our attention to interpolations having the same topological representation; thus we obtain a well-defined notion of simplification for finite data, and are able to extend the continuous theory to the discrete case, including the lower bound on approximation error.

(2) Breadth and generality: Extremal Simplification is insensitive to the dimension and homological complexity of the Peano space X upon which the scalar field $f : X \to \mathbb{R}$ is defined. Restriction to "piecewise monotone" functions tames the complexity and floods the potentially fractal characteristics of Peano spaces. Any piecewise monotone $f : X \to \mathbb{R}$ can be simplified; "degenerate" functions are not an issue. Patch and sample geometries for X are not restricted to polygonal meshes having samples at the vertices.

1.3 Related Work

Simplification of sampled scalar fields has appeared recently in work by Carr [4], Carr et al. [5], Weber et al. [28], Bremer et al. [1], and Edelsbrunner et al. [11, 9]. Each of these papers describe simplification of sampled scalar fields defined on either two-dimensional manifolds or three-dimensional volumes. Edelsbrunner et. al [9] state: "Use of the simplified complex together with the original data may be tolerable for visualization purposes, but it is not satisfactory when the simplified data is used in the subsequent data analysis stage". This is in alignment with Extremal Simplification, which is application-neutral, being concerned only with data simplification but not the use to which it is put.

Each of the the papers mentioned above ([4, 5, 28, 1, 11, 9]) uses the topological structure of the scalar field to guide simplification. Carr et al. [4, 5] and Weber et al. [28] use the Reeb graph [24], Bremer et al. [1] use the Morse-Smale complex [10], and Edelsbrunner et al. [11, 9] use the persistence diagram. The Reeb-based techniques are concerned with removing extrema; the Morse-Smale and persistence-diagram methods may also remove critical points related to the genus of isosurfaces. Extremal Simplification represents scalar field topology as the function's middle space and light factor, an augmentation of the Reeb graph.

Each of the papers [4, 5, 28, 1, 11, 9] includes computational considerations, including data structures and runtime complexity. As presented herein, Extremal Simplification theory does not explicitly address computation; however, many of the methods referenced in the literature are applicable, and the intent of Extremal Simplification is to provide the basis for computational methods.

Detailed analysis of the papers [4, 5, 28, 1, 11, 9] is presented in section 7.

Computational methods for simplification of three-dimensional geometry have been a topic of interest in the research literature for a decade, almost entirely in connection with visualization. The primary simplification mechanism is edge contraction in a triangular mesh [14]. Some approaches include topological considerations based on the Reeb graph, e.g. Takahashi et al. [25], or Morse-Smale complex, e.g. Gyulassy et al. [13]. These works differ from Extremal Simplification, because they focus on simplifying the geometry of triangulated surfaces rather than scalar fields.

2 Peano Spaces & Monotone-Light Factorization

A topological space X is a Peano space when it is a compact, connected, locally connected, metric space. Peano spaces include disks and compact manifolds in $\mathbb{R}^n$, as well as non-manifold surfaces resulting from gluing together compact manifolds. Peano spaces are not necessarily smooth; they include polygons, simplexes, graphs and fractals. Peano spaces are also called Peano continua; continuum theory had its heyday in the mid-twentieth century, e.g. [31], although there are some more recent treatments [22].

Throughout this paper all spaces are Peano and all functions are continuous.

For Peano spaces X, Y and continuous $f : X \to Y$, when $W \subset X$ is, respectively, connected, locally connected, closed or compact, then fW has this same property; and when $W \subset Y$ is, respectively, open, closed or compact, then $f^{-1}W$ has this same property.

$f : X \to Y$ is monotone when for every connected $W \subset Y$, $f^{-1}W$ is connected. f is light when for every discrete $W \subset Y$, $f^{-1}W$ is discrete. The Eilenberg-Whyburn monotone-light factorization [31, 29] states that there exists a unique Peano space M, called f's middle space, such that $f = \mu \circ \lambda$, where $\mu : X \to M$ is monotone and $\lambda : M \to Y$ is light. We specify f's monotone-light factorization by simply listing $\mu M \lambda$. See Appendix A for historical discussion of the monotone-light factorization.

Suppose $f : X \to Y$. f's middle space M is defined exactly as is the Reeb graph [24], but is not in general a graph. The middle space is a quotient of the domain X, where $x, y \in X$ are identified if and only if they both lie in a connected component of a level set $f^{-1}z$. f's monotone factor μ is the natural map from X to M; f's light factor λ assigns each point $p \in M$ the value $f(\mu^{-1}p) \in Y$. Thus $f = \mu \circ \lambda$.

Standard results [31, 29] state that the middle space is a Peano space, and that the monotone-light factorization is unique.

3 Piecewise Monotone Functions

The monotone-light factorization provides the basis for a generalization of “piecewise monotone”. We give a general definition, followed by focus on real-valued piecewise monotone functions.

Definition 3.1. *Suppose $f : X \to Y$ has monotone-light factorization $\mu M \lambda$.*
λ is locally monotone *at $p \in M$ when p has a neighborhood upon which λ is monotone.*
M^* denotes the set of all points of M at which λ is locally monotone.
f is piecewise monotone *when:*
(1) M^* is dense in M;
(2) M^* has finitely many components; and
(3) λ is monotone on the closure of each component of M^*.

The closures of the the components of M^* are the *monotone pieces* referred to in the name “piecewise monotone”. λ is a homeomorphism on each monotone piece.

We will see that the middle space of a real-valued function is its Reeb graph.

$f : X \to \mathbb{R}$ is piecewise monotone whenever the set of points not locally monotone, $M \smallsetminus M^*$, is finite. Condition (2) of definition 3.1 then follows from compactness of M; condition (3) follows from lightness of λ and separability of M.

When $f : [0\ 1] \to \mathbb{R}$, definition 3.1 is exactly the usual meaning of “piecewise monotone”.

Throughout this paper all real-valued functions on Peano space X will be piecewise monotone. Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$.

f's middle space M is partially ordered, with $p < q$ whenever there exists a path $P \subset M$ from p to q with λ numerically monotone increasing along P. A point $p \in M$ is a *maximum* (resp. *minimum*) when there does not exist $q > p$ (resp. $q < p$). When $p \in M$ is both a maximum and a minimum, then it must be the case that $M = \{p\}$ and f is constant, in which case we count f as having no monotone pieces. Assuming f not constant, each maximum and minimum is an *extremum.* Two extrema are *same-sense* when they are both maxima or both minima. Extremum $p \in M$ is *global* when λp is an endpoint of the interval fX.

Throughout this paper all functions $f : X \to \mathbb{R}$ will be assumed to be non-constant.

The extrema of M's partial ordering correspond exactly to the intuitive notion of f's extrema. Working in f's domain X, a closed connected subset $K \subset X$ is a maximum (minimum) of f if and only if f is constant on K, K is a component of $f^{-1}fK$, and there exists an open set $V \supset K$ such that $fx < fK$ ($fx > fK$) for all $x \in V \smallsetminus K$. Working in f's middle space M, an extremum $p \in M$ is a maximum (minimum) if and only if p has a neighborhood U with $q < p$ (resp. $q > p$) for all $q \in U \smallsetminus \{p\}$. Therefore, $K \subset X$ is a maximum (minimum) if and only if $K = \mu^{-1}p$ for maximum (minimum) p in M. Throughout this paper we focus on extrema in the middle space; when $p \in M$ is an extremum we interchangeably refer to p as an extremum of M and as an extremum of f.

Every monotone piece $S \subset M$ is mapped homeomorphicly by λ to the real interval λS. Monotone pieces may only intersect at their endpoints. Thus we recognize the middle space of a real-valued piecewise monotone function as its Reeb graph. When X is simply connected then M is a tree; see Appendix A for proof of this well-known result.

The approach taken in this paper is to focus on the topological properties of the middle space. However, we adopt some graph-related terminology: Any $p \in M \smallsetminus M^*$ lies in the boundaries of at least two monotone pieces; we call p a *branch point.* A *saddle* is a branch point that is not an extremum. The union of the extrema and saddles comprise M's *nodes.* Graph-theoretically, the paths connecting M's nodes comprise M's arcs.

3.1 Scale

This section defines the peak-to-valley vertical extent of an extremum is its *scale*.

For any set S and point $p \in S$, we denote by $\mathcal{C}_p(S)$ the connected component of S containing p.

Definition 3.2. *Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$; let $p \in M$ be an extremum.*
Then p's scale interval*, denoted $\mathbf{I}_p$, is the shortest closed interval I containing λp such that either:*
There exists $q \neq p \in \mathcal{C}_p(\lambda^{-1} I)$ with $\lambda q = \lambda p$; or
$I = fX$.

The length of $\mathbf{I}_p$ is p's scale*, denoted $\sigma(p)$.*
The smallest of f's extremas' scales is f's least significant scale*, denoted $\sigma(f)$.*

$\mathbf{I}_p = fX$ implies p is a global extremum; the converse implication does not hold. $\sigma(p) = 0$ if and only if f is constant, i.e. $M = \{p\}$.

For maximum $p \in M$, note that $\lambda p = \max \mathbf{I}_p$. We also see that when there exists at least one other maximum then there exists a branch point $q < p \in M$ with $\lambda q = \min \mathbf{I}_p$ such that there exists $r \neq p \in M$ with $q < r$ and $\lambda r = \lambda p$. The symmetric statements hold for minimum $p \in M$. We call each such q a *turnaround* of p. Every non-global extremum has at least one turnaround.

When f is a Morse function, every branch point is a saddle, and the light factor λ takes unique values on the extrema and saddles of M. Therefore every non-global extremum has exactly one turnaround, and no saddle is the turnaround for more than one extremum. The pairing of non-global extrema and their turnarounds is exactly the pairing of critical points used to determine persistence by Edelsbrunner et al. [11, 10], and each extremum's scale is equal to its persistence. To see this, let $p \in M$ be a non-global minimum having turnaround q; we consider the components of sublevel sets $\mathcal{S}(z) = \{r \in M \mid \lambda r \leq z\}$. The point p comprises a component of $\mathcal{S}(\lambda p)$. As z increases from λp, the component of $\mathcal{S}z$ containing p remains distinct from all components containing other points of $\mathcal{S}(\lambda p)$ until $z = \lambda q$.

3.2 Approximation Error

Functions $f, g : X \to \mathbb{R}$ will be compared in L_∞. If we think of g as approximating f, then $e(f, g)$ denotes the approximation error, with $e(f, g) = \max_{x \in X} |fx - gx|$.

The following theorem captures the relationship between scale and approximation error; proved in Appendix B.

Theorem 3.3. *Suppose $f, g : X \to \mathbb{R}$ are piecewise monotone, where g has fewer extrema than f. Then $e(f, g) \geq \sigma(f)/2$.*

Theorem 3.3 is a generalization of the well-known result of Ubhaya for isotone approximation [26]. Bremer et al. [1] refer without proof to this bound for simplification of Morse-Smale functions on two-dimensional manifolds. Edelsbrunner et al. [9] show that their ϵ-simplification cannot always achieve this lower bound. The next section includes discussion of the achievability of this bound for piecewise monotone functions.

