



Near-Linear Algorithms for Geometric Hitting Sets and Set Covers

Pankaj K. Agarwal¹ · Jiangwei Pan^{1,2}

Received: 20 January 2018 / Revised: 10 April 2019 / Accepted: 23 April 2019 / Published online: 9 May 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Given a finite range space $\Sigma = (X, \mathcal{R})$, with $N = |X| + |\mathcal{R}|$, we present two simple algorithms, based on the multiplicative-weight method, for computing a small-size hitting set or set cover of Σ . The first algorithm is a simpler variant of the Brönnimann–Goodrich algorithm but more efficient to implement, and the second algorithm can be viewed as solving a two-player zero-sum game. These algorithms, in conjunction with some standard geometric data structures, lead to near-linear algorithms for computing a small-size hitting set or set cover for a number of geometric range spaces. For example, they lead to $O(N \text{polylog}(N))$ expected-time randomized $O(1)$ -approximation algorithms for both hitting set and set cover if X is a set of points and \mathcal{R} a set of disks in \mathbb{R}^2 .

Keywords Geometric set cover · Near-linear algorithms · Multiplicative weight method · Disks · Rectangles

Mathematics Subject Classification 68U05 · 52C17

1 Introduction

Let $\Sigma = (X, \mathcal{R})$ be a finite range space where X is a finite set of *objects* and \mathcal{R} is a family of subsets of X called *ranges*. Set $n = |X|$, $m = |\mathcal{R}|$, and $N = n + m$. A

Editor in Charge: Kenneth Clarkson

Work on this paper is supported by NSF under Grants CCF-11-61359, IIS-14-08846, CCF-15-13816, and ISS-14-47554, by an ARO Grant W911NF-15-1-0408, and by Grant 2012/229 from the U.S.-Israel Binational Science Foundation. A preliminary version of the paper appeared in Proc. 30th Annual Symp. Comput. Geom. (SoCG'14).

Pankaj K. Agarwal
pankaj@cs.duke.edu

Jiangwei Pan
panjiangwei@gmail.com

¹ Department of Computer Science, Duke University, Box 90129, Durham, NC 27708-0129, USA

² Present Address: Netflix, Inc, 100 Winchester Cir, Los Gatos, CA 95032, USA

subset $H \subseteq X$ is called a *hitting set* of Σ if H intersects every range of \mathcal{R} , and a subset $\mathcal{C} \subseteq \mathcal{R}$ is called a *set cover* of Σ if the union of ranges in \mathcal{C} is equal to X . It is well known that a set cover of Σ is a hitting set of the dual range space Σ^\perp (see Sect. 2 for the definition of Σ^\perp). The hitting-set (resp. set-cover) problem is to find the smallest hitting set (resp. set cover) of Σ . Both problems are well-known to be NP-complete [27] and have been extensively studied. Let $\kappa := \kappa(\Sigma)$ and $\chi := \chi(\Sigma)$ denote the size of an optimal hitting set and set cover, respectively, of Σ . Sometimes, we also use OPT to denote either κ or χ and the meaning will be clear from the context.

In this paper, we are primarily interested in the hitting-set and set-cover problems for geometric range spaces, where X is a finite set of points in \mathbb{R}^d and \mathcal{R} is a finite family of simply-shaped regions chosen from some infinite class (e.g., rectangles, balls, simplices, halfspaces). In this case, ranges are $X \cap R$ for $R \in \mathcal{R}$. With a slight abuse of notation, we denote the range space $\Sigma = (X, \{X \cap R \mid R \in \mathcal{R}\})$ as (X, \mathcal{R}) . The hitting-set and set-cover problems are NP-complete even for simple geometric range spaces, e.g., when \mathcal{R} is a set of unit disks or unit squares in \mathbb{R}^2 [25]. This has led to the development of polynomial-time approximation algorithms for these problems. Traditionally, the focus has been on developing polynomial-time algorithms with the smallest possible approximation ratio. Many recent applications (e.g., in sensor networks, database systems, computer vision) call for repeated computation of hitting sets and set covers. For such applications, it is desirable to have near-linear-time algorithms for these problems even if it means sacrificing a little on the approximation ratio. Motivated by these applications, we study the problem of computing, in near linear time, near-optimal hitting sets and set covers for geometric range spaces.

Related work The well-known greedy algorithm gives a polynomial-time $O(\log n)$ -approximation for a hitting set or set cover [21]. The known lower bound results imply that this is the best one can hope for, within a constant factor, assuming certain conjectures in complexity theory [24]. However, by exploiting the underlying geometry, polynomial-time algorithms with better approximation factors can be obtained for many geometric range spaces; see [6, 18, 20, 39] and the references therein for various results of this kind. These algorithms employ and adapt a wide range of novel techniques, including the usage of ε -nets [30].

Given a parameter $\varepsilon \in (0, 1]$, an ε -net for range space $\Sigma = (X, \mathcal{R})$ is a subset $Q \subseteq X$ that intersects every range $R \in \mathcal{R}$ whose size is at least $\varepsilon|X|$. In other words, Q is a hitting set for all the “heavy” ranges. Haussler and Welzl [30] proved that a range space of VC-dimension¹ δ has an ε -net of size $O((\delta/\varepsilon) \log(\delta/\varepsilon))$. The bounds on the size of ε -nets have been improved for several geometric range spaces. For example, $O(1/\varepsilon)$ -size ε -nets exist when X is a set of points and \mathcal{R} is a set of halfspaces in \mathbb{R}^2 or \mathbb{R}^3 , pseudo-disks \mathbb{R}^2 , or translates of a fixed convex polytope in \mathbb{R}^3 [33, 36, 42]. Aronov et al. [6] gave an $O(n \log^{d-1} n)$ -expected-time randomized algorithm to construct an ε -net of size $O((1/\varepsilon) \log \log(1/\varepsilon))$ for the case when \mathcal{R} is a set of axis-parallel d -rectangles for $d = 2, 3$. Bus et al. [11] further improved the upper-bound to $13.4/\varepsilon$ for disks in the plane, based on a construction using Delaunay triangulations. Recent

¹ The VC-dimension of a range space $\Sigma = (X, \mathcal{R})$ is the size of the largest subset $A \subseteq X$ such that $|\{A \cap R \mid R \in \mathcal{R}\}| = 2^{|A|}$.

results of Pach and Tardos [40] and Kupavskii et al. [32] show that if \mathcal{R} is a set of axis-parallel d -rectangles, then the size of an ε -net in the worst-case is $\Omega(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$ for $d = 2, 3$ and $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ for $d \geq 4$. See also [4,23] for other results on ε -nets.

Building on an iterative-reweighting technique by Clarkson [18], Brönnimann and Goodrich [10] showed that if an ε -net of a range space Σ of size $O((1/\varepsilon)g(1/\varepsilon))$ can be computed in polynomial time, where $g(\cdot)$ is a monotonically-nondecreasing sublinear function, then a hitting set of size $O(\kappa g(\kappa))$ also can be computed in polynomial time. Since a set cover of Σ is a hitting set of the dual range space of Σ , their algorithm also extends to set cover. If Σ has finite VC-dimension, then the dual range space also has finite VC-dimension, and thus a hitting set (resp. set cover) of Σ of size $O(\kappa \log \kappa)$ (resp. $O(\chi \log \chi)$) can be computed in polynomial time. The size of hitting set reduces to $O(\kappa)$ if Σ admits an ε -net of size $O(1/\varepsilon)$ (e.g. when \mathcal{R} is a set of halfspaces in \mathbb{R}^3 , or a set of pseudo-disks in \mathbb{R}^2), and to $O(\kappa \log \log \kappa)$ if \mathcal{R} is a set of rectangles in \mathbb{R}^2 or \mathbb{R}^3 . The Brönnimann–Goodrich algorithm can be viewed as first solving the LP relaxation of the hitting-set problem and then obtaining an approximate hitting set from the fractional LP solution via ε -net. A more general LP solver can also be used to obtain the same results [34]. See [6,20,44] for improved approximation ratios for set covers of geometric range spaces.

The algorithms by Clarkson [18] and Brönnimann–Goodrich [10] are instances of the so-called *multiplicative weight* (MW) method. The MW method, which has been repeatedly discovered, goes back to the 1950's and has been used in numerous fields including machine learning, linear programming, semidefinite programming, graph algorithms, game theory, on-line algorithms, and computational geometry. In computational geometry, besides hitting-set/set-cover, the MW method has been used for linear programming [19,31], constructing spanning trees of low stabbing number [16], and constructing partition trees for range searching [35]. See the survey by Arora et al. [8] for a comprehensive review of this method.

Approximation algorithms for hitting set and set cover are known that do not rely on ε -nets. For instance, the local-search technique by Mustafa and Ray [39] computes $(1+\varepsilon)$ -approximate hitting sets for halfspaces in \mathbb{R}^3 and pseudo-disks in \mathbb{R}^2 in $N^{O(\varepsilon^{-2})}$ time. Mustafa et al. [38] developed a quasi-polynomial-time algorithm for the case where the halfspaces and pseudo-disks are weighted.

Both the greedy and the Brönnimann–Goodrich algorithms work in $O(\kappa \log N)$ stages. A straightforward implementation of both algorithms takes $O(mn)$ time per stage, which can be improved to $O(N \text{polylog}(N))$ in some cases using geometric data structures. Thus these algorithms run in time $\Omega(N\kappa)$.

