*Journal of Global Optimization* **10:** 439–465, 1997.



# Performance of Neural Net Heuristics for Maximum Clique on Diverse Highly Compressible Graphs

ARUN K. JAGOTA*
*Department of Computer Science, University of California, Santa Cruz, U.S.A.*

KENNETH W. REGAN**
*Department of Computer Science, State University of New York at Buffalo, U.S.A.*

**Abstract.** The problem of finding the size of the largest clique in an undirected graph is NP-hard, even to well-approximate, in the worst case. Simple algorithms, including some we study here, work quite well however, on graphs sampled from $\mathbf{u}(n)$, the uniform distribution on $n$-vertex graphs. It is felt by many, however, that $\mathbf{u}(n)$ does not accurately reflect the nature of instances that come up in practice. It is argued that when the actual distribution of instances is unknown, it is more appropriate to suppose that instances come from the Solomonoff–Levin or *universal* distribution $\mathbf{m}(x)$ instead, which assigns higher weight to instances with shorter descriptions (i.e., to those that are structured or compressible). We extend a theorem of Li and Vitanyi to show that the average-case performance ratio of any approximation algorithm on random instances drawn from $\mathbf{m}(x)$ has the same asymptotic order as its worst-case performance ratio. Because $\mathbf{m}(x)$ is neither computable nor samplable, we employ a realistic analogue $\mathbf{q}(x)$ which lends itself to efficient empirical testing. We experimentally evaluate how well certain neural network algorithms for Maximum Clique perform on graphs drawn from $\mathbf{q}(x)$, as compared to those drawn from $\mathbf{u}(n)$. The experimental results are as follows. All nine algorithms we evaluated performed roughly equally-well on $\mathbf{u}(n)$, where as three of them — the simplest ones — performed markedly poorer than the other six on $\mathbf{q}(x)$. Our results suggest that $\mathbf{q}(x)$, while postulated as a more realistic distribution to test the performance of algorithms than $\mathbf{u}(n)$, also discriminates their performance better. Our $\mathbf{q}(x)$ sampler can be used to generate compressible instances of any discrete problem.

**Key words:** Universal distribution, compressible data, heuristic algorithms.

## 1. Introduction

The MAX-CLIQUE problem is to compute the size $\omega(G)$ of the largest clique (i.e., complete subgraph) in a given graph $G$, and further to find a clique of that size. It has long been known that $\omega(G)$ is NP-hard to compute exactly, so interest has centered on approximating $\omega(G)$ closely enough to suit the many applications which can be formulated in terms of MAX-CLIQUE. These include constraint-satisfaction, object-recognition, and other real-world problems. Many approximation heuristics have resulted from the study of neural networks for solving hard combinatorial optimization problems (see [2; 17]) and are applicable to MAX-CLIQUE. For an extensive survey on MAX-CLIQUE, its applications, and algorithms for it, see [16].

* The first author was involved in this research at the Dept of Computer Science, SUNY-Buffalo
** Supported in part by NSF Grant CCR-9011248.

Recently it has been shown that even calculating $\omega(G)$ to within a given constant factor is NP-hard [1]. In fact, the main result of [1] gives a fixed $\epsilon > 0$ such that if there is a polynomial-time algorithm $A$ which gives $\omega(G)/A(G) \leq n^{\epsilon}$ for all $G$ of sufficient size $n$, then NP = P. (Here $A(G)$ denotes the clique size returned by algorithm A.) Nevertheless, there are simple heuristics which come to within a factor of 2 of optimal for "most" graphs, in the sense of the uniform distribution $\mathbf{u}(n)$. This distribution is defined on $n$-vertex undirected graphs, by letting each edge $(i, j)$, $(1 \leq i < j \leq n)$, exist independently at random with probability 1/2. When the probability is $p$ instead of 1/2, we denote the distribution as $\mathbf{u}_p(n)$ instead. For example, Theorem 8 of R. Karp [11] implies that for any $\epsilon > 0$ and sufficiently large $n$, the heuristic we call SD-$\emptyset$ gives an expected performance ratio $E[\omega(G)/SD - \emptyset(G) : \mathbf{u}_p(n)]$ which is less than $2 + \epsilon$.

However, it is felt in many quarters that the uniform distribution does not accurately reflect the nature of instances which come up in practice. M. Li and P. Vitanyi [12] argue that when the actual distribution of instances is unknown, it is most appropriate to suppose that they come from the non-computable Solomonoff–Levin or *universal* distribution $\mathbf{m}(x)$. One reason is that every computable distribution is majorized by a constant multiple of $\mathbf{m}(x)$. Another reason related to *Occam's Razor* is that the objects occurring most often in nature or in practice have short descriptions. A *description* of $x$ is a program $P$ and an argument $y$ such that $P(y) = x$; the description is *short* if the total bit length of $Py$ is appreciably less than the bit length $|x|$ of $x$. The shorter descriptions $x$ has, the more weight is given to $\mathbf{m}(x)$. By contrast, uniform distributions favor instance strings $x$ which are incompressible; i.e. whose shortest descriptions are essentially "PRINT $x$". Li and Vitanyi [13] show that with respect to $\mathbf{m}(x)$, the average-case running time of any algorithm whatsoever has the same order as its worst-case running time. In Section 2 we obtain the same result with performance ratios by approximation algorithms in place of running time.

The main purpose of our work is to test whether certain algorithms perform noticeably worse under distributions weighted toward compressible inputs. We also wish to ascertain which algorithms perform poorly on such distributions and which perform well. Section 3 describes our efficient approximation $\mathbf{q}(x)$ to $\mathbf{m}(x)$. Section 4 describes the algorithms we tested, ranging from simple to sophisticated neural network heuristics. All these algorithms were tested extensively on random graphs in [9], and several of them on a variety of graphs with different kinds of structures in [10], in which paper extensive comparisons were made with other algorithms. Section 5 describes the experimental methodology, which includes a description of the test graphs, a discussion of the criterion for evaluating performance, and a description of the parameter values of certain parametrized algorithms. In Section 6 we present the experimental results and their analysis.

The main contributions of this paper are:

– Theorem 1, extending Li and Vitanyi's result to approximation performance ratio.

– A method for testing algorithms under an efficient approximation to the universal distribution. To our knowledge, this is the first concerted effort at testing under such non-uniform distributions. The resulting generator, $\mathbf{q}(x)$, provides a reasonably efficient means of generating a diverse set of application-independent graphs, amongst which lie some quite hard graphs for MAX-CLIQUE – harder, as experiments and conceptual arguments indicate, than random graphs.
– Experimental results (Section 6) which show that, while all nine heuristics performed roughly equally-well under $\mathbf{u}(n)$, three of the simpler ones performed markedly poorer than the other six under $\mathbf{q}(x)$. For example, SD-$\emptyset$, the simplest heuristic, retrieved on average a clique size within a factor of 1.35 of the one retrieved by the best heuristic, on graphs drawn from $\mathbf{u}(n)$ (see Table 4). By contrast, SD-$\emptyset$ was on average poorer by a factor of 2.42 and 7.56 than the best heuristic on graphs drawn from $\mathbf{q}(100)$ and $\mathbf{q}(400)$ respectively (see Table 4).
– Experimental results (Section 6) which reveal some interesting characteristics of certain individual algorithms. For example, a simple variant of SD-$\emptyset$, only a very little more sophisticated than SD-$\emptyset$, ended up performing much better on $\mathbf{q}(x)$. One of the continuous neural network methods, MFA, which has been claimed in the neural network literature to work well in practice, did indeed work very well on $\mathbf{q}(x)$. The other continuous method, CHD, was discernibly poorer than MFA but significantly better than SD-$\emptyset$ and the other two simplest algorithms, on $\mathbf{q}(x)$.

## 2. Theoretical Work

Consider any maximization problem $\pi$, on a graph, which employs a quality measure $q : \mathcal{G}_n \rightarrow Z^+$. For any algorithm A for problem $\pi$, define

$$wc_\pi(A, m) = \max_{x \in \mathcal{G}_n} \frac{\pi_{opt}(x)}{A(x)}$$

where $m = \binom{n}{2}$, $x \in \{0, 1\}^m$ is the usual bitstring representation of a graph in $\mathcal{G}_n$, $\pi_{opt}(x)$ is the quality of the optimum solution in $x$ and $A(x)$ is the quality of the solution in $x$ found by algorithm A.

Clearly $wc_\pi(A, m)$ is recursive.

Define the average-case performance ratio of algorithm A with respect to the $\mathbf{m}(x)$ distribution on inputs of length $m$ as

$$ac_\pi^{\mathbf{m}}(A, m) = \frac{\sum_{l(x)=m} \mathbf{m}(x)\pi_{opt}(x)/A(x)}{\sum_{l(x)=m} \mathbf{m}(x)}$$

where $l(x)$ denotes the length of $x$.

Let $\mathcal{M}$ denote the set $\{\binom{2}{2}, \binom{3}{2}, \ldots\}$.

