Optimizing Spacecraft Risk Management Using Heuristic Search

Martin S. Feather, Steven L. Cornford, Julia Dunphy, Jose Salcedo, Tim Menzies

300 char summary:

Risk reduction is critical to mission success, yet reduction measures cost time, $, etc. We show the role of heuristic search techniques (Simulated Annealing, Genetic Algorithms and Machine Learning) in determining an optimal risk-reduction strategy, balancing value of risk reduction against cost.

Full abstract:

The Problem:

Spacecraft designers seek to achieve ambitious mission goals while constrained by severely limited resources. In this context, design success is critically dependent on risk management. The many forms of risk that threaten mission success must be tamed. However, risk-reducing options incur resource costs, and so must be selected judiciously. This is fundamentally an optimization problem: given resources to apply to risk reduction, find the way to apply them to maximize likelihood of mission success and science value.

In recent years JPL has been pioneering the quantitative management of risk through the Defect Detection and Prevention (DDP) process http://ddptool.jpl.nasa.gov. This has been used in design of spacecraft technologies and, ongoing, for assisting mission assurance planning for MSL.

The quantitative nature of DDP's risk management makes it amenable to treatment as an optimization problem. However, it is a challenging optimization problem, because:

(1) The design space is very large. For example, Hicks&Cornford et al's study of holographic memory technology identified 99 risk reduction options. Each option is an independent choice of something to apply or not. Thus the number of combinations of choices is 2^99, which is approximately 10^30.
(2) The design space is very intertwined. In the above study, there were 440 links between the 99 risk reduction options and the 69 risks, and in turn there were 352 links between those 69 risks and the 32 requirements defining the mission goals and constraints.
Manual exploration of such large and complicated design spaces is unsatisfactory. The outcome is often selection of an inferior design, because whole areas of superior designs go unexplored. It is also an inefficient process, very expensive in terms of experts' time.

The Solution:

We have had success using heuristic search techniques able to explore a large space of design options to locate (near) optimal designs within that space:

Genetic Algorithm (GA). We used GA to optimize risk reduction for a given level of resources. The GA has the desirable property of exploring a whole swathe of solutions at once, so the net result is not just a single solution, but a set of such, which can include variants among which the experts can then choose.

Machine Learning. We collaborated with Prof. Menzies (WVA/NASA IV&V) to apply his machine-learning algorithm. This both identifies how to move towards an optimal design, and does so by identifying the most critical decisions first (e.g., on the holographic memory, it identified the 33 most critical options).

Simulated Annealing (SA). We used SA to quickly optimize for several forms of desiderata: maximal risk reduction for given resources, minimal resources to get down to a given risk level, or a combination of the two. Studies with this method have also served to confirm that Menzies machine learning technique indeed lead to (near) optimal solutions.

California
Institute of
Technology

Optimizing Spacecraft Risk Management Using Heuristic Search

Martin S. Feather, Steven L. Cornford,
Julia Dunphy & Jose Salcedo

Jet Propulsion Laboratory
California Institute of Technology

Tim Menzies
West Virginia University

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Funded by NASA's:
- Code Q FDPP program
- Code Q/IV&V ARRT task, and
- Code R ECS program

