An Adaptive Cost Function for Multi-Fidelity Optimisation

Mark Pellowe, NCCAT, Loughborough University

Supervised by Prof. Gary Page and Prof. Adrian Spencer

Loughborough University

Rolls-Royce

Introduction to Multi-Fidelity Optimisation

- Uses a hierarchy of data sources
- Each source has a different “cost”

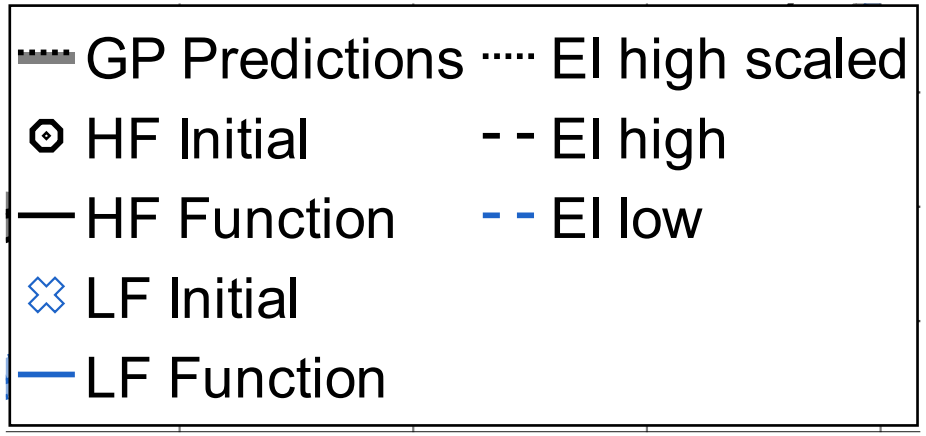

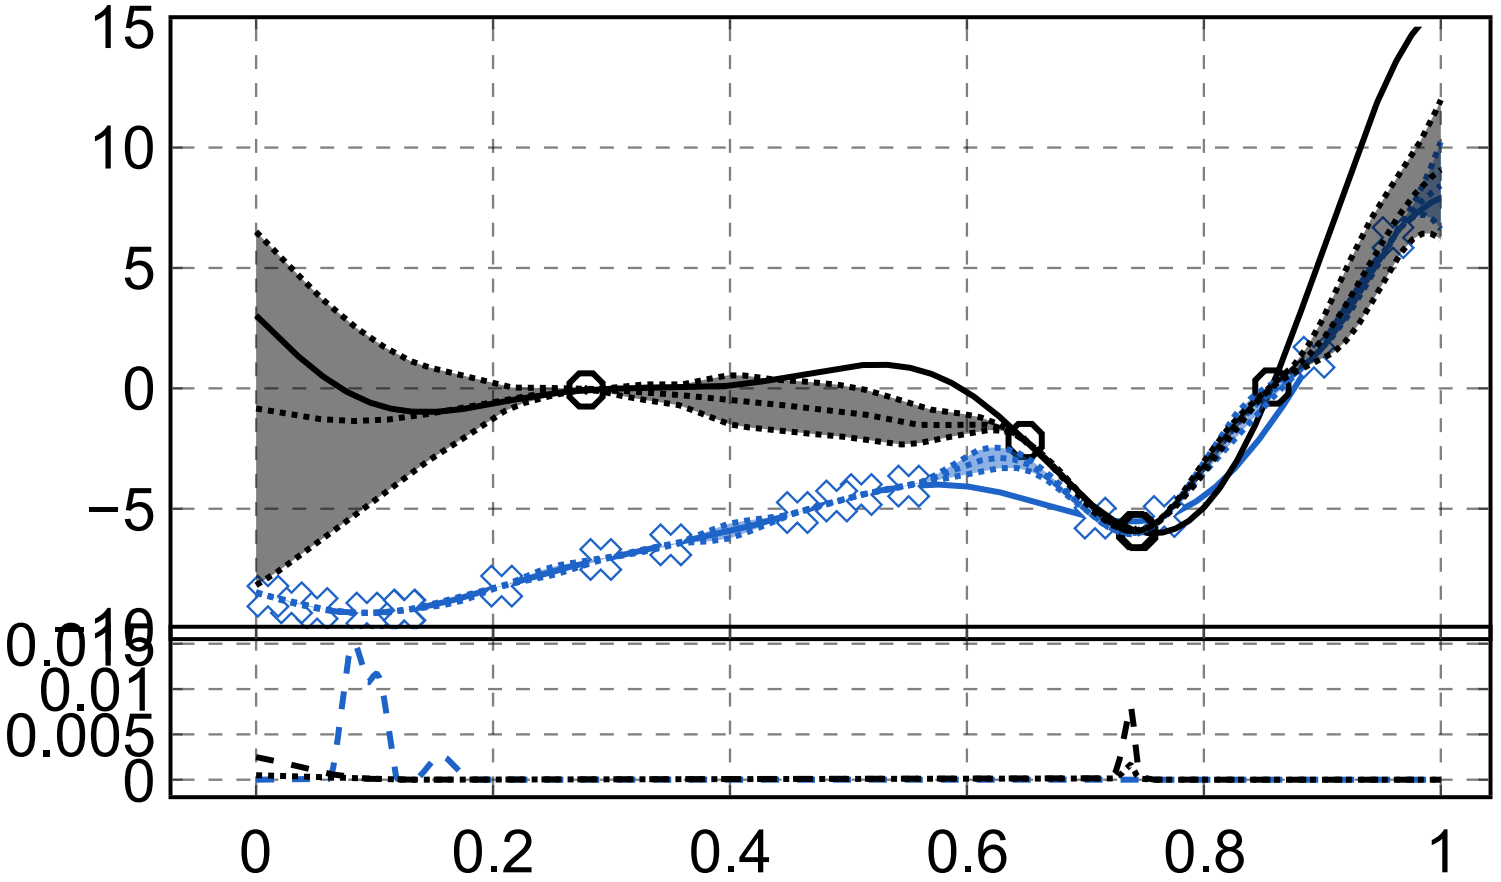

What is Fidelity?

