

# Understanding the Empirical Hardness of Random Optimisation Problems

Anonymous

An Undisclosed Location  
anonymous@anonymous

## Abstract

We look at the empirical complexity of the maximum clique problem, the graph colouring problem, and the maximum satisfiability problem, in randomly generated instances. Although each is NP-hard, we encounter exponential behaviour only with certain choices of instance generation parameters. To explain this, we link the difficulty of optimisation to the difficulty of a small number of decision problems, which are already better-understood through phenomena like phase transitions with associated complexity peaks. However, our results show that individual decision problems can interact in very different ways, leading to different behaviour for each optimisation problem. Finally, we uncover a conflict between any-time and overall behaviour in algorithm design, and discuss the implications for the design of experiments and of search strategies such as variable- and value-ordering heuristics.

## 1 Introduction

The gap between the best theoretical understanding we have of what makes problems hard and the behaviour witnessed in practice from modern solvers remains vast. For many decision problems in random instances, we have a good general understanding of what happens: as a key parameter is altered, there is often sharp phase transition from satisfiable to unsatisfiable instances, and associated with this is a complexity peak, where instances near the transition are much harder to solve than those far from it on either side [Cheeseman *et al.*, 1991; Mitchell *et al.*, 1992]. (However, this behaviour is not universal—for example, problems involving more than one kind of constraint can exhibit much more complicated behaviour [Dudek *et al.*, 2017; McCreesh *et al.*, 2018]).

This paper looks at optimisation problems. An optimisation problem can be viewed as a sequence of decision problems—but is that all that is needed to understand their behaviour? Previous small-scale experiments (e.g. Prosser [2012], Mann [2018]) have only been able to provide an incomplete picture. In this paper we perform experiments on




Figure 1: On top, the difficulty of solving the maximum clique problem in random graphs  $G(50, x)$ ,  $G(100, x)$ ,  $G(150, x)$ , and  $G(200, x)$ . Underneath, the mean size of an optimal solution. Density is increased in steps of 0.001 with 100,000 samples per step for the three smaller families, and 1,000 per step for the largest.




Figure 2: Does  $G(150, x)$  contains a clique of 20 vertices? Mean search effort for satisfiable and unsatisfiable instances are also shown separately. Density is increased in steps of 0.001, with 100,000 samples per step.

tens of billions of problem instances, which is finally sufficient to comprehensively answer the question: yes, there is a link between individual decision and optimisation problems, but these decision problems can interact in many different ways, leading to complex emergent behaviour. Along the way, we uncover interesting implications for the design of search algorithms, and provide lessons for future experimenters. Most interestingly, we identify a trade-off between anytime behaviour and overall behaviour, which may encourage a rethink of the entire branch and bound paradigm.

## 1.1 Experimental Setup

Our experiments are performed on an Anonymous HPC facility, on systems with dual Intel Xeon E5-2695 v4 CPUs and 256GB of RAM, running Centos 7.3.1611, with GCC 7.2.0 as the compiler. These machines are optimised for providing throughput rather than consistent timing measurements, so we avoid measuring runtimes, and instead use whichever natural measure of work each solver provides. Our results therefore do not allow for comparisons between different solvers.

In Section 2, we use the Glasgow Subgraph Solver implementation of Prosser’s [2012] MCSa1<sup>1</sup>. This is a bit-parallel branch and bound algorithm, which uses a greedy colouring as its bound [Tomita *et al.*, 2010; Segundo *et al.*, 2013]; it can easily be modified to solve the decision problem, rather than the optimisation problem. We measure instance difficulty by counting the number of recursive calls carried out by the algorithm. In Section 3 we use Trick’s implementation<sup>2</sup> of the classic DSATUR branch and bound algorithm [Br  laz, 1979], and Zhou *et al.*’s [2014] Color6 solver<sup>3</sup> (which solves only the decision problem). For the Trick solver we measure the number of recursive calls made, whilst for Color6 we measure the number of backtracks. In Section 4 we use the Clasp solver<sup>4</sup> version 3.3.4 [Gebser *et al.*, 2012], and we measure the number of decisions made.

