

Publications of Andrew M. Sutton

This document lists all peer-reviewed publications of Andrew M. Sutton, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on October 18, 2021. An up-to-date version is available online at hpi.de/friedrich/docs/publist/sutton.pdf.

Journal articles

- [1] Dang, D.-C., Friedrich, T., Kötzing, T., Krejca, M. S., Lehre, P. K., Oliveto, P. S., Sudholt, D., Sutton, A. M., [Escaping Local Optima Using Crossover with Emergent Diversity](#). In: *Transactions on Evolutionary Computation* 22, pp. 484–497, 2018.

Population diversity is essential for the effective use of any crossover operator. We compare seven commonly used diversity mechanisms and prove rigorous run time bounds for the $(\mu + 1)$ GA using uniform crossover on the fitness function $Jump_k$. All previous results in this context only hold for unrealistically low crossover probability $p_c = O(k/n)$, while we give analyses for the setting of constant $p_c < 1$ in all but one case. Our bounds show a dependence on the problem size n , the jump length k , the population size μ , and the crossover probability p_c . For the typical case of constant $k > 2$ and constant p_c , we can compare the resulting expected optimisation times for different diversity mechanisms assuming an optimal choice of μ : $O(n^{k-1})$ for duplicate elimination/minimisation, $O(n^2 \log n)$ for maximising the convex hull, $O(n \log n)$ for det. crowding (assuming $p_c = k/n$), $O(n \log n)$ for maximising the Hamming distance, $O(n \log n)$ for fitness sharing, $O(n \log n)$ for the single-receiver island model. This proves a sizeable advantage of all variants of the $(\mu + 1)$ GA compared to the $(1+1)$ EA, which requires $\Theta(n^k)$. In a short empirical study we confirm that the asymptotic differences can also be observed experimentally.

- [2] Doerr, B., Neumann, F., Sutton, A. M., [Time Complexity Analysis of Evolutionary Algorithms on Random Satisfiable k-CNF Formulas](#). In: *Algorithmica* 78, pp. 561–586, 2017.

We contribute to the theoretical understanding of randomized search heuristics by investigating their optimization behavior on satisfiable random k-satisfiability instances both in the planted solution model and the uniform model conditional on satisfiability. Denoting the number of variables by n , our main technical result is that the simple $(1+1)$ evolutionary algorithm with high probability finds a satisfying assignment in time $O(n \log n)$ when the clause-variable density is at least logarithmic. For low density instances, evolutionary algorithms seem to be less effective, and all we can show is a subexponential upper bound on the runtime for densities below $1/(k(k-1))$. We complement these mathematical results with numerical experiments on a broader density spectrum. They indicate that, indeed, the $(1+1)$ EA is less efficient on lower densities. Our experiments also suggest that the implicit constants hidden in our main runtime guarantee are low. Our main result extends and considerably improves the result obtained by Sutton and Neumann (Lect Notes Comput Sci 8672:942-951, 2014) in terms of runtime, minimum density, and clause length. These improvements are made possible by establishing a close fitness-distance correlation in certain parts of the search space. This approach might be of independent interest and could be useful for other average-case analyses of randomized search heuristics. While the notion of a fitness-distance correlation has been around for a long time, to the best of our knowledge, this is the first time that fitness-distance correlation is explicitly used to rigorously prove a performance statement for an evolutionary algorithm.

- [3] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [The Compact Genetic Algorithm is Efficient under Extreme Gaussian Noise](#). In: *Transactions on Evolutionary Computation* 21, pp. 477–490, 2017.

Practical optimization problems frequently include uncertainty about the quality measure, for example due to noisy evaluations. Thus, they do not allow for a straightforward application of traditional optimization techniques. In these settings, randomized search heuristics such as evolutionary algorithms are a popular choice because they are often assumed to exhibit some kind of resistance to noise. Empirical evidence suggests that some algorithms, such as estimation of distribution algorithms (EDAs) are robust against a scaling of the noise intensity, even without resorting to explicit noise-handling techniques such as resampling. In this paper, we want to support such claims with mathematical rigor. We introduce the concept of graceful scaling in which the run time of an algorithm scales polynomially with noise intensity. We study a monotone fitness function over binary strings with additive noise taken from a Gaussian distribution. We show that myopic heuristics cannot efficiently optimize the function under arbitrarily intense noise without any explicit noise-handling. Furthermore, we prove that using a population does not help. Finally we show that a simple EDA called the compact Genetic Algorithm can overcome the shortsightedness of mutation-only heuristics to scale gracefully with noise. We conjecture that recombinative genetic algorithms also have this property.

- [4] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [Robustness of Ant Colony Optimization to Noise](#). In: *Evolutionary Computation* 24, pp. 237–254, 2016.

Recently Ant Colony Optimization (ACO) algorithms have been proven to be efficient in uncertain environments, such as noisy or dynamically changing fitness functions. Most of these analyses focus on combinatorial problems, such as path finding. We analyze an ACO algorithm in a setting where we try to optimize the simple OneMax test function, but with additive posterior noise sampled from a Gaussian distribution. Without noise the classical $(\mu + 1)$ -EA outperforms any ACO algorithm, with smaller μ being better; however, with large noise, the $(\mu + 1)$ -EA fails, even for high values of μ (which are known to help against small noise). In this paper we show that ACO is able to deal with arbitrarily large noise in a graceful manner, that is, as long as the evaporation factor p is small enough dependent on the parameter δ^2 of the noise and the dimension n of the search space ($p = o(1/(n(n + \delta \log n)^2 \log n))$), optimization will be successful.

- [5] Sutton, A. M. [Superpolynomial Lower Bounds for the \$\(1+1\)\$ EA on Some Easy Combinatorial Problems](#). In: *Algorithmica* 75, pp. 507–528, 2016.

The (1+1) EA is a simple evolutionary algorithm that is known to be efficient on linear functions and on some combinatorial optimization problems. In this paper, we rigorously study its behavior on two easy combinatorial problems: finding the 2-coloring of a class of bipartite graphs, and constructing satisfying assignments for a class of satisfiable 2-CNF Boolean formulas. We prove that it is inefficient on both problems in the sense that the number of iterations the algorithm needs to minimize the cost functions is superpolynomial with high probability. Our motivation is to better understand the influence of problem instance structure on the runtime character of a simple evolutionary algorithm. We are interested in what kind of structural features give rise to so-called metastable states at which, with probability $1 - o(1)$, the (1+1) EA becomes trapped and subsequently has difficulty leaving. Finally, we show how to modify the (1+1) EA slightly in order to obtain a polynomial-time performance guarantee on both problems.

