
Hierarchical Knowledge for Heuristic Problem Solving — A Case Study on the Traveling Salesperson Problem

Alexandra Kirsch

ALEXANDRA.KIRSCH@UNI-TUEBINGEN.DE

Department of Computer Science, University of Tübingen

Abstract

Abstraction is a fundamental feature of human-level intelligence. But it is not clear how to combine knowledge on different levels of abstraction. This paper examines the use of hierarchical knowledge for heuristic problem-solving algorithms, regarding three options for integrating hierarchical knowledge into heuristic search: as a state evaluation heuristic, as a search-guiding heuristic and as a hierarchical search strategy. The Traveling Salesperson Problem serves as an example of a problem-solving task and the different strategies are evaluated with respect to tour length, robustness against misleading hierarchy assignments and acceptability of results by humans. It turns out that the most effective and stable results can be achieved with hierarchical knowledge as a search-guiding heuristic combined with other heuristics.

1. Motivation

Many problems that humans are confronted with, such as planning a tour through a given number of cities, are NP hard problems. Even though the space of possible solutions is huge, humans usually find remarkably good solutions. What is more, human solution strategies are stable when the problems are complicated by uncertainty, ignorance of facts or several optimization criteria. The human ability to find satisficing (Simon, 1956) rather than optimal solutions seems to be at the core of this fascinating ability, as Cassimatis (2012) states: “it may be the case that incorrectness and non-optimality are essential to human-level intelligence.”

In early AI research, most notably that of Newell and Simon (1972), research on artificial and natural cognition were pursued in parallel and findings from psychology were used as a basis for intelligent algorithms. This paper follows a similar approach: Greedy Expert Search (GES), the basic problem-solving strategy used, is based on heuristics, a heavily used mechanism in human problem solving (Shah & Oppenheimer, 2008). Heuristics do not only enhance efficiency: Gigerenzer and Gaissmaier (2011) point out a range of real-world examples, in which heuristic problem solving leads to better results than “rational” decision-making.

Another important concept underlying intelligent behavior is hierarchical abstraction. Typically hierarchical algorithms first find a solution in an abstracted state space, which is then refined in the original state space (Knoblock, 1991; Pang & Holte, 2011). In this paper I propose to treat abstract knowledge as a form of heuristic and show how this approach improves solutions for a specific problem, the Traveling Salesperson Problem.

A Traveling Salesperson Problem consists of finding the shortest possible path that visits all points in a given set, here in 2D space. Psychological research has produced evidence that humans employ some kind of hierarchical strategy when solving Traveling Salesperson Problems (Kong & Schunn, 2007; Graham, Joshi, & Pizlo, 2000). Explicit regions (such as membership of a city to a country) also have an influence on the human solution (Wiener, Ehbauer, & Mallot, 2009). Similar to most of the psychological models, hierarchical search algorithms in AI typically use a strict top-down hierarchy (Knoblock, 1991; Pang & Holte, 2011) with the challenge of how to make decisions on the lower level with the result from the more abstract level. However, Wiener and Mallot (2003) suggest a model for human wayfinding skills with a more flexible interaction of the knowledge on different levels of abstraction. In the same line, Hayes-Roth and Hayes-Roth (1979) have shown that different levels of abstraction interact constantly when humans solve everyday problems.

Four sets of configurations for problem solving are compared: 1) a flat approach ignoring hierarchical knowledge, 2) the flat approach complemented with a state evaluation function that uses knowledge about regions, 3) a mixed approach using hierarchical knowledge together with other heuristics to determine the direction of the search, and 4) an emulation of a hierarchical search algorithm, in which decisions on the abstract level restrict the possibilities on the lower level. The configurations are evaluated with four test sets in which the region definitions are more or less appropriate — from very helpful to distracting. It turns out that the third configuration, which uses hierarchical knowledge in addition to other heuristics, is the most robust with respect to misleading region definitions and returns good results.

The paper makes the following contributions:

- introduction of a novel way of using hierarchical knowledge in a greedy heuristic search algorithm,
- implementation of different paradigms of using hierarchical knowledge for the Traveling Salesperson Problem,
- comparison of the different strategies with respect to tour length and user acceptance.

The paper first gives an overview of related work, covering heuristic search, hierarchical approaches to Traveling Salesperson Problem solving and hierarchical search. Then the underlying heuristic search algorithm is introduced with its extension to using hierarchical knowledge in several ways. The different configurations are then evaluated on different data sets. The paper ends with a discussion and conclusion.

2. Related Work

The GES algorithm used here is most related to the decision-making procedure in the cognitive architecture FORR (Epstein, 2004). In contrast to GES, FORR uses different types of experts. In addition to heuristics, it uses reactive and deliberative strategies. And FORR includes a learning mechanism to learn appropriate configurations of experts, something which is done manually in this work, partly because the classification of Traveling Salesperson Problems is an open problem. FORR has been used for simulated robot navigation (Epstein, 1998) as an alternative to typical top-down methods with a global planner to determine a path and a local planner to determine the motor commands.

