SAND2006-3370P

A Fitness Landscape Analysis of the Vehicle Routing Problem and Implications for Metaheuristic Behavior

Jean-Paul Watson
jwatson@sandia.gov
Discrete Algorithms and Mathematics
Sandia National Laboratories
Albuquerque, NM USA

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

Talk Caveats

- I *don't* care about "beating benchmarks"
  - Operations Research modus operandi for metaheuristics papers
  - Hooker (1994,1995)
  - Engineering (in the best case), not science

- I *do* care about analyzing existing state-of-the-art algorithms
  - Watson et al. (JAIR, 2005)
  - Watson et al. (C&OR, 2006)

- Why?
  - To build behavioral models of metaheuristic algorithms…
  - which can be used to inform the design process…
  - and build a theory of metaheuristics…
  - to move beyond "iterative hacking"

Talk Outline

- Background
  - Capacitated Vehicle Routing (CVRP)
  - Tabu Search for the CVRP
- Motivation
- Analyzing the Run-Time Dynamics of Tabu Search on the CVRP
  - Where is tabu search spending time in the search space?
  - How can the behavior of tabu search be modeled?
- Fitness-Distance Correlation and the “Big Valley”
  - Background and introduction
  - Prior investigations on the CVRP
  - Re-considering fitness-distance correlation in the CVRP
  - Implications for pool-based metaheuristics
- Conclusions

Background: Tabu Search and the CVRP

- Experimental context
  - The standard capacitated vehicle routing problem (CVRP)
  - Start simple, because we don’t yet understand simple

- Tabu search algorithms are very effective metaheuristics for the CVRP
  - Taillard et al.
  - Gendreau et al. (TabuRoute)
  - Cordeau et al. (Unified Tabu Search)
  - Ho and Gendreau

- Tabu search state-of-the-art through early 2000, but no longer
  - Evolutionary strategies (Mester and Braysy, 2004)
  - Memetic algorithms (Prins, 2004)

Motivation(s)

- The Central Dogma of Metaheuristic Theory
  - Performance=f(fitness landscape, search strategy)
- Implications
  - A metaheuristic is effective if and only if the search strategy is strongly "correlated" with, i.e., exploits, the fitness landscape
  - If you don't understand the fitness landscape structure, how can you design effective metaheuristics?

- How many fitness landscape analyses have you found on the CVRP?
  - 1.5 – despite the central role of landscape structure in performance

- To better understand what tabu search algorithms for the CVRP are really doing, why they are so effective, and how they can be further improved

- Disclosure: I don't have all the answers, and I raise more questions than answers

Tabu Search CVRP Run-Time Dynamics: Preliminaries

- One advantage of benchmarks and benchmark-beating papers
  - Optimal and/or presumed-optimal solutions are widely available

- Experimental question
  - Search is *supposed* to be driven toward the optimal solution…
  - but where is search really spending most of its time?

- Experimental methodology
  - Grab an optimal/presumed-optimal solution from the literature
  - Implementation of Taillard-like tabu search with (1,0) neighborhood
    - For now, without frequency-based long-term memory
  - Run algorithm for 100,000 iterations
  - Record "distance" to reference/optimal solution every iteration
  - Distance principle: Are customers in the same tour in both solutions?

Tabu Search CVRP Run-Time Dynamics: Results (1)

•Test instance: CMT 2 (n=75)

•Search appears to be largely stuck near solutions that are 0.5 distance away from the optimal solution

•There are periodic "bursts" of progress toward the optimal solution – and this is when new best-so-far solutions are identified

•99+% of search is expended in regions of the space that are nowhere near the optimal solution

