

On the Height of a Random Set of Points in a d -Dimensional Unit Cube

Eric Breimer, Mark Goldberg, Brian Kolstad, and Malik Magdon-Ismail

CONTENTS

- 1. Introduction
- 2. Location of Maximal Chains
- 3. Boosting
- 4. Coconvergence
- 5. Algorithms
- 6. Experiments
- 7. Some Technical Details
- 8. Conclusions
- Acknowledgement
- References

We investigate, through numerical experiments, the asymptotic behavior of the length $H_d(n)$ of a maximal chain (longest totally ordered subset) of a set of n points drawn from a uniform distribution on the d -dimensional unit cube $\mathbf{V}_D = [0, 1]^d$. For $d \geq 2$, it is known that $c_d(n) = H_d(n)/n^{1/d}$ converges in probability to a constant $c_d < e$, with $\lim_{d \rightarrow \infty} c_d = e$. For $d = 2$, the problem has been extensively studied, and it is known that $c_2 = 2$; c_d is not currently known for any $d \geq 3$. Straightforward Monte Carlo simulations to obtain c_d have already been proposed, and shown to be beyond the scope of current computational resources. In this paper, we present a computational approach which yields *feasible* experiments that lead to estimates for c_d . We prove that $H_d(n)$ can be estimated by considering only those chains close to the diagonal of the cube. A new conjecture regarding the asymptotic behavior of $c_d(n)$ leads to even more efficient experiments. We present experimental support for our conjecture, and the new estimates of c_d obtained from our experiments, for $d \in \{3, 4, 5, 6\}$.

1. INTRODUCTION

Erdős and Szekeres [1935] proved that every permutation of $\{1, 2, \dots, p^2+1\}$ has either an increasing or a decreasing subsequence of length $p+1$. This result inspired Ulam to investigate the expected value L_n of the length of a longest increasing subsequence in a random permutation of n numbers. The Monte Carlo simulations he conducted [Ulam 1961] led him to conjecture that the answer is asymptotic to $c\sqrt{n}$, for some constant $c > 1$. Based on more extensive computations, Baer and Brock [1968] conjectured that $c = 2$. Hammersley [1972] proved that $\pi/2 \leq c \leq e$, and presented heuristic arguments for the value $c = 2$. Logan and Shepp [1977] proved that $c \leq 2$, and Pilpel [1990] presented an elementary proof of this inequality. Finally, Versík and Kerov [1977] settled the conjecture by proving that $c \geq 2$. The structure of the distribution of L_n has

been studied in [Baik et al. 1999; Borodin 1999; Odlyzko and Rains 2000].

The multidimensional generalization of the monotone subsequence problem was first considered by Steel [1977], who conjectured the existence of a constant that generalizes the Hammersley's constant c to the d -dimensional case. The correct generalization was found by Bollobás and Winkler [1988]. Let $\mathbf{V}_d = [0, 1]^d$ be the unit cube in d dimensions and let n random points $\mathbf{x}(1), \mathbf{x}(2), \dots, \mathbf{x}(n)$ be chosen independently from the uniform distribution on \mathbf{V}_d . These points form the underlying set of a *random order* $\mathbf{P}_d(n)$ with a partial ordering given by $\mathbf{x}(i) \leq \mathbf{x}(j)$ if and only if $x_k(i) \leq x_k(j)$ for all $k = 1, \dots, d$. Let the height $H_d(n)$ be the number of elements in a longest chain (totally ordered subset) of $\mathbf{P}_d(n)$. We are interested in the asymptotic behavior of the random variable

$$c_d(n) = \frac{H_d(n)}{n^{1/d}}.$$

Bollobás and Winkler [1988] showed that, for each $d \geq 0$, $c_d(n)$ converges in probability to a constant c_d . For convenience, we restate their theorem here.

Theorem 1.1 [Bollobás and Winkler 1988]. *For $d = 1, 2, \dots$, define $c_d = \limsup_{n \rightarrow \infty} E[c_d(n)]$. Then:*

1. $c_d < e$.
2. For all $\varepsilon > 0$ and $\delta > 0$, there exists $M(\varepsilon, \delta)$ such that $n > M(\varepsilon, \delta)$ implies

$$P \left[\left| \frac{H_d(n)}{n^{1/d}} - c_d \right| \leq \varepsilon \right] > 1 - \delta.$$

3. $\lim_{d \rightarrow \infty} c_d = e$.

Except for $c_2 = 2$, no other c_d is currently known for $d \geq 2$. It is not even known whether the sequence $\{c_d\}$ is monotonically increasing in d . It is not hard to conceive a computational experiment that would, hopefully, lead to approximate values of c_d , for some $d > 2$:

Experiment 1 (Naïve determination of c_d).

1. For a large integer n , randomly generate a set S of n points from the uniform distribution on $\mathbf{V}_d = [0, 1]^d$.
2. Compute the height $H_d(S)$ (i.e., the length of a maximal chain in S) and $c_d(S) = H_d(S)/n^{1/d}$.
3. Repeat these computations K times to obtain $c_d(S_1), c_d(S_2), \dots, c_d(S_K)$.

4. Compute the mean, $\bar{c}_d(n)$, and the standard deviation, $\sigma_d(n)$, of the $c_d(S_i)$'s.

For K sufficiently large, one expects that $\bar{c}_d(n)$ will be close to $E[c_d(n)]$ (to within $\sigma_d(n)/\sqrt{K}$), and if n is sufficiently large, then $E[c_d(n)]$ should be close to c_d . The determination of whether n is sufficiently large is a difficult problem. One usually runs an experiment for the largest value of n that is practically feasible, given the current computational resources. Unfortunately, however, it has been discovered through experimental investigations (Catherine McGeoch, private communication, 2000) that the values of n that need to be considered in order to get a good approximation for c_d are not computationally feasible. In particular, according to McGeoch, $n = 10^9$ points for $d = 3$ is not sufficiently large to suggest an estimate for c_3 . Thus, we would need to consider n at least as large as 10^{10} . Since the standard dynamic programming algorithm for computing the length of a maximal chain is quadratic in n , solving even one example with $n = 10^{10}$ points would typically take a few weeks on a modern workstation. In order to develop reasonable statistics, one would need several years. An equally acute problem is with the memory requirements for such an experiment: storing 10^{10} d -dimensional vectors in RAM is impossible on almost any advanced workstation, and storing them on disk would substantially slow down the algorithm.

Hence, the main challenge for an experimental approach to estimating c_d lies in the design of feasible computational experiments that estimate $c_d(n)$ for very large n . In this paper, we propose an approach that addresses both the computational efficiency and the memory issues. The core idea behind our approach is the observation that a maximal chain "close" to the diagonal of the unit cube $\mathbf{V}_d = [0, 1]^d$ must exist as $n \rightarrow \infty$. We present a theoretical justification of this observation and use this fact to construct sequences that converge to c_d . Based upon geometrical considerations, we *conjecture* that each of these sequences converges at essentially the same rate. We present experimental data to support this conjecture, and then use this conjecture to derive, from the data, estimates for $\{c_d\}$, for $d = 3, 4, 5, 6$. The intervals for c_d for those values of d are:

d	c_d
2	[1.998, 2.002]
3	[2.363, 2.366]
4	[2.514, 2.521]
5	[2.583, 2.589]
6	[2.607, 2.617]

Our estimates suggest that $\{c_d\}$ is a monotonically increasing sequence in d , a fact that has previously not been suggested by any experiments.

The outline of the paper is as follows. The next section describes two approaches to the design of experiments that build upon Experiment 1. The presentation is informal, and we do not describe every detail of all attempts that were made. The only exception is for the approach that actually yielded estimates for c_d . In Section 3, we present the data supporting our conjecture, in addition to the data used for deriving the estimates for c_d . Technical details are left to an appendix. Preliminary experimental results were presented in [Breimer et al. 2001].

2. LOCATION OF MAXIMAL CHAINS

Since a longest chain is only a very small subset of the set of n random points in $\mathbf{V}_d(r)$, we can expect to gain an advantage over the straightforward simulation in Experiment 1 by restricting our search for such a chain to a small region where we are likely to find it. Thus, our first task is to identify a region in the cube which with probability tending to 1 as $n \rightarrow \infty$ contains a maximal chain. It is intuitively apparent that a good candidate for such a region is the set of points that are “close” to the diagonal, say, at most a distance r from the diagonal. The justification of this observation is given in Theorem 2.1, which we prove using a modification of Bollobás and Winkler’s proof of Theorem 1.1.

