Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 8 (2007) 192–207

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Global vs. local information processing in visual/spatial problem solving: The case of traveling salesman problem

Action editor: Fabio del Missier

Xiaohui Kong, Christian D. Schunn *

Learning Research and Development Center, University of Pittsburgh, PA 15260, United States

Received 21 February 2007; accepted 11 June 2007
Available online 3 July 2007

Abstract

Human visual/spatial problem solving often requires both global and local information to be processed. But the relationship between those two kinds of information and the way in which they interact with one another during problem solving has not been thoroughly discussed. In the particular setting of solving the traveling salesman problem (TSP), we investigated into the relative roles of global and local information processing. An experiment was conducted to measure the importance of global information and the possible constraints of global information processing on search. A model was built to simulate human TSP performance and was used to investigate further the relationship between global information processing and local information processing. Our model was compared with the human data we collected and with other models of human TSP solving.

Keywords: Visual/spatial problem solving; Traveling salesman problem; Global/local information processing

1. Introduction

Visual/spatial information in the world exists at many grain-sizes. One strategy for studying visual/spatial processing is to focus on a particular grain-size. But, as suggested by Gestalt psychologists, information processing at one grain-size is often influenced by more global patterns (gestalts) at the next grain size up. From an information processing perspective, core questions about global–local information interactions concern storage and processing limitations: (a) how can global patterns be stored efficiently to effectively influence processing of local information?, and (b) how can global patterns be constructed without first invoking complex or large scale local pattern processes?

In this paper, we test two broad predictions about the nature of global information that is stored and used to influence local information. The first prediction is that global information consists of low spatial frequency information because it is easily processed in peripheral vision and because it contains few bits of information that are processed quickly. This prediction is in contrast to global models that use contour information or isolated feature maps, which can contain many, many bits of information.

The second prediction is that the stored information must be sufficient to usefully guide local search. In other words, the global information must be generally effective for reducing the scope of local search throughout visual problem solving. This prediction is in contrast to models that emphasize contour maps because the contour provides guidance for only the points near the contour rather than generally throughout local search.

We examine these hypotheses in the context of the traveling salesman problem because it is a well-studied problem that highlights the importance of integrating global visual/spatial information into local search.

The (Euclidean) traveling salesman problem (ETSP) is to find a path of minimum Euclidean distance between

* Corresponding author. Tel.: +1 412 624 8807.
E-mail addresses: xik2+@pitt.edu (X. Kong), schunn@pitt.edu (C.D. Schunn).

1389-0417/$ - see front matter
doi:10.1016/j.cogsys.2007.06.002

points in a plane, which includes each point exactly once and returns to its starting point. As a NP-hard combinatory optimization problem, the traveling salesman problem (TSP) is believed to be "intractable" in computer science for large inputs as long as exact optimal path is concerned.

Recently there has been some interest in the traveling salesman problem in cognitive science because humans are able to easily find near-optimal solutions of small size ETSP with little apparent effort (MacGregor & Ormerod, 1996). This result seems to be contradictory with the general assumption that a large problem space implies a difficult problem (Newell & Simon, 1972), since even a small size ($n < 30$) ETSP has a huge problem space. Simple trial-and-error search in the original problem space could not explain human performance on this problem.

Several hypotheses have been proposed to account for the human strategy on TSP solving from either a quantitative explanation (MacGregor & Ormerod, 1996; Victers, Lee, Dry, & Hughes, 2003; Van Rooij, Stege, & Schactman, 2003) or qualitative modeling perspective (Best, 2005; Graham, Joshi, & Pizlo, 2000; MacGregor, Ormerod, & Chronicle, 2000). Although different explanations and models use different measures and heuristics, the global information processing vs. local information processing has been one of the central issues under discussion.

Ormerod and Chronicle (1999) first provided support for the hypothesis that human cognition is capable of perceiving and utilizing global information in the identification of TSP solutions. MacGregor et al. (2000) then proposed a model using the convex hull as the global information and developing a TSP path from it, where the convex hull is the smallest convex containing all the points in it. In the model by Graham et al. (2000), several layers of global information are perceived and developed in a cascade to approximate the final solution, which they called them a pyramid.

By contrast, Van Rooij et al. (2003) argued that purely local search based on a nearest neighbor approach could help to form some kind of global clustering information by eliminating the majority of potential intersections. However, some recent studies by Best (2005) suggest that after a global information-processing phase, human participants only perform local search in the rest of the TSP solving procedure.

Two important questions about global information used in human TSP solving remained unaddressed. First, what kind of global information is perceived and utilized? Second, how important is the global information? To answer those questions, we conducted an experiment, built a model, and evaluate the fit to human data of this model against other models.

2. Previous models of TSP

We begin by considering the rest of models previously proposed for human performance on the TSP.

2.1. Nearest neighbor

The most basic model of TSP is the nearest neighbor model (Rosenkrantz, Stearns, & Lewis, 1977) in which the problem solver always selects the closest next point to the current point, i.e., simply following a hill-climbing heuristic. The model is elegant in that it only assumes a single heuristic that is already known to be part of the human information-processing repertoire (Newell & Simon, 1972). However the model makes no use at all of global information and tends to produce solutions that are not as good as those found by humans (Ormerod & Chronicle, 1999).

2.2. Convex hull

The next simplest model of TSP is the convex hull model (Golden, Bodin, Doyle, & Stewart, 1980), which assumes that people compute a traversal around the perimeter points, including inner points opportunistically along the way using a minimal insertion rule. The global information used by this model is the Convex Hull contour, which may be rather complex, and thus require significant working memory. The minimal insertion rule is applied globally at each point in time during path computation, and points are added that cause the smallest increase in total path length. It is somewhat implausible that people would be able to compute these minimal insertions (a local processing task) at the global level.

2.3. Sequential convex hull model

MacGregor et al. (2000) adapted the convex hull model to more plausible incremental local search version of the convex hull model. In support of this adaptation, they found that humans perform better on problems with fewer interior points within the convex hull (MacGregor & Ormerod, 1996). Second, their experiments provided support for their hypothesis that human participants are sensitive to global information (Ormerod & Chronicle, 1999). We would call this model the sequential convex hull model. The outline of the model is as follows (MacGregor et al., 2000):

1. Sketch the connections between adjacent boundary points of the convex hull.
2. Select a starting point and a direction randomly.
3. If the starting point is on the boundary, the starting node is the current node. The arc connecting the current node to the adjacent boundary node in the direction of travel is referred to as the current arc. Proceed immediately to Step 4. If the starting point is not on the boundary, apply the insertion rule to find the closest arc on the boundary. Connect the starting point to the end node of the closest arc, which is in the direction of travel. This node becomes the current node.

4. Apply the insertion criterion to identify which unconnected interior point is closest to the current arc. Apply the insertion criterion to check whether the closest node is closer to any other arc. If not, proceed to Step 5. If it is, move to the end node of the current arc. This becomes the current node. Repeat Step 4.
5. Insert the closest node. The connection between the current node and the newly inserted node becomes the current arc. Retaining the current node, return to Step 4 and repeat Steps 4 and 5 until a complete tour is obtained.

2.4. Pyramid model

Graham et al.'s model (2000) of traveling salesman problem was inspired by a hierarchical architecture of human visual and spatial perception. Their model first Gaussian-blurs the original set of points into a variety of degrees and stores those blurred images in different layers of hierarchy with the most blurred image on the top. The more blurred images serve as the global information for the less blurred images. Each layer directly guides the next layer below it each time the model develops a node into the path. So layers in the hierarchy change in a repeatedly cascaded process. The pyramid model computes TSP solutions in the following steps:

1. Gaussian-blur the original n-points TSP image into $k-1$ different degrees and store them in a k-layer pyramid with the original TSP image on the bottom and the most blurred image on the top.
2. Calculate L_i modes of the image in each layer i. Consider those modes in each layer as nodes in a reduce-sized TSP problem. The top layer has 3 nodes and the bottom layer has n nodes. Layer k has $\frac{n}{b^k}$ nodes. (The parameter b is the reduction ratio. Bottom layer is layer 1.)
3. Layer n (top layer) has 3 nodes and forms a unique tour.
4. Generate a tour of the TSP in each layer by inserting them into the tour on the previously higher layer with the following rules: (a) Sort the intensity level of the mode locations in each layer. (b) Insert these modes into the tour in descending order of their intensity, so as to produce the minimum increase in tour length. Repeat Step 4 until the algorithm generates a tour in the bottom layer.

