

Practice makes proficient: pigeons (*Columba livia*) learn efficient routes on full-circuit navigational traveling salesperson problems

Danielle M. Baron · Alejandro J. Ramirez ·
Vadim Bulitko · Christopher R. Madan ·
Ariel Greiner · Peter L. Hurd · Marcia L. Spetch

Received: 7 March 2014/Revised: 6 June 2014/Accepted: 11 June 2014/Published online: 26 June 2014
© Springer-Verlag Berlin Heidelberg 2014

Abstract Visiting multiple locations and returning to the start via the shortest route, referred to as the traveling salesman (or salesperson) problem (TSP), is a valuable skill for both humans and non-humans. In the current study, pigeons were trained with increasing set sizes of up to six goals, with each set size presented in three distinct configurations, until consistency in route selection emerged. After training at each set size, the pigeons were tested with two novel configurations. All pigeons acquired routes that were significantly more efficient (i.e., shorter in length) than expected by chance selection of the goals. On average, the pigeons also selected routes that were more efficient than expected based on a local nearest-neighbor strategy and were as efficient as the average route generated by a crossing-avoidance strategy. Analysis of the routes taken indicated that they conformed to both a nearest-neighbor and a crossing-avoidance strategy significantly more often than expected by chance. Both the time taken to visit all goals and the actual distance traveled decreased from the first to the last trials of training in each set size. On the first trial with novel configurations, average efficiency was higher than chance, but was not higher than expected from a nearest-neighbor or crossing-avoidance strategy. These results indicate that pigeons can learn to select efficient routes on a TSP problem.

Keywords Traveling salesman problem · Pigeon · Route learning · Problem solving · Nearest-neighbor strategy · Crossing avoidance · Planning · Foraging

Introduction

The traveling salesperson problem (TSP) is straightforward—given a set of cities or goals, a traveling salesperson must visit all of the cities once and then return to the start point via the shortest route possible (e.g., Applegate et al. 2007). Despite the simplicity of this description, the solution to the problem is computationally hard: Although the optimal path for a given configuration of goals (hereafter called “array”) can be obtained by computing the set of all possible routes and selecting the one with shortest total distance, this calculation becomes more taxing as the number of goals (set size) increases because the number of possible routes increases factorially with set size. For example, when required to travel to four distinct goals and return to the start, there are 24 potential routes one could take, but if set size is increased to 10, the number of potential routes is over three million. The development of a tractable algorithm that would determine the optimal route for any arbitrary TSP problem has eluded computer scientists, although approximation algorithms have been developed (e.g., Golden et al. 1980).

Finding the most efficient route between multiple goals can be important for numerous everyday activities, whether it be a business scheduling deliveries of their product or an individual navigating between stores or attractions in a new city. Interest in the TSP and the cognitive processes underlying efficient route selection by humans has intensified recently because of the surprising finding that untrained humans perform remarkably close to optimal on

D. M. Baron · C. R. Madan · A. Greiner ·
P. L. Hurd · M. L. Spetch (✉)
Department of Psychology, BSP-217, University of Alberta,
Edmonton, AB T6G 2E9, Canada
e-mail: mspetch@ualberta.ca

A. J. Ramirez · V. Bulitko
Department of Computing Science, University of Alberta,
Edmonton, AB, Canada

pictorial two-dimensional (2D) versions of the TSP (e.g., “paper-and-pencil” tasks), even with a fairly large set size (MacGregor and Ormerod 1996). In fact, the suboptimality of human solutions seems to increase only linearly with increases in the set size, even though the number of possible routes increases factorially with number of nodes (see MacGregor and Chu 2011 for a review). Optimal or near-optimal performance on TSP tasks has been suggested to indicate higher-order problem-solving and future-planning abilities, and in the field of neuropsychology, TSP tasks have been used to assess cognitive function and problem-solving abilities in humans (e.g., Foti et al. 2011; Vickers et al. 2001).

Several potential strategies used by humans to determine efficient routes between multiple goals have been proposed and evaluated, largely through the use of 2D pictorial tasks (MacGregor and Chu 2011), although some studies have used navigational tasks, in which people visit multiple goal locations in a laboratory room (e.g., Blaser and Wilber 2013; Wiener et al. 2009). One simple and widely tested strategy is the “nearest-neighbor” (NN) heuristic, in which the participant repeatedly chooses the nearest unvisited goal. Humans generally exceed the average efficiency levels generated by the NN strategy (Graham et al. 2000; Ormerod and Chronicle 1999; Wiener et al. 2009). In contrast to NN, which is considered to be a local strategy, global strategies are based on the spatial layout of the array for route selection (MacGregor and Ormerod 1996; MacGregor 2012). One such strategy is “crossing avoidance” (CA), in which participants select routes that visit all goals without crossing any paths. Consistent with a CA strategy, studies have found that human-generated TSP paths rarely self-intersect (MacGregor et al. 2000; van Rooij et al. 2003; Vickers et al. 2003).

TSP problems are important not only for humans but also for non-human animals. Many animals encounter TSP-like problems in the wild when navigating away from home to forage for food. Finding an efficient route to and from various foraging sites is beneficial for reducing energy expenditure and exposure to predation and to conserve time available for other crucial activities such as mating or tending one’s young (Stephens et al. 2007). To illustrate the importance of efficient route selection, consider that starlings (*Sturnus vulgaris*) spend most of their time foraging for food to bring back to the nest for their young, sometimes making as many as 400 round trips to and from their nest daily (Kacelnik 1984). The large number of potential configurations of goal locations an individual might encounter, coupled with the fact that the locations of food sources, shelters, and homes may change over time or with the seasons, suggests that learning and memorizing a single route would typically not suffice. Thus, a strategy

that is flexible enough to support efficient search among a variety of goal configurations would be valuable.

The first reported animal experiment with the TSP was conducted by Menzel (1973) with chimpanzees (*Pan troglodytes*). One experimenter carried a young chimpanzee around a large outdoor enclosure on his back, while a second experimenter hid 18 food items. The chimpanzee was allowed to watch where the food was being hidden and then, after being returned to the start location, was able to go out and search for the hidden food. The routes taken by the chimpanzee were more efficient than the route originally taken by the experimenters, suggesting that while being carried around, the chimpanzee formed a “cognitive map” (Gallistel 1990; Tolman 1948) of the enclosure and was able to determine an efficient route between the goal locations. Efficient travel on a multi-goal task was also reported for yellow-nosed monkeys (MacDonald and Wilkie 1990) and vervet monkeys (Cramer and Gallistel 1997). Interestingly, Janson (2014) specifically analyzed the data from these studies (Cramer and Gallistel 1997; MacDonald and Wilkie 1990; Menzel 1973) and concluded that performance did not exceed the efficiency that could be obtained by using an NN strategy. Janson (2014) further proposed that the actual movement sequences were consistent with a “gravity” rule, in which animals respond on the basis of summed attraction to unused resources; in this model, proximity, value and clustering of goals can all affect choices, but no complex planning is required.