Motivational Insights

California Institute of Technology

Assurance activities "filter out" risk - Dr. Steve Cornford

"Risk as a Resource" - Dr. Michael Greenfield

overfiltered risk *unfiltered risk*

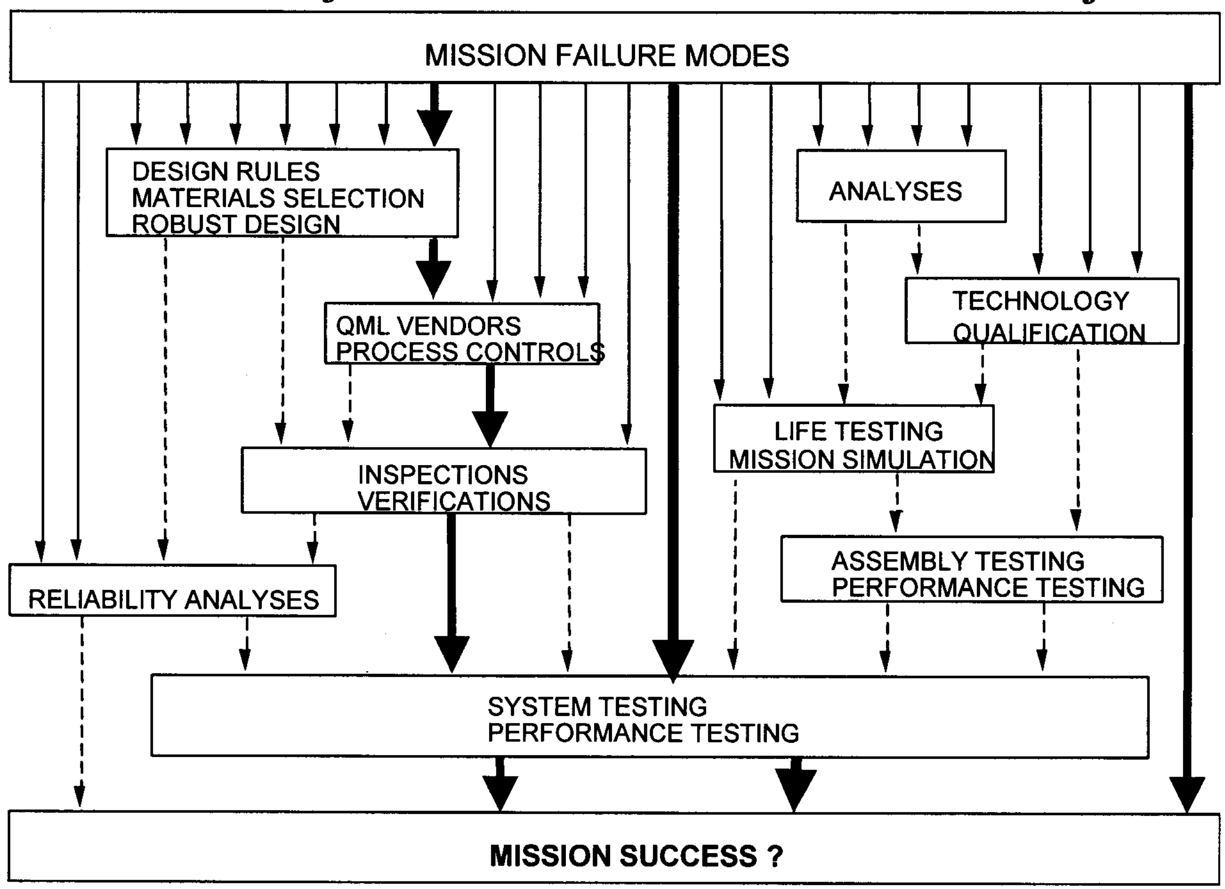

Optimizing ***singly filtered risk***

Risk Management Optimization Goals

The selection of activities such that:

For a **given set of resources**
(time, budget, personnel, test beds, mass, power, ...)
benefits are maximized

or

For a given set of objectives
(science return goals; on-time and in-budget development; 99+% expectation of successful landing)
costs are minimized.

JPL
California Institute of Technology

What's Needed to do Risk Management Optimization

1. A model to calculate assurance **costs** & **benefits**- we use Defect Detection and Prevention (DDP)

2. Data to populate the model - we populate with metrics from experience (when available) augmented with experts' best estimates

3. **Optimization** over the model - we have explored three heuristic search techniques: Genetic Algorithms, Machine Learning, Simulated Annealing

Costs & Benefits

California
Institute of
Technology

Activities have **costs**:

- Requirements inspections take skilled peoples' **time**
- Test-what-you-fly takes **high-fidelity testbeds**
- Radiation shielding takes **mass** and **volume**

Activities have **benefits**:

- Requirements inspections may **catch problems early**, when it is inexpensive to fix them
- Test-what-you-fly may **catch~~es~~ problems that would jeopardize the mission**
- Bounds checking may **decrease the frequency of switching into safe mode**