THEOREM 1. *For any algorithm A for any problem $\pi$ as defined above, and for any $m \in \mathcal{M}$,*

$$ac_\pi^{\mathbf{m}}(A, m) = \Theta(wc_\pi(A, m))$$

*Proof.* Following [13], define a probability distribution $P(x)$ that assigns high probability to some inputs for which the worst-case complexity is reached, and zero probability in other cases. Specifically,

$$P(x) := \begin{cases} 0 & \text{if } l(x) \notin \mathcal{M} \\ \sum_{l(x)=m} \mathbf{m}(x) & \text{if } l(x) \in \mathcal{M} \text{ and } x \text{ is the lexicographically} \\ & \text{first graph in } \mathcal{G}_n \text{ with } \pi_{opt}(x)/A(x) = wc_\pi(A, m) \\ 0 & \text{if } l(x) \in \mathcal{M} \text{ and } x \text{ is not as described above} \end{cases}$$

The distribution $P$ is defined on $S = Z^+ \cup \{u\}$ where $u$ is a symbol not in $Z^+$ used to cover the remaining probability uncovered by $Z^+$ [13]. We set

$$P(u) := \mathbf{m}(u) + \sum_{x:l(x)\notin\mathcal{M}} \mathbf{m}(x)$$

which makes

$$\sum_{x\in S} P(x) = \sum_{x\in S} \mathbf{m}(x) = 1.$$

The distribution $P(x)$ is enumerable since $\mathbf{m}(x)$ is enumerable and because the above construction preserves this property.

Therefore there is a constant $c_P > 0$ such that for all $x \in S : c_P \times \mathbf{m}(x) \geq P(x)$ [13]. Hence

$$\begin{aligned} ac_\pi^{\mathbf{m}}(A, m) &= \frac{\sum_{l(x)=m} \mathbf{m}(x)\pi_{opt}(x)/A(x)}{\sum_{l(x)=m} \mathbf{m}(x)} \\ &\geq \frac{1}{c_P} \frac{\sum_{l(x)=m} P(x)\pi_{opt}(x)/A(x)}{\sum_{l(x)=m} \mathbf{m}(x)} \\ &= \frac{1}{c_P} \frac{\sum_{l(x)=m} \mathbf{m}(x) wc_\pi(A, m)}{\sum_{l(x)=m} \mathbf{m}(x)} \\ &= \frac{wc_\pi(A, m)}{c_P}. \end{aligned}$$

The proof is completed because, trivially,

$$wc_\pi(A, m) \geq ac_\pi^{\mathbf{m}}(A, m).$$

Theorem 1, while a rigorous result, has two limitations as far as the experimental part of this paper is concerned. First, the proof hinges on the distribution $\mathbf{m}(x)$,

which is not computable, in particular, as Li and Vitanyi explain, on the unique dominance property of $\mathbf{m}(x)$. Second, the constant $c_P$ used in the proof is quite large [13], which in turn implies that the proof holds only for large problem instances.

This raises the question: while the current proof of Theorem 1 does not go through when the hypotheses are changed as noted above, does the effect captured by Theorem 1 still persist? One way to address this question is via experimentation. This is the approach taken in the remainder of this paper. Here, we present Li and Vitanyi's intuitive explanation of the effect that gives rise to Theorem 1 [13], adapting it to performance ratio.

> Let algorithm A have average-case and worst-case performance ratios of $\Theta(f(n))$ and $\Theta(g(n))$ respectively, where $f(n) = o(g(n))$. This implies that only on a sparse subset of inputs of length $n$ does A achieve its worst-case performance ratio of $\Theta(g(n))$. Owing to its sparseness, given $n$, every such input is compressible.

This conclusion, that "all worst-case inputs are compressible" coupled with the fact that the set of compressible inputs is, like the set of worst-case inputs, a sparse subset of the set of all inputs of length $n$ offers hope that if one samples sufficiently-many compressible inputs (which is not a lot), one will find some worst-case ones. It is this hope that drives the experimentation in the remainder of this paper.

## 3. The $\mathbf{q}(x)$ Sampler

We use a functional programming notation $\mathcal{L}$ due to Y. Gurevich and S. Shelah [5] which captures all programs which run in *nearly linear time*, viz. time bounded by $c \cdot n(\log n)^k$ for some fixed $c, k > 0$. We wrote in the C programming language a program $D$ which decodes any binary string $z$ into an $\mathcal{L}$-program $P$ and an argument $y$. Details of $\mathcal{L}$ and $D$ are given below. We also wrote in C a program $U$ which takes $P$, $y$, and the target graph size $n$ as arguments, and simulates $P$ on input $y$. If $x := P(y)$ does not have length $n$ or greater, the output is discarded. Else we define $U(z, n) = U(P, y, n)$ to be the graph $G_x$ constructed by taking the first $n$ bits of $x$.

With reference to our specific decoder $D$ and simulator $U$, we define, for all $n$ and $x \in \{0,1\}^n$:

$$\begin{aligned} w_{\mathbf{q}}(x) &:= \sum 2^{-|z|} : |z| \leq |x| + 1, U(z, |x|) = x, \\ W_{\mathbf{q}}(n) &:= \sum w_{\mathbf{q}}(x) : x \in \{0,1\}^n, \\ \mathbf{q}(x) &:= w_{\mathbf{q}}(x)/W_{\mathbf{q}}(n). \end{aligned} \tag{1}$$

The distribution $\mathbf{q}$ is computable and samplable by generating strings $z$ of length $\leq n + 1$ uniformly at random. Except in relatively rare cases, the sampling and decoding of each $z$ takes nearly linear time as a function of $n$.

In the series of tests reported in this paper we did not do this, but rather restricted our sampling to strings $z$ of length approximately $\sqrt{n}$. Limitations of

time and hardware led us to avoid working with seed strings of length close to $n$, as sampling according to **q** would require. We felt that if the phenomenon raised above were true, we could detect it more readily by limiting the sample to graphs known to have relatively high weight under **q**, and comparing that to samples drawn from the uniform distribution. Our choice of quadratic compression was partly motivated by the hard-to-approximate graphs in [4; 1]. These graphs are described by a fixed oracle for an NP-complete language, and an instance $x$ of some length $m$. The graph $G_{M,x}$ has $m^{O(1)}$ vertices which represent accepting *transcripts* of the protocol (see [4]), two of which are joined by an edge iff the oracle answers in the respective transcripts do not contradict each other. It is an open question whether the "O(1)" can be reduced to nearly-linear; even if so, the bit-size of $G_{M,x}$ is still bounded below by $\frac{1}{2}m^2$ (with judicious padding just in case $M$ rejects $x$).

The decoding strategy we used was to regard $z$ as $l \cdot y \cdot P'$, where $l$ is a self-delimiting description of the length of the argument $y$ and $P'$ includes the bits for $P$ as well its self-delimiting description (number of functions in $P$; number of bits for each of their parameters). The advantage of this decoding strategy is that $y$ and $P$ scale well with length of $z$. The seeds were recorded to make the experiments repeatable. See the Appendix for implementation details.

The following description of the eight basic string functions from [5] assumes that the shown occurrences of substrings meeting the 'if' conditions are leftmost in $x$, and for R2, R3, and $E$, that the "parameter strings" $u, v, \ldots$, all have the same length.

$R0_{u,y}(x)$ : If $x = ur$ then $yr$, else $x$.
$R1_{u,y}(x)$ : If $x = tur$ then $tyr$, else $x$.
$R2_{u,v,y,z}(x)$ : If $x = sutvr$ then $sytzr$, else $x$.
$R3_{u,v,w,y,z}(x)$ : If $x = sut_1vt_2wr$ with $|t_1| = |t_2|$ then $sut_1yt_2zr$, else $x$.
$E_{u,v}(x)$ : Simultaneously replace every 0 in $x$ with $u$ and every 1 with $v$.
$C_{u,v}(x)$ : If $x = E_{u,v}(y)$ for some $y$ then $y$, else $x$.
$A_u(x)$ : Add a tail of $|x|$ log $|x|$-many copies of $u$ to $x$.
$D_u(x)$ : Delete the maximal tail of $u$'s in $x$.

There are two constructors: functional composition, and "iterated replacements" of the form $(\mathbf{R})^*(x)$, where **R** is a composition of any number of $R0 \ldots R3$ functions, and **R** is applied $|x|$-many times. The main theorem of [5] states that every function computed by a random access machine in nearly linear time (NLT) is computed by some program in $\mathcal{L}$, and vice-versa. Thus $\mathcal{L}$ is universal for NLT computation. This justifies regarding **q** as an efficiently computable analog of **m**. However, it should be pointed out that whereas the parameter strings $u, v, \ldots$, are fixed in individual $\mathcal{L}$-programs, the definition of **q** effectively quantifies over

them. This allows for quadratic and greater expansions. Except for cases where an occurrence of $C_{u,v}(\cdot)$ or $D_u(\cdot)$ causes a large contraction of an expanded string, the time remains nearly linear in the length of the output.

The main practical reason for using $\mathcal{L}$ is its simplicity and ease of implementability. Also, while the expansion operations $E$ and $A$ always apply, the contraction operations $C$ and $D$ most often have no effect. Hence $\mathcal{L}$ has a bias toward expansion which is not unnatural, and which reduces the sampling time. Indeed, we were surprised to find that no fewer than one out of every six randomly chosen seeds expanded out to a large enough graph.