4 Simplification of Piecewise Monotone Functions

Roughly speaking, a function is simplified by taking a certain type of quotient of its middle space. This section defines simplification, introduces several types of simplification, and defines simplification sequences.

Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$.

A subset $K \subset M$ is a *collapse set* for f when K has finitely many components $K_1 \dots K_n$, where each K_i is closed, has nonempty interior, and λ is constant on each ∂K_i. We indicate collapse set K's components by writing $K = \sum K_i$.

Every collapse set contains an extremum of f in each component of its interior.

Every collapse set $K = \sum K_i$ defines a quotient M_K of M by identifying the points within each component K_i. This quotient has natural map denoted $\phi_K : M \to M_K$, where for any $q \in M_K$, $\phi_K^{-1} q$ is either singleton or equal to one of the K_i, and if $q \neq q'$ then $\phi_K^{-1} q$ and $\phi_K^{-1} q'$ are disjoint. Note that ϕ_K is monotone. We call M_K the *monotone quotient* of M by K.

M_K admits a light function $\lambda_K : M_K \to \mathbb{R}$ defined by $\lambda_K q = \lambda(\partial \phi_K^{-1} q)$. Note that λ_K is well-defined because λ is constant on $\partial \phi_K^{-1} q$. λ_K is light because ϕ_K is one-to-one on $M \setminus K$ and K has finitely many components. We call λ_K the *monotone quotient* of λ by K

The points of M_K are partially ordered: For $p, q \in M_K$, $p < q$ when there exists a path P from p to q with λ_K monotone increasing on P. Consequently, we can speak of M_K's extrema.

K is an *extremal* collapse set when for each component K_i such that $q = \phi_K K_i$ is an extremum of M_K, ∂K_i is comprised entirely of extrema of f of having the same sense as q. It follows that for every extremum $q \in M_K$, the set $\partial \phi_K^{-1} q$ is comprised entirely of of extrema of f of having the same sense as q.

Definition 4.1. *Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$, and suppose $K \subset M$ is an extremal collapse set for f. Let M_K, λ_K be the monotone quotients of M, λ by K. Then any function $g : X \to \mathbb{R}$ having middle space M_K and light factor λ_K is an* extremal simplification of f generated by K.

Throughout this paper we abbreviate "extremal simplification" to "simplification".

Suppose $K = \sum K_i$ is an extremal collapse set for f, let M_K, λ_K be the monotone quotients of M, λ by K, and let ϕ_K be the natural map $M \to M_K$.

Section 4.1.1 will show that there always exists at least one simplification of f generated by K. Every simplification of f is piecewise monotone and has fewer monotone pieces than f; this is proved in appendix C.

There may be many different simplifications of f generated by K. Each simplification has its unique monotone factor; all share the same middle space and light factor. When we say that functions "have the same middle space", we mean "up to homeomorphism commuting with the light factor".

K *removes* some of f's extrema: For each component K_i, when $\phi_K K_i$ an extremum of M_K, then the extrema of f lying in K_i° are removed; otherwise, when $\phi_K K_i$ is not an extremum of M_K, all extrema lying in K_i are removed. The extrema not removed by K *survive* K.

Suppose g is a simplification of f generated by K.

ϕ_K maps f's surviving extrema onto g's extrema, with $\lambda p = \lambda_K \phi_K p$ for each surviving extremum p. Distinct surviving extrema of f are mapped to the same extremum q of g only if they lie in the boundary a component K_i having $\phi_K K_i = q$.

The number of extrema of g is less than that of f by exactly the number of extrema contained in K minus the number of components $K_i \subset K$ that are mapped by ϕ_K to extrema of g. This difference is always nonzero.

4.1 Special Types of Simplification

Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$. The following sections discuss a variety of simplifications of f.

4.1.1 Flat Simplification

Suppose $K = \sum K_i$ is an extremal collapse set for f, let M_K, λ_K be the monotone quotients of M, λ by K, and let ϕ_K be the natural map $M \to M_K$.

K generates at least one simplification, called K's *flat simplification*, denoted f_K, constructed by defining the monotone factor $\mu_K = \mu \circ \phi_K$. Thus $f_K(x) = \mu_K \circ \lambda_K \ (x) = \lambda(\partial\phi_K^{-1}(\phi_K(\mu x)))$. Stated in words, f_K is computed by first mapping $x \in X$ to $p = \mu x \in M$, pulling p back to the set $C = \phi_K^{-1} p$, and then – noting that λ is constant on ∂C – taking the value of λ on ∂C. In other words, f_K flattens each of the connected sets $\mu^{-1} K_i$ to the constant value of $f(\partial K_i)$, and is otherwise equal to f on $X \setminus \mu^{-1} K$.

The approximation error $e(f, f_K)$ is easily determined: $e(f, f_K) = \max\limits_i \max\limits_{p \in K_i} |\lambda p - \lambda \partial K_i|$. Furthermore, there exists an extremum $p \in K^\circ$ such that $e(f, f_K) = |\lambda p - \lambda_K p|$.

4.1.2 Standard Simplification

Suppose $p \in M$ is an extremum of f such that p has a turnaround q. Recalling the notation of definition 3.2, let $\mathbf{I}_p$ be p's scale interval, and define half-open interval $\widetilde{\mathbf{I}}_p = \mathbf{I}_p \setminus \lambda q$. Define p's *standard collapse set*, denoted $\mathbf{C}_p$, by $\mathbf{C}_p = \overline{\mathcal{C}_p(\lambda^{-1}\widetilde{\mathbf{I}}_p)}$. Then $\mathbf{C}_p$ is an extremal collapse set having p in its interior. Now define f's *standard simplification removing* p, denoted f_p, as $\mathbf{C}_p$'s flat simplification. Clearly $e(f, f_p) = \sigma(p) = |\lambda p - \lambda_{\mathbf{C}_p} p|$; note that this approximation error is twice the lower bound of theorem 3.3.

When $p, q \in M$ are same-sense extrema, then $\sigma(p) \le \sigma(q)$ implies that either $\mathbf{C}_p \subset \mathbf{C}_q$ or $\mathbf{C}_p^\circ \cap \mathbf{C}_q^\circ = \emptyset$. However, this statement is not true when $p, q \in M$ are opposite-sense extrema.

Suppose extremum $q \in M$ survives the standard simplification removing p; let $q' \in M_{\mathbf{C}_p}$ be the extremum to which q is mapped by natural map $\phi_{\mathbf{C}_p}$. Then it is possible that $\sigma(q)$, q's scale in M, is not equal to $\sigma(q')$, q''s scale in $M_{\mathbf{C}_p}$. When p and q are opposite-sense, then this situation arises if and only if $\mathbf{C}_p^\circ \cap \mathbf{C}_q^\circ \ne \emptyset$ and $\mathbf{C}_p \not\subset \mathbf{C}_q$; in this case $(\sigma(q) - \sigma(p)) \le \sigma(q') < \sigma(q)$. When p and q are same-sense, this situation may arise when $\sigma(p) = \sigma(q)$, $\lambda p = \lambda q$, and $\mathbf{C}_p^\circ \cap \mathbf{C}_q^\circ = \emptyset$ but $\partial \mathbf{C}_p \cap \partial \mathbf{C}_q^\circ \ne \emptyset$; in this case $\sigma(q') > \sigma(q)$, and $\sigma(q')$ is bounded above only in relation to the values of M's global extrema.

4.1.3 Optimal Simplification

Suppose g is a simplification of f removing extrema $p_1 \dots p_m$. Then g is *optimal* when $e(f, g) = \max\limits_i \sigma(p_i)/2$. The theory of optimal simplification is left for future research. Two examples are presented here.

Example: Let X denote the graph comprised of two copies of the real interval $[-1\ +1]$ glued together at 0. X is a Peano space. Suppose $f : \mathsf{X} \to \mathbb{R}$ maps each point to its corresponding point in $[-1\ +1]$. Then f is light, so its monotone factor is the identity function, its middle space is X, and its light factor is f. Denote the two maxima as p, q, the two minima as r, s, and the saddle as t. Thus $\sigma(p) = \sigma(q) = \sigma(r) = \sigma(s) = 1$. The standard collapse set for p, $\mathbf{C}_p$, is the closed arc from p to t, and similarly for $\mathbf{C}_q$, $\mathbf{C}_r$ and $\mathbf{C}_s$. It is obvious that any simplification g removing both the maximum p and minimum r will have $e(f, g) \geq 1$, and therefore no such simplification may be optimal.

Example: Suppose X is any Peano space and $f : X \to \mathbb{R}$ a piecewise monotone function having monotone-light factorization $\mu M \lambda$, where M is a graph having the form of the letter Y. M has two maxima, p, q, one minimum r and one saddle t, with arcs from t to each of p, q, r. Suppose $\lambda p = \lambda q = 1$, $\lambda t = 0$, and $\lambda r = -1$; define λ linearly between these nodes. Thus $\sigma(p) = 1$ and $\mathbf{C}_p$ is the closed arc from p to t. It may not be obvious that there exist simplifications g removing the maximum p such that $e(f, g) = 1/2$. However, such a simplification may be constructed as follows. First, construct f_p, the standard simplification removing p; let $\mu_p M_p \lambda_p$ be f_p's monotone-light factorization. Note that f_p is monotone, and thus μ_p is constant on $f_p^{-1} z$ for any $z \in [-1\ 1]$. Next, choose any $0 < \epsilon \leq 1/2$, let $v : [-\epsilon\ 0] \to [-\epsilon\ 1/2]$ and $w : [0\ (1/2 + \epsilon)] \to [1/2\ (1/2 + \epsilon)]$ denote linear maps between the intervals, and define $\mu' : X \to M_p$ by the following rules:

$\mu' x = \mu_p x$ when $f_p x \leq -\epsilon$ or $f_p x \geq 1/2 + \epsilon$
$\mu' x = \mu_p(f_p^{-1}(v(f_p x)))$ when $f_p x \in [-\epsilon\ 0]$
$\mu' x = \mu_p(f_p^{-1}(w(f_p x)))$ when $f_p x \in [0\ (1/2 + \epsilon)]$

Then μ' is monotone. Let $g = \mu' \circ \lambda_p$; then g has monotone-light factorization $\mu' M_p \lambda_p$, and is thus a simplification of f removing p. It is easily checked that $e(f, g) = 1/2$, and therefore g is optimal.

4.1.4 One-Dimensional Simplification

Function $f : X \to \mathbb{R}$ is *one-dimensional* when its middle space has no saddles. Thus the middle space is homeomorphic to the real unit interval or the circle. When X is a real interval or the circle then f is one-dimensional. However, f may be one-dimensional for other domains; for example when f is a wave train across a two-dimensional surface.

