

Preprocessing Ambiguous Imprecise Points

Ivor van der Hoog

Department of Information and Computing Sciences, Utrecht University, The Netherlands
i.d.vanderhoog@uu.nl

Irina Kostitsyna

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
i.kostitsyna@tue.nl

Maarten Löffler

Department of Information and Computing Sciences, Utrecht University, The Netherlands
m.loffler@uu.nl

Bettina Speckmann

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
b.speckmann@tue.nl

Abstract

Let $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ be a set of regions and let $X = \{x_1, x_2, \dots, x_n\}$ be an (unknown) point set with $x_i \in R_i$. Region R_i represents the uncertainty region of x_i . We consider the following question: how fast can we establish order if we are allowed to preprocess the regions in \mathcal{R} ? The *preprocessing model* of uncertainty uses two consecutive phases: a preprocessing phase which has access only to \mathcal{R} followed by a reconstruction phase during which a desired structure on X is computed. Recent results in this model parametrize the reconstruction time by the *ply* of \mathcal{R} , which is the maximum overlap between the regions in \mathcal{R} . We introduce the *ambiguity* $\mathcal{A}(\mathcal{R})$ as a more fine-grained measure of the degree of overlap in \mathcal{R} . We show how to preprocess a set of d -dimensional disks in $\mathcal{O}(n \log n)$ time such that we can sort X (if $d = 1$) and reconstruct a quadtree on X (if $d \geq 1$ but constant) in $\mathcal{O}(\mathcal{A}(\mathcal{R}))$ time. If $\mathcal{A}(\mathcal{R})$ is sub-linear, then reporting the result dominates the running time of the reconstruction phase. However, we can still return a suitable data structure representing the result in $\mathcal{O}(\mathcal{A}(\mathcal{R}))$ time.

In one dimension, \mathcal{R} is a set of intervals and the ambiguity is linked to interval entropy, which in turn relates to the well-studied problem of sorting under partial information. The number of comparisons necessary to find the linear order underlying a poset P is lower-bounded by the graph entropy of P . We show that if P is an interval order, then the ambiguity provides a constant-factor approximation of the graph entropy. This gives a lower bound of $\Omega(\mathcal{A}(\mathcal{R}))$ in all dimensions for the reconstruction phase (sorting or any proximity structure), independent of any preprocessing; hence our result is tight. Finally, our results imply that one can approximate the entropy of interval graphs in $\mathcal{O}(n \log n)$ time, improving the $\mathcal{O}(n^{2.5})$ bound by Cardinal et al.

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases preprocessing, imprecise points, entropy, sorting, proximity structures

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.42

Related Version Find the full version on <https://arxiv.org/abs/1903.08280>.

Funding *Ivor van der Hoog*: Supported by the Netherlands Organisation for Scientific Research (NWO); 614.001.504.

Maarten Löffler: Partially supported by the Netherlands Organisation for Scientific Research (NWO); 614.001.504.

Bettina Speckmann: Partially supported by the Netherlands Organisation for Scientific Research (NWO); 639.023.208.

Acknowledgements The authors would like to thank Jean Cardinal for insightful discussions.



© Ivor van der Hoog, Irina Kostitsyna, Maarten Löffler, and Bettina Speckmann;
licensed under Creative Commons License CC-BY


35th International Symposium on Computational Geometry (SoCG 2019).

Editors: Gill Barequet and Yusu Wang; Article No. 42; pp. 42:1–42:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany




1 Introduction

A fundamental assumption in classic algorithms research is that the input data given to an algorithm is exact. Clearly this assumption is generally not justified in practice: real-world data tends to have (measurement or labeling) errors, heterogeneous data sources introduce yet other type of errors, and “big data” is compounding the effects. To increase the relevance of algorithmic techniques for practical applications, various paradigms for dealing with uncertain data have been introduced over the past decades. Many of these approaches have in common that they represent the uncertainty, imprecision, or error of a data point as a *disk* in a suitable distance metric which we call an uncertainty region. We focus on a fundamental problem from the realm of computation with uncertainties and errors: given a set of imprecise points represented by uncertainty regions, how much proximity information do the regions contain about the imprecise points?

Preprocessing model. We study this problem within the preprocessing framework initially proposed by Held and Mitchell [11]. In this framework we have a set $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ of regions and an point set $X = \{x_1, x_2, \dots, x_n\}$ with $x_i \in R_i$. This model has 2 consecutive phases: a preprocessing phase followed by a reconstruction phase. In the preprocessing phase we have access only to \mathcal{R} and we typically want to preprocess \mathcal{R} in $\mathcal{O}(n \log n)$ time to create some linear-size auxiliary data structure which we will denote by Ξ . In the reconstruction phase, we have access to X and we want to construct a desired output on X using Ξ faster than would be possible otherwise. Löffler and Snoeyink [18] were the first to use this model as a way to deal with data uncertainty: one may interpret the regions \mathcal{R} as *imprecise* points, and the points in X as their true (initially unknown) locations. This interpretation of the preprocessing framework has been successfully applied to various problems in computational geometry [2, 3, 6, 7, 16, 23]. Several results restrict \mathcal{R} to be a set of disjoint (unit) disks in the plane, while others consider partially overlapping disks. Traditionally, the *ply* $\Delta(\mathcal{R})$ of \mathcal{R} , which measures the maximal number of overlapping regions, has been used to measure the degree of overlap, leading, for example, to reconstruction times of $\mathcal{O}(n \log \Delta(\mathcal{R}))$.

The ply is arguably a somewhat coarse measure of the degree of overlap of the regions. Consider the following example: suppose that we have a collection of \sqrt{n} disks in the plane that overlap in one point and that the remainder of \mathcal{R} is mutually disjoint (see Figure 1 left). Then $\Delta(\mathcal{R}) = \sqrt{n}$ and the resulting time complexity of the reconstruction phase is $\mathcal{O}(n \log n)$ even though it might be possible to achieve better bounds (\mathcal{R} is arguably not in a worst-case configuration for that given ply, see Figure 1 right).

Ambiguity. We introduce the *ambiguity* $A(\mathcal{R})$ as a more fine-grained measure of the degree of overlap in \mathcal{R} . The ambiguity is based on the number of regions each individual region intersects (see Figure 1). We count this number with respect to particular permutations of



■ **Figure 1** Two sets of 16 disks each in the plane, both with a ply of 4. The ambiguity of the set on the right is four times as large as the ambiguity of the set on the left.

the regions: for each region we count only the overlap with regions that appear earlier in the permutation. A proper technical definition of ambiguity can be found in Section 2. We also show how to compute a 3-approximation of the ambiguity in $\mathcal{O}(n \log n)$ time.

Ambiguity and entropy. In one dimension, \mathcal{R} is a set of intervals and the ambiguity is linked to interval (and graph) entropy (refer to the full version for a definition), which in turn relates to the well-studied problem of sorting under partial information. Fredman [9] shows that if the only information we are given about a set of values is a partial order P , and $e(P)$ is the number of *linear extensions* (total orders compatible with) of P , then we need at least $\Omega(\log e(P))$ comparisons to sort the values. Brightwell and Winkler prove that computing the number of linear extensions $e(P)$ is $\#P$ -complete [1]. Hence efforts have concentrated on computing approximations, most notably via the concept of graph entropy as introduced by Körner [14]. Specifically, Khan and Kim [13] prove that $\log e(P) = \Theta(n \cdot H(G))$ where $H(G)$ denotes the entropy of the *incomparability graph* G of the poset P . To the best of our knowledge there is currently no exact algorithm to compute $H(G)$. Cardinal et al. [4] describe the fastest known algorithm to approximate $H(G)$, which runs in $\mathcal{O}(n^{2.5})$ time. Refer to the full version for a more in-depth discussion of sorting and its relation to graph entropy.


