



Invited Review

Metaheuristics “In the Large”

Jerry Swan^{a,*}, Steven Adriaensen^b, Alexander E. I. Brownlee^c, Kevin Hammond^d,
Colin G. Johnson^e, Ahmed Kheiri^f, Faustyna Krawiec^g, J. J. Merelo^h, Leandro L. Minkuⁱ,
Ender Özcan^e, Gisele L. Pappa^j, Pablo García-Sánchez^h, Kenneth Sörensen^k, Stefan Voß^l,
Markus Wagner^m, David R. Whiteⁿ

^a University of York, Heslington, York YO10 5DD, UK

^b Albert-Ludwigs-Universität Freiburg, 79110 Freiburg, Germany

^c University of Stirling, Stirling, FK9 4LA, UK

^d Hylomorph Solutions, Glasgow, G2 4JR, UK

^e University of Nottingham, Nottingham, NG8 1BB, UK

^f Lancaster University, Lancaster, LA1 4YW, UK

^g University of Cambridge, Cambridge, CB3 0FD, UK

^h Universidad de Granada, 18071, Granada, ES

ⁱ University of Birmingham, Birmingham, B15 2TT, UK

^j Universidade Federal de Minas Gerais, Belo Horizonte, 31270-010, BR

^k Universiteit Antwerpen, Prinsstraat 13, 2000 Antwerp, NL

^l Universität Hamburg, 20146 Hamburg, DE

^m University of Adelaide, Adelaide, SA 5005, AU

ⁿ University of Sheffield, Sheffield S10 2TN, UK



ARTICLE INFO

Article history:

Received 7 November 2019

Accepted 25 May 2021

Available online 6 June 2021

Keywords:

Evolutionary Computation

Operational Research

Heuristic design

Heuristic methods

Architecture

Frameworks

Interoperability

ABSTRACT

Following decades of sustained improvement, metaheuristics are one of the great success stories of optimization research. However, in order for research in metaheuristics to avoid fragmentation and a lack of reproducibility, there is a pressing need for stronger scientific and computational infrastructure to support the development, analysis and comparison of new approaches. To this end, we present the vision and progress of the Metaheuristics “In the Large” project. The conceptual underpinnings of the project are: truly extensible algorithm templates that support reuse without modification, white box problem descriptions that provide generic support for the injection of domain specific knowledge, and remotely accessible frameworks, components and problems that will enhance reproducibility and accelerate the field’s progress. We argue that, via such principled choice of infrastructure support, the field can pursue a higher level of scientific enquiry. We describe our vision and report on progress, showing how the adoption of common protocols for all metaheuristics can help liberate the potential of the field, easing the exploration of the design space of metaheuristics.

© 2021 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

* Corresponding author.

E-mail addresses: jerry.swan@york.ac.uk (J. Swan), adriaens@cs.uni-freiburg.de (S. Adriaensen), sbr@cs.stir.ac.uk (A.E.I. Brownlee), kevin.hammond@hylomorph-solutions.com (K. Hammond), Colin.Johnson@nottingham.ac.uk (C.G. Johnson), a.kheiri@lancaster.ac.uk (A. Kheiri), fmk31@cam.ac.uk (F. Krawiec), jmerelo@kal-el.ugr.es (J.J. Merelo), l.l.minku@cs.bham.ac.uk (L.L. Minku), ender.ozcan@nottingham.ac.uk (E. Özcan), glpappa@dcc.ufmg.br (G.L. Pappa), pablogarcia@ugr.es (P. García-Sánchez), kenneth.sorensen@uantwerpen.be (K. Sörensen), stefan.voss@hamburg.de (S. Voß), markus.wagner@adelaide.edu.au (M. Wagner), d.r.white@sheffield.ac.uk (D.R. White).

1. Introduction

Optimization problems have myriad real-world applications (Hoos & Stützle, 2005) and have motivated a wealth of research since before the advent of the digital computer (Dantzig, 1990). Recent decades have seen enormous progress in the discipline of metaheuristic optimization¹. In contrast to exact approaches that

¹ In this article (and in the spirit of Sörensen and Glover, Sörensen & Glover (2013)), we reserve the term “metaheuristic” for the generic, cross domain framework and the term “heuristic” for a customization of such a framework to one or more specific domains.

guarantee optimality, “a metaheuristic is an iterative master process that guides and modifies the operations of subordinate heuristics to efficiently produce high-quality solutions. At each iteration, it manipulates either a complete (or partial) single solution or else a collection of such solutions. The subordinate heuristics may be high or low level procedures, or a simple local search, or just a construction method. The family of metaheuristics includes, but is not limited to: adaptive memory procedures, tabu search, swarm intelligence, greedy randomized adaptive search, variable neighborhood search, evolutionary methods, genetic algorithms, scatter search, neural networks, simulated annealing, and their hybrids” (taken from Voß et al., 1999, p. IX; see also Pellerin, Perrier, & Berthaut, 2020; Sörensen & Glover, 2013). One of the major advantages of metaheuristics is that they are abstract search methods (Senington & Duke, 2013): the underlying search logic can be applied to any problem which can be decomposed into a few elementary aspects, namely solution representation, solution quality evaluation and some notion of locality. The latter denotes the ability to generate neighboring solutions via a heuristically-informed function of one or more incumbent solutions.

Very broadly speaking, one might distinguish between *classical* OR and metaheuristic approaches with respect to the former's emphasis on *analytic* methods and the latter's emphasis on *empirical* ones (Swan, De Causmaecker, Martin, & Özcan, 2018). An analytic approach uses problem domain information (and typically also *a priori* human ingenuity) to derive effective algorithms for search components — edge-assembly (EAX) crossover for the Travelling Salesman Problem (TSP) being one such example (Nagata & Kobayashi, 1997). Indeed, the pre-eminent successes of OR tend to arise from a direct match between the solver (e.g. linear programming) and the analytic characteristics of the problem. In contrast, the use of analytic problem characteristics to choose solvers is not part of mainstream metaheuristics. The empirical approach performs configuration tuning (by hand, using statistical design or some machine learning technique) to create a metaheuristic biased either offline by a target distribution of problem instances and/or online by the search trajectory.

Despite the significant progress in metaheuristic optimization research, it is increasingly acknowledged within the scientific community that the field suffers from a duplication of effort and siloing of research between papers, groups, and software frameworks. This lack of re-use is evident at both *conceptual* and *implementation* levels:

- **Conceptual:** an over-reliance on reasoning by metaphor (Sörensen, 2013) hides commonalities between algorithms, leading to the repeated discovery of the same ideas and heuristics, and widespread duplication of research effort.
- **Implementation:** despite past efforts in developing software frameworks, there is a tendency to re-implement metaheuristics from scratch, hindering reproducibility and replicability (Pamparà & Engelbrecht, 2015; Swan, Adriaensen, Barwell, Hammond, & White, 2019; Swan et al., 2015).

This duplication of effort limits scientific progress; instead of building a cohesive body of knowledge consisting of robust scientific conclusions, accumulated wisdom in the field is more akin to “folklore”: observations over individual algorithms and optimization problems, without a structured underlying narrative. In the words of Pamparà, there is a “throw away” research culture (Pamparà & Engelbrecht, 2015): it is difficult to locate and compare against prior-art, and there is a lack of understanding as to which heuristics work best and in what context. Despite the availability of many software libraries, it is difficult to reuse existing implementations — in particular, it is extremely difficult to combine heuristics from different libraries or incorporate domain-specific

knowledge in a general manner, which has hindered the creation of easily testable and deployable metaheuristic pipelines.

While many authors have identified and critiqued a lack of rigor and weak empirical method in the field (e.g. (Collberg, Proebsting, & Warren, 2015; Johnson, 2002; Kendall et al., 2016)), we are more concerned by the lack of generality of enquiry, and consequently the generality of conclusions. We believe that, instead of examining individual datapoints concerning a particular algorithm implementation and a problem set of an author's choosing, we should be investigating deeper scientific questions, such as:

- Q1 Why do our methods work? In particular, how can we assign credit to individual components and eliminate those that do not matter?
- Q2 How can we *analytically* arrive at a solution method given a problem description? Conversely, is it *practically useful* to map an optimisation problem of a given type into another? Is there a ubiquitous “SAT-like” problem to which many optimisation problems can usefully be reduced to and solved in practice?
- Q3 Is it possible to more fully automate the exploration of the space of metaheuristics, allowing researchers to focus on improving that automation?

To these ends, this paper describes the vision and progress of ‘Metaheuristics in the Large’ (MitL). MitL is a research community initiative (first introduced in Swan et al. (2015)) that seeks to address the lack of re-use at both conceptual and implementation level. MitL is both a synthesis and extension of existing ideas dispersed throughout the literature, and simultaneously a project producing new software tools and exemplars to show how these problems can be overcome. We draw on many contributions previously made in this direction: hyper-heuristics, constraint programming and the fundamental principles of substitutability of software components. We rely heavily on functional programming constructs to express metaheuristic components in a truly reusable way. In constructing this synthesis, we have exposed gaps in the literature that we are now closing with new contributions: in particular the MitL initiative has introduced a) the “Automated Open Closed Principle” (Swan et al., 2019), which shows how to express algorithm frameworks as ‘closed’ design spaces which can nonetheless be configured in an open-ended manner via combinatorial assembly, and b) the removal of the domain barrier from hyper-heuristics (Swan et al., 2018), essential in raising the level of genericity with respect to problem domains. These address Q3, making easier automation of the space of metaheuristics, and hence progress towards answers to questions Q1 and Q2.

This paper provides a survey, progress report and roadmap of our attempts to reduce the fragmentation of metaheuristics research, improve reproducibility, and accelerate progress through infrastructure improvement. We have made several concrete steps forward and can see the road ahead, but there are many problems left to be solved.

2. Contemporary Research Practice in Metaheuristics

One may characterize *researchers* as being broadly concerned with the scientific practice of obtaining concise explanations of empirical observations (Popper, 1963). In contrast, for *practitioners* (e.g. in industry), the goal is to maximally exploit the information obtained via research, with very little expert knowledge. We now present in more detail some of the challenges facing metaheuristic research, drawing on previous discussions in the literature. These observations motivate the MitL approach described in Section 3.

2.1. Replication and Reuse

Scientific progress in any discipline requires ready determination of the nature and merits of previous contributions, and the ability to build on the work of others to make further progress. Algorithm descriptions in many papers on metaheuristics are far from precise enough to allow independent re-implementation, and public access to the associated source code is rarely mandated by editors or programme committees. As a consequence, replication studies are very uncommon: a recent such paper, the only replication study in metaheuristics of which the authors are aware, obtained results which were an order of magnitude worse than those originally claimed (Sörensen, Arnold, & Palhazi Cuervo, 2019)².