Suppose $f : X \to \mathbb{R}$ is one-dimensional with monotone-light factorization $\mu M \lambda$, and $p \in M$ is an extremum having turnaround q. Then q is an opposite-sense extremum.

There may exist optimal simplifications for a one-dimensional function. We consider then case where extremum p has turnaround q such that p is also a turnaround for q. Note that this will be the case whenever p lies between a global maximum and a global minimum, e.g. as will be the case when M is homeomorphic to the circle. Then $|\lambda p - \lambda q| = \sigma(p) = \sigma(q)$ and $\mathbf{I}_p = \mathbf{I}_q$. Bisect $\mathbf{I}_p$ into closed intervals J_p, J_q, where $\lambda p \in J_p$ and $\lambda q \in J_q$. Let $K = \mathcal{C}_p(\lambda^{-1} J_p) \cup \mathcal{C}_q(\lambda^{-1} J_q)$; then K is a connected extremal collapse set containing p and q in its interior, and K's flat simplification f_K is an optimal simplification removing p and q.

One-dimensional simplification admits efficient computation. Brooks et al. [3] provide a linear time algorithm that pairs each extremum with its turnaround.

4.2 Simplification Sequences

In this section we consider sequences of functions $f_1 \ldots f_n$ that demonstrate successively increased simplification of an initial function f_0.

A sequence $f_1 \ldots f_n$ might be constructed in one of two ways, where for each $i > 0$:

f_i is a simplification of f_{i-1}; or
f_i is a simplification of f_0 generated by extremal collapse set K_i, and $K_i \supset K_{i-1}$.

In appendix D we show these are equivalent.

Our goal of successively increased simplification is captured as:

Definition 4.2. *Suppose $f_0 : X \to \mathbb{R}$ is piecewise monotone. A sequence of functions $f_1 \dots f_n : X \to \mathbb{R}$ is a* simplification sequence *for f_0 when for each $i > 0$:*
(1) f_i is a simplification of f_{i-1};
(2) f_n has only global extrema; and
(3) $\sigma(f_i) > \sigma(f_{i-1})$.

Suppose $f_1 \dots f_n$ is a simplification sequence for f_0. Then given any $\epsilon > 0$, we can identify a simplification g of f_0 such all of g's extrema have scale greater that ϵ, by choosing the least index j such that f_j has this property.

We show that for any f_0 there exists at least one simplification sequence. Suppose f_{i-1} has already been constructed for some $i > 0$. If f_{i-1} has only global extrema, then we are done. Otherwise, determine f_i by constructing a sequence of simplifications $g_0\ g_1 \dots$ of f_{i-1}, as follows. Let $g_0 = f_{i-1}$. Suppose g_j has been constructed for some $j \geq 0$. If $\sigma(g_j) > \sigma(f_{i-1})$, then define $f_i = g_j$. Otherwise, construct g_{j+1} as follows: Choose any extremum p of g_j such that $\sigma(p) \leq \sigma(f_{i-1})$, and let g_{j+1} be the standard simplification of g_j that removes p.

Note that this construction has the possibility of choice of extremum p when simplifying g_j to g_{j+1}. Any simplification sequence constructed as described is called a *standard* simplification sequence.

It follows from the construction that for each $i > 0$, $e(f_{i-1}, f_i) \geq \sigma(f_{i-1})/2$. This information is not particularly useful, since the scale of a surviving extremum of f_{i-1} is not necessarily equal to the scale of the extremum of f_i to which it maps, as discussed in section 4.1.2.

Ideally, we would like analysis of function f_0's middle space M_0 to result in a *simplification sequence map*, identifying extremal collapse sets $K_1 \dots K_n$ generating the simplifications $f_1 \dots f_n$, perhaps together with the approximation errors $e(f_i, f_0)$ and or scales $\sigma(f_i)$, $i = 1 \dots n$.

For example, when f_0 is one-dimensional and such that no same-sense extrema have equal values, then each non-global extremum p_i has a unique turnaround q_j, which is an opposite-sense non-global extremum. For each such pair of extrema p_i, q_j, we identify a collapse set K_{ij} as follows. When p_i is q_j's turnaround, then we define K_{ij} be the collapse set of section 4.1.4 which when flattened gives the optimal simplification removing p_i and q_j. Otherwise, we let $K_{ij} = \mathbf{C}_{p_i}$, p_i's standard collapse set. It follows that the interiors of these collapse sets are pairwise either disjoint or nested. Let $K_1 \dots K_n$ be any ordering of the collection of collapse sets K_{ij}, such that for indices $a < b$, $K_a \not\supset K_b$. Then any sequence $f_1 \dots f_n$ of simplifications generated, respectively, by $K_1 \dots K_n$ constitutes a simplification sequence for f_0. This one-dimensional situation has been studied by Brooks [2] and Brooks et al. [3].

5 Scalar Data

Scalar data is a collection of real-valued measurements on a finite collection D of points, i.e. a finite function $F : D \to \mathbb{R}$. We refer to D as *sample locations*.

Scalar data may be interpolated between the sample locations D, resulting in a continuous function $f : X \to \mathbb{R}$, where $f|D = F$. In this section we provide a topological uniqueness condition guaranteeing that certain interpolations all have the same middle space and light factor, thereby allowing assignment of this middle space and light factor to the scalar data.

5.1 Patches

Not all Peano spaces can be triangulated; it is an open question as to whether all Peano spaces may be decomposed into a collection of patches each of which is itself a Peano space. Brick partitions (ref Bing) suggest that they might. In any case, we restrict our attention to spaces that can be decomposed into patches.

Definition 5.1. *A finite collection $\mathcal{P} = P_1 \dots P_n$ of subsets of X is a* patch collection *for X when:*

(1) each $P_i \in \mathcal{P}$ is connected, and $P_i = \overline{P_i^\circ}$;
(2) $\mathcal{P}$ covers X;
(3) $P_i^\circ \cap P_j^\circ = \emptyset$ when $i \neq j$; and
(4) each intersection $P_i \cap P_j$ has finitely many components.

Condition (1) implies each $P_i \in \mathcal{P}$ is a Peano subspace of X, since the closure of an open subset of a locally connected space is locally connected.

Suppose $\mathcal{H} = h_1 \dots h_n$ is a collection of functions $h_i : P_i \to \mathbb{R}$. Then $\mathcal{H}$ comprises *patch functions* for $\mathcal{P}$ when $h_i|(P_i \cap P_j) = h_j|(P_i \cap P_j)$ for all i, j. We denote by $f_{\mathcal{PH}}$ the unique function such that $f_{\mathcal{PH}}|P_i = h_i$.

5.2 $\mathcal{P}D$ Interpolation

Scalar data may be interpolated many ways; different interpolations may have distinct piecewise monotone structure. A priori, no one interpolation is "right". Definition 5.2, below, relates patches, patch functions, and sample locations, providing a uniqueness condition for interpolations' piecewise monotone structure.

Suppose $f : X \to \mathbb{R}$, $D \subset X$ are sample locations, and $z \in \mathbb{R}$. We say that D *witnesses* z when there exists $x \in D$ such that $fx = z$. Denote f's monotone-light factorization as $\mu M \lambda$, and let $p \in M$. We may also say that D *witnesses* p when $D \cap \mu^{-1}p$ witnesses λp.

Definition 5.2. *Given patch collection $\mathcal{P} = P_1 \dots P_n$ for X and sample locations $D \subset X$, a function $f : X \to \mathbb{R}$ is a $\mathcal{P}D$* function *when $f = f_{\mathcal{PH}}$ for patch functions $\mathcal{H} = h_1 \dots h_n$ such that:*
(1) each h_i is monotone;
(2) $D \cap P_i$ witnesses $\min h_i P_i$ and $\max h_i P_i$ for each $i = 1 \dots n$; and
(3) for each component K of $P_i \cap P_j$, $D \cap K$ witnesses $\min h_i K$ and $\max h_i K$ for each $i, j \in 1 \dots n$.

Suppose $f : X \to \mathbb{R}$ is $\mathcal{P}D$. Then f is piecewise monotone. Let $\mu M \lambda$ denote f's monotone-light factorization and let $p \in M$ be any node; then D witnesses p. These results are proved in appendix F. The following theorem is proved in Appendix G.

Theorem 5.3. *Suppose $\mathcal{P}$ is a patch collection for X and $F : D \to \mathbb{R}$ is scalar data. Then any two $\mathcal{P}D$ functions f, f' interpolating F have identical middle spaces and light factors.*

When $\mathcal{P}$ is a triangular mesh and D comprises the triangles' vertices, then the piecewise linear interpolation is $\mathcal{P}D$.

When $\mathcal{P}$ is an n-dimensional cubic mesh and D comprises the cubes' vertices, then n-linear interpolation is $\mathcal{P}D$ only when each cube's interpolation is monotone. Cubes that are not monotone may be triangularly subdivided and perhaps additional samples defined on the triangles' vertices, resulting in patch collection $\mathcal{P}'$ and samples locations D'. Then n-linear interpolation on the monotone cubes together with linear interpolation on the triangles gives a $\mathcal{P}'D'$ interpolation.

We now consider the special case of X being a smooth manifold and f being a Morse-Smale function. In two dimensions, each cell of the Morse-Smale complex [10] is a quadrilateral having a critical point at each corner. We note that the scalar field is topologically monotone on each cell and numerically monotone on each cell edge. This statement generalizes to Morse-Smale complexes of higher dimension. Thus, when D includes all critical points of f then f's Morse-Smale complex constitutes a patch collection $\mathcal{P}$ such that f is a $\mathcal{P}D$ interpolation of the critical points.

5.3 Piecewise Monotone Data

There are four principal ingredients upon which Extremal Simplification positions the definition of simplification for scalar data:

1. A Peano space X.
2. A patch collection $\mathcal{P}$ covering X.
3. A finite set of sample locations $D \subset X$.
4. Scalar data $F : D \to \mathbb{R}$.

In differing application contexts these four ingredients may arise in various orders. For example, one might start with D and X, choose $\mathcal{P}$ to be a particular triangulation of D, and then consider the data F. Alternatively, one might start with X, and then be given both D and F, and then choose a suitable patch collection $\mathcal{P}$. In any case, the four ingredients are bound together by:

Definition 5.4. *Suppose X is a Peano space, $\mathcal{P}$ is a patch collection covering X, and $D \subset X$ is a finite set of sample locations. Scalar data $F : D \to \mathbb{R}$ is* piecewise monotone *when F has a $\mathcal{P}D$ interpolation $f : X \to \mathbb{R}$.*

Note that use of the term "piecewise monotone" assumes a context where X, $\mathcal{P}$ and D are given.

Theorem 5.3 states that all $\mathcal{P}D$ interpolations of piecewise monotone F share the same middle space and light factor. Therefore, we may speak of F's middle space and light function without reference to a particular interpolation. Similarly, we may speak of F's extrema, including the number of extrema, and the scale of each extremum. We may define F's scale $\sigma(F)$ as the scale of F's least significant extremum. Likewise, F's extremal collapse sets are well-defined.