Let $\mathbf{V}_d(r)$ be the region of $\mathbf{V}_d = [0, 1]^d$ obtained by translating a cube of side r along the diagonal (Figure 1). The volume $V_d(r)$ of $\mathbf{V}_d(r)$ is given by

$$V_d(r) = r^d + dr^{d-1}(1 - r);$$

to see this, note that $\mathbf{V}_d(r)$ can be divided into disjoint parts as follows:

1. Polygons p_i , for $1 \leq i \leq d$, defined by a $(d-1)$ -dimensional cube surface of side r parallel to the plane $x_i = 0$ and translated along the diagonal

- from $(0, 0, \dots, 0)$ to $(1-r, 1-r, \dots, 1-r)$; each such polygon has volume $r^{d-1}(1-r)$.
2. A d -dimensional cube of side r defined by $[1-r, 1]^d$, whose volume is r^d .




FIGURE 1. Diagonal volume element obtained by translation of a cube of side r .

Let n points be sampled uniformly in $\mathbf{V}_d(r)$ to form a random partially ordered set $\mathbf{P}_d(n, r)$. Denote the height of this set by $H_d(n, r)$ and define the random variable $c_d(n, r)$ by

$$c_d(n, r) = \frac{H_d(n, r)}{(n/V_d(r))^{1/d}}.$$

Intuitively, we expect that if one were to generate $n^* = \lceil n/V_d(r) \rceil$ points in the whole cube \mathbf{V}_d , then about n of them would fall in $\mathbf{V}_d(r)$. If n^* is large, a maximal chain should exist in $\mathbf{V}_d(r)$; thus $H_d(n, r)$ should approximate $H_d(n^*)$, the height of $\mathbf{P}_d(n^*)$. Therefore, for n sufficiently large,

$$c_d(n, r) \approx \frac{H_d(n^*)}{n^{*1/d}} = c_d(n^*).$$

Hence, for fixed r , $c_d(n, r)$ should converge to c_d . This is exactly the claim of the next theorem, which is essentially an analog of Theorem 1.1.

Theorem 2.1. Fix $d \in \{1, 2, \dots\}$ and $r > 0$. For all $\varepsilon > 0$ and $\delta > 0$, there exists $M_r(\varepsilon, \delta)$ such that if $n > M_r(\varepsilon, \delta)$ then

$$P[|c_d(n, r) - c_d| \leq \varepsilon] > 1 - \delta,$$

where the constants c_d are defined in Theorem 1.1.

Proof. An equivalent way to generate n points uniformly in $\mathbf{V}_d(r)$ is to generate n^* points uniformly

in V_d until n of them are in $V_d(r)$. Thus n^* is a random variable with a distribution given by

$$P[n^*=n+k] = p^n(1-p)^k \binom{n-1+k}{k}, \quad k = 0, 1, \dots,$$

where we have let $p = V_d(r)$. The expected value and standard deviation of n^* are then given by

$$E[n^*] = n/p,$$

$$\text{std}[n^*] = \sqrt{n(1-p)/p^2}.$$

Given $\varepsilon > 0$ and for any $\gamma > 0$, a straightforward application of Markov's inequality gives

$$P\left[\left(1 - \frac{\varepsilon p}{n^{1/2-\gamma}}\right)^{1/d} \leq \left(\frac{n^*}{E[n^*]}\right)^{1/d} \leq \left(1 + \frac{\varepsilon p}{n^{1/2-\gamma}}\right)^{1/d}\right]$$

$$> 1 - \frac{1-p}{(\varepsilon p n^\gamma)^2}.$$

Thus, $(n^*/E[n^*])^{1/d}$ converges in probability to 1. First we show that the inequality $c_d(n, r) \leq c_d$ holds in probability as $n \rightarrow \infty$. Since $H_d(n^*) \geq H_d(n, r)$, we have

$$c_d(n, r) = \frac{H_d(n, r)}{(E[n^*])^{1/d}} \leq c_d(n^*) \left(\frac{n^*}{E[n^*]}\right)^{1/d}.$$

Since the product of two sequences that converge in probability also converges in probability (to the product of the two limits), the right-hand side of the inequality converges in probability to c_d .

In order to show the reverse inequality, we proceed as in [Bollobás and Winkler 1988] and introduce subcubes of side $1/t$ along the diagonal, given by $D_i = [(i-1)/t, i/t]^d$ (Figure 2). Suppose that t is chosen large enough so that each subcube is in the region $V_d(r)$, i.e., $1/t < r$. Chains in these subcubes combine to form chains in $V_d(r)$. Let a maximal




FIGURE 2. Diagonal subcubes of side $1/t < r$.

chain in the subcube i be denoted $H_d(n^*, i)$. It is then clear that $H_d(n, r) \geq \sum_i H_d(n^*, i)$; therefore

$$c_d(n, r) \geq \frac{\sum_{i=1}^t H_d(n^*, i)}{E[n^*]^{1/d}} = \frac{\sum_{i=1}^t H_d(n^*, i)}{(n^*)^{1/d}} \left(\frac{n^*}{E[n^*]}\right)^{1/d}.$$

Bollobás and Winkler [1988] have shown that

$$\sum_i \frac{H_d(n^*, i)}{(n^*)^{1/d}}$$

converges in probability to c_d ; hence the right-hand side of the last equality converges in probability to c_d , proving that the inequality $c_d(n, r) \geq c_d$ holds in probability as $n \rightarrow \infty$. This concludes the proof that $c_d(n, r)$ converges in probability to c_d . \square

Essentially, Theorem 2.1 states that for every $r > 0$, there is a sufficiently large n such that a set of n points generated uniformly in a region close to the diagonal “represents” a much larger “virtual set” generated uniformly in the whole cube. Thus, computing the height of this smaller set would be equivalent (asymptotically) to computing the height of the larger virtual set. This result has several interesting implications for our experiments. The first is, of course, that we can restrict our search for a longest chain to points close to the diagonal. The second is that there is no need to generate points in the whole cube. Since we are interested exclusively in the points that are close to the diagonal, we only need to generate these diagonal points in a way that would be consistent with having generated a larger number of points in the entire cube. A third, and crucial, implication is that if one fixes an $r > 0$, then eventually, for large enough n , a maximal chain should fall in this region close to the diagonal.

Thus, for a fixed $r > 0$, a Monte Carlo simulation to determine $\bar{c}_d(n, r)$ can be used to suggest a value for c_d , provided n is sufficiently large. One might hope that the experiment with a large and infeasible n^* that yields a good estimate for c_d would be equivalent to an experiment with a feasible $n \approx n^* V_d(r)$, for some $r > 0$. Thus, our general approach will be to estimate $E[c_d(n, r)]$ for various n and r by Monte Carlo simulation, and use these values to construct a final estimate for c_d .

3. BOOSTING

Not every combination of n and r is computationally equivalent to an $n^* = \lceil n/V_d(r) \rceil$. In fact, for fixed n , as $r \rightarrow 0$, $c_d(n, r) \rightarrow (rn)^{1-1/d}d^{1/d}$, which approaches zero. This simply indicates that r is too small for the given n , and it is unlikely that a maximal chain of a random set with n^* points in \mathbf{V}_d is located inside $\mathbf{V}_d(r)$. On the other hand, if n is fixed and r is sufficiently large but less than 1, it is likely that about n out of n^* points fall in $\mathbf{V}_d(r)$ and a maximal chain exists within even a smaller distance of the diagonal than r . These observations suggest that given n , there exists an optimal $r = r_{\text{opt}}(n)$, which yields a maximal value for $c_d(n, r_{\text{opt}})$. (In fact, our experiments suggest that $c_d(n, r)$ has a single maximum as a function of r .) In the next section, we present an algorithm for computing $c_d(n, r)$ which requires $O(rnd)$ memory and runs in $O(rn^2d)$ time. Thus, it is of practical importance to select as small an r as possible for which $c_d(n, r)$ approximates c_d . We are thus led to the following experiment

Experiment 2 (Boosting).