As noted above, metaheuristics, and in particular Evolutionary Computation research, has developed something of a ‘throw away’ culture (Collberg et al., 2015; Pamparà & Engelbrecht, 2015), in which a large percentage of researchers neither build upon the research implementations of their peers nor create such reusable software artifacts themselves. This inability to consolidate is in contrast to other research areas that have successfully embraced re-use: for example the SBML standard³ in systems biology, which allows the researcher to easily create test and deployment pipelines (König, 2020); or the Taverna framework⁴ used for workflow construction in a variety of other scientific disciplines. One might wonder why metaheuristics, which enjoy a small and ubiquitous set of abstract components such as acceptance or perturbation, has seen relatively little progress in large-scale re-use. It is possible that the very simplicity of metaheuristics at the component level is in part responsible for this culture of Babel-like proliferation. Metaheuristic researchers or practitioners often choose the creation of *ad hoc* solutions to using pre-existing resources, perhaps because implementing baseline versions of (say) Simulated Annealing or Genetic Algorithms from scratch is relatively simple – provided they are not intended for reuse by others.

The last few decades have seen the development of many popular metaheuristic libraries, implemented in a variety of programming languages, some of which feature components that are (in principle) reusable at the framework level. A first collection can be found in Voß and Woodruff (2002). Taking Evolutionary Algorithms as an example, Parejo, Ruiz-Cortés, Lozano, and Fernandez (2012) provide an overview of commonly used libraries along with their performances on benchmarks, including HeuristicLab⁵ Wagner and Affenzeller (2005), ECJ⁶ Luke (2017), FOM⁷ Parejo, Racero, Guerrero, Kwok, and Smith (2003), Opt4J⁸ Lukasiewicz, Glaß, Reimann, and Teich (2011), jMetal⁹ Durillo and Nebro (2011) and JAMES¹⁰ De Beukelaer, Davenport, De Meyer, and Fack (2017). The majority of these libraries support component interoperability within their frameworks. However, a component implemented in a specific framework cannot readily be reused within, or hybridized with, another framework. Recognizing this problem, an early attempt (Merelo-Guervós, Castillo-Valdivieso, Romero-López, & García-Arenas, 2003) sought to achieve interoperability through the use of a common description language based in XML, albeit restricted in focus to evolutionary algorithms. PISA (Bleuler,

Laumanns, Thiele, & Zitzler, 2003) was another early attempt to achieve interoperability across frameworks¹¹. In PISA, the problem domain component is separated from the metaheuristic component, and implementations of those components are reusable and interoperable, communicating via a file-based textual description.

In practice the extensibility of these frameworks is limited (though uniquely, to our knowledge, a progression towards mechanisms that enable wider re-use can be seen in Cilib (Cloete, Engelbrecht, & Pamparà, 2008; Pamparà & Engelbrecht, 2015; Pamparà, Engelbrecht, & Cloete, 2008; Peer, Engelbrecht, Pamparà, & Masiye, 2005)), and, crucially, implementation often requires the modification of internal source code, presenting a barrier to distribution, reuse, and understanding for other practitioners. We are left with a fragmented set of implementations that are incapable of representing an extensible design space for metaheuristics, without requiring modification to the frameworks themselves.

2.2. Transparency

Metaphorically-inspired approaches have recently suffered strong criticism for their lack of rigor. Where the use of metaphor obscures specific solution-domain mechanisms (Sörensen, 2013) the novelty of the metaphorical contribution becomes difficult to determine. At worst, this can lead to the re-invention or renaming of mechanisms that are already well-understood. For example, it has been argued that the popular ‘Harmony search’ metaheuristic can be formulated as a simple variant of the foundational ‘Evolution Strategies’ approach (Weyland, 2010), and it has recently been claimed (Camacho-Villalón, Dorigo, & Stützle, 2019) that the ‘Intelligent Water Drops’ algorithm is similarly not novel. Such ‘explanation by metaphor’ unnecessarily obfuscates the field and makes it appear impenetrable to outsiders. In de Armas, Lalla-Ruiz, Tilahun, and Voß (2021) a pool template is proposed and used to categorize metaheuristic algorithm components permitting to analyze them in a structured way.

This problem is at least partly cultural: the ‘reward’ of publications and citations in metaheuristics is often more readily achieved by producing a method that ‘beats the competition’ than one that makes the additional effort to be transparent about the contribution of its mechanisms (Adriaensen, Brys, & Nowé, 2014). While improving on the state-of-the-art should always be a key driver for a research community, the relentless pursuit of (apparent) novelty and the ‘up-the-wall game’ (Sörensen, 2013) is counter-productive. For as long as researchers continue to labor in relative isolation, the risks of overfitting and misidentifying novelty remain present. In contrast, we propose in the following sections a more ‘bottom up’ approach. With such an approach, new solution methods can be grounded in the principled decomposition of existing ones (López-Ibáñez, Mascia, Marmion, & Stützle, 2014), thereby allowing ready identification of potential novelty.

2.3. Knowledge Discovery

Hooker (1995) and Sörensen (2013) argue that there needs to be more scientific analysis of how metaheuristics solve problems. If a metaheuristic claims to work in some way – say, for example, it is claimed that a particular operator works by moving the search out of local minima – then experiments should be performed that test this, or (even better) a theoretical justification provided. This is a particular problem for complex metaheuristics, where a number of innovations are often introduced in tandem. Compounding this issue, the existence of ‘No Free Lunch’ theorems for optimization (Wolpert & Macready, 1997) implies that metaheuristics often

² It should be emphasized that this problem is not restricted to metaheuristics: a recent study (Collberg et al., 2015) showed that in the computer science papers considered, 34.65% of them were not repeatable and the authors could not conclusively determine repeatability in another 20.87% of cases.

³ <http://sbml.org>

⁴ <https://taverna.incubator.apache.org>

⁵ <https://dev.heuristiclab.com/trac.fcgi/wiki/>

⁶ <https://cs.gmu.edu/~eclab/projects/ecj/>

⁷ <http://www.isa.us.es/fom/>

⁸ <http://opt4j.sourceforge.net/>

⁹ <http://jmetal.sourceforge.net/>

¹⁰ <http://www.jamesframework.org/>

¹¹ <http://www.tik.ee.ethz.ch/pisa/>

require domain-specific heuristics for success (Swan et al., 2018), and current practices mean that any expert knowledge on which mechanisms work well on a given problem must typically be reverse engineered from publications on a per-case basis. To move to a more generalized level of enquiry, it is necessary to combine, exchange, and reason about metaheuristics and their component parts (such as acceptance, selection or perturbation) on a far larger scale than has been possible to date. Significantly, we believe that this requires a shift in community culture from ‘individual competition’ to ‘collective knowledge discovery’ and the development of a large pool of shared experimental data from which to draw general conclusions.

One approach to deriving such general conclusions from a large pool of data is the application of data mining and machine learning (ML) techniques successfully used by other communities (e.g. meta-learning (Pappa et al., 2014)). Individual publications have applied ML methods to selecting or constructing heuristics, and provide some evidence for their efficacy: Xu, Hutter, Hoos, & Leyton-Brown (2008) proposed a portfolio solver that won several SAT competitions by automatically selecting between various state-of-the-art SAT solvers based on a learned model of their relative performance conditioned on problem properties; Thabtah and Cowling (Thabtah & Cowling, 2008) show associative classification can indicate which heuristic to use in each iteration of a personnel scheduling problem; Miranda, Prudêncio, and Pappa (2017) used fitness landscape information to decide whether to build or select a new particle swarm optimization algorithm; Nallaperuma, Wagner, and Neumann (2014, 2015) generated predictive models of the best parameters from ant colony optimization methods based on features of previously evolved instances; Malan and Engelbrecht (2014) used landscape characteristics to predict the success of a collection of PSO algorithms on unseen continuous optimization problems; Consoli, Minku, and Yao (2014) used online learning and features extracted from the fitness landscape of the problem to choose the most appropriate genetic operator and Asta and Özcan (2015); Asta, Özcan, and Curtois (2016) integrated knowledge discovery directly into the search algorithm.

More generally, Smith-Miles, Baatar, Wreford, and Lewis (2014) proposed a methodology where instances of a problem are represented by a set of features in an instance space, and machine learning algorithms used to classify the regions of the space where algorithms are expected to perform well or poorly, given many insights on algorithms strengths and weaknesses. Such feature-based approaches provide a baseline for generating and mining knowledge of relevance to metaheuristic research and practice.

The re-use of such knowledge will first require a knowledge-base and the associated effort to constantly update it. Initial efforts towards a schema for such a database was presented by Scheibenpflug, Wagner, Pitzer, and Affenzeller (2012): their Optimization Knowledge Database (OKD), contained data about algorithms, the problems they were used to solve and their parameters. The authors emphasize that populating the database is time-consuming and requires the effort of the whole community. Other works in the literature have created specific instances of such datasets (Miranda et al., 2017; Smith-Miles et al., 2014). Such a community effort to effectively create and populate a knowledge base is paramount for the success of metaheuristics mining. Given a sufficiently rich representation for components, such analysis could be carried out in a semi-automated way.

2.4. Automated Design

Contemporary scientific and engineering disciplines rely heavily on standardization and automated tools. The design of these tools and their underlying algorithms tends to be an *ad hoc* process, often regarded as an art rather than a science (Hunt & Thomas,

2001). As a consequence, the design of an algorithm is time-consuming and costly. Furthermore, the process itself is rarely documented, making it untraceable, i.e. it is often unclear what motivated certain design decisions (e.g. expert knowledge, experimentation, intuition) and which alternatives were considered. Not only do we lose potentially interesting information and insights which can be used to design algorithms in the future, it also makes the process susceptible to accidental human bias.

The automated design of algorithms has significant potential to address these issues. Unsurprisingly, attempts to (partially) automate algorithm design, in one form or the other, are ubiquitous and can be traced back to the origins of computational intelligence (e.g. program synthesis (Manna & Waldinger, 1980), genetic programming (Koza, 1992), swarm algorithms (Khichane, Albert, & Solnon, 2008), algorithm selection (Rice, 1976), algorithm configuration (Birattari, Stützle, Paquete, & Varrentrapp, 2002)). However, the application of these techniques has thus far been largely a privilege of experts, restricted to isolated case studies, and is far from a standard practice in algorithmics.

The metaheuristics community is no exception. While metaheuristics are most commonly designed manually, the idea of automating this process is hardly new, and has been actively pursued for almost two decades in the hyper-heuristics (Burke et al., 2013) and the algorithm configuration (Stützle & López-Ibáñez, 2019) communities. An important aspect of our vision (see Section 4) is to facilitate the integration and further development of these design automation techniques.

2.5. Scalability

Historically, computing systems have tended to get faster at an exponential rate. Software performance automatically scaled along, without requiring any additional efforts from the developer. The situation is no longer so simple: contemporary systems, rather than getting faster and faster, are able to do more and more work in parallel (Sutter, 2005). To take advantage of increasing parallel processing capabilities, computations must be subdivided into a set of interdependent tasks to be executed efficiently in parallel across multiple cores and/or across networked machines. In computer science in general, much human effort has been invested in algorithm-specific parallelization strategies.

Scalability is also an issue when solving ever larger problem instances: despite the increase in computing power, it is hard to solve large instances of many practical optimization problems. This will of course always be the case – contrary to other computational domains, the field of combinatorial optimization will never have “enough” computing power. Fortunately, many popular metaheuristics are ‘embarrassingly parallel’: for example, determining the fitness of each population member in evolutionary approaches can be readily parallelized; strictly from the performance point of view and depending on latency and throughput, this simplistic approach might not be the most efficient; however, the fact that it can be done at all shows that there are parallel approaches which are functionally equivalent to sequential ones and that have a straightforward implementation.