Example: Suppose X is a smooth n-manifold, $\mathcal{P}$ is an n-dimensional triangular mesh on X, and D comprises all vertices of the mesh. Then any scalar data F is piecewise monotone, since linear interpolation within each patch results in a $\mathcal{P}D$ function.

5.4 Approximation Error for Piecewise Monotone Data

Given two sets of piecewise monotone data F, G, and thinking of G as approximating F, we measure the approximation error as $e(F, G) = \max_{x \in D} |Fx - Gx|$. Appendix H proves the following theorem:

Theorem 5.5. *Let Peano space X, patch collection $\mathcal{P}$, and finite sample locations $D \subset X$ be given. Suppose $F, G : X \to \mathbb{R}$ are piecewise monotone data, where G has fewer extrema than F. Then $e(F, G) \geq \sigma(F)/2$.*

6 Simplification of Piecewise Monotone Data

We simplify piecewise monotone data F by simplifying a $\mathcal{P}D$ interpolation.

Definition 6.1. *Suppose X is a Peano space, $\mathcal{P}$ is a patch collection covering X, $D \subset X$ is a finite set of sample locations, and $F : D \to \mathbb{R}$ is piecewise monotone data. Let F's middle space be denoted by M, and suppose $K \subset M$ is an extremal collapse set. Then piecewise monotone data $G : D \to \mathbb{R}$ is an* extremal simplification *of F* generated by *K when there exist $\mathcal{P}D$ interpolations $f, g : X \to \mathbb{R}$ of, respectively, F, G such that g is a simplification of f generated by K.*

Theorem 5.3 makes the choices of f and g irrelevant in definition 6.1.

Suppose f is a $\mathcal{P}D$ interpolation of F. Given an arbitrary simplification g of f, it may not always be the case that g is a $\mathcal{P}D$ function. Only simplifications of f that are $\mathcal{P}D$ give rise to simplifications of F.

There always exists at least one simplification of F generated by extremal collapse set K, namely K's flat simplification, denoted F_K, and defined as $f_K|D$, where f_K is K's flat simplification of f (section 4.1.1). We prove f_K is $\mathcal{P}D$ in appendix I. It follows that for any extremum p of F, the standard simplification of f removing p is $\mathcal{P}D$; we denote $f_p|D$ as F_p.

To see the practical significance of this result, suppose $\mathcal{P}$ and D comprise a triangulation of X having samples at the vertices. Let f be the piecewise linear interpolation of data F, having middle space M. Choose any extremal collapse set $K \subset M$. Then the flat simplification f_K is typically not linear on each triangle of $\mathcal{P}$. Nevertheless, the piecewise linear interpolation of F_K has the same middle space and light factor as f_K.

In general, when G is a simplification of F, with $\mathcal{P}D$ interpolations g and f, respectively, then approximation error $e(F, G) \neq e(f, g)$ since $e(f, g)$ may vary over the choices of f and g. However, for the flat simplification F_K, $e(F, F_K) = e(f, f_K)$, since there exists an extremum $p \in K^\circ$ such that $e(f, f_K) = |fp - f_K p|$ and p is witnessed by D. Similarly, for F_p, the standard simplification removing p, $e(F, F_p) = \sigma(p)$.

6.1 Simplification Sequences for Piecewise Monotone Data

We may construct simplification sequences in exact analogy to the continuous case:

Definition 6.2. *Suppose X is a Peano space, $\mathcal{P}$ is a patch collection covering X, and $D \subset X$ is a finite set of sample locations, and $F_0 : D \to \mathbb{R}$ is piecewise monotone data. A sequence of piecewise monotone data $F_1 \ldots F_n : D \to \mathbb{R}$ is a* simplification sequence *for F_0 when for each $i > 0$:*
(1) F_i is a simplification of F_{i-1};
(2) F_n has only global extrema; and
(3) $\sigma(F_i) > \sigma(F_{i-1})$.

Construction of standard simplification sequences for piecewise monotone data goes through in exact analogy to continuous functions (section 4.2).

7 Analysis of Related Methods

In this section we analyze and comment on the Reeb graph simplification method of Carr [4], Carr et al. [5] and Weber et al. [28], the Morse-Smale simplification method of Bremer et al. [1], and the persistence diagram simplification of Edelsbrunner et al. [11, 9].

7.1 Reeb Graph Simplification

Carr [4] describes how to compute the Reeb graph, represented as a contour tree, from an arbitrarily interpolated mesh of any dimension. In practice, Carr [4], Carr et al. [5] and Weber et al. [28] interpolate triangulated and cubic two- and three-dimensional meshes. A simple rule set simplifies the contour tree, a proper subset of Extremal Simplification, as as will be shown shortly. The flat simplification is used to generate the sampled scalar field for a simple running example ([4], Chapter 11); however, implementation-dependent methods are used when visualizing isosurfaces. Weber et al. [28] prefer a smooth alternative to the flat simplification, with the flat simplification being used only when this alternative is not possible. The order in which extrema are removed is determined by pruning contour tree leaves in preference order, using any one of a variety of local geometric measures or persistence. We note that local geometric measures could be straightforwardly and beneficially introduced into Extremal Simplification sequences (section 4.2).

7.1.1 Failure to be Piecewise Monotone

Because the trilinear interpolant may generate a non-monotone function on a cubic patch, the sampled scalar field may fail to be piecewise monotone in the sense of definition 5.4. This means that if one generates a sampled scalar field from the simplified Reeb graph, then in fact the simplified data may not have this same Reeb graph.

This phenomenon is easily illustrated in two dimensions. Consider a square patch having the following sample values at the corners, in clockwise order from top-left: +1, -1, +1, -1. The bilinear interpolation is non-monotone, having a saddle in the centre with value 0. Thus the Reeb graph, i.e. middle space, has the form of the letter X. Carr's branch prune rule [4] simplifies the middle space by removing one of the minima; the simplified middle space has the form of the letter Y, where the maxima have light-factor value +1, the minimum value -1, and the saddle value 0. When we create a new sampled function by flattening, the samples now read, in clockwise order: +1, -1, +1, 0. Note that the middle space of these samples' bilinear interpolation is again an X, not the desired Y.

This problem can be avoided by further subdividing cubes into tetrahedrons. Only those cubes having non-monotone interpolation need to be subdivided; theorem 5.3 guarantees that the flat simplification used by Carr [5] will not require subdivision of other cubes.

Carr's [4] running simple example, which also appears in Carr et al. [5], does not suffer from this problem, because it is built on a triangulation. In his thesis' "Future Work" section, Carr [4] reflects on the problem: " ... (for the simple example) we constructed equivalent surfaces to the simplified surface by hand. This is straightforward for simplicial meshes, where we can change the isovalues at vertices without altering the contour tree. It is less trivial to do this for non-simplicial meshes with complex interpolants, and we would like to examine this problem in more detail."

7.1.2 Carr's Simplification Rules

Rules for Reeb graph simplification are defined by Carr [4]. Since their domain is simply connected, the Reeb graph is a tree [7]. The basic operations are: prune a leaf; and, remove a vertex having order two. Furthermore, a leaf is defined as *prunable* only when the vertex to which it attaches also has another branch going in the same direction.

Pruning a leaf is clearly a quotient having collapse set as kernel. Prunability ensures that this collapse set is extremal. Therefore, Carr's rules provide a subset of the simplifications allowed by Extremal Simplification. Two examples illustrate that the subset is proper:

Example W: Consider a one-dimensional Reeb graph having the form of the letter W. The middle maximum is not a leaf, and therefore cannot be removed by Carr's rules, whereas Extremal Simplification can remove the middle maximum, simplifying the W to a V.

Example K: Consider any function with middle space having the form of the letter K, with light factor as follows. The left maximum and minimum have, respectively, light-factor values +10 and -10; the right maximum has light-factor value +2, the right minimum -2, the left saddle 0, and the right saddle +1. Carr's rules do not allow removal of the top-right maximum, whereas Extremal Simplification does, resulting in a middle space having form of the letter λ with saddle light-factor value 0.

7.2 Morse-Smale Complex Simplification

Bremer et al. [1] use piecewise linear interpolation of samples at triangulated sample locations on a 2-manifold. The continuous function's Morse-Smale complex is simplified as per Edelsbrunner et al. [10]. Morse-Smale simplification in two dimensions is extrema removal. Although Bremer et al. [1] state that extrema are removed in persistence order [11], they do not guarantee that the persistence of surviving extrema is preserved; their statement must be understood accordingly. Heuristic application of smoothing techniques are used to fit data to the simplified Morse-Smale complex while maintaining a specified target error, which they state (without proof) must be greater than half the removed extremum's persistence.

7.3 Persistence Diagram Simplification

Edelsbrunner et al. [11, 9] use the persistence diagram [11] to guide simplification of a piecewise linear scalar field defined by interpolation of triangulated sample locations on a two-dimensional manifold. Direct manipulation of the triangulation is used to the simplify scalar field f to a scalar field g such that g's persistence diagram is a subset of f's, missing exactly those critical points having persistence no more than any given constant ϵ.

In the method of [11], all surviving extrema have persistence reduced by ϵ, whereas in the ϵ-simplification method of [9], all surviving extrema have unchanged persistence. Extremal Simplification does not typically admit an ϵ-simplification; however, neither does it typically change the scale of all surviving extrema (section 4.1.2).

To see that Extremal Simplification cannot do ϵ-simplification, consider Example K from section 7.1.2. Because any collapse set must have boundary upon which f's light factor is constant, removal of the top-right maximum necessarily results in a middle space having a single saddle with light-factor value 0. Therefore, the persistence in the simplified middle space of the right minimum must be 2, as opposed to its original value 3. We leave it to future research to determine how Extremal

Simplification might be extended to encompass ϵ-simplification.

Edelsbrunner et al. [9] show that approximation error for ϵ-simplification must in some cases exceed half the removed extremum's persistence. The function used to demonstrate this result in Part 5 of [9] has the middle space of Example K (see section 7.1.2).

8 Appendices

The remainder of this paper consists of nine appendices, providing proofs for all the theorems. Each proof can be read individually. It is noteworthy that all results of Extremal Simplification follow straightforwardly from simple definitions of point-set topology, with the exception of Appendix A.

A Appendix: Historical Context: Monotone-Light Factorizations

This work considers continuous functions defined on a compact, connected manifold[1] or metric space X. $f : X \to Y$ is called *monotone* whenever $f^{-1}(y)$ is connected, and *light* whenever $f^{-1}(y)$ is discrete (equivalently, $\dim\left(f^{-1}(y)\right) = 0$).

In 1934, Eilenberg [12] and Whyburn [30] introduced monotone-light factorizations independently[2]. A complete proof of the following theorem can be found in Whyburn's book *Analytic Topology* [31].

