In *Artificial Intelligence*, **81** (19996) 17–29.

Generating Hard Satisfiability Problems

Bart Selman[*] David Mitchell[†] and Hector Levesque[‡]

Abstract

We report results from large-scale experiments in satisfiability testing. As has been observed by others, testing the satisfiability of random formulas often appears surprisingly easy. Here we show that by using the right distribution of instances, and appropriate parameter values, it is possible to generate random formulas that are hard, that is, for which satisfiability testing is quite difficult. Our results provide a benchmark for the evaluation of satisfiability-testing procedures.

1 Introduction

Many computational tasks of interest to AI, to the extent that they can be precisely characterized at all, can be shown to be NP-hard in their most general form. However, there is fundamental disagreement, at least within the AI community, about the implications of this. It is claimed on the one hand that since the performance of algorithms designed to solve NP-hard tasks degrades rapidly with small increases in input size, something will need to be given up to obtain acceptable behavior. On the other hand, it is argued that this analysis is irrelevant to AI since it is based

[*] AT&T Bell Laboratories, Murray Hill, NJ 07974, selman@research.att.com.

[†] Dept. of Computing Science, Simon Fraser University, Burnaby, Canada V5A 1S6, mitchell@cs.sfu.ca. This research was carried out while visiting AT&T Bell Laboratories.

[‡] Dept. of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, hector@ai.toronto.edu. Fellow of the Canadian Institute for Advanced Research.

Correspondence to: Bart Selman, AT&T Bell Laboratories, Murray Hill, NJ 07974. Telephone: (908) 582-2221. E-mail: selman@research.att.com.

on worst-case scenarios, and that what is really needed is a better understanding of how these procedures perform "on average".

The first computational task shown to be NP-hard, by Cook (1971), was propositional satisfiability or SAT: Given a formula of the propositional calculus, decide if there is an assignment to its variables that makes the formula true according to the usual rules of interpretation. Subsequent tasks have been shown to be NP-hard by proving they are at least as hard as SAT. Roughly, a task is NP-hard if a good algorithm for it would entail a good algorithm for SAT. Unlike many other NP-hard tasks (see Garey and Johnson (1979) for a catalogue), SAT is of special concern to AI because of its direct relationship to deductive reasoning (*i.e.*, given a collection of base facts Σ, a sentence α may be deduced iff $\Sigma \cup \{\neg\alpha\}$ is not satisfiable). Many other forms of reasoning, including default reasoning, diagnosis, planning and image interpretation, also make direct appeal to satisfiability. The fact that these usually require much more than the propositional calculus simply highlights the fact that SAT is a fundamental task, and that developing SAT procedures that work well in AI applications is essential.

We might ask when it is reasonable to use a sound and complete procedure for SAT, and when we should settle for something less. Do hard cases come up often, or are they always a result of strange encodings tailored for some specific purpose? One difficulty in answering such questions is that there appear to be few applicable *analytical* results on the expected difficulty of SAT (although see below). It seems that, at least for the time being, we must rely largely on empirical results.

A number of papers (some discussed below) have claimed that the difficulty of SAT on randomly generated problems is not so daunting. For example, an often–quoted result (Goldberg 1979) suggests that SAT can be readily solved "on average" in $O(n^2)$ time. This does not settle the question of how well the methods will work *in practice*, but at first blush it does appear to be more relevant to AI than contrived worst cases.

The big problem is that to examine how well a procedure does on average one must assume a distribution of instances. Indeed, as we will discuss below, Franco and Paull (1983) refuted the Goldberg result by showing that it was a direct consequence of the choice of distribution. It is not that Goldberg had a clever algorithm, or that the problem is easy, but that he had used a distribution with a preponderance of easy instances. That is, from the space of all problem instances, they sampled in a way that produced almost no hard cases.

Nevertheless, papers continue to appear purporting to empirically demonstrate

the efficacy of some new procedure, but using just this distribution (*e.g.*, Hooker, 1988; Kamath *et al.* 1990), or presenting data suggesting that very large satisfiability problems — with thousands of propositional variables — can be solved. How are we to evaluate these empirical results, given the danger of biasing the sample to suit the procedure in question, or of simply using easy problems (even if unwittingly)?

In this paper, we present empirical results showing that random instances of satisfiability can be generated in such a way that easy and hard sets of instances (for a particular SAT procedure, anyway) are predictable in advance. If we care about the *robustness* of the procedures we develop, we will want to consider their performance on a wide spectrum of examples. While the easy cases we have found can be solved by almost *any* reasonable method, it is the hard cases ultimately that separate the winners from the losers. Thus, our data is presented as challenging test material for developers of SAT procedures (see Selman *et al.*, 1992, for example).

The SAT procedure we used for our tests is the Davis-Putnam procedure, which we describe below. We believe this was a good choice for two reasons: First, it is basically a variant of resolution (Vellino 1989; Galil 1977), the most widely used general reasoning method in AI; second, almost all empirical work on SAT testing has used one or another refinement of this method, which facilitates comparison. We suspect that our results on hard and easy areas generalize to *all* SAT procedures.