- [6] Paixão, T., Badkobeh, G., Barton, N. H., Çörüş, D., Dang, D.-C., Friedrich, T., Lehre, P. K., Sudholt, D., Sutton, A., Trubenová, B., [Toward a unifying framework for evolutionary processes](#). In: *Journal of Theoretical Biology* 383, pp. 28–43, 2015.

The theory of population genetics and evolutionary computation have been evolving separately for nearly 30 years. Many results have been independently obtained in both fields and many others are unique to its respective field. We aim to bridge this gap by developing a unifying framework for evolutionary processes that allows both evolutionary algorithms and population genetics models to be cast in the same formal framework. The framework we present here decomposes the evolutionary process into its several components in order to facilitate the identification of similarities between different models. In particular, we propose a classification of evolutionary operators based on the defining properties of the different components. We cast several commonly used operators from both fields into this common framework. Using this, we map different evolutionary and genetic algorithms to different evolutionary regimes and identify candidates with the most potential for the translation of results between the fields. This provides a unified description of evolutionary processes and represents a stepping stone towards new tools and results to both fields.

- [7] Chicano, F., Sutton, A. M., Whitley, L. D., Alba, E., [Fitness Probability Distribution of Bit-Flip Mutation](#). In: *Evolutionary Computation* 23, pp. 217–248, 2015.

Bit-flip mutation is a common mutation operator for evolutionary algorithms applied to optimize functions over binary strings. In this paper, we develop results from the theory of landscapes and Krawtchouk polynomials to exactly compute the probability distribution of fitness values of a binary string undergoing uniform bit-flip mutation. We prove that this probability distribution can be expressed as a polynomial in p , the probability of flipping each bit. We analyze these polynomials and provide closed-form expressions for an easy linear problem Onemax, and an NP-hard problem, MAX-SAT. We also discuss a connection of the results with runtime analysis.

- [8] Nguyen, A. Q., Sutton, A. M., Neumann, F., [Population size matters: Rigorous runtime results for maximizing the hypervolume indicator](#). In: *Theoretical Computer Science* 561, pp. 24–36, 2015.

Using the hypervolume indicator to guide the search of evolutionary multi-objective algorithms has become very popular in recent years. We contribute to the theoretical understanding of these algorithms by carrying out rigorous runtime analyses. We consider multi-objective variants of the problems OneMax and LeadingOnes called OMM and LOTZ, respectively, and investigate hypervolume-based algorithms with population sizes that do not allow coverage of the entire Pareto front. Our results show that LOTZ is easier to optimize than OMM for hypervolume-based evolutionary multi-objective algorithms which is contrary to the results on their single-objective variants and the well-studied (1+1)-EA.

- [9] Whitley, D., Sutton, A. M., Ochoa, G., Chicano, F., [The component model for elementary landscapes and partial neighborhoods](#). In: *Theoretical Computer Science* 545, pp. 59–75, 2014.

Local search algorithms exploit moves on an adjacency graph of the search space. An "elementary landscape" exists if the objective function f is an eigenfunction of the Laplacian of the graph induced by the neighborhood operator; this allows various statistics about the neighborhood to be computed in closed form. A new component based model makes it relatively simple to prove that certain types of landscapes are elementary. The traveling salesperson problem, weighted graph (vertex) coloring and the minimum graph bisection problem yield elementary landscapes under commonly used local search operators. The component model is then used to efficiently compute the mean objective function value over partial neighborhoods for these same problems. For a traveling salesperson problem over n cities, the 2-opt neighborhood can be decomposed into $\lfloor n/2 - 1 \rfloor$ partial neighborhoods. For graph coloring and the minimum graph bisection problem, partial neighborhoods can be used to focus search on those moves that are capable of producing a solution with a strictly improving objective function value.

- [10] Sutton, A. M., Neumann, F., Nallaperuma, S., [Parameterized Runtime Analyses of Evolutionary Algorithms for the Planar Euclidean Traveling Salesperson Problem](#). In: *Evolutionary Computation* 22, pp. 595–628, 2014.

Parameterized runtime analysis seeks to understand the influence of problem structure on algorithmic runtime. In this paper, we contribute to the theoretical understanding of evolutionary algorithms and carry out a parameterized analysis of evolutionary algorithms for the Euclidean traveling salesperson problem (Euclidean TSP). We investigate the structural properties in TSP instances that influence the optimization process of evolutionary algorithms and use this information to bound their runtime. We analyze the runtime in dependence of the number of inner points k . In the first part of the paper, we study a $(\mu + \lambda)$ EA in a strictly black box setting and show that it can solve the Euclidean TSP in expected time $O(n \cdot A(\epsilon) \cdot \max\{(\mu/\lambda) \cdot n^2, 1\} + (\mu/\lambda) \cdot n^{4k}(2k - 1)!)$ where A is a function of the minimum angle ϵ between any three points. Based on insights provided by the analysis, we improve this upper bound by introducing a mixed mutation strategy that incorporates both 2-opt moves and permutation jumps. This strategy improves the upper bound to $O(n \cdot A(\epsilon) \cdot \max\{(\mu/\lambda) \cdot n^2, 1\} + (\mu/\lambda) \cdot n^{2k}(k - 1)!)$. In the second part of the paper, we use the information gained in the analysis to incorporate domain knowledge to design two fixed-parameter tractable (FPT) evolutionary algorithms for the planar Euclidean TSP. We first develop a $(\mu + \lambda)$ EA based on an analysis by M. Theile, 2009, "Exact solutions to the traveling salesperson problem by a population-based evolutionary algorithm," Lecture notes in computer science, Vol. 5482 (pp. 145-155), that solves the TSP with k inner points in $O(\max\{2^k k^2 n^2 \lambda^{-1}, n\})$ generations with probability $1 - e^{-\Omega(n)}$. We then design a $(\mu + \lambda)$ EA that incorporates a dynamic programming step into the fitness evaluation. We prove that a variant of this evolutionary algorithm using 2-opt mutation solves the problem after $O(\max\{(k - 2)! k^{2k-2} \lambda^{-1}, 1\})$ steps in expectation with a cost of $O(nk)$ for each fitness evaluation.

- [11] Kötzing, T., Sutton, A. M., Neumann, F., O'Reilly, U.-M., [The Max problem revisited: The importance of mutation in genetic programming](#). In: *Theoretical Computer Science* 545, pp. 94–107, 2014.

This paper contributes to the rigorous understanding of genetic programming algorithms by providing runtime complexity analyses of the well-studied Max problem. Several experimental studies have indicated that it is hard to solve the Max problem with crossover-based algorithms. Our analyses show that different variants of the Max problem can provably be solved using simple mutation-based genetic programming algorithms. Our results advance the body of computational complexity analyses of genetic programming, indicate the importance of mutation in genetic programming, and reveal new insights into the behavior of mutation-based genetic programming algorithms.

