

Greedy Monochromatic Island Partitions

Steven van den Broek¹, Wouter Meulemans^{*1}, and Bettina Speckmann¹

¹ TU Eindhoven, the Netherlands

[s.w.v.d.broek|w.meulemans|b.speckmann]@tue.nl

Abstract

Constructing partitions of colored points is a well-studied problem in discrete and computational geometry. We study the problem of creating a minimum-cardinality partition into monochromatic islands. Our input is a set S of n points in the plane where each point has one of $k \geq 2$ colors. A set of points is monochromatic if it contains points of only one color. An island I is a subset of S such that $\mathcal{CH}(I) \cap S = I$, where $\mathcal{CH}(I)$ denotes the convex hull of I . We identify an island with its convex hull; therefore, a partition into islands has the additional requirement that the convex hulls of the islands are pairwise-disjoint. We present three greedy algorithms for constructing island partitions and analyze their approximation ratios.

Related Version arXiv:2402.13340

1 Introduction

Constructing partitions of colored points is a well-studied problem in discrete [8, 12] and computational geometry [1, 4, 5, 16]. The colors of the points can be present in the constraints and the optimization criterion in different ways. For example, one may require the partition to be *balanced*—see the survey by Kano and Urrutia [12] for many such instances—or *monochromatic* [1, 4, 5, 8]. Alternatively, one may want to minimize or maximize the *diversity* [16] or *discrepancy* [3, 7] of the partition. Furthermore, one can use different geometries to partition the points, such as triangles [1], disks [5], or lines [4].

We study the problem of creating a minimum-cardinality partition into *monochromatic islands* [2]. Our input is a set S of n points in the plane where each point has one of $k \geq 2$ colors. A set of points is *monochromatic* if it contains points of only one color. An *island* I is a subset of S such that $\mathcal{CH}(I) \cap S = I$, where $\mathcal{CH}(I)$ denotes the convex hull of I . We identify an island with its convex hull; therefore, a partition into islands has the additional requirement that the convex hulls of the islands are pairwise-disjoint.

Related work. Bautista-Santiago et al. [2] study islands and describe an algorithm that can find a monochromatic island of maximum cardinality in $O(n^3)$ time, improving upon an earlier $O(n^3 \log n)$ algorithm [10]. Dumitrescu and Pach [8] consider monochromatic island partitions and prove how many islands are sufficient and sometimes necessary for different types of input. Bereg et al. [3] use island partitions to define a notion of *coarseness* that captures how blended a set of red and blue points are. Agarwal and Suri [1] study the following problem: given red and blue points, cover the blue points with the minimum number of pairwise-disjoint monochromatic triangles. They prove that this problem is NP-hard and describe approximation algorithms. Their NP-hardness reduction can be used to prove that covering and partitioning points of only one color into the minimum number of

* W. Meulemans is partially supported by the Dutch Research Council (NWO) under project number VI.Vidi.223.137.

3:2 Greedy Monochromatic Island Partitions





Figure 1 Left: optimal island partition; middle-left: disjoint-greedy island partition; middle-right: overlap-greedy island cover; right: line-greedy separating lines.

monochromatic islands—the points of the other colors serving only as obstacles—is NP-hard, as observed by Bautista-Santiago et al. [2]. We suspect that the problem we study is NP-hard as well, which motivates us to focus on approximation algorithms.


Overview. In the remainder, we consider only monochromatic islands. We denote by Opt_P the minimum cardinality of an island partition of S . In the following sections, we use three greedy algorithms—*disjoint-greedy*, *overlap-greedy*, and *line-greedy*—to construct island partitions. Disjoint-greedy creates an island partition by iteratively picking the island that covers most uncovered points and does not intersect any island chosen before. We sketch a proof that shows that disjoint-greedy has an approximation ratio of $\Omega(n/\log^2 n)$. The overlap-greedy algorithm greedily constructs an $O(\log n)$ -approximation of the minimum-cardinality island cover. We prove that any algorithm that transforms an island cover returned by overlap-greedy into an island partition has approximation ratio $\Omega(\sqrt{n})$, and describe one such algorithm that has approximation ratio $O(\text{Opt}_P^2 \log^2 n)$. Lastly, we investigate the relation between constructing a minimum-cardinality island partition and finding the minimum number of lines that separate the points into monochromatic regions. In particular, we show that greedily choosing the line that separates most pairs of points of different color induces an $O(\text{Opt}_P \log^2 n)$ -approximation to the minimum-cardinality island partition. Figure 1 illustrates the greedy algorithms. The full paper contains all technical details.