Agarwal et al. [3] studied near-linear algorithms for the hitting-set problem. They proposed a variant of the greedy algorithm that performs $O(\log n)$ stages and chooses, in near-linear time, $O(\kappa \phi(\kappa))$ points in each stage for the hitting set if the union of any subset $\mathcal{F} \subseteq \mathcal{R}$ has complexity $O(|\mathcal{F}| \phi(|\mathcal{F}|))$, where $\phi(\cdot)$ is a sublinear function. They also presented an efficient implementation of the Brönnimann–Goodrich algorithm for the special case when \mathcal{R} is a set of d -rectangles for $d = 2, 3$. Their algorithm is rather complex and computes a hitting set of size $O(\kappa \log \log \kappa)$ in $O((N + \kappa^{d+1}) \text{polylog}(n))$ time. Not only is this algorithm inefficient for large values of κ , but it also does not extend to other range spaces.

Our results We present two simple algorithms, both based on the MW method, for computing a small-size hitting set or set cover of a range space Σ of finite VC-dimension. They first compute a near-optimal fractional solution using the MW method and then transform this fractional solution to an integral solution using ε -nets (e.g., as in [34]). As mentioned earlier, a set cover of Σ is a hitting set of the dual range space Σ^\perp , so we only describe our algorithms in terms of computing a hitting set.

The first algorithm, described in Sect. 3, is a simpler but more efficient variant of the Brönnimann–Goodrich algorithm [10]. It divides the $O(\kappa \log n)$ stages of their algorithm into $O(\log n)$ rounds so that each range is processed only once in each round, and it computes an ε -net twice, instead of $O(\kappa \log n)$ times as in [10]. Assuming that an ε -net of Σ can be computed in $O(n \text{polylog}(n))$ time and range queries on Σ can be answered in $\text{polylog}(n)$ time, the expected running time of the algorithm is $O(N \text{polylog}(N))$. See Theorem 3.2 for a more precise bound.

The second algorithm, described in Sect. 4, is a Monte Carlo algorithm that computes a small-size hitting set or set cover with high probability (i.e., with probability at least $1 - 1/N^{\Omega(1)}$). It can be viewed as solving a two-player zero-sum game. In this game, one player Alice has the points in X as her pure strategies and the other player Bob has the ranges in \mathcal{R} as his pure strategies. If Alice plays $x \in X$ and Bob plays $R \in \mathcal{R}$, then Bob pays Alice 1 if $x \in R$ and 0 otherwise. An optimal mixed strategy for Alice gives an optimal fractional solution to the hitting-set, from which one can use ε -net to obtain a hitting set of small size. Using the MW method, the algorithm computes near-optimal mixed strategies for both Alice and Bob, which gives a near-optimal fractional solution to the hitting-set problem. We remark that the MW method has been used to solve the zero-sum game approximately [26,28], but our algorithm is somewhat different, tailored to computing a hitting set.

We next describe in Sect. 5 consequences of these algorithms for a number of geometric range spaces. The results are presented in terms of the second algorithm and thus achieve the approximation ratios with high probability. We can obtain Las Vegas algorithms for these problems using the first algorithm, possibly paying an additional logarithmic factor in the running time, that are guaranteed to obtain these approximation ratios. See Table 1 for a summary of the results.

Either no near-linear algorithms with guaranteed approximation ratio were known for these problems, or near-linear algorithms attained a worse approximation ratio. For example, the algorithm by Agarwal et al. [3] computes only an $O(\log n)$ -approximate hitting set for disks in $O(N \log^3 N)$ time, and it does not extend to computing a set cover. We remark that our algorithms rely on standard range searching data structures. It might be possible to improve polylog factors in the running time by optimizing these data structures, but we feel it is not worth the effort.

Finally, in Sect. 6, we extend the second algorithm to compute, in near-linear time, an $O(1)$ -approximate (discrete) *maximum independent set* for disks. In this problem, we are also given a discrete range space $\Sigma = (X, \mathcal{R})$, where X is a set of points and \mathcal{R} a set of disks in \mathbb{R}^2 , and the goal is to compute the largest subset \mathcal{J} of disks such that no point of X is contained in more than one disk of \mathcal{J} . We first use our second algorithm to compute a fractional solution to this problem, and then round the fractional solution to an integral one using the algorithm by Chan and Har-Peled [13]. The algorithm in [13] employs an existing LP solver and runs in quadratic time.

Table 1 Summary of results for different geometric range spaces using the second algorithm, so these bounds hold with probability at least $1 - \frac{1}{N^{\Omega(1)}}$

Ranges	Approximation ratio		Running time
	Hitting set	Set cover	
Rectangles ($d = 2, 3$)	$O(\log \log \kappa)$	$O(\log \chi)$	$O(N \log^{d+1} N \log \log \text{OPT})$
Halfspaces ($d = 3$)	$O(1)$	$O(1)$	$O(N \log^3 N)$
Disks ($d = 2$)	$O(1)$	$O(1)$	$O(N \log^3 N)$
Fat triangles ($d = 2$)	$O(\log \kappa)$	$O(\log \chi)$	$O(N \log^7 N \log \log \text{OPT})$

The first algorithm gives the same approximation ratios with probability 1, but the running time might have an extra $\log N$ factor

2 Preliminaries

Range space and ε -nets Let $\Sigma = (X, \mathcal{R})$ be a finite range space, as defined above. The dual range space of Σ , denoted by $\Sigma^\perp = (X^\perp, \mathcal{R}^\perp)$ is defined as follows. There is an object in X^\perp for each range $R \in \mathcal{R}$, and \mathcal{R}^\perp contains a range R_x for each point $x \in X$, namely, $R_x = \{R \in \mathcal{R} \mid R \ni x\}$. As mentioned above, a set cover of Σ is a hitting set of Σ^\perp and vice-versa. It is also known that if Σ has finite VC-dimension, then so does Σ^\perp [29].

A hitting set $H \subseteq X$ is called *c-approximate*, for $c \geq 1$, if $|H| \leq c\kappa(\Sigma)$. Similarly, a set cover $\mathcal{C} \subseteq \mathcal{R}$ is called *c-approximate* if $|\mathcal{C}| \leq c\chi(\Sigma)$.

Given a weight function $w : X \rightarrow \mathbb{R}_{\geq 0}$, for a subset $A \subseteq X$, we use $w(A)$ to denote the total weight of points in A . Given w and a parameter $\varepsilon \in (0, 1]$, we call a range $R \in \mathcal{R}$ *ε -heavy* if $w(R) \geq \varepsilon w(X)$ and *ε -light* otherwise. A subset $Q \subseteq X$ is called an *ε -net* with respect to weight function w if $Q \cap R \neq \emptyset$ for every ε -heavy range R of \mathcal{R} .

Hitting-set algorithm We first briefly describe the main idea behind the algorithm of Brönnimann–Goodrich [10] (and Clarkson [18]), which uses a similar framework as other more general multiplicative-weight algorithms such as the Plotkin–Shmoys–Tardos algorithm [41] for solving packing and covering LPs.

Let k be an integer such that $k/2 < \kappa \leq k$. Initialize the weight of each point to 1, i.e., $w(x) = 1$ for all $x \in X$, and repeat the following weight-doubling step until every range is $\frac{1}{2k}$ -heavy:

Find a $\frac{1}{2k}$ -light range R and double the weights of all points in R .

When the process stops, return a $\frac{1}{2k}$ -net Q of Σ with respect to the final weights. Since each range in Σ is $\frac{1}{2k}$ -heavy, Q is a hitting set of Σ . Hence, if a $\frac{1}{2k}$ -net of size $O(kg(k))$ can be computed efficiently, the above algorithm computes a hitting set of size $O(\kappa g(\kappa))$. The following lemma, by now a well-known argument (see e.g. [8, 10, 16, 18, 19]), is the key to the performance of the algorithm. For completeness, we also give the proof here.²

² Throughout this paper, we use $\log x$ to denote $\log_2 x$.

Lemma 2.1 *If Σ has a hitting set of size at most k , then the algorithm performs at most $\mu_k := 4k \log(n/k)$ weight-doubling steps. The final weight of X is at most n^4/k^3 .*

Proof Let H be a hitting set of size k . Since each weight-doubling step is performed on a $\frac{1}{2k}$ -light range R , the total weight of X increases by a factor of at most $(1 + \frac{1}{2k})$. Thus, after z weight-doubling steps,

$$w(X) \leq n \left(1 + \frac{1}{2k}\right)^z \leq n \exp\left(\frac{z}{2k}\right).$$

On the other hand, $H \cap R \neq \emptyset$, thus at least one point of H gets its weight doubled after the weight-doubling step on R . Suppose each element of $h \in H$ is doubled z_h times after z weight-doubling steps, we have

$$w(H) = \sum_{h \in H} 2^{z_h} \geq k 2^{z/k}.$$

Since $w(H) \leq w(X)$, we get

$$k 2^{\frac{z}{k}} \leq n e^{\frac{z}{2k}} \leq n 2^{\frac{3z}{4k}},$$

from which $z \leq \mu_k = 4k \log(n/k)$ follows. The final weight of X follows from the first inequality. □

By Lemma 2.1, if the algorithm does not terminate within μ_k steps, we can conclude that Σ does not have a hitting set of size at most k . An exponential search is used to guess the value of k such that $k/2 < \kappa \leq k$.

In each step, the Brönnimann–Goodrich algorithm computes a $\frac{1}{2k}$ -net H and identifies a range that does not intersect H as a light range. However, verifying whether H hits all ranges takes $\Omega(N)$ time, and thus the algorithm can take $\Omega(N^2)$ time in the worst case; see the original paper for details. In the next section, we group every $O(k)$ weight-doubling steps into one round, and find light ranges for these steps together by looking at each range in R only once. This optimization makes it possible to finish all weight-doubling steps in $\tilde{O}(N)$ time for many geometric instances.

