

The Average-Case Area of Heilbronn-Type Triangles*

Tao Jiang,^{1,†} Ming Li,^{2,‡} Paul Vitányi^{3,§}

¹Department of Computer Science, University of California, Riverside, CA 92521, USA

²Department of Computer Science, University of California, Santa Barbara, CA 93106, USA

³CWI, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands

Received 1 August 2000; revised 7 May 2001; accepted 13 October 2001

ABSTRACT: From among $\binom{n}{3}$ triangles with vertices chosen from n points in the unit square, let T be the one with the smallest area, and let A be the area of T . Heilbronn's triangle problem asks for the maximum value assumed by A over all choices of n points. We consider the average-case: If the n points are chosen independently and at random (with a uniform distribution), then there exist positive constants c and C such that $c/n^3 < \mu_n < C/n^3$ for all large enough values of n , where μ_n is the expectation of A . Moreover, $c/n^3 < A < C/n^3$, with probability close to one. Our proof uses the incompressibility method based on Kolmogorov complexity; it actually determines the area of the smallest triangle for an arrangement in "general position." © 2002 John Wiley & Sons, Inc. *Random Struct. Alg.*, 00, 1–14, 2002

1. INTRODUCTION

From among $\binom{n}{3}$ triangles with vertices chosen from among n points in the unit circle, let T be the one of least area, and let A be the area of T . Let Δ_n be

Correspondence to: Paul Vitányi; e-mail: paulv@cwi.nl.

*A preliminary version of this work was presented at the 14th Computational complexity Conference held in Atlanta in 1999.

†Supported in part by the NSERC Research Grant OGP0046613 and a CITO grant.

‡Supported in part by the NSERC Research Grant OGP0046506, a CITO grant, and the Steacie Fellowship.

§Partially supported by the European Union through ESPRIT BRA IV NeuroCOLT II Working Group EP 27150.

© 2002 John Wiley & Sons, Inc.

DOI 10.1002/rsa.10024

the maximum assumed by A over all choices of n points. H.A. Heilbronn (1908–1975) asked for the exact value or approximation of Δ_n . The list [1–4, 7, 8, 13, 14, 17–27] is a selection of papers dealing with the problem. Obviously, the value of Δ_n will change only by a small constant factor for every unit area convex shape, and it has become customary to consider the unit square [22]. A brief history is as follows. Heilbronn observed the trivial upper bound $\Delta_n = O(1/n)$ and conjectured that $\Delta_n = O(1/n^2)$, and Erdős proved that this conjecture—if true—would be tight since $\Delta_n = \Omega(1/n^2)$ [18]. The first nontrivial result due to Roth in 1951 established the upper bound $\Delta_n = O(1/(n\sqrt{\log \log n}))$ [18], which was improved in 1972 by Schmidt to $O(1/(n\sqrt{\log n}))$ [23] and in the same year by Roth first to $O(1/n^{1.105\dots})$ [19] and then to $\Delta_n = O(1/n^{1.117\dots})$ [20]. Roth simplified his arguments in 1973 and 1976 [21, 22]. Exact values of Δ_n for $n \leq 15$ were studied in [7, 25–27]. In 1981, Komlós, et al. [13] improved Roth’s upper bound to $O(1/n^{8/7-\epsilon})$, using the simplified arguments of Roth. The really surprising news came in 1982 when the same authors [14] derived a lower bound $\Omega(\log n/n^2)$, narrowly refuting Heilbronn’s original conjecture. Some believe that this lower bound is perhaps the best possible [5, 6]. In 1997, Bertram-Kretzberg et al. [3] gave an algorithm that finds a specific set of n points in the unit square whose Δ_n (as defined above) is $\Omega(\log n/n^2)$ for every fixed n , using a discretization of the problem. In 1999, Barequet [1] derived lower bounds on d -dimensional versions of Heilbronn’s problem where $d > 2$. All of this work concerns the *worst-case* value of the minimal triangle area.

1.1. Results

Here, we consider the *expected* value: If the n points are chosen independently and at random (uniform distribution), then there exist positive constants c and C such that $c/n^3 < \mu_n < C/n^3$ for all large enough n , where μ_n is the expectation of the area A of the smallest triangle formed by any three points. Moreover, with probability close to one, $c/n^3 < A < C/n^3$. This follows directly from corollaries 2 and 4 of Theorems 1 and 2. Our technique is to discretize the problem and show that all Kolmogorov-random arrangements (see below) of n points in the unit square satisfy this range of area of the smallest triangle, where the constants c, C are functions of the “randomness deficiency” of the arrangement—that is, how far the Kolmogorov complexity of the arrangement falls short of the maximum attainable Kolmogorov complexity. A Kolmogorov-random arrangement is a rigorous way to say that the arrangement is in “general position” or “typical”: there are no simple describable properties that can distinguish any such arrangement from another one [15]. As a consequence, every arrangement in which the smallest triangle has area outside this range—smaller or larger—cannot be Kolmogorov random. According to a recent article [16], this result can act as a mathematical guarantee of the efficacy of certain pseudo Monte Carlo methods to determine the fair market value of derivatives (on the stock market)—these methods give a sequence of points satisfying certain pseudo-randomness properties but having less clustering and larger smallest triangles than to be expected from truly random sequences. For its use in geometrical modeling see [1].