2.5. Global/local TSP solver

In the global/local TSP solver (Best, 2004), global information-processing and local information-processing phases are clearly separated. The outline of the GL-TSP solver is as follows.

2.5.1. Global information-processing phase

Using the CODE theory of human perceptual clustering (Compton & Logan, 1993; Van Oeffelen & Vos, 1982) by setting a threshold parameter, a certain number of clusters of points were generated. Global planned path among clusters is calculated using the convex hull heuristic.

2.5.2. Local information-processing phase

Starting from the current point, the next point to visit on the exact path is chosen from the current cluster by using a 6-points look-ahead rule and the global planned path as constraints.

The global information posited by this model is much smaller in size than that posited by the convex hull models.

In all but the nearest neighbor model, some kind of global information processing was engaged and extracted global information was used to guide the local search. The sequential convex hull model used the convex hull and convex contour; the pyramid model used layers of pyramid; the GL-TSP solver used the clustering result and planned path between clusters. So what kind of global information is perceived and utilized by human cognition, due to its own constraints? How important is the global information? To answer those questions we did the following experiment.

3. Experiment

3.1. Method

3.1.1. Participants

Twenty eight graduate students participated in the experiment.

3.1.2. Materials and methods

The materials were 20 TSPs. Ten are real world problems borrowed from TSPLIB (http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/) ranging in size from 16 points to 100 points. Those real world problems are generally more structured. An example would be the cities on a map where they tend to form some dense clusters (Fig. 1a). The remaining 10 were randomly pregenerated according to a uniform distribution ranging from 10 to 80 points. Fig. 1b shows an example of randomly generated TSP. Note that all participants saw the exact same 20 TSP problems but in a random order, which allows us to examine how well the models predict the performance on particular TSP problems rather than just general trends for the effect of number of points.

The problems were displayed in an 800 * 800 pixels window on a 17-in. computer screen with resolution 1440 * 900 pixels. Participants sit about 17–20 in. away from the computer screen. So all the problem lies in the human visual field with maximum angle of 10–13°. Participants were asked to find the shortest possible path and indicate the path using mouse-clicks. The program recorded all the click data. Participants were randomly

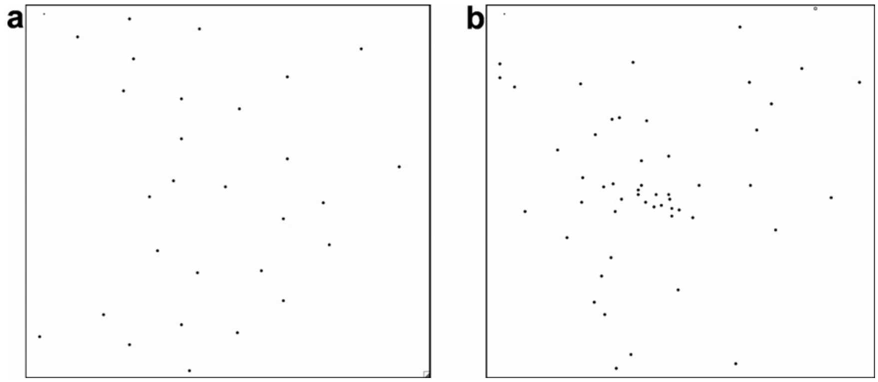

Fig. 1. Examples of a natural TSP (right) and a random TSP (left).

assigned into one of the following three groups. The groups were designed to examine the influence of the global and local information.

3.1.3. Control (10 participants)

Each participant was asked to solve the TSP problems while all point locations remained on the screen throughout.

3.1.4. Global preview (9 participants)

Each participant was asked to solve the same TSPs as in the control condition, with three distinct phases for each TSP.

1. The full TSP is shown, but paths cannot yet be clicked. Each participant was given a pen and a piece of paper to draw the global information they would need in the later phases. Participants were also asked to pick a start point to begin their TSP trip (Fig. 2a).
2. The TSP problem points were clustered into 5–12 clusters using a K-Means algorithm (MacQueen, 1967). The *k*-means centroids (geometric centers) were displayed as larger dots. Participants were asked to pick a path through just the centroids to determine the order in which the clusters show up in phase 3 (Fig. 2b).
3. All points were hidden. Then subsets of points were presented one cluster at a time, and participants had to pick a path through all the points within a cluster. When all the points in the current cluster were visited, the next cluster of points would become visible (Fig. 2c).

3.1.5. No global preview (9 participants)

This condition was identical to the Global Preview condition in that only one cluster worth of points is shown at a time during the path selection process (i.e., phase 3), except that participants did not first see the full set of points (i.e., phases 1 and 2 were skipped). So there was no global information available during any part of the process.

The No Global Preview vs. Global Preview comparison tests the effect of access to global information on local search quality, and whether the global information can fit in working memory (as opposed to it being important that global information be externally available). If those two conditions do not differ in solution performance, then the control condition assesses whether even continuously available global information is helpful. If the Global Preview and No Global Preview conditions do differ, then the Control condition assesses whether to what extent continuously available global information further shapes local search.

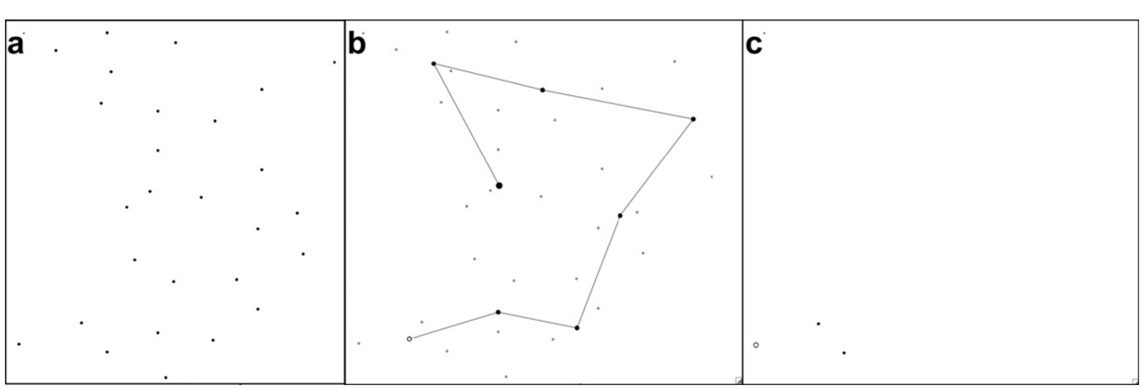

Fig. 2. Illustrations of Global Preview condition phase 1, 2, 3 (left, mid, right).

Finally, the Control condition also provides baseline TSP problem solving data against which the computational models can be compared.

After each participant finished all 20 TSPs, there was a post-experiment measurement on how fast the participant clicked the mouse. This step involved re-presenting all 20 TSPs, but instead of finding the shortest path, participants were asked to click through all the points as fast as possible in an arbitrary order. From this data, we will estimate participants' thinking time by subtracting mouse-clicking time from solution time.

3.2. Results and discussion

Accuracy and reaction time were calculated as our measurements of performance. Accuracy was calculated as the ratio of participant path length over the optimal path length. Reaction time was calculated as difference between the time to finish the TSP and the time to click through all the points. So accuracy is a number larger than 1. The closer the value is to 1, the better the performance is. The reaction time is an approximation of participant thinking time.

ANOVAs on accuracy and reaction time revealed significant effect on both accuracy ($p < .0001$) and on reaction time ($p = .0001$). But the condition effect of reaction time ($f = 9.0$) is much weaker than that of accuracy ($f = 172.1$), while they have the same degree of freedom. The control group had the best accuracy ($\overline{X} = 1.05$) but highest RT ($\overline{X} = 76$ s). The global preview group had middle levels on both (accuracy = 1.11, RT = 54 s). The No Global preview group had the worst accuracy ($\overline{X} = 1.16$) but fastest RT ($\overline{X} = 42$ s). Post-hoc Tukey comparisons found significant pair-wise difference between all groups on accuracy ($p < .0001$). That the control condition is significantly slower than the no global preview condition ($p < .0001$) and the global preview condition ($p = .0074$) suggests that processing global information does take time. That the condition effects are very strong on accuracy ($f = 172.1$) and much weaker on RT ($f = 9.0$) suggest that a simple speed-accuracy tradeoff could not explain the overall condition effect.