3 Hitting Set in Rounds

Algorithm As in [10], assume we have an integer k such that $\kappa \in (k/2, k]$; we perform an exponential search to find such a value of k . The algorithm is similar to [10] except that it works in rounds. Initially, it sets $w(x) = 1$ for all $x \in X$. Each round of the algorithm performs at most $2k$ weight-doubling steps, as follows. It processes each range $R \in \mathcal{R}$ one by one. If R is $\frac{1}{2k}$ -light, it doubles the weights of all points in R . This weight-doubling step on R is performed repeatedly until R becomes $\frac{1}{2k}$ -heavy or $2k$ weight-doubling steps have been performed in the current round. Once R becomes

$\frac{1}{2k}$ -heavy, it is not processed again in the current round, even though it may become $\frac{1}{2k}$ -light again later in the current round while other ranges are being processed.

If $2k$ weight-doubling steps have been performed in the current round, the algorithm aborts the current round and moves to the next round. On the other hand, if all ranges have been processed with less than $2k$ weight-doubling steps, the algorithm stops and returns a $\frac{1}{2ke}$ -net Q of Σ as a hitting set of Σ .

By Lemma 2.1, if Σ has a hitting set of size at most k , then the algorithm performs at most $\mu_k = 4k \log(n/k)$ weight-doubling steps. Since each except the last round performs $2k$ weight-doubling steps, the number of rounds is at most $\frac{\mu_k}{2k} + 1 = 2 \log(n/k) + 1$.

Correctness The correctness of the algorithm follows from the following lemma.

Lemma 3.1 *All ranges are $\frac{1}{2ke}$ -heavy when the algorithm terminates.*

Proof Suppose the algorithm stops in round i . Let W_i be the total weight $w(X)$ in the beginning of round i , and let W_f be the total weight when the algorithm stops. Since at most $2k$ weight-doubling steps are performed in round i and each of them increases the total weight by a factor of at most $1 + \frac{1}{2k}$,

$$W_f \leq \left(1 + \frac{1}{2k}\right)^{2k} W_i \leq e W_i.$$

After a range R has been processed in round i , it is $\frac{1}{2k}$ -heavy with respect to the current weight of X , which is at least W_i . Therefore

$$w(R) \geq \frac{1}{2k} W_i \geq \frac{1}{2ke} W_f,$$

implying that R is $\frac{1}{2ke}$ -heavy. Since the algorithm terminates in round i only after all ranges have been processed in that round, all ranges are $\frac{1}{2ke}$ -heavy. □

If an ε -net of Σ of size $O\left(\frac{1}{\varepsilon} g\left(\frac{1}{\varepsilon}\right)\right)$ can be computed, then by Lemma 3.1, the above algorithm returns a hitting set of size $O(\kappa g(\kappa))$.

Running time To expedite the weight-doubling step, we perform the following preprocessing step before running the algorithm: compute a $\frac{1}{2k}$ -net Q_0 of Σ , and set $X = X \setminus Q_0$ and $\mathcal{R} = \{R \mid R \cap Q_0 = \emptyset\}$. After this preprocessing, each range in \mathcal{R} contains at most $\frac{n}{2k}$ points. The algorithm returns $Q_0 \cup Q$ as a hitting set of Σ .

We assume the existence of the following three procedures:

- (P1) A net finder that can compute an ε -net of size $O\left(\frac{1}{\varepsilon} g\left(\frac{1}{\varepsilon}\right)\right)$ in $\varphi(N)$ time.
- (P2) A range-counting data structure that, given a range R , returns YES if R is $\frac{1}{2k}$ -light and NO otherwise, and that can also update the weight of a point. Let $\tau(n) = \Omega(\log n)$ be the time taken by each of these two operations. We assume that the data structure can be built in $O(n\tau(n))$ time.

(P3) A range-reporting data structure that reports all s points of a range in $O(\tau(n) + s)$ time, and that can be constructed in $O(n\tau(n))$ time.

Using (P2) and (P3), a range $R \in \mathcal{R}$ can be processed in each round as follows. First, using (P2) we check in $\tau(n)$ time whether R is $\frac{1}{2k}$ -light. If the answer is YES, we use (P3) to report all points of R in time $\tau(n) + |R| \leq \tau(n) + \frac{n}{2k}$. Recall that at most $2k$ weight-doubling steps are performed in each round, so $O(n)$ points are reported in a round. For each reported point, it takes $\tau(n)$ time to double its weight in the range-counting data structure. We thus conclude that a round takes $O(N\tau(n))$ time. Summing over all $2 \log(n/2k) + 1$ rounds, the total time spent by the algorithm, including the time spent in the preprocessing phase, is $O(N\tau(n) \log n + \varphi(N))$. We repeat the algorithm $O(\log \kappa)$ times to perform the exponential search. Putting everything together, we obtain the following.

Theorem 3.2 *Let $\Sigma = (X, \mathcal{R})$ be a finite range space with $|X| + |\mathcal{R}| = N$. Suppose procedures (P1), (P2) and (P3) exist for Σ , then an $O(g(\kappa))$ -approximate hitting set of Σ can be computed in $O((N\tau(N) \log N + \varphi(N)) \log \kappa)$ time.*

Recalling that a set cover of Σ is a hitting set of the dual range space Σ^\perp , we obtain the following.

Theorem 3.3 *Let $\Sigma = (X, \mathcal{R})$ be a finite range space with $|X| + |\mathcal{R}| = N$. Suppose procedures (P1), (P2) and (P3) exist for Σ^\perp , then an $O(g(\chi))$ -approximate set cover of Σ can be computed in $O((N\tau(N) \log N + \varphi(N)) \log \chi)$ time.*

Remarks (i) An approximate range counting data structure can be used in (P2), namely, a data structure that always returns YES if R is $\frac{1}{4k}$ -light and always returns NO if R is $\frac{1}{2k}$ -heavy. It may return YES or NO if $w(R) \in [\frac{1}{4k}w(X), \frac{1}{2k}w(X)]$. So if the algorithm returns YES, then R must be $\frac{1}{2k}$ -light; on the other hand, if the algorithm returns NO, then R must be $\frac{1}{4k}$ -heavy. A weight-doubling step is performed if the procedure returns YES. The same argument as in Lemma 3.1 implies that each range is at least $\frac{1}{4ke}$ -heavy when the algorithm terminates, after $O(\log(n/k))$ rounds. Hence, the size of hitting set is affected by a factor of at most 2. This is convenient because faster data structures are known for approximate range counting for some range spaces (e.g., halfspace range counting).

(ii) By Lemma 2.1, the final weight of X is bounded by $O(n^4/k^3)$, so $O(\log n)$ bits suffice for maintaining the weight of each point during the algorithm.

4 Hitting Set as 2-Player Game

Algorithm We now describe the second algorithm for computing a hitting set. As earlier, suppose we have an integer k such that $k/2 < \kappa \leq k$; we perform a backward exponential search to find such a value of k as described later. The algorithm now maintains weights on both X and \mathcal{R} , and it also works in rounds. For $i \geq 1$, let $\pi^i : X \rightarrow \mathbb{R}_{\geq 0}$ and $\omega^i : \mathcal{R} \rightarrow \mathbb{R}_{\geq 0}$ be the weights of points and ranges, respectively, in the beginning of round i . We also use $\pi^i(X)$ and $\omega^i(\mathcal{R})$ to denote the total weight

of all points in X and all ranges in \mathcal{R} , respectively. Initially, $\pi^1(x) = 1$ for all $x \in X$ and $\omega^1(R) = 1$ for all $R \in \mathcal{R}$. Let Π^i (resp. Ω^i) denote the probability distribution induced by π^i (resp. ω^i), i.e.,

$$\Pi^i = \left\langle \frac{\pi^i(x_j)}{\pi^i(X)} \mid 1 \leq j \leq n \right\rangle, \quad \Omega^i = \left\langle \frac{\omega^i(R_j)}{\omega^i(\mathcal{R})} \mid 1 \leq j \leq m \right\rangle.$$

We set

$$\mu := \frac{2}{\ln 2} k \ln(m^2 n).$$

The algorithm performs μ rounds of the following two steps: for $1 \leq i \leq \mu$,

- (i) Sample a point $\bar{x}_i \in X$ from the distribution Π^i and a range $\bar{R}_i \in \mathcal{R}$ from the distribution Ω^i .
- (ii) For each point $x \in \bar{R}_i$, double its weight, i.e., $\pi^{i+1}(x) = 2\pi^i(x)$, and for each range R that contains \bar{x}_i , halve its weight, i.e., $\omega^{i+1}(R) = \omega^i(R)/2$.

Let $\tilde{X} = \langle \bar{x}_1, \dots, \bar{x}_\mu \rangle$ be the multi-set of points chosen by the algorithm. Let $\tilde{\Pi}$ be the distribution on X induced by \tilde{X} , i.e., if a point $x \in X$ appears μ_x times in \tilde{X} , then set $\Pr(x \sim \tilde{\Pi}) = \mu_x/\mu$. We compute a $\frac{1}{8k}$ -net Q of Σ with respect to $\tilde{\Pi}$. If Q is a hitting set of Σ , we return Q ; otherwise, we repeat the above algorithm.

Intuitively, the weights of points that lie in many ranges are doubled more frequently, and thus more likely to be in the optimal hitting set; on the other hand, the weights of ranges that contain few points are halved less frequently, so that they can be sampled to ensure points contained in them have high weights.

Correctness We view the hitting-set problem for Σ as a two-player zero-sum game, and the above algorithm computes a near-optimal mixed strategy for each of the two players. More precisely, let Alice and Bob be two players who play the following game: Alice chooses a point $x \in X$, Bob chooses a range $R \in \mathcal{R}$ and Bob pays $I(x, R)$ to Alice, where

$$I(x, R) = \begin{cases} 1 & \text{if } x \in R, \\ 0 & \text{if } x \notin R. \end{cases}$$

The algorithm can be interpreted as Alice and Bob playing the game for μ rounds. Initially, both players have uniform distributions over their pure strategies. In each round, Alice samples a point x from the current distribution over points and Bob samples a range R from the current distribution over ranges. According to the definition of the game, points in R are good for Alice because had she played one of them, she would have won 1 dollar. Therefore, Alice doubles the weights of points in R so that she can choose them more often in the future. Similarly, ranges containing x are bad for Bob, thus he halves the weights of those ranges.