2 Disjoint-Greedy


We sketch our lower bound construction that shows that disjoint-greedy has an approximation ratio of $\Omega(n/\log^2 n)$. Consider a family of problem instances that have the form of two opposing complete binary trees of height ℓ (Figure 2). Sets of points are placed close together at the nodes of these trees. The idea is that by placing sufficiently many points at the nodes, and by placing obstacle points appropriately, the disjoint-greedy algorithm iteratively picks points of two opposing nodes such that the problem instance is split into two symmetric nearly independent parts that have nearly the same structure as the original instance. This results in disjoint-greedy returning a partition into $\Omega(2^\ell)$ islands (Figure 3). However, there exists a partition such that each layer in the tree consists of a constant number of islands, resulting in $O(\ell)$ islands in total (Figure 4). In our construction, the number of red points at a node at height $i \in \{0, \dots, \ell - 1\}$ is 2^{i+6} and the number of blue points at a node is constant. Hence, the problem instance contains $\Theta(\ell \cdot 2^\ell)$ points in total and the approximation ratio of disjoint-greedy is $\Omega(2^\ell/\ell) = \Omega(n/\log^2 n)$.



■ **Figure 2** The problem instance for $\ell = 5$. The lines in the figure are not part of the problem instance, but illustrate its structure. The purple squares represent red and blue points lying close together inside a square. The red disk inside the square represents many red points placed together inside a disk. The centers of the purple squares lie within the strip bounded by the two dashed lines.



■ **Figure 3** The solution returned by disjoint-greedy.



■ **Figure 4** An alternative solution, serving as an upper bound for the optimal solution.


3 Overlap-Greedy

A greedy algorithm that iteratively picks the island that covers most uncovered points results in an $O(\log n)$ approximation to the minimum-cardinality island cover. This follows immediately from viewing the problem as a set cover problem, where islands form the sets. We refer to this greedy algorithm as *overlap-greedy*. Below, we explore how to transform the island cover returned by overlap-greedy into an island partition. We assume that the greedy algorithm breaks ties by choosing an island that covers the fewest previously covered points.

We first define the relation between island covers and partitions based on them. Intuitively, islands that intersect can be transformed into a set of pairwise-disjoint islands by splitting them. In the transformation, each island has a corresponding family of islands into which it is split. The union of such a family should be a subset of the original island—a subset, not equal, because a point that was originally covered by multiple islands should be part of exactly one island after the transformation. This motivates the following definition.

- **Definition 1 (Compatible).** Families $\mathcal{I}'_1, \dots, \mathcal{I}'_m$ are *compatible* with islands I_1, \dots, I_m if:
- Families $\mathcal{I}'_1, \dots, \mathcal{I}'_m$ cover the same points as I_1, \dots, I_m : $\bigcup_k \bigcup \mathcal{I}'_k = \bigcup_i I_i$;
 - For every i , we have $\bigcup \mathcal{I}'_i \subseteq I_i$;
 - Islands $\bigcup_k \mathcal{I}'_k$ are pairwise-disjoint.
- Islands $I'_1, \dots, I'_{m'}$ are compatible with islands I_1, \dots, I_m if there exists a partition of $\{I'_1, \dots, I'_{m'}\}$ into families that are compatible with I_1, \dots, I_m .

3:4 Greedy Monochromatic Island Partitions



■ **Figure 5** A lower bound on the cardinality of solutions compatible with islands returned by overlap-greedy. Left: overlap-greedy cover; middle: a partition compatible with the overlap-greedy cover; right: an optimal island partition.

