Optimizing and “pessimizing” – human performance with instructional variants of the traveling salesperson problem

Ed Chronicle

Department of Psychology

Acknowledgements

- Ulrike Stege
- Jim MacGregor & Tom Ormerod
- Andy Dewald, Yun Chu & Michelle Lopez (graduate students)
- Takashi Nakamura, Brooke Pearlman & Joli Malone (undergraduate students)

Waikiki beach

Location of my house

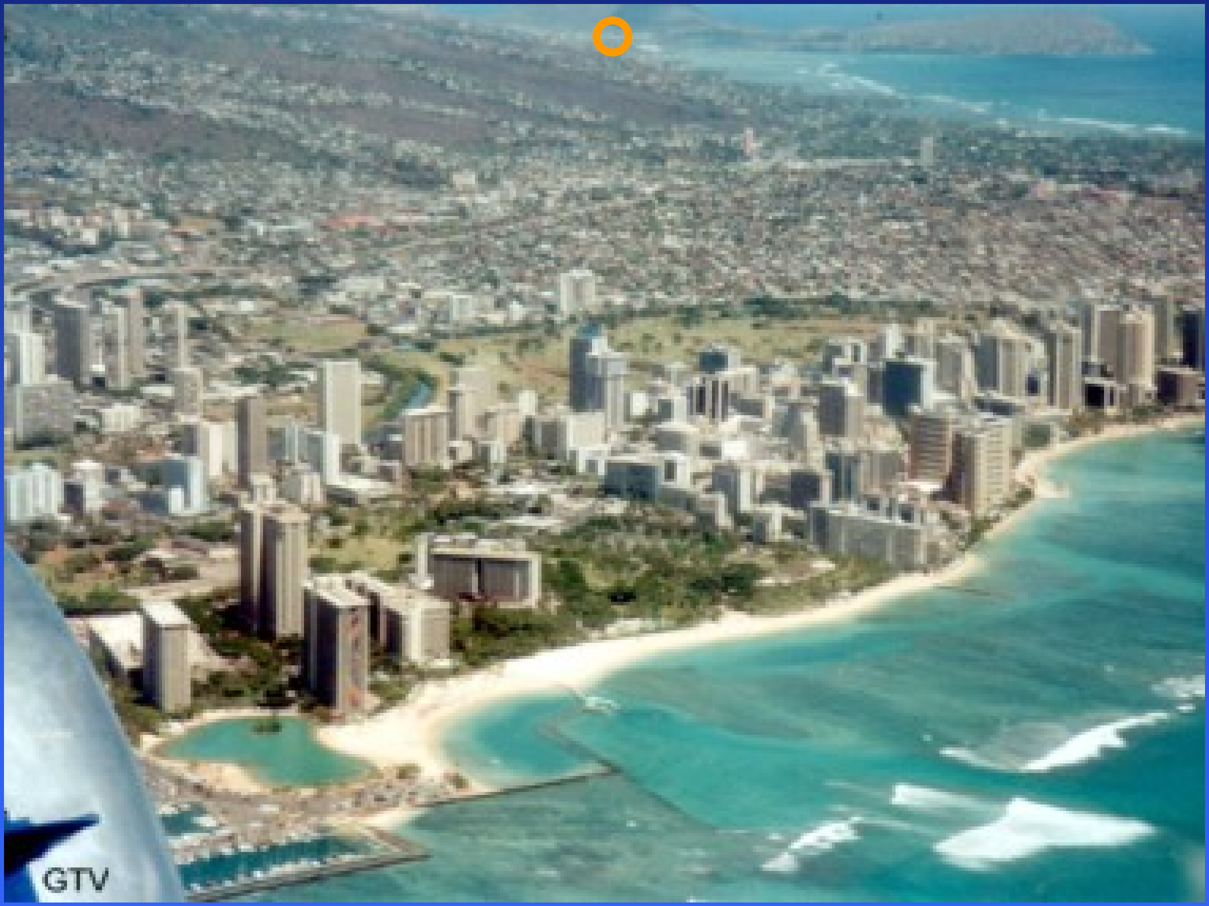

The perils of the tropics…

Human optimization

- Human subjects are pretty good at finding close-to-optimal solutions to TSPs
- Roughly speaking, an overview of the published literature suggests that humans get worse linearly with the number of points in the problem.