- [12] Sutton, A. M., Chicano, F., Whitley, L. D., [Fitness Function Distributions over Generalized Search Neighborhoods in the \$q\$ -ary Hypercube](#). In: *Evolutionary Computation 21*, pp. 561–590, 2013.

The frequency distribution of a fitness function over regions of its domain is an important quantity for understanding the behavior of algorithms that employ randomized sampling to search the function. In general, exactly characterizing this distribution is at least as hard as the search problem, since the solutions typically live in the tails of the distribution. However, in some cases it is possible to efficiently retrieve a collection of quantities called moments that describe the distribution. In this paper, we consider functions of bounded epistasis that are defined over length- n strings from a finite alphabet of cardinality q . Many problems in combinatorial optimization can be specified as search problems over functions of this type. Employing Fourier analysis of functions over finite groups, we derive an efficient method for computing the exact moments of the frequency distribution of fitness functions over Hamming regions of the q -ary hypercube. We then use this approach to derive equations that describe the expected fitness of the offspring of any point undergoing uniform mutation. The results we present provide insight into the statistical structure of the fitness function for a number of combinatorial problems. For the graph coloring problem, we apply our results to efficiently compute the average number of constraint violations that lie within a certain number of steps of any coloring. We derive an expression for the mutation rate that maximizes the expected fitness of an offspring at each fitness level. We also apply the results to the slightly more complex frequency assignment problem, a relevant application in the domain of the telecommunications industry. As with the graph coloring problem, we provide formulas for the average value of the fitness function in Hamming regions around a solution and the expectation-optimal mutation rate.

- [13] Sutton, A. M., Whitley, L. D., Howe, A. E., [Computing the moments of \$k\$ -bounded pseudo-Boolean functions over Hamming spheres of arbitrary radius in polynomial time](#). In: *Theoretical Computer Science* 425, pp. 58–74, 2012.

We show that given a k -bounded pseudo-Boolean function f , one can always compute the c th moment of f over regions of arbitrary radius in Hamming space in polynomial time using algebraic information from the adjacency structure (where k and c are constants). This result has implications for evolutionary algorithms and local search algorithms because information about promising regions of the search space can be efficiently retrieved, even if the cardinality of the region is exponential in the problem size. Finally, we use our results to introduce a method of efficiently calculating the expected fitness of mutations for evolutionary algorithms.

Conference papers

- [14] Friedrich, T., Neumann, F., Rothenberger, R., Sutton, A. M., [Solving Non-Uniform Planted and Filtered Random SAT Formulas Greedily](#). In: *Theory and Applications of Satisfiability Testing (SAT)*, pp. 188–206, 2021.

Recently, there has been an interest in studying non-uniform random k -satisfiability (k -SAT) models in order to address the non-uniformity of formulas arising from real-world applications. While uniform random k -SAT has been extensively studied from both a theoretical and experimental perspective, understanding the algorithmic complexity of heterogeneous distributions is still an open challenge. When a sufficiently dense formula is guaranteed to be satisfiable by conditioning or a planted assignment, it is well-known that uniform random k -SAT is easy on average. We generalize this result to the broad class of non-uniform random k -SAT models that are characterized only by an ensemble of distributions over variables with a mild balancing condition. This balancing condition rules out extremely skewed distributions in which nearly half the variables occur less frequently than a small constant fraction of the most frequent variables, but generalizes recently studied non-uniform k -SAT distributions such as power-law and geometric formulas. We show that for all formulas generated from this model of at least logarithmic densities, a simple greedy algorithm can find a solution with high probability. As a side result we show that the total variation distance between planted and filtered (conditioned on satisfiability) models is $o(1)$ once the planted model produces formulas with a unique solution with probability $1 - o(1)$. This holds for all random k -SAT models where the signs of variables are drawn uniformly and independently at random.

- [15] Bläsius, T., Friedrich, T., Sutton, A. M., [On the Empirical Time Complexity of Scale-Free 3-SAT at the Phase Transition](#). In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 117–134, 2019.

The hardness of formulas at the solubility phase transition of random propositional satisfiability (SAT) has been intensely studied for decades both empirically and theoretically. Solvers based on stochastic local search (SLS) appear to scale very well at the critical threshold, while complete backtracking solvers exhibit exponential scaling. On industrial SAT instances, this phenomenon is inverted: backtracking solvers can tackle large industrial problems, where SLS-based solvers appear to stall. Industrial instances exhibit sharply different structure than uniform random instances. Among many other properties, they are often heterogeneous in the sense that some variables appear in many while others appear in only few clauses. We conjecture that the heterogeneity of SAT formulas alone already contributes to the trade-off in performance between SLS solvers and complete backtracking solvers. We empirically determine how the run time of SLS vs. backtracking solvers depends on the heterogeneity of the input, which is controlled by drawing variables according to a scale-free distribution. Our experiments reveal that the efficiency of complete solvers at the phase transition is strongly related to the heterogeneity of the degree distribution. We report results that suggest the depth of satisfying assignments in complete search trees is influenced by the level of heterogeneity as measured by a power-law exponent. We also find that incomplete SLS solvers, which scale well on uniform instances, are not affected by heterogeneity. The main contribution of this paper utilizes the scale-free random 3-SAT model to isolate heterogeneity as an important factor in the scaling discrepancy between complete and SLS solvers at the uniform phase transition found in previous works.

- [16] Friedrich, T., Kötzing, T., Quinzan, F., Sutton, A. M., [Improving the Run Time of the \(1+1\) Evolutionary Algorithm with Luby Sequences](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 301–308, 2018.

In the context of black box optimization, the most common way to handle deceptive attractors is to periodically restart the algorithm. In this paper, we explore the benefits of combining the simple (1 + 1) Evolutionary Algorithm (EA) with the Luby Universal Strategy - the (1 + 1) EA_M , a meta-heuristic that does not require parameter tuning. We first consider two artificial pseudo-Boolean landscapes, on which the (1 + 1) EA exhibits exponential run time. We prove that the (1 + 1) EA_M has polynomial run time on both instances. We then consider the Minimum Vertex Cover on two classes of graphs. Again, the (1 + 1) EA yields exponential run time on those instances, and the (1 + 1) EA_M finds the global optimum in polynomial time. We conclude by studying the Makespan Scheduling. We consider an instance on which the (1 + 1) EA does not find a $(4/3 - \epsilon)$ -approximation in polynomial time, and we show that the (1 + 1) EA_M reaches a $(4/3 - \epsilon)$ -approximation in polynomial time. We then prove that the (1 + 1) EA_M serves as an Efficient Polynomial-time Approximation Scheme (EPTAS) for the Partition Problem, for a $(1 + \epsilon)$ -approximation with $\epsilon > 4/n$.

