

# Anytime Behavior of Inexact TSP Solvers and Perspectives for Automated Algorithm Selection

Jakob Bossek  
*Optimisation and Logistics*  
*The University of Adelaide*  
Adelaide, Australia  
jakob.bossek@adelaide.edu.au

Pascal Kerschke  
*Information Systems and Statistics*  
*University of Münster*  
Münster, Germany  
kerschke@uni-muenster.de

Heike Trautmann  
*Information Systems and Statistics*  
*University of Münster*  
Münster, Germany  
trautmann@wi.uni-muenster.de

**Abstract**—The Traveling-Salesperson-Problem (TSP) is arguably one of the best-known  $\mathcal{NP}$ -hard combinatorial optimization problems. The two sophisticated heuristic solvers LKH and EAX and respective (restart) variants manage to calculate close-to optimal or even optimal solutions, also for large instances with several thousand nodes in reasonable time. In this work we extend existing benchmarking studies by addressing anytime behaviour of inexact TSP solvers based on empirical runtime distributions leading to an increased understanding of solver behaviour and the respective relation to problem hardness. It turns out that performance ranking of solvers is highly dependent on the focused approximation quality. Insights on intersection points of performances offer huge potential for the construction of hybridized solvers depending on instance features. Moreover, instance features tailored to anytime performance and corresponding performance indicators will highly improve automated algorithm selection models by including comprehensive information on solver quality.

**Index Terms**—anytime behavior, traveling salesperson problem, automated algorithm selection, performance assessment, hybridization

## I. INTRODUCTION

The Traveling-Salesperson-Problem (TSP) is an intriguing fundamental and well-studied  $\mathcal{NP}$ -hard optimization problem. Given a complete graph the TSP asks for a Hamiltonian cycle of minimum length, i.e., a round-trip salesperson tour that visits each node exactly once before ending at the start node. In the Euclidean TSP (E-TSP) nodes are associated with point coordinates in the Euclidean space and pairwise (symmetric) inter-node distances are given by the Euclidean distance; the E-TSP remains  $\mathcal{NP}$ -hard.

Since its introduction in 1930 a body of knowledge has been built around the TSP. As a consequence, a plethora of methods has been developed ranging from sophisticated exact solvers (guarantee to find an optimum) to fast heuristic algorithmic approaches with no performance guarantees at all. In the domain of exact TSP solving, the branch-and-cut based Concorde solver by [1] is the state of the art. However, even though instances with hundreds of nodes can be solved within seconds [2], no guarantees for reasonable runtime can be given for large instances.

For the E-TSP, distances adhere to the triangle inequality induced by the Euclidean metric. This property can be leveraged to come up with approximation algorithms. For a

minimization problem, an  $(1 + \alpha)$ -approximation algorithm  $A$  guarantees that the tour length of  $A$  on a problem instance is at most  $(1 + \alpha) \cdot OPT$ , where  $OPT$  is the optimal tour length. Christofides [3] introduced an algorithm that achieves an approximation factor of  $3/2$ , which is the best constant approximation factor known for the E-TSP. Celebrated work by Arora [4] and – independently – by Mitchell [5] introduced a *Polynomial Time Approximation Schema* (PTAS) algorithm for the metric TSP which guarantees to produce solutions of quality at most  $(1 + 1/\alpha) \cdot OPT$  for each constant  $\alpha > 0$  in polynomial time. However, the algorithms are highly sophisticated and to the best of our knowledge no practical implementation of this PTAS is available. Moreover, PTAS naturally suffers from impractical runtimes if  $\alpha$  is increased – in other words: a reduction of  $\alpha$  goes hand in hand with a dramatic increase of the polynomial degree.

In recent years tremendous advances in heuristic TSP solving have been made where no formal performance guarantee can be given. Nevertheless the two best performing heuristics, LKH by [6] (based on sophisticated  $k$ -opt moves and the Lin-Kernighan heuristic) and the genetic algorithm EAX by [7] (a  $(\mu + \lambda)$  genetic algorithm adopting the eponymous edge assembly crossover operator and sophisticated diversity preservation) solve instances with thousands of nodes to optimality within reasonable time limits [8]. The respective restart versions LKH+r and EAX+r, which trigger a restart once the internal stopping conditions of the respective vanilla versions have been satisfied, pose the state of the art in inexact TSP solving [9], [10]. Recent endeavours by [11] extend both solvers by a sophisticated crossover operator – the generalized partition crossover (GPX2) – which has shown superior performance over the vanilla version of EAX on large instances with node numbers in the five-digit range.

In the field of per-instance algorithm selection (AS) – see, e.g., the survey of [12] for further details – the goal is to build a model which automatically selects the best-performing algorithm from a portfolio of algorithms with complementary performances. In case of the E-TSP, the complementary behavior of EAX(+r) and LKH(+r) across a wide range of problem instances has been leveraged in several studies in recent years [9], [10]. Both of these works focused on optimality of the found solutions, i.e., the runtime until a solver found a tour of

optimal (i.e., minimal) length was measured, and runs which did not succeed (within the given time of one hour) were penalized.

Despite their success in solving E-TSP instances up to optimality, little is known about the empirical approximation qualities, i.e., the anytime behaviour, of LKH+r and EAX+r building on the concept of empirical runtime distributions [13]. Our work will thus shed light on the relationship between runtimes and respective approximation qualities. This is conceptually similar to the commonly accepted benchmarking practice in single-objective continuous optimization on the Black-Box Optimization Benchmark (BBOB, [14]). It should be noted, however, that despite simultaneous consideration of runtime and solution quality, the herein considered analysis of the anytime behavior of a solver (w.r.t. its solution quality) differs substantially from the multi-objective approach that was taken in [15]. Our analysis is supported by an extensive study on a wide range of TSP instances from the literature with 500, 1 000 and 2 000 nodes, respectively. More precisely, we pursue three research questions in this work:

- R1** Given different  $\alpha$ -values and time limits  $T$ , what is the probability to calculate an  $(1 + \alpha)$ -approximate solution for different variants of LKH+r and EAX+r on a large and diverse set of E-TSP instances in time  $T$ ?
- R2** Which approximation quality  $(1 + \alpha)$  can we expect given a time limit  $T$ ?
- R3** How can automated algorithm selection approaches make use of information on anytime behaviour of TSP solvers?

