Cost Versus Distance In the Traveling Salesman Problem*

Kenneth D. Boese

UCLA Computer Science Dept., Los Angeles, CA 90024-1596 USA

Abstract

This paper studies the distribution of good solutions for the traveling salesman problem (TSP) on a well-known 532-city instance that has been solved optimally by Padberg and Rinaldi [16]. For each of five local search heuristics, solutions are obtained from 2,500 different random starting points. Comparisons of these solutions show that lower-cost solutions have a strong tendency to be both closer to the optimal tour and closer to other good solutions. (Distance between two solutions is defined in terms of the number of edges they have in common.) These results support the conjecture of Boese, Kahng and Muddu [3] that the solution spaces of TSP instances have a "globally convex" or "big valley" character. This observation was used by [3] to motivate a new multi-start strategy for global optimization called *Adaptive Multi-Start* (AMS).

1 Introduction

Local search is probably the most successful approach to finding heuristic solutions to combinatorial global optimization problems. In global optimization, objective is to find a solution s in solution space S which minizes a cost function $f(s)$ defined on S. Local search moves iteratively from a solution s_i to some "nearby" solution s_{i+1} in the neighborhood of s_i, $N(s_i)$. The definition of neighborhoods $N(s) \subset S$ for each $s \in S$, together with solution costs $f(s)$, give rise to a *cost surface* for the particular problem instance. Understanding this cost surface can help both to explain the success of previous heuristics (e.g., simulated annealing) and to motivate new, more effective heuristics (e.g., multi-start strategies or better annealing schedules). Our results indicate that cost surfaces for the traveling salesman problem (TSP) exhibit a "globally convex" [6] or what we call a "big valley" structure. Figure 1 gives an intuitive picture of the big valley, in which the set of local minima appears convex with one central global minimum.

In this paper, we discuss experimental results obtained by running five different local search heuristics many times on a single, well-known TSP instance called "ATT532". ATT532 was compiled by AT&T Bell Laboratories and is based on locations of 532 cities in the continental United States. It has been used in a

*This work was performed under support from the UCLA Dissertation Year Fellowship.

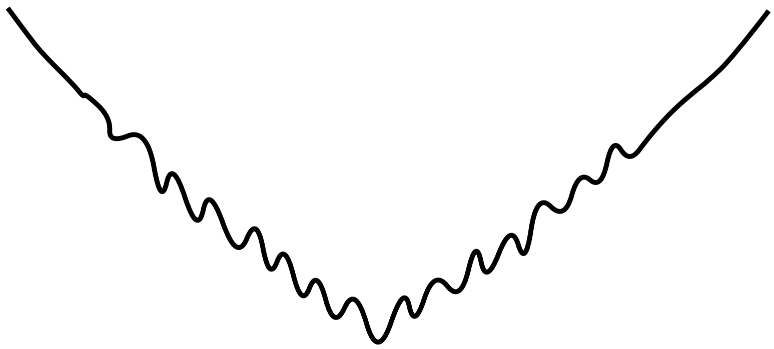

Figure 1: Intuitive picture of the "big valley" solution space structure.

number of other studies including [12, 13] and was solved to optimality by Padberg and Rinaldi in 1987 [16]. We have chosen this instance because (i) it represents a real-world geometric TSP instance; (ii) it is large enough to prove difficult for most heuristics to solve optimally; and (iii) its optimal tour is known, allowing us to compare heuristic solutions to the optimal solution.

In [3] we presented similar results for two random geometric TSP instances with 100 and 500 cities. The plots in [3] were over local minima obtained by a randomized implementation of the 2-Opt local search heuristic. The current study augments [3] by using four additional local search heuristics for an instance with a known globally optimal tour. We also note that other authors such as Mühlenbein et al. [15] and Sourlas [18] have used similar plots to justify their heuristics. However, our results in [3] and in this report (i) involve more solutions and use better local search heuristics; (ii) compare mean distances to other solutions, in addition to distances to the optimal solution; (iii) lead to the observation that the optimal solution is more central among good solutions; and (iv) motivate a different heuristic (Adaptive Multi-Start or AMS) for global optimization.

