

Phase transitions in random satisfiability problems

Eli Chertkov

May 8, 2017

Abstract

Boolean satisfiability is a difficult computational problem that is actively studied due to its many practical applications and its deep connections to computational complexity theory. In this essay, we demonstrate how techniques from the study of phase transitions can be used to better understand random instances of boolean satisfiability problems. We discuss the relation of satisfiability problems to Ising models and spin glasses and discuss how finite-size scaling can be applied to study satisfiable-unsatisfiable transitions in random instances of the k -satisfiability problem. We perform numerical calculations using a modern satisfiability problem solving algorithm to demonstrate the finite-size scaling approach in detail. We numerically show that these transitions appear continuous and discontinuous in the order parameter for $k = 2$ and $k \geq 3$ respectively. We also empirically demonstrate that problem instances near the critical region of the phase transition are the most difficult to solve, which is useful for benchmarking satisfiability algorithms.

1 Introduction

Methods from statistical physics have been successfully applied to solve many challenging problems in engineering and applied science. Famous examples include Markov chain Monte Carlo — which was originally developed to study corrections to the ideal gas law for interacting particles, simulated annealing — which is based on the physical process of annealing a material, and belief propagation — which has connections to the cavity method developed to study spin glasses [1, 2, 3, 4]. These techniques have been widely used to solve problems in artificial intelligence, optimization, and error-correcting codes.

The concept of phase transitions from statistical physics has also been useful for understanding fundamental questions in computational science. Complexity, for example, is an important problem in computer science relevant to many areas of mathematics and to practical applications of algorithms in computational science and engineering. Studying the complexity of algorithms allows us to characterize the difficulty of various computational tasks and to create better algorithms to tackle them. Phase transitions provide a unique perspective on the difficulty of solving certain problems that is not captured in the worst-case algorithmic analysis typical in computer science.

Hundreds of computational problems have been put into a category known as NP-complete. Intuitively, for problems in this category, solutions are easy to check but hard to find. By “easy” (“hard”), we mean that in the worst-case an algorithm takes a time polynomial (superpolynomial) in the input size. Problems in NP-complete can be mapped onto one another in polynomial time, making them equivalent to one another in terms of worst-case complexity. By this equivalence, if a problem in NP-complete would turn out to be solvable in polynomial time in the worst-case, then all NP-complete complete problems would be in the polynomial (P) complexity class. This would imply that $P = NP$. It is generally believed that algorithms for solving NP-complete problems must be superpolynomial in the worst-case, or, equivalently, that $P \neq NP$. Nonetheless, while in the worst-case NP-complete problems can be difficult to solve, it can be the case that many practical instances of an NP-complete problem are easy to solve.

The boolean satisfiability problem (SAT) is one of the first problems proven to be NP-complete and has been a useful workhorse for studying the properties of NP-complete problems. The problem consists of determining whether a boolean formula of many binary variables can evaluate to TRUE, or be “satisfied.” As of yet, no algorithm that can solve all instances of the SAT problem in polynomial time exists. Nonetheless, efficient exact and heuristic SAT-solvers have been developed to solve practical problem instances that occur in engineering and industry applications.

SAT can be usefully characterized by techniques from statistical physics. In particular, properties of random SAT problem instances can be described using the language of phase transitions and critical phenomena. Moreover, as we will discuss, random SAT problems are closely related to disordered physical systems known as spin glasses. This makes the study of SAT using statistical physics interesting for two reasons: (1) computer scientists can learn about SAT and other NP-complete problems using concepts from statistical physics and (2) physicists can learn about spin glasses using results from computer science on the satisfiability problem.

In this essay, we will analyze the properties of random instances of the k -SAT problem

using techniques developed to study phase transitions and critical phenomena. First, in section 1.1, we provide a detailed description of the k -SAT problem. In sections 1.2 and 1.3, we describe the relation of k -SAT and other NP-complete problems to the Ising model and spin glass systems. In sections 2.1 and 2.2, we discuss a SAT solving algorithm and the finite-size scaling technique, both of which we will use in our numerical analysis. In section 3.1, we present and discuss our numerical results on the satisfiable-unsatisfiable transition for the k -SAT problem for $k = 2, 3, 4$. Finally, we present our conclusions in section 4.