The rest of the paper is organized as follows. In the next two sections we describe the Davis-Putnam procedure, and consider its performance on one distribution of formulas, the fixed clause-length model. We show that with the right parameter values, it produces computationally challenging SAT instances. In the following section we consider a second distribution, the constant-density model, and argue that it is not useful in the evaluation of satisfiability-testing procedures. We then briefly review related work, and summarize our results.

2 The Davis-Putnam Procedure

The Davis-Putnam (DP) procedure (Davis and Putnam 1960) is sketched in Figure 1. It takes as input a set of clauses Σ over a set of variables V, and returns either "satisfiable" or "unsatisfiable." (A clause is a disjunction of literals. A set of clauses represents a conjunction of disjunctions, *i.e.*, a formula in conjunctive normal form (CNF).) DP performs a backtracking depth-first search in the

PROCEDURE DP
Given a set of clauses Σ defined over a set of variables V:

- If Σ is empty, return "satisfiable".
- If Σ contains an empty clause, return "unsatisfiable".
- (Unit-Clause Rule) If Σ contains a unit clause c, assign to the variable mentioned the truth value which satisfies c, and return the result of calling DP on the simplified formula.
- (Splitting Rule) Select from V a variable v which has not been assigned a truth value. Assign it a value, and call DP on the simplified formula. If this call returns "satisfiable", then return "satisfiable". Otherwise, set v to the opposite value, and return the result of calling DP on the re-simplified formula.

Figure 1: The DP procedure with unit propagation.

space of all truth assignments, incrementally assigning truth values to variables and simplifying the formula. If no new variable can be assigned a value without producing an empty clause, it backtracks by changing a previously made assignment. The performance of simple backtracking is greatly improved by employing *unit propagation*: Whenever a unit clause (one containing a single literal) arises, the variable occurring in that clause is immediately assigned the truth value that satisfies it. The formula is then simplified, which may lead to new unit clauses, and so on. This propagation process can be executed in time linear in the total number of literals. DP combined with unit propagation is one of the most widely used methods for propositional satisfiability testing. (It is also common to include the "pure literal rule", which we excluded from this implementation because it is relatively expensive, and seemed to provide only a modest improvement on the formulas we consider here.)

3 The fixed clause-length model

In this section, we study formulas generated using the fixed clause-length model, which we call *Random K-SAT*. There are three parameters: the number of variables N, the number of literals per clause K, and the number of clauses L. To keep the volume of data presented manageable and yet give a detailed picture, we limit

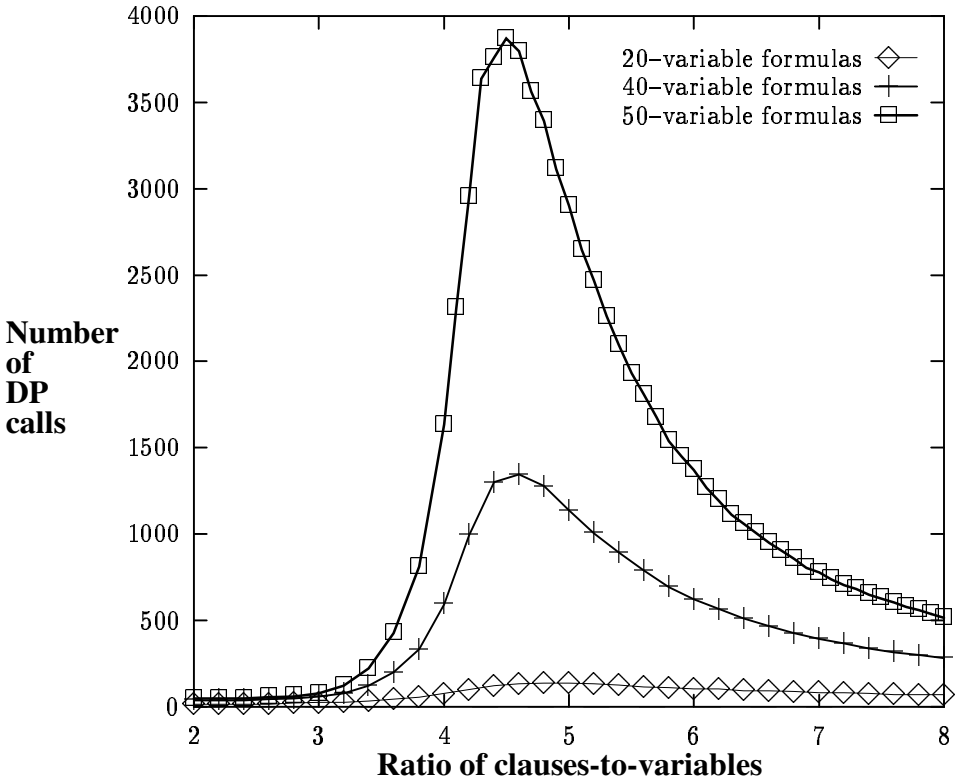

Figure 2: Median number of recursive DP calls for Random 3-SAT formulas, as a function of the ratio of clauses–to–variables.

our attention to formulas with $K = 3$, that is, random 3-SAT. For a given N and L, an instance of random 3-SAT is produced by randomly generating L clauses of length 3. Each clause is produced by randomly choosing a set of 3 variables from the set of N available, and negating each with probability 0.5.