Theorem (Eilenberg-Whyburn). *Every continuous function $f : X \to Y$ admits a factorization*

$$\begin{tikzcd} X \arrow[rr, "f"] \arrow[dr, "\mu"'] & & Y \\ & M_f \arrow[ur, "\lambda"'] & \end{tikzcd}$$

where μ is monotone and λ is light. This factorization is unique in the sense that the middle space M_f is unique up to homeomorphism.

It is often convenient to denote the monotone-light factorization of f as a triple (M_f, μ, λ) where $f = \lambda\mu$. It should be noted that there are no restrictions on the topological space Y (f's target), however if we require Y to be a manifold then the monotone-light factorization gives a unique factorization system on the category of manifolds and continuous maps[3] [survey reference needed, Joyal].

Both authors used these factorizations for arguments involving dimension (see [18] for a survey of results). For example, the following result is found in [12] (see also [18]):

Theorem (Eilenberg). *Suppose $f : \mathbb{S}^2 \to Y$ is non constant and $\pi_1\left(f^{-1}(y)\right) = 1$ for each $y \in Y$. Then* $\dim\left(f(\mathbb{S}^2)\right) \geq 2$.

In 1946, Reeb [24] gives another point of view in the case $Y = \mathbb{R}$, athough he was not making use of monotone-light factorizations, or any of the existing literature cited above. Assuming $Y = \mathbb{R}$ and $f \in C^2$ we have:

[1] In fact, we could work more generally on Peano spaces. For the most general setting see [31].
[2] It is noted in [6] that it had already been observed in some cases by Kerékjártó [15].
[3] Or, more generally, on the category of Peano spaces and continuous maps.

Theorem (Reeb). *M_f is a graph where $\pi_1 M_f$ is a quotient of $\pi_1 X$.*

Corollary A.1. *In particular, whenever $\pi_1 X = 1$, M_f is a tree.*

It is interesting to note that Kronrod [16] (see [17, 21]) made use of the fact that real valued functions on $\mathbb{R}^2$ and $\mathbb{S}^2$ could be studied by way of a tree structure.

Monotone-light factorizations have been applied in the study of surface area [6, 23]. For example (see [19]) it is known that two parametrizations $f, g : I^2 \to \Sigma$ of a surface Σ give rise to the same surface area whenever f and g share the same light factor[4]. Another nice and very explicit application can be found in Micheal's paper[5] *Cuts* [20].

More recently, the study of monotone-light factorizations has spread to more general categories [find a survey reference], while the graph from Reeb's Theorem (known as the Reeb graph) has found applications among computational geometers (in man cases, without the C^2 requirement). In fact, de Berg and van Krevald [8] prove that M_f is a tree for piecewise linear functions $f : \mathbb{R}^2 \to \mathbb{R}$, while Bajaj et al. [27] state that continuous $f : \mathbb{R}^n \to \mathbb{R}$ gives a tree structure to M_f in general (in in both cases f is defined on compact, simply connected regions of $\mathbb{R}^n$).

For completeness, we include the following.

Proposition A.2. *For continuous $f : X \to \mathbb{R}$, $H_1 X = 0$ implies M_f is a tree.*

To prove this fact we will apply the following lemma.

Lemma A.3. *Suppose $Y \underset{closed}{\subset} U \underset{open}{\subset} X$ where $H_1 X = 0$. Then if $U \smallsetminus Y$ is disconected, $X \smallsetminus Y$ must be disconected as well.*

proof of lemma A.3. By assumption, $H_1 X = 0$. There is a long exact sequence

$$\cdots \longrightarrow H_1 X \longrightarrow H_1(X, U) \longrightarrow H_0 U \longrightarrow H_0 X$$

which gives

$$\cdots \longrightarrow 0 \longrightarrow H_1(X, U) \longrightarrow \mathbb{Z} \xrightarrow{\cong} \mathbb{Z}$$

and hence $H_1(X, U) = 0$. Excision tells us that

$$H_1(X, U) \cong H_1(X \smallsetminus Y, U \smallsetminus Y),$$

and applying this in the long exact sequence

$$\cdots \longrightarrow H_1(X \smallsetminus Y, U \smallsetminus Y) \longrightarrow H_0(U \smallsetminus Y) \longrightarrow H_0(X \smallsetminus Y)$$

we have

$$\cdots \longrightarrow 0 \longrightarrow \underbrace{\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}}_{m} \longrightarrow H_0(X \smallsetminus Y)$$

and hence $H_0(X \smallsetminus Y)$ contains $\underbrace{\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}}_{m}$ as a subgroup. In particular, if $U \smallsetminus Y$ has m components then $X \smallsetminus Y$ has at least m components. ☐

[4]This is known as Kerékjártó equivalence or K-equivalence.

[5]See theorem 1-1 concerning nowhere cutting subsets in Tychonoff spaces.

proof of proposition A.2. Suppose for a contradiction that M_f is not a tree. In particular, let $v_1, \ldots, v_n$ be a sequence of one or more distinct vertices of M_f such that v_i is connected to v_{i+1} by an edge e_i of (M_f) where $v_{n+1} = v_1$. In particular, $\gamma = \amalg_i^n v_i \cup \amalg_i^n e_i$ is an embedding $\mathbb{S}^1 \hookrightarrow M_f$ (a cycle in the graph).

For any v among the v_i, consider the open set U containing v consisting of v together with the set of edges U_i of M_f that meet v (it the case that the cycle contains only one edge we should be careful to choose this open set as some small ϵ-neighborhood of v). Now $\mu^{-1}U$ is open and connected, while $\mu^{-1}U \setminus \mu^{-1}v = \{\mu^{-1}U_i\}$ is disconnected.

Now $v \underset{\text{closed}}{\subset} U \underset{\text{open}}{\subset} M_f$ however $M_f \setminus v$ is necessarily connected due to $\gamma \subset M_f$. Therefore $\mu^{-1}(M_f \setminus v) = \mu^{-1}(M_f) \setminus \mu^{-1}(v)$ is connected which by lemma A.3 is impossible since $H_1 X = 0$.

Therefore no such γ exists, and M_f is a tree. □

In particular, since the abelianization of $\pi_1 X$ gives $H_1 X$ proposition A.2 gives an alternative proof of corollary A.1 without making use of the C^2 hypothesis of Reeb's theorem.

B Approximation Error for Piecewise Monotone Functions

Theorem 3.3. *Suppose $f, g : X \to \mathbb{R}$ are piecewise monotone, where g has fewer extrema than f. Then $e(f, g) \geq \sigma(f)/2$.*

Proof. Consider the special case where g is constant. Then $e(f, g) \geq (\max fX - \min fX)/2 \geq \sigma(f)/2$. We continue with the assumption that g is not constant.

Suppose g has fewer maxima than f. Since g is not constant, f has at least two maxima. Let $p_1 \ldots p_m \in M_f$ be f's maxima, and let $\delta = \min_i \sigma(p_i)$. We show $e(f, g) \geq \delta/2$. Since $\delta \geq \sigma(f)$, the theorem follows.

Denote the monotone-light factorizations of f, g, respectively, as $\mu_f M_f \lambda_f$, $\mu_g M_g \lambda_g$.

For each of f's maxima p_i, define $U_i \subset M_f$ to be the largest neighborhood of p_i such that $p_i \neq q \in U_i$ implies $p_i > q$. The sets $U_1 \ldots U_m$ are pairwise disjoint. Each $\partial U_i \neq \emptyset$, because $\partial U_i = \emptyset$ only when $U_i = M_f$, in which case p_i would be f's only maximum. Note that each $q \in \partial U_i$ is a branch point of M_f such that for some $j \neq i$ the maximum $p_j > q$.

Choose any maximum $p_i \in M_f$. We claim $\lambda_f p_i - \lambda_f q \geq \delta$ for all $q \in \partial U_i$. Choose a $q \in \partial U_i$ such that $\lambda_f p_i - \lambda_f q$ is minimal over all such choices, and choose any $j \neq i$ such that $p_j > q$. Then $\lambda_f p_i \leq \lambda_f p_j$ implies q is a turnaround of p_i and hence $\lambda_f p_i - \lambda_f q = \sigma(p_i)$, whereas $\lambda_f p_i > \lambda_f p_j$ implies that $\lambda_f p_i - \lambda_f q \geq \sigma(p_j)$. In either case, $\lambda_f p_i - \lambda_f q \geq \delta$.

For each $i = 1 \ldots m$, define $Y_i = \mu_f^{-1} U_i$. Each Y_i is open and connected, and has $\partial Y_i \neq \emptyset$. The sets $Y_1 \ldots Y_m$ are pairwise disjoint.

μ_f maps each set ∂Y_i onto ∂U_i. Therefore, $fx - fy \geq \delta$ for any $x \in \mu_f^{-1} p_i$ and $y \in \partial Y_i$.

Suppose that for each $i = 1 \ldots m$ there exists a maximum $q_i \in M_g$ such that $\mu_g^{-1} q_i \subset Y_i$. Then g would have at least m maxima, a contradiction. Therefore we may choose an index k such that for each of g's maxima $q \in M_g$, $\mu_g^{-1} q \not\subset Y_k$.

We now identify points $x, y \in X$ such that $|fx - fy| \geq \delta/2$. First, for f's maximum p_k choose any $x \in \mu_f^{-1} p_k$. Next, let $q \in M_g$ be any maximum of g such that $\mu_g x \leq q$; then there exists $r \in M_g$,

$\mu_g x \leq r \leq q$ such that $(\mu_g^{-1} r \cap \partial Y_k) \neq \emptyset$. Finally, choose any $y \in (\mu_g^{-1} r \cap \partial Y_k)$; then $gx \leq gy$. But $fx - fy \geq \delta$, so at least one of $|fx - gx| \geq \delta/2$ or $|fy - gy| \geq \delta/2$. □

C Appendix: Every Simplification of a Piecewise Monotone Function is Piecewise Monotone

Theorem C.1. *Suppose $f : X \to \mathbb{R}$ is piecewise monotone with monotone-light factorization $\mu M \lambda$. Suppose $K \subset M$ is an extremal collapse set, and suppose g is a simplification of f generated by K. Then g is piecewise monotone. Furthermore, g has fewer monotone pieces than f.*

Proof. Notation: g's middle space and light factor are denoted M_K and λ_K; ϕ_K is the natural map from M to M_K.

We show that g satisfies definition 3.1, using the notation from that definition. In particular, we show that $M_K \setminus M_K^*$ is finite, i.e. that λ_K fails to be locally monotone at only finitely many points.

ϕ_K is one-to-one on $M \setminus K$. Since $M \setminus K$ is open, and since $\lambda q = \lambda_K \phi_K q$ for all $q \in (M \setminus K)$, it follows that λ is locally monotone at $q \in (M \setminus K)$ if and only if λ_K is locally monotone at $\phi_K q$. ϕ_K maps each component $K_i \subset K$ to a unique point $p_i \in M_K$; λ_K may or may not be locally monotone at p_i. In any case, $M_K \setminus M_K^*$ is finite, since K has only finitely many components.