These questions are addressed in the remainder of the paper which is organized as follows. Section II introduces the methodology underlying the experimental study, including problem instances and algorithms. Results are presented and analyzed in Section III, and the paper concludes with a discussion of promising future research directions in Section IV.

## II. METHODOLOGY

In this section we detail the setup of our experimental study.

### A. Problem instances

Several AS-studies revealed – which is in line with intuition – that characteristics of TSP problems, such as clustering properties or the depth of a minimum spanning tree, may have a strong impact on the running time until an optimal solution is found [10]. It is legitimate to assume that this will also be true for the  $(1 + \alpha)$ -approximation case. However, feature impacts and relations may change. Thus, to investigate – and ideally support – our assumptions empirically, we conducted an experimental study across a wide range of different E-TSP instances. For better comparability of our results, our setup is aligned with previous studies [9], [10], [16] and hence covers the following TSP sets:

**rue** Classical Random Uniform Euclidean instances where point coordinates are spread uniformly at random in the bounded Euclidean space  $[0, 10^6] \times [0, 10^6]$ .

**netgen** In [17] this type of strongly clustered instances of size  $n$  was proposed. For a given number of clusters  $n_c \in$

$\{2, 3, 4, 5, 10\}$ , respective cluster centers are placed well-spread in the Euclidean plane (in  $[0, 10^6] \times [0, 10^6]$ ) by Latin-Hypercube-Sampling (LHS). Subsequently,  $n/n_c$  cities are sampled around each cluster center assuring cluster segregation.

**morphed** A morphed instance originates from combining a rue with a netgen instance of equal size. First, an optimal weighted point matching is calculated between the point coordinates of both instances. Next, the matched points are used to calculate new points by convex combination of the coordinates of the matched points. This approach was introduced in [18] and later improved in [17].

**tspgen** Instances were generated by sequential application of mutation to an initial rue instance as proposed by [19]. For each mutation, a random subset of points is selected and rearranged by means of “creative” mutation operators, e.g., a grid-mutation, which aligns a random subset of points in a grid structure. These operators are inspired by observations on real-world instances (e.g., from Very Large Scale Integration, VLSI) and meant to produce instances that are structurally heterogeneous.

**evolved** TSP instances evolved by means of an evolutionary algorithm which minimizes the ratio of Penalized-Average-Runtime (PAR, [20]) scores<sup>1</sup> of solvers EAX+r and LKH+r producing instances that are easy for one and hard(er) for the competitor. The set of evolved instances considered within this work is taken from [19].

For instance sets rue, netgen, morphed and tsgen we each consider 150 instances of size  $n \in \{500, 1\,500, 2\,000\}$ . Subsets of 30 instances thus contain  $n_c \in \{2, 3, 4, 5, 10\}$  clusters for instances of type netgen. The evolved instances – taken from [19] – are restricted to  $n = 500$  nodes<sup>2</sup>. There are each 100 instances which are easy for EAX+r and LKH+r respectively. Summing up, in total, our benchmark set constitutes 2 000 E-TSP instances. Note that we intentionally do not include instances from the well-known TSPLIB [21] benchmark set. To make well-founded statements about the research questions addressed in this work we require a large and systematic set of instances from different classes of equal size. However, TSPLIB instances are very heterogeneous in both size and structure which does not allow for proper evaluation.

### B. Considered Algorithms

In total our study considers six different solvers for the E-TSP. The first four are restart variants of LKH [6] while the latter two are restart variants of EAX [7]<sup>3</sup>. In particular these variants incorporate generalized partition crossover (GPX2) into the algorithms [11].

<sup>1</sup>The PAR score is the average running time of  $m$  independent solver runs until an optimal solution was found. Runs which are not successful within time  $T$  are penalized with  $f \cdot T$  where  $f$  is a penalty factor usually set to 10.

<sup>2</sup>Those instances were generated by an evolutionary algorithm where a single fitness function evaluation requires (1) a call of the exact Concorde solver and (2) multiple runs of LKH+r and EAX+r respectively. This becomes computationally very expensive for  $n \in \{1\,000, 2\,000\}$ .

<sup>3</sup>Restart variants trigger a cold restart once the internal stopping conditions are hit. This modification to LKH and EAX was introduced in [9].

**LKH variants:** The LKH algorithm is an iterated local search algorithm that heuristically generates  $k$ -opt moves. A powerful improvement of LKH was the introduction of multi-trial LKH, where several solutions originating from soft restarts of the Lin-Kernighan heuristic are recombined by a partition crossover operator named Iterative Partial Transcription (IPT). A recent proposal replaces IPT by the alternative crossover operator GPX2. Additionally, LKH v2.0.9 allows to use both IPT and GPX2 in sequence. Therefore, we consider the four restarts variants LKH+r (IPT), i.e., the vanilla version of LKH+r, LKH+r (GPX), LKH+r (IPT+GPX) and LKH+r (GPX+IPT).

**EAX variants:** EAX is a powerful genetic algorithm which uses the Edge Assembly Crossover (EAX) operator to combine two parents. The operator is designed to keep as many edges from the parents as possible and introduces only a few short edges to complete the tour. The EAX algorithm is a  $(\mu + \lambda)$ -strategy with a sophisticated diversity preservation technique based on edge entropy to prevent the algorithm from premature convergence. We use the restart version EAX+r and additionally consider a modified version where individuals created by EAX+r are further recombined by applying GPX2. It should be noted that our modification is more straight-forward than the different variants introduced in [11].