1.2. Technique

Our analysis uses the *incompressibility method* based on Kolmogorov complexity. The argument proceeds by using some property to be contradicted to obtain a short

encoding for some object. In the present article, the object concerned is usually an arrangement of n pebbles on a $K \times K$ grid. The Kolmogorov complexity of the object is a lower bound on the length of an encoding of the object. A contradiction arises by the short encoding having length below the Kolmogorov complexity. We have found that thinking in terms of coding is often helpful to solve our problems. Afterwards, there may arise alternative proofs using counting, as in the case of [11], or the probabilistic method with respect to the present result.¹ In some cases [9], no other proof methods seem to work. Thinking in terms of code length and Kolmogorov complexity enabled advances in problems that were open for decades, like for example [9, 11]. Although the technique has been widely used in a plethora of applications, see the survey [15], it is not yet as familiar as the counting method or the probabilistic method. One goal of the present article is to widen acquaintance with it by giving yet another nontrivial example of its application.

2. KOLMOGOROV COMPLEXITY AND THE INCOMPRESSIBILITY METHOD

We give some definitions to establish notation. For introduction, details, and proofs, see [15]. We write *string* to mean a finite binary string. Other finite objects can be encoded into strings in natural ways. The set of strings is denoted by $\{0, 1\}^*$. The *length* of a string x is denoted by $l(x)$, distinguishing it from the *area* $|PQR|$ of a triangle on the points P, Q, R in the plane.

Let $x, y, z \in \mathcal{N}$, where \mathcal{N} denotes the set of natural numbers. Identify \mathcal{N} and $\{0, 1\}^*$ according to the correspondence

$$(0, \epsilon), (1, 0), (2, 1), (3, 00), (4, 01), \dots$$

Here ϵ denotes the *empty word* with no letters. The *length* $l(x)$ of x is the number of bits in the binary string x .

The emphasis is on binary sequences only for convenience; observations in any alphabet can be so encoded in a way that is ‘theory neutral’.

2.1. Self-delimiting Codes

A binary string y is a *proper prefix* of a binary string x if we can write $x = yz$ for $z \neq \epsilon$. A set $\{x, y, \dots\} \subseteq \{0, 1\}^*$ is *prefix-free* if for any pair of distinct elements in the set neither is a proper prefix of the other. A prefix-free set is also called a *prefix code*. Each binary string $x = x_1x_2 \dots x_n$ has a special type of prefix code, called a *self-delimiting code*,

$$\bar{x} = 1^n 0 x_1 x_2 \dots x_n.$$

This code is self-delimiting because we can determine where the code word \bar{x} ends by reading it from left to right without backing up. Using this code, we define

¹ John Tromp has informed us in December 1999 that, following a preliminary version [10] of this work, he has given an alternative proof of the main result based on the probabilistic method.

the standard self-delimiting code for x to be $x' = \overline{l(x)}x$. It is easy to check that $l(\bar{x}) = 2n + 1$ and $l(x') = n + 2 \log n + 1$.

Let $\langle \cdot, \cdot \rangle$ be a standard one-one mapping from $\mathcal{N} \times \mathcal{N}$ to \mathcal{N} , for technical reasons chosen such that $l(\langle x, y \rangle) = l(y) + l(x) + 2l(l(x)) + 1$, for example $\langle x, y \rangle = 1^{l(l(x))}0l(x)xy$.

2.2. Kolmogorov Complexity

Informally, the Kolmogorov complexity, or algorithmic entropy, $C(x)$ of a string x is the length (number of bits) of a shortest binary program (string) to compute x on a fixed reference universal computer (such as a particular universal Turing machine). Intuitively, $C(x)$ represents the minimal amount of information required to generate x by any effective process, [12]. The conditional Kolmogorov complexity $C(x | y)$ of x relative to y is defined similarly as the length of a shortest program to compute x , if y is furnished as an auxiliary input to the computation. The functions $C(\cdot)$ and $C(\cdot | \cdot)$, though defined in terms of a particular machine model, are machine-independent up to an additive constant (depending on the particular enumeration of Turing machines and the particular reference universal Turing machine selected). They acquire an asymptotically universal and absolute character through Church's thesis, and from the ability of universal machines to simulate one another and execute any effective process, see for example [15]. Formally:

Definition 1. *Let T_0, T_1, \dots be a standard enumeration of all Turing machines. Choose a universal Turing machine U that expresses its universality in the following manner:*

$$U(\langle \langle i, p \rangle, y \rangle) = T_i(\langle p, y \rangle),$$

for all i and $\langle p, y \rangle$, where p denotes a Turing program for T_i and y an input. We fix U as our reference universal computer and define the conditional Kolmogorov complexity of x given y by

$$C(x | y) = \min_{q \in \{0,1\}^*} \{l(q) : U(\langle q, y \rangle) = x\},$$

for every q (for example, $q = \langle i, p \rangle$ above) and auxiliary input y . The unconditional Kolmogorov complexity of x is defined by $C(x) = C(x | \epsilon)$. For convenience we write $C(x, y)$ for $C(\langle x, y \rangle)$, and $C(x | y, z)$ for $C(x | \langle y, z \rangle)$.