## 4. The Neural Network Algorithms

All neural network algorithms evaluated in this paper are based on the Hopfield model [7; 8], and are described in detail in [9]. Here we describe them only briefly, without explaining their neural network implementation in much detail. It is worth noting that all these algorithms arise as manifestations of essentially a single meta-algorithm: one that minimizes the usual *energy function* in the Hopfield model [7; 8].

### 4.1. Discrete Algorithms

**Steepest Descent**. Steepest Descent (SD) is a discrete serial-update neural network heuristic that minimizes energy in greedy fashion. In each time step, the unit to switch decreases energy by the maximum amount. We use the notation $SD(V_0)$ to denote the Steepest Descent starts initially from some subset $V_0 \subseteq V$ of vertices. SD iteratively transforms $V_0$ into a maximal clique $C$, terminating efficiently within $2n$ iterations [9]. Let $V_i$ denote the vertex set in iteration $i$ and assume that it is not a maximal clique. SD emulates the following heuristic in iteration $i$:

If $V_i$ is not a clique then
    remove a minimum degree vertex in induced subgraph $G[V_i]$ from $V_i$
else if $V_i$ is a clique then
    add to $V_i$ a vertex in $V \backslash V_i$ that is adjacent to every vertex in $V_i$.

    Ties are broken lexicographically.

**$\rho$-annealing**. $\rho$-annealing is another discrete serial-update neural network heuristic, which works by carrying out annealing while minimizing energy. More precisely, a certain parameter of the network, called $\rho$, is varied while the network minimizes energy by steepest descent. This is analogous to varying the temperature $T$ in simulated annealing. Like $T$ in simulated annealing, increasing the parameter $\rho$ has the effect of progressively tightening the constraints until, eventually, the solution becomes feasible (a clique in our case). We omit the precise description

of $\rho$-annealing here, for which the reader is referred to [9]. An intuitive description is as follows.

1. Start with small $\rho$ and with the initial state $V_0 := V$.
2. Run SD($V_0$) with this value of $\rho$ to transform $V_0$ into $U$.
3. Increase $\rho$, set $V_0 := U$, and go to step 2.

The algorithm is terminated when $\rho$ becomes sufficiently large. It turns out that when $\rho$ is small, the set $U$ retrieved in step 2 is not required to be a clique; however as $\rho$ is increased, certain constraints get ever tighter, ultimately forcing $U$ to be a clique. In other words, like simulated annealing, this algorithm starts with loose constraints – allowing an unbiased exploration of the search space – and progressively tightens them until the final solution $U$ forms a clique. A precise characterization of the behavior of this algorithm is in [9].

**Stochastic Steepest Descent.** Stochastic Steepest Descent (SSD) is a randomized variant of SD. The deterministic moves of SD are replaced by energy-minimizing moves that favor the steepest direction, but probabilistically. More precisely,

> The unit to switch is picked with probability proportional to the amount of energy its switch would decrease. (The probability is zero if the switch would keep the energy same or increase it.)

The algorithm is motivated by the desire to randomize the choice of unit to switch, which allows one to use repeated runs of the algorithm to boost the size of the clique found, while not totally relinquishing the greedy heuristic emulated by SD, which often works well (see Tables 1 and 2, and [9]).

Let $SSD(V_0, i)$ denote $i$ runs of SSD on a given graph, with $V_0$ as the initial state (vertex set) in each run. (Note that the initial state is the same in each run.) The largest clique found in a run is the output of the algorithm. One run of SSD terminates within $2n$ unit-switches (iterations) [9], which keeps one run as efficient as SD.

### 4.2. CONTINUOUS ALGORITHMS

The description of the continuous algorithms assumes familiarity with the continuous Hopfield model [8].

**Continuous Hopfield Dynamics.** This algorithm, called the continuous Hopfield dynamics (CHD) [8; 6], is described by a system of $n$ coupled nonlinear differential equations, presented here in discretized form:

$$S(t+1) := S(t) + \gamma(-S(t) + \bar{g}_\lambda(W S(t) + I)). \tag{2}$$

Here $S_i \in [0, 1]$ is the state of the $i^{th}$ neuron, $I_i$ the external bias of the $i^{th}$ neuron, $W$ the $n \times n$ symmetric weight matrix, $g_\lambda(x) = \frac{1}{1+e^{-\lambda x}}$ a sigmoid with gain $\lambda$, $\bar{g}(\bar{x})$ notational shorthand for $(g(x_i))$, and $\gamma$ the Euler step size. The

Table 1. 100-vertex graphs sampled from $\mathbf{q}(x)$. The column MAX gives the size of the maximum clique when known, computed using dfmax, an exact semi-exhaustive search algorithm.

| Graph | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset$, 1) | SSD($V$, 1) | SSD($\emptyset$, $n$) | SSD($V$, $n$) | CHD | MFA | MAX |
|---|---|---|---|---|---|---|---|---|---|---|
| G1 | 5 | 8 | 8 | 7 | 6 | 9 | 11 | 10 | 11 | 11 |
| G2 | 10 | 29 | 29 | 29 | 29 | 29 | 29 | 29 | 28 | 29 |
| G3 | 8 | 21 | 21 | 11 | 14 | 21 | 21 | 16 | 21 | 21 |
| G4 | 4 | 9 | 9 | 4 | 7 | 8 | 9 | 4 | 8 | 9 |
| G5 | 9 | 8 | 8 | 5 | 10 | 10 | 10 | 7 | 10 | 10 |
| G6 | 17 | 40 | 40 | 36 | 30 | 40 | 40 | 40 | 40 | 40 |
| G7 | 7 | 12 | 13 | 7 | 6 | 11 | 13 | 9 | 13 | 13 |
| G8 | 5 | 6 | 6 | 5 | 5 | 6 | 6 | 6 | 5 | 6 |
| G9 | 4 | 3 | 3 | 4 | 2 | 4 | 4 | 4 | 3 | 4 |
| G10 | 5 | 10 | 10 | 8 | 5 | 10 | 10 | 4 | 8 | 10 |
| G11 | 7 | 22 | 22 | 11 | 16 | 21 | 22 | 16 | 22 | 22 |
| G12 | 3 | 27 | 27 | 6 | 8 | 27 | 26 | 27 | 27 | 27 |
| G13 | 8 | 13 | 13 | 9 | 9 | 14 | 14 | 7 | 13 | 15 |
| G14 | 5 | 7 | 7 | 4 | 4 | 7 | 8 | 4 | 7 | 8 |
| G15 | 19 | 26 | 26 | 21 | 17 | 27 | 26 | 22 | 27 | 27 |
| G16 | 4 | 3 | 3 | 4 | 2 | 4 | 4 | 4 | 3 | 4 |
| G17 | 6 | 9 | 9 | 7 | 8 | 10 | 11 | 8 | 10 | 11 |
| G18 | 15 | 31 | 31 | 4 | 23 | 31 | 31 | 31 | 31 | 31 |
| G19 | 6 | 9 | 8 | 7 | 5 | 8 | 9 | 7 | 8 | 9 |
| G20 | 15 | 31 | 31 | 31 | 31 | 31 | 31 | 31 | 30 | 31 |
| G21 | 4 | 5 | 5 | 4 | 3 | 5 | 5 | 4 | 4 | 5 |
| G22 | 9 | 26 | 26 | 13 | 16 | 26 | 26 | 21 | 26 | 26 |
| G23 | 5 | 8 | 9 | 4 | 8 | 11 | 10 | 9 | 10 | 11 |
| G24 | 5 | 9 | 9 | 8 | 7 | 10 | 10 | 8 | 9 | 10 |