DDP Cost/Benefit Model

California Institute of Technology

Benefits = Σ attainment of objectives

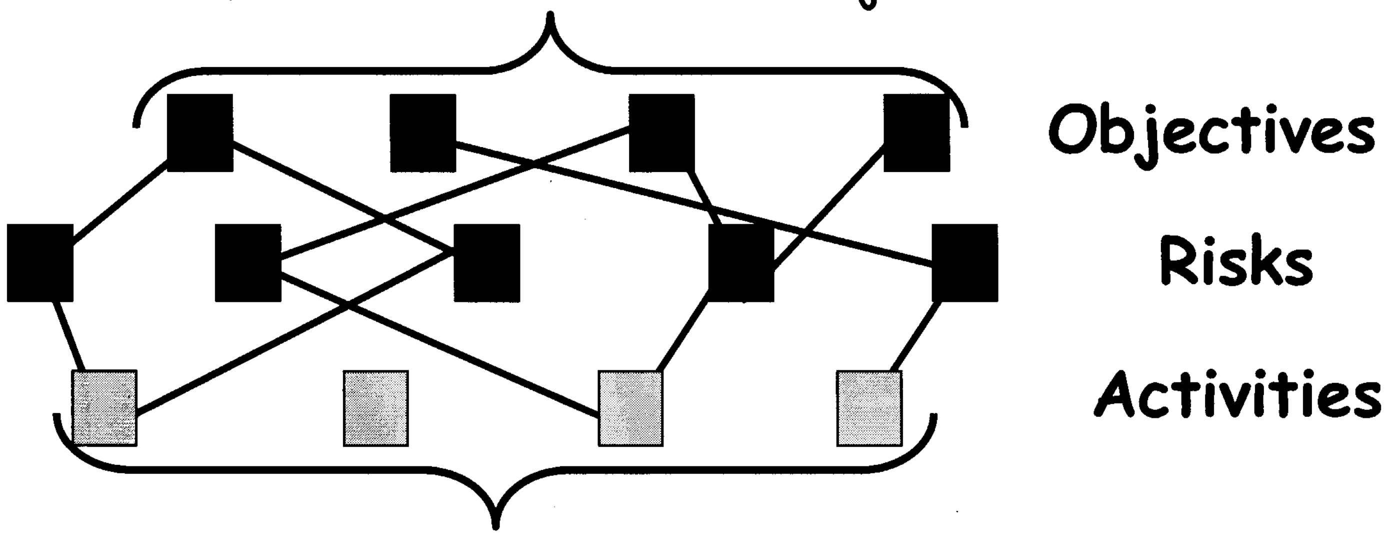

Costs = Σ costs of selected activities and the repairs they induce

Model holds *quantitative* measures of:
How much each risk impacts each objective, and
How much each activity reduces each risk.

Risks are crucial intermediaries in the model -
objectives impacted by **risks** to differing extents
activities reduce **risks** to differing extents

A Populated DDP Dataset (Real Data from Experts)