2.3. Incompressibility

Since there is a Turing machine, say T_i , that computes the identity function $T_i(x) \equiv x$, it follows that $U(\langle i, p \rangle) = T_i(p)$. Hence, $C(x) \leq l(x) + c$ for fixed $c \leq 2 \log i + 1$ and all x .^{2,3}

² “ $2 \log i$ ” and not “ $\log i$ ” since we need to encode i in such a way that U can determine the end of the encoding. One way to do that is to use the code $1^{l(l(i))}0l(i)i$ which has length $2l(l(i)) + l(i) + 1 < 2 \log i$ bits.

³ In what follows, “ \log ” denotes the binary logarithm. “[r]” is the greatest integer q such that $q \leq r$.

It is easy to see that there are also strings that can be described by programs much shorter than themselves. For instance, the function defined by $f(1) = 2$ and $f(i) = 2^{f(i-1)}$ for $i > 1$ grows very fast, $f(k)$ is a “stack” of k twos. Yet for every k , it is clear that $f(k)$ has complexity at most $C(k) + O(1)$. What about incompressibility? For every n there are 2^n binary strings of length n , but only $\sum_{i=0}^{n-1} 2^i = 2^n - 1$ descriptions in binary string format of length less than n . Therefore, there is at least one binary string x of length n such that $C(x) \geq n$. We call such strings *incompressible*. The same argument holds for conditional complexity: since for every length n there are at most $2^n - 1$ binary programs of length $< n$, for every binary string y , there is a binary string x of length n such that $C(x | y) \geq n$. Strings that are incompressible are patternless, since a pattern could be used to reduce the description length. Intuitively, we think of such patternless sequences as being random, and we use “random sequence” synonymously with “incompressible sequence.” Since there are few short programs, there can be only few objects of low complexity: the number of strings of length n that are compressible by at most δ bits is at least $2^n - 2^{n-\delta} + 1$.

Lemma 1. *Let δ be a positive integer. For every fixed y , every set S of cardinality m has at least $m(1 - 2^{-\delta}) + 1$ elements x with $C(x | y) \geq \lfloor \log m \rfloor - \delta$.*

Proof. There are $N = \sum_{i=0}^{n-1} 2^i = 2^n - 1$ binary strings of length less than n . A fortiori there are at most N elements of S that can be computed by binary programs of length less than n , given y . This implies that at least $m - N$ elements of S cannot be computed by binary programs of length less than n , given y . Substituting n by $\lfloor \log m \rfloor - \delta$ together with Definition 1 yields the lemma. ■

If we are given S as an explicit table, then we can simply enumerate its elements (in, say, lexicographical order) using a fixed program not depending on S or y . Such a fixed program can be given in $O(1)$ bits. Hence the complexity satisfies $C(x | S, y) \leq \log |S| + O(1)$.

2.4. Incompressibility Method

In a typical proof using the incompressibility method, one first chooses an incompressible object from the class under discussion. The argument invariably says that if a desired property does not hold, then in contrast with the assumption, the object can be compressed. This yields the required contradiction. Since most objects are almost incompressible, the desired property usually also holds for almost all objects, and hence on average.

3. GRID AND PEBBLES

In the analysis of the triangle problem, we first consider a discrete version based on an equally spaced $K \times K$ grid in the unit square. The general result for the continuous situation is then obtained by taking the limit for $K \rightarrow \infty$. Call the resulting axis-parallel $2K$ lines *grid lines* and their crossing points *grid points*. We place n points on grid points. These n points will be referred to as *pebbles* to avoid confusion with grid points or other geometric points arising in the discussion.

There are $\binom{K^2}{n}$ ways to put n *unlabeled* pebbles on the grid where at most one pebble is put on every grid point. We count only distinguishable arrangements without regard for the identities of the placed pebbles. Clearly, the restriction that no two pebbles can be placed on the same grid point is no restriction anymore when we let K grow unboundedly.

Erdős [18] demonstrated that for the special case of $p \times p$ grids, where p is a prime number, there are necessary arrangements of p pebbles with every pebble placed on a grid point such that no three pebbles are collinear. The least area of a triangle in such an arrangement is at least $1/(2p^2)$. This implies that the triangle constant $\Delta_n = \Omega(1/n^2)$ as $n \rightarrow \infty$ through the special sequence of primes.

We now give some detailed examples—used later—of the use of the incompressibility method. By Lemma 1, for every integer δ independent of K , every arrangement X_1, \dots, X_n (locations of pebbles), out of at least a fraction of $1 - 1/2^\delta$ of all arrangements of n pebbles on the grid, satisfies

$$C(X_1, \dots, X_n \mid n, K) \geq \log \binom{K^2}{n} - \delta. \quad (1)$$

Notation 1. For convenience, we abbreviate the many occurrences of the phrase “Let X_1, \dots, X_n be an arrangement of n pebbles on the $K \times K$ grid, let n be fixed and K be sufficiently large, and let δ be a positive integer constant such that (1) holds” to “If (1) holds” in the rest of the article.

Note that, for every arrangement X_1, \dots, X_n of n pebbles on a $K \times K$ grid, we have $C(X_1, \dots, X_n \mid n, K) \leq \log \binom{K^2}{n} + O(1)$ —there is a fixed program of $O(1)$ bits for the reference universal computer that reconstructs the X_1, \dots, X_n from n, K , and its index in the lexicographical ordering of all possible arrangements. That (1) holds with δ small means that the arrangement X_1, \dots, X_n of pebbles on the grid has no regularity that can be used to prepare a description that is significantly shorter than simply giving the index in the lexicographical ordering of all possible choices of n positions from the available $K \times K$ grid positions. We can view such an arrangement as being “random” or “in general position.”