2 Preliminaries

Suppose that t_1 and t_2 are TSP tours over the same set of n cities. We define the *distance* $d(t_1, t_2)$ to be n minus the number of edges contained in both t_1 and t_2. This measure of distance has been used in a number of previous studies of TSP solution spaces (e.g., [9, 14, 18]). In [3], we showed that this distance approximates the number of 2-Opt operations required to transform one tour into another, to within a factor of two.[1]

Each of the heuristics used in this report is based on the k-Opt local search strategy, which iteratively transforms tours into lower-cost tours by performing a sequence of k-Opt moves. Each k-Opt move replaces k edges in a tour with k new edges to form a new tour. We believe that $d(t_1, t_2)$ is closely related to the k-Opt "distance" between tours for general k, in addition to $k = 2$. Thus, we believe $d(t_1, t_2)$ is a good measure of proximity between solutions produced by k-Opt-based heuristics. The five local search heuristics we study include:

1. **Random 2-Opt**. At each iteration, we test all $\Theta(n^2)$ possible 2-Opt moves in random order, until an improving move is found or the current tour is shown to be a local minimum.

2. **Fast 2-Opt**. At each iteration, we perform the 2-Opt search proposed by Bentley [2]. This reduces the time complexity of 2-Opt from $\Omega(n^2)$ for Random 2-Opt to approximately $\Theta(n \log n)$ on average.

3. **Fast 3-Opt**. We follow Bentley's [2] efficient implementation of the 3-Opt heuristic originally described by Lin [10].[2]

4. **Lin-Kernighan**. We have implemented, as accurately and completely as possible, Lin and Kernighan's [11] variation of k-Opt that searches a small but effective subset of all k-Opt moves for $2 \leq k \leq n$.

5. **Large-Step Markov Chains (LSMC)** Finally, we use the heuristic of Martin et al. [12] [13] which iteratively applies 3-Opt to find a sequence of local minima; the starting tour for each 3-Opt descent is obtained by applying a random 4-Opt move to the most recent 3-Opt local minimum. Our implementation returns the best tour visited after a sequence of 1,000 3-Opt descents.[3]

We include Random 2-Opt to provide continuity with our original paper [3]. Interestingly, Random 2-Opt returns solutions with significantly higher cost than those obtained by Fast 2-Opt. Heuristics 2 through 4 have been compared to other heuristics by Johnson [7] and Bentley [1] and appear to be among the most effective TSP heuristics. For example, 3-Opt and Lin-Kernighan return tours even better than simulated annealing

[1]The same result was proved independently by Kececioglu and Sankoff [8] in the context of computing the number of chromosome inversions required to evolve one organism into another.

[2]Note that our implementations of Fast 2-Opt and Fast 3-Opt differ slightly from Bentley's in that we precompute a "nearest neighbor" matrix of the 25 closest cities to each city in the instance.

[3]Note that before applying a random 4-Opt move, LSMC sometimes returns to the previous 3-Opt local minimum if the current one has higher cost. This decision is based on a Metropolis criterion for which we have found a good temperature to be 10.0 for this instance.

when applied in a multi-start regime. Heuristic 5 is perhaps the *best* TSP heuristic available for returning solutions very close to optimal, although it does require more computation time than the other heuristics considered here.

3 Experimental Results

We ran each of the heuristics 2,500 times from random starting tours. We then computed the distance of each solution to the optimal tour and to each of the other solutions found by the same heuristic. Our results are plotted in Figures 2 through 6 and summarized in Table 1. Our experiments resulted in 2,500 unique tours for each of the heuristics except LSMC, which found 1,884 unique tours. LSMC also found an optimal tour six times, four times finding the tour published in [16] and twice finding a tour with equal cost (27,686) at distance two from the published optimal. None of the other heuristics found an optimal tour in any of its 2,500 runs. Our results show a very clear relationship between cost and distance: better heuristic tours are both closer to the optimal tour and to other heuristic tours. Moreover, the optimal tour is located at a more central position within the subspace of good solutions: the optimal tour is closer on average to the heuristic tours than are most of the heuristic tours themselves. This suggests a "globally convex" [6] or "big valley" structure for the TSP solution space, with the optimal solution near the center of a single valley of low-cost solutions.