- A hierarchy of different representations of a system
- Quality depends on several factors
  - Spatial, temporal, different physical assumptions
- Generally, higher fidelity = higher cost

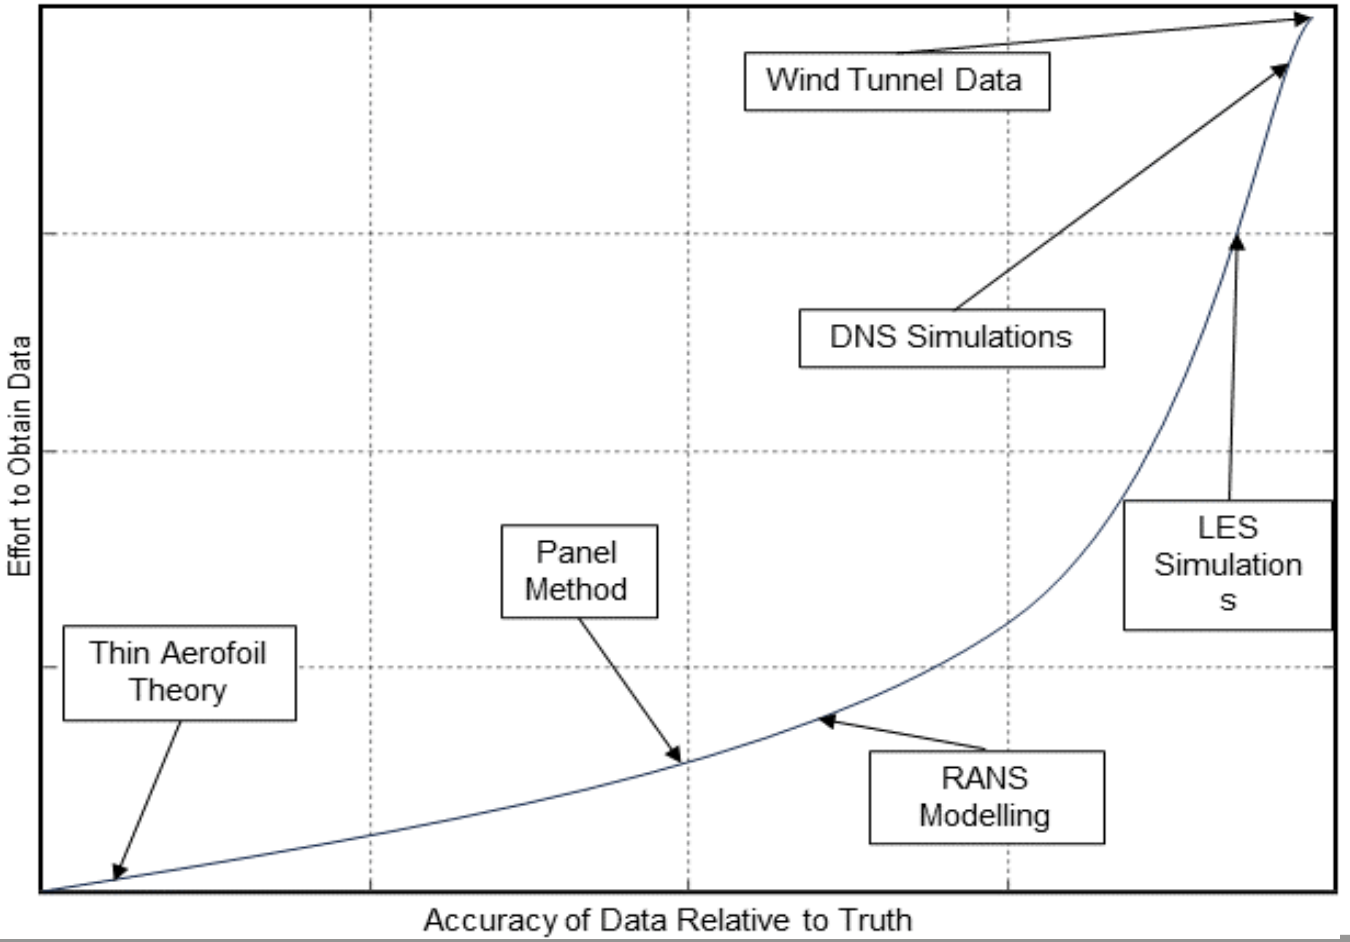

DEFINING COST AND CHEAPER ROUTES

Adapting Cost Based on States

- Cheaper dimensions exist for the next observation based on current location

- Next sample is determined by “utility” or “relative value”