### C. Estimation of Probabilities

Next, we describe the process of probability estimation. Given a value  $\alpha \geq 0$ , a time-limit  $T$ , a stochastic algorithm  $A$ , and an instance  $I$  with optimal tour length OPT, we denote the probability to reach a solution of desired quality within the time-limit  $T$  as  $p_{\alpha,T}^A(I) = P(A(I) \leq (1 + \alpha) \cdot \text{OPT})$ . Given  $A$  is stochastic and the trajectories of  $m$  independent runs of that algorithm on instance  $I$  are available, the probability  $p_{\alpha,T}^A(I)$  can be estimated by the relative number of runs that succeeded in finding an  $(1 + \alpha)$ -approximation within  $T$ , i.e.,

$$\hat{p}_{\alpha,T}^A(I) = \frac{1}{m} \sum_{i=1}^m \mathbf{1}(A_T^i(I) \leq (1 + \alpha) \cdot \text{OPT}).$$

Here,  $A_T^i(I)$  is the incumbent solution of  $A$  on  $I$  in the  $i$ -th run after time  $T$  and  $\mathbf{1}$  is the indicator function evaluating to 1 if its argument is true. Given a set  $\mathcal{I}$  of instances, an estimator for the success probability  $p_{\alpha,T}^A(\mathcal{I})$  on the set  $\mathcal{I}$  is the average probability over all instances in  $\mathcal{I}$ , i.e.,

$$\hat{p}_{\alpha,T}^A(\mathcal{I}) = \frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} \hat{p}_{\alpha,T}^A(I).$$

## III. EXPERIMENTS

### A. Experimental Setup

Each of the six algorithms was run on each instance  $m = 10$  times with different random seeds in order to account for stochasticity. Throughout those experiments, we used a cutoff time of  $T = 3600$  seconds (i.e., one hour). The algorithms log the incumbent, i.e., best-so-far, solution in every run.

As all results for the TSP sets netgen, morphed and tspgen were qualitatively similar, we combined the respective information into a single set called “structured”.

### B. Perspective 1: First hitting times

In practical applications, often an upper bound on solver performance is desired in order to realistically assess the worst case scenario. Therefore, it is of interest how long it will take a solver *at most* to find a  $(1 + \alpha)$ -approximation of the true optimum. The chosen solvers were executed for a variety of approximation gaps  $\alpha$  and for each combination with TSP set and instance sizes  $n \in \{500, 1000, 2000\}$ . Note that in this part of our study, we analyze the algorithms’ performances across a wide range of approximation factors,  $\alpha \in \{0.5, 0.1, 0.05, 0.01, \dots, 5 \cdot 10^{-5}, 10^{-5}\}$ . Thereby, we get a comprehensive overview of the behavior of the different algorithms, and later on can zoom in on more relevant areas. We take a pessimistic perspective and estimate the first hitting time as the maximum time needed by each algorithm to find a solution of the corresponding quality  $(1 + \alpha)$  for the first time across all instances and runs of each instance set and size respectively.

As depicted in Fig. 1, both EAX variants are extremely fast for large values of  $\alpha$ , i.e., their solutions at early stages of the runs are already of very good quality. In fact, for  $\alpha = 0.5$ , EAX+r (GPX) had the lowest first hitting times across all eight considered scenarios, which supports the effectiveness of the sophisticated crossover operator GPX. More precisely, for small ( $n = 500$ ) and medium-sized ( $n = 1000$ ) instances, all runs of EAX+r (GPX) found a tour that is at most 50% longer than the optimal tour within less than a second. On larger instances ( $n = 2000$ ), this optimizer needed just slightly more than a second. Yet, it is also observable that the classical EAX+r performs better than its GPX-counterpart for decreasing  $\alpha$ -values and even outperforms it for all approximation gaps  $\alpha \leq 0.01$ .

Further noticeable findings can be derived for LKH and its variants. First, the trajectories of all four LKH variants are almost identical across all the investigated scenarios, implying that no substantial differences among the considered versions regarding latest first hitting times could be detected. Moreover, with the exception of the medium-sized and large structured TSP instances with  $n \in \{1000, 2000\}$  nodes, LKH performs at least as well as EAX within the mid-range approximation factors  $\alpha \in [0.0005, 0.01]$ . However, for the small approximation factors, LKH is again inferior to EAX – except for the problems that were specifically tailored in favor of LKH.

Thus, depending on the desired approximation quality, one could provide a three-fold recommendation: if the acceptable approximation-quality is rather large ( $\alpha > 0.1$ ) EAX+r (ideally its GPX-enhanced version) is preferable, for mid-range values of  $\alpha$  one should rather use one of the LKH variants, and for very small approximation factors ( $\alpha \leq 0.0001$ ) one should consider the classical EAX+r. However, it has to be kept in mind that these findings solely focus the worst case scenario.




Fig. 1. Maximum first hitting times for different  $\alpha$ -values. That is, we report for each approximation factor  $\alpha$  ( $x$ -axis) the maximum time needed by EAX+r and LKH+r to find a solution of the corresponding quality for the first time across all instances and runs. It is therefore a maximally pessimistic view on approximation quality. Splits are done by instance class and size.

### C. Perspective 2: Probability of Success

While the previous subsection focused on the latest first hitting time, i.e., a worst case analysis, we are now interested in the solvers’ average case performances. Of course, previous studies already addressed the average case as well, but usually with a focus on (penalized) aggregations of runtimes. Unfortunately, those runtimes only considered whether an algorithm found a tour of optimal length – which corresponds to  $\alpha = 0$ . Further information on the solver’s performance, such as whether it failed (to find an optimal tour) by orders of magnitude or just by an infinitesimal deviation was neglected. In the following, we will overcome this “knowledge gap” by comparing the performances of a total of six versions of EAX+r and LKH+r across different TSP sets and approximation factors. More precisely, within this subsection we investigate the change of an algorithm’s average success probability over time for different approximation gaps. For small instances ( $n = 500$ ) we used  $\alpha \in [0, 10^{-4}]$  and otherwise  $\alpha \in [0, 10^{-3}]$ . This choice of  $\alpha$ -values was considered sufficiently large, as for the largest values in those intervals the average success probability of most algorithms already converged to 1.0 within the investigated time.