Having shown that g is piecewise monotone, we now show that g has fewer monotone pieces than f.

Suppose $U \subset M$ is a component of M^* such that $U \cap (M \setminus K) \neq \emptyset$. Since ϕ_K is one-to-one on $(M \setminus K)$, there exists exactly one component $U_K \subset M_K^*$ such that $\phi_K^{-1} U_K \supset U$. Furthermore, since ϕ_K is monotone, for every component $U_K \subset M_K^*$ there exists at least component $U \subset M$ such that $\phi_K^{-1} U_K \supset U$. Therefore, M_K^* has no more components than M^*.

Consider any component $K_i \subset K$. We claim that K_i contains at least two nodes of M; therefore K_i contains at least one component of M^*, and consequently M_K^* has strictly fewer components than M^*. Suppose to the contrary that K_i contains only one node $q \in M$. Since λ is constant on ∂K_i, q must be an extremum and must lie in K_i°. But then $\phi_K K_i$ is an extremum of M_K, which means that ∂K_i must be comprised entirely of extrema, contradicting the assumption that q is the only node in K_i. □

D Appendix: Transitivity of Simplification

Theorem D.1. *Suppose $f_0, f_1, f_2 : X \to \mathbb{R}$ are piecewise monotone with, respectively, monotone-light factorizations $\mu_i M_i \lambda_i$ for $i = 0, 1, 2$. Suppose also that $K_1 \subset M_0$ and $K_2 \subset M_0$ are extremal collapse sets such that f_1 is a simplification of f_0 generated by K_1, and f_2 is a simplification of f_0 generated by K_2. Then f_2 is a simplification of f_1 if and only if $K_2 \supset K_1$.*

The theorem follows from the two lemmas below. Lemma D.2 implies that when f_2 is a simplification of f_1 then $K_2 \supset K_1$; lemma D.3 states the converse.

Lemma D.2. *Suppose $f_0, f_1, f_2 : X \to \mathbb{R}$ are piecewise monotone with, respectively, monotone-light factorizations $\mu_i M_i \lambda_i$ for $i = 0, 1, 2$. Suppose also that $J \subset M_0$ and $K \subset M_1$ are extremal collapse sets such that f_1 is a simplification of f_0 generated by J, and f_2 is a simplification of f_1 generated by K. Then $J \cup \phi_J^{-1} K$ is an extremal collapse set of M_0, and f_2 is a simplification of f_0 generated by $J \cup \phi_J^{-1} K$.*

Proof. Suppose $J = \sum J_i$ and $K = \sum K_j$. Let $L = \phi_J^{-1}K$ and $H = J \cup L$. Let $\psi = \phi_J \circ \phi_K : M_0 \to M_2$. We prove the lemma by showing that H is an extremal collapse set; $M_H = M_2$, and $\psi = \phi_H$.

For any component $K_j \subset K$, denote $L_j = \phi_J^{-1}K_j$. L_j has nonempty interior since K_j does, and L_j is connected since ϕ_J is monotone. Thus the components of L are the L_j, in one-to-one correspondence with the components of K.

For each component $J_i \subset J$, let $p_i = \phi_J J_i \in M_1$. Then, either: (1) there exists no component $K_j \subset K$ such that $p_i \in K_j$; or, (2) there exists a unique K_j such that $p_i \in K_j^\circ$; or, (3) there exists a unique K_j such that $p_i \in \partial K_j$. In cases 2 & 3, $L_j \supset J_i$. Thus the components of H comprise all the L_j plus those J_i such that case 1 obtains.

To see that H is a collapse set, we show that λ_0 is constant on the boundary of each component $H_k \subset H$. This follows immediately when $H_k = J_i$ such that case 1 holds, and when $H_k = L_j$ such that no $p_i \in \partial K_j$. So, suppose $H_k = L_j$ such that some $p_i \in \partial K_j$. Note that λ_0 is constant on ∂J_i, with $\lambda_0 \partial J_i = \lambda_1 p_i$. Since λ_1 is constant on ∂K_j, that constant value must be $\lambda_1 p_i$. Therefore, λ_0 is constant on ∂L_j.

Since H is a collapse set, we may form the monotone quotient M_H, identifying all points within each component $H_k \subset H$; denote the natural map $\phi_H : M_0 \to M_H$.

Let $h : M_2 \to M_H$ be defined as follows: For $p \in M_2$, note that ϕ_H is constant on $\psi^{-1}p$; define $h(p) = \phi_H \psi^{-1} p$. Then h is a homeomorphism. Thus, we may identify $M_H = M_2$; this implies $\phi_H = \psi$.

To complete the proof, we must show that H is an extremal collapse set. Suppose component $H_k \subset H$ is such that $p = \phi_H H_k$ is an extremum of M_2. Then every $q \in \partial \phi_K^{-1} p$ is an extremum of M_1 having the same sense as p, and for each such q, every $r \in \partial \phi_J^{-1} q$ is an extremum of M_0 having this same sense. □

Lemma D.3. *Suppose $f_0, f_1, f_2 : X \to \mathbb{R}$ are piecewise monotone with, respectively, monotone-light factorizations $\mu_i M_i \lambda_i$ for $i = 0, 1, 2$. Suppose also that $K_1 \subset M_0$ and $K_2 \subset M_0$ are extremal collapse sets such that f_1 is a simplification of f_0 generated by K_1, and f_2 is a simplification of f_0 generated by K_2. Then f_2 is a simplification of f_1 when $K_2 \supset K_1$.*

Proof. Let $\phi : M_0 \to M_1$ be the natural map, and let $K = \phi K_2$. It follows that K is an extremal collapse set of M_1, and hence f_2 is a simplification of f_1. □

E Appendix: Iterative Construction of Monotone-Light Factorization for $\mathcal{P}D$ Functions

Given patch collection $\mathcal{P} = P_1 \ldots P_N$ for X, and given sample locations $D \subset X$, suppose $f : X \to \mathbb{R}$ is $\mathcal{P}D$; i.e. $f = f_{\mathcal{P}\mathcal{H}}$ for patch functions $\mathcal{H} = h_1 \ldots h_N$ satisfying the conditions of definition 5.2. We construct f's monotone-light factorization in N steps, incorporating one patch per step.

Denote the monotone-light factorization of each h_i as $\mu_{h_i} M_{h_i} \lambda_{h_i}$. For each $n \le N$ denote $\mathcal{P}_n = P_1 \ldots P_n$, $X_n = \cup \mathcal{P}_n$, and $f_n = f | X_n$. Assume the P_i are indexed so that each X_n is connected. Thus X_n is a Peano space; denote the monotone-light factorization of f_n as $\mu_{f_n} M_{f_n} \lambda_{f_n}$.

Step 1: The monotone-light factorization of f_1 is $\mu_{h_1} M_{h_1} \lambda_{h_1}$.

Step $n > 1$: Assume we have already constructed f_{n-1}'s monotone-light factorization $\mu_{f_{n-1}} M_{f_{n-1}} \lambda_{f_{n-1}}$. Our goal is to construct f_n's monotone-light factorization $\mu_{f_n} M_{f_n} \lambda_{f_n}$.

A space M_n – which eventually proves to be f_n's middle space M_{f_n} – may be constructed from the two middle spaces $M_{f_{n-1}}$ and M_{h_n} in two steps. First, construct the topological disjoint union $M_{f_{n-1}} \oplus M_{h_n}$. Second, define M_n as a quotient of this disjoint union: For every $x \in P_n \cap X_{n-1}$, identify the points $\mu_{f_{n-1}} x \in M_{f_{n-1}}$ and $\mu_{h_n} x \in M_{h_n}$. Note that $P_n \cap X_{n-1}$ is nonempty, since we assume $X_n = P_n \cup X_{n-1}$ connected.

Each point $p \in M_n$ corresponds to a unique equivalence class $\tilde{p}$ of points from $M_{f_{n-1}} \oplus M_{h_n}$. The equivalence class $\tilde{p}$ may be singleton, containing a point from either $M_{f_{n-1}}$ or from M_{h_n}; or, the equivalence class may be non-singleton, containing one point from M_{h_n} – only one, since h_n is monotone – and one or more points from $M_{f_{n-1}}$.

Consider any set $S \subset M_n$; denote $\tilde{S} = \bigcup_{p \in S} \tilde{p}$. Then S is open in M_n if and only if both $M_{f_{n-1}} \cap \tilde{S}$ is open in $M_{f_{n-1}}$ and $M_{h_n} \cap \tilde{S}$ is open in M_{h_n}.

We now construct f_n's monotone and light factors $\mu_{f_n} : X_n \to M_n$ and $\lambda_{f_n} : M_n \to \mathbb{R}$. By uniqueness of the monotone-light factorization it follows that M_n is in fact f_n's middle space.

Construct f_n's monotone factor μ_{f_n} as follows. Define $\mu_{f_n} : X_n \to M_n$ by letting $\mu_{f_n} x$ be the unique $p \in M_n$ such that when $x \in X_{n-1}$ then $\mu_{f_{n-1}} x \in \tilde{p}$, and/or when $x \in P_n$ then $\mu_{h_n} x \in \tilde{p}$. Now choose any $p \in M_n$; then $\mu_{f_n}^{-1} p$ is connected, and hence μ_{f_n} is monotone.

Construct f_n's light factor λ_{f_n} as follows. For each $p \in M_n$:
- When $\tilde{p} = \{q\}$ with $q \in M_{h_n}$, then $\lambda_{f_n} p = \lambda_{h_n} q$.
- When $\tilde{p} = \{r\}$ with $r \in M_{f_{n-1}}$, then $\lambda_{f_n} p = \lambda_{f_{n-1}} r$.
- When $\tilde{p} = \{q\ r_1 \ldots r_m\}$, where $q \in M_{h_n}$ and each $r_i \in M_{f_{n-1}}$, then $\lambda_{f_n} p = \lambda_{h_n} q = \lambda_{f_{n-1}} r_1 = \cdots = \lambda_{f_{n-1}} r_m$.

It follows that λ_{f_n} is light.

F Appendix: $\mathcal{P}D$ Interpolation

We use the iterative construction of appendix E to prove the two results referred to in section 5.2.

Theorem F.1. *Given patch collection $\mathcal{P} = P_1 \ldots P_N$ for X, and given sample locations $D \subset X$, suppose $f : X \to \mathbb{R}$ is $\mathcal{P}D$. Then:*
(1) f is piecewise monotone.
(2) Every node of f's middle space is witnessed by D.

Proof. For result (1), we show by induction that f satisfies definition 3.1, using notation from both that definition and from the iterative construction. In particular, for each $n = 1 \ldots N$ we show that $M_{f_n} \setminus M^*_{f_n}$ is finite, i.e. that λ_{f_n} fails to be locally monotone at only finitely many points. Result (2) also utilizes this induction.

Suppose $f = f_{\mathcal{PH}}$ for patch functions $\mathcal{H} = h_1 \ldots h_N$.