The basis of human ability

- Why should people be good at the TSP?
- Two broad types of explanation seem to suggest themselves:
  - generic spatial cognitive ability (e.g. ability to judge proximity relations or pairwise distances)
  - some cognitive process that is (for some reason) specialized for optimization

Longest and shortest tours

- If it is a generic spatial ability that is letting people find shortest tours, then they should also be good at finding longest tours

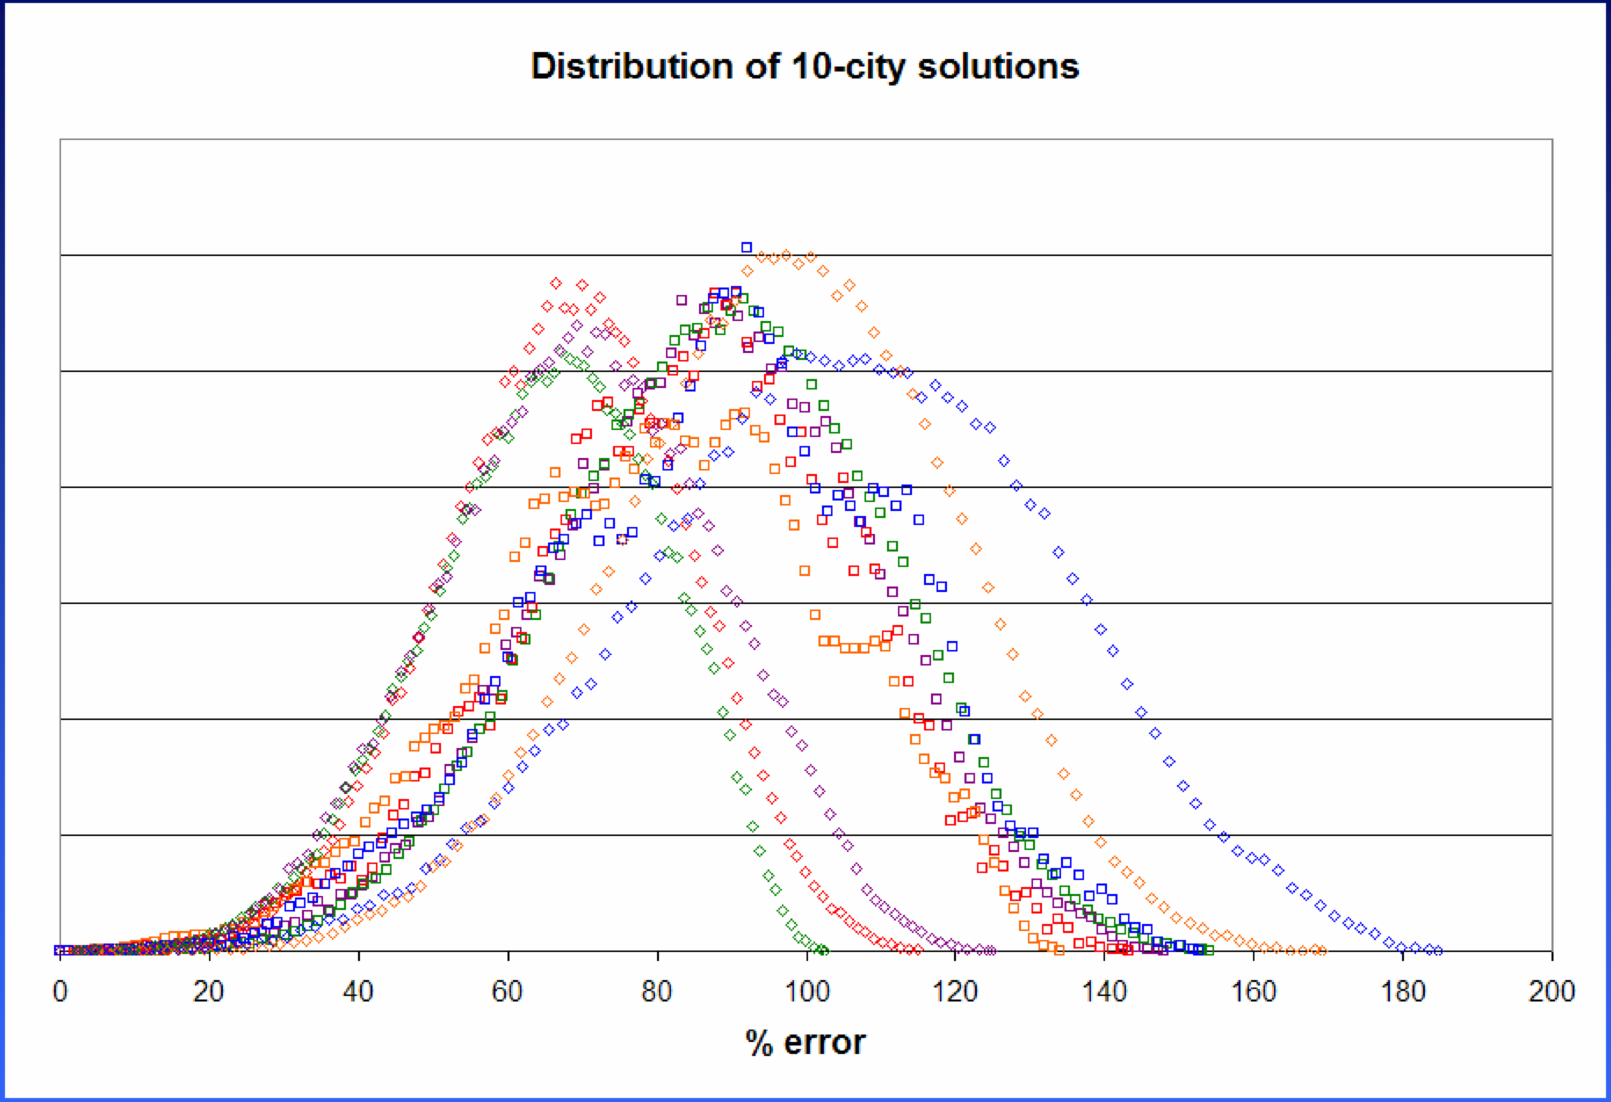

With thanks to Jack Saalweachter

Experiment 1

- Contrasted performance on standard TSPs with performance on “longest-path” TSPs
- In these latter, the instruction is to visit all the locations once and return to the starting point in the longest possible route
- The longest-path version is as combinatorially complex as the standard version