TSP-like problems have now been studied in a wide range of additional organisms, including ants (Dorigo and Gambardella 1997), bumblebees (Lihoreau et al. 2010; Ohashi et al. 2007), pigeons (Gibson et al. 2007, 2012; Miyata and Fujita 2008, 2010), rats (Blaser and Ginchansky 2012; Bures et al. 1992; de Jong et al. 2011; Reid and Reid 2005), lemurs (Lühns et al. 2009) and other species of monkeys (Janson 2007; MacDonald et al. 1994). The procedures used, questions addressed and results obtained have varied considerably across studies. In most cases, performance by non-humans has exceeded chance level, but it has often not exceeded the efficiency that could be obtained with an NN strategy.

Gibson et al. (2007) tested humans and pigeons on a one-way TSP task with clusters of three, four, or five nodes on a computer screen. The subjects were required to select each of the nodes, but unlike the full-circuit TSP, they were not required to return to the point of origin, and they could start at any node. Both humans and pigeons performed better than random, but significantly worse than routes created by an NN strategy. When the pigeons were explicitly reinforced for selecting efficient routes, their efficiency became closer to the level that could be achieved by NN. Blaser and Ginchansky (2012) tested rats and

humans in a one-way version of a navigational TSP, which required them to visit ten goal locations. Subjects were presented with each of four possible configurations once, and in all configurations, the goals were organized into a roughly circular formation. Again, results indicated that both rats and humans were more efficient than random, but were generally not as efficient as an NN strategy.

Interestingly, recent studies have suggested that bees can learn to perform well on the TSP. Lihoreau et al. (2012a) tested bumblebees in a flight-optimization task between several flower locations, in which following an NN strategy would result in taking a sub-optimal route. After extensive training, bees converged on the optimal route between flowers, and not the NN strategy. However, the bees were not tested on novel flower configurations, so their ability to select efficient routes for unfamiliar arrays was not assessed.

Recently, Gibson et al. (2012) tested pigeons on both one-way and full-circuit TSPs in a navigational task with two or three goals. An interesting feature of their study was that after training, they tested pigeons with configurations that were specifically designed to pit the use of a local NN strategy against selection of a more efficient route. With some configurations, the pigeons made NN choices even though these led to longer than optimal routes. However, the birds chose more optimally in some 3-goal configurations in which the NN strategy would produce a large increase in path length.

Our study extended previous investigations by testing pigeons' performance on full-circuit TSPs in a navigational task with more than three goals. Like the study by Gibson et al. (2012), we used a navigational task, rather than a computer-based task, to more closely simulate some of the properties of natural foraging. Unlike most studies of TSP, we measured not only the sequence of choices made (which we refer to as the *route*) but also the actual travel paths of the birds (which we refer to as the *path*). Our study was designed to answer four specific questions. First, if pigeons were repeatedly exposed to fixed sets of goal locations, would they converge on routes that were not only more efficient than expected by random choice but also more efficient than predicted by an NN strategy? Second, would the stable routes they developed indicate use of an NN or a CA strategy? Third, would the pigeons perform above chance, and at or above the level of an NN strategy when first presented with novel arrays? Fourth, would the pigeons improve with training in their efficiency as measured by the actual distance traveled (i.e., their path lengths) and the time taken to complete the circuit (their path duration)?

Several specific features of our study design were geared toward answering these questions. To simulate the pressures that are presumably inherent in real-world foraging

situations and to encourage efficient route selection, the pigeons were given a time limit to complete the circuit. Over the course of the experiment, set size was systematically increased from three to six goals. At each set size, the birds were trained repeatedly on three distinct arrays of goal locations. Training at each set size continued until each bird satisfied a route-consistency criterion (see Methods). The birds were then tested with two novel arrays of the same set size to examine the routes taken by the pigeons when they were first exposed to unfamiliar arrays.

We hypothesized that pigeons would learn to select efficient routes between goal locations in a navigational task. Although information about pigeons' foraging patterns in the wild primarily comes from studies of flock visits to different sites, a study using global positioning system (GPS) tracking by Rose et al. (2006) suggests that pigeons should be able to learn to forage efficiently among multiple locations. They studied the travel paths to urban locations by feral pigeons in Switzerland and found that individual pigeons visited several locations, with some birds visiting up to nine different places in a day. Thus, the TSP should be a meaningful problem for pigeons.

Methods

Subjects

Two female and four male adult pigeons (*Columba livia*; 3 Silver King and 3 Homing) participated in the experiment. All birds had previous experience in spatial orientation tasks, but none had previously served in foraging tasks with multiple goals. The birds were maintained at 85 percent of their free-feeding body weight on a diet of commercial pigeon pellets (Mazuri Ex Gamebird Breeder Diet, PMI International). They were housed individually in metal cages (42 cm high × 47 cm wide × 42 cm deep), and the colony room was maintained on a 12:12 h light:dark cycle with light onset at 7:00am. Grit and water were provided ad libitum.

Materials

Enclosure

All experimental sessions took place in a laboratory open-field arena (2 m width × 2.25 m length). White plastic walls (92 cm high) lined the space. Above the walls, white translucent curtains hung around the perimeter of the enclosure. A small start box (34 cm width × 23 cm depth × 34 cm height) was separated from the experimental arena via guillotine door (17 cm width × 20 cm height) that could be raised to allow the bird in and out of

the arena. The experimental room was lit by four 40-watt fluorescent light bulbs located on the ceiling, which were hidden from view by a ceiling curtain. An overhead video camera was centered above the arena with the lens protruding through the curtain. The arena floor was lined with approximately 4 cm of aspen chip bedding. A 5×5 grid marked on the floor under the bedding was used by the experimenter to select among the 25 possible locations for the goals (food containers, 7 cm diameter \times 4 cm height). Grid points were spaced 32 cm apart on both x- and y-axes; the diagonal distance between the nearest grid points was approximately 45.25 cm.