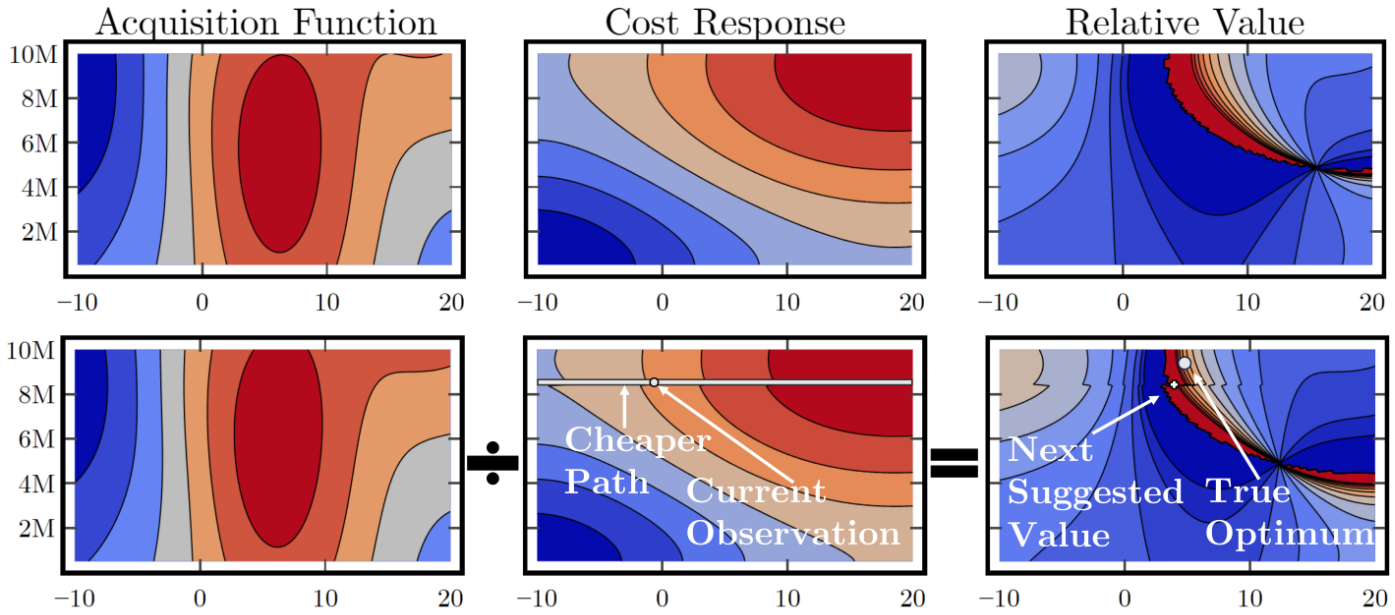

Cheaper to change flap angle than the width of the wings

Cost function

- Current cost function:
  - Defined as an integer per fidelity or a function
  - No consideration of already observed data

$$low\ fidelity = 1 \quad \text{OR} \quad high\ fidelity = 5$$

$$low\ fidelity = f_l(\theta_1, \dots, \theta_n) \quad high\ fidelity = f_h(\theta_1, \dots, \theta_n)$$

- New cost function:
  - Defined as the level of change from a previous observation
  - Weighted to show some parameters are cheaper to change

$$cost_{intrinsic} = function(\theta_1, \dots, \theta_n) > 0$$

$$cost_{change} = \sum_{i=1}^{n} \lambda_i \delta(\theta_n - \theta_{nprev})$$

$$cost_{total} = cost_{intrinsic}(\theta_1, \dots, \theta_n) + cost_{change}(\theta_1, \dots, \theta_n, \theta_{1prev}, \dots, \theta_{nprev})$$

MODEL PERFORMANCE ON TEST FUNCTIONS

Models used in Comparison

- SF-OC: Single fidelity with a non-dynamic cost model
- SF-AC: Single fidelity with novel cost model
- MF-AC: Multi fidelity with novel cost model
- MF-OC: Multi fidelity with a non-dynamic cost model