The results for the TSP instances with  $n = 500$  nodes are depicted in Fig. 2 and also listed in Tab. I. As the performances of all four LKH variants are very similar only the results for LKH+r’s default version, i.e., LKH+r (IPT) are provided as a representative. Interestingly, LKH+r outperforms both EAX variants on the TSP set “easy for LKH+r” across all considered combinations of time steps and approximation gaps, although those instances have been generated for just one specific pair of maximum runtime ( $T = 5$  minutes) and approximation quality ( $\alpha = 0$ ) as detailed in [19]. These findings are reflected

by the median probability curves depicted in Fig. 2. As shown in the second row of this plot, the curve of LKH+r is located in the very top-left corner, implying that it has achieved a high success probability (w.r.t. the respective approximation quality) within seconds. In contrast, the “tubes” associated with the EAX+r variants are rather broad and spread diagonally across all four images. That is, their median success probabilities show a high variance and both solvers require much more time to achieve feasible success probabilities (closer to 1.00) on those instances. The superiority of LKH+r over the two EAX+r versions is also confirmed by pairwise Wilcoxon-tests to a significance level of 5% (as indicated by the  $1^+$  and  $2^+$  in the last column of Tab. I).

On the other hand, the two EAX versions are absolutely superior to LKH+r on the “easy for EAX+r” problems. In fact, the tube of LKH+r basically covers the majority of each of the images in the first row of Fig. 2.

We further noticed that while EAX+r and its variant performs well on all sets except for “easy for LKH”, LKH+r exhibits clear preferences – ranging from very poor performances on the EAX-tailored instances, via mediocre behavior on the structured instances, up to good performances on rue, and (of course) the LKH-tailored problems. This clearly indicates that the structural properties of an instance, i.e., its node alignment, strongly affect the optimization behavior of LKH.

In contrast to our a priori expectations, we could not detect strong differences of a solver’s success probabilities across the different approximation qualities for the small instances ( $n = 500$ ) at hand. However, for larger instances with  $n \in \{1000, 2000\}$  nodes the picture slightly changes as shown in Fig. 3 and Tab. II. First, with increasing instance size  $n$  and decreasing  $\alpha$ , the EAX-variant EAX+r (GPX) is

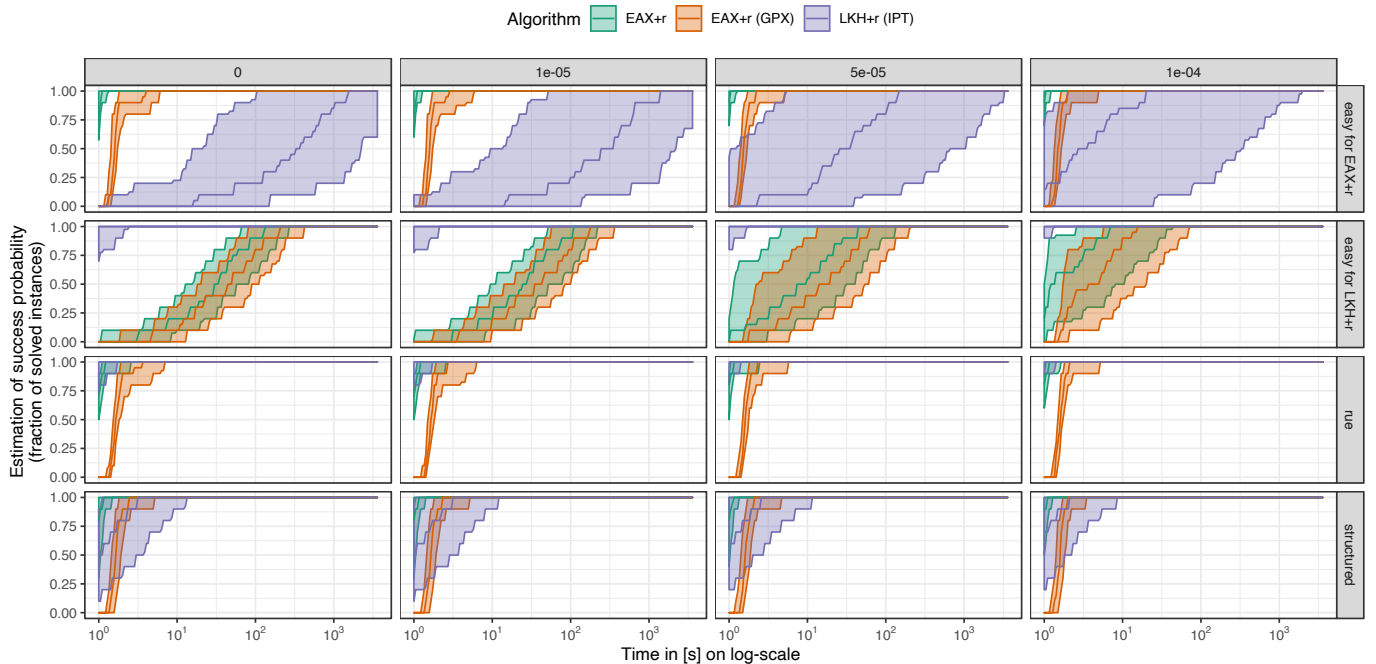


Fig. 2. This plot shows the median success probabilities to locate a solution of quality  $(1 + \alpha) \cdot \text{OPT}$  (columns show different  $\alpha$  values) for instances with  $n = 500$  nodes. The tubes are defined by the corresponding 0.25-quantile and 0.75-quantile, respectively.

more and more outperformed by its contenders. This effect is reflected by the large number of  $2^+$ 's in Tab. II, which indicate a significantly better performance of the respective solver compared to EAX+r (GPX). Secondly, across all solvers one can observe an increase in variation of the success probabilities (i.e., wider tubes), as well as slightly increasing runtimes (i.e., shift to the right) for decreasing approximation factors  $\alpha$ . However, those findings are quite intuitive as diminishing gaps correspond to more accurate results – which are harder to find.

Observing Tab. II another pattern becomes visible: EAX+r performs significantly better than LKH+r on the structured instances. An even more interesting pattern can be observed for the (unstructured) rue instances. If the accepted approximation factor is rather small ( $\alpha \leq 0.0001$ ) and a sufficiently large time ( $T \geq 50$ s) is given, then EAX+r performs significantly better than LKH+r. However, if the instance size is large ( $n = 2000$ ) and only a short amount of time is given ( $T = 10$ s), LKH+r is superior to both (considered) versions of EAX.

As a result, the ideal choice of (inexact) TSP optimization heuristic depends on a combination of (i) the magnitude of the allowed approximation gap, (ii) the size of the TSP instance, (iii) the given cutoff time for the solver, as well as (iv) the structure (i.e., node placement) of the instance itself. With the exception of (i), these findings are in line with the automated algorithm selection studies by [9], [10], in which the authors showed (for  $\alpha = 0$ ) that an instance's structural information can be efficiently exploited to select the best performing optimization algorithm for the instance at hand.