We now consider the performance of DP on such random formulas. Figure 2 shows the total number of recursive calls by DP to find one satisfying assignment, or to determine that the formula is unsatisfiable. There are three curves, for formulas with 20, 40, and 50 variables. Along the horizontal axis is the number of clauses in the formulas tested, normalized through division by the number of variables. Each data point gives the median number of calls for a sample size of

10,000.[1]

In Figure 2, we see the following pattern: For formulas that are either relatively short or relatively long, DP finishes quickly, but the formulas of medium length take much longer. Since formulas with few clauses are *under-constrained* and have many satisfying assignments, an assignment is likely to be found early in the search. Formulas with very many clauses are *over-constrained* (and usually unsatisfiable), so contradictions are found easily, and a full search can be completed quickly. Finally, formulas in between are much harder because they have relatively few (if any) satisfying assignments, but the empty clause will only be generated after assigning values to many variables, resulting in a deep search tree. Similar under- and over-constrained areas have been found for random instances of other NP-complete problems (Cheeseman *et al.* 1991; Williams and Hogg 1992).

The curves in Figure 2 are for *all* formulas of a given size, that is they are composites of satisfiable and unsatisfiable subsets. In Figure 3 the median number of calls for 50-variable formulas is factored into satisfiable and unsatisfiable cases, showing that the two sets are quite different. The extremely rare unsatisfiable short formulas are very hard, whereas the rare long satisfiable formulas remain moderately difficult. Thus, the easy parts of the composite distribution appear to be a consequence of a relative abundance of short satisfiable formulas or long unsatisfiable ones.

To understand the hard area in terms of the likelihood of satisfiability, we experimentally determined the probability that a random 50-variable instance is satisfiable (Figure 4). There is a remarkable correspondence between the peak on our recursive calls curve and the point where the probability that a formula is satisfiable is 0.5. The main empirical conclusion we draw from this is that *the hardest area for satisfiability is near the point where 50% of the formulas are satisfiable*.

This "50% satisfiable" point seems to occur at a fixed ratio of the number of clauses to the number of variables: when the number of clauses is about 4.3 times the number of variables. There is a boundary effect for small formulas: for formulas with 20 variables, the point occurs at 4.55; for 50 variables, at 4.31; and for 140 variables, at 4.3. While we conjecture that this ratio approaches about 4.25 for very large numbers of variables, it remains a challenging open problem to *analytically* determine the "50% satisfiable" point as a function of the number

[1]The means of the number of calls are influenced by a very small number of very large values. As the median is less sensitive to such "out-liers", it appears to be a more informative statistic.

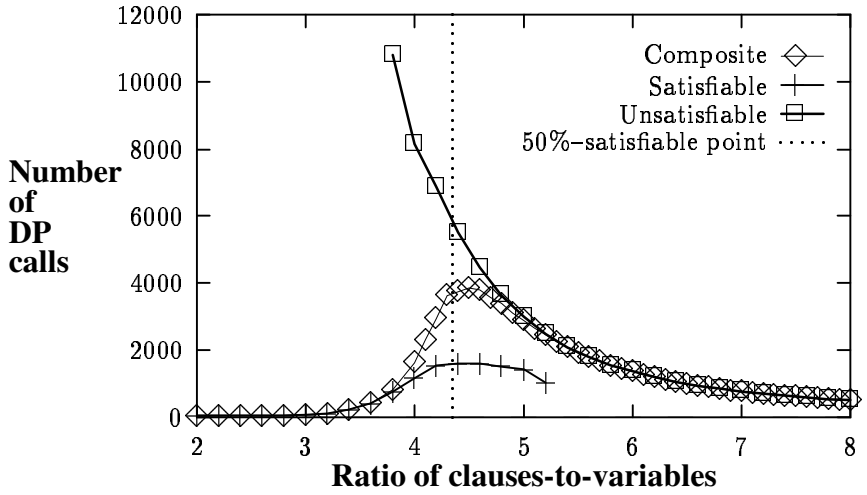

Figure 3: Median DP calls for 50–variable Random 3-SAT as a function of the ratio of clauses–to–variables.

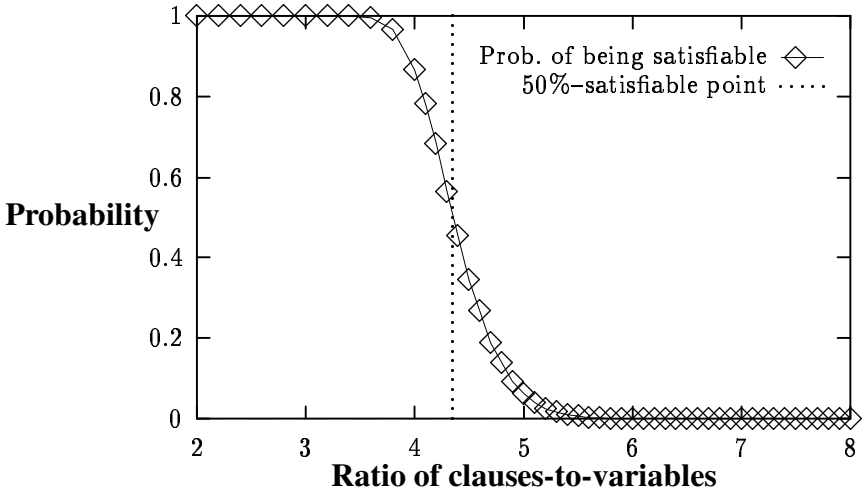

Figure 4: Probability of satisfiability of 50–variable formulas, as a function of the ratio of clauses–to–variables.

