Answering Hamming



Joshua Knowles - FOGA XVIII, 27 August 2025

Abstract

The story goes that while working at Bell Labs in the 1950s, the mathematician and computer scientist Richard Hamming would ask colleagues, “What’s the most important problem in your field?” ... and then follow up with, “So, why aren’t you working on it?” Both questions have many possible answers, even for just one person at one time, but they are certainly provocative, tough and uncomfortable. In the talk, I will reflect on my personal answers at various times, some answers for evolutionary computation (EC) and evolutionary multiobjective optimization (EMO) more broadly, as well as for industrial research & innovation. My particular answers (or anyone’s) are almost certainly not as important as the struggle behind them to grapple with the questions.

CCS Concepts

- Mathematics of computing → Evolutionary algorithms;
- Applied computing → Multi-criterion optimization and decision-making.

Keywords Evolutionary computation, multiobjective optimization, industrial research.

Author Bio

Based in the UK, Joshua Knowles works as a scientific advisor for the multinational energy technology company SLB, holds an honorary professorship in the decision sciences group of Alliance Manchester Business School at The University of Manchester, and is a former professor of natural computation at the University of Birmingham (now an hon. senior fellow). Publishing in evolutionary multiobjective optimization (EMO) since the late 90s, his work includes fundamental research on archiving with diversity, local search, performance assessment, hypervolume-as-selection, machine decision makers, heterogeneous objectives, and “multiobjectivization”. In 2004-5, he developed the influential multiobjective Bayesian optimization method, ParEGO, for expensive problems. More broadly, Josh is interested in and has published (joint work) on the evolution of evolvability, the evolution of cooperation, neutral evolution, and symbiogenesis (including Deep Optimization). He has also described evolutionary and ML applications work in premier journals in astrophysics, analytical chemistry, theoretical biology, bioinformatics, and operations research.

slb

Acknowledgments

David Corne (Heriot-Watt)
Marco Dorigo (ULB)
Richard Allmendinger, Julia Handl, Manuel López-Ibáñez (Manchester)
Richard A. Watson (Southampton)
Douglas B. Kell, CBE (Liverpool)
Ruth Misener (Imperial)
Michael J. Williams, Maria Fernanda Vargas Izquierdo, Mahdi Davoodi (SLB)
… and numerous other colleagues, collaborators, students, & interns.

MANCHESTER 1824
The University of Manchester

UNIVERSITY OF BIRMINGHAM

slb

EMO 2027

Evolutionary Multi-criterion Optimization

Exeter, UK. 5^{th}-8^{th} April 2027

- 5^{th} April: Tutorials
- 6^{th}-8^{th} April: Single track scientific presentations (inc. MCDM track, Industry track and a poster session)
- *Outstanding EMO Paper and Outstanding EMO Student Paper awards*
- *Papers published in Springer LNCS series*
- *Conference to receive an iCORE ranking early 2026*
- **Daily direct flights from Amsterdam Schiphol to Exeter with KLM**
- Conference will be held at the University of Exeter Business School

Venue

Banquet venue

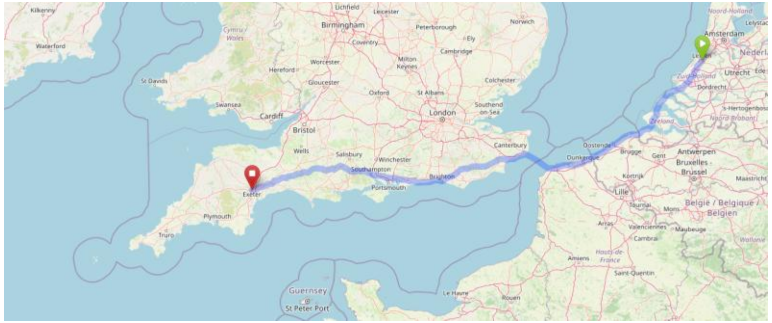

Purpose

→ To give you a nudge to reflect on *the struggle*, and on the importance of problems, in your continuing to do great work.

General Disclaimer

The views expressed in this talk are mine and do not necessarily represent those of SLB.