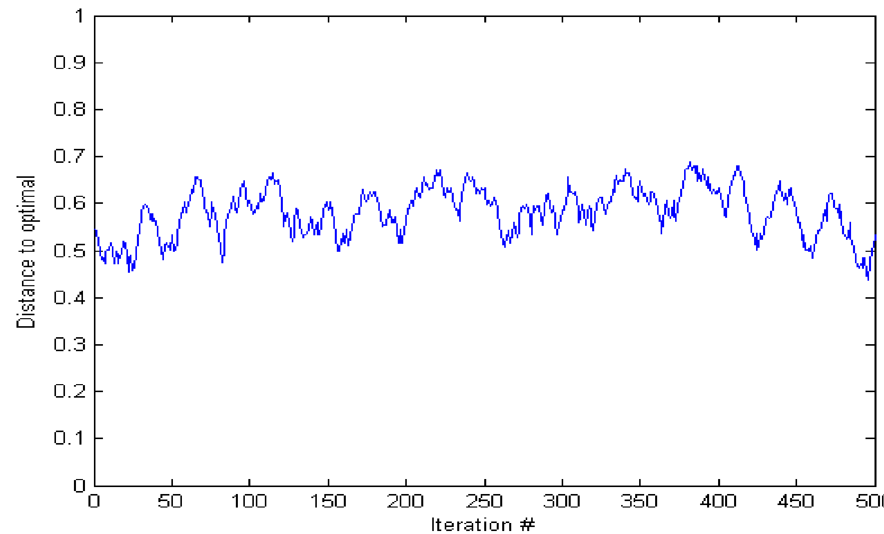

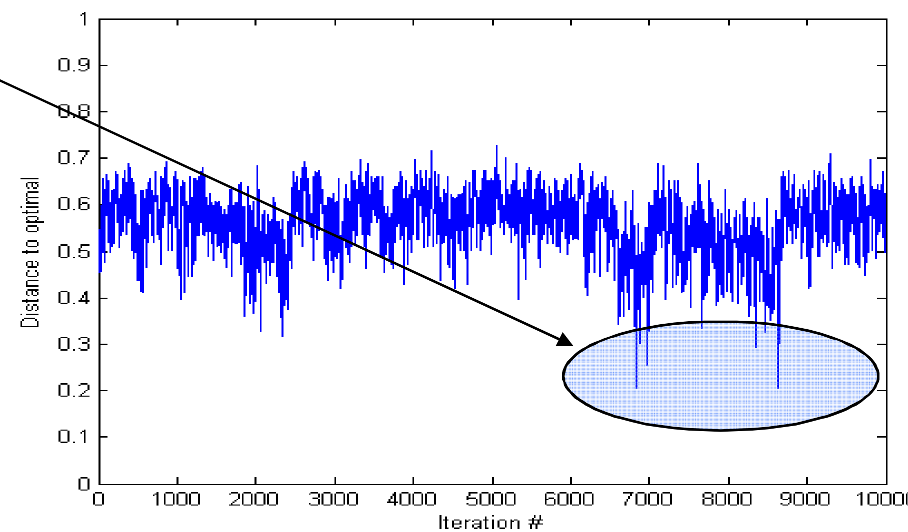

Sandia National Laboratories

Tabu Search CVRP Run-Time Dynamics: Results (2)