1. Select a maximal feasible n .
2. Set $r = 1$, and obtain an estimate for $\bar{c}_d(n, r)$.
3. Decrease r ($r \leftarrow r - \varepsilon$, for a suitably chosen ε) and re-estimate $c_d(n, r)$. Repeat until $\bar{c}_d(n, r)$ stops increasing.
4. Find $r_{\text{opt}}(n)$, the r after which $\bar{c}_d(n, r)$ started to decrease.
5. If $r_{\text{opt}}(n)n^2d$ is much less than the available resources, increase n ($n \leftarrow n + n_0$, for a suitably chosen n_0). Compute $\bar{c}_d(n, r)$. Go back to step 3.
6. Once a maximal n has been reached and its associated $r_{\text{opt}}(n)$ has been determined, compute $\bar{c}_d(n, r_{\text{opt}})$ to the desired accuracy. Output this value as the estimate for c_d .

As the name of Experiment 2 suggests, given a maximal n for Experiment 1 (determined by computational resources), we can *boost* up to a higher n^* by going to a smaller r_{opt} and, perhaps, further increasing n as described in Experiment 2, while still using the same amount of computational resources. Unfortunately, boosting does not quite yield a computationally feasible experiment for estimating c_d . The root of the problem is that $r_{\text{opt}}(n)$ is not sufficiently small, thus the corresponding n^* is not sufficiently

large. Additionally, due to statistical fluctuations, $c_d(n, r)$ needs to be computed many times to obtain a reliable estimate for $\bar{c}_d(n, r)$. Furthermore, in order to obtain $r_{\text{opt}}(n)$ to a reasonable accuracy, $\bar{c}_d(n, r)$ needs to be computed for many r , hence making the determination of $r_{\text{opt}}(n)$ computationally intensive, even for a feasible n . As an example, for $d = 3$, our simulations indicate that $r_{\text{opt}}(n)$ lies in the range $[0.08, 0.11]$; thus, to reach $n^* = 10^{10}$, we would need to process sets of size close to 10^8 or larger. The result is that boosting to an acceptable n^* is still “close” to infeasible.

4. COCONVERGENCE

Theorem 2.1 suggests yet another approach to estimating c_d . Select a set $\{r_i\}_{i=1}^T$ of values for r and consider T sequences $\{\bar{c}_d(n, r_i)\}$ for $i = 1, \dots, T$. By Theorem 2.1, each of these sequences converges to the same limit c_d as $n \rightarrow \infty$. The problem now is one of estimating the common limit of these T sequences. Suppose $\bar{c}_d(n)$ ($= c_d(n, 1)$) converges at some rate $f(n)$ to c_d :

$$\bar{c}_d(n) = c_d - f(n).$$

Figure 2 suggests that chains in $\mathbf{V}_d(r)$ can be constructed by summing chains in the subcubes of side r along the diagonal. Hence the convergence of the chains in $\mathbf{V}_d(r)$ should be closely related to the convergence of the sum of chains in the diagonal subcubes, which in turn should be closely related to the convergence in each subcube. This convergence is governed by $f(n)$ and, thus, we are led to *conjecture* that the order of convergence of $c_d(n, r)$ for all $r > 0$ is essentially the same.

Conjecture. *There exists a function $\mu(r) > 0$, such that, for all $r > 0$,*

$$\bar{c}_d(n, r) = c_d - \mu(r)f(n) + o(f(n)). \quad (4-1)$$

The goal of Experiment 2 was to obtain the smallest r , namely $r_{\text{opt}}(n)$, for which $\mu(r_{\text{opt}})$ is a minimum. According to our conjecture, $r_{\text{opt}}(n)$ should be independent of n , at least asymptotically, and this could be the first test of the conjecture. The conjecture suggests a strong interdependence between the sequences, and we might be able to exploit this interdependence in order to get a more accurate estimate of c_d . The traditional approach to obtaining

the convergence point of the sequence $c_d(n)$ would be to assume that $f(n)$ has a certain form and then obtain a value for c_d consistent with this assumption and with the observed values $\bar{c}_d(n)$. The success of this kind of approach depends largely on the validity of the assumption on $f(n)$. Our conjecture allows us to estimate c_d without estimating $f(n)$. Given n and distinct r_i, r_j , (4-1) implies that the equalities

$$\begin{aligned} \bar{c}_d(n, r_i) &= c_d - \mu(r_i)f(n) + o(f(n)), \\ \bar{c}_d(n, r_j) &= c_d - \mu(r_j)f(n) + o(f(n)) \end{aligned}$$

hold simultaneously. Resolving this system with respect to $f(n)$, we have

$$\begin{aligned} \bar{c}_d(n, r_i) &= \\ (1 - A(r_i, r_j))c_d + A(r_i, r_j)\bar{c}_d(n, r_j) + o(f(n)), \end{aligned} \quad (4-2)$$

where $A(r_i, r_j) = \mu(r_i)/\mu(r_j)$. In other words, for all $i, j \in [1, T]$, $\bar{c}_d(n, r_i)$ linearly depends on $\bar{c}_d(n, r_j)$, up to $o(f(n))$. Further, from the functions $\bar{c}_d(n, r_i)$ and $\bar{c}_d(n, r_j)$, one can estimate $A(r_i, r_j)$ from the slope of this dependence, and $(1 - A(r_i, r_j))c_d$ from the intercept. c_d can then be obtained by dividing this intercept by $(1 - \text{slope})$. We have thus constructed an estimate for c_d without having to make any statements about $f(n)$. Further, from (4-2), we see that the convergence to linear behavior is at a rate $o(f(n))$, whereas the convergence of the sequences themselves is at the rate $f(n)$. Thus, the linear behavior will materialize at smaller n than the actual convergence. Hence, we expect to extract more accurate estimates for c_d in this way, given the computational resources. We are thus led to the following experiment.

Experiment 3 (Coconvergence).

1. Select a set $\mathcal{T} = \{r_1, \dots, r_T\}$ of values for r .
2. Select a set $\mathcal{N} = \{n_1, \dots, n_L\}$ of values for n . Let $\mathcal{N}_h = \{n_1, \dots, n_h\}$ for $h = 1, \dots, L$.
3. Compute $\bar{c}_d(n, r)$, for all $r \in \mathcal{T}$ and all $n \in \mathcal{N}$.
4. For all $i, j \in \{1, 2, \dots, T\}$ with $i \neq j$, perform an analysis on the pair of sequences $\{\bar{c}_d(n, r_i)\}$ and $\{\bar{c}_d(n, r_j)\}$ for $n \in \mathcal{N}_h$ to obtain the slope, $A_d(h, i, j)$, and the intercept, $B_d(h, i, j)$, for $h = 2, \dots, L$.
5. Evaluate $c_d(h, i, j) = B_d(h, i, j)/(1 - A_d(h, i, j))$ for $h = 1, \dots, L$. (Here $c_d(h, i, j)$ is a sequence of estimates for c_d ; this sequence should converge to

c_d at the rate at which the linear behavior arises, i.e., $o(f(n))$.)

6. Compute $e_d(i, j)$, the value to which $c_d(h, i, j)$ converges with respect to h .
7. Repeat steps 4-6 to compute $e_d(i, j)$ for all $\binom{T}{2}$ distinct pairs $i, j \in \mathcal{T}$.
8. Set $L_d = \min_{i, j \in \mathcal{T}} e_d(i, j)$, $U_d = \max_{i, j \in \mathcal{T}} e_d(i, j)$.
9. Output the interval $[L_d, U_d]$ as an estimate for c_d .

