Painting with Evolutionary Algorithms: the Effects of Brush Stroke Sparsity

Karim Nasr
Department of Computer Science
Vrije Universiteit Amsterdam
The Netherlands
karimnasr93@gmail.com

Okke van Eck
Department of Computer Science
Vrije Universiteit Amsterdam
The Netherlands
okke.van.eck@gmail.com

Daan van den Berg
Department of Computer Science
Vrije Universiteit Amsterdam
The Netherlands
daan@yamasan.nl

***Abstract*—This paper investigates the use of hill climbing, simulated annealing, and tabu search algorithms for approximating target paintings with low number of brush strokes. The resulting paintings are obviously of lower resolution, but surprisingly enough, the distribution of brush strokes also changed. Whereas denser constellations have roughly uniform distributions of brush strokes, sparse constellations show clear brush stroke preference in non-trivial distributions. Furthermore, the algorithmic trajectory of just one painting consistently deviated from all the others. Its brush stroke sizes *shrunk* where all others grew. From an optimization perspective, it tells us that optimization metaheuristics can behave substantially different on sparse instances. Moreover, with a bit of fantasy and goodwill, the results can give an aspiring painter a bit of advise as well.**

***Keywords*-Evolutionary Algorithms; Evolutionary Art; Painting; Computational Creativity.**

I. Introduction

Computational creativity and evolutionary optimization have been merrily married for quite some time now. Having their own conference exactly on the intersection (EvoMUSART within EvoSTAR), it hosts a variety of creative endeavours [3], [4], but is by no means the only conference that currently supports the topic [5], [6], and the more rigorous approaches have even reached journals nowadays [7], [8]. So, optimization in computational creativity is on the rise, like everything in artificial intelligence, and an often seen avenue is the various approximation methods on classical paintings by means of brush strokes, transparent polygons, or other geometric shapes [9], [10].

Many of these studies select their target paintings from a fairly small set, and although this might be boring from a computationally-creative perspective, it is useful from an optimization perspective, as comparisons can more easily be made. This study is about sparse brush stroke constellations, and we indeed use a subset of earlier approximated paintings, namely seven 240×180 pixel target bitmaps in portrait or landscape orientation: *Mona Lisa* (1503) by Leonardo da Vinci, *The Starry Night* (1889) by Vincent van Gogh, *The Kiss* (1908) by Gustav Klimt, *Composition with Red, Yellow and Blue* (1930) by Piet Mondriaan, *The Persistence of Memory* (1931) by Salvador Dali, *Convergence* (1952) by Jackson Pollock, and the only known portrait of Leipzig-based composer Johann Sebastian Bach (1746) by Elias Gottlieb Haussmann (Figure 2 shows three portrait-oriented paintings from the set). Target bitmaps come from the public domain, are slightly cropped or rescaled if necessary, and are comprised of 8-bit RGB-pixels.

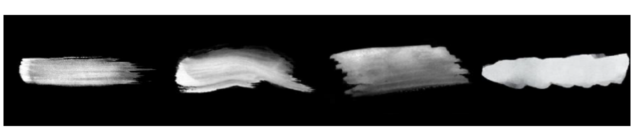

Figure 1. Left to right: brush types 1, 2, 3 and 4 as used in this study. These often used images come from the open repository of Anastasia Opara [1].

Much like earlier work on brush strokes (e.g. [1], [11], [12]), we will approximate these paintings by optimizing brush strokes' parameters. The novelty here is, that we deliberately use *sparse* sets of brush strokes, causing some unexpected behaviour for the trajectories in general, and one painting in particular. According to the explanation put forth in the conclusion of the text, it might serve as a stark reminder that when problem instances get sparser, the contextual background gains a lot of influence, which might be relevant for optimization problems in a wider sense In short, the main research question of this paper is:

- How does brush stroke sparsity influence the characteristics (the type and size of the brush strokes) of the approximated paintings when using hill climbing, simulated annealing and tabu search?