## 2 Maximum Clique

We begin by looking at the maximum clique problem. A clique in a graph is a subset of vertices, each of which is adjacent to every other within the subset, and a maximum clique is one with as many vertices as possible. For random graphs, we use the Erd  s-R  nyi model: by  $G(n, p)$  we mean a graph with  $n$  vertices, and an edge between every distinct pair of vertices with probability  $p$ . Clique-finding in Erd  s-R  nyi graphs is known to be exponentially difficult for any resolution-type solver [Atserias *et al.*, 2018].

**Maximum cliques in random graphs.** In Figure 1 we show the difficulty of solving the maximum clique problem as we vary the edge probability in Erd  s-R  nyi graphs with a fixed number of vertices, as well the mean size of an optimal solution. In extremely sparse and extremely dense graphs, the algorithm finds all instances extremely easy, whilst at around densities of 0.8 to 0.96, instances are particularly hard—and

unsurprisingly, as the number of vertices increases, all densities get exponentially harder.

These rough trends match up with those presented by Prosser [2012]. However, we are using a *much* larger number of instances: we increase density in steps of 0.001, and take 100,000 samples per density step. This scale of experiments reveals a new interesting feature of the plots: the lines are, for lack of a better term, wiggly. This is most readily apparent towards the left of the graph, where several slight peaks and troughs are easily visible by eye, but in fact the wiggles are present throughout the entire plot, with a decreasing “wavelength” as density increases. The remainder of this section shows that these wiggles are not an experimental artifact or sampling error, but instead illustrate an important aspect of the algorithm’s behaviour.

**The clique decision problem.** To understand what is going on, we first revert to the clique decision problem. In Figure 2 we ask whether  $G(150, x)$  contains a clique of twenty vertices. For very sparse graphs, the answer is obviously no, and the solver can establish this with no search effort. For very dense graphs, the answer is obviously yes, and the solver similarly finds all instances easy. For densities in between 0.691 and 0.782, there is a mix of satisfiable and unsatisfiable instances, but these instances are hard for the solver. For unsatisfiable instances, the higher the density the harder the instance, and the hardest density is 0.780, where all but one of the 100,000 instances sampled are satisfiable. Unexpectedly, for satisfiable instances, we do not get a hard—easy curve, but rather a medium—hard—easy peak, with the hardest density being 0.756 where 62,587 instances were satisfiable. Instances in the “medium” region are extremely rare, however.

**Decision and optimisation.** In Figure 3 we simultaneously plot the difficulty of every decision problem, and show how this correlates with the total search effort seen in Figure 1. The “total” line is usually only slightly above whichever decision line is the hardest at a particular density—even at the hardest density of 0.905, the mean gap between the hardest decision problem and the overall cost of solving is only a factor of 2.2. This explains the wiggly lines: they are the result of the gaps between the complexity peaks of different decision problems.

**Time to find versus prove optimality.** Figure 3 also breaks down the runtimes to show the mean time to find an optimal solution but not prove its optimality, and the time to prove optimality once an optimal solution has already been found. These two lines are perfectly out of phase with each other: densities where finding a solution is relatively easy are the hardest for proving optimality, and vice-versa.

**Difficulty by actual solution size.** Another way of grouping results is presented in Figure 4. Alongside a plot of mean search effort, we also show mean search effort only considering instances where the maximum clique has  $\omega$  vertices, for each value of  $\omega$ . The darkness of each line indicates the relative sample size. The plot shows that at any given density, there are several common solution sizes, and the difficulty varies considerably depending upon what the optimal solution size actually is. It also shows that, for any particular maximum clique size  $\omega$ , there are unusually low densities

<sup>1</sup><https://github.com/ciaranm/glasgow-subgraph-solver>

<sup>2</sup><https://mat.gsia.cmu.edu/COLOR/color.html>

<sup>3</sup><https://home.mis.u-picardie.fr/~cli/EnglishPage.html>

<sup>4</sup><https://potassco.org/clasp/>




Figure 3: A more detailed picture of difficulty of solving the clique optimisation problem for  $G(150, x)$ . Also plotted is the mean search effort to find the optimal solution but not prove its optimality, and the mean search effort needed to prove optimality after the optimal solution is found. Finally, each light line shows the mean search effort for a single decision problem. For each line, density is increased in steps of 0.001, with 100,000 samples per step.

where occasionally this is the optimum, and these instances are very easy. There are also rare unusually high densities where this is the optimum, and these instances are very hard. Finally, for densities in the middle, instances with solution size  $\omega$  are common, and are of moderate difficulty. Alternatively, for a given instance, if the maximum clique size is unexpectedly large, the instance will be relatively easy, whilst if it is unexpectedly small, it will be unusually hard.