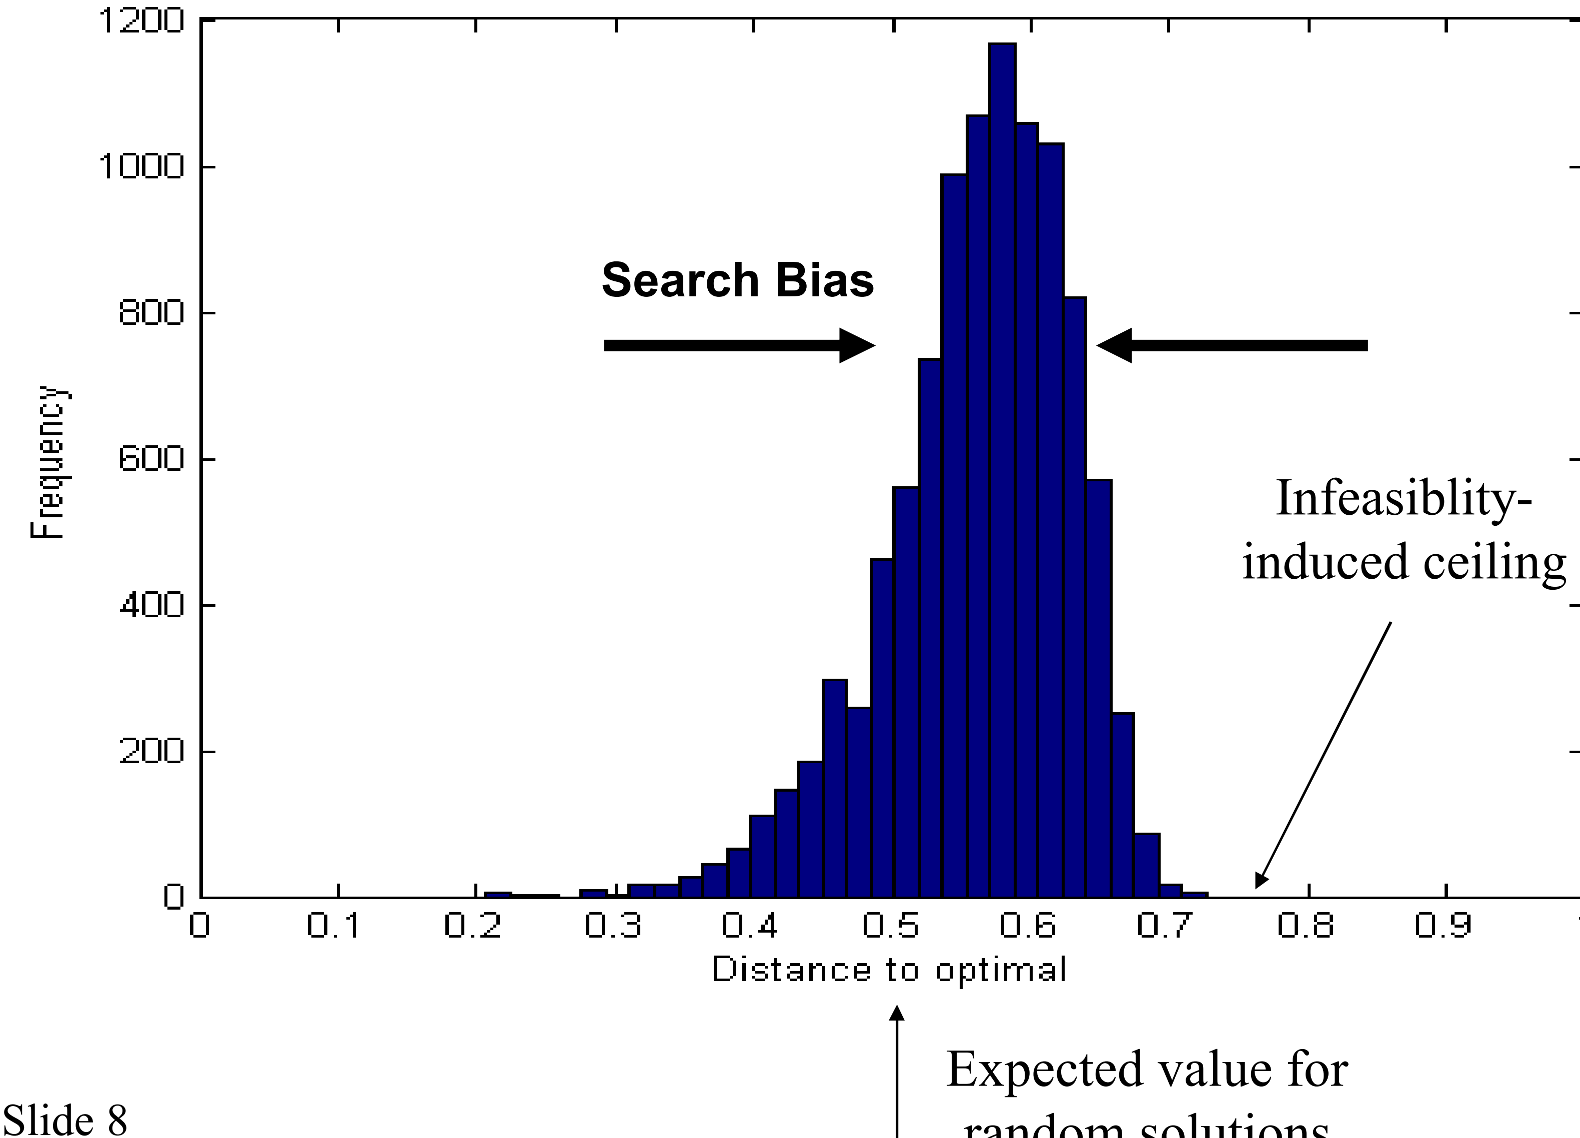

Tabu Search CVRP Run-Time Dynamics: Commentary (1)

- The long-term dynamics of basic tabu search in the CVRP can be modeled as a random walk with a strong bias toward "average-distance" solutions
  - Ehrenfest diffusion model with a restorative force (Feller, Volume 1)
  - Extension of prior results for job shop scheduling (Watson et al, 2005)
- What causes the restorative force?
  - The underlying representation
  - Given that search is near a reference solution, there are always more moves away from the reference solution that toward it
- Improving search efficiency and effectiveness is easy
  - "Just" spend more time in the left tail of the distance-to-optimal distribution
  - If you're not doing this, then you're likely wasting your time
  - If you're not doing something qualitatively different than existing algorithms, you're wasting everyone's time…

Tabu Search CVRP Run-Time Dynamics: Commentary (2)

- In case you were wondering…

- What happens when you add in long-term memory?
  - Fundamental nature of the histograms does not change
  - Discernible (order-of-magnitude) growth in the left-tail mass

- What happens when you allow search in infeasible space?
  - Early results indicate further increase in the left-tail mass
  - Not claiming solid results yet, as replicability of some algorithms is more difficult than expected

Fitness-Distance Correlation: Introduction

- What is Fitness-Distance Correlation (FDC)?
  - Take an <u>individual</u> problem instance
  - Generate a large number of local optima
  - Compute the quality of each optimum
  - Compute the distance between each local optimum and the globally optimal solution
  - FDC is the (Pearson's) correlation coefficient for distance vs. quality

- Another variant of the FDC methodology
  - Compute average distance between a local optimum and all other local optima in the sample

- So what?

How to Exploit FDC? (1)

- High FDC => the existence of a "Big-Valley" local optima structure
  - Boese et al. (1994)
- Also found in other NP-hard problems
  - Traveling Salesman Problem
  - Flow-Shop Scheduling Problem
  - Graph Bi-Partitioning
  - …
- Two exploitation mechanisms
  1. Intensification
  2. Path relinking