We consider the special case where the partial order is induced by uncertainty intervals. We define the entropy $H(\mathcal{R})$ of a set of intervals as the entropy of their intersection graph (which is also an incomparability graph) using the definition of graph entropy given by Körner. In this setting we prove that the ambiguity $A(\mathcal{R})$ provides a constant-factor approximation of the interval entropy (see Section 2). Since we can compute a constant-factor approximation of the ambiguity in $\mathcal{O}(n \log n)$ time, we can hence also compute a constant-factor approximation of the entropy of interval graphs in $\mathcal{O}(n \log n)$ time, thereby improving the result by Cardinal et al. [4] for this special case.

Ambiguity and reconstruction. Since $\Omega(\log e(P))$ is a lower bound for the number of comparisons needed to complete P into a total order, $\Omega(A(\mathcal{R}))$ is a lower bound for the reconstruction phase in the preprocessing model when \mathcal{R} is a set of intervals and the goal is to sort the unknown points in X . This lower bound extends to higher dimensions and to proximity structures in general, independent of any preprocessing.

The ambiguity $A(\mathcal{R})$ ranges between 0 and $\Theta(n \log n)$ for a set of n regions \mathcal{R} . If the value of $A(\mathcal{R})$ lies between $\Theta(n)$ and $\Theta(n \log n)$ then we can preprocess \mathcal{R} in $\mathcal{O}(n \log n)$ time and sort in $\mathcal{O}(A(\mathcal{R}))$ time (in one dimension for arbitrary intervals) or build a quadtree in $\mathcal{O}(A(\mathcal{R}))$ time (in all dimensions for unit disks).

If the ambiguity lies between 0 and $\Theta(n)$, then reporting the results explicitly in $\Omega(n)$ time dominates the reconstruction time. But the ambiguity suggests that the information-theoretic amount of work necessary to compute the results should be lower than $\Theta(n)$. To capture this, we hence introduce a new variant of the preprocessing model, which allows us to return a pointer to an implicit representation of the results.

Specifically, in one dimension, \mathcal{R} is a set of intervals and we aim to return the sorted order of the unknown points in X . If, for example, all intervals are mutually disjoint, then $A(\mathcal{R}) = \mathcal{O}(1)$ and we have essentially no time for the reconstruction phase. However, a binary search tree T on \mathcal{R} , which we can construct in $\mathcal{O}(n \log n)$ time in the preprocessing phase, actually captures all necessary information. In the reconstruction phase we can hence return a pointer to T as an implicit representation of the sorted order. In Section 3 we show how to handle arbitrary sets of intervals in a similar manner. That is, we describe how to construct in $\mathcal{O}(n \log n)$ time an auxiliary data structure Ξ on \mathcal{R} in the preprocessing phase (without access to X), such that, in the reconstruction phase (using X), we can construct a linear-size AVL-tree T on X in $\mathcal{O}(A(\mathcal{R})) = \mathcal{O}(\log e(\mathcal{R}))$ time, which is tight.



■ **Figure 2** A set of overlapping intervals with a permutation. In this figure $\Gamma_4^\pi = \{R_1, R_2, R_3, R_4\}$. In all figures, bottom intervals are indicated in blue (in this case this is only R_1).

In all dimensions, we consider \mathcal{R} to be a set of unit disks and our aim is to return a quadtree T on the points in X where each point in X lies in a unique quadtree cell. Note that in 2 dimensions, T also allows us to construct e.g. the Delaunay triangulation of X in linear time [2]. However, we show that constructing such a quadtree explicitly in $\mathcal{O}(A(\mathcal{R}))$ time is not possible, and the work necessary to distinguish individual points could dominate the running time and overshadow the detail in the analysis brought by the ambiguity measure. We hence follow Buchin et al. [2] and use so-called λ -deflated *quadtrees* which contain up to a constant λ points in each leaf. From T one can construct a quadtree on X where each point lies in a unique quadtree cell in linear time. In Section 4 we describe how to reconstruct a linear-size λ -deflated quadtree T (with a suitable constant λ) in $\mathcal{O}(A(\mathcal{R})) = \mathcal{O}(\log e(\mathcal{R}))$ time, which is tight (in fact, in one dimension our result also extends to non-unit intervals).

2 Ambiguity

We introduce a new measure on a set of regions \mathcal{R} to reflect the degree of overlap, which we call the *ambiguity*. The sequence in which we process regions matters (refer to Section 2.1), thus we distinguish between the π -ambiguity defined on a given permutation of the regions in \mathcal{R} , and the minimum ambiguity defined over all possible permutations. We demonstrate several properties of the ambiguity, and discuss its relation to graph entropy when \mathcal{R} is a set of intervals in one dimension.

Processing permutation. Let \mathcal{R} be a set of n regions and let $\mathcal{R}^\pi = \langle R_1, R_2, \dots, R_n \rangle$ (note that for all i , the region R_i could be any region depending on the permutation π) be the sequence of elements in \mathcal{R} according to a given permutation π . Then we say that π is a *processing permutation* of \mathcal{R} . Furthermore, let $\mathcal{R}_{\leq i}^\pi := \{R_j \mid j \leq i\}$ be the prefix of \mathcal{R}^π , that is, the first i elements in the sequence \mathcal{R}^π . A permutation π is *containment-compatible* if $R_i \subset R_j$ implies $i < j$ for all i and j [8]. When π is clear from context, we denote \mathcal{R}^π by \mathcal{R} .

Contact set (for a permutation π). For a region $R_i \in \mathcal{R}^\pi$ we define its *contact set* Γ_i^π to be the set of regions which precede or are equal to R_i in the order π , and which intersect R_i : $\Gamma_i^\pi := \{R_j \in \mathcal{R}_{\leq i}^\pi \mid R_j \cap R_i \neq \emptyset\}$. Note that a region is always in its own contact set. A region R_i whose contact set Γ_i^π contains only R_i itself is called a *bottom region* (refer to Figure 2).

Ambiguity. For a set of regions \mathcal{R} and a fixed permutation π we define the π -ambiguity $A^\pi(\mathcal{R}) := \sum_i \log |\Gamma_i^\pi|$ (with the logarithm to the base 2). Observe that bottom regions do not contribute to the value of the π -ambiguity. The ambiguity of \mathcal{R} is now the minimal π -ambiguity over all permutations π , $A(\mathcal{R}) := \min_{\pi \in \Pi} A^\pi(\mathcal{R})$.



$$A^{\pi_1}(\mathcal{R}) = \log 1 + \log 1 + \log 1 + \log 1 + \log 5 \quad A^{\pi_2}(\mathcal{R}) = \log 2 + \log 2 + \log 2 + \log 2 + \log 1$$

■ **Figure 3** An example of the π -ambiguity induced by two permutations π_1 (on the left) and π_2 (on the right) of the same five intervals. The π_1 -ambiguity is $\log 5$ and the π_2 -ambiguity is 4.