- [17] Friedrich, T., Krohmer, A., Rothenberger, R., Sutton, A. M., [Phase Transitions for Scale-Free SAT Formulas](#). In: *Conference on Artificial Intelligence (AAAI)*, pp. 3893–3899, 2017.

Recently, a number of non-uniform random satisfiability models have been proposed that are closer to practical satisfiability problems in some characteristics. In contrast to uniform random Boolean formulas, scale-free formulas have a variable occurrence distribution that follows a power law. It has been conjectured that such a distribution is a more accurate model for some industrial instances than the uniform random model. Though it seems that there is already an awareness of a threshold phenomenon in such models, there is still a complete picture lacking. In contrast to the uniform model, the critical density threshold does not lie at a single point, but instead exhibits a functional dependency on the power-law exponent. For scale-free formulas with clauses of length $k = 2$, we give a lower bound on the phase transition threshold as a function of the scaling parameter. We also perform computational studies that suggest our bound is tight and investigate the critical density for formulas with higher clause lengths. Similar to the uniform model, on formulas with $k \geq 3$, we find that the phase transition regime corresponds to a set of formulas that are difficult to solve by backtracking search.

- [18] Friedrich, T., Krohmer, A., Rothenberger, R., Sauerwald, T., Sutton, A. M., [Bounds on the Satisfiability Threshold for Power Law Distributed Random SAT](#). In: *European Symposium on Algorithms (ESA)*, pp. 37:1–37:15, 2017.

Propositional satisfiability (SAT) is one of the most fundamental problems in computer science. The worst-case hardness of SAT lies at the core of computational complexity theory. The average-case analysis of SAT has triggered the development of sophisticated rigorous and non-rigorous techniques for analyzing random structures. Despite a long line of research and substantial progress, nearly all theoretical work on random SAT assumes a uniform distribution on the variables. In contrast, real-world instances often exhibit large fluctuations in variable occurrence. This can be modeled by a scale-free distribution of the variables, which results in distributions closer to industrial SAT instances. We study random k -SAT on n variables, $m = \Theta(n)$ clauses, and a power law distribution on the variable occurrences with exponent β . We observe a satisfiability threshold at $\beta = (2k - 1)/(k - 1)$. This threshold is tight in the sense that instances with $\beta < (2k - 1)/(k - 1) - \epsilon$ for any constant $\epsilon > 0$ are unsatisfiable with high probability (w.h.p.). For $\beta \geq (2k - 1)/(k - 1) + \epsilon$, the picture is reminiscent of the uniform case: instances are satisfiable w.h.p. for sufficiently small constant clause-variable ratios m/n ; they are unsatisfiable above a ratio m/n that depends on β .

- [19] Friedrich, T., Kötzing, T., Quinzan, F., Sutton, A. M., [Resampling vs Recombination: a Statistical Run Time Estimation](#). In: *Foundations of Genetic Algorithms (FOGA)*, pp. 25–35, 2017.

Noise is pervasive in real-world optimization, but there is still little understanding of the interplay between the operators of randomized search heuristics and explicit noise-handling techniques, such as statistical resampling. In this paper, we report on several statistical models and theoretical results that help to clarify this reciprocal relationship for a collection of randomized search heuristics on noisy functions. We consider the optimization of pseudo-Boolean functions under additive posterior Gaussian noise and explore the trade-off between noise reduction and the computational cost of resampling. We first perform experiments to find the optimal parameters at a given noise intensity for a mutation-only evolutionary algorithm, a genetic algorithm employing recombination, an estimation of distribution algorithm (EDA), and an ant colony optimization algorithm. We then observe how the optimal parameter depends on the noise intensity for the different algorithms. Finally, we locate the point where statistical resampling costs more than it is worth in terms of run time. We find that the EA requires the highest number of resamples to obtain the best speed-up, whereas crossover reduces both the run time and the number of resamples required. Most surprisingly, we find that EDA-like algorithms require no resampling, and can handle noise implicitly.

- [20] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [The Benefit of Recombination in Noisy Evolutionary Search](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 161–162, 2016.

Practical optimization problems frequently include uncertainty about the quality measure, for example due to noisy evaluations. Thus, they do not allow for a straightforward application of traditional optimization techniques. In these settings, randomized search heuristics such as evolutionary algorithms are a popular choice because they are often assumed to exhibit some kind of resistance to noise. Empirical evidence suggests that some algorithms, such as estimation of distribution algorithms (EDAs) are robust against a scaling of the noise intensity, even without resorting to explicit noise-handling techniques such as resampling. In this paper, we want to support such claims with mathematical rigor. We introduce the concept of graceful scaling in which the run time of an algorithm scales polynomially with noise intensity. We study a monotone fitness function over binary strings with additive noise taken from a Gaussian distribution. We show that myopic heuristics cannot efficiently optimize the function under arbitrarily intense noise without any explicit noise-handling. Furthermore, we prove that using a population does not help. Finally we show that a simple EDA called the Compact Genetic Algorithm can overcome the shortsightedness of mutation-only heuristics to scale gracefully with noise. We conjecture that recombinative genetic algorithms also have this property.

- [21] Friedrich, T., Kötzing, T., Quinzan, F., Sutton, A. M., [Ant Colony Optimization Beats Resampling on Noisy Functions](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 3–4, 2016.

Despite the pervasiveness of noise in real-world optimization, there is little understanding of the interplay between the operators of randomized search heuristics and explicit noise-handling techniques such as statistical resampling. Ant Colony Optimization (ACO) algorithms are claimed to be particularly well-suited to dynamic and noisy problems, even without explicit noise-handling techniques. In this work, we empirically investigate the trade-offs between resampling and the noise-handling abilities of ACO algorithms. Our main focus is to locate the point where resampling costs more than it is worth.

- [22] Dang, D.-C., Friedrich, T., Krejca, M. S., Kötzing, T., Lehre, P. K., Oliveto, P. S., Sudholt, D., Sutton, A. M., [Escaping Local Optima with Diversity Mechanisms and Crossover](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 645–652, 2016.

Population diversity is essential for the effective use of any crossover operator. We compare seven commonly used diversity mechanisms and prove rigorous run time bounds for the $(\mu + 1)$ GA using uniform crossover on the fitness function $Jump_k$. All previous results in this context only hold for unrealistically low crossover probability $p_c = O(k/n)$, while we give analyses for the setting of constant $p_c < 1$ in all but one case. Our bounds show a dependence on the problem size n , the jump length k , the population size μ , and the crossover probability p_c . For the typical case of constant $k > 2$ and constant p_c , we can compare the resulting expected optimisation times for different diversity mechanisms assuming an optimal choice of μ : $O(n^{k-1})$ for duplicate elimination/minimisation, $O(n^2 \log n)$ for maximising the convex hull, $O(n \log n)$ for det. crowding (assuming $p_c = k/n$), $O(n \log n)$ for maximising the Hamming distance, $O(n \log n)$ for fitness sharing, $O(n \log n)$ for the single-receiver island model. This proves a sizeable advantage of all variants of the $(\mu + 1)$ GA compared to the $(1+1)$ EA, which requires $\Theta(n^k)$. In a short empirical study we confirm that the asymptotic differences can also be observed experimentally.

