Sampling Combinatorial Spaces Using Biased Random Walks

Jordan Erenrich and Bart Selman*
Dept. of Computer Science
Cornell University
{erenrich, selman}@cs.cornell.edu

Abstract

For probabilistic reasoning, one often needs to sample from a combinatorial space. For example, one may need to sample uniformly from the space of all satisfying assignments. Can state-of-the-art search procedures for SAT be used to sample effectively from the space of all solutions? And, if so, how uniformly do they sample? Our initial results find that on the positive side, one can find all solutions to a given instance. Nevertheless, sampling can be highly biased. This research provides a starting point for the development of more balanced procedures.

1 Introduction

There has been much recent progress in the area of algorithms for solving Boolean satisfiability problems. Current methods can handle problem instances with up to a million variables. The emphasis in this work has been on the decision version of the satisfiability problem. A problem instance is either shown satisfiable by producing a satisfiable assignment or by proving that no such solution exists. A natural extension of this problem, for the satisfiable case, is to consider whether one can sample from the set of all satisfying assignments. An effective sampling procedure would be of use in a large range of probabilistic reasoning tasks. In the worst-case, the sampling problem is even harder than the satisfiability problem [Roth, 1996]. However, this leaves open the question how of hard it is to sample solutions from practical problem instances. In particular, can state-of-the-art random walk style procedures for SAT, such as WalkSAT [Selman et al., 1994], be used to sample from the space of satisfying assignments?

We investigated this question, and our key findings are as follows. On reasonable size problem instances, the WalkSAT procedure does reach all satisfying assignments. However, the sampling can be highly non-uniform (several orders of magnitude in the probability of reaching solutions). Quite unexpectedly, we found that sampling on our structured problem instance is much more uniform than on our hard random problem instance. This shows that, in a sense, the search on a random problem instance is more complex. We also considered the question of how to improve the uniformity of our sampling approach. Our results show, again somewhat surprisingly, that coverage does not become more uniform with more noise injected in the random walk. (Of course, the effectiveness of finding solutions degrades, but apparently, this degradation is proportional for all solutions.)

Because of space limitations, we present only some of our key data. We hope this work will lead to a better understanding of random walk style procedures for sampling solutions, including possible ways for getting more uniform sampling.

Our work should be contrasted with the more standard Markov Chain Monte Carlo (MCMC) approaches, such as simulated annealing (SA). For MCMC methods, one can show that the underlying Markov Chain eventually reaches an equilibrium distribution (although this can take exponential time). Unfortunately, formal guarantees are of little practical value. For example, even for finding a single satisfiable assignment, SA has been shown to be orders of magnitude slower than WalkSAT [Selman et al., 1994] (even ignoring the issue of whether SA reaches the stationary distribution). The new class of local search methods for SAT are based on ideas from biased random walks. These methods are designed to reach some solution as quickly as possible. Such strategies have proven to be quite effective, but they inherently exploit the fact that they operate far away from equilibrium. Nevertheless, as our results show, they can reach all solutions on non-trivial problem instances. Moreover, our work may suggest ways of making the sampling more uniform, possibly leading to a new more powerful class of sampling methods.

2 Sampling Experiments

We consider a hard random 3-SAT instance and a more structured problem from a planning application. In order to generate a random 3-SAT instance, we first created 100 uniform random 3-SAT instances, with 70 variables and 301 clauses. We chose an instance with 2531 solutions for our detailed

Table 1: Summary Statistics from Coverage Experiments

INSTANCE	RUNS	HITS RAREST	HITS COMMON	COM.-RARE RATIO
RANDOM	$50*10^6$	53	$9*10^5$	$1.7*10^4$
LOGISTIC	$1*10^6$	84	$4*10^3$	50

*Research supported in part by AFOSR, Darpa, and NSF

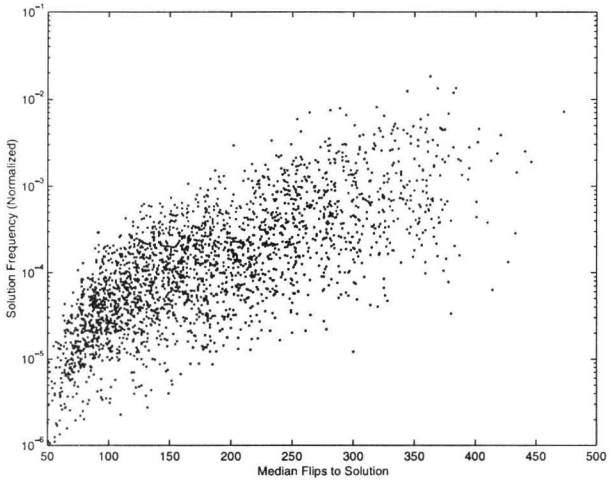

Figure 1: Frequency of finding a specific soln. vs. the median number of flips to find this solution for a 3-SAT problem.