2.1 Properties of ambiguity

We show the following properties of ambiguity: (1) the π -ambiguity may vary significantly with the choice of the processing permutation π , (2) in one dimension, the π -ambiguity for any containment-compatible permutation π on a set of intervals \mathcal{R} implies a 3-approximation on the entropy of the interval graph of \mathcal{R} , and (3) the permutation that realizes the ambiguity is containment-compatible. Therefore in one dimension, the ambiguity of a set of intervals \mathcal{R} implies a 3-approximation of the entropy of the interval graph of \mathcal{R} .

We start with the first property: it is easy to see that the processing permutation π has a significant influence on the value of the π -ambiguity (refer to Figure 3). Even though π -ambiguity can vary considerably, we show that if we restrict the permutations to be containment-compatible, their π -ambiguities lie within a constant factor of the ambiguity.

Interval entropy. The entropy of a graph G was first introduced by Körner [14]. Since then several equivalent definitions appeared [22]. We define the *interval entropy* $H(\mathcal{R})$, for a set of intervals \mathcal{R} , as the entropy of the intersection graph of \mathcal{R} . While investigating the question of sorting an arbitrary poset, Cardinal et al. [4] found an interesting geometrical interpretation of the poset entropy, which applies to our interval entropy: let a poset P describe a set of (open) intervals \mathcal{R} combinatorially, that is, for each R_i we know which intervals intersect R_i , are contained in R_i , contain R_i , and are disjoint from R_i . Denote by $E(\mathcal{R})$ the infinite set of sets of intervals on the domain $(0, 1)$ (that is, each $\mathcal{I} \in E(\mathcal{R})$ is a set of intervals, where each interval $I_i \in \mathcal{I}$ has endpoints in $(0, 1)$) which induce the same poset as \mathcal{R} . Then Cardinal et al. prove the following lemma (see Figure 4 for an illustration):

► **Lemma 1** ([4], Lemma 3.2 paraphrased).


$$H(\mathcal{R}) = \log n - \min_{\mathcal{I} \in E(\mathcal{R})} \left\{ \frac{1}{n} \sum_{I_i \in \mathcal{I}} -\log |I_i| \right\}.$$

We show that the π -ambiguity for any containment-compatible π is a 3-approximation of $n \cdot H(\mathcal{R})$. To achieve this we rewrite the lemma from Cardinal et al. in the following way,

$$H(\mathcal{R}) = \log n - \min_{\mathcal{I} \in E(\mathcal{R})} \left\{ \frac{1}{n} \sum_{I_i \in \mathcal{I}} (\log n - \log(n|I_i|)) \right\} = \max_{\mathcal{I} \in E(\mathcal{R})} \left\{ \frac{1}{n} \sum_{I_i \in \mathcal{I}} \log(n|I_i|) \right\}.$$

An embedding \mathcal{I} gives each interval I_i a size between 0 and 1. To simplify the algebra later, we re-interpret this size as the *fraction* (weight) of the domain $(0, 1)$ that I_i occupies. We associate with each $\mathcal{I} \in E(\mathcal{R})$ a set of weights W such that for all i , $w_i = |I_i|$; we write $W \sim E(\mathcal{R})$. From now on we consider embeddings on the domain $(0, n)$: an interval then has a size $n|I_i| = nw_i$. The formula for the entropy becomes:

$$H(\mathcal{R}) = \frac{1}{n} \max_{W \sim E(\mathcal{R})} \left\{ \log \left(\prod_{w_i \in W} n w_i \right) \right\}. \tag{1}$$



■ **Figure 4** Let \mathcal{R} be a set of five intervals, where four intervals are mutually disjoint and contained in one larger interval. We show three embeddings $\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3 \in E(\mathcal{R})$ of these intervals on the domain $(0, 1)$ with the same combinatorial properties. Embedding \mathcal{I}_1 shows that $H(\mathcal{R}) \geq \log 5 - \frac{1}{5} \log(1 \cdot \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{4} \cdot \frac{1}{4})$. \mathcal{I}_2 shows that $H(\mathcal{R}) \geq \log 5 - \frac{1}{5} \log(\frac{1}{2} \cdot \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{8})$ and \mathcal{I}_3 is the optimal embedding which shows that $H(\mathcal{R}) = \log 5 - \frac{1}{5} \log(1 \cdot \frac{1}{4} \cdot \frac{1}{4} \cdot \frac{1}{4} \cdot \frac{1}{4})$.

Ambiguity and entropy. Next, we show that the interval entropy gives an upper bound on the ambiguity. The entropy of \mathcal{R} is the maximum over all embeddings on $(0, n)$, so any embedding of \mathcal{R} on the domain $(0, n)$ gives a lower bound on $H(\mathcal{R})$. We will create an embedding with a corresponding weight assignment W such that:

$$A^\pi(\mathcal{R}) = \log \left(\prod_{R_i \in \mathcal{R}} |\Gamma_i^\pi| \right) \leq \log \left(\prod_{w_i \in W} (n w_i)^2 \right) \leq 2nH(\mathcal{R}). \quad (2)$$

We start with the original input embedding of \mathcal{R} and we sort the coordinates of all the endpoints (both left- and right-). To each endpoint p we assign a new coordinate $\frac{k}{2}$ if p is the k th endpoint in the sorted order (indexing from 0). Thus, we obtain an embedding of \mathcal{R} on $(0, n - \frac{1}{2})$. For any containment-compatible permutation π , the length of each interval R_i in this embedding is at least $\frac{1}{2}|\Gamma_i^\pi|$, as each interval R_i contains at least $|\Gamma_i^\pi| - 1$ endpoints of the intervals from its contact set in its interior. Also note that the distance between every right endpoint and the consecutive endpoint to the right is $\frac{1}{2}$. Thus, we can increase the coordinate of every right endpoint by $\frac{1}{2}$ and obtain an embedding of \mathcal{R} on $(0, n)$ with a corresponding weight assignment W , such that the length of each interval R_i is at least $\frac{1}{2}(|\Gamma_i^\pi| + 1)$. This allows us to prove the following lemma:

▶ **Lemma 2.** *For any containment-compatible permutation π of a set of intervals \mathcal{R} ,*

$$A^\pi(\mathcal{R}) \leq 2nH(\mathcal{R}).$$

Proof. Consider the embedding and corresponding weight assignment W constructed above. Consider any containment-compatible permutation π . We split the intervals of \mathcal{R} into four sets depending on the size of their contact set: let $A := \{R_i \mid |\Gamma_i^\pi| = 1\}$, $B := \{R_i \mid |\Gamma_i^\pi| = 2\}$, $C := \{R_i \mid |\Gamma_i^\pi| = 3\}$ and $D := \mathcal{R} \setminus \{A, B, C\}$. Let these sets contain a, b, c and d intervals respectively. Then, using Equation (1) for the entropy,

$$\begin{aligned} 2^{nH(\mathcal{R})} &\geq \prod_{R_i \in A} \frac{|\Gamma_i^\pi| + 1}{2} \prod_{R_i \in B} \frac{|\Gamma_i^\pi| + 1}{2} \prod_{R_i \in C} \frac{|\Gamma_i^\pi| + 1}{2} \prod_{R_i \in D} \frac{|\Gamma_i^\pi| + 1}{2} \\ &\geq \left(\frac{2}{2}\right)^a \left(\frac{3}{2}\right)^b \left(\frac{4}{2}\right)^c \left(\frac{4}{2}\right)^d. \end{aligned} \quad (3)$$