- [23] Dang, D.-C., Lehre, P. K., Friedrich, T., Kötzing, T., Krejca, M. S., Oliveto, P. S., Sudholt, D., Sutton, A. M., [Emergence of Diversity and its Benefits for Crossover in Genetic Algorithms](#). In: *Parallel Problem Solving From Nature (PPSN)*, pp. 890–900, 2016.

Population diversity is essential for avoiding premature convergence in Genetic Algorithms (GAs) and for the effective use of crossover. Yet the dynamics of how diversity emerges in populations are not well understood. We use rigorous runtime analysis to gain insight into population dynamics and GA performance for a standard $(\mu + 1)$ GA and the $Jump_k$ test function. By studying the stochastic process underlying the size of the largest collection of identical genotypes we show that the interplay of crossover followed by mutation may serve as a catalyst leading to a sudden burst of diversity. This leads to improvements of the expected optimisation time of order $\Omega(n/\log n)$ compared to mutation-only algorithms like the $(1+1)$ EA.

- [24] Friedrich, T., Kötzing, T., Sutton, A. M., [On the Robustness of Evolving Populations](#). In: *Parallel Problem Solving From Nature (PPSN)*, pp. 771–781, 2016.

Most theoretical work that studies the benefit of recombination focuses on the ability of crossover to speed up optimization time on specific search problems. In this paper, we take a slightly different perspective and investigate recombination in the context of evolving solutions that exhibit mutational robustness, i.e., they display insensitivity to small perturbations. Various models in population genetics have demonstrated that increasing the effective recombination rate promotes the evolution of robustness. We show this result also holds in the context of evolutionary computation by proving crossover promotes the evolution of robust solutions in the standard $(\mu + 1)$ GA. Surprisingly, our results show that the effect is present even when robust solutions are at a selective disadvantage due to lower fitness values.

- [25] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [Graceful Scaling on Uniform versus Steep-Tailed Noise](#). In: *Parallel Problem Solving From Nature (PPSN)*, pp. 761–770, 2016.

Recently, different evolutionary algorithms (EAs) have been analyzed in noisy environments. The most frequently used noise model for this was additive posterior noise (noise added after the fitness evaluation) taken from a Gaussian distribution. In particular, for this setting it was shown that the $(\mu + 1)$ -EA on OneMax does not scale gracefully (higher noise cannot efficiently be compensated by higher μ). In this paper we want to understand whether there is anything special about the Gaussian distribution which makes the $(\mu + 1)$ -EA not scale gracefully. We keep the setting of posterior noise, but we look at other distributions. We see that for exponential tails the $(\mu + 1)$ -EA on OneMax does also not scale gracefully, for similar reasons as in the case of Gaussian noise. On the other hand, for uniform distributions (as well as other, similar distributions) we see that the $(\mu + 1)$ -EA on OneMax does scale gracefully, indicating the importance of the noise model.

- [26] Doerr, B., Neumann, F., Sutton, A. M., [Improved Runtime Bounds for the \$\(1+1\)\$ EA on Random 3-CNF Formulas Based on Fitness-Distance Correlation](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 1415–1422, 2015.

With this paper, we contribute to the theoretical understanding of randomized search heuristics by investigating their behavior on random 3-SAT instances. We improve the results for the $(1+1)$ EA obtained by Sutton and Neumann [PPSN 2014, 942-951] in three ways. First, we reduce the upper bound by a linear factor and prove that the $(1+1)$ EA obtains optimal solutions in time $O(n \log n)$ with high probability on asymptotically almost all high-density satisfiable 3-CNF formulas. Second, we extend the range of densities for which this bound holds to satisfiable formulas of at least logarithmic density. Finally, we complement these mathematical results with numerical experiments that summarize the behavior of the $(1+1)$ EA on formulas along the density spectrum, and suggest that the implicit constants hidden in our bounds are low. Our proofs are based on analyzing the run of the algorithm by establishing a fitness-distance correlation. This approach might be of independent interest and we are optimistic that it is useful for the analysis of randomized search heuristics in various other settings. To our knowledge, this is the first time that fitness-distance correlation is explicitly used to rigorously prove a performance statement for an evolutionary algorithm.

- [27] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [Robustness of Ant Colony Optimization to Noise](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 17–24, 2015. **Best-Paper Award (ACO/SI Track)**.

Recently Ant Colony Optimization (ACO) algorithms have been proven to be efficient in uncertain environments, such as noisy or dynamically changing fitness functions. Most of these analyses focus on combinatorial problems, such as path finding. We analyze an ACO algorithm in a setting where we try to optimize the simple OneMax test function, but with additive posterior noise sampled from a Gaussian distribution. Without noise the classical $(\mu + 1)$ -EA outperforms any ACO algorithm, with smaller μ being better; however, with large noise, the $(\mu + 1)$ -EA fails, even for high values of μ (which are known to help against small noise). In this paper we show that ACO is able to deal with arbitrarily large noise in a graceful manner, that is, as long as the evaporation factor μ is small enough dependent on the parameter δ^2 of the noise and the dimension n of the search space ($p = o(1/(n(n + \delta \log n)^2 \log n))$), optimization will be successful.

- [28] Friedrich, T., Kötzing, T., Krejca, M. S., Sutton, A. M., [The Benefit of Recombination in Noisy Evolutionary Search](#). In: *International Symposium of Algorithms and Computation (ISAAC)*, pp. 140–150, 2015.

Practical optimization problems frequently include uncertainty about the quality measure, for example due to noisy evaluations. Thus, they do not allow for a straightforward application of traditional optimization techniques. In these settings meta-heuristics are a popular choice for deriving good optimization algorithms, most notably evolutionary algorithms which mimic evolution in nature. Empirical evidence suggests that genetic recombination is useful in uncertain environments because it can stabilize a noisy fitness signal. With this paper we want to support this claim with mathematical rigor. The setting we consider is that of noisy optimization. We study a simple noisy fitness function that is derived by adding Gaussian noise to a monotone function. First, we show that a classical evolutionary algorithm that does not employ sexual recombination (the $(\mu + 1)$ -EA) cannot handle the noise efficiently, regardless of the population size. Then we show that an evolutionary algorithm which does employ sexual recombination (the Compact Genetic Algorithm, short: cGA) can handle the noise using a graceful scaling of the population.