Array computation

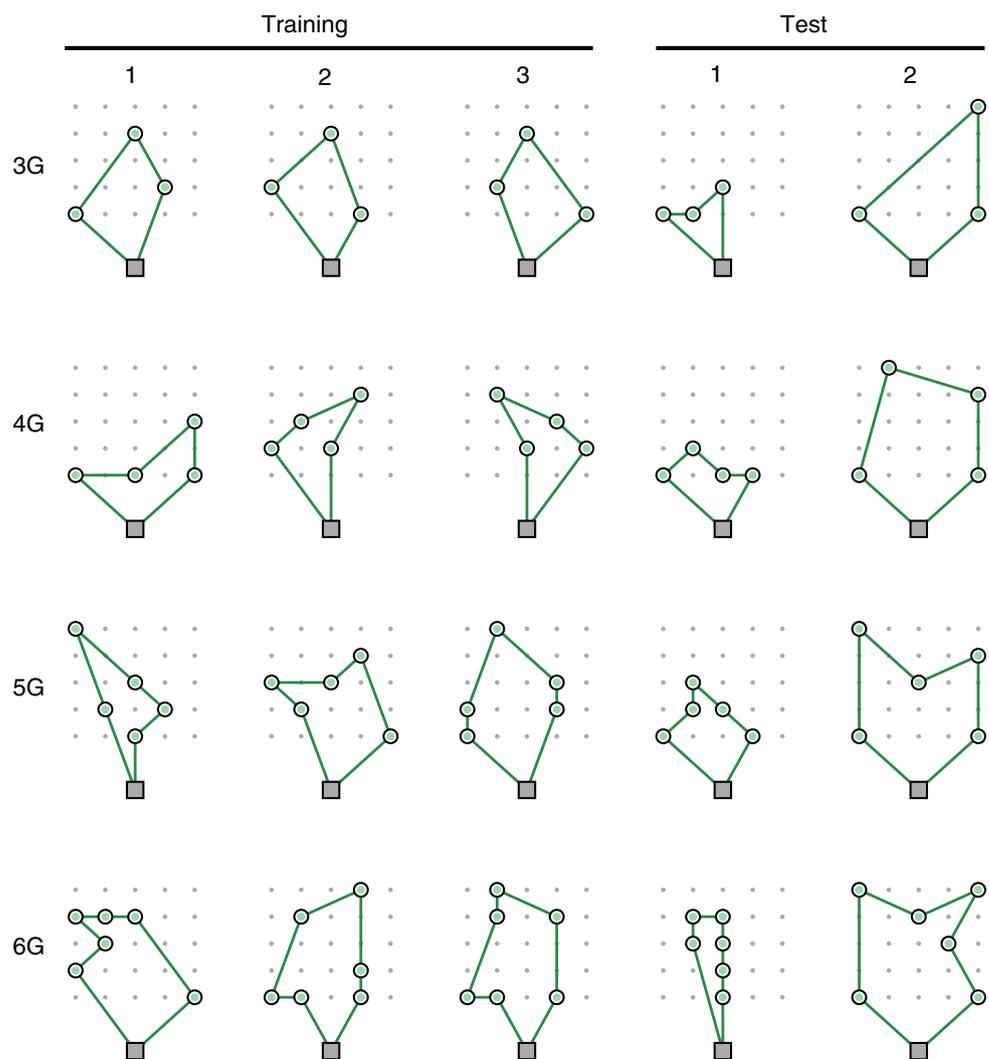
To compute the arrays for the goal settings to be used in training and testing, a series of scripts for the mathematical/analytical software MATLAB (The MathWorks Inc., Natick, MA) were created. These scripts sampled from all

possible configurations in the 25 potential locations for each set size (i.e., three, four, five, and six goals; hereafter referred to as 3G, 4G, 5G and 6G). For each configuration, the script calculated the length of the optimal route, and the length of a randomly selected NN route (note that when two goals are equally far apart, more than one NN route is possible). For each set size (number of goals), three arrays were randomly selected for training, and two arrays were selected for testing that differed from each other in terms of the ratio of the lengths of the NN route to the optimal route. Specifically, the ratio was smaller for the “easy” test arrays (“Test 1”) than the “hard” test arrays (“Test 2”). A schematic of the arrays used is shown in Fig. 1.

Scoring and measures

Each point on the 5×5 grid in the enclosure was assigned a number from 1 to 25. The choices were scored from the video recording and the birds’ routes documented in

Fig. 1 Schematic of the three training arrays and two test arrays for each set size. Test array 1 was selected to be easier according to an NN strategy than test array 2. *G* refers to number of goals (i.e., set size). The start box is indicated by the gray box, and the goal locations are indicated by circles. The line connecting each dot shows an optimal full-circuit route for each set size



sequential order by the numbers corresponding to the goal locations. Only complete (i.e., full-circuit) routes were considered valid and included in the reported data.

Studies of TSP in animals typically report only the straight line distance between the sequence of visited locations, which is variously referred to as the path length or route length. We recorded both the actual path traveled and the sequence of choices. We use the term *route* to indicate the sequence of unique goal locations visited by the pigeons. We use the term *path* to indicate the actual nonlinear path traveled by the birds as they foraged for food in the arena. Figure 2 shows the distinction between route and path.

To measure the pigeons' routes, we determined the coordinates (x, y) for each visited goal location and calculated the straight line distance between sequentially visited locations. The *route length* was the sum of these straight line distances starting at the origin (i.e., the start box) and ending with the return to the start box. Route length provided an inverse measure of efficiency.

To measure the pigeons' paths, we used a spectral time-lapse algorithm recently developed by Madan and Spetch (2014) to determine the actual paths traveled by the pigeons and to extract the length and duration of these paths from the video recordings of the trials. Path length and path duration provided inverse measures of actual foraging efficiency. Path length is influenced not only by the route taken but also by the directness of travel between goals. For example, path length is increased if the pigeons take a meandering path between goals (cf. Wystrach et al. 2011). In addition, on rare occasions, the birds flew around while in the arena, particularly after visiting the last food cup, which would inflate the path length measure, but not the route length. To assess the directness of the path between goals and back to the start box, we calculated ratios of the path length to the route length. Higher ratios indicated less direct paths.

To assess the pigeons' routes in terms of strategies, we first determined the set of all possible full-circuit routes in each specific array that would be generated by the NN strategy, the CA strategy or random choice (R) among all possible routes. NN routes included all routes in which each subsequent choice was to the closest unvisited goal; if the distance to two or more goals was equal, multiple routes conformed to the NN strategy. CA routes were all possible full-circuit that had no path intersections; analogous to NN, multiple CA routes are possible for an array. R routes are the full set of possible solutions to a TSP instance. To determine whether the routes taken by the pigeons conformed to the NN or CA strategies, we classified each pigeon route in terms of whether it matched an NN route or a CA route (note that these are not mutually exclusive). We calculated the proportion of the birds' routes that matched these NN or CA routes and the proportion expected by chance (i.e., the number of NN or CA routes divided by the full set of possible routes for the array).

Procedures

Each bird participated in one experimental session per day, which consisted of one to four individual trials, depending on the experimental phase. At the start of each trial, the bird was placed in the start box for 1 min of habituation and then the guillotine door was opened until the bird exited into the arena; the door was then closed until the end of the trial. Time constraints and criteria to complete a trial varied across experimental stages, as described below.