Example problem

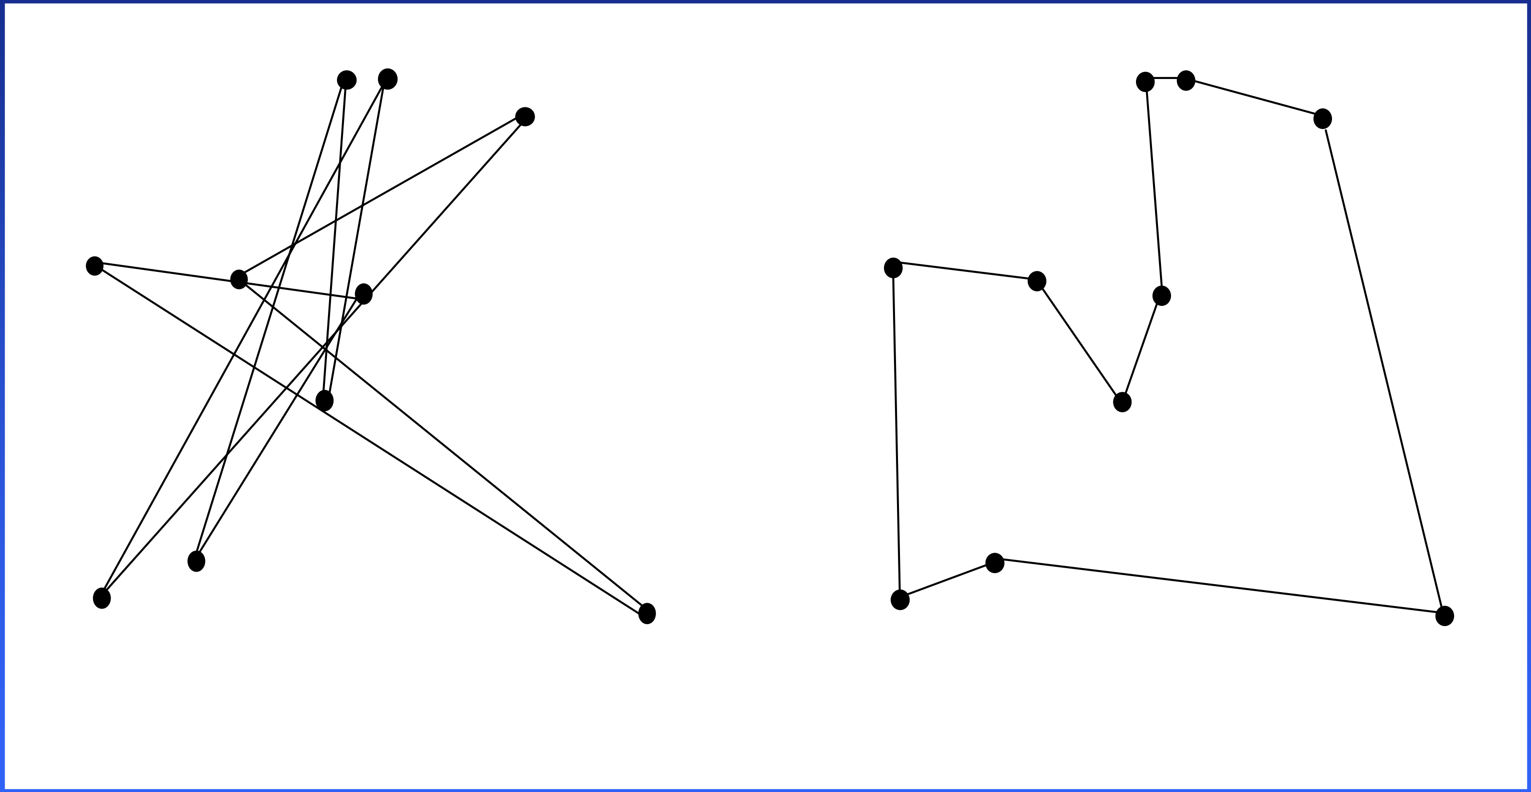

Longest tour

Shortest tour

Method

- Twenty PSY100 students, naïve to the TSP
- Each participant produced both shortest and longest tours to the same five 10-point problems
- Pencil-and-paper procedure, classroom setting
- Order of problems randomized within task blocks; task block order counterbalanced across participants
- Timing data self-recorded by participants

Experiment 1 results

- 31 of 100 tours in the standard condition were minimal, none was maximal in the longest-path condition.
- On average, tour lengths were 6.02% above the minimal for short tours, and 16.23% below the maximal for long tours, with short tour performance significantly superior to long, $t(19)=4.30$, $p<.001$.
- On average short tours took significantly less time than long tours, at 5.67s versus 12.76s, $t(19)=9.67$, $p<.001$.
- After removing estimated drawing time, the respective means were 0.93s and 6.63s, $t(19)=8.06$, $p<.001$.

Comparison with simple heuristics

- In the longest-path condition, humans perform significantly less well than a Furthest Neighbor heuristic

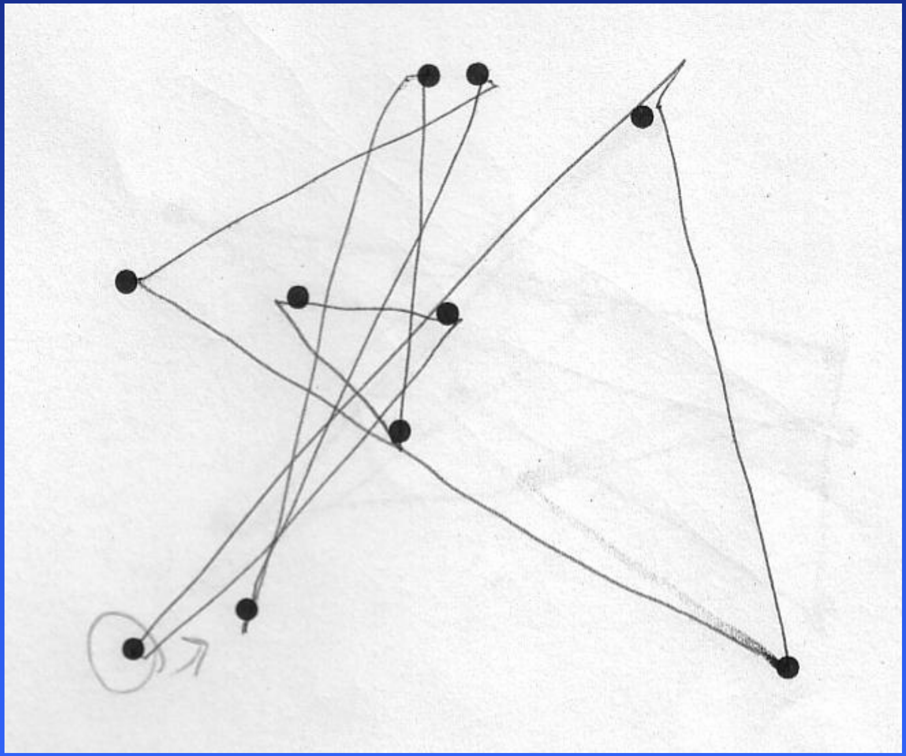

Optimizing and "pessimizing"

- Experiment 1 suggests that humans are much better at optimizing than they are at pessimizing
- Furthermore, they outperform the nearest neighbour heuristic when optimizing, but significantly underperform the furthest neighbour heuristic when pessimizing
- Experiment 2 explored these issues with more participants

Experiment 2

- Factorial design: Type of tour production x Type of instruction
  - participants optimized
  - participants pessimized
  - participants followed NN instructions
  - participants followed FN instructions

Example of heuristic-following instruction

- Draw a line starting at any circle connecting through **all** the other circles on the page and end your line at the same circle you started at.
- <u>**Visit each circle only once.**</u>
- Start and finish at the **same** circle.
- <u>Construct your lines as follows: from your chosen starting circle, **connect to the FURTHEST AWAY unconnected circle**. Then, from that circle, again connect to the furthest away unconnected circle. Continue this method until you get back to the circle you started at.</u>

Participants and procedure

- 114 psychology undergraduates participated for class credit. They had no knowledge of TSP research from the class.
- Each participant randomly assigned to one of the four conditions, and completed a booklet of the same 6 15-point problems in random order.
- Most participants finished within 15 minutes, but no individual timing data available.