Table 1. *Continued*

| Graph | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset$, 1) | SSD($V$, 1) | SSD($\emptyset$, $n$) | SSD($V$, $n$) | CHD | MFA | MAX |
|---|---|---|---|---|---|---|---|---|---|---|
| G25 | 5 | 4 | 4 | 4 | 4 | 4 | 5 | 4 | 5 | 5 |
| G26 | 19 | 42 | 42 | 32 | 40 | 42 | 42 | 36 | 41 | 42 |
| G27 | 100 | 100 | 100 | 100 | 100 | 100 | 100 | 100 | 100 | 100 |
| G28 | 5 | 7 | 9 | 6 | 6 | 8 | 8 | 7 | 7 | 9 |
| G29 | 16 | 21 | 21 | 11 | 16 | 23 | 23 | 13 | 21 | 23 |
| G30 | 5 | 9 | 9 | 11 | 8 | 12 | 12 | 9 | 10 | 13 |
| G31 | 3 | 3 | 3 | 2 | 3 | 4 | 3 | 3 | 2 | 4 |
| G32 | 15 | 39 | 39 | 18 | 21 | 39 | 39 | 38 | 38 | 39 |
| G33 | 60 | 60 | 60 | 38 | 57 | 65 | 60 | 60 | 60 | 60 |
| G34 | 8 | 8 | 8 | 6 | 4 | 13 | 13 | 5 | 7 | 13 |
| G35 | 5 | 35 | 35 | 25 | 35 | 35 | 35 | 35 | 35 | 35 |
| G36 | 8 | 10 | 11 | 8 | 10 | 11 | 11 | 7 | 11 | 12 |
| G37 | 3 | 2 | 2 | 2 | 2 | 3 | 3 | 2 | 2 | 3 |
| G38 | 8 | 13 | 13 | 7 | 9 | 13 | 13 | 7 | 13 | 13 |
| G39 | 4 | 17 | 17 | 10 | 6 | 16 | 17 | 17 | 17 | 17 |
| G40 | 8 | 29 | 29 | 27 | 11 | 27 | 29 | 29 | 29 | 29 |
| G41 | 5 | 11 | 11 | 8 | 4 | 13 | 12 | 8 | 11 | 13 |
| G42 | 4 | 5 | 5 | 5 | 3 | 5 | 5 | 5 | 4 | 5 |
| G43 | 7 | 21 | 21 | 8 | 11 | 20 | 20 | 18 | 21 | 21 |
| G44 | 2 | 26 | 26 | 26 | 24 | 26 | 26 | 26 | 25 | 26 |
| G45 | 12 | 37 | 37 | 21 | 32 | 37 | 37 | 37 | 36 | 37 |
| G46 | 3 | 6 | 6 | 2 | 3 | 5 | 6 | 5 | 6 | 6 |
| G47 | 6 | 12 | 13 | 7 | 10 | 13 | 14 | 9 | 14 | 14 |
| G48 | 4 | 6 | 6 | 4 | 6 | 7 | 7 | 5 | 6 | 7 |
| G49 | 4 | 11 | 11 | 4 | 9 | 10 | 11 | 8 | 11 | 11 |
| G50 | 6 | 10 | 10 | 7 | 8 | 12 | 13 | 10 | 13 | 13 |
| Avgs. | 10.20 | 18.20 | 18.42 | 12.96 | 14.26 | 18.76 | 19.00 | 16.62 | 18.38 | 19.20 |

Table 2. 400-vertex graphs sampled from $\mathbf{q}(x)$. The column labeled MAX contains the maximum clique size, computed as described in Table 1. A $\geq$ in the MAX column indicates the maximum clique was not found for that graph – either because dfmax did not (or it was anticipated it would not) terminate in reasonable time.

| Graph | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset$, 1) | SSD($V$, 1) | SSD($\emptyset$, $n$) | SSD($V$, $n$) | CHD | MFA | MAX |
|---|---|---|---|---|---|---|---|---|---|---|
| G1 | 6 | 16 | 16 | 9 | 6 | 12 | 15 | 16 | 15 | 17 |
| G2 | 12 | 90 | 90 | 16 | 12 | 90 | 90 | 90 | 90 | $\geq$90 |
| G3 | 12 | 36 | 39 | 10 | 27 | 34 | 40 | 39 | 41 | 41 |
| G4 | 16 | 166 | 166 | 130 | 144 | 163 | 166 | 166 | 166 | $\geq$166 |
| G5 | 20 | 130 | 130 | 100 | 40 | 130 | 130 | 130 | 130 | $\geq$130 |
| G6 | 10 | 31 | 33 | 20 | 22 | 27 | 34 | 26 | 29 | 38 |
| G7 | 8 | 26 | 26 | 19 | 17 | 27 | 26 | 23 | 27 | 30 |
| G8 | 4 | 43 | 43 | 11 | 21 | 43 | 43 | 41 | 43 | $\geq$43 |
| G9 | 4 | 23 | 23 | 9 | 5 | 22 | 23 | 22 | 23 | 23 |
| G10 | 8 | 25 | 30 | 9 | 13 | 29 | 30 | 20 | 23 | 30 |
| G11 | 11 | 74 | 74 | 28 | 35 | 72 | 74 | 74 | 74 | $\geq$74 |
| G12 | 3 | 68 | 68 | 47 | 23 | 68 | 68 | 67 | 65 | $\geq$68 |
| G13 | 8 | 35 | 35 | 17 | 6 | 32 | 34 | 34 | 34 | 35 |
| G14 | 10 | 25 | 25 | 14 | 11 | 25 | 25 | 15 | 24 | 26 |
| G15 | 11 | 93 | 93 | 85 | 65 | 93 | 93 | 63 | 62 | $\geq$93 |
| G16 | 21 | 51 | 51 | 35 | 55 | 65 | 65 | 65 | 64 | $\geq$65 |
| G17 | 47 | 115 | 115 | 64 | 115 | 114 | 115 | 83 | 114 | $\geq$115 |
| G18 | 41 | 141 | 141 | 92 | 100 | 139 | 138 | 122 | 141 | $\geq$141 |
| G19 | 4 | 11 | 11 | 5 | 10 | 13 | 14 | 10 | 14 | 15 |
| G20 | 7 | 27 | 27 | 11 | 11 | 27 | 22 | 16 | 28 | 29 |
| G21 | 12 | 80 | 82 | 24 | 39 | 82 | 82 | 61 | 80 | $\geq$82 |
| G22 | 15 | 40 | 58 | 39 | 52 | 76 | 76 | 76 | 57 | $\geq$76 |
| G23 | 6 | 59 | 59 | 15 | 29 | 53 | 59 | 56 | 59 | $\geq$59 |
| G24 | 2 | 132 | 132 | 75 | 131 | 132 | 132 | 132 | 132 | $\geq$132 |

Table 2. *Continued*

| Graph | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset$, 1) | SSD($V$, 1) | SSD($\emptyset$, $n$) | SSD($V$, $n$) | CHD | MFA | MAX |
|---|---|---|---|---|---|---|---|---|---|---|
| G25 | 4 | 35 | 35 | 7 | 7 | 34 | 35 | 34 | 33 | 35 |
| G26 | 5 | 14 | 20 | 20 | 12 | 25 | 24 | 21 | 27 | 27 |
| G27 | 6 | 18 | 18 | 9 | 8 | 14 | 15 | 8 | 16 | 18 |
| G28 | 135 | 135 | 135 | 6 | 97 | 135 | 135 | 135 | 135 | $\geq$135 |
| G29 | 60 | 71 | 71 | 28 | 67 | 72 | 72 | 72 | 72 | $\geq$72 |
| G30 | 5 | 14 | 14 | 9 | 12 | 14 | 14 | 11 | 12 | 14 |
| G31 | 192 | 226 | 226 | 222 | 223 | 225 | 226 | 205 | 225 | $\geq$226 |
| G32 | 3 | 125 | 125 | 111 | 105 | 125 | 125 | 124 | 123 | $\geq$125 |
| G33 | 11 | 13 | 13 | 12 | 10 | 13 | 13 | 13 | 13 | 13 |
| G34 | 9 | 16 | 17 | 10 | 9 | 16 | 16 | 16 | 17 | 19 |
| G35 | 8 | 25 | 24 | 11 | 27 | 25 | 27 | 23 | 25 | 29 |
| G36 | 6 | 27 | 27 | 22 | 10 | 29 | 27 | 12 | 26 | 29 |
| G37 | 25 | 56 | 56 | 56 | 28 | 56 | 56 | 20 | 57 | $\geq$57 |
| G38 | 4 | 22 | 22 | 5 | 10 | 21 | 21 | 20 | 22 | 22 |
| G39 | 166 | 193 | 188 | 197 | 177 | 200 | 203 | 174 | 194 | $\geq$203 |
| G40 | 9 | 28 | 28 | 51 | 10 | 51 | 51 | 28 | 27 | 51 |
| G41 | 11 | 92 | 93 | 62 | 78 | 93 | 93 | 80 | 93 | $\geq$93 |
| G42 | 17 | 56 | 42 | 30 | 20 | 55 | 56 | 40 | 43 | 56 |
| G43 | 13 | 17 | 29 | 17 | 19 | 33 | 30 | 17 | 29 | 34 |
| G44 | 19 | 57 | 57 | 43 | 41 | 59 | 58 | 46 | 58 | $\geq$59 |
| G45 | 6 | 52 | 62 | 17 | 30 | 62 | 62 | 58 | 62 | $\geq$62 |
| G46 | 8 | 50 | 70 | 40 | 32 | 69 | 69 | 43 | 69 | $\geq$70 |
| G47 | 4 | 7 | 7 | 7 | 3 | 7 | 7 | 7 | 6 | 7 |
| G48 | 5 | 58 | 58 | 6 | 12 | 58 | 58 | 58 | 58 | $\geq$58 |
| G49 | 2 | 59 | 59 | 42 | 29 | 59 | 59 | 59 | 58 | $\geq$59 |
| G50 | 14 | 58 | 76 | 56 | 54 | 74 | 75 | 58 | 77 | $\geq$77 |
| Avgs. | 20.90 | 61.22 | 62.78 | 39.6 | 42.38 | 63.84 | 64.42 | 56.58 | 62.22 | $\geq$65.36 |

continuous-time version of (2) minimizes an energy function during its evolution [8], into which the MAX-CLIQUE problem can be encoded [9]. With sufficiently large $\lambda$ and sufficiently small $\gamma$, if (2) is started from any initial state $S(0) \in [0,1]^n$ and iterated sufficiently-many times, it provably terminates at a fixed point $S$ from which a maximal clique of the encoded graph can be recovered [9].