Currently many metaheuristic methods rely on parallelism at a specific level of abstraction, typically by parallelizing either fitness evaluation or part of a population of solutions using an island model. Both these approaches are limiting in the assumptions they make about the complexity of metaheuristics: to achieve sophistication beyond previous applications may, for example, require much more complex and involved search operators and we may wish to parallelize a metaheuristic not just at these fixed and algorithm-specific levels, but to the greatest extent possible, i.e. at the level of individual components; in this sense, it might also be convenient to simply use underlying concurrent or parallel mod-

els, such as communicating sequential processes (Guervós, Valdez, & Galeano, 2020).

3. The ‘Metaheuristics in the Large’ Approach

The challenges of the previous section motivate the creation of infrastructure support for community-wide sharing of problems, metaheuristic frameworks and heuristic components. In general, there are many good tools and libraries already in existence. We do not propose to reinvent or replace them; we are not proposing “just another library” but rather a different way of structuring and implementing metaheuristics research.

Our proposed approach has three conceptual underpinnings:

- **Extensible and re-usable framework templates** To support open-ended innovation and provide true reusability, these templates (or any other suitable problem description language) must be configurable via a palette of components that is extensible *without requiring the modification of existing code*. To support such extensibility and the automated configuration of these templates, we require a stronger notion of interoperability than existing software: there must be infrastructure support for *state-threading*, i.e. passing framework or component-specific state (such as the temperature in simulated annealing) via a dedicated mechanism. This strong notion of extensibility is described in Section 3.1. Our approach facilitates reuse through this extensibility, and transparency and automated design by explicitly encapsulating component behavior (rather than relying on metaphorical description) and allowing machine-inspection of behaviors.
- **White box problem descriptions** Having embraced the necessity of state threading, it follows that we can thread more than merely empirically-obtained data relating to the search trajectory. In particular, threaded state can include *analytic* information, such as declarative/white box problem descriptions. This analytic information can be used to guide algorithm selection or construction in a more informed manner than has traditionally been embraced by the hyper-heuristics community (Swan et al., 2018). We discuss this further in Section 3.2.
- **Remotely accessible frameworks, components and problems** By building upon the two concepts above, it is possible to configure pre-existing, remotely-hosted, algorithm frameworks with some (potentially newly-devised) collection of heuristic components. The practical obstacle to further progress is then the relatively procedural one of community agreement on definitions for component interfaces and communication protocols. For inspiration, we look to work on (Service-Oriented Architecture), which we discuss in Section 3.3. This enables widespread reuse, replicability, and shared knowledge discovery.

3.1. Re-usable Framework Templates

The main obstacle to the open-ended extension and automated composition of existing implementations of metaheuristic components is that they suffer from an intrinsic lack of modularity. In this section, we illustrate why this is an issue for research ‘in the large’ and describe the proposed solution. In part, this is due to the lack of adoption of best practice from software engineering (Guervós et al., 2020). In the case of metaheuristics, there is an added complexity due to state dependencies between different algorithms components, which we now describe.

Framework configuration can be defined in general terms by expressing frameworks as higher-order functions that take components as parameters. For example, a possible function signature¹²

for acceptance for some generic candidate solution Sol is:

$accept : incumbent : Sol \times incoming : Sol \rightarrow Boolean$

Listing 1 gives a simple local search framework that allows for three design decisions, viz. the choice of perturbation, acceptance and termination conditions. In order to support alternative designs, local search is a *higher-order function*: it takes as arguments separate *functions* for *perturb*, *accept* and *finished* and returns a candidate solution of type Sol . As described above, *accept* takes as argument a pair of candidate solutions (i.e. the incumbent and incoming solutions) and returns the preferred one, denoted in the following listing as:

$accept : (Sol, Sol) \Rightarrow Sol$

perturb and *finished* are defined in the correspondingly obvious manner.

Each specific triple of components (*perturb*, *accept*, *isFinished*) used to configure the framework corresponds to a specific local search algorithm. This allows us to concisely specify a combinatorial design space of alternative component configurations, and also makes design space commonalities explicit. In order for a framework to permit the substitution of different choices for each component, components must ultimately conform to some well-defined interface, e.g. the higher-order function arguments to the framework have some signature that is fixed *a priori*. In our example, for candidate solution type Sol (e.g. a list of cities in the Traveling Salesperson Problem) perturbation is assumed to have type $Sol \rightarrow Sol$.

However, such fixed signatures are problematic if we wish such frameworks to be ‘closed to modification’, i.e. be able to accommodate unanticipated component dependencies without requiring changes to framework code. The need for such modification is clearly incompatible with the MitL goal of frameworks that can be both shared across the research community and be configured (perhaps automatically) with new components. As a concrete example: suppose we now wish to incorporate a further heuristic that requires information about the search trajectory, e.g. a tabu list of solutions (Glover & Laguna, 1997) that promotes search diversification. We are therefore required to change the implementation of local search to keep track of the trajectory. Listing 2, gives a revised version in which the history list of previous incumbent solutions is denoted by $[Sol]$.

The modified implementation now supports solution-based tabu mechanisms, but the issue of course persists if we wish to incorporate components which require new state dependencies, for example Metropolis-Hastings acceptance, which requires some measure of ‘temperature’ (Kirkpatrick et al., 1983) to be statefully maintained. In the general case, we clearly cannot anticipate in advance what information will be required by some component yet to be devised. These are examples of *environmental state*, which provides the context for decisions made by the search process. For extensibility, it is therefore necessary for support for environmental state to be open-ended, i.e. for frameworks to be configurable with components that access aspects of environmental state that are not known at the time of framework implementation. Principled handling of environment state is key to metaheuristic modularization, and is therefore essential for both component interoperability and scalability in automated construction of metaheuristics. The technical specifics of MitL support for this approach are described in an associated publication (Swan et al., 2019) and summarized in Appendix A. Software exemplars of the proposed infrastructure support are publicly available, as described in Appendix B.

3.2. White box Problem Representations

It is well-known that the exploitation of problem information is key in rendering optimization problems tractable (Wolpert &

¹² The *signature* of a function is the formal description of its parameter and return types.

Macready, 1997). As such, a principal challenge lies in devising frameworks and solvers that support injection of problem information to drive the search process, without incurring loss of genericity (i.e. they can be applied to many different problems).

By including *white box problem descriptions* as part of the environmental state, it is possible to define frameworks in terms of rich domain information, allowing the aforementioned challenge to be tackled using an open-ended combination of human ingenuity and automation.

What Kind of Information Can Be Exploited?

In principle, any machine-readable knowledge could be exploited to bias the search, to synthesize feasible operators, etc. As discussed in the introduction, this knowledge can be split into two categories:

Analytic knowledge about intrinsic features of the problem.
Empirical knowledge gained through experience, i.e. experimentation.

Many examples of successfully exploiting a combination of analytic and empirical knowledge can be found in the literature:

- Reactive tabu search (Battiti & Tecchiolli, 1994) exploits knowledge about the presence of specific substructures in candidate solutions to diversify the search, and uses trajectory information to adapt the tabu tenure parameter dynamically.
- Variable neighborhood descent (Hansen & Mladenović, 2001) exploits knowledge about the relative sizes of multiple domain-specific neighbor relations to (local) search them more efficiently, and uses empirical information (a candidate solution being locally optimal/improving) to switch between neighbor relations.
- Matheuristics (e.g. Addis, Carello, & Ceselli, 2013; Nikzad, Bashiri, & Abbasi, 2021) typically combine analytical approaches such as ILP to solve sub-problems, with higher-level searches using empirical feedback on solution quality to build the results into solutions for a larger-scale problem.
- Solution spaces can be decomposed using techniques with their origins in mathematical programming, to increase efficiency at the metaheuristic level (Raidl, 2015).
- Constraint relaxation (e.g. (Fuellerer, Doerner, Hartl, & Iori, 2010)) typically uses analytical knowledge of the acceptable constraint bounds to allow a metaheuristic to search across infeasible regions of the space using empirical feedback on solution quality to determine when the relaxation should be reduced. Similarly, different heuristics can be targeted at different constraints, driven by analytical knowledge of the constraints themselves (Goh, Kendall, & Sabar, 2017).
- Portfolio solvers (e.g. SATzilla (Xu et al., 2008)), select between multiple solvers based on analytic features of the problem instance to be solved. The mapping from features to solvers is generated empirically, using machine learning.

Historical Development Towards White Box Approaches

While the importance of exploiting problem structure is widely recognized, arguably there is a historical aversion to do so at the hyper-heuristic level, leaving this task up to the domain-specific instantiations or low-level heuristics (Drake, Kheiri, Özcan, & Burke, 2020). Maintaining generality is often cited as the motivation for this information hiding practice. For example, Chakhlevitch & Cowling (2008) argue for the importance of limiting problem domain information in achieving cross-domain generality in selection hyper-heuristics. They argue that a framework can be applied to any problem that shares the “lowest common denominator” characteristics. While sufficient for generality, information hiding is not necessary. It is easy to see that a framework can exploit arbitrary information without loss of generality, as long as it is also capable of solving the problem without it. For instance, a general optimizer could use gradient information when available (e.g. when training neural networks) and default to a derivative-free approach otherwise.

The progression of hyper-heuristic research demonstrates an increased acknowledgment that use of white box problem descriptions is both possible and desirable. Following the pattern set by initial work (Cowling, Kendall, & Soubeiga, 2001), most of the selection hyper-heuristics studies maintain a black-box interface between the hyper-heuristic and problem domain known as the *domain barrier*. The original rationale for the domain barrier, which disallows a hyper-heuristic from retrieving any problem-specific information, was thought to be necessary for cross-domain generality. However, it has been recognized that the domain barrier might be more a problem than a feature: Ross (2014) argued that an explicit domain barrier that enforces a strict separation between the hyper-heuristic and the problem-specific aspects makes hyper-heuristics undesirable for use in large real-world applications. Furthermore, (Parkes, Özcan, & Karapetyan, 2015) and Pappa et al. (2014) suggested an increased exchange of information between the problem domain and the higher search level which could then be analyzed via data science techniques and machine learning. More advanced learning for heuristic selection has progressively been introduced (Agarwal, Colak, & Eryarsoy, 2006; Ahmed, Mumford, & Kheiri, 2019; Bengio, Lodi, & Prouvost, 2020; Burke, McCollum, Meisels, Petrovic, & Qu, 2007; Kheiri & Özcan, 2016; Qu, Burke, & McCollum, 2009; Soria-Alcaraz, Ochoa, Sotelo-Figeroa, & Burke, 2017).