- [29] Chicano, F., Whitley, D., Sutton, A. M., [Efficient identification of improving moves in a ball for pseudo-boolean problems](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 437–444, 2014.

Hill climbing algorithms are at the core of many approaches to solve optimization problems. Such algorithms usually require the complete enumeration of a neighborhood of the current solution. In the case of problems defined over binary strings of length n , we define the r -ball neighborhood as the set of solutions at Hamming distance r or less from the current solution. For $r \ll n$ this neighborhood contains $\Theta(n^r)$ solutions. In this paper efficient methods are introduced to locate improving moves in the r -ball neighborhood for problems that can be written as a sum of a linear number of subfunctions depending on a bounded number of variables. NK-landscapes and MAX- k -SAT are examples of these problems. If the number of subfunctions depending on any given variable is also bounded, then we prove that the method can explore the neighborhood in constant time, despite the fact that the number of solutions in the neighborhood is polynomial in n . We develop a hill climber based on our exploration method and we analyze its efficiency and efficacy using experiments with NK-landscapes instances.

- [30] Sutton, A. M., Neumann, F., [Runtime Analysis of Evolutionary Algorithms on Randomly Constructed High-Density Satisfiable 3-CNF Formulas](#). In: *Parallel Problem Solving from Nature (PPSN)*, pp. 942–951, 2014.

We show that simple mutation-only evolutionary algorithms find a satisfying assignment on two similar models of random planted 3-CNF Boolean formulas in polynomial time with high probability in the high constraint density regime. We extend the analysis to random formulas conditioned on satisfiability (i.e., the so-called filtered distribution) and conclude that most high-density satisfiable formulas are easy for simple evolutionary algorithms. With this paper, we contribute the first rigorous study of randomized search heuristics from the evolutionary computation community on well-studied distributions of random satisfiability problems.

- [31] Nallaperuma, S., Sutton, A. M., Neumann, F., [Fixed-parameter evolutionary algorithms for the Euclidean Traveling Salesperson problem](#). In: *Congress on Evolutionary Computation (CEC)*, pp. 2037–2044, 2013.

We contribute to the theoretical understanding of evolutionary algorithms and carry out a parameterized analysis of evolutionary algorithms for the Euclidean traveling salesperson problem (Euclidean TSP). We exploit structural properties related to the optimization process of evolutionary algorithms for this problem and use them to bound the runtime of evolutionary algorithms. Our analysis studies the runtime in dependence of the number of inner points k and shows that simple evolutionary algorithms solve the Euclidean TSP in expected time $O(n^{4k}(2k-1)!)$. Moreover, we show that, under reasonable geometric constraints, a locally optimal 2-opt tour can be found by randomized local search in expected time $O(n^{2k!})$.

- [32] Nallaperuma, S., Sutton, A. M., Neumann, F., [Parameterized complexity analysis and more effective construction methods for ACO algorithms and the euclidean traveling salesperson problem](#). In: *Congress on Evolutionary Computation (CEC)*, pp. 2045–2052, 2013.

We propose a new construction procedure for ant colony optimization (ACO) algorithms working on the Euclidean traveling salesperson problem (TSP) that preserves the ordering on the convex hull of the points in the instance. The procedure is inspired by theoretical analyses for simple evolutionary algorithms that are provably more efficient on instances where the number of inner points of the instance is not too large. We integrate the construction procedure into the well-known MaxMin Ant System (MMAS) and empirically show that it leads to more efficient optimization on instances where the number of inner points is not too high.

- [33] Sutton, A. M., Neumann, F., [A Parameterized Runtime Analysis of Evolutionary Algorithms for the Euclidean Traveling Salesperson Problem](#). In: *Conference on Artificial Intelligence (AAAI)*, 2012.

Parameterized runtime analysis seeks to understand the influence of problem structure on algorithmic runtime. In this paper, we contribute to the theoretical understanding of evolutionary algorithms and carry out a parameterized analysis of evolutionary algorithms for the Euclidean traveling salesperson problem (Euclidean TSP). We investigate the structural properties in TSP instances that influence the optimization process of evolutionary algorithms and use this information to bound the runtime of simple evolutionary algorithms. Our analysis studies the runtime in dependence of the number of inner points k and shows that $(\mu + \lambda)$ evolutionary algorithms solve the Euclidean TSP in expected time $O((\mu/\lambda) \cdot n^3\gamma(\epsilon) + n\gamma(\epsilon) + (\mu/\lambda) \cdot n^{4k}(2k-1)!)$ where γ is a function of the minimum angle *epsilon* between any three points. Finally, our analysis provides insights into designing a mutation operator that improves the upper bound on expected runtime. We show that a mixed mutation strategy that incorporates both 2-opt moves and permutation jumps results in an upper bound of $O((\mu/\lambda) \cdot n^3\gamma(\epsilon) + n\gamma(\epsilon) + (\mu/\lambda) \cdot n^{2k}(k-1)!)$ for the $(\mu + \lambda)$ EA.

- [34] Sutton, A. M., Day, J., Neumann, F., [A parameterized runtime analysis of evolutionary algorithms for MAX-2-SAT](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 433–440, 2012.

We investigate the MAX-2-SAT problem and study evolutionary algorithms by parameterized runtime analysis. The parameterized runtime analysis of evolutionary algorithms has been initiated recently and reveals new insights into which type of instances of NP-hard combinatorial optimization problems are hard to solve by evolutionary computing methods. We show that a variant of the (1+1) EA is a fixed-parameter evolutionary algorithm with respect to the standard parameterization for MAX-2-SAT. Furthermore, we study how the dependencies between the variables affect problem difficulty and present fixed-parameter evolutionary algorithms for the MAX-(2,3)-SAT problem where the studied parameter is the diameter of the variable graph.

- [35] Kötzing, T., Sutton, A. M., Neumann, F., O'Reilly, U.-M., [The max problem revisited: the importance of mutation in genetic programming](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 1333–1340, 2012.

This paper contributes to the rigorous understanding of genetic programming algorithms by providing runtime complexity analyses of the well-studied Max problem. Several experimental studies have indicated that it is hard to solve the Max problem with crossover-based algorithms. Our analyses show that different variants of the Max problem can provably be solved using simple mutation-based genetic programming algorithms. Our results advance the body of computational complexity analyses of genetic programming, indicate the importance of mutation in genetic programming, and reveal new insights into the behavior of mutation-based genetic programming algorithms.

- [36] Sutton, A. M., Neumann, F., [A Parameterized Runtime Analysis of Simple Evolutionary Algorithms for Makespan Scheduling](#). In: *Parallel Problem Solving from Nature (PPSN)*, pp. 52–61, 2012.