For a discussion of the significance of CHD from the point of view of neural implementation and optimization applications, see [8; 6; 9]. CHD is especially interesting because it may be viewed as the essential special case of the algorithm presented next – a continuous optimization method developed only recently, but one that is already beginning to make its mark on optimization as it occurs in practice.

**Mean Field Annealing.** The second continuous heuristic, called *Mean Field Annealing* (MFA) [2; 17], may be described as a generalization of CHD in which the sigmoidal gain $\lambda$ is varied during the evolution of (2). This is done by employing an annealing schedule, a sequence $\{\lambda_i, \mu_i\}$ of $k$ elements, where $\lambda_i$ is the value of the sigmoidal gain and $\mu_i$ the number of times (2) is to be iterated with the sigmoidal gain set at $\lambda_i$. Usually $\lambda_i$ is a monotonically increasing function of $i$. The detailed algorithm is as follows.

$S := S(0)$
for $i := 1$ to $k$ do
  for $j := 1$ to $\mu_i$ do
    $S := S + \gamma(-S + \bar{g}_{\lambda_i}(WS + I)$

With sufficiently small $\gamma$ and sufficiently large $\mu_i$, $S$ converges to a fixed point at each value of $i$ [8; 6]. Additionally, with sufficiently large $k$, and with $\lambda_i$ growing sufficiently slowly with $i$, MFA is known to deterministically approximate simulated annealing during its evolution [2; 6], while being more efficient.

## 5. Experimental Methodology

All experiments on the neural network algorithms and their evaluation on $\mathbf{u}(n)$ and $\mathbf{q}(x)$ were performed on a SUN SparcStation I.

**Details of the Sampling Process.** Experiments were performed on 100-vertex graphs and on 400-vertex graphs. The bitstrings of the 100-vertex graphs had length 4950, and those of the 400-vertex graphs had length 79,800. For $n = 100$ and $n = 400$, three sets of fifty $n$-vertex graphs were generated. One set was drawn from the uniform distribution with $p = 0.5$, and one with $p = 0.9$. All seed strings were generated using the standard UNIX pseudorandom number generator, and recorded to make the experiments repeatable.

The third set was generated using seed strings of lengths 65..85 for the fifty 100-vertex graphs and eleven of the 400-vertex graphs, and 270..285 for thirty-

nine of the 400-vertex graphs. When we compiled the 100-vertex set we found that eleven seeds expanded to strings of length greater than 79,800. Rather than truncate them to length 4950, we decided to discard them from the 100-vertex sample but include them into the 400-vertex sample. For hardware reasons we also set a limit of 700,000 on the number of bits produced at any stage of the decoding, and discarded those seeds which broke it from the 400-vertex sample. We believe that these practical decisions did not bias our results in any significant way. It took about 12 hours of computing time to assemble this set.

The long strings of 0s and 1s were truncated to length 4950 or 79,800 and formed into an adjacency matrix for the graph in the order (1,2), (1,3), (2,3), (1,4),... of edges. Over three-fourths of these strings were generated by $\mathcal{L}$-programs whose final instruction was "Add a tail of $|x|$-many copies of $u$ to $x$", where $u$ was fairly long, and so ended with many repetitions of $u$. We do not have an intuitive idea of the extent to which this yielded repeated patterns in the graph.

Nine heuristics were tested on each of the six sets, giving 2700 runs in all. For each 400-vertex graph, it took about two hours to run all nine. The MFA heuristic was by far the slowest of the lot.

**Structural Comparisons of the Graphs.** It is instructive to compare the $\mathbf{q}(n)$ graphs versus the $\mathbf{u}(n)$ graphs on certain structural properties. The three properties we adopted for comparison are: $m(G)$, the number of edges in a graph $G$, $(\Delta - \delta)(G)$ the difference between the maximum degree and the minimum degree in a graph $G$, and $\omega(G)$, the order of a maximum clique in a graph $G$. The numbers are presented below.

| Property | $\mathbf{q}(100)$ | $\mathbf{q}(400)$ | $\mathbf{u}_{0.5}(100)$ | $\mathbf{u}_{0.5}(400)$ |
|---|---|---|---|---|
| | min, avg, max | min, avg, max | min, avg, max | min, avg, max |
| $m$ | 343,2313.1,4950 | 3958,38189.9,68491 | 2402,2477.4,2553 | 39528,39867.4,40161 |
| $\Delta - \delta$ | 0,39.8,96 | 24,186.1,389 | 20,24.7,30 | 50,59.34,73 |
| $\omega$ | 3,19,100 | 7,–,$\geq$226 | $9^*$,$9.99^*$,$10^*$ | $12^*$,$12.99^*$,$13^*$ |

Each group of three columns is associated with a particular class of graphs. In any particular row, the three values associated with a particular class of graphs are the minimum, the maximum, and the average values of the property associated with that row, taken over the fifty graphs of that class. The $^*$-ed values are very sharp estimates taken from random graph theory.

From the above numbers we see that the $\mathbf{q}(n)$ graphs spanned a wide range of densities – from the very sparse to the very dense – while their average density was very close to that of a $\mathbf{u}_{0.5}(n)$ graph. The difference between the maximum degree and the minimum degree of a $\mathbf{q}(n)$ graph was also spread over a wide range, with a large average value. (Notice that this says two things: (i) in most graphs, there was a large gap between maximum and minimum vertex degree and (ii) this gap

itself had a large variance.) Finally, the maximum clique sizes of $\mathbf{q}(n)$ graphs were also in a wide range.

In striking contrast, for graphs drawn from $\mathbf{u}_{0.5}(n)$ all three numbers – $m$: the number of edges, $\Delta - \delta$: the maximum degree minus the minimum degree, and $\omega$: the order of the maximum clique – were confined in a very narrow range. In addition, the average value of $\Delta - \delta$ was itself very small, indicating that most degrees are bunched together. This gives some indication of the diversity of the $\mathbf{q}(n)$ graphs in comparison with those drawn from $\mathbf{u}_{0.5}(n)$.

**Sample Sizes.** Each set of test graphs contained fifty graphs. It is reasonable to ask if this sample size is adequate. For graphs drawn from $\mathbf{u}_p(n)$, several arguments lead to the conclusion that a sample size of fifty graphs is more than adequate for our purposes. First, the expected size of the maximum clique in a graph drawn from $\mathbf{u}_p(n)$ has a sharp threshold [15] and the range of sizes of maximal cliques in such graphs is also quite narrow. Thus, any maximal-clique finding algorithm, for example most of the ones in the current paper, is guaranteed to find a clique in a narrow range. This argument is reinforced by experimental results reported in [9], which give the distribution of clique sizes found in fifty graphs drawn from $\mathbf{u}_p(400)$, $p = 0.5, 0.9$, which turns out to have a very small variance.

For graphs drawn from $\mathbf{q}(x)$, however, it was not clear a priori when an adequate sample size should be. We decided to start with a sample size of fifty. On this sample size, the results reported in Tables 1 and 2 (see Section 6 for their presentation and analysis) displayed certain trends so clearly and consistently that we felt confident that our observations were sound and the trends would remain similar on larger sample sizes.

**Evaluating Performance.** The main hurdle in analyzing the results is that there is no easy way of calculating the size of the largest clique in a graph. We could have used some exponential-time algorithm to find the exact answer, but this would have been quite time-consuming on the 2700 runs. Therefore, instead of comparing the *absolute* performance of these algorithms on $\mathbf{u}(n)$ versus those on $\mathbf{q}(x)$, we decided to compare their *relative* performances, in particular how well or poorly certain algorithms performed relative to others, on $\mathbf{u}(n)$ versus $\mathbf{q}(x)$.

**Parameter Settings of the Continuous Algorithms.** The continuous algorithms – CHD and MFA – use certain free parameters whose values needed to be set. The values that we used are described below to make the experiments independently repeatable. One needs to refer to [9] in order to understand some of the parameters.

CHD was operated at $\rho = -10n$, $\lambda = 1$, $\gamma = 0.1$, $I_i = |\rho|/4$ for all $i$, and with the number of iterations of (2) fixed in advance to $n$. The initial state to CHD was set to $S(0) := (0.5 + \delta)^n$, where $\delta$ was a random value in $[-0.05, 0.05]$. The settings are the same as in [9], and are motivated there.

MFA was operated with the same settings for $\rho, \lambda, \gamma, I$, and the initial state $S(0)$ as was CHD, and with the following geometric annealing schedule:

$$T_i = a_{i-1}T_{i-1}; T_1 = \frac{2}{6}n|\rho|$$

where $a_i = 0.9$ for $i \leq 4$ and $a_i = 0.5$ for $i > 4$. Here $T_i = 1/\lambda_i$. The settings are essentially the same as in [9], and are motivated there. Experimental results in [9] also reveal that CHD and MFA continue to work well on the graphs tested in [9], on parameter settings in a reasonably large neighborhood of those described above.

