Proceedings of the Joint Austrian Computer Vision and Robotics Workshop 2020 DOI: 10.3217/978-3-85125-752-6-31

Border Propagation: A Novel Approach To Determine Slope Region Decompositions

Florian Bogner
TU Wien
e1225415@student.tuwien.ac.at
Co-First Author

Alexander Palmrich
TU Wien
apalmrich@gmail.com
Co-First Author

Walter G. Kropatsch
TU Wien
krw@prip.tuwien.ac.at
Supervisory Author

Abstract. *Slope regions are a useful tool in pattern recognition. We review theory about slope regions and prove a theorem linking monotonic paths and the connectedness of levelsets. Unexpected behavior of slope regions in higher dimensions is illustrated by two examples. We introduce the* border propagation *(BP) algorithm, which decomposes a d-dimensional array ($d \in \mathbb{N}$) of scalar values into slope regions. It is novel as it allows more than 2-dimensional data.*

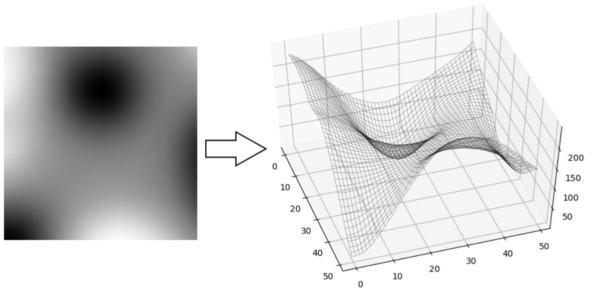

Figure 1. gray-scale to height-map conversion

1. Introduction

In this section we develop an intuitive understanding of the term *slope region* [3] and its generalization to higher dimensions. The concise definition of the terms already employed here is reserved for the next section.

Consider an image, either gray-scale or in color. If it is a color image, it can be decomposed into its color channels (red-green-blue), which can individually be read as gray-scale images. We consider pixel intensity of one such gray-scale image as the height of a landscape, yielding a 2D surface in 3D space. The surface will have peaks in areas where the image is bright, and will have dales in dark areas.

Now our aim is to partition the surface into *regions* (i.e. subsets) in a particular way: We require each region to consist only of a single slope, by which we mean that we can ascend (or descend) from any given point of the region, to any other given point of the region, along a path that runs entirely within the region. Such a decomposition is not unique, but we can at least try to get a partition *as coarse as possible*, meaning that we merge slope regions if the resulting subset is still a slope region, and we iterate this until no further change occurs. There might be many different coarsest slope decompositions.

The criterion we used to describe slopes, any two points being connected by either an ascending or a descending path, can easily be used in higher dimensions. Think of a computed tomography scan, which will yield gray-scale data, but not just on a 2D image, but rather on a 3D volume. We want to partition the 3D volume, such that any two points in a region can be connected via an either ascending or descending path within the region. Recall that *ascending* and *descending* refers to the intensity value of the tomography scan as we move in the volume. For piecewise linear functions on a volume, decompositions were introduced in [1].

By abstracting from image and tomography to a real function $f : \Omega \to \mathbb{R}$ defined on some subset of $\mathbb{R}^n$ (think of it as the pixel intensity function), and by rigorously defining a coarsest slope decomposition, we can lift the concept to arbitrary dimensions in a mathematically concise fashion.

2. Defining Slope Regions

In this and the following chapters we will consider a topological space $(\Omega, \mathcal{T})$ with a continuous function $f : \Omega \to \mathbb{R}$. In practice or for ease of imagination, $(\Omega, \mathcal{T})$ will typically be a rectangle or cuboid subset of $\mathbb{R}^2$ or $\mathbb{R}^3$ equipped with the euclidean topology and f will describe a continuous image or 3D-scan.

Definition 2.1. A *path* is a continuous function from the real interval $[a, b]$ (with $a < b$) into a topological space Ω.

Definition 2.2. Two points $x \neq y$ in a topological space Ω are called *path-connected* if and only if there exists a path $\gamma : [a, b] \to \Omega$ with $\gamma(a) = x$ and $\gamma(b) = y$.