We consider simple multi-start evolutionary algorithms applied to the classical NP-hard combinatorial optimization problem of Makespan Scheduling on two machines. We study the dependence of the runtime of this type of algorithm on three different key hardness parameters. By doing this, we provide further structural insights into the behavior of evolutionary algorithms for this classical problem.

- [37] Sutton, A. M., Whitley, D., Howe, A. E., [Approximating the distribution of fitness over hamming regions](#). In: *Foundations of Genetic Algorithms (FOGA)*, pp. 93–104, 2011.

The distribution of fitness values across a set of states sharply influences the dynamics of evolutionary processes and heuristic search in combinatorial optimization. In this paper we present a method for approximating the distribution of fitness values over Hamming regions by solving a linear programming problem that incorporates low order moments of the target function. These moments can be retrieved in polynomial time for select problems such as MAX- k -SAT using Walsh analysis. The method is applicable to any real function on binary strings that is epistatically bounded and discrete with asymptotic bounds on the cardinality of its codomain. We perform several studies on the ONE-MAX and MAX- k -SAT domains to assess the accuracy of the approximation and its dependence on various factors. We show that the approximation can accurately predict the number of states within a Hamming region that have an improving fitness value.

- [38] Sutton, A. M., Whitley, D., Howe, A. E., [Mutation rates of the \(1+1\)-EA on pseudo-boolean functions of bounded epistasis](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 973–980, 2011.

When the epistasis of the fitness function is bounded by a constant, we show that the expected fitness of an offspring of the (1+1)-EA can be efficiently computed for any point. Moreover, we show that, for any point, it is always possible to efficiently retrieve the "best" mutation rate at that point in the sense that the expected fitness of the resulting offspring is maximized. On linear functions, it has been shown that a mutation rate of $1/n$ is provably optimal. On functions where epistasis is bounded by a constant k , we show that for sufficiently high fitness, the commonly used mutation rate of $1/n$ is also best, at least in terms of maximizing the expected fitness of the offspring. However, we find for certain ranges of the fitness function, a better mutation rate can be considerably higher, and can be found by solving for the real roots of a degree- k polynomial whose coefficients contain the nonzero Walsh coefficients of the fitness function. Simulation results on maximum k -satisfiability problems and NK-landscapes show that this expectation-maximized mutation rate can cause significant gains early in search.

- [39] Sutton, A. M., Howe, A. E., Whitley, L. D., [Directed Plateau Search for MAX- \$k\$ -SAT](#). In: *Symposium on Combinatorial Search (SOCS)*, 2010.

Local search algorithms for MAX- k -SAT must often explore large regions of mutually connected equal moves, or plateaus, typically by taking random walks through the region. In this paper, we develop a surrogate plateau "gradient" function using a Walsh transform of the objective function. This function gives the mean value of the objective function over localized volumes of the search space. This information can be used to direct search through plateaus more quickly. The focus of this paper is on demonstrating that formal analysis of search space structure can direct existing algorithms in a more principled manner than random walks. We show that embedding the gradient computation into a hill-climbing local search for MAX- k -SAT improves its convergence profile.

- [40] Whitley, L. D., Sutton, A. M., [Partial neighborhoods of elementary landscapes](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 381–388, 2009.

This paper introduces a new component based model that makes it relatively simple to prove that certain types of landscapes are elementary. We use the model to reconstruct proofs for the Traveling Salesman Problem, Graph Coloring and Min-Cut Graph Partitioning. The same model is then used to efficiently compute the average values over particular partial neighborhoods for these same problems. For Graph Coloring and Min-Cut Graph Partitioning, this computation can be used to focus search on those moves that are most likely to yield an improving move, ignoring moves that cannot yield an improving move. Let x be a candidate solution with objective function value $f(x)$. The mean value of the objective function over the entire landscape is denoted f . Normally in an elementary landscape one can only be sure that a neighborhood includes an improving move (assuming minimization) if $f(x) > f$. However, by computing the expected value of an appropriate partial neighborhood it is sometimes possible to know that an improving move exists in the partial neighborhood even when $f(x) < f$.

- [41] Sutton, A. M., Whitley, L. D., Howe, A. E., [A polynomial time computation of the exact correlation structure of \$k\$ -satisfiability landscapes](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 365–372, 2009.

The autocorrelation function and related correlation length are statistical quantities that capture the ruggedness of the fitness landscape: a measure that is directly related to the hardness of a problem for certain heuristic search algorithms. Typically, these quantities are estimated empirically by sampling along a random walk. In this paper, we show that a polynomial-time Walsh decomposition of the k -satisfiability evaluation function allows us to compute the exact autocorrelation function and correlation length for any given k -satisfiability instance. We also use the decomposition to compute a theoretical expectation for the autocorrelation function and correlation length over the ensemble of instances generated uniformly at random. We find that this expectation is invariant to the constrainedness of the problem as measured by the ratio of clauses to variables. However, we show that filtered problems, which are typically used in local search studies, have a bias that causes a significant deviation from the expected correlation structure of unfiltered, uniformly generated problems.