Thus, we arrive at the following problem: given m islands I_1, \dots, I_m obtained by overlap-greedy, find compatible families $\mathcal{I}'_1, \dots, \mathcal{I}'_m$ with $|\bigcup_k \mathcal{I}'_k|$ minimum. Solving this problem optimally is non-trivial. A natural approach to tackle the problem is to create an arrangement of the islands I_1, \dots, I_m and extract compatible families from that arrangement. However, two minor issues arise: the faces in the arrangement may not be convex, and the quality of the solution is not immediately clear as the number of faces in the arrangement may be arbitrarily greater than m . To resolve these issues, we build an arrangement iteratively. Before describing this process, we give a lower bound on the approximation ratio of algorithms that return islands compatible with those returned by overlap-greedy.


► **Lemma 2.** *Any algorithm that returns islands compatible with those returned by overlap-greedy has approximation ratio $\Omega(\max\{\text{Opt}_P, \sqrt{n}\})$.*

Proof. Let $k \in \mathbb{N}_{\geq 1}$. Consider a problem instance that consists of k evenly spaced vertical blue lines each formed by $k + 1$ evenly spaced blue points, and symmetrically k evenly spaced horizontal red lines (Figure 5). The cover returned by overlap-greedy has cardinality $2k$ and is induced by exactly those lines that were just described (Figure 5, left). Any partition that is compatible with the overlap-greedy cover has cardinality at least $2k + k^2$ (Figure 5, middle). Indeed, each intersection between two lines in the overlap-greedy cover forces an additional island in a compatible partition. An optimal island partition has cardinality $\text{Opt}_P = 2k + 1$ and consists of either all horizontal or all vertical islands (Figure 5, right). The number of points $n = O(k^2)$. Thus, the approximation ratio of an algorithm that produces solutions compatible with that of overlap-greedy is at least $\frac{2k+k^2}{2k+1} = \Omega(k) = \Omega(\text{Opt}_P) = \Omega(\sqrt{n}) = \Omega(\max\{\text{Opt}_P, \sqrt{n}\})$. ◀

3.1 Upper Bound

As mentioned earlier, our algorithm for creating a compatible island partition from an island cover works in an iterative manner. Throughout the iterations, we keep track of a restricted planar subdivision, which we call an *island arrangement*, to bound the cardinality of the island partition constructed by the algorithm. To simplify the arguments, we assume all islands in the island cover have cardinality at least three and that no three points are collinear. Then, a compatible island partition can be created from the faces of the island arrangement. See Figure 6 for an overview of the transformation from island cover to island partition.

We now define the notion of an *island arrangement*. In the following, vertices, edges, and faces of a planar subdivision are collectively referred to as *features*.



■ **Figure 6** Left: an island cover I_1, \dots, I_m returned by overlap-greedy; middle: an island arrangement of I_1, \dots, I_m ; right: an island partition compatible with I_1, \dots, I_m induced by the arrangement.

► **Definition 3** (Island arrangement). An island arrangement of islands I_1, \dots, I_i is a planar subdivision with the following additional requirements:

- Bounded faces are convex;
- Every bounded feature is a subset of $\mathcal{CH}(I_j)$ for some $1 \leq j \leq i$;
- For every $1 \leq j \leq i$, $\mathcal{CH}(I_j)$ is covered by bounded features.

Let I_1, \dots, I_m be the islands chosen by overlap-greedy for some set of points S . Let $U_i = I_i \setminus \bigcup_{j < i} I_j$ be the set of uncovered points island I_i covers. Because islands are chosen greedily by overlap-greedy, they satisfy $|U_i| \geq |U_{i+1}|$ for $i \in \{1, \dots, m-1\}$ and for all i island I_i is such that $|U_i|$ is maximum. By using these properties, the following lemma can be proven to hold for islands I_1, \dots, I_m .

► **Lemma 4.** *Let δ denote the boundary operator on sets. For distinct $1 \leq i < j \leq m$, the number of intersections between $\delta(\mathcal{CH}(I_i))$ and $\delta(\mathcal{CH}(I_j))$ is at most 2Opt_P .*

Using this lemma we can prove that an island arrangement of I_1, \dots, I_{i-1} with $1 \leq i \leq m$ can be modified into an island arrangement of I_1, \dots, I_i such that the increase in the number of faces is bounded in terms of i and Opt_P . We call this modification an *augmentation* of the arrangement. The following lemma makes one face for the new island I_i and modifies any existing features to make room. We refer to this as a *bold augmentation* of the arrangement.