Six-Hump Camel Back

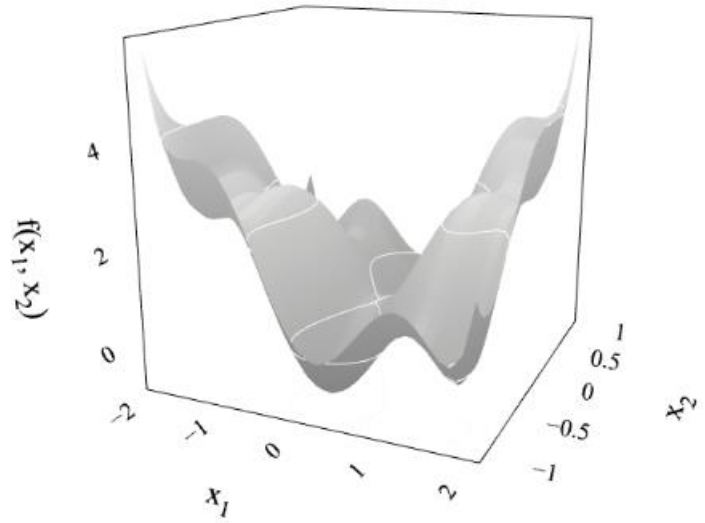

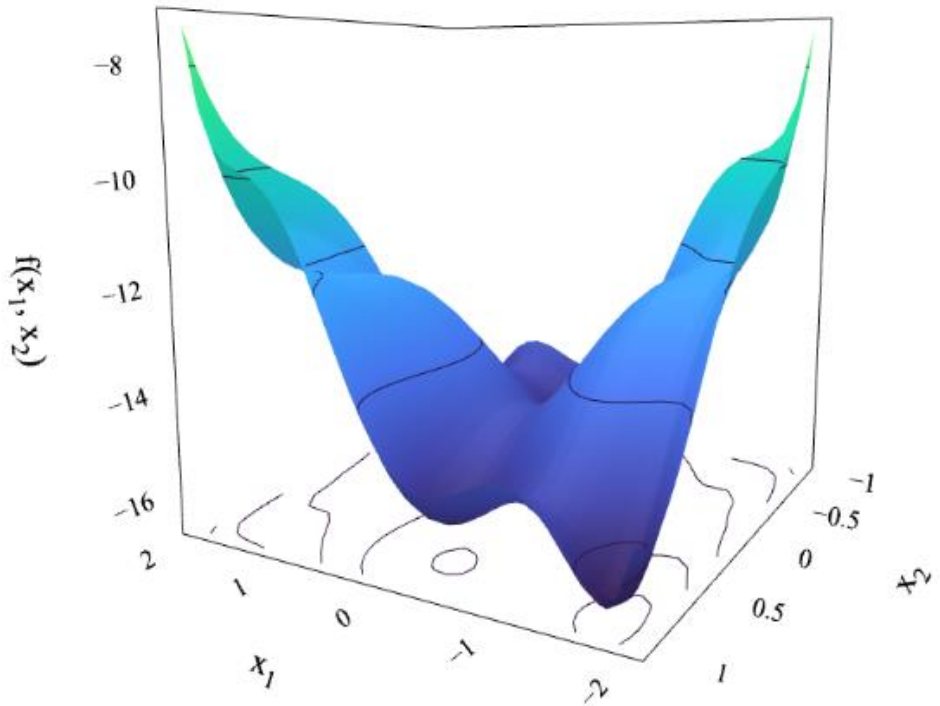

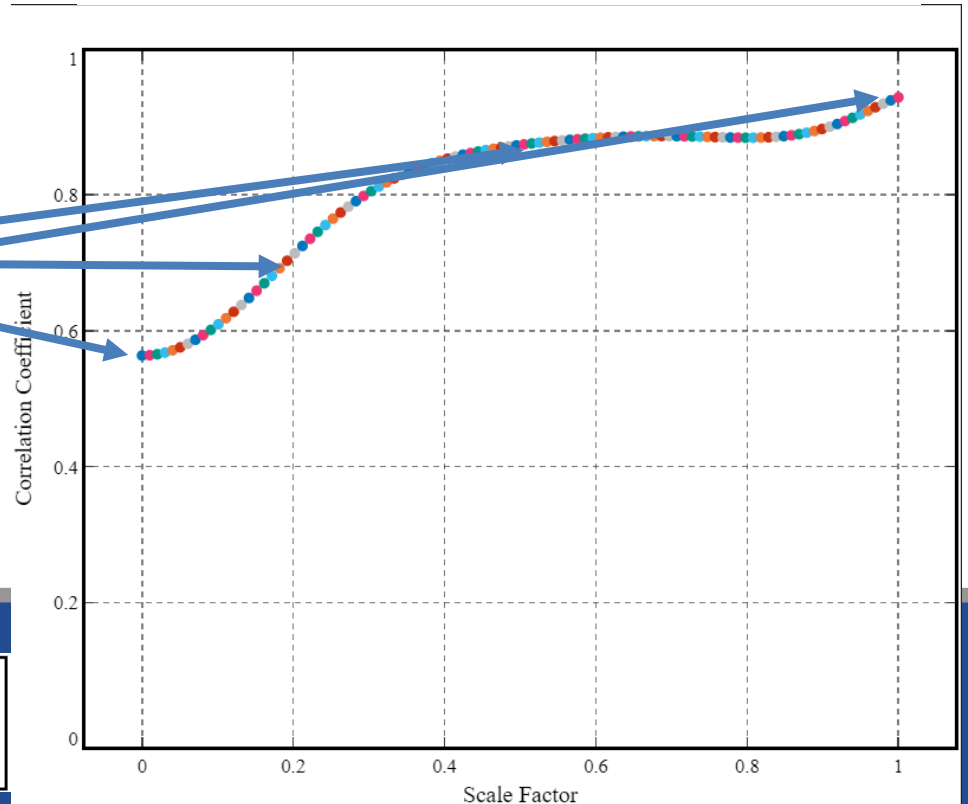

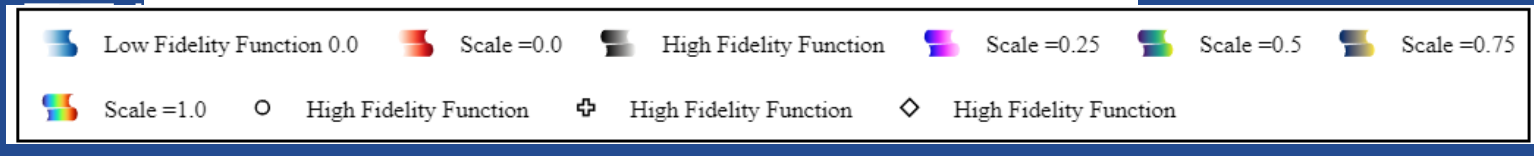

What a Single Search Looks Like

- Multi-Fidelity searches always begin more expensive
- Adaptive costs allow for faster convergence