1.1 The k -satisfiability problem

The k -satisfiability problem (k -SAT) is an NP-complete decision problem that is frequently used as a test bed for the development of new heuristic algorithms [5]. The goal of k -SAT is to determine whether it is possible to assign boolean TRUE and FALSE values to satisfy a boolean formula of a particular form.

A problem instance of k -SAT is specified by a set of N boolean variables x_1, \dots, x_N , whose values can be TRUE or FALSE (which we indicate by 1 or 0), and a set of M clauses each made up of k variables.

An example of a 3-SAT problem instance with $N = 4$ variables and $M = 6$ clauses is

$$\begin{aligned} & (x_1 \vee x_2 \vee \neg x_4) \wedge (x_1 \vee \neg x_3 \vee x_4) \wedge (x_2 \vee x_3 \vee \neg x_4) \wedge \\ & (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_3 \vee \neg x_4) \end{aligned} \quad (1)$$

This problem instance is a boolean formula, where \vee is a logical OR, \wedge is a logical AND, and \neg is a logical NOT. In particular, it is in conjunctive normal form (CNF), so that the entire formula is a logical AND of many clauses and each clause is represented by a logical OR of its boolean variables (or their logical negations). It turns out that this particular 3-SAT problem instance is satisfiable: the logical assignment $(x_1, x_2, x_3, x_4) = (1, 1, 1, 0)$ evaluates the entire boolean expression to 1.

Interestingly, any problem instance in 2-SAT is solvable in at-most polynomial time, i.e., is in the P complexity class. On the other hand, k -SAT for $k \geq 3$ is NP-complete. For example, there exist many 3-SAT problem instances where modern-day SAT solvers run in time exponential in N . Evidently, there is a sharp distinction between the $k = 2$ and $k = 3$ cases putting the problems into entirely different complexity classes.

1.2 Spin glasses and their relation to k -SAT

Spin glasses are physical systems with magnetic disorder. They are named in analogy to typical glasses and amorphous solids, which possess positional disorder. Spin glasses are often described as systems of Ising spins $S_i = \pm 1$ interacting under the Hamiltonian

$$H = - \sum_{ij} J_{ij} S_i S_j - \sum_i H_i S_i \quad (2)$$

where the J_{ij}, H_i are independent random variables drawn from a probability distribution such as a Gaussian distribution. For example, the Edwards-Anderson model and the Sherrington-Kirkpatrick models of spin glasses are of this form [6]. The random J_{ij}, H_i , which

can be positive or negative, induce frustration between spins, similar to the frustration seen in certain magnetic lattices, such as the triangular lattice.

The usual statistical mechanics analyses can be performed on spin glasses by computing partition functions Z and free energies f , but with an additional complication: these functions need to be averaged over different disordered realization of the J_{ij}, H_i . Exactly determining the disorder-averaged quantities \bar{Z} and \bar{f} can be difficult to do analytically. Techniques such as the replica method and mean-field theory can be applied to analytically determine the properties of some simple spin glass models [6].

Instances of the k -SAT problem can be seen as Hamiltonians of the form Eq. (2). First, by the relation $S_i = 2(x_i - 1/2)$, each boolean variable $x_i = 0, 1$ can be mapped onto an Ising spin $S_i = \pm 1$. Next, we can define a clause matrix W such that $W_{ji} = +1$ if clause j includes the boolean variable x_i , $W_{ji} = -1$ if clause j includes the boolean variable $\neg x_i$, and $W_{ji} = 0$ otherwise. Then, we can write down the indicator expression

$$V_j = \frac{1}{2^k} \prod_{i=1}^N (1 - W_{ji} S_i) \quad (3)$$

which is 0 when clause j is satisfied and 1 when it is violated [7]. By summing over all such clause indicators, we can write down a Hamiltonian for k -SAT whose energy corresponds to the number of violated clauses [7]:

$$H_{k\text{-SAT}} = \sum_{j=1}^M V_j = \frac{1}{2^k} \sum_{j=1}^M \prod_{i=1}^N (1 - W_{ji} S_i). \quad (4)$$

The k -SAT decision problem can now be framed as checking if the ground state energy of $H_{k\text{-SAT}}$ is 0, which corresponds to the k -SAT instance being satisfiable.