The rest of this paper is organized as follows: in Section 2, we will describe the experiments, giving the exact parameters, the measure of performance (being the Mean Squared Error), and describing the hill climber, simulated annealing and tabu search algorithms. In Section 3, the results of these algorithms' runs will be discussed for the different paintings under the settings of the algorithms described in Section 2. Finally in Section 4, we will discuss *why* the results are the way they are, and we conjecture a bit on the results within the broader context of optimization problems.

 ISBN: 978-1-68558-213-5

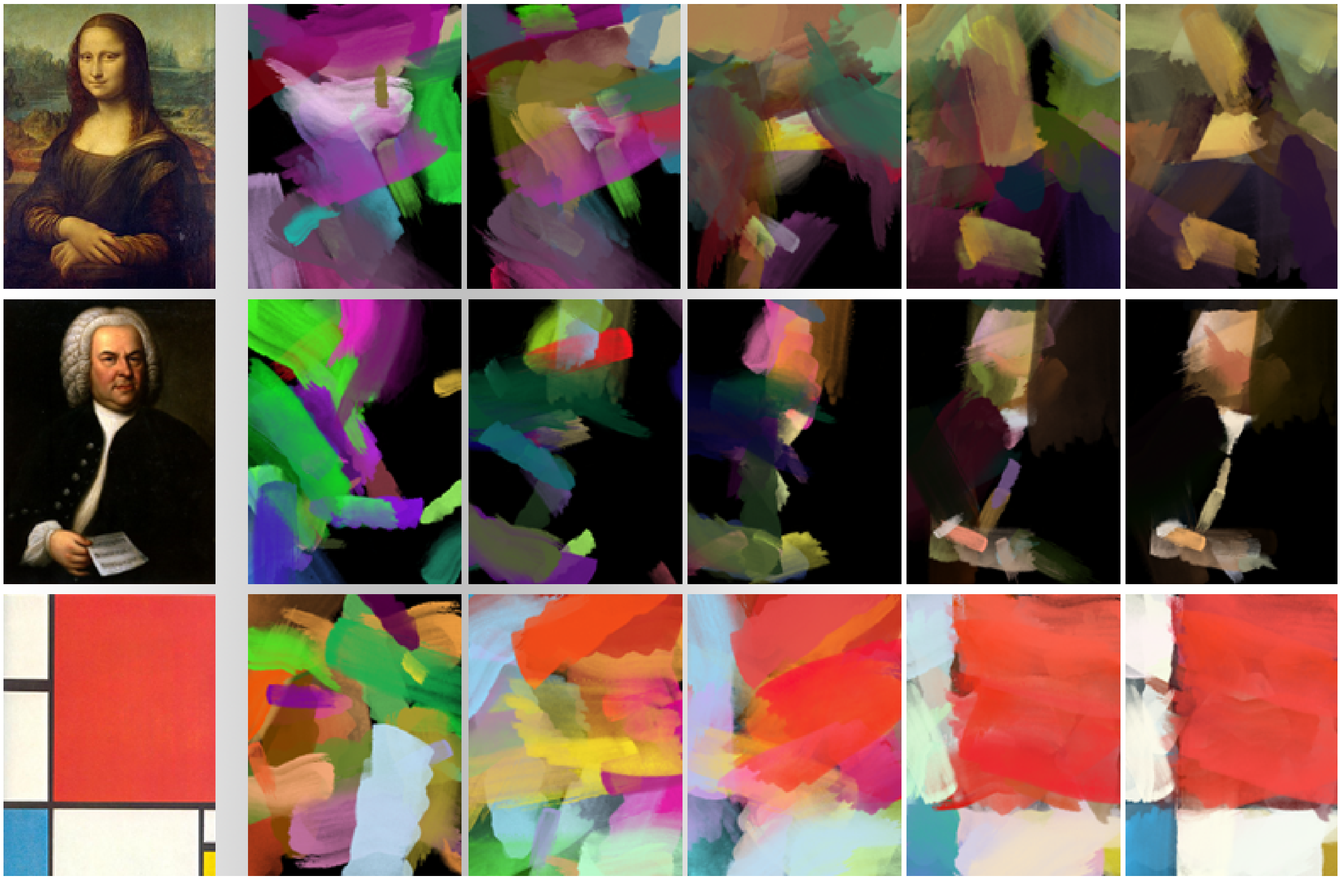

Figure 2. Starting from a black canvas with a constellation of 25 randomly initialized brush stroke, three evolutionary algorithms (hill climbing, simulated annealing, tabu search) iteratively improved the constellation for 1 million evaluations. These three paintings are representative for seven paintings in total, all of which can be found in our repository [2].