#### IV. CONCLUSION AND FUTURE WORK

Taking an “anytime perspective” in TSP solver benchmarking, i.e., addressing research questions R1 and R2 stated in the introduction, results in detailed insights into solver performances together with respective approximation speed and indicates structural relationships with instance properties. Specifically, our sophisticated evolutionary approach [19] for generating instances which are very hard for EAX+r and easy for LKH+r proves to perform extremely well reflecting the huge impact of instance properties on problem hardness.

Next steps will incorporate the insights in anytime performance of TSP solvers gained from our empirical study into automated algorithm selection models (ref. to research question R3). Building on existing high-performing approaches conditioned on the necessity of solving the problem to optimality (e.g., [10]) an extension to the anytime scenario would be highly desirable. For this purpose, several ingredients need to be developed such as instance features characterizing anytime performance. Here, initial work on so-called probing features exists [9] and needs to be extended. Moreover, one of course needs an anytime performance indicator which is not straightforward to design as it has to incorporate different aspects of quality and hence is multi-objective in nature. Interesting concepts in the context of automated algorithm configuration [22] will be tested and improved upon.

Secondly, the derived insights offer very promising potential in terms of hybridization of inexact TSP solvers as performance rankings differ along the optimization runs with increasing approximation quality. Comparing, e.g., maximum first hitting times, EAX+r generates substantial improvements

TABLE I

MAXIMUM GAP (MAX), MEAN SUCCESS PROBABILITY (MEAN), STANDARD DEVIATION (STD) AND RESULTS OF PAIRWISE WILCOXON-TESTS FOR EAX+r, EAX+r (GPX) AND LKH+r (IPT). A VALUE  $X^+$  IN THE STAT COLUMN INDICATES THAT THE RESULTS OF THE ALGORITHM ARE STATISTICALLY SIGNIFICANT IN COMPARISON TO ALGORITHM  $X$ . RESULTS ARE SPLIT BY INSTANCE GROUP,  $\alpha$  AND TIME FOR INSTANCES OF SIZE  $n = 500$ . BEST MEAN VALUES PER ROW ARE HIGHLIGHTED IN **BOLD FACE**.

Group	$n$	$\alpha$	$T$	EAX+r (1)			EAX+r (GPX) (2)				LKH+r (IPT) (3)				
				max	mean	std	stat	max	mean	std	stat	max	mean	std	stat
easy for EAX+r	500	0.00010	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04	3 <sup>+</sup>	0.00	0.55	0.46	
	500	0.00010	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.66	0.44	
	500	0.00010	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.69	0.42	
	500	0.00005	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04	3 <sup>+</sup>	0.00	0.42	0.46	
	500	0.00005	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.54	0.45	
	500	0.00005	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.59	0.44	
	500	0.00001	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.06	3 <sup>+</sup>	0.00	0.27	0.40	
	500	0.00001	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.01	3 <sup>+</sup>	0.00	0.38	0.43	
	500	0.00001	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.43	0.43	
	500	0.00000	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.06	3 <sup>+</sup>	0.00	0.24	0.38	
	500	0.00000	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.01	3 <sup>+</sup>	0.00	0.35	0.41	
	500	0.00000	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.40	0.41	
	500	0.00010	10	0.00	0.77	0.34	2 <sup>+</sup>	0	0.68	0.37		0.00	<b>1.00</b>	0.01	1 <sup>+</sup> , 2 <sup>+</sup>
	500	0.00010	50	0.00	0.90	0.20		0	0.85	0.28		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>
	500	0.00010	100	0.00	0.95	0.14		0	0.92	0.18		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>
500	0.00005	10	0.00	0.57	0.39	2 <sup>+</sup>	0	0.48	0.38		0.00	<b>1.00</b>	0.01	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00005	50	0.00	0.81	0.27	2 <sup>+</sup>	0	0.73	0.31		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00005	100	0.00	0.89	0.21		0	0.86	0.22		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00001	10	0.00	0.32	0.31		0	0.26	0.27		0.00	<b>1.00</b>	0.01	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00001	50	0.00	0.67	0.32	2 <sup>+</sup>	0	0.57	0.32		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00001	100	0.00	0.79	0.27	2 <sup>+</sup>	0	0.73	0.29		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00000	10	0.00	0.26	0.24		0	0.21	0.22		0.00	<b>1.00</b>	0.01	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00000	50	0.00	0.64	0.31	2 <sup>+</sup>	0	0.53	0.32		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00000	100	0.00	0.78	0.27	2 <sup>+</sup>	0	0.69	0.31		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>	
500	0.00010	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04		0.00	0.99	0.07		
500	0.00010	50	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.03		
500	0.00010	100	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00		
500	0.00005	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.05		0.00	0.99	0.07		
500	0.00005	50	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.03		
500	0.00005	100	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00		
500	0.00001	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.05		0.00	0.98	0.09		
500	0.00001	50	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.04		
500	0.00001	100	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.01		
500	0.00000	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.98	0.06		0.00	0.98	0.09		
500	0.00000	50	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.04		
500	0.00000	100	0.00	<b>1.00</b>	0.00		0	<b>1.00</b>	0.00		0.00	1.00	0.01		
500	0.00010	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	1.00	0.03	3 <sup>+</sup>	0.01	0.91	0.22		
500	0.00010	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.00	3 <sup>+</sup>	0.00	0.97	0.13		
500	0.00010	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.98	0.11		
500	0.00005	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04	3 <sup>+</sup>	0.01	0.89	0.24		
500	0.00005	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.00	3 <sup>+</sup>	0.00	0.96	0.15		
500	0.00005	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.98	0.13		
500	0.00001	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04	3 <sup>+</sup>	0.01	0.87	0.26		
500	0.00001	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.00	3 <sup>+</sup>	0.00	0.96	0.16		
500	0.00001	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.97	0.13		
500	0.00000	10	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0	0.99	0.04	3 <sup>+</sup>	0.01	0.87	0.26		
500	0.00000	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	1.00	0.00	3 <sup>+</sup>	0.00	0.96	0.16		
500	0.00000	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.97	0.13		

very fast, is then overtaken by LKH+r (or respective LKH variants) while it clearly is the single-best solver referring to final approximation qualities. Of course, this behaviour is dependent on the instances' structural properties such that a hybridized TSP solver variant that is capable of processing instance features could vastly improve approximation speed and final quality.