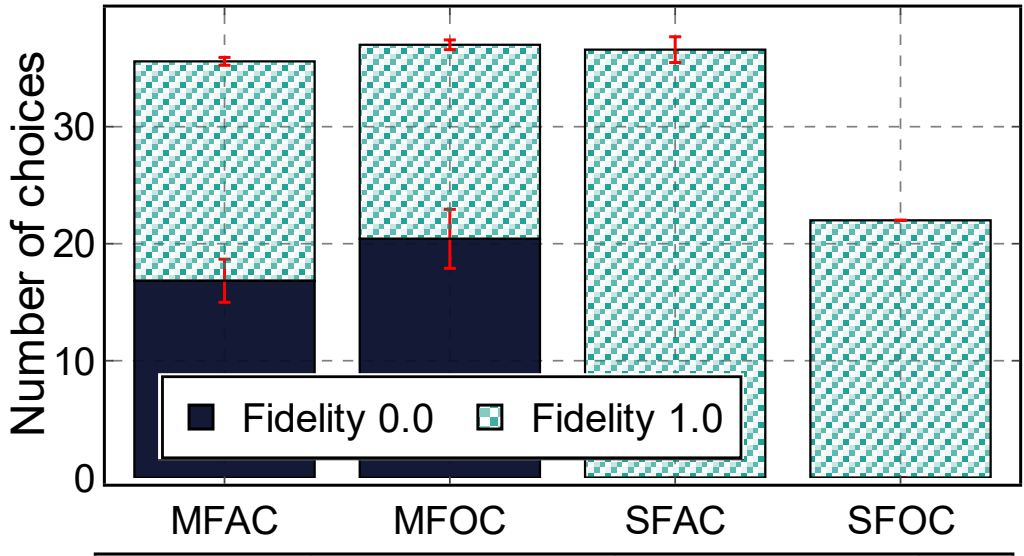

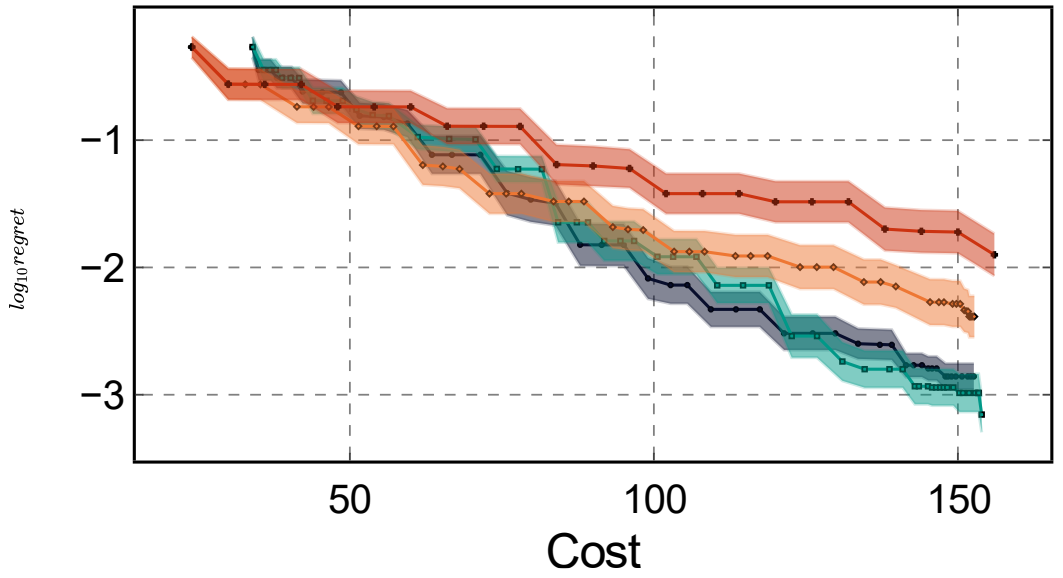

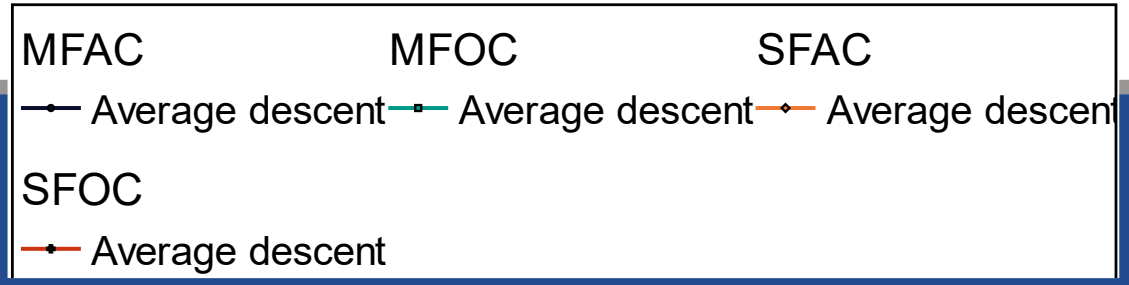

Optimum Searches with Scale 0

- Correlation = 0.56
- Higher ratios mean less search space within budget
- Also mean less utility “value” for MF optimisers