- [42] Matthews, D. C., Sutton, A. M., Hains, D., Whitley, L. D., [Improved Robustness through Population Variance in Ant Colony Optimization](#). In: *Stochastic Local Search Algorithms (SLS)*, pp. 145–149, 2009.
- Ant Colony Optimization algorithms are population-based Stochastic Local Search algorithms that mimic the behavior of ants, simulating pheromone trails to search for solutions to combinatorial optimization problems. This paper introduces Population Variance, a novel approach to ACO algorithms that allows parameters to vary across the population over time, leading to solution construction differences that are not strictly stochastic. The increased exploration appears to help the search escape from local optima, significantly improving the robustness of the algorithm with respect to suboptimal parameter settings.
- [43] Sutton, A. M., Howe, A. E., Whitley, L. D., [A Theoretical Analysis of the \$k\$ -Satisfiability Search Space](#). In: *Stochastic Local Search Algorithms (SLS)*, pp. 46–60, 2009.
- Local search algorithms perform surprisingly well on the k -satisfiability (k -SAT) problem. However, few theoretical analyses of the k -SAT search space exist. In this paper we study the search space of the k -SAT problem and show that it can be analyzed by a decomposition. In particular, we prove that the objective function can be represented as a superposition of exactly k elementary landscapes. We show that this decomposition allows us to immediately compute the expectation of the objective function evaluated across neighboring points. We use this result to prove previously unknown bounds for local maxima and plateau width in the 3-SAT search space. We compute these bounds numerically for a number of instances and show that they are non-trivial across a large set of benchmarks.
- [44] Sutton, A. M., Howe, A. E., Whitley, L. D., [Estimating Bounds on Expected Plateau Size in MAXSAT Problems](#). In: *Stochastic Local Search Algorithms (SLS)*, pp. 31–45, 2009.
- Stochastic local search algorithms can now successfully solve MAXSAT problems with thousands of variables or more. A key to this success is how effectively the search can navigate and escape plateau regions. Furthermore, the solubility of a problem depends on the size and exit density of plateaus, especially those closest to the optimal solution. In this paper we model the plateau phenomenon as a percolation process on hypercube graphs. We develop two models for estimating bounds on the size of plateaus and prove that one is a lower bound and the other an upper bound on the expected size of plateaus at a given level. The models’ accuracy is demonstrated on controlled random hypercube landscapes. We apply the models to MAXSAT through analogy to hypercube graphs and by introducing an approach to estimating, through sampling, a key parameter of the models. Using this approach, we assess the accuracy of our bound estimations on uniform random and structured benchmarks. Surprisingly, we find similar trends in accuracy across random and structured problem instances. Less surprisingly, we find a high accuracy on smaller plateaus with systematic divergence as plateaus increase in size.
- [45] Whitley, D., Sutton, A. M., Howe, A. E., [Understanding elementary landscapes](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 585–592, 2008.
- The landscape formalism unites a finite candidate solution set to a neighborhood topology and an objective function. This construct can be used to model the behavior of local search on combinatorial optimization problems. A landscape is elementary when it possesses a unique property that results in a relative smoothness and decomposability to its structure. In this paper we explain elementary landscapes in terms of the expected value of solution components which are transformed in the process of moving from an incumbent solution to a neighboring solution. We introduce new results about the properties of elementary landscapes and discuss the practical implications for search algorithms.
- [46] Lunacek, M., Whitley, D., Sutton, A. M., [The Impact of Global Structure on Search](#). In: *Parallel Problem Solving from Nature (PPSN)*, pp. 498–507, 2008.
- Population-based methods are often considered superior on multimodal functions because they tend to explore more of the fitness landscape before they converge. We show that the effectiveness of this strategy is highly dependent on a function’s global structure. When the local optima are not structured in a predictable way, exploration can misguide search into sub-optimal regions. Limiting exploration can result in a better non-intuitive global search strategy.
- [47] Sutton, A. M., Lunacek, M., Whitley, L. D., [Differential evolution and non-separability: using selective pressure to focus search](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 1428–1435, 2007.
- Recent results show that the Differential Evolution algorithm has significant difficulty on functions that are not linearly separable. On such functions, the algorithm must rely primarily on its differential mutation procedure which, unlike its recombination strategy, is rotationally invariant. We conjecture that this mutation strategy lacks sufficient selective pressure when appointing parent and donor vectors to have satisfactory exploitative power on non-separable functions. We find that imposing pressure in the form of rank-based differential mutation results in a significant improvement of exploitation on rotated benchmarks.
- [48] Sutton, A. M., Howe, A. E., Whitley, L. D., [Using Adaptive Priority Weighting to Direct Search in Probabilistic Scheduling](#). In: *International Conference on Automated Planning and Scheduling (ICAPS)*, pp. 320–327, 2007.
- Many scheduling problems reside in uncertain and dynamic environments - tasks have a nonzero probability of failure and may need to be rescheduled. In these cases, an optimized solution for a short-term time horizon may have a detrimental impact over a broader time scale. We examine a scheduling domain in which time and energy on a phased array radar system is allocated to track objects in orbit around the earth. This domain requires probabilistic modeling to optimize the expected number of successful tasks on a particular day. Failed tasks must be attempted again on subsequent days. Given a set of task requests, we study two long-term objectives: percentage of requests initially successful, and the average time between successful request updates. We investigate adaptive priority weighting strategies that directly influence the short-term objective function and thus indirectly influence the long-term goals. We find that adapting priority weights based on when individual tasks succeed or fail allows a catalog of requests to be filled more quickly. Furthermore, with adaptive priorities, we observe a Pareto-front effect between the two long-term objectives as we modify how priorities are weighted, but an inverse effect of weighting when the priorities are not adapted.
- [49] Smith, J., Briceno, L. D., Maciejewski, A. A., Siegel, H. J., Renner, T., Shestak, V., Ladd, J., Sutton, A. M., Janovy, D. L., Govindasamy, S., Alqudah, A., Dewri, R., Prakash, P., [Measuring](#)

the Robustness of Resource Allocations in a Stochastic Dynamic Environment. In: *International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 1–10, 2007.

Heterogeneous distributed computing systems often must operate in an environment where system parameters are subject to uncertainty. Robustness can be defined as the degree to which a system can function correctly in the presence of parameter values different from those assumed. We present a methodology for quantifying the robustness of resource allocations in a dynamic environment where task execution times are stochastic. The methodology is evaluated through measuring the robustness of three different resource allocation heuristics within the context of a stochastic dynamic environment. A Bayesian regression model is fit to the combined results of the three heuristics to demonstrate the correlation between the stochastic robustness metric and the presented performance metric. The correlation results demonstrated the significant potential of the stochastic robustness metric to predict the relative performance of the three heuristics given a common objective function.

- [50] Sutton, A. M., Whitley, D., Lunacek, M., Howe, A. E., [PSO and multi-funnel landscapes: how cooperation might limit exploration](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 75–82, 2006.

Particle Swarm Optimization (PSO) is a population-based optimization method in which search points employ a cooperative strategy to move toward one another. In this paper we show that PSO appears to work well on "singlefunnel" optimization functions. On more complex optimization problems, PSO tends to converge too quickly and then fail to make further progress. We contend that most benchmarks for PSO have classically been demonstrated on singlefunnel functions. However, in practice, optimization tasks are more complex and possess higher problem dimensionality. We present empirical results that support our conjecture that PSO performs well on single-funnel functions but tends to stagnate on more complicated landscapes.

- [51] Sutton, A. M., Howe, A. E., Whitley, L. D., [Spacetrack: Trading off Quality and Utilization in Oversubscribed Schedules](#). In: *International Conference on Automated Planning and Scheduling (ICAPS)*, pp. 430–433, 2006.

Many scheduling problems are posed as optimization problems where the goal is to find a feasible schedule that maximizes the utilization of some resource. In some domains it is also necessary to consider the quality of the resulting schedule. In most research these two quantities are independent. This paper introduces a real world problem in which radar tasks must be allocated to track objects in space. We explore the trade-off between off-line task resource utilization and a measure of task quality that correlates to whether tasks are actually successfully executed. We develop two general types of algorithms that differ in the way they reason about quality and explore the trade-off between high quality solutions and solutions with high resource utilization.