GECCO

Genetic and Evolutionary Computation Conference

2015

Madrid, Spain
July 11-15, 2015

acm

Metaheuristic Design Pattern: Surrogate Fitness Functions

Alexander E.I. Brownlee[1]
John R. Woodward[1]
Jerry Swan[2]

[1] University of Stirling, UK
[2] University of York, UK

CHORDS
Computational Heuristics Operational Research Decision-Support

Outline

- Problem statement
- Solution
- Consequences
- Implementation
- Examples

Problem statement

- Metaheuristics need some notion of “fitness”

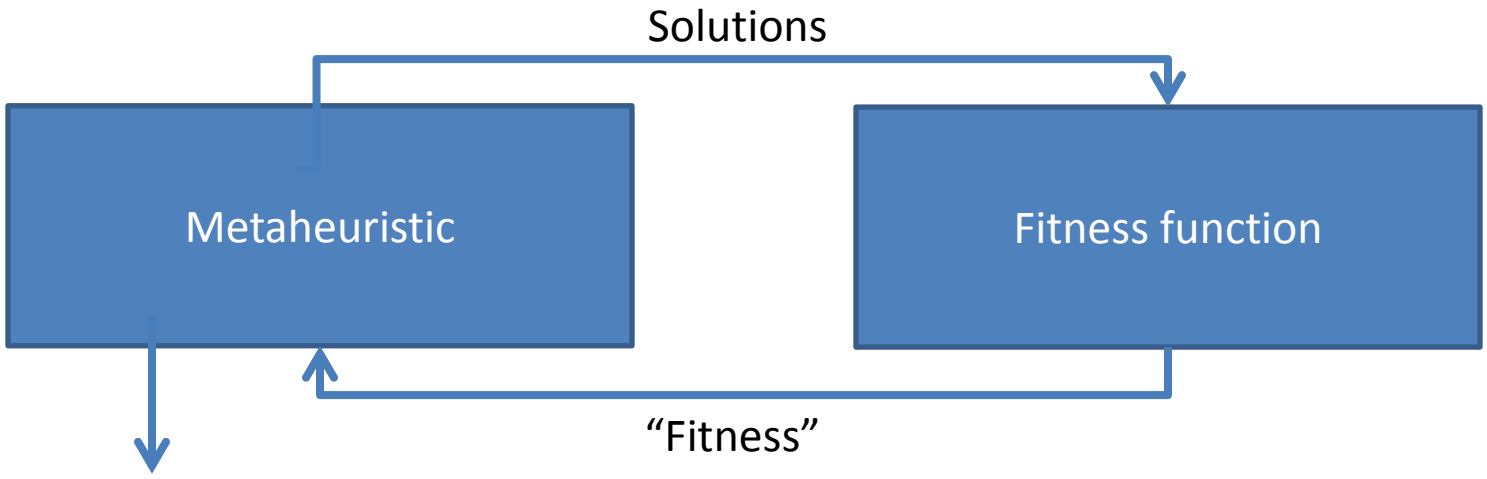

- Two purposes:
  1. Measure quality (wrt the *objectives+constraints*)
  2. Guide the search
- (1) is not necessarily suitable for (2)…

Problem statement

The *true* fitness function might:

1. Be costly
2. Be noisy
3. Not have a useful search gradient

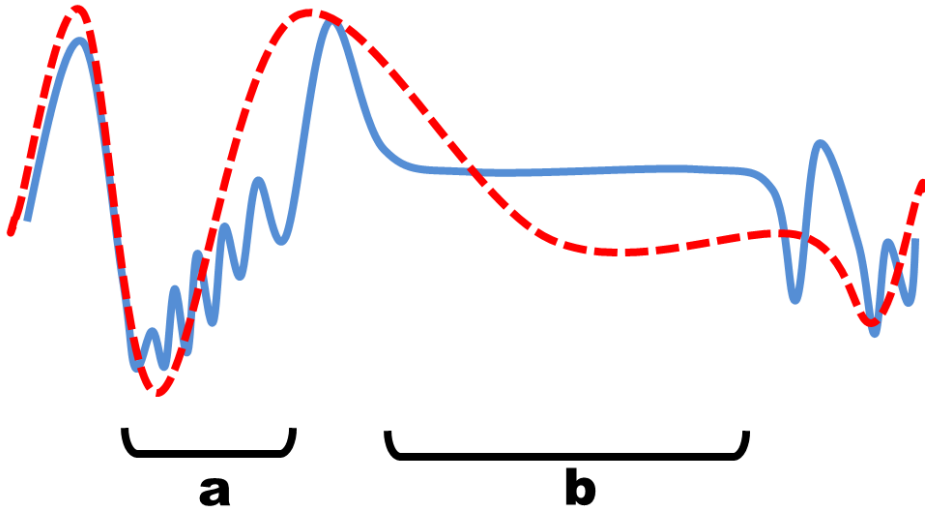

Solution

- *Surrogate* fitness function in place of “true” FF
  - Still need to refer to “true” fitness occasionally
- a.k.a. meta-model, proxy, fitness model or approximation
  - typically one of the above for costly problems, but less so for noisy problems or reshaping landscape
- Two types:
  - Static
  - Dynamic

Solution

- Static surrogates
  - part of problem definition
  - can include domain knowledge
  - typically guides search towards partial solutions
- Constraint relaxation, multi-objective weights
  - *might* be classed as surrogates
- Often used already!
  - We don't usually directly search the real world

Solution

- Dynamic surrogates
  - Regression or machine learning: polynomials, Kriging, artificial neural networks, interpolations
  - Fitness inheritance
- Trained using samples of “true” fitness
- Updated or replaced over time
  - bridge / handle / body pattern
- Ensembles combine strengths of many
  - “composite” pattern

Consequences

- Search landscape altered
- Approximation errors
  - Must make reference to objective function
  - e.g. Surrogate filters new solutions before full evaluation, or switch between surrogate & true
- Can offer speed up – but balance with overhead
- Surrogate explicitly models fitness: mine it to support decision making

Example 1

- Long-running simulations of building energy performance (mins to hrs)
- RBFN surrogate uses population as training data
- Filters offspring before evaluation with full simulation
- Many similar examples

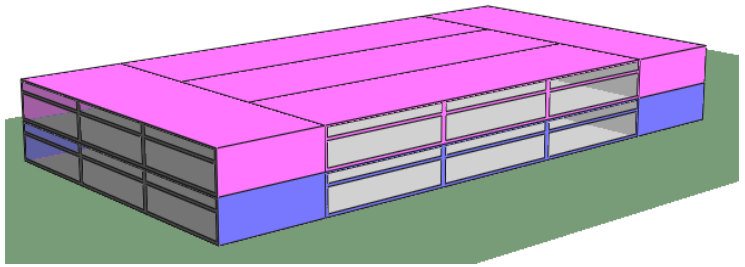

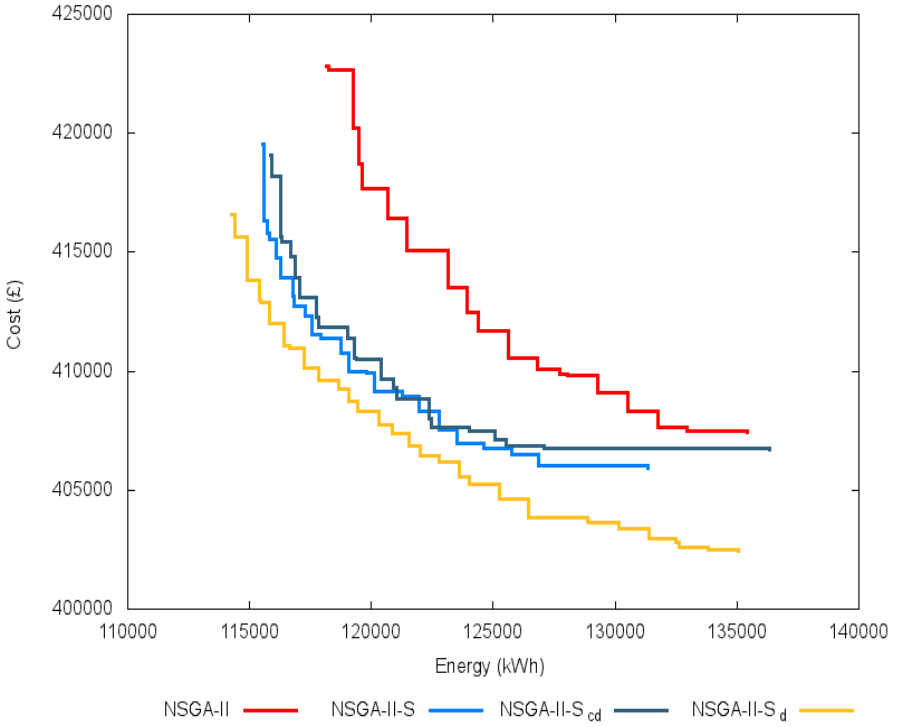

Brownlee, A. E. and Wright, J. A. (2015) *Constrained, mixed-integer and multi-objective optimisation of building designs by NSGA-II with fitness approximation.* Applied Soft Computing, vol.33, pp 114-126