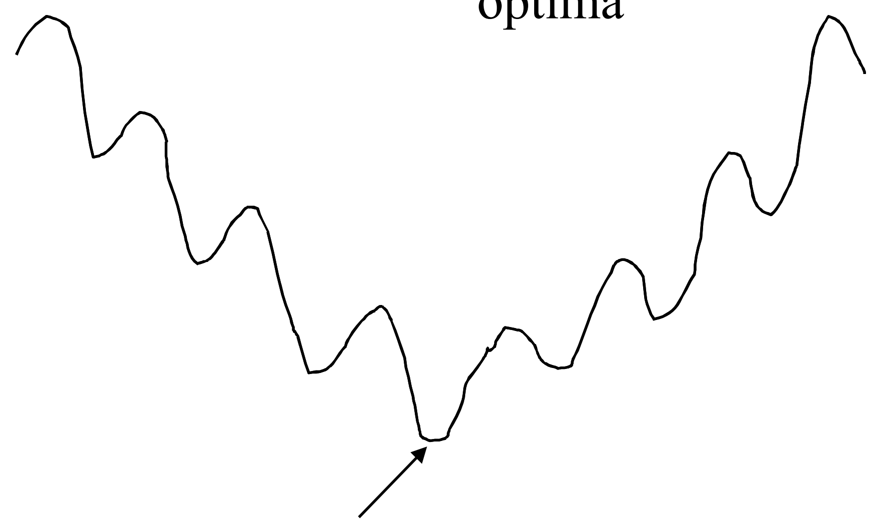

Sandia National Laboratories

How to Exploit FDC? (2)

- Maintain a pool *P* of elite solutions
  - Small, typically $|P| < 10$
  - Initialize using $|P|$ best solutions encountered by a moderate-length run of tabu search
- General design philosophy
  - Short bursts of tabu search from carefully selected initial solutions
- Intensification
  - Pick an elite solution *p* at random
  - Perform a short run of tabu search from *p*
  - If better solution found, replace *p* in *P*
- Diversification
  - Pick two elite solutions *p1* and *p2* at random
  - Perform path relinking to generate a solution “in between” *p1* and *p2*
  - Perform a short run of tabu search from p
  - If better solution found, replace worst of *p1,p2*
- Perform reintensification and diversification with equal probability
  - Only diversification => might miss good nearby solutions
  - Only intensification => you might not be “deep” enough in the big-valley