**How common are optimal solutions?** Recall that typically, proving optimality is many times harder than finding an optimal solution. If an instance has an unusually large optimal solution, this should make the proof of optimality much easier. But what about *finding* this unusually large optimum? We might expect that there will only be one optimal solution, if the optimum is unusually large, whilst if the optimum is unusually small, perhaps there are many witnesses to choose from?

In the top plot of Figure 5 we show that this is the case, looking only at instances where twenty is the optimal solution. We plot the frequency of optimal solutions (how common they are, by density), as well as the effort required to find a first optimal solution but not prove its optimality, and the effort to both find and prove optimality. Finally, we also solve the *maximum* clique enumeration problem, and count how many such optimal solutions exist.

Towards the left of this plot, with densities up to 0.72, instances with a maximum clique size of twenty are rare. Furthermore, the total number of optimal solutions (witnesses) in any given instance is very low, often being one or only a few—and nearly all of our search effort is spent finding the optimal, with optimality proofs being easy. As the density rises, the typical number of optimal solutions per instance also rises, and the time to find but not prove optimality makes up smaller and smaller portions of the overall runtime.

Interestingly, the instances with the higher solution counts are *not* the instances where the time taken to find a solution is lowest. Rather, finding the unique optimal solution in a lower density graph is somewhat easier than finding any one of several optimal solutions in a medium density graph, and it is not until much higher densities that finding becomes easier again. This is similar to the “medium–hard–easy” complexity peak seen in Figure 2. One might think that this is because higher densities are harder overall than lower densities. However, the second plot of Figure 5 looks at instances where sixty is the optimal solution, with densities between 0.94 and 0.975. At this stage, higher densities are easier overall—but the same pattern occurs.

**Anytime behaviour.** To explain this behaviour, we now demonstrate that the algorithm is in fact not optimised for anytime behaviour, but rather aims to make the proof of optimality as short as possible. McCreesh and Prosser [2014] observe that the branching strategy used by this algorithm approximates “smallest domain first” [Haralick and Elliott, 1980], and that (contrary to the claims of the algorithm’s designers) it is not good at finding a strong incumbent quickly. So what if we reverse the branching strategy used by the algorithm? Figure 6 compares the behaviour of the heuristic and the anti-heuristic, showing that the anti-heuristic performs much worse except on the easiest of instances. However, in Figure 7, we plot the mean size of the first solution found by both heuristics, as a proportion of the optimal: despite being much worse overall, the anti-heuristic finds a better first solution in nearly all cases.

To understand this seemingly contradictory behaviour, we compare the size of the incumbent as a function of time for the two algorithms. In Figure 8 we select all the instances of  $G(150, x)$  where the optimal solution had twenty vertices, and for both heuristics, record a timepoint for each time the




Figure 4: The mean difficulty of solving the clique optimisation problem for  $G(150, x)$ , also showing the search effort for each actual optimal size. On each of the individual value lines, darker colours represent exponentially larger sample sizes. Density is increased in steps of 0.001, with 100,000 samples per step.