of variables.[2]

As an aside, we would like to mention some of the recent progress on a theoretical analysis of the 50% point. The phenomenon that we see in Figure 4 is called a *threshold phenomenon*. It is quite common in random graphs. For example, it has been shown that randomly generated graphs with n nodes will almost certainly have cliques of size $2logn$ or smaller but almost certainly will not have any larger cliques. In other words, the probability of a random graph having a clique of size k is represented by a threshold function with the threshold at $2logn$ (Matula 1972; Bollobas 1985). It appears much harder to show formally that there exists a threshold phenomenon for SAT, even though our experimental results certainly do suggest so. A recent breakthrough was the establishment of a threshold for 2SAT (Goerdt 1991; Chvatal and Reed 1992; de la Vega 1992). The threshold lies at a ratio of 1, *i.e.*, for a n variable formula, if one generates less then n binary clauses the formula is almost certainly satisfiable, whereas when generating more than n clauses the formula is almost certainly unsatisfiable. For kSAT, with $k \geq 3$, no such result has yet been obtained. Some preliminary results are known for 3SAT. Broder *et al.* (1993) show that for random 3SAT instances with a ratio of variables to clauses of less than 1.7 almost all formulas are satisfiable. Moreover, for a ratio of 5.2 or higher, it has been shown that almost all formulas are unsatisfiable (Dubois 1991; Chvatal and Szemeredi 1988). Note that these ratios are still quite far removed from our experimentally determined threshold at 4.3.

We now return to the discussion of our experimental results. To get some further insight into the search performed by DP, we consider two more statistics. This time we let DP search the full space, *i.e.,* we do *not* stop the procedure as soon as one assignment is found. Figure 5 gives the total number of recursive calls. (Again, each data point gives the median value of a sample 10,000 formulas.) We see that the search space monotonically increases for decreasing ratios of clauses to variables. This is consistent with our expectation that the fewer clauses in the formula the longer it takes before DP will run into a contradiction, and the deeper the search tree. We also see that the search tree is somewhat larger for satisfiable formulas, at least on average. This can be intuitively explained by the fact that when an assignment is found DP tends to assign many variables in order to satisfy all clauses, so it generates a relatively long branches for satisfying assignments.

[2]Recent, more detailed experiments confirm our observations about the 50% point (Crawford and Auton 1993; Larrabee and Tsuji 1993).

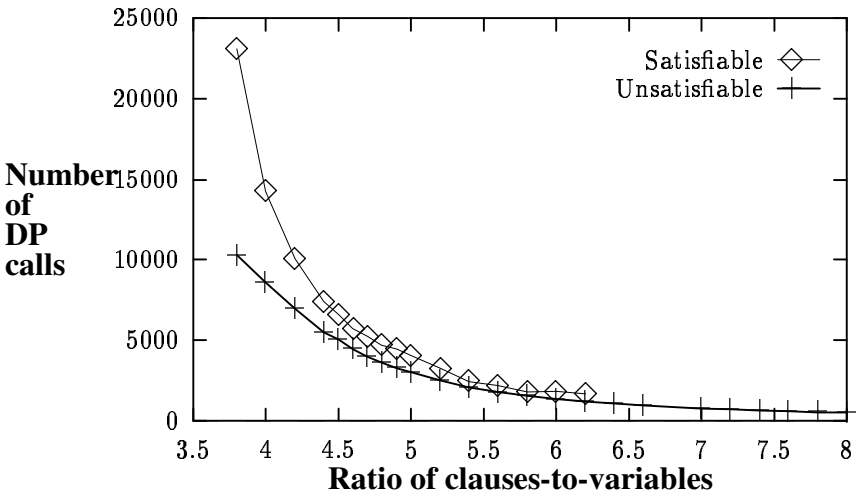

Figure 5: Searching for all assignments. Median DP calls for 50–variable Random 3-SAT as a function of the ratio of clauses–to–variables.