II. EXPERIMENTS

For this study, each run starts with a black canvas (180×240 pixels) with 25 brush strokes, which all have their parameters type, colour, size (0.1 to 0.7, corresponding to 30×30 to 210×210 pixels) position (between [0,0] and [180×240]) and rotation (-180 to +180) uniformly randomly assigned. The brush types (Figure 1) for this project are taken from Opara's GitHub [1]. Every brush stroke image has its unique opacity properties, and rendering of the constellation is done by drawing stroke by stroke, pixel by pixel on the canvas, which is a relatively costly procedure. After rendering, the proximity to the target painting is assessed by calculating its pixel-by-pixel Mean Squared Error (MSE):

$$\sum_{i=1}^{180\cdot 240\cdot 3} \frac{(Rendered_i - Target_i)^2}{180 \cdot 240} \tag{1}$$

which is identical to earlier studies [9], [11], [12]. The objective for the evolutionary algorithms is to minimize the MSE, and as such it can be viewed as a purely combinatorial optimization study with computationally creative byproducts, as much as a study in computational creativity alone.

After initialization, the algorithm at hand randomly mutates one parameter (size, position, color, rotation, brush type, and drawing index) of one brush stroke by assigning it a new random value from its domain. In this study, we deploy three evolutionary algorithms. First, a very basic **hill climber** works by applying a single mutation to the incumbent brush stroke constellation, and retains it if the new constellation's MSE is lower, or reverts the mutation otherwise. Second, **simulated annealing** works identical to the hill climber, but besides accepting improving mutations, it can also accept mutations that lead to a *worse* MSE. The probability thereof depends on two factors: the magnitude of the deterioration and the 'temperature' parameter as

$$P(accept) = e^{-(\frac{\Delta MSE}{temp})} \tag{2}$$

in which e is Euler's number ($\approx$ 2.718 ...), and the temperature is a constant that declines with the number of iterations i and a constant $c = 1$ as

$$temp = \frac{c}{log(i)}, \tag{3}$$

which is know as the Geman&Geman cooling schedule, the only cooling schedule for which convergence to a global