Definition 2.3. The set of all points which are path-connected to a point $x \in \Omega$ is the *connected component of* x:

$$[x] := \{y \in \Omega | x \text{ is path-connected to } y\}$$

Any subset of Ω which can be written in above way (for a suitable choice of x) is called a *connected component*.

Definition 2.4. A path $\gamma : [a, b] \to \Omega$ is called *monotonic* if and only if the whole path is ascending or the whole path is descending, meaning the first or second formula below has to hold, respectively:

$$\forall s, t \in [a, b] : s < t \Rightarrow f(\gamma(s)) \leq f(\gamma(t))$$
$$\forall s, t \in [a, b] : s < t \Rightarrow f(\gamma(s)) \geq f(\gamma(t))$$

Definition 2.5. Let $R \subset \Omega$. R is called *slope region* or *monotonically connected* if and only if for all $x, y \in R$ there exists a monotonic path $\gamma : [a, b] \to R$ with $\gamma(a) = x$ and $\gamma(b) = y$.

Definition 2.6. A family of sets $\{A_i \subset \Omega \mid i \in I\}$ is called a *slope region decomposition* if and only if:

- A_i is a slope region for all $i \in I$
- $\forall i, j \in I : i \neq j \Rightarrow A_i \cap A_j = \emptyset$.
- $\bigcup_{i \in I} A_i = \Omega$

Definition 2.7. Consider two slope region decompositions $\mathcal{A} = \{A_i \subset \Omega \mid i \in I\}$ and $\mathcal{B} = \{B_j \subset \Omega \mid j \in J\}$. We call $\mathcal{A}$ *coarser than* $\mathcal{B}$, alternatively $\mathcal{B}$ *finer than* $\mathcal{A}$, in Symbols $\mathcal{A} \succeq \mathcal{B}$ if and only if

$$\forall j \in J \, \exists i \in I : B_j \subset A_i.$$

Theorem 2.8. *$\succeq$ is a partial order, i.e. fulfills reflexivity, antisymmetry and transitivity.*

Proof: Straight forward. Antisymmetry follows from the decomposition property. □

Definition 2.9. A slope region decomposition $\mathcal{A}$ is called *maximally coarse* or simply *coarse* if and only if there is no other coarser slope region decomposition.

We can apply Zorn's lemma [6] to the partial order $\succeq$, which yields the existence of maximal elements. For this we need to show that chains have upper bounds.

Theorem 2.10. *For any ascending chain of slope region decompositions $(\mathcal{A}_i)_{i \in I}$, that is $t \geq s \Rightarrow \mathcal{A}_t \succeq \mathcal{A}_s$, there is a slope region decomposition $\mathcal{A}_\infty$ satisfying $\forall i \in I : \mathcal{A}_\infty \succeq \mathcal{A}_i$.*

Proof: We consider the equivalence relation "*connected in $\mathcal{A}_i$*" for two points $x, y \in \Omega$:

$$x \sim_i y :\Leftrightarrow \exists A \in \mathcal{A}_i : x \in A \wedge y \in A$$

The equivalence relation is a subset of Ω^2, and $\mathcal{A}_t \succeq \mathcal{A}_s$ implies $\sim_t \supset \sim_s$. This suggests the use of $\sim_\infty := \bigcup_{i \in I} \sim_i$ to get an upper bound. Indeed the equivalence classes of $\sim_\infty$ yield a partition $\mathcal{A}_\infty$ of Ω, which is coarser than any $\mathcal{A}_i$. But do they form a slope region decomposition? Yes: For any two fixed points x, y to be $\sim_\infty$-connected, they need to be $\sim_i$-connected for some $i \in I$. So there is a monotonic path linking x and y in $A = [x]_{\sim_i} \subset [x]_{\sim_\infty}$, by which they are monotonically connected in $\mathcal{A}_\infty$. Therefor $\mathcal{A}_\infty$ is a slope region decomposition. □

Hence every set Ω has a coarse decomposition.