The combination of many simple heuristics instead of one monolithic strategy is supported by observations of human behavior. Tenbrink and Seifert (2011) asked participants to plan a holiday trip and analyzed verbal reports from this task. They found that humans combined spatial knowledge with knowledge about travel modalities, for example the physical effort when riding a bike. This behavior can be reproduced by GES or FORR with a combination of experts containing different kinds of knowledge.

Wiener and Mallot (2003) have also studied human route planning. They have suggested a ‘fine-to-coarse’ strategy that takes into account region knowledge in path planning, but unlike ‘coarse-to-fine’ strategies that promote a strict top-down approach, they suggest that in working memory a simplified copy of the hierarchical organization is held and modified during the planning process. In this paradigm the membership of a point to a region can be adapted according to the situation. Although in the presented work the regions are fixed, a similar effect occurs by treating regions as knowledge that can be used, combined or ignored.

For the Traveling Salesperson Problem several hierarchical solution approaches have been suggested, which were also inspired and compared to human problem solving skills. Kong and Schunn (2007) describe a classical ‘coarse-to-fine’ approach: first a set of regions is defined using k-means clustering. The centroids of the clusters define a more abstract problem, which is solved first. The points of the original problem are then added along the global path. For this kind of approach Best (2006) has proposed to use a distorted model for the distances between points, based on observations that humans underestimate intra-cluster distances and overestimate inter-cluster distances. Graham et al. (2000) present another hierarchical approach, in which the hierarchy is based on perceptual properties and the number of layers is determined by the algorithm. Also in this case the lower levels are bound to the solutions of the higher levels.

In AI hierarchical knowledge is usually also exploited in hierarchical, top-down algorithms to enhance search efficiency. Knoblock (1991) has shown that hierarchical decomposition returns optimal solutions if the subproblems are independent, but that the solution quality is still acceptable when this assumption is slightly violated, as it is in most real-world problems. The abstraction strategies in AI planning and abstraction systems, in which states are clustered into state sets, have been analyzed by Pang and Holte (2011).

In hierarchical A* search (Holte et al., 1996) solutions of abstracted versions of the problem are used as heuristics for the underlying flat A* search. The use of the abstract search result is similar to the method proposed here to use hierarchical information as additional knowledge. However, the underlying search paradigms are very different: while A* uses one heuristic function, GES is designed for a wide variety of knowledge sources and has the additional filtering of operators, while discarding the goal of reaching optimal solutions. Leighton et al. (2011) propose a hierarchical A* strategy to trade off search quality against computational efficiency, in accordance with the traditional view with the goal of optimal solutions, which sometimes has to be loosened for practical applicability.

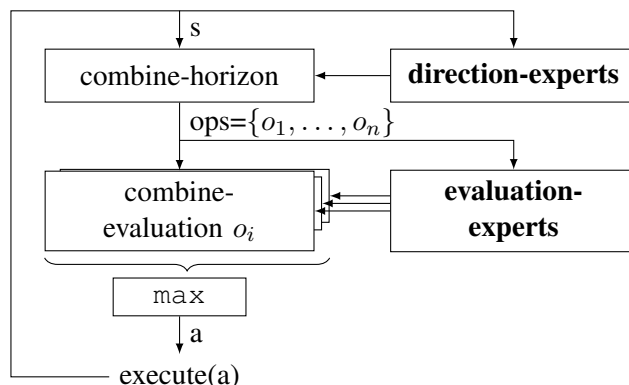


Figure 1. Schema of the GES algorithm.

3. Approach

Greedy Expert Search has been used for Traveling Salesperson Problem solving with a set of problem-dependent heuristics (Kirsch, 2011). In the following Greedy Expert Search is briefly described with the heuristics that have been used before. Three new heuristics are included for hierarchical knowledge.

3.1 Greedy Expert Search (GES)

In the standard definition of search problems as described by Newell and Simon (1972) a search problem is defined by an initial state, a goal test and a set of operators with a successor function, which returns the expected next state when an operator is applied in a specific state. In addition, a problem definition usually contains a cost function, which is replaced in GES by a set of experts.

Figure 1 shows a schema of the GES algorithm. In each search step GES determines promising operators, called *horizon*, by consulting the *direction experts*. Using the successor function, the expected resulting next states are evaluated by the *evaluation experts*. The operator that is rated most highly by the evaluation experts is executed, resulting in the next state. For the Traveling Salesperson Problem a state is a path from the Traveling Salesperson Problem start point to the current decision point and the set of still unvisited points. An operator removes a point from the set of unvisited points and adds it to the path to become the new decision point.