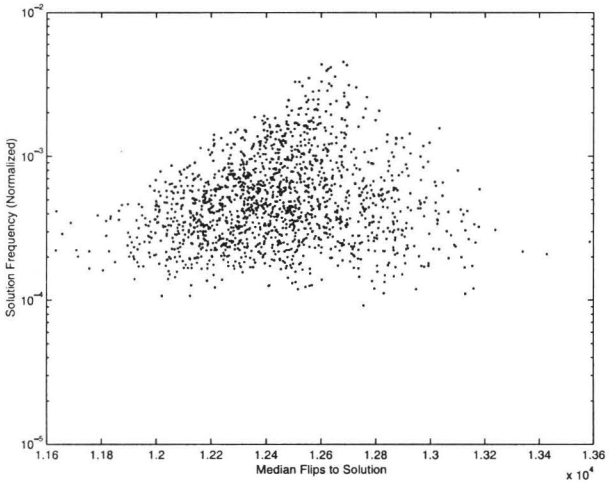

Figure 2: Frequency of finding a specific soln. vs. the median number of flips to find this solution for a logistic problem.

experiments. (Relsat [Bayardo and Pehoushek 2000] was used to systematically generate all solutions.) For the logistics planning formula, we started with *logistics.d.cnf* available from Satlib with over 10e10 solutions. To facilatate solution set sampling experiments, we added 52 unit clauses, constraining the instance to 1600 solutions.

For the 3-SAT and logistic instances, we performed 50e6 and 1e6 runs of WalkSAT, respectively. All of these runs successfully found a solution, and WalkSAT ultimately found every solution for both instances. Summary statistics for these experiments are available in Table 1.

Figures 1 and 2 display the relationship between "solution frequency" (how often a particular solution is reached) and "speed" (the median number of flips to reach that solution). Each point represents a unique solution. The figure further confirms the results from Table 1: WalkSAT's solution sampling can be highly non-uniform. Notice the scale on the y-axis. For the random instance we see a range over 4 orders of magnitude. For the logistic problem, we see that the solution sampling is also biased, albeit much less than on the 3-SAT problem. The figure also indicates that on the 3-SAT problem, rare solutions are typically found with fewer flips than common solutions. No such relation is evident for the logistic problem in Figure 2B.

3 Solution Space Structure

In this section, we attempt to gain insight into the structure of the solution space by analyzing hamming distances between solutions. For both the logistics and 3-SAT problems, we computed the hamming distance between all solutions. Thus, we generated $O(n^2)$ distances per problem, where n is the number of problem solutions.

Figures 3A and B show histograms of these distances for the 3-SAT and logistics problem, respectively. The 3-SAT hamming distance distribution is trimodal, while the logistics problem gives a unimodal distribution. These distributions suggest that the 3-SAT instance's solution space has multiple clusters of solutions, while the logistic problem's solution space has only a single cluster of solutions.

Of course, it is inherently difficult to analyze structure within these 70+ dimension solution spaces, and many other underlying structures could explain the results seen in Figure 3. Nevertheless, we believe that these figures are highly suggestive.

4 Obtaining More Uniform Coverage

In the previous experiment with the 3-SAT instance, we observed that rare solutions are generally found with fewer flips than common solutions. Intuitively, one needs to start relatively close to such a rare solution to find it. Following up on this observation, we decreased the cutoff value (number of flips until restart) for Walksat. As shown in Figure 4A, our experiments indicate that a lower cutoff indeed makes the sampling more uniform on the 3-SAT instance. Unfortunately, as a side effect, the median number of flips to reach a solution goes up. Thus, there appears to be a tradeoff between speed and uniformity on the 3-SAT instance. Note that theoretically, a cutoff of 0 (no flips) gives us purely uniform sampling. However, such a strategy is clearly infeasible for any instance of practical interest. (A cutoff below 40 makes it practically infeasible to find all solutions for our random problem instance.) Varying the cutoff used for the logistic problem did not have an observable effect on uniformity of sampling, as expected.

Intuitively, increasing WalkSAT's noise level would lead to a less greedy and more "erratic" search. It would therefore be reasonable to expect more uniform sampling of solutions. Quite surprisingly however, as shown in Figure 4B, the frequency of the most common and the rarest solution is hardly affected by the noise level. This unexpected finding requires further study. It also means that we need to develop other ways to improve WalkSAT's sampling. We obtained similar results for the logistic problem, which are not shown due to space constraints.