Habituation and procedural training

Training began with habituation to the experimental arena, in which the pigeon was given the opportunity to explore and find food. Four food containers filled with grit and each

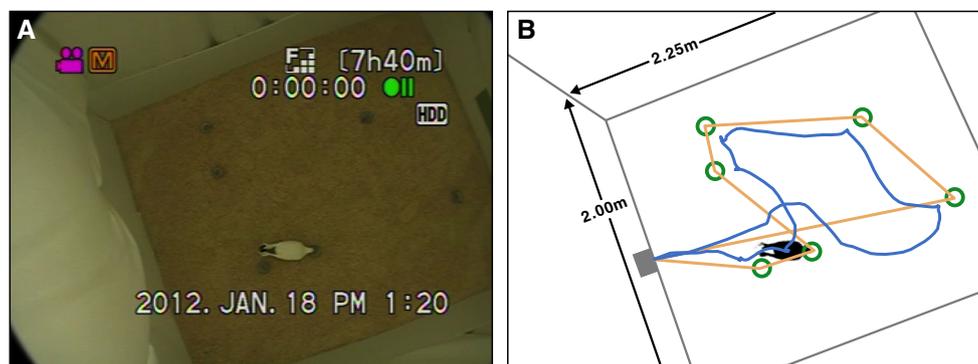


Fig. 2 **a** Frame from a video showing a pigeon performing the task with six goals. **b** Calculated travel path of the pigeon (blue curvy line) and route connecting the selected goals (orange straight lines). The

start box is indicated by the gray box, and the goal locations are indicated by green circles (color figure online)

containing fifteen food pellets formed a square array in the center of the room. The sliding door was opened, and the bird received a maximum of 30 min to explore and eat the food before the lights were extinguished and the trial ended. If the bird consumed all pellets from all four containers in less than 30 min, it was given additional trials to a maximum of four, or until the 30 min session time elapsed, whichever occurred first.

Once the bird left the start box within 20 s of the door opening for three consecutive trials in 1 day, training to return to the start box began. After the bird entered the arena and the door to the start box was closed, a container of food was placed inside the start box near the entrance. Once the bird consumed the pellets from all four containers, the door to the start box was opened. The trial ended when the bird consumed the pellets from the start box container. Across trials, the container was gradually shifted further inside of the start box, forcing the bird to enter the box to consume the food. In the final stage, the bird had to fully enter the start box and the door was closed behind it while it consumed the food.

Efficiency training

The next training phase was designed to encourage pigeons to forage efficiently. First, the bird was required to visit all four containers (still arranged in the square array) and return to the start box within 1 min for three consecutive trials in 1 day. The mean time to complete a trial (i.e., from when the start box door was opened to when it was closed behind the bird at the end of the trial) across these three trials was calculated and used as the time limit for each individual bird. On subsequent trials, the bird needed to complete all three trials within their time limit or the lights were extinguished and the pigeon sat in the dark for 2 min before being retrieved by the experimenter. The bird continued this phase of training until it successfully completed three consecutive trials without exceeding its time limit for three consecutive days.

Task training

The next phase used a set size of three goals and was designed to train the pigeons to search among goal locations arranged in different arrays. In this phase, and in all subsequent training and testing phases, each goal contained four pellets. The grit in each container was gradually reduced across trials until it only lined the bottom, making the food sitting on top of the grit only visible within close proximity to the container. On each trial, the pigeon was required to visit all three goals and then return to the holding box within its time limit or else the lights were extinguished for 2 min and the trial ended. Each pigeon

was trained with three distinct arrays of three goal locations (see Fig. 1) until it met a stability criterion of taking the same route for a given array for three consecutive days; this consistency was required for all arrays (i.e., consistency in 9/9 cases). Traveling the exact same route (sequence of choices) but in the opposite direction was classified as the same route. Once this criterion was met, the bird was tested with two new arrays, one trial of each per day, until it completed five trials with each array.

Array training and testing

The general design of the rest of the experiment involved a series of training and testing phases that started with a set size of four goal locations and increased across phases to six goal locations. All training phases used three distinct arrays (see Fig. 1). Each session consisted of one trial with each training array, with the order randomized across sessions. Failure to visit each goal and then return to the holding box within the bird's individual time limit resulted in the lights being extinguished for 2 min and the trial ended. The same individual time limits were used in all phases, except for one bird (Bird 48), which sustained a minor foot injury part way through 6G training. The bird received 2 weeks of recovery before continuing with the experiment, and the time limit was increased by 15 s for the remainder of the experiment.

The consistency criteria for advancing from training to testing varied with the number of goal locations, to reflect the increased number of potential routes at each set size. In 4G training, one route deviation in one of the arrays was permitted across the 3 days (i.e., 8/9). In 5G training, birds were allowed one deviation on up to two arrays across the 3 days (i.e., 7/9), and in 6G training, they were allowed to deviate once for each array (i.e., 6/9). These consistency criteria remain very strict given the dramatic increase in the number of potential routes as set size increases (from 12 routes at 4G to 360 routes at 6G, given that travel among the same sequence of goal locations, but in the opposite direction was counted as the same route).

A test phase followed each of the training phases. Each test phase consisted of two trials per day, one with the array selected to be easy and one with the more difficult array. Birds were presented with the same two test arrays daily in random order. Testing lasted until birds successfully completed five trials in each of the two arrays without exceeding their time limit or for a maximum of 10 days.

Analyses

We conducted several analyses on subsets of the data to address our specific questions. When multiple t tests were conducted on a set of data, we corrected for family-wise

error rate using the Holm–Bonferroni method. All ratio data were log-transformed for statistical analyses.

Stable routes To assess the efficiency of the stable routes we compared the pigeons' route lengths on the last two training trials with each array to the average lengths of the NN, CA and R routes. To assess whether the stable routes suggested use of an NN or CA strategy, we compared the proportion of routes from the last two training trials that matched each strategy to the proportion expected by chance.

Improvement over training To assess whether efficiency improved over training, we compared the first two and last two training trials with each array at the 4G, 5G and 6G set sizes (one bird completed training at one set size within 3 sessions; in this one case, we used only the first and last training trial with each array). These comparisons were made on all measures.

Routes on novel arrays To assess whether pigeons developed efficient routes on their first exposure to novel arrays, we compared route length on the first trial with each array in both training and testing to the average lengths of NN, CA and R. To assess efficiency on the easy and hard test routes, we also separately compared route length for the first four test trials (the minimum number completed by all birds) with each array to the average lengths of NN, CA and R for those arrays.

Results

The mean and standard error of the mean (\pm SEM) for the number of sessions required to establish stable paths for the six birds was 11.3 (\pm 1.8) for the initial 3G training. For the subsequent 4G, 5G and 6G training, the number of sessions was 19.3 (\pm 4.1), 20.7 (\pm 8.4) and 21.3 (\pm 8.5), respectively. The lights-off procedure (for failure to complete the trial in the designated time limit) was imposed on an average of 6 % of the total trials across the experiment (range of 1–16 % across the six birds). These incomplete trials are not included in any of the analyses. All of the following results are based on data from set sizes four to six.