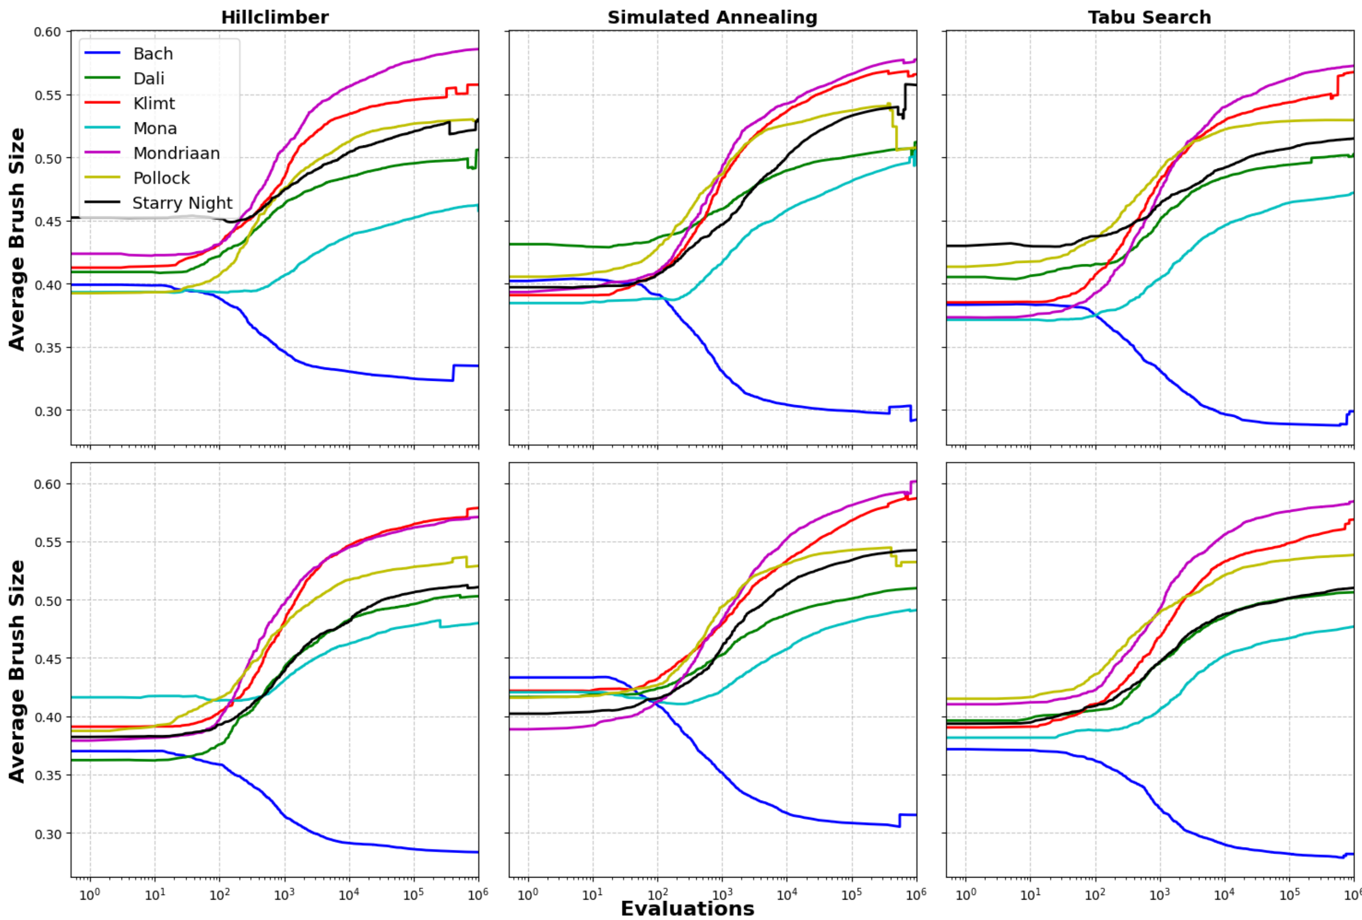

Figure 3. As can be expected (retrospectively) for sparse constellations, all algorithms increase the average brush sizes enormously, both when 3 types are used (top row) or when 4 types are used (bottom row). But not for all paintings; Bach is the odd man out. This might have to do with its unmutably black canvas background, whose influence increases significantly for sparse constellations.

optimum is proven (but practically unattainable) [12]–[14]. The third algorithm of this short study, **tabu search**, keeps a 'tabu list' in which the last 50 evaluated constellations are retained. It mutates the incumbent constellation, accepts the new constellation only if better *and not in its tabu list* [15].

With these algorithms, we ran two experiments: one with brush types 1, 2 and 3, and one with brush types 1, 2, 3 and 4. In both experiments, all three algorithms did 5 runs of 1,000,000 iterations (eq.: 'evaluations') on all 7 paintings, amounting to 210,000,000 evaluations in total. Computation was done on the Snellius supercomputer service provided at the VU University and all source code, experimental results and extra figures are publicly available [2].

III. RESULTS

When it comes to results, the surprises were not in the convergence of the algorithms. The end MSEs were a little less optimal then in Dijkzeul et al.'s earlier study [11], which can easily be explained by the higher number of brush strokes in their study. The convergence patterns were even similar, with the unstructured Pollock painting's approximation remaining relatively high in its MSE values, and constellations for Mondriaan's painting making by far the largest MSE improvement for all three algorithms. So far nothing new, except for the observation that tabu search and simulated annealing performed roughly equal.

The real surprise was in the brush stroke distributions (Figure 4). Whereas Dijkzeul et al., with their far higher number of brush strokes achieved an almost uniform distribution for all types, our algorithms showed a clear preference for brush stroke type 2 in the 3-type experiment (44% to 48% averaged over all runs and all paintings, against an expected 33.3%). For the 4-stroke experiment however, brush stroke type 3 was preferred (40% to 44% averaged over all runs and all paintings, against an expected 25%). In short: adding a fourth brush stroke type favoured use of the third, counterintuitively enough.