Lemma 2. *If (1) holds, then no three pebbles can be collinear, and so the area of a smallest triangle is at least $1/(2(K-1)^2)$.*

Remark 1. This is the first proof of the article using the incompressibility argument. Let us explain the proof idea in detail: On the one hand, we construct a description d such that the arrangement X_1, \dots, X_n can be reconstructed from d by a fixed program p for the universal reference computer, given also n and K . If p is in self-delimiting format, then the universal reference computer can parse pd into its constituent parts p and d , and then execute p to reconstruct X_1, \dots, X_n from the auxiliary information n, K , together with the description d . On the other hand, by definition, the Kolmogorov complexity of an object is the length of its *shortest* program for the reference universal computer and we have assumed a lower bound on the Kolmogorov complexity. Since the description pd is a program for the reference universal computer, its length $l(pd)$ must be at least as large as the Kolmogorov complexity (the auxiliary information n, K being the same in both cases). By the

lower bound (1), this shows that $l(pd) \geq \log \binom{K^2}{n} - \delta$. Since $l(p)$ is independent of n, K , we can set $l(p) = O(1)$ in this context, and obtain $l(d) \geq \log \binom{K^2}{n} - \delta - O(1)$. By exploiting collinearity of pebbles in the description d , to make it as compact as possible, this inequality will yield the required contradiction for n fixed and K large enough.

Proof. Place $n - 1$ pebbles at positions chosen from the total of K^2 grid points—there are $\binom{K^2}{n-1}$ choices. Choose two pebbles, P and Q , from among the $n - 1$ pebbles—there are $\binom{n-1}{2}$ choices. Choose a new pebble R on the straight line determined by P, Q . The number of grid points on this line between P (or Q) and R , which number is $< K$, identifies R uniquely in $\leq \log K$ bits. There is a fixed algorithm that, on input n and K , decodes a binary description consisting of the items above—each encoded as the logarithm of the number of choices—and computes the positions of the n pebbles. By (1) this implies

$$\log \binom{K^2}{n-1} + \log \binom{n-1}{2} + \log K + O(1) \geq \log \binom{K^2}{n} - \delta.$$

Using the asymptotic expression

$$\log \binom{a}{b} - b \log \frac{a}{b} \rightarrow b \log e - \frac{1}{2} \log b + O(1), \quad (2)$$

for b fixed and $a \rightarrow \infty$, one obtains $3 \log n \geq \log K - \delta + O(1)$, which is a contradiction for n fixed and K sufficiently large. ■

Lemma 3. *If (1) holds, then no two pebbles can be on the same (horizontal) grid line.*

Proof. Place $n - 1$ pebbles at positions chosen from the total of K^2 grid points—there are $\binom{K^2}{n-1}$ choices. Choose one pebble P from among the $n - 1$ pebbles—there are $n - 1$ choices. Choose a new pebble R on the (horizontal) grid line determined by P —there are $K - 1$ choices. There is a fixed algorithm that, on input n and K , reconstructs the positions of all n pebbles from a description of these choices. By (1) this implies

$$\log \binom{K^2}{n-1} + \log(n-1) + \log K + O(1) \geq \log \binom{K^2}{n} - \delta.$$

Using (2) with fixed n and $K \rightarrow \infty$ we obtain $2 \log n \geq \log K - \delta + O(1)$, which is a contradiction for large enough K . ■

4. LOWER BOUND

Our strategy is to show that if we place n pebbles on a $K \times K$ grid, such that the arrangement has high Kolmogorov complexity, then every three pebbles form a triangle of at least a certain size area. If the area is smaller, then this can be used to compress the description size of the arrangement to below the assumed Kolmogorov complexity.

Theorem 1. *If (1) holds, then there is a positive constant c_1 such that the least area of every triangle formed by three pebbles on the grid is at least $c_1/(2^\delta n^3)$.*

Proof. Place $n - 1$ pebbles at positions chosen from the total of K^2 grid points—there are $\binom{K^2}{n-1}$ choices. Choose two pebbles, P and Q , from among the n pebbles—there are $\binom{n}{2}$ choices. Place a new pebble R at one of the remaining grid points. Without loss of generality, let the triangle PQR have PQ as the longest side. Center the grid coordinates on $P = (0, 0)$ with $Q = (q_1, q_2)$ and $R = (r_1, r_2)$ in units of $1/(K - 1)$ in both axes directions. Then R is one of the grid points on the two parallel line segments of length $L = |PQ| = \sqrt{q_1^2 + q_2^2}/(K - 1)$ at distance $H = |q_2 r_1 - q_1 r_2|/((K - 1)\sqrt{q_1^2 + q_2^2})$ from the line segment PQ , as in Figure 1. The number of grid points on each of these line segments (including one endpoint and excluding the other endpoint) is a positive integer $g = \gcd(q_1, q_2)$ —the line $q_2 x = q_1 y$ has g integer coordinate points between $(0, 0)$ and (q_1, q_2) including one of the endpoints. This implies that f defined by $LH(K - 1)^2 = fg$ is a positive integer as well.