An Example of High Fitness-Distance Correlation

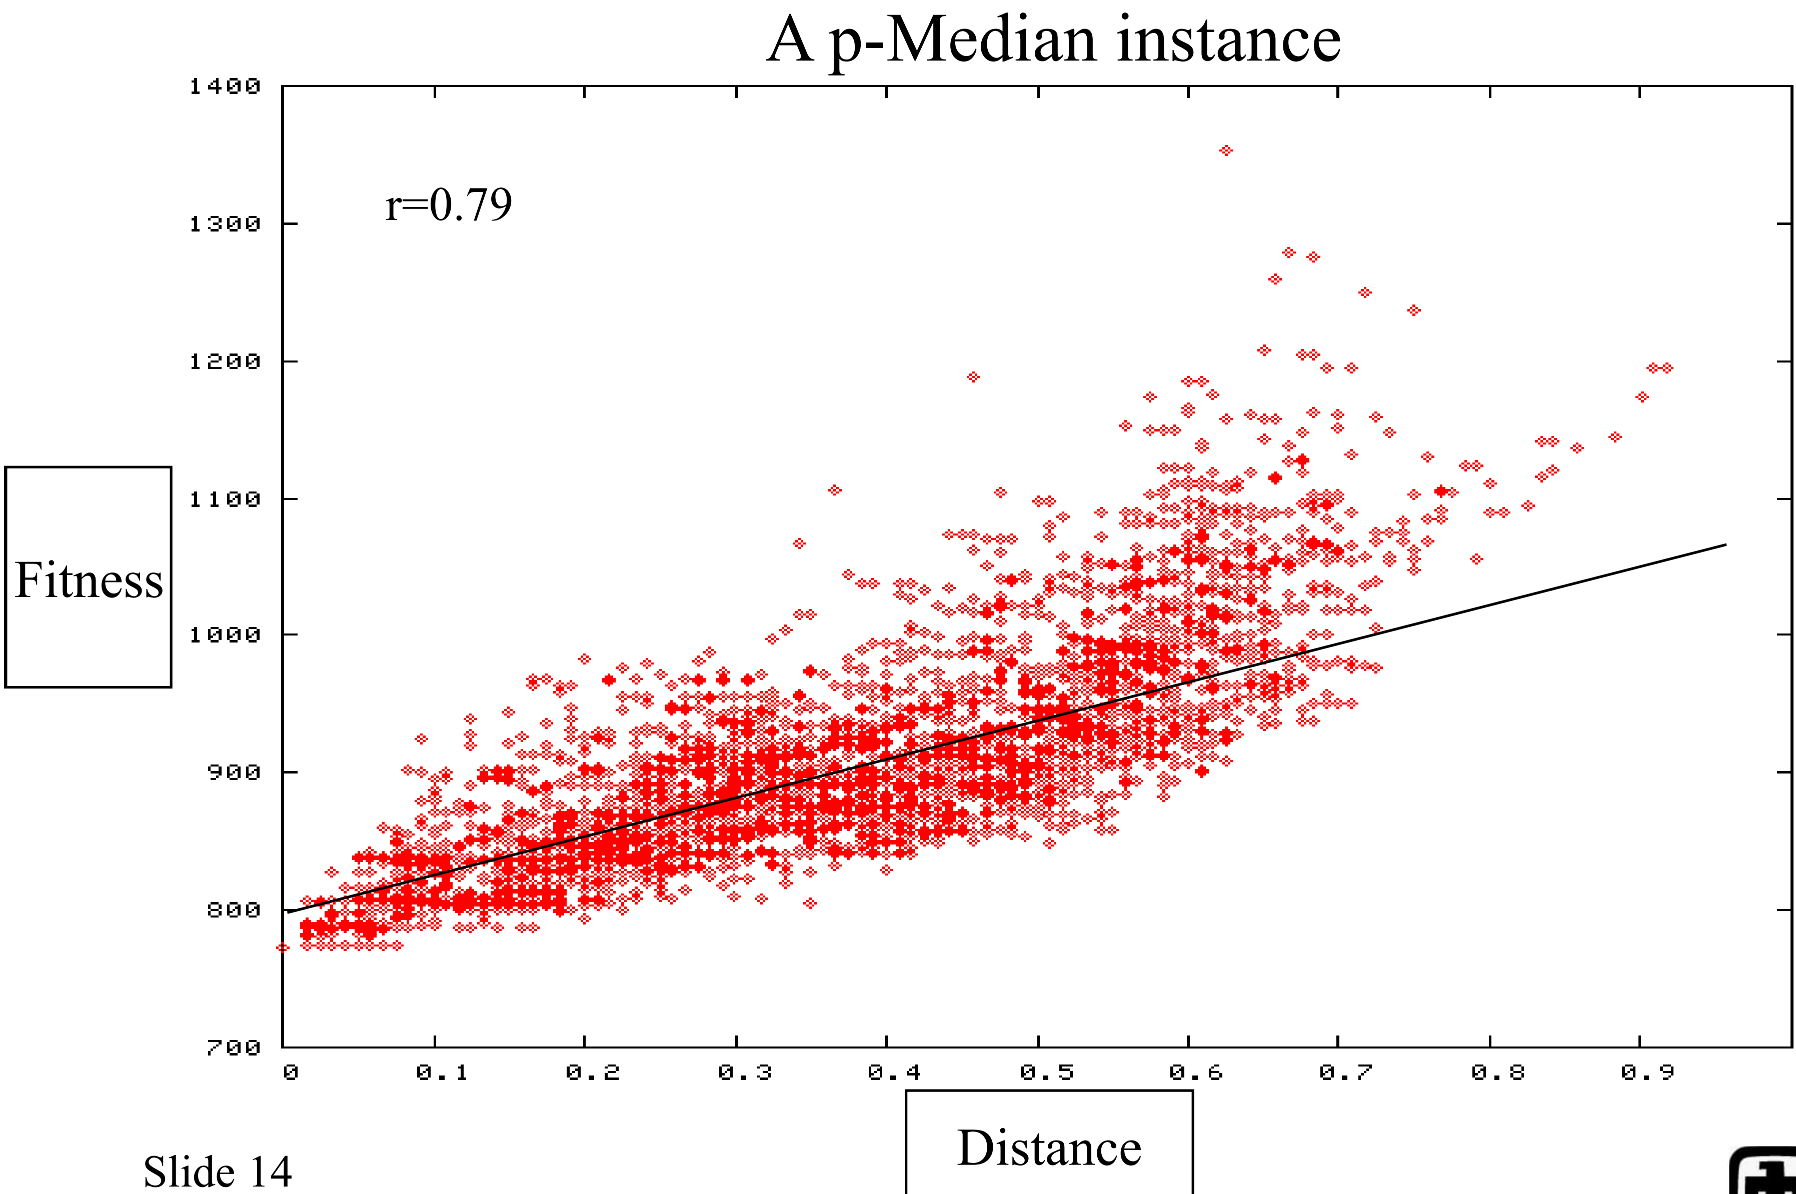

An Example of Low Fitness-Distance