Theorem 2.11. *Let $A \subset \Omega$ be a path-connected set. A is a slope region if and only if all levelsets of f in A are path-connected, i.e.*

$$\forall c \in \mathbb{R} : f^{-1}(\{c\}) \cap A \text{ is path-connected.}$$

Proof: "$\Rightarrow$" via contraposition:

Suppose there exists a $c \in \mathbb{R}$ with $L := f^{-1}(c) \cap A$ not path-connected. We decompose L in its components and pick x and y from different components. Since $f(x) = f(y) = c$ a monotonic path between x and y would have to lie completely in L. However, since x and y are from different components, they cannot be connected by a path in L and therefore cannot be connected with a monotonic path. Therefore, A is not a slope region.

"$\Leftarrow$" via *ironing out* an arbitrary path:

Given $x, y \in A$ we have to find a monotonic path γ. Without loss of generality suppose $f(x) \geq f(y)$. Since A is path-connected, there exists an (not necessarily monotonic) path $\gamma_0 : [a, b] \to A$ from x to

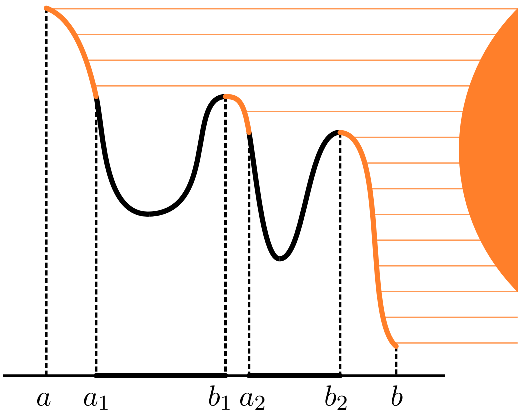

Figure 2. applying the Rising Sun lemma

y. Using the Rising Sun lemma [5] on the continuous function $f \circ \gamma_0$ we get the *shadow* $S = \bigcup_{i \in I}(a_i, b_i)$ consisting of at most countably many intervals.

S consists of the points which contradict the monotonicity of $f \circ \gamma_0$, thus we want to *iron out* these points.

Let $c_n := f(a_n) = f(b_n)$. Since the levelset of c_n is path-connected, we can connect $\gamma_0(a_n)$ and $\gamma_0(b_n)$ with a level path $\gamma_n^* : [a_n, b_n] \to A$.

Finally, we define:

$$\gamma(\sigma) := \begin{cases} \gamma_n^*(\sigma) & \sigma \in [a_n, b_n] \\ \gamma_0(\sigma) & \text{elsewhere} \end{cases}$$

γ is a monotonic path from x to y, thus A is a slope region. □

3. Results In The Plane And Counterexamples In Higher Dimensions

There are two theorems ([3] Lemma 1 and [3] Lemma 2) that are useful, but only hold if $\Omega \subset \mathbb{R}^2$, not in general if $\Omega \subset \mathbb{R}^d$ for $d > 2$. But first, we prove a lemma.

Theorem 3.1. *Let A be a slope region. Then the closure $\bar{A}$ is also a slope region.*

Proof: Follows from continuity of f. □

The following theorem is only formulated for closed slope regions, but because of the above theorem this is not a big restriction.

Theorem 3.2. *Let $d = 2$ and $A \subset \mathbb{R}^d$ be a closed and bounded slope region. Let $(\partial A_i)_{i \in I}$ be an enumeration of the connected components of the boundary ∂A. For $i \in I$, if ∂A_i is homeomorphic to a circle, then $f|_{\partial A_i}$ has at most one local minimum and one local maximum (but the extrema might be spread out in a connected plateau).*

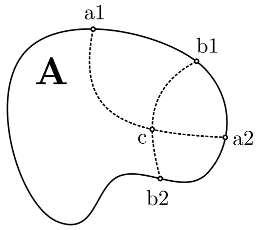

Figure 3. a sketch of the situation in Theorem 3.2

Proof: Assume there are two local minima $a_1, a_2 \in \partial A_i$ with $f(a_1) \leq f(a_2)$. Since ∂A_i is homeomorphic to a circle, there have to be local maxima b_1 and $b_2 \in \partial A_i$ between them with $f(a_2) < f(b_1) \leq f(b_2)$, one on each arc.