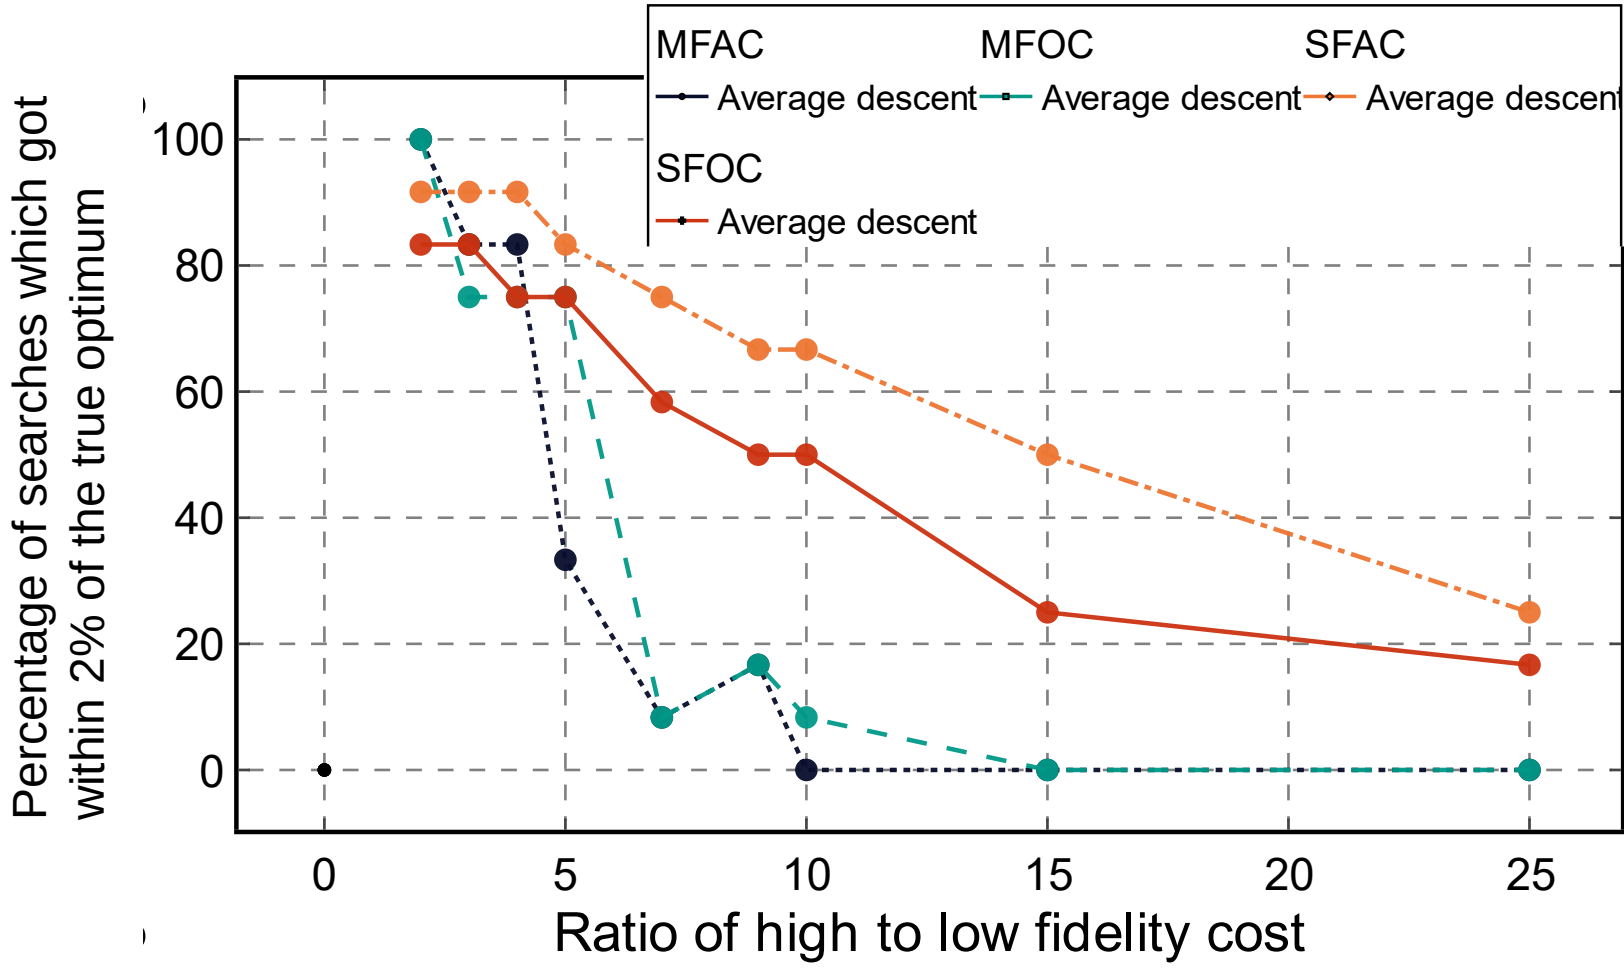

Optimum Searches with Scale 1.0

- Correlation = 0.94
- Higher degree of correlation = Better MF models

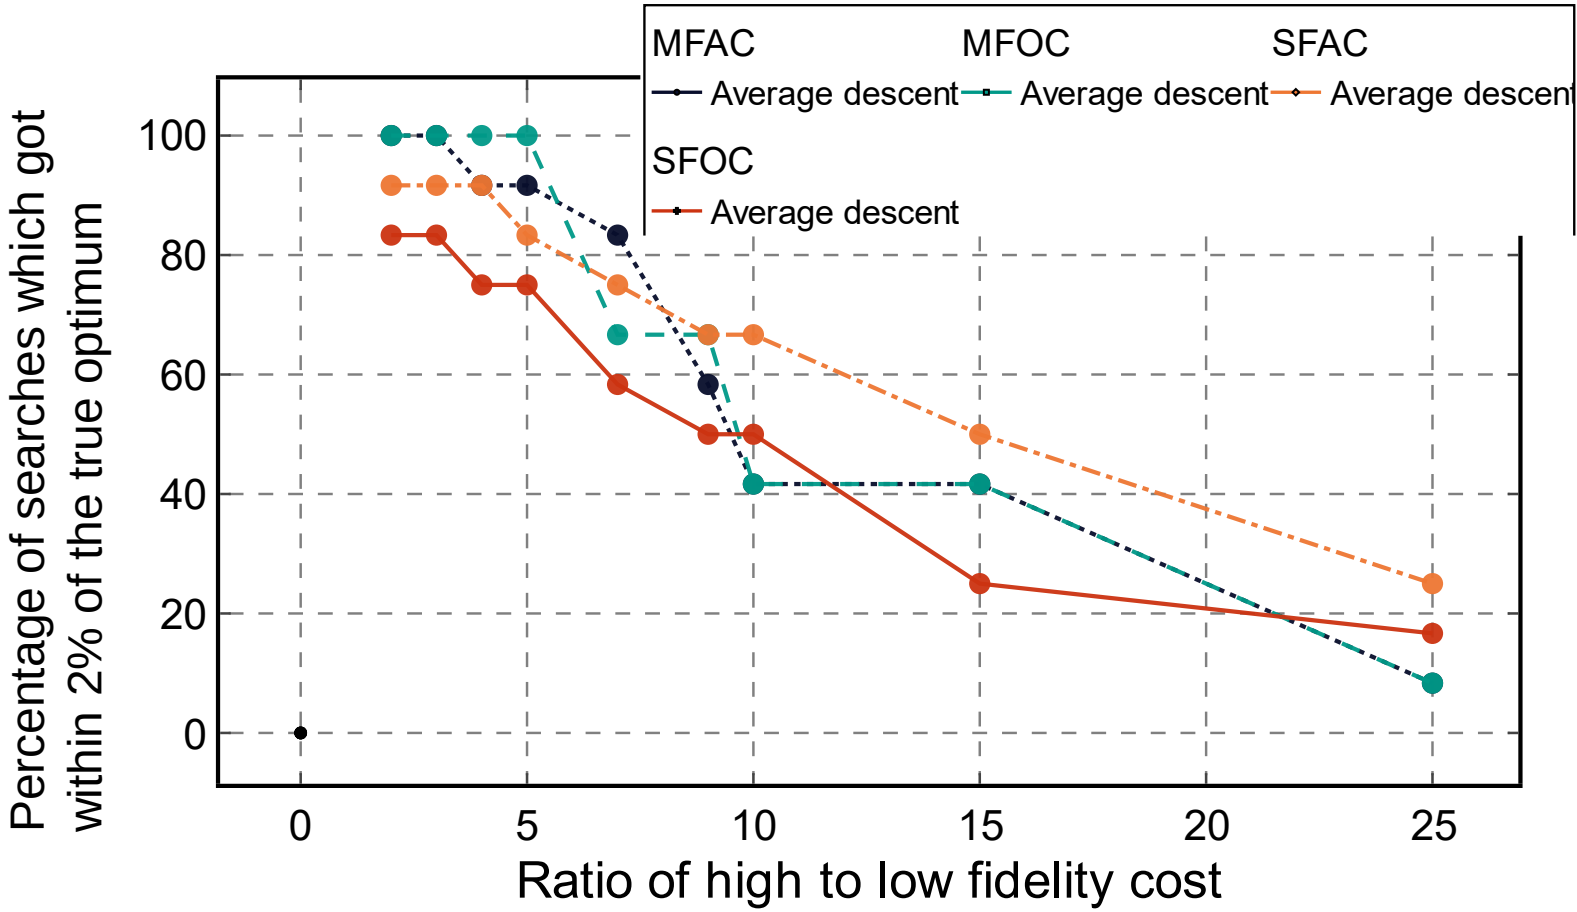

Optimising Conditions for Aerofoils

TESTS ON REAL DATA

Aerofoils

- Wind tunnel
  - Experimental measurements
  - The “most true” but expensive
- SU2
  - Steady flow CFD solver