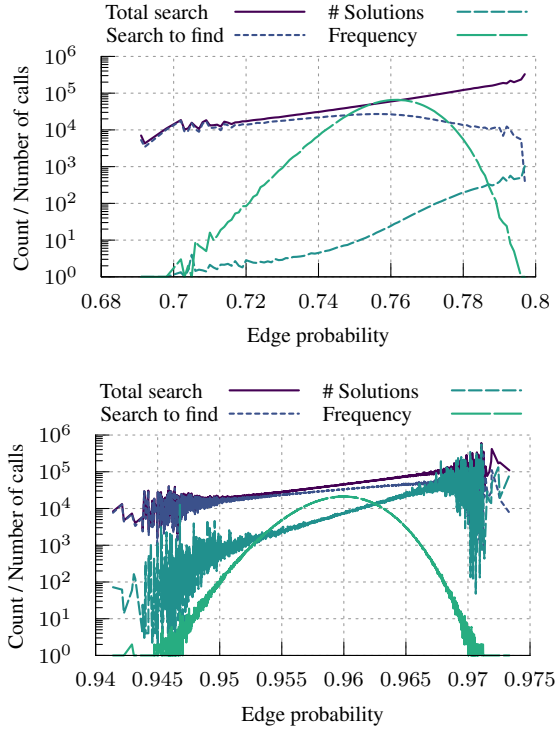


Figure 5: Above, looking at only instances where the maximum clique has twenty vertices in  $G(150, x)$ , and showing the mean search effort, mean time to find but not prove optimality, the frequency of such instances, and the mean number of times a clique of that size occurs in any selected instance. Below, the same, for a maximum clique of sixty vertices. Density is increased in steps of 0.001 (twenty) or 0.00001 (sixty), with 100,000 samples per step.




Figure 6: The mean difficulty of solving the clique optimisation problem for  $G(150, x)$ , using both the standard search heuristic order for the algorithm, and the opposite search order. For each line, density is increased in steps of 0.001, with 100,000 samples per step.




Figure 7: The mean size of the first solution found expressed as a proportion of the optimal, for  $G(100, x)$ , using both the standard search heuristic order for the algorithm, and the opposite search order. For each line, density is increased in steps of 0.001, with 100,000 samples per step.




Figure 8: Comparing the solution quality over time for the two different search orders. We look at instances of  $G(100, x)$  where the maximum clique has twenty vertices. Each time a new incumbent is found, we record the search effort so far; each grid point shows the relative frequency of new incumbents of that size during that time window, with darker colours being exponentially more common. We also show termination time, in the final column. Instances drawn from a run with density increased in steps of 0.001, with 100,000 samples per step.

incumbent is improved. We also record when the algorithm terminates, representing this as an incumbent of twenty-one. We then convert this to a heatmap by bucketing, using darker colours to represent exponentially larger buckets. The plot shows us that with the good heuristic, the initial solution size is lower (most commonly fourteen or fifteen) compared to the anti-heuristic (most commonly sixteen to eighteen), but that the anti-heuristic then becomes slower to advance, and slower still to finally prove optimality. This suggests that the anti-heuristic’s branching choices cause it to become trapped in larger subproblems before it can advance to a better region of the search space. In contrast, the good heuristic tries to eliminate as many subproblems as possible, even at the expense of much less favourable anytime behaviour.

This observation also explains Figure 5: the algorithm does not spend nearly all of its time attempting to find an unusually large optimal solution in a sparser instance because this solution is rare, but rather because it is instead spending all of its time eliminating the remaining portions of the search space. As density increases, the remaining portion of the search space increases, explaining the increase in difficulty despite the higher solution counts.

**Implications for future algorithms.** These results show that the algorithm has been optimised for proofs of optimality, at the expense of worse anytime behaviour—and also that, if the algorithm were better at finding strong solutions quickly, then its performance would improve considerably on certain instances. It is therefore worth considering whether it is possible to modify the algorithm for both good anytime behaviour, and good overall performance. However, adapting its search order heuristics does not appear to help: although doing so can help it find stronger solutions faster, it then quickly becomes stuck in a subproblem that is hard to eliminate.

Other alternatives may be possible. For example, Maslov *et al.* [2014] apply an iterated local search (ILS) heuristic to generate an initial solution, rather than starting from zero. This technique was also adopted by Tomita *et al.* [2016], who use a different form of local search to prime the incumbent. Both papers describe this as assistance, rather than recognising that their exact algorithms are not optimised for finding strong solutions quickly; both papers also have difficulties selecting a principled amount of time to spend running local search before starting the exact algorithm. Both papers also claim large successes (sometimes being thousands or millions of times faster), particularly on certain families from the standard DIMACS benchmark suite. However, a close inspection of the instances where this happens shows that all come from crafted families that are designed to have unusually large hidden optimal solutions [Brockington and Culberson, 1993; Sanchis, 1992; Sanchis, 1995], rather than from application instances.