The second surprise came from Bach. As it turns out, the dominant pattern shows a significant increase in brush size in the optimized constellations for all algorithms and all paintings – *except* Haussmann's Bach (Figure 3). The only explanation we can think of is that the Bach painting holds large swaths of black, and the algorithms unintendedly sacrificed coverage of the black area in favour of detailing smaller brush strokes

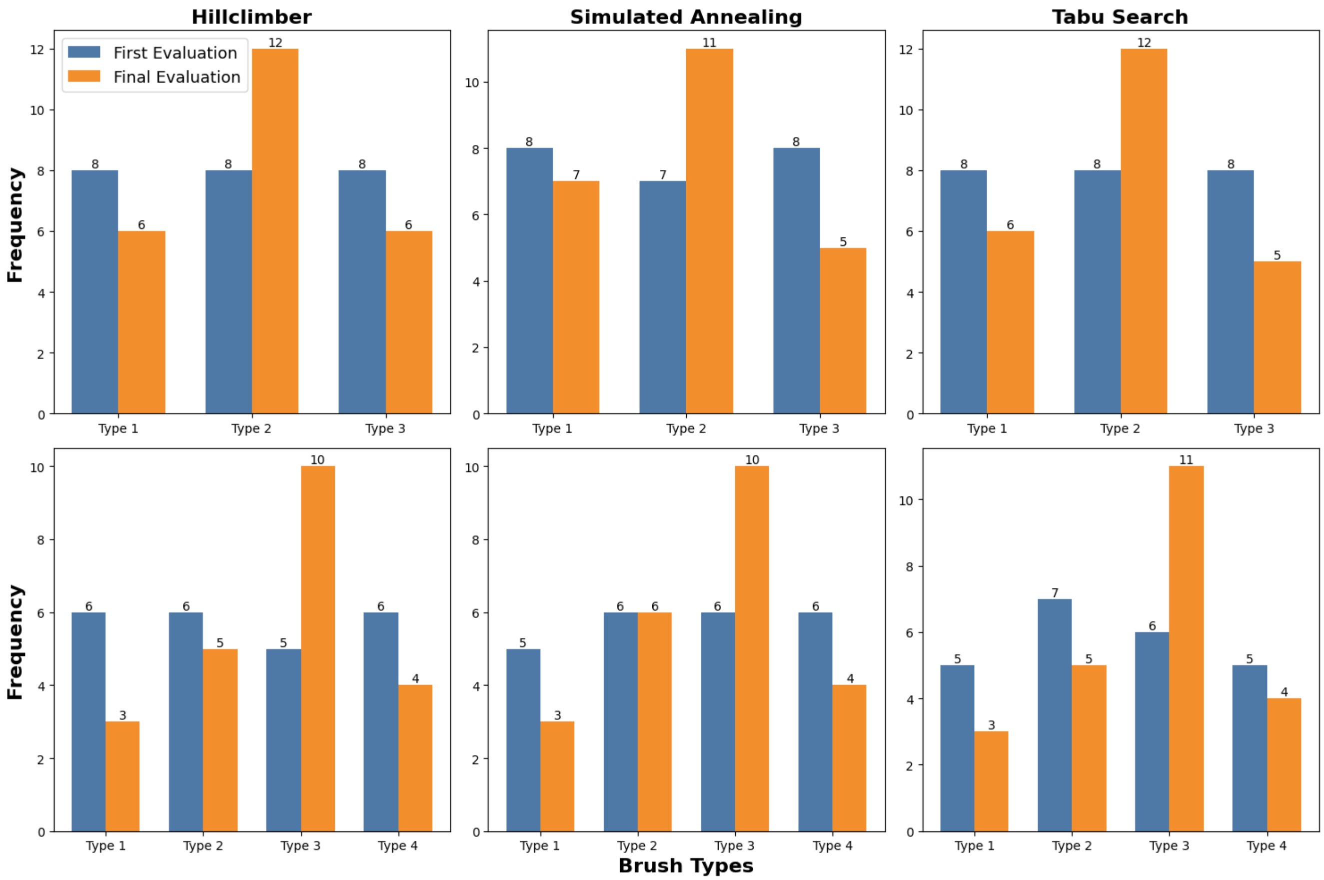

Figure 4. Evolution of brush distributions for three brush types (top row) tend towards preferring brush type 2. When a fourth brush type is introduced however (bottom row), brush type 3 becomes the preferred type. This shift is consistent throughout the experiments, but remains unexplained for now.