For a probability distribution Π over X and for a range $R \in \mathcal{R}$, let $I(\Pi, R)$ denote the expected payoff to Alice if Bob chooses R , i.e.,

$$I(\Pi, R) = \sum_{x \in X} \Pr(x \sim \Pi) I(x, R).$$

Similarly, we define $I(x, \Omega)$ for a point $x \in X$ and a distribution Ω over \mathcal{R} . Let λ^* be the value of the above game. By the min-max theorem [37],

$$\lambda^* = \max_{\Pi} \min_{R \in \mathcal{R}} I(\Pi, R) = \min_{\Omega} \max_{x \in X} I(x, \Omega), \tag{1}$$

where Π, Ω are probability distributions over X and \mathcal{R} , respectively.

Let $H^* \subseteq X$ be an optimal hitting set of Σ of size κ , and let Π_{H^*} be the distribution where $\pi(x) = 1/\kappa$ if $x \in H^*$ and 0 otherwise. Then

$$\min_{R \in \mathcal{R}} I(\Pi_{H^*}, R) \geq \frac{1}{\kappa}$$

because $R \cap H^* \neq \emptyset$ for all $R \in \mathcal{R}$. Hence $\lambda^* \geq 1/\kappa \geq 1/k$. Let

$$\Pi^* = \arg \max_{\Pi} \min_{R \in \mathcal{R}} I(\Pi, R)$$

be the optimal (mixed) strategy for Alice. If we can compute Π^* , then we can simply return a $(1/k)$ -net Q of Σ under the distribution Π^* : Q is a hitting set of Σ because the weight of any $R \in \mathcal{R}$ under Π^* is at least $\min_{R \in \mathcal{R}} I(\Pi^*, R) = \lambda^* \geq 1/k$. We show that $\tilde{\Pi}$ is an approximation of Π^* , in the sense that with constant probability,

$$\min_{R \in \mathcal{R}} I(\tilde{\Pi}, R) \geq \frac{1}{8k},$$

and thus a $\frac{1}{8k}$ -net of Σ under $\tilde{\Pi}$ is a hitting set of Σ .

We begin by proving two lemmas, which follow from standard arguments for the MW method. The first one states that the total expected payoff to Alice over μ rounds is not much less than the payoff she would get by the best pure strategy.

Lemma 4.1 *For every $x \in X$,*

$$\sum_{t=1}^{\mu} I(\Pi^t, \bar{R}_t) \geq -\ln n + \ln 2 \sum_{t=1}^{\mu} I(x, \bar{R}_t).$$

Proof Let $\pi^t(X)$ be the total weight of X in the beginning of round t . Then by construction,

$$\begin{aligned} \pi^{t+1}(X) &= \sum_{x \in X} \pi^t(x)(1 + I(x, \bar{R}_t)) \\ &= \pi^t(X) \left(1 + \sum_{x \in X} \frac{\pi^t(x)}{\pi^t(X)} I(x, \bar{R}_t) \right) \end{aligned}$$

$$\begin{aligned}
 &= \pi^t(X)(1 + I(\Pi^t, \bar{R}_t)) \leq \pi^t(X) \exp(I(\Pi^t, \bar{R}_t)) \\
 &\leq n \exp\left(\sum_{i=1}^t I(\Pi^i, \bar{R}_i)\right). \tag{2}
 \end{aligned}$$

The last inequality follows because $\pi^1(X) = n$. However, for every $x \in X$,

$$\pi^{\mu+1}(X) \geq \pi^{\mu+1}(x) = \exp\left(\ln 2 \sum_{t=1}^{\mu} I(x, \bar{R}_t)\right). \tag{3}$$

The lemma follows from (2) and (3). □

Since Lemma 4.1 holds for every x , it also holds for the optimal mixed strategy Π^* .

Corollary 4.2

$$\sum_{t=1}^{\mu} I(\Pi^t, \bar{R}_t) \geq -\ln n + \ln 2 \sum_{t=1}^{\mu} I(\Pi^*, \bar{R}_t).$$

A similar argument proves the following lemma, which states that the expected cost to Bob is not much worse than that by his best pure strategy.

Lemma 4.3 *For every $R \in \mathcal{R}$,*

$$\sum_{t=1}^{\mu} I(\bar{x}_t, \Omega^t) \leq 2 \ln m + \ln 4 \sum_{t=1}^{\mu} I(\bar{x}_t, R).$$

Proof Let $\omega^t(\mathcal{R})$ be the total weight of \mathcal{R} in the beginning of round t . Then by construction,

$$\begin{aligned}
 \omega^{t+1}(\mathcal{R}) &= \sum_{R \in \mathcal{R}} \omega^t(R) \left(1 - \frac{1}{2} I(\bar{x}_t, R)\right) \\
 &= \omega^t(\mathcal{R}) \left(1 - \frac{1}{2} \sum_{R \in \mathcal{R}} \frac{\omega^t(R)}{\omega^t(\mathcal{R})} I(\bar{x}_t, R)\right) \\
 &= \omega^t(\mathcal{R}) \left(1 - \frac{1}{2} I(\bar{x}_t, \Omega^t)\right) \leq \omega^t(\mathcal{R}) \exp\left(-\frac{1}{2} I(\bar{x}_t, \Omega^t)\right) \\
 &\leq m \exp\left(-\frac{1}{2} \sum_{i=1}^t I(\bar{x}_i, \Omega^i)\right). \tag{4}
 \end{aligned}$$

The last inequality follows because $\omega^1(\mathcal{R}) = m$. However, for every $R \in \mathcal{R}$,

$$\omega^{\mu+1}(\mathcal{R}) \geq \omega^{\mu+1}(R) = \exp\left(-\ln 2 \sum_{t=1}^{\mu} I(\bar{x}_t, R)\right). \tag{5}$$

The lemma follows from (4) and (5). □

The following lemma follows from the fact that the game Alice and Bob play is a zero-sum game, but we sketch a proof for completeness.

Lemma 4.4

$$\mathbb{E}\left[\sum_{t=1}^{\mu} I(\Pi^t, \bar{R}_t)\right] = \mathbb{E}\left[\sum_{t=1}^{\mu} I(\bar{x}_t, \Omega^t)\right],$$

where the expectation is taken over the random sequence of points and ranges chosen in all rounds.

Proof For $t \geq 1$, let S^t be the set of all point-range-pair sequences of length t , i.e.,

$$S^t = \{(x_1, R_1), \dots, (x_t, R_t) \mid x_i \in X, R_i \in \mathcal{R}, 1 \leq i \leq t\}.$$

If we fix a sequence $S \in S^{t-1}$, then the distributions Π^t, Ω^t are fixed, which we denote by Π^t_S, Ω^t_S .

$$\begin{aligned} \mathbb{E}[I(\Pi^t, \bar{R}_t)] &= \sum_{S \in S^{t-1}} \Pr(S) \sum_{R_t \in \mathcal{R}} \Pr(R_t|S) I(\Pi^t_S, R_t) \\ &= \sum_{S \in S^{t-1}} \Pr(S) \sum_{\substack{x_t \in X \\ R_t \in \mathcal{R}}} \Pr(x_t, R_t|S) I(x_t, R_t) \\ &= \sum_{S \in S^{t-1}} \Pr(S) \sum_{x_t \in X} \Pr(x_t|S) I(x_t, \Omega^t_S) \\ &= \mathbb{E}[I(x_t, \Omega^t)]. \end{aligned}$$

The lemma now follows from linearity of expectation. □

We are now ready to prove that the distribution $\tilde{\Pi}$ on X implied by \tilde{X} is close to Π^* . For a range $R \in \mathcal{R}$, let

$$\lambda_R = \frac{1}{\mu} \sum_{t=1}^{\mu} I(\bar{x}_t, R) = I(\tilde{\Pi}, R)$$

and

$$\tilde{\lambda} = \min_{R \in \mathcal{R}} \lambda_R.$$

Note that λ_R , for every $R \in \mathcal{R}$, and $\tilde{\lambda}$ are random variables.

Lemma 4.5

- (i) $\mathbb{E}[\tilde{\lambda}] \geq \frac{\lambda^*}{4} \geq \frac{1}{4k}$, and
(ii) $\Pr[\tilde{\lambda} > \mathbb{E}[\tilde{\lambda}]/2] \geq \frac{1}{7}$.

Proof Using Lemmas 4.3 and 4.4 and Corollary 4.2,

$$\begin{aligned} \mathbb{E}[\tilde{\lambda}] &\geq \frac{1}{\mu \ln 4} \mathbb{E} \left[\sum_{t=1}^{\mu} I(\bar{x}_t, \bar{\Omega}^t) \right] - \frac{\ln m}{\mu \ln 2} \quad (\text{By Lemma 4.3}) \\ &= \frac{1}{\mu \ln 4} \mathbb{E} \left[\sum_{t=1}^{\mu} I(\Pi^t, \bar{R}_t) \right] - \frac{\ln m}{\mu \ln 2} \quad (\text{By Lemma 4.4}) \\ &\geq \frac{1}{\mu \ln 4} \mathbb{E} \left[\ln 2 \sum_{t=1}^{\mu} I(\Pi^*, \bar{R}_t) - \ln n \right] - \frac{\ln m}{\mu \ln 2} \\ &\quad (\text{By Corollary 4.2}) \\ &= \frac{1}{2\mu} \mathbb{E} \left[\sum_{t=1}^{\mu} I(\Pi^*, \bar{R}_t) \right] - \frac{1}{\mu \ln 4} \ln(m^2 n). \end{aligned}$$

However,

$$I(\Pi^*, \bar{R}_t) \geq \min_{R \in \mathcal{R}} I(\Pi^*, R) = \lambda^*$$

and

$$\mu = \frac{2}{\ln 2} k \ln(m^2 n).$$

Therefore,

$$\mathbb{E}[\tilde{\lambda}] \geq \frac{\lambda^*}{2} - \frac{1}{4k} \geq \frac{1}{4k}$$

because $\lambda^* \geq 1/k$. This proves part (i) of the lemma.