Recent work throws further doubt on the necessity of the domain barrier. Swan et al. (2018) state that work in constraint-satisfaction provides abundant evidence that problems can be described in a domain-independent manner without loss of solver generality. The lack of necessity for the domain barrier was further evidenced by Kheiri (2020), who designed a hyper-heuristic utilizing extended domain information that nonetheless manages low-level heuristics in a domain-independent manner. Martin et al. (2016) designed a hyper-heuristic controlling the parameter set-

```
def localSearch(
    incumbent: Sol,
    perturb: Sol => Sol,
    accept: (Sol, Sol) => Sol,
    finished: Sol => Boolean
): Sol {

    while( not finished( incumbent ) )
        incumbent = accept( incumbent, perturb( incumbent ) )

    return incumbent;
}
```

Listing 1. Local Search framework parameterized by design choices

```

def localSearch(
    current: Sol,
    history: [Sol],
    perturb: (Sol, [Sol]) => Sol,
    accept: (Sol, Sol, [Sol]) => (Sol, [Sol]),
    finished: (Sol, [Sol]) => Boolean
): (Sol, [Sol]) {

    while( not finished( current, history ) )
        (current, history) = accept( current, perturb( current, history ),
                                   history );

    return (current, history);
}

```

Listing. 2. Explicit incorporation of solution history

tings of randomised heuristics based on an agent-based cooperative search framework. An ontology is used to translate problem-specific elements into problem-independent abstract objects.

White Box Descriptions and Automation

In order to communicate the required information to the solver and its heuristics, it is necessary to move beyond the bespoke approaches above: it must be possible to communicate arbitrary domain knowledge to an algorithm framework template. In a black box setting, life is simple: solvers can choose from a closed set of interfaces and can solve any problem that implements it. In contrast, the white box setting is completely open-ended: solvers can require whatever information they deem fit. Clearly, this presents an array of novel challenges. Most notably:

1. Deducing (analytically) which solvers can be applied to which problems.
2. Overcoming limited applicability due to interface mismatches.

This is clearly a rich topic for further research. However, we believe that some foundational aspects can be identified:

1. The use of a declarative, machine-readable language to express the information problems provide, solvers require, and their relations.
2. Automated algorithm selection and problem (re)formulation, as facilitated by white box descriptions (Swan et al., 2018).

The practical choice for such a language is rightfully an open-ended research question, but the MitL proposal (Swan et al., 2018) is that prior art in Constraint Programming provides a suitably generic baseline. Existing standards within Constraint Programming, such as XCSP3 (Boussemart, Lecoutre, & Piette, 2016) support white box descriptions of constraints that capture a very wide range of common problems. As research advances, the set of supported constraints can further expand, without limiting applicability. Ultimately, interface boundaries will fade, causing a paradigm shift, where human researchers specify components using problem information and computers assemble them into a single framework, automatically selecting the component believed to work best, based on all available analytic and empirical information.

3.3. Service-Oriented Architecture

A large-scale solution to lack of re-use lies in “Service-Oriented Science”, which applies the increasingly-widely adopted practice of service-oriented architectures (Hackney, Xu, & Ranchhod, 2006) to scientific computing. The concept is defined as “the pursuit of scientific research using distributed and interoperable services, the accessibility of these interfaces being the key to success” (Foster, 2005). By such means, researchers can discover and access services

without developing specific programmatic clients for each data source, or program. Such an approach clearly has the potential to increase scientific productivity via public and distributed services, and also to increase data analysis automation. There are many examples that attempt to boost this paradigm, such as the Open Science Grid (Altunay et al., 2011) and GLOBUS (Foster, 2005). These projects include scientific communities and globally distributed infrastructures that support scientific and integrated applications of different domains.

As we have argued above, it is highly advantageous for meta-heuristic researchers and practitioners to converge on a standard machine-readable language for problem description, experimental configuration and results. Service-Oriented Architectures (SOAs) offer several ways to build a research workflow from these elements. SOA is a computational paradigm in which agents interact using loosely coupled, coarse-grained, and autonomous components called *services* (Rotem-Gal-Oz, Bruno, & Dahan, 2012). A *service* is a distributed entity, such as a node, program or function, used to obtain a result, increasing the integration of systems that are heterogeneous in respect of operating systems, protocols or languages.

The SOA perspective promotes the creation of services that are discoverable and dynamically-bound, self-contained/modular, loosely-coupled, location-transparent and composable (Valipour, Amirzafari, Maleki, & Daneshpour, 2009). As such, SOA is clearly therefore a good fit for the process of “devolved community research” which we advocate here. Lately, SOA has seen a trend towards “microservice architectures”: distributed, cloud-based and cloud-native, these architectures follow the principle of separation of concern to create applications that are easily scalable and deployable, with a stable response and maximum availability. Several frameworks, such as the one proposed by Khalloof et al. (2018), exploit the capabilities of microservices to create scalable systems that can be used at different levels (from the desktop to the web) for optimization. However, at the time of writing there are no generally accepted standards for microservice discovery, and although they offer some advantages in terms of composability and scalability, they lack the service representation feature that would make it amenable to use within a large-scale metaheuristics framework.

SOA for Metaheuristics

Previous work on SOA for metaheuristics has mainly been concerned with the application of a specific metaheuristic, such as Genetic Algorithms, to optimize a service selection or composition based on the QoS (Quality of Service) of their execution (Rosenberg et al., 2010).

Different SOA technologies, such as web services, have been proposed for solving optimization problems via grid computing (Cox, Fairman, Xue, Wason, & Keane, 2001; Song, Keane, & Cox, 2003; Song et al., 2004), where services are defined using WSDL

(Web Services Description Language) interfaces and other transmission mechanisms (such as Remote Procedure Call (Lim, Ong, Jin, Sendhoff, & Lee, 2007) or Globus Toolkit (Imade, Morishita, Ono, Ono, & Okamoto, 2004)). ROS (Remote Optimization Service) (García-Nieto, Alba, & Chicano, 2007) was one of the first attempts to allow remote execution of metaheuristics, with inputs and outputs described via XML specification. Other metaheuristic frameworks such as HeuristicLab include plug-ins to allow parallelism and interoperability using web services (Wagner et al., 2014). GridUFO is a service-oriented framework (Munawar, Wahib, Munetomo, & Akama, 2010), but it only allows the modification of the objective function and the addition of whole algorithms, without combining existing services.

García-Sánchez, González, Castillo, Arenas, and Guervós (2013) have previously proposed a SOA for Evolutionary Algorithms. Several suggestions on different concerns about the design and development of the elements of an evolutionary algorithm using SOA were presented, such as the operator behavior, dynamism or solution representation. A specific SOA technology, OSGi, was used as an example of implementation. More recently, MOSES (Parejo, 2016) was proposed as the design of a global architecture based on service contracts, allowing the automation of the experimentation process in a metaheuristic optimization context. This architecture is based on different tools, such as specific automatic experimental description languages and statistical services, forming a contract-based chain of software components for experimental execution. These global architectures have already been proposed for other fields, such as distributed simulation (Taylor, 2019).

Most of these approaches are more or less direct mappings from the original implementation to a SOA framework; however, one key way in which the proposed approach facilitates knowledge discovery is the ability to add arbitrary instrumentation to components via the generic environment representation. In particular, this allows for data mining on metaheuristic traces. In addition, by employing our generic notion of state (which denotes one or more solutions, together with any environmental information required to represent the current algorithmic state of the search), the same framework can instantiate metaheuristics operating at different scales. For example, a composite recombination operator can choose from different types of recombination strategies, putting meta and hyper-heuristics under the same framework (e.g. in the manner of Woodward, Swan, & Martin (2014)). Having these different types of algorithms under a common framework greatly facilitates their extension and comparison.

More generally, we envision the emergence of a distributed, community driven suite of tools, providing an expanded repository of interoperable frameworks and components, bringing together researchers and practitioners across domains, unifying the field and closing the gap between scientific research and empirical practice.

4. Use-Cases for MitL

As we have intimated, a purely technical solution to the issues of metaheuristics research is insufficient: community-level engagement is also required. While we have proposed a means by which extensible algorithm templates can integrate with other frameworks, the ultimate arbiter for success is the enthusiasm of the wider research community to embrace such initiatives. In this section, we describe some of the prospective benefits of doing so.

A recent paper by Kendall et al. (2016) has, importantly, emphasized the need for good laboratory practice in optimization research. The set of practices that they advocate includes making datasets available in a standardized format, reporting the results from the individual components of a hybrid approach, describing

in a reproducible way the evaluation function and the metaheuristic used, clearly presenting computational times, the use of appropriate statistical tests, etc. We disagree with none of this. However, a driving philosophy for MitL is that it is necessary to go beyond the mere *advocacy* of good practice, to making it easy – indeed, almost inevitable – that good practice can happen. The MitL proposal is that it is possible to embed foundational support for good practice directly into the software that is used by metaheuristics researchers, consequently making the fruits of that research available to other practitioners as via their default workflow. Below, we describe some specific use cases in metaheuristic research which are facilitated by the proposed approach.

4.1. Comparison between Metaheuristics

Many papers in the metaheuristics literature compare the performance of a new metaheuristic against a small sample of other metaheuristics. Sometimes, it is made clear that the comparators chosen have been specifically selected because they represent the state-of-the-art for that particular problem area: however, in many cases this is not made clear. Furthermore, many papers simply compare the new metaheuristic against other metaheuristics of the same broad type, for example, comparing a new variant on Particle Swarm optimization (PSO) against other PSO variants.

Providing some evidence for the effectiveness of a new method is clearly important. However, as the number of metaheuristics continues to expand, comparing against a few other metaheuristics seems weak; even where an assertion is made that the comparators represent the state-of-the-art, this is usually presented as an assertion to be taken on trust, and the method used to choose the comparators is unstated. In particular, there is no guarantee that the chosen problem instances actually exhibit different landscape characteristics. We therefore propose that creators of new metaheuristics should test their metaheuristics against *all* other appropriate metaheuristics.

The overall aim is therefore to ensure that comparisons are both thorough and fair. This is clearly only realistic at a large scale if the process of comparisons is automated. Some initial progress in this direction was made several years ago: Taillard (2005) suggests that iterative metaheuristics should be compared not only for a single computational effort (e.g. giving the best solution found after a fixed number of iterations), but also continuously at each iteration¹³. Some recent work has also started to address the issue of fair comparison of algorithms by providing statistical testing frameworks which ensure that the preconditions for the various tests applied are actually met (Neumann, Swan, Harman, & Clark, 2014). This is particularly important for metaheuristics, since common assumptions (e.g. of normality) are not in general true. There is also a need to ground reported results in terms of “effect magnitude” (Neumann et al., 2014): for example, an improvement of 0.1% on the state-of-the-art may have more practical relevance for the Traveling Salesman Problem than for Bin-packing. In addition to statistical considerations, the specifics of the termination condition are obviously also a vital aspect of fair comparisons. We claim that the transparency afforded by the proposed approach is vital in ensuring that comparisons are commensurate.¹⁴

¹³ <http://mistic.heig-vd.ch/taillard/qualopt/>

¹⁴ For a community effort that promotes best practices in benchmarking, we refer the curious reader to Bartz-Beielstein et al. (2020). It discusses eight topics: clearly stated goals, well-specified problems, suitable algorithms, adequate performance measures, thoughtful analysis, effective and efficient designs, comprehensive presentations, and guaranteed reproducibility.

4.2. Testing Against Problem Instances

Establishing the effectiveness of a metaheuristic requires two components: other metaheuristics against which to compare, and a range of problem instances upon which to compare them. Again, many papers in the metaheuristics literature are considerably less than comprehensive in the range of problems which are investigated. While certain benchmark test suites are available, the application of a specific new method to these benchmarks is often done in a seemingly *ad hoc* way, with a number of examples from the benchmark chosen without any justification (Hooker, 1995).