Beside the problem, this algorithm has as input a set of direction experts, a set of evaluation experts and combination functions for the horizons (i.e. the combination of the single results of the direction experts) and the state evaluation. The combination functions will not be considered closely in this paper. For the combination of horizons the union of the single horizons is used or a sequential combination in which a direction expert is only used when other experts have failed to return any candidate points. The evaluation combination function is a weighted sum of the results of the evaluation experts, each returning a value indicating the quality of the proposed state and a confidence (which corresponds to the weight).

3.2 Experts for Traveling Salesperson Problems

Kirsch (2011) describes a set of experts for Traveling Salesperson Problems, which are motivated by research in psychology and have been developed as simple, efficient heuristics. A short summary of the experts is given below. Three additional experts are added: two direction and one evaluation expert to integrate hierarchical knowledge, which are described in more detail. All the experts are listed in Table 1.

The most straightforward candidate points for the next expansion step are the points closest to the current point on the path (NEIGHBORHOOD direction expert). In the psychology literature strategies for following the convex hull are often mentioned (Golden et al., 1980; MacGregor, Ormerod, & Chronicle, 2000). Using the CONVEX-HULL direction expert together with the INDENTATION, INNER-POINTS and/or CHEAPEST-INSERTION evaluation experts results in a strategy that combines convex hull algorithms with the observation that humans build paths stepwise and not globally (MacGregor, Ormerod, & Chronicle, 2000). The PINWHEEL direction expert is motivated by the work of Best and Simon (2000), which proposes points that follow in a circle around the center of gravity of the Traveling Salesperson Problem.

For the evaluation of candidate points, the closeness of a candidate to the last point is one criterion (POINT-DISTANCE evaluation expert). The PROBLEM-DIAMETER expert assumes that small problems (determined by the problem diameter) are easier to solve than larger problem and thus favors points that lead to a small remaining problem. Since solutions with intersections are never optimal and also hardly produced by humans (Golden et al., 1980; MacGregor & Chu, 2011), the AVOID-INTERSECTIONS expert penalizes the addition of points that would lead to an intersection. Since intersections are often caused by early decisions, the AVOID-SPLITTING expert uses a heuristic to avoid lines through the Traveling Salesperson Problem, which would lead to two parts that would later have to be combined by an intersecting line (this strategy also favors a “going-in-rounds” or convex hull strategy). Finally, the FOLLOW-LINES expert is a simple attempt to add some gestalt heuristics by favoring the following of straight paths.

Hierarchical knowledge is added with two direction and one evaluation expert. The REGIONS evaluation expert favors candidate points that lie in the same region as the last point on the path. The direction experts return all remaining points in the current region as candidate points for the next expansion step. Different behavior is only implemented for the case when no points are left in the region. The REGION-STRICT expert simply returns no points at all. When using this direction expert, at least one other direction expert must be specified to return candidate points for the case that no points in the last region are left. In this way, the algorithm chooses the next region without hierarchical knowledge. In contrast, the REGION-PW expert always returns candidate points by using a global pinwheel strategy: when there are no points left in the current region, the next region on an imaginary circle around the center of gravity of the problem is chosen.

4. Evaluation

The goal is to examine the influence of hierarchical knowledge in a greedy, heuristic search algorithm such as GES on the solution quality. Since humans use region knowledge for solving Traveling Salesperson Problems, an algorithm using regions should be superior to one ignoring such knowl-

Table 1. Used combinations of experts. The numbers in the evaluation expert fields indicate the confidences. Subscripts on check marks indicate the order in which the direction experts are used.

	direction experts					evaluation experts								
	NEIGHBORHOOD	CONVEX-HULL	PINWHEEL	REGION-STRICT	REGION-PW	POINT-DISTANCE	INDENTATION	INNER-POINTS	CHEAPEST-INSERTION	PROBLEM-DIAMETER	AVOID-INTERSECTIONS	FOLLOW-LINES	AVOID-SPLITTING	REGIONS
nn-f	✓					1								
ch-f		✓					.5	.5	.5					
pw-f			✓			1				1			1	
nn-r	✓					1								1
ch-r		✓					.5	.5	.5					1
pw-r			✓			1				1			1	1
nn-hk	✓			✓		1			.5	.5	.5	.5		
ch-hk	✓	(✓ ₂)		✓ ₁		1			.5	.5	.5	.5		
pw-hk	✓				✓	1			.5	.5	.5	.5		
nn-ha	(✓ ₂)			✓ ₁		1			.5	.5	.5	.5		
ch-ha		(✓ ₂)		✓ ₁		1			.5	.5	.5	.5		
pw-ha					✓	1			.5	.5	.5	.5		

edge, if not in terms of tour length, at least in terms of user acceptance. But humans do not slavishly follow regions when this would be inappropriate. This is why a combined strategy using region knowledge as a special kind of heuristic can be expected to outperform a strict top-down approach. With the Traveling Salesperson Problem as an example, two hypotheses are tested:

Hypothesis 1: Using knowledge about regions improves the problem solution.