### 3 Graph Colouring

Having looked in detail at the maximum clique problem, we now repeat some of our experiments using solvers for the graph colouring problem: we must give a colour to each vertex in a graph, giving adjacent vertices different colours, and




Figure 9: The five-colouring phase transition in  $G(60, x)$ , using the color6 solver. Density is increased in steps of 0.001, with 100,000 samples per step.




Figure 10: The difficulty of each colouring decision problem in  $G(60, x)$ , using Color6. Outliers removed. For each line, density is increased in steps of 0.001, with 100,000 samples per step. The top line is the minimisation problem, using Trick’s DSATUR.




Figure 11: The difficulty of the minimum colouring problem in  $G(60, x)$ , using the Trick solver, also showing the means only for instances with each individual optimal solution. Density is increased in steps of 0.001, with 100,000 samples per step.



using as few colours as possible.

**A phase transition, and outliers.** In Figure 9 we show the difficulty of five-colouring  $G(60, x)$ , for varying values of  $x$ , using the Color6 solver. For densities in between 0.16 and 0.23, we encounter a mix of satisfiable and unsatisfiable instances, and the solver finds the instances more difficult than those outside of this density range. However, at much lower densities, comfortably inside the “satisfiable” region, the mean search effort is extremely variable, and is sometimes far higher than at the complexity peak. Looking more closely at the data shows that for densities between 0.11 and 0.15, between one in ten thousand and one in a hundred thousand instances that we generate are tens of millions of times harder than typical (and this rarity explains why Mann’s [2018] experiments did not uncover them). Furthermore, rather than being entirely satisfiable, these instances are a mix of satisfiable and unsatisfiable. Such instances also occur for other values of the decision problem, although it appears to be even less common as the objective value increases. A similar phenomenon occurs with other random satisfaction problems [Smith and Grant, 1997; Achlioptas *et al.*, 2001], and it could potentially be alleviated by the use of restarts and randomisation [Gomes *et al.*, 1998].

**Branch and bound.** In Figure 10 we plot the difficulty of each decision problem together, but exclude these outliers from calculating the means. As for the clique problem, we observe wiggles, with the problem getting easier then harder then easier then harder and so on as we pass successive complexity peaks. The Color6 solver only supports the decision problem. Thus, we also plot the classic DSATUR branch and bound algorithm (whose performance is somewhat worse overall). As with the maximum clique algorithm, the mean complexity line goes from easy to hard to easy over the full range of densities, but this peak has wiggles that line up with the objective values changing.

We also break down the behaviour of the DSATUR solver by optimal solution size: we show this in Figure 11, in the same style as Figure 4. Because we are dealing with a minimisation problem, instances that are relatively sparse for their solution size are now found to be harder, rather than easier. And, as with Color6, DSATUR also occasionally finds very sparse instances very hard.

## 4 Maximum Satisfiability

We finish with a brief look at random maximum satisfiability, or MaxSAT. To generate random MaxSAT instances, we use 40 variables, and a varying number of clauses. Each clause contains three distinct variables chosen uniformly at random, and the polarity of each variable in each clause is also set uniformly at random; all clauses are soft with equal weight. We plot our results in Figure 12, showing both Clasp’s default branch and bound mode, and core-guided optimisation [Andres *et al.*, 2012], which performs better. Although harder to see, the mean search effort lines for both configuration do exhibit wiggles. Interestingly, the relative difficulty of different instances depends upon the search strategy used—we believe this warrants further experimentation.