Since the success of Experiment 3 largely depends on the validity of Conjecture (4-1), we discuss the motivation for the conjecture and how one might experimentally verify it. Consider the quantity $c_d(n, r)$. Keeping in mind that “ n points in $\mathbf{V}_d(r)$ ” is “probabilistically” equivalent to “ n^* points in \mathbf{V}_d ”, and that $H_d(n, r) \leq H_d(n^*)$, we conclude that

$$\bar{c}_d(n, r) \leq c_d(n^*) = c_d - f(n/V_d(r)).$$

A lower bound can be obtained by considering the $1/r$ diagonal subcubes inside $\mathbf{V}_d(r)$ (see Figure 2). There are about $n_c = nr^d/V_d(r)$ points in each subcube, with heights given by $H_c = (n_c)^{1/d}(c_d - f(n_c))$. Since $H_d(n, r) \geq H_c/r$, we have

$$c_d(n, r) \geq c_d - f(nr^d/V_d(r)).$$

Thus, setting $\lambda_r = 1/V_d(r)$ and $\gamma_r = r^d\lambda_r$, we conclude (heuristically) that

$$c_d - f(\gamma_r n) \leq c_d(n, r) \leq c_d - f(\lambda_r n),$$

motivating the observation that $c_d(n, r)$ converges at essentially the same rate for all r . For large r , we expect the upper bound to be somewhat tight, and for small r , we expect the lower bound to be somewhat tight. Hence, starting from large r , by decreasing r , we initially expect to get better values for c_d , and then continuing to decrease r leads to worse performance, indicating that an optimal r exists. Further, the bounds indicate that $f(\lambda_r n) \leq f(\gamma_r n)$, suggesting that $f(\cdot)$ is monotonically decreasing (or that $c_d(n)$ is increasing in n). These were exactly the observations that led to boosting (Experiment 2).

We now discuss the experimental verification of the conjecture. We have seen that the conjecture implies a linear relationship, up to $o(f(n))$. Suppose that we observe such a relationship; can we then conclude that the conjecture holds? The following theorem tells us that we can.

Theorem 4.1. *Define a set of sequences $c(n, r)$ for every $r \in (0, 1]$ and suppose that each sequence converges with respect to n to a value c independent of r . Suppose also that $c(n, r)$ is differentiable with respect to r . The following two statements are equivalent.*

1. $c(n, r) = c - \mu(r)f(n)$ for some functions $\mu(r) > 0$ and $f(n)$ and for all $r \in (0, 1]$.
2. $c(n, r_1) = c(1 - A(r_1, r_2)) + A(r_1, r_2)c(n, r_2)$ for some differentiable function $A(r, r')$ and for all distinct pairs $r_1, r_2 \in (0, 1]$.

Proof. That $1 \implies 2$ has been demonstrated with $A(r_1, r_2) = \mu(r_1)/\mu(r_2)$ (the assumptions imply that μ is differentiable). We show that $2 \implies 1$. Suppose that the second statement holds. Then

$$\begin{aligned} c_d(n, r_1) - c &= A(r_1, r_2)(c_d(n, r_2) - c) \\ &= A(r_1, r_2)Q(n, r_2), \end{aligned} \tag{4-3}$$

where $Q(n, r_2) = c_d(n, r_2) - c$. The left-hand side is independent of r_2 , so taking the derivative of the log of both sides, we get

$$\frac{d}{dr_2} \log A(r_1, r_2) = -\frac{d}{dr_2} Q(n, r_2)$$

for all $r_1, r_2 \in (0, 1]$. The right-hand side is independent of r_1 ; therefore so is the left-hand side. The left-hand side is independent of n , and thus so is the right-hand side. Hence, they are both equal to some function $\zeta(r_2)$. Let $\alpha(r_2) = \int \zeta(r_2)$. After an integration, we conclude that for some functions $\kappa(r_1)$ and $g(n)$,

$$\begin{aligned} \log A(r_1, r_2) &= \alpha(r_2) + \log \kappa(r_1), \\ \log Q(n, r_2) &= -\alpha(r_2) + \log g(n), \end{aligned}$$

from which $A(r_1, r_2)Q(n, r_2) = \kappa(r_1)g(n)$. Substituting into (4-3) we find

$$c(n, r_1) = c + \kappa(r_1)g(n).$$

Taking $\mu = -\kappa$ and $f = g$ concludes the proof. \square

Thus, not only does the conjecture imply a very specific linear relationship between all pairs of sequences, but, if for some value of c , that same linear relationship holds for all pairs of sequences (i.e., all $r_i \neq r_j$), then the conjecture holds. Experimentally, this means that $c(h, r_i, r_j)$ in Experiment 3 should converge (with respect to h) to the same value for all distinct pairs r_i, r_j , or, equivalently, the interval $[L_d, U_d]$ should be small. If this were to happen,

then not only would we provide experimental justification for the conjecture, but we would also obtain an accurate estimate for c_d .

5. ALGORITHMS

Both experimental approaches described in the previous section try to get more accurate estimates for c_d by effectively “accessing” higher n without actually computing with the higher n . Both techniques rely on efficient algorithms for computing $c_d(n, r)$.

The first task is to generate n points chosen independently from a uniform distribution in $V_d(r)$. A trivial algorithm is the one that generates random points in V_d and keeps only those that fall in $V_d(r)$, continuing until n points in $V_d(r)$ have been generated. This can be highly inefficient, especially if r is small, as the acceptance rate will be extremely small. A more efficient approach is to generate random points in $V_d(r)$ itself. This can be done quite efficiently, and further, it is possible to generate the points in a sequential manner so that the dynamic programming algorithm for computing $H_d(n, r)$ keeps in memory only a small portion of the total of n points, those that are necessary for executing the algorithm.

The second task is to compute $H_d(n, r)$. The standard dynamic programming algorithm has computational complexity $O(dn^2)$ [Cormen et al. 1990]. However, since the points are generated in a sequential manner, it is possible to design an algorithm that maintains a *working set* that takes advantage of the sequential point generation. Operations need only be performed on this working set, resulting in a factor of r reduction in both computational complexity and memory requirements.

5A. Generation of Input Points

The generation of points is illustrated in Figure 3, for the case of $d = 2$. First, an “origin” point, $t(i) = (t(i), t(i), \dots, t(i))$, is generated along the diagonal. Then a point is generated from a uniform distribution on one of the hypercubes $p_k(i)$, $k = 1, \dots, d$, where $p_k(i)$ is the $(d-1)$ -dimensional cube with origin at $t(i)$ and the k -th coordinate set to zero. In this way one can generate n points in $V_d(r)$. The probability density of the origin point coordinate $t(i)$ is uniform up to $t(i) = 1 - r$ and

then decaying like $(1 - t(i))^{d-1}$ for $t(i) > 1 - r$. This density is shown in Figure 4, for the case $d = 3$, $r = 0.5$. Further, instead of generating the origin points in a random order, we generate the n order statistics for the origin points.




FIGURE 3. The sequential generation of points in $V_d(r)$ illustrated for the case $d = 2$.

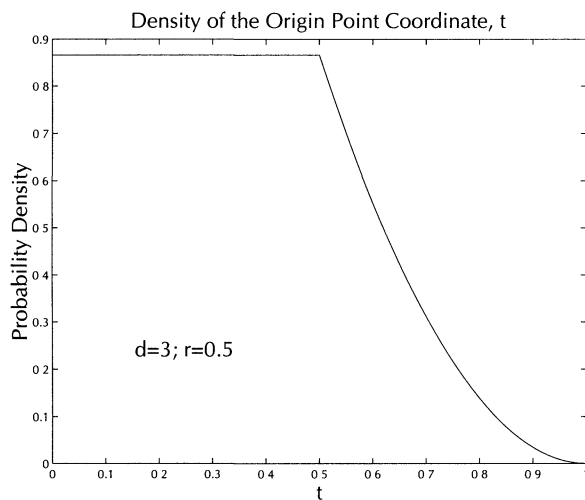


FIGURE 4. Probability density of origin points for $d = 3$, $r = 0.5$.

Algorithm 1 (Sequential generation of input points).

1. Set $i = 1$ and $t_{\text{prev}} = 0$.
2. Generate the i -th origin point

$$t(i) = (t(i), t(i), \dots, t(i)),$$

where $t(i) \in [t_{\text{prev}}, 1]$ is the i -th order statistic (given that the $(i-1)$ -th order statistic was at t_{prev}) of n points generated from the distribution of origin points such as the one in Figure 4.

3. Generate a random vector $=\{v_1, v_2, \dots, v_d\}$ with each v_i in $[0, \min\{r, 1-t_i\}]$.
4. Generate a random integer k from $\{1, 2, \dots, d\}$ and set $v_k = 0$.
5. Generate input point

$$x(i) = \{t(i) + v_1, t(i) + v_2, \dots, t(i) + v_d\}.$$

6. Set $t_{\text{prev}} = t(i)$, $i = i+1$ and go back to step 2 if $i \leq n$.

The detailed implementation of step 2 is given in Section 7A. The following considerations prove the correctness of the algorithm. One can show that step 2 generates the order statistics for the distribution of the origin coordinate, where the origin coordinate distribution is uniform from 0 to $1-r$ and then decaying like $(1-x)^{d-1}$ from $1-r$ to 1, at which point it is 0. The remainder of the algorithm then takes that as origin point and generates a point on one of the $(d-1)$ -dimensional faces of side r given by $x_j = 0$.