Three Kinds (or Levels) of Problem

1. **Optimization problems –** travelling salesperson; knapsack; Rastrigin function; sphere; black-box optimization; convex-continuous; non-linear constrained; ...
2. **Applied problems** in industry, science, society,...
3. **Meta-problems* –** generalizations of problems (1. & 2.), theory, tools, measurements,...

   Examples of meta-problems
   - Does the building block hypothesis help to design "competent" EAs?
   - What is the value of recombination in EAs?
   - What information about a problem is useful in selecting/tuning an algorithm?
   - Are large language models a threat to evolutionary computation?

   ...And there are many more.

*We might also call these simply "scientific problems".

Many important scientific problems here, but *how important*?

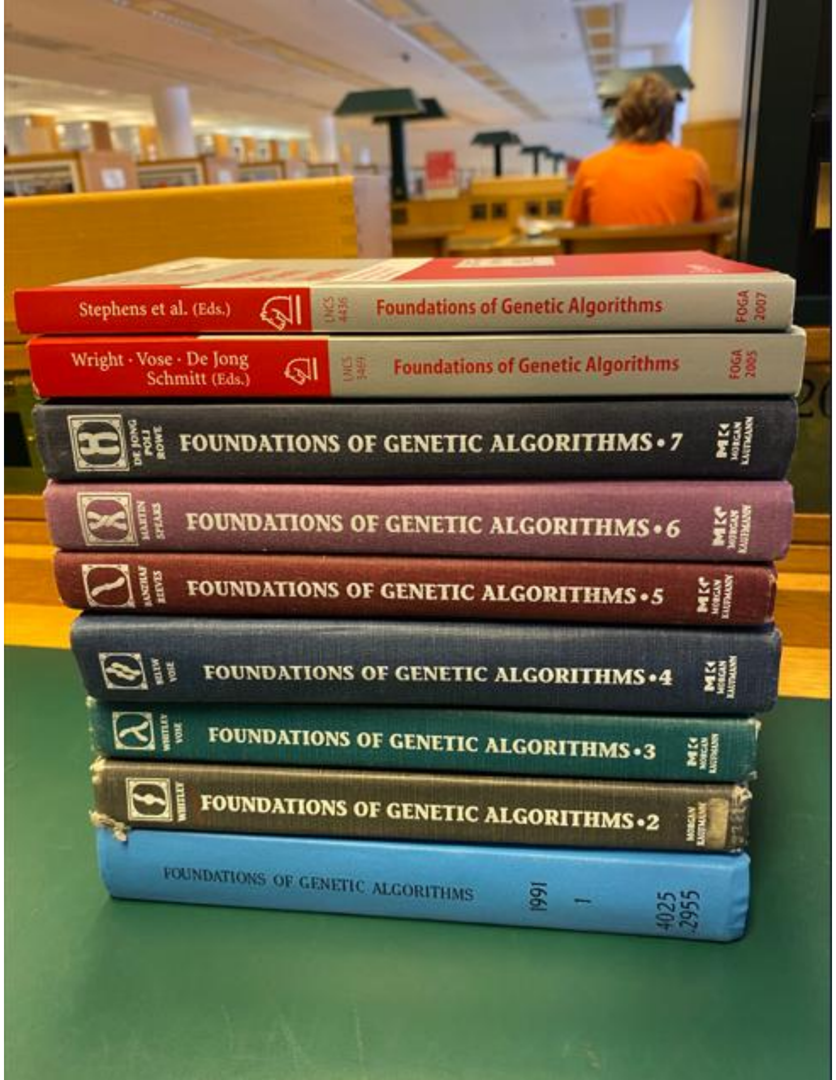

→ See refs [4-12].

Hamming code | Hamming window | Hamming numbers | Hamming distance | Hamming weight

Hamming, Richard W. (1915-'98)

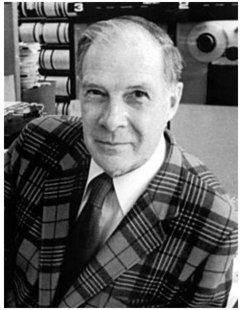

image source: http://cm.bell-labs.com/cm/cs/alumni/hamming/

CV

'44-'45	Manhattan Project
'46-'76	Bell Labs
'76-	Naval Postgrad School (teaching)
'97	*The Art of Doing Science & Engineering* | *Learning to Learn*

"You ought to try to make significant contributions to humanity [...] It has often been observed the true gain is in the struggle and not in the achievement".