Hypothesis 2: Using hierarchical knowledge in a flat algorithm rather than a hierarchical algorithm leads to better results, in particular when the region assignment is not adequate.

Table 1 shows four classes of configurations composed of different direction and evaluation experts. Each class contains three types of configurations: nn-x configurations are mainly based on the nearest-neighbor heuristic, ch-x configurations take their knowledge mainly from convex hull heuristics, and pw-x configurations use some form of pinwheel strategy. The class xx-f contains “flat” configurations without any regional knowledge, xx-r configurations use hierarchical knowledge in the form of the REGIONS evaluation expert, otherwise being identical to the respective xx-f configuration. The xx-ha configurations emulate strictly hierarchical algorithms: the direction experts only return points of the same region as the last point in the tour. If no such point exists, the

nn-ha and ch-ha configurations use a second direction expert to guide the search to another region. pw-ha chooses the next region on an imaginary circle around the center of mass of the problem. The xx-hk configurations are identical to the xx-ha configurations but for an additional NEIGHBORHOOD direction expert returning the two closest points. This expert is used in addition to the hierarchical region experts and thus loosens the restrictions of regional knowledge.

4.1 Data Sets

Wiener et al. (2009) have studied human Traveling Salesperson Problem solving skills in a real-world setting. Their data sets contain 18 *Region-Strategy adequate (rs-adequate) tasks*, problems which can be solved optimally when all points in a region are visited before moving to the next region and 18 *Region-Strategy inadequate (rs-inadequate) tasks*, problems with optimal solutions that leave a region and re-enter it later. The problems in both data sets have 5–10 points.

In addition, 10 problems with 10 points each were generated with a uniform random spatial distribution. Using the WEKA data mining software (Witten, Frank, & Hall, 2011), the problems were clustered into three regions with the Expectation-Maximization (EM) algorithm. This data set is called in the following *clustered-regions*. These problems mostly fall into the class of rs-inadequate problems, because the clusters are only built with respect to spatial closeness, not simplification of the problem solution. Additionally a data set, *random-regions* was created from the same randomly generated problems. This time the points were assigned to one of three regions randomly, thus the region information can be expected not to be helpful for the solution process.

4.2 Tour Length Evaluation

The standard measure for Traveling Salesperson Problem solution quality is the percentage above the optimal tour length (PAO). The two pairs of configuration classes xx-f and xx-r are compared for verifying Hypothesis 1 and xx-ha and xx-hk for Hypothesis 2 for each of the four data sets.

Since the PAO values cannot be expected to be normally distributed for the algorithms, a Wilcoxon Signed Ranks test was used to assess the strength of the findings. For the case that there is indeed a normal distribution underlying the data, paired T-tests reveal similar results as the Wilcoxon Signed Ranks tests.

Hypothesis 1: Hierarchical vs. Flat Knowledge Figure 2 shows the development of the average PAO for the three configuration types (nn-x, ch-x, pw-x). Contrary to Hypothesis 1, the results do not always get better and sometimes they degrade when adding region information. None of the changes has turned out to be statistically significant, except for the change for the nn-configurations in the random-regions condition.

The left two graphs of Figure 2 show that there is a higher tendency for improvement in the rs-adequate cases. In contrast, for the rs-inadequate cases the tendency is rather a worsening of results. This is not surprising, when the regions are not helpful for finding a good solution, they can mislead an algorithm depending on region knowledge. The attempt to use knowledge about regions in this case can be compared to that of the strictly xx-ha configurations (see discussion of Hypothesis 2).

This observation is confirmed clearly in the random test set where the regions are assigned at random and thus have no meaning at all. In this case the results with regional knowledge are always