Algorithm	Ave. Cost: Percent Above Optimal	Ave. Distance to Optimal	Max. Distance to Optimal	Fraction of Solution Space	Ave. Mean Distance to Other Solutions	Ave. Running Time (seconds)
Random 2-Opt	11.8	196	233	10^{-569}	232	11.5
Fast 2-Opt	6.7	152	194	10^{-670}	176	0.28
Fast 3-Opt	2.3	110	153	10^{-779}	129	0.27
Lin-Kernighan	1.2	96	142	10^{-809}	110	6.2
LSMC (3-Opt)	0.14	59	97	10^{-935}	65	33.8

Table 1: Summary of solutions from 2,500 runs each of five different TSP heuristics on ATT532. All heuristics except LSMC found 2,500 unique tours; LSMC found 1,884 unique tours. Running times are for an HP Apollo 9000-735.

The relationship between cost and distance is most striking for Random 2-Opt and for Lin-Kernighan, in Figures 2 and 5. For Fast 2-Opt and Fast 3-Opt, the relationship is somewhat obscured by a relatively small number of local minima with high cost. (Note that high-cost local minima are ignored by our AMS heuristic in [3].) For LSMC, the relationship appears to be quite strong again, although there is perhaps a second

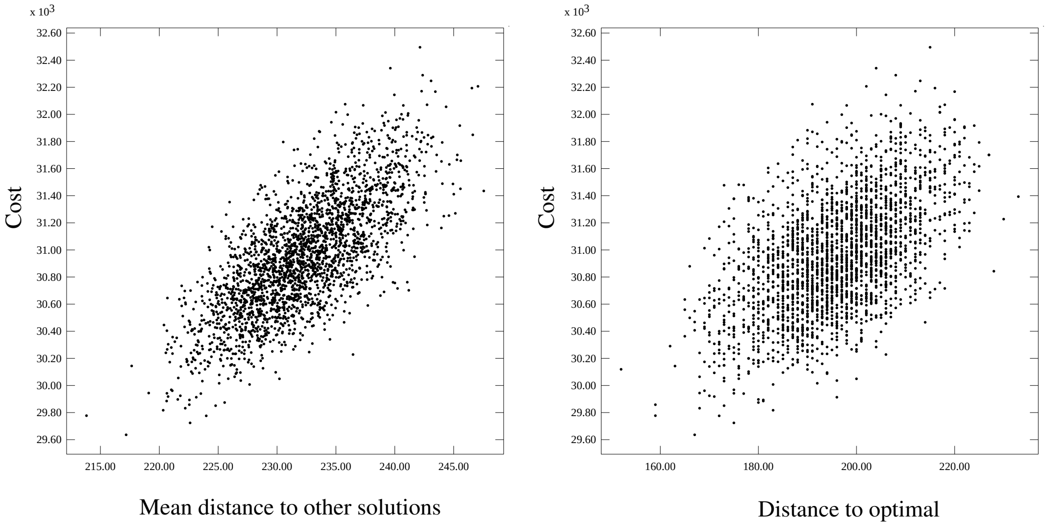

(a) (b)

Figure 2: 2,500 Random 2-Opt local minima for ATT532. Tour cost (vertical axis) is plotted against (a) mean distance to the 2,499 other local minima and (b) distance to the global minimum.

valley at a distance between 60 and 80 from the optimal tour.

Our results indicate that studies of TSP solution spaces should concentrate on a very small subspace. Define a *ball* $b(t, k)$ to be the subset of tours within distance k of a tour t. From the third column in Table 1, we see that all the tours found by the five heuristics are contained in $b(t_{opt}, 233)$. In Appendix B of [3], we described how to calculate the number of tours within $b(t, k)$ for any $k \leq n$. We used this calculation to obtain the fourth column of Table 1, which gives the fraction of the solution space contained in a ball centered at the optimal tour and containing all tours obtained by the heuristic. For instance, all solutions found by Fast 2-Opt lie within a ball containing a fraction $1/10^{670}$ of the solution space, while all of the LSMC solutions lie in $1/10^{935}$ of the solution space.[4]

Finally, in Table 2 we analyze the relationship between cost and distance more formally. For each of the five heuristics, we compute the correlations between cost and the distance to optimal and also the correlations between cost and the mean distance to other solutions. The table confirms that the relationship between cost and distance is strongest for the Random 2-Opt and Lin-Kernighan heuristics. The t-Statistics reported in Table 2 indicate whether each correlation is statistically significant (i.e., could not occur merely by chance):

[4] Because there are $(531!)/2 \approx 10^{1218}$ possible tours, these balls contain approximately 10^{648} and 10^{283} tours, respectively.