Since A is a slope region, a_1 and a_2, as well as b_1 and b_2 have to be connected by a monotonic path. Because of the Jordan Curve Theorem [2, p.169], these paths have to cross in a point $c \in A$. But this yields a contradiction: $f(c) \leq f(a_2) < f(b_1) \leq f(c)$. Thus the assumption of the existence of two local minima has to be false. □

Note: The circle assumption is actually unnecessary and the proof without it remains the same in spirit, but becomes inhibitivly technical, which is why we omit it here.

Example 3.3. Let $\Omega = \mathbb{R}^3$ and $A = B_1(0,0,0)$ be the closed unit ball. Let f be the distance to the x-Axis.

$$f : \mathbb{R}^3 \to \mathbb{R} : (x, y, z) \mapsto \sqrt{y^2 + z^2}$$

The levelsets of f in A are either the x-Axis for $f \equiv 0$ or the sides of cylinders for $f > 0$. In any case, they are connected. Thus, by Theorem 2.11, A is a slope region. ∂A has one connected component, which is the unit sphere. $f|_{\partial A}$ has two local minima, which are the intersections with the x-Axis, $(1, 0, 0)$ and $(-1, 0, 0)$.

Thus, the previous theorem does not hold in $\mathbb{R}^3$. In fact, it does not hold in any $\mathbb{R}^d$ for $d > 2$. There is also no limit on the number of local minima on the surface of a slope region.

Theorem 3.4. *Let $d = 2$ and $A \subset \mathbb{R}^d$ be a slope region. Let $s \in A$ be a saddle point. Then, $s \in \partial A$.*

Proof: Assume s is an interior point of A, which means there is a open set U with $s \in U \subset A$. s being a saddle point means there is a neighborhood $V \subset U$ so that $V_- := V \cap [f < f(s)]$ as well as $V_+ := V \cap [f > f(s)]$ decompose into two or more connected components.

Pick a_1, a_2 from different components of V_- as well as b_1, b_2 from different components of V_+. a_1 and a_2 have to be connected by a monotonic path, but this path has to move outside of V since the points are from different components of V_- and by virtue of being monotonic, the path cannot go through V_+. Analogue for b_1 and b_2.

Again by the Jordan Curve Theorem, these two paths have to cross in some point c, which again yields a contradiction.

$$f(c) \leq max(f(a_1), f(a_2)) < \\ < f(s) < min(f(b_1), f(b_2)) \leq f(c)$$

Thus the assumption that s is a interior point has to be false. □

Example 3.5. Let $\Omega = A = \mathbb{R}^3$. Let f be the distance from the unit circle laying in the x-y-plane.

$$f : \mathbb{R}^3 \to \mathbb{R} : (x, y, z) \mapsto \\ \begin{cases} \left\| \left(x - \frac{x}{||x,y||_2}, y - \frac{y}{||x,y||_2}, z \right) \right\|_2 & \|(x,y)\|_2 \neq 0 \\ \|(1,0,z)\|_2 & \|(x,y)\|_2 = 0 \end{cases}$$

Again, let us look at the levelsets to show A is a slope region. The levelset of $f \equiv 0$ is the unit circle. For $0 < f < 1$ the levelsets are tori. $f \equiv 1$ marks a transition and the levelset is a torus with its hole closed. Then, for $f > 1$ the levelsets look like the exterior surface of a self intersecting torus, topologically equivalent to a sphere. All these levelsets are connected. Thus, A is indeed a slope region.

Now consider the point $(0, 0, 0)$. Along the x and y-direction it is a local maximum, however along the z-direction it is a local minimum. Thus, it is a saddle point. Therefore, theorem 3.4 does not hold in higher dimensions.