"If you do not work on important problems, how can you expect to do important work? Yet direct observation and direct questioning of people show most scientists spend most of their time working on things they believe are not important and are not likely to lead to important things."

See Ref [1]

Hamming and me (some similarities)

On this slide, I do something unwise: compare myself to others. But it's just to underline the relevance of Hamming to me (or you) today

Hamming	**Knowles (me)**
computer scientist	computer scientist
industrial scientist (Bell Labs)	industrial scientist (SLB, …)
teacher (Naval college)	teacher (UK universities)
cross-disciplinary interactions	cross-disciplinary interactions
rapidly-changing contexts (electronic & digital revolutions)	rapidly-changing contexts (AI, ML, automation)
interested in how problems are identified & clarified	interested in how problems are identified & clarified
interested in the *how* and *why* of doing research – *style, art*	interested in the *how* and *why* of doing research – *style, art*
interested in ethics – how to serve the public good.	interested in ethics – how to serve the public good.

→ Differences are rather obvious!

What's an important problem?

→ We see it when we look back, not so easy looking ahead (fortune telling may be required!)

→ Is timely and "in the air" but perhaps invisible

→ Or is in need of clarification

→ Provides building blocks for much further productive work (generality)

→ Is attackable, addressable now.

Hamming advice: don't forget why you thought something was *not* attackable! (FFT anecdote).

Engineering vs science/research – some modern views!

Yann LeCun · Following
VP & Chief AI Scientist at Meta
2w ·

There is a difference between research and engineering in (1) modus operandi, (2) methodology, (3) openness, (4) evaluation criteria.

Research uses the methodology of science to discover new principles, demonstrate that they can work in practice, analyze their advantages and limitations, and interact with the wider research community to criticize, validate, reproduce, compare, and improve. The criteria are conceptual simplicity, theoretical beauty/explainability, clear performance advantage over prior art on some accepted metrics. This is true for research in academia as well as in industry.

Engineering integrates methods, often developed in a research mode, to build working systems. The philosophy is to go with the first set of methods that work well enough for the task. It generally involves a lot of tinkering, tweaking, fine-tuning, and an occasional kludge to get the performance up on a real task. Whether the method is the absolute best matters less than whether it is good enough for the tasks at hand.

Researchers are evaluated largely on intellectual impact. Research evaluation is a difficult task because the product impact may occur years (sometimes decades) after the work. For that reason, evaluation must often rely on the collective opinion of the research community through proxies such as publications, citations, invited talks, awards, etc. That's one reason research must be published.

Engineers are evaluated largely on product impact, sometimes through proxy metrics such as pull requests, lines of code, etc.

By operating in engineering mode, researchers are incentivize to do incremental work. If you make no distinction between the two activities, if you don't evaluate researchers and engineers with different criteria, you run the risk of killing breakthrough innovation. True breakthroughs require teams with a long horizon and minimal constraints from product development and management.

The industry research labs of yore that have left an indelible mark on scientific and technological progress (Bell Labs Area 11, IBM Research, Xerox PARC, etc) were all research divisions that were clearly separate from engineering divisions.

Elon Musk @elonmusk
Subscribe

This false nomenclature of "researcher" and "engineer", which is a thinly-masked way of describing a two-tier engineering system, is being deleted from @xAI today.

There are only engineers.

Researcher is a relic term from academia.

→ Hamming's view: they are different but often feed each other. (More in-line with LeCun).

Important applied problems?