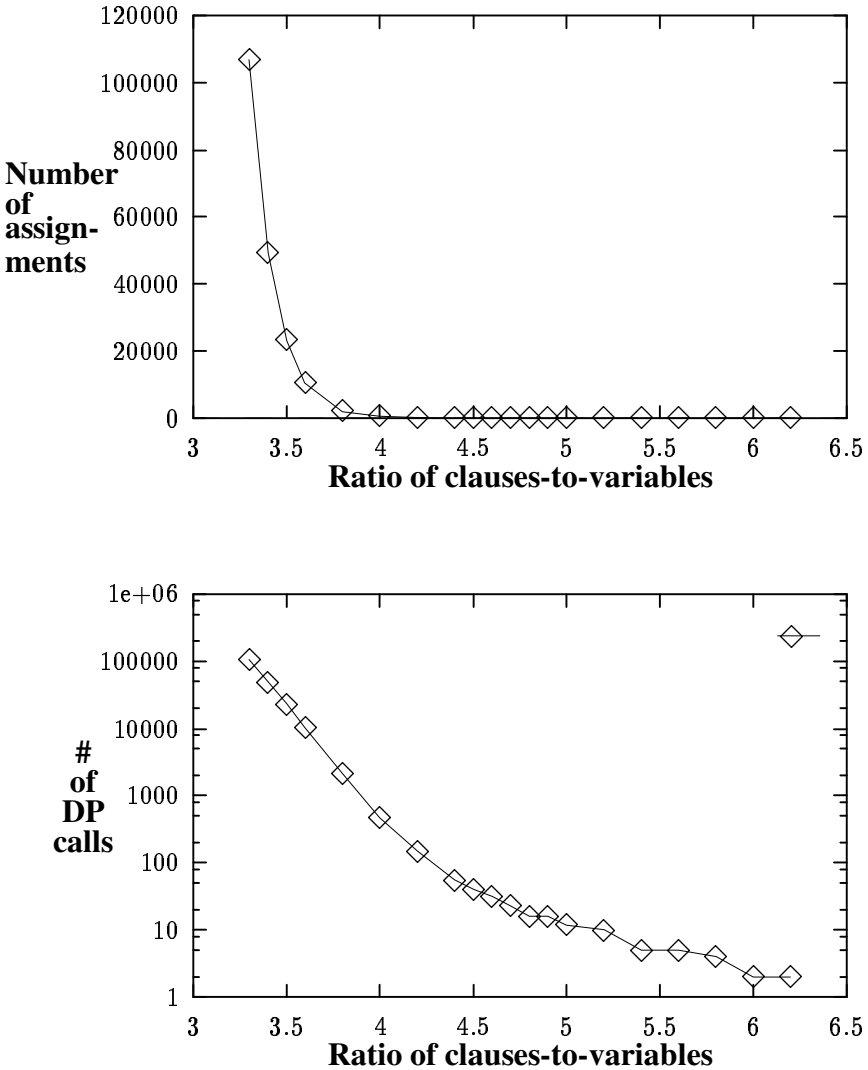

Figure 6: Median number of satisfying assignments for satisfiable 50–variable Random 3-SAT formulas as a function of the ratio of clauses–to–variables. Note the logarithmic scale in the bottom panel.

The full search space for under-constrained formulas quickly becomes extremely large. For example, at a ratio of 3.3, we have over 200,000 recursive calls, for satisfiable instances. Note however that such a large search space is no problem when searching for a *single* assignment. The reason for this is that there are a tremendous number of satisfying assignments for the under-constrained formulas. In fact, below the 50% satisfiable point the number of satisfying assignments grows exponentially when decreasing the ratio of variables to clauses. This is shown in Figure 6. In this figure, we give the median number of satisfying assignments for satisfiable random 3-SAT formulas as a function of the ratio of clauses to variables. (Each data point is based on 10,000 instances.) For example, at a ratio of 3.3, the formulas have around 110,000 satisfying assignments. Therefore, when searching for a single satisfying assignments, DP will often find one early on in the search.

Finally, note that we did not specify a method for choosing which variable to guess in the "splitting" step of DP. In our implementation, we simply set variables in lexical order (except when there are unit clauses.) DP can be made faster by using clever selection strategies (*e.g.*, Zabih and McAllester 1988; Kleine-Büning 1992), but it seems unlikely that such heuristics will *qualitatively* alter the easy-hard-easy pattern. The formulas in the hard area appear to be the most challenging for the strategies we have tested, and we conjecture that they will be for every (heuristic) method.[3]

4 The constant-density model

We now examine formulas generated using the constant-density model. The model has three parameters: the number of variables N, and number of clauses L as before; but instead of a fixed clause length, clauses are generated by including a variable in a clause with some fixed probability P, and then negating it with probability 0.5. Large formulas generated this way very often have at least one empty clause and several unit clauses, so that they tend to be either trivially unsatisfiable, or easily shown satisfiable. Thus, the more interesting results are for the modified version in which empty and unit clauses are disallowed. This distribution we call *Random P-SAT*.

Analytic results by Franco and Paull (1983) suggest that one probably cannot generate computationally challenging instances from this model, and our experi-

[3]Recent results by Larrabee and Tsuji (1993) support this conjecture. They study a satisfiability procedure very different from the Davis-Putnam procedure, but find the same hard and easy areas.