As it can be seen in Fig. 3, the accuracy of each group slowly goes up (less accurate) when the size of the problem goes up. The accuracy of control condition fits well to a linear trend ($R^2 = .689$). The accuracies of the other two conditions basically follow linear trends ($R^2 = .33$ and .21). The control group has the highest accuracy performance. This result is consistent with our hypothesis that human participants utilized both global information and local information to solve the problem. The control group has all points visible during the entire problem solving procedure; the points on the screen appear to help them to retain the global information through some kind of active memory during the solution process.

The global preview group has better accuracy than the group w/o global preview. This result confirmed the importance of global information in the human TSP solving procedure. A typical example of the scratch notes of participants in the global preview group is in Fig. 4. Not all participants sketched a Spline-curve. Some participants just recorded the relative position of each cluster and some just left the scratch paper blank. But when connections between clusters were drawn, they tended to resemble splines.

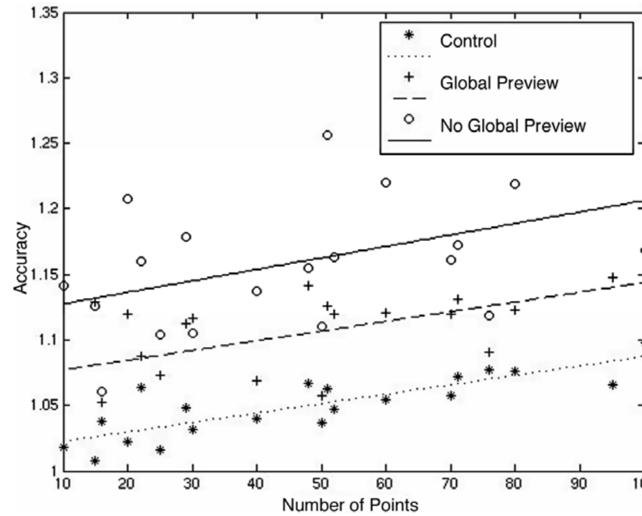

Fig. 3. Mean accuracy of each individual problem within each condition.

In sum, it appears that global information stored only mentally does help local search. Global information presented throughout problem solving helps even more. Thus, global information computed in global preview condition, either slightly exceeds capacity limits, and/or is not stored with the same fidelity as global information that is supported with continual visual input. One could interpret the results as a support that human TSP problem solving relies heavily on compact global representations suggested by the notes of global preview condition. However, one

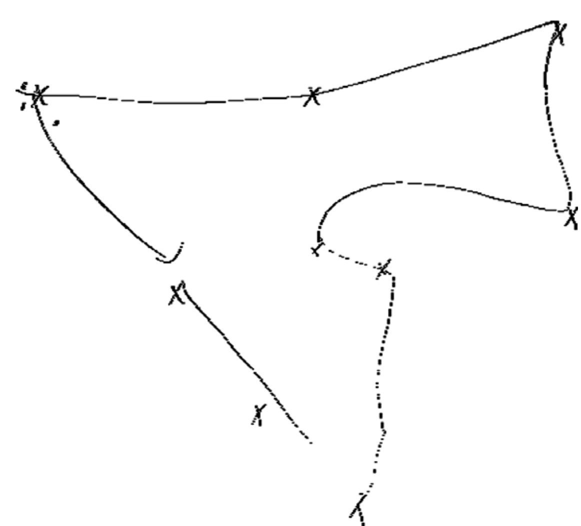

Fig. 4. A typical scratch note from a participant in the group with global preview.

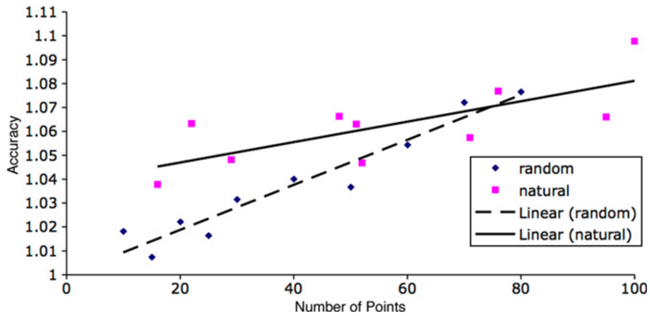

Fig. 5. Effects of problem size on accuracy separately for random and natural problems.

could also argue that human TSP may use elaborate global information that is dependent upon constant peripheral visual input, since control condition has better accuracy than global preview condition. So the question is whether the recorded global information in global preview condition is close to the one used in control condition or it's only an abstract of it. Precise modeling of the exact human data in the control condition may help to resolve this.

3.2.1. Natural vs. random

We used two categories of traveling salesman problems in our experiment. Ten problems were randomly generated according to uniform distribution (random problems). The other 10 were borrowed from TSPLIB (natural problems), most of which are data collected from real world (such as cities in a country, etc.).

The overall accuracy performance of the control group between natural problems and random problems is the only significant difference we found ($p < .001$). There is no significant difference between natural points and random points on the other two groups' accuracy ($p > .5$) or RT ($p > .8$) values.

The control group has better accuracy on random problems than natural problems and there is an interaction with number of points (Fig. 5). For the natural TSP problems, the accuracy value (the larger the worse) was higher than random problems for small problem size. As problem size increases, accuracy value of natural problem increases at a slower rate than that of random problem. That this effect only appeared in the control group suggests that the added structure contained in natural problems is subtle and not easily stored in working memory.

4. TSP global vs. local information processing

Having conducted an experiment to test the role of global information in TSP performance, we now turn to a more detailed theoretical examination of global vs. local information processing in TSP models. More advanced existing TSP models used a certain kind of global information to guide their local information processing. In the sequential convex hull model, the local information processing of individual nodes relies on the global information of the convex hull contour developed through Step 1 to Step 3 and the previous iterations of Steps 4 and 5. In the pyramid model, the global information is contained in all layers of the pyramid except the bottom layer where the exact node locations are stored. The top-to-bottom cascade iteration of the pyramid model is to specify the global information layer by layer until a tour of the original TSP is constructed on the bottom.

A recent study provided support for the hypothesis that separated processing phases of global and local information are an important characteristic of human TSP solving (Best, 2005). This naturally leads to the question of what kind of global information is used in the global information-processing phase? There are several criterions that global information should meet:

1. The global information should be sufficient to guide the local search. It should remain fixed or change only locally in the local search phase.Since the global information processing and local information processing are in two separate phases, the global information is basically fixed after the global information-processing phase. Therefore, the global information should be sufficient to guide the local information processing. Local information processing should only help to maintain or slightly change the local part of the global information.
2. The global information should be compact in size to fit well in active human visual/spatial model representations.

Since the global information is processed during the first several seconds of the problem solving procedure, it should be compact in size. Also it should have the form of the representation that fits human working memory well, so that it would remain in the working memory during the local information-processing phase.

In the model developed by MacGregor et al. (2000), the convex hull contour serves as the global information. In the global information-processing phase, the original convex hull is perceived as the global information. But this convex hull does not contain enough information to guide each local search, so the convex hull contour is modified each time in the local search phase. The modification of the contour could change it a lot. Before connecting to the next node, the model would need to generate some number of temporary edges that constitute the temporary state of the convex hull contour. Those temporary edges are invisible to human participants, so they have to be stored in visual/spatial working memory. But sometimes the number of temporary edges could be quite large and keeping all those edges in visual/spatial working memory may not be feasible. So the convex hull itself does not fit to our criterion of sufficiency, and the convex hull contour does not fit to our criterion of compactness.