#### ACKNOWLEDGMENT

P. Kerschke and H. Trautmann acknowledge support by the European Research Center for Information Systems (ERCIS). J. Bossek was supported by the Australian Research Council (ARC) through grant DP160102401.

TABLE II

MAXIMUM GAP (**MAX**), MEAN SUCCESS PROBABILITY (**MEAN**), STANDARD DEVIATION (**STD**) AND RESULTS OF PAIRWISE WILCOXON-TESTS FOR EAX+r, EAX+r (GPX) AND LKH+r (IPT). A VALUE  $X^+$  IN THE STAT COLUMN INDICATES THAT THE RESULTS OF THE ALGORITHM ARE STATISTICALLY SIGNIFICANT IN COMPARISON TO ALGORITHM  $X$ . RESULTS ARE SPLIT BY INSTANCE GROUP,  $\alpha$  AND TIME FOR INSTANCES OF SIZE  $n > 500$ . BEST MEAN VALUES PER ROW ARE HIGHLIGHTED IN **BOLD FACE**.

Group	$n$	$\alpha$	$T$	EAX+r (1)				EAX+r (GPX) (2)				LKH+r (IPT) (3)			
				max	mean	std	stat	max	mean	std	stat	max	mean	std	stat
rue	1000	0.00100	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup>	0.00	0.92	0.13		0.00	<b>1.00</b>	0.00	2 <sup>+</sup>
	1000	0.00100	50	0.00	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00	
	1000	0.00100	100	0.00	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00		0.00	<b>1.00</b>	0.00	
	1000	0.00010	10	0.00	<b>0.97</b>	0.08	2 <sup>+</sup>	0.00	0.27	0.24		0.00	0.93	0.18	2 <sup>+</sup>
	1000	0.00010	50	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.82	0.24		0.00	0.98	0.09	2 <sup>+</sup>
	1000	0.00010	100	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.93	0.16		0.00	0.99	0.07	2 <sup>+</sup>
	1000	0.00001	10	0.00	<b>0.89</b>	0.18	2 <sup>+</sup>	0.00	0.16	0.20		0.00	0.84	0.27	2 <sup>+</sup>
	1000	0.00001	50	0.00	<b>1.00</b>	0.02	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.66	0.32		0.00	0.96	0.15	2 <sup>+</sup>
	1000	0.00001	100	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.81	0.26		0.00	0.97	0.11	2 <sup>+</sup>
	1000	0.00000	10	0.00	<b>0.88</b>	0.19	2 <sup>+</sup>	0.00	0.15	0.20		0.00	0.83	0.27	2 <sup>+</sup>
	1000	0.00000	50	0.00	<b>1.00</b>	0.03	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.64	0.32		0.00	0.96	0.15	2 <sup>+</sup>
	1000	0.00000	100	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.79	0.27		0.00	0.97	0.11	2 <sup>+</sup>
	2000	0.00100	10	0.03	0.43	0.17	2 <sup>+</sup>	0.04	0.00	0.00		0.00	<b>1.00</b>	0.00	1 <sup>+</sup> , 2 <sup>+</sup>
	2000	0.00100	50	0.00	<b>1.00</b>	0.00	2 <sup>+</sup>	0.00	0.05	0.09		0.00	<b>1.00</b>	0.00	2 <sup>+</sup>
	2000	0.00100	100	0.00	<b>1.00</b>	0.00	2 <sup>+</sup>	0.00	0.49	0.21		0.00	<b>1.00</b>	0.00	2 <sup>+</sup>
	2000	0.00010	10	0.03	0.18	0.16	2 <sup>+</sup>	0.04	0.00	0.00		0.00	<b>0.59</b>	0.30	1 <sup>+</sup> , 2 <sup>+</sup>
	2000	0.00010	50	0.00	<b>0.99</b>	0.05	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.00	0.00		0.00	0.89	0.18	2 <sup>+</sup>
	2000	0.00010	100	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.01	0.03		0.00	0.95	0.12	2 <sup>+</sup>
	2000	0.00001	10	0.03	0.06	0.10	2 <sup>+</sup>	0.04	0.00	0.00		0.00	<b>0.34</b>	0.28	1 <sup>+</sup> , 2 <sup>+</sup>
	2000	0.00001	50	0.00	<b>0.83</b>	0.23	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.00	0.00		0.00	0.71	0.32	2 <sup>+</sup>
	2000	0.00001	100	0.00	<b>0.94</b>	0.13	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.00	0.01		0.00	0.83	0.28	2 <sup>+</sup>
	2000	0.00000	10	0.03	0.04	0.08	2 <sup>+</sup>	0.04	0.00	0.00		0.00	<b>0.31</b>	0.27	1 <sup>+</sup> , 2 <sup>+</sup>
	2000	0.00000	50	0.00	<b>0.78</b>	0.25	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.00	0.00		0.00	0.67	0.32	2 <sup>+</sup>
	2000	0.00000	100	0.00	<b>0.91</b>	0.17	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.00	0.01		0.00	0.80	0.29	2 <sup>+</sup>
	1000	0.00100	10	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.99	0.06		0.03	0.97	0.11	
	1000	0.00100	50	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	0.99	0.06	
	1000	0.00100	100	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	<b>1.00</b>	0.00	3 <sup>+</sup>	0.00	1.00	0.04	
	1000	0.00010	10	0.00	<b>0.99</b>	0.07	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.66	0.31		0.03	0.80	0.29	2 <sup>+</sup>
	1000	0.00010	50	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.96	0.13		0.00	0.94	0.17	
	1000	0.00010	100	0.00	<b>1.00</b>	0.01	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.98	0.09	3 <sup>+</sup>	0.00	0.97	0.13	
1000	0.00001	10	0.00	<b>0.95</b>	0.14	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.48	0.34		0.03	0.68	0.35	2 <sup>+</sup>	
1000	0.00001	50	0.00	<b>1.00</b>	0.03	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.88	0.22		0.00	0.87	0.26		
1000	0.00001	100	0.00	<b>1.00</b>	0.03	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.94	0.17		0.00	0.92	0.22		
1000	0.00000	10	0.00	<b>0.94</b>	0.15	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.45	0.34		0.03	0.65	0.36	2 <sup>+</sup>	
1000	0.00000	50	0.00	<b>1.00</b>	0.04	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.87	0.24		0.00	0.86	0.27		
1000	0.00000	100	0.00	<b>1.00</b>	0.03	2 <sup>+</sup> , 3 <sup>+</sup>	0.00	0.93	0.18		0.00	0.91	0.23		
2000	0.00100	10	0.02	0.78	0.26	2 <sup>+</sup>	0.06	0.00	0.00		0.09	<b>0.89</b>	0.21	1 <sup>+</sup> , 2 <sup>+</sup>	
2000	0.00100	50	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.45	0.35		0.04	0.99	0.06	2 <sup>+</sup>	
2000	0.00100	100	0.00	<b>1.00</b>	0.00	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.87	0.19		0.04	1.00	0.03	2 <sup>+</sup>	
2000	0.00010	10	0.02	<b>0.51</b>	0.28	2 <sup>+</sup> , 3 <sup>+</sup>	0.06	0.00	0.00		0.09	0.40	0.33	2 <sup>+</sup>	
2000	0.00010	50	0.00	<b>0.99</b>	0.06	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.03	0.10		0.04	0.80	0.29	2 <sup>+</sup>	
2000	0.00010	100	0.00	<b>1.00</b>	0.02	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.15	0.18		0.04	0.88	0.24	2 <sup>+</sup>	
2000	0.00001	10	0.02	<b>0.32</b>	0.26	2 <sup>+</sup> , 3 <sup>+</sup>	0.06	0.00	0.00		0.09	0.19	0.25	2 <sup>+</sup>	
2000	0.00001	50	0.00	<b>0.93</b>	0.16	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.01	0.05		0.04	0.60	0.35	2 <sup>+</sup>	
2000	0.00001	100	0.00	<b>0.98</b>	0.09	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.04	0.10		0.04	0.72	0.34	2 <sup>+</sup>	
2000	0.00000	10	0.02	<b>0.28</b>	0.25	2 <sup>+</sup> , 3 <sup>+</sup>	0.06	0.00	0.00		0.09	0.15	0.24	2 <sup>+</sup>	
2000	0.00000	50	0.00	<b>0.90</b>	0.18	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.01	0.05		0.04	0.55	0.36	2 <sup>+</sup>	
2000	0.00000	100	0.00	<b>0.97</b>	0.11	2 <sup>+</sup> , 3 <sup>+</sup>	0.01	0.03	0.09		0.04	0.68	0.35	2 <sup>+</sup>	