## 6. Experimental Results

Tables 1 and 2 give the size of the clique retrieved by each of the nine algorithms, on fifty 100-vertex and fifty 400-vertex graphs sampled from $\mathbf{q}(x)$ respectively.

SD($\emptyset$) is the steepest descent algorithm whose initial state is the empty set. It emulates the naive heuristic:

> Start from the empty set and extend it, by adding, in each step, one suitable vertex selected lexicographically, until it forms a maximal clique.

SSD($\emptyset, 1$) is a randomized version of SD($\emptyset$) in which the vertex to be added is selected randomly, from the feasible choices, instead of in lexicographic fashion. SD($V$) is the steepest descent algorithm whose initial state is the entire vertex set $V$. It turns out that SD($V$) emulates the following algorithm [9]:

$S := V$
while $S$ is not a clique do
    Pick a vertex $v \in S$ with minimum degree in $S$
    Delete v from S
endwhile
while S is not a *maximal* clique do
    Pick smallest-numbered vertex $v \notin S$ adjacent to every vertex in S
    Add v to S
endwhile

SSD($V, 1$) is a randomized version of SD($V$) in which the vertex to be deleted in an iteration of the first loop is picked with probability proportional to $S - \text{degree}_S(v)$ (the smaller the degree, the higher the probability), and the vertex to be added in an iteration of the second loop is picked at random from the feasible choices. SSD($\emptyset, n$) and SSD($V, n$) are multiple restart versions of SSD($\emptyset, 1$) and SSD($V, 1$) respectively – the largest clique found in the $n$ runs is output.

From Tables 1 and 2, the following observations can be made:

– SD($\emptyset$) works the poorest. SSD($\emptyset, 1$) and SSD($V, 1$) work moderately better but remain significantly poorer than the best algorithms. Thus, randomization alone helps but not as much as one might expect.
– SD(V) works much better than SD($\emptyset$) and nearly as well as the best algorithms. Thus replacing the initial state of SD($\emptyset$) by $V$, which makes the SD algorithm greedier, improves the performance much more dramatically than by randomizing SD($\emptyset$) alone. Randomizing SD(V), to get SSD(V,1), in fact worsens the performance significantly.
– The $\rho$-annealing algorithm consistently works just slightly better than SD(V).
– The multiple restart algorithms – SSD($\emptyset, n$) and SSD($V, n$) – work the best, with SSD($V, n$) working just very slightly better. This shows that the real benefit of randomization is that it allows multiple restarts, which boosts performance immensely.
– The continuous algorithm CHD works moderately poorer than the continuous algorithm MFA, which was anticipated, but also works discernibly poorer than the discrete algorithm SD($V$), which was not anticipated.

These observations hold for both Tables 1 and 2 – if anything, the effects are more pronounced in Table 2 than in Table 1.

From these results we may cluster the algorithms into four groups, using the size of the clique found as the measure. In order of decreasing performance, the clusters are:

1. SSD($V, n$), SSD($\emptyset, n$), $\rho$-annealing, MFA, and SD($V$).
2. CHD.
3. SSD($V, 1$) and SSD($\emptyset, 1$).
4. SD($\emptyset$).

### 6.1. COMPUTING THE MAXIMUM CLIQUE SIZE

To compute exact performance ratios of our neural network algorithms on the $\mathbf{q}(x)$ graphs, we ran an exact semi-exhaustive search program, called `dfmax`, to find a maximum clique on all fifty graphs from $\mathbf{q}(100)$, and twenty three graphs from $\mathbf{q}(400)$. The program `dfmax` is a C version of the code described in [3], and was implemented by D.S. Johnson and used earlier as a benchmark in the Second DIMACS Challenge.

In addition to allowing computation of exact performance ratios, these experiments also provided an unanticipated insight. In particular, they revealed that `dfmax` slowed down tremendously on several of the $\mathbf{q}(400)$ graphs. We ran `dfmax` on only 23 graphs – the ones with the smallest cliques – on which it found a maximum clique quite quickly. We tried `dfmax` on several of the remaining 27 graphs with larger cliques. In all cases we tried, `dfmax` ran for several days – in some cases for two weeks – without finding a maximum clique. We were forced to terminate these runs, and did not even attempt to evaluate `dfmax` on the graphs we hadn't tried on (the ones with even larger cliques).

For comparison, we ran dfmax on random graphs, on which it was able to find a maximum clique in each of fifty graphs drawn from $\mathbf{u}_{0.5}(400)$ in time ranging from 43.9 to 53 seconds. These findings reinforce the indication from the performance ratio results (see Tables 1 and 2) that graphs drawn from $\mathbf{q}(n)$ are harder than those drawn from $\mathbf{u}(n)$, and that, as $n$ increases, the $\mathbf{q}(n)$ graphs get harder more quickly than do the $\mathbf{u}(n)$ ones.

It is to be noted that when we ran dfmax on a graph from $\mathbf{u}_{0.9}(400)$ it ran for fifteen hours without giving an answer, at which point the run was terminated. Although this does seem to indicate that dense random graphs are harder for dfmax than uniform random graphs, it does not give any indication of whether the $\mathbf{u}(0.9)$ graphs are as hard as the $\mathbf{q}(400)$ ones for dfmax.

It is reasonable to suspect that, because dfmax is a semi-exhaustive search algorithm, the reason it ran efficiently on the $\mathbf{u}_{0.5}(400)$ graphs is that they all had very small maximum cliques, while the reason that it ran so poorly on the $\mathbf{q}(400)$ graphs with large cliques (and also on the one dense random graph) is precisely because they have large cliques. That graphs with large maximum cliques arise frequently in $\mathbf{q}(x)$ is to be expected (they almost never arise in $\mathbf{u}_{0.5}(x)$) and if this causes problems for dfmax, then so be it. It is worth noting that the relative performance ratio of SD($\emptyset$) to the best heuristic was 10.93 and 7.57, averaged over these twenty seven and the entire set of fifty $\mathbf{q}(400)$ graphs respectively. Thus, these twenty seven graphs not only seem hard for dfmax but are also harder on average for SD($\emptyset$) than the remaining ones in $\mathbf{q}(400)$.

Finally, these dfmax experiments revealed the practical difficulty of conducting experiments on the $\mathbf{q}(n)$ graphs for all but small $n$.

### 6.2. THE $\mathbf{q}(n)$ VERSUS $\mathbf{u}_{0.5}(n)$ COMPARISONS

Table 3 gives the clique sizes found by these nine algorithms on random graphs, i.e. graphs drawn from $\mathbf{u}_p(n)$. The relative rankings of the algorithms in Tables 1, 2, and 3 are mostly the same, though there is one notable exception, explained later in this paragraph. The performance differentials between these algorithms are however far wider on graphs drawn from $\mathbf{q}(x)$ than on graphs drawn from $\mathbf{u}_p(n)$. On graphs drawn from $\mathbf{u}_{1/2}(n)$, the clique sizes obtained by all nine algorithms are in a small range. SD($\emptyset$) remains the poorest working algorithm, but only marginally so. On graphs drawn from $\mathbf{u}_{0.9}(n)$, the range of clique sizes gets larger and so do the performance differentials, although most of the relative rankings remain unchanged. The one notable exception is the MFA algorithm which performs significantly better than all the others. Although MFA worked very slightly poorer than the multiple restart algorithms on graphs drawn from $\mathbf{q}(x)$ it works significantly better on graphs drawn from $\mathbf{u}_{0.9}(n)$. Another interesting feature of the graphs drawn from $\mathbf{u}_{1/2}(n)$ versus those drawn from $\mathbf{q}(x)$ revealed by the algorithms is that, whereas both types of graphs have roughly the same density, the algorithms, on average, retrieve much larger cliques on graphs drawn from $\mathbf{q}(x)$ than on graphs

Table 3. Average performance on $p$-random graphs. This table is excerpted from Table I in [?].

| $n$ | $p$ | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset, 1$) | SSD($V, 1$) | SSD($\emptyset, n$) | SSD($V, n$) | CHD | MFA |
|---|---|---|---|---|---|---|---|---|---|---|
| 100 | 0.5 | 6.34 | 7.98 | 8.06 | 6.48 | 6.42 | 8.36 | 8.60 | 7.44 | 8.50 |
| 100 | 0.9 | 23.86 | 28.16 | 28.34 | 23.40 | 24.82 | 27.60 | 28.76 | 27.92 | 30.02 |
| 400 | 0.5 | 8.30 | 9.88 | 10.34 | 8.44 | 8.24 | 10.80 | 11.04 | 9.16 | 10.36 |
| 400 | 0.9 | 36.12 | 43.80 | 44.58 | 35.84 | 36.82 | 41.86 | 43.20 | 43.24 | 49.94 |

Table 4. Average performance ratio SSD($V, n$)($G$)/$A(G)$ on all graphs. The numbers for the **q** graphs are obtained by averaging over the performance ratios, not by taking the performance ratio of the average.

| Source | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset, 1$) | SSD($V, 1$) | SSD($\emptyset, 1$) | SSD($V, n$) | CHD | MFA |
|---|---|---|---|---|---|---|---|---|---|
| **q**(100) | 2.42 | 1.10 | 1.09 | 1.79 | 1.60 | 1.03 | — | 1.31 | 1.10 |
| $\mathcal{G}_{.5}(100)$ | 1.35 | 1.08 | 1.07 | 1.37 | 1.32 | 1.02 | — | 1.15 | 1.01 |
| $\mathcal{G}_{.9}(100)$ | 1.20 | 1.02 | 1.01 | 1.23 | 1.15 | 1.04 | — | 1.03 | 0.96 |
| **q**(400) | 7.56 | 1.09 | 1.03 | 2.64 | 2.19 | 1.02 | — | 1.24 | 1.05 |
| $\mathcal{G}_{.5}(400)$ | 1.33 | 1.12 | 1.07 | 1.31 | 1.33 | 1.02 | — | 1.21 | 1.02 |
| $\mathcal{G}_{.9}(400)$ | 1.19 | 0.99 | 0.97 | 1.21 | 1.17 | 1.03 | — | 1.00 | 0.87 |

drawn from $\mathbf{u}_{1/2}(n)$. Certainly this is not surprising as it is reasonable to expect a correlation between the compressibility of graphs drawn from $\mathbf{q}(x)$ and the fact that they contain large cliques.