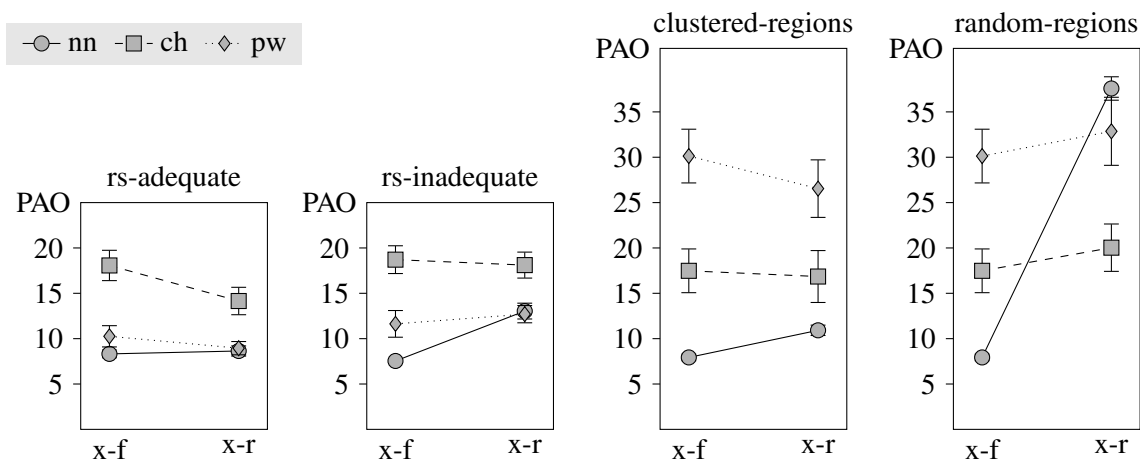


Figure 2. Comparison of tour lengths for x-f and x-r configurations.

worse, for the nn-f/nn-r configurations with statistical significance ($p = 0.005, W_+ = 0.0$). The fact that the deterioration is not too pronounced for ch-f/ch-r and pw-f/pw-r is an indication that the combination of a wider range of experts can counterbalance the region effects (for better or worse).

Figures 4 and 5 show examples of how region knowledge can influence the search result. In the rs-adequate example in Figure 4 the region information overrides the nearest-neighbor choice of the next point and leads to a better solution. In this case, the distance to the next in-region point is not much larger than that to a point in another region. In other problems, the influence of the REGIONS expert is too weak to make any difference. Figure 5 shows how region knowledge can degrade a previously good solution when the regions are not appropriately chosen.

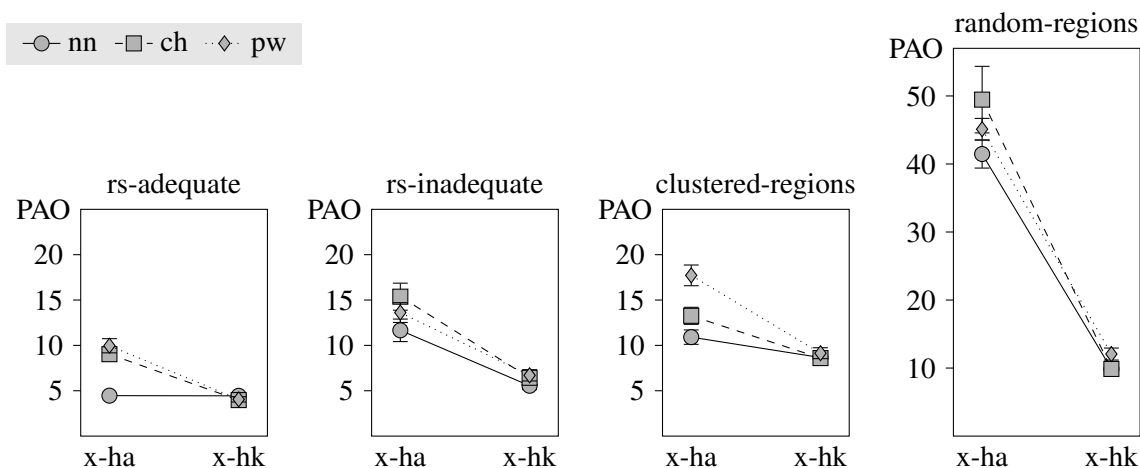


Figure 3. Comparison of tour lengths for x-ha and x-hk configurations. Note that the scale for the bad-regions data is compressed.




Figure 4. Example of two solutions for an rs-adequate problem.

Hypothesis 2: Hierarchical Knowledge vs. Hierarchical Algorithm Hypothesis 2 claims that the use of hierarchical knowledge in a flat algorithms leads to better results, in particular with a non-ideal hierarchy, than a hierarchical algorithm using the same hierarchical knowledge.

Figure 3 shows the changes in the performance when using the xx-ha configurations compared to xx-hk. In the rs-adequate data set the results of both approaches are almost the same for the nn-based heuristics, but for the ch- and pw configuration pairs a tendency towards a better result for the flat solution is visible, although failing to be statistically significant (ch: $p = 0.033$, $W_+ = 9.0$ pw: $p = 0.013$, $W_+ = 10.0$). This improvement is rather surprising since the problems were designed to work well with a hierarchical approach. It is explicable by the fact that the algorithms do not necessarily find the optimal solution and a solution not fully keeping to the region information may still be better than one that exploits this information fully. So even in situations that are well-suited to a hierarchical solution algorithm, the use of hierarchical knowledge in a flat algorithm rather improves the result and does not degrade it.