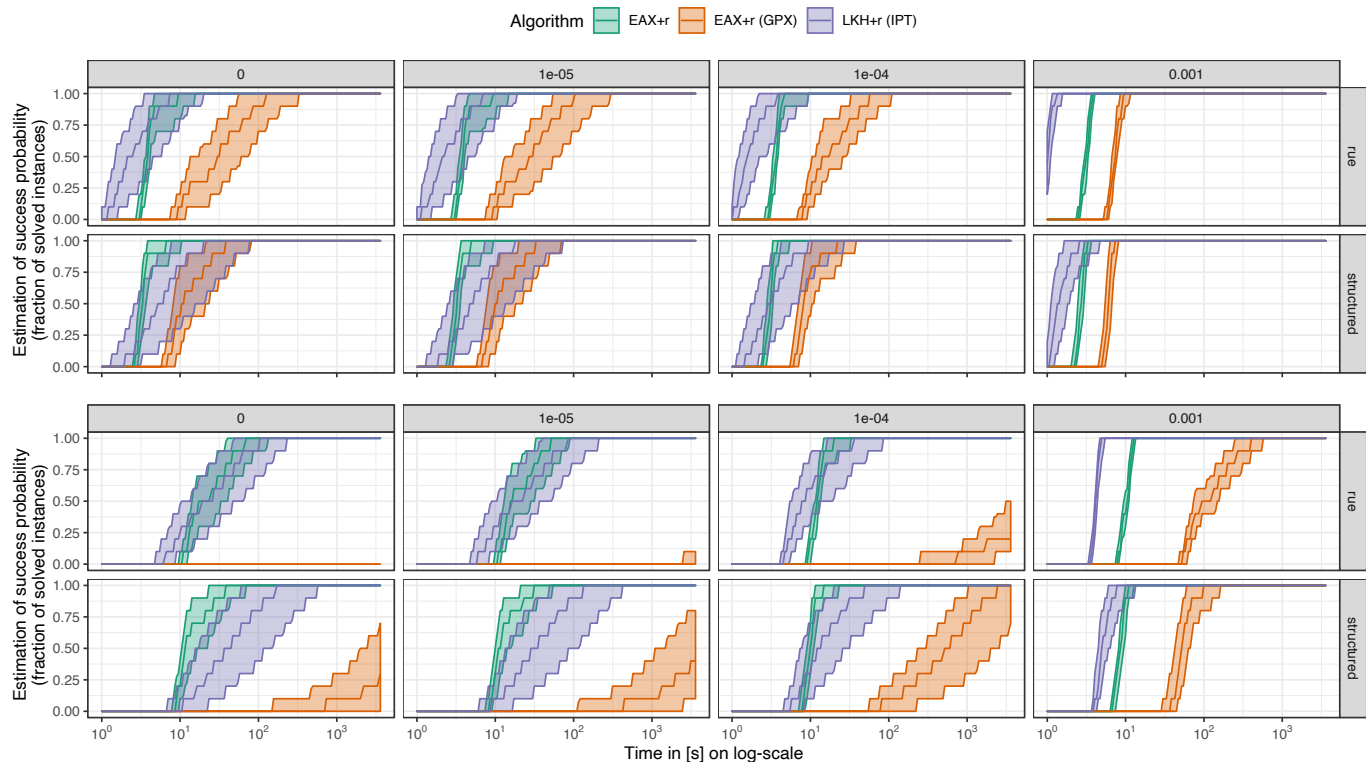


Fig. 3. This plot shows the median success probabilities to locate a solution of quality  $(1 + \alpha) \cdot \text{OPT}$  (columns show different  $\alpha$  values) for instances with  $n = 1000$  nodes (two top rows) and  $n = 2000$  nodes (two bottom rows). The tubes are defined by the 0.25-quantile and 0.75-quantile, respectively.