On the other hand,

$$\begin{aligned} 2^{nH(\mathcal{R})} &\geq \prod_{R_i \in \mathcal{R}} \frac{|\Gamma_i^\pi| + 1}{2} \geq \prod_{R_i \in A} |\Gamma_i^\pi| \prod_{R_i \in B} \frac{3}{4} |\Gamma_i^\pi| \prod_{R_i \in C} \frac{2}{3} |\Gamma_i^\pi| \prod_{R_i \in D} \frac{1}{2} |\Gamma_i^\pi| \\ &= \left(\frac{3}{4}\right)^b \left(\frac{2}{3}\right)^c \left(\frac{1}{2}\right)^d 2^{A^\pi(\mathcal{R})}, \end{aligned}$$

as

$$\frac{|\Gamma_i^\pi| + 1}{2} \begin{cases} 1 = |\Gamma_i^\pi|, & \text{if } R_i \in A, \\ \frac{3}{2} = \frac{3}{4}|\Gamma_i^\pi|, & \text{if } R_i \in B, \\ 2 = \frac{2}{3}|\Gamma_i^\pi|, & \text{if } R_i \in C, \end{cases} \quad \frac{|\Gamma_i^\pi| + 1}{2} \geq \frac{1}{2}|\Gamma_i^\pi|, \quad \text{if } R_i \in D.$$

Then, using Equation (3) we get

$$2^{nH(\mathcal{R})} \cdot 2^{nH(\mathcal{R})} \geq \left(\frac{3}{2}\right)^b \left(\frac{4}{2}\right)^c \left(\frac{4}{2}\right)^d \cdot \left(\frac{3}{4}\right)^b \left(\frac{2}{3}\right)^c \left(\frac{1}{2}\right)^d 2^{A^\pi(\mathcal{R})} \geq 2^{A^\pi(\mathcal{R})},$$

and therefore

$$2nH(\mathcal{R}) \geq A^\pi(\mathcal{R}). \quad \blacktriangleleft$$

We continue by showing that the ambiguity also gives an upper-bound for the interval entropy. Starting with a helper lemma:

► **Lemma 3.** *Suppose \mathcal{R} is partitioned into two sets X and Y such that for each $R \in X, R' \in Y$, R and R' are disjoint. In any weight assignment W that realizes $H(\mathcal{R})$, the intervals in X together have length $|X|$ and the intervals in Y together have length $|Y|$ on the domain $(0, n)$.*

Proof. In Equation (1) we rewrote the formula for entropy in terms of weights: for any weight assignment $W \sim E(\mathcal{R})$, w_i is the proportion that R_i occupies on the domain, and we embedded \mathcal{R} on the domain $(0, n)$. We can similarly embed \mathcal{R} on the domain $(0, \lambda)$ for an arbitrary scalar λ . We define the *relative entropy* of \mathcal{R} (refer to Figure 5 (top)) as:

$$H(\mathcal{R}, \lambda) := \frac{1}{n} \max_{W \sim E(\mathcal{R})} \left\{ \log \left(\prod_{w_i \in W} \lambda w_i \right) \right\}.$$

Observe that $H(\mathcal{R}, n) = H(\mathcal{R})$ and that:

$$\forall \lambda, \mu, \quad \mu w_i = \left(\frac{\mu}{\lambda}\right) \lambda w_i \Rightarrow 2^{nH(X, \mu)} = \left(\frac{\mu}{\lambda}\right)^{|X|} 2^{nH(X, \lambda)}. \quad (4)$$

If the intervals in X can occupy a width of at most λ , then it is always optimal to give the intervals in Y a total width of $n - \lambda$ (since the entropy maximizes the product of the lengths of intervals in X and Y). This implies:

$$2^{nH(X \cup Y)} = \max_{\lambda \in [0, n]} \{2^{nH(X, \lambda)} \cdot 2^{nH(Y, n-\lambda)}\}.$$

See Figure 5 (bottom) for an illustration of the argument. If we now substitute Equation (4) into this equation we get that the maximum is realized if $\lambda = |X|$ which proves the lemma. ◀

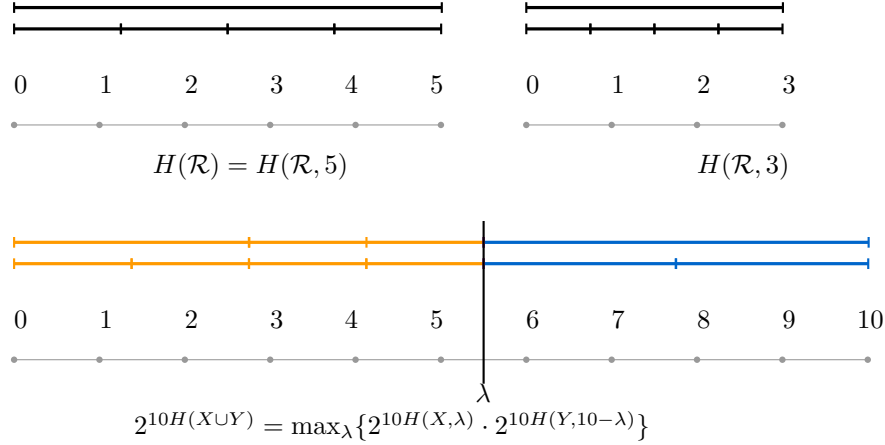
► **Lemma 4.** *Let π be any containment-compatible permutation, then $nH(\mathcal{R}) \leq 3A^\pi(\mathcal{R})$.*

Proof. We defined $\mathcal{R}_{\leq i}$ as the prefix of \mathcal{R} . We prove the lemma with induction on i .

$$\text{Induction Hypothesis: } \forall j \leq i \quad jH(\mathcal{R}_{\leq j}, j) \leq 3A^\pi(\mathcal{R}_{\leq j}).$$

For $i = 1$ both the lefthand and the righthand side are 0. So we assume that the lemma holds for all $j \leq i$ and we prove it for $j = i + 1$. $H(\mathcal{R}_{\leq i+1}, i + 1)$ is the relative entropy of $\mathcal{R}_{\leq(i+1)}$ on the domain $(0, i + 1)$. We know that $3A^\pi(\mathcal{R}_{\leq(i+1)}) = 3A^\pi(\mathcal{R}_{\leq i}) + 3 \log |\Gamma_{i+1}|$.

42:8 Preprocessing Ambiguous Imprecise Points



■ **Figure 5** (top left) A set \mathcal{R} of five intervals and their optimal embedding for the entropy relative to $\lambda = 5$. (top right) The optimal embedding of \mathcal{R} for the entropy relative to $\lambda = 3$. Observe that the proportion that each interval obtains of the domain is the same in both embeddings. (bottom) An illustration of the argument for Lemma 3: we see a set X of 7 intervals and a set Y of 3 intervals with the intervals in X disjoint from the intervals in Y . If we vary λ , we vary the total width on which X and Y are embedded. The entropy is given by the maximal embedding and therefore found by optimizing λ .