As described in previous MitL work (Swan et al., 2018), progress in this area depends on the ability to express a wide variety of problem types in a common format. This further facilitates the creation of systems that could apply a new metaheuristic systematically over a wide range of problems. Researchers would not be limited by the amount of time it would take to set up experiments with a large number of problem instances; an automated script can work through a repository of problems, automatically applying the new metaheuristic to each appropriate example.

Of course, benchmarking metaheuristics cannot be reduced to simply counting the number of problems for which a particular metaheuristic is “better” than another one (Bartz-Beielstein et al., 2020). Instead, the aim of this would be to gather a rich set of data about the performance of each metaheuristic on a wide variety of problem types and instances. Such a knowledge database, in combination with white box descriptions of problems (and introspectable methods), could then be mined to gain deeper insights into which methods work best when (and why). In particular, it allows more general metaheuristics to be constructed automatically from more specialized ones (e.g., using algorithm selection portfolios as in Xu et al. (2008)).

4.3. Hybridization

The hybridization of solution methods has been a successful approach for combining the complementary strengths of different optimization paradigms and to reduce their individual weaknesses with the aim of obtaining more effective algorithms (see e.g., Milano and van Hentenryck (2011) for a review from the perspective of the CP-AI-OR community). Hybrid approaches can be classified according to many dimensions (Puchinger & Raidl, 2005), e.g. whether the components involved in the hybridization come from different search paradigms (usually constructive methods or exact methods such as Constraint Programming or Integer Linear Programming) or whether they are homogeneous (e.g., local search or evolutionary methods). Indeed, presenting a specific hybridization of two or more metaheuristics is a common source of novelty in metaheuristic research. Unfortunately, the way in which these hybridizations are evaluated is often a very simplistic comparison in terms of accuracy or error measures, without any attempt to attribute specific behaviors in a run of the metaheuristic to particular components, or indeed to perform any elimination of “accidental complexity” (Adriaensen et al., 2014).

Often, the way in which metaheuristics are combined is the strongest contribution of a method. For example, ‘Fair Share ILS’ (Adriaensen et al., 2014) performs well because of the synergistic interaction between its acceptance and perturbation heuristics. More generally, it is of relatively little use for perturbation to be operating strongly as a “search intensifier” if the acceptance criterion only permits large increases in solution quality. Such decisions are best informed via large-scale studies (as supported by combinatorial assembly over a range of problems and algorithm configurations) and component instrumentation (as supported via the proposed environmental state threading).

4.4. Matching Metaheuristics to Problem Types

Although there has long been interest in relating problem characteristics to solution strategies (Rice, 1976), we claim that this is an area that has been particularly hindered by the lack of re-use. One problem with these studies, valuable as they are, is that they represent a single sample point in time. As new metaheuristics are created, the value of that cross-cutting analysis becomes weaker, as researchers present new, *ad hoc* evidence for the value of a particular metaheuristic on a particular problem. Given the importance of such cross-cutting studies, we propose that it is key for the community to support a constantly-updated repository of metaheuristics and experiments, and subsequently so that the most effective metaheuristic for a particular problem area can be identified and kept up-to-date. Importantly, this would need more than just the creation of such a repository. Analysis tools would also be needed, which would mine the ever-expanding repository to find features that best predict which kind of metaheuristic is well-suited to a novel dataset. This would ideally involve the automated application of metaheuristics to problems, with the repository constantly being updated as new problems and (meta)heuristics are added.

5. Conclusion

Metaheuristics in the Large (MitL) is a community project that addresses some of the cultural and technical issues we believe are impediments to progress in metaheuristic research:

- Through the MitL component-based architecture and explicit state threading outlined in Section 3, heuristics can be described from a behavioral standpoint, moving away from an over-reliance on metaphor and the accidental re-invention of established heuristics.
- It is vital for the research community to eliminate the need to modify existing framework source code when implementing new heuristics. Whilst this is of great benefit to a practitioner, it is *essential* for the open-ended combinatorial assembly of metaheuristics. Design automation can raise the abstraction level of research from manual labor such as parameter tuning and selecting and combining heuristics, towards answering more general scientific questions.
- Within the space of a few years, Deep Learning approaches have changed the perspective on what is possible in Machine Learning. By following MitL’s approach to defining heuristics, it should be easier for practitioners to recursively define very large-scale metaheuristic architectures (e.g. exploiting parallelism) without undue concern for low-level implementation details, enabling exploitation of the large-scale parallelism of modern compute platforms.

The approaches described in Section 3 combine to provide a basis for extensible Software as a Service implementations of metaheuristics provided via stateless web-services, supporting shared framework templates which allow combinatorial assembly and comparison of metaheuristics. The language and platform agnosticism of this approach in turn addresses issues of reproducibility and scalability.

Metaheuristics are one of the great contributions to practical computer science of the last few decades. However, without interoperable frameworks for analyzing, comparing and hybridizing them, advances in the *science* of metaheuristics are few and far between. Once such frameworks are in place, we will be able to put metaheuristics on a much more experimentally rigorous footing, to advance the science of metaheuristics, and to build a communal resource that is of benefit to both practitioners and researchers in this important area of computational intelligence.

In this article, we have described the key functionality that supporting infrastructure requires. What is now needed is community consensus on the relatively procedural aspects of interoperability protocols. Editors and reviewers can then insist on a thorough and systematic application of new metaheuristics to a wide range of problems, with the attendant rich analysis possibilities that are opened up.

Acknowledgments

Many people have generously given their time to the various activities of the MitL initiative. Particular gratitude is due to Adam Barwell, John A. Clark, Patrick De Causmaecker, Emma Hart, Zoltan A. Kocsis, Ben Kovitz, Krzysztof Krawiec, John McCall, Nelishia Pillay, Kevin Sim, Jim Smith, Thomas Stützle, Eric Taillard and Stefan Wagner. J. Swan acknowledges the support of UK EPSRC grant EP/J017515/1 and the EU H2020 SAFIRE Factories project. P. García-Sánchez and J. J. Merelo acknowledges the support of TIN2017-85727-C4-2-P by the Spanish Ministry of Economy and Competitiveness. M. Wagner acknowledges the support of the Australian Research Council grants DE160100850 and DP200102364.

Appendix A. The ‘Automated Open-Closed Principle’

Existing metaheuristic implementations handle *environmental state* in different ways. Many make *ad hoc* use of non-local variables to share information between components. This runs counter to automated assembly requirements: in order to present them to a configuration tool, such dependencies must be specified manually, i.e. the configuration space cannot be derived from the implementation automatically. That being said, some implementations treat environmental state in a more principled manner (e.g. (Cahon, Melab, & Talbi, 2004; Di Gaspero & Schaefer, 2003)). Here, metaheuristics are typically described as compositions of generically typed, stateful components having a well-defined interface controlling access to the encapsulated state. A framework then explicitly passes the shared state between subordinate components. While it has been demonstrated in prior-art that such implementations can be successfully coupled to configuration tools to perform bottom-up automated assembly (Marmion, Mascia, López-Ibáñez, & Stützle, 2013; Stützle & López-Ibáñez, 2019) *in the small*, they do not support the open-ended extension we propose is required *in the large* (as explained in Section 3.1).

In software engineering, a framework which can be configured from an open-ended palette of components while remaining unchanged is said to conform to the “Open-Closed principle” (‘a framework should be *closed* to modification, but *open* to extension by new components’). We have extended this principle (Swan et al., 2019) to incorporate the behavior required to support automated design, yielding the “Automated Open Closed Principle” (AOCP); that paper discusses the issue in greater technical detail and describes experiments with a suitably equipped algorithm configurator. The adoption of the AOCP provides open-ended reuse of components, and thus a systematic approach to the automated exploration of the metaheuristic design space. A key aspect is that components must be “pure functional”, as a first approximation¹⁵ this can be interpreted as meaning:

- They do not rely on hidden state.
- For the same argument, they always return the same result.

Note that aforementioned metaheuristic implementations using global variables and/or stateful components clearly violate the

AOCP. An alternative to *state encapsulation* is the use of *state threading* (Kocsis, Brownlee, Swan, & Senington, 2015; Merelo Guervós & Valdez, 2018), as can be seen in Listing 2 where *perturb*, *accept*, and *finished* each take an additional parameter that is used to *thread* the environmental state through the search algorithm. More generally, using state threading, the desired signature for *accept* is of the form:

$$accept_{Sol, Env} : Sol \times Sol \times Env \rightarrow Sol \times Env$$

However, manually threading the state through the algorithm (as in Listing 2) of Section 3 is error prone, since the framework implementer must ensure that the correct value of the state is passed to the correct stage of the algorithm. It is instead desirable to use a mechanism that *implicitly* performs state threading in a *well-defined* and *consistent* manner.

In functional programming, the problem of state propagation is addressed via a well-known design pattern: the *State monad*. For the purposes of this article, we can simply consider a monad to be a principled means of sequencing computations whilst abstracting over possible side-effects (in this case, state manipulation).

Functional languages such as Haskell and Scala provide syntactic sugar for monads. In particular, they allow monad operations to be chained together using syntax that looks like a traditional *for* loop. This is illustrated in Listing 3, a re-formulation of our local search example in Listing 1 that uses the State monad. The state, in this case an integer representing the number of iterations, is implicitly threaded through each stage of the computation. This can be extended to yield a re-formulation of Listing 2, by defining that *Env* allows access to [*Sol*] the search trajectory. Although the algorithm in Listing 3 looks similar to its imperative counterpart, the internal state is fully encapsulated within the definition of *LocalSearch*.

By virtue of open-ended support for state dependencies, the proposed approach therefore supports bottom-up automated assembly. Such an approach is less subject to human bias than the *a priori* prescription of a particular metaheuristic and therefore has relevance to foundational knowledge discovery efforts. In other areas of design (e.g. manufacturing), standardization has allowed a shift from the design of integrated systems to the design of individual components within the system. In metaheuristics, this reflects the natural trend for incorporating specialized problem- or solution-domain knowledge, i.e. a researcher can specialize in a particular kind of component such as acceptance criteria and determine their cross-domain ubiquity.

It might be thought that a monadic workflow requires metaheuristic researchers to become expert functional programmers, so it should be emphasized that this workflow is a consequence of our proposed formulation, rather than a mandatory aspect. In particular, the intention is that core metaheuristic templates can be written monadically ‘once and for all’, allowing non-expert users of these frameworks to obtain the benefits. Another minor but pleasing property is that the explicit denotation of state makes the parameter space of a component explicit, facilitating configuration via automated tools such as *Trace* (López-Ibáñez, Dubois-Lacoste, Pérez Cáceres, Birattari, & Stützle, 2016).

This pure functional perspective also provides a number of other advantages (Hughes, 1989), of particular relevance to large-scale and automated design of metaheuristics: they make it easy to reason about behavioural equivalence and coupling between components, hence improving transparency. Determinism and lack of side-effects yields reproducibility of behavior. Furthermore, a functional treatment of metaheuristics greatly facilitates architectures which can take advantage of abundant computing resources, e.g. thread-safe parallelism (Hammond & Michaelson, 1999) or ‘Service-Oriented Architecture’ (SOA) implementation via stateless web-services, as subsequently described in Section 3.3.