4. Motivating The Border Propagation (BP) Algorithm

Now we will work our way to the central insights on which the border propagation algorithm (*BP*) hinges. Let us develop ideas for smooth

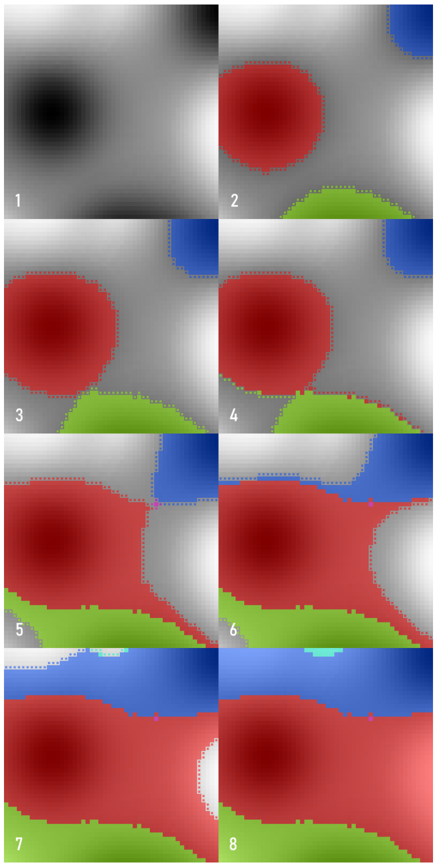

Figure 4. evolution of discretized regions during the algorithm

(hyper-)surfaces first, and deal with discrete variants in the next section.

Slope regions can be constructed and grown in a straight-forward iterative manner by sweeping through the function values from lowest to highest. This is similar to the intuition employed in Morse theory[4, Section 1.4]. Visualize a smooth, compact 2D surface in 3D space. We want to decompose this surface into slope regions. Initially, our decomposition is empty, i.e. there are no slope regions (thus we don't have an actual *decomposition* yet). This is shown in Figure 4, Image 1.

Imagine a water level rising from below the surface, up to the point of first contact. Starting at this global minimum, we add a new region, containing only the argmin (i.e. a single point on the 2D image where the minimal value is taken).

Now, there might be many points where the global minimum is taken. This will either be due to a connected region (*plateau*) on the surface, which we want to include into the single existing region, or it will be due to individual dales, which all have their lowest point at the same height. In this case, we can't put the points into the existing region, because we would not be able to get from one argmin to another via a monotonic path. Instead, we need to add a new region for each individual dale.

Both cases can be dealt with by contracting connected points into their connected component, and creating a new region for each resulting component. This will ensure that plateaus are assigned to a single region.

As the water rises, we can add points to an existing region growing it upwards if they are just outside the region[1]. Otherwise they correspond to a distant local minimum and have to be dealt with as before, by opening a new region for each point (or rather, each connected component). This is shown in Figure 4, Image 2.

With the water rising further still, the regions will grow upwards to a point where they meet (Figure 4, Image 3). Any such point is a saddle point, and we have to account for it the next time we want to grow any one of the touching regions. The saddle point connects the edges of the regions which meet in it, at the current height of the water level. It might be the case that the not-yet-assigned points (the ones above the water) get separated into multiple connected components, or they might remain connected.

If the points remain connected, then we decide for a single one of the involved regions to be allowed to grow upwards from the component. This means that one region effectively inherits the growth directions from the other region(s). The other region(s) lose their potential for expansion and remain frozen in their current state.

If the unassigned points have multiple components (as in Figure 4, Image 3: the unassigned grey points are separated into the lower left and upper right areas), then we may assign one component to each involved region. The regions will then grow only in the directions determined by the assigned components as the water rises. This can be observed in Figure 4, Image 4: The green region is allowed to grow to the lower left, while the red region floods the upper right. The same procedure of swapping areas of expansion also happens as we move from Image 5 to 6. Any region without an assigned component remains frozen.

[1]Why can we do that? By adding only points which are connected to the region we ensure path-connectedness, and by growing the region upwards, we can construct ascending paths from old points to new ones. The smoothness of the surface guarantees that while moving at a fixed height, we can reach all points of the region with that height.

Should there be more components than regions, then we open up a new region for each surplus component, as in the top of Image 7.

An oddity which can occur are self-loops: A region might grow into a ”C”-shape, and then proceed to close up into an ”O”-shape. This case can be treated similarly as above, the only difference is that the saddle is found by recognizing that the region collides with itself, not with another region.