Fig. 6 visualizes the convex hull contour used by the sequential convex hull model as the global information in

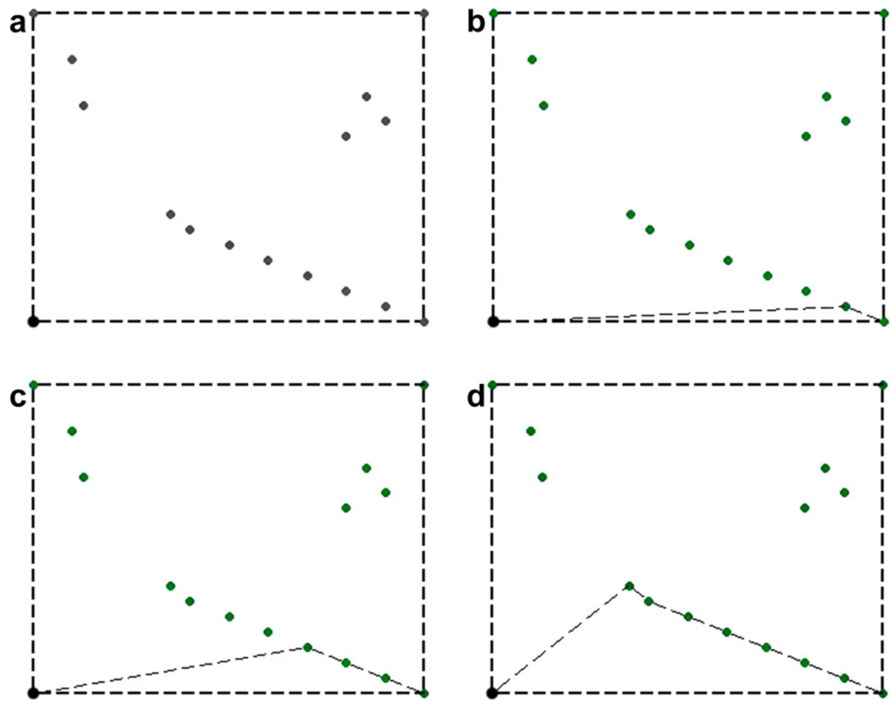

Fig. 6. The global information used by the sequential convex hull model.

some steps of the model execution. Fig. 6a shows the convex hull of this simple artificial TSP. The rectangle resembles the convex hulls. Now suppose the model chooses the lower-right corner as the starting point and picks counterclockwise as the direction to travel. Fig. 6b describes the convex contour after the first iteration of the model when a node is inserted into the contour using the cheapest insertion criterion. Correspondingly, Fig. 6c is the convex contour after the third iteration and Fig. 6d is the convex contour after the seventh iteration. Notice the model has to keep all those dashed lines in working memory, since none of those edges has been actually drawn at this stage. There could be a larger number of them in a more complicated problem.

In the model developed by Graham et al. (2000), the upper part of the "pyramid" (all the layers except the base layer) serves as the global information. Since the upper part of the pyramid may contain many layers and those layers need to be retrieved and manipulated during problem solving process, this might be too much information for human working memory to carry. So the pyramid does not fit well to the second criterion of global information.

Fig. 7 illustrates the solution process of the Pyramid model. Fig. 7a–e visualizes the global information kept in

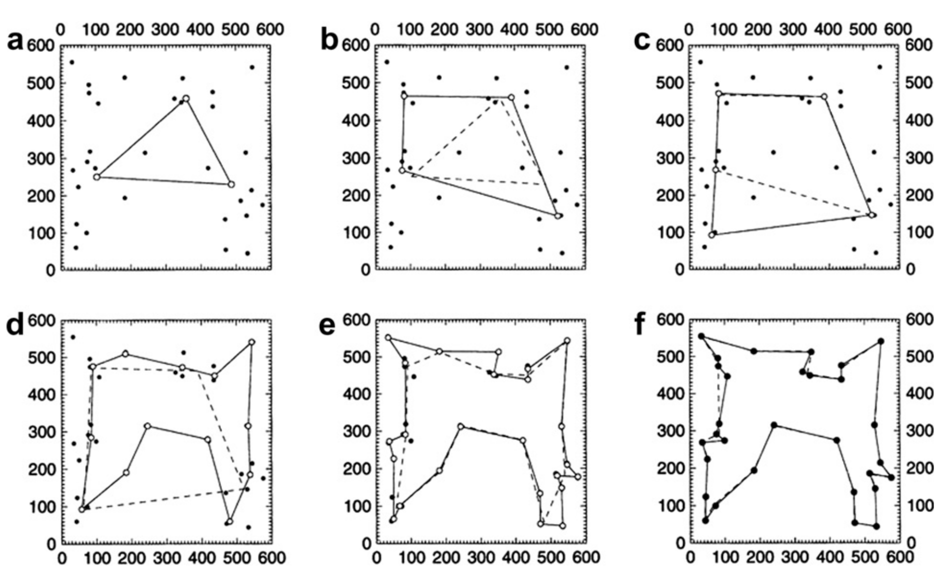

Fig. 7. Global information in different layers of the pyramid. Adapted from Graham et al. (2000).

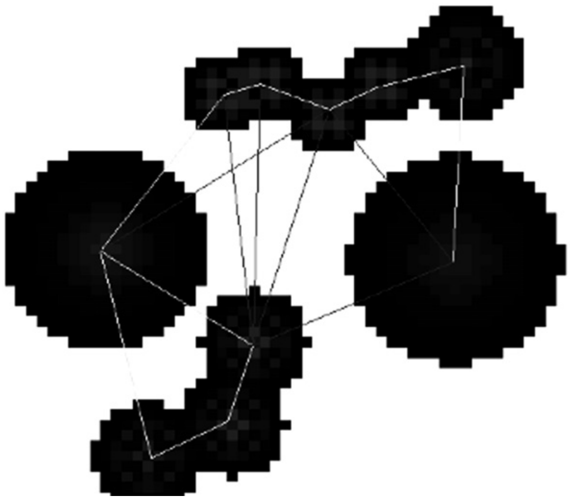

Fig. 8. Global information used by GL-TSP solver. Adapted from Best (2004).

various layers of the pyramid and how they were developed in a cascaded manner. In each of those figures, the empty dots and the lines connecting them are actually invisible to human subject and needs to be loaded in spatial/visual working memory until the contour in the adjacent lower layer is developed. So the amount of global information could be quite large and hard to be kept in verbal or visual working memory (Cowan, 2001; Pylyshyn, 1989).

In the GL-TSP solver, the global information is the clustering results and the global planned path. Fig. 8 shows that the global information the GL-TSP solver generated on a particular TSP problem after its global information-processing phase. The clustering information is three dimensional with lighter gray represents higher z-coordinate. Although the GL-TSP solver successfully characterized the separated global local information processing, its global information may be too much to carry in human spatial/visual working memory during the entire solution procedure (Cowan, 2001; Pylyshyn, 1989).

5. K-Means TSP model

Based on this theoretical analysis and observations of human behavior in the global preview condition, we propose a new model for TSP problem solving.

Our K-Means TSP model is based on the following three steps:

1. Clusters are identified.
   In this step, points are grouped according to visual density. Points constructing a higher visual density are more likely to be grouped together.
   Our model approximates this clustering identification process using a K-Means clustering algorithm, because it is available in standard software packages. The K-Means clustering algorithm clusters N data points into K disjoint subsets S_j containing N_j data points so as to minimize the sum of squares criterion:
   $$J = \sum_{j=1}^{K} \sum_{n \in S_j} |x_n - \mu_j|^2,$$
   where x_n is a vector representing the nth point and μ_n is the geometric centroid of the points in S_j. Now the original problem is reduced to the problem to find the shortest path among all μ_n.
2. A sketch of the path is conceived.
   Here by sketch of the path, we mean the path visiting all the groups and returning to the starting group. Using this strategy, human cognition reduces the original problem to a main problem of much smaller size with simple sub-problems. Here we use a Spline-curve of all the centroids to model this sketched path.
3. Connect all the points along the sketched path.
   We model this step using a projection rule. We project all the points to nearest point on the Spline-curve. Then we construct the final solution by connecting all the points in the same order as their projection on the Spline-curve.

Steps 1 and 2 of our model are the global information processing part, and Step 3 is the local information processing part. The global information perceived in Steps 1 and 2 will guide the local information processing in Step 3. The Spline-curve is the global information developed after Steps 1 and 2. The clusters and centroids are no longer needed after the Spline-curve is sketched. So in the local search phase, the cluster and centroids information can be discarded, since the Spline-curve itself is enough to guide the local information processing in Step 3.

The Spline-curve plotted fits both of our criteria for global information. First, it is sufficient to guide the local search in the third step of the model, where the model only need to project the points onto their nearest curve. Second, because clustering result and centroids information can be discarded after Step 2, the Spline-curve itself is compact in size and has a visual representation that may fit well to human visual/spatial working memory capacity.