Stable routes

Our first main finding is that by the end of training, the pigeons converged on routes that were more efficient than both random choice and by use of an NN strategy. Figure 3 shows each bird's mean route length from the last two training trials for each array in each set size condition. The average lengths of R, NN and CA routes are indicated by

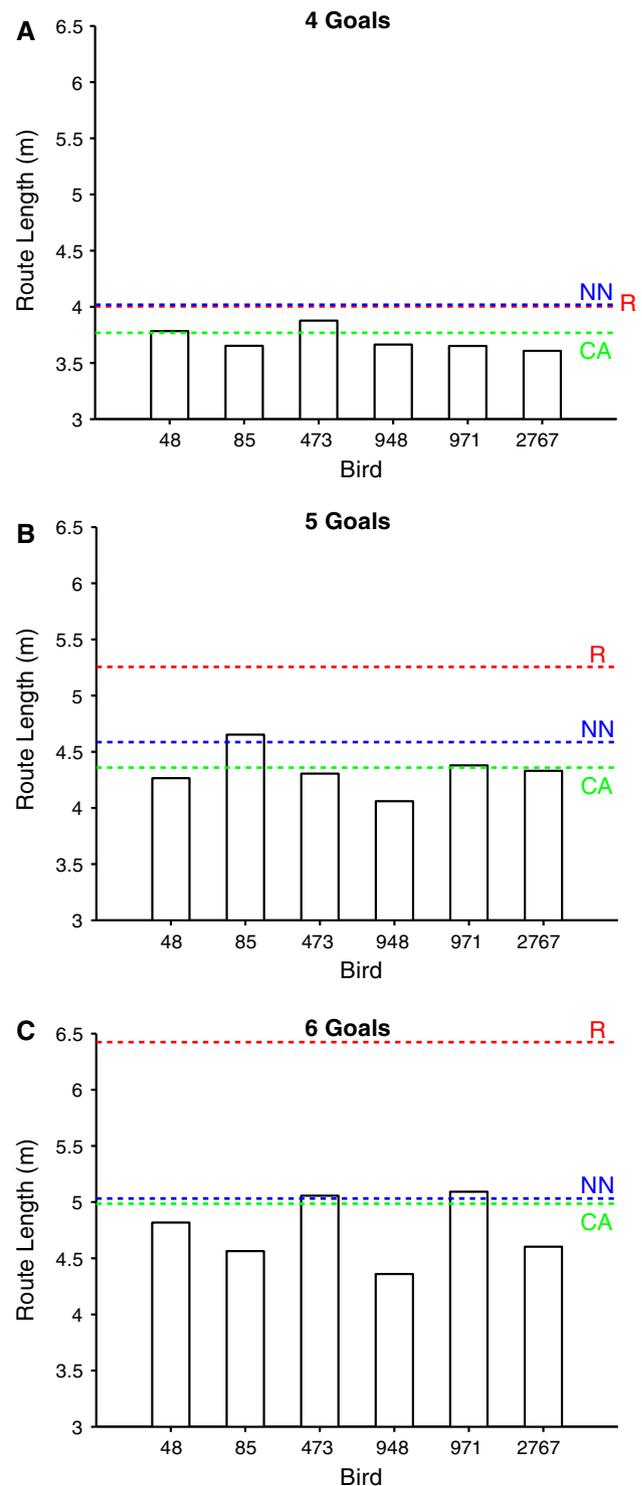


Fig. 3 Mean route length for each bird at each set size. The horizontal lines indicate the efficiency level expected based on random selection of the locations (R), or by following the nearest-neighbor (NN) and crossing-avoidance (CA) strategies. See text for details

the horizontal lines. Averaged across set sizes, the birds' route lengths (4.26 ± 0.06 m) were significantly shorter than the average of R (5.26 m; $t(5) = 17.37$, $P < .001$, *Cohen's* $d = 7.09$) and NN (4.52 m; $t(5) = 4.52$, $P = .006$, $d = 1.85$). The birds' route lengths were shorter than the average route generated by the CA strategy (4.37 m), but the difference did not reach statistical significance ($t(5) = 2.05$, $P = .096$, $d = .837$). Thus, the stable routes pigeons developed over training were more efficient than expected based on a local, NN strategy and at least as efficient as those expected by use of a CA strategy.

Figure 4 shows that the stable routes developed by the pigeons also fit those generated by NN and by CA strategies significantly more often than expected by chance. Averaged across set sizes for the last two training trials with each array, the proportion of choices that fit the NN

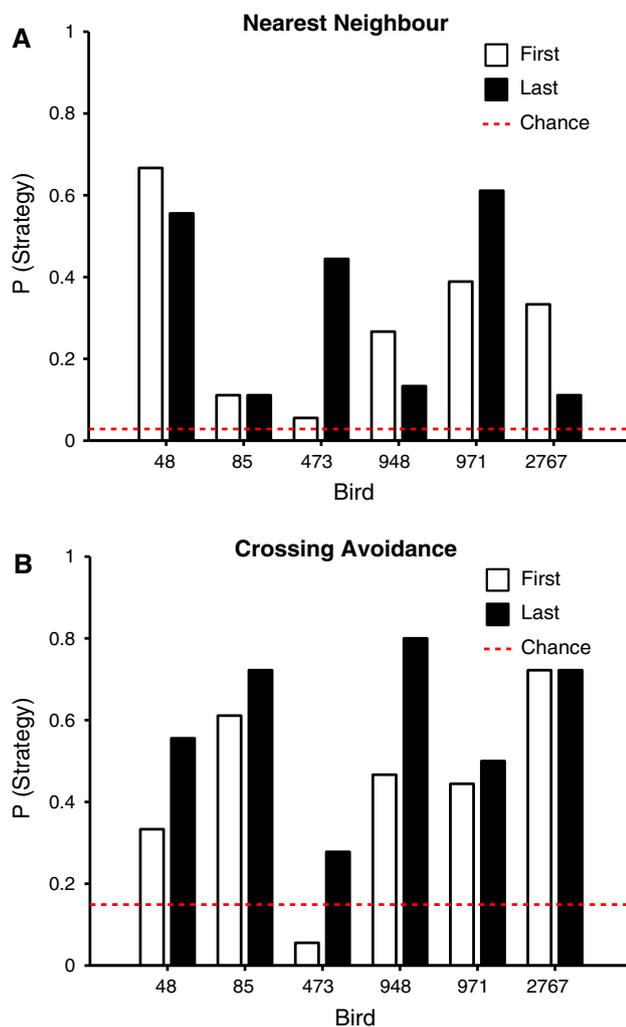


Fig. 4 Proportion of routes for each bird that matched routes generated by the nearest-neighbor (a) and crossing-avoidance (b) strategies for the first two and last two trials with each training array, averaged over set sizes four to six

strategy (0.33 ± 0.10) were significantly higher than expected by chance (0.029 ; $t(5) = 3.11$, $P = .026$, $d = 1.27$) and the proportion of choices that fit the CA strategy (0.60 ± 0.08) were also significantly higher than expected by chance (0.149 ; $t(5) = 5.69$, $P = .002$, $d = 2.321$). Thus, pigeons clearly did not show a uniformly random selection of locations, but instead appeared to show a preference both for selecting the nearest unvisited locations and for avoiding routes that required intersecting paths.