¹⁵ The interested reader is referred to the wealth of literature on ‘referential transparency’.

```

type Perturb[Env,Sol]    = Sol => State[Env,Sol]
type Accept[Env,Sol]     = (Sol,Sol) => State[Env,Sol]
type IsFinished[Env,Sol] = Sol => State[Env,Boolean]

class LocalSearch[Env,Sol] {

  def apply[Sol](incumbent : Sol,
                 perturb : Perturb[Env,Sol],
                 accept : Accept[Env,Sol],
                 finished : IsFinished[Env,Sol]) : State[Env,Sol] = {
    def until(s : Sol) : State[Env,Sol] = {
      for {
        perturbed <- perturb(s)
        accepted <- accept(s, perturbed)
        c <- finished(accepted)
        result <- if (c) {
          State.pure[Env,Sol](accepted)
        } else {
          until(accepted)
        }
      } yield result
    }

    for {
      result <- until(incumbent)
    } yield result
  }
}

```

Listing 3. Local Search in Scala with State monad

To our knowledge, the first proposed use of monads for state threading in metaheuristics was as part of the MitL initiative (Swan et al., 2015), whilst the first concrete implementation subsequently appeared in CLib (Pamparà & Engelbrecht, 2015) and has since been further developed (Pamparà & Engelbrecht, 2019). Although there have been many frameworks and publications that describe ‘modular decompositions’ of metaheuristics, to the best of our knowledge only MitL and CLib employ this principled approach to open-ended state dependencies. The additional contribution of the MitL initiative in this respect is the use of the monadic approach to explicitly support automated assembly (Swan et al., 2019): of particular value in this respect is the fact that a strict type-system can be used to discriminate between stateless and stateful operations and to provide information about *which* aspects of component behaviour contribute to solution quality, this being vital for the elimination of accidental complexity.

Appendix B. MitL Software Libraries

Realizing the MitL vision of community-level research based on shared scientific infrastructure requires the development of three central building blocks:

1. Support for modular, extensible metaheuristic frameworks.
2. Machine-readable descriptions of problems, heuristic components and results.
3. A two-tier architecture defining both *Programmatic* and a *Service Oriented* interfaces, the latter being in direct correspondence with the former.

Taken together, these building blocks provide necessary support for the construction of a community knowledge base, in which fixed ‘reference versions’ of metaheuristic templates can be configured with problems and components in an open-ended manner. Although the main purpose of this paper is to describe the motivation and vision for MitL, the project has nonetheless made concrete implementation progress. The MitL repository ([https://](https://github.com/MitLware)

github.com/MitLware) contains various software libraries providing infrastructure support, together with a number of examples of how the proposed approach can be applied in practice. The infrastructure support libraries are:

- **MitLware-java**
This library contains Java interfaces for the ‘Metaheuristics in the Large’ components (Perturb, Evaluate etc), as motivated by the discussion in Section 3.1.
- **mitl-support**
This library contains general utilities, metaheuristic-specific and otherwise. The former includes random selection and sampling.
- **mitl-problem**
Example problem domains, defined in terms of the MitLware-java interfaces. The problem domains include: various bitvector problems, such as Checkerboard, Royal Road, Trap and HIFF; blocksworld; tower of Hanoi; the Iterated Prisoner’s Dilemma; Magic Square; the n -puzzle; the De Jong suite of real-vector problems; SAT; the TSP; the Travelling-Thief Problem; Windfarm placement.
- **mitl-solution**
Representations for ubiquitous candidate solution types (e.g. permutations, bit vectors and polynomials), as used in mitl-problem.

Although the following examples happen to be implemented in Java/Scala, adoption at the Service-Oriented Architecture level means that components written in other languages can nonetheless interoperate via standard serialization protocols (such as JSON or XML). For example, either or both of client or server in mitl-soa-example could be written in any language, as long as it is capable of serializing candidate solutions in JSON. The example applications include:

- **mitl-whitebox-hyper-heuristics**
As an elementary example of the approach described in “A Re-characterization of Hyper-Heuristics”, this demonstrates a white box analog of the hyflex hyper-heuristic framework which

takes as input any problem domain (examples used are SAT, bin-packing, TSP, the vehicle routing problem) generically described via the XCSP constraint programming format. It then uses heuristic pattern matching to determine if the problem constraints are isomorphic to the TSP: if so, then the problem is rewritten on the fly to TSPLib format and a dedicated TSP solver (Concorde's 'LINKERN' (Applegate, Bixby, Chvatal, & Cook, 2007)) is used, if not then the generic Choco Solver (Prud'homme, Fages, & Lorca, 2016) is invoked instead.

- **mitl-soa-example** This provides a concrete demonstration of the 'two tier architecture' described above: the MitL component interfaces defined in MitLware-java are 'lifted' to the service level via RPC (Remote-Procedure Call) support. There is thus a 1-1 correspondence between local and remote component interfaces. A metaheuristic framework can therefore be transparently configured with components that happen to be hosted remotely. A simple client-server example is provided, with remote invocation of a perturbation heuristic via JSON-RPC. The server-side implementation of perturb is actually achieved via a constraint solver, thereby giving another example of how one may freely mix between analytic 'OR-style' and empirical 'metaheuristic-style' approaches.
- **mitl-ecj-jmetal-interopability-example** The ECJ (Luke, 2010) and JMetal frameworks (Durillo & Nebro, 2011) are both popular and widely used. However, it is not an easy task to achieve interoperability between them. This application shows how both can be represented as a MitL Perturb operator, allowing either to be interoperably invoked.
- **mitl-aocp** This provides an example application of our proposed *Automated Open-Closed Principle* to automated algorithm configuration of the Traveling Salesperson Problem over a fixed algorithm framework. It uses ant-programming as a generative hyper-heuristic (Kocsis & Swan, 2017) to automatically configure a local search framework with components which have different state dependencies. Further technical specifics of enabling communal research via extensible algorithm templates are described in detail in Swan et al. (2019).
- **mitl-hyperion** This provides extensible algorithm templates for several of the evolutionary algorithms described in 'Essentials of Metaheuristics' (Luke, 2010), as combinatorially instantiated in mitl-aocp, above.

References

- Addis, B., Carello, G., & Ceselli, A. (2013). Combining very large scale and ILP based neighborhoods for a two-level location problem. *European Journal of Operational Research*, 231(3), 535–546. <https://doi.org/10.1016/j.ejor.2013.06.010>.
- Adriaenssen, S., Brys, T., & Nowé, A. (2014). Fair-share ILS: A Simple State-of-the-art Iterated Local Search Hyperheuristic. In *Proceedings of the 2014 annual conference on genetic and evolutionary computation*. In *GECCO '14* (pp. 1303–1310). New York, NY, USA: ACM. <https://doi.org/10.1145/2576768.2598285>.
- Agarwal, A., Colak, S., & Eryarsoy, E. (2006). Improvement heuristic for the flow-shop scheduling problem: An adaptive-learning approach. *European Journal of Operational Research*, 169(3), 801–815. <https://doi.org/10.1016/j.ejor.2004.06.039>.
- Ahmed, L., Mumford, C., & Kheiri, A. (2019). Solving urban transit route design problem using selection hyper-heuristics. *European Journal of Operational Research*, 274(2), 545–559.
- Altunay, M., Avery, P., Blackburn, K., Bockelman, B., Ernst, M., Fraser, D., ... Open Sci Grid Executive Board (2011). A Science Driven Production Cyberinfrastructure the Open Science Grid. *Journal of GRID Computing*, 9(2, Sp. Iss. SI), 201–218.
- Applegate, D. L., Bixby, R. E., Chvatal, V., & Cook, W. J. (2007). *The Traveling Salesman Problem: A Computational Study* (Princeton Series in Applied Mathematics). USA: Princeton University Press.
- Asta, S., & Özcan, E. (2015). A tensor-based selection hyper-heuristic for cross-domain heuristic search. *Information Sciences*, 299, 412–432. <https://doi.org/10.1016/j.ins.2014.12.020>.
- Asta, S., Özcan, E., & Curtois, T. (2016). A tensor based hyper-heuristic for nurse rostering. *Knowledge-Based Systems*, 98, 185–199. <https://doi.org/10.1016/j.knsys.2016.01.031>.
- Bartz-Beielstein, T., Doerr, C., Bossek, J., Chandrasekaran, S., Eftimov, T., Fischbach, A., ... Weise, T. (2020). Benchmarking in optimization: Best practice and open issues. *arXiv preprint arXiv:2007.03488*.
- Battiti, R., & Tecchioni, G. (1994). The reactive tabu search. *Journal on Computing*, 6(2), 126–140.
- Bengio, Y., Lodi, A., & Prouvost, A. (2020). Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2020.07.063>.
- Birattari, M., Stützle, T., Paquete, L., & Varrenttrapp, K. (2002). A racing algorithm for configuring metaheuristics. In *Proceedings of the genetic and evolutionary computation conference*. In *GECCO '02* (pp. 11–18). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc..
- Bleuler, S., Laumanns, M., Thiele, L., Zitzler, E., et al. (2003). PISA—A Platform and Programming Language Independent Interface for Search Algorithms. In C. M. Fonseca, et al. (Eds.), *Conference on evolutionary multi-criterion optimization (EMO 2003)*. In *LNCS*: 2632 (pp. 494–508). Berlin: Springer.
- Boussemart, F., Lecoutre, C., & Piette, C. (2016). XCSP3: an integrated format for benchmarking combinatorial constrained problems. *CoRR*, abs/1611.03398.
- Burke, E. K., Gendreau, M., Hyde, M., Kendall, G., Ochoa, G., Özcan, E., & Qu, R. (2013). Hyper-heuristics: a survey of the state of the art. *Journal of the Operational Research Society*, 64(12), 1695–1724. <https://doi.org/10.1057/jors.2013.71>.
- Burke, E. K., McCollum, B., Meisels, A., Petrovic, S., & Qu, R. (2007). A graph-based hyper-heuristic for educational timetabling problems. *European Journal of Operational Research*, 176(1), 177–192.
- Cahon, S., Melab, N., & Talbi, E.-G. (2004). ParadisEO: A framework for the reusable design of parallel and distributed metaheuristics. *Journal of Heuristics*, 10(3), 357–380.
- Camacho-Villalón, C. L., Dorigo, M., & Stützle, T. (2019). The intelligent water drops algorithm: why it cannot be considered a novel algorithm - A brief discussion on the use of metaphors in optimization. *Swarm Intelligence*, 13(3–4), 173–192. <https://doi.org/10.1007/s11721-019-00165-y>.
- Chakhlevitch, K., & Cowling, P. (2008). *Hyperheuristics: Recent developments*. In C. Cotta, M. Sevaux, & K. Sörensen (Eds.) (pp. 3–29). Berlin Heidelberg: Springer.
- Cloete, T., Engelbrecht, A. P., & Pamparà, G. (2008). Cilib: A collaborative framework for computational intelligence algorithms - part II. In *Proceedings of the international joint conference on neural networks, IJCNN 2008, part of the IEEE world congress on computational intelligence, WCCI 2008, Hong Kong, China, June 1–6, 2008* (pp. 1764–1773). IEEE. <https://doi.org/10.1109/IJCNN.2008.4634037>.
- Collberg, C., Proebsting, T., & Warren, A. M. (2015). Repeatability and beneficence in computer systems research. *Technical Report*. University of Arizona.
- Consoli, P., Minku, L. L., & Yao, X. (2014). Dynamic selection of evolutionary algorithm operators based on online learning and fitness landscape metrics. In *10th international conference on simulated evolution and learning, volume 8886 of LNCS*. Springer.
- Cowling, P., Kendall, G., & Soubeiga, E. (2001). A hyperheuristic approach to scheduling a sales summit. In E. Burke, & W. Erben (Eds.), *Practice and theory of automated timetabling iii*. In *Lecture Notes in Computer Science*: 2079 (pp. 176–190). Berlin Heidelberg: Springer.
- Cox, S. J., Fairman, M. J., Xue, G., Wason, J. L., & Keane, A. J. (2001). The grid: Computational and data resource sharing in engineering optimisation and design search. In *30th international workshops on parallel processing (icpp 2001 workshops), 3–7 september 2001, valencia, spain* (pp. 207–212). IEEE Computer Society.
- Dantzig, G. B. (1990). *A history of scientific computing* (pp. 141–151). New York, NY, USA: ACM. <https://doi.org/10.1145/87252.88081>.
- De Beukelaer, H., Davenport, G. F., De Meyer, G., & Fack, V. (2017). James: An object-oriented java framework for discrete optimization using local search metaheuristics. *Software: Practice and Experience*, 47(6), 921–938. <https://doi.org/10.1002/spe.2459>.
- de Armas, J., Lalla-Ruiz, E., Tilahun, S. L., & Voß, S. (2021). Similarity in metaheuristics: a gentle step towards a comparison methodology. *Natural Computing*. <https://doi.org/10.1007/s11047-020-09837-9>.
- Di Gaspero, L., & Schaerf, A. (2003). Easylocal++: an object-oriented framework for the flexible design of local-search algorithms. *Software: Practice and Experience*, 33(8), 733–765. <https://doi.org/10.1002/spe.524>.
- Drake, J. H., Kheiri, A., Özcan, E., & Burke, E. K. (2020). Recent advances in selection hyper-heuristics. *European Journal of Operational Research*, 285(2), 405–428.
- Durillo, J. J., & Nebro, A. J. (2011). jMetal: A Java framework for multi-objective optimization. *Advances in Engineering Software*, 42(10).
- Foster, I. (2005). Globus Toolkit version 4: Software for service-oriented systems. In Jin, H and Reed, D and Jiang, W (Ed.), *Network and Parallel Computing Proceedings*. In *Lecture Notes in Computer Science*: 3779 (pp. 2–13).
- Foster, I. (2005). Service-oriented science. *Science*, 308(5723), 814–817. <https://doi.org/10.1126/science.1110411>.
- Fuellerer, G., Doerner, K. F., Hartl, R. F., & Iori, M. (2010). Metaheuristics for vehicle routing problems with three-dimensional loading constraints. *European Journal of Operational Research*, 201(3), 751–759. <https://doi.org/10.1016/j.ejor.2009.03.046>.
- García-Nieto, J., Alba, E., & Chicano, J. F. (2007). Using metaheuristic algorithms remotely via ROS. In *Proceedings of the 9th annual conference on genetic and evolutionary computation*. In *GECCO '07*. New York, NY, USA: ACM. <https://doi.org/10.1145/1276958.1277239>. 1510–1510
- García-Sánchez, P., González, J., Castillo, P. A., Arenas, M. G., & Guervós, J. J. M. (2013). Service oriented evolutionary algorithms. *Soft Comput.*, 17(6), 1059–1075.