The next proposition states two (almost trivial) ordering properties that the points $x(1), \dots, x(n)$ have. While these properties might seem trivial, they are of key importance to the algorithm for computing the height of the n points in $V_d(r)$, as we shall see in the next section.

Proposition 5.1 (Ordering properties). *Suppose points $x(1), \dots, x(n)$ are generated according to the algorithm above. Then:*

1. If $t(i) \geq x(j)$, then $x(k) \geq x(j)$ for all $k \geq i$.
2. If $i < j$, it cannot be that $x(i) \geq x(j)$. More precisely, $P[x(i) \geq x(j)] = 0$.

Proof. The first claim follows because $x(k) \geq t(i)$ for $k \geq i$. Let the origin point of $x(i)$ be $t(i)$ and let $j > i$. One of the components of $x(i)$ is $t(i)$; therefore that component of $x(i)$ must be less than the corresponding component of $x(j)$ with probability 1, proving the second statement. \square

5B. Computing $H_d(n, r)$

Although the ordering of the origin points does not guarantee that the projected input points $\{x(1), x(2), \dots, x(n)\}$ are ordered, it guarantees the ordering properties stated in Proposition 5.1. Property 2 ensures that a newly generated point can never be below a previously generated point. Therefore, the

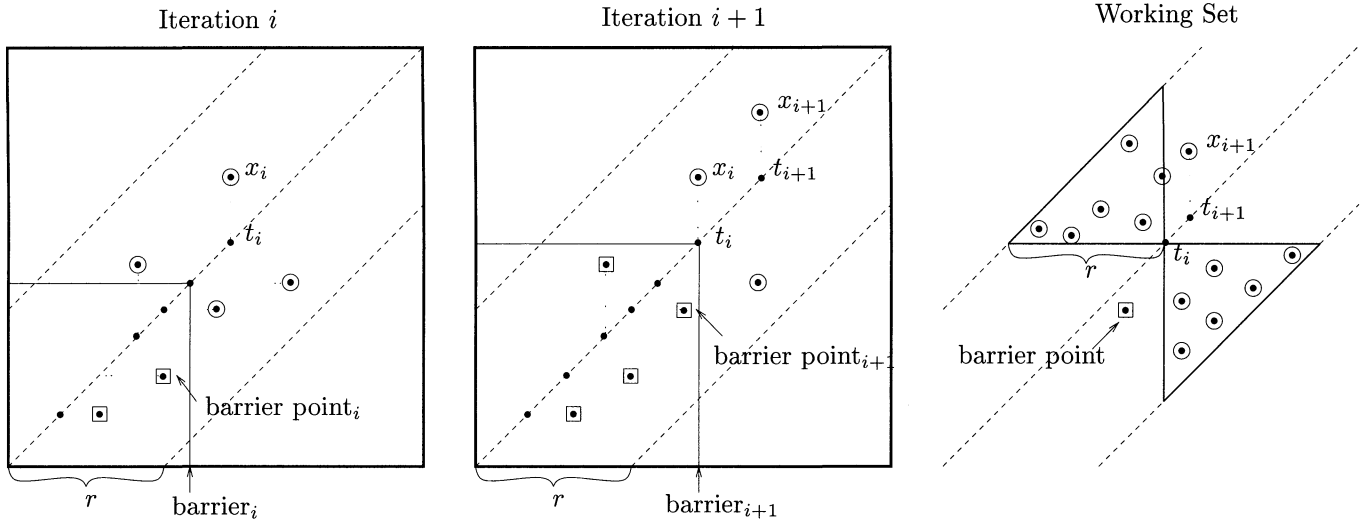


FIGURE 5. The *working set* at iteration i and $i + 1$. Key: \square represents a discarded input point, \odot represents a working set input point, and \bullet an origin point.

height of newly generated points can be immediately computed without considering future points. At iteration i , property 1 ensures $\mathbf{x}(i)$ has a height of at least $h(\mathbf{x}(j)) + 1$ for all $\mathbf{x}(j) < \mathbf{t}(i)$. Among these input points, the $\mathbf{x}(j)$ with maximum height is identified as the *barrier point* and the rest are discarded. The remaining input points (all $\mathbf{x}(j) \geq \mathbf{t}(i)$ plus the *barrier point*) define the *working set* for iteration $i + 1$. The height of $\mathbf{x}(i + 1)$ can be determined by inspecting only the points in the *working set*.

$H_d(n, r)$ is computed using the following procedure where t and x are the origin point and projected input point as generated by sequential generator described above. $h(x)$ is the height of the maximal chain ending at point x , and x_{barrier} is the *barrier point*.

Algorithm 2 (Computation of $H_d(n, r)$).

1. Set $h_{\text{max}} = 1$.
2. Using Algorithm 1, generate $t(1)$ and $\mathbf{x}(1)$ and add $\mathbf{x}(1)$ to the *working set*.
3. Set the height of $\mathbf{x}(1)$ to 1 ($h(\mathbf{x}(1)) = 1$).
4. For $i = 2$ to n
 - a. Using Algorithm 1, generate $t(i)$ and $\mathbf{x}(i)$.
 - b. Initialize the height of $\mathbf{x}(i)$ to 1 ($h(\mathbf{x}(i)) = 1$).
 - c. For every point $\mathbf{x}(j)$ in the *working set*
 - i. If $\mathbf{x}(i)$ is above $\mathbf{x}(j)$ and $h(\mathbf{x}(j)) \geq h(\mathbf{x}(i))$ then set $h(\mathbf{x}(i)) = h(\mathbf{x}(j)) + 1$.
 - ii. If $\mathbf{x}(j)$ is below $\mathbf{t}(i)$, if the *barrier point* is defined, and if $h(\mathbf{x}(j)) > h(x_{\text{barrier}})$, then

remove the current *barrier point* and replace it with $\mathbf{x}(j)$.

- iii. If $h(\mathbf{x}(j)) \leq h(x_{\text{barrier}})$ then remove $\mathbf{x}(j)$ from the *working set*.
 - iv. If $\mathbf{x}(j)$ is below $\mathbf{t}(i)$ and the *barrier point* is not defined then set $\mathbf{x}(j)$ to be the *barrier point*.
 - d. Add $\mathbf{x}(i)$ to the *working set*.
 - e. If $h(\mathbf{x}(i)) > h_{\text{max}}$ then set $h_{\text{max}} = h(\mathbf{x}(i))$.
5. Return $h_{\text{max}}/n^{1/d}$.

The algorithm computes the height of each $\mathbf{x}(i)$ as the input points are generated. $\mathbf{x}(i)$ is given a default height of 1. Every input point forms a chain of height 1 with itself as the beginning and end points. The new input point is compared with every $\mathbf{x}(j)$ in the *working set* to find the highest chain which can be continued by $\mathbf{x}(i)$. Each $\mathbf{x}(j)$ is also compared with $\mathbf{t}(i)$ to determine if it is below the *barrier*. If $\mathbf{x}(j)$ is below the *barrier*, it either is removed from the working set or (in the case where $h(\mathbf{x}(j)) > h(x_{\text{barrier}})$) replaces the current *barrier point*. After iteration i , the *working set* includes only the *barrier point* and the points that are not below the *origin point* $\mathbf{t}(i)$ (see Figure 5). This defines a subset of the $\mathbf{V}_d(r)$ which has volume $O(\mathbf{V}_d(r))r$. Since n points are uniformly distributed in $\mathbf{V}_d(r)$, it is expected that the *working set* contains rn points. Since the algorithm must iterate over the *working set* for every newly generated point, the expected number of point comparisons is $O(rn^2)$. Therefore,

using a very small r not only increases the number of virtual points n^* but decreases the memory requirements and the number of computations by a factor of r .