  - Cheaper but heavily biased by simulation parameters
- Xfoil
  - Potential flow + boundary layer solver
  - Very cheap and reliable in parts
  - Still biased

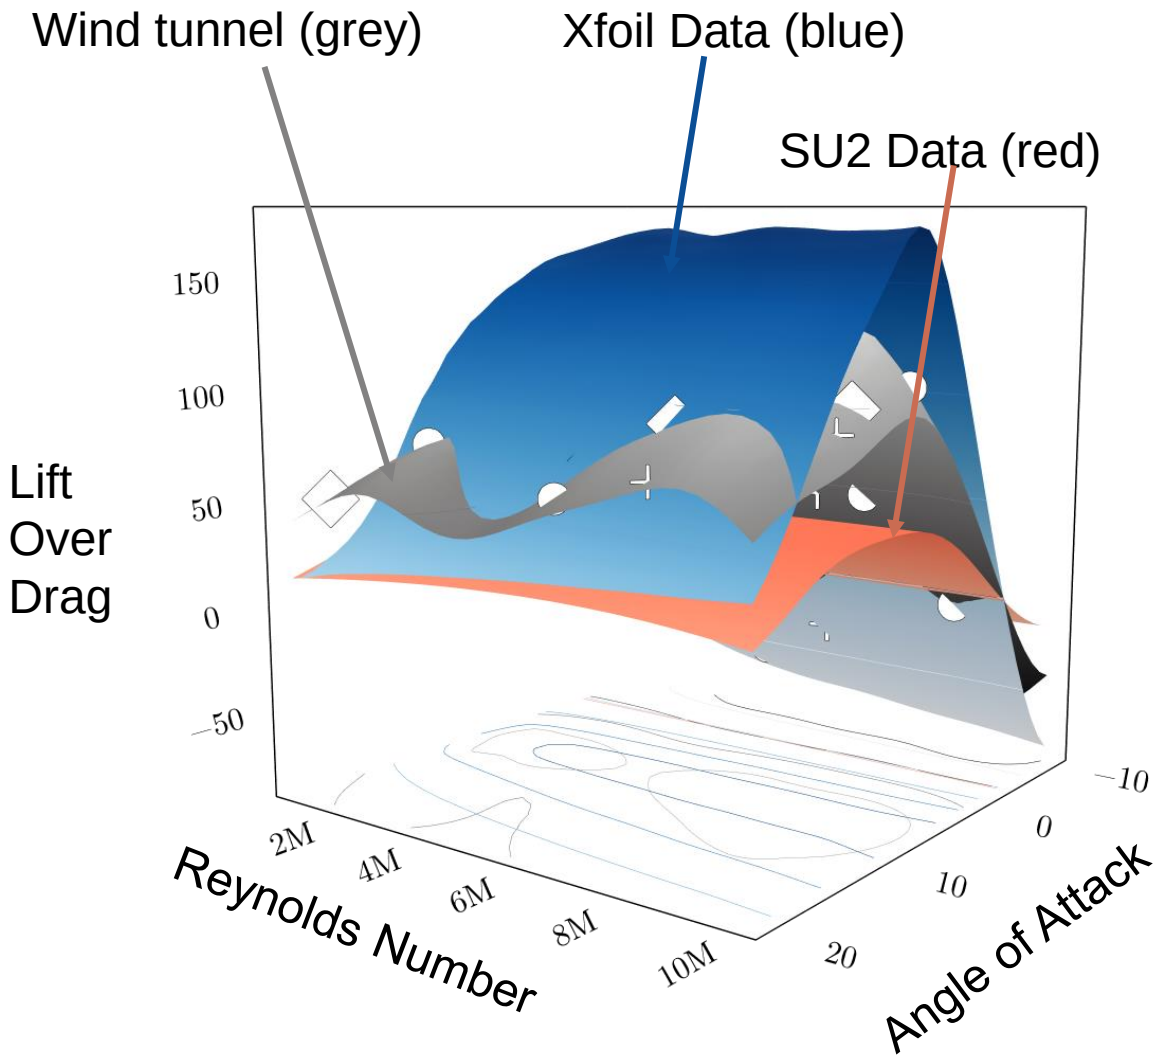

Aerofoils

- Multi-fidelity approaches are superior to single fidelity ones
- The scaling factors have large effects on single fidelity.
  - Smaller effects on multi-fidelity approaches