grasping the details of Bach's face, hair and clothes. But contrary to this explanation, Bach's MSE was not outstandingly good, so a new experiment, perhaps enabling a mutation of the canvas background colour, or cluster analysis of brush stroke pixels (rather than just locations) could shed some light on this question. However for now, it was somewhat surprising that fewer brush stroke problems can elicit substantially different patterns of convergence, simply because the role of the canvas' backround might relatively increase.

IV. Conclusion & Discussion

Looking back through the eye of an optimization specialist, a particular quality concern is that none of the three algorithms hill climber, simulated annealing or tabu search outperformed the others. An explanation might be given from two directions, and one might not exclude the other.

First, the parameterization of the algorithms might have been very poor. For simulated annealing, the Geman&Geman cooling schedule is indeed the only one (so far) that has some proof of optimality, and the $c = 1$ parameter makes a lot of sense, but also not. In literature, the value of c should be equal to "the highest energy barrier that possibly needs to be crossed" [13], [14]. Translated into optimization terms, it means c should the quantify depth of the deepest local minimum, which is very unlikely to be just 1. Or is it? In a very similar study by Dahmani et al, it was the best performing out of 9 cooling schedules, which justifies the choice somewhat.

For tabu search, a list of 50 taboo instances might be very short. The combinatorial state space contains more than 10^{284} states, even for just 25 brush strokes, and the number of possible mutations from any state exceeds 10^{10}. This means that despite the ominous birthday paradox, the chances of encountering the exact same state is well below 1%, rendering the tabu list effectively useless.

Eventually, all these algorithmic assessments tie into the age old question: how convex is the optimization landscape? Possibly not very much so, or on a very small scale (small undulations on a huge valley) but we just do not know. An explorative landscape analysis could theroretically shed some light on it, but even these studies have state spaces not nearly as big as ours [16], [17].

One interesting idea comes from the possible generality of these findings. In HP protein folding, it is known that the required number of random samples to obtain a valid conformation follow a straight line on a log scale, *except* for the very small ('sparse') instances [18], [19]. In the optimization variant of the partition problem, the distribution of bits ('background') of the set appears to have a critical

 ISBN: 978-1-68558-213-5

influence on the distribution of solutions across the fitness landscape [20]–[22]. For the traveling tournament problem, even non-optimal but just *valid* solutions appear to be sparse [23]–[26]. As such, paintings-from-paintstrokes might not be so different from from other optimization problems. But like in the realm of real paintings, those tiny details matter. They ultimately separate the very good paintings from the global optimum of a masterpiece.

References

[1] A. Opara, *Anastasia opara's "Genetic Drawing"*, https://github.com/anopara/genetic-drawing. Last accessed 9th of September 2024, 2022.

[2] K. Nasr, *Repo with publicly accessible material: Https://github.com/karimn7/painting-generation*, Last accessed 9th of September 2024, 2024.

[3] Y. Tian and D. Ha, "Modern evolution strategies for creativity: Fitting concrete images and abstract concepts", in *International conference on computational intelligence in music, sound, art and design (part of evostar)*, Springer, 2022, pp. 275–291.

[4] Y. Tian, "Evolving three dimension (3d) abstract art: Fitting concepts by language", *arXiv preprint arXiv:2304.12932*, 2023.

[5] P. Williams, K. Li, and G. Min, "A surrogate assisted evolutionary strategy for image approximation by density-ratio estimation", in *2023 IEEE Congress on Evolutionary Computation (CEC)*, IEEE, 2023, pp. 1–8.