For the rs-inadequate examples, the improvement when using the xx-hk configurations compared to xx-ha is more visible, still the results are not statistically significant (nn: $p = 0.079$, $W_+ = 34.0$, ch: $p = 0.012$, $W_+ = 16.0$, pw: $p = 0.043$, $W_+ = 39.0$). However, the lack of statistical significance can also be explained by the relatively small data set and the high variability in the data. A Wilcoxon Signed Ranks test over the combination of rs-adequate and rs-inadequate data reveals a




Figure 5. Example of two solutions of an rs-inadequate problem.




Figure 6. Example of two solutions of a random-region problem.

significant improvement for the ch ($p = 0.001$, $W_+ = 49.0$) and pw pair ($p = 0.001$, $W_+ = 83.0$), but not for the nn pair ($p = 0.107$, $W_+ = 85.0$).

The results for good-region problems are similar to those of the rs-inadequate data set, as was to be expected. The results are not statistically significant, but again there is a tendency towards better results with xx-hk.

The improvement is more pronounced when the clusters are inappropriately designed in the random-regions data set. In this case the improvement of the flat solution compared to the hierarchical algorithm is statistically significant for all pairs of configurations ($p = 0.005$, $W_+ = 0.0$ for all three pairs). The example in Figure 6 illustrates the differences in the results. While the hierarchical algorithm strictly follows the given structure, the combined one can better exploit all available knowledge.

4.3 User-Centered Evaluation

Even though the examples in this paper are well-defined, small-world problems, the goal of this research is the applicability in complex, large-world domains. Since humans are well-adapted to the latter kind of environment, a small user study was performed to 1) obtain human solutions for comparison with the computation ones, and 2) get subjective ratings for the acceptability of the computational results. The latter aspect was included following the argumentation of Tak et al. (2008) that tour lengths are not necessarily the most appropriate measure for Traveling Salesperson Problems, especially when comparing with human skills.

4.3.1 Experimental Setup

In this study three of the randomly generated instances (shown in Figure 7) were used in three variants: without clusters, with three clusters found by the EM algorithm and with a random assignment of points to “regions”, which resulted in 9 problems.

The ten participants¹ received written instructions and six problems each (all three non-region problems and either the three problems with clusters or the three problems with random assignments; the two groups were equally distributed). The order of the problems was randomized and re-

1. The participants were one computer scientist and nine students of information management. One participant was female, nine were male.




Figure 7. TSPs used in the user study. Top: clustered-regions problems, bottom: random-regions problems.

regions were indicated by different symbols as in Figure 7. The participants had no information about the algorithm that may have produced the solutions. For each problem, the participants were first presented with the problem on one page and had to connect the points. After that the participants were given eight tours, which had been generated by different configurations². The participants were told that some of the solutions had been found by humans, others were computer-generated. They were asked to rate each solution along these criteria:

- How do you estimate the quality of the solution? (possible answers: good–medium–bad)
- Could this solution have been proposed by a person? (possible answers: yes–no)
- Is this the optimal, i.e. shortest, solution? (possible answers: yes–no)

4.3.2 Data Analysis and Results

Almost all participants forgot to provide answers for some given problem solutions. These were ignored and the valid answers were averaged. For the subjective answers (Figure 9) the evaluations of the non-region problems were used, because for those the most data was available. The results for the two other variants were mostly comparable, with pw-ha being rated worse in the random-regions condition.

Figure 8 compares the PAO of the solutions found by the participants with those generated by the xx-hk and xx-ha algorithm configurations. For the algorithm, the data was aggregated in the same way as if nn-hk, ch-hk and pw-hk had been three study participants.

Not surprisingly, the participants were not distracted by the bad region assignment — they seem to have ignored it. The solutions in the clustered-regions condition were slightly worse than those for the non-region and badly clustered condition (average PAO: no regions: 2.91, clustered regions: 4.36, random regions: 0.93). This difference is not statistically significant and may be due to natural variation considering the small number of problems and participants. The variance between participants is quite high, most pronounced in problem 2, where optimal solutions were proposed as well as one with 23.18% above the optimum.

2. For the non-region problems, the participants were shown the solutions obtained with clustered regions. The motivation is that an algorithm could generate the clusters with a standard algorithm and then apply the search strategies. For the other two variants the solutions of the respective region assignments were shown.




Figure 8. Average percentage above the optimum for three problems solved by participants of the user study and the x-hk and x-ha configurations. The bars with brackets indicate the best and worst result for each group.

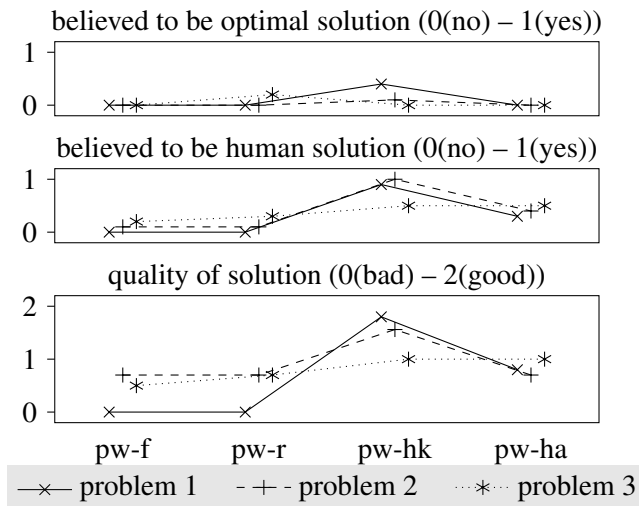


Figure 9. Subjective average ratings of the study participants.