Our hypothesis is that there are some visual operators for human cognition that enable it to do the first two steps within a near constant time and the third step in a linear time. Fig. 9 illustrates the three steps of our model when solving a 70-points TSP.

5.1. Model simulation

We used a fixed the k-means centroids in the upper right plot set of 20 problems across participants in our experiment. The negative consequence of this experimental design choice is that we do not have a pure estimate of the effect of problem size because of the small idiosyncrasies of our chosen problems. However, the positive consequence is that we have enough data for each exact problem to evaluate how well each model can explain performance on those particular problems, in trends across problems and exact fit to problem performance.

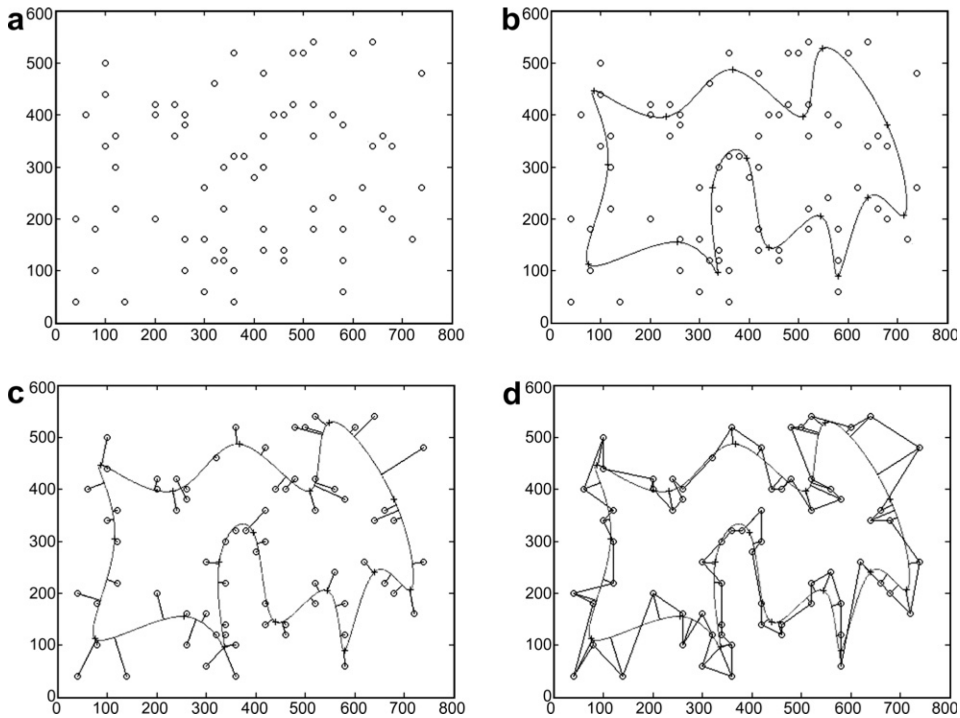

Fig. 9. Three steps of the Kmeans-TSP model on solving a 70 points TSP. "+" indicate the locations of the Kmeans centroids in the upper right plot.

The number of clusters is the only parameter setting in this model. In our simulation we set it to $\#Clusters = 2 \times \sqrt{\#points}$. This setting was based on the intuition that we do not want points to be too far away from its cluster centriod to avoid too much error. If recursion has to happen for large TSPs, we want the depth of recursion to be no more than one. In order to draw the Spline-curve around all the centroids in step two of the model, we need a TSP path around all centroids. In our simulation we used a recursive call to our model until the size of the problem is below 6, when we can easily use an exhaustive search function to find the shortest path around the remaining points. For our current set of problems, the depth of recursive calls is at most two. Since the K-Means clustering algorithm may converge to a local minimum and may yield different clustering results on different runs, we ran through our model on all the 20 problems 40 times.

The mean accuracy of the 40 runs is pretty close to the accuracy performance of participants in the global preview group. Since our model employed a naïve local search strategy and the group with global preview had incomplete local information, the closeness of accuracy between them is what we expected. However, the minimum accuracy generated by the 40 runs of our model is very close to the accuracy performance of the control group. The reason might be sometimes the naïve projection rule in our local search fits to the generated global information very well, so it produced a similar result as the more sophisticated local search strategy used by human cognition.

6. Model evaluation

To evaluate our model in depth, we did a comparison between our model and other models against the human performance in the control condition along four dimension of TSP behavior across the 20 TSP problems used in our experiment. We focused on the control condition because that condition best represents full use of global information. We calculated three results on the set of solutions that each model produced on the twenty traveling salesman problems used in our experiment: number of intersections, mean accuracy, standard deviation of accuracy and exact path chosen. Both trends and exact values are important measurements on how well a model fits human performance data (Schunn & Wallach, 2001). Pearson correlations between model and human data are used to measure fit to trends; average signed errors are used to measure fit to exact value.

We compared our model with the following models and heuristics: the Pyramid model[1] (Graham et al., 2000), Sequential Convex-Hull model (MacGregor et al., 2000), Nearest Neighbor and Convex Hull. Table 1 summarizes the global and local information strategies used by each of the models.

[1] The code for the pyramid model was downloaded from http://www2.psych.purdue.edu/tsp/workshop/downloads.html in November 2005. The code is an improved version of the model described in Graham et al. (2000).

Table 1
Global and local information/strategies of different models

	Global information processing strategy	Global information	Local information processing strategy
Nearest neighbor heuristic	No global information processing	No global information	Find nearest point to the current point
Convex hull heuristic	Find then develop the convex hull contour	Convex hull contour	Jumping around the contour edges to insert the points that will yield the minimum increase in total length. (Search through all points.)
Sequential convex hull model CHSQ	Find then develop the convex hull contour in a clock-wise or counter-clockwise sequence	Convex hull contour	Apply the insertion rule to the current edge. (Search through all points close to the current edge.)
Pyramid model	Apply Gaussian filters to build a pyramid. Use insertion rules to update the pyramid	Pyramid: A hierarchy of convex contours	No local information processing
Kmeans model	Generate Kmeans cluster centroids Sketch Spline curve around the cluster centroids	Spline-curve	Project points onto their nearest Spline-curve

6.1. Number of intersections

For human and models, we computed the number of times the selected final path crossed itself (called intersection). For data from humans and models with a random factor (Human, NN, CHSQ, Kmeans), we computed means (see Table 2).

Fig. 10 plots the number of intersections generated by human and each models on the 20 TSP problems used in our experiment. Because NN has no global information, it generates many more intersections than the rest of the models and human data. Pyramid and convex hull have deterministic algorithms, so they generated certain high peeks on particular problems and zero values on others. Both Kmeans and CHSQ are close to human data in value of number of intersections. None of the correlations with human performance were statistically significant, although the CHSQ correlation was marginally significant ($p < .1$).

Table 2
The sum (over all 20 problems) of means (over different participants or model runs) on number of intersections generated

	Human	NN	Convex Hull	Pyramid	Kmeans	CHSQ
Number of intersections	2.3	110	6	14	2.8	2.3
Correlation		0.30	−0.13	0.23	0.31	0.39 (*)
Average signed error		4.98	0.15	0.63	0.03 (*)	−0.03 (*)

Correlation and ASE on number of intersections between each model and human participants; "*" indicates best matched model on each dimension.

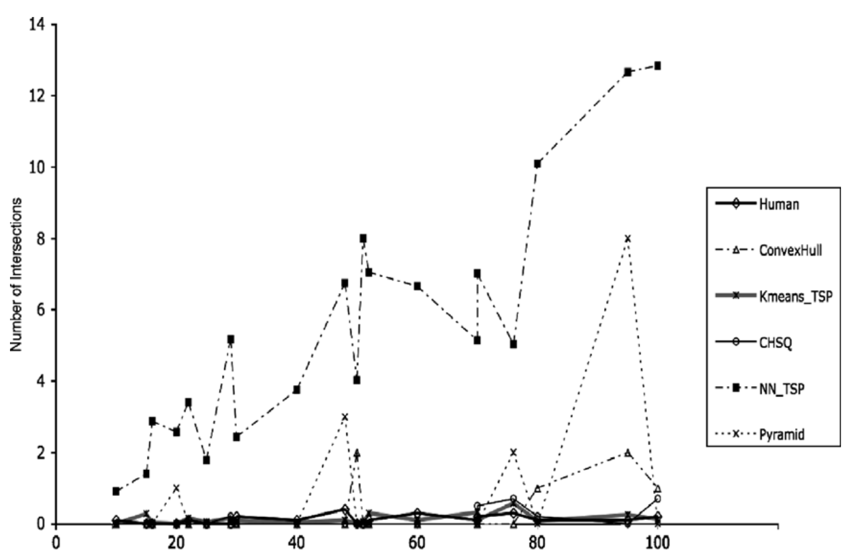

Fig. 10. Average number of intersections generated by each model for each problem.