Example 2

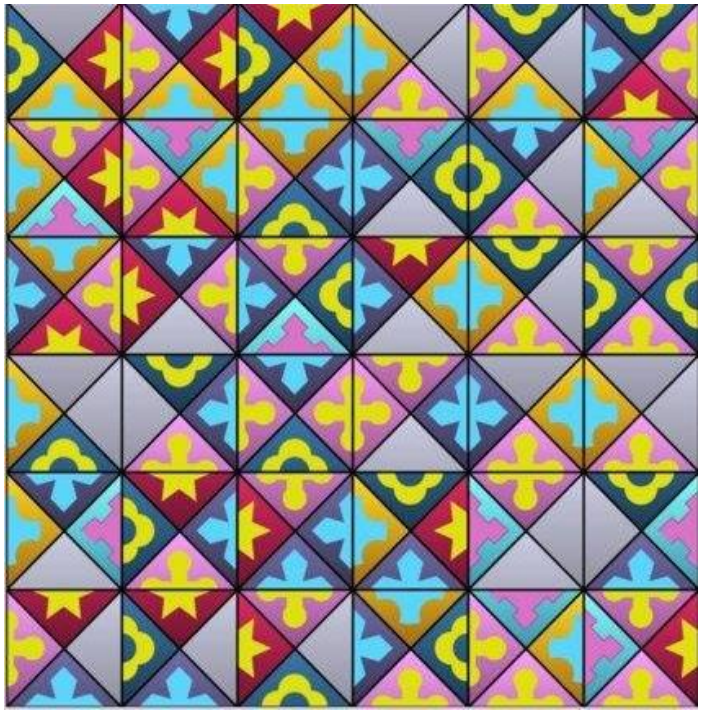

- Eternity II puzzle
- Objective: maximise matched adjacent edges
- *Surrogate* objectives:
  - Completed 2x2 squares
  - Completed 3x3 squares
  - Completed 4x4 squares
  - Tiles with all 4 edges matched
- Search iterates over two stages: surrogate, then objective

Tony Wauters, Wim Vancroonenburg and Greet Vanden Berghe. *A Guide-and-Observe Hyper-Heuristic Approach to the Eternity II Puzzle.* J Math Model Algor (2012) 11:217–233

Example 3

- MFM-GA uses undirected PGM (Markov network) to approximate fitness
- PGM initialised with dependencies between 5-bit blocks in problem, coefficients estimated using randomly generated population
- Fewer evals wrt GA, but more overhead

A. E. I. Brownlee, O. Regnier-Coudert, J. A. W. McCall, and S. Massie. *Using a Markov network as a surrogate fitness function in a genetic algorithm.* IEEE CEC 2010. pp 4525-4532, Barcelona, Spain.