## REFERENCES

- [1] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, 2007.
- [2] W. J. Cook, *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press, 2012.
- [3] N. Christofides, “The Vehicle Routing Problem,” *Revue française d’automatique, d’informatique et de recherche opérationnelle (RAIRO). Recherche opérationnelle*, vol. 10, no. 1, pp. 55 – 70, 1976.
- [4] S. Arora, “Polynomial Time Approximation Schemes for Euclidean Traveling Salesman and Other Geometric Problems,” *Journal of the ACM (JACM)*, vol. 45, no. 5, pp. 753 – 782, 1998.
- [5] J. S. B. Mitchell, “Guillotine Subdivisions Approximate Polygonal Subdivisions: A Simple Polynomial-Time Approximation Scheme for Geometric TSP, k-MST, and Related Problems,” *SIAM Journal on Computing*, vol. 28, no. 4, pp. 1298 – 1309, 1999.
- [6] K. Helsgaun, “General k-opt submoves for the Lin-Kernighan TSP heuristic,” *Mathematical Programming Computation*, vol. 1, no. 2-3, pp. 119 – 163, 2009.
- [7] Y. Nagata and S. Kobayashi, “A Powerful Genetic Algorithm Using Edge Assembly Crossover for the Traveling Salesman Problem,” *INFORMS Journal on Computing*, vol. 25, no. 2, pp. 346 – 363, 2013.
- [8] F. Hutter, L. Xu, H. H. Hoos, and K. Leyton-Brown, “Algorithm Runtime Prediction: Methods & Evaluation,” *International Joint Conference on Artificial Intelligence*, pp. 4197–4201, 2015.
- [9] L. Kotthoff, P. Kerschke, H. Hoos, and H. Trautmann, “Improving the state of the art in inexact tsp solving using per-instance algorithm selection,” in *Learning and Intelligent Optimization*, C. Dhaenens, L. Jourdan, and M.-E. Marmion, Eds. Springer, 2015, pp. 202–217.
- [10] P. Kerschke, L. Kotthoff, J. Bossek, H. H. Hoos, and H. Trautmann, “Leveraging TSP Solver Complementarity through Machine Learning,” *Evolutionary Computation (ECJ)*, vol. 26, pp. 597 – 620, 2018.
- [11] D. Sanches, L. D. Whitley, and R. Tinós, “Building a Better Heuristic for the Traveling Salesman Problem: Combining Edge Assembly Crossover and Partition Crossover,” in *Proc. of the 19th Genetic and Evolutionary Computation Conference (GECCO)*. ACM, July 2017, pp. 329 – 336.
- [12] P. Kerschke, H. H. Hoos, F. Neumann, and H. Trautmann, “Automated Algorithm Selection: Survey and Perspectives,” *Evolutionary Computation (ECJ)*, vol. 27, pp. 3 – 45, 2019.
- [13] H. H. Hoos and T. Stützle, *Stochastic Local Search: Foundations and Applications*. Elsevier, 2004.
- [14] N. Hansen, A. Auger, O. Mersmann, T. Tušar, and D. Brockhoff, “COCO: A platform for comparing continuous optimizers in a black-box setting,” *ArXiv e-prints*, vol. arXiv:1603.08785, 2016.
- [15] J. Bossek, P. Kerschke, and H. Trautmann, “A Multi-Objective Perspective on Performance Assessment and Automated Selection of Single-Objective Optimization Algorithms,” *Applied Soft Computing (ASOC)*, vol. 88, p. 105901, 2020.
- [16] P. McMenemy, N. Veerapen, J. Adair, and G. Ochoa, “Rigorous Performance Analysis of State-of-the-Art TSP Heuristic Solvers,” in *Evolutionary Computation in Combinatorial Optimization*, A. Liefvooghe and L. Paquete, Eds. Springer, 2019, pp. 99 – 114.
- [17] S. Meisel, C. Grimme, J. Bossek, M. Wölck, G. Rudolph, and H. Trautmann, “Evaluation of a Multi-Objective EA on Benchmark Instances for Dynamic Routing of a Vehicle,” in *Proceedings of the 17th Genetic and Evolutionary Computation Conf. (GECCO)*. ACM, 2015, pp. 425 – 432.
- [18] O. Mersmann, B. Bischl, H. Trautmann, M. Wagner, J. Bossek, and F. Neumann, “A Novel Feature-Based Approach to Characterize Algorithm Performance for the Traveling Salesperson Problem,” *Annals of Math. and Artificial Intelligence*, vol. 69, no. 2, pp. 151–182, 2013.
- [19] J. Bossek, P. Kerschke, A. Neumann, M. Wagner, F. Neumann, and H. Trautmann, “Evolving Diverse TSP Instances by Means of Novel and Creative Mutation Operators,” in *Proc. of the 15th ACM/SIGEVO Workshop on Found. of Genetic Alg. (FOGA)*. ACM, 2019, pp. 58–71.
- [20] B. Bischl, P. Kerschke, L. Kotthoff, M. Lindauer, Y. Malitsky, A. Fréchette, H. H. Hoos, F. Hutter, K. Leyton-Brown, K. Tierney, and J. Vanschoren, “ASlib: A Benchmark Library for Algorithm Selection,” *Artificial Intelligence Journal (AIJ)*, vol. 237, pp. 41 – 58, 2016.
- [21] G. Reinelt, “TSPLIB – A Traveling Salesman Problem Library,” *ORSA Journal on Computing*, vol. 3, no. 4, pp. 376–384, 1991.
- [22] M. López-Ibáñez, Manuel and T. Stützle, “Automatically Improving the Anytime Behaviour of Optimisation Algorithms,” *European Journal of Operational Research (EJOR)*, vol. 235, no. 3, pp. 569 – 582, 2014.