We make a distinction between two cases: $|\Gamma_{i+1}| = 1$ or otherwise. If $|\Gamma_{i+1}| = 1$ then R_{i+1} is disjoint from $R_{\leq i}$. Lemma 3 guarantees, that if we want to embed $R_{\leq i} \cup \{R_{i+1}\}$ on $(0, i + 1)$ that R_{i+1} gets a size of 1. The remaining intervals get embedded with a total width of i which they already had in the previous iteration. So:

$$(i + 1)H(\mathcal{R}_{\leq(i+1)}, i + 1) = iH(\mathcal{R}_i, i) + \log 1 \leq 3A^\pi(\mathcal{R}_{\leq i}) + 3 \log |\Gamma_{i+1}| = 3A^\pi(\mathcal{R}_{\leq(i+1)}).$$

In the second case $|\Gamma_{i+1}|$ is at least 2. The other intervals used to be optimally embedded on $(0, i)$ and are now embedded on $(0, i + 1)$. So each of them expands with at most a factor $\frac{i+1}{i}$ or algebraically:

$$\begin{aligned} nH(\mathcal{R}_{\leq(i+1)}, i + 1) &\leq nH(\mathcal{R}_{\leq i}, i) + \log \left(\left(\frac{i+1}{i} \right)^i \right) + \log((i + 1)w_{i+1}) \leq \\ &nH(\mathcal{R}_{\leq i}, i) + \log e + \log((i + 1)w_{i+1}). \end{aligned}$$

There are $i - |\Gamma_{i+1}|$ intervals disjoint from R_{i+1} so Lemma 3 guarantees that $(i + 1)w_{i+1} \geq |\Gamma_{i+1}| \geq 2$. It follows that:


$$nH(\mathcal{R}_{\leq(i+1)}, i + 1) \leq nH(\mathcal{R}_{\leq i}, i) + 3 \log |\Gamma_{i+1}|$$

which implies the Lemma. ◀

Lemmas 2 and 4 imply the following theorem.

► **Theorem 5.** For any set of intervals \mathcal{R} in one dimension, for any containment-compatible permutation π on \mathcal{R} , $A^\pi(\mathcal{R})$ is a 3-approximation of $nH(\mathcal{R})$.

► **Corollary 6.** For any set of intervals \mathcal{R} in one dimension, the ambiguity $A(\mathcal{R})$ is a 3-approximation of $nH(\mathcal{R})$.



■ **Figure 6** A set of intervals with a containment graph with quadratic complexity.

Proof. The permutation which realizes the ambiguity of \mathcal{R} must always be containment-compatible. This is because swapping a region R with a region R' that contains R in the permutation π always improves the π -ambiguity. ◀

Let $e(\mathcal{R})$ be the number of linear extensions of the poset induced by \mathcal{R} . In the proof of Lemma 3.2 [4] Cardinal et al. show that $\log e(\mathcal{R}) \leq nH(\mathcal{R}) \leq 2 \log e(\mathcal{R})$. This implies that the interval graph entropy is a lower-bound for constructing any unique linear order underlying a poset. Proximity structures depend on sorting [5]. Thus, we conclude:

▶ **Theorem 7.** *Reconstructing a proximity structure on \mathcal{R} is lower-bounded by $\Omega(A(\mathcal{R}))$.*

3 Sorting

Let $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ be a set of intervals and let $X = \{x_1, x_2, \dots, x_n\}$ be a set of points (values) with $x_i \in R_i$. We show how to construct an auxiliary structure Ξ on \mathcal{R} in the preprocessing phase without using X , such that, in the reconstruction phase, we can construct a linear-size binary search tree T on X in $\Theta(A(\mathcal{R}))$ time. To achieve this, we first construct a specific containment-compatible permutation π of \mathcal{R} , and then show how to maintain Ξ when we process the intervals in this order.

3.1 Level permutation

We need a processing permutation π of \mathcal{R} with the following conditions:


- (i) π is containment-compatible,
- (ii) intervals containing no interval of \mathcal{R} come first and are ordered from right to left and
- (iii) we can construct π in $\mathcal{O}(n \log n)$ time.

In Section 2.1 we showed that if condition (i) holds, the π -ambiguity is a lower-bound for sorting X . In Section 3.2 we show that condition (ii) is useful to reconstruct an AVL-tree on X in $\mathcal{O}(A^\pi(\mathcal{R}))$ time. Condition (iii) bounds the time used in the preprocessing phase.

Below, we define two natural partitions of \mathcal{R} based on the containment graph of \mathcal{R} : the *height partition* and the *depth partition*. However, a permutation compatible with the height partition satisfies conditions (i) and (ii) but not (iii), and a permutation compatible with the depth partition satisfies conditions (i) and (iii) but not (ii). Therefore, we define a hybrid partition, which we call the *level partition*, which implies a permutation which does satisfy all three conditions, below.

Containment graph. For a set of intervals \mathcal{R} , its containment graph $G(\mathcal{R})$ represents the containment relations on \mathcal{R} . $G(\mathcal{R})$ is a directed acyclic graph where R_i contains R_j if and only if there is a directed path from R_i to R_j and all intervals $R \in \mathcal{R}$ that are contained in no other interval of \mathcal{R} share a common root. The bottom intervals are a subset of the leaves of this graph. Note that $G(\mathcal{R})$ can have quadratic complexity (Figure 6).

42:10 Preprocessing Ambiguous Imprecise Points



■ **Figure 7** (left) A set of intervals \mathcal{R} and the corresponding containment graph $G(\mathcal{R})$, leaves of $G(\mathcal{R})$ are purple. (middle) The height partition. (right) The depth partition.

Height and Depth partition. We define the *height partition* as the partition of \mathcal{R} into m levels $\mathcal{H} = H_1 \dots H_m$, $H_i \subseteq \mathcal{R}$ where all $R \in H_j$ have *height* (minimal distance from R to a leaf) $j + 1$ in $G(\mathcal{R})$ or equivalently: the intervals in H_{j+1} contain no intervals in $\mathcal{R} \setminus H_{\leq j}$ (Figure 7). We analogously define the *depth partition* as the partition of \mathcal{R} into m levels $\mathcal{D} = D_1 \dots D_m$, $D_i \subseteq \mathcal{R}$ where all $R \in D_j$ have *depth* (maximal distance from the root to R) $(m - j)$ in $G(\mathcal{R})$. Clearly any permutation compatible with \mathcal{H} or \mathcal{D} satisfies condition (i). All leaves of $G(\mathcal{R})$ have height 1 so per definition are all in H_1 and thus any permutation compatible with \mathcal{H} that sorts H_1 satisfies condition (ii). Clearly the same is not true for \mathcal{D} . On the other hand, in Lemma 8 we show how to construct \mathcal{D} in $\mathcal{O}(n \log n)$ time. It is unknown whether the height partition can be created in $\mathcal{O}(n \log n)$ time (refer to the full version).

► **Lemma 8.** *For any set of intervals \mathcal{R} we can construct \mathcal{D} in $\mathcal{O}(n \log n)$ time.*

Proof. We iteratively insert intervals from left to right; refer to the full version. ◀

Level partition. We now define the *level partition*: a hybrid between \mathcal{H} and \mathcal{D} : $\mathcal{L} = L_1 \dots L_m$, where all $R \in L_j$ have depth $(m - j)$ in $G(\mathcal{R})$ except for the leaves of $G(\mathcal{R})$, which are in L_1 regardless of their depth. We can compute the level partition from \mathcal{D} in $\mathcal{O}(n \log n)$ time by identifying all leaves of $G(\mathcal{R})$ with a range query. The *level permutation* is the permutation where intervals in L_i precede intervals in L_j and where within each level the intervals are ordered from right to left. It can be constructed from \mathcal{L} in $\mathcal{O}(n \log n)$ time by sorting.