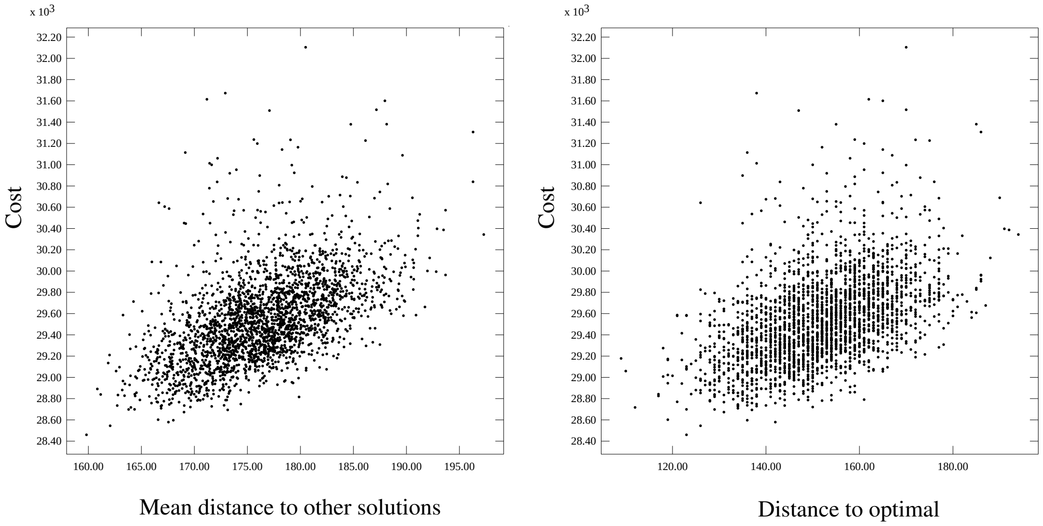

(a) (b)

Figure 3: 2,500 Fast 2-Opt local minima for ATT532.

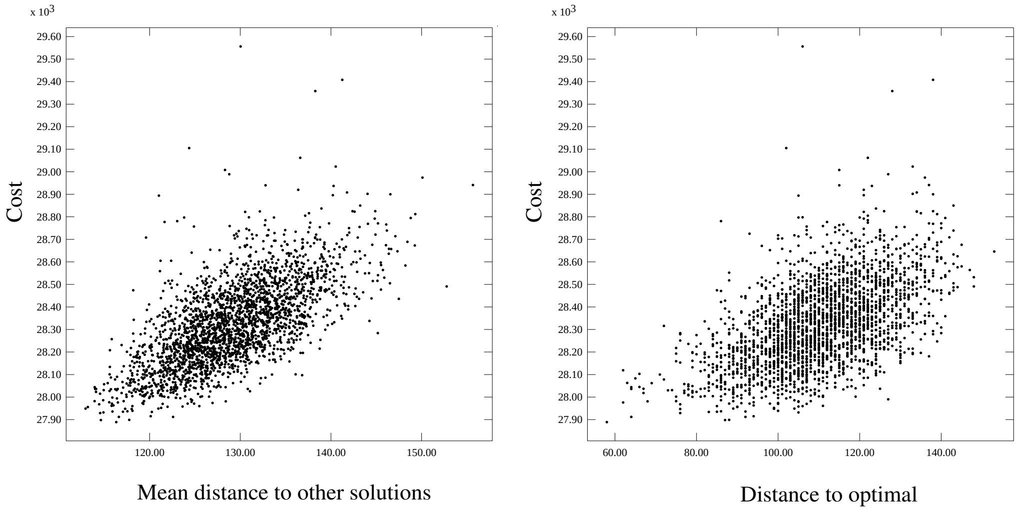

(a) (b)

Figure 4: 2,500 Fast 3-Opt local minima for ATT532.

a value of approximately 2.0 or greater indicates a correlation significant at the 95% confidence level, and a value of 2.6 or greater indicates significance at the 99% confidence level [17]. With t-Statistics ranging from 19 to 54, the correlations between distance and cost are highly significant statistically.