Enumerating the grid points concerned in lexicographical order, the index of R takes at most $\log(2gf) = \log(2g) + \log f = \log(4|PQR|(K - 1)^2)$ bits, where $|PQR|$ denotes the area of the triangle PQR . Altogether this constitutes an effective description of the arrangement of the n pebbles. By the assumption in the theorem the arrangement satisfies (1), that is, the number of bits involved in any effective description of the arrangement is lower bounded by the right-hand side. Then,

$$\log \binom{K^2}{n-1} + \log \binom{n}{2} + \log(4|PQR|(K - 1)^2) + O(1) \geq \log \binom{K^2}{n} - \delta.$$

By approximation (2),

$$\log \binom{K^2}{n} - \log \binom{K^2}{n-1} \rightarrow \log \frac{K^2}{n} + O(1),$$

for large enough fixed n and $K \rightarrow \infty$. Therefore, $\log |PQR| + O(1) \geq -3 \log n - \delta + O(1)$, $K \rightarrow \infty$. Consequently, there exists a positive constant c_1 , independently of the particular triangle PQR , such that $|PQR| > c_1/(n^3 2^\delta)$ for all large enough n and K . Since this holds for every triangle PQR , constructed as above, it holds in particular for a triangle of least area A . ■




Fig. 1. Smallest triangle based on pebbles P, Q .

By Lemma 1, the probability concentrated on the set of arrangements satisfying (1) is at least $1 - 1/2^\delta$:

Corollary 1. *If n points are chosen independently and at random (uniform distribution) in the unit square, and A is the least area of a triangle formed by three points, then there is a positive constant c_1 such that for every positive δ we have $A > c_1/(2^\delta n^3)$ with probability at least $1 - 1/2^\delta$.*

In the particular case of $\delta = 1$, the probability concentrated on arrangements satisfying (1) is at least $\frac{1}{2}$ which immediately implies:

Corollary 2. *If n points are chosen independently and at random (uniform distribution) in the unit square, then there is a positive constant c such that the least area of some triangle formed by three points has expectation $\mu_n > c/n^3$.*

5. UPPER BOUND

Every pair of pebbles out of an incompressible arrangement of n pebbles on a $K \times K$ grid defines a distinct line by Lemma 2. The two pebbles defining such a line together with any other pebble forms a triangle. If A is the least area of a triangle formed by three pebbles, then this constrains the possibilities of placing a third pebble close to a line defined by two pebbles. Thus, every such line defines a forbidden strip on both sides of the line where no pebbles can be placed. It is easy geometry to see that every forbidden strip covers an interval of length $2A$ of every grid line on both sides of the intercept of the “forbidding” line concerned. Our strategy is as follows: Divide the pebbled unit square by a straight line parallel to the horizontal sides into two parts containing about one half of the n pebbles each. Show that the pebbles in the larger half (the halves may not have equal area) of the unit square define $\Omega(n^2)$ distinct “forbidding lines”, that cross both the dividing line and the opposite parallel side of the unit square. While the associated forbidden grid point positions can overlap, we show that they do not overlap too much. As a consequence, the set of grid points allowed to place the remaining $n/2$ pebbles in the smaller remaining half of the unit square, gets restricted to the point that the description of the arrangement can be compressed too far. This argument is so precise that for small δ in (1), the upper bound is the same order of magnitude as the previously proven lower bound.

Theorem 2. *If (1) holds with $\delta < (2 - \epsilon) \log n$ for some positive constant ϵ , then there is a positive constant C_1 such that the least area of some triangle formed by three pebbles on the grid is at most*

$$A(\delta) = \frac{14\delta + O(1)}{4C_1 n^3 \log e}. \quad (3)$$

Proof. Choose n pebbles at positions chosen from the total of K^2 grid points such that (1) is satisfied. Divide the unit square by a horizontal grid line into an upper and a lower half, each of which contains $n/2 \pm 1$ pebbles—there are no grid lines

containing two pebbles by Lemma 3. We write *forbidding line* for a line determined by two pebbles in the upper half that intersects all horizontal grid lines in the lower half of the unit square.

Claim 1. *If (1) holds, then there is a positive constant C_1 such that there are at least $C_1 n^2$ forbidding lines.*

Proof. Take the top half to be the larger half so that it has area at least $1/2$. Divide the top half into five vertical strips of equal width of $1/5$ and five horizontal strips of equal width $1/10$ starting from the top—ignore the possibly remaining horizontal strip at the bottom of the top half. Clearly, a forbidding line determined by a pebble in the upper rectangle and a pebble in the lower rectangle of the middle vertical strip intersects the bottom horizontal grid line. We show that these rectangles contain at least $n/100$ points each, and hence the claim holds with $C_1 = 1/10,000$.