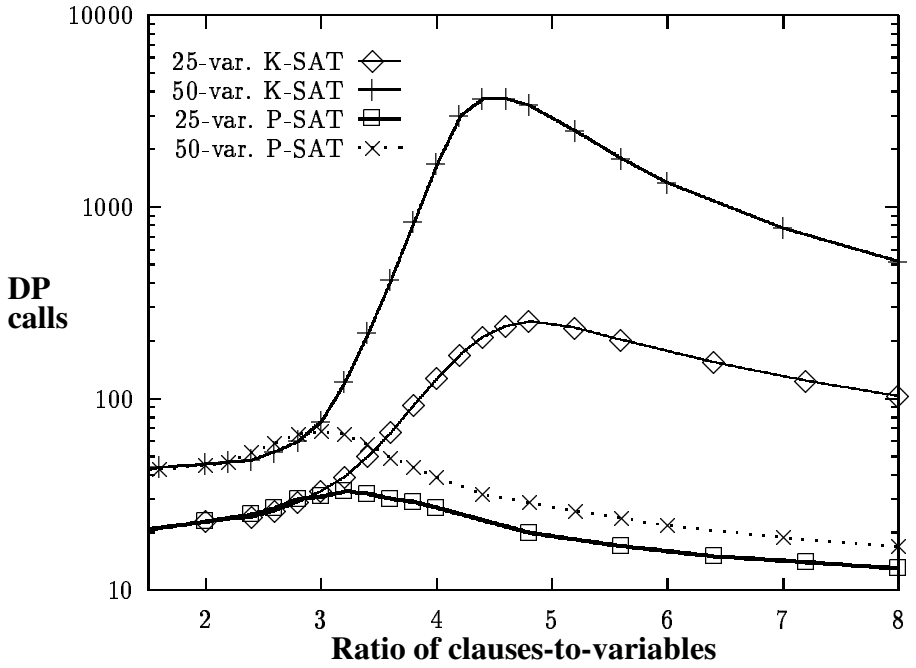

Figure 7: Comparison of median DP calls for fixed–length (3-SAT) and constant–density formulas (average clause length 3), with 25 and 50 variables.

ments confirm this prediction. In Figure 7, we compare the number of recursive DP calls to solve instances of random P-SAT with the same numbers for random 3-SAT. In this case, we have set the probability P to give an average clause length of 3. Although we see a slight easy-hard-easy pattern (previously noted by Hooker and Fedjki, 1989), the hard area is not nearly as pronounced as that for random 3-SAT formulas of similar size, and in absolute terms the random P-SAT formulas are much easier. Note that we are using a logarithmic scale for the number of DP calls, so the greater height of the 3-SAT curve represents much greater difficulty. Further, the rate of growth in difficulty of the random 3-SAT formulas is much higher. Table 1 gives the median number of DP calls at the hardest point on the curve for random 3-SAT and random P-SAT formulas, as a function of the number of variables (For comparison, we also include here the numbers for random P-SAT with expected clause length of 4). Again, the table clearly shows a large difference in growth rate for the two random formula models.

	formula distribution		
vars	K-SAT (3)	P-SAT (3)	P-SAT (4)
25	253	33	53
50	3,683	68	157
75	46,915	111	392

Table 1: Number of DP calls at hardest point, for fixed–length (K-SAT) and constant–density (P-SAT) formulas, with 25, 50, and 75 variables. The numbers in brackets indicate the number of literals per clause (for the P-SAT model, this is the average number of literals per clause).

In conclusion, our experimental results and Franco and Paull's analysis (discussed below) strongly suggest that the constant-density model is not suitable for the evaluation of satisfiability testing procedures.

5 Related work

There is a large body of literature on testing the satisfiability of random formulas, reporting mostly analytic results. The main impetus for this research was early experimental results by Goldberg (1979), which suggested that SAT might in fact be easily solvable, on average, using DP. Franco and Paull (1983) showed that Goldberg's positive results were a direct consequence of the distribution used — a variant of the constant-density model — and thus overly optimistic. Goldberg's formulas were so easily satisfiable that an algorithm which simply tried randomly generated assignments would, with probability 1, finds a satisfying assignment in a constant number of guesses. Further analytic results for the constant-density model can be found in Franco (1986), Franco and Ho (1988), and Purdom and Brown (1987).

Franco and Paull (1983) also investigated the performance of DP on random formulas with a fixed clause length of 3, and suggested that it might be more useful for generating random instances, an hypothesis we have confirmed experimentally here. They showed that for any fixed ratio of clauses to variables, if DP is forced to find *all* satisfying truth assignments, its expected time will be exponential in the number of variables, with probability approaching 1 as the number of

variables approaches infinity while keeping the ratio of clauses to variables fixed. Unfortunately, this result does not directly tell us much about the expected time to find a single assignment.

A recent result by Chvatal and Szemeredi (1988) can be used to obtain some further insights. Extending a ground-breaking result by Haken (1985), they showed that any resolution strategy requires exponential time with probability approaching 1 on formulas where the ratio of clauses to variables is a constant greater than 5.2. (They also show that with probability approaching 1 such formulas are unsatisfiable.) Given that DP corresponds to a particular resolution strategy, as mentioned above, it follows that on such formulas the average time complexity of DP is exponential.