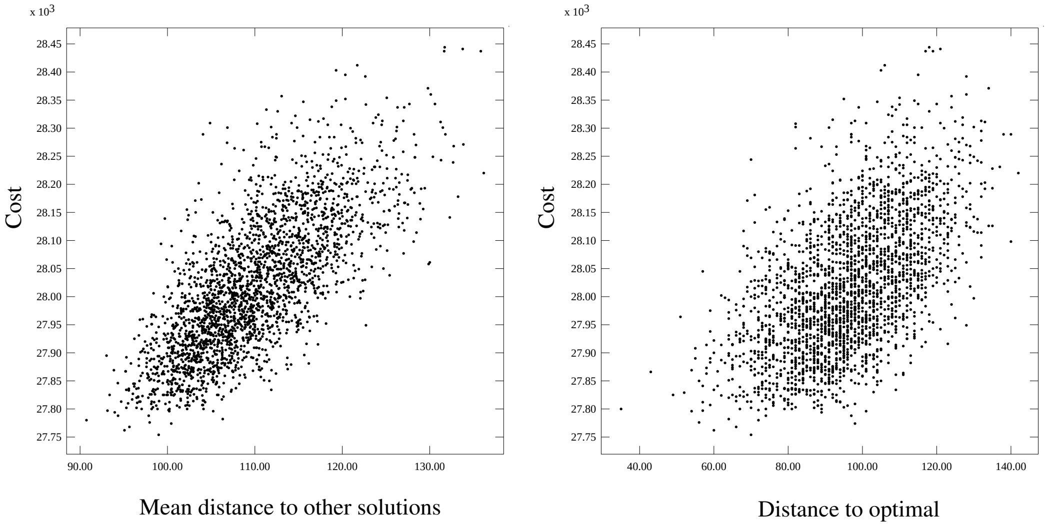

(a) (b)

Figure 5: 2,500 Lin-Kernighan local minima for ATT532.

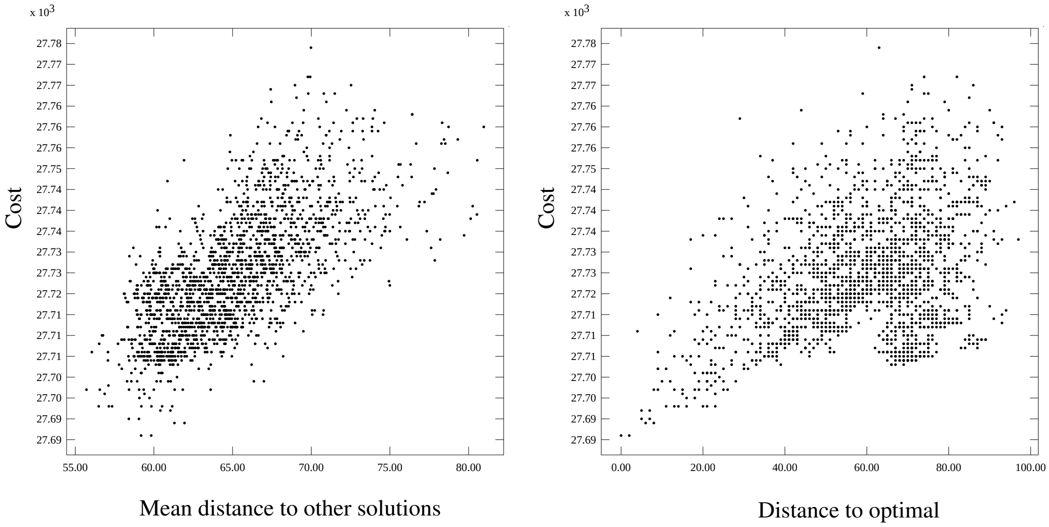

(a) (b)

Figure 6: 1,884 unique solutions found by Large-Step Markov Chains (LSMC) in 2,500 runs.