The *relative performance ratio* of algorithm A to algorithm B on graph G is defined as the clique size found by B divided by the clique size found by A. Table 4 gives the relative performance ratio of each algorithm to SSD($V, n$), averaged over graphs drawn from $\mathbf{q}(x)$ and over graphs drawn from $\mathbf{u}_p(n)$. SSD($V, n$) is chosen as the reference algorithm because it works best on graphs drawn from $\mathbf{q}(x)$. The results reported in Table 4 are drawn from the earlier tables. The results viewed in this fashion tend to support our earlier observations in more dramatic fashion. For example, SD($\emptyset$) has a poor relative performance ratio, 2.42, on graphs drawn from $\mathbf{q}(100)$, which worsens markedly, to 7.56, on graphs drawn from $\mathbf{q}(400)$. By contrast, SD($\emptyset$) has much better relative performance ratios on graphs drawn from $\mathbf{u}_p(100)$, $p = 0.5, 0.9$, which remain unchanged on graphs drawn from $\mathbf{u}_p(400)$, $p = 0.5, 0.9$.

The *absolute performance ratio* of algorithm A on graph G is defined as the maximum clique size in G divided by the clique size found by A. Table 5 gives the exact, estimated, or lower-bounded absolute performance ratio. For the three poorest algorithms, the Table 5 results reveal essentially no new information (the poorest algorithms are revealed to be just slightly poorer than known from Table 4). However, for five of the remaining six algorithms they reveal good news – all except CHD perform near-optimally on the $\mathbf{q}(n)$ graphs, with CHD not far behind.

## 7. Discussion

In this wrap-up section, we have chosen to adopt a rather unusual format – question and answer – which we feel justified given the nature of this paper. These questions originate from the anonymous reviewers.

> Q1: *This combination of theory and experiment might be misleading, favoring an overestimation of the applicability of theoretical results on the universal distribution to the practical problem of test instance generation.*
>
> This is a good point and our response to that takes the following form:
>
> The experimental results neither confirm nor refute the validity of the precise theoretical result in practical situations. They do give evidence of the following, somewhat weaker, phenomenon associated with the theoretical result.
>
> The theoretical result states that graphs drawn from $\mathbf{m}(n)$ draw out worst-case performance ratio from any given algorithm. The experimental results show that graphs drawn from $\mathbf{q}(n)$ – a heuristic analogue of $\mathbf{m}(n)$ – draw out poorer performance ratio than on graphs drawn from $\mathbf{u}(n)$, from three of the nine evaluated heuristic algorithms. Moreover, these three are the ones that also happen to work poorest on graphs drawn from $\mathbf{u}(n)$ (although significantly better on $\mathbf{u}(n)$ than on $\mathbf{q}(\mathbf{n})$).

Table 5. Average absolute performance ratios $\omega(G)/A(G)$ on all graphs. The numbers reported in the rows for $\mathbf{q}(100)$ and $\mathbf{q}(400)$ are exact ones and lower bounds respectively, computed using the MAX columns of Tables 1 and 2. The numbers reported in the rows for $\mathcal{G}_{.5}(100)$, $\mathcal{G}_{.9}(100)$, $\mathcal{G}_{.5}(400)$, and $\mathcal{G}_{.9}(400)$ are very sharp estimates, found using the very sharp estimate on the maximum clique size in a random graph.

| Source | SD($\emptyset$) | SD($V$) | $\rho$-A | SSD($\emptyset, 1$) | SSD($V, 1$) | SSD($\emptyset, n$) | SSD($V, n$) | CHD | MFA |
|---|---|---|---|---|---|---|---|---|---|
| $\mathbf{q}(100)$ | 2.46 | 1.12 | 1.11 | 1.82 | 1.63 | 1.04 | 1.02 | 1.34 | 1.12 |
| $\mathcal{G}_{.5}(100)$ | 1.58 | 1.25 | 1.24 | 1.54 | 1.56 | 1.20 | 1.16 | 1.34 | 1.18 |
| $\mathcal{G}_{.9}(100)$ | 1.47 | 1.24 | 1.23 | 1.50 | 1.41 | 1.27 | 1.22 | 1.25 | 1.17 |
| $\mathbf{q}(400)$ | $\geq$7.68 | $\geq$1.13 | $\geq$1.07 | $\geq$2.71 | $\geq$2.26 | $\geq$1.05 | $\geq$1.04 | $\geq$1.29 | $\geq$1.09 |
| $\mathcal{G}_{.5}(400)$ | 1.57 | 1.32 | 1.26 | 1.54 | 1.58 | 1.20 | 1.18 | 1.42 | 1.25 |
| $\mathcal{G}_{.9}(400)$ | 1.52 | 1.26 | 1.23 | 1.53 | 1.49 | 1.31 | 1.27 | 1.27 | 1.10 |

In a sense the results – that three particular algorithms perform poorly, while the others don't – are not unexpected. We claim that these three algorithms are intrinsically weaker than the others. In support of this claim, we note that these exact three algorithms perform poorest on $\mathbf{u}_{0.5}(n)$. It was also shown in [9] that on a certain distribution of graphs quite different from both $\mathbf{u}_p(n)$ and $\mathbf{q}(n)$, these same three algorithms performed the poorest. Viewed this way, $\mathbf{q}(n)$ serves mainly to amplify the differences amongst these algorithms. A plausible explanation for why the others seem to remain robust on $\mathbf{q}(n)$ is that our choices of $n$ are too small. It is reasonable to suspect that some of them might perform poorer on larger $n$. As an extreme example, consider an exact algorithm. Obviously then, no matter how large $n$ is, it never performs poorly, in terms of solution-quality (though a solution might not be forthcoming in reasonable time).

Q2: *It is easy to see that it is possible to make the choices of language $\mathcal{L}$, decoder $D$ and the simulator $U$ in such a way that an arbitrary given optimization heuristic is favored or disfavored by the resulting test instances.*

This might be so, but this is not what we did. Our design of the **q**-sampler was not influenced in any way by any optimization heuristic. Indeed, if one examines the **q**-sampler algorithm, it is hard to imagine where a bias for or against any particular algorithm could have crept in even by accident.

Q3: *The approach seems to be less suited for comparing heuristics in an objective way than e.g. the uniform distribution.*

This is certainly the case. It is unfortunate that $\mathbf{m}(x)$ is not computable, let alone tractable. One has no choice then, but to resort to a heuristic approximation, which certainly introduces some degree of arbitrariness. The best one can do is to minimize the arbitrariness exhibited by the heuristic approximation chosen. We have tried to achieve this.

Consider arguing this another way. View the situation as a trade-off between objectivity and relevance. At one extreme is the uniform distribution – efficiently computable and objective – but, we argue, not very relevant (if one is to believe that the real world is highly structured, i.e., compressible). At the other extreme are distributions highly relevant to particular applications, but not objective. The distribution $\mathbf{q}(x)$, we argue, is in the middle; it is more relevant than the uniform distribution while being more objective than application-specific distributions.

Q4: *Is $\mathbf{q}(x)$, at all, an "approximation" to $\mathbf{m}(x)$?*

In a rigorous sense, it seems to us, no. The terminology has been revised to call it a *heuristic* analogue. We think it serves this role reasonably well, as argued earlier.

Q5: *The replacement of arbitrary programs by programs in nearly linear time might essentially change the situation (note that this replacement also obstructs the proof of Theorem 1).*

That it obstructs the proof is clear. That it essentially changes the situation is not clear. In particular, as explained at the end of Section 2, a reasonable phenomenon capturing the spirit of Theorem 1 is not ruled out if **m** is replaced by an efficiently samplable analog, provided it is reasonably diverse. Indeed, the experimental results may be viewed as a means of testing this question (given that testing under $\mathbf{m}(x)$ is infeasible). The results give convincing (even if partial) answers in this regard.