Research and engineering at SLB (including my present work)

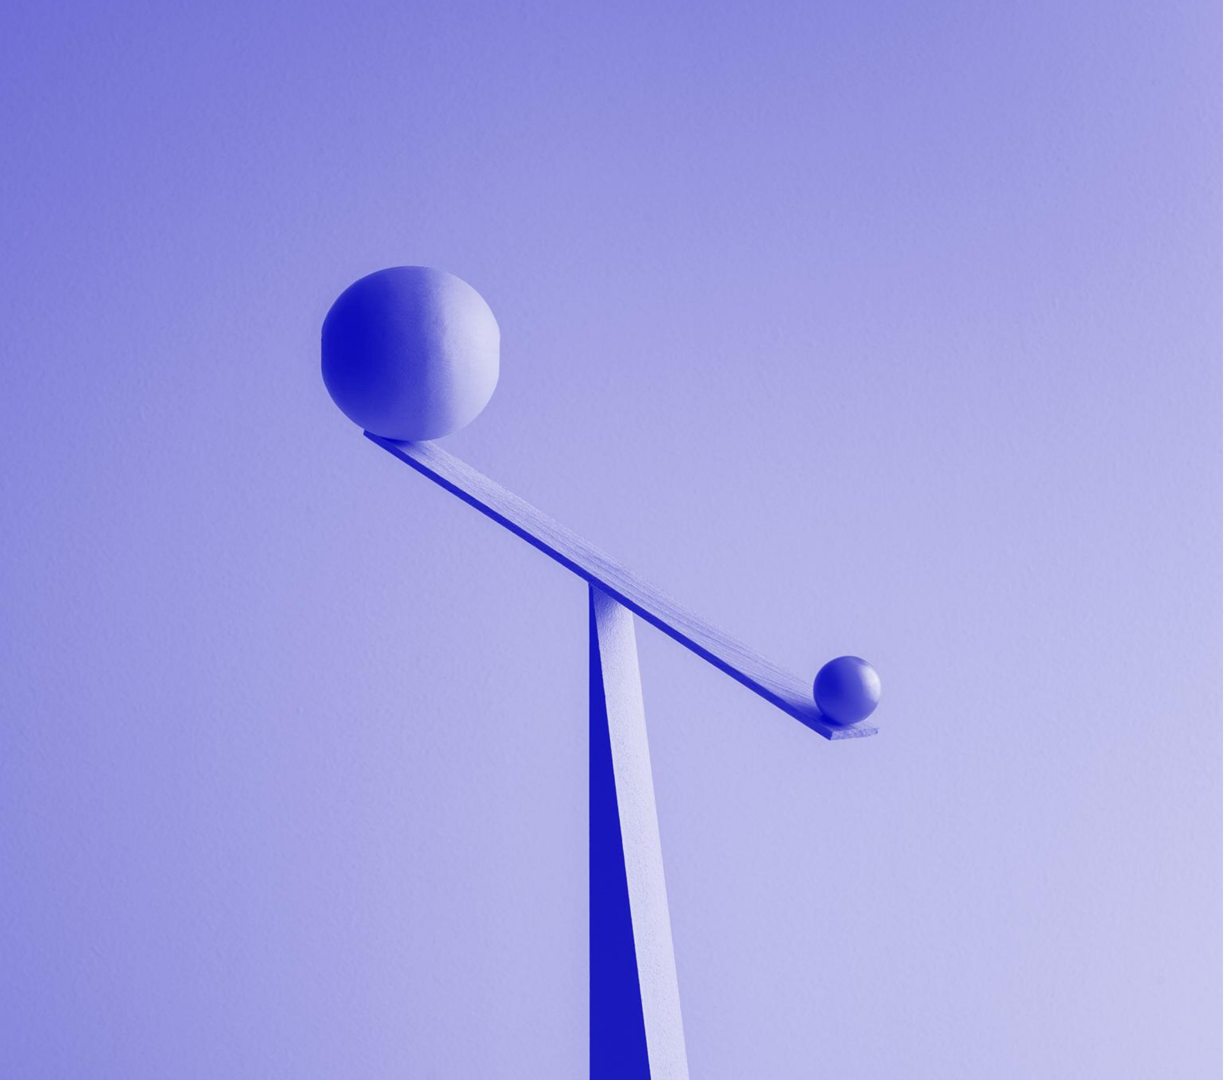

For a balanced planet

"... SLB is committed to helping deliver the world's greatest balancing act: enabling secure, accessible, sustainable energy to meet growing demand"

(SLB Mission Statement)

SLB Innovations (very selected)

On Concorde

In oil & gas (See refs. [2] and [3])

Subsurface mapping (1912; Schlumberger brothers)

Electrical well logging (1927; Doll, Henri-Georges)

Tubing conveyed perforating (1952; with Humble Oil)

Description logics / semantic networks (1983; Brachman et al.)