This result may appear inconsistent with our claim that over-constrained formulas are easy, but it is not. Our results show that there is an easy-hard-easy pattern for formulas with a given (fixed) number of variables when varying the number of clauses and thus the ratio of clauses to variables. Chvatal and Szemeredi's applies to the case where we increase *both* the number of variables and clauses, while keeping the ratio of clauses to variables fixed.[4] They show that growth of the DP tree is an exponential function in the number of variables n. Our results suggest that the exponential for ratios in the hard area growths much faster than in the easy (over-constrained) area. This has been confirmed by recent results for a highly-optimized variant of DP developed by Crawford and Auton (1993). Experiments show that in the hard area, at a ratio of 4.3, their procedure scales with $2^{(n/17)}$, whereas at a ratio of 10, the scaling is with $2^{(n/57)}$. This difference in growth rate has important practical consequences. For example, our DP consistently takes only several seconds to determine the unsatisfiability of 1000-variable, 50,000-clause instances of 3-SAT, even though there are some 300-variable 1290-clause formulas that it cannot practically solve. So, what we have called the "easy area," is definitely easy compared to the formulas in the hard area (around the 50% point), and for lower values of n is often much too easy to be used as test formulas. But for formulas with larger numbers of variables, eventually the truly easy area will occur only at ever higher ratios of clauses to variables.

Turning to under-constrained formulas, the behavior of DP can be explained by the fact that they tend to have very many satisfying truth assignments — see Figure 6 — so that the procedure almost always finds one early in the search.

[4] To state this differently, consider Figure 2. Chvatal and Szemeredi analyze the scaling behavior going in the vertical direction at a given fixed ratio of clauses to variables.

Other analytic results for random 3-SAT are reviewed in Chao and Franco (1990).

Not only are our experimental results consistent with the analytic results, they also provide a much more fine-grained picture of how 3-SAT behaves in practice. One reason for the limitations of analytic results is the complexity of the analysis required. Another is that they are asymptotic, *i.e.*, they hold in the limit as the number of variables goes to infinity, so they do not necessarily tell us much about formulas that have only a modest number of variables (say, up to a few thousand) as encountered in practice.

Finally, we would like to mention the valuable contribution of a recent paper by Cheeseman *et al.* (1991). This paper explores the hardness of random instances of various NP-complete problems, their main results being for graph coloring and Hamiltonian circuit problems. They observe a similar easy-hard-easy pattern as a function of one or more parameters of instance generation. They also give some preliminary results for satisfiability. Unfortunately, they do not describe exactly how the formulas are generated, and their findings are based on relatively small formulas (up to 25 variables). Also, their backtrack search procedure does not appear to incorporate unit resolution, which would limit its ability to find quick cutoff points. Possibly because of the preliminary nature of their investigation, they observe that they do not know how to generate hard SAT instances except via transformation of hard graph coloring problems. Their hard area appears to be at a different place than in our data, suggesting that the mapping process generates a distribution somewhat different than random 3-SAT. Note that when formulas are not generated randomly, but encode some other NP-complete problem, such as graph coloring, the ratio of clauses to variables may not be a good indicator of expected difficulty. In fact, it is possible to arbitrarily change the clause-to-variable ratio of a formula by artificially "padding" it, without substantially affecting its difficulty.

6 Conclusions

There has been much debate in AI on the importance of worst-case complexity results, such as NP-hardness results. In particular, it has been suggested that satisfiability testing might be quite easy on average.

We have carried out a detailed study of the average-case difficulty of SAT testing for random formulas. We confirmed previous observations that many instances are quite easy, but we also showed how hard instances can be generated.

The fixed clause-length model with roughly 4.3 times as many clauses as variables gives computationally challenging instances which have about a 0.5 probability of being satisfiable. Randomly generated formulas with many more or fewer clauses are quite easy.

Our data provide two important lessons. The first is that the constant-density model is inappropriate for evaluating satisfiability procedures, since it seems to be dominated by easy instances for all values of the parameters. The second is that it is not necessarily the case that generating larger formulas provides harder formulas. For example, our DP can solve 1000 variable 3000 clause (under-constrained) and 1000 variable 50000 clause (over-constrained) random 3-SAT instances in seconds. On the other hand, it cannot solve random 3-SAT instances with 300 variables and 1290 clauses.

Because random 3-SAT instances can be readily generated, those from the hard area can be very useful in the evaluation of satisfiability testing procedures, and algorithms for related tasks such as Boolean constraint satisfaction. We hope that our results will help prevent further inaccurate or misleading reports on the average case performance of SAT procedures.

Acknowledgments

We thank Jim Crawford and Larry Auton for providing us with their very efficient tableau code, which we modified to obtain a fast implementation of DP. We thank Henry Kautz for many useful discussions and comments, and Fahiem Bacchus for helpful comments on an earlier draft. The first and third authors were funded in part by the Natural Sciences and Engineering Research Council of Canada, and the Institute for Robotics and Intelligent Systems.

References

Bollobas, B. (1985). *Random Graphs.* Academic Press, London, 1985.

Broder, A., Frieze, A., and Upfal, E. (1993). On the Satisfiability and Maximum Satisfiability of Random 3-CNF Formulas. *Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1993, 322–330.