Theorem 9 follows directly from the preceding discussion.

► **Theorem 9.** *The level permutation satisfies conditions (i), (ii) and (iii).*

3.2 Algorithm

We continue to describe a preprocessing and reconstruction algorithm to preprocess a set of intervals \mathcal{R} in $\mathcal{O}(n \log n)$ time such that we can sort X in $\Theta(A(\mathcal{R}))$ time.

Anchors. Let π be the level permutation of \mathcal{R} . In the preprocessing phase we build an AVL-tree T on the bottom intervals. In the reconstruction phase, we insert each remaining $x_i \in X$ into T in the order π in $\mathcal{O}(A^\pi(\mathcal{R}))$ time. This implies that for bottom intervals we are not allowed to spend even constant time and for each non-bottom interval R_i , we want to locate x_i in T in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. To achieve this, we supply every non-bottom interval R with an *anchor* denoted by $\perp^\pi(R_i)$. For a non-bottom interval $R \notin L_1$, we define its anchor



■ **Figure 8** The auxiliary structure Ξ . In the level L_1 all non-bottom intervals are shown their anchor. (top) A schematic representation of intervals in the level permutation π (from bottom to top). (bottom) The Fibonacci tree T containing the subset X^b corresponding to the bottom intervals. Note that we added one dummy node in red.

as an arbitrary interval contained in R . All intervals in L_1 are ordered from right to left, so for any non-bottom interval $R \in L_1$, its right endpoint is contained in the interval preceding it and we make this interval the anchor of R (refer to Figure 8).

Preprocessing phase. The auxiliary structure Ξ is an AVL-tree T on the bottom intervals, augmented with a set of pointers leading from intervals to their anchors. We will implement T as a leaf-based AVL-tree, i.e., where values are stored in the leaves, and inner nodes are decision nodes. Finally, we will use a doubly linked list to connect the leaves of the tree.

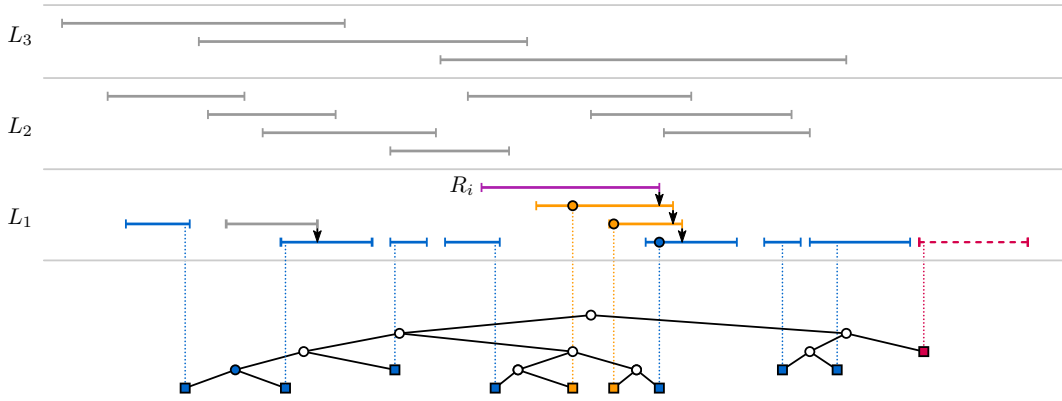
Let $X^b \subset X$ be the points corresponding to bottom intervals. Bottom intervals are mutually disjoint and we can build an AVL-tree T on X^b without knowing their true values. Recall that a Fibonacci tree is a tree binary where for every inner node, its left subtree has a depth 1 greater than its right subtree. A Fibonacci tree is a valid AVL-tree and we construct the AVL-tree over X^b as a Fibonacci tree where we add at most $|X^b|$ dummy leaves with value ∞ to ensure that the total number of nodes is a Fibonacci number. Refer to Figure 8 for an example. We remove the bottom intervals from \mathcal{R} and for each non-bottom interval R we identify its anchor $\perp^\pi(R)$ and we supply R with a pointer to $\perp^\pi(R)$. As the final step of the preprocessing phase we connect the leaves of T in a doubly linked list. To summarize: Ξ consists of a graph of intervals connected by anchor pointers and an AVL-tree T . Each bottom interval is in T and each non-bottom interval has a directed path to a node in T .

► **Lemma 10.** *We can construct the auxiliary structure Ξ in $\mathcal{O}(n \log n)$ time.*

Proof. The level partition and permutation can be constructed in $\mathcal{O}(n \log n)$ time and with it we get access to the intervals in L_1 sorted from right to left. We scan L_1 from right to left and for each interval $R \in L_1$ we either identify it as a bottom interval or to supply it with its anchor. We identify for each $R \notin L_1$ its anchor in logarithmic time using a range query. We construct the Fibonacci tree on X^b with leaf pointers in $\mathcal{O}(n \log n)$ time [20]. ◀

Reconstruction phase. During the reconstruction phase, we need to maintain the balance of T when we insert new values. T contains bottom intervals which we are not allowed to change even constant time, so the classical amortized-constant analysis [19] of AVL-trees does not immediately apply. Nonetheless we show in the full version:

42:12 Preprocessing Ambiguous Imprecise Points



■ **Figure 9** The tree T from Figure 8 after two iterations in the reconstruction phase. We inserted the true values of the two orange intervals. Note that an orange interval requested the true value of a bottom interval. At this iteration we want to insert the point x_i of R_i into T . R_i is a non-bottom interval in W_1 so its anchor must be the interval preceding it.

► **Lemma 11.** *Let T be an AVL-tree where each inner node has two subtrees with a depth difference of 1. We can dynamically maintain the balance of T in amortized $\mathcal{O}(1)$ time.*


► **Theorem 12.** *Given Ξ , we can reconstruct an AVL-tree on X in $\Theta(A^\pi(\mathcal{R}))$ time.*

Proof. Given Ξ and the level permutation π we want to sort the points in X (insert them into T) in $\mathcal{O}(A^\pi(\mathcal{R}))$ time. Because T starts as a Fibonacci tree, Lemma 11 guarantees that we can dynamically maintain the balance of T with at most $\mathcal{O}(A^\pi(\mathcal{R}))$ operations. The bottom intervals are already in T , thus we need to insert only the remaining $x_i \in X \setminus X^b$, in the order π , into T in $\log |\Gamma_i^\pi|$ time plus some additional time which we charge to the anchor (each anchor will only get charged once).