The solutions returned by the xx-hk configurations are on average worse than the human results, but not dramatically. In most cases they can at least compete with the worst human solutions. The results returned by the three algorithm configurations (nn-hk, ch-hk, pw-hk) were in most cases identical, which can be explained by the large overlap of used experts. The solution quality for the random-regions problem variants are very similar to those for the good region assignment, which corresponds qualitatively to the human skills.

The xx-ha configurations return comparable, but slightly worse results than the xx-hk configurations in the clustered-regions condition. This can be explained by the clustered-regions data falling into the “rs-inadequate” class of problems, which means that the cluster in which the start point lies has to be visited at the beginning and end of the tour for an optimal solution and thus is not ideal for the hierarchical configurations. For the random-regions condition, the advantage of the xx-hk combinations becomes very obvious. Here the PAO values for the xx-ha configurations are about four times as high as for the xx-hk combinations and in direct contrast to the human performance, which does not degrade at all.

Figure 9 shows the outcome of the evaluations on the proposed solutions of all four configurations that use the pinwheel strategy in some way. The two groups (pw-f/pw-r and pw-hk/pw-ha) are not directly comparable since they differ largely in the evaluation experts, but they do show different degrees of including hierarchical knowledge.

The subjective quality of the solution and the belief that the solution was produced by a person correlate strongly. In both cases the pw-f and pw-r solutions were rated very low, the pw-hk solutions were rated very good and the pw-ha solutions were ranked in between. For problem 3 the solutions of pw-hk and pw-ha were rated as almost equally mediocre. Even though pw-hk was rated highly, the participants did not have a strong belief that the solution was the optimal one. This rating is similarly low for pw-hk as for the other configurations.

5. Discussion

This paper has examined different possibilities to use hierarchical knowledge in a greedy, heuristic search algorithm for one specific problem: the Traveling Salesperson Problem. The basic idea is that hierarchical abstraction may be better employed as a type of knowledge or heuristic than as a hierarchical solution approach. The efficiency in this paradigm comes from using GES, which reduces the branching factor of the search with direction experts. Hierarchical knowledge contributes to this paradigm in the form of a very valuable direction expert, which is however most effective when combined with other direction experts, slightly increasing the branching factor, but still being efficient.

In contrast to hierarchical algorithms, where abstraction rather reduces solution quality if the subproblems are not independent, hierarchical knowledge in a greedy heuristic search algorithm such as GES was claimed to improve the results (this is not directly comparable since the first approach starts from an optimal solution perspective, while GES is built on the assumption that flexibility and efficiency are more important for real-world problems). This claim could not be confirmed by only adding region knowledge as an evaluation expert. It is possible that this evaluation expert was not weighted strong enough to show an effect. With a higher weight, however, the

algorithm would resemble more the strict hierarchical approaches of the xx-ha configurations and suffer from the same problem of dependence on good region definitions. Still, when regarding the differences between the xx-r and xx-hk configurations, also with respect to subjective evaluations, a clear improvement is visible. Since the xx-r and xx-hk configurations are not directly comparable, this should be examined more closely in future research.

The second hypothesis, claiming that hierarchical knowledge in combination with other knowledge in a flat algorithm is more effective than a strict hierarchical solution approach, could be confirmed for the Traveling Salesperson Problem instances used, most clearly for the random-regions condition with distracting region information. The approach of combining hierarchical and other knowledge does not rely strongly on the quality of the region assignment and is thus a lot more robust than hierarchical algorithms, even if the regions are not assigned badly deliberately, but also if the hierarchy generation is not adapted for the particular problem as in the clustered-regions data set.

For route planning in general the construction of regional knowledge has still not been fully understood. Even though the approach proposed here is relatively robust against bad region assignments, it can still not compete with human skills. A problem in the region assignment with standard clustering algorithms is that these are not developed for route planning. For example for problem 1 in Figure 7 a regionalization into a left and right “wing” would be more helpful than the linear arrangement of clusters generated by EM. Possibly a tighter coupling of clustering and problems solving can remedy this. Also the assignment of points to regions may change during the course of the planning process.