Table 3
Correlation and ASE on accuracies between models and human participants

	NN	Convex hull	Pyramid	Kmeans	CHSQ
Correlation	0.70	0.62	0.41	0.84 (*)	0.67
Average signed error	0.17	−0.01	0.02	0.04	0.00 (*)

"*" indicates best matched model on each dimension.

6.2. Mean accuracy

As shown in Table 3, NN is much worse than the human performance in term of mean accuracy. Convex Hull, Pyramid, Kmeans and CHSQ are close to human accuracy levels as has been found in the part. Kmeans model did depart from the human data and the other three models as the number of points got larger. One reason for this is Kmeans model employed a naïve local search strategy that project points onto their nearest Spline-curve. As the number of points going bigger, the ratio of centroids to points is smaller. So the Spline-curve is more inaccurate in characterizing the detail local information. In this situation a more sophisticated local search strategy should be employed.

All but the Pyramid model led statistically significant correlation with (human data). The Kmeans model correlated with the trend of human performance best among the models we compared. Our hypothesis is that the global information Kmeans model utilizes is the best approximation to the global information human use, so it generate a similar trend with human performance. Fig. 11 plots the means of accuracy of human performance and each model.

6.3. Standard deviation

In addition to accounting for overall and problem-specific differences in mean accuracy, a model could also try account for overall and problem-specific differences in the variability across participants in accuracy (as measured by standard deviations). These differences in variability might reflect the degree of garden path effects from different start points (i.e., small choices made early have large down-the road consequences). A Levene Test shows that there are statistically significant differences in the standard deviations of human accuracies on different problems.

At the level of overall standard deviations, the pyramid and Convex Hull models fail outright because they are deterministic, and thus predict standard deviations of zero. The nearest neighbor model predicts standard deviations that are too large. The Kmeans and CHSQ models are close to observed human levels overall (Fig. 12). In terms of predicting problems specific differences in variability, none were statistically significant. It may be because that a few participants who were using different strategies than others (see Table 4).

6.4. Natural vs. random performance

We also tested the performance of the models on the two different set of points: natural and random. Recall that natural paths start with worse accuracy but have smaller decrease in accuracy as the number of points increase. None of the current models could predict the main effect of problem type nor the interaction with numbers of cities. However, a variation of the Kmeans model was successful: The minimum error in the 40 runs of Kmeans model for a given TSP problem generated the same pattern shown in the human performance. As we discussed earlier, there is no significant different between natural and random problems in term of accuracy performance in the other two participant groups. So only when both global and local information are available, the participants will generate this pattern of difference between natural and random problems. When we took the shortest path generated by our model across 40 runs, the accuracy performance for the two kinds of problems generated by our best model run (Fig. 13a) matched the solution accuracy generated by human participants (Fig. 13b). This is additional evi-

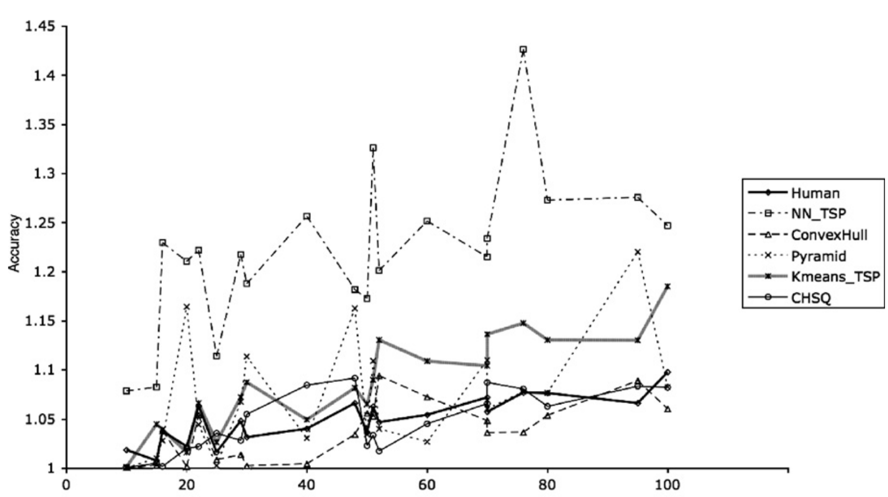

Fig. 11. Accuracy performance of models and humans on each problem.

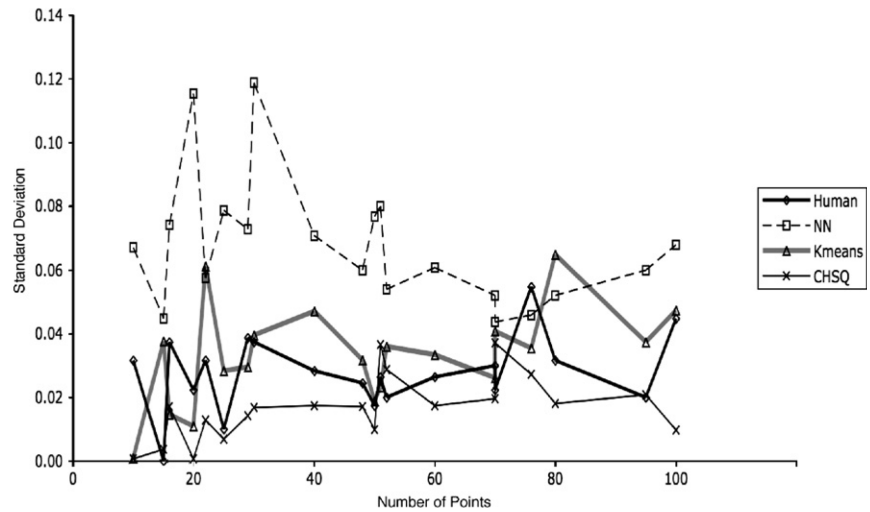

Fig. 12. Standard deviation on accuracy of human participants and the three models with any variability.

Table 4
Mean standard deviation, correlation and ASE of the standard deviations between human participants and models

	Human	NN	Kmeans	CHSQ
Mean	0.03	0.07	0.03	0.02
Correlation		0.06	0.11	0.22
Average signed error		0.04	0.01 (*)	−0.01 (*)

"*" indicates best matched models.

dence that there exists a spline-curve generated by our model that captured the global information human uses. The reason that the Kmeans model in average did not capture this phenomenon may be because it generated too many "bad" clusterings. In other words, only the "best" clustering generated by K-Means clustering algorithm and the spline-curve it follows would resemble the global information used by human cognition.

As the number of points goes up, the average number of points in each cluster becomes bigger ($\#Clusters = 2 \times \sqrt{\#points}$). The global information generated by our model becomes less accurate in guiding the local search, which results in the decrease in accuracy. However, in mid-size natural TSPs, there are some dense clusters. Those dense clusters contain much more points than average. Since the clusters are so dense, that accuracy is not much affected by the possible detours around the points inside those dense clusters. The number of points in the rest of the clusters remains small. So the global information is still quite accurate. In other words, because of an uneven distribution, the complexity of global information for natural points increased at a slower rate as the number points increases.

6.5. Exact path correlations

A good model of human TSP problem solving should not only predict the accuracy of the total path length that a human would generate on a TSP problem but also should be able to predict the likelihood of human participant taking a particular path. We used the following method to calculate the exact path correlation between human-generated and model-generated solutions. For each TSP problem with n cities, build a matrix of $n \times n$, where each cell $M(i,j)$ equals the numbers of observed paths between city

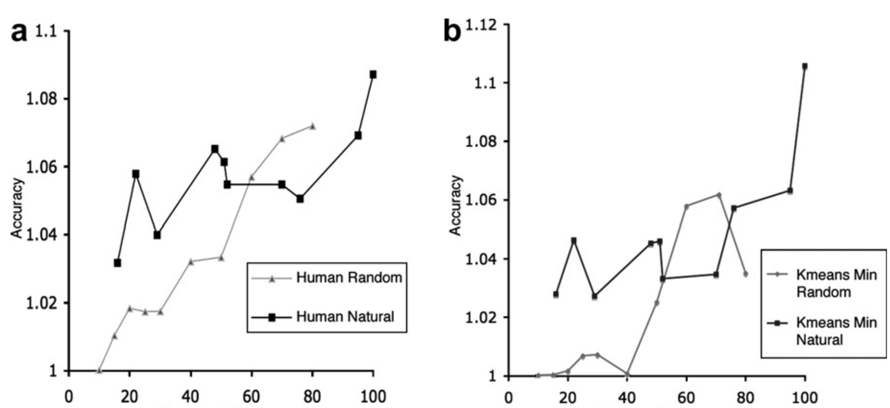

Fig. 13. Accuracy trends on natural/random problems for human participants (left) and Kmeans best run (right).