Whenever we process a non-bottom interval R_i we know that its anchor is already inserted in T . By construction, there are at most $\mathcal{O}(|\Gamma_i^\pi|)$ leaves in T which have coordinates on the domain of R_i (because these values can come only from intervals in the contact set of R_i). We know that we must insert x_i next to one of these $\mathcal{O}(|\Gamma_i^\pi|)$ leaves in T . This means that if we have a pointer to any leaf on the domain of R_i , then we locate x_i in T with at most $\mathcal{O}(\log |\Gamma_i^\pi|)$ edge traversals. During these traversals, we collapse each interval we encounter to a point. We obtain such a pointer from $\perp^\pi(R_i)$. Assume $\perp^\pi(R_i) \subset R_i$. Then the leaf corresponding to $\perp^\pi(R_i)$ must lie on the domain of R_i . Otherwise, R_i and $\perp^\pi(R_i)$ are both in the level L_1 (illustrated in Figure 9) and $\perp^\pi(R_i) = R_{i-1}$ and must contain the right endpoint of R_i . With a similar analysis, R_{i-1} can locate the right endpoint of R_i in T in $\mathcal{O}(\log |\Gamma_{i-1}^\pi|)$ time. In both cases we found a leaf of T in R_i and locate x_i in T in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. Each interval in L_1 has a unique anchor, so each anchor in L_1 is charged this extra work once. ◀

4 Quadrees

Let $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ be a set of unit intervals in a bounding box (interval) \mathcal{B} (we discuss how to extend the approach later) and let $X = \{x_1, x_2, \dots, x_n\}$ be a set of points (values) with $x_i \in R_i$. We show how to construct an auxiliary structure Ξ on \mathcal{R} in the preprocessing phase without using X , such that, in the reconstruction phase, we can construct a linear-size quadtree T on X in $\Theta(A(\mathcal{R}))$ time. We recall several standard definitions.



■ **Figure 10** A set of points \mathcal{R} where the quadtree on \mathcal{R} has linear depth. If the blue points lie very close, the quadtree on \mathcal{R} needs unbounded complexity.


Point quadtrees. Suppose that we have a d -dimensional point set X in a bounding hypercube \mathcal{B} . A quadtree on (\mathcal{B}, X) is defined as follows: *split operator* is an operator that splits any d -dimensional hypercube into 2^d equal-sized hypercubes called *cells*. We recursively split \mathcal{B} until each point $p \in P$ lies within a unique cell [21]. A λ -deflated quadtree is a more relaxed quadtree where \mathcal{B} is split until each leaf cell contains at most λ points [3].

Region quadtrees. Let \mathcal{R} be a set of d -dimensional disks in a bounding hypercube \mathcal{B} . Let $\mathcal{T}(\mathcal{B})$ be the infinite set of possible quadtree cells on \mathcal{B} . For each $R_i \in \mathcal{R}$, we define its *storing cell* denoted by C_i as the largest cell in $\mathcal{T}(\mathcal{B})$ that is contained in R_i and contains the center of R_i [17]. T_i is the subtree induced by C_i . The *neighborhood* of R_i is the set of possible cells $C \in \mathcal{T}(\mathcal{B})$ with size $|C_i|$ that are intersected by R_i . We consider the quadtree T on \mathcal{R} to be the unique compressed quadtree where for each $R_i \in \mathcal{R}$, its neighborhood is in T .

Edge oracle tree. Depending on \mathcal{B} and X , the quadtree on (\mathcal{B}, X) does not necessarily have logarithmic depth (Figure 10) thus, point location in T is non-trivial. Har-Peled [10] introduced a fast point-location structure (later dubbed *edge-oracle tree* [17]) for any quadtree T . The edge-oracle tree E is created through *centroid decomposition*. Any tree with bounded degree δ has at least one *centroid edge* which separates a tree of n nodes into two trees with at least $\frac{n}{\delta}$ and at most $n - \frac{n}{\delta}$ nodes each. Moreover, one of these 2 trees is a *subtree* of T (a tree induced by a node as a root). For any subtree T' of T , we define its *corresponding node* in E (edge in T) as the lowest node in E which splits T into two parts, one of which contains T' and the other contains the root of T . This node must exist, is unique and the subtree containing T' has $\mathcal{O}(|T'|)$ nodes (refer to Figure 11).

Given a query point q , we can find the leaf cell C_q that contains q in the following way: each decision node v of E has 2 children where 1 child node w corresponds to a subtree T_w of T . We test whether q is contained in w in $\mathcal{O}(1)$ time by checking the bounding box of T_w .

We wish to preprocess \mathcal{R} such that we can reconstruct a linear-size λ -deflated quadtree T for X with pointers between leaves. However, T does not necessarily have linear size and dynamically maintaining pointers between leaves is non-trivial. To achieve this, one needs to maintain a *compressed* and *smooth* quadtree T (refer to the full version for details) and Hoog et al. [12] show how to dynamically maintain a smooth compressed quadtree with constant update time. We will build such a quadtree augmented with an edge-oracle tree initialized as a Fibonacci tree. We proceed analogously to the approach in Section 3.



■ **Figure 11** (left) A tree T with recursive centroid edges. (right) The corresponding edge-oracle tree E . The orange leaf is a subtree of T and its corresponding node in E is e_3 .

4.1 1-dimensional quadtrees on unit-size intervals

We show how to construct an auxiliary structure Ξ on \mathcal{R} without using X , such that we can construct a 2-deflated quadtree T on (\mathcal{B}, X) in $\Theta(A(\mathcal{R}))$ time.

Preprocessing phase. The auxiliary structure Ξ will be a smooth compressed quadtree T on the intervals \mathcal{R} augmented with an edge-oracle tree E on T , anchor pointers, and a containment-compatible processing permutation π of \mathcal{R} . Given T , we initialize E as a Fibonacci tree, possibly adding dummy leaves¹. We supply each R_i with a pointer to the node in E corresponding to T_i and we call this its *anchor* $\perp^\pi(R_i)$.

► **Lemma 13.** *The auxiliary structure Ξ can be constructed in $\mathcal{O}(n \log n)$ time.*

Proof. Hoog et al. [12] show that for any set of d -dimensional disks \mathcal{R} , its smooth compressed quadtree T on \mathcal{R} with corresponding edge-oracle tree E can be constructed in $\mathcal{O}(n \log n)$ time and that this tree has a worst-case constant update time. We turn E into a Fibonacci tree by inserting at most $\mathcal{O}(n)$ dummy leaves in $\mathcal{O}(n \log n)$ time in total. ◀

Reconstruction phase. By construction, each leaf in T intersects at most 2 bottom intervals of \mathcal{R} (since these are mutually disjoint). Therefore, we can construct a 2-deflated quadtree on X by inserting each $x_i \in X \setminus X^b$ in the order π into T . We observe the following:

► **Lemma 14.** *When we process an interval $R_i \in \mathcal{R}$, R_i intersects $\mathcal{O}(|\Gamma_i^\pi|)$ leaf cells of T .*

Proof. There can be at most 2 bottom intervals (left and right) of R_i whose neighborhood intersects R_i . All the other leaves on the domain of R_i are caused by either already processed points on the domain of R_i or are dummy nodes. For each dummy node there is a corresponding non-dummy node also on the domain of R_i . ◀

► **Lemma 15.** *When we process an interval R_i , we can locate, for any point $q \in R_i$, the leaf $C_q \in T$ which contains q in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time.*

Proof. If $C_q \in T_i$ then R_i has an anchor to T_i and from this anchor we locate C_q in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. Suppose C_q is to the left of T_i . We locate the left-most leaf of T_i in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time and traverse its neighbor pointer. The neighboring cell must lie in a subtree T_q neighboring T_i with $\mathcal{O}(|\Gamma_i^\pi|)$ nodes and this tree must contain C_q (Lemma 14). We now have a pointer to a node in T_q and from this node we locate C_q in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. ◀

► **Theorem 16.** *Given Ξ , we can construct a 2-deflated quadtree on X in $\Theta(A^\pi(\mathcal{R}))$ time.*

Proof. Given Ξ and any containment-compatible permutation π , we want to insert X into T in $\mathcal{O}(A^\pi(\mathcal{R}))$ time. An insertion in T creates 2 additional leaves in T (and therefore also in E) and Lemma 11 guarantees that we can dynamically maintain the balance of E with at most $\mathcal{O}(A^\pi(\mathcal{R}))$ operations. If we only consider the point set $X^b \subset X$ corresponding to the bottom intervals then T is already a 2-deflated quadtree on X^b independent of where the points of X^b lie in their uncertainty intervals. Therefore, we only need to insert the remaining $x_i \in X \setminus X^b$, in the order π , into T in $\log |\Gamma_i^\pi|$ time (potentially collapsing some of the bottom intervals when necessary). Using Lemma 15 we can locate the quadtree leaf C_{x_i} that contains x_i in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. This leaf is intersected by at most 2 bottom intervals, which we collapse into points whose location we locate in constant time using the leaf pointers. Thus each non-bottom interval inserts at most 3 points into T in $\mathcal{O}(\log |\Gamma_i^\pi|)$ time. ◀

¹ We may need to allow parents of leaves of T to have a single dummy leaf.