[6] J. Berg *et al.*, "Evolved art with transparent, overlapping, and geometric shapes", in *Nordic Artificial Intelligence Research and Development: Third Symposium of the Norwegian AI Society, NAIS 2019, Trondheim, Norway, May 27–28, 2019, Proceedings 3*, Springer, 2019, pp. 3–15.

[7] A. Neumann, B. Alexander, and F. Neumann, "Evolutionary image transition and painting using random walks", *Evolutionary computation*, vol. 28, no. 4, pp. 643–675, 2020.

[8] H. Richter, "Designing color symmetry in stigmergic art", *Mathematics*, vol. 9, no. 16, p. 1882, 2021.

[9] M. Paauw and D. Van den Berg, "Paintings, polygons and plant propagation", in *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*, Springer, 2019, pp. 84–97.

[10] F. Nolte, A. Melnik, and H. Ritter, "Stroke-based rendering: From heuristics to deep learning", *arXiv preprint arXiv:2302.00595*, 2022.

[14] Y. Nourani and B. Andresen, "A comparison of simulated annealing cooling strategies.", *Journal of Physics A: Mathematical and General*, vol. 31, pp. 8373–8385, 1998. DOI: 10.1088/0305-4470/31/41/011.

[15] S. Verel, S. Thomson, and O. Rifki, "Where the really hard quadratic assignment problems are: The qap-sat instances", *EvoSTAR 2024*, 2024.

[11] D. Dijkzeul, N. Brouwer, I. Pijning, L. Koppenhol, and D. Van den Berg, "Painting with evolutionary algorithms", in *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*, Springer, 2022, pp. 52–67.

[12] R. Dahmani, S. Boogmans, A. Meijs, and D. Van den Berg, "Paintings-from-polygons: Simulated annealing", in *International Conference on Computational Creativity (ICCC'20)*, 2020.

[13] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images", *IEEE Transactions on pattern analysis and machine intelligence*, no. 6, pp. 721–741, 1984.

[16] S. L. Thomson and G. Ochoa, "On funnel depths and acceptance criteria in stochastic local search", in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 287–295.

[17] S. L. Thomson, G. Ochoa, and S. Verel, "The fractal geometry of fitness landscapes at the local optima level", *Natural Computing*, vol. 21, no. 2, pp. 317–333, 2022.

[18] J. Koutstaal *et al.*, "Frequency fitness assignment for untangling proteins in hp", *EvoSTAR 2024*, 2024.

[19] R. Jansen *et al.*, "Can hp-protein folding be solved with genetic algorithms? maybe not.", in *Proceedings of the 15th International Joint Conference on Computational Intelligence ECTA - Volume 1*, 2023, pp. 131–140. DOI: 10 . 5220 / 0012248500003595.

[20] D. Van den Berg and P. Adriaans, "Subset Sum and the Distribution of Information", in *Proceedings of the 13th International Joint Conference on Computational Intelligence*, Jan. 2021, pp. 134–140. DOI: 10.5220/0010673200003063.

[21] N. Sazhinov, R. Horn, P. Adriaans, and D. van den Berg, "The partition problem, and how the distribution of input bits affects the solving process", in *Proceedings of the 15th International Joint Conference on Computational Intelligence*, SCITEPRESS - Science and Technology Publications, 2023. DOI: 10.5220/0012143600003595.

[22] R. Horn, S. L. Thomson, D. van den Berg, and P. Adriaans, "The easiest hard problem: Now even easier", 2024.

[23] K. Verduin, S. L. Thomson, and D. van den Berg, "Too constrained for genetic algorithms. too hard for evolutionary computing. the traveling tournament problem.", 2023.

[24] K. Verduin, T. Weise, and D. van den Berg, "Why is the traveling tournament problem not solved with genetic algorithms?", 2023.

[25] K. Verduin *et al.*, "The traveling tournament problem: Rows-first versus columns-first", *ICEIS 2024*, pp. 447–455, 2024.

[26] C. Xiang, Z. Wu, D. van den Berg, and T. Weise, "Randomized local search vs. nsga-ii vs. frequency fitness assignment on the traveling tournament problem", *ECTA 2024*, (to appear).