6. Conclusion

The ability to form and use abstract representations is a core feature of human intelligence and thus an important topic for artificial cognitive systems. Many approaches integrate abstraction in a strict top-down fashion. The work presented in this paper has shown that this approach suffers from a lack of flexibility, especially when the abstraction is imperfect (as is often the case). An alternative view of using abstraction as a special kind of heuristic is proposed. The results in this paper indicate that, at least in the Traveling Salesperson Problem domain, the most promising approach is to use hierarchical information as a search-guiding heuristic in combination with other heuristics.

The presented work is intended to open new directions for getting adequate solutions for real problems. In the end, the motivation for most computational problems is to assist people in some task of their life. Heuristics for the chosen example, the Traveling Salesperson Problem, are not completely understood (as heuristics of human problem solving in general (Gigerenzer & Gaissmaier, 2011; Shah & Oppenheimer, 2008)), so that the solutions measured in tour length are not yet competitive to human performance or other computational methods. But the approach is more general than other algorithms used for Traveling Salesperson Problem solving. It offers the possibility to solve a wide range of problems and problem variants. For example, a shopping tour is almost a Traveling Salesperson Problem, but has additional constraints such as the weight or the risk of deterioration of the goods to buy. GES is a first step towards a general problem-solving framework that needs to be explored further and with different applications.

Acknowledgments

With the support of the Technische Universität München — Institute for Advanced Study, funded by the German Excellence Initiative, and of the Bavarian Academy of Sciences and Humanities. (?)

References

- Best, B. J. (2006). Cognitive approaches to the traveling salesperson problem: Perceptual complexity that produces computational simplicity. *Papers from the American Artificial Intelligence Association (AAAI) Spring Symposium*.
- Best, B. J., & Simon, H. A. (2000). Simulating human performance on the traveling salesman problem. *Proceedings of the Third International Conference on Cognitive Modeling*.
- Cassimatis, N. L. (2012). Human-level artificial intelligence must be an extraordinary science. *Advances in Cognitive Systems, 1*, 37–45.
- Epstein, S. (2004). Metaknowledge for autonomous systems. *AAAI Spring Symposium on Knowledge Representation and Ontology for Autonomous Systems*.
- Epstein, S. L. (1998). Pragmatic navigation: Reactivity, heuristics, and search. *Artificial Intelligence, 100*, 275–322.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology, 62*, 451–482.
- Golden, B., Bodin, L., Doyle, T., & Stewart, W. (1980). Approximate travelling salesman algorithms. *Operations Research, 28*, 694–711.
- Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition, 28*, 1191–1204.
- Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science, 3*, 275–310.
- Holte, R., Perez, M. B., Zimmer, R. M., & Macdonald, A. J. (1996). Hierarchical A*: Searching abstraction hierarchies efficiently. *Proceedings of the National Conference on Artificial Intelligence*.
- Kirsch, A. (2011). Humanlike problem solving in the context of the traveling salesperson problem. *AAAI Fall Symposium on Advances in Cognitive Systems*.
- Knoblock, C. A. (1991). Search reduction in hierarchical problem solving. *Ninth International Conference on Artificial Intelligence*. AAAI Press.
- Kong, X., & Schunn, C. D. (2007). Global vs. local information processing in visual/spatial problem solving: The case of traveling salesman problem. *Cognitive Systems Research, 8*, 192–207.
- Leighton, M. J., Ruml, W., & Holte, R. C. (2011). Faster optimal and suboptimal hierarchical search. *Proceedings of the Fourth Symposium on Combinatorial Search*.
- MacGregor, J., Ormerod, T., & Chronicle, E. (2000). A model of human performance on the travelling salesperson problem. *Memory & Cognition, 7*, 1183–1190.
- MacGregor, J. N., & Chu, Y. (2011). Human performance on the traveling salesman and related problems: A review. *The Journal of Problem Solving, 3*.

- Newell, A., & Simon, H. (1972). *Human problem solving*. Upper Saddle River, New Jersey: Prentice Hall.
- Pang, B., & Holte, R. C. (2011). State-set search. *Proceedings of the Fourth Symposium on Combinatorial Search*.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*, 207–222.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*, 129–138.
- Tak, S., Plaisier, M., & van Rooij, I. (2008). Some tours are more equal than others: The convex-hull model revisited with lessons for testing models of the traveling salesperson problem. *The Journal of Problem Solving*, *2*.
- Tenbrink, T., & Seifert, I. (2011). Conceptual layers and strategies in tour planning. *Cognitive Processing*, *12*, 109–125.
- Wiener, J., Ehbauer, N., & Mallot, H. (2009). Planning paths to multiple targets: memory involvement and planning heuristics in spatial problem solving. *Psychological Research*, *73*, 644–658.
- Wiener, J. M., & Mallot, H. A. (2003). ‘Fine-to-coarse’ route planning and navigation in regionalized environments. *Spatial Cognition and Computation*, *3*, 331–358.
- Witten, I. H., Frank, E., & Hall, M. A. (2011). *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann. Third edition.