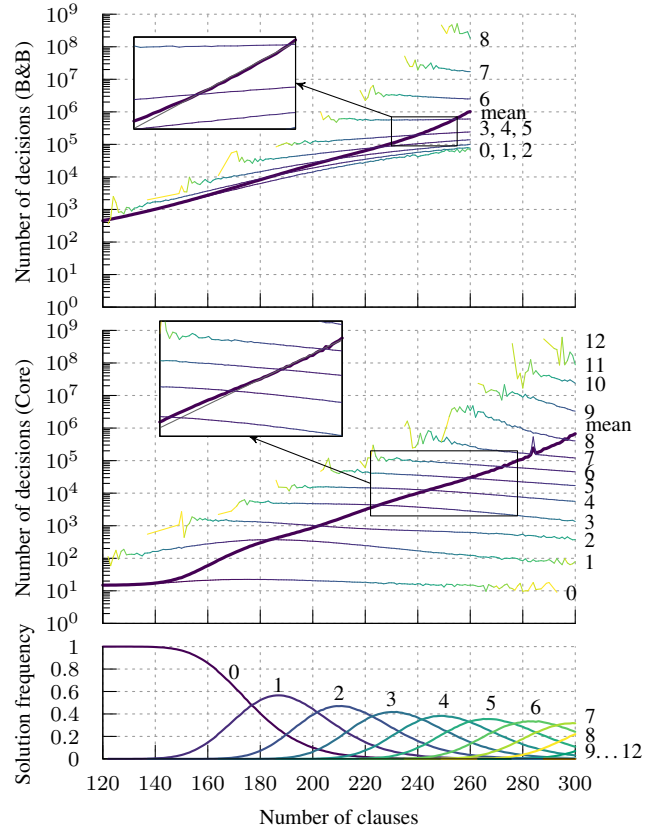


Figure 12: The difficulty of the maximum 3-satisfiability problem in random instances with 40 variables using Clasp in branch and bound mode (top) and in core-guided mode (middle). The number of clauses is increased in steps of one, and there are 100,000 samples per step; the core-guided plot omits six instances with 290 or more clauses that timed out after one day. Results only for each particular objective value (i.e. the number of unsatisfiable clauses) are also shown as smaller lines in both plots. On the bottom, we show how common each objective value is.

## 5 Discussion and Conclusion

By using very large sample sizes, we have demonstrated that the behaviour of solvers on hard optimisation problems is indeed influenced by the behaviour on individual decision problems—but that these decision problems can interact in many different ways. We also uncovered several interesting phenomena that happened only for one in every ten thousand instances (or even fewer). We therefore encourage future experiments to use similarly large sample sizes if possible, and to consider running many relatively easy experiments instead of a small number of experiments on instances that are as large as possible.

We also uncovered a conflict between designing search order heuristics for anytime behaviour or for overall performance in branch and bound algorithms, which explains why recent exact clique algorithms are using priming with local search algorithms, and which has implications for the design of future solvers. This conflict should also be recognised by experimenters when comparing algorithms in the future.