Proposition 5.2 (Correctness of Algorithm 2). *Suppose that $h(\mathbf{x}(i))$, for $i = 1, \dots, n$, is computed according to Algorithm 2. Then:*

1. $h(\mathbf{x}(i))$ is equal to the height of the longest chain ending at $\mathbf{x}(i)$.
2. h_{\max} is equal to the height of the longest chain in $V_d(r)$.

Proof. At the first iteration of the algorithm, $\mathbf{x}(1)$ is assigned the correct height of 1. Each $\mathbf{x}(i + 1)$ is compared with all $\mathbf{x}(j)$, $j \leq i$, with the exception of the $\mathbf{x}(j) < \mathbf{t}(i)$ which are guaranteed by Proposition 5.1 to be below $\mathbf{x}(i + 1)$. The maximum height of these points is represented by the *barrier point*. Therefore, all possible maximum length chains ending at $\mathbf{x}(i + 1)$ are evaluated and $h(\mathbf{x}(i))$ is set to the maximum height. The second statement follows directly from the first by taking the maximum of all $h(\mathbf{x}(i))$. \square

6. EXPERIMENTS

We now summarize our results and the calculations used to obtain them. We discuss the case $d = 3$; for the other dimensions we just provide graphs and tabular data. We have performed computer simulations to support the following claims:

1. The coconvergence of $\{c_d(n, r)\}$, leading to estimates of $\{c_d\}$.
2. The (asymptotic) independence (and uniqueness) of $r_{\text{opt}}(n)$ with respect to n , for $d = 2, 3, 4, 5, 6$.

The coconvergence simulations aim to confirm that for every pair of radii, $r_i \neq r_j$, and sufficiently large n , there exist constants A and B , depending on r_i and r_j , such that

$$\bar{c}(n, r_i) = B + A\bar{c}(n, r_j) + \varepsilon.$$

The ε signifies an error term composed of the $o(f(n))$ term, which governs the rate of coconvergence, and the statistical fluctuations in measuring \bar{c} . These fluctuations can be made arbitrarily small by taking a large enough sample for computing \bar{c} . Having obtained A and B , $B/(1-A)$ is an estimate for c_d . We selected n and r according to the scheme

$$n \in \{ \lfloor 10^{2+0.1k} \rfloor : k \in [0, 40] \},$$

$$r \in \{0.01, 0.005, 0.001, 0.0005, 0.0001\}.$$

The number of trials used for computing $\bar{c}_d(n, r)$ was heuristically selected to obtain roughly the same degree of statistical accuracy (within the given computational requirements) for each n and r . In principle, the larger the number of trials for a given n and r , the more accurate the estimates will be.

Figure 7, left, exhibits the convergence behavior of $c_3(n, r)$, for the five values of r . Even for $n = 10^6$, it is not clear to what value these sequences are converging.





FIGURE 7. Convergence of $\bar{c}_3(n, r)$ (left) and of $c_3(k; i, j)$ (right).

(r_j, r_k)	$d = 2$	$d = 3$	$d = 4$	$d = 5$	$d = 6$
(.01, .005)	1.9985	2.3636	2.5144	2.5887	2.6105
(.01, .001)	1.9978	2.3627	2.5139	2.5864	2.6145
(.01, .0005)	1.9982	2.3629	2.5151	2.5868	2.6153
(.01, .0001)	1.9985	2.3628	2.5165	2.5892	2.6156
(.005, .001)	1.9999	2.3631	2.5207	2.5867	2.6158
(.005, .0005)	2.0001	2.3638	2.5195	2.5866	2.6156
(.005, .0001)	2.0006	2.3632	2.5192	2.5866	2.6167
(.001, .0005)	2.0010	2.3658	2.5204	2.5827	2.6118
(.001, .0001)	2.0015	2.3643	2.5186	2.5837	2.6096
(.0005, .0001)	2.0018	2.3663	2.5168	2.5832	2.6074

TABLE 1. Estimates of c_d obtained from various pairs of sequences for $d = 2, 3, 4, 5, 6$.

Given $c_d(n, r_i)$ and $c_d(n, r_j)$ for $n = n_1, \dots, n_h$, we can estimate $A_d(h, i, j)$ and $B_d(h, i, j)$ for $h = 2, \dots, L$ as in Experiment 3. The estimate given by $B(h, i, j)/(1 - A(h, i, j))$ should converge with respect to h to c_d , independent of which particular pair i, j is used. The behavior of this ratio is demonstrated in Figure 7, right, for $d = 3$. Section 7B explains how $A(h, i, j)$ and $B(h, i, j)$ were estimated and shows some sample behavior for the case $d = 3$.

It is clear that these curves are all converging to the same value, supporting our conjecture via Theorem 4.1. Further, the convergence occurs earlier as compared with the convergence of $c_3(n, r)$, in accordance with the expected $o(f(n))$ behavior. Thus, we propose a range for c_3 by taking the range of values to which these curves are converging. Similar plots are shown in Figure 9, for $d = 2, 4, 5, 6$. Table 1

summarizes the results, displaying the estimates for c_d obtained from each of the ten pairs of sequences for each $d = 2, 3, 4, 5, 6$.

As further evidence for our conjecture, we conducted experiments to investigate whether $r_{\text{opt}}(n)$ is independent of n . Our second series of simulations used these values of n and r :

$$n \in \{50000, 100000, 200000, 300000\},$$

$$r \in \{0.01, 0.02, 0.06, 0.08, 0.10, 0.14, 0.20, 0.30, 1.00\}.$$




FIGURE 8. Dependence of $\bar{c}_d(n, r)$ on r for various values of n .

For each pair (n, r) , over 350 samples were used to obtain $\bar{c}_d(n, r)$. The results are plotted in Figure 8. The figure shows that $r_{\text{opt}}(n)$, the radius that maximizes $c_d(n, r)$, appears to be independent of n , and, further, since the curves appear very nearly parallel, the function $\mu(r)$ itself appears to be independent of n , as required by the conjecture.





FIGURE 9. Convergence of $\bar{c}_d(n, r)$ and $c_d(k; i, j)$, for $d = 2$ (this page) and $d = 4, 5, 6$ (overleaf).