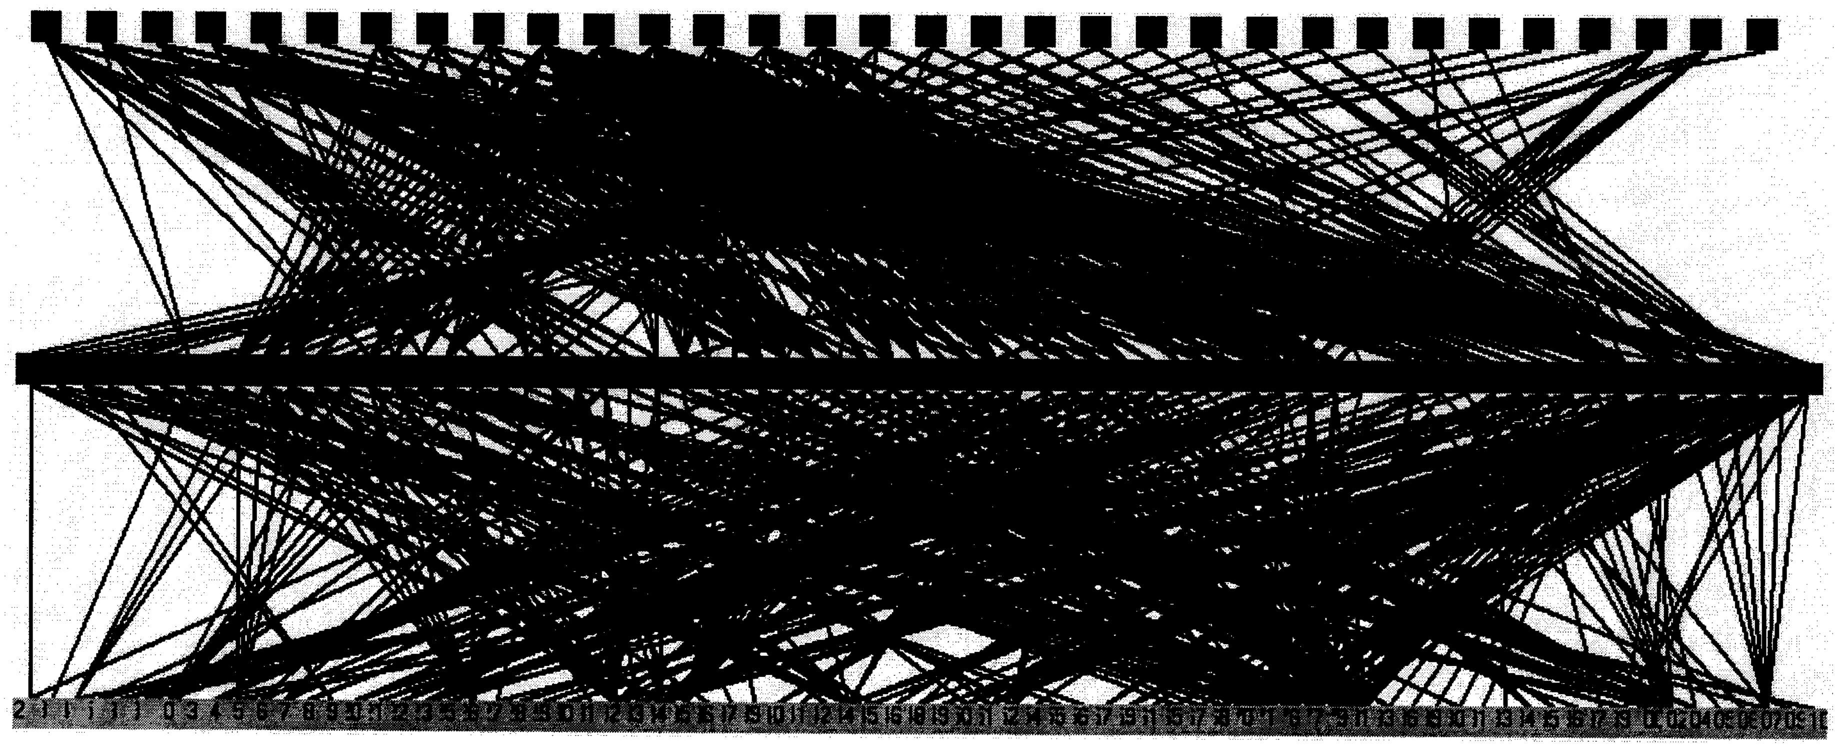

32 objectives, 69 risks, 99 activities
352 non-zero quantitative objective-risk links
440 non-zero quantitative activity-risk links

JPL
California Institute of Technology

Heuristic Search Studies

I) Genetic Algorithms

II) Machine Learning

III) Simulated Annealing

I) Genetic Algorithm Studies

California Institute of Technology

GA key ideas and **DDP risk optimization**

Work with a **population** of candidate solutions.

For DDP, a solution is a **selection of activities**.

Mutate each candidate, and score the results.

For DDP, a mutation is achieved by **changing the selection of activities**. A score is computed from the DDP calculated **cost and benefit combined into a single numeric score**
E.g., if cost < max then benefit else 0
Remember, cost is of activities and the repairs they induce, benefit is attainment of objectives

Favor the better scoring candidates for inclusion in next generation's population.

Repeat until converges.

I) Genetic Algorithm Studies

California Institute of Technology

Studies on optimizing for maximum benefit within cost ceiling

...Work with a **population** of candidate solutions.

Result is a set of alternative solutions (each a selection of DDP activities), not just a single one.

... **Mutate** each candidate, and score the results.

Customized the mutation step to efficiently generate ONLY candidates that were within the cost ceiling.

... **Repeat** until converges.

Rapid convergence observed.

Strengths:

Rapid convergence

Set of alternative solutions

Weaknesses:

If instead optimizing for min cost to get above a benefit floor, hard to *efficiently* generate candidates above that floor

Hard to maintain as DDP cost-benefit model gets more elaborate

Many obscure ways to "tune" GAs

II) Optimization Using Menzies' (*) Machine Learning based approach

California Institute of Technology

DDP

TAR2 (Menzies)