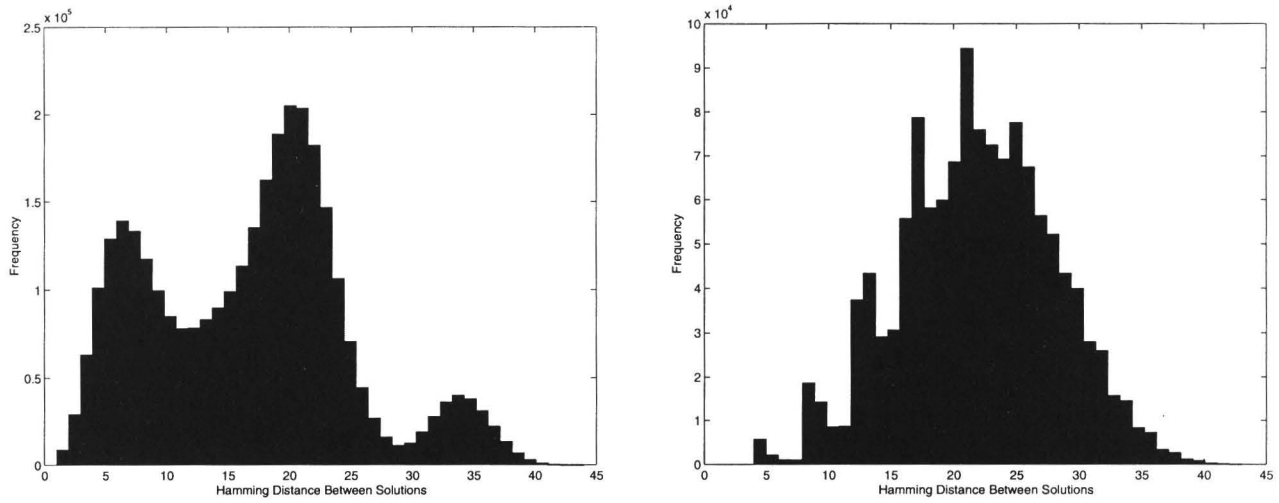

Figure 3: Hamming distances between all solutions for a (A) 3-SAT and (B) logistics problem

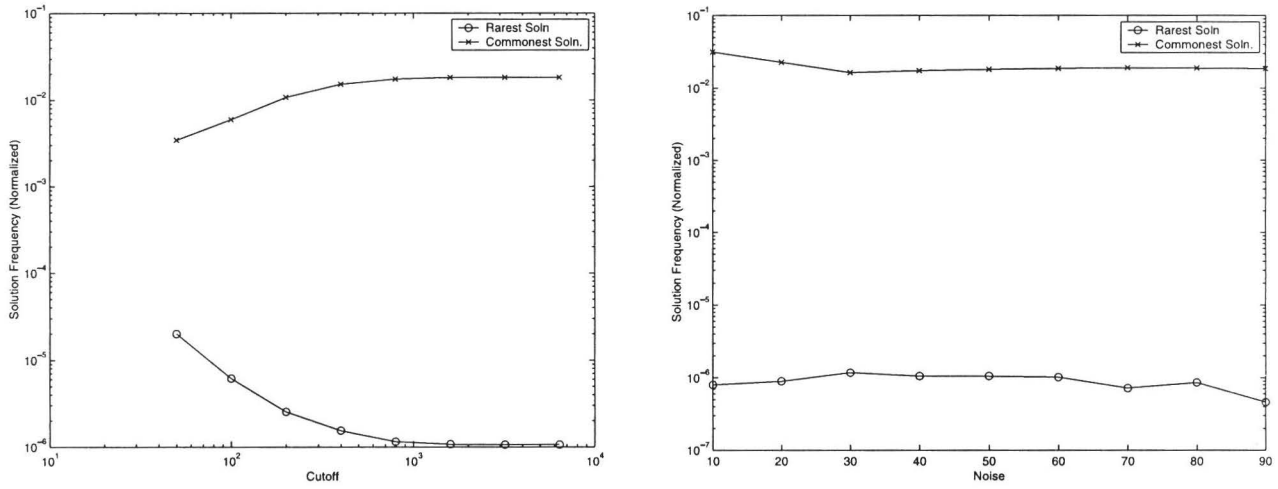

Figure 4: Uniformity of Sampling for a 3-SAT problem: (A) Soln. Frequency vs. Cutoff (B) Soln. Frequency vs. Noise

5 Summary

1. WalkSAT is able to find all solutions for the 3-SAT and constrained logistic problem, within a reasonable number of trials.
2. Solution sampling can be highly non-uniform, particularly for the 3-SAT problem.
3. The 3-SAT instances appear to have a more complicated solution space than the logistic problem.
4. Increasing WalkSAT's noise does not appear to increase the solution sampling uniformity, while it degrades WalkSAT's speed.
5. Challenge: Devise ways to obtain more uniform sampling.

References

[Bayardo and Pehoushek, 2000] R.J. Bayardo and J.D. Pehoushek. Counting models using connected components. In *Proc. AAAI-00.*

[McAllester*et al.*, 1997] D. McAllester, B. Selman, and H. Kautz. Evidence for invariants in local search. In *Proc. of AAAI-97*, pp 321-326, 1997.

[Roth, 1996] D. Roth, On the Hardness of Approximate Reasoning, In *Artificial Intelligence Journal*, Vol. 82, 1-2, pp 273-302, 1996.

[Selman *et al.*, 1994] B. Selman, H. Kautz, and B. Cohen. Noise Strategies for Local Search. In *Proc. AAAI-94*, pp 337-343, 1994.