Buro, M. and Kleine-Büning, H. (1992). Report on a SAT competition. Technical Report # 110, Dept. of Mathematics and Informatics, University of Paderborn, Germany, November 1992

Chao, Ming-te and Franco, John (1990). Probabilistic Analysis of a Generalization of the Unit-Clause Literal Selection Heuristics for the k Satisfiability Problem. *Inform. Sci.*, 51, 1990, 23–42.

Chvatal,V. and Reed, B. (1992). Mick gets some (the odds are on his side). *Proc. FOCS–92,* 1992, 620–627.

Chvatal,V. and Szemeredi, E. (1988). Many hard examples for resolution. *JACM*, val. 35, no. 4, 1988, 759–208.

Cheeseman, Peter and Kanefsky, Bob and Taylor, William M. (1991). Where the Really Hard Problems Are. *Proceedings IJCAI–91,* 1991, 163–169.

Cook, S.A. (1971). The complexity of theorem-proving procedures. *Proceedings of the 3rd Annual ACM Symposium on the Theory of Computing,* 1971, 151–158.

Crawford, J.M. and Auton, L.D. (1993). Experimental results on the cross-over point in satisfiability problems. *Proc. AAAI93*, to appear.

Davis, M. and Putnam, H. (1960). A computing procedure for quantification theory. *J. Assoc. Comput. Mach.*, 7, 1960, 201–215.

de la Vega, W.F. (1992). On random 2SAT. Manuscript 1992.

Dubois, O. (1991). Counting the number of solutions for instances of satisfiability. *Theoretical Computer Science*, 81, 1991, 65–75.

Franco, J. (1986). Probabilistic analysis of algorithms for NP-complete problems. United States Air Force Annual Scientific Report, 1986.

Franco, J. and Ho, Y.C. (1988). Probabilistic performance of a heuristic for the satisfiability problem. *Discrete Applied Math.*, 22, 1988, 35-51.

Franco, J. and Paull, M. (1983). Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem. *Discrete Applied Math.* 5, 1983, 77–87.

Galil, Z. (1977). On the complexity of regular resolution and the Davis-Putnam procedure. *Theoretical Computer Science*, 4, 1977, 23–46.

Gallo, G. and Urbani, G. (1989). Algorithms for Testing the Satisfiability of Propositional Formulae. *J. Logic Programming*, 7, 1989, 45–61.

Garey, M.R. and Johnson, D.S. (1979). *Computers and Intractability, A Guide to the Theory of NP-Completeness.* New York, NY: W.H. Freeman, 1979.

Goerdt, A. (1991). A threshold for unsatisfiability. Manuscript 1991.

Goldberg, A. (1979). *On the complexity of the satisfiability problem.* Courant Computer Science Report, No. 16, New York University, NY, 1979.

Goldberg, A., Purdom, Jr. P.W., and Brown, C.A. (1982). Average time analysis of simplified Davis-Putnam procedures. *Information Process. Lett.*, 15, 1982, 72–75; see also "Errata", vol. 16, 1983, p. 213.

Haken, A. (1985). The intractability of resolution. *Theoretical Computer Science*, 39, 1985, 297–308.

Hooker, J.N. (1988). Resolution vs. cutting plane solution of inference problems: Some computational experience. *Operations Research Letter*, 7(1), 1988, 1–7.

Hooker, J.N. and Fedjki, C. (1989). *Branch-and-cut solution of inference problems in propositional logic*, Working paper 77-88-89, Graduate School of Industrial Administration, Carnegie Mellon University, 1989.

Kamath, A.P., Karmarker, N.K., Ramakrishnan, K.G., and Resende, M.G.C. (1990). Computational experience with an interior point algorithm on the satisfiability problem. *Proceedings, Integer Programming and Combinatorial Optimization*, Waterloo, Canada: Mathematical Programming Society, 1990, 333–349.

Larrabee, T. and Tsuji, Y. (1993) Evidence for a Satisfiability Threshold for Random 3CNF Formulas, *Proc. of the AAAI Spring Symposium on AI and NP-hard problems*, Palto Alto, CA, 1993.

Matula, D.W. (1972). The employee party problem. *Notices of the American Mathematical Society*, 19, 1972, A-382.

Mitchell, D., Selman, B., and Levesque, H.J. (1992). Hard and easy distributions of SAT problems. *Tenth National Conference on Artificial Intelligence (AAAI-92)*, San Jose, CA, July 1992, 459–465.

Purdom, Jr. P.W., and Brown, C.A. (1987). Polynomial average-time satisfiability problems. *Inform. Sci.*, 41, 1987, 23–42.

Selman, B., Levesque, H.J., Mitchell, D. (1992) GSAT: A new method for solving hard satisfiability problems. *Tenth National Conference on Artificial Intelligence (AAAI-92)*, San Jose, CA, July 1992, 440–446.

Vellino, A. (1989). The complexity of automated reasoning. Ph.D. thesis, Dept. of Philosophy, University of Toronto, Toronto, Canada, 1989.

Williams, C.P. and Hogg, T. (1992) Using deep structure to locate hard problems. *Proc. AAAI-92,* San Jose, CA, July 1992, 472–277.

Zabih, R. and McAllester, D. (1988). A rearrangement search strategy for determining propositional satisfiability. *Proc. AAAI-88*, 1988, 155–160.