- Glover, F., & Laguna, M. (1997). *Tabu Search*. Norwell, MA, USA: Kluwer Academic Publishers.
- Goh, S. L., Kendall, G., & Sabar, N. R. (2017). Improved local search approaches to solve the post enrolment course timetabling problem. *European Journal of Operational Research*, 261(1), 17–29. <https://doi.org/10.1016/j.ejor.2017.01.040>.
- Guervós, J. J. M., Valdez, M. G., & Galeano, S. R. (2020). Implementation matters, also in concurrent evolutionary algorithms. In C. A. C. Coello (Ed.), *GECCO '20: Genetic and evolutionary computation conference, companion volume, Cancún, Mexico, July 8–12, 2020* (pp. 1591–1598). ACM. <https://doi.org/10.1145/3377929.3398120>.
- Hackney, R., Xu, H., & Ranchhod, A. (2006). Evaluating web services: Towards a framework for emergent contexts. *European Journal of Operational Research*, 173(3), 1161–1174. <https://doi.org/10.1016/j.ejor.2005.07.010>.
- Hammond, K., & Michaelson, G. (1999). *Research Directions in Parallel Functional Programming*. Springer.
- Hansen, P., & Mladenović, N. (2001). Variable neighborhood search: Principles and applications. *European Journal of Operational Research*, 130(3), 449–467.
- Hooker, J. (1995). Testing heuristics: We have it all wrong. *Journal of Heuristics*, 1(1), 33–42. <https://doi.org/10.1007/bf02430364>.
- Hoos, H., & Stützle, T. (2005). *Stochastic Local Search: Foundations & Applications*. Morgan Kaufmann.
- Hughes, J. (1989). Why functional programming matters. *The Computer Journal*, 32(2), 98–107. <https://doi.org/10.1093/comjnl/32.2.98>.
- Hunt, A., & Thomas, D. (2001). The art in computer programming. *The Pragmatic Programmers, LLC*.
- Imade, H., Morishita, R., Ono, I., Ono, N., & Okamoto, M. (2004). A grid-oriented genetic algorithm framework for bioinformatics. *New Gen. Comput.*, 22(2), 177–186.
- Johnson, D. S. (2002). A theoretician's guide to the experimental analysis of algorithms. *Data structures, near neighbor searches, and methodology: fifth and sixth DIMACS implementation challenges*, 59, 215–250.
- Kendall, G., Bai, R., Blazewicz, J., De Causmaecker, P., Gendreau, M., John, R., ... Qu, R., et al. (2016). Good laboratory practice for optimization research. *Journal of the Operational Research Society*, 67(4), 676–689.
- Khalloof, H., Jakob, W., Liu, J., Braun, E., Shahoud, S., Duepmeier, C., & Hagenmeyer, V. (2018). A generic distributed microservices and container based framework for metaheuristic optimization. In *Proceedings of the genetic and evolutionary computation conference companion* (pp. 1363–1370). ACM.
- Kheiri, A. (2020). Heuristic sequence selection for inventory routing problem. *Transportation Science*, 54(2), 302–312.
- Kheiri, A., & Özcan, E. (2016). An iterated multi-stage selection hyper-heuristic. *European Journal of Operational Research*, 250(1), 77–90.
- Khichane, M., Albert, P., & Solnon, C. (2008). Integration of ACO in a constraint programming language. *Ant Colony Optimization and Swarm Intelligence*, 84–95.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671–680.
- Kocsis, Z. A., Brownlee, A. E. I., Swan, J., & Senington, R. (2015). Haiku- a Scala combinator toolkit for semi-automated composition of metaheuristics. In M. Barros, & Y. Labiche (Eds.), *Search-based software engineering. In Lecture Notes in Computer Science*: 9275 (pp. 125–140). Springer International Publishing. https://doi.org/10.1007/978-3-319-22183-0_9.
- Kocsis, Z. A., & Swan, J. (2017). Dependency injection for programming by optimization. *CoRR*, abs/1707.04016.
- König, M. (2020). Executable simulation model of the liver. *bioRxiv*. <https://doi.org/10.1101/2020.01.04.894873>.
- Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection*. MIT press.
- Lim, D., Ong, Y.-S., Jin, Y., Sendhoff, B., & Lee, B.-S. (2007). Efficient hierarchical parallel genetic algorithms using grid computing. *Future Generation Computer Systems*, 23(4), 658–670.
- López-Ibáñez, M., Mascia, F., Marmion, M.-E., & Stützle, T. (2014). A template for designing single-solution hybrid metaheuristics. In *Gecco comp '14* (pp. 1423–1426). New York, USA.
- López-Ibáñez, M., Dubois-Lacoste, J., Pérez Cáceres, L., Birattari, M., & Stützle, T. (2016). The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3, 43–58. <https://doi.org/10.1016/j.orp.2016.09.002>.
- Lukasiewicz, M., Glaß, M., Reimann, F., & Teich, J. (2011). Opt4j - a modular framework for meta-heuristic optimization. In *Proceedings of the genetic and evolutionary computing conference (gecco 2011)* (pp. 1723–1730). Dublin, Ireland.
- Luke, S. (2010). The ECJ owner's manual. <http://www.cs.gmu.edu/~eclab/projects/ecj>.
- Luke, S. (2017). ECJ then and now. In *Proceedings of the genetic and evolutionary computation conference companion*. In *GECCO '17* (pp. 1223–1230). New York, NY, USA: ACM.
- Malan, K. M., & Engelbrecht, A. P. (2014). *Fitness landscape analysis for metaheuristic performance prediction*. In H. Richter, & A. Engelbrecht (Eds.) (pp. 103–132)). Berlin, Heidelberg: Springer.
- Manna, Z., & Waldinger, R. (1980). A deductive approach to program synthesis. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 2(1), 90–121.
- Marmion, M.-E., Mascia, F., López-Ibáñez, M., & Stützle, T. (2013). Automatic design of hybrid stochastic local search algorithms. In *International workshop on hybrid metaheuristics* (pp. 144–158). Springer.
- Martin, S., Ouelhadj, D., Beullens, P., Özcan, E., Juan, A. A., & Burke, E. K. (2016). A multi-agent based cooperative approach to scheduling and routing. *European Journal of Operational Research*, 254(1), 169–178. <https://doi.org/10.1016/j.ejor.2016.02.045>.
- Merelo-Guervós, J. J., Castillo-Valdivieso, P. Á., Romero-López, G., & García-Arenas, M. (2003). Specifying evolutionary algorithms in XML. In *International work-conference on artificial neural networks* (pp. 502–509). Springer.
- Merelo Guervós, J. J., & Valdez, J. M. G. (2018). Mapping evolutionary algorithms to a reactive, stateless architecture: using a modern concurrent language. In H. E. Aguirre, & K. Takadama (Eds.), *Proceedings of the genetic and evolutionary computation conference companion, GECCO 2018, Kyoto, Japan, July 15–19, 2018* (pp. 1870–1877). ACM. <https://doi.org/10.1145/3205651>.
- Miranda, P. B., Prudêncio, R. B., & Pappa, G. L. (2017). H3ad: A hybrid hyper-heuristic for algorithm design. *Information Sciences*, 414(Supplement C), 340–354. <https://doi.org/10.1016/j.ins.2017.05.029>.
- (pp. 1–9). (2011). Berlin, Germany: Springer.
- Munawar, A., Wahib, M., Munetomo, M., & Akama, K. (2010). The design, usage, and performance of GridUFO: A grid based unified framework for optimization. *Future Generation Computer Systems*, 26(4), 633–644.
- Nagata, Y., & Kobayashi, S. (1997). Edge assembly crossover: A high-power genetic algorithm for the travelling salesman problem. In *Proceedings of the 7th international conference on genetic algorithms, MI, USA*.
- Nallaperuma, S., Wagner, M., & Neumann, F. (2014). *Parameter prediction based on features of evolved instances for ant colony optimization and the traveling salesman problem*. In T. Bartz-Beielstein, J. Branke, B. Filipič, & J. Smith (Eds.) (pp. 100–109). Cham: Springer International Publishing.
- Nallaperuma, S., Wagner, M., & Neumann, F. (2015). Analyzing the effects of instance features and algorithm parameters for maxmin ant system and the traveling salesman problem. *Frontiers in Robotics and AI*, 2, 18. <https://doi.org/10.3389/frobt.2015.00018>.
- Neumann, G., Swan, J., Harman, M., & Clark, J. A. (2014). The executable experimental template pattern for the systematic comparison of metaheuristics: Extended abstract. In *Proceedings of the companion publication of the 2014 annual conference on genetic and evolutionary computation*. In *GECCO Comp '14* (pp. 1427–1430). New York, NY, USA: ACM. <https://doi.org/10.1145/2598394.2609850>.
- Nikzad, E., Bashiri, M., & Abbasi, B. (2021). A math heuristic algorithm for stochastic home health care planning. *European Journal of Operational Research*, 288(3), 753–774. <https://doi.org/10.1016/j.ejor.2020.06.040>.
- Pamparà, G., & Engelbrecht, A. (2015). Towards a generic computational intelligence library: Preventing insanity. In *2015 IEEE symposium series on computational intelligence: IEEE workshop on computational intelligence tools (2015 IEEE WCIT)*. Cape Town, South Africa.
- Pamparà, G., & Engelbrecht, A. P. (2019). Evolutionary and swarm-intelligence algorithms through monadic composition. In *Proceedings of the genetic and evolutionary computation conference companion*. In *GECCO '19* (pp. 1382–1390). New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3319619.3326845>.
- Pamparà, G., Engelbrecht, A. P., & Cloete, T. (2008). Cilib: A collaborative framework for computational intelligence algorithms - part I. In *Proceedings of the international joint conference on neural networks, IJCNN 2008, part of the IEEE world congress on computational intelligence, WCCI 2008, Hong Kong, China, June 1–6, 2008* (pp. 1750–1757). IEEE. <https://doi.org/10.1109/IJCNN.2008.4634035>.
- Pappa, G. L., Ochoa, G., Hyde, M. R., Freitas, A. A., Woodward, J., & Swan, J. (2014). Contrasting meta-learning and hyper-heuristic research: the role of evolutionary algorithms. *Genetic Programming and Evolvable Machines*, 15(1), 3–35.
- Parejo, J. A. (2016). MOSES: A metaheuristic optimization framework ecosystem. *AI Commun.*, 29(1), 223–225. <https://doi.org/10.3233/AIC-140646>.
- Parejo, J. A., Racero, J., Guerrero, F., Kwok, T., & Smith, K. A. (2003). *Fom: A framework for metaheuristic optimization*. In P. M. A. Slood, D. Abramson, A. V. Bogdanov, Y. E. Gorbachev, J. J. Dongarra, & A. Y. Zomaya (Eds.) (pp. 886–895). Berlin, Heidelberg: Springer.
- Parejo, J. A., Ruiz-Cortés, A., Lozano, S., & Fernandez, P. (2012). Metaheuristic optimization frameworks: a survey and benchmarking. *Soft Computing*, 16(3), 527–561.
- Parkes, A. J., Özcan, E., & Karapetyan, D. (2015). A software interface for supporting the application of data science to optimisation. In *International conference on learning and intelligent optimization* (pp. 306–311). Springer.
- Peer, E. S., Engelbrecht, A. P., Pamparà, G., & Masiye, B. S. (2005). Ciclops: computational intelligence collaborative laboratory of pantological software. In *Proceedings 2005 IEEE swarm intelligence symposium. SIS 2005* (pp. 130–137).
- Pellerin, R., Perrier, N., & Berthaut, F. (2020). A survey of hybrid metaheuristics for the resource-constrained project scheduling problem. *European Journal of Operational Research*, 280(2), 395–416. <https://doi.org/10.1016/j.ejor.2019.01.063>.
- Popper, K. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge classics. Routledge.
- Prud'homme, C., Fages, J.-G., & Lorca, X. (2016). Choco solver documentation. TASC, INRIA Rennes, LINA CNRS UMR 6241, COSLING S.A.S.
- Puchinger, J., & Raidl, G. (2005). Combining metaheuristics and exact algorithms in combinatorial optimization: A survey and classification. In *Artificial intelligence and knowledge engineering applications: a bioinspired approach* (pp. 113–124). Springer.
- Qu, R., Burke, E. K., & McCollum, B. (2009). Adaptive automated construction of hybrid heuristics for exam timetabling and graph colouring problems. *European Journal of Operational Research*, 198(2), 392–404.
- Raidl, G. R. (2015). Decomposition based hybrid metaheuristics. *European Journal of Operational Research*, 244(1), 66–76. <https://doi.org/10.1016/j.ejor.2014.12.005>.
- Rice, J. R. (1976). The algorithm selection problem. In M. Rubinoff, & M. C. Yovits (Eds.), *Advances in computers*: 15 (pp. 65–118). Elsevier.