Consider either rectangle (the same argument will hold for the other rectangle). Let it contain $m \leq n$ pebbles. Since the area of the rectangle is $1/5 \times 1/10$, it contains $K^2/50$ grid points (plus or minus the grid points on the circumference of length $3K/5$ which we ignore). Place $n - m$ pebbles at positions chosen from $49K^2/50$ grid points outside the rectangle—there are $\binom{49K^2/50}{n-m}$ choices—and place m pebbles at positions chosen from the total of $K^2/50$ grid points in the rectangle—there are $\binom{K^2/50}{m}$ choices. Given n and K , the n pebble positions are determined by m , the position of the rectangle and an index number i of $\log i$ bits with

$$\begin{aligned} \log i &= \log \binom{49K^2/50}{n-m} \binom{K^2/50}{m} \\ &\rightarrow (n-m) \log \frac{49K^2/50}{n-m} + m \log \frac{K^2/50}{m} + n \log e - \frac{1}{2} \log nm + O(1), \end{aligned}$$

for $K \rightarrow \infty$ with n, m fixed, by (2). Given n , we can describe m in $\log n$ bits. Thus, given n and K , the total description length of the description of the arrangement of the n pebbles is $\log n + \log i + O(1)$ bits. This must be at least the Kolmogorov complexity of the arrangement. Then, by (1),

$$(n-m) \log \frac{49K^2/50}{n-m} + m \log \frac{K^2/50}{m} - \frac{1}{2} \log m + O(1) \geq n \log \frac{K^2}{n} - \delta.$$

This implies

$$\begin{aligned} \delta &\geq (n-m) \log \frac{50(n-m)}{49} + m \log 50m + \frac{1}{2} \log m - n \log n - O(1) \\ &> (n-m) \log \frac{50(n-m)}{49} + m \log 50m - n \log n - O(1). \end{aligned}$$

Assume, by way of contradiction, $m \leq n/100$. Then,

$$\begin{aligned} \delta &\geq \frac{99}{100} n \log \frac{4950}{4900} n + \frac{1}{100} n \log \frac{50}{100} n - n \log n - O(1) \\ &= n \left(\frac{99}{100} \log \frac{4950}{4900} + \frac{1}{100} \log \frac{50}{100} \right) - O(1) \\ &> n(0.0145 - 0.01) - O(1), \end{aligned}$$

which contradicts $\delta = O(\log n)$ in the statement of the theorem. Hence the top rectangle and the bottom rectangle of the middle strip in the top half contain at least $n/100$ pebbles each. Each pair of pebbles, one in the top rectangle and one in the bottom rectangle, determine a distinct forbidding line by Lemma 2 (no three pebbles can be collinear under assumption (1)). The claim is proven with $C_1 = (1/100) \cdot (1/100) = 1/10^4$. ■

Claim 2. *Let w_1, w_2, w_3, w_4, w_5 be the spacings between the six consecutive intercepts of a sextuplet of forbidding lines with a horizontal grid line in the bottom half containing a pebble, and let $D = w_1 + w_2 + w_3 + w_4 + w_5$. If (1) holds, then there is a positive C_2 such that $D > C_2/n^{3-\epsilon/5}$ with ϵ as in the statement of the theorem.*

Proof. Place $n - 5$ pebbles at positions chosen from the total of K^2 grid points—there are $\binom{K^2}{n-5}$ choices. Choose eight pebbles, P_i ($i = 0, 1, 2, 3, 5, 7, 9, 11$) from among the $n - 5$ pebbles—there are at most $\binom{n-5}{8}$ choices—and five new pebbles P_j ($j = 4, 6, 8, 10, 12$) such that $P_1P_2, P_3P_4, P_5P_6, P_7P_8, P_9P_{10}, P_{11}P_{12}$ is the sextuplet of forbidding lines in the claim, and P_0 is a pebble in the lower half. Without loss of generality, we assume that the “middle” pebbles of unknown position P_j ($j = 4, 6, 8, 10, 12$), as well as P_2 in known position, are in between the other defining pebble of the forbidding line concerned and its intercept with the lower grid line containing P_0 . That is, the top-to-bottom order a forbidding line is P_1, P_2 , intercept₁, P_3, P_4 , intercept₂, and so on. Then, a forbidding line determined by an outermost pebble and an intercept, together with the grid line containing the middle pebble, enables us to determine the grid point on which the middle pebble is located. An error in the position of the intercept leads to a smaller error in the position of the middle pebble. Thus, a precision of the position of the intercept up to $1/(4(K - 1))$, together with the precise position of the outermost pebble, enables us to determine the grid point containing the middle pebble as the unique grid point in a circle with radius $1/(4(K - 1))$ centered on the computed geometric point. The coordinates of the five unknown P_j s are determined by (i) the locations of the five intercepts of the associated quintuplet of forbidding lines with the lower half horizontal grid line on which P_0 is located, and (ii) the five unknown distances between these intercepts and the P_j s along the five associated forbidding lines. The grid point positions of the P_j s are uniquely determined if we know the latter distances up to precision $1/4(K - 1)$. All six intercepts in the statement of the theorem are in an interval of length D which contains DK grid points (rounded to the appropriate close entire value). We can describe every intercept in this interval (up to the required precision) in $\log DK + O(1)$ bits. Relative to the intersection of the known forbidding line P_1P_2 , therefore, item (i) uses $5 \log DK + O(1)$ bits. Item (ii) uses $5 \log K + O(1)$ bits. Given n, K , we can describe the placement of the $n - 5$ pebbles in $\log \binom{K^2}{n-5}$ bits; the choice of the eight pebbles among them in $\log \binom{n-5}{8}$ bits; and we have shown that the placement of the five unknown pebbles can be reconstructed from an additional $5 \log DK + 5 \log K + O(1)$ bits. Together, this forms a description of the complete arrangement. By (1) this implies:

$$\log \binom{K^2}{n-5} + 8 \log n + 5 \log DK + 5 \log K + O(1) \geq \log \binom{K^2}{n} - \delta.$$