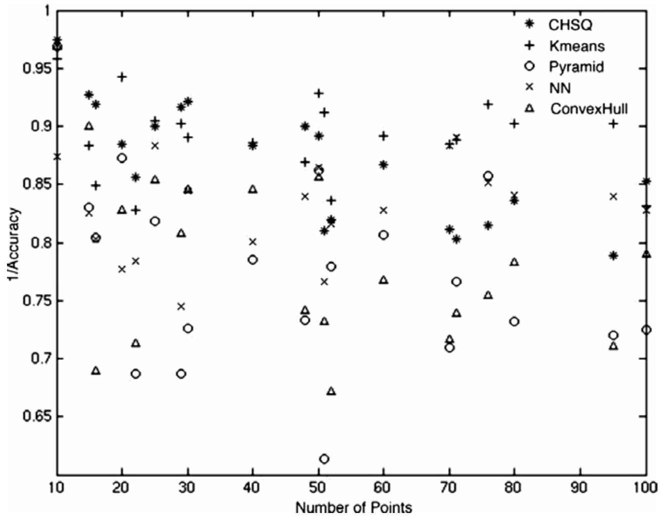

Fig. 14. Exact path correlation between human participants and each model.

i and city *j*. Then compare the similarity between the models and the participants at the individual path level by linearizing the matrix and compute the correlation between the two resulting vectors. Fig. 14 shows the mean correlation between the participant solutions and model solutions on each problem.

Table 5
Mean correlations between participant and model solutions on each of the 20 problems

	NN	CH	Pyramid	Kmeans	CHSQ
Mean correlation	0.83	0.78	0.77	0.89*	0.87

"*" indicates best matched models.

As Table 5 shows, the Kmeans model correlates best with the participants' choices on exact paths. The Kmeans model outperforms other models especially on larger size problems ($n > 50$), as Fig. 14 shows. One possible reason for this phenomenon is that as the size of the problem grows larger, human participants display a larger diversity of possible path choices. Our model captured this characteristic of human TSP solving by generating different paths on each run.

We visualized the characteristic by plotting the frequency of an edge selection by participants or each model as its thickness. Fig. 15 shows the solutions of a 50-points TSP generated by participants, Kmeans, CHSQ and Pyramid models. The arrows point to areas where participant-generated paths and Kmeans-generated paths displayed a great similarity in both pattern (path choices) and thickness (path frequency). The same kind of similarity could not be found in other models.

This visualization technique also helps us to identify reasons for why the Kmeans model departed from human data. As we can see in Fig. 16, the main outside contour of the Kmeans model displayed a large number of zigzags (as the arrows identifies) while the participants, the Pyramid and CHSQ models did not. The zigzags in the Kmeans model are the result of its naïve local search rule of projection. Since the Kmeans model connects points that have the nearest projections on the Spline-curve, those points themselves could be far away if one point is inside the Spline-curve and the next one is outside it. Through repeated crossing of the Spline-curve, the zigzag pattern is generated.

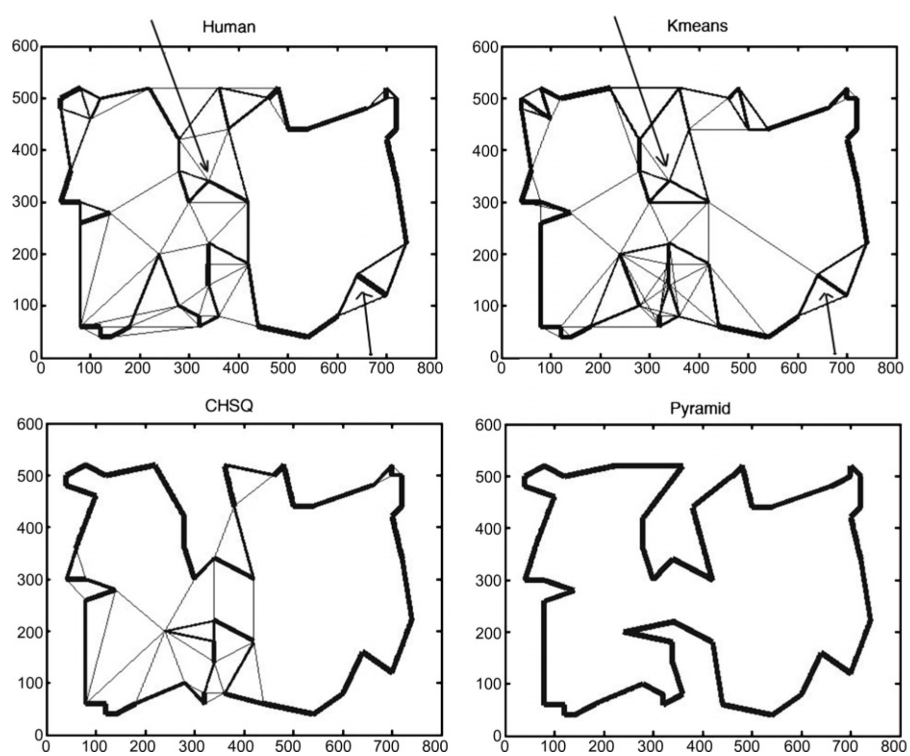

Fig. 15. Chosen paths for a 50-point TSP generated by participants and models.

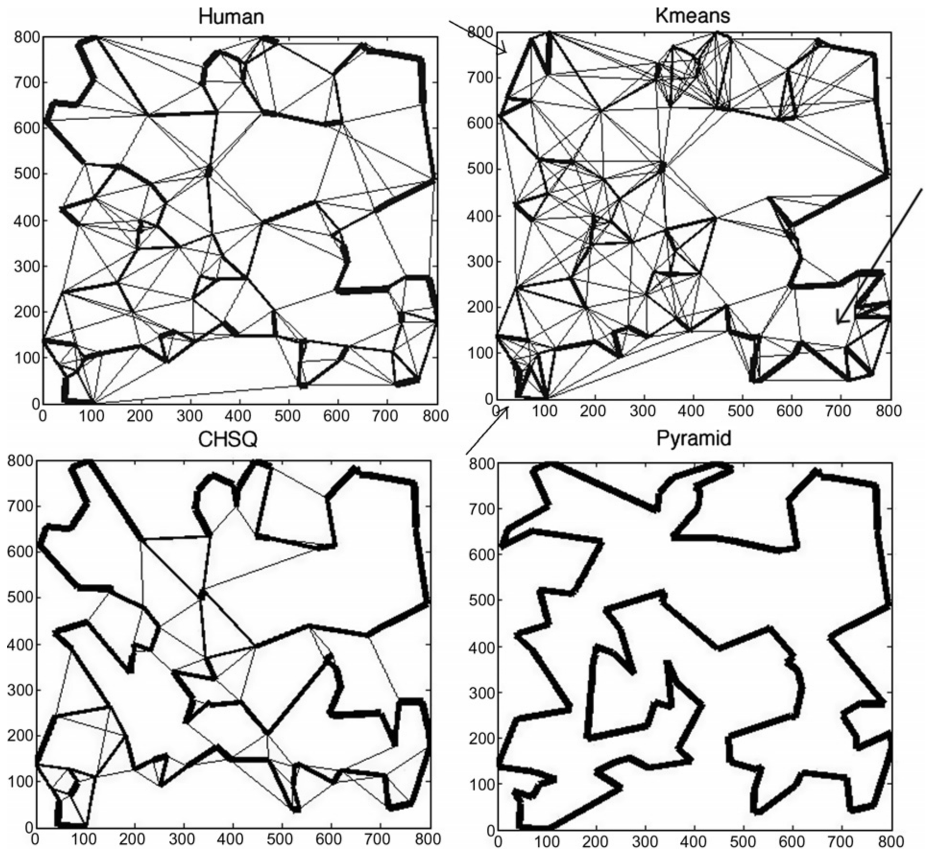

Fig. 16. Chosen paths for a 100-point TSP generated by participants and models.