Correlation

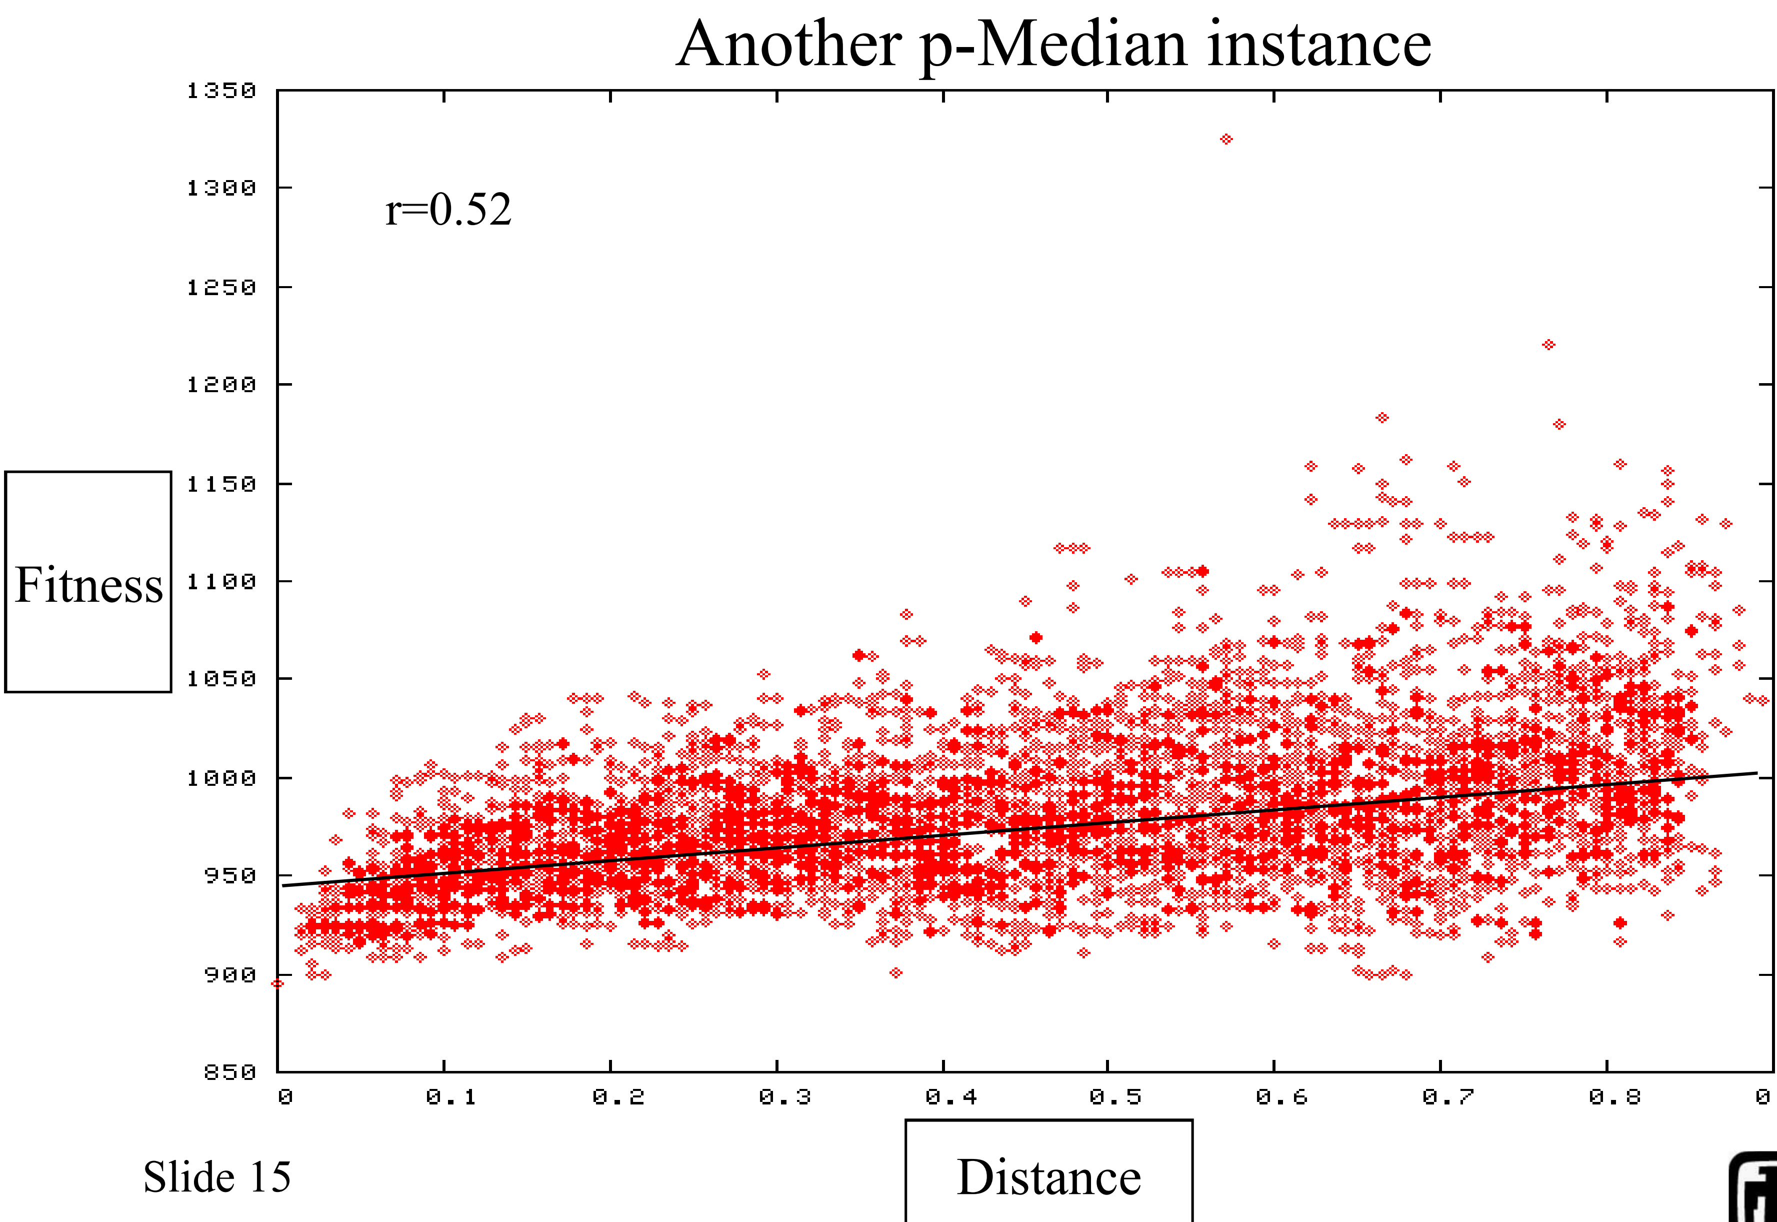

"Online" versus "Offline" Fitness-Distance Correlation (1)