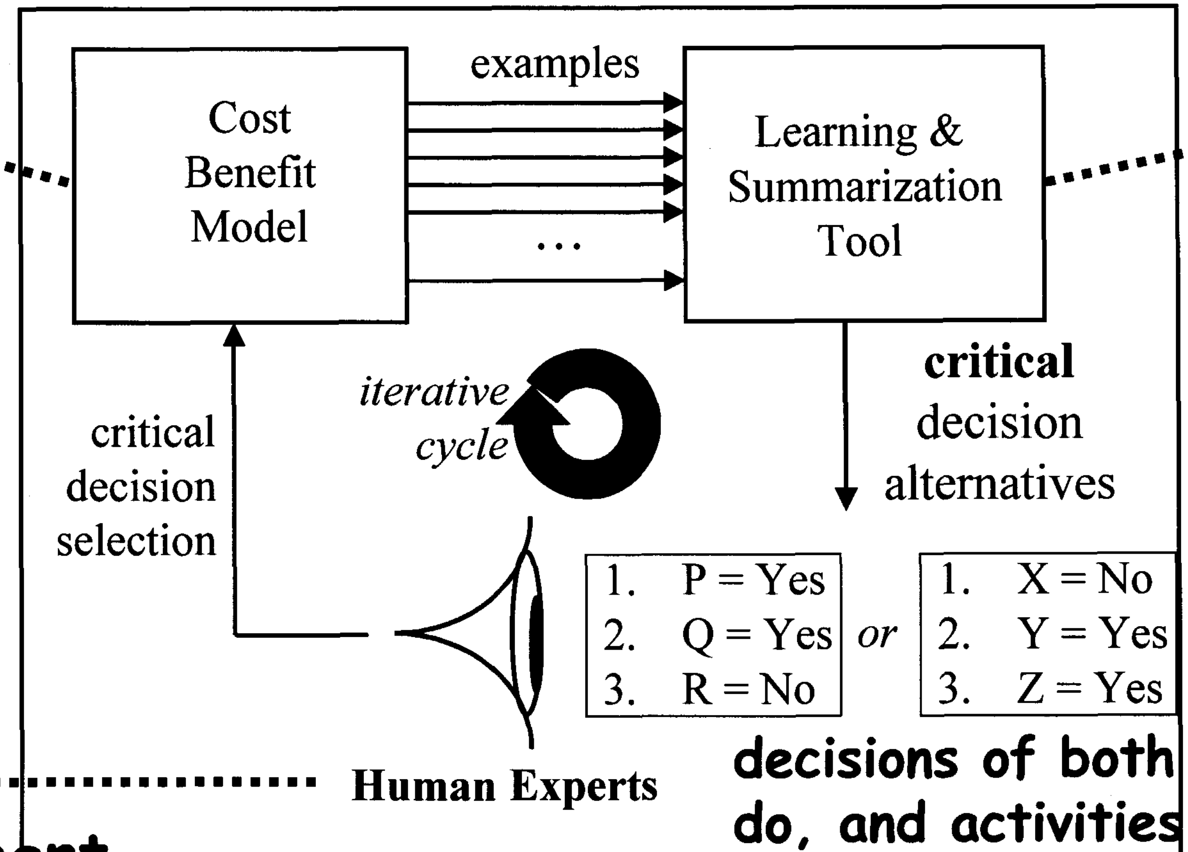

retains expert involvement

decisions of both activities to do, and activities to *not* do

*http://tim.menzies.com

Dataset *before* Optimization

California Institute of Technology

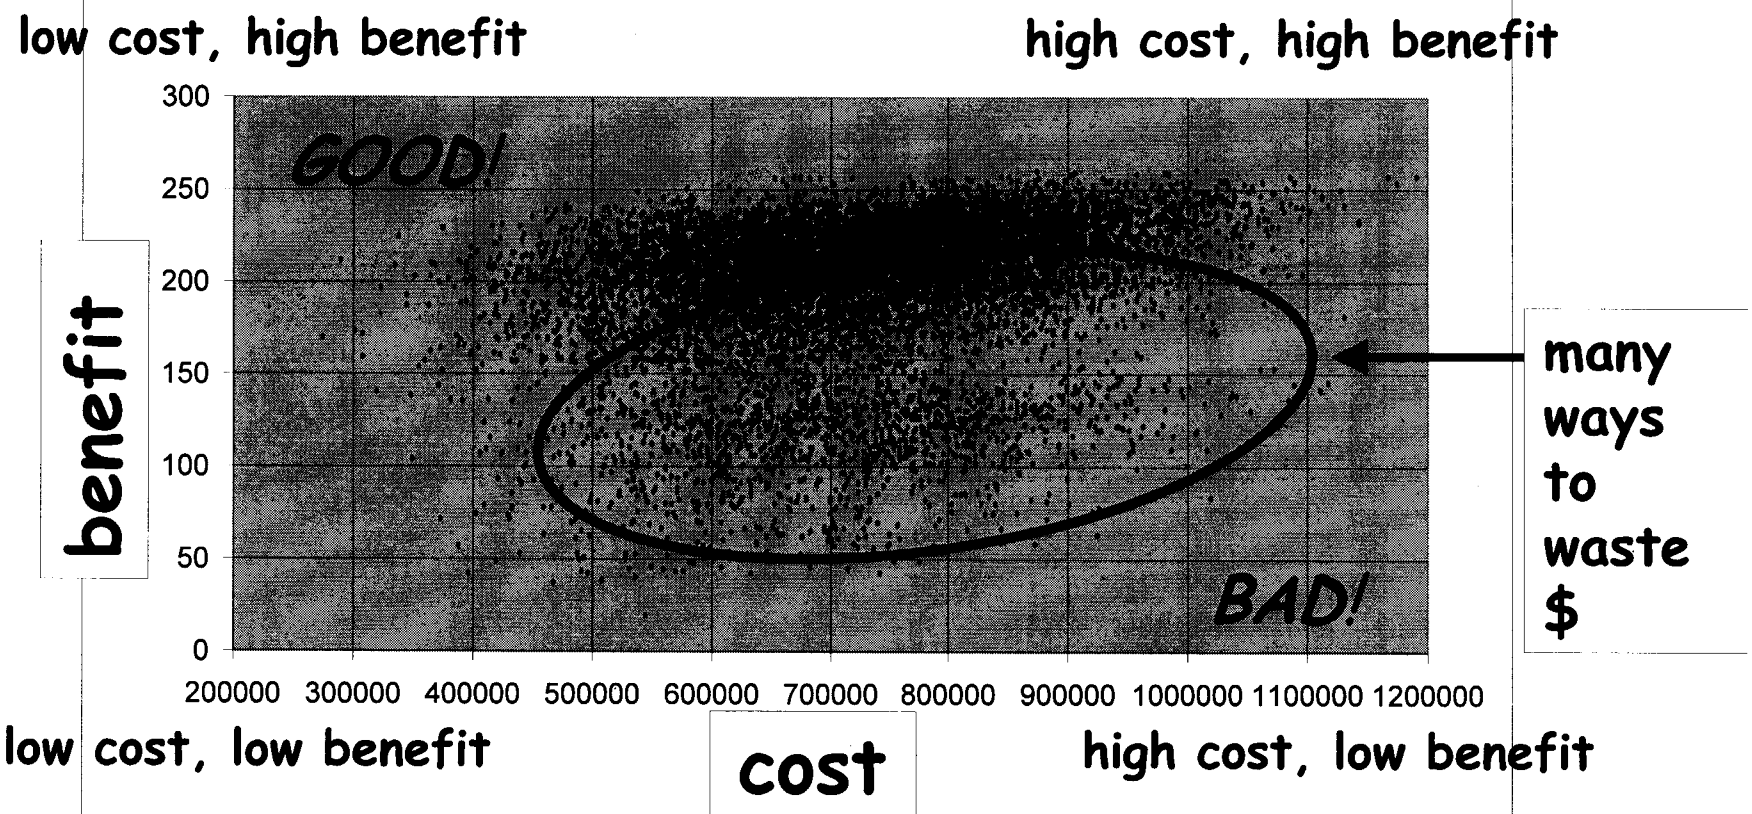

Each black point a randomly chosen selection of activities. DDP used to calculate **cost** and **benefit** of each such selection.

Dataset *after* Optimization

Each white point is an optimized selection of activities (33 critical ones are as directed by TAR2, other 66 chosen at random).