Note that as written the Hamiltonian $H_{k\text{-SAT}}$ has k -body spin interactions of the form $S_{i_1} S_{i_2} \cdots S_{i_k}$. There exist techniques for introducing auxiliary spins to couple with the original spins that reduce the k -body interactions into 2-body terms as in Eq. (2) [8]. However, this can introduce many auxiliary spins, sometimes many more than the original number of boolean variables N .

Just as the J_{ij}, H_i can be random variables, so too can the clause matrix W_{ji} be generated randomly. A different random instance of W_{ji} corresponds to a different random k -SAT Hamiltonian $H_{k\text{-SAT}}[W_{ji}]$. Therefore, randomly generated instances of k -SAT can be interpreted as random realizations of disorder for a spin glass. This correspondence has allowed physicists to use tools designed for the analysis of spin glasses, such as the replica method, to analyze the critical behavior of k -SAT [6]. We note this correspondence to provide some rationale for considering k -SAT as analogous to a physical system and thus capable of exhibiting phase transitions. However, we do not discuss the spin glass theory predictions made for k -SAT. For more details, refer to [6].

1.3 Other NP-complete problems as Ising models

As detailed in the pedagogical article [8], many famous NP-complete problems can be described in terms of spin degrees of freedoms and spin-glass Hamiltonians such as Eq. (2).

In general, the procedure of mapping involves writing down the decision problem as an optimization problem, with an objective function and constraints. We can associate the objective function with a Hamiltonian H_A , which when equal to 0 solves the given decision problem, and the constraints with a Hamiltonian H_B , which when equal to 0 satisfies all of the constraints of the problem. The entire NP-complete problem can then be encoded in a Hamiltonian $H = H_A + H_B$. Finding the ground state energy and checking if it is zero is equivalent to solving the original decision problem.

2 Methods

We now present the tools that we will use to analyze the satisfiable-unsatisfiable transition seen in random k -SAT problem instances. The relevant variable to tune across the transition is the clause density $\alpha = M/N$. In the thermodynamic limit $N \rightarrow \infty$ with clause density α fixed, a phase transition appears at a critical clause density α_C , which depends on k .

We study the phase transition by examining finite-size systems and employing finite-size scaling techniques on $P_{SAT}(\alpha)$, the probability that a random k -SAT instance with clause density α is satisfiable. We compute $P_{SAT}(\alpha)$ by numerically solving many samples of random k -SAT problem instances at a fixed α with a SAT solving algorithm.

In this section, we first discuss the DPLL algorithm for solving SAT problem instances, which we used to generate the data in Section 3. We then discuss how finite-size scaling can be applied to k -SAT.

2.1 The DPLL algorithm

An important SAT solver is the Davis-Putnam-Logemann-Loveland (DPLL) algorithm developed in the early 1960s. The DPLL algorithm is a non-heuristic backtracking search algorithm [9]. The method consists of performing a binary search, while avoiding large regions of the search space by using two operations known as unit propagation and pure literal evaluation.

The basis idea of how DPLL works is as follows. Consider a brute force binary search. At each step of the algorithm, you consider a boolean variable x_j . You try setting $x_j = 0$, then $x_j = 1$, and do so recursively until you reach a configuration of $\{x_i\}$ that satisfy the boolean formula or until you have exhausted all possible configurations. Instead of the naive binary search, DPLL performs unit propagation and pure literal evaluation at each step to deduce whether $x_j = 0$ or $x_j = 1$ must hold. Empirically, this modification often results in exploring a much smaller number of configurations than the 2^N configurations explored in the brute force search.

Many SAT solvers are extensions or modifications of the underlying DPLL method. We make use of one such state-of-the-art solver, called zChaff [10], to determine $P_{SAT}(\alpha)$, the fraction of random k -SAT instances for clauses of length $M = \alpha N$ that are satisfiable.

2.2 Finite-size scaling

Finite-size physical systems at non-zero temperatures do not have phase transitions. It is only in the thermodynamic limit, when the system size is taken to infinity, that singularities in the free energy density emerge and phase transitions become well-defined. Nevertheless, information about the scaling laws near a phase transition can be obtained by analyzing finite-size systems. The procedure for doing this for continuous phase transitions is known as finite-size scaling.