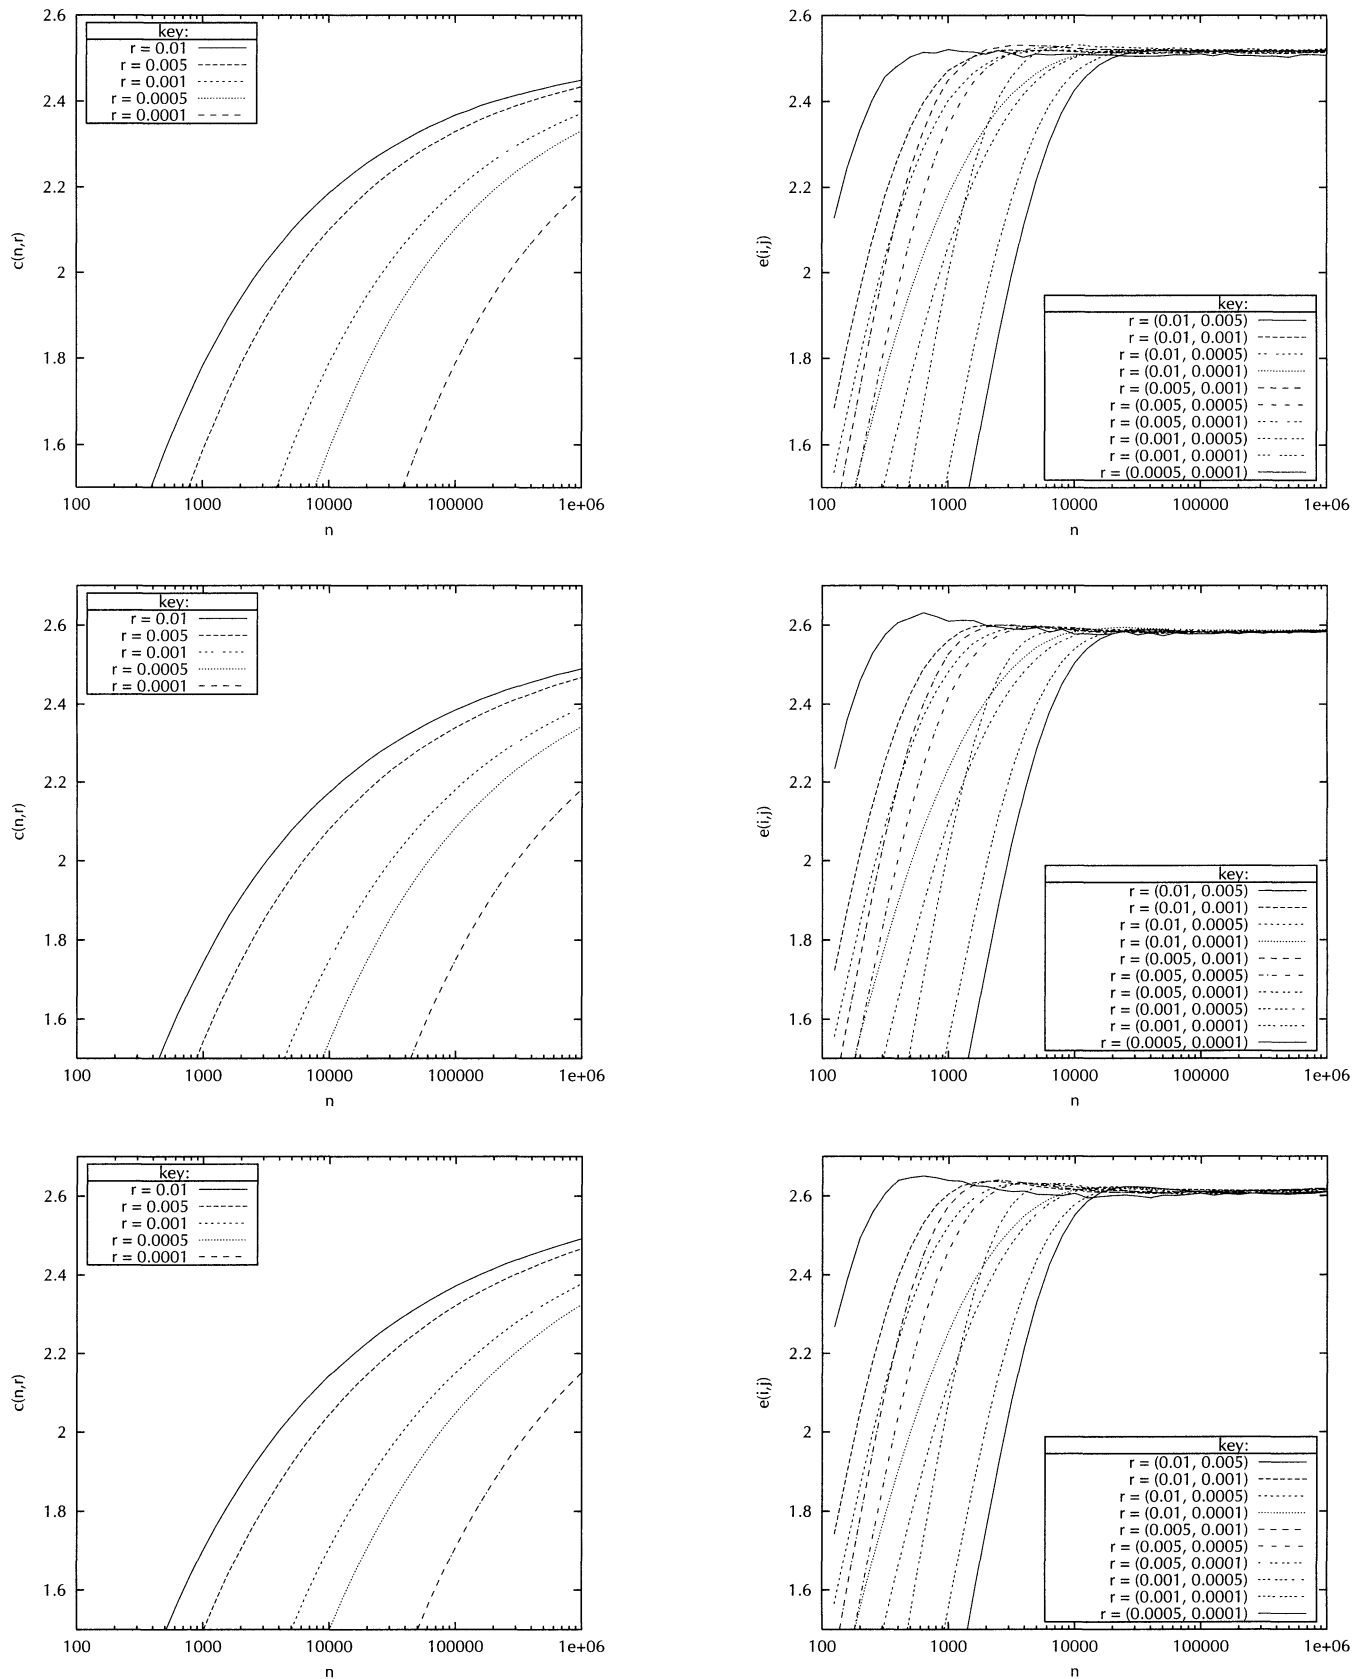


FIGURE 9 (continued). Convergence of $\bar{c}_d(n,r)$ and $c_d(k;i,j)$, for $d = 4, 5, 6$.

7. SOME TECHNICAL DETAILS

7A. Generating the Order Statistics for the Origin Points

Here we provide the details for the generation of the order statistics of the origin points. Suppose we wish to generate the i -th order statistic. Let t_{prev} be the coordinate for the $(i-1)$ -th order statistic. If $i = 1$ then $t_{\text{prev}} = 0$. Define

$$V_d(r) = r^d + dr^{d-1}(1-r);$$

$$F(l, r, d) = \begin{cases} \frac{dr^{d-1}l}{V_d(r)}, & 0 \leq l \leq 1-r, \\ 1 - \frac{(1-l)^d}{V_d(r)}, & 1-r \leq l \leq 1; \end{cases}$$

$$\Phi(u, r, d) = \begin{cases} \frac{V_d(r)u}{dr^{d-1}}, & 0 \leq u \leq 1 - \frac{r^d}{V_d(r)}, \\ 1 - (V_d(r)(1-u))^{1/d}, & 1 - \frac{r^d}{V_d(r)} \leq u \leq 1; \end{cases}$$

$$G^{-1}(u, l, n, k, r, d) = \Phi(1 - (1 - F(l, r, d))(1 - u)^{1/(n-k+1)}, r, d).$$

(Note that $\Phi(\cdot, r, d)$ is inverse to $F(\cdot, r, d)$.) Now generate uniformly a random number $u \in [0, 1]$. Then the i -th origin point, $\mathbf{t}(i) = (t(i), \dots, t(i))$, is given by

$$t(i) = G^{-1}(u, t_{\text{prev}}, n, i, r, d),$$

where n is the total number of points to be generated, r is the size of $V_d(r)$ and d is its dimension. A tedious but an elementary exercise (see [Billingsley 1986]) shows that this procedure does indeed generate the correct order statistics, and we do not include that analysis here.

7B. Method used to obtain $A_d(h, i, j)$ and $B_d(h, i, j)$

We briefly elaborate on step 4 of Experiment 3 (co-convergence). For a pair of sequences given by r_i, r_j , let

$$\Delta_d(k, i) = \bar{c}_d(n_k, r_i) - \bar{c}_d(n_{k-1}, r_i),$$

$$a_d(k; i, j) = \frac{\Delta_d(k, i)}{\Delta_d(k, j)} \quad (k = 1, \dots, 40).$$

Assuming that constants $A_d(i, j)$ and $B_d(i, j)$ exist, $a_d(k; i, j)$ would approximate $A_d(i, j)$. However, the $o(f(n))$ term and the statistical noise require that we take a more robust approach to get more reliable results. We, therefore, use the following procedure

to approximate $A_d(i, j)$ by a sequence $A_d(h, i, j)$ for $h = 2, \dots, L$.

Obtaining $A_d(h, i, j)$. For every $g = 2, \dots, h$, compute the average $A^{(g)}$ of $\{a_d(k; i, j)\}_{k=g}^h$. Compute the standard deviation $\delta^{(h)}$ of $\{a_d(k; i, j)\}_{k=g}^h$. Select the smallest $g = g_0$ for which $\delta^{(g)}$ is the minimum to within 0.01. Output $A^{(g_0)}$ as an approximation for $A_d(h, i, j)$.