## References

- [Achlioptas *et al.*, 2001] Dimitris Achlioptas, Paul Beame, and Michael S. O. Molloy. A sharp threshold in proof complexity. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 337–346, 2001.
- [Andres *et al.*, 2012] Benjamin Andres, Benjamin Kaufmann, Oliver Matheis, and Torsten Schaub. Unsatisfiability-based optimization in clasp. In *Technical Communications of the 28th International Conference on Logic Programming, ICLP 2012, September 4-8, 2012, Budapest, Hungary*, pages 211–221, 2012.
- [Atserias *et al.*, 2018] Albert Atserias, Ilario Bonacina, Susanna F. de Rezende, Massimo Lauria, Jakob Nordström, and Alexander A. Razborov. Clique is hard on average for regular resolution. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 866–877, 2018.
- [Brélaz, 1979] Daniel Brélaz. New methods to color vertices of a graph. *Commun. ACM*, 22(4):251–256, 1979.
- [Brockington and Culberson, 1993] Mark Brockington and Joseph C. Culberson. Camouflaging independent sets in quasi-random graphs. In *Cliques, Coloring, and Satisfiability, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, October 11-13, 1993*, pages 75–88, 1993.
- [Cheeseman *et al.*, 1991] Peter C. Cheeseman, Bob Kanefsky, and William M. Taylor. Where the really hard problems are. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence, Sydney, Australia, August 24-30, 1991*, pages 331–340, 1991.
- [Dudek *et al.*, 2017] Jeffrey M. Dudek, Kuldeep S. Meel, and Moshe Y. Vardi. The hard problems are almost everywhere for random CNF-XOR formulas. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 600–606, 2017.
- [Gebser *et al.*, 2012] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Conflict-driven answer set solving: From theory to practice. *Artif. Intell.*, 187:52–89, 2012.
- [Gomes *et al.*, 1998] Carla P. Gomes, Bart Selman, and Henry A. Kautz. Boosting combinatorial search through randomization. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, Madison, WI., USA., pages 431–437, 1998*.
- [Haralick and Elliott, 1980] Robert M. Haralick and Gordon L. Elliott. Increasing tree search efficiency for constraint satisfaction problems. *Artif. Intell.*, 14(3):263–313, 1980.
- [Mann, 2018] Zoltán Ádám Mann. Complexity of coloring random graphs: An experimental study of the hardest region. *ACM J. of Experimental Algorithmics*, 23, 2018.
- [Maslov *et al.*, 2014] Evgeny Maslov, Mikhail Batsyn, and Panos M. Pardalos. Speeding up branch and bound algorithms for solving the maximum clique problem. *J. Global Optimization*, 59(1):1–21, 2014.
- [McCreesh and Prosser, 2014] Ciaran McCreesh and Patrick Prosser. Reducing the branching in a branch and bound algorithm for the maximum clique problem. In *Principles and Practice of Constraint Programming - 20th International Conference, CP 2014, Lyon, France, September 8-12, 2014. Proceedings*, pages 549–563, 2014.
- [McCreesh *et al.*, 2018] Ciaran McCreesh, Patrick Prosser, Christine Solnon, and James Trimble. When subgraph isomorphism is really hard, and why this matters for graph databases. *J. Artif. Intell. Res.*, 61:723–759, 2018.
- [Mitchell *et al.*, 1992] David G. Mitchell, Bart Selman, and Hector J. Levesque. Hard and easy distributions of SAT problems. In *Proceedings of the 10th National Conference on Artificial Intelligence, San Jose, CA, USA, July 12-16, 1992.*, pages 459–465, 1992.
- [Prosser, 2012] Patrick Prosser. Exact algorithms for maximum clique: A computational study. *Algorithms*, 5(4):545–587, 2012.
- [Sanchis, 1992] Laura A. Sanchis. Test case construction for the vertex cover problem. In *Computational Support for Discrete Mathematics, Proceedings of a DIMACS Workshop, Piscataway, New Jersey, USA, March 12-14, 1992*, pages 315–326, 1992.
- [Sanchis, 1995] Laura A. Sanchis. Generating hard and diverse test sets for NP-hard graph problems. *Discrete Applied Mathematics*, 58(1):35–66, 1995.
- [Segundo *et al.*, 2013] Pablo San Segundo, Fernando Matía, Diego Rodríguez-Losada, and Miguel Hernando. An improved bit parallel exact maximum clique algorithm. *Optimization Letters*, 7(3):467–479, 2013.
- [Smith and Grant, 1997] Barbara M. Smith and Stuart A. Grant. Modelling exceptionally hard constraint satisfaction problems. In *Principles and Practice of Constraint Programming - CP97, Third International Conference, Linz, Austria, October 29 - November 1, 1997, Proceedings*, pages 182–195, 1997.
- [Tomita *et al.*, 2010] Etsuji Tomita, Yoichi Sutani, Takanori Higashi, Shinya Takahashi, and Mitsuo Wakatsuki. A simple and faster branch-and-bound algorithm for finding a maximum clique. In *WALCOM: Algorithms and Computation, 4th International Workshop. Proceedings*, pages 191–203, 2010.
- [Tomita *et al.*, 2016] Etsuji Tomita, Kohei Yoshida, Takuro Hatta, Atsuki Nagao, Hiro Ito, and Mitsuo Wakatsuki. A much faster branch-and-bound algorithm for finding a maximum clique. In *Frontiers in Algorithmics, 10th International Workshop, FAW 2016, Qingdao, China, June 30-July 2, 2016, Proceedings*, pages 215–226, 2016.
- [Zhou *et al.*, 2014] Zhaoyang Zhou, Chu Min Li, Chong Huang, and Ruchu Xu. An exact algorithm with learning for the graph coloring problem. *Computers & OR*, 51:282–301, 2014.