The main tool for understanding scaling laws near a continuous phase transition is the Renormalization Group (RG) [11]. The RG is an iterative procedure that involves an ‘‘RG transformation’’ from a system to another with a length scale increased by a factor $\ell > 1$. Iteratively applying the RG procedure changes the coupling constants of a Hamiltonian, causing them to flow towards fixed points in coupling constant-space. Certain fixed points, called critical fixed points, correspond to phase transition boundaries.

Consider a finite-size physical system of size L with a single coupling constant $K \propto 1/T$ tuned to be near a critical fixed point. The RG transformation can be applied near this point and linearized to obtain various scaling laws [11]. One such scaling law is for the correlation length

$$\frac{\xi(t, L^{-1})}{L} = F(L/\xi_\infty) = F(Lt^\nu) \quad (5)$$

where F is a scaling function, $t \equiv (T - T_C)/T_C$ is the reduced temperature, and ν is the critical exponent of the correlation length in the infinite-size system ($\xi_\infty(t) = \lim_{L \rightarrow \infty} \xi(t, L^{-1}) \sim t^{-\nu}$).

Since for a finite L there is no singular behavior for ξ , the expression is analytic and can be safely Taylor expanded about $t = 0$ to give

$$\frac{L}{\xi(t, L^{-1})} = A + BtL^{1/\nu} + O(t^2) \quad (6)$$

where A and B are constants [11]. To determine the critical exponent, one can take a derivative of this expression with respect to the coupling constant K to obtain

$$\left. \frac{\partial}{\partial K} \left(\frac{L}{\xi(K, L^{-1})} \right) \right|_{T=T_C} \propto L^{1/\nu}. \quad (7)$$

This logic can be applied to random k -SAT problem instances with little modification. In the case of k -SAT, the analog of system size L is N and the analog of the coupling constant K is the clause density $\alpha = M/N$. Finally, in the k -SAT literature, $L/\xi(t, L^{-1})$ corresponds to the order parameter $P_{SAT}(\bar{\alpha}, N^{-1})$, where $\bar{\alpha} \equiv (\alpha - \alpha_C)/\alpha_C$ [12].

In summary, the critical scaling for the k -SAT phase transition about $\alpha = \alpha_C$ is

$$P_{SAT}(\bar{\alpha}, N^{-1}) = G(N\bar{\alpha}^\nu) \quad (8)$$

where G is a scaling function. The critical exponent ν can be determined by the relation

$$\left. \frac{\partial P_{SAT}(\alpha, N^{-1})}{\partial \alpha} \right|_{\alpha=\alpha_C} \propto N^{1/\nu}. \quad (9)$$

3 Results and Discussion

In our study of the k -SAT transition, we generate random k -SAT problem instances in systems with $N = 10, 20, 40$ and 100 boolean variables. Each problem instance is solved by the zChaff SAT solver [10]. This allows us to estimate the disorder-averaged quantity $P_{SAT}(\alpha)$, which we use for finite-size scaling.

3.1 Phase transitions in k -SAT

The critical behavior of k -SAT was numerically described by Scott Kirkpatrick and Bart Selman in 1994 [13]. We compare our results for the k -SAT phase transitions for $k = 2, 3, 4$ with their results.

For $k = 2$, it has been analytically shown [6] that a satisfiable-unsatisfiable phase transition occurs at the critical clause density $\alpha_C(2) = 1$ with an estimated critical exponent of $\nu(2) = 2.6$ [13]. For $k = 3, 4$, numerical estimates of the critical points and exponents are $\alpha_C(3) = 4.17, \alpha_C(4) = 9.75$ and $\nu(3) = 1.5, \nu(4) = 1.25$ as described in [13].