4 Continuing Research

Our continuing research has produced similar plots for a number of other combinatorial optimization problems, including circuit/graph partitioning, satisfiability, number partitioning, and job shop scheduling. In [3] we also presented two plots for random graph partitioning instances, which again showed a strong relationship

	Mean Dist. to Other Solutions		Distance to Optimal	
Algorithm	Correlation	T-Statistic	Correlation	T-Statistic
Random 2-Opt	0.73	54	0.55	32
Fast 2-Opt	0.53	31	0.47	27
Fast 3-Opt	0.66	44	0.54	32
Lin-Kernighan	0.73	54	0.57	34
LSMC (3-Opt)	0.69	41	0.40	19

Table 2: Correlations between distance and cost for the five heuristics applied to ATT532. (Based on the unique minima resulting from 2,500 runs of each heuristic.)

between cost and distance. However, Hagen and Kahng [5] have shown for circuit partitioning at least, that this relationship deteriorates for lower-cost solutions (i.e., those produced by more powerful heuristics such as Fiduccia-Mattheyses [4]). In other problem formulations we also find weaker cost-distance relationships than in the TSP, although in some of them (e.g., job shop scheduling) the relationship becomes more apparent when we use better heuristics. Finally, we are testing multi-start heuristics for the TSP that constrain edges in later descents if they are common to all of the best tours in earlier descents. This strategy is very similar to a multi-start approach suggested by Lin and Kernighan in their 1973 paper, except that we now "freeze" edges common to the *best* previous solutions (cf. [5]) rather than only the edges common to *all* previous solutions.

References

[1] J. L. Bentley, "Experiments on Traveling Salesman Heuristics" in *First Annual ACM-SIAM Symposium on Discrete Algorithms* (January 1990), pp. 187-197.

[2] J. L. Bentley, "Fast Algorithms for Geometric Traveling Salesman Problems", *ORSA Journal on Computing* 4 (4) (Fall 1992), pp. 387-410.

[3] K. D. Boese, A. B. Kahng and S. Muddu, "A New Adaptive Multi-Start Technique for Combinatorial Global Optimizations", *Operations Research Letters*, 16(2), Sept. 1994, pp. 101-113.

[4] C. M. Fiduccia and R. M. Mattheyses, "A Linear-Time Heuristic for Improving Network Partitions", in *ACM IEEE Nineteenth Design Automation Conference*, June 1982, pp. 175-181.

[5] L. Hagen and A. B. Kahng, "Combining Problem Reduction and Adaptive Multi-Start: A New Technique For Superior Iterative Partitioning", to appear in *IEEE Trans. Computer Aided Design*, 1995.

[6] T. C. Hu, V. Klee and D. Larman, "Optimization of Globally Convex Functions", *SIAM J. on Control and Optimization* 27(5), 1989, pp. 1026-1047.

[7] D. S. Johnson, "Local Optimization and the Traveling Salesman Problem", in *Proceedings of the 17th International Colloquium on Automata, Languages and Programming*, July 1990, pp. 446-460.

[8] J. Kececioglu and D. Sankoff, "Exact and Approximation Algorithms for the Inversion Distance Between Two Chromosomes", in *Proceedings of the 4th Annual Symposium on Combinatorial Pattern Matching*, July 1993, pp. 87-105.

[9] S. Kirkpatrick and G. Toulouse, "Configuration Space Analysis of Traveling Salesman Problems", *Journal de Physique* 46, 1985, pp. 1277-1292.

[10] S. Lin, "Computer Solutions of the Traveling Salesman Problem", *Bell System Technical Journal* 44, 1965, pp. 2245-2269.

[11] S. Lin and B. W. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem", *Operations Research* 31, 1973, pp. 498-516.

[12] O. Martin, S. W. Otto and E. W. Felten, "Large-Step Markov Chains For the Traveling Salesman Problem" *Complex Systems* 5(3), June 1991, pp. 299-326.

[13] O. Martin, S. W. Otto, and E. W. Felten, "Large-Step Markov Chains for the TSP Incorporating Local Search Heuristics", *Operations Res. Letters* 11, 1992, pp. 219-224.

[14] M. Mezard and G. Parisi, "A Replica Analysis of the Travelling Salesman Problem", *Journal de Physique* 47, 1986, 1285-1296.

[15] H. Mühlenbein, M. Georges-Schleuter, and O. Krämer, "Evolution Algorithms in Combinatorial Optimization," *Parallel Computing* 7, 1988, pp. 65–85.

[16] M. Padberg and G. Rinaldi, "Optimization of a 532-city symmetric traveling salesman problem by branch and cut", *Operations Res. Letters* 6, 1987, pp. 1-7.

[17] S. M. Ross, *Introduction to Probability and Statistics for Engineers and Scientists*, (Wiley, New York, 1987).

[18] N. Sourlas, "Statistical Mechanics and the Travelling Salesman Problem", *Europhysics Letters* 2(12), 1986, pp. 919-923.