A now familiar calculation using (2) yields $5 \log D + O(1) \geq -13 \log n - \delta$, for fixed n and $K \rightarrow \infty$. This shows $D > C_2 2^{(2 \log n - \delta)/5} / n^3$ for some positive constant C_2 . Substituting $\delta < (2 - \epsilon) \log n$ proves the claim. ■

We have now established that there are $C_1 n^2$ distinct forbidding lines (with C_1 as in Claim 1) determined by pairs of pebbles in the upper half, and by construction every such forbidding line intersects every lower half horizontal grid line. Moreover, every D -length interval (with D as in Claim 2) on a lower half horizontal grid line—that contains a pebble—contains at most six intercepts of forbidding lines. This means that we can select $C_1 n^2 / 7$ consecutive intercepts on such a grid line that are separated by intervals of at least length D . The two pebbles P, Q defining the forbidding line l_1 , together with any pebble R on a lower half horizontal grid line l_2 , determine a triangle. If d is the distance between the intercept point of l_1 with l_2 and the pebble R , and α is the angle between the forbidding line l_1 and grid line l_2 , then the triangle side located on the forbidding line has length $\leq 1 / \cos \alpha$ while the height of the triangle with respect to that side is $d \cos \alpha$. Thus, if A is the area of the smallest triangle formed by any three pebbles, then $d \geq 2A$. Consequently, all grid positions in intervals of length $2A$ on both sides of an intercept of a forbidding line with a lower half grid line—that contains a pebble—are forbidden for pebble placement. As long as $2d \leq D$, or

$$4A \leq D, \quad (4)$$

this means that the $C_1 n^2 / 7$ consecutive intercepts exclude $4AC_1 n^2 / 7$ grid positions from pebble placement on the horizontal lower grid line concerned. If (4) does not hold, that is, $4A > D$, then at least $DC_1 n^2 / 7$ grid positions are excluded. Given the pebbles in the upper half, and therefore the forbidding lines, the excluded grid points in the lower half are determined. Therefore, with

$$B = \min\{4A, D\}, \quad (5)$$

and also given the horizontal lower half grid line concerned, we can place a pebble on the grid line in at most positions.

$$K(1 - C_1 n^2 B / 7) \quad (6)$$

We now use this fact to construct a short encoding of the total arrangement of the n pebbles satisfying (1): Select n horizontal grid lines (there can be only one pebble per grid line by Lemma 2) chosen from the total of K grid lines—there are $\binom{K}{n}$ choices. Select on everyone of the upper $n/2$ horizontal grid lines a grid point to place a pebble—there are $K^{n/2}$ choices. Finally, select in order from top to bottom on the lower $n/2$ horizontal grid lines $n/2$ grid points to place the pebbles—there are only $(K(1 - C_1 n^2 B / 7))^{n/2}$ choices by (6). Together, these choices form a description of the arrangement. Given the values of n, K we can encode these choices in self-delimiting items, and by (1) this implies:

$$\log \binom{K}{n} + \frac{n}{2} \log K + \frac{n}{2} \log K(1 - C_1 n^2 B / 7) + O(1) \geq \log \binom{K^2}{n} - \delta.$$

Using (2) with n fixed yields

$$\frac{n}{2} \log(1 - C_1 n^2 B / 7) \geq -\delta - O(1), \quad K \rightarrow \infty.$$

The left-hand side

$$\log\left(1 - \frac{C_1 n^3 B/14}{n/2}\right)^{n/2} = \log e^{-C_1 n^3 B/14}, \quad n \rightarrow \infty,$$

so that

$$B \leq \frac{14\delta + O(1)}{C_1 n^3 \log e} \quad (7)$$

Since $\delta < 2 \log n$ in the right-hand side, Claim 2 shows that $D > B$. Therefore, (5) implies $B = 4A$ so that (7) establishes the theorem. ■

Together with Lemma 1, Theorem 2 implies that the smallest triangle in an arrangement has an area below a particular upper bound with a certain probability.

Corollary 3. *If n points are chosen independently and at random (uniform distribution) in the unit square, and A is the least area of a triangle formed by three points, then for every positive $\delta < (2 - \epsilon) \log n$ ($\epsilon > 0$), we have*

$$A < A(\delta)$$

with probability at least $1 - 1/2^\delta$.

That is, the probability that $A < A(1)$ at least $\frac{1}{2}$ ($\delta = 1$), the probability that $A < A(2)$ is at least $\frac{3}{4}$ ($\delta = 2$), and so on. Since $A(\delta + 1) \geq A(\delta)$, we can upper bound the expectation μ_n of A by upper bounding the probability of A with $A(\delta) < A \leq A(\delta + 1)$ by $2^{-\delta+1} = [(1 - 2^{-\delta}) - (1 - 2^{-\delta-1})]$. We do this for $\delta \leq 1.9 \log n$. The remaining probability is $1/n^{1.9}$ or slightly less (because δ is integer). This probability is so small that, even if we assume the known worst-case upper bound on A for the remaining cases, known to be $C_3/n^{8/7-\epsilon'}$ for some positive constant C_3 for every $\epsilon' > 0$, [13], the result is insignificant. There is a positive constant C such that:

$$\mu_n \leq \sum_{\delta=1}^{1.9 \log n} 2^{-\delta} A(\delta) + \frac{1}{n^{1.9}} \frac{C_3}{n^{8/7-\epsilon'}} < \frac{C}{n^3}.$$

Corollary 4. *If n points are chosen independently and at random (uniform distribution) in the unit square, then there is a positive constant C such that the least area of some triangle formed by three points has expectation $\mu_n < C/n^3$.*

ACKNOWLEDGMENTS

We thank John Tromp for help with the proof of Theorem 1. We also thank him and the anonymous referees for valuable comments on drafts of this article.