The approach is designed to accommodate the statistical noise and the fact that the convergence of the $a_d(k, i, j)$'s might not have occurred at the beginning of the sequence.

The procedure for approximating $B_d(i, j)$ is similar to the one for $A_d(i, j)$:

Obtaining $B_d(h, i, j)$. Compute $b_d(k; i, j) = c_d(n_k, r_i) - A_d(k, i, j)c_d(n_k, r_j)$. For every $g = 2, \dots, h$, compute the average $B^{(g)}$ of $\{b_d(k; i, j)\}_{k=g}^h$. Compute the standard deviation $\delta^{(h)}$ of $\{b_d(k; i, j)\}_{k=g}^h$. Select the smallest $g = g_0$ for which $\delta^{(g)}$ is the minimum to within 0.01. Output $B^{(g_0)}$ as an approximation for $B_d(h, i, j)$.

Figure 10 demonstrates the behavior of $A(h, i, j)$ and $B(h, i, j)$ for $d = 3$. As can be seen, convergence occurs quite soon. The cases $d = 2, 4, 5, 6$ are similar and do not convey additional insight.

8. CONCLUSIONS

We hope that the observations we have made, based upon our experiments, can be rigorously proved. Our data seems to provide ample support for the Conjecture on page 587, but a rigorous justification is certainly desirable. The translated cube $V_d(r)$ was chosen on the grounds of simplicity. The analog of the Bollobás and Winkler's theorem can be proved for similar diagonal volume elements with slightly different shapes. It is not clear what the optimal such shape is, in terms of rate of convergence to c_d with respect to n . The behavior of c_d with respect to d (other than being monotonic and converging to e) does not appear to display a simple enough behavior that would lead to an immediate conjecture for the constant c_d as a function of d . We hope that further experiments will provide some guidelines toward a theoretical understanding

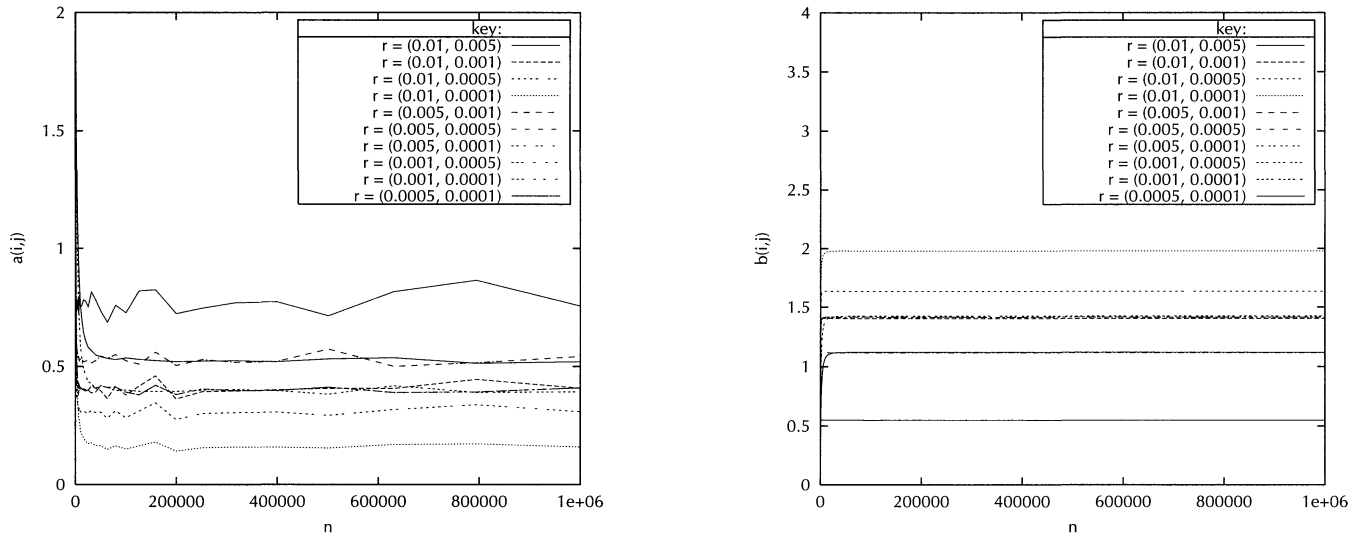


FIGURE 10. Convergence of the estimates for $A_d(k; i, j)$ (left) and $B_d(k; i, j)$ (right) for $d = 3$.

of the behavior of the maximal chains, and of the constants c_d .

ACKNOWLEDGEMENT

We are grateful to Catherine McGeogh for introducing us to the problem and sharing the information about her experiments.

REFERENCES

[Baer and Brook 1968] R. M. Baer and P. Brook, “Natural sorting over permutation spaces”, *Mat. Comp.* **22** (1968), 385–510.

[Baik et al. 1999] J. Baik, P. Deift, and K. Johansson, “On the distribution of the length of the longest increasing subsequence of random permutations”, *J. Amer. Math. Soc.* **12**:4 (1999), 1119–1178.

[Billingsley 1986] P. Billingsley, *Probability and measure*, Wiley, New York, 1986.

[Bollobás and Winkler 1988] B. Bollobás and P. Winkler, “The longest chain among random points in euclidean space”, *Proc. Amer. Math. Soc.* **103**:2 (1988), 347–353.

[Borodin 1999] A. Borodin, “Longest increasing subsequences of random colored permutations”, *Electr. J. Comb.*, **6**:1 (1999), R13.

[Breimer et al. 2001] E. Breimer, M. Goldberg, B. Kolstad, and M. Magdon-Ismail, “Experimental evaluation of the height of a random set of points in a d-dimensional cube”, *3rd Workshop on Algorithm Engineering and Experiments (ALENEX 01)*, Jan. 5–6, 2001.

[Cormen et al. 1990] T. Cormen, C. Leiserson, R. Rivest, *Introduction to Algorithms*, MIT Press, Cambridge (MA), 1990.

[Erdős and Szekeres 1935] P. Erdős and G. Szekeres, “A combinatorial problem in geometry”, *Compositio Math.* **2** (1935), 463–470.

[Hammersley 1972] J. M. Hammersley, “A few seedlings of research”, *Proc. Sixth Berkeley Sympos. Math. Stat. Prob.*, Univ. of California Press, Berkeley, 345–394.

[Logan and Shepp 1977] B. Logan and L. Shepp, “A variational problem for random Young tableaux”, *Adv. Math.* **26** (1977), 206–222.

[Odlyzko and Rains 2000] A. M. Odlyzko and E. M. Rains, “On longest increasing subsequences in random permutations”, pp. 439–451 in *Analysis, Geometry, Number Theory: The Mathematics of Leon Ehrenpreis* (Philadelphia, 1998), Contemp. Math. **251**, edited by E. Grinberg et al., AMS, Providence, 2000.

[Pilpel 1990] S. Pilpel, “Descending subsequences of random permutations”, *J. Comb Theory A*, **53**:1 (1990), 96–116.

[Steele 1977] J. M. Steele, “Limit properties of random variables associated with a partial order in \mathbb{R}^d ”, *Ann. Probab.* **5** (1977), 179–180.

[Ulam 1961] S. M. Ulam, “Monte Carlo calculations in problems of mathematical physics”, pp. 261–281 in *Modern mathematics for engineers*, edited by E. F. Beckenbach, McGraw-Hill, New York, 1961.

[Veršik 1977] A. M. Veršik and S. V. Kerov, “Asymptotics of the Plancherel measure of the symmetric group and limiting form of Young tableaux” (Russian), *Dokl. Acad. Nauk SSSR* **233**:6 (1977), 1024–1028.

Eric Breimer, Computer Science Department, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, United States (breime@cs.rpi.edu)

Mark Goldberg, Computer Science Department, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, United States (goldberg@cs.rpi.edu)

Brian Kolstad, Computer Science Department, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, United States (kolstb@cs.rpi.edu)

Malik Magdon-Ismail, Computer Science Department, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, United States (magdon@cs.rpi.edu)

Received September 27, 2000; accepted May 2, 2001