Improvement over training

Pigeons became more efficient with training, both in the actual distance traveled (path length) and the duration of travel to complete a full circuit. Arguably, these are the two most important measures from a biological perspective. As shown in Fig. 5, path length and path duration decreased from the first two training trials with each training array to the last two training trials. The data were analyzed with repeated-measures multivariate analysis of variance with training (first two trials vs. last two trials) and set size (4G, 5G, and 6G) as factors. For path length, there was a significant effect of training ($F(1,5) = 9.28$, $P = .029$, $\eta_p^2 = .650$), the effect of set size just missed significance ($F(2,4) = 6.70$, $P = .053$, $\eta_p^2 = .770$), and there was no significant interaction between training and set size ($F(2,4) = 2.46$, $P = .201$, $\eta_p^2 = .552$). Path length is determined both by the route and by the directness of the path between consecutive locations on the route. Subsequent analyses suggested that neither of these accounted fully for the improvement in path length. An analysis of route length showed that the effect of training failed to reach significance ($F(1,5) = 5.35$, $P = .069$, $\eta_p^2 = .517$), but there was a significant effect of set size ($F(2,4) = 83.43$, $P = .001$, $\eta_p^2 = .977$). There was no significant interaction between training and set size ($F(2,4) = 0.37$, $P = .715$, $\eta_p^2 = .154$). An analysis of the ratio of path length to route length showed that the effect of training was near the threshold of significance ($F(1,5) = 6.38$, $P = .053$, $\eta_p^2 = .561$), and there was no significant effect of set size ($F(2,4) = 0.84$, $P = .495$, $\eta_p^2 = .296$), and no significant interaction between training and set size ($F(2,4) = 3.51$, $P = .132$, $\eta_p^2 = .637$). Thus, the significant improvement in path length with training may reflect a combination of selecting more efficient routes and taking more direct paths between locations, neither of which quite reached significance on their own.

For path duration, there was a significant effect of training ($F(1,5) = 15.220$, $P = .011$, $\eta_p^2 = .753$), but no significant effect of set size ($F(2,4) = 3.93$, $P = .11$, $\eta_p^2 = .663$), and no significant interaction ($F(2,4) = 0.10$, $P = .990$, $\eta_p^2 = .005$).

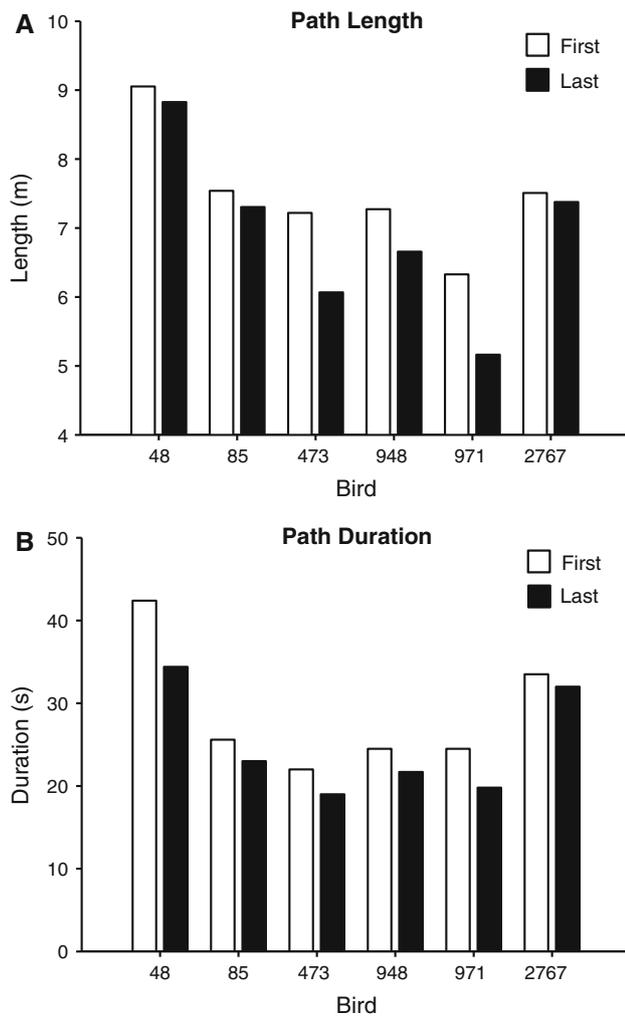


Fig. 5 Mean path length (a) and path duration (b) of each pigeon for the first two trials and last two trials of training averaged across arrays and set sizes four to six

Performance on test arrays

On the first test trial with new goal arrangements, the birds were more efficient than R, but did not exceed the

efficiency of either NN or CA (see Table 1). Averaged across the arrays and set sizes, the birds had significantly shorter route lengths (4.56 ± 0.178 m) than expected by random choice (5.36 m; $t(5) = 4.45$, $P = .006$, $d = 1.652$). The length of the birds' routes did not differ significantly from those generated by the NN strategy (4.49 m; $t(5) = 0.43$, $P = .689$, $d = 0.173$), or the CA strategy (4.31 m; $t(5) = 1.41$, $P = .217$, $d = .576$).

The route lengths for the easy and hard arrays could not be directly compared because the distance of the optimal path differed. Therefore, to compare them we computed ratios of the birds' route lengths over the optimal route lengths. With this measure, higher ratios indicate less efficient routes relative to optimal. Averaged across the first 4 trials with each test array (the minimum number completed by any bird at each set size), the ratios were higher for the difficult arrays than for the easier arrays, but the difference failed to reach statistical significance ($t(5) = 2.33$, $P = .067$, $d = .951$). Table 1 shows the ratios for individual birds.