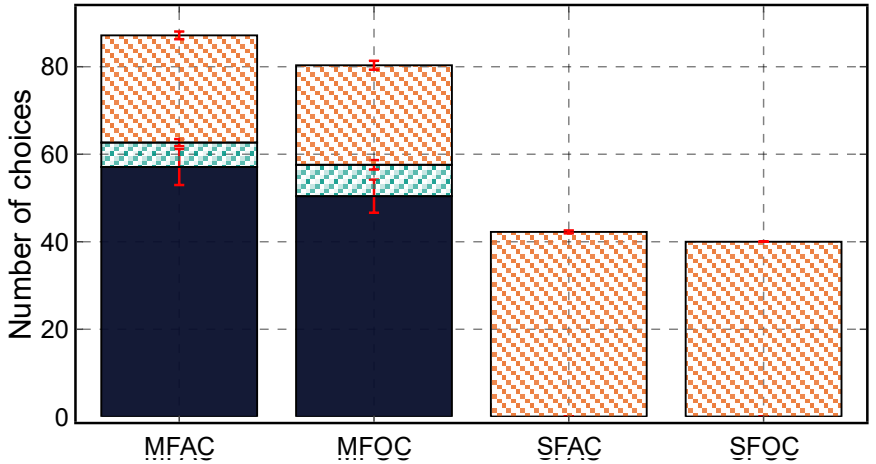

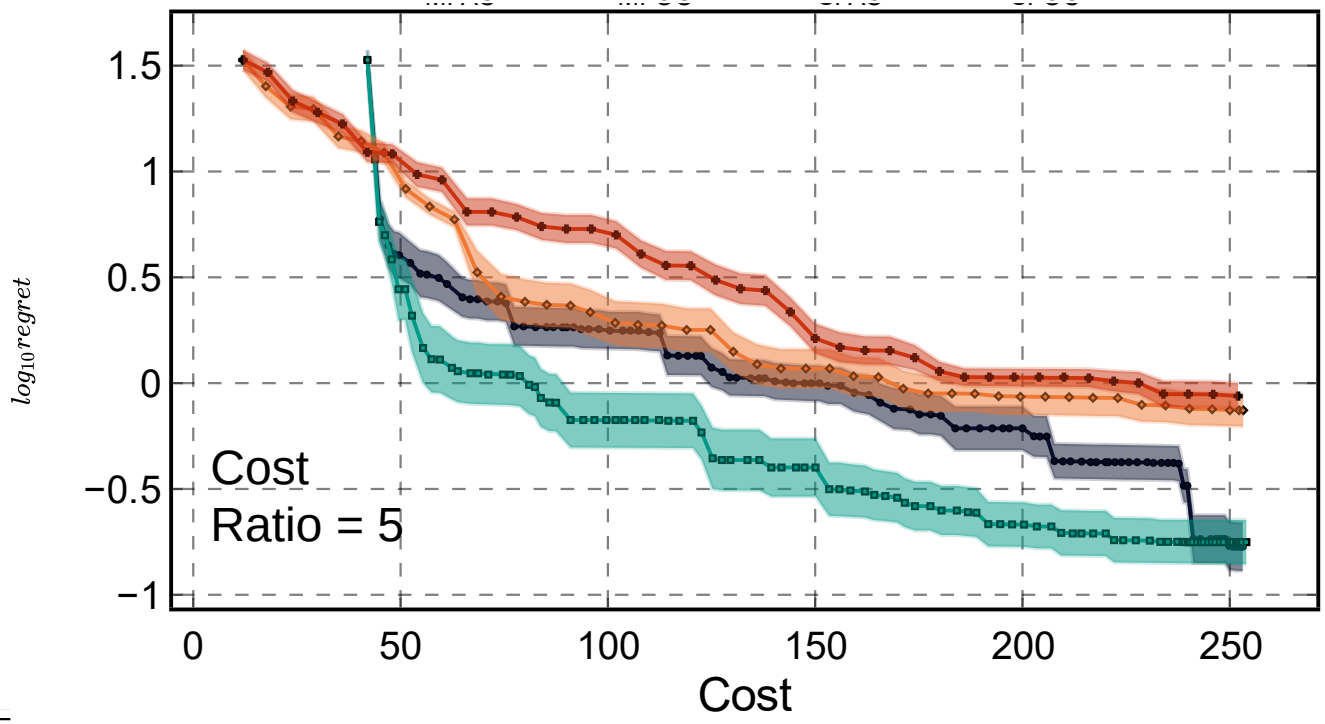

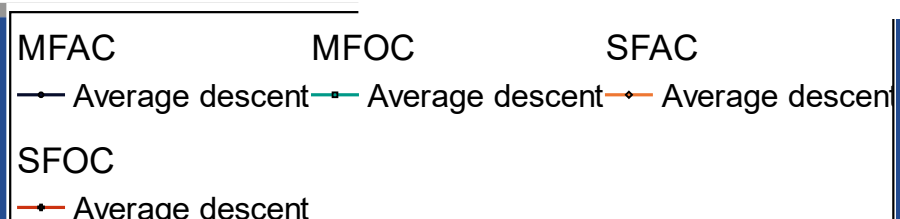

Conclusions

- Multi-fidelity optimisation outperforms single fidelity within the same budget
  - Assuming lower fidelity has a suitable degree of correlation with the “truth”

- Adaptive cost treatment does not harm optimisation
  - A cheaper exploitative dimension can be efficient and useful, even if the solution does not exist down the path

THANK YOU!

Loughborough University

Rolls-Royce®