6.6. Efficiency of global information representation

As the simulation results shown, the Kmeans model and CHSQ model are the best two in predicting human performance. But we have argued that the Kmeans model should predict human performance better than other models in part because it has a more compact representation of global information. Here we formulize this intuition, specifically by comparing the efficiency of the spline-curve of the Kmeans model with the traditional convex hull for representing global information in TSP.

Fig. 17 shows the convex hull and a spline-curve generated by our model on the same 70 points TSP. Both have 11 turning points, by which one could argue that they occupy approximately the same amount of visual working memory. But the spline-curve contains more information of the original problem so that it would be a better guide to the local search. Less information is required to be processed to build a path in the local information-processing phase using the spline-curve than the convex hull as the global information. This idea can be formalized as follows:

1. $I(S) \approx I(C)$.
2. $I(Path|S) < I(Path|C)$.
3. $\frac{I(S)}{I(Path|S)} > \frac{I(C)}{I(Path|C)}$.

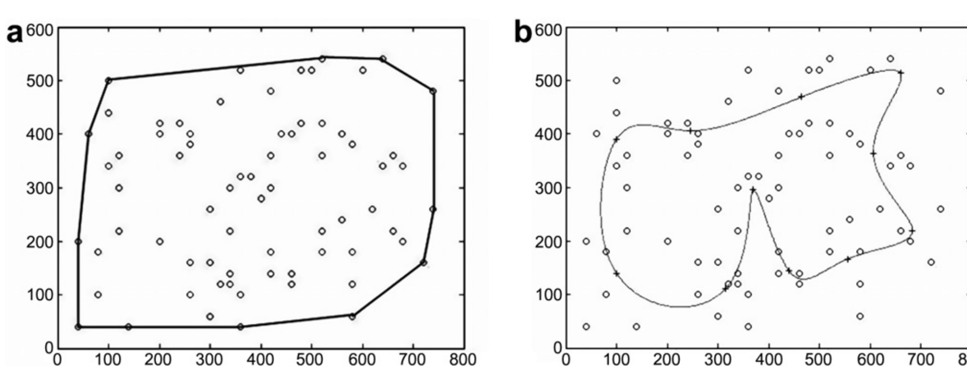

Fig. 17. Convex hull (left) and spline-curve (right) of 70 points TSP.

$I(S)$ is the information contained in the spline-curve, and $I(C)$ is the information contained in the convex hull, which are approximations of their size in human visual/spatial working memory. $I(Path|S)$ is the conditional information of the TSP path given the spline-curve as the global information, which is an approximation of the amount of information-processing in the local search phase. $I(Path|C)$ is the same information conditional on the given convex hull. As we can see, the spline-curve yields a higher ratio of information efficiency.

7. Conclusion

In this paper, we proposed an alternative model of human performance on the TSP, to test general predictions about the nature of global visual information that guides local visual information search. It has been previously argued that human cognition utilizes both global information and local information to solve the traveling salesman problem. Our model has separate phases of global information processing and local information processing, consistent with that of Best (2005). Result of our experiment and model simulations has shown that global information is important for human TSP solving and humans tend to use spline-like low-frequency curves around clusters to represent the global information.

We proposed two criteria for the global information: compactness and sufficiency. We defined the global information in our model to be a Spline-curve generated from the Kmeans cluster centroids. This global information is compact in size to be able to easily fit human visual/spatial working memory constrains. It is also sufficient to guide the local search without high search costs. Our local search strategy is very simple. By doing a simple projection onto the Spline-curve, our local information-processing phase guarantees the linearity on reaction time as a function of number of points. This local information processing part of the Kmeans model also processes points in a sequential way as suggested in previous research (Best, 2005; MacGregor et al., 2000).

To evaluate our model, we compared the models with human performance data that we collected. Our model fits the human performance well in number of intersections and the mean standard deviations of accuracy. The accuracy performance of the Kmeans model departs did a little bit from the human data when the number of points was large. This departure might be a result of an overly simplistic local search strategy, and/or our assumption that the global information is fully processed in the beginning no longer holds for large TSPs due to its complexity and working memory requirement. But our model generates a high correlation with the trend of human performance on different problems, suggesting that we have captured important elements of the global information that human use.

The presented ideas in modeling the global information processing in the TSP solving could also be adopted in other problem solving domains. When a certain factor in the problem solving process can not be directly observed or can only be weakly measured, it might be helpful to build a computational model of the problem solving process and compare its fitness in different dimensions as the uncertain factor is manipulated.

8. Future work

Currently our model is based on the hypothesis that global information is fully processed in the beginning of TSP solving. Furthermore, the global information used in control condition is not directly observable. It would be interesting to explore the possibility of interleaved global/local information processing when the size of TSP is large. An eye-tracking experiment would be helpful to further investigate these issues.

A more sophisticated local search strategy should also be developed to substitute the naïve projection rule. To further test our hypothesis about the nature of the global information processing, some other visual problem solving tasks, such as map reading and navigation, would be interesting to look at.

MacGregor and Ormerod (1996) and MacGregor, Ormerod, and Chronicle (1999) suggest that the difficulty of a TSP is affected by the number of points inside the convex hull and the layout of points. Our result in this paper further points to the direction that the complexity of the global information, instead of merely the number of points inside, caused the difficulty. So a qualitative analysis of the relationship between the complexity of global information and the difficulty of the problem would be interesting too.

Acknowledgement

Work on this project was supported by grants N00014-02-1-0113 and N00014-03-1-0061 to the second author from the Office of Naval Research.

References

Best, B. J. (2004). Modeling human performance on the traveling salesperson problem: Empirical studies and computational simulations (Doctoral dissertation, Carnegie Mellon University, 2004). *Dissertation Abstracts International*, 66, 6302.

Best, B. J. (2005). A model of fast human performance on a computationally hard problem. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 256–262).

Compton, B. J., & Logan, G. D. (1993). Evaluating a computational model of perceptual grouping by proximity. *Perception and Psychophysics, 53*, 403–421.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24*, 87–185.

Golden, B. L., Bodin, L. D., Doyle, T., & Stewart, W. Jr., (1980). Approximate traveling salesman algorithms. *Operations Research, 28*, 694–711.

Graham, S., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory and Cognition, 28*, 1191–1204.

MacGregor, J. N., & Ormerod, T. C. (1996). Human performance on the traveling salesman problem. *Perception and Psychophysics, 62*, 527–539.

MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (1999). Spatial and contextual factors in human performance on the traveling salesperson problem. *Perception, 28*, 1417–1427.

MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the traveling salesperson problem. *Memory and Cognition, 28*, 1183–1190.

MacQueen, J. B. (1967). Some methods for classification and analysis of multivariate observations. In *The Proceedings of the fifth Berkeley symposium on mathematical statistics and probability* (pp. 281–297). Berkeley, CA: University of California Press.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Ormerod, T. C., & Chronicle, E. P. (1999). Global perceptual processes in problem solving: The case of the traveling salesperson. *Perception and Psychophysics, 61*, 1227–1238.

Pylyshyn, Z. W. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition, 32*, 65–97.

Rosenkrantz, J. D., Stearns, R. E., & Lewis, M. P. II, (1977). An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing, 6*, 563–581.

Schunn, D. C., & Wallach, D. (2001). *Evaluating goodness-of-fit in comparison of models to data*. http://www.lrdc.pitt.edu/schunn/gof/.

Van Oeffelen, M. P., & Vos, P. G. (1982). Configurational effects on the enumerations of dots: Counting by groups. *Memory and Cognition, 10*, 296–404.

Van Rooij, I., Stege, U., & Schactman, A. (2003). Convex hull and tour crossings in the Euclidean traveling salesperson problem: Implications for human performance studies. *Memory and Cognition, 31*, 215–220.

Victers, D., Lee, M., Dry, M., & Hughes, P. (2003). The roles of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems. *Memory and Cognition, 31*, 1094–1104.