Discussion

This experiment revealed several interesting results that advance knowledge of TSP behavior in pigeons. First, with repeated experience in a full-circuit TSP task, pigeons learned to select efficient routes for familiar arrays even up to a set size of six goals: By the end of training, the route lengths (based on straight line distances between selected locations) were significantly shorter (i.e., more efficient) than expected by random selection of the goals, and they were also significantly shorter than expected based on the NN strategy. On average, the routes were at least as short as expected based on a crossing-avoidance strategy. This high level of efficiency is interesting in light of previous evidence that pigeons performed less efficiently than expected based on an NN strategy (Gibson et al. 2007) when tested on a computer task with arrays that continually changed across trials. Our results clearly indicate that

Table 1 Results of test trials with novel arrays

	Strategy		Bird						
	<i>First test trial, both arrays</i>								
Route length (m)	R	NN	CA	48	85	473	948	971	2767
	5.36	4.49	4.31	4.05	4.30	4.66	5.32	4.44	4.62
	<i>First four test trials, easy array</i>								
Ratio (bird route/optimal route)	48	85	473	948	971	2767			
	1.14	1.12	1.11	1.22	1.02	1.13			
	<i>First four test trials, hard array</i>								
Ratio (bird route/optimal route)	48	85	473	948	971	2767			
	1.13	1.12	1.23	1.25	1.11	1.20			

R Random, NN nearest neighbor, CA crossing avoidance

pigeons can *learn* efficient routes in a navigational TSP task. Similar conclusions regarding efficient route learning have been proposed with other species. Lihoreau et al. (2012a) presented bumblebees with an array of six artificial flowers designed to provide a suboptimal route if an NN strategy was employed. Results showed that bees significantly reduced their travel distance and rarely chose an NN route once they gained experience with the array.

Second, the stable routes developed by the pigeons conformed to both NN and CA significantly more often than expected by chance. Thus, the birds showed a bias to select the nearest unvisited goal location, but they improved their efficiency over that expected by the NN strategy alone by also showing a preference for routes that avoided path intersections.

When tested with novel arrays, the pigeons performed at efficiency levels well above that expected by random choice, but they did not exceed the level expected by an NN strategy. This suggests that processes that allow the pigeons to exceed the efficiency levels generated by the NN strategy for training arrays may depend on experience with a particular array. It is worth noting, however, that the pigeons' efficiency on novel arrays was also not below that expected by NN. Our study did not provide a sufficient number of novel test arrays at each set size to analyze the strategies used by the pigeons with the novel arrays. Tests with a larger set of novel arrays and tests, in which an NN strategy produces highly inefficient routes, (such as in Gibson et al. 2012) would be useful in future studies to investigate the processes used by pigeons when confronted with novel arrays of four or more goals.

Several specific aspects of our training procedure may have contributed to the efficiency of the pigeons' performance in our study. First, we used a full-circuit navigational foraging task designed to mimic real-world foraging conditions. Navigational problems differ from pictorial 2D TSP problems in several ways, including the perspective on the array and the energetic costs of taking longer than necessary routes; these differences may have encouraged more efficient responding than found in previous studies with 2D TSP tasks (e.g., Gibson et al. 2007). Second, we attempted to explicitly encourage efficient foraging during initial training by terminating trials in which the pigeon foraged too slowly. Third, we continued training at each set size until consistent routes developed. Although consistency of routes rather than route length was the criterion for completion of each training phase, the time limit may have encouraged the birds to converge on stable routes that were efficient. Any of these procedural factors may have contributed to the pigeons' efficient performance.

The pigeons in our study clearly established efficient routes, but further research is needed to determine the

specific learning processes that allowed the pigeons to select efficient routes. For example, familiarity with the configurations may have allowed the pigeons to use global mapping processes to plan out an efficient route for those arrays, such as routes that avoid path intersections. That is, the pigeons may have learned to plan their routes among the locations. Alternatively, the learning to select efficient routes and to avoid crossing paths could reflect processes other than global planning. For example, the learning could reflect reinforcement processes in which routes that lead to efficient collection of food are differentially strengthened. Research using learning algorithms for travel paths over training (e.g., Lihoreau et al. 2012b) may be useful for addressing these questions.

Our third interesting result is that the pigeons' actual path lengths and the time taken to complete the circuit decreased with experience at each set size during training. Although actual path lengths are typically not reported in the TSP literature, from a biological point of view, these measures of TSP performance are important because the benefit of adopting an efficient path between goals in nature is the reduction in time and energy required to obtain food. Thus, confirmation that performance on these measures improved during training provides confidence that the pigeons were sufficiently motivated and sensitive to the requirements of the task and that the improvement in efficiency is meaningful.

Finally, it is worth noting that some pigeons in our study clearly foraged more efficiently than others. For example, Birds 48 and 948 consistently chose routes that exceeded the efficiency of CA whereas bird 473 was consistently less efficient than expected by the CA strategy. This kind of individual variation in foraging is perhaps not surprising given that GPS tracking data on foraging by feral pigeons (Rose et al. 2006) also revealed large individual differences, with some pigeons flying longer distances to reach a stable food source, and others foraging at several food sources nearer to the loft.

In summary, over time pigeons acquired routes that were more efficient than those expected based on a local NN strategy and as efficient as those generated by a more global CA strategy. With novel routes their performance was comparable to that generated by a NN strategy. Thus, with sufficient experience and motivation to respond efficiently, pigeons can perform as well or better than expected based on an NN strategy in a full-circuit navigational TSP task with up to six goals.

Acknowledgments We thank Nicole Savignac for assistance with subject running and data scoring and Isaac Lank for assistance with constructing the apparatus. This study was funded by a Natural Sciences and Engineering Research Council of Canada Discovery grant awarded to MLS. Some of these data were presented at the 19th Annual International Conference on Comparative Cognition in

March, 2012. All research was conducted in accordance with Canadian Council on Animal Care guidelines and with approval from the University of Alberta Animal Care and Use Committee. The experiments comply with the current laws of Canada.

Conflict of interest The authors declare that they have no conflict of interest.