4.2 Generalization

If we stay in one dimension, then the result of Theorem 16 in fact generalizes to the case where \mathcal{R} is a set of arbitrary intervals since Lemma 14 and 15 do not depend on the intervals being unit size. However, the result also generalizes to the case where \mathcal{R} is a set of unit-size disks in d (constant) dimensions: first of all, any permutation of \mathcal{R} is containment-compatible. If the disks are unit size then each disk intersects at most K_d bottom disks where K_d is the kissing number so Lemma 14 generalizes. For any disk $R_i \in \mathcal{R}$, recall that T_i was the subtree of the storing cell of R_i . Any point $q \in R_i$ must lie in the *perimeter* of T_i which consists of at most $\mathcal{O}(5^d)$ subtrees of size $\mathcal{O}(|\Gamma_i^\pi|)$ therefore, Lemma 15 also generalizes. The result is even more general: this approach works for any collection \mathcal{R} of unit-size *fat* convex regions similar to, e.g. [2]. Interestingly, generalizing the result of Theorem 16 both to higher dimensions and to non-unit regions at the same time is not possible: in the full version we show that, independent of preprocessing, reconstructing a λ -deflated quadtree has a lower bound of $\Omega(\log n)$, which could be more than $A(\mathcal{R})$.

5 Conclusion

We introduced the ambiguity $A(\mathcal{R})$ of a set of regions \mathcal{R} as a more fine-grained measure of the degree of their overlap. We applied this concept to uncertainty regions representing imprecise points. In the preprocessing model we show that the ambiguity is a natural lower bound for the time complexity of the reconstruction of any proximity structure. We achieved these results via a link to the entropy of partial orders which is of independent interest. If the regions are intervals in 1D we show how to sort in $\Theta(A(\mathcal{R}))$ time, if the regions are unit balls in any dimension we show how to reconstruct quadtrees $\Theta(A(\mathcal{R}))$ time.

In the future we plan to investigate if our results can be generalized to other proximity structures such as Delaunay triangulations, minimum spanning trees, and convex hulls. In principle it is possible to convert quadtrees into all of these structures in linear time [15]. However, it is not clear how to do so, when working with an implicit representation of the results in the case that $A(\mathcal{R})$ is sub-linear.

References

- 1 Graham Brightwell and Peter Winkler. Counting linear extensions. *Order*, 8(3):225–242, 1991.
- 2 Kevin Buchin, Maarten Löffler, Pat Morin, and Wolfgang Mulzer. Delaunay triangulation of imprecise points simplified and extended. *Algorithmica*, 61:674–693, 2011. doi:10.1007/s00453-010-9430-0.
- 3 Kevin Buchin and Wolfgang Mulzer. Delaunay triangulations in $O(\text{sort}(n))$ time and more. *Journal of the ACM (JACM)*, 58(2):6, 2011.
- 4 Jean Cardinal, Samuel Fiorini, Gwenaël Joret, Raphaël M Jungers, and J Ian Munro. Sorting under partial information (without the ellipsoid algorithm). *Combinatorica*, 33(6):655–697, 2013.
- 5 Mark De Berg, Otfried Cheong, Marc Van Kreveld, and Mark Overmars. *Computational Geometry: Introduction*. Springer, 2008.
- 6 Olivier Devillers. Delaunay triangulation of imprecise points, preprocess and actually get a fast query time. *Journal of Computational Geometry*, 2(1):30–45, 2011.
- 7 Esther Ezra and Wolfgang Mulzer. Convex hull of points lying on lines in $o(n \log n)$ time after preprocessing. *Computational Geometry*, 46(4):417–434, 2013.
- 8 P.C. Fishburn and W.T. Trotter. Geometric containment orders: a survey. *Order*, 15:167–182, 1998.

- 9 Michael L Fredman. How good is the information theory bound in sorting? *Theoretical Computer Science*, 1(4):355–361, 1976.
- 10 Sariel Har-Peled. *Geometric approximation algorithms*. Number 173 in Mathematical Surveys and Monographs. American Mathematical Soc., 2011.
- 11 Martin Held and Joseph SB Mitchell. Triangulating input-constrained planar point sets. *Information Processing Letters*, 109(1):54–56, 2008.
- 12 Ivor Hoog vd, Elena Khramtcova, and Maarten. Löffler. Dynamic Smooth Compressed Quadtrees. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 99. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 13 Jeff Kahn and Jeong Han Kim. Entropy and sorting. *Journal of Computer and System Sciences*, 51(3):390–399, 1995.
- 14 János Körner. Coding of an information source having ambiguous alphabet and the entropy of graphs. In *6th Prague conference on information theory*, pages 411–425, 1973.
- 15 Maarten Löffler and Wolfgang Mulzer. Triangulating the square and squaring the triangle: quadtrees and Delaunay triangulations are equivalent. *SIAM Journal on Computing*, 41(4):941–974, 2012.
- 16 Maarten Löffler and Wolfgang Mulzer. Unions of onions: Preprocessing imprecise points for fast onion decomposition. *Journal of Computational Geometry*, 5:1–13, 2014.
- 17 Maarten Löffler, Joseph A Simons, and Darren Strash. Dynamic planar point location with sub-logarithmic local updates. In *Workshop on Algorithms and Data Structures*, pages 499–511. Springer, 2013.
- 18 Maarten Löffler and Jack Snoeyink. Delaunay triangulation of imprecise points in linear time after preprocessing. *Computational Geometry*, 43(3):234–242, 2010.
- 19 Kurt Mehlhorn and Athanasios Tsakalidis. An amortized analysis of insertions into AVL-trees. *SIAM Journal on Computing*, 15(1):22–33, 1986.
- 20 Jürg Nievergelt and Edward M Reingold. Binary search trees of bounded balance. *SIAM journal on Computing*, 2(1):33–43, 1973.
- 21 Hanan Samet. The quadtree and related hierarchical data structures. *ACM Computing Surveys (CSUR)*, 16(2):187–260, 1984.
- 22 Gábor Simonyi. Graph entropy: a survey. *Combinatorial Optimization*, 20:399–441, 1995.
- 23 Marc Van Kreveld, Maarten Löffler, and Joseph SB Mitchell. Preprocessing imprecise points and splitting triangulations. *SIAM Journal on Computing*, 39(7):2990–3000, 2010.