> Q6: *Dealing with the uniform distributions* $\mathbf{u}_{0.5}(x)$ *and* $\mathbf{u}_{0.9}(x)$*, in reality pseudorandom numbers have been used. In the given investigation context, this makes an essential difference, since a long sequence of pseudorandom numbers can be considered as a set of highly compressible data, while real random numbers are non-compressible by definition. Thus, what has been done in fact is that one set of compressible problem instances has been compared with another such set.*

This is a good argument. We respond to different aspects of it individually.

First, using pseudorandom numbers instead of true random numbers does not make an essential difference in this application. It is well known that the structural properties of graphs drawn from $\mathbf{u}_{0.5}(x)$ and $\mathbf{u}_{0.9}(x)$ in practice via pseudorandom numbers are in excellent agreement with those calculated from theory (which assumes perfect random numbers). Our own experience is consistent with this fact. The properties that our own experiments (and those of others in the published literature) reveal to be in excellent agreement are: average vertex degree, graph density, and maximum clique size. For instance, the maximum clique sizes found by exact algorithms on such random graphs are in excellent agreement with sharp estimates known from random graph theory derived by Matula and more rigorously by Bollobas & Erdős.

Second, the bitstrings we obtained (from our pseudorandom number generator) to generate the $\mathbf{u}(n)$ graphs were entirely unstructured as far as we could tell while the bitstrings we obtained from the **q**-sampler revealed themselves to be highly structured, even on casual visual inspection. It is reasonable to suspect that the bitstrings derived from pseudorandom numbers would pass simple tests for randomness; there is no doubt that the bitstrings from the **q**-sampler would not.

Finally, the graphs produced by pseudorandom generators are compressible. However, in any polynomial size family of such graphs, almost surely only a tiny fraction of those produced by the universal decompressor $\mathbf{q}(x)$ would arise. Said in a more down-to-earth way: if one repeatedly samples from $\mathbf{u}(x)$, one repeatedly gets graphs with the same structure (i.e., lack of it); in striking contrast, if one repeatedly samples from $\mathbf{q}(x)$ one gets graphs with a wide variety of structures.

**Acknowledgements**

The authors thank Marcus Peinado for helpful discussions, and for running preliminary experiments on some graphs drawn from $\mathbf{q}(1000)$ revealing the impracticality

of computing their maximum clique size in reasonable time. The authors thank the anonymous reviewers, especially reviewer #1, for their excellent comments that have greatly improved this paper.

## Appendix: A Practical Guide to the $\mathbf{q}(x)$ Sampler

In this appendix we describe the $\mathbf{q}(x)$ sampler at a level appropriate for users who might consider evaluating their algorithms on test data sampled from $\mathbf{q}(x)$. At the end of this section, we also give instructions on how to fetch the software.

A macro-level description of the sampling procedure is as follows:

Algorithm $\mathbf{q}(x)$-`Sampler`

1. Generate a random string $s$ in $\{0,1\}^m$.
2. Decode $s$ to give a pair $(P, x)$, where $P$ is a program in the formalism explained in Section 3 and $x$ is a binary string.
3. Run program $P$ on binary string $x$ and output binary string $y$.

The string $s$ is then a description of the string $y$.

In practice, the $\mathbf{q}(x)$-`Sampler` is called repeatedly and only those strings $y$ whose lengths exceed that of $s$ by a certain threshold are kept, all others are thrown away. This ensures that the remaining strings are compressible by a certain minimum amount.

### The Implementation of Step 2

The algorithm to decode a string $s$ into a pair $(P, x)$ was designed to ensure that almost every program that could be the target of some string $s'$ of length $m$ had nonzero probability to be the target of a random string $s$. The decoding procedure is sketched below.

1. Interpret the first $\frac{\lfloor \log m \rfloor}{2}$ bits of $s$ as storing the length $|x|$ of $x$.
2. Interpret the next $|x|$ bits as storing $x$.
3. Let $r$ denote the number of bits remaining at this point. Interpret the first $\frac{\lfloor \log r \rfloor}{2}$ bits of these remaining $r$ bits as storing the number of noniterated operators in the program $P$.
4. Let $o$ denote the number of operators found above. Interpret the next $\frac{\lfloor \log o \rfloor}{2}$ bits as storing the number of iterated operators. These are obtained by placing balanced parentheses around subsequences of non-iterated operators.
5. Interpret the next set of bits to indicate where the parentheses should be placed.
6. Interpret the next set of $o \times 3$ bits as the actual codes of the non-iterated operators. There are 8 such operators, so 3 bits suffice to identify individuals amongst them.

7. Interpret the next set of bits to give the length of every parameter to every operator.
8. Interpret the remaining bits as the actual parameters for all the operators, using the knowledge of the syntax of every operator and the knowledge of the length of every parameter computed in the previous step.

### AN EXAMPLE

This example was constructed from an actual run using the software discussed later. The following command was invoked at the `Unix` command line.

```
rsb 1 30 10 | nlt_dec 30 | nlt 100
```

The program `rsb` generates a random string of length 30 using 10 as the seed. The program `nlt_dec` decodes this string, taking the string length as a command-line argument, and outputs the string $x$, followed by the program P. The operators in P are identified by name and followed by their parameter strings. The operators in P are sequenced under the assumption that P is to be executed from left to right, with the first operator taking $x$ as the input and producing an output string, which becomes an input to the second operator and so on. The program `nlt` interprets P, and executes it on the string $x$.

The actual strings and program produced in this example were:

```
s   =    01001011110000001001101110010 0
x,P =    00 D 011 E 10 01 R0 0 0 .
y   =    1010
```

The output $y$ in this instance turned out to be shorter than $s$.

### HOW TO FETCH THE $\mathbf{q}(x)$ SAMPLER

The $\mathbf{q}(x)$ sampler may be retrieved by anonymous ftp as follows:

```
ftp ftp.cs.buffalo.edu
Name: anonymous
> cd users/jagota
> get nlt.README
> quit
```

The file `nlt.README` provides further instructions.

OTHER IMPORTANT USAGE DETAILS

The $\mathbf{q}(x)$ sampler can be used to generate compressible instances of any discrete problem, whether on binary strings, or on structured objects such as graphs. One needs to keep in mind the following issues however.

1. The good news is that the generation of compressible strings is quite efficient. In our case, one of every six seed strings turned out to generate a highly compressible string. To ensure this efficiency, however, the seed strings need to be longer than a certain minimum length. We found that seed strings of length less than 70 almost never generated compressible strings $y$, perhaps because the programs P that result from such strings are too short and trivial to generate compressible strings.
2. The bad news is that our implementation of `nlt` is quite inefficient because the code that performs certain internal searches that need to be performed in order to execute certain operators in P is not optimized.

## References


1. S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. In *The Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science*, pages 14–23, 1992.
2. G. Bilbro, R. Mann, T.K. Miller, W.E. Snyder, D.E. Van den Bout, and M. White. Optimization by mean field annealing. In D.S. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 1, pages 91–98, San Mateo, 1989. (Denver 1988), Morgan Kaufmann.
3. R. Carraghan and P.M. Pardalos. An exact algorithm for the maximum clique problem. *Operations Research Letters*, 9: 375–382, 1990.
4. U. Feige, S. Goldwasser, L. Lovasz, S. Safra, and M. Szegedy. Approximating clique is almost np-complete. In *The Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*, pages 2–12, 1991.
5. Y. Gurevich and S. Shelah. Nearly linear time. In *Proceedings, Logic at Botik*, Lecture Notes in Computer Science No. 363, pages 108–118. Springer-Verlag, 1989.
6. J. Hertz, A. Krogh, and R.G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991.
7. J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, USA*, 79, 1982.
8. J.J. Hopfield. Neurons with graded responses have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, USA*, 81, 1984.
9. A. Jagota, Approximating maximum clique in a Hopfield-style network. *IEEE Transactions on Neural Networks* **6**(3): 724–735, 1995.
10. A. Jagota, L. Sanchis, and R. Ganesan. Approximating maximum clique using neural network and related heuristics. In D.S. Johnson and M. Trick, editors, *DIMACS Series: Second DIMACS Implementation Challenge* **26**: 169–204. AMS, 1996.
11. R.M. Karp. The probabilistic analysis of some combinatorial search algorithms. In J.F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 1–19. Academic Press, New York, 1976.
12. M. Li and P.M.B. Vitanyi. Kolmogorov complexity and its applications. In J. vanLeeuwen, editor, *Handbook of Theoretical Computer Science*, pages 187–254. Elsevier and MIT Press, Amsterdam/New York, 1990.
13. M. Li and P.M.B. Vitanyi. Average case complexity under the universal distribution equals worst-case complexity. *Information Processing Letters* **42**: 145–149, May 1992.

14. P. Miltersen. The complexity of malign ensembles. In *The Proceedings of the 6th Annual IEEE Conference on Structure in Complexity Theory*, pages 164–171, 1991.
15. E.M. Palmer. *Graphical Evolution*. Wiley, New York, 1985. Matula's theorem on page 76.
16. P.M. Pardalos and J. Xue. The maximum clique problem. *Journal of Global Optimization* **4**: 301–328, 1994.
17. C. Peterson and B. Söderberg. A new method for mapping optimization problems onto neural networks. *International Journal of Neural Systems* **1**: 3–22, 1989.