For $n = 1$, both results follow immediately from definition 5.2. For $n > 1$, assume f_{n-1} is piecewise monotone and every node of $M_{f_{n-1}}$ is witnessed by $D \cap X_{n-1}$.

$p \in M_{f_n}$ is an extremum if and only if every point in $\tilde{p}$ is an extremum of M_{h_n} or $M_{f_{n-1}}$. It follows that each extremum of M_{f_n} is witnessed by either $D \cap P_n$ or $D \cap X_{n-1}$.

Choose any $p \in M_{f_n}$ such that $\tilde{p} = \{q\}$ with $q \in M_{h_n}$. Then λ_{f_n} is locally monotone at p, since h_n is monotone.

Choose any $p \in M_{f_n}$ such that $\tilde{p} = \{r\}$ with $r \in M_{f_{n-1}}$. Then λ_{f_n} is locally monotone at p if and only if $\lambda_{f_{n-1}}$ is locally monotone at r. Since f_{n-1} is piecewise monotone, there exist only finitely many such p for which λ_{f_n} fails to be locally monotone. Note that each such failure is witnessed by $D \cap X_{n-1}$.

Let $p \in M_{f_n}$ be such that $\tilde{p} = \{q\ r_1 \ldots r_m\}$, where $q \in M_{h_n}$ and each $r_i \in M_{f_{n-1}}$. Then λ_{f_n} fails to be locally monotone at p when $\lambda_{f_{n-1}}$ fails to be locally monotone at one or more of $r_1 \ldots r_m$. Note that there are only finitely many $p \in M_{f_n}$ such this failure occurs; note also that each such failure is witnessed by $D \cap X_{n-1}$. We continue the analysis with the assumption that $\lambda_{f_{n-1}}$ is locally monotone at each $r_1 \ldots r_m$.

Because $\tilde{p} = \{q\ r_1 \ldots r_m\}$, it must be the case that $\mu_{h_n}^{-1} q \cap (P_n \cap X_{n-1}) \neq \emptyset$. Let $\mathcal{Q} \subset \mathcal{P}$ be all patches P_j with $j < n$ such that $\mu_{h_n}^{-1} q \cap (P_n \cap P_j) \neq \emptyset$, and let $\mathcal{K}$ be the collection of all components $K \subset (P_n \cap P_j)$ for any and all $P_j \in \mathcal{Q}$ such that $\mu_{h_n}^{-1} q \cap K \neq \emptyset$. Then $\mathcal{K}$ is finite; denote its elements as $K_1 \ldots K_k$.

For each K_i, let I_i be the real closed interval $f K_i$. When $\lambda_{h_n} q \in I_i^\circ$ for each $i = 1 \ldots k$, then it follows that λ_{f_n} is locally monotone at p. Therefore λ_{f_n} may fail to be locally monotone at p only if there exists an index j such that $\lambda_{h_n} q = \min I_j$ or $\lambda_{h_n} q = \max I_j$. Note that this condition is not sufficient for failure of local monotonicity: additionally, we would need that for every neighborhood $U \subset M_{h_n}$ of q, $\mu_{h_n}^{-1} q \not\subset (P_n \cap X_{n-1})$. However, the condition is sufficient to conclude that there are only finitely many points $p \in M_{f_n}$ upon which λ_{f_n} may fail to be locally monotone, and that each failure is witnessed by $D \cap X_n$. $\square$

G All $\mathcal{P}D$ Interpolations Have Identical Middle Spaces and Light Factors

Theorem 5.3. *Let $D \subset X$ be sample locations, let $F : D \to \mathbb{R}$ be scalar data, and let $\mathcal{P}$ be a patch collection for X. Then any two $\mathcal{P}D$ functions f, f' interpolating F have identical middle spaces and light factors.*

Proof. Suppose $f = f_{\mathcal{P}\mathcal{H}}$ and $f' = f_{\mathcal{P}\mathcal{H}'}$, where $\mathcal{P} = P_1 \ldots P_N$, $\mathcal{H} = h_1 \ldots h_N$ and $\mathcal{H}' = h'_1 \ldots h'_N$. Denote the monotone-light factorization of each h_i and h'_i as, respectively, $\mu_{h_i} M_{h_i} \lambda_{h_i}$ and $\mu_{h'_i} M_{h'_i} \lambda_{h'_i}$. For each $n \leq N$ denote $\mathcal{P}_n = P_1 \ldots P_n$, $X_n = \cup \mathcal{P}_n$, and $f_n = f | X_n$ and $f'_n = f' | X_n$. Assume the P_i are indexed so that each X_n is connected. X_n is a Peano space and f_n, f'_n are piecewise monotone; denote the monotone-light factorization of f_n and f'_n as, respectively, $\mu_{f_n} M_{f_n} \lambda_{f_n}$ and $\mu_{f'_n} M_{f'_n} \lambda_{f'_n}$.

The theorem is proved by exhibiting a homeomorphism between f's and f''s middle spaces that commutes with their light factors. The proof proceeds by induction, using the notation from the iterative construction in Appendix E. We show that Property Z, below, holds for each of $\mathcal{P}_1 \ldots \mathcal{P}_N$. Since $\mathcal{P}_N = \mathcal{P}$, this completes the proof.

Property Z. *Suppose $\mathcal{Q} \subset \mathcal{P}$. Denoting $X_{\mathcal{Q}} = \cup \mathcal{Q}$, suppose $X_{\mathcal{Q}}$ is connected. Then $X_{\mathcal{Q}}$ is a Peano space; denote the monotone-light factorization of $f | X_{\mathcal{Q}}$ and $f' | X_{\mathcal{Q}}$ as, respectively, $\mu M \lambda$ and $\mu' M' \lambda'$. Then $\mathcal{Q}$ has* property Z *when there exists a homeomorphism $\phi : M \to M'$ such that:*

- *(1)* $\lambda' = \phi \circ \lambda$
- *(2) Suppose patch $P \in (\mathcal{P} \setminus \mathcal{Q})$ and patch $Q \in \mathcal{Q}$ are such that $P \cap Q \neq \emptyset$:*
  - *(2a)* $f(P \cap Q) = f'(P \cap Q)$

(2b) $\mu' x' = \phi \mu x$ *for each* $x, x' \in (P \cap Q)$ *such that* $fx = f'x'$.

We claim that for every $P_i \in \mathcal{P}$ the singleton patch collection $\{P_i\}$ has Property Z. This follows directly from the definition 5.2, because $f|D = f'|D$.

In particular, $\mathcal{P}_1$ has Property Z, starting the induction with $n = 1$.

Suppose $n > 1$, and assume by induction that $\mathcal{P}_{n-1}$ has Property Z; let $\phi_{\mathcal{P}(n-1)} : M_{f_{n-1}} \to M_{f'_{n-1}}$ be the relevant homeomorphism. $\{P_n\}$ also has Property Z; let $\phi_{\{P_n\}} : M_{h_n} \to M_{h'_n}$ be the relevant homeomorphism.

We define homeomorphism $\phi_{\mathcal{P}n} : M_{f_n} \to M_{f'_n}$ using the two homeomorphisms $\phi_{\mathcal{P}(n-1)}$ and $\phi_{\{P_n\}}$. Let $p \in M_{f_n}$; then $\phi_{\mathcal{P}n} p = p'$, where:

When $\tilde{p} = \{q\}$ with $q \in M_{h_n}$, then $p' \in M_{h'_n}$ is the unique point having $\tilde{p}' = \{\phi_{\{P_n\}} q\}$.
When $\tilde{p} = \{r\}$ with $r \in M_{f_{n-1}}$, then $p' \in M f_{n-1}$ is the unique point having $\tilde{p}' = \{\phi_{\mathcal{P}(n-1)} r\}$.
When $\tilde{p} = \{q\ r_1 \ldots r_m\}$, where $q \in M_{h_n}$ and each $r_i \in M_{f_{n-1}}$, then $p' \in M f_{n-1}$ is the unique point having $\tilde{p}' = \{\phi_{\{P_n\}} q\ \phi_{\mathcal{P}(n-1)} r_1 \ldots \phi_{\mathcal{P}(n-1)} r_m\}$.

We complete the proof by showing that $\phi_{\mathcal{P}n}$ satisfies the conditions of Property Z.

(Z.1) By construction.
(Z.2) Suppose $m > n$ such that patch P_m has nonempty intersection with patch P_i, where $1 \leq i \leq n$.
  (Z.2a) Since Property Z.2a holds for $\{P_i\}$, $f(P_m \cap P_i) = f'(P_m \cap P_i)$.
  (Z.2b) Choose any $x, x' \in (P_m \cap P_i)$ such that $fx = f'x'$. When $i = n$ then $\phi_{\{P_n\}} \mu_{h_n} x = \mu_{h'_n} x'$, since Property Z.2b holds for $\{P_n\}$. When $i < n$ then $\phi_{\mathcal{P}(n-1)} \mu_{f_{n-1}} x = \mu_{f'_{n-1}} x'$, since Property Z.2b holds for $\{P_{n-1}\}$. In either case, when $\mu_{f_n} x = p \in M_{f_n}$, then $\mu_{f'_n} x = \phi_{\mathcal{P}n} p$, where $p' = \phi_{\mathcal{P}n} p \in M_{f'_n}$ is the unique point such that $\phi_{\{P_n\}} \mu_{h_n} x \in \tilde{p}'$ and/or $\phi_{\mathcal{P}(n-1)} \mu_{f_{n-1}} x \in \tilde{p}'$.

□

H Appendix: Approximation Error Bound for Piecewise Monotone Data

Theorem 5.5. *Suppose* $F, G : X \to \mathbb{R}$ *are piecewise monotone data, where* G *has fewer extrema than* F*. Then* $e(F, G) \geq \sigma(F)/2$.

Proof. The proof is similar to the proof of theorem 3.3. Denote F's middle space as M_F. Suppose G has fewer maxima than F. Let $p_1 \ldots p_m \in M_F$ be F's maxima, and let $\delta = \min_i \sigma(p_i)$. We show $e(F, G) \geq \delta/2$. Since $\delta \geq \sigma(F)$, the theorem follows.

Let f, g be $\mathcal{PD}$ interpolations of, respectively, F, G. Denote the monotone-light factorizations of f, g, respectively, as $\mu_f M_f \lambda_f$, $\mu_g M_g \lambda_g$. Note that $M_f = M_F$, and so f's maxima are $p_1 \ldots p_m$.

For each of f's maxima p_i, define $U_i \subset M_f$ to be the largest neighborhood of p_i such that $p_i \neq q \in U_i$ implies $p_i > q$, and define $Y_i = \mu_f^{-1} U_i$. As in the proof of theorem 3.3, the sets $Y_1 \ldots Y_m$ are pairwise disjoint, and for each index i: $\partial Y_i \neq \emptyset$; f is constant on each component of ∂Y_i; and $fx - fy \geq \delta$ for any $x \in \mu_f^{-1} p_i$ and $y \in \partial Y_i$.