Eventually this procedure will arrive at the global maximum, and the entire surface will be divided into regions. Since we proceeded with the necessary care and attention along the way, we ensured that the regions remained slope regions, and we also only created additional regions when we absolutely had to, showing that the resulting composition is maximally coarse.

The same algorithm can be applied in higher dimensions. We deal with iso-hyper-surfaces as level sets, but the topological considerations about connectedness remain the same as in the illustrative 2D case.

5. Discrete BP

The somewhat vague description of BP in the previous section assumed a continuous surface. In most applications, however, the data will be provided in a discrete raster format. Some intricacies arise from this discretization, most notably iso-surfaces of a smooth function f will not have a straight-forward representation in the discrete grid obtained from rasterizing f. The data structure we use is a set of indices, representing the positions in the discrete array already assigned to a region.

Each region also has a set of (yet) unassigned points, which determine where the region might grow in the next iteration, called the *border*. This effectively models the smooth levelsets in the discrete representation.

The pseudo-code for the algorithm is printed below, the executable python code can be accessed in our github repository: *https://github.com/SirFloIII/MustererkennungLVA*

6. Further Potential Development

The result of BP is satisfactory, but we assume that improvements can be made in running time. The code was profiled multiple times and has been adapted to run faster with significant gains in many instances.

Algorithm 1 Border Propagation

Enumerate all values of f and collect points into levelsets.
for each levelset in bottom to top order **do**
  Add points to regions if they are in the border of a region
  if an added point is in the border of different region **then**
    Find the union of the borders of the involved regions
    Find the connected components thereof
    Assign these to the regions in an arbitrary way
  end if
  if an added point splits the border of the region in two **then**
    Reduce the border the region to one component
    for each other component **do**
      Create a new region containing the component as border
    end for
  end if
  for leftover points that cannot be added to any regions **do**
    Create a new region containing only that point
  end for
end for

Additional features we consider:

- Providing a *tolerance* parameter, which governs how steep a continuous function might get, before an iso-surface is deemed disconnected in the discrete data. This would allow for a trade-off between continuous connectedness and discrete connectedness. Modeling continuous connectedness creates fewer slope regions and yields pleasing results on smooth data, but the resulting regions are not monotonically connected (in the discrete sense of *connected*) in general. Discrete connectedness guarantees monotonic connectedness, but it necessarily creates significantly more and smaller slope regions. On smooth data the latter tends to produce too fine of a decomposition.
- Using established data structures that model smooth level sets from discrete data. There might be performance gains in employing such a data structure.

7. Conclusion

In this paper we have shown that slope regions of continuous functions in high dimensions ($n \geq 3$) do not have the same critical point properties well-established in 2D. Hence previous graph-based methods of building slope region decompositions by merging regions according to their border extrema will fail in high dimensions. Instead we developed a new, levelset-based method of growing regions, which yields slope region decompositions on discrete data of arbitrary dimension.

Acknowledgements

The BP algorithm as well as this paper are the result of a pattern recognition course held at the *Vienna University Of Technology* from Oct 2019 till Jan 2020. Professor W. KROPATSCH introduced us to the concept of slope regions and posed the challenge to compute them in high dimensions (> 2). We wish to thank him as well as DARSHAN BATAVIA and the anonymous reviewers for their valuable input.

References

[1] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci. Morse-smale complexes for piecewise linear 3-manifolds. In *Proceedings of the nineteenth annual symposium on Computational geometry*, pages 361–370, 2003.

[2] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[3] W. G. Kropatsch, R. M. Casablanca, D. Batavia, and R. Gonzalez-Diaz. Computing and reducing slope complexes. In *International Workshop on Computational Topology in Image Context*, pages 12–25. Springer, 2019.

[4] Y. Matsumoto. *An introduction to Morse theory*. Iwanami series in modern mathematics. American Math. Soc., Providence, RI, 2002.

[5] F. Riesz. Sur un Théoreme de Maximum de Mm. Hardy et Littlewood. *Journal of the London Mathematical Society*, 1(1):10–13, 1932.

[6] M. Zorn. A remark on method in transfinite algebra. *Bulletin of the American Mathematical Society*, 41(10):667–670, 1935.