Marine surveying positioning control (1996; Bittlestone)

Operations planning & scheduling (2015; Long & Fox).

"Retina" at-bit imaging system (2024; Cook et al.).

In "new" energy

Lithium - sustainable production plant in NV, USA (2024; with partners)

Carbon capture on a cement plant (2024; SLB Capturi)

Simulation of non-Newtonian fluids & elastic turbulence effects (2020; Davoodi, Davoodi & Clarke) – active in various applications.

Character of problems at SLB

→ Global impact: energy technologies with high levels of

- Availability
- Cost-efficiency
- Sustainability
- Equity

to meet growing energy demand, world-wide.

→ Complexity

- Multiple disciplines involved, distributed teams, diverse stakeholders, complex business landscape, uncertain future, very high technologies developed over decades.

→ Yet, simple solutions are possible to discover (and are often best!)

MOO problems everywhere!

	Problem	#Objectives	#Constraints
Advice	Fluids advisor	5 or 6	several
	Drill bit advisor	3	several
	Directional drilling advisor	>20 currently	>10 currently
	Optimal operating window - drilling	~12	many
Design	Drill bit design	6 or 7	>10
	Cutter design	3 or 4	many
	Heat exchangers - data centers	3	several
	Hole-cleaning drill pipes	3	several
Strategy	Oilfield scheduling	2 or 3	100s or 1000s
	Site-location (e.g., nuclear waste)	12	18
	Technology portfolio planning	2 or 3	several

The broader perspective

Important problems in EC, EMO, and my work (past, present and future)

Three important open problems in evolutionary computation

- → What is the best algorithm for my problem?* Is it time for more graybox and "tailoring"?

- → Is natural selection necessary and sufficient for learning/adaptation/optimization? (cf. Watson); what about symbiogenesis, what about other elements from the extended evolutionary synthesis, e.g., niche construction?

- → How can EC *principles* contribute to answering scientific questions in other disciplines? – Machine learning, biology, complexity, neuroscience, physics?

*Entails defining *problem* and *algorithm*, solution concepts, and performance measures. But it's a central question.

Three important problems from my past work (in EMO)

→ What is the equivalent of the random mutation hillclimber for MOO? (PAES). Led to related work on nondominated solution archives (with diversity) and hypervolume maximization.

→ How should we assess the performance of EMOAs, including those with preferences?

→ Is searching with more objective dimensions easier or harder? When/why? ("Multiobjectivization" supported by theory)

→ Several more examples are given in the bibliography

Three important open problems in evolutionary multiobjective optimization (EMO)

- → Demonstrate with a killer app or otherwise why MOO is better than SOO…including solution selection from the Pareto front*
- → Alignment of MOO solution concepts with the origin/meaning of the objectives (cf. "modes of problem solving", Handl & Knowles)
- → Develop effective cooperative solution methods for MOO problems exploiting single-objective EAs & single-objective "classical" methods. (See next slide for current division of SOO and MOO-capable software).

* This seems central but is very hard to do quantitatively.

Software

Table: Optimization software libraries and details of their licensing agreements plus notes on functionality offered for multiobjective problems.