- Previously reported FDC for the CVRP range from ~0.7 to ~0.8
  - Random local optima (Kubiak, 2004)
- But: Random local optima aren't likely to enter an elite pool
  - Look at local optima *actually visited* by tabu search over 25K iterations

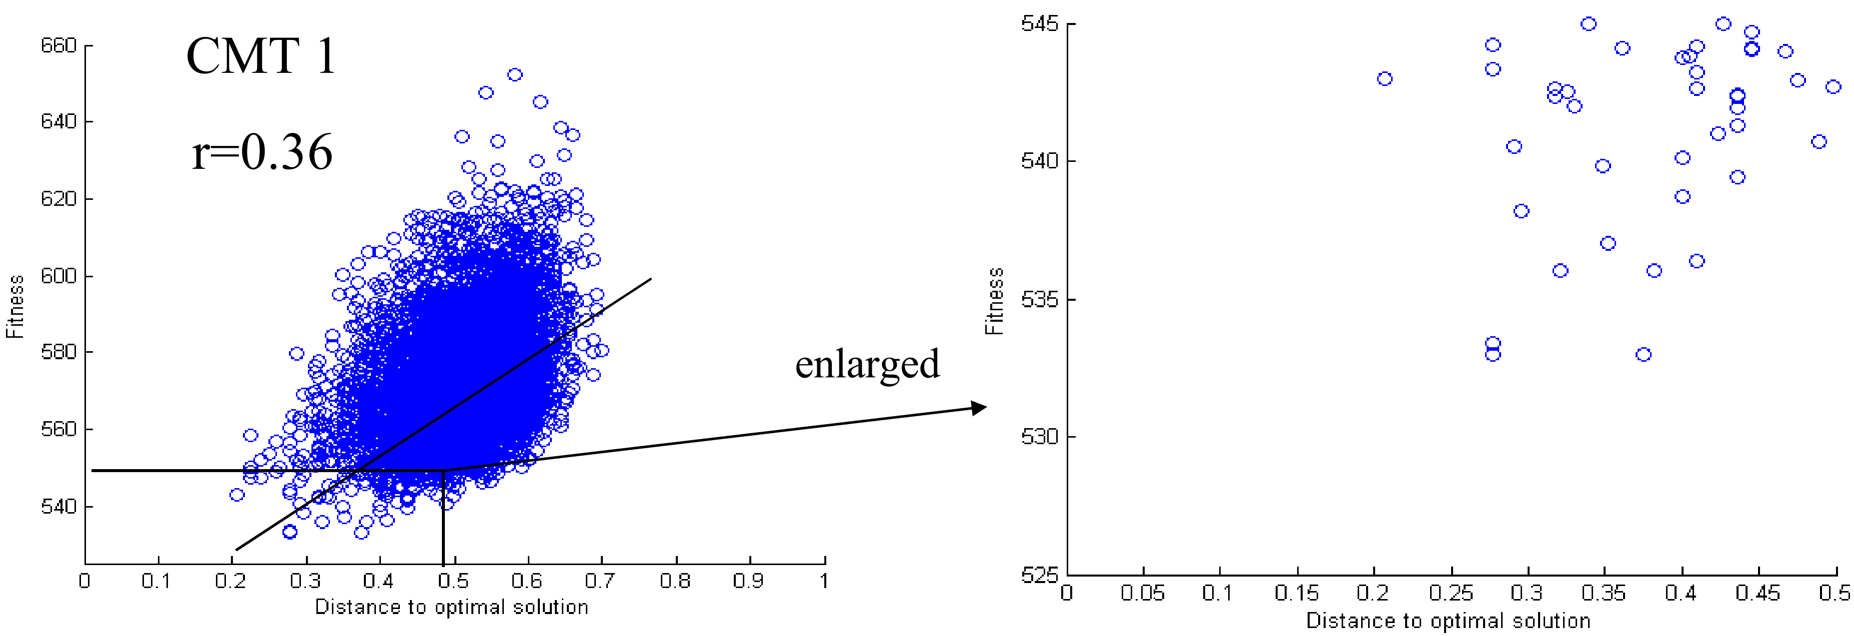

"Online" versus "Offline" Fitness-Distance Correlation (2)