We now prove (ii). By (1),

$$\tilde{\lambda} = \min_{R \in \mathcal{R}} \lambda_R = \min_{R \in \mathcal{R}} I(\tilde{\Pi}, R) \leq \lambda^*.$$

Also, according to part (i),

$$\mathbb{E}[\tilde{\lambda}] \geq \frac{\lambda^*}{4}.$$

Let $p = \Pr[\tilde{\lambda} \leq \frac{1}{2} E[\tilde{\lambda}]]$. Then

$$E[\tilde{\lambda}] \leq \frac{p}{2} E[\tilde{\lambda}] + \lambda^*(1 - p) \leq \frac{p}{2} E[\tilde{\lambda}] + 4E[\tilde{\lambda}](1 - p),$$

and we obtain $p \leq 6/7$. This proves part (ii) of the lemma. □

Lemma 4.5 implies that, with constant probability, $\tilde{\lambda} > \frac{1}{8k}$, i.e., every range is $\frac{1}{8k}$ -heavy with respect to $\tilde{\Pi}$. Thus, the algorithm indeed returns a hitting set of Σ .

Fast implementation and running time We perform two preprocessing steps to expedite the algorithm. In the first preprocessing step, as in Sect. 3, we compute a $\frac{1}{k}$ -net Q_0 of Σ , set $X = X \setminus Q_0$ and $\mathcal{R} = \{R \mid R \cap Q_0 = \emptyset\}$, i.e., remove the ranges hit by Q_0 . In the second preprocessing step, we choose a set $Q_1 \subseteq X$ of $O(k)$ points so that any point in $X \setminus Q_1$ intersects at most m/k ranges of $\{R \mid R \cap Q_1 = \emptyset\}$. Since Σ has a hitting set of size at most k , such a set Q_1 always exists. We now set $X = X \setminus Q_1$, $\mathcal{R} = \{R \mid R \cap Q_1 = \emptyset\}$. After the preprocessing, each point of X lies in at most m/k ranges, and each range of \mathcal{R} contains at most n/k points. The algorithm returns $Q_0 \cup Q_1 \cup Q$ as the hitting set of Σ .

For this algorithm, we need the net finder of (P1), the range-reporting data structure of (P3) and two other procedures:

- (P4) A dual range-reporting data structure that reports all s ranges containing a query point in $O(\tau(m) + s)$ time, and that can be constructed in $O(m\tau(m))$ time.
- (P5) An algorithm that performs the second preprocessing step in $\varphi(N)$ time.

Finally, we build a balanced binary search tree \mathcal{T}_X on X and another one $\mathcal{T}_{\mathcal{R}}$ on \mathcal{R} so that the weights of a point in X and of a range in \mathcal{R} can be updated in $O(\log n)$ and $O(\log m)$ time, respectively, and so that a random element of X or \mathcal{R} can be chosen within the same time; see e.g. [3].

In round i , we report all points of \bar{R}_i in $O(\tau(n) + n/k)$ time using (P3) and update the weights of these points in the binary tree \mathcal{T}_X in $O(\frac{n}{k} \log n)$ time. Similarly we report the ranges of \mathcal{R} that contain \bar{x}_i in $O(\tau(m) + m/k)$ time using (P4) and update their weights in $\mathcal{T}_{\mathcal{R}}$ in $O(\frac{m}{k} \log m)$ time. Hence, round i takes $O(\tau(N) + k^{-1}N \log(N))$ time. Summing over all μ rounds, adding the preprocessing time, the algorithm takes $O(\varphi(N) + k \log(N)\tau(N) + N \log^2(N)) = O(\varphi(N) + N \log(N)\tau(N))$ time. Since the algorithm succeeds with probability at least $1/7$, it is repeated $O(1)$ expected times.

We perform backward exponential search on the value of k as follows. Initially, we set $k = n$. Given the current value of k , we run the algorithm for $\Theta(\log N)$ times. If the algorithm returns a hitting set H , we remember H , halve the value of k , and continue. Otherwise, we stop and return the smallest hitting set remembered. At most $\log n$ values of k are examined during the backward exponential search. If $k \geq \kappa$, the expected number of times the algorithm is run for a fixed value of k is $O(1)$ because each iteration returns a valid hitting set with constant probability. If $k < c\kappa/g(\kappa)$ for some constant c , no hitting set of size $O(kg(k)) < \kappa$ exists, and thus the algorithm will terminate in $\Theta(\log N)$ iterations. For $k \in [c\kappa/g(\kappa), \kappa)$, a hitting set may be found

at the end of the $\Theta(\log N)$ iterations in the worst case. Thus, the algorithm is repeated $O(\log(N) \log(g(\kappa)))$ times in the worst case, for this range of k .

Unlike Sect. 3, the weight of points can become quite large and those of ranges can become quite small, so algebraic complexity of the algorithm can be large if we store all weights explicitly. One possible method is to scale down all the weights by a factor of 2 periodically and set the minimum weight to 1, i.e., remove the points whose weights become less than 1. If we keep the ratio of the maximum and minimum weights polynomial (e.g., N^4), the approximation of the small weights will have negligible impact to the correctness analysis. We conclude the following.

Theorem 4.6 *Let $\Sigma = (X, \mathcal{R})$ be a finite range space with $|X| + |\mathcal{R}| = N$.*

- (i) *Suppose procedures (P1) and (P3)–(P5) exist for Σ , then an $O(g(\kappa))$ -approximate hitting set of Σ can be computed in expected time $O((\varphi(N) + N\tau(N) \log N) \log N \log(g(\kappa)))$, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*
- (ii) *Suppose procedures (P1) and (P3)–(P5) exist for Σ^\perp , then an $O(g(\chi))$ -approximate set cover of Σ can be computed in expected time $O((\varphi(N) + N\tau(N) \log N) \log N \log(g(\chi)))$, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

Remarks (i) The algorithm described here can be viewed as a “kinder and gentler” version of the greedy algorithm. A point lying in many “uncovered” ranges is likely to have higher weight and thus higher probability of being chosen. Instead of removing the ranges containing a chosen point, we simply halve their weights. By the time algorithm stops, the weight of every range is very small.

(ii) It is worth contrasting this algorithm with the previous one. This one is simpler and does not require a dynamic data structure for range counting to verify whether a range is light. Also, this algorithm is slightly faster for some geometric instances with one less logarithmic factor in the running time. However, it requires a range-reporting data structure for both Σ and Σ^\perp , and an additional preprocessing step to compute Q_1 so that no point lies in more than m/k ranges.

(iii) This algorithm falls under the primal-dual framework for solving linear programs (LPs). A very similar approach was used by Koufogiannakis and Young [31] to give near-linear-time PTAS for fractional packing and covering LPs. Their running time depends on the number of non-zero elements in the constraint matrix, which can be quadratic for geometric covering problems. We achieve near-linear running time using geometric data structures and careful preprocessing.

5 Fast Algorithms for Geometric Instances

In this section, we show that the algorithms described in Sects. 3 and 4 yield near-linear hitting-set and set-cover algorithms for a number of geometric range spaces, for which the required procedures (P1)–(P5) can be performed efficiently. Let X be a set of n points in \mathbb{R}^d and \mathcal{R} a set of m geometric shapes such as rectangles, balls, and simplices. As mentioned in Introduction, we will use $R \in \mathcal{R}$ to denote a shape as well as the subset $X \cap R$. It will be clear from the context which of the two we are referring to.

We mainly describe the implementations of the second algorithm, which require procedures (P1) and (P3)–(P5) and give the desired approximation ratios with high probability. One can also use the first algorithm, which may add a logarithmic factor in the running time in some cases, but it always returns a hitting set or set cover with the approximation ratio stated in the corresponding theorems.

Rectangles in 2D and 3D Let \mathcal{R} be a set of m axis-parallel rectangles in \mathbb{R}^d for $d = 2, 3$. Aronov et al. [6] have shown that an ε -net of Σ of size $O((1/\varepsilon) \log \log(1/\varepsilon))$ can be constructed in $O(n \log^{d-1} n)$ randomized expected time, thus $g(\kappa) = \log \log(\kappa)$ and $\varphi(N) = O(N \log^{d-1} N)$ in procedure (P1).

For the range reporting data structure in (P3), we construct a d -dimensional range tree on X in $O(n \log^{d-1} n)$ time and a range reporting query can be answered in $O(\log^{d-1} n + s)$ time, where s is the output size. For (P4), we need a dual range-reporting data structure that reports all rectangles of \mathcal{R} containing a query point. This can be done in $O(\log^{d-1} m + s)$ time, where s is the output size, after $O(m \log^{d-1} m)$ preprocessing [15]. Hence, $\tau(N) = O(\log^{d-1} N)$.

Finally, we also need an algorithm for (P5) that computes a set Q_1 of $O(k)$ points so that no point in $X \setminus Q_1$ lies in more than m/k rectangles of $\{R \in \mathcal{R} \mid R \cap Q_1 = \emptyset\}$. This can be accomplished using a d -dimensional segment tree on \mathcal{R} , which can be constructed in $O(m \log^d m)$ time. For each point $p \in X$, the segment tree is queried in $O(\log^d m)$ time to check whether p lies in more than m/k rectangles of \mathcal{R} . If so, we add p to Q_1 and delete all the rectangles of \mathcal{R} that contain p from the segment tree. Obviously, at most k points are reported, and the time spent is $\varphi(N) = O(N \log^d N)$.³

Plugging the bounds of $\tau(\cdot)$, $g(\cdot)$ and $\varphi(\cdot)$ in Theorem 4.6, we obtain the following.