Recommended:
- pyMultiobjective
- pymoo

Name	Language(s)	Licence	Location / Notes
PlatEMO	MATLAB	Matlab Licence	200+ open-source evolutionary algorithms 400+ open-source benchmark problems PlatEMO - File Exchange - MATLAB Central (mathworks.com)
pymoo	Python	Apache v2.0	pymoo: Multi-objective Optimization in Python
pyMultiobjective	Python	GNU PGL 3.0	pyMultiobjective · PyPI
jMetal	Java	GNU LPGL	jMetal Web site (sourceforge.net)
jMetalCpp	C++	GNU LPGL	jMetalCpp Web site (sourceforge.net)
Pareto.py	Python	GNU LPGL 3.0	GitHub - pareto.py
CPLEX	Concert / IPL interfacing with C, C#, Java & Python	Proprietary	Little or no support for multiobjective problems IBM ILOG CPLEX Optimization Studio
Gurobi	C back-end; APIs for most languages	Proprietary	Little or no support for multiobjective problems gurobi.com/
GAMS	GAMS with support for Python	Proprietary	Little or no support for multiobjective problems gams.com
ModeFrontier	Workflow-based IDE	Proprietary	Multiobjective and multi-disciplinary design package with extensive functionality engineering.esteco.com

Tenth important problem

→ Whatever you’re working on, presenting this week

→ Whatever you’re thinking about today, tomorrow, next year…

References

[1] Hamming, R. W. (1997). *The art of doing science and engineering: Learning to learn*. Gordon & Breach.

[2] Mau, M., & Edmundsen, H. (2015). *Groundbreakers: The story of oilfield technology and the people who made it happen.* FastPrint Publishing.

[3] Smith, R. G., Schoen, E. J., & Tenenbaum, J. M. (2022). Early AI applications at Schlumberger. *IEEE Annals of the History of Computing,* 44(1), 88-102. doi: 10.1109/MAHC.2022.3149469.

FOGA (early editions)

[4] Rawlins, G. J. (Ed.). (1991). *Foundations of Genetic Algorithms 1991 (FOGA 1)* (Vol. 1). Elsevier.

[5] Whitley, L. D. (Ed.). (1993). *Foundations of Genetic Algorithms 2*. Morgan Kaufmann.

[6] Vose, M. D., & Whitley, L. D. (Eds.). (1995). *Foundations of Genetic Algorithms 3*. Morgan Kaufmann.

[7] Belew, R.K., & Vose, M. D. (Eds.). (1997). *Foundations of Genetic Algorithms 4*. Morgan Kaufmann.

[8] Banzhaf, W., & Reeves, C. (Eds.). (1999). *Foundations of Genetic Algorithms 5*. Morgan Kaufmann.

[9] Worthy, M. & Spears W. (2001). *Foundations of Genetic Algorithms 2001 (FOGA 6)*. Elsevier.

[10] De Jong, K., Poli, R., & Rowe, J. (2003). *Foundations of Genetic Algorithms 7*. Morgan Kaufmann.

[11] Wright, A. H., Vose, M. D., De Jong, K. A., & Schmitt, L. M. (Eds.). (2005). *Foundations of Genetic Algorithms: 8th International Workshop, FOGA 2005*, Revised Selected Papers (Vol. 3469). Springer.

[12] Stephens, C. R., Toussaint, M., Whitley, D., & Stadler, P. F. (Eds.). (2007). *Foundations of Genetic Algorithms: 9th International Workshop, FOGA 2007*, Revised Selected Papers (Vol. 4436). Springer.

Bibliography (my work, selected)

TITLE	CITED BY	YEAR	**Problem**
Approximating the nondominated front using the Pareto archived evolution strategy JD Knowles, DW Corne Evolutionary computation 8 (2), 149-172	3190	2000	What's the simplest effective algorithm for Pareto multiobjective optimization? Can we define something like a simple hillclimber? Do current population-based EMO algorithms compare favourably?
The pareto archived evolution strategy: A new baseline algorithm for pareto multiobjective optimisation J Knowles, D Corne Proceedings of the 1999 congress on evolutionary computation-CEC99 (Cat. No ...	2054	1999	
PESA-II: Region-based selection in evolutionary multiobjective optimization DW Corne, NR Jerram, JD Knowles, MJ Oates Proceedings of the 3rd annual conference on genetic and evolutionary ...	1585	2001	Diversity maintenance in objective space.
ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems J Knowles IEEE transactions on evolutionary computation 10 (1), 50-66	1568	2006	Multiobjective generalization of EGO (a Bayesian optimization algorithm in modern parlance) for expensive functions.
On metrics for comparing nondominated sets J Knowles, D Corne Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat ...	834	2002	Desirable properties of performance measures in MOO.
Reducing local optima in single-objective problems by multi-objectivization JD Knowles, RA Watson, DW Corne International conference on evolutionary multi-criterion optimization, 269-283	517	2001	Are more objectives harder or easier? Demonstration that they can be easier.
Quality assessment of pareto set approximations E Zitzler, J Knowles, L Thiele Multiobjective optimization: Interactive and evolutionary approaches, 373-404	453	2008	More fundamentals on performance measures over multiple runs.