Caveat

- My lab has only just finished coding the huge amount of data from this experiment
- I am presenting a preview of the findings, without statistical analysis at this point
- However, the differences appear very clear

Results – long tours

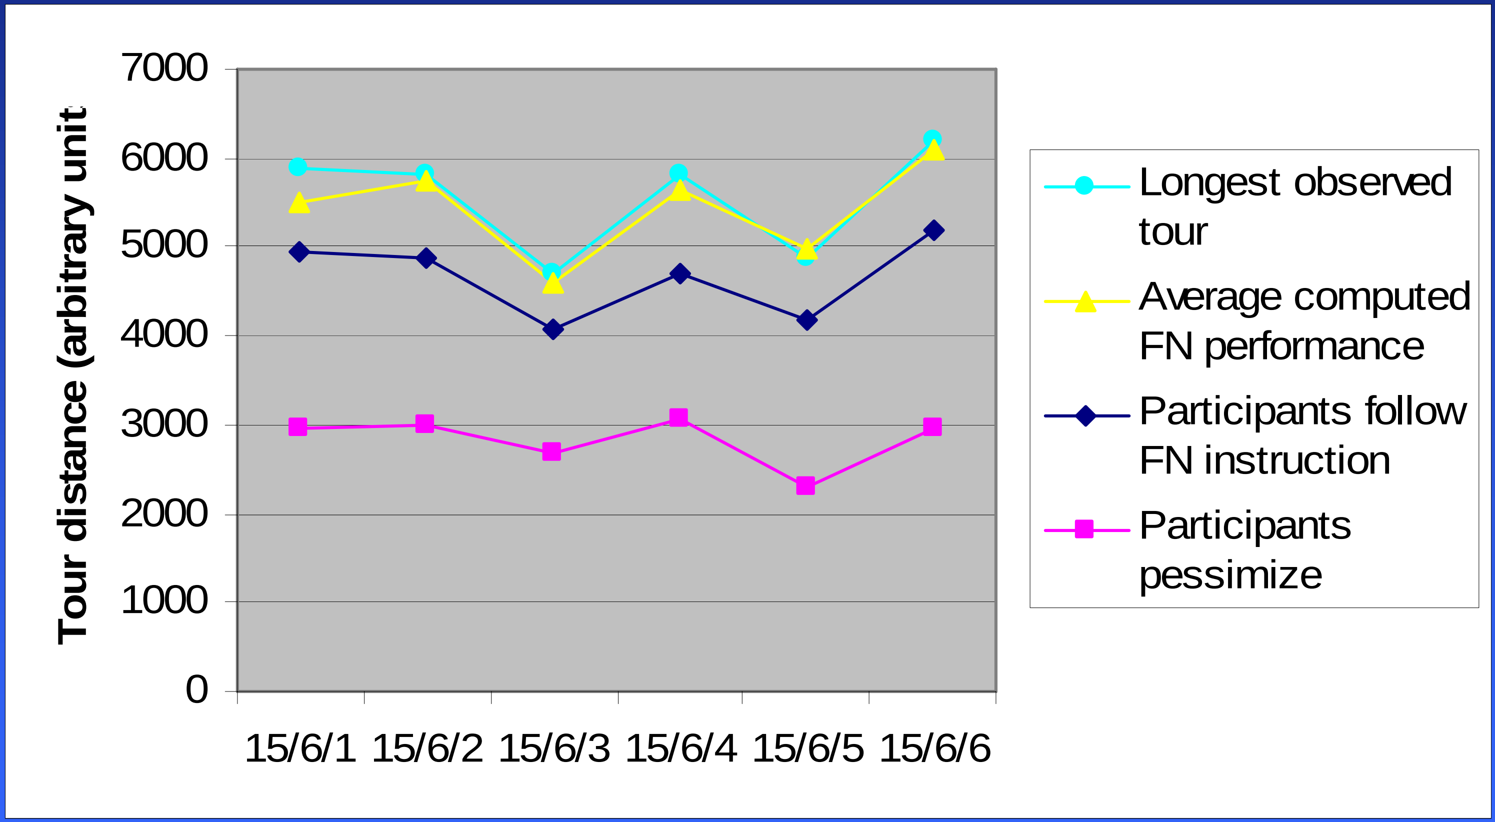

Results – short tours, same scale

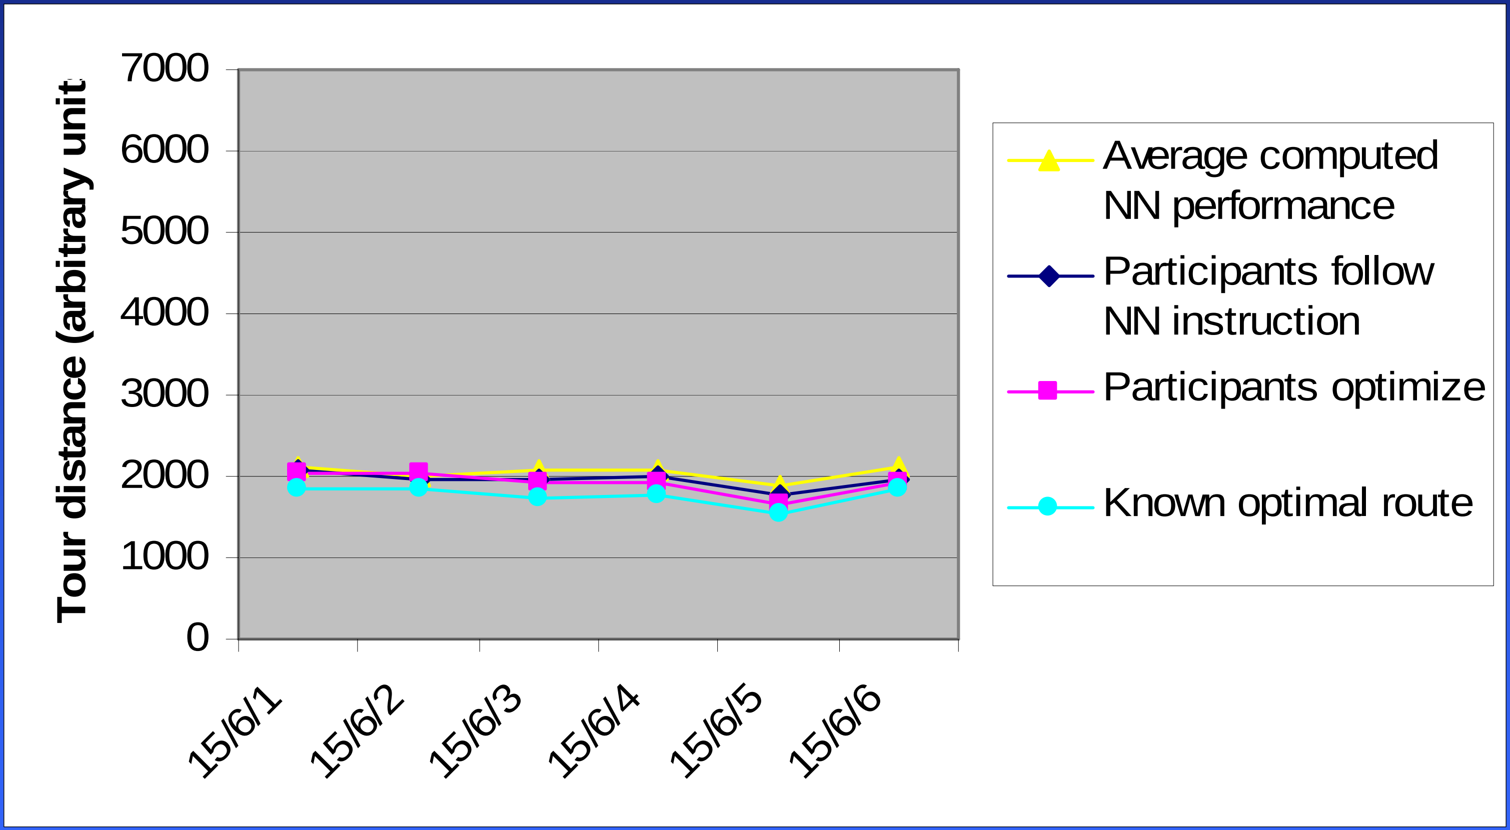

Results – short tours, expanded scale

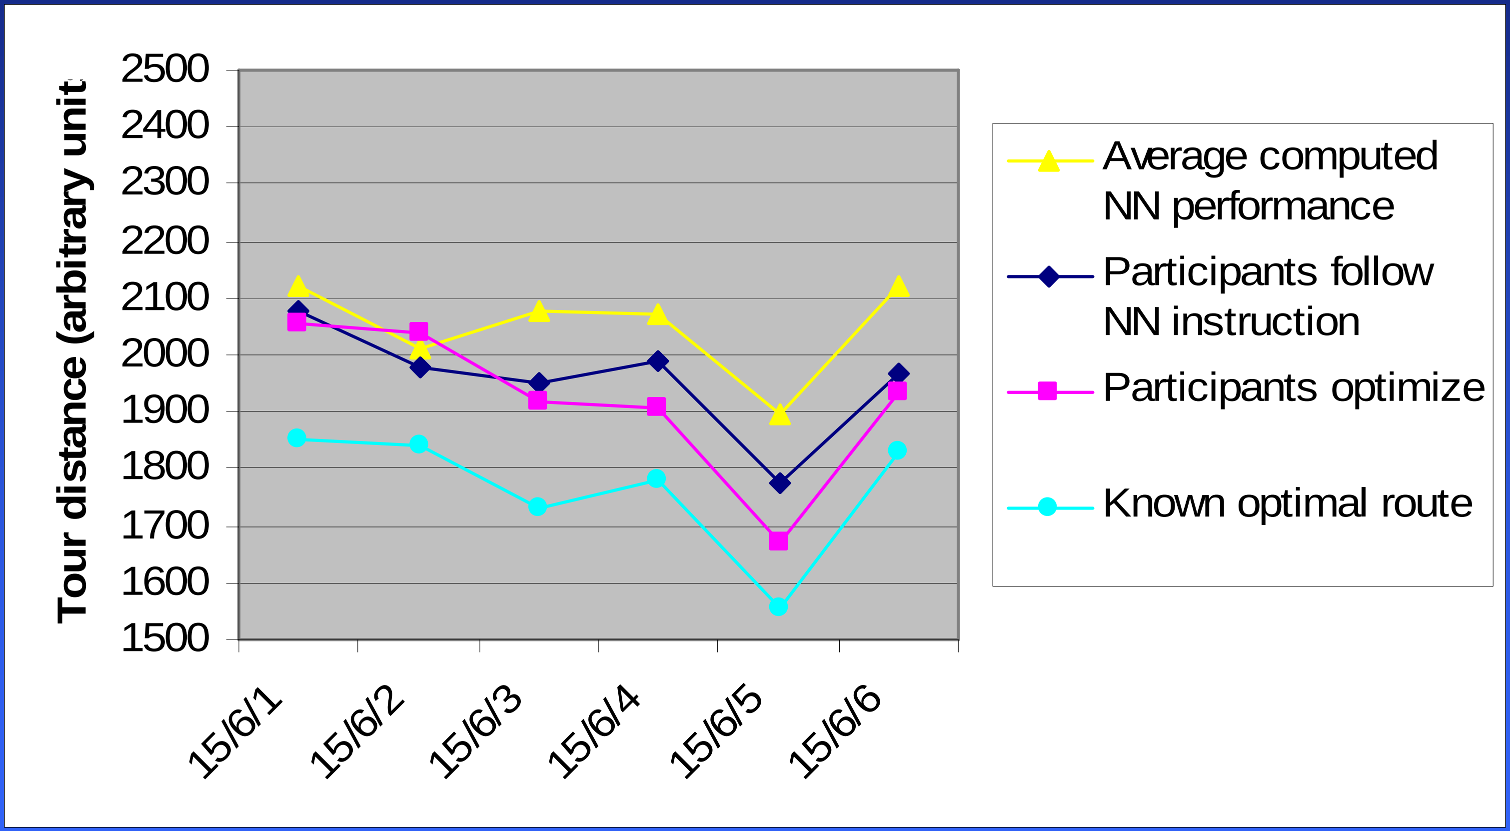

Summary of data

- Human participants performed well with the optimization task: tours were 6% above optimal in Experiment 1 and 8% above optimal in Experiment 2
- However, they were poor at pessimizing! ~16% below maximal in Experiment 1, and 49% below maximal in Experiment 2

Implications

- Seems unlikely that the optimizing and pessimizing tasks reflect the operation of identical generic spatial cognitive processes
- The striking dissociation between optimizing/pessimizing and heuristic-following performance is suggestive of separate processes

Where do human optimization heuristics originate?

- Tempting to speculate that optimization ability might confer some evolutionary advantage, à la Gigerenzer.
- By the same token, it seems unlikely that we would have evolved heuristics that let us do “as badly as possible”

- This does not really tell us much about the processes behind optimization heuristics

Optimization in other human systems

- Intriguingly, we are beginning to see evidence that other human systems are capable of giving close-to-optimal performance

- In recent work with John Findlay and Simon Liversedge (University of Durham, UK) we have begun to look at eye-movement patterns