► **Lemma 5** (Bold augmentation). *Given an island arrangement A of I_1, \dots, I_{i-1} with f faces, there exists an island arrangement A' of I_1, \dots, I_i with at most $f + 2\text{Opt}_P \cdot (i-1)$ faces such that there is exactly one face whose closure equals $\mathcal{CH}(I_i)$.*

By repeatedly applying Lemma 5, an island cover returned by overlap-greedy can be transformed into a compatible island partition. Let *bold overlap-greedy* be the algorithm that first runs overlap-greedy to obtain islands I_1, \dots, I_m , then repeatedly applies the bold augmentation step to create an island arrangement A of I_1, \dots, I_m , and finally extracts an island partition from the faces of A . The following bound holds on its approximation ratio.


► **Corollary 6.** *Bold overlap-greedy has approximation ratio $O(\text{Opt}_P^2 \log^2 n)$.*

4 Line-Greedy

In this section, we explore the relation between our problem and that of separating colors with the minimum number of lines. In particular, we show that greedily chosen separating lines induce an $O(\text{Opt}_P \log^2 n)$ -approximation to the minimum-cardinality island partition.

A set of lines L *separates* a set of colored points S if each face in the arrangement $\mathcal{A}(L)$ is monochromatic. The problem of finding the minimum-cardinality set of such separating lines is W[1]-hard with the parameter being the solution size [4]. Furthermore, the problem

3:6 Greedy Monochromatic Island Partitions



■ **Figure 7** Left: islands; right: expanded islands and their contact graph.

is NP-hard [14] and APX-hard [6], even when allowing only axis-parallel lines. The problem can be viewed as a set cover problem where lines are used to cover line segments between pairs of points of different color. Thus, the corresponding greedy algorithm, which we refer to as *line-greedy*, yields a $O(\log n)$ -approximation [11, 13]. Line-greedy can be implemented to run in $O(k \text{Opt}_L n^2 \log n)$ time [11], where Opt_L is the optimal number of lines and k is the number of colors of the input set.

If L separates S , then the faces of the arrangement $\mathcal{A}(L)$ induce a partition of S into $O(|L|^2)$ islands. Conversely, an island partition \mathcal{P} of S , with $|\mathcal{P}| \geq 3$, induces a set of $O(|\mathcal{P}|)$ lines that separates S . This can be shown using a construction by Edelsbrunner, Robison, and Shen [9]. We sketch their construction, adapted slightly for our use; see their paper for details and proofs. Circumscribe a rectangle around all the polygons—the convex hulls of the islands in \mathcal{P} . Grow the polygons, by moving their sides, until they are maximal. Extend each shared polygon side to obtain a set of lines. This set of lines separates the input points. Furthermore, each line corresponds to an edge of the contact graph of the expanded polygons (Figure 7). Because the contact graph is planar, there are at most $3|\mathcal{P}| - 6$ lines, yielding the desired result. While the exact running time of the construction is unclear, it is clearly polynomial. Pocchiola and Vegter [15] provide an alternative construction that makes use of a pseudo-triangulation of the polygons. Their algorithm runs in $O(n + |\mathcal{P}| \log n)$ time.

Thus, an optimal island partition induces an $O(\text{Opt}_L)$ -approximation to the optimal set of separating lines. Conversely, an optimal set of separating lines induces an $O(\text{Opt}_P)$ -approximation to the optimal island partition. There is an analogous relation between approximation algorithms of the two problems. In particular, we have the following result.

► **Lemma 7.** *Line-greedy induces an $O(\text{Opt}_P \log^2 n)$ -approximation to the minimum-cardinality island partition.*

For a lower bound instance, place points in a square grid of $k = 2^\ell$ rows and columns and color them alternately as in a checkerboard. In addition, place points on the corners of thin axis-parallel rectangles on the sides of the grid to encourage the line-greedy algorithm to use axis-parallel lines (Figure 8). We suspect that for any $\ell \geq 1$ line-greedy returns horizontal and vertical separating lines that separate the rows and columns of the grid as shown on the left in Figure 8. However, a formal proof eludes us. If this were true, then the island partition induced by the line-greedy solution would have cardinality $\Omega(k^2)$. Because an island partition of cardinality $O(k)$ exists (Figure 8, right), this would result in an $\Omega(\sqrt{n}) = \Omega(\text{Opt}_P)$ lower bound on the approximation that is attained by an island partition induced by line-greedy.