REFERENCES

- [1] G. Barequet, A lower bound for Heilbronn's triangle problem in d dimensions, SIAM J Discrete Math 14(2) (2001), 230–236.
- [2] J. Beck, Almost collinear triples among N points on the plane, in A Tribute to Paul Erdős, A. Baker, B. Bollobas and A. Hajnal (Editors), Cambridge University Press, Cambridge, 1990, pp. 39–57.

- [3] C. Bertram-Kretzberg, T. Hofmeister, and H. Lefmann (Editors), An algorithm for Heilbronn's problem, *SIAM J Comput* 30(2) (2000), 383–390.
- [4] G. Cairns, M. McIntyre, and J. Strantzen, Geometric proofs of some recent results of Yang Lu, *Math Magazine* 66 (1993), 263–265.
- [5] P. Erdős, “Problems and results in combinatorial geometry,” *Discrete geometry and convexity*, Ann New York Acad Sci 440 (1985), 1–11.
- [6] P. Erdős and G. Purdy, “Extremal problems in combinatorial theory,” *Handbook of combinatorics*, R.L. Graham, M. Grötschel, L. Lovász, (Editors), Elsevier/MIT Press, 1995, pp. 861–862.
- [7] M. Goldberg, Maximizing the smallest triangle made by N points in a square, *Math Magazine* 45 (1972), 135–144.
- [8] R.K. Guy, *Unsolved problems in number theory*, 2nd ed., Springer-Verlag, Berlin, 1994, pp. 242–244.
- [9] T. Jiang, J.I. Seiferas, and P.M.B. Vitányi, Two heads are better than two tapes, *J Assoc Comput Mach* 44(2) (1997), 237–256.
- [10] T. Jiang, M. Li, and P.M.B. Vitányi, The expected size of Heilbronn's triangles, *Proc 14th IEEE Conf Comput Complexity*, IEEE Comp Soc, Los Alamitos, CA, 1999, pp. 105–113.
- [11] T. Jiang, M. Li, and P.M.B. Vitányi, A lower bound on the average-case complexity of Shellsort, *J Assoc Comput Mach* 47(5) (2000), 905–911.
- [12] A.N. Kolmogorov, Three approaches to the quantitative definition of information, *Problems Inform Transmission* 1(1) (1965), 1–7.
- [13] J. Komlós, J. Pintz, and E. Szemerédi, On Heilbronn's triangle problem, *J London Math Soc* 24(2) (1981), 385–396.
- [14] J. Komlós, J. Pintz, and E. Szemerédi, A lower bound for Heilbronn's problem, *J London Math Soc* 25 (1982), 13–24.
- [15] M. Li and P.M.B. Vitányi, *An introduction to Kolmogorov complexity and its applications*, 2nd ed., Springer-Verlag, New York, 1997.
- [16] D. Mackenzie, On a roll, *New Scientist* 6 (1999), 44–48.
- [17] A.M. Odlyzko, J. Pintz, and K.B. Stolarsky, Partitions of planar sets into small triangles, *Discrete Math* 57 (1985), 89–97.
- [18] K.F. Roth, On a problem of Heilbronn, *J London Math Soc* 26 (1951), 198–204.
- [19] K.F. Roth, On a problem of Heilbronn II, *Proc. London Math Soc* 25(3) (1972), 193–212.
- [20] K.F. Roth, On a problem of Heilbronn III, *Proc. London Math Soc* 25(3) (1972), 543–549.
- [21] K.F. Roth, Estimation of the area of the smallest triangle obtained by selecting three out of n points in a disc of unit area, *Proc. Symp. Pure Mathematics*, Vol. 24, AMS, Providence, 1973, pp. 251–262.
- [22] K.F. Roth, Developments in Heilbronn's triangle problem, *Adv Math* 22 (1976), 364–385.
- [23] W.M. Schmidt, On a problem of Heilbronn, *J London Math Soc* 4(2) (1972), 545–550.
- [24] T. Z. Ping, On the problem of Heilbronn type, *Northeast Math J* 10 (1994), 215–216.
- [25] L. Yang, J.Z. Zhang, and Z.B. Zeng, Heilbronn problem for five points, *Int'l Centre Theoret Physics preprint IC/91/252* (1991).
- [26] L. Yang, J.Z. Zhang, and Z.B. Zeng, A conjecture on the first several Heilbronn numbers and a computation, *Chinese Ann Math Ser A* 13 (1992), 503–515.
- [27] L. Yang, J.Z. Zhang, and Z.B. Zeng, On the Heilbronn numbers of triangular regions, *Acta Math Sinica* 37 (1994), 678–689.