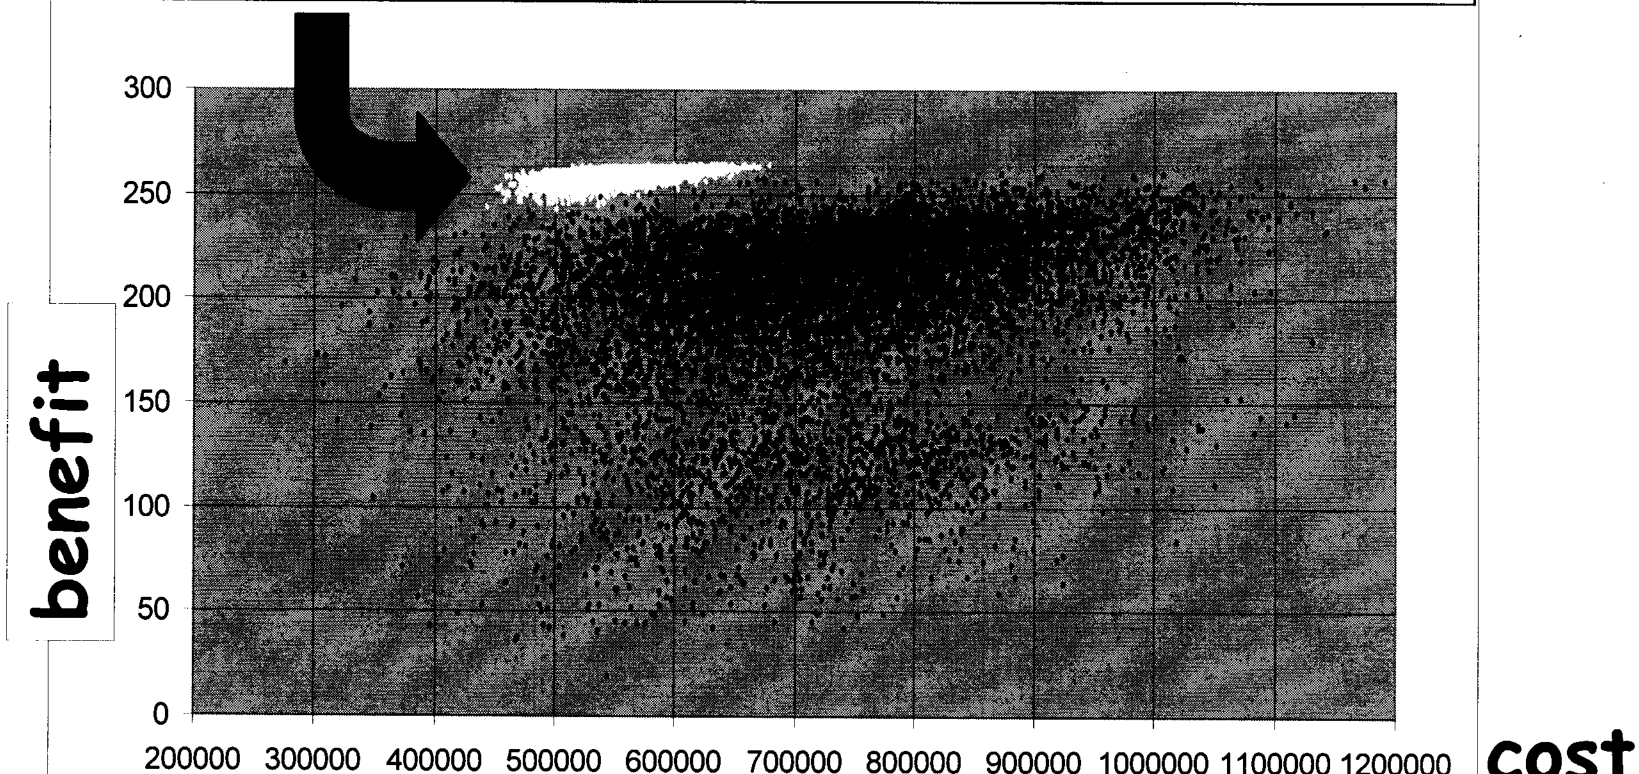

Menzies' TAR2 identified 33 most critical decisions:
21 of them activities to perform
12 of them activities to *not* perform.

II) Machine Learning Studies

California Institute of Technology

Strengths:

Identifies most critical decisions

Offers alternative solutions at each iteration (opportunity to introduce additional knowledge)

Robust – search mechanism need not know details how cost-benefit computations are performed

Optimize for min cost, max benefit, or combination of both

Weaknesses:

Somewhat slow – needs large number of examples, several iterations

Some manual control of TAR2 required

Additional comments:

Pilot study reported in: "Converging on the Optimal Attainment of Requirements" by Martin S. Feather & Tim Menzies, in *Proceedings, IEEE Joint International Conference on Requirements Engineering*, Essen, Germany, 9-13 Sep. 2002, IEEE Computer Society, pp. 263-270.

III) Simulated Annealing Studies

JPL
California Institute of Technology

SA key ideas and **DDP risk optimization**

Work with a single candidate solution; mutate this, and score the results.

For DDP, a solution is a **selection of activities** and score is DDP calculated **cost and benefit combined into a single numeric score**

If the mutation is an improvement,
then **continue from that superior mutation**,
else ***maybe*(*) continue from that inferior mutation**.

(*) likelihood of doing so depends on:

- how much inferior (the more inferior, the less likely)
- where in search process (the later on, the less likely) - intuition is of the "temperature" cooling over time