References

- Applegate DL, Bixby RE, Chvatal V, Cook WJ (2007) The traveling salesman problem: a computational study. Princeton series in applied mathematics, Princeton University Press, Princeton
- Blaser RE, Ginchansky RR (2012) Route selection by rats and humans in a navigational traveling salesman problem. *Anim Cogn* 15:239–250. doi:10.1007/s10071-011-0449-7
- Blaser RE, Wilber J (2013) A comparison of human performance in figural and navigational versions of the traveling salesman problem. *Psychol Res* 77:761–772. doi:10.1007/s00426-012-0470-8
- Bures J, Buresova O, Nerad L (1992) Can rats solve a simple version of the traveling salesman problem? *Behav Brain Res* 52:133–142. doi:10.1016/S0166-4328(05)80223-2
- Cramer AE, Gallistel CR (1997) Vervet monkeys as travelling salesmen. *Nature* 38:464. doi:10.1038/387464a0
- de Jong LW, Gereke B, Martin GM, Fellous JM (2011) The traveling salesrat: insights into the dynamics of efficient spatial navigation in the rodent. *J Neural Eng* 8: 65010–65020. doi:10.1088/1741-2560/8/6/065010
- Dorigo M, Gambardella LM (1997) Ant colonies for the travelling salesman problem. *BioSystems* 43:73–81. doi:10.1016/S0303-2647(97)01708-5
- Foti F, Petrosini L, Cutuli D, Menghini D, Chiarotti F, Vicari S, Mandolesi L (2011) Explorative function in Williams syndrome analyzed through a large-scale task with multiple rewards. *Res Dev Disabil* 32:972–985. doi:10.1016/j.ridd.2011.02.001
- Gallistel CR (1990) The organization of learning. MIT Press, Cambridge
- Gibson B, Wasserman EA, Kamil AC (2007) Pigeons and people select efficient routes when solving a one-way travelling salesperson task. *J Exp Psychol Anim B* 33:244–261. doi:10.1037/0097-7403.33.3.244
- Gibson B, Wilkinson M, Kelly D (2012) Let the pigeon drive the bus: pigeons can plan future routes in a room. *Anim Cogn* 15:379–391. doi:10.1007/s10071-011-0463-9
- Golden B, Bodin L, Doyle T, Stewart W (1980) Approximate travelling salesman algorithms. *Oper Res* 28:694–711
- Graham SM, Joshi A, Pizlo Z (2000) The traveling salesman problem: a hierarchical model. *Mem Cogn* 28:1191–1204. doi:10.3758/BF03211820
- Janson CH (2007) Experimental evidence for route integration and strategic planning in wild capuchin monkeys. *Anim Cogn* 10:341–356. doi:10.1007/s10071-007-0079-2
- Janson C (2014) Death of the (traveling) salesman: primates do not show clear evidence of multi-step route planning. *Am J Primatol* 76(5):410–420. doi:10.1002/ajp.22186
- Kacelnik A (1984) Central place foraging in starlings (*Sturnus vulgaris*). I. Patch residence time. *J Anim Ecol* 53:283–299. doi:10.2307/4357
- Lihoreau M, Chittka L, Raine N (2010) Travel optimization by foraging bumblebees through readjustments of traplines after discovery of new feeding locations. *Am Nat* 176:744–757. doi:10.1086/657042
- Lihoreau M, Chittka L, Le Comber SC, Raine NE (2012a) Bees do not use nearest-neighbour rules for optimization of multi-location routes. *Biol Lett* 8:13–16. doi:10.1098/rsbl.2011.0661
- Lihoreau M, Raine NE, Reynolds AM, Stelzer RJ, Lim KS, Smith AD, Osborne JL, Chittka L (2012b) Radar tracking and motion sensitive cameras on flowers reveal the development of pollinator multi-destination routes over large spatial scales. *PLoS Biol*. doi:10.1371/journal.pbio.1001392
- Lühns M, Dammhahn M, Kappeler PM, Fichtel C (2009) Spatial memory in the grey mouse lemur (*Microcebus murinus*). *Anim Cogn* 12:599–609. doi:10.1007/s10071-009-0219-y
- MacDonald SE, Wilkie DM (1990) Yellow-nosed monkeys (*Cerco-pithecus ascanius whitesidei*) spatial memory in a simulated foraging environment. *J Comp Psychol* 104(4):382–387. doi:10.1037/0735-7036.104.4.382
- MacDonald SE, Pang JC, Gibeault S (1994) Marmoset (*Callithrix jacchus*) spatial memory in a foraging task: win-stay versus win-shift strategies. *J Comp Psychol* 108:328–334. doi:10.1037/0735-7036.108.4.328
- MacGregor JN (2012) Indentations and starting points in traveling sales tour problems: implications for theory. *J Probl Solving* 5(1):3. doi:10.7771/1932-6246.1140
- MacGregor JN, Chu Y (2011) Human performance on the traveling salesman and related problems: a review. *J Probl Solving* 3(2):2. doi:10.7771/1932-6246.1090
- MacGregor JN, Ormerod T (1996) Human performance on the traveling salesman problem. *Percept Psychophys* 58:527–539. doi:10.3758/BF03213088
- MacGregor JN, Ormerod TC, Chronicle EP (2000) A model of human performance on the traveling salesperson problem. *Mem Cogn* 28(7):1183–1190. doi:10.3758/BF03211819
- Madan CR, Spetch ML (2014) Visualizing and quantifying movement from pre-recorded videos: the spectral time-lapse (STL) algorithm. *F1000Research* 3:19. doi:10.12688/f1000research.3-19.v1
- Menzel EW (1973) Chimpanzee spatial memory organization. *Science* 182:943–945. doi:10.1126/science.182.4115.943
- Miyata H, Fujita K (2008) Pigeons (*Columba livia*) plan future moves on computerized maze tasks. *Anim Cogn* 11:505–516. doi:10.1007/s10071-008-0141-8
- Miyata H, Fujita K (2010) Route selection by pigeons (*Columba livia*) in “traveling salesperson” navigation tasks presented on an LCD screen. *J Comp Psychol* 124:433–446. doi:10.1037/a0019931
- Ohashi K, Thomson JD, D’Souza D (2007) Trapline foraging by bumblebees. IV. Optimization of route geometry in the absence of competition. *Behav Ecol* 18:1–11. doi:10.1093/beheco/ar1053
- Ormerod TC, Chronicle EP (1999) Global perceptual processing in problem solving: the case of the traveling salesperson. *Percept Psychophys* 61(6):1227–1238. doi:10.3758/BF03207625
- Reid RA, Reid AK (2005) Route finding by rats in an open arena. *Behav Proc* 68:51–67. doi:10.1016/j.beproc.2004.11.004
- Rose E, Haag-Wackernagel D, Nagel P (2006) Practical use of GPS-localization of feral pigeons *Columba livia* in the urban environment. *Ibis* 148(2):231–239. doi:10.1111/j.1474-919X.2006.00499.x
- Stephens DW, Brown JS, Ydenberg RC (eds) (2007) Foraging: behavior and ecology. University of Chicago Press, Chicago
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208. doi:10.1037/h0061626
- van Rooij I, Stege U, Schactman A (2003) Convex hull and tour crossings in the Euclidean traveling salesperson problem: implications for human performance studies. *Mem Cogn* 31:215–220. doi:10.3758/BF03194380
- Vickers D, Butavicius M, Lee M, Medvedev A (2001) Human performance on visually presented traveling salesman problems. *Psychol Res* 65:34–45. doi:10.1007/s004260000031

- Vickers D, Lee MD, 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. *Mem Cogn* 31:1094–1104. doi:[10.3758/BF03196130](https://doi.org/10.3758/BF03196130)
- Wiener JM, Ehbauer NN, Mallot HA (2009) Planning paths to multiple targets: memory involvement and planning heuristics in spatial problem solving. *Psychol Res* 73(5):644–658. doi:[10.1007/s00426-008-0181-3](https://doi.org/10.1007/s00426-008-0181-3)
- Wystrach A, Schwarz S, Schultheiss P, Beugnon G, Cheng K (2011) Views, landmarks, and routes: how do desert ants negotiate an obstacle course? *J Comp Physiol A* 197(2):167–179. doi:[10.1007/s00359-010-0597-2](https://doi.org/10.1007/s00359-010-0597-2)