- Rosenberg, F., Müller, M. B., Leitner, P., Michlmayr, A., Bouguettaya, A., & Dustdar, S. (2010). Metaheuristic optimization of large-scale qos-aware service compositions. In *2010 IEEE international conference on services computing, SCC 2010, Miami, FL, USA, July 5–10, 2010* (pp. 97–104). IEEE Computer Society.
- Ross, P. (2014). *Search methodologies: Introductory tutorials in optimization and decision support techniques* (pp. 611–638). Springer US.
- Rotem-Gal-Oz, A., Bruno, E., & Dahan, U. (2012). *SOA patterns*. Manning.
- Scheibenpflug, A., Wagner, S., Pitzer, E., & Affenzeller, M. (2012). Optimization Knowledge Base: An open database for algorithm and problem characteristics and optimization results. In *Gecco '12*. NY, USA: ACM.
- Senington, R., & Duke, D. (2013). Decomposing metaheuristic operations. In R. Hinze (Ed.), *Implementation and application of functional languages*. In *LNCS 8241* (pp. 224–239). Berlin Heidelberg: Springer.
- Smith-Miles, K., Baatar, D., Wreford, B., & Lewis, R. (2014). Towards objective measures of algorithm performance across instance space. *Computers & Operations Research*, 45, 12–24.
- Song, W., Keane, A., & Cox, S. (2003). Cfd-based shape optimisation with grid-enabled design search toolkits. In *Uk e-science all hands meeting 2003* (pp. 619–627). EPSRC.
- Song, W., Ong, Y. S., Ng, H. K., Keane, A., Cox, S., & Lee, B. S. (2004). A service-oriented approach for aerodynamic shape optimisation across institutional boundaries. In *Control, automation, robotics and vision conference, 2004. icarcv 2004 8th: 3* (pp. 2274–2279).
- Sörensen, K. (2013). Metaheuristics—the metaphor exposed. *International Transactions in Operational Research*, 22(1), 3–18.
- Sörensen, K., & Glover, F. W. (2013). *Metaheuristics*. In S. I. Gass, & M. C. Fu (Eds.) (pp. 960–970). Boston, MA: Springer US.
- Sörensen, K., Arnold, F., & Palhazi Cuervo, P. (2019). A critical analysis of the “improved Clarke and Wright savings algorithm”. *International Transactions in Operational Research*, 26(1), 54–63. <https://doi.org/10.1111/itor.12443>.
- Soria-Alcaraz, J. A., Ochoa, G., Sotelo-Figeroa, M. A., & Burke, E. K. (2017). A methodology for determining an effective subset of heuristics in selection hyper-heuristics. *European Journal of Operational Research*, 260(3), 972–983.
- Stützle, T., & López-Ibáñez, M. (2019). Automated design of metaheuristic algorithms. In *Handbook of metaheuristics* (pp. 541–579). Springer.
- Sutter, H. (2005). The free lunch is over: A fundamental turn toward concurrency in software. *Dr. Dobbs Journal*, 30(3).
- Swan, J., Adriaensen, S., Barwell, A. D., Hammond, K., & White, D. R. (2019). Extending the ‘open-closed principle’ to automated algorithm configuration. *Evolutionary Computation*, 27(1), 173–193.
- Swan, J., Adriaensen, S., Bishr, M., Burke, E. K., Clark, J. A., Causmaecker, P. D., ... Yao, X. (2015). A research agenda for metaheuristic standardization. In *Proceedings of the eleventh Metaheuristics International Conference (mic)*, Agadir, Morocco.
- Swan, J., De Causmaecker, P., Martin, S., & Özcan, E. (2018). *A re-characterization of hyper-heuristics*. In L. Amodeo, E.-G. Talbi, & F. Yalaoui (Eds.) (pp. 75–89). Cham: Springer International Publishing.
- Taillard, E. D. (2005). Tutorial: Few guidelines for analyzing methods. In *Metaheuristic interantional conference (mic'05) proceedings*.
- Taylor, S. J. (2019). Distributed simulation: state-of-the-art and potential for operational research. *European Journal of Operational Research*, 273(1), 1–19. <https://doi.org/10.1016/j.ejor.2018.04.032>.
- Thabtah, F., & Cowling, P. (2008). Mining the data from a hyperheuristic approach using associative classification. *Expert Systems with Applications*, 34(2), 1093–1101.
- Valipour, M. H., Amirzafari, B., Maleki, K. N., & Daneshpour, N. (2009). A brief survey of software architecture concepts and service oriented architecture. In *2nd IEEE international conference on computer science and information technology. ICCSIT 2009* (pp. 34–38). <https://doi.org/10.1109/ICCSIT.2009.5235004>.
- Voß, S., Martello, S., Osman, I. H., & Roucairol, C. (Eds.). (1999). *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*. Kluwer, Boston. <https://doi.org/10.1007/978-1-4615-5775-3>.
- Voß, S., & Woodruff, D. L. (Eds.). (2002). *Optimization Software Class Libraries*. Kluwer, Boston. https://doi.org/10.1007/0-306-48126-X_1.
- Wagner, S., & Affenzeller, M. (2005). *Adaptive and natural computing algorithms: Proceedings of the international conference in Coimbra, Portugal, 2005* (pp. 538–541). Vienna: Springer.
- Wagner, S., Kronberger, G., Beham, A., Kommenda, M., Scheibenpflug, A., Pitzer, E., ... Affenzeller, M. (2014). *Advanced methods and applications in computational intelligence* (vol. 6, pp. 197–261). Springer.
- Weyland, D. (2010). A rigorous analysis of the harmony search algorithm: How the research community can be misled by a “novel” methodology. *Int. J. Appl. Metaheuristic Comput.*, 1(2), 50–60. <https://doi.org/10.4018/jamc.2010040104>.
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *Trans. Evol. Comp.*, 1(1), 67–82. <https://doi.org/10.1109/4235.585893>.
- Woodward, J., Swan, J., & Martin, S. (2014). The ‘Composite’ Design Pattern in Metaheuristics. In *Gecco '14 companion*. New York, USA.
- Xu, L., Hutter, F., Hoos, H. H., & Leyton-Brown, K. (2008). SATzilla: portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research*, 32, 565–606.