Theorem 5.1 *Let X be a set of n points in \mathbb{R}^d and \mathcal{R} a set of m rectangles in \mathbb{R}^d , for $d = 2, 3$, with $|X| + |\mathcal{R}| = N$. An $O(\log \log \kappa)$ -approximate hitting set of (X, \mathcal{R}) can be computed in $O(N \log^{d+1} N \log \log \log \kappa)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

For the set-cover problem, an ε -net of Σ^\perp required in (P1) can be computed simply by choosing a random sample of size $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$, with probability at least $1/2$ [30]. Pach and Tardos [40] showed that the size of ε -net for Σ^\perp , even for $d = 2$, is $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ in the worst case. The procedures (P3) and (P4) remain the same as (P4) and (P3) in the hitting-set case. The set $Q_1 \subseteq \mathcal{R}$ in procedure (P5) can be computed in a similar way, except that a d -dimensional range tree, instead of a segment tree, is used. The running time is $\varphi(N) = O(N \log^d N)$. Hence, we obtain the following.

Theorem 5.2 *Let X be a set of points in \mathbb{R}^d and \mathcal{R} a set of rectangles in \mathbb{R}^d , for $d = 2, 3$, with $|X| + |\mathcal{R}| = N$. An $O(\log \chi)$ -approximate set cover of (X, \mathcal{R}) can be computed in $O(N \log^{d+1} N \log \log \chi)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

³ Using a sweep-plane technique, the running time of the algorithm for (P5) can be improved by a logarithmic factor. However, since this step is not the bottleneck, we use the above approach, which is simpler.

Halfspaces in 3D and disks in 2D Let X be a set of n points and \mathcal{R} a set of m halfspaces in \mathbb{R}^3 . Building on the algorithm by Matoušek [36], Chan and Tsakalidis [14] have described an $O(n \log \frac{1}{\varepsilon})$ -time deterministic algorithm to compute an ε -net of size $O(1/\varepsilon)$ for the range space $\Sigma = (X, \mathcal{R})$. Therefore $g(\cdot) = 1$ and $\varphi(N) = O(N \log N)$ in (P1).

Using the data structure by Afshani and Chan [1], the halfspace range-reporting query in (P3) can be answered in $O(\log n + s)$ time, after $O(n \log n)$ expected-time preprocessing. By duality transform, the dual range-reporting query in (P4) can be answered using the same data structure in $O(\log m + s)$ time, after $O(m \log m)$ preprocessing. Hence $\tau(n) = O(\log n)$.

We describe an algorithm for computing Q_1 in the preprocessing step of the second algorithm. The algorithm is based on constructing *shallow cuttings* of a set of hyperplanes, defined by Matoušek [36]. Given parameters $\ell, r \in [1, n]$, an ℓ -shallow $\frac{1}{r}$ -cutting of a collection \mathcal{H} of n hyperplanes in \mathbb{R}^d is a set of interior-disjoint simplices such that the interior of every simplex intersects at most $\frac{n}{r}$ hyperplanes of \mathcal{H} , and the union of the simplices cover all points that lie above ℓ or fewer hyperplanes in \mathcal{H} .

The algorithm works in $\lceil \log k \rceil$ rounds.⁴ At the beginning of round i , we have a set $X_i \subseteq X$ of points and a subset $\mathcal{R}_i \subseteq \mathcal{R}$ of halfspaces. Initially, in round 0, $X_i = X$ and $\mathcal{R}_i = \mathcal{R}$. The maximum depth of a point in X_i is at most $\ell_i = |\mathcal{R}_{i-1}|/2^i$. We set $r_i = 2^{i+1}$ and compute a ℓ_i -shallow $(1/r_i)$ -cutting Ξ_i of the halfspaces of \mathcal{R}_i ; $|\Xi_i| = O(2^i)$. Since the maximum depth of a point of X_i is ℓ_i , each point of X_i lies in a simplex of Ξ_i . For each simplex $\Delta \in \Xi_i$, if $X_i \cap \Delta \neq \emptyset$, we choose one point from Δ and add it to Q_1 . Let P_i be the set of points chosen in round i , and let $\overline{\mathcal{R}}_i = \{R \mid R \cap P_i \neq \emptyset\}$. We set $X_{i+1} = X_i \setminus P_i$ and $\mathcal{R}_{i+1} = \mathcal{R}_i \setminus \overline{\mathcal{R}}_i$. Since each simplex $\Delta \in \Xi_i$ intersects the boundary of at most $|\mathcal{R}_i|/2^{i+1}$ halfspaces, which are the only possible halfspaces remaining in \mathcal{R}_{i+1} that can cover points in Δ , the maximum depth of a point in X_{i+1} is $|\mathcal{R}_i|/2^{i+1} \leq m/2^{i+1}$ in \mathcal{R}_{i+1} .

By construction, $|Q_1| = \sum_i O(2^i) = O(k)$. We now analyze the running time of the above algorithm. Using the algorithm by Chan and Tsakalidis [14], Ξ_i can be computed in $O(|\mathcal{R}_i| \log r_i) = O(m \cdot i)$ time. A property of their algorithm is that each cell in Ξ_i is a semi-bounded vertical prism that is unbounded from below and whose boundary consists of vertical walls and a top triangle. Furthermore, the upper boundary of these prisms forms an xy -monotone triangulated surface. The xy -projection of the upper boundary is a triangulation Ξ_i^* of \mathbb{R}^2 . By preprocessing Ξ_i^* for planar point-location queries (see e.g. [22]), in time $O(r_i \log r_i)$, we can determine in $O(\log r_i) = O(i)$ time whether a query point $q \in \mathbb{R}^3$ lies in a prism of Ξ_i , and also the prism of Ξ_i that contains q if the answer is yes. Thus, P_i can be computed by going through every point $q \in X_i$ and adding q to P_i if it is the first point visited in some simplex of Ξ_i . The time for computing P_i is $O(m \cdot i + n \cdot i)$. Finally, $\overline{\mathcal{R}}_i$ is computed using $|\mathcal{R}_i|$ range-emptiness queries on P_i , each taking $O(\log |\mathcal{R}_i|)$ time. Therefore, summing over all $\lceil \log k \rceil$ rounds, the expected time for computing Q_1 is $O(N \log^2 N)$.

Lemma 5.3 *A set Q_1 of size $O(k)$ required by (P5) can be computed in $O(N \log^2 N)$ time.*

⁴ Recall that k is an integer satisfying $k/2 < \kappa \leq k$.

Hence $\varphi(N) = O(N \log^2 N)$ in (P5). Using the duality transform, the set-cover problem in the primal space becomes the hitting-set problem in the dual space. Plugging the bounds in Theorem 4.6, we obtain the following.

Theorem 5.4 *Let X be a set of points in \mathbb{R}^3 and \mathcal{R} a set of halfspaces in \mathbb{R}^3 , with $|X| + |\mathcal{R}| = N$. An $O(1)$ -approximate hitting set or set cover of (X, \mathcal{R}) can be computed in $O(N \log^3 N)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

By a standard lifting transform [22], we have the following result for disks in \mathbb{R}^2 .

Corollary 5.5 *Let X be a set of points in \mathbb{R}^2 and \mathcal{R} a set of disks in \mathbb{R}^2 , with $|X| + |\mathcal{R}| = N$. An $O(1)$ -approximate hitting set or set cover of (X, \mathcal{R}) can be computed in $O(N \log^3 N)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

Fat triangles Let $X \subseteq \mathbb{R}^2$ be a set of n points and \mathcal{R} a set of α -fat triangles in \mathbb{R}^2 , for some constant $\alpha \geq 1$, i.e., the aspect ratio of every triangle in \mathcal{R} is at most α . It is well known that $\Sigma = (X, \mathcal{R})$ has constant VC-dimension and by the ε -net theorem [30], an ε -net of Σ in procedure (P1) can be obtained by randomly sampling $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ points from X . A range-reporting query (P3) for Σ can be answered in $O(\log^3 n + s)$ time, after $O(n \log^3 n)$ preprocessing [43, Sect. 4.5], and a dual range-reporting query (P4) for Σ can also be answered in $O(\log^3 m + s)$ time [2]. The set Q_1 in (P5) for Σ can be computed using the same method as that for 3D halfspaces described in Lemma 5.3. We use the results in [3, Lems. 2.1 and 2.2] for computing the shallow cutting and point location data structure, and the running time is $O(N \log^3 N \log^* N)$ using the improved bound on the union of fat triangles [5]. Thus, $g(1/\varepsilon) = \log(1/\varepsilon)$, $\tau(N) = O(\log^3 N)$, $\varphi(N) = O(N \log^3 N \log^* N)$, and we obtain the following.

Theorem 5.6 *Let X be a set of points in \mathbb{R}^2 and \mathcal{R} a set of α -fat triangles, with $|X| + |\mathcal{R}| = N$. An $O(\log \kappa)$ -approximate hitting set of (X, \mathcal{R}) can be computed in $O(N \log^5 N \log \log \kappa)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

For the set cover problem, an ε -net is computed by randomly sampling $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ triangles from \mathcal{R} , since the dual range space Σ^\perp also has constant VC-dimension.