Simulated Annealing Optimization

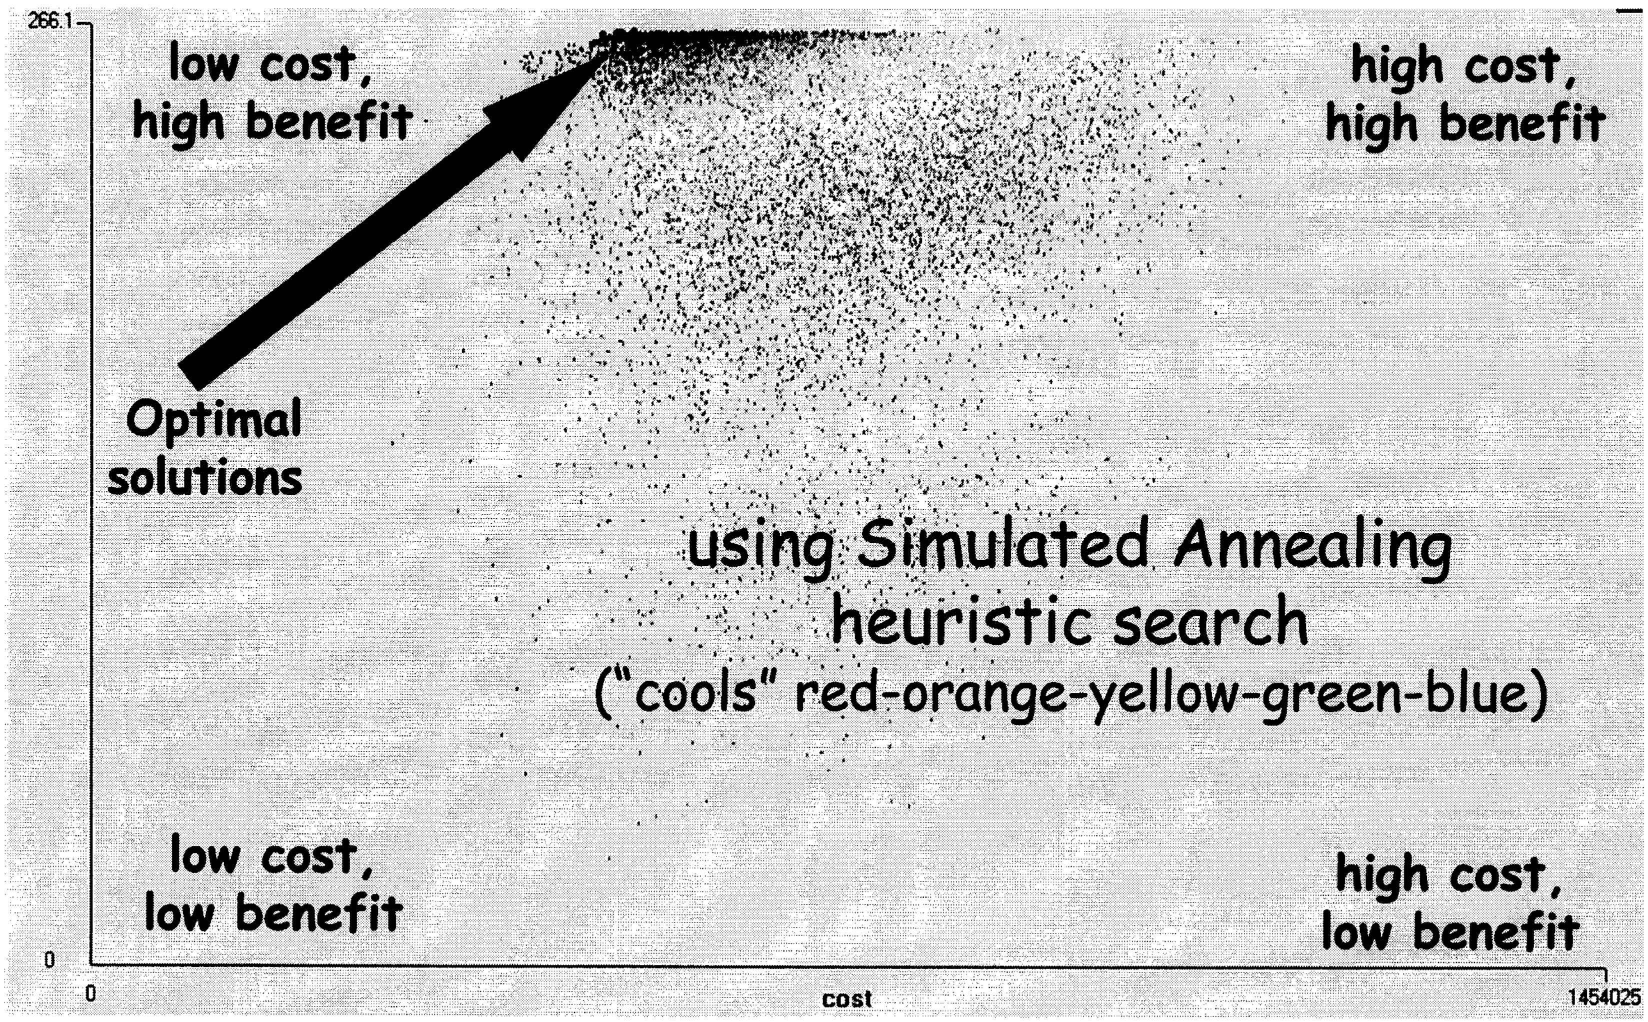

III) Simulated Annealing Studies

Strengths:

Fairly rapid convergence

Robust w.r.t. elaborations of DDP cost-benefit model

Optimize for min cost, max benefit or combination of both

Automatic (little or no tuning required)

Weaknesses:

Can't distinguish critical decisions

Works with singleton solution, not set of such (extended to track N-best solutions found so far)

Additional comments:

SA and Menzies' TAR2 compared on same dataset – confirm each getting to the same near-optimal region of solutions

SA now built in to DDP tool

Menzies has ideas on how to blend with his TAR2 Machine Learning (best of both worlds)

SA used successfully to reveal cost/benefit trade space for Ken Hicks & Ken Johnson's study of Chip-On-Board Technology Utilization

Optimizing Spacecraft Risk Management

California Institute of Technology

for more information:

Defect Detection and Prevention (DDP):

http://ddptool.jpl.nasa.gov

The lead for this is

Steven.L.Cornford@Jpl.Nasa.Gov

Tim Menzies & his "TAR2" system:

http://tim.menzies.com

tim@menzies.com

Advice, guidance, feedback, ideas, applications: all welcome!

Martin.S.Feather@Jpl.Nasa.Gov