One subject’s scan path

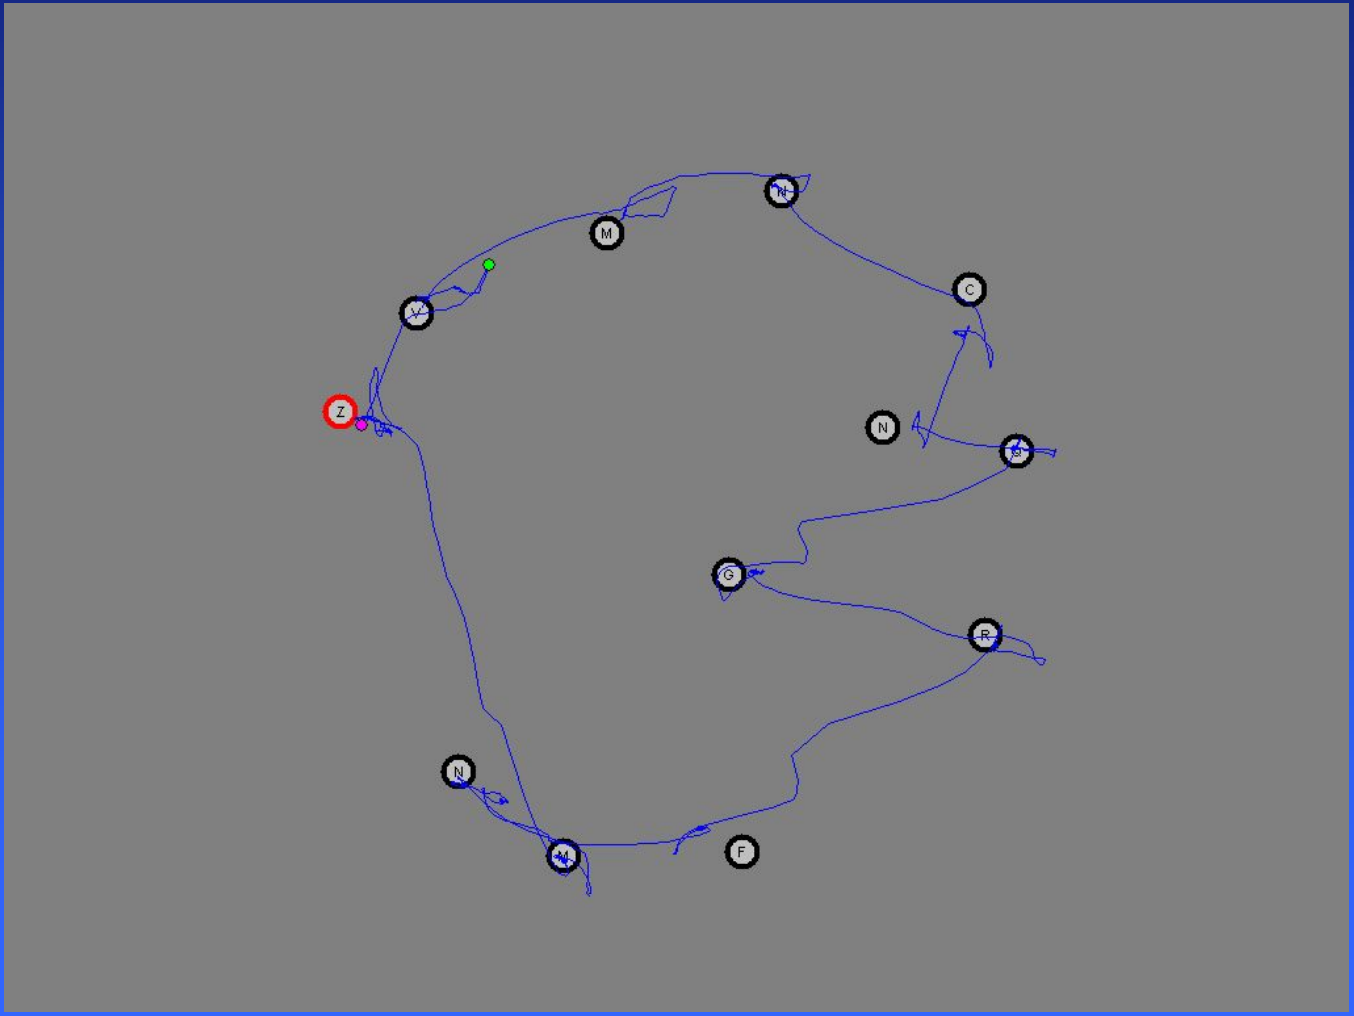

Contour/boundary detection in human vision

- It seems plausible to suggest that optimization heuristics may have their roots in the competencies of the visual system
- Low level visual processes have as major tasks:
  - boundary detection
  - contour finding
- Under certain circumstances, curvilinear contours are preferred by the visual system
- A convex-hull heuristic for TSPs may reflect a natural ability of vision.

Conclusions

- Humans seem to optimize better than they pessimize
- These early data will require follow-up over a range of experimental settings
- In my view, human optimization heuristics are likely to be related to some of the fundamental low-level competencies of the visual system