Next, we describe an algorithm for computing $Q_1 \subseteq \mathcal{R}$ for Σ^\perp , i.e., each range in $\mathcal{R} \setminus Q_1$ contains at most n/k points of $X \setminus \bigcup_{R \in Q_1} R$. Q_1 is constructed in $\lceil \log_2 k \rceil$ steps. At the beginning of step i , we have a set $\mathcal{R}_i \subseteq \mathcal{R}$ of triangles and a set $X_i \subseteq X$ of points such that for any $\Delta \in \mathcal{R}_i$, $|\Delta \cap X_i| \leq \frac{n}{2^{i-1}}$. Initially, for $i = 1$, $X_1 = X$ and $\mathcal{R}_1 = \mathcal{R}$. Step i chooses a subset $T_i \subseteq \mathcal{R}_i$ of $O(2^i)$ triangles such that each triangle of T_i contains $\Omega(\frac{n}{2^i})$ points of X_i , and no triangle of $\mathcal{R}_{i+1} = \mathcal{R}_i \setminus T_i$ contains more than $\frac{n}{2^i}$ points of $X_{i+1} = X_i \setminus \bigcup_{\Delta \in T_i} \Delta$. We construct a $(9/8)$ -approximate range-counting data structure on X_i with respect to triangles in \mathcal{R} . That is, for a triangle Δ , it returns a value μ_Δ such that

$$|X_i \cap \Delta| \leq \mu_\Delta \leq \frac{9}{8} |X_i \cap \Delta|.$$

Such a data structure can be built in $O(n \log^3 n)$ time by applying the range-emptiness data structure [43] as a black-box to the data structure in [7, Sect. 5], and can answer a query in $O(\log^4 n)$ time. We also preprocess X_i for range-reporting queries with respect to \mathcal{R} — the data structure can be built in $O(n \log^3 n)$ time and a query can be answered in $O(\log^3 n + t)$ time where t is the output size.

We also maintain a set $Y_i \subset X_i$ in step i . Initially $Y_i = \emptyset$, and points will be added to Y_i as step i progresses. We maintain a (9/8)-approximate range counting data structure on Y_i with respect to \mathcal{R} . Using the so-called logarithmic method [9] on the static range-counting data structure mentioned above, a point can be inserted in $O(\log^4 n)$ time and a query can be answered in $O(\log^5 n)$ time.

We are now ready to describe step i . We process the triangles of \mathcal{R}_i one by one. For each $\Delta \in \mathcal{R}_i$, we query the two range-counting data structures with Δ . Suppose they return μ_Δ and ν_Δ as (9/8)-approximate values of $|\Delta \cap X_i|$ and $|\Delta \cap Y_i|$, respectively. If $\mu_\Delta - \nu_\Delta \geq \frac{n}{2^{i+1}}$, we add Δ to T_i , report $\Delta \cap X_i$, and insert each point of $\Delta \cap X_i$ to Y_i . (Note that a point p may be inserted into Y_i multiple times; we keep only one copy of p in Y_i .)

After $\lceil \log_2 k \rceil$ steps, we return $\bigcup T_i$ as Q_1 . We now prove the correctness of the procedure and analyze its running time.

Fix a triangle $\Delta \in \mathcal{R}_i$. In what follows, we refer to Y_i as the set just before Δ was processed. We observe that $\mu_\Delta - \nu_\Delta \geq \frac{n}{2^{i+1}}$ implies that

$$|\Delta \cap (X_i \setminus Y_i)| \geq \frac{n}{2^{i+1}} - \frac{1}{8} |\Delta \cap X_i|.$$

By the invariant maintained by the algorithm, $|\Delta \cap X_i| \leq \frac{n}{2^{i-1}}$. Therefore

$$|\Delta \cap (X_i \setminus Y_i)| \geq \frac{n}{2^{i+1}} - \frac{n}{2^{i+2}} = \frac{n}{2^{i+2}}.$$

Hence Δ contains at least $\frac{n}{2^{i+2}}$ “new” points of X_i , i.e., the ones that are not already in Y_i . Since all points of $\Delta \cap X_i$ are inserted into Y_i and a point is never deleted from Y_i , we can conclude that $|T_i| = O(2^i)$. Furthermore,

$$|\Delta \cap X_i| \leq \frac{n}{2^{i-1}} \leq 8|\Delta \cap (X_i \setminus Y_i)|.$$

Summing over all triangles of T_i , $\sum_{\Delta \in T_i} |\Delta \cap X_i| = O(n)$. The total time spent in reporting $\Delta \cap X_i$ and adding these points to the data structure is thus $O(n \log^4 n)$. We spend $O(\log^5 n)$ time in computing μ_Δ and ν_Δ . Summing over all triangles of \mathcal{R}_i , and summing over all $i \leq \lceil \log_2 k \rceil$, the total time spent in computing Q_1 is $O(N \log^6 N)$.

Finally, if $\Delta \notin T_i$, i.e., $\mu_\Delta - \nu_\Delta \leq \frac{n}{2^{i+1}}$, then

$$|\Delta \cap X_{i+1}| \leq |\Delta \cap (X_i \setminus Y_i)| \leq \frac{n}{2^{i+1}} + \frac{1}{8} |Y_i \cap \Delta| \leq \frac{n}{2^{i+1}} + \frac{1}{8} \cdot \frac{n}{2^{i-1}} < \frac{n}{2^i}.$$

Therefore, $|\Delta \cap X_{i+1}| \leq \frac{n}{2^i}$, for all $\Delta \in \mathcal{R}_{i+1}$, as desired.

Hence, after $u = \lceil \log_2 k \rceil$ steps,

$$|\Delta \cap X_{u+1}| \leq \frac{n}{2^u} \leq \frac{n}{k}.$$

Putting everything together, we conclude the following:

Theorem 5.7 *Let X be a set of points in \mathbb{R}^2 and \mathcal{R} a set of α -fat triangles, with $|X| + |\mathcal{R}| = N$. An $O(\log \chi)$ -approximate set cover of (X, \mathcal{R}) can be computed in $O(N \log^7 N \log \log \chi)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

Remarks Better bounds on the sizes of primal and dual ε -nets exist for fat triangles [6]. We use the simple random-sampling approach to compute an ε -net since it can be done in near-linear time. With a careful implementation, it might be possible to implement the ε -net construction algorithms in [6] in near-linear time, in which case the approximation ratios of the hitting-set and set-cover problems for fat triangles can be improved.

6 Independent Set for Disks

Given a finite range space $\Sigma = (X, \mathcal{R})$, an *independent set* is a subset $\mathcal{J} \subseteq \mathcal{R}$ of ranges such that for any $R_1, R_2 \in \mathcal{J}$, $R_1 \cap R_2 \cap X = \emptyset$, i.e., any point of X lies in at most one range of \mathcal{J} . The goal is to find a maximum-size independent set (MIS).

The main observation is that, an optimal mixed strategy for Bob in the second algorithm gives a fractional solution for the independent-set problem. So we proceed as follows. Suppose we have an integer k such that $\delta/2 < k \leq \delta$, where δ is the size of an optimal independent set. We set $\mu = 3k \log N$ and run the second algorithm for μ rounds. Let $\mathcal{F} = \langle \bar{R}_1, \dots, \bar{R}_\mu \rangle$ be the sequence of ranges chosen by Bob. Let $\tilde{\Omega}$ be the distribution of \mathcal{R} induced by \mathcal{F} , i.e., $\Pr(R \sim \tilde{\Omega})$ is proportional to the number of times R appears in \mathcal{F} . Using Lemmas 4.1, 4.3 and 4.4, we can prove that

$$E \left[\max_{x \in X} I(x, \tilde{\Omega}) \right] \leq \frac{3}{k}.$$

By the Markov inequality, with probability at least $1/2$,

$$\max_{x \in X} I(x, \tilde{\Omega}) \leq \frac{6}{k}.$$

Let y_i be the variable corresponding to the range $R_i \in \mathcal{R}$. Then by setting $y_i = \frac{k}{6} \Pr(R_i \sim \tilde{\Omega})$, with probability at least $1/2$, we obtain a fractional solution for the independent-set problem with objective value $\sum y_i = k/6$. Next, we convert the fractional solution into an integral solution using the approach of Chan and Har-Peled [13]: we process the ranges of \mathcal{R} in an arbitrary order, and maintain an independent set \mathcal{J} of ranges. A range R_i is processed as follows: if R_i does not contain any of the points lying in a range of the current \mathcal{J} , then R_i is added to \mathcal{J} with probability y_i . Otherwise, it is discarded.

Chan and Har-Peled [13] have shown that if \mathcal{R} is a set of disks in \mathcal{R}^2 and the fractional solution has value ν , then the above rounding scheme returns an independent set of size $\Omega(\nu)$ with constant probability. Hence, the set \mathcal{J} has size $\Omega(k)$ with constant probability.

We perform the same preprocessing as in Sect. 4 to compute a set Q of at most $2k/5$ points such that in the range space $\bar{\Sigma} = (X \setminus Q, \{R \in \mathcal{R} \mid R \cap Q = \emptyset\})$, each

range of $\bar{\Sigma}$ contains $O(n/k)$ points and each point of $X \setminus Q$ lies in $O(m/k)$ ranges. Furthermore, one can argue that if Σ has an independent set of size δ , then $\bar{\Sigma}$ has one of size at least $3\delta/5$. Hence, we run the above algorithm on $\bar{\Sigma}$, and the total time for computing the fractional solution, including preprocessing, is $O(N \log^3 N)$.

To convert the fractional solution into an integral one, a dynamic insertion-only disk-emptiness data structure \mathcal{D} is used, which can be implemented by using a static disk-emptiness data structure [2] with the logarithmic method [9]. Initially, \mathcal{D} contains no point. For each range R_i , we perform a disk-emptiness query in \mathcal{D} in $O(\log^2 n)$ time to determine whether R_i can be added to \mathcal{J} ; if it is added, every point in R_i is inserted into \mathcal{D} in $O(\log^2 n)$ amortized time. The total running time is $O(N \log^2 N)$. Hence, we obtain the following.