Choose any Y_i and suppose there exists a patch $P \in \mathcal{P}$ is such that $Y_i \cap P^\circ \neq \emptyset$ and $\partial Y_i \cap P^\circ \neq \emptyset$. We note the following properties: Choose any component $C \subset (\partial Y_i \cap P)$ such that $C \cap P^\circ \neq \emptyset$. f is constant on C. Since $f|P$ is monotone, $P \smallsetminus f^{-1} fC$ comprises at least one component K_1 and perhaps a second component K_2, where $x \in K_1$ implies $fx > fC$ and $x \in K_2$ implies $fx \leq fC$. Note that $Y_i \cap P^\circ \subset K_1$. It follows that the component C is unique, i.e. no other component of $\partial Y_i \cap P$ intersects P°.

For each i, define Z_i as the interior of the union of all patches $P \in \mathcal{P}$ such that $Y_i \cap P^\circ \neq \emptyset$.

Choose any maximum p_i of f. We claim $fx - fy \geq \delta$ for any $x \in \mu_f^{-1} p_i$ and any $y \in \partial Z_i$. Let P be any patch such that $Y_i \cap P^\circ \neq \emptyset$ and $\partial Z_i \cap P \neq \emptyset$. Then either $Y_i \supset P^\circ$ or $\partial Y_i \cap P^\circ \neq \emptyset$. In the first case $\partial Y_i \supset (\partial Z_i \cap P)$, so $y \in (\partial Z_i \cap P)$ implies $y \in \partial Y_i$ and so $fx - fy \geq \delta$. In the second case, let C be the unique component of $\partial Y_i \cap P$ such that $C \cap P^\circ \neq \emptyset$; then $y \in (\partial Z_i \cap P)$ implies $fy \leq fC$, and thus $fx - fy \geq \delta$.

We claim the sets $Z_1 \dots Z_m$ are pairwise disjoint. To see this, suppose that for some patch $P \in \mathcal{P}$ there exist indices $i \neq j$ such that $Y_i \cap P^\circ \neq \emptyset$ and $Y_j \cap P^\circ \neq \emptyset$. It follows that $\partial Y_i \cap P^\circ \neq \emptyset$ and $\partial Y_j \cap P^\circ \neq \emptyset$. Let C_i be the unique component of $\partial Y_i \cap P$ such that $C_i \cap P^\circ \neq \emptyset$ and let K_i be the component of $P \smallsetminus C_i$ containing $Y_i \cap P^\circ$; define C_j and K_j similarly. Then $fC_i \leq fC_j$ implies $K_i \supset K_j$, and $fC_j \leq fC_i$ implies $K_j \supset K_i$, neither of which is possible since $Y_i \cap Y_j = \emptyset$.

Suppose that for each $i = 1 \dots m$ there exists maximum $q_i \in M_g$ such that $\mu_g^{-1} q \subset Z_i$. Then g would have at least m maxima, a contradiction. Therefore we may choose an index k such that for each of g's maxima $q \in M_g$, $\mu_g^{-1} q \not\subset Z_k$.

We now identify points $x, y \in D$ such that $|fx - fy| \geq \delta/2$. First choose any $x \in (D \cap \mu_f^{-1} p_k)$. Next, let $q \in M_g$ be any maximum of g such that $\mu_g x \leq q$; then there exists $r \in M_g$, $\mu_g x \leq r \leq q$ such that $\mu_g^{-1} r \cap \partial Z_k \neq \emptyset$. Choose any $z \in \mu_g^{-1} r \cap \partial Z_k$; then there exist patches $P, Q \in \mathcal{P}$ such that $z \in P \cap Q$, $Y_k \cap P^\circ \neq \emptyset$ and $Y_k \cap Q^\circ = \emptyset$. Note that $P \cap Q \subset \partial Z_k$. Let K be the component of $P \cap Q$ containing z, and let $y \in (D \cap K)$ be such that $gy = \max gK$. Then $gx \leq gy$. But $fx - fy \geq \delta$, so at least one of $|fx - gx| \geq \delta/2$ or $|fy - gy| \geq \delta/2$. $\square$

I Appendix: Every Flat Simplification of Piecewise Monotone Data is Piecewise Monotone

Theorem I.1. *Let $D \subset X$ be sample locations, and let $\mathcal{P}$ be a patch collection covering X. Choose any $\mathcal{P}D$ function f, and let K be any extremal collapse set for f. Then K's flat simplification f_K is $\mathcal{P}D$.*

Proof. Suppose $K = \sum_{i=1 \dots n} K_i$. Then the flat simplification f_K can be sequentially derived by flattening the components K_i one at a time: Letting $K' = \sum_{i=2 \dots n} K_i$, then $f_K = (f_{K_1})_{K'}$. Therefore it suffices to prove the theorem for K connected.

Denote f and f_K's monotone-light factorizations, respectively, as $\mu M \lambda$ and $\mu_K M_K \lambda_K$.

Consider any patch $P \in \mathcal{P}$; we show that $f_K|P$ satisfies definition 5.2, utilizing three cases regarding P's intersection with $\mu^{-1} K$.

When $P \cap \mu^{-1}(K) = \emptyset$ then $f_K|P = f|P$, so definition 5.2 is satisfied.

When $P \subset \mu^{-1}(K^\circ)$ then f_K is constant on P, so definition 5.2 is satisfied.

When $P \subset \mu^{-1}(\partial K)$ is nonempty, then $f|P$ monotone and f constant on $\mu^{-1}(\partial K)$ imply $P \setminus \mu^{-1}(\partial K)$ has either one or two components. Since $f|P \neq f_K|P$, f and f_K differ on exactly one these components; denote this component as R. f is non-constant on R, whereas f_K is constant on R. Therefore $f_K|P$ is monotone, and the points of $D \cap P$ that witnessed $f|P$'s minimum and maximum also witness $f_K|P$'s minimum and maximum, and similarly for the components of P's intersections with other patches in $\mathcal{P}$. □

References

[1] P.-T. Bremer, B. Hamann, H. Edelsbrunner, and V. Pascucci. A topological hierarchy for functions on triangulated surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 10(4):385 – 396, 2004.

[2] M Brooks. Approximation complexity for piecewise monotone functions real data. *Computers and Mathematics with Applications*, 27(8):47 – 58, 1994.

[3] M. Brooks, Y. Yan, and D Lemire. Scale-based monotonicity analysis in qualitative modelling with flat segments. In *Proc. International Joint Conf. on Artificial Intelligence (IJCAI)*, August 2005.

[4] Hamish Carr. Topological manipulation of isosurfaces. PhD Thesis, University of British Columbia, 2004.

[5] Hamish Carr, Jack Snoeyink, and Michiel van de Panne. Simplifying flexible isosurfaces using local geometric measures. In *IEEE Visualization*, pages 497–504. IEEE Computer Society, 2004.

[6] Lamberto Cesari. *Surface area.* Annals of Mathematics Studies, no. 35. Princeton University Press, Princeton, N. J., 1956.

[7] K. Cole-McLaughlin, H. Edelsbrunner, V. Natarajan J. Harer, and V. Pascucci. Loops in reeb graphs of 2-manifolds. *Discrete Comput. Geom.*, 32:231 – 244, 2004.

[8] M. de Berg and M. van Kreveld. Trekking in the Alps without freezing or getting tired. *Algorithmica*, 18(3):306–323, 1997. First European Symposium on Algorithms (Bad Honnef, 1993).

[9] H. Edelsbrunner, D. Morozov, and V. Pascucci. Persistence-sensitive simplification of functions on 2-manifolds. *Proc. 22nd ACM Symp. Computational Geometry*, 2006.

[10] Herbert Edelsbrunner, John Harer, and Afra Zomorodian. Hierarchical morse-smale complexes for piecewise linear 2-manifolds. *Discrete Comput. Geom.*, 28:87 – 107, 2003.

[11] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete Comp. Geo.*, 28:511 – 533, 2002.

[12] S. Eilenberg. Sur les transformations continues d'espaces métriques compacts. *Fundam. Math.*, 22:292–296, 1934.

[13] Attila Gyulassy, Vijay Natarajan, Valerio Pascucci, Peer Timo Bremer, and Bernd Hamann. A topological approach to simplification of three-dimensional scalar fields. *IEEE Transactions on Visualization and Computer Graphics*, pages 474 – 484, 2006.

[14] H. Hoppe. Progressive meshes. *SIGGRAPH*, 1996.

[15] B. V. Kerékjártó. On parametric representation of continuous surfaces. *Proc. Nat. Acad. Sci.*, 10:267–271, 1924.

[16] A. S. Kronrod. On functions of two variables. *Uspehi Matem. Nauk (N.S.)*, 5(1(35)):24–134, 1950.

[17] E. M. Landis and I. M. Yaglom. Remembering A. S. Kronrod. *Math. Intelligencer*, 24(1):22–30, 2002. Translated from the Russian by Viola Brudno and edited by Walter Gautschi.

[18] Harriet Lord. Monotone-light factorizations: a brief history. preprint.

[19] Harriet Lord. An application of abstract nonsense to surface area. *J. Interdisciplinary Studies*, 12:179–188, 1999.

[20] E. Michael. Cuts. *Acta Math.*, 111:1–36, 1964.

[21] H.P. Mulholland. AMS review of Kronrod's paper"On functions of two variables".

[22] Sam B. Nadler, Jr. *Continuum Theory: An Introduction.* Pure and Applied Mathematics. Marcel Decker, New York, 1992.

[23] Tibor Radó. *Length and Area.* American Mathematical Society Colloquium Publications, vol. 30. American Mathematical Society, New York, 1948.

[24] Georges Reeb. Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique. *C. R. Acad. Sci. Paris*, 222:847–849, 1946.

[25] Shigeo Takahashi, Yuriko Takeshima, and Issei Fujishiro. Topological volume skeletonization and its application to transfer function design. *Graphical Models*, 66(1), January 2004.

[26] V. A. Ubhaya. Isotone approximation i. *J. Approximation Theory*, 12:146–159, 1974.

[27] M. van Kreveld, R. van Oostrum, C.L. Bajaj, V. Pascucci, and D. Schikiore. Contour trees and small seed sets for isosurface traversal. In *Proceedings of the 13th Annual Symposium on Computational Geometry*, pages 212–220, 1997.

[28] G.H. Weber, S.E. Dillard, H. Carr, V. Pascucci, and B. Hamann. Topology-controlled volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 13(2), 2007.

[29] Gordon Whyburn and Edwin Duda. *Dynamic topology.* Springer-Verlag, New York, 1979. Undergraduate Texts in Mathematics, Kelley.

[30] Gordon Thomas Whyburn. Non-alternating transformations. *Amer. J. Math.*, 56:294–302, 1934.

[31] Gordon Thomas Whyburn. *Analytic Topology.* American Mathematical Society Colloquium Publications, v. 28. American Mathematical Society, New York, 1942.