■ **Figure 8** The figure shows an idea of a lower bound on the approximation that is attained by an island partition induced by line-greedy.

References

- 1 Pankaj K. Agarwal and Subhash Suri. Surface approximation and geometric partitions. *SIAM Journal on Computing*, 27(4):1016–1035, 1998. doi:10.1137/S0097539794269801.
- 2 Crevel Bautista-Santiago, José Miguel Díaz-Báñez, Dolores Lara, Pablo Pérez-Lantero, Jorge Urrutia, and Inmaculada Ventura. Computing optimal islands. *Operations Research Letters*, 39(4):246–251, 2011. doi:10.1016/j.orl.2011.04.008.
- 3 Sergey Bereg, José Miguel Díaz-Báñez, Dolores Lara, Pablo Pérez-Lantero, Carlos Seara, and Jorge Urrutia. On the coarseness of bicolored point sets. *Computational Geometry*, 46(1):65–77, 2013. doi:10.1016/J.COMGEO.2012.04.003.
- 4 Edouard Bonnet, Panos Giannopoulos, and Michael Lampis. On the parameterized complexity of red-blue points separation. *Journal of Computational Geometry*, 10(1):181–206, 2019. doi:10.20382/jocg.v10i1a7.
- 5 Sergio Cabello, José Miguel Díaz-Báñez, and Pablo Pérez-Lantero. Covering a bichromatic point set with two disjoint monochromatic disks. *Computational Geometry*, 46(3):203–212, 2013. doi:10.1016/j.comgeo.2012.06.002.
- 6 Gruia Călinescu, Adrian Dumitrescu, Howard J. Karloff, and Peng-Jun Wan. Separating points by axis-parallel lines. *International Journal of Computational Geometry & Applications*, 15(6):575–590, 2005. doi:10.1142/S0218195905001865.
- 7 José Miguel Díaz-Báñez, Ruy Fabila Monroy, Pablo Pérez-Lantero, and Inmaculada Ventura. New results on the coarseness of bicolored point sets. *Information Processing Letters*, 123:1–7, 2017. doi:10.1016/J.IPL.2017.02.007.
- 8 Adrian Dumitrescu and János Pach. Partitioning colored point sets into monochromatic parts. *International Journal of Computational Geometry & Applications*, 12(5):401–412, 2002. doi:10.1142/S0218195902000943.
- 9 Herbert Edelsbrunner, Arch D. Robison, and Xiaojun Shen. Covering convex sets with non-overlapping polygons. *Discrete Mathematics*, 81(2):153–164, 1990. doi:10.1016/0012-365X(90)90147-A.
- 10 Paul Fischer. Sequential and parallel algorithms for finding a maximum convex polygon. *Computational Geometry*, 7:187–200, 1997. doi:10.1016/0925-7721(95)00035-6.
- 11 Hossein Jowhari and Mohsen Rezapour. Monochromatic partitioning of colored points by lines. *Information Processing Letters*, 182:106402, 2023. doi:10.1016/j.ipl.2023.106402.
- 12 Mikio Kano and Jorge Urrutia. Discrete geometry on colored point sets in the plane - A survey. *Graphs and Combinatorics*, 37(1):1–53, 2021. doi:10.1007/S00373-020-02210-8.
- 13 Bing Lu, Hongwei Du, Xiaohua Jia, Yinfeng Xu, and Binhai Zhu. On a minimum linear classification problem. *Journal of Global Optimization*, 35(1):103–109, 2006.
- 14 Nimrod Megiddo. On the complexity of polyhedral separability. *Discrete & Computational Geometry*, 3:325–337, 1988. doi:10.1007/BF02187916.
- 15 Michel Pocchiola and Gert Vegter. On polygonal covers. *Contemporary Mathematics*, 223:257–268, 1999.

3:8 Greedy Monochromatic Island Partitions

- 16 Marc van Kreveld, Bettina Speckmann, and Jérôme Urhausen. Diverse partitions of colored points. In *17th Algorithms and Data Structures Symposium (WADS)*, volume 12808 of *Lecture Notes in Computer Science*, pages 641–654, 2021. doi : 10.1007/978-3-030-83508-8_46.