Theorem 6.1 *Let X be a set of points in \mathbb{R}^2 and \mathcal{R} a set of disks in \mathbb{R}^2 , with $|X|+|\mathcal{R}| = N$. An $O(1)$ -approximate independent set of (X, \mathcal{R}) can be computed in $O(N \log^3 N)$ expected time, with probability at least $1 - \frac{1}{N^{\Omega(1)}}$.*

7 Conclusion

In this paper, we presented two simple, efficient algorithms, based on the MW method, to compute small-size hitting sets and set covers for range spaces with finite VC-dimension. The first algorithm is Las Vegas and the second one is Monte Carlo. They yield near-linear time algorithms for many geometric range spaces. Our second algorithm can be used to compute an $O(1)$ -approximate solution for the (discrete) maximum independent set problem for disks in near-linear time.

The two algorithms require slightly different data structures, so depending on the application, one algorithm can be more efficient than the other. One advantage of the second algorithm is its simplicity. In distributed settings where two parties want to collaboratively compute nearoptimal hitting-set (or more generally, solve a zero-sum game), the second algorithm only requires the two parties to exchange the randomly sampled point/range in each iteration. Suppose each point/range can be represented by $O(1)$ bits, then the second algorithm needs $O(\text{OPT} \log N)$ bits of communication in total. On the other hand, using the first algorithm, one party eventually needs to send all the ranges to the other party to determine whether all ranges are ε -heavy, and thus $\Omega(N)$ bits of communication is required. We also remark that both algorithms can achieve $(1 + \varepsilon)$ approximations in the fractional solution by updating the weights by a factor of $(1 + \varepsilon)$ or $(1 - \varepsilon)$ in each step (see e.g. [31]). However, since the final step of rounding using epsilon net already loses a constant factor, we chose to double or halve the weights in each step.

We conclude by mentioning several open problems:

- (i) Can our approach be extended to compute a multi-cover of range spaces [17]?
- (ii) Is there an $O(\log \log \chi)$ -approximate algorithm for computing the set cover when \mathcal{R} is a set of rectangles in \mathbb{R}^2 ?
- (iii) Let X be a set of points and \mathcal{R} a set of unit squares in \mathbb{R}^2 . Can a small-size hitting set of \mathcal{R} be maintained under insertion/deletion of points and squares?

- (iv) Is there a fast $O(1)$ -approximate algorithm for the independent set problem when \mathcal{R} is a set of rectangles in \mathbb{R}^2 . The best known algorithm computes an $O(\log \log n)$ -approximate independent set [12].

Acknowledgements The authors thank Kamesh Munagala for useful discussions and the two reviewers for their helpful comments.

References

1. Afshani, P., Chan, T.M.: Optimal halfspace range reporting in three dimensions. In: Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09), pp. 180–186. SIAM, Philadelphia (2009)
2. Agarwal, P.K., Erickson, J.: Geometric range searching and its relatives. In: Chazelle, B., Goodman, J., Pollack, R. (eds.) *Advances in Discrete and Computational Geometry*. Contemporary Mathematics, vol. 223, pp. 1–56. American Mathematical Society, Providence (1999)
3. Agarwal, P.K., Ezra, E., Sharir, M.: Near-linear approximation algorithms for geometric hitting sets. *Algorithmica* **63**(1–2), 1–25 (2012)
4. Alon, N.: A non-linear lower bound for planar epsilon-nets. *Discrete Comput. Geom.* **47**(2), 235–244 (2012)
5. Aronov, B., de Berg, M., Ezra, E., Sharir, M.: Improved bounds for the union of locally fat objects in the plane. *SIAM J. Comput.* **43**(2), 543–572 (2014)
6. Aronov, B., Ezra, E., Sharir, M.: Small-size ε -nets for axis-parallel rectangles and boxes. *SIAM J. Comput.* **39**(7), 3248–3282 (2010)
7. Aronov, B., Har-Peled, S.: On approximating the depth and related problems. *SIAM J. Comput.* **38**(3), 899–921 (2008)
8. Arora, S., Hazan, E., Kale, S.: The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.* **8**(6), 121–164 (2012)
9. Bentley, J.L., Saxe, J.B.: Decomposable searching problems. I. Static-to-dynamic transformation. *J. Algorithms* **1**(4), 301–358 (1980)
10. Brönnimann, H., Goodrich, M.T.: Almost optimal set covers in finite VC-dimension. *Discrete Comput. Geom.* **14**(1), 463–479 (1995)
11. Bus, N., Garg, S., Mustafa, N.H., Ray, S.: Tighter estimates for ε -nets for disks. *Comput. Geom. Theory Appl.* **53**(C), 27–35 (2016)
12. Chalmersook, P., Chuzhoy, J.: Maximum independent set of rectangles. In: Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09), pp. 892–901. SIAM, Philadelphia (2009)
13. Chan, T.M., Har-Peled, S.: Approximation algorithms for maximum independent set of pseudo-disks. In: Proceedings of the 25th Annual Symposium on Computational Geometry (SoCG'09), pp. 333–340. ACM, New York (2009)
14. Chan, T.M., Tsakalidis, K.: Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. In: Proceedings of the 31th Annual Symposium on Computational Geometry (SoCG'15). LIPIcs. Leibniz International Proceedings in Informatics, vol. 34, pp. 719–732. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2015)
15. Chazelle, B.: Filtering search: a new approach to query answering. *SIAM J. Comput.* **15**(3), 703–724 (1986)
16. Chazelle, B., Welzl, E.: Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.* **4**(5), 467–489 (1989)
17. Chekuri, C., Clarkson, K.L., Har-Peled, S.: On the set multicover problem in geometric settings. *ACM Trans. Algorithms* **9**(1), 9:1–9:17 (2012)
18. Clarkson, K.L.: Algorithms for polytope covering and approximation. In: Proceedings of the 3rd Workshop on Algorithms and Data Structures (WADS'93), pp. 246–252. Springer, Berlin (1993)
19. Clarkson, K.L.: Las Vegas algorithms for linear and integer programming when the dimension is small. *J. ACM* **42**(2), 488–499 (1995)
20. Clarkson, K.L., Varadarajan, K.R.: Improved approximation algorithms for geometric set cover. *Discrete Comput. Geom.* **37**(1), 43–58 (2007)

21. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms, 3rd edn. McGraw-Hill, Boston (2009)
22. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.: Computational Geometry: Algorithms and Applications, 3rd edn. Springer, Berlin (2008)
23. Ezra, E.E.: Small-size relative (p, ε) -approximations for well-behaved range spaces. In: Proceedings of the 29th Annual Symposium on Computational Geometry (SoCG'13), pp. 233–242. ACM, New York (2013)
24. Feige, U.: A threshold of $\ln n$ for approximating set cover. *J. ACM* **45**(4), 634–652 (1998)
25. Fowler, R.J., Paterson, M.S., Tanimoto, S.L.: Optimal packing and covering in the plane are NP-complete. *Inf. Process. Lett.* **12**(3), 133–137 (1981)
26. Freund, Y., Schapire, R.E.: Adaptive game playing using multiplicative weights. *Games Econ. Behav.* **29**(1–2), 79–103 (1999)
27. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, San Francisco (1979)
28. Grigoriadis, M.D., Khachiyan, L.G.: A sublinear-time randomized approximation algorithm for matrix games. *Oper. Res. Lett.* **18**(2), 53–58 (1995)
29. Har-Peled, S.: Geometric Approximation Algorithms. Mathematical Surveys and Monographs, vol. 173. American Mathematical Society, Providence (2011)
30. Haussler, D., Welzl, E.: ε -nets and simplex range queries. *Discrete Comput. Geom.* **2**(2), 127–151 (1987)
31. Koufogiannakis, C., Young, N.E.: A nearly linear-time PTAS for explicit fractional packing and covering linear programs. *Algorithmica* **70**(4), 648–674 (2014)
32. Kupavskii, A., Mustafa, N., Pach, J.: New lower bounds for epsilon-nets. In: Proceedings of the 32nd International Symposium on Computational Geometry. LIPIcs. Leibniz International Proceedings in Informatics, vol. 51, pp. 54:1–54:16. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2016)
33. Lauen, S.: Geometric set cover and hitting sets for polytopes in R^3 . In: Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science. LIPIcs. Leibniz International Proceedings in Informatics, vol. 1, pp. 479–490. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2008)
34. Long, P.M.: Using the pseudo-dimension to analyze approximation algorithms for integer programming. In: Proceedings of the 7th Workshop Algorithms Data Structures. Lecture Notes in Computer Science, vol. 2125, pp. 26–37. Springer, Berlin (2001)
35. Matoušek, J.: Efficient partition trees. *Discrete Comput. Geom.* **8**(3), 315–334 (1992)
36. Matoušek, J.: Reporting points in halfspaces. *Comput. Geom.* **2**(3), 169–186 (1992)
37. Motwani, R., Raghavan, P.: Randomized Algorithms. Cambridge University Press, Cambridge (1995)
38. Mustafa, N.H., Raman, R., Ray, S.: Settling the APX-hardness status for geometric set cover. In: Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science, pp. 541–550. IEEE, Los Alamitos (2014)
39. Mustafa, N.H., Ray, S.: Improved results on geometric hitting set problems. *Discrete Comput. Geom.* **44**(4), 883–895 (2010)
40. Pach, J., Tardos, G.: Tight lower bounds for the size of epsilon-nets. *J. Am. Math. Soc.* **26**(3), 645–658 (2013)
41. Plotkin, S.A., Shmoys, D.B., Tardos, É.: Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.* **20**(2), 257–301 (1995)
42. Pyrga, E., Ray, S.: New existence proofs for ε -nets. In: Proceedings of the 24th Annual Symposium on Computational Geometry (SoCG'08), pp. 199–207. ACM, New York (2008)
43. Shaul, H.: Range Searching: Emptiness, Reporting, and Approximate Counting. PhD thesis, Tel Aviv University (2011)
44. Varadarajan, K.: Epsilon nets and union complexity. In: Proceedings of the 25th Annual Symposium on Computational Geometry (SoCG'09), pp. 11–16. ACM, New York (2009)