Bibliography (my work, selected)

TITLE	CITED BY	YEAR	Problem
Properties of an adaptive archiving algorithm for storing nondominated vectors J Knowles, D Corne IEEE Transactions on Evolutionary Computation 7 (2), 100-116	430	2003	Theory concerning archiving convergence with diversity. Hypervolume for selection also proposed for first time.
Techniques for highly multiobjective optimisation: some nondominated points are better than others DW Corne, JD Knowles Proceedings of the 9th annual conference on Genetic and evolutionary ...	394	2007	Revisiting simple methods for handling many objectives.
Multiobjective optimization in bioinformatics and computational biology J Handl, DB Kell, J Knowles IEEE/ACM Transactions on computational biology and bioinformatics 4 (2), 279-292	368	2007	Application survey paper establishing MOO in the bio-sciences.
Fifty years of pulsar candidate selection: from simple filters to a new principled real-time classification approach RJ Lyon, BW Stappers, S Cooper, JM Brooke, JD Knowles Monthly Notices of the Royal Astronomical Society 459 (1), 1104-1123	331	2016	Pulsar detection in astrophysics.
Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization JD Knowles Reading, UK, Diss	324	2002	Thesis containing 21 publications including non-refereed works.
Multiobjectivization by decomposition of scalar cost functions J Handl, SC Lovell, J Knowles International Conference on Parallel Problem Solving from Nature, 31-40	141	2008	Drift theory first used in MOO to show exponential/polynomial divide on different formulations of problems.
No free lunch and free leftovers theorems for multiobjective optimisation problems DW Corne, JD Knowles International Conference on Evolutionary Multi-Criterion Optimization, 327-341	108	2003	No free lunch theorems considered from MOO perspective

Bibliography (my work, selected)

TITLE	CITED BY	YEAR	Problem
On sequential online archiving of objective vectors M López-Ibáñez, J Knowles, M Laumanns International Conference on Evolutionary Multi-Criterion Optimization, 46-60	87	2011	Theory of size-bounded archives with diversity.
Multiobjective optimization: When objectives exhibit non-uniform latencies R Allmendinger, J Handl, J Knowles European Journal of Operational Research 243 (2), 497-513	59	2015	Asynchronous MOO algorithms because objectives do not have uniform latency.
Machine Decision Makers as a Laboratory for Interactive EMO M López-Ibánez, J Knowles Evolutionary Multi-Criterion Optimization 9019, 295-309	36	2015	How to assess preference-learning or preference-controlled EMO algorithms.
On Handling Ephemeral Resource Constraints in Evolutionary Search R Allmendinger, J Knowles Evolutionary Computation 21 (3), 497-531	22	2013	Transient (or ephemeral) constraints exist in many real problems – theory & framework.
Modes of problem solving with multiple objectives: Implications for interpreting the pareto set and for decision making J Handl, J Knowles Multiobjective Problem Solving from Nature: From Concepts to Applications ...	20	2007	How do multiple objectives arise in a problem, and what solutions concepts are appropriate for these different origins?
Deep optimisation: Solving combinatorial optimisation problems using deep neural networks JR Caldwell, RA Watson, C Thies, JD Knowles arXiv preprint arXiv:1811.00784	19	2018	Optimization method based on symbiogenetic theory of evolution – using deep learning to induce cooperation.
Policy learning in resource-constrained optimization R Allmendinger, J Knowles Proceedings of the 13th annual conference on Genetic and evolutionary ...	7	2011	Reinforcement learning-based parameter control (state-based vs simple bandits).