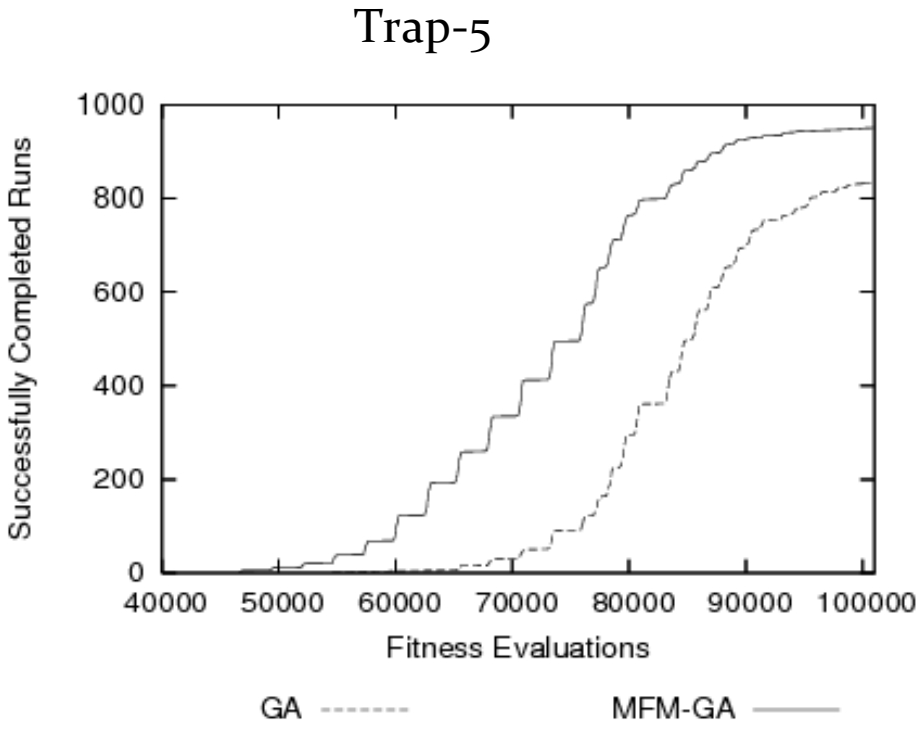

	GA	MFM-GA
Evals	83421	70839 (surr.) 14520 (true)
Run time	1.44 s	24.6 s

Mining a surrogate model

- Examine the surrogate model to gain insight into the problem
- Model here shows where glass is preferred (blue) on the façade

0.017	0.017	0.016	0.016	0.015	0.015	0.015	0.015	0.016	0.015	0.015	0.015	0.016	0.016	0.016
0.016	0.016	0.016	0.016	0.016	0.016	0.016	0.016	0.015	0.015	0.016	0.015	0.016	0.016	0.016
0.017	0.017	0.016	0.016	0.015	0.016	0.016	0.016	0.016	0.016	0.015	0.016	0.016	0.016	0.016
0.017	0.016	0.017	0.017	0.016	0.016	0.015	0.015	0.016	0.015	0.016	0.017	0.016	0.017	0.017
0.017	0.017	0.017	0.016	0.016	0.016	0.016	0.016	0.015	0.016	0.016	0.016	0.017	0.017	0.017
0.017	0.018	0.018	0.017	0.018	0.017	0.017	0.016	0.016	0.017	0.016	0.017	0.016	0.017	0.016
0.018	0.017	0.018	0.018	0.018	0.017	0.017	0.017	0.018	0.017	0.017	0.017	0.016	0.017	0.017
0.017	0.018	0.017	0.018	0.018	0.018	0.018	0.019	0.018	0.019	0.017	0.018	0.018	0.017	0.018

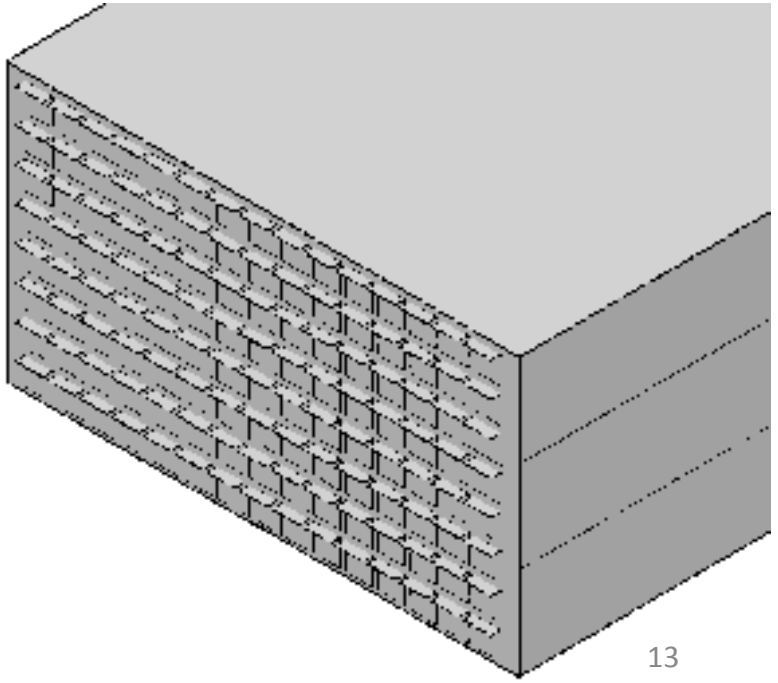

Summary

- "true" fitness not always suited to guiding search
- Use surrogates to improve search efficiency
- Static surrogates often used already!
- More reading…
  - Yaochu Jin (2005). *A comprehensive survey of fitness approximation in evolutionary computation.* Soft Computing, 9(1):3-12.
  - Yaochu Jin (2011). *Surrogate-assisted evolutionary computation: Recent advances and future challenges.* Swarm & Evolutionary Computation 1(2):61-70.

www.cs.stir.ac.uk/~sbr sbr@cs.stir.ac.uk

Class diagram