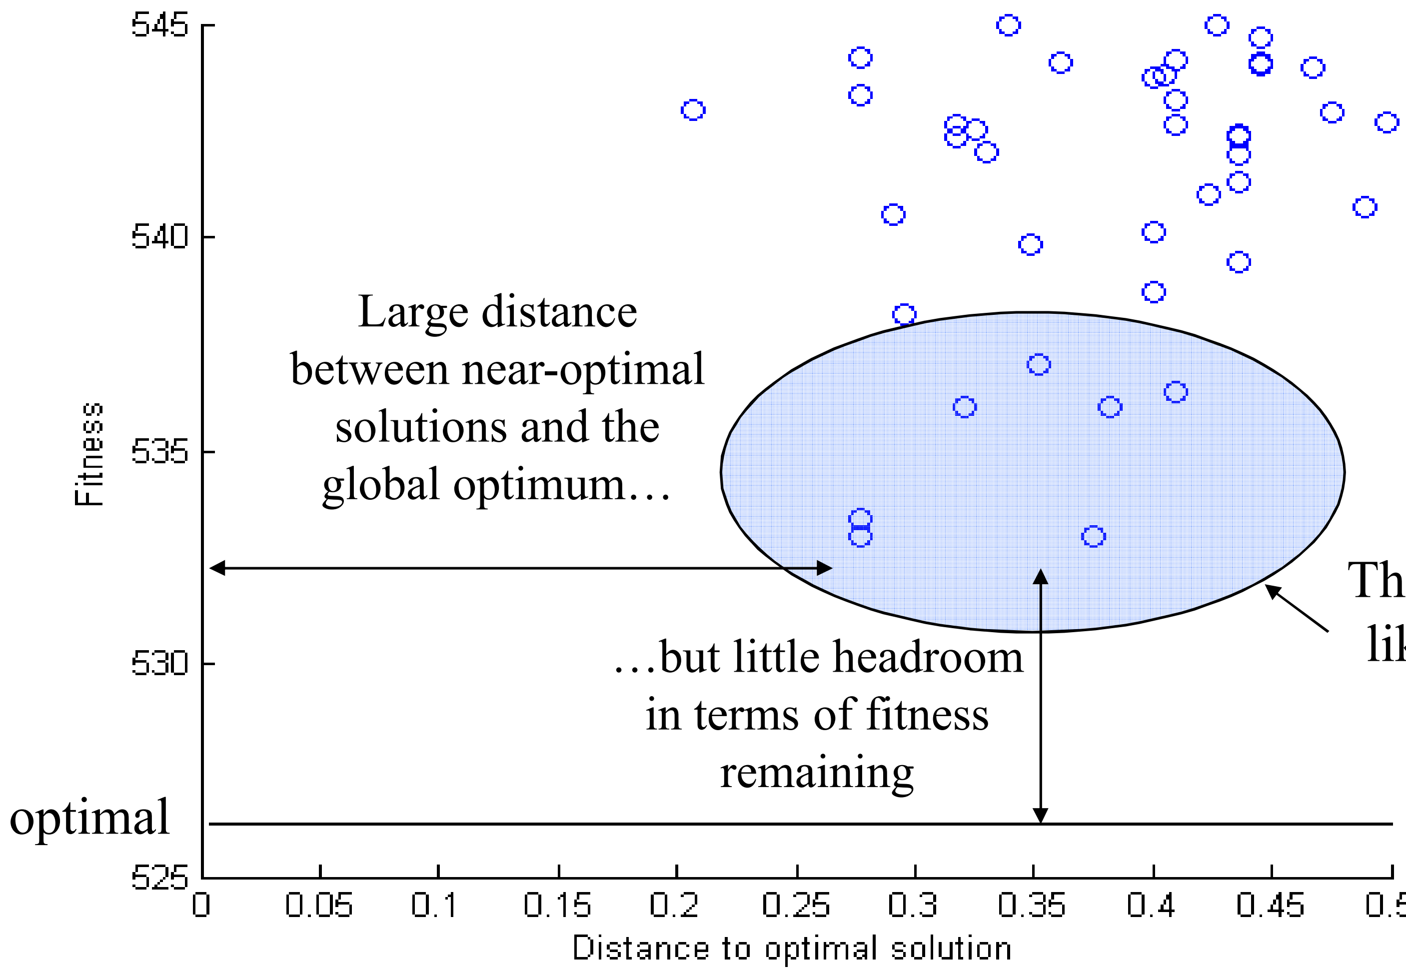

The Implications of Low FDC for Pool-Based Metaheuristics

- Key observation for low FDC in the CVRP
  - Near-optimal solutions are frequently very far from optimal solutions

- Implication for (re)intensification
  - (1) High-quality solutions are very distant from optimal solutions
  - (2) Naïve pool maintenance schemes
  - (3) Search is quickly driven toward “0.5” distance solutions
  - => Intensification is likely to fail far more often than not

- Implication for diversification
  - (1) High-quality solutions are very distant from optimal solutions
  - (2) High probability that two elite solutions are not “deep” in the valley
  - => Path relinking is likely to search sub-optimal regions of the space

Conclusions

- Empirical analysis of fitness landscape – metaheuristic interactions can lead to significant insights into the behavior of these algorithms
  - Even state-of-the-art metaheuristics are wasting 99+% of their time
  - Fundamentally new mechanisms are needed to overcome this "feature"
  - Simple analyses can determine whether a newly proposed mechanism is qualitatively changing the way search is being performed
- Fitness-distance correlation based on random local optima is flawed
  - If you consider local optima likely to be visited by an algorithm, FDC in the CVRP is substantially lower than previously reported
  - Low FDC causes unique difficulties for pool-based metaheuristics

Questions?

- Thanks!