Figure 1 shows our calculations of $P_{SAT}(\alpha)$ for $k = 2, 3, 4$ in finite-size systems. For large N , finite-size effects are diminished and the limiting behavior of $P_{SAT}(\alpha)$ in the infinite-size limit becomes evident. For $k = 2$, we can see that $P_{SAT}(\alpha)$ continuously changes from 0 to 1 as α is changed from $\alpha_C = 1$ to 0. For $k \geq 3$ on the other hand, when $N \rightarrow \infty$ the transition becomes discontinuous such that $P_{SAT}(\alpha_C^-) = 1$ and $P_{SAT}(\alpha_C^+) = 0$.

Our finite-size scaling results for k -SAT are shown in Figure 2. For each k , the $P_{SAT}(\alpha)$ curves rescaled according to Eq. (8) collapse well onto a universal scaling function, even for α far away from the critical clause density α_C . Moreover, Figure 3 demonstrates how the scaling relation from Eq. (9) can be used to compute the critical exponent ν . Our result for $k = 3$ produces an estimate $\nu = 1.64 \pm 0.11$ for the critical exponent that agrees well with the established result of $\nu = 1.5 \pm 0.1$ from [13].

In [5], the authors argue that the behavior of the order parameter P_{SAT} for 2-SAT and 3-SAT (as displayed in Figure 1) demonstrates that the 2-SAT transition can be thought of as a continuous ‘second-order’ transition and that the 3-SAT transition as a discontinuous ‘first-order’ transition. However, we disagree with this interpretation due to the fact that finite-size scaling, which only applies for continuous phase transitions, works well for all k .

Finally, it is interesting to consider the average running time of the DPLL algorithm, as shown in Figure 4. Evidently, the “hardest” k -SAT problem instances, which take the longest for the DPLL method to solve, occur near the critical region. However, for most clause densities, problem instances are solved quickly. Interestingly, this empirical easy-hard-easy transition indicates that, despite the fact that k -SAT is NP-complete, most k -SAT problems can actually be solved efficiently.

4 Conclusions

By performing numerical analyses on random instances of the k -SAT problem for $k = 2, 3, 4$, we demonstrated how techniques for describing phase transitions could be used to characterize important properties of k -SAT. We applied finite-size scaling to the system and

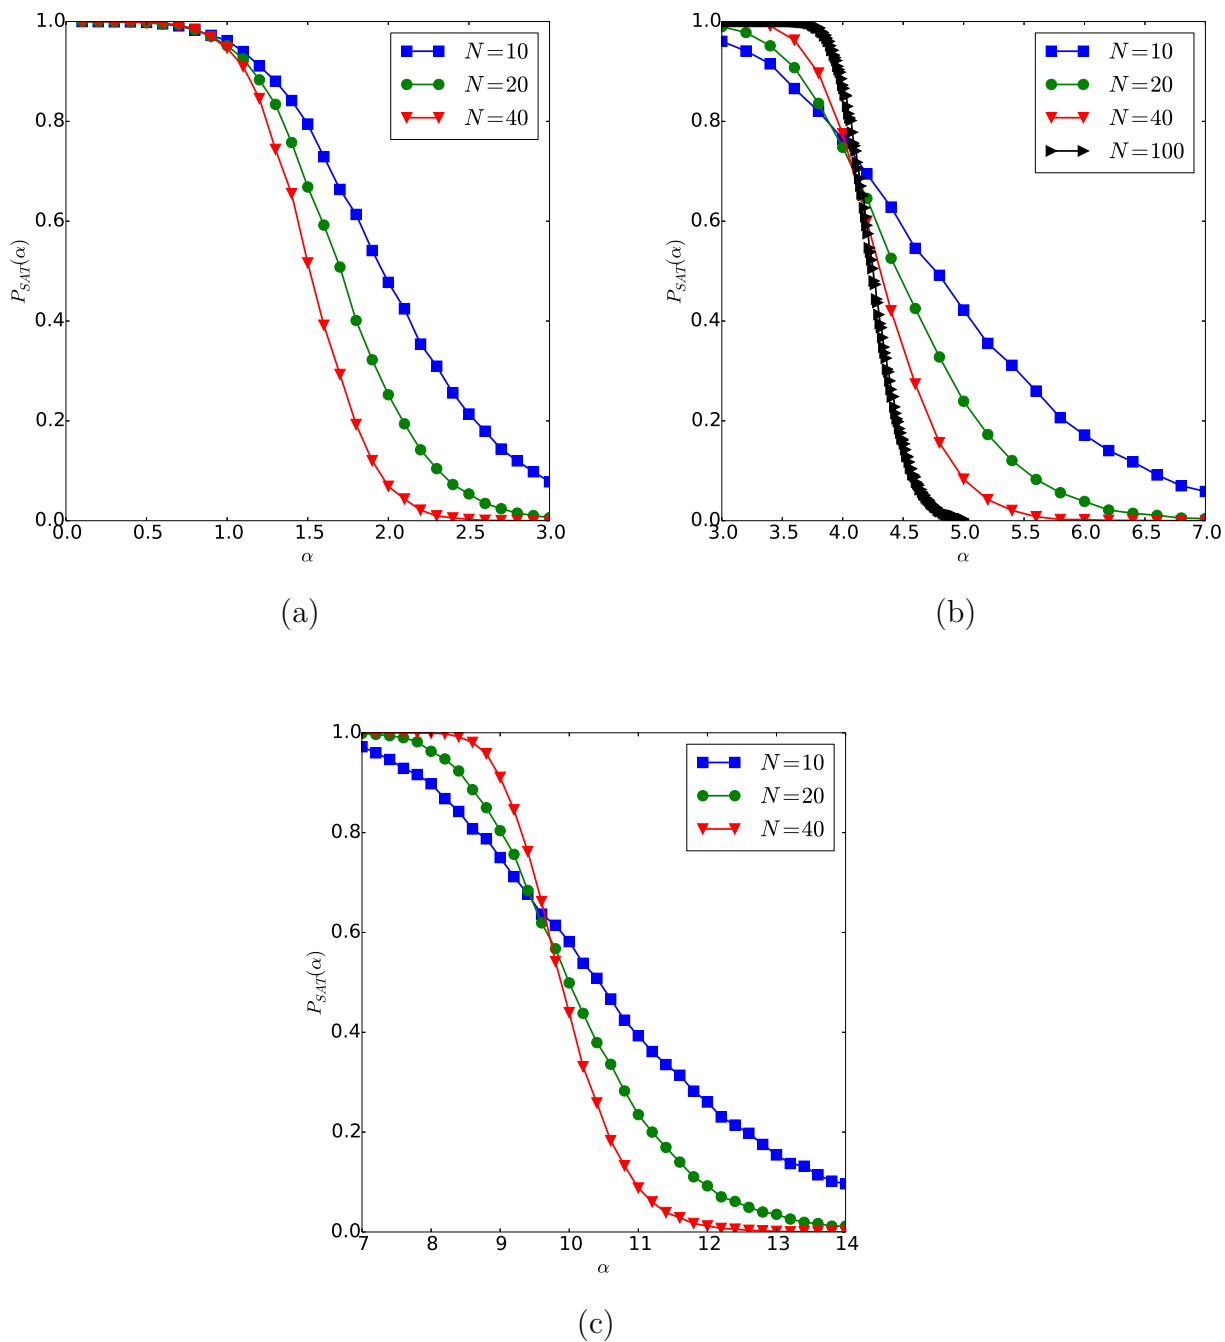


Figure 1: Our numerical estimates of $P_{SAT}(\alpha)$, the probability that a random k -SAT problem instance with clause density α is satisfiable, for (a) $k = 2$, (b) $k = 3$, and (c) $k = 4$. Each estimate was made by solving 10,000 randomly generated k -SAT instances of $N = 10, 20, 40$ boolean variables (and $N = 100$ for 3-SAT) using the zChaff algorithm [10].

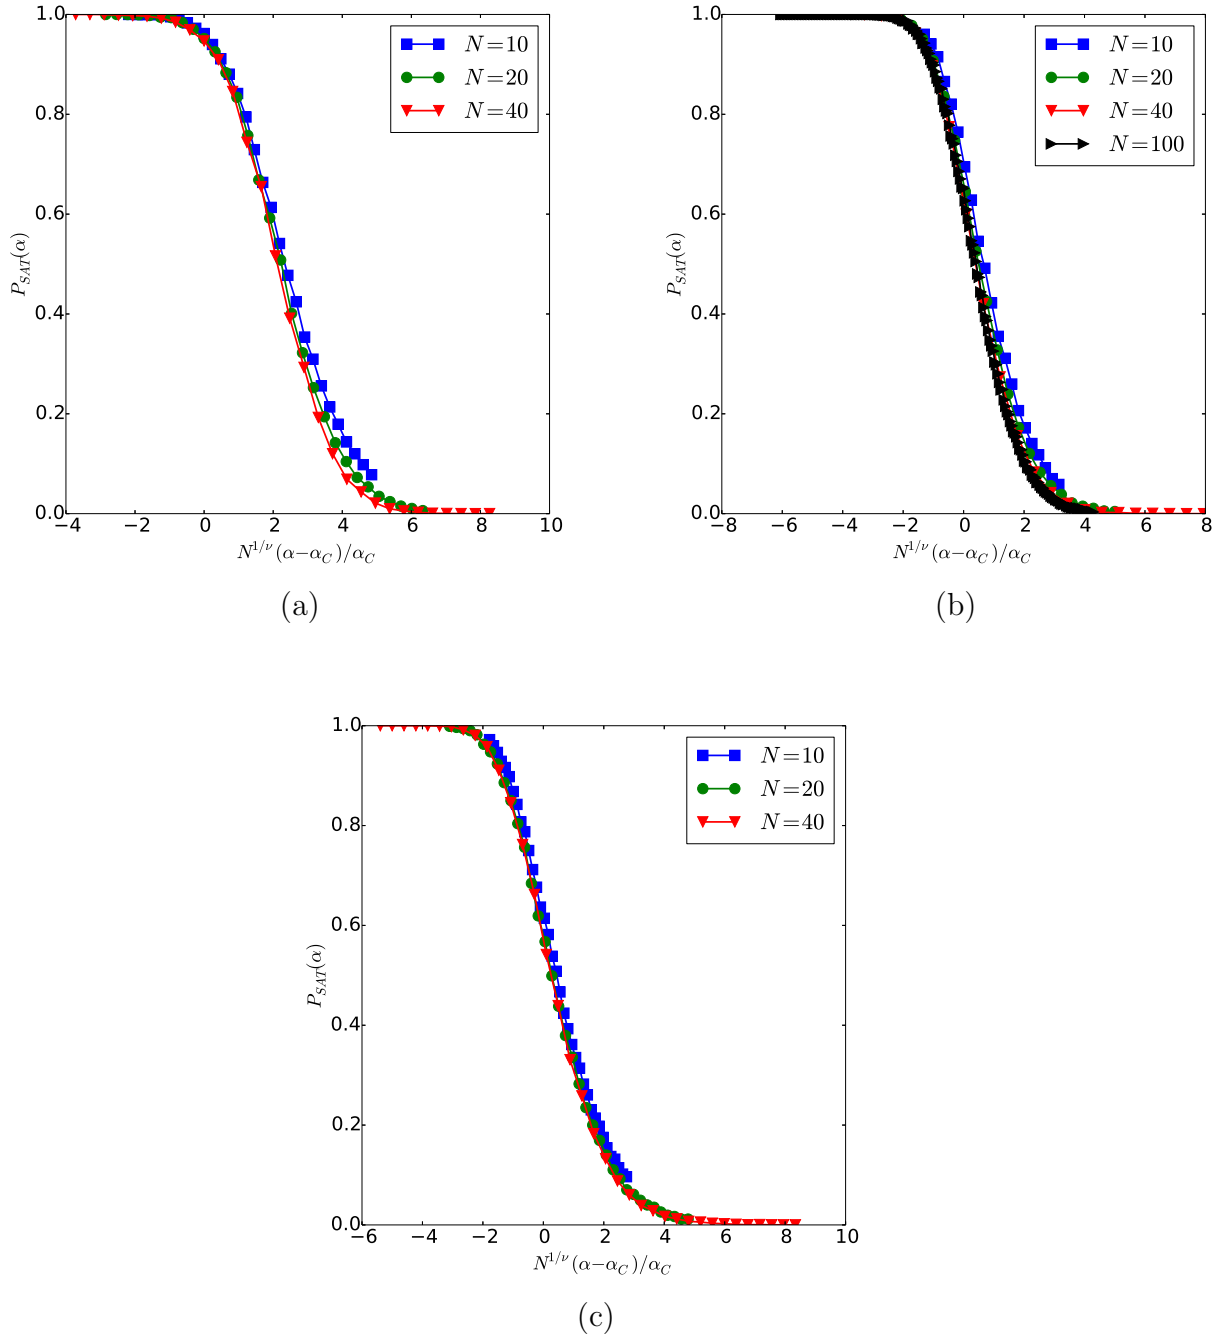


Figure 2: Finite-size scaling of k -SAT for (a) $k = 2$, (b) $k = 3$, and (c) $k = 4$. For each k , the $P_{SAT}(N^{1/\nu}(\alpha - \alpha_C)/\alpha_C)$ curves for different N collapse onto a universal scaling function for α near α_C . Each point is generated from 10,000 random k -SAT problem instances with $N = 10, 20, 40$ variables (and $N = 100$ for 3-SAT). The critical clause densities and critical exponents used in the scaling were $\alpha_C = 1, 4.17, 9.75$ and $\nu = 2.6, 1.5, 1.25$ for $k = 2, 3, 4$ respectively as obtained in [13].




Figure 3: Log-log plot of the α -derivative of $P_{SAT}(\alpha, N^{-1})$ at the critical point $\alpha = \alpha_C = 4.17$ for different system sizes N for 3-SAT. According to Eq. (9), this curve should be a straight line with slope $1/\nu$. By performing a least-squares linear fit, we estimate the critical exponent $\nu = 1.64 \pm 0.11$, which matches with the value $\nu = 1.5 \pm 0.1$ measured in the literature [13].




Figure 4: The average running time of the zChaff algorithm (a modified DPLL method [10]) in milliseconds for 3-SAT with $N = 100$. Each point was estimated by measuring the running time of zChaff on 10,000 randomly generated 3-SAT problem instances with $M = \alpha N$ clauses. The red dashed line indicates the position of the phase transition in the infinite-size limit, which is at $\alpha = \alpha_C = 4.17$ for 3-SAT [13].

demonstrated that it could accurately describe the behavior near the satisfiable-unsatisfiable transition. In this analysis, we showed data collapse and were able to compute the critical exponent for 3-SAT, which agreed well with results in the literature. We also saw that for 2-SAT, which is in the P complexity class, the phase transition appeared continuous in the order parameter P_{SAT} and that for k -SAT with $k \geq 3$, which is NP-complete, the transition displayed a discontinuity in P_{SAT} . Moreover, we noted that despite the apparent discontinuities in P_{SAT} for $k \geq 3$, the phase transitions for $k \geq 2$ are still likely continuous transitions since finite-size scaling holds very accurately. Finally, empirically, we showed that the hardest 3-SAT problem instances occurred in the critical region around the transition. This suggests that problem instances in the critical region of k -SAT can be used as useful benchmarks for the performance of SAT solving algorithms.

References

- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087, 1953.
- [2] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671, 1983.
- [3] J. Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In *Proceedings of the Second AAAI Conference on Artificial Intelligence*, AAAI'82, page 133. AAAI Press, 1982.
- [4] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Exploring artificial intelligence in the new millennium. chapter Understanding Belief Propagation and Its Generalizations, page 239. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.
- [5] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400:133, 1999.
- [6] M. Mezard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, Inc., New York, NY, USA, 2009.
- [7] A. G. Percus, G. Istrate, and C. Moore. *Computational Complexity and Statistical Physics*. Oxford University Press, Inc., Oxford, 2006.
- [8] A. Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2:5, 2014.
- [9] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Commun. ACM*, 5:394, 1962.
- [10] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient sat solver. *39th Design Automation Conference (DAC 2001)*, 2001.
- [11] N. Goldenfeld. *Lectures on phase transitions and the renormalization group*. Addison-Wesley, Advanced Book Program, Reading, 1992.

- [12] M. Leone, F. Ricci-Tersenghi, and R. Zecchina. Phase coexistence and finite-size scaling in random combinatorial problems. *Journal of Physics A: Mathematical and General*, 34:4615, 2001.
- [13] S